

Network Working Group
Internet-Draft
Obsoletes: [6834](#) (if approved)
Intended status: Standards Track
Expires: August 19, 2019

L. Iannone
Telecom ParisTech
D. Saucez
INRIA Sophia Antipolis
O. Bonaventure
Universite catholique de Louvain
February 15, 2019

Locator/ID Separation Protocol (LISP) Map-Versioning
draft-ietf-lisp-6834bis-03

Abstract

This document describes the LISP (Locator/ID Separation Protocol) Map-Versioning mechanism, which provides in-packet information about Endpoint ID to Routing Locator (EID-to-RLOC) mappings used to encapsulate LISP data packets. The proposed approach is based on associating a version number to EID-to-RLOC mappings and the transport of such a version number in the LISP-specific header of LISP-encapsulated packets. LISP Map-Versioning is particularly useful to inform communicating Ingress Tunnel Routers (ITRs) and Egress Tunnel Routers (ETRs) about modifications of the mappings used to encapsulate packets. The mechanism is optional and transparent to implementations not supporting this feature, since in the LISP-specific header and in the Map Records, bits used for Map-Versioning can be safely ignored by ITRs and ETRs that do not support or do not want to use the mechanism.

This document obsoletes [RFC 6834](#) "Locator/ID Separation Protocol (LISP) Map-Versioning", which is the initial experimental specifications of the mechanisms updated by this document.

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 <https://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 19, 2019.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://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. | Requirements Notation | 4 |
| 3. | Definitions of Terms | 4 |
| 4. | EID-to-RLOC Map-Version Number | 4 |
| 4.1. | The Null Map-Version | 5 |
| 5. | Dealing with Map-Version Numbers | 6 |
| 5.1. | Handling Destination Map-Version Number | 7 |
| 5.2. | Handling Source Map-Version Number | 9 |
| 6. | LISP Header and Map-Version Numbers | 10 |
| 7. | LISP Generic Protocol Encapsulation (GPE) Header and Map-Version Numbers | 10 |
| 8. | Map Record and Map-Version | 11 |
| 9. | Benefits and Case Studies for Map-Versioning | 12 |
| 9.1. | Map-Versioning and Unidirectional Traffic | 12 |
| 9.2. | Map-Versioning and Interworking | 13 |
| 9.2.1. | Map-Versioning and Proxy-ITRs | 13 |
| 9.2.2. | Map-Versioning and LISP-NAT | 14 |
| 9.2.3. | Map-Versioning and Proxy-ETRs | 14 |
| 9.3. | RLOC Shutdown/Withdraw | 14 |
| 9.4. | Map-Version Additional Use Cases | 15 |
| 10. | Security Considerations | 15 |
| 10.1. | Map-Versioning against Traffic Disruption | 15 |
| 10.2. | Map-Versioning against Reachability Information DoS | 16 |
| 11. | Deployment Considerations | 16 |
| 12. | IANA Considerations | 18 |
| 13. | Acknowledgments | 18 |
| 14. | References | 18 |
| 14.1. | Normative References | 18 |
| 14.2. | Informative References | 19 |

| | |
|------------------------------|--------------------|
| Authors' Addresses | 19 |
|------------------------------|--------------------|

1. Introduction

This document describes the Map-Versioning mechanism used to provide information on changes in the EID-to-RLOC (Endpoint ID to Routing Locator) mappings used in the LISP (Locator/ID Separation Protocol [[I-D.ietf-lisp-rfc6830bis](#)][[I-D.ietf-lisp-rfc6833bis](#)]) context to perform packet encapsulation. The mechanism is totally transparent to xTRs (Ingress and Egress Tunnel Routers) not supporting or not using such functionality.

This document obsoletes [[RFC6834](#)], which is the initial experimental specifications of the mechanisms updated by this document.

The basic mechanism is to associate a Map-Version number to each LISP EID-to-RLOC mapping and transport such a version number in the LISP-specific header. When a mapping changes, a new version number is assigned to the updated mapping. A change in an EID-to-RLOC mapping can be a change in the RLOCs set, by adding or removing one or more RLOCs, but it can also be a change in the priority or weight of one or more RLOCs.

