Network Working Group
Internet-Draft
Intended status: Experimental
Expires: September 4, 2009

L. Iannone
TU Berlin - Deutsche Telekom Laboratories AG
D. Saucez
O. Bonaventure
Universite catholique de Louvain
March 3, 2009

LISP Mapping Versioning
draft-iannone-lisp-mapping-versioning-00.txt

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on September 4, 2009.

Copyright Notice

Copyright (c) 2009 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 in effect on the date of publication of this document (http://trustee.ietf.org/license-info). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

The present document sketches an alternative approach to provide information about changes to EID-to-RLOC mappings in the context of LISP. The proposed approach is based on a versioning system for the EID-to-RLOC mapping itself. When there is a change in the mapping (where change could mean adding/removing an RLOC or just a modification in the priority or weight of one or more RLOCs) a new version number is generated and propagated in the LISP data packet. In the LISP context, ETRs do not keep state that allows to know when an ITR changes a mapping. The versioning system is a data-driven mechanism to annonce those changes.

In order to support such an approach, the LISP encapsulation need to be modified. In particular LISP-encapsulated data packets have to contain the version number of the mapings used to select the RLOCs in the outer header. These version numbers are contained in a "new" LISP header.

The mappings are distributed as usual through the mapping distribution system (e.g., CONS, ALT); versioning is only a mean to announce that something has changed in the mapping. The infrastructure built by each specific mapping protocol does not change anyhow. Nevertheless, two modifications are needed. The first modification consist in including version number in the Map-Reply messages. The second modification consist in the introduction of a new message, the "Map-Update-Notification" message used by ETRs to notify ITRs that the mapping used to encapsulate the packet is old and needs to be updated. This message does not contain the mapping, it just suggests ITRs to perform a Map-Request in order to retrieve the updated mapping.

Table of Contents

1. Requirements notation 4
2. Introduction 5
3. EID-to-RLOC Mapping Version Number 7
4. Version Numbers wrap-around 8
5. Dealing with Version Numbers 9
5.1. Handling Destination Mapping Version Number 9
5.2. Handling Source Mapping Version Number 10
6. Proposed changes to the LISP header 12
7. Proposed changes to the Map-Reply Packet format 14
8. Map-Update-Notification Message 15
9. Further Observations 16
9.1. Mapping Versioning and RLOCs reachability 16
9.2. Mapping Versioning to simplify LISP implementation 16
9.3. Mapping Versioning and original LISP cohexistence 17
10. Security Considerations 18
10.1. Robustness against reachability information spoofing. . . 18
10.2. Robustness against traffic disruption 18
10.3. Robustness of the Map-Update-Notification message 18
10.4. About robusteness of Reachability bits 19
11. Aknoledgements 20
12. References 21
12.1. Normative References 21
12.2. Informative References 21
Authors' Addresses 22

1. Requirements notation

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].

2. Introduction

The present document introduces the use of version numbers in order to provide information on a change in the EID-to-RLOC mappings used in the LISP ([I-D.farinacci-lisp]) context to perform encapsulation. The approach is based on the idea that version numbers are associated to each mapping. When a mapping is changes, a new version number is assigned to the updated mapping. A change in a 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. The change can even consist in the reachability of one or more RLOCs. Reachability information is intended from the local-domain perspective, and can be obtained for instance monitoring IGP's routing tables. This should not be confused with the intra-domain "Locator Path Liveness problem" described in [I-D.meyer-loc-id-implications].

With this approach, LISP-encapsulated data packets have to contain the version number of the mapings used to select the RLOCs in the outer header. These version numbers are contained in a "new" LISP header (described in Section 6). When an ITR encapsulates a packet, 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 receiving the packet to know if the ITR that sent it is using the latest mapping for the destination EID. If it is not the case the ETR will send a "Map-Update-Notification" message (newly defined in Section 8) to notify the ITR that a more recent version of the mapping is available and can be retrieved as usual from the mapping distribution system. Due to security issues (discussed in Section 10.3), this message does not contain the mapping. On the other hand, it enables the ETR receiving the packet to know if it has in its cache the latest mapping for the source EID.

