                      LISP Threats Analysis
                    draft-ietf-lisp-threats-13.txt

Abstract

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

Status of this Memo

Copyright Notice

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 three main parts: the first defines the
   general threat model that attackers can follow to mount attacks. The
   second describes the techniques based on the LISP protocol and
   architecture that attackers can use to construct attacks. The third
   discusses mitigation techniques and general solutions to protect the
   LISP protocol and architecture from attacks.

   This document does not consider all the possible uses of LISP as
   discussed in [RFC6830] and [RFC7215]. 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. Threat model

   This document assumes that attackers can be located anywhere in the
   Internet (either in LISP sites or outside LISP sites) and that
   attacks can be mounted either by a single attacker or by the
   collusion of several attackers.

   An attacker is a malicious entity that performs the action of
   attacking a target in a network where LISP is (partially) deployed by
   leveraging the LISP protocol and/or architecture.

   An attack is the action of performing an illegitimate action on a
   target in a network where LISP is (partially) deployed.

   The target of an attack is the entity (i.e., a device connected to
   the network or a network) that is aimed to undergo the consequences
   of an attack. Other entities can potentially undergo side effects of
   an attack, even though they are not directly targeted by the attack.
   The target of an attack can be selected specifically, i.e., a
   particular entity, or arbitrarily, i.e., any entity. Finally, an
   attacker can aim at attacking one or several targets with a single
   attack.

   Section 2.1 specifies the different modes of operation that attackers

can follow to mount attacks and [Section 2.2](#) specifies the different
categories of attacks that attackers can build.

2.1. Attacker's Operation Modes

Attackers can be classified according to the following four modes of
operation, i.e., the temporal and spacial diversity of the attacker.

2.1.1. On-path vs. Off-path Attackers

On-path attackers, also known as Men-in-the-Middle, are able to
intercept and modify packets between legitimate communicating
entities. On-path attackers are located either directly on the
normal communication path (either by gaining access to a node on the
path or by placing themselves directly on the path) or outside the
location path but manage to deviate (or gain a copy of) packets sent
between the communication entities. On-path attackers hence mount
their attacks by modifying packets initially sent legitimately
between communication entities.

An attacker is called off-path attacker if it does not have access to
packets exchanged during the communication or if there is no
communication. In order for their attacks to succeed, off-path
attackers must hence generate packets and inject them in the network.

2.1.2. Internal vs. External Attackers

An internal attacker launches its attack from a node located within a
legitimate LISP site. Such an attacker is either a legitimate node
of the site or it exploits a vulnerability to gain access to a
legitimate node in the site. Because of their location, internal
attackers are trusted by the site they are in.

On the contrary, an external attacker launches its attacks from the
outside of a legitimate LISP site.

2.1.3. Live vs. Time-shifted attackers

A live attacker mounts attacks for which it must remain connected as
long as the attack is mounted. In other words, the attacker must
remain active for the whole duration of the attack. Consequently,
the attack ends as soon as the attacker (or the used attack vector)
is neutralized.

On the contrary, a time-shifted attacker mounts attacks that remain
active after it disconnects from the Internet.

2.1.4. Control-plane vs. Data-plane attackers

   A control-plane attacker mounts its attack by using control-plane
   functionalities, typically the mapping system.

   A data-plane attacker mounts its attack by using data-plane
   functionalities.

   As there is no complete isolation between the control-plane and the
   data-plane, an attacker can operate in the control-plane (resp. data-
   plane) to mount attacks targeting the data-plane (resp. control-
   plane) or keep the attacked and targeted planes at the same layer
   (i.e., from control-plane to control-plane or from data-plane to
   data-plane).

2.1.5. Cross mode attackers

   The attacker modes of operation are not mutually exclusive and hence
   attackers can combine them to mount attacks.

   For example, an attacker can launch an attack using the control-plane
   directly from within a LISP site to which it got temporary access
   (i.e., internal + control-plane attacker) to create a vulnerability
   on its target and later on (i.e., time-shifted + external attacker)
   mount an attack on the data plane (i.e., data-plane attacker) that
   leverages the vulnerability.

2.2. Threat categories

   Attacks can be classified according to the nine following categories.

2.2.1. Replay attack

   A replay attack happens when an attacker retransmits at a later time,
   and without modifying it, a packet (or a sequence of packets) that
   has already been transmitted.

2.2.2. Packet manipulation

   A packet manipulation attack happens when an attacker receives a
   packet, modifies the packet (i.e., changes some information contained
   in the packet) and finally transmits the packet to its final
   destination that can be the initial destination of the packet or a
   different one.

2.2.3. Packet interception and suppression

   In a packet interception and suppression attack, the attacker
   captures the packet and drops it before it can reach its final
   destination.