When Map-Versioning is used, LISP-encapsulated data packets contain the version number of the two mappings used to select the RLOCs in the outer header (i.e., both source and destination). These version numbers are encoded in the 24 low-order bits of the first longword of the LISP header and indicated by a specific bit in the flags (first 8 high-order bits of the first longword of the LISP header). Note that not all packets need to carry version numbers.

When an ITR (Ingress Tunnel Router) encapsulates a data packet, with a LISP header containing the Map-Version numbers, it puts in the LISP-specific header two version numbers:

1. The version number assigned to the mapping (contained in the EID-to-RLOC Database) used to select the source RLOC.
2. The version number assigned to the mapping (contained in the EID-to-RLOC Cache) used to select the destination RLOC.

This operation is two-fold. On the one hand, it enables the ETR (Egress Tunnel Router) receiving the packet to know if the ITR is using the latest mapping version that any ETR at the destination EID site would provide to the ITR in a Map-Reply. If this is not the case, the ETR can send to the ITR a Map-Request containing the updated mapping or solicit a Map-Request from the ITR (both cases are already defined in [[I-D.ietf-lisp-rfc6833bis](#)]). In this way, the ITR

can update its EID-to-RLOC Cache. On the other hand, it enables an ETR receiving such a packet to know if it has in its EID-to-RLOC Cache the latest mapping for the source EID. If this is not the case, a Map-Request can be sent.

Considerations about the deployment of LISP Map-Versioning are discussed in [Section 11](#).

2. Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

3. Definitions of Terms

This document uses terms already defined in the main LISP specification ([I-D.ietf-lisp-rfc6830bis](#) [[I-D.ietf-lisp-rfc6833bis](#)]). Here, we define the terms that are specific to the Map-Versioning mechanism. Throughout the whole document, Big Endian bit ordering is used.

Map-Version number: An unsigned 12-bit integer is assigned to an EID-to-RLOC mapping, not including the value 0 (0x000).

Null Map-Version: The 12-bit null value of 0 (0x000) is not used as a Map-Version number. It is used to signal that no Map-Version number is assigned to the EID-to-RLOC mapping.

Source Map-Version number: This Map-Version number of the EID-to-RLOC mapping is used to select the source address (RLOC) of the outer IP header of LISP-encapsulated packets.

Destination Map-Version number: This Map-Version number of the EID-to-RLOC mapping is used to select the destination address (RLOC) of the outer IP header of LISP-encapsulated packets.

4. EID-to-RLOC Map-Version Number

The EID-to-RLOC Map-Version number consists of an unsigned 12-bit integer. The version number is assigned on a per-mapping basis, meaning that different mappings have a different version number, which is also updated independently. An update in the version number (i.e., a newer version) consists of incrementing by one the older version number.

The space of version numbers has a circular order where half of the version numbers are greater (i.e., newer) than the current Map-Version number and the other half of the version numbers are smaller (i.e., older) than the current Map-Version number. In a more formal way, assuming that we have two version numbers V1 and V2 and that the numbers are expressed on N bits, the following steps MUST be performed (in the same order as shown below) to strictly define their order:

1. $V1 = V2$: The Map-Version numbers are the same.

2. $V2 > V1$: if and only if

$$V2 > V1 \text{ AND } (V2 - V1) \leq 2^{N-1}$$

OR

$$V1 > V2 \text{ AND } (V1 - V2) > 2^{N-1}$$

3. $V1 > V2$: otherwise.

Using 12 bits, as defined in this document, and assuming a Map-Version value of 69, Map-Version numbers in the range [70; 69 + 2048] are greater than 69, while Map-Version numbers in the range [69 + 2049; (69 + 4096) mod 4096] are smaller than 69.

Map-version numbers are assigned to mappings by configuration. The initial Map-Version number of a new EID-to-RLOC mapping SHOULD be assigned randomly, but it MUST NOT be set to the Null Map-Version value (0x000), because the Null Map-Version number has a special meaning (see [Section 4.1](#)).

Upon reboot, an ETR will use mappings configured in its EID-to-RLOC Database. If those mappings have a Map-Version number, it will be used according to the mechanisms described in this document. ETRs MUST NOT automatically generate and assign Map-Version numbers to mappings in the EID-to-RLOC Database.

