

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 19, 2013

D. Saucez
INRIA
L. Iannone
Telecom ParisTech
O. Bonaventure
Universite catholique de Louvain
October 16, 2012

LISP Threats Analysis
draft-ietf-lisp-threats-03.txt

Abstract

This document analyzes the threats against the security of the Locator/Identifier Separation Protocol (LISP) and proposes a set of recommendations to mitigate some of the identified security risks.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 19, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as

described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Definition of Terms	3
3.	On-path Attackers	4
4.	Off-Path Attackers: Reference Environment	4
5.	Data-Plane Threats	6
5.1.	EID-to-RLLOC Database Threats	6
5.2.	EID-to-RLLOC Cache Threats	7
5.2.1.	EID-to-RLLOC Cache poisoning	7
5.2.2.	EID-to-RLLOC Cache overflow	9
5.3.	Attacks not leveraging on the LISP header	9
5.4.	Attacks leveraging on the LISP header	10
5.4.1.	Attacks using the Locator Status Bits	10
5.4.2.	Attacks using the Map-Version bit	11
5.4.3.	Attacks using the Nonce-Present and the Echo-Nonce bits	12
5.4.4.	Attacks using the ID Instance bits	13
6.	Control Plane Threats	13
6.1.	Attacks with Map-Request messages	13
6.2.	Attacks with Map-Reply messages	15
6.3.	Gleaning Attacks	16
7.	Threats concerning Interworking	17
8.	Threats with Malicious xTRs	18
9.	Security of the Proposed Mapping Systems	21
9.1.	LISP+ALT	21
9.2.	LISP-DDT	22
10.	Threats concerning LISP-MS	23
10.1.	Map Server	23
10.2.	Map Resolver	24
11.	Suggested Recommendations	26
12.	IANA Considerations	28
13.	Security Considerations	28
14.	Acknowledgments	28
15.	References	28
15.1.	Normative References	28
15.2.	Informative References	29
Appendix A.	Document Change Log	30
	Authors' Addresses	31

1. Introduction

The Locator/ID Separation Protocol (LISP) is defined in [\[I-D.ietf-lisp\]](#). The present document aims at assessing the security level and identifying security threats in the LISP specification. As a result of the performed analysis, the document also proposes some recommendations aiming at improving LISP's resiliency against off-path attackers.

The document is composed of two main parts: the first discussing the LISP data-plane; while the second discussing the LISP control-plane.

The LISP data-plane consists of LISP packet encapsulation, decapsulation, and forwarding and includes the EID-to-RLLOC Cache and EID-to-RLLOC Database data structures used to perform these operations.

The LISP control-plane consists in the mapping distribution system, which can be one of the mapping distribution systems proposed so far (e.g., [\[I-D.ietf-lisp\]](#), [\[I-D.fuller-lisp-ddt\]](#), [\[I-D.ietf-lisp-alt\]](#), [\[I-D.ietf-lisp-ms\]](#), [\[I-D.meyer-lisp-cons\]](#), and [\[I-D.lear-lisp-nerd\]](#)), and the Map-Request, Map-Reply, Map-Register, and Map-Notification messages.

This document does not consider all the possible uses of LISP as discussed in [\[I-D.ietf-lisp\]](#). The document focuses on LISP unicast, including as well LISP Interworking [\[I-D.ietf-lisp-interworking\]](#), LISP-MS [\[I-D.ietf-lisp-ms\]](#), LISP Map-Versioning [\[I-D.ietf-lisp-map-versioning\]](#), and briefly considering the ALT mapping system described in [\[I-D.ietf-lisp-alt\]](#) and the Delegated Database Tree mapping system described in [\[I-D.fuller-lisp-ddt\]](#). The reading of these documents is a prerequisite for understanding the present document.

Unless otherwise stated, the document assumes a generic IP service and does not discuss the difference, from a security viewpoint, between using IPv4 or IPv6.

2. Definition of Terms

The present document does not introduce any new term, compared to the main LISP specification. For a complete list of terms please refer to [\[I-D.ietf-lisp\]](#).

3. On-path Attackers

On-path attackers are attackers that are able to capture and modify all the packets exchanged between an Ingress Tunnel Router (ITR) and an Egress Tunnel Router (ETR). To cope with such an attacker, cryptographic techniques such as those used by IPSec ([\[RFC4301\]](#)) are required. We do not consider that LISP has to cope with such kind of attackers.

Mobile IP has also considered time-shifted attacks from on-path attackers. A time-shifted attack is an attack where the attacker is temporarily on the path between two communicating hosts. While it is on-path, the attacker sends specially crafted packets or modifies packets exchanged by the communicating hosts in order to disturb the packet flow (e.g., by performing a man in the middle attack). An important issue for time-shifted attacks is the duration of the attack once the attacker has left the path between the two communicating hosts. We do not consider time-shifted attacks in this document.

4. Off-Path Attackers: Reference Environment

Throughout this document we consider the reference environment shown in the figure below. There are two hosts attached to LISP routers: HA and HB. HA is attached to the two LISP xTRs LR1 and LR2, which in turn are attached to two different ISPs. HB is attached to the two LISP xTRs LR3 and LR4. HA and HB are the EIDs of the two hosts. LR1, LR2, LR3, and LR4 are the RLOCs of the xTRs.

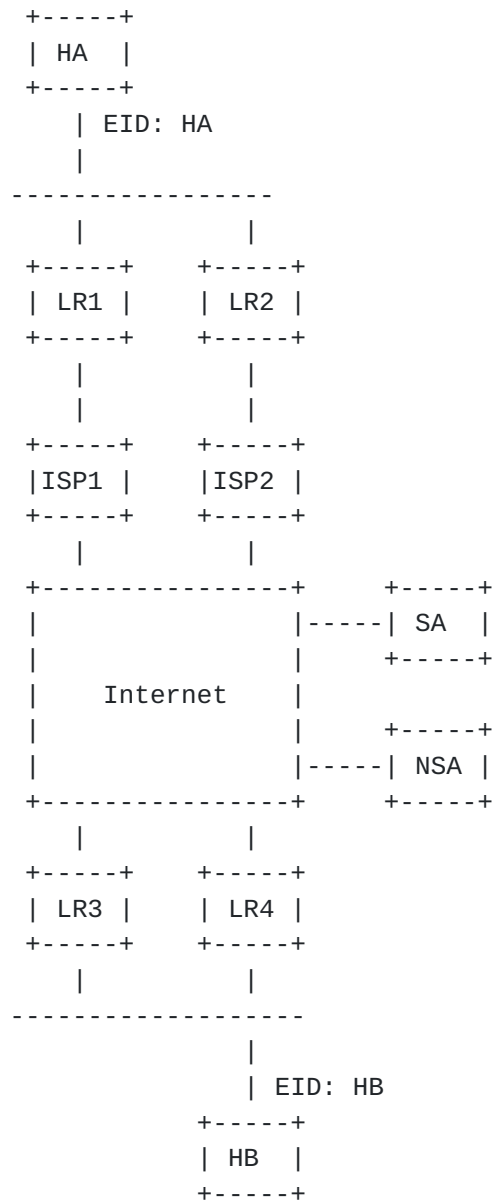


Figure 1: Reference Network

We consider two off-path attackers with different capabilities:

SA is an off-path attacker that is able to send spoofed packets, i.e., packets with a different source IP address than its assigned IP address. SA stands for Spoofing Attacker.

