

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: March 20, 2016

M. Boucadair
C. Jacquenet
France Telecom
September 17, 2015

Mapping System-Assisted Forwarding for Inter-Domain LISP Deployments
draft-boucadair-lisp-ms-assisted-forwarding-00

Abstract

One of the issues with current LISP operation is that some packets are likely to be lost when there is no matching mapping entry maintained by the Ingress Tunnel Router (ITR). This document proposes a solution to address this issue with a particular focus on LISP deployments at large scale.

This document introduces the concept of Implicit Map-Request and Unsolicited Map-Reply messages. Also, it describes a solution to disable data gleaning for a given flow at the upstream Egress Tunnel Router (ETR).

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](#)].

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 March 20, 2016.

Internet-Draft

MS-Assisted Forwarding

September 2015

Copyright Notice

Copyright (c) 2015 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.	Proposed Solution	3
3.	Disable Data Gleaning	6
4.	Security Considerations	8
5.	IANA Considerations	8
6.	Acknowledgments	8
7.	References	8
7.1.	Normative references	8
7.2.	Informative references	8
	Authors' Addresses	9

[1.](#) Introduction

Locator/ID Separation Protocol (LISP, [[RFC6830](#)]) operation relies upon a mapping mechanism that is used by ingress/egress Tunnel Routers (xTR) to forward traffic over the LISP network. It is commonly acknowledged that some packets are likely to be lost in some specific situations, such as the absence of a mapping entry in the mapping cache maintained by an ITR. This phenomenon is usually denoted as the "first packet loss" issue.

This document suggests a solution that relies upon the assistance of the Mapping System for the forwarding of the first packet(s) of a flow, while corresponding mapping resolution is in progress.

Deploying LISP at the scale of the Internet heaviness relies upon the

reliability of the LISP Mapping service. In particular, LISP deployments at large scale must not degrade the level of quality as currently provided by several decades of inter-domain routing practices.

This document makes the following assumptions:

- o Various LISP players (network operators, service providers, etc.) are likely to deploy and operate LISP Mapping Systems. Multiple Mapping Systems will therefore coexist for various reasons, e.g., avoid country-centric governance, allow for distinct technologies to implement the mapping service, business opportunities, service innovation, etc.
- o Interconnection between these Mapping Systems is required for the sake of global connectivity and also to minimize the risk of fragmenting the Internet.
- o The entries of the mapping tables are exchanged between these Mapping systems so that Map-Request messages can be processed as close to the LISP leaf networks as possible.
- o A leaf LISP-enabled network subscribes to the Mapping Service provided by one or several Mapping Service operators.
- o The contribution of each player involved in the provisioning and the operation of a LISP-based connectivity forwarding service should be rationalized so that clear interfaces are defined and adequate mechanisms for troubleshooting, diagnosis and repair purposes can be easily implemented and adopted. The inability of identifying what is at the origin of the degradation of a LISP connectivity service is seen as one of the hurdles that are likely to jeopardize LISP deployments at large scale.
- o ITRs are configured with a list of one or more Map-Resolvers locators. Whether the provisioning of MR-related information to ITRs is achieved using a configuration interface or a specific discovery mechanism is out of scope of this document.
- o Like [[RFC6830](#)], this document does not make any assumption of the Mapping Service other than it supports the primitives defined in [[RFC6833](#)].

This document focuses on the sole LISP inter-domain use case. As such, it is out of scope of this document to discuss the applicability of the proposed solution to other LISP use cases (e.g., LISP-based data center networking) .

Within this document, "Unsolicited Map-Reply" is used to refer to a Map-Reply that is not associated with an (explicit) Map-Request message. An unsolicited Map-Reply is sent to an ITR without receiving a Map-Request from that ITR.

2. Proposed Solution

The rationale adopted by this document is that, instead of dropping packets that do not match an existing mapping entry in a local cache maintained by an ITR, these packets are used as Implicit Map-Requests. In particular, the LISP-based forwarding of the first

packet(s) can be delegated to the Mapping System (MS) that is supposed to maintain the required information to process the packet (preferably close to the LISP leaf network or at least without inducing severe path stretch). This mode is called MS-assisted LISP forwarding.

