Network Working Group
Internet-Draft
Intended status: Informational
Expires: October 10, 2014

D. Saucez
INRIA
L. Iannone
Telecom ParisTech
O. Bonaventure
Universite catholique de Louvain

April 8, 2014

LISP Threats Analysis
draft-ietf-lisp-threats-09.txt

Abstract

This document proposes a threat analysis of the Locator/Identifier Separation Protocol (LISP) if deployed in the Internet.

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 October 10, 2014.

Table of Contents

1. Introduction

The Locator/ID Separation Protocol (LISP) is specified in [RFC6830]. The present document assess the potential security threats identified in the LISP specifications if LISP is deployed in the Internet (i.e., a public non-trustable environment).

The document is composed of two main parts: the first discussing the techniques that can be used by attackers to succeed attacks based on the LISP protocol and architecture; the second discussing the main categories of attacks and how to construct them.

This document does not consider all the possible uses of LISP as discussed in [RFC6830] and [I-D.ietf-lisp-deployment]. The document focuses on LISP unicast, including as well LISP Interworking [RFC6832], LISP-MS [RFC6833], and LISP Map-Versioning [RFC6834]. The reading of these documents is a prerequisite for understanding the present document.

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

2. 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. As with IP, LISP relies on higher layer cryptography to secure packet payloads from on path attacks, so this document does not consider on-path attackers.

Similarly, 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. This documents does not consider time-shifted attacks.

3. Off-Path Attackers: Reference Environment

The reference environment shown in the figure below is considered throughout this document. 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. PxTR is a proxy xTR and MR/MS plays the roles of Map Server and/or Map Resolver.

```
             +-----+
             | HA  |
             +-----+
                | EID: HA
                |
             -----------------
                |          |
             +-----+    +-----+
             | LR1 |    | LR2 |
             +-----+    +-----+
                |          |
                |          |
             +-----+    +-----+
             |ISP1 |    |ISP2 |
             +-----+    +-----+
                |          |
+------+     +----------------+     +-----+
| PxTR |-----|                |-----| SA  |
+------+     |                |     +-----+
             |    Internet    |
+-------+    |                |     +-----+
| MR/MS |----|                |-----| NSA |
+-------+    +----------------+     +-----+
                |          |
             +-----+    +-----+
             | LR3 |    | LR4 |
             +-----+    +-----+
                |          |
             -------------------
                           |
                           | EID: HB
                        +-----+
                        | HB  |
                        +-----+
```

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. To perform some of the attacks described in this document SA needs to be in a non-LISP site.

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 could 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 could 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.

4. Attack vectors

This section presents techniques that can be used by attackers to succeed attacks leveraging the LISP protocol and architecture. This section focuses on the techniques while Section 5 presents the attacks that can be succeeded while using these techniques.

4.1. Configured EID-to-RLOC mappings

Each xTR maintains a set of configured mappings related to the EID-Prefixes that are "behind" the xTR [RFC6830]. Where "behind" means that at least one of the xTR's globally visible IP addresses is a RLOC for those EID-Prefixes.

As these mappings are 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.

4.2. EID-to-RLOC Cache

The EID-to-RLOC Cache (also called the Map-Cache) is the data structure that stores a copy of the mappings retrieved from a remote ETR's mapping via the LISP control-plane. Attacks against this data structure could happen either when the mappings are first installed in the cache or by corrupting (poisoning) the mappings already present in the cache.

Cache poisoning attacks are used to alter (any combination of) the following parts of the mappings installed in the EID-to-RLOC Cache:

- EID prefix
- RLOC list
- RLOC priority
- RLOC weight
- RLOC reachability
- Mapping TTL
- Mapping version
- Mapping Instance ID

Cache poisoning attacks can be performed by attackers from the outside of the attacked LISP network but also directly from the inside. As a matter of fact, end-hosts behind an ITR can use the data-plane to overflow the ITR's EID-to-RLOC Cache by sending packets to non-popular EID prefixes (similar to scan attack but with a different goal). In such a scenario the ITR may evict popular (in-use) entries from the map-cache disrupting the normal operation of

the network by forcing cache miss [Florin13].

4.3. Attacks using the data-plane

The data-plane is constituted of the operations of encapsulation, decapsulation, and forwarding as well as the content of the EID-to-RLOC Cache and configured EID-to-RLOC mappings as specified in the original LISP document ([RFC6830]).

4.3.1. Attacks not leveraging on the LISP header

An attacker can inject packets into flows without using the LISP header, i.e., with the N, L, E, V, and I bits ([RFC6830]).

