

Internet Engineering Task Force (IETF)
Request for Comments: 9302
Obsoletes: 6834
Category: Standards Track
ISSN: 2070-1721

L. Iannone
Huawei Technologies France
D. Saucez
Inria
O. Bonaventure
Universite catholique de Louvain
October 2022

Locator/ID Separation Protocol (LISP) Map-Versioning

Abstract

This document describes the Locator/ID Separation Protocol (LISP) Map-Versioning mechanism, which provides in-packet information about Endpoint-ID-to-Routing-Locator (EID-to-RLOC) mappings used to encapsulate LISP data packets. This approach is based on associating a version number to EID-to-RLOC mappings and transporting 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, which is the initial experimental specifications of the mechanisms updated by this document.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9302>.

Copyright Notice

Copyright (c) 2022 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 (<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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1.	Introduction
2.	Requirements Notation
3.	Definitions of Terms
4.	LISP-Specific Header and Map-Version Numbers
5.	Map Record and Map-Version
6.	EID-to-RLOC Map-Version Number
6.1.	The Null Map-Version
7.	Dealing with Map-Version Numbers
7.1.	Handling Dest Map-Version Number
7.2.	Handling Source Map-Version Number
8.	Security Considerations
9.	Deployment Considerations
10.	IANA Considerations
11.	References
11.1.	Normative References
11.2.	Informative References
Appendix A.	Benefits and Case Studies for Map-Versioning
A.1.	Map-Versioning and Unidirectional Traffic
A.2.	Map-Versioning and Interworking
A.2.1.	Map-Versioning and Proxy-ITRs
A.2.2.	Map-Versioning and LISP-NAT
A.2.3.	Map-Versioning and Proxy-ETRs
A.3.	RLOC Shutdown/Withdraw
Authors'	Addresses

1. Introduction

This document describes the Map-Versioning mechanism used to provide information on changes in the Endpoint-ID-to-Routing-Locator (EID-to-RLOC) mappings used in the Locator/ID Separation Protocol (LISP) [RFC9300] [RFC9301] context to perform packet encapsulation. The mechanism is totally transparent to Ingress and Egress Tunnel Routers (xTRs) not supporting or not using such functionality. The architecture of LISP is described in [RFC9299]. The reader is expected to be familiar with this introductory document.

This document obsoletes [RFC6834], which is the initial experimental specification that describes 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 modification in the RLOCs set, such as addition of, removal of, or 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 RLOCs). This information has two uses:

1. Map-Versioning enables the Egress Tunnel Router (ETR) receiving the packet to know if the Ingress Tunnel Router (ITR) is using the latest mapping version for the destination EID. If this is not the case, the ETR can directly send a Map-Request containing the updated mapping to the ITR to notify it of the latest version. The ETR can also solicit the ITR to trigger a Map-Request to obtain the latest mapping by sending a Solicit Map-Request (SMR) message. Both options are defined in [RFC9301].
2. Map-Versioning enables an ETR receiving the packet to know if it has in its EID-to-RLOC Map-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 9.

The benefits of Map-Versioning in some common LISP-related use cases are discussed in Appendix A.

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 specifications ([RFC9300] and [RFC9301]). 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-RLLOC mapping, indicating its version number (Section 6).

Null Map-Version: A Map-Version number with a value of 0x000 (zero), which is used to signal that the Map-Version feature is not used and no Map-Version number is assigned to the EID-to-RLOC mapping (Section 6.1).

Dest Map-Version number: Map-Version of the mapping in the EID-to-RLOC Map-Cache used by the ITR to select the RLOC present in the 'Destination Routing Locator' field of the outer IP header of LISP-encapsulated packets (Section 7.1).

Source Map-Version number: Map-Version of the mapping in the EID-to-
RLOC Database used by the ITR to select the RLOC present in the
'Source Routing Locator' field of the outer IP header of LISP-
encapsulated packets (Section 7.2).