4.1. The Null Map-Version

The value 0x000 (zero) is not a valid Map-Version number indicating the version of the EID-to-RLOC mapping. Such a value is used for special purposes and is named the Null Map-Version number.

The Null Map-Version MAY appear in the LISP-specific header as either a Source Map-Version number (cf. [Section 5.2](#)) or a Destination Map-Version number (cf. [Section 5.1](#)). When the Source Map-Version number is set to the Null Map-Version value, it means that no map

version information is conveyed for the source site. This means that if a mapping exists for the source EID in the EID-to-RLOC Cache, then the ETR MUST NOT compare the received Null Map-Version with the content of the EID-to-RLOC Cache. When the Destination Map-Version number is set to the Null Map-Version value, it means that no map version information is conveyed for the destination site. This means that the ETR MUST NOT compare the value with the Map-Version number of the mapping for the destination EID present in the EID-to-RLOC Database.

The other use of the Null Map-Version number is in the Map Records, which are part of the Map-Request, Map-Reply, and Map-Register messages (defined in [[I-D.ietf-lisp-rfc6833bis](#)]). Map Records that have a Null Map-Version number indicate that there is no Map-Version number associated with the mapping. This means that LISP-encapsulated packets destined to the EID-Prefix referred to by the Map Record MUST either not contain any Map-Version numbers (V bit set to 0) or, if they contain Map-Version numbers (V bit set to 1), then the destination Map-Version number MUST be set to the Null Map-Version number. Any value different from zero means that Map-Versioning is supported and MAY be used.

The fact that the 0 value has a special meaning for the Map-Version number implies that, when updating a Map-Version number because of a change in the mapping, if the next value is 0, then the Map-Version number MUST be incremented by 2 (i.e., set to 1, which is the next valid value).

5. Dealing with Map-Version Numbers

The main idea of using Map-Version numbers is that whenever there is a change in the mapping (e.g., adding/removing RLOCs, a change in the weights due to Traffic Engineering policies, or a change in the priorities) or a LISP site realizes that one or more of its own RLOCs are not reachable anymore from a local perspective (e.g., through IGP, or policy changes) the LISP site updates the mapping, also assigning a new Map-Version number.

To each mapping, a version number is associated and changes each time the mapping is changed. Note that Map-Versioning does not introduce new problems concerning the coordination of different ETRs of a domain. Indeed, ETRs belonging to the same LISP site must return for a specific EID-prefix the same mapping, including the same Map-Version number. This is orthogonal to whether or not Map-Versioning is used. The synchronization problem and its implications on the traffic are out of the scope of this document.

In order to announce in a data-driven fashion that the mapping has been updated, Map-Version numbers used to create the outer IP header of the LISP-encapsulated packet are embedded in the LISP-specific header. This means that the header needs to contain two Map-Version numbers:

- o The Source Map-Version number of the EID-to-RLOC mapping in the EID-to-RLOC Database used to select the source RLOC.
- o The Destination Map-Version number of the EID-to-RLOC mapping in the EID-to-RLOC Cache used to select the destination RLOC.

By embedding both the Source Map-Version number and the Destination Map-Version number, an ETR receiving a LISP packet with Map-Version numbers can perform the following checks:

1. The ITR that has sent the packet has an up-to-date mapping in its EID-to-RLOC Cache for the destination EID and is performing encapsulation correctly.
2. In the case of bidirectional traffic, the mapping in the local ETR EID-to-RLOC Cache for the source EID is up to date.

If one or both of the above conditions do not hold, the ETR can send a Map-Request either to make the ITR aware that a new mapping is available (see [Section 5.1](#)) or to update the mapping in the local EID-to-RLOC Cache (see [Section 5.2](#)).

