

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: August 17, 2014

F. Brockners
S. Bhandari
F. Maino
D. Lewis
Cisco
February 13, 2014

LISP Extensions for Segment Routing
draft-brockners-lisp-sr-01

Abstract

Segment Routing (SR) combines source routing and tunneling to steer traffic through the transit network. The Locator/ID Separation Protocol (LISP) separates IP addresses into Endpoint Identifiers (EIDs) and Routing Locators (RLOCs) and also leverages tunneling mechanisms. Mapping between EIDs and RLOCs is facilitated by the LISP mapping system. Combining LISP and SR enables the LISP mapping system to provide SR information to encapsulating routers so that traffic can be steered in the transit network or the list of segments a particular packet traverses is recorded in the packet header.

This document describes extensions required to the Locator/ID Separation Protocol (LISP) to enable a LISP mapping system to communicate list of segment identifiers or the request to record the list of segments a particular packet traverses to the encapsulating router.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on August 17, 2014.

Copyright Notice

Copyright (c) 2014 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	3
2. Conventions	3
3. Use cases that combine LIISP and SR	4
3.1. Traffic steering/traffic engineering	4
3.2. Traffic tracing	5
4. LIISP extensions to support SR	5
4.1. Deployment Scenario	7
4.2. Example ELPs	7
4.2.1. Example: ELP with only SR used	7
4.2.2. Example: ELP with SR and reencapsulating routers combined	8
5. IANA Considerations	9
6. Manageability Considerations	9
7. Security Considerations	9
8. Acknowledgements	10
9. Change log	10
10. References	10
10.1. Normative References	10
10.2. Informative References	10
Authors' Addresses	11

1. Introduction

Segment Routing (SR) allows for a flexible definition of end-to-end paths within network topologies by encoding paths as sequences of topological sub-paths, called "segments" as described in [I-D.filsfils-rtgwg-segment-routing]. Segment routing can be applied to IPv6 with a new type of routing extension header. The Locator/ID Separation Protocol [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, 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.

When LISP is combined with SR, the EID to RLOC mapping information can be extended with segment routing information. This allows for a closer correlation between the transit network, that is sometimes also referred to as the underlay network, and the overlay network. It is beyond the scope of this document to describe how the LISP mapping system obtains a segment list for a particular EID-to-RLOC mapping. This draft outlines use-cases for combining LISP and SR as well as extensions to the LISP Canonical Address Format (LCAF) for traffic engineering (LCAF type 10) [I-D.ietf-lisp-lcaf]. These extensions are to be integrated into a future revision of [I-D.ietf-lisp-lcaf].

2. Conventions

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].

This document uses the Terminology as defined in [I-D.filsfils-rtgwg-segment-routing] and [I-D.ietf-lisp-lcaf].

Abbreviations used in this document:

AFI: Address Family Identifier

EID: Endpoint Identifier

ELP: Explicit Locator Path

ETR: Egress Tunnel Router

ITR: Ingress Tunnel Router

LCAF: LISP Canonical Address Format

LISP: Locator/ID Separation Protocol

OAM: Operation Administration Maintenance

RLOC: Routing Locator

SR: Segment Routing

SID: Segment Identifier

Segment List: Ordered list of segment identifiers

3. Use cases that combine LISP and SR

Use-cases that combine LISP and SR include traffic steering/traffic engineering as well as traffic tracing in the underlay network.

3.1. Traffic steering/traffic engineering

LISP combined with SR can be used to steer traffic in the underlay network: The mapping system communicates a segment list to the LISP ingress tunnel router (ITR) when resolving the EID-to-RLOC mapping as part of a LISP Map-Reply. This extension allows the LISP mapping system to provide a list of segment identifiers to encapsulating routers so that traffic can be steered in the transit network. In a typical setup the LISP ingress tunnel router would retrieve the segment list from the mapping system along with the associated RLOC using the EID as the lookup key. The ITR encapsulates the packet to the RLOC, also including the segment list in the segment routing extension header. The packet is forwarded to the ETR using segment routing techniques. The ETR decapsulates the packet and delivers the packet to the destination EID. Given that in SR with IPv6 transport the entire segment list is available in the SR-specific extension header of the outer IPv6 header, the LISP egress tunnel router, which is the tunnel endpoint is also informed about the path a particular packet took in the transport network.

LISP with SR for traffic engineering adds to the LISP traffic engineering use-cases described in [I-D.farinacci-lisp-te]. LISP combined with SR offers traffic engineering without using reencapsulating tunnels [RFC6830]. Reencapsulating tunnels and SR with LISP are complementary traffic engineering techniques and could be combined. SR could for example be used in an explicit locator

path (ELP) to further traffic engineer a path between two reencapsulating routers.