The versioning approach does not need to re-design the mapping distribution infrastructure, which is always done through the mapping distribution protocol (e.g., CONS [I-D.meyer-lisp-cons], ALT [I-D.fuller-lisp-alt]). The mappings are distributed as usual through the Map-Request/Map-Reply message exchange. There are only the following two modifications that need to be introduced:

1. Map-Reply messages need to include the mapping version (described in Section 7).

2. Support has to be added for the new Map-Update-Notification message.

3. EID-to-RLOC Mapping Version Number

The EID-to-RLOC Mapping Version Number consist in an unsigned 15-bit integer. The version number is assigned in a per-mapping fashion, meaning that different mappings will have assigned a different version number, which is also updated independently. An update in the version number (i.e., a newer version) consist in incrementing by one the older version number.

The space of version numbers has a circular order where half of the version numbers is greater than the current Mapping Version Number and the other half is smaller than Mapping Version Number. In a more formal way, assuming we have two version numbers V1 and V2 and that the numbers are expressed on N bits, the following three cases may happen:

V1 = V2 : This is the exact match case.

V1 < V2 : True if and only if V1 < V2 < (V1 + 2**(N-1)).

V1 > V2 : True if and only if V1 > V2 > (V1 - 2**(N-1)).

As an example, using 24 bits, if the Mapping Version Number is 0, versions in]1; (2**14)-1[are greater and versions in [2**14; (2**15)-1[are smaller.

4. Version Numbers wrap-around

The proposed 15 bits size for the Mapping Version Number based on the assumption that Map-Requests are rate limited with a granularity of seconds. Using a granularity of seconds and assuming that a new version is issued each second, we have around 9 hours of autonomy before the versions wraps around, which is a reasonable time. Alternatively a granularity of minutes can also be used, as for the TTL of the Map-Reply([I-D.farinacci-lisp]). Nevertheless, using a granularity of minutes leads to very long (pointless) wrap-around periods. Hereafter there is a table with a rough estimation of the obtained autonomy with different sizes of the version number and different time granularity.

Version Number Size (bits)	Time before wrap around	
	Granularity: Minutes	Granularity: Seconds
32	8171 Years	136 Years
30	2042 Years	34 Years
24	31 Years	194 Days
16	45 Days	18 Hours
15	22 Days	9 Hours
14	11 Days	4 Hours

Figure 1

5. Dealing with Version Numbers

The main idea of using Mapping Version Numbers is that whenever there is a change in the mapping (e.g., adding/removing RLOCs, a change in the weights due to new TE policies, or a change in the priorities) or an ISP realizes that one or more of its own RLOCs are not reachable anymore from a local perspective (e.g., through IGP, due to route flap, or policy changes) the ISP updates the mapping with a new mapping version number.

In order to announce in a data-driven fashion that the mapping has been updated, the version numbers used to encapsulate the packet are embedded in the packets encapsulation. This means that the header needs to contain two mapping version numbers. A first one from the EID-to-RLOC mapping in the EID-to-RLOC Database used to select the source RLOC, and called Source Mapping Version Number. A second one from the EID-to-RLOC mapping in the EID-to-RLOC Cache used to select the destination RLOC, and called Destination Mapping Version Number. The purpose of carring these version numbers is two-fold, allowing the ETR receiving the encapsulated packet to check if i) the ITR has an up-to-date mapping; ii) the mapping in the local cache for the source EID is up-to-date. If the first condition does not hold the Map-Update-Notification (see Section 8) is used to make the ITR aware that a newer mapping is available. The ITR will retrieve the mapping with the normal Map-Request/Map-Reply mechanism. If the second condition does not hold the ETR sends a Map-Request for the source EID (obtained from the inner header of the packet) in order to retrieve the up-to-date mapping.

Further details on how to handle the Source and Destination Mapping Version Numbers are provided hereafter, while the proposed new LISP header format is detailled in Section 6 (Figure 2).