[5.1](#). Handling Destination Map-Version Number

When an ETR receives a packet, the Destination Map-Version number relates to the mapping for the destination EID for which the ETR is an RLOC. This mapping is part of the ETR EID-to-RLOC Database. Since the ETR is authoritative for the mapping, it has the correct and up-to-date Destination Map-Version number. A check on this version number can be done, where the following cases can arise:

1. The packet arrives with the same Destination Map-Version number stored in the EID-to-RLOC Database. This is the regular case. The ITR sending the packet has in its EID-to-RLOC Cache an up-to-date mapping. No further actions are needed.
2. The packet arrives with a Destination Map-Version number greater (i.e., newer) than the one stored in the EID-to-RLOC Database. Since the ETR is authoritative on the mapping, meaning that the Map-Version number of its mapping is the correct one, this implies that someone is not behaving correctly with respect to the specifications. In this case, the packet carries a version

number that is not valid; otherwise, the ETR would have the same number, and the packet SHOULD be silently dropped.

3. The packets arrive with a Destination Map-Version number smaller (i.e., older) than the one stored in the EID-to-RLOC Database. This means that the ITR sending the packet has an old mapping in its EID-to-RLOC Cache containing stale information. The ETR MAY choose to normally process the encapsulated datagram according to [\[I-D.ietf-lisp-rfc6830bis\]](#); however, the ITR sending the packet has to be informed that a newer mapping is available. This is done with a Map-Request message sent back to the ITR. The Map-Request will either trigger a Map-Request back using the Solicit-Map-Request (SMR) bit or it will piggyback the newer mapping. These are not new mechanisms; how to use the SMR bit or how to piggyback mappings in Map-Request messages is already described in [\[I-D.ietf-lisp-rfc6833bis\]](#). One feature introduced by Map-Version numbers is the possibility of blocking traffic not using the latest mapping. Indeed, after a certain number of retries, if the Destination Map-Version number in the packets is not updated, the ETR MAY drop packets with a stale Map-Version number while strongly reducing the rate of Map-Request messages. This is because either the ITR is refusing to use the mapping for which the ETR is authoritative, or (worse) it might be some form of attack.

The rule in the third case MAY be more restrictive. If the mapping has been the same for a period of time as long as the Time To Live (TTL) (defined in [\[I-D.ietf-lisp-rfc6833bis\]](#)) of the previous version of the mapping, all packets arriving with an old Map-Version SHOULD be silently dropped right away without issuing any Map-Request. Such action is permitted because if the new mapping with the updated version number has been unchanged for at least the same time as the TTL of the older mapping, all the entries in the EID-to-RLOC Caches of ITRs must have expired. Hence, all ITRs sending traffic should have refreshed the mapping according to [\[I-D.ietf-lisp-rfc6833bis\]](#). If packets with old Map-Version numbers are still received, then either someone has not respected the TTL or it is a form of spoof/attack. In both cases, this is not valid behavior with respect to the specifications and the packet SHOULD be silently dropped.

LISP-encapsulated packets with the V-bit set, when the original mapping in the EID-to-RLOC Database has the version number set to the Null Map-Version value, MAY be silently dropped. As explained in [Section 4.1](#), if an EID-to-RLOC mapping has a Null Map-Version, it means that ITRs, using the mapping for encapsulation, MUST NOT use a Map-Version number in the LISP-specific header.

For LISP-encapsulated packets with the V-bit set, when the original mapping in the EID-to-RLOC Database has the version number set to a value different from the Null Map-Version value, a Destination Map-Version number equal to the Null Map-Version value means that the Destination Map-Version number MUST be ignored.

5.2. Handling Source Map-Version Number

When an ETR receives a packet, the Source Map-Version number relates to the mapping for the source EID for which the ITR that sent the packet is authoritative. If the ETR has an entry in its EID-to-RLOC Cache for the source EID, then a check can be performed and the following cases can arise:

1. The packet arrives with the same Source Map-Version number as that stored in the EID-to-RLOC Cache. This is the correct regular case. The ITR has in its EID-to-RLOC Cache an up-to-date copy of the mapping. No further actions are needed.
2. The packet arrives with a Source Map-Version number greater (i.e., newer) than the one stored in the local EID-to-RLOC Cache. This means that the ETR has in its EID-to-RLOC Cache a mapping that is stale and needs to be updated. A Map-Request SHOULD be sent to get the new mapping for the source EID. This is a normal Map-Request message sent through the mapping system and MUST respect the specifications in [[I-D.ietf-lisp-rfc6833bis](#)], including rate-limitation policies.
3. The packet arrives with a Source Map-Version number smaller (i.e., older) than the one stored in the local EID-to-RLOC Cache. Such a case is not valid with respect to the specifications. Indeed, if the mapping is already present in the EID-to-RLOC Cache, this means that an explicit Map-Request has been sent and a Map-Reply has been received from an authoritative source. Assuming that the mapping system is not corrupted, the Map-Version in the EID-to-RLOC Cache is the correct one, while the one carried by the packet is stale. In this situation, the packet MAY be silently dropped.