2.2.4. Spoofing

   With a spoofing attack, the attacker injects packets in the network
   pretending being another node. Spoofing attacks are made by forging
   source addresses in packets.

   It should be noted that with LISP, packet spoofing is similar to any
   other existing tunneling technology currently deployed in the
   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. Hence, 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.

   Inner address spoofing: the attacker uses encapsulation and uses a
         spoofed source address in the inner packet. In case of data-
         plane LISP encapsulation, that corresponds to spoof the source
         EID address of the encapsulated packet.

   Outer address spoofing: the attacker does not use encapsulation and
         spoofs the source address of the packet. In case of data-plane
         LISP encapsulation, that corresponds to spoof the source RLOC
         address of the encapsulated packet.

   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. For example, an attacker outside a LISP
   site can generate a packet with a forged source IP address (i.e.,
   outer address spoofing) and forward it to a LISP destination. The
   packet is then eventually encapsulated by a PITR so that once
   encapsulated the attack corresponds to a inner address spoofing. One
   can also imagine an attacker forging a packet with encapsulation
   where both inner an outer source addresses are spoofed.

   It is important to notice that the combination of inner and outer
   spoofing makes the identification of the attacker complex as the
   packet may not contain information that allows to detect the origin
   of the attack.

2.2.5. Rogue attack

   In a rogue attack the attacker manages to appear as a legitimate
   source of information, without faking its identity (as opposed to a
   spoofing attacker).

2.2.6. Denial of Service (DoS) attack

   A Denial of Service (DoS) attack aims at disrupting a specific
   targeted service to make it unable to operate properly.

2.2.7. Performance attack

   A performance attacks aims at exploiting computational resources
   (e.g., memory, processor) of a targeted node so to make it unable to
   operate properly.

2.2.8. Intrusion attack

   In an intrusion attack the attacker gains remote access to a resource
   (e.g., a host, a router, or a network) or information that it
   normally doesn't have access to. Intrusion attacks can lead to
   privacy leakages.

2.2.9. Amplification attack

   In an amplification attack, the traffic generated by the target of
   the attack in response to the attack is larger than the traffic that
   the attacker must generate.

   In some cases, the data-plane can be several order of magnitude
   faster than the control-plane at processing packets. This difference
   can be exploited to overload the control-plane via the data-plane
   without overloading the data-plane.

2.2.10. Multi-category attacks

   Attacks categories are not mutually exclusive and any combination can
   be used to perform specific attacks.

   For example, one can mount a rogue attack to perform a performance
   attack starving the memory of an ITR resulting in a DoS on the ITR.

3. Attack vectors

   This section presents techniques that can be used by attackers in
   order to succeed attacks leveraging the LISP protocol and/or

architecture.

3.1. Gleaning

To reduce the time required to obtain a mapping, the optional
gleaning mechanism allows an xTR to directly learn a mapping from the
LISP data encapsulated packets and the Map-Request packets that it
receives. LISP encapsulated data packets contain a source RLOC,
destination RLOC, source EID and destination EID. When an xTR
receives an encapsulated data 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.
The same technique can be used when an xTR 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 xTR
sends a Map-Request to retrieve the actual mapping for the gleaned
EID from the mapping system.

If a packet injected by an off-path attacker and with a spoofed inner
address is gleaned by an xTR then the attacker may divert the traffic
meant to be delivered to the spoofed EID as long as the gleaned entry
is used by the xTR. This attack can be used as part of replay,
packet manipulation, packet interception and suppression, or DoS
attacks as the packets are sent to the attacker.

If the packet sent by the attacker contains a spoofed outer address
instead of a spoofed inner address then it can achieve a DoS or a
performance attack as the traffic normally destined to the attacker
will be redirected to the spoofed source RLOC. Such traffic may
overload the owner of the spoofed source RLOC, preventing it from
operating properly.

If the packet injected uses both inner and outer spoofing, the
attacker can achieve a spoofing, a performance, or an amplification
attack as traffic normally destined to the spoofed EID address will
be sent to the spoofed RLOC address. If the attacked LISP site also
generates traffic to the spoofed EID address, such traffic may have a
positive amplification factor.

A gleaning attack does not only impact the data-plane but can also
have repercussions on the control-plane as a Map-Request is sent
after the creation of a gleaned entry. The attacker can then achieve
DoS and performance attacks on the control-plane. For example, if an
attacker sends a packet for each address of a prefix not yet cached
in the EID-to-RLOC cache of an xTR, the xTR will potentially send a
Map-Request for each such packet until the mapping is installed which
leads to an over-utilisation of the control-plane as each packet
generates a control-plane event. In order for this attack to

succeed, the attacker may not need to use spoofing. This issue can
occur even if gleaning is turned off since whether or not gleaning is
used the ITR may need to send a Map-Request in response to incoming
packets whose EID is not currently in the cache.