NSA is an off-path attacker that is only able to send packets whose source address is its assigned IP address. NSA stands for Non Spoofing Attacker.

It should be noted that with LISP, packet spoofing is slightly

different than in the current Internet. Generally the term "spoofed packet" indicates a packet containing a source IP address that is not the one of the actual originator of the packet. Since LISP uses encapsulation, the spoofed address can be in the outer header as well as in the inner header, this translates in two types of spoofing:

EID Spoofing: the originator of the packet puts in it a spoofed EID. The packet will be normally encapsulated by the ITR of the site (or a PITR if the source site is not LISP enabled).

RLOC Spoofing: the originator of the packet generates directly a LISP-encapsulated packet with a spoofed source RLOC.

Note that the two types of spoofing are not mutually exclusive, rather all combinations are possible and can be used to perform different kind of attacks.

In the reference environment, both SA and NSA attackers are capable of sending LISP encapsulated data packets and LISP control packets. This means that SA is able to perform both RLOC and EID spoofing while NSA can only perform EID spoofing. They may also send other types of IP packets such as ICMP messages. We assume that both attackers can query the LISP mapping system (e.g., through a public Map Resolver) to obtain the mappings for both HA and HB.

5. Data-Plane Threats

This section discusses threats and attacks related to the LISP data-plane. More precisely, we discuss the operations of encapsulation, decapsulation, and forwarding as well as the content of the EID-to-RLOC Cache and EID-to-RLOC Database as specified in the original LISP document ([[I-D.ietf-lisp](#)]).

We start considering the two main data structures of LISP, namely the EID-to-RLOC Database and the EID-to-RLOC Cache. Then, we look at the data plane attacks that can be performed by a spoofing off-path attacker (SA) and discuss how they can be mitigated by LISP xTRs. In this analysis, we assume that the LR1 and LR2 (resp. LR3 and LR4) xTRs maintain an EID-to-RLOC Cache that contains the required mapping entries to allow HA and HB to exchange packets.

5.1. EID-to-RLOC Database Threats

The EID-to-RLOC Database on each xTR maintains the set of mappings related to the EID-Prefixes that are "behind" the xTR. Where "behind" means that at least one of the xTR's globally visible IP addresses is a RLOC for those EID-Prefixes.

As described in [[I-D.ietf-lisp](#)], the EID-to-RLOC Database content is determined by configuration. This means that the only way to attack this data structure is by gaining privileged access to the xTR. As such, it is out of the scope of LISP to propose any mechanism to protect routers and, hence, it is no further analyzed in this document.

5.2. EID-to-RLOC Cache Threats

A key component of the overall LISP architecture is the EID-to-RLOC Cache. The EID-to-RLOC Cache is the data structure that stores the bindings between EID and RLOC (namely the "mappings") to be used later on. Attacks against this data structure can happen either when the mappings are first installed in the cache (see also [Section 6](#)) or by corrupting (poisoning) the mappings already present in the cache.

5.2.1. EID-to-RLOC Cache poisoning

The content of the EID-to-RLOC Cache can be poisoned by spoofing LISP encapsulated packets. Examples of EID-to-RLOC Cache poisoning are:

Fake mapping: The cache contains entirely fake mappings that do not originate from an authoritative mapping server. This can be achieved either through gleaning as described in [Section 6.3](#) or by attacking the control-plane as described in [Section 6](#).

EID Poisoning: The EID-Prefix in a specific mapping is not owned by the originator of the entry. Similarly to the previous case, this can be achieved either through gleaning as described in [Section 6.3](#) or by attacking the control-plane as described in [Section 6](#).

EID redirection/RLOC poisoning: The EID-Prefix in the mapping is not bound to (located by) the set of RLOCs present in the mapping. This can result in packets being redirected elsewhere, eavesdropped, or even blackholed. Note that not necessarily all RLOCs are fake/spoofed. The attack works also if only part of the RLOCs, the highest priority ones, is compromised. Again, this can be achieved either through the gleaning as described in [Section 6.3](#) or by attacking the control-plane as described in [Section 6](#).

Reachability poisoning: The reachability information stored in the mapping could be poisoned, redirecting the packets to a subset of the RLOCs (or even stopping it if locator status bits are all set to 0). If reachability information is not verified through the control-plane this attack can be simply achieved by sending a spoofed packet with swapped or all locator status

bits reset. The same result can be obtained by attacking the control-plane as described in [Section 6](#). Depending on how the RLOC reachability information is stored on the router, the attack can impact only one mapping or all the mappings that share the same RLOC.

Traffic Engineering information poisoning: The LISP protocol defines two attributes associated to each RLOC in order to perform inbound Traffic Engineering (TE), namely priority and weight. By injecting fake TE attributes, the attacker is able to break load balancing policies and concentrate all the traffic on a single RLOC or put more load on a RLOC than what is expected, creating congestion. It is even possible to block the traffic if all the priorities are set to 255 (special value indicating not to use the RLOC). Corrupting the TE attributes can be achieved by attacking the control-plane as described in [Section 6](#).

Mapping TTL poisoning: The LISP protocol associates a Time-To-Live to each mapping that, once expired, allows to delete a mapping from the EID-to-RLOC Cache (or forces a Map-Request/Map-Reply exchange to refresh it if still needed). By injecting fake TTL values, an attacker can either shrink the EID-to-RLOC Cache (using very short TTL), thus creating an excess of cache miss causing a DoS on the mapping system, or it can increase the size of the cache by putting very high TTL values, up to a cache overflow (see [Section 5.2.2](#)). Corrupting the TTL can be achieved by attacking the control-plane as described in [Section 6](#). Long TTL can be use in fake mappings to increase attack duration.

Instance ID poisoning: The LISP protocol allows using a 24-bit identifier to select the forwarding table to use on the decapsulating ETR to forward the decapsulated packet. By spoofing this attribute the attacker is able to redirect or blackhole inbound traffic. Corrupting the Instance ID attribute can be achieved by attacking the control-plane as described in [Section 6](#).

Map-Version poisoning: The LISP protocol allows associating a version number to mappings ([\[I-D.ietf-lisp-map-versioning\]](#)). The LISP header can transport source and destination map-versions, describing which version of the mapping have been used to select the source and the destination RLOCs of the LISP encapsulated packet. By spoofing this attribute the attacker is able to trigger Map-Request on the receiving ETR. Corrupting the Map-Version attribute can be achieved either by attacking the control-plane as described in [Section 6](#) or by

using spoofed packets as described in [Section 5.4.2](#).

If the above listed attacks succeed, the attacker has the means of controlling the traffic.

5.2.2. EID-to-RLOC Cache overflow

Depending on how the EID-to-RLOC Cache is managed (e.g., Least Recently Used - LRU vs. Least Frequently Used - LFU) and depending on its size, an attacker can try to fill the cache with fake mappings. Once the cache is full, some mappings will be replaced by new fake ones, causing traffic disruption.

This can be achieved either through gleaning as described in [Section 6.3](#) or by attacking the control-plane as described in [Section 6](#).