If the ETR does not have an entry in the EID-to-RLOC Cache for the source EID, then the Source Map-Version number can be ignored.

For LISP-encapsulated packets with the V-bit set, if the Source Map-Version number is the Null Map-Version value, it means that the Source Map-Version number MUST be ignored.

6. LISP Header and Map-Version Numbers

In order for the versioning approach to work, the LISP-specific header has to carry both the Source Map-Version number and Destination Map-Version number. This is done by setting the V-bit in the LISP-specific header as defined in [I-D.ietf-lisp-rfc6830bis]. When the V-bit is set and the P bit is reset (0), the low-order 24 bits of the first longword are used to transport both the source and destination Map-Version numbers. In particular, the first 12 bits are used for the Source Map-Version number and the second 12 bits for the Destination Map-Version number.

Below is an example of a LISP header carrying version numbers.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
/ |N|L|E|V|I|P|K|K|  Source Map-Version   |Destination Map-Version|
LISP+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
\ |                               Instance ID/Locator Status Bits   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Source Map-Version number (12 bits): Map-Version of the mapping used by the ITR to select the RLOC present in the "Source Routing Locator" field. [Section 5.2](#) describes how to set this value on transmission and handle it on reception.

Destination Map-Version number (12 bits): Map-Version of the mapping used by the ITR to select the RLOC present in the "Destination Routing Locator" field. [Section 5.1](#) describes how to set this value on transmission and handle it on reception.

Not all of the LISP-encapsulated packets need to carry version numbers. When Map-Version numbers are carried in these packets, the V-bit MUST be set to 1. All permissible combinations of the flags when the V-bit is set to 1 are described in [I-D.ietf-lisp-rfc6830bis] and [I-D.ietf-lisp-gpe].

7. LISP Generic Protocol Encapsulation (GPE) Header and Map-Version Numbers

[I-D.ietf-lisp-gpe] extends the Locator/ID Separation Protocol (LISP) Data-Plane, changing the LISP header, to support multi-protocol encapsulation. A flag in the LISP header, called the P-bit, is used to signal the presence of the Next Protocol field in the low-order 8 bits of the first longword. When the V-bit and P-bit are both set, the middle-order 16 bits of the first longword are used to transport both the source and destination Map-Version numbers. In particular,

the first 8 bits are used for the Source Map-Version number and the second 8 bits for the Destination Map-Version number.

Below is an example of a LISP header carrying version numbers.

```

      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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
/ |N|L|E|V|I|P|K|K|Src Map-Version|Dst Map-Version| Next Protocol |
LISP+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
\ |                               Instance ID/Locator Status Bits      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

The Source Map-Version number and the Destination Map-Version number are used exactly in the same way previously described. There are only three differences:

- o When filling the LISP-GPE-specific header only the low order 8 bits are copied in the Source and Destination Map-Version Number (out of the original 12 bits).
- o When comparing a Map-Version retrieved from the LISP-GPE-specific header (either Source or Destination Map-Version number) with the version number of a mapping (stored in the LISP Cache or LISP Database) only the low-order 8 bits of the latter are used for the comparison.
- o When trimming a Map-Version number from 12 to 8 bits it may happen that it is converted to a Null Map-Version number, which will change the way Map-Version number is interpreted as described in [Section 4.1](#). To avoid such wrong behavior, Map-Version number with the low-order 8 bits all equal to zero SHOULD be avoided on xTRs using LISP-GPE.

8. Map Record and Map-Version

To accommodate the proposed mechanism, the Map Records that are transported in Map-Request/Map-Reply/Map-Register messages need to carry the Map-Version number as well. For this purpose, the 12 bits before the EID-AFI field in the Record that describes a mapping are used (see [[I-D.ietf-lisp-rfc6833bis](#)] and reported here as an example.