Although this feature can be supported by an upstream transit provider, this document focuses on the deployment of the MS-assisted LISP forwarding solution at the Mapping System side.

The detailed procedure that aims at minimizing the risk of the aforementioned "first packet loss" issue is specified hereafter (see Figure 1 for a typical flow example):

- 1 The Mapping System is configured with a list of networks that are allowed to invoke the MS-assisted forwarding service. The corresponding authorization procedure may rely upon the service subscription procedure itself (using static or dynamic means such as [[I-D.boucadair-connectivity-provisioning-protocol](#)]).
 - * Also, the Mapping System provides a leaf LISP network with the appropriate RLOC (referred to as MS_RLOC) so that it can use the MS-assisted forwarding feature.
 - * MS_RLOC may be identical or distinct from the locator assigned to one of the Map-Resolvers that can be solicited by the leaf LISP network.
- 2 ITRs MUST support a configuration parameter to enable/disable this procedure. The default value of this parameter is "Disabled".
- 3 ITRs MAY be configured with a dedicated RLOC for this feature.

This RLOC MAY NOT be the same locator as the one used to contact a Map-Resolver. If no dedicated RLOC is explicitly configured on an ITR for which the MS-assisted forwarding procedure is enabled, the ITR MUST use the locator of its Map-Resolver (i.e., MS_RLOC=ITR_Locator).

- 4 When an ITR receives a packet to be forwarded outside a given LISP domain, it MUST proceed to a lookup of the local mapping cache to check whether an entry matches this packet.
 - 4.1 If a mapping entry is found, the ITR MUST proceed as in [\[RFC6830\]](#).
 - 4.2 If no mapping entry is found and the MS-assisted forwarding feature is enabled, the ITR MUST use the MS_RLOC to forward the packet. That is, the origin packet is forwarded using a LISP encapsulation header; the destination IP address of the outer header is set to MS_RLOC (instead of the remote ETR's RLOC associated with the destination EID).

- 4.2.1 The ITR MUST set the nonce-present and echo-nonce-request bits.
 - 4.2.2 Once forwarded, the ITR MUST listen, using port 4343, to Unsolicited Map-Reply messages that will be received from the Map-Resolver.
 - 4.2.3 The ITR MUST follow the same behavior for packets that belong to the same flow until a mapping is retrieved from the Mapping System side. The packet will be used as an "implicit Map-Request" from a downstream ITR.
- 5 Upon receipt of the encapsulated packet, the Mapping System:
 - 5.1 MUST extract the destination EID and proceed to the lookup in its global mapping table to retrieve the corresponding entry. If a mapping entry is found, it MUST rewrite the source RLOC to be set to the destination RLOC of the encapsulated packet received from the leaf LISP network and the destination RLOC to the RLOC retrieved from the mapping table. Then, the packet is forwarded to the next hop.
 - 5.1.1 Using the initial source RLOC to forward the packet might be tempting, but this behavior is discouraged as upstream networks implementing ingress filtering may

- consider this packet as a spoofing attack.
- 5.1.2 The Mapping System may decide to reset the nonce-present and echo-nonce-request bits. The setting of these bits can be part of the service agreement contracted between the leaf LISP network and the Mapping Service provider.
 - 5.1.3 Because upstream ETRs may use the outer LISP header if it implemented information "gleaning", the LISP packet may provide an explicit indication to the ETR to not rely upon the outer header to create a "gleaned" Map-Cache entry but rather proceed with an explicit Map-Request. instead [Section 3](#) proposes a solution for carrying such indication in the LISP header.
- 5.2 In the meantime, an unsolicited Map-Reply message that carries records associated with the destination EID, MUST be sent to the ITR so that it can handle the packets locally without any assistance from the Mapping System.
- 5.2.1 The Map-Reply message MUST use the same Nonce that was included in the LISP-encapsulated packet received from the downstream ITR.
 - 5.2.2 A timer (or a maximum number of the packets) MAY be used so that the assistance mode is deactivated at the Mapping System side for this leaf LISP network/EID. Discarding subsequent packets and associated settings