Taking notation of the reference environment (Figure 1), to inject a packet in the HA->HB flow, a spoofing off-path attacker (SA) could 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. A non-spoofing off-path attacker (NSA) could only send a packet whose source address is set to its assigned IP address. The destination address of the encapsulated packet could be LR3 or LR4.

4.3.1.1. Gleaning Attacks

In order to reduce the time required to obtain a mapping, [RFC6830] proposes the gleaning mechanism that 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 could 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. [RFC6830] recommends storing the gleaned entries for only a few seconds.

An attacker can send LISP encapsulated packets to host HB with host HA's EID and if the xTRs that serve host HB do not store a mapping for host HA at that time. The xTR will store the gleaned entry and use it to return the packets sent by host HB. In parallel, the ETR will send a Map-Request to retrieve the mapping for HA but until the reception of the Map-Reply, host HB will exchange packets with the attacker instead of HA.

Similarly, if an off-path attacker knows that hosts HA and HB that resides in different sites will exchange information at a given time the attacker could send 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. If host HA is willing to establish a flow with host HB at that time, the packets that they exchange will pass through the attacker as long as the gleaned entry is active on the xTRs.

By itself, an attack made solely using gleaning cannot last long, however it should be noted that with current network capacities, a large amount of packets might be exchanged during even a small fraction of time.

4.3.1.2. Threats concerning Interworking

[RFC6832] defines Proxy-ITR and Proxy-ETR network elements to allow LISP and non-LISP sites to communicate. The Proxy-ITR has functionality similar to the ITR, however, its main purpose is to encapsulate packets arriving from the DFZ in order to reach LISP sites. A Proxy-ETR has functionality similar to the ETR, however, its main purpose is to inject de-encapsulated packets in the DFZ in order to reach non-LISP Sites from LISP sites. As a PITR (resp. PETR) is a particular case of ITR (resp. ETR), it is subject to same attacks than ITRs (resp. ETR).

PxTRs can be targeted by attacks aiming to influence traffic between LISP and non-LISP sites but also to launch relay attacks.

It is worth to notice that when PITR and PETR functions are separated, attacks targeting nodes that collocate PITR and PETR functionality are ineffective.

4.3.2. Attacks leveraging on the LISP header

The main LISP document [RFC6830] 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.

4.3.2.1. 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 [RFC6830].

An attacker can send a data packet with the L bit set to 1 and some or all Locator Status Bits set to zero. Therefore, by blindly trusting the Locator Status Bits communication going on can be altered or forced to go through a particular set of locators.

4.3.2.2. Attacks using the Map-Version bit

The optional 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:

- It compares the Destination Map-Version found in the header with the current version of its own configured EID-to-RLOC mapping, 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 [RFC6830] and send a Map-Request with the SMR bit set.

- If a mapping exists in the EID-to-RLOC 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.

An off-path attacker could use the Map-Version bit to force an ETR to send Map-Request messages. The attacker could retrieve the current source and destination Map-Version for both HA and HB. Based on this information, it could 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 4.4.1) so that more bandwidth is consumed on the target (because of the larger packets) than the bandwidth necessary at the attacker side.

4.3.2.3. 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.

The second type of packet could be exploited to attack 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 depends on the ratio between the amount of packets sent by the legitimate ITR and the spoofing off-path attacker (SA).

4.3.2.4. Attacks using the Instance ID bits

LISP allows to carry in its header a 24-bits value called "Instance ID" and used on the ITR to indicate which local 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.

The Instance ID increases exposure to attacks ([RFC6169]) as if an

off-path attacker can randomly guess a valid Instance ID value to get access to network that might not been accessible in normal conditions. However, such attacks target end-systems, which is out of the scope of this document.

4.4. Attacks using the control-plane

In this section, we discuss the different types of attacks that could 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 the different LISP indexing sub-system.

4.4.1. 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 [RFC6830] contains several particularities that could be exploited by an off-path attacker.

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

A spoofing off-path attacker (SA) could use the RLOC record 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. Map-Requests are usually smaller than a hundred bytes while the maximum size of a Map-Reply (without considering any MTU constrain) can be above 1 MB, largely bigger than the message sent by the attacker. 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, limiting so the amplification factor.

Similarly, if a non-spoofing off-path attacker (NSA) sends a Map-Request with the RLOC record P bit set, it will receive a Map-Reply with the RLOC record P bit set.

An 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, 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 back to the victim ITR.

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. Note that according to [RFC6830] the SMR-triggered Map-Request might be sent through the mapping system, depending on the number of RLOCs in the locators set. This raises similar problems as the RLOC record 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 allows an off-path attacker to 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.