3.2. Traffic tracing

LISP combined with SR can be used to get more information about the path a packet took in the underlay network without sending extra probe traffic. When SR is applied to IPv6, the segment list describing the path that a packet takes through the network can be recorded in the SR-specific extension header of the outer IPv6 packet header. This activity is referred to as segment tracing. Segment tracing can be performed independently from steering traffic using SR techniques. It can also be used in a transit network which performs normal IPv6 routing. When tracing is enabled, the segment ID of every segment that a packet traverses is recorded in the SR-specific extension header. This means that the egress tunnel router receives information about the path, represented by the segment list, a particular packet has taken in the underlay network. Different from OAM mechanisms which send active probe packets, tracing information can be made available for production traffic. The egress tunnel router can choose to provide the traced segment list back to the mapping system, for example through a LISP Map-Register. This information can be used to ensure path symmetry send/receive traffic in the transit network, or can serve other OAM or statistical purposes.

4. LISP extensions to support SR


Segment routing information can be contained within the LISP mapping system. A segment identifier is a 32-bit identification either for a topological instruction or a service instruction. See [I-D.filsfils-rtgwg-segment-routing] for details.

An EID can be associated with one or multiple ordered lists of segment identifiers, also referred to as "segment lists", encoding the topological and service source route of a packet. The segment list can serve either traffic engineering or operational purposes. In case of traffic engineering purposes, the segment list describes the set of segments a packet visits when traversing the transit network. The segment list enables the ITR to steer traffic using segment routing techniques. For operations and maintenance use, the segment list documents the set of segments a packet visited on its way through the transit network. It is beyond the scope of this document to describe the detailed procedures how the LISP mapping system obtains a segment list for a particular EID-to-RLLOC mapping.

Segment routing extensions for LISP extend the Explicit Locator Path

(ELP) Canonical Address Format, which is LCAF type 10, see[I-D.ietf-lisp-lcaf] for details. A new Address Family Identifier (AFI) in LISP Canonical Address Format (LCAF) type [I-D.ietf-lisp-lcaf] is required to carry the 32-bit segment identifier. For a given EID lookup in the mapping database, the segment routing list in ELP LCAF type can be returned to provide a segment list to each locator in the Map-Reply locator set. The ELP LCAF type can also be used to send the segment list that a particular packet traversed to the LISP mapping system using a Map-Register message defined in [RFC6833].

The segment identification AFI to be allocated is described below:



AFI=TBD_SID: TBD_SID is a value allocated from [AFI] for segment identifiers.

Rsvd: should be set to zero and ignored.

SID: 4 byte segment identifier

The explicit path to be followed in the underlay is then encoded using the AFI for segment ID in the LISP LCAF type 10 described in [I-D.ietf-lisp-lcaf].

T-Bit: An additional bit in the Rsvd3 field is to be allocated in LCAF type 10. The T-bit (T=1) is used by the LISP mapping system to indicate to an ITR that for particular EID-to-RLOC mapping the segments traversed by packets SHOULD be recorded as a segment list in the SR IPv6 extension header. This bit is ignored if present in a Map-Register message. A Map-Register message could be used by the ETR to inform the mapping system about the segments that a packet visited in the transit network.

S-Bit: The S-bit SHOULD be set when AFI = TBD_SID.

P-Bit: The P-bit SHOULD be ignored when AFI = TBD_SID.

L-Bit: The L-bit SHOULD be ignored when AFI = TBD SID.

4.1. Deployment Scenario

As described in [RFC6833] LISP Mapping Service defines: the Map-Resolver, which accepts Map-Requests from an Ingress Tunnel Router (ITR) and "resolves" the EID-to-RLOC mapping using a mapping database; and the Map-Server, which learns authoritative EID-to-RLOC mappings from an Egress Tunnel Router (ETR) and publishes them in a database. The LISP Extensions for Segment Routing described in this document primarily apply to deployment scenarios where MAP-Server and MAP-Resolver have visibility into or interface with a system that has knowledge of the network topology and can determine paths from source to destination RLOCs. Implementations of the LISP mapping systems which complement Software Defined Networking (SDN) architectures, such as the implementation as part of the OpenDaylight project[ODLLISP] fall into this category. In these deployments the LISP mapping system can retrieve the necessary information related to topology and path selection to implement the extensions defined in this document. It allows the mapping system to provide the required information in the MAP resolve response to correlate overlay with underlay network and offer solutions to control the path taken in the underlay network.