Another way to generate an EID-to-RLOC Cache overflow is by injecting mapping with a fake and very large TTL value. In this case the cache will keep a large amount of mappings ending with a completely full cache. This type of attack can also be performed through the control-plane.

5.3. Attacks not leveraging on the LISP header

We first consider an attacker that sends packets without exploiting the LISP header, i.e., with the N, L, E, V, and I bits reset ([\[I-D.ietf-lisp\]](#)).

To inject a packet in the HA-HB flow, a spoofing off-path attacker (SA) can send a LISP encapsulated packet whose source is set to LR1 or LR2 and destination LR3 or LR4. The packet will reach HB as if the packet was sent by host HA. This is not different from today's Internet where a spoofing off-path attacker may inject data packets in any flow. Several existing techniques can be used by hosts to prevent such attacks from affecting established flows, e.g., [\[RFC4301\]](#) and [\[I-D.ietf-tcpm-tcp-security\]](#).

On the other hand, a non-spoofing off-path attacker (NSA) can only send a packet whose source address is set to its assigned IP address. The destination address of the encapsulated packet can be LR3 or LR4. When the LISP ETR that serves HB receives the encapsulated packet, it can consult its EID-to-RLOC Cache and verify that NSA is not a valid source address for LISP encapsulated packets containing a packet sent by HA. This verification is only possible if the ETR already has a valid mapping for HA. Otherwise, and to avoid such data packet injection attacks, the LISP ETR should reject the packet and possibly query the mapping system to obtain a mapping for the encapsulated

source EID (HA).

5.4. Attacks leveraging on the LISP header

The main LISP document [[I-D.ietf-lisp](#)] defines several flags that modify the interpretation of the LISP header in data packets. In this section, we discuss how an off-path attacker could exploit this LISP header.

5.4.1. Attacks using the Locator Status Bits

When the L bit is set to 1, it indicates that the second 32-bits longword of the LISP header contains the Locator Status Bits. In this field, each bit position reflects the status of one of the RLOCs mapped to the source EID found in the encapsulated packet. In particular, a packet with the L bit set and all Locator Status Bits set to zero indicates that none of the locators of the encapsulated source EID are reachable. The reaction of a LISP ETR that receives such a packet is not clearly described in [[I-D.ietf-lisp](#)].

A spoofing off-path attacker (SA) can send a data packet with the L bit set to 1, all Locator Status Bits set to zero, a spoofed source RLOC (e.g. LR1), destination LR3, and containing an encapsulated packet whose source is HA. If LR3 blindly trusts the Locator Status Bits of the received packet it will set LR1 and LR2 as unreachable, possibly disrupting ongoing communication.

Locator Status Bits can be blindly trusted only in secure environments. In the general unsecured Internet environment, the safest practice for xTRs is to confirm the reachability change through the control plane (e.g., RLOC probing). In the above example, LR3 should note that something has changed in the Locator Status Bits and query the mapping system (assuming it is trusted) in order to confirm status of the RLOCs of the source EID.

A similar attack could occur by setting only one Locator Status Bit to 1, e.g., the one that corresponds to the source RLOC of the packet.

If a non-spoofing off-path attacker (NSA) sends a data packet with the L bit set to 1 and all Locator Status Bits set to zero, this packet will contain the source address of the attacker. Similarly as in [Section 5.3](#), if the xTR accepts the packet without checking the EID-to-RLOC Cache for a mapping that binds the source EID and the source RLOC of the received packet, then the same observation like for the spoofing attacker (SA) apply with the difference that instead of complete disruption, the traffic will flow through only one RLOC, possibly resulting in a DoS attack.

Otherwise, if the xTR does make the check through the EID-to-RLLOC Cache, it should reject the packet because its source address is not one of the addresses listed as RLLOCs for the source EID. Nevertheless, in this case a Map-Request should be sent, which can be used to perform Denial of Service attacks. Indeed an attacker can frequently change the Locator Status Bits in order to trigger a large amount of Map-Requests. Rate limitation, as described in [\[I-D.ietf-lisp\]](#), does not allow sending high number of such a request, resulting in the attacker saturating the rate with these spoofed packets.

[5.4.2.](#) Attacks using the Map-Version bit

The Map-Version bit is used to indicate whether the low-order 24 bits of the first 32 bits longword of the LISP header contain a Source and Destination Map-Version. When a LISP ETR receives a LISP encapsulated packet with the Map-Version bit set to 1, the following actions are taken:

- o It compares the Destination Map-Version found in the header with the current version of its own mapping, in the EID-to-RLLOC Database, for the destination EID found in the encapsulated packet. If the received Destination Map-Version is smaller (i.e., older) than the current version, the ETR should apply the SMR procedure described in [\[I-D.ietf-lisp\]](#) and send a Map-Request with the SMR bit set.
- o If a mapping exists in the EID-to-RLLOC Cache for the source EID, then it compares the Map-Version of that entry with the Source Map-Version found in the header of the packet. If the stored mapping is older (i.e., the Map-Version is smaller) than the source version of the LISP encapsulated packet, the xTR should send a Map-Request for the source EID.

A spoofing off-path attacker (SA) could use the Map-Version bit to force an ETR to send Map-Request messages. The attacker can retrieve the current source and destination Map-Version for both HA and HB. Based on this information, it can send a spoofed packet with an older Source Map-Version or Destination Map-Version. If the size of the Map-Request message is larger than the size of the smallest LISP-encapsulated packet that could trigger such a message, this could lead to amplification attacks (see [Section 6.1](#)). However, [\[I-D.ietf-lisp\]](#) recommends to rate limit the Map-Request messages that are sent by an xTR. This prevents the amplification attack, but there is a risk of Denial of Service attack if an attacker sends packets with Source and Destination Map-Versions that frequently change. In this case, the ETR could consume its entire rate by sending Map-Request messages in response to these spoofed packets.

A non-spoofing off-path attacker (NSA) cannot success in such an attack if the destination xTR rejects the LISP encapsulated packets that are not sent by one of the RLOCs mapped to the included source EID. If it is not the case, the attacker can be able to perform attacks concerning the Destination Map Version number as for the spoofing off-path attacker (SA).

5.4.3. Attacks using the Nonce-Present and the Echo-Nonce bits

The Nonce-Present and Echo-Nonce bits are used when verifying the reachability of a remote ETR. Assume that LR3 wants to verify that LR1 receives the packets that it sends. LR3 can set the Echo-Nonce and the Nonce-Present bits in LISP data encapsulated packets and include a random nonce in these packets. Upon reception of these packets, LR1 will store the nonce sent by LR3 and echo it when it returns LISP encapsulated data packets to LR3.

A spoofing off-path attacker (SA) could interfere with this reachability test by sending two different types of packets:

1. LISP data encapsulated packets with the Nonce-Present bit set and a random nonce and the appropriate source and destination RLOCs.
2. LISP data encapsulated packets with the Nonce-Present and the Echo-Nonce bits both set and the appropriate source and destination RLOCs. These packets will force the receiving ETR to store the received nonce and echo it in the LISP encapsulated packets that it sends.