4. LISP-Specific 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 Dest Map-Version number. This is done by setting the V-bit in the LISP-specific header as specified in [RFC9300] and shown in the example in Figure 1. All permissible combinations of the flags when the V-bit is set to 1 are described in [RFC9300]. Not all of the LISP-encapsulated packets need to carry version numbers. When the V-bit is set, the LISP-specific header has the following encoding:




Figure 1: LISP-Specific Header Example When Map-Versioning Is in Use

Source Map-Version number (12 bits): See Section 3.

Dest Map-Version number (12 bits): See Section 3.

5. Map Record and Map-Version

To accommodate the 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 reference, the Map Record (specified in [RFC9301]) is reported here as an example in Figure 2. This memo does not change the operation of Map-Request/Map-Reply/Map-Register messages; they continue to be used as specified in [RFC9301].

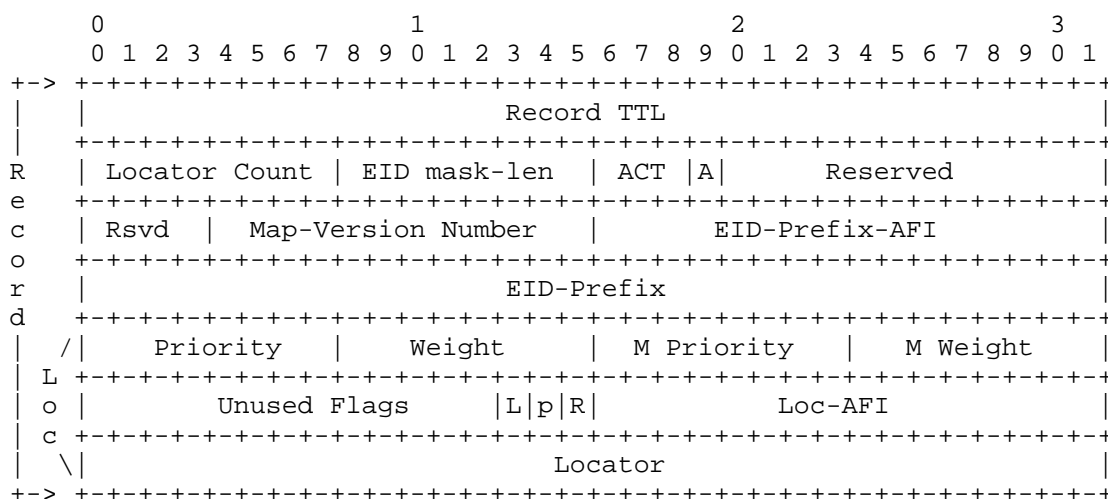


Figure 2: Map-Record Format Example

Map-Version Number: Map-Version of the mapping contained in the Record. As explained in Section 6.1, this field can be zero (0), meaning that no Map-Version is associated to the mapping.

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

A Map-Server receiving a message with an unexpected Map-Version number, for instance an old one, MUST silently drop the message and an appropriate log action SHOULD be taken.

6. 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 different version numbers, which are updated independently. An update in the version number (i.e., a newer version) MUST consist of an increment of the older version number (the only exception is for the Null Map-Version as explained in at the end of Section 6.1).

The space of version numbers has a circular order where half of the version numbers are considered greater (i.e., newer) than the current Map-Version number and the other half of the version numbers are considered smaller (i.e., older) than the current Map-Version number. This is basically a serial number on which the arithmetic described in [RFC1982] applies. The ordering enables different reactions to "older" and "newer" Map-Version numbers, whereby "older" numbers are discarded and "newer" numbers trigger Map-Requests (see Section 7 for further details). In a formal way, assuming that we have two version numbers (V1 and V2), both different from the special value Null Map-Version (see Section 6.1), and that the numbers are expressed on 12 bits, the following steps MUST be performed (in the same order shown below) to strictly define their order:

1. $V1 = V2$: The Map-Version numbers are the same.
2. $V2 > V1$: if and only if

$$V2 > V1 \text{ AND } (V2 - V1) \leq 2^{(12-1)}$$

OR

$V1 > V2 \text{ AND } (V1 - V2) > 2^{(12-1)}$

3. $V1 > V2$: otherwise.