4.2. Example ELPs

4.2.1. Example: ELP with only SR used

This example shows the Explicit Locator Path (ELP) Canonical Address Format in a setup where segment routing is used in the transit network between ITR and ETR. Traffic engineering using reencapsulating routers is not used.

The reply to an EID-to-RLOC lookup contains the SIDs to be visited in the underlay network to reach the RLOC address returned in AFI=x. In the example below SID_1,...,SID_p are to be used for segment routing towards the "Address" RLOC. SID_p is the identifier of the last segment which takes the packet to the "Address" RLOC.

```


      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|               AFI = 16387               |   Rsvd1   |   Flags   |
+-----+-----+-----+-----+-----+-----+-----+-----+
|   Type = 10   |   Rsvd2   |               n               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               AFI = TBD_SID               |   Rsvd3   | T | L | P | S |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               SID_1               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               AFI = TBD_SID               |   Rsvd3   | T | L | P | S |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               SID_p               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               AFI = x               |   Rsvd3   | T | L | P | S |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               Address ...               |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

4.2.2. Example: ELP with SR and reencapsulating routers combined

This example shows the Explicit Locator Path (ELP) Canonical Address Format when using SR combined with reencapsulation routers.

Segment routing and traffic engineering using reencapsulating routers can be combined. In the example below, segment routing is used to steer traffic in the underlay between reencapsulating routers "f" and "g". There is no segment routing used between any of the other reencapsulating router hops.



5. IANA Considerations

TBD.

6. Manageability Considerations

Manageability considerations will be addressed in a later version of this document..

7. Security Considerations

Security considerations will be addressed in a later version of this document.

8. Acknowledgements

The authors would like to thank Dino Farinacci, Erik Nordmark and participants of the LISP wg for their input on this document.

9. Change log

Changes from 00 - 01

- o Added a section on deployment scenario to clarify the applicability of the extension described in this draft.

10. References

10.1. Normative References

- [AFI] "IANA, Address Family Identifier (AFIs), <http://www.iana.org/assignments/address-family-numbers/address-family-numbers.xhtml>", July 2013.
- [I-D.filsfils-rtgwg-segment-routing]
Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., Horneffer, M., Milojevic, I., Shakir, R., Ytti, S., Henderickx, W., Tantsura, J., and E. Crabbe, "Segment Routing Architecture", draft-filsfils-rtgwg-segment-routing-00 (work in progress), June 2013.
- [I-D.ietf-lisp-lcaf]
Farinacci, D., Meyer, D., and J. Snijders, "LISP Canonical Address Format (LCAF)", draft-ietf-lisp-lcaf-02 (work in progress), March 2013.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

10.2. Informative References

- [I-D.farinacci-lisp-te]
Farinacci, D., Lahiri, P., and M. Kowal, "LISP Traffic Engineering Use-Cases", draft-farinacci-lisp-te-03 (work in progress), July 2013.
- [I-D.filsfils-rtgwg-segment-routing-use-cases]
Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., Horneffer, M., Milojevic, I., Shakir, R.,

Ytti, S., Henderickx, W., Tantsura, J., and E. Crabbe,
"Segment Routing Use Cases",
draft-filsfils-rtgwg-segment-routing-use-cases-00 (work in
progress), June 2013.

[I-D.sivabalan-pce-segment-routing]
Sivabalan, S., Filsfils, C., Medved, J., Crabbe, E., and
R. Raszuk, "PCE-Initiated Traffic Engineering Path Setup
in Segment Routed Networks",
draft-sivabalan-pce-segment-routing-00 (work in progress),
June 2013.

[ODLLISP] "Open Day Light Lisp Flow Mapping, [https://
wiki.opendaylight.org/view/
OpenDaylight_Lisp_Flow_Mapping:Architecture](https://wiki.opendaylight.org/view/OpenDaylight_Lisp_Flow_Mapping:Architecture)", Feb 2014.

[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.

Authors' Addresses

Frank Brockners
Cisco
Hansaallee 249, 3rd Floor
DUESSELDORF, NORDRHEIN-WESTFALEN 40549
Germany

Email: fbrockne@cisco.com

Shwetha Bhandari
Cisco
Cessna Business Park, Sarjapura Marathalli Outer Ring Road
Bangalore, KARNATAKA 560 087
India

Email: shwethab@cisco.com

Fabio Maino
Cisco
San Jose
USA

Email: fmaino@cisco.com

Darrel Lewis
Cisco
San Jose
USA

Email: darlewis@cisco.com