5.1. Handling Destination Mapping Version Number

When an ETR receives a packet, the destination mapping version number relates to the mapping of the domain the ETR is part of, thus the ETR has the correct and up-to-date Version Number for the destination mapping. A check on this version number is done, where the following cases can arise:

- The packets arrive with the same destination mapping version number stored in the LEID-to-RLOC Database. This is the correct regular case. The ITR sending the packet has in its EID-to-RLOC Cache an up-to-date mapping. No further actions are needed.

- The packets arrive with an destination mapping 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. Further actions are needed. The ITR sending the packet should be informed that a newer mapping is available. This is done with a "Map-Update-Notification" message sent back to the ITR, soliciting an update of the mapping. This message should be rate limited and if after a certain amount of retries the mapping is not updated packets coming from that ITR with smaller mapping version number can be silently dropped, since most likely there is a spoof or the ITR is not behaving correctly. Note that the rule can be even more restrictive. If the mapping has been the same for a period of time as long as the TTL (defined in LISP [I-D.farinacci-lisp]) of the previous version of the mapping, all packets arriving with an old mapping version can be silently dropped right away. Indeed, if the new mapping with the updated version number has been stable for at least the same time as the TTL of the older mapping, all the entries in the caches of ITRs must have expired. If packets with old mapping version number are still received, the reason is that either someone has not respected the TTL, or it is a spoof, or a not valid behaviour w.r.t. the specifications. In all those cases the packet can be silently dropped.

- The packet arrives with a destination mapping version number greater (i.e., newer) than the one stored in the EID-to-RLOC Database. Since the ETR is in the domain owning the mapping, this means that someone is not behaving correctly w.r.t. the specifications, thus the packets carries a not valid version number and can be silently dropped.

5.2. Handling Source Mapping Version Number

When an ETR receives a packet, the source mapping version number relates to the mapping owned by the domain the ITR is part of and whose copy should be present in the EID-to-RLOC Cache of the ETR (except for the first packet that generates a cache miss triggering a Map-Request message). A check on this version number is done, where the following cases can arise:

- The packet arrives with the same source mapping version number

stored in the EID-to-RLOC Cache. This is the correct regular case. The ETR has in its EID-to-RLOC Cache an up-to-date copy of the mapping. No further actions are needed.

- The packet arrives with a source mapping version number smaller (i.e., older) than the one stored in the local EID-to-RLOC Cache. Such a case is not valid w.r.t. the specifications and hence the packet is silently dropped. If the mapping is already present in the EID-to-RLOC Cache, this means that an explicit Map-Request has

been send and a Map-Reply has been received. The latter sent by the "owner" of the mapping. Assuming the mapping system is not corrupted anyhow, the mapping version in the EID-to-RLOC Cache is the correct one, hence the packet is not valid.

- The packet arrives with a source mapping version number greater (i.e., newer) than the one stored in the local EID-to-RLOC Cache. The ETR has in its EID-to-RLOC Cache a mapping that is stale and needs to be updated. The packet is considered valid but further actions are needed. In particular a Map-Request must be sent to the owner of the source mapping to retrieve the new mapping. This should be rate limited.

6. Proposed changes to the LISP header

In order for the versioning approach to work, the LISP encapsulation format needs to be changed. In particular the LISP header is modified to include the source mapping version number and the destination mapping version number.

Here is the new packet format, changes occur only on the LISP header:

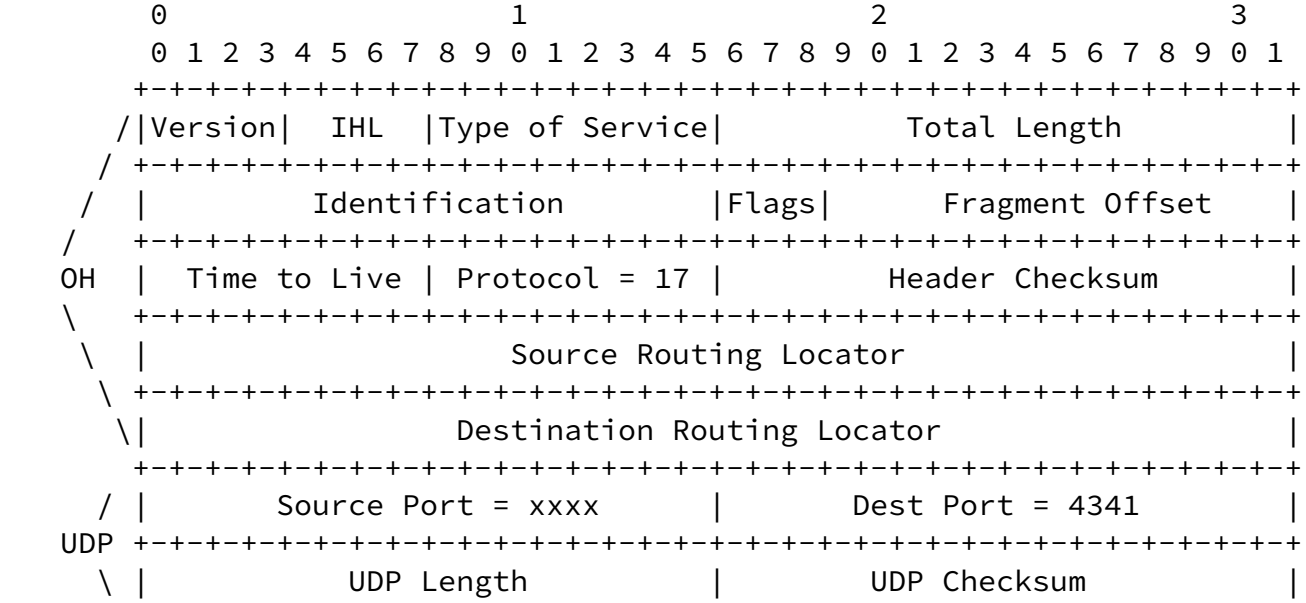

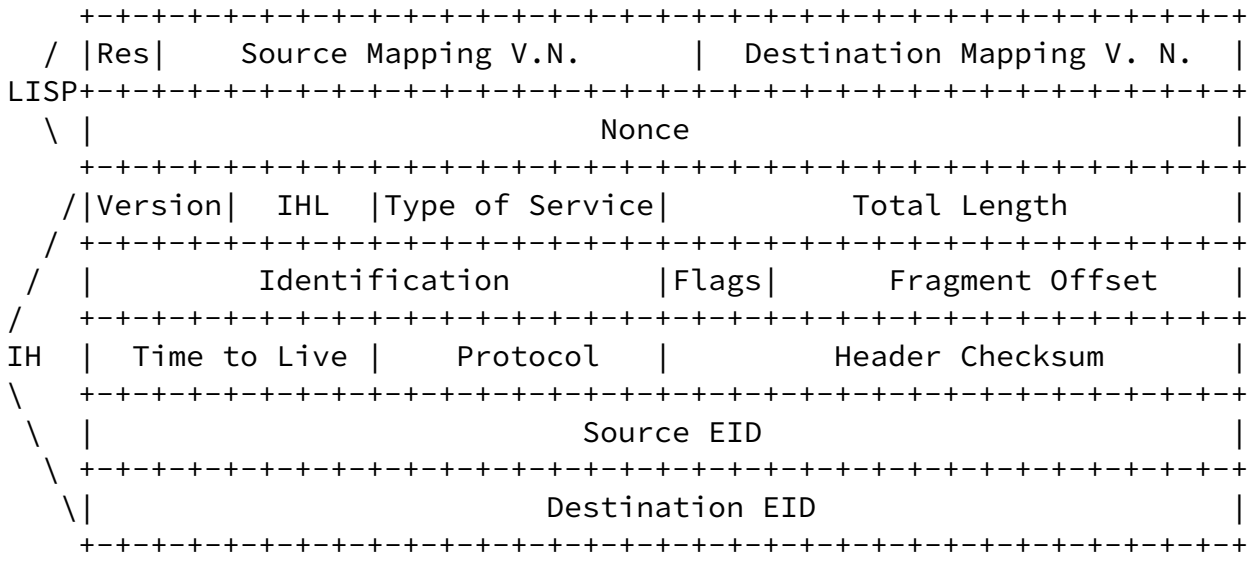

Figure 2

Reserved (2 bits): Reserved for future use. Sent as 0 and ignored on reception. A possible use of these two bits is proposed in Section 9.3.