Moreover, attackers can use Map Resolver and/or Map Server network elements to perform relay attacks. Indeed, on the one hand, a Map Resolver is used to dispatch Map-Request to the mapping system and, on the other hand, a Map Server is used to dispatch Map-Requests coming from the mapping system to ETRs that are authoritative for the EID in the Map-Request.

4.4.2. 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 indicate non-LISP prefixes. ITRs can, if needed, be configured to send all traffic destined for

non-LISP prefixes to a Proxy-ETR.

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. 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). However, the nonce only confirms that the Map-Reply received was sent in response to a Map-Request sent, it does not validate the contents of that Map-Reply.

In addition, an attacker could 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.

In presence of malicious ETRs, overclaiming attacks are possible. Such an attack happens when and ETR replies to a legitimate Map-Request message it received with a Map-Reply message that contains an EID-Prefix that is larger than the prefix owned by the site that encompasses the EID of the Map-Request. For instance if the prefix owned by the site is 192.0.2.0/25 but the Map-Reply contains a mapping for 192.0.2.0/24, then the mapping will influence packets destined to other EIDs than the one the LISP site has authority on.

A malicious ETR might also fragment its configured EID-to-RLOC mappings so that ITR's might have to install much more mappings than really necessary. This attack is called de-aggregation attack.

4.4.3. Attacks with Map-Register messages

Map-Register messages are sent by ETRs to indicate to the mapping system the EID prefixes associated to them. The Map-Register message provides an EID prefix and the list of ETRs that are able to provide Map-Replies for the EID covered by the EID prefix.

As Map-Register messages are protected by an authentication mechanism, only a compromised ETR can register itself to its allocated Map Server.

A compromised ETR can perform an overclaiming attack in order to influence the route followed by Map-Requests for EIDs outside the scope of its legitimate EID prefix.

A compromised ETR can also perform a deaggregation attack in order to register more EID prefixes than necessary to its Map Servers.

Similarly, a compromised Map Server can accept invalid registration or advertise invalid EID prefix to the indexing sub-system.

4.4.4. Attacks with Map-Notify messages

Map-Notify messages are sent by a Map Server to an ETR to acknowledge the good reception and processing of a Map-Register message.

A compromised ETR using EID that it is not authoritative for can send a Map-Register with the M-bit set and a spoofed source address to force the Map Server to send a Map-Notify message to the spoofed address and then succeed a relay attack. Similarly to the pair Map-Request/Map-Reply, the pair Map-Register/Map-Notify is protected by a nonce making it hard for an attacker to inject a falsified notification to an ETR to make this ETR believe that the registration succeeded while it has not.

5. Attack categories

5.1. Intrusion

5.1.1. Description

With an intrusion attack an attacker gains remote access to some resources (e.g., a host, a router, or a network) that are normally denied to her.

5.1.2. Vectors

Intrusion attacks can be mounted using:

- Spoofing EID or RLOCs

- Instance ID bits

5.2. Denial of Service (DoS)

5.2.1. Description

A Denial of Service (DoS) attack aims at disrupting a specific targeted service either by exhausting the resources of the victim up to the point that it is not able to provide a reliable service to legit traffic and/or systems or by exploiting vulnerabilities to make the targeted service unable to operate properly.

5.2.2. Vectors

Denial of Service attacks can be mounted using

- o Gleaning
- o Interworking
- o Locator Status Bits
- o Map-Version bit
- o Nonce-Present and Echo-Nonce bits
- o Map-Request message
- o Map-Reply message
- o Map-Register message
- o Map-Notify message

5.3. Subversion

5.3.1. Description

With subversion an attacker can gain access (e.g., using eavesdropping or impersonation) to restricted or sensitive information such as passwords, session tokens, or any other confidential information. This type of attack is usually carried out in a way such that the target does not even notice the attack. When the attacker is positioned on the path of the target traffic, it is called a Man-in-the-Middle attack. However, this is not a requirement to carry out and eavesdropping attack. Indeed the attacker might be able, for instance through an intrusion attack on a weaker system, either to duplicate or even re-direct the traffic, in both cases having access to the raw packets.

5.3.2. Vectors

Subversion attacks can be mounted using

- o Gleaning
- o Locator Status Bits
- o Nonce-Present and the Echo-Nonce bits
- o Map-Request messages
- o Map-Reply messages

6. Note on Privacy

As presented by [RFC6973], universal privacy considerations are impossible to establish as the privacy definition may vary from one to another. As a consequence, this document does not aim at identifying privacy issues related to the LISP protocol but it is necessary to highlight that security threats identified in this document could play a role in privacy threats as defined in section 5 of [RFC6973].