The first type of packet should not cause any major problem to ITRs. As the reachability test uses a 24 bits nonce, it is unlikely that an off-path attacker could send a single packet that causes an ITR to believe that the ETR it is testing is reachable while in reality it is not reachable. To increase the success likelihood of such attack, the attacker should create a massive amount of packets carrying all possible nonce values. However, "flood attack" can be easily detected and blocked.

The second type of packet could be exploited to create a Denial of Service attack against the nonce-based reachability test. Consider a spoofing off-path attacker (SA) that sends a continuous flow of spoofed LISP data encapsulated packets that contain the Nonce-Present and the Echo-Nonce bit and each packet contains a different random nonce. The ETR that receives such packets will continuously change the nonce that it returns to the remote ITR. If the remote ITR starts a nonce-reachability test, this test may fail because the ETR has received a spoofed LISP data encapsulated packet with a different random nonce and never echoes the real nonce. In this case the ITR

will consider the ETR not reachable. The success of this test will of course depend on the ratio between the amount of packets sent by the legitimate ITR and the spoofing off-path attacker (SA).

Packets sent by a non-spoofing off-path attacker (NSA) can cause similar problem if no check is done with the EID-to-RLOC Cache (see [Section 5.3](#) for the EID-to-RLOC Cache check). Otherwise, if the check is performed the packets will be rejected by the ETR that receives them and cannot cause problems.

[5.4.4](#). Attacks using the ID Instance bits

LISP allows to carry in its header a 24-bits value called "Instance ID" and used on the ITR to indicate which private Instance ID has been used for encapsulation, while on the ETR can be used to select the forwarding table used for forwarding the decapsulated packet.

Even if an off-path attacker could randomly guess a valid Instance ID value, there is no LISP specific problem. Obviously the attacker could be now able to reach hosts that are only reachable through the routing table identified by the attacked Instance ID, however, end-system security is out of the scope of this document. Nevertheless, access lists can be configured to protect the network from Instance ID based attacks.

[6](#). Control Plane Threats

In this section, we discuss the different types of attacks that can occur when an off-path attacker sends control plane packets. We focus on the packets that are sent directly to the ETR and do not analyze the particularities of a LISP mapping system. The LISP+ALT and LISP-DDT mapping systems are discussed in [Section 9](#).

[6.1](#). Attacks with Map-Request messages

An off-path attacker could send Map-Request packets to a victim ETR. In theory, a Map-Request packet is only used to solicit an answer and as such it should not lead to security problems. However, the LISP specification [[I-D.ietf-lisp](#)] contains several particularities that could be exploited by an off-path attacker.

The first possible exploitation is the P bit. The P bit is used to probe the reachability of remote ETRs in the control plane. In our reference environment, LR3 could probe the reachability of LR1 by sending a Map-Request with the P bit set. LR1 would reply by sending a Map-Reply message with the P bit set and the same nonce as in the Map-Request message.

A spoofing off-path attacker (SA) could use the P bit to force a victim ETR to send a Map-Reply to the spoofed source address of the Map-Request message. As the Map-Reply can be larger than the Map-Request message, there is a risk of amplification attack.

Considering only IPv6 addresses, a Map-Request can be as small as 40 bytes, considering one single ITR address and no Mapping Protocol Data. The Map-Reply instead has a size of $O(12 + (R * (28 + N * 24)))$ bytes, where N is the maximum number of RLOCs in a mapping and R the maximum number of records in a Map-Reply. Since up to 255 RLOCs can be associated to an EID-Prefix and 255 records can be stored in a Map-Reply, the maximum size of a Map-Reply is thus above 1 MB showing a size factor of up to 39,193 between the message sent by the attacker and the message sent by the ETR. These numbers are however theoretical values not considering transport layer limitations and it is more likely that the reply will contain only one record with at most a dozen of locators, giving an amplification factor around 8.

Any ISP with a large number of potential RLOCs for a given EID-Prefix should carefully ponder the best trade-off between the number of RLOCs through which it wants that the EID is reachable and the consequences that an amplification attack can produce.

It should be noted that the maximum rate of Map-Reply messages should apply to all Map-Replies and also be associated to each destination that receives Map-Reply messages. Otherwise, a possible amplification attack could be launched by a spoofing off-path attacker (SA) as follows. Consider an attacker SA and EID-Prefix 192.0.2.0/24 and a victim ITR. To amplify a Denial of Service attack against the victim ITR, SA could send spoofed Map-Request messages whose source EID addresses are all the addresses inside 192.0.2.0/24 and source RLOC address is the victim ITR. Upon reception of these Map-Request messages, the ETR would send large Map-Reply messages for each of the addresses inside p/P back to the victim ITR.

If a non-spoofing off-path attacker (NSA) sends a Map-Request with the P bit set, it will receive a Map-Reply with the P bit set. This does not raise security issues besides the usual risk of overloading a victim ETR by sending too many Map-Request messages.

The Map-Request message may also contain the SMR bit. Upon reception of a Map-Request message with the SMR bit, an ETR must return to the source of the Map-Request message a Map-Request message to retrieve the corresponding mapping. This raises similar problems as the P bit discussed above except that as the Map-Request messages are smaller than Map-Reply messages, the risk of amplification attacks is reduced. This is not true anymore if the ETR append to the Map-Request messages its own Map-Records. This mechanism is meant to

reduce the delay in mapping distribution since mapping information is provided in the Map-Request message.

Furthermore, appending Map-Records to Map-Request messages represents a major security risk since an off-path attacker could generate a (spoofed or not) Map-Request message and include in the Map-Reply portion of the message mapping for EID prefixes that it does not serve. This could lead to various types of redirection and denial of service attacks. An xTR should not process the Map-Records information that it receives in a Map-Request message.

6.2. Attacks with Map-Reply messages

In this section we analyze the attacks that could occur when an off-path attacker sends directly Map-Reply messages to ETRs without using one of the proposed LISP mapping systems.

There are two different types of Map-Reply messages:

Positive Map-Reply: These messages contain a Map-Record binding an EID-Prefix to one or more RLOCs.

Negative Map-Reply: These messages contain a Map-Record for an EID-Prefix with an empty locator-set and specifying an action, which may be either Drop, Natively forward, or Send Map-Request.

Positive Map-Reply messages are used to map EID-Prefixes onto RLOCs. Negative Map-Reply messages are used to support PTR and interconnect the LISP Internet with the legacy Internet.

Most of the security of the Map-Reply messages depends on the 64 bits nonce that is included in a Map-Request and returned in the Map-Reply. An ETR must never accept a Map-Reply message whose nonce does not match one of the pending Map-Request messages. If an ETR does not accept Map-Reply messages with an invalid nonce, the risk of attack is acceptable given the size of the nonce (64 bits).

Note, however, that the nonce only confirms that the Map-Reply was sent by the ETR that received the Map-Request. It does not validate the content of the Map-Reply message.

In addition, an attacker can perform EID-to-RLOC Cache overflow attack by de-aggregating (i.e., splitting an EID prefix into artificially smaller EID prefixes) either positive or negative mappings.