Source Mapping Version Number (15 bits): Version of the mapping used by the ITR to select the RLOC to put in the "Source Routing Locator" field. Note that the mapping used for such a selection is determined by the Source EID, used as lookup key in the EID-to-RLOC Database of the ITR.

Destination Mapping Version Number (15 bits): Version of the mapping used by the ITR to select the RLOC to put in the "Destination Routing Locator" field. Note that the mapping used for such a selection is determined by the Destination EID, used as lookup key in the EID-to-RLOC Cache of the ITR.

Note that the change proposed concerns only the part of the LISP specific header originally containing the reachability bits.

7. Proposed changes to the Map-Reply Packet format

To accommodate the proposed mechanism, also the Map-Reply packet format needs to be changed (original definition of the Map-Reply packet format can be found in [I-D.farinacci-lisp]). This is a simple change to carry the version number assigned to the mapping in this message. It does not mean a change in the mapping distribution system from an architectural point of view.

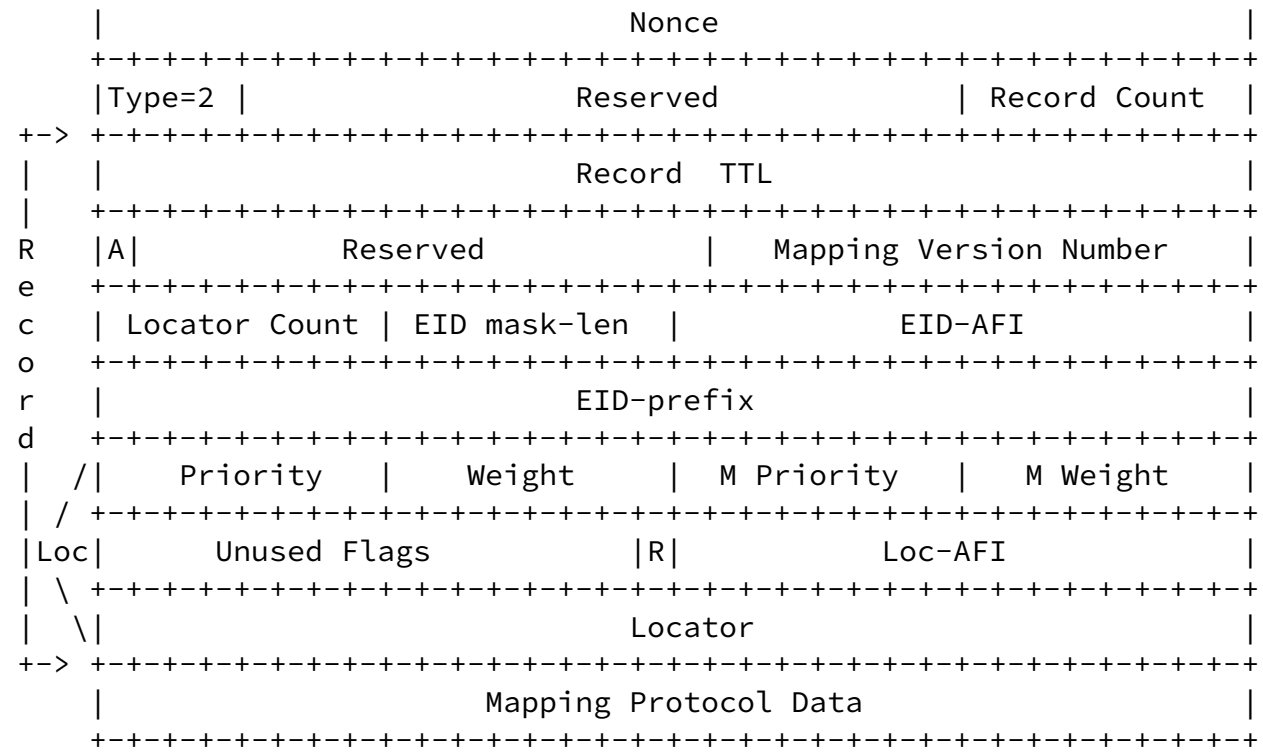

Mapping Version Number: Version Number of the mapping in the Record.