Using 12 bits 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 + 4095) mod 4096] are smaller than 69.

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 6.1). Optionally, the initial Map-version number may be configured.

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.

6.1. The Null Map-Version

The value 0x000 (zero) is a special Map-Version number indicating that there is actually no version number associated to the EID-to-RLOC mapping. Such a value is used for special purposes and is named the Null Map-Version number.

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 NOT contain any Map-Version numbers (V-bit set to 0). If an ETR receives 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, then those packets MUST be silently dropped.

The Null Map-Version may appear in the LISP-specific header as a Source Map-Version number (Section 7.2). 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 Map-Cache, then the ETR MUST NOT compare the received Null Map-Version with the content of the EID-to-RLOC Map-Cache (Section 7.2).

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 (0x001), which is the next valid value).

7. 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 no longer reachable from a local perspective (e.g., through IGP or policy changes), the LISP site updates the mapping and also assigns a new Map-Version number. Only the latest Map-Version number has to be considered valid. Mapping updates and their corresponding Map-Version Number must be managed so that a very old version number will not be confused as a new version number (because of the circular numbering space). To this end, simple measures can be taken, like

updating a mapping only when all active traffic is using the latest version, or waiting a sufficient amount of time to be sure that the mapping in LISP caches expires, which means waiting at least as long as the mapping Time To Live (TTL) (as defined in [RFC9301]).

An ETR receiving a LISP packet with Map-Version numbers checks the following predicates:

1. The ITR that has sent the packet has an up-to-date mapping in its EID-to-RLOC Map-Cache for the destination EID and is performing encapsulation correctly. See Section 7.1 for details.
2. In the case of bidirectional traffic, the mapping in the local ETR EID-to-RLOC Map-Cache for the source EID is up to date. See Section 7.2 for details.

7.1. Handling Dest Map-Version Number

When an ETR receives a packet, the Dest 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 Dest Map-Version number. A check on this version number MUST be done, where the following cases can arise:

1. The packet arrives with the same Dest 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 Map-Cache, an up-to-date mapping. No further actions are needed.
2. The packet arrives with a Dest Map-Version number newer (as defined in Section 6) 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, the packet carries a version number that is not considered valid. Therefore, the packet MUST be silently dropped and an appropriate log action SHOULD be taken.
3. The packet arrives with a Dest Map-Version number older (as defined in Section 6) 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 Map-Cache containing stale information. The ETR MAY choose to normally process the encapsulated datagram according to [RFC9300]; however, the ITR sending the packet MUST be informed that a newer mapping is available, respecting rate-limitation policies described in [RFC9301]. This is done with a Map-Request message sent back to the ITR, as specified in [RFC9301]. One feature introduced by Map-Version numbers is the possibility of blocking traffic not using the latest mapping. This can happen if an ITR is not updating the mapping for which the ETR is authoritative, or it might be some form of attack. According to the rate-limitation policy defined in [RFC9301] for Map-Request messages, after 10 retries, Map-Requests are sent every 30 seconds; if after the first 10 retries the Dest Map-Version number in the packets is not updated, the ETR SHOULD drop packets with a stale Map-Version number. Operators can configure exceptions to this recommendation, which are outside the scope of this document.

The rule in the third case MAY be more restrictive. If the Record TTL of the previous mapping has already expired, all packets arriving with an old Map-Version MUST 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 amount of time as the Record TTL of the older mapping, all the entries in the EID-to-RLOC Map-Caches of ITRs must have

expired. Indeed, all ITRs sending traffic should have refreshed the mapping according to [RFC9301].

It is a protocol violation for LISP-encapsulated packets to contain a Dest Map-Version number equal to the Null Map-Version number (see Section 6.1).