6.3. Gleaning Attacks

A third type of attack involves the gleaning mechanism proposed in [I-D.ietf-lisp] and discussed in [Saucez09]. In order to reduce the time required to obtain a mapping, [I-D.ietf-lisp] allows an ITR to learn a mapping from the LISP data encapsulated packets and the Map-Request packets that it receives. LISP data encapsulated packet contains a source RLOC, destination RLOC, source EID and destination EID. When an ITR receives a data encapsulated packet coming from a source EID for which it does not already know a mapping, it may insert the mapping between the source RLOC and the source EID in its EID-to-RLOC Cache. Gleaning can also be used when an ITR receives a Map-Request as the Map-Request also contains a source EID address and a source RLOC. Once a gleaned entry has been added to the EID-to-RLOC cache, the LISP ITR sends a Map-Request to retrieve the mapping for the gleaned EID from the mapping system. [I-D.ietf-lisp] recommends storing the gleaned entries for only a few seconds.

The first risk of gleaning is the ability to temporarily hijack an identity. Consider an off-path attacker that wants to temporarily hijack host HA's identity and send packets to host HB with host HA's identity. If the xTRs that serve host HB do not store a mapping for host HA, a non-spoofing off-path attacker (NSA) could send a LISP encapsulated data packet to LR3 or LR4. The ETR will store the gleaned entry and use it to return the packets sent by host HB to the attacker. In parallel, the ETR will send a Map-Request to retrieve the mapping for HA. During a few seconds or until the reception of the Map-Reply, host HB will exchange packets with the attacker that has hijacked HA's identity. Note that the attacker could in parallel send lots of Map-Requests or lots of LISP data encapsulated packets with random sources to force the xTR that is responsible for host HA to send lots of Map-Request messages in order to force it to exceed its rate limit for control plane messages. This could further delay the arrival of the Map-Reply message on the requesting ETR.

Gleaning also introduces the possibility of a man-in-the-middle attack. Consider an off-path attacker that knows that hosts HA and HB that resides in different sites will exchange information at time t. An off-path attacker could use this knowledge to launch a man-in-the-middle attack if the xTRs that serve the two hosts do not have mapping for the other EID. For this, the attacker sends to LR1 (resp. LR3) a LISP data encapsulated packet whose source RLOC is its IP address and contains an IP packet whose source is set to HB (resp. HA). The attacker chooses a packet that will not trigger an answer, for example the last part of a fragmented packet. Upon reception of these packets, LR1 and LR3 install gleaned entries that point to the attacker. As explained above, the attacker could, at the same time, send lots of packets to LR1 and LR3 to force them to exhaust their

control plane rate limit. This will extend the duration of the gleaned entry. If host HA establishes a flow with host HB at that time, the packets that they exchange will first pass through the attacker.

In both cases, the attack only lasts for a few seconds (unless the attacker is able to exhaust the rate limitation). However it should be noted that today a large amount of packets might be exchanged during even a small fraction of time.

7. Threats concerning Interworking

[I-D.ietf-lisp-interworking] defines two network elements to allow LISP and non-LISP sites to communicate, namely the Proxy-ITR and the Proxy-ETR. The Proxy-ITR encapsulates traffic from non-LISP sites in order to forward it toward LISP sites, while the Proxy-ETR decapsulates traffic arriving from LISP sites in order to forward it toward non-LISP sites. For these elements some of the attack based on the LISP specific header are not possible, for the simple reason that some of the fields cannot be used due to the unidirectional nature of the traffic.

The Proxy-ITR has functionalities similar to the ITR, however, its main purpose is to encapsulate packets arriving from the DFZ in order to reach LISP sites. This means that it is not bound to any particular EID-Prefix, hence no mapping exists and no mapping can be configured in the EID-to-RLOC Database. This means that the Proxy-ITR element itself is not able to check whether or not the arriving traffic has the right to be encapsulated or not. To limit such an issue it is recommended to use the current practice based on firewalls and ACLs on the machine running the Proxy-ITR service. On the other side, the Proxy-ITR is meant to encapsulate only packets that are destined to one of the LISP sites it is serving. This is the case for instance for a service provider selling Proxy-ITR services. For this purpose a static EID-to-RLOC Cache can be configured in order to encapsulate only valid packets. In case of a cache-miss no Map-Request needs to be sent and the packet can be silently dropped.

The Proxy-ETR has functionalities similar to the ETR, however, its main purpose is to inject un-encapsulated packet in the DFZ in order to reach non-LISP-Sites. This means that since there is no specific EID-Prefix downstream, it has no EID-to-RLOC Database that can be used to check whether or not the destination EID is part of its domain. In order to avoid for the Proxy-ETR to be used as relay in a DoS attack it is preferable to configure the EID-to-RLOC Cache with static entries used to check if an encapsulated packet coming from a

specific RLOC and having a specific source EID is actually allowed to transit through the Proxy-ETR. This is also important for services provider selling Proxy-ETR service to actually process only packets arriving from its customers. However, in case of cache-miss no Map-Request needs to be sent, rather the packet can be silently dropped since it is not originating from a valid site. The same drop policy should be used for packets with an invalid source RLOC or a valid source RLOC but an invalid EID.

8. Threats with Malicious xTRs

In this section, we discuss the threats that could be caused by malicious xTRs. We consider the reference environment below where EL1 is a malicious or compromised xTR. This malicious xTR serves a set of hosts that includes HC. The other xTRs and hosts in this network play the same role as in the reference environment described in [Section 4](#).




Figure 2: Malicious xTRs' Reference Environment

Malicious xTRs are probably the most serious threat to the LISP control plane from a security viewpoint. To understand the problem, let us consider the following scenario. Host HC and HB exchange packets with host HA. As all these hosts reside in LISP sites, LR1 and LR2 store mappings for HB and HC. Thus, these xTRs may need to exchange LISP control plane packets with EL1, e.g., to perform reachability tests or to refresh expired mappings (e.g., if HC's mapping has a small TTL).

A first threat against the LISP control plane is when EL1 replies to a legitimate Map-Request message sent by LR1 or LR2 with a Map-Reply

message that contains an EID-Prefix that is larger than the prefix owned by the site attached to EL1. For instance if the prefix owned by EL1 is 192.0.2.0/25 but the Map-Reply contain a mapping for 192.0.2.0/24. This could allow EL1 to attract packets destined to other EIDs than the EIDs that are attached to EL1. This attack is called an "overclaiming" attack.

Overclaiming attack can be combined with de-aggregation to succeed a LISP Cache poisoning attack and prefix hijacking. For example, if the EID prefix of the attacker is 192.0.2.0/25, it cannot provide a mapping for the EID prefix 192.0.2.128/25 (i.e., it cannot hijack the prefix). However, since a Map-Reply can contain several map records, it is possible to hijack such a prefix by providing as well a mapping for it. To this end, the attacker can send a Map-Reply with an EID prefix that covers at the same time the requested EID and the hijacked target prefix. Continuing the previous example, if the requested mapping is for EID 192.0.2.1, and the target hijack prefix is 192.0.2.128/25, the Map-Reply will contain a map record for 192.0.2.0/24 and a map record for 192.0.2.128/25. Such a reply is considered legitimate according to the requested EID, while the map record of the hijacked prefix may lead to traffic redirection/disruption and ITR's Cache poisoning.

Another variant of the overclaiming attack is a Denial of Service attack by sending a Negative Map-Reply message for a larger prefix without any locator and with the Drop action. Such a Negative Map-Reply indicates that the ETR that receives it should discard all packets.