```

      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
+-> +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  |                                     Record TTL                             |
|  +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
R  | Locator Count | EID mask-len | ACT |A|      Reserved      |
e  +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
c  | Rsvd  |  Map-Version Number  |      EID-Prefix-AFI      |
o  +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
r  |                                     EID-Prefix                             |
d  +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  /|      Priority  |      Weight      | M Priority  |  M Weight  |
|  L +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  o |      Unused Flags      |L|p|R|      Loc-AFI      |
|  c +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  \|                                     Locator                             |
+-> +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Map-Version Number: Map-Version of the mapping contained in the Record. As explained in [Section 4.1](#), this field can be zero (0), meaning that no Map-Version is associated to the mapping; hence, packets that are LISP encapsulated using this mapping MUST NOT contain Map-Version numbers in the LISP-specific header, and the V-bit MUST be set to 0.

This packet format works perfectly with xTRs that do not support Map-Versioning, since they can simply ignore those bits.

9. Benefits and Case Studies for Map-Versioning

In the following sections, we provide more discussion on various aspects and uses of Map-Versioning. Security observations are grouped in [Section 10](#).

9.1. Map-Versioning and Unidirectional Traffic

When using Map-Versioning, the LISP-specific header carries two Map-Version numbers, for both source and destination mappings. This can raise the question on what will happen in the case of unidirectional flows, for instance, in the case presented in Figure 1, since the LISP specification does not mandate that the ETR have a mapping for the source EID.




Figure 1: Unidirectional traffic between LISP domains.

In the case of the ITR, the ITR is able to put both the source and destination version number in the LISP header, since the Source Map-Version number is in the ITR's database, while the Destination Map-Version number is in the ITR's cache.

In the case of the ETR, the ETR simply checks only the Destination Map-Version number in the same way as that described in [Section 5](#), ignoring the Source Map-Version number.

9.2. Map-Versioning and Interworking

Map-Versioning is compatible with the LISP interworking between LISP and non-LISP sites as defined in [\[RFC6832\]](#). LISP interworking defines three techniques to make LISP sites and non-LISP sites, namely Proxy-ITR, LISP-NAT, and Proxy-ETR. The following text describes how Map-Versioning relates to these three mechanisms.

9.2.1. Map-Versioning and Proxy-ITRs

The purpose of the Proxy-ITR (PITR) is to encapsulate traffic originating in a non-LISP site in order to deliver the packet to one of the ETRs of the LISP site (cf. Figure 2). This case is very similar to the unidirectional traffic case described in [Section 9.1](#); hence, similar rules apply.




Figure 2: Unidirectional traffic from non-LISP domain to LISP domain.

The main difference is that a Proxy-ITR does not have any mapping, since it just encapsulates packets arriving from the non-LISP site,

and thus cannot provide a Source Map-Version. In this case, the proxy-ITR will just put the Null Map-Version value as the Source Map-Version number, while the receiving ETR will ignore the field.

With this setup, LISP Domain A is able to check whether or not the PITR is using the latest mapping.

9.2.2. Map-Versioning and LISP-NAT

The LISP-NAT mechanism is based on address translation from non-routable EIDs to routable EIDs and does not involve any form of encapsulation. As such, Map-Versioning does not apply in this case.

9.2.3. Map-Versioning and Proxy-ETRs

The purpose of the Proxy-ETR (PETR) is to decapsulate traffic originating in a LISP site in order to deliver the packet to the non-LISP site (cf. Figure 3). One of the main reasons to deploy PETRs is to bypass uRPF (Unicast Reverse Path Forwarding) checks on the provider edge.




Figure 3: Unidirectional traffic from LISP domain to non-LISP domain.

A Proxy-ETR does not have any mapping, since it just decapsulates packets arriving from the LISP site. In this case, the ITR will just put the Null Map-Version value as the Destination Map-Version number, while the receiving Proxy-ETR will ignore the field.

With this setup, the Proxy-ETR is able to check whether or not the mapping has changed.

9.3. RLOC Shutdown/Withdraw

Map-Versioning can also be used to perform a graceful shutdown or withdraw of a specific RLOC. This is achieved by simply issuing a new mapping, with an updated Map-Version number where the specific RLOC to be shut down is withdrawn or announced as unreachable (via the R bit in the Map Record; see [[I-D.ietf-lisp-rfc6833bis](#)]), but without actually turning it off.