The proposed change does not make the whole packet bigger. Indeed, the proposed solution is to put the fields "Locator Count", "EID mask-len" and "EID-AFI" in the same 32 bits block. In [I-D.farinacci-lisp] "locator Count" and "EID mask-len" are in the same 32 bits block but "EID-AFI" is in a different 32 bits block. The first two fields are alligned left, while the latter is alligned right, with a gap of 32 bits, from which only one is used (the authoritative bit "A") while the other 31 are just marked as "Reserved". After compacting the fields, the "A" bit is now placed in the same 32 bits block like the mapping version number. Even, if fields have been displaced, they keep their original definition described in [I-D.farinacci-lisp].

8. Map-Update-Notification Message

From the LISP packet that triggered a Map-Update-Notification, it is known the ITR has to be notified about the existence of a newer

mapping. It is not possible to send the mapping directly with a Map-Reply message, since this can introduce security threats. This is why the Map-Update-Notification message is needed and must only be a notification that a new version of the mapping is available. The message is meant to trigger a Map-Request from the ITR, in order to update its copy of the mapping.

In [I-D.farinacci-lisp], the equivalent of the Map-Update-Notification message is the SMR bit set in the LISP header of the data packets. The SMR bit approach does not work in the case of unidirectional traffic (e.g., due to unidirectional flows or Traffic Engineering). The ETR has no mean to send back to the ITR that a new mapping is available. Moreover,with the SMR bit only there is not enough information for understanding if ITRs have updated the mapping so that the SMR bit can be reset.

The format for the Map-Update-Notification message will be provided in future versions of this draft.

9. Further Observations

In the following sections we provide more discussion on various aspects of the mapping versioning. Security observations are instead grouped in Section 10.

9.1. Mapping Versioning and RLOCs reachability

When the reachability problem occurs on the path between two RLOCs of different LISP sites (this is called path-liveness problem in the recent draft by D. Meyer and D. Lewis [I-D.meyer-loc-id-implications]), the versioning approach does not help. In this case other mechanisms are necessary, however, such an issue is not new and is part of all loc/ID split solutions, thus versioning does not introduce a new issue.

9.2. Mapping Versioning to simplify LISP implementation

The use of versioning numbers for mapping has also the effect of simplifying operations. The set of RLOCs can always be maintained ordered, since no consistency must be preserved with the reachability bits.

In other words it is not necessary to "append" new locators to the existing ones as explained in Section 6.5 of the LISP draft. A new mapping with a new version number will be issued, and since the old locators are still valid the transition will be disruptionless. The same applies for the case a RLOC is withdrawn. There is no need to maintain holes in the list of locators, as previously, for sites that are not using the RLOC that has been withdrawn, the transition will be disruptionless.

It is even possible to perform a graceful shutdown. This is obtained by simply issuing a mapping where the specific RLOC to be shut down is withdrawn, but without actually turning it off. Once no more traffic is received by the RLOC, because all sites have updated the mapping the RLOC can be shut down safely.

All of these operations, as already stated, do not need to maintain any consistency among reachability bits, and the way RLOC are stored in the cache. This eases implementation.

Finally, with the versioning approach there is no need to perform a "clock sweep" as described in Section 6.5.1 of the LISP draft. Every LISP site communicating to a specific LISP site that has updated the mapping will be informed of the available new mapping in a data-driven manner.

9.3. Mapping Versioning and original LISP cohexistence

The solution proposed in this draft includes changes in the LISP header. However, and for experimentation purpose, it could be worth instead of replacing the original LISP header, to extend the original LISP header by appending the one proposed in this draft.

Alternatively, the first two bits of the LISP specific header can be used to select the type of header. For instance the second leftmost bit can be used to state if the following bits have to be interpreted as reachability bits or version numbers.