- are deployment-specific. It is out of scope of this document to elaborate on such design considerations.
- 6 Upon receipt of the Unsolicited Map-Reply message, the ITR MUST proceed to Nonce validation checks.
 - 6.1 If no error is found, it MUST retrieve the record carried in the Map-Reply message.
 - 6.2 The ITR MUST stop using the MS-assisted mode (i.e., for forthcoming packets matching this mapping entry).
 - 7 Subsequent packets that belong to the same flow are handled locally (i.e., normal LISP operation is in progress).

```

+-----+
| Mapping|
+-----+
| System |
+-----+

```

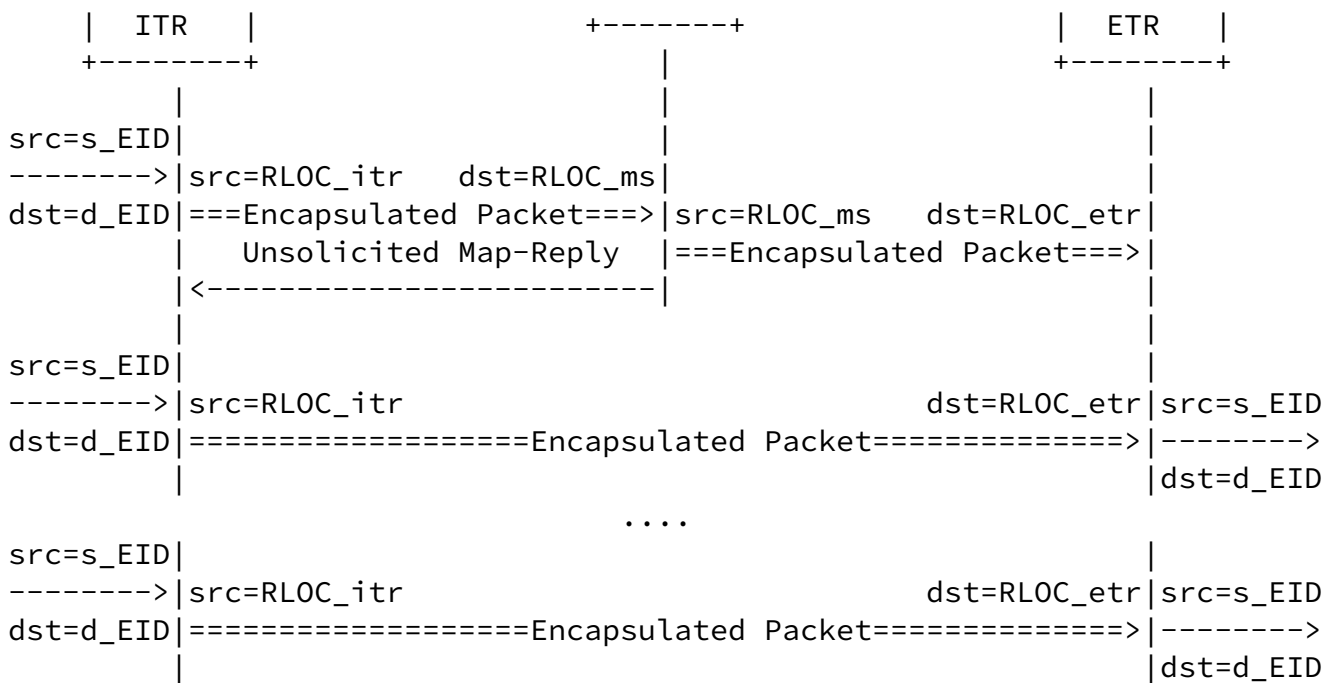


Figure 1: Flow Example

3. Disable Data Gleaning

[RFC6830] indicates the following:

[Section 4](#): "In order to defer the need for a mapping lookup in the reverse direction, an ETR MAY create a cache entry that maps the source EID (inner-header source IP address) to the source RLOC (outer-header source IP address) in a received LISP packet. Such a cache entry is termed a "gleaned" mapping and only contains a single RLOC for the EID in question."