The current LISP specification briefly discusses the overclaiming problem [[I-D.ietf-lisp](#)], but does not propose any specific solution to solve the problem. Nevertheless, [[I-D.ietf-lisp-sec](#)] proposes a solution to protect LISP against overclaiming attacks under the assumption that the mapping system can be trusted.

Another concern with malicious xTRs is the possibility of Denial of Service attacks. A first attack is the flooding attack that was described in [[I-D.bagnulo-lisp-threat](#)]. This attack allows a malicious xTR to redirect traffic to a victim. The malicious xTR first defines a mapping for HC with two RLOCs: its own RLOC (EL1) and the RLOC of the victim (e.g., LR3). The victim's RLOC is set as unreachable in the mapping. HC starts a large download from host HA. Once the download starts, the malicious xTR updates its Locator Status Bits, changes the mapping's version number or sets the SMR bit such that LR1 updates its EID-to-RLOC Cache to send all packets destined to HC to the victim's RLOC. Instead of downloading from HA, the attacker could also send packets that trigger a response (e.g., ICMP, TCP SYN, DNS request, ...) to HA. HA would then send its

response and its xTR would forward it to the victim's RLOC.

An important point to note about this flooding attack is that it reveals a limitation of the LISP architecture. A LISP ITR relies on the received mapping and possible reachability information to select the RLOC of the ETR that it uses to reach a given EID or block of EIDs. However, if the ITR made a mistake, e.g., due to misconfiguration, wrong implementation, or other types of errors and has chosen a RLOC that does not serve the destination EID, there is no easy way for the LISP ETR to inform the ITR of its mistake. A possible solution is to enforce an ETR to perform a reachability test with the selected ITR as soon as there is LISP encapsulated traffic between the two.

Note that the attacks discussed in this section are for documentation purpose only. Malicious xTRs are either somehow directly deployed by attackers or the result of attackers gaining privileged access to existing xTRs. As such, it is out of the scope of LISP to propose any mechanism to protect routers or to avoid their deployments with malicious intentions.

9. Security of the Proposed Mapping Systems

9.1. LISP+ALT

One of the assumptions in [[I-D.ietf-lisp](#)] is that the mapping system is more secure than sending Map-Request and Map-Reply messages directly. We analyze this assumption in this section by analyzing the security of the ALT mapping system.

The ALT mapping system is basically a manually configured overlay of GRE tunnels between ALT routers. BGP sessions are established between ALT routers that are connected through such tunnels. An ALT router advertises the EID prefixes that it serves over its BGP sessions with neighboring ALT routers and the EID-Prefixes that it has learned from neighboring ALT routers.

The ALT mapping system is in fact a discovery system that allows any ALT router to discover the ALT router that is responsible for a given EID-Prefix. To obtain a mapping from the ALT system, an ITR sends a packet containing a Map-Request on the overlay. This Map-Request is sent inside a packet whose destination is the requested EID. The Map-Request is routed on the overlay until it reaches the ALT router that advertised initially the prefix that contains the requested EID. This ALT router then replies directly by sending a Map-Reply to the RLOC of the requesting ITR.

The security of the ALT mapping system depends on many factors, including:

- o The security of the intermediate ALT routers.
- o The validity of the BGP advertisements sent on the ALT overlay.

Unfortunately, experience with BGP on the global Internet has shown that BGP is subject to various types of misconfiguration problems and security attacks. The SIDR working group is developing a more secure inter-domain routing architecture to solve this problem ([[RFC6480](#)]).

The security of the intermediate ALT routers is another concern. A malicious intermediate ALT router could manipulate the received BGP advertisements and also answer to received Map-Requests without forwarding them to their final destination on the overlay. This could lead to various types of redirection attacks. Note that in contrast with a regular IP router that could also manipulate in transit packets, when a malicious or compromised ALT router replies to a Map-Request, it can redirect legitimate traffic for a long period of time by sending an invalid Map-Reply message. Thus, the impact of a malicious ALT router could be much more severe than a malicious router in today's Internet.

[9.2.](#) LISP-DDT

The LISP Delegated Database Tree (LISP-DDT) mapping system is proposed as an alternative for LISP+ALT [[I-D.fuller-lisp-ddt](#)]. LISP-DDT is a hierarchical distributed database for EID-to-RLOC mappings. Each DDT node is configured with an EID prefix it is authoritative for, as well as the RLOC addresses and EID prefixes of the LISP-DDT nodes that are authoritative for more specific EID prefix. In LISP-DDT, mappings are retrieved iteratively. A Map Resolver that needs to locate a mapping traverses the tree of DDT nodes contacting them one after another until the leaf of the DDT tree that is authoritative for the longest matching EID prefix for the mapping's EID is reached. The Map Resolver traverses the hierarchy of LISP-DDT nodes by sending Map-Requests, with the LISP-DDT-originated bit set, to LISP-DDT nodes. The Map Resolver first contacts the root of the hierarchy. When a LISP-DDT node receives a Map-Request, it replies to the Map Resolver with a Map-Referral that contains the list of the locators of its children that are authoritative of a prefix that covers the EID in the Map-Request. The Map Resolver then contacts one of these children that will return, at its turn, a Map-Referral. This procedure is iteratively executed until a Map-Referral marked with the done flag is received. The locators that appear in a referral with the done flag are those of the authoritative ETRs for the EID in the Map-Request. At that moment, the Map Resolver falls

back to its normal behavior and sends a Map-Request to the ETR in order for the ITR to obtain the mapping. It is worth to mention that the Map Resolver can cache the referrals to avoid traversing all the whole hierarchy for all mapping retrievals.

The operation in LISP-DDT is different from ALT and thus it does not present the same threats as LISP+ALT. As a first difference, LISP-DDT natively includes security specification providing data origin authentication, data integrity protection and secure EID prefix delegation. Hence, these aspects are no further explored in this document.

However, threats exist for LISP-DDT as well. For instance, a DoS attack can be performed on the mapping infrastructure by asking to retrieve a large amount of mappings at the same time, hence, the importance of carefully dimensioning the topology of the DDT hierarchy.

If an attacker manages to compromise a LISP-DDT node it can send fake referrals to the Map Resolver and then control the mappings delivered to the ITRs. Furthermore, the effects of such an attack can be longer than the attack itself if the Map Resolver caches the referrals.

10. Threats concerning LISP-MS

LISP-MS ([\[I-D.ietf-lisp-ms\]](#)) specifies two network elements, namely the Map Server and the Map Resolver, that are meant to be used by xTRs to access the mapping system. The advantage is clearly the fact that even if the mapping system changes in time xTRs do not need to change anything since they deal only with Map Servers and Map Resolvers. This includes the security aspects, since no change in the local security policies is needed.

10.1. Map Server

Map Server is used to dispatch Map-Request coming from the mapping system to ETRs that are authoritative for the EID in the request. To this end it is necessary that ETRs register their mappings to the Map Server. This allows the Map Server to know toward which ETR to forward Map-Requests and also to announce the EID-prefixes of the registered mappings in the mapping system.