Once no more traffic is received by the RLOC, it can be shut down gracefully, because all sites actively using the mapping have updated it.

9.4. Map-Version Additional Use Cases

The use of Map-Versioning can help in developing a lightweight implementation of LISP. However, this comes with the price of not supporting the Loc-Status-Bit, which may be useful in some contexts.

In the current LISP specifications, the set of RLOCs must always be maintained ordered and consistent with the content of the Loc Status Bits ([[I-D.ietf-lisp-rfc6830bis](#)]). With Map-Versioning, such types of mechanisms can be avoided. When a new RLOC is added to a mapping, it is not necessary to "append" new locators to the existing ones as explained in [[I-D.ietf-lisp-rfc6830bis](#)]. A new mapping with a new Map-Version number will be issued, and since the old locators are still valid, the transition will occur with no disruptions. The same applies for the case where an RLOC is withdrawn. There is no need to maintain holes in the list of locators, as is the case when using Locator Status Bits, for sites that are not using the RLOC that has been withdrawn; in this case, the transition will occur with no disruptions.

All of these operations, as already stated, do not need to maintain any consistency among Locator Status Bits and in the way that the RLOCs are stored in the EID-to-RLOC Cache.

10. Security Considerations

Map-Versioning does not introduce any security issues concerning both the data plane and the control plane. On the contrary, as described below, if Map-Versioning may also be used to update mappings in the case of change in the reachability information (i.e., instead of the Locator Status Bits), it is possible to reduce the effects of some DoS or spoofing attacks that can happen in an untrusted environment.

Robustness of the Map-Versioning mechanism leverages on a trusted Mapping Distribution System. A thorough security analysis of LISP is documented in [[RFC7835](#)].

10.1. Map-Versioning against Traffic Disruption

An attacker can try to disrupt ongoing communications by creating LISP-encapsulated packets with wrong Locator Status Bits. If the xTR blindly trusts the Locator Status Bits, it will change the encapsulation accordingly, which can result in traffic disruption.

This does not happen in the case of Map-Versioning. As described in [Section 5](#), upon a version number change the xTR first issues a Map-Request. The assumption is that the mapping distribution system is sufficiently secure that Map-Request and Map-Reply messages and their content can be trusted. Security issues concerning specific mapping distribution systems are out of the scope of this document. In the case of Map-Versioning, the attacker should "guess" a valid version number that triggers a Map-Request as described in [Section 5](#); otherwise, the packet is simply dropped. Nevertheless, guessing a version number that generates a Map-Request is easy; hence, it is important to follow the rate-limitation policies described in [\[I-D.ietf-lisp-rfc6833bis\]](#) in order to avoid DoS attacks.

Note that a similar level of security can be obtained with Loc Status Bits by simply making it mandatory to verify any change through a Map-Request. However, in this case Locator Status Bits lose their meaning, because it does not matter anymore which specific bits have changed; the xTR will query the mapping system and trust the content of the received Map-Reply. Furthermore, there is no way to perform filtering as in Map-Versioning in order to drop packets that do not carry a valid Map-Version number. In the case of Locator Status Bits, any random change can trigger a Map-Request (unless rate limitation is enabled, which raises another type of attack as discussed in [Section 10.2](#)).

[10.2](#). Map-Versioning against Reachability Information DoS

Attackers can try to trigger a large number of Map-Requests by simply forging packets with random Map-Versions or random Locator Status Bits. In both cases, the Map-Requests are rate-limited as described in [\[I-D.ietf-lisp-rfc6833bis\]](#). However, in contrast to the Locator Status Bit, where there is no filtering possible, in the case of Map-Versioning it is possible to filter invalid version numbers before triggering a Map-Request, thus helping to reduce the effects of DoS attacks. In other words, the use of Map-Versioning enables a fine control on when to update a mapping or when to notify someone that a mapping has been updated.

It is clear that Map-Versioning does not protect against DoS and DDoS attacks, where an xTR loses processing power when doing checks on the LISP header of packets sent by attackers. This is independent of Map-Versioning and is the same for Loc Status Bits.