7.2. 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 Map-Cache for the source EID, then a check MUST 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 Map-Cache. This is the regular case. The ETR has in its EID-to-RLOC Map-Cache an up-to-date copy of the mapping. No further actions are needed.
2. The packet arrives with a Source Map-Version number newer (as defined in Section 6) than the one stored in the local EID-to-RLOC Map-Cache. This means that the ETR has in its EID-to-RLOC Map-Cache a mapping that is stale and needs to be updated. A Map-Request MUST be sent to get the new mapping for the source EID, respecting rate-limitation policies described in [RFC9301].
3. The packet arrives with a Source Map-Version number older (as defined in Section 6) than the one stored in the local EID-to-RLOC Map-Cache. Note that if the mapping is already present in the EID-to-RLOC Map-Cache, this means that an explicit Map-Request has been sent and a Map-Reply has been received from an authoritative source. In this situation, the packet SHOULD be silently dropped. Operators can configure exceptions to this recommendation, which are outside the scope of this document.

If the ETR does not have an entry in the EID-to-RLOC Map-Cache for the source EID, then the Source Map-Version number MUST be ignored. See Appendix A.1 for an example of when this situation can arise.

8. Security Considerations

This document builds on the specification and operation of the LISP control and data planes. The Security Considerations of [RFC9300] and [RFC9301] apply. As such, Map-Versioning MUST NOT be used over the public Internet and MUST only be used in trusted and closed deployments. A thorough security analysis of LISP is documented in [RFC7835].

Attackers can try to trigger a large number of Map-Requests by simply forging packets with random Map-Versions. The Map-Requests are rate limited as described in [RFC9301]. With Map-Versioning, it is possible to filter packets carrying invalid version numbers before triggering a Map-Request, thus helping to reduce the effects of DoS attacks. However, it might not be enough to really protect against a DDos attack.

The present memo includes log action to be taken upon certain events. It is recommended that implementations include mechanisms (which are beyond the scope of this document) to avoid log resource exhaustion attacks.

The specifications in the present memo are relatively conservative in the sense that, in several cases, the packets are dropped. Such an approach is the outcome of considerations made about the possible risks that control plane actions that are triggered by the data plane

can be used to carry out attacks. There exists corner cases where, even with an invalid Map-Version number, forwarding the packet might be potentially considered safe; however, system manageability has been given priority with respect to having to put in place more machinery to be able to identify legitimate traffic.

9. Deployment Considerations

LISP requires multiple ETRs within the same site to provide identical mappings for a given EID-Prefix. Map-Versioning does not require additional synchronization mechanisms. Clearly, all the ETRs have to reply with the same mapping, including the same Map-Version number; otherwise, there can be an inconsistency that creates additional control traffic, instabilities, and traffic disruptions.

There are two ways Map-Versioning is helpful with respect to synchronization. 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 3 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.

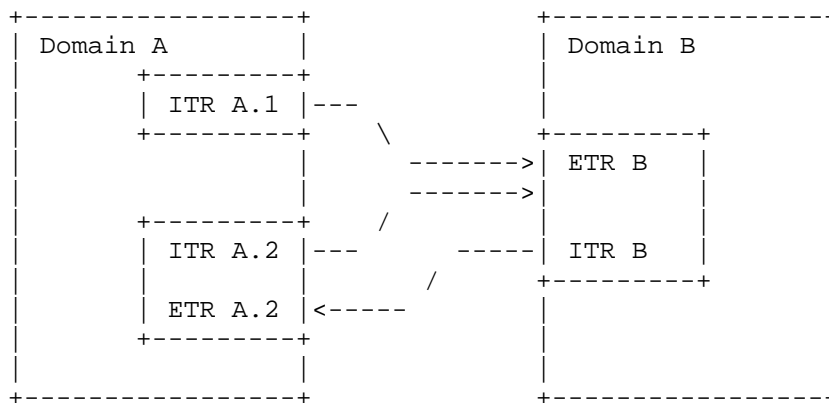


Figure 3: Example Topology

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 does not verify reachability or if ETR B 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 Map Record format, as described in Section 5.