LISP uses a shared key approach in order to protect the Map Server and grant registration rights only to ETRs that have a valid key. Shared key must be used to protect both the registration message and the Map-Notify message when used. The mechanism used to share the

key between a Map Server and an ETRs must be secured to avoid that a malicious nodes catch the key and uses it to send forged Map-Register message to the Map Server. A forged Map-Register message can be use to attract Map-Request and thus provide invalid Map-Replies or the redirect Map-Requests to a target to mount a DoS attack.

More subtle attacks can be carried out only in the case of malicious ETRs. A malicious ETR can register an invalid RLOC to divert Map-Requests to a target ETR and succeed a DoS attack on it. To avoid this kind of attack, the Map Server must check that the registered RLOCs belong to ETRs authoritative for the registered EID prefix. Such check can be done by sending an explicit Map-Request for the EID to the ETRs in the mapping and check that replies with a Map-Reply. If the ETRs return a valid Map-Reply, the RLOC belongs to an authoritative ETR. Note that this does not protect against malicious ETRs that create forged Map-Replies. Stronger techniques for RLOC check are presented in [[I-D.saucez-lisp-mapping-security](#)].

Similarly to the previous case, a malicious ETR can register an invalid EID-prefix to attract Map-Requests or to redirect them to a target to mount a DoS attack. To avoid this kind of attack, the Map Server must check that the prefixes registered by an ETR belong to that ETR. One method could be to manually configure EID-prefix ranges that can be announced by ETRs.

[[I-D.saucez-lisp-mapping-security](#)] present alternative techniques to verify the prefix claimed by an ETR.

[10.2.](#) Map Resolver

Map Resolvers receive Map-Requests, typically from ITRs, and use the mapping system to find a mapping for the EID in the Map-Request. It can work in two modes:

Non-Caching Mode: The resolver just forwards the Map-Request to the mapping system, which will take care of delivering the request to an authoritative ETR. The latter will send back a Map-Reply directly to the ITR that has originally issued the request.

Caching Mode: The resolver will generate a new Map-Request and send it to the mapping system. In this way it will receive the corresponding reply, store a local copy in a cache, and send back a reply to the original requester. Since all requested mappings are locally cached, before actually making a request to the mapping system it performs a lookup in the local cache and in case of an hit, it send back a reply without querying the mapping system.

In its basic mode, i.e., non-caching mode, the Map Resolver does not

keep state, hence, the only direct form of attack is a DoS attack, where an attacker (or a group of attackers) can try to exhaust computational power by flooding the resolver with requests. Common filtering techniques and BCP against DoS attacks can be applied in this case.

Nonetheless, attackers can use resolvers as relay for DoS attacks against xTRs. An off-path spoofing attacker can generate a high load of requests to a set of resolvers, hence distributing the load in order to avoid to be blocked. All this requests can use a specific EID that makes all the requests to be forwarded to a specific ETR, which, as a result, will be victim of a DDoS attack. Similarly, the attacker can use a spoofed source address making all the replies to converge to one single ITR, which, as a result, will be victim of a DDoS attack. Such scenarios are not specific to LISP, but rather a common problem of every query infrastructure, hence the same BCP can be applied in order to limit the attacks.

When functioning in caching-mode, the resolver will use the same type of cache than ITRs. Due to its similarity with the ITRs' cache the analysis provided in [Section 5.2](#) holds also in this case. However, an important difference exists: this cache is not used for packet encapsulation but only for quick replies when new requests arrive. Therefore, as the caching-mode is only an optimization, the attacks that aim at filling the Map Resolver cache have a less severe impact on the traffic.

When Map Resolvers are used as front-end of the LIS-DDT mapping system they may be exposed to another variant of DoS. Indeed, the iterative operation of the Map Resolver on the DDT hierarchy implies that it has to maintain state about the ITR that requested the mapping, this in order to send the final Map-Request to the ETR on behalf of the ITR. An attacker might leverage on this to fill the Map Resolver memory and then cause a DoS.

The question may arise on whether a Kaminsky-like attack is possible for an off-path attacker against ITRs sending requests to a certain resolver. The 64-bits nonce present in every message has been introduced in the LISP specification to avoid such kind of attack. There has been discussion within the LISP Working Group on the optimal size of the nonce, and it seems that 64-bits provides sufficient protection.

A possible way to limit the above-described attacks is to introduce strong identification in the Map-Request/Map-Reply by using the Encapsulated Control Message with authentication enabled.

11. Suggested Recommendations

To mitigate the impact of attacks against LISP, the following recommendations should be followed.

First, the use of some form of filtering can help in avoid or at least mitigate some types of attacks.

- o On ETRs, packets should be decapsulated only if the destination EID is effectively part of the EID-Prefix downstream the ETR. Further, still on ETRs, packets should be decapsulated only if a mapping for the source EID is present in the EID-to-RLLOC Cache and has been obtained through the mapping system (not gleaned).
- o On ITRs, packets should be encapsulated only if the source EID is effectively part of the EID-Prefix downstream the ITR. Further, still on ITRs, packets should be encapsulated only if a mapping obtained from the mapping system is present in the EID-to-RLLOC Cache (no Data-Probing).

Note that this filtering, since complete mappings need to be installed in both ITRs and ETRs, can introduce higher connection setup latency and hence potentially more packets drops due to the lack of mappings in the EID-to-RLLOC Cache.

While the gleaning mechanism allows starting encapsulating packets to a certain EID in parallel with the Map-Request to obtain a mapping when a new flow is established, it creates important security risks since it allows attackers to perform identity hijacks. Although the duration of these identity hijacks is limited (except the case of rate limitation exhaustion), their impact can be severe. A first option would be to disable gleaning until the security concerns are solved. A second option would be to strictly limit the number of packets that can be forwarded via a gleaned entry. Overall the benefits of gleaning, i.e., avoiding the loss of the first packet of a flow, seems very small compared to the associated security risks. Furthermore, measurements performed in data centers show that today's Internet often operate with packet loss ratio of 1 or 2 percentage ([[Chu](#)]). These packet loss ratios are probably already orders of magnitude larger than the improvement provided by the gleaning mechanism.

With the increasing deployment of spoofing prevention techniques such as [[RFC3704](#)] or SAVI [[SAVI](#)], it can be expected that attackers will become less capable of sending packets with a spoofed source address. To prevent packet injection attacks from non-spoofing attackers (NSA), ETRs should always verify that the source RLLOC of each received LISP data encapsulated packet corresponds to one of the

RLOCs listed in the mappings for the source EID found in the inner packet. An alternative could be to use existing IPSec techniques [[RFC4301](#)] and when necessary including perhaps [[RFC5386](#)] to establish an authenticated tunnel between the ITR and the ETR.

[I-D.ietf-lisp] recommends to rate limit the control messages that are sent by an xTR. This limit is important to deal with denial of service attacks. However, a strict limit, e.g., implemented with a token bucket, on all the Map-Request and Map-Reply messages sent by an xTR is not sufficient. An xTR should distinguish between different types of control plane packets:

1. The Map-Request messages that it sends to refresh expired mapping information.
2. The Map-Request messages that it sends to obtain mapping information because one of the served hosts tried to contact an external EID.
3. The Map-Request messages that it sends as reachability probes.
4. The Map-Reply messages that it sends as response to reachability probes.
5. The Map-Request messages that it sends to support gleaning.