Gleaning attacks are fundamentally involving a time-shifted mode of
operation as the attack may last as long as the gleaned entry is kept
by the targeted xTR. RFC 6830 [RFC6830] recommends to store the
gleaned entries for only a few seconds which limits the duration of
the attack.

Gleaning attacks always involve external data-plane attackers but
results in attacks on either the control-plane or data-plane.

It is worth to notice that the outer spoofed address does not need to
be the RLOC of a LISP site an may be any address.

3.2. Locator Status Bits

When the L bit in the LISP header 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. The reaction of a LISP xTR that receives such a packet is
left as operational choice in [RFC6830].

When an attacker sends a LISP encapsulated packet with a crafted LSB
to an xTR, it can influence the xTR's choice of the locators for the
prefix associated to the source EID. In case of an off-path
attacker, the attacker must inject a forged packet in the network
with a spoofed inner address. An on-path attacker can manipulate the
LSB of legitimate packets passing through it and hence does not need
to use spoofing. Instead of manipulating the LSB field, an on-path
attacker can also obtain the same result of injecting packets with
invalid LSB values by replaying packets.

The LSB field can be leveraged to mount a DoS attack by either
declaring all RLOCs as unreachable (all LSB set to 0), or by
concentrating all the traffic to one RLOC (e.g., all but one LSB set
to 0) and hence overloading the RLOC concentrating all the traffic
from the xTR, or by forcing packets to be sent to RLOCs that are
actually not reachable (e.g., invert LSB values).

The LSB field can also be used to mount a replay, a packet
manipulation, or a packet interception and suppression attack.
Indeed, if the attacker manages to be on the path between the xTR and
one of the RLOCs specified in the mapping, forcing packets to go via
that RLOC implies that the attacker will gain access to the packets.

Attacks using the LSB are fundamentally involving a time-shifted mode
of operation as the attack may last as long as the reachability
information gathered from the LSB is used by the xTR to decide the
RLOCs to be used.

3.3. Map-Version

When the Map-Version bit of the LISP header is set to 1, it indicates
that 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 xTR
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 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.

o 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.

A cross-mode attacker can use the Map-Version bit to mount a DoS
attack, an amplification attack, or a spoofing attack. For instance
if the mapping cached at the xTR is outdated, the xTR will send a
Map-Request to retrieve the new mapping which can yield to a DoS
attack (by excess of signalling traffic) or an amplification attack
if the data-plane packet sent by the attacker is smaller, or
otherwise uses fewer resources, than the control-plane packets sent
in response to the attacker's packet. With a spoofing attack and if
the xTR considers that the spoofed ITR has an outdated mapping, it
will send an SMR to the spoofed ITR which can result in performance,
amplification, or DoS attack as well.

Map-Version attackers are inherently cross mode as the Map-Version is
a method to put control information in the data-plane. Moreover,
this vector involves live attackers. Nevertheless, on-path attackers
do not take specific advantage over off-path attackers.

3.4. Routing Locator Reachability

The Nonce-Present and Echo-Nonce bits in the LISP header are used to
verify the reachability of an xTR. A testing xTR sets the Echo-Nonce

and the Nonce-Present bits in LISP data encapsulated packets and
include a random nonce in the LISP header of packets. Upon reception
of these packets, the tested xTR stores the nonce and echo it
whenever it returns a LISP encapsulated data packets to the testing
xTR. The reception of the echoed nonce confirms that the tested xTR
is reachable.

An attacker can interfere with the reachability test by sending two
different types of packets:

1. LISP data encapsulated packets with the Nonce-Present bit set and
    a random nonce. Such packets are normally used in response to a
    reachability test.

2. LISP data encapsulated packets with the Nonce-Present and the
    Echo-Nonce bits both set. These packets will force the receiving
    ETR to store the received nonce and echo it in the LISP
    encapsulated packets that it sends. These packets are normally
    used as trigger for a reachability test.

The first type of packets is used to make xTRs think that an other
xTR is reachable while it is not. It is hence a way to mount a DoS
attack (i.e., the ITR will send its packet to a non-reachable ETR
while it should use another one).

The second type of packets could be exploited to attack the nonce-
based reachability test. If the attacker sends a continuous flow of
packets that each have a different random nonce, the ETR that
receives such packets will continuously change the nonce that it
returns to the remote ITR, which can yield to a performance attack.
If the remote ITR tries a nonce-reachability test, this test may fail
because the ETR may echo an invalid nonce. This hence yields to a
DoS attack.

In the case of an on-path attacker, a packet manipulation attack is
necessary to mount the attack. To mount such an attack, an off-path
attacker must mount an outer address spoofing attack.

If an xTR chooses to periodically check with active probes the
liveness of entries in its