10. IANA Considerations

This document has no IANA actions.

11. References

11.1. 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, <<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>>.
- [RFC9300] Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos, Ed., "The Locator/ID Separation Protocol (LISP)", RFC 9300, DOI 10.17487/RFC9300, October 2022, <<https://www.rfc-editor.org/info/rfc9300>>.
- [RFC9301] Farinacci, D., Maino, F., Fuller, V., and A. Cabellos, Ed., "Locator/ID Separation Protocol (LISP) Control Plane", RFC 9301, DOI 10.17487/RFC9301, October 2022, <<https://www.rfc-editor.org/info/rfc9301>>.

11.2. Informative References

- [RFC1982] Elz, R. and R. Bush, "Serial Number Arithmetic", RFC 1982, DOI 10.17487/RFC1982, August 1996, <<https://www.rfc-editor.org/info/rfc1982>>.
- [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>>.
- [RFC9299] Cabellos, A. and D. Saucez, Ed., "An Architectural Introduction to the Locator/ID Separation Protocol (LISP)", RFC 9299, DOI 10.17487/RFC9299, October 2022, <<https://www.rfc-editor.org/info/rfc9299>>.

Appendix A. 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 8.

A.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 4, since the LISP specifications do not mandate that the ETR have a mapping from the source EID.

+-----+

+-----+




Figure 4: Unidirectional Traffic between LISP Domains

An ITR is able to put both the source and destination version numbers in the LISP-specific header since the Source Map-Version number is in its database, while the Dest Map-Version number is in its cache.

The ETR checks only the Dest Map-Version number, ignoring the Source Map-Version number as specified in the final sentence of Section 7.2.

A.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 allow communication 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.

A.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 5). This case is very similar to the unidirectional traffic case described in Appendix A.1; hence, similar rules apply.




Figure 5: 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 the PITR is using the latest mapping. In the Dest Map-Version Number of the LISP-specific header, the Proxy-ITR will put the version number of the mapping it is using for encapsulation; the ETR A can use such value as defined in Section 7.1.

A.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.

A.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 6). One of the main reasons to deploy PETRs is to bypass Unicast Reverse Path Forwarding checks on the domain.




Figure 6: 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 can interchangeably put a Map-Version value or the Null Map-Version value as the Dest Map-Version number, since the receiving Proxy-ETR will ignore the field.

With this setup, the Proxy-ETR, by looking at the Source Map-Version Number, is able to check whether the mapping of the source EID has changed. This is useful to perform source RLOC validation. In the example above, traffic coming from the LISP domain has to be LISP encapsulated with a source address being an RLOC of the domain. The Proxy-ETR can retrieve the mapping associated to the LISP domain and check if incoming LISP-encapsulated traffic is arriving from a valid RLOC. A change in the RLOC-Set that can be used as source addresses can be signaled via the version number, with the Proxy-ETR able to request the latest mapping if necessary as described in Section 7.2.

A.3. RLOC Shutdown/Withdraw

Map-Versioning can also be used to perform a graceful shutdown or to withdraw 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 [RFC9301]) but without actually turning it off.

Upon updating the mapping, the RLOC will receive less and less traffic because remote LISP sites will request the updated mapping and see that it is disabled. At least one TTL, plus a little time for traffic transit, after the mapping is updated, it should be safe to shut down the RLOC gracefully, because all sites actively using the mapping should have been updated.

Note that a change in ETR for a flow can result in the reordering of the packet in the flow just as any other routing change could cause reordering.

Authors' Addresses

Luigi Iannone
Huawei Technologies France
Email: luigi.iannone@huawei.com

Damien Saucez
Inria
2004 route des Lucioles - BP 93
Sophia Antipolis
France
Email: damien.saucez@inria.fr

Olivier Bonaventure
Universite catholique de Louvain
Email: olivier.bonaventure@uclouvain.be