These control plane messages are used for different purposes. Fixing a global rate limit for all control plane messages increases the risk of Denial of Service attacks if a single type of control plane message can exceed the configured limit. This risk could be mitigated by either specifying a rate for each of the five types of control plane messages. Another option could be to define a maximum rate for all control plane messages, and prioritize the control plane messages according to the list above (with the highest priority for message type 1).

In [[I-D.ietf-lisp](#)], there is no mechanism that allows an xTR to verify the validity of the content a Map-Reply message that it receives. Besides the attacks discussed earlier in the document, a time-shifted attack where an attacker is able to modify the content of a Map-Reply message but then needs to move off-path could also create redirection attacks. The nonce only allows an xTR to verify that a Map-Reply responds to a previously sent Map-Request message. In order to allow verifying the validity and integrity of bindings between EID-Prefixes and their RLOCs solutions proposed in [[I-D.saucez-lisp-mapping-security](#)] and [[I-D.ietf-lisp-sec](#)] should be deployed. Having such kind of mechanisms would allow ITRs to ignore non-verified mappings, thus increasing security.

Finally, there is also the risk of Denial of Service attack against the EID-to-RLOC Cache. We have discussed these attacks when considering external attackers with, e.g., the gleaning mechanism and in [Section 5.2](#). If an ITR has a limited EID-to-RLOC Cache, a malicious or compromised host residing in the site that it serves could generate packets to random destinations to force the ITR to issue a large number of Map-Requests whose answers could fill its cache. Faced with such misbehaving hosts, LISP ITR should be able to limit the percent of Map-Requests that it sends for a given source EID.

In order to mitigate flooding attacks it would be worth consider developing secure mechanisms to allow an ETR to indicate to an ITR that it does not serve a particular EID or block of EIDs.

[12.](#) IANA Considerations

This document makes no request to IANA.

[13.](#) Security Considerations

Security considerations are the core of this document and do not need to be further discussed in this section.

[14.](#) Acknowledgments

This document builds upon the draft of Marcelo Bagnulo ([\[I-D.bagnulo-lisp-threat\]](#)), where the flooding attack and the reference environment were first described.

We would like to thank Jeff Wheeler for his comments.

This work has been partially supported by the INFSO-ICT-216372 TRILOGY Project (www.trilogy-project.org).

[15.](#) References

[15.1.](#) Normative References

[I-D.fuller-lisp-ddt]
Fuller, V., Lewis, D., Ermagan, V., and A. Jain, "LISP Delegated Database Tree", [draft-fuller-lisp-ddt-04](#) (work in progress), September 2012.

[I-D.ietf-lisp]

Farinacci, D., Fuller, V., Meyer, D., and D. Lewis,
"Locator/ID Separation Protocol (LISP)",
[draft-ietf-lisp-23](#) (work in progress), May 2012.

[I-D.ietf-lisp-alt]

Fuller, V., Farinacci, D., Meyer, D., and D. Lewis, "LISP
Alternative Topology (LISP+ALT)", [draft-ietf-lisp-alt-10](#)
(work in progress), December 2011.

[I-D.ietf-lisp-interworking]

Lewis, D., Meyer, D., Farinacci, D., and V. Fuller,
"Interworking LISP with IPv4 and IPv6",
[draft-ietf-lisp-interworking-06](#) (work in progress),
March 2012.

[I-D.ietf-lisp-map-versioning]

Iannone, L., Saucez, D., and O. Bonaventure, "LISP Map-
Versioning", [draft-ietf-lisp-map-versioning-09](#) (work in
progress), March 2012.

[I-D.ietf-lisp-ms]

Fuller, V. and D. Farinacci, "LISP Map Server Interface",
[draft-ietf-lisp-ms-16](#) (work in progress), March 2012.

15.2. Informative References

[Chu] Jerry Chu, H., "Tuning TCP Parameters for the 21st
Century", 75th IETF, Stockholm, July 2009,
<<http://tools.ietf.org/wg/savi/>>.

[I-D.bagnulo-lisp-threat]

Bagnulo, M., "Preliminary LISP Threat Analysis",
[draft-bagnulo-lisp-threat-01](#) (work in progress),
July 2007.

[I-D.ietf-lisp-sec]

Maino, F., Ermagan, V., Cabellos-Aparicio, A., Saucez, D.,
and O. Bonaventure, "LISP-Security (LISP-SEC)",
[draft-ietf-lisp-sec-04](#) (work in progress), October 2012.

[I-D.ietf-tcpm-tcp-security]

Gont, F., "Survey of Security Hardening Methods for
Transmission Control Protocol (TCP) Implementations",
[draft-ietf-tcpm-tcp-security-03](#) (work in progress),
March 2012.

[I-D.lear-lisp-nerd]

Lear, E., "NERD: A Not-so-novel EID to RLOC Database", [draft-lear-lisp-nerd-09](#) (work in progress), April 2012.

[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.saucez-lisp-mapping-security]

Saucez, D. and O. Bonaventure, "Securing LISP Mapping replies", [draft-saucez-lisp-mapping-security-00](#) (work in progress), February 2011.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", [BCP 84](#), [RFC 3704](#), March 2004.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", [RFC 4301](#), December 2005.

[RFC5386] Williams, N. and M. Richardson, "Better-Than-Nothing Security: An Unauthenticated Mode of IPsec", [RFC 5386](#), November 2008.

[RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", [RFC 6480](#), February 2012.

[SAVI] IETF, "Source Address Validation Improvements Working Group", <<http://tools.ietf.org/wg/savi/>>.

[Saucez09]

Saucez, D. and L. Iannone, "How to mitigate the effect of scans on mapping systems", Submitted to the Trilogy Summer School on Future Internet.

[Appendix A](#). Document Change Log

o Version 03 Posted October 2012.

- * Dropped Reference to [RFC 2119](#) notation because it is not actually used in the document.
- * Deleted future plans section.
- * Updated References
- * Deleted/Modified sentences referring to the early status of the LISP WG and documents at the time of writing early versions of

- the document.
- * Further editorial polishing.
- * Fixed all ID nits.
- o Version 02 Posted September 2012.
 - * Added a new attack that combines overclaiming and de-aggregation (see [Section 6.2](#)).
 - * Editorial polishing.
- o Version 01 Posted February 2012.
 - * Added discussion on LISP-DDT in [Section 9.2](#).
- o Version 00 Posted July 2011.
 - * Added discussion on LISP-MS in [Section 10](#).
 - * Added discussion on Instance ID in [Section 5.4](#).
 - * Editorial polishing of the whole document.
 - * Added "Change Log" appendix to keep track of main changes.
 - * Renamed "[draft-saucez-lisp-security-03.txt](#)".

Authors' Addresses

Damien Saucez
INRIA
2004 route des Lucioles BP 93
06902 Sophia Antipolis Cedex
France

Email: damien.saucez@inria.fr

Luigi Iannone
Telecom ParisTech
23, Avenue d'Italie, CS 51327
75214 PARIS Cedex 13
France

Email: luigi.iannone@telecom-paristech.fr

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

Email: olivier.bonaventure@uclouvain.be