[11](#). Deployment Considerations

Even without Map-Versioning, LISP requires ETRs to announce the same mapping for the same EID-Prefix to a requester. Map-Versioning does not require additional synchronization mechanisms as compared to the

normal functioning of LISP without Map-Versioning. Clearly, all the ETRs have to reply with the same Map-Version number; otherwise, there can be an inconsistency that creates additional control traffic, instabilities, and traffic disruptions. It is the same without Map-Versioning, with ETRs that have to reply with the same mapping; otherwise, the same problems can arise.

There are two ways Map-Versioning is helpful with respect to the synchronization problem. On the one hand, assigning version numbers to mappings helps in debugging, since quick checks on the consistency of the mappings on different ETRs can be done by looking at the Map-Version number. On the other hand, Map-Versioning can be used to control the traffic toward ETRs that announce the latest mapping.

As an example, let's consider the topology of Figure 4 where ITR A.1 of Domain A is sending unidirectional traffic to Domain B, while A.2 of Domain A exchanges bidirectional traffic with Domain B. In particular, ITR A.2 sends traffic to ETR B, and ETR A.2 receives traffic from ITR B.




Figure 4: Example topology.

Obviously, in the case of Map-Versioning, both ITR A.1 and ITR A.2 of Domain A must use the same value; otherwise, the ETR of Domain B will start to send Map-Requests.

The same problem can, however, arise without Map-Versioning, for instance, if the two ITRs of Domain A send different Locator Status Bits. In this case, either the traffic is disrupted if ETR B trusts the Locator Status Bits, or if ETR B does not trust the Locator Status Bits it will start sending Map-Requests to confirm each change in reachability.

So far, LISP does not provide any specific synchronization mechanism but assumes that synchronization is provided by configuring the different xTRs consistently. The same applies for Map-Versioning. If in the future any synchronization mechanism is provided, Map-Versioning will take advantage of it automatically, since it is included in the Record format, as described in [Section 8](#).

[12.](#) IANA Considerations

This document includes no request to IANA.

[13.](#) Acknowledgments

This work benefited support from NewNet@Paris, Cisco's Chair "Networks for the Future" at Telecom ParisTech (<http://newnet.telecom-paristech.fr>). Any opinions, findings or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of partners of the Chair.

[14.](#) References

[14.1.](#) Normative References

- [I-D.ietf-lisp-gpe]
Maino, F., Lemon, J., Agarwal, P., Lewis, D., and M. Smith, "LISP Generic Protocol Extension", [draft-ietf-lisp-gpe-06](#) (work in progress), September 2018.
- [I-D.ietf-lisp-rfc6830bis]
Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos-Aparicio, "The Locator/ID Separation Protocol (LISP)", [draft-ietf-lisp-rfc6830bis-26](#) (work in progress), November 2018.
- [I-D.ietf-lisp-rfc6833bis]
Fuller, V., Farinacci, D., and A. Cabellos-Aparicio, "Locator/ID Separation Protocol (LISP) Control-Plane", [draft-ietf-lisp-rfc6833bis-24](#) (work in progress), February 2019.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

14.2. Informative References

- [RFC6832] Lewis, D., Meyer, D., Farinacci, D., and V. Fuller, "Interworking between Locator/ID Separation Protocol (LISP) and Non-LISP Sites", [RFC 6832](#), DOI 10.17487/RFC6832, January 2013, <<https://www.rfc-editor.org/info/rfc6832>>.
- [RFC6834] Iannone, L., Saucez, D., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Map-Versioning", [RFC 6834](#), DOI 10.17487/RFC6834, January 2013, <<https://www.rfc-editor.org/info/rfc6834>>.
- [RFC7835] Saucez, D., Iannone, L., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Threat Analysis", [RFC 7835](#), DOI 10.17487/RFC7835, April 2016, <<https://www.rfc-editor.org/info/rfc7835>>.

Authors' Addresses

Luigi Iannone
Telecom ParisTech

EMail: ggx@gigix.net

Damien Saucez
INRIA Sophia Antipolis

EMail: damien.saucez@inria.fr

Olivier Bonaventure
Universite catholique de Louvain

EMail: olivier.bonaventure@uclouvain.be