[Section 6.2](#): "Either side (more likely the server-side ETR) decides not to send a Map-Request. For example, if the server-side ETR does not send Map-Requests, it gleans RLOCs from the client-side ITR, giving the client-side ITR responsibility for bidirectional RLOC reachability and preferability. Server-side ETR gleaning of the client-side ITR RLOC is done by caching the inner-header source EID and the outer-header source RLOC of received packets. The client-side ITR controls how traffic is

returned and can alternate using an outer- header source RLOC, which then can be added to the list the server-side ETR uses to return traffic. Since no Priority or Weights are provided using this method, the server-side ETR MUST assume that each client-side ITR RLOC uses the same best Priority with a Weight of zero. In addition, since EID-Prefix encoding cannot be conveyed in data packets, the EID-to-RLOC Cache on Tunnel Routers can grow to be very large."

But the LISP specification does not describe any means for an ITR to explicitly inform an ETR that it MUST NOT rely upon the data gleaning but, instead, privilege the sending of an explicit Map-request.

For the particular case covered in this document, the lack of such capability may lead to the involvement of an intermediate node for both traffic directions. This behavior may not be suitable in some deployment situations (e.g., mis-use the relay in the MS domain to forward traffic, abuse, denial-of-service, etc.). In order to solve this issue, this document proposes to associate a meaning with one of the reserved flag bits (see [Section 5.3 of \[RFC6830\]](#)) to explicitly indicate that, when this bit is set, data gleaning must be deactivated. This bit is called the G-bit ("Gleaning" flag bit).

Figure 2 shows the required change to the LISP header.

OLD:

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
L  |N|L|E|V|I|flags|                               Nonce/Map-Version          |
I  \ +-----+-----+-----+-----+-----+-----+-----+-----+-----+
S  / |                               Instance ID/Locator-Status-Bits          |
P  +-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

NEW:

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
L  |N|L|E|V|I|G|flg|                               Nonce/Map-Version          |
I  \ +-----+-----+-----+-----+-----+-----+-----+-----+-----+
S  / |                               Instance ID/Locator-Status-Bits          |
P  +-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 2: G-bit in the LISP Header

The "flg" bits are reserved bits for future assignment as additional

flag bits. These additional flag bits MUST each be set to zero and MUST be ignored upon receipt.

The description of the remaining fields is the same as in [[RFC6830](#)].

[4.](#) Security Considerations

Security considerations discussed in [[RFC6833](#)] and [[RFC6830](#)] should be taken into account.

[5.](#) IANA Considerations

This document does not make any request to IANA.

[6.](#) Acknowledgments

This work is partly funded by ANR LISP-Lab project #ANR-13-INFR-009-X.

[7.](#) References

[7.1.](#) Normative references

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

[RFC6833] Fuller, V. and D. Farinacci, "Locator/ID Separation Protocol (LISP) Map-Server Interface", [RFC 6833](#), DOI 10.17487/RFC6833, January 2013, <<http://www.rfc-editor.org/info/rfc6833>>.

[7.2.](#) Informative references

[I-D.boucadair-connectivity-provisioning-protocol]
Boucadair, M., Jacquenet, C., Zhang, D., and P. Georgatsos, "Connectivity Provisioning Negotiation Protocol (CPNP)", [draft-boucadair-connectivity-provisioning-protocol-10](#) (work in progress), September 2015.

[RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", [RFC 6830](#), DOI 10.17487/RFC6830, January 2013, <<http://www.rfc-editor.org/info/rfc6830>>.

Internet-Draft

MS-Assisted Forwarding

September 2015

Authors' Addresses

Mohamed Boucadair
France Telecom
Rennes 35000
France

EMail: mohamed.boucadair@orange.com

Christian Jacquenet
France Telecom
Rennes 35000
France

EMail: christian.jacquenet@orange.com