Note that the reference document for LISP implementation and interoperability tests remains [I-D.farinacci-lisp]. The present draft proposes an alternative approach that needs experimental validation and should not be considered as a permanent design of the LISP protocol.

10. Security Considerations

Mapping Versioning present the same reactivity of Reachability bits and SMR bit, however, provides more robustness to possible attacks.

10.1. Robustness against reachability information spoofing.

Compared to the reachability bits solution, since the new mapping is obtained through the mapping system the data plane results robust to reachability information spoofing. Such a statement is true assuming that the mapping distribution system is secure. Security issues concerning specific mapping distribution system are described in the documents specific to each proposal.

10.2. Robustness against traffic disruption

Attackers can try to use the Mapping version number to trigger either Map-Update-Notification messages or Map-Request messages, by simply sending packets for all the possible version numbers. Nevertheless, as described in Section 5 it is possible to easily filter a large part of the packets containing a wrong version number. If the version number is sufficiently large, exploring all the possible numbers will take too much time to be really feasible.

Furthermore, even if the attacker is able to "guess" the correct version number to trigger a Map-Request, the traffic is not stopped. The normal behavior will be that the xTR continue to use the mapping, while asking in parallel for a new one, in a rate limited fashion. This is not the case with explicit reachability bits where an attacker can set all RLOCs to down with one single packet, with disruptive consequences on the traffic.

It is clear, that mapping versioning does not protect against DoS and

DDoS attacks, where an xTR looses processing power doing version number checks on packets sent by attackers.

10.3. Robustness of the Map-Update-Notification message

The Map-Update-Notification should never include the new mapping inside the packet itself, otherwise a security threat would be introduced, where attackers send fake Map-Update-Notifications messages poisoning the cache.

The mapping could be included in the Map-Update-Notification only in the case where the sender can be clearly identified (e.g., using certificates or a PKI).

10.4. About robusteness of Reachability bits

The scenarii presented in the previous sections are correct if ETR react to Reachability bits change without performing a further check with the ITR. In other ways the ETR blindly trusts the content of the Reachability bits. If ETR does not trust such a content and before changing the reachability state of the RLOCs it can send a Map-Request in order to confirm the change. On the one hand, having such a confirmation improves the robusteness of the reachability bits mechanism. On the other hand, this is very close to the mapping versioning system, with the only difference that the use of version numbers enables a fine control on when to update a mapping or when to notify that a mapping has been updated.

11. Aknoledgements

The authors would like to thank Pierre Francois and Dino Farinacci for their comments and review.

12. References

12.1. Normative References

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

12.2. Informative References

[I-D.farinacci-lisp] Farinacci, D., Fuller, V., Oran, D., Meyer, D., and S. Brim, "Locator/ID Separation Protocol (LISP)",

draft-farinacci-lisp-11 (work in progress), December 2008.

[I-D.fuller-lisp-alt]
Fuller, V., Meyer, D., and D. Farinacci, "LISP Alternative Topology (LISP+ALT)", draft-fuller-lisp-alt-03 (work in progress), October 2008.

[I-D.meyer-lisp-cons]
Brim, S., "LISP-CONS: A Content distribution Overlay Network Service for LISP", draft-meyer-lisp-cons-04 (work in progress), April 2008.

[I-D.meyer-loc-id-implications]
Meyer, D. and D. Lewis, "Architectural Implications of Locator/ID Separation", draft-meyer-loc-id-implications-00 (work in progress), December 2008.

Authors' Addresses

Luigi Iannone
TU Berlin - Deutsche Telekom Laboratories AG
Ernst-Reuter Platz 7

Berlin
Germany

Email: luigi@net.t-labs.tu-berlin.de

Damien Saucez
Universite catholique de Louvain
Place St. Barbe 2
Louvain la Neuve
Belgium

Email: damien.saucez@uclouvain.be

Olivier Bonaventure
Universite catholique de Louvain
Place St. Barbe 2
Louvain la Neuve
Belgium

Email: olivier.bonaventure@uclouvain.be