7. IANA Considerations

This document makes no request to IANA.

8. Security Considerations

This document is devoted to threat analysis of the Locator/Identifier Separation Protocol and is then a piece of choice to understand the security risks at stake while deploying LISP in non-trustable environment.

The purpose of this document is not to provide recommendations to protect against attacks, however most of threats can be prevented with careful deployment and configuration (e.g., filter) and also by applying the general rules in security that consist in activating only features that are necessary in the deployment and verifying the validity of the information obtained from third parties. More detailed recommendations are given in [Saucez13].

The control-plane is probably the most critical part of LISP from a security viewpoint and it is worth to notice that the specifications already offer authentication mechanism for Map-Register messages ([RFC6833]) and that [I-D.ietf-lisp-sec] and [I-D.ietf-lisp-ddt] are clearly going in the direction of a secure control-plane.

9. 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.

The authors would like to thank Ronald Bonica, Albert Cabellos, Noel Chiappa, Florin Coras, Vina Ermagan, Dino Farinacci, Stephen Farrell, Joel Halpern, Emily Hiltzik, Darrel Lewis, Edward Lopez, Fabio Maino, Terry Manderson, and Jeff Wheeler for their comments.

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

The work of Luigi Iannone has been partially supported by the ANR-13-INFR-0009 LISP-Lab Project (www.lisp-lab.org) and the EIT KIC ICT-Labs SOFNETS Project.

10. References

10.1. Normative References

[RFC6169] Krishnan, S., Thaler, D., and J. Hoagland, "Security Concerns with IP Tunneling", RFC 6169, April 2011.

[RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, January 2013.

[RFC6832] Lewis, D., Meyer, D., Farinacci, D., and V. Fuller, "Interworking between Locator/ID Separation Protocol (LISP) and Non-LISP Sites", RFC 6832, January 2013.

[RFC6833] Fuller, V. and D. Farinacci, "Locator/ID Separation Protocol (LISP) Map-Server Interface", RFC 6833, January 2013.

[RFC6834] Iannone, L., Saucez, D., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Map-Versioning", RFC 6834, January 2013.

[RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, July 2013.

10.2. Informative References

[Florin13]
Coras, F., Domingo-Pascual, J., Lewis, D., and A. Cabellos-Aparicio, "An Analytical Model for Loc/ID Mappings Caches", Technical Report. arXiv:1312.1378v2, 2013, <http://arxiv.org/pdf/1312.1378v2.pdf>.

[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-ddt]
Fuller, V., Lewis, D., Ermagan, V., and A. Jain, "LISP Delegated Database Tree", draft-ietf-lisp-ddt-01 (work in progress), March 2013.

[I-D.ietf-lisp-deployment]
Jakab, L., Cabellos-Aparicio, A., Coras, F., Domingo-Pascual, J., and D. Lewis, "LISP Network Element Deployment Considerations", draft-ietf-lisp-deployment-12 (work in progress), January 2014.

[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-05 (work in progress), October 2013.

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

[Saucez13]
Saucez, D., Iannone, L., and O. Bonaventure, "The Map-and-Encap Locator/Identifier separation paradigm: a Security Analysis", Solutions for Sustaining Scalability in Internet Growth, IGI Global, 2013.

Appendix A. Document Change Log

- o Version 09 Posted March 2014.
  - * Updated document according to the review of A. Cabellos.
- o Version 08 Posted October 2013.
  - * Addition of a privacy consideration note.
  - * Editorial changes
- o Version 07 Posted October 2013.
  - * This version is updated according to the thorough review made during October 2013 LISP WG interim meeting.
  - * Brief recommendations put in the security consideration section.
  - * Editorial changes

- o Version 06 Posted October 2013.
  - * Complete restructuration, temporary version to be used at October 2013 interim meeting.
- o Version 05 Posted August 2013.
  - * Removal of severity levels to become a short recommendation to reduce the risk of the discussed threat.
- o Version 04 Posted February 2013.
  - * Clear statement that the document compares threats of public LISP deployments with threats in the current Internet architecture.
  - * Addition of a severity level discussion at the end of each section.
  - * Addressed comments from V. Ermagan and D. Lewis' reviews.
  - * Updated References.
  - * Further editorial polishing.
- 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 4.4.2).

* Editorial polishing.

o Version 01 Posted February 2012.

* Added discussion on LISP-DDT.

o Version 00 Posted July 2011.

* Added discussion on LISP-MS>.

* Added discussion on Instance ID in Section 4.3.2.

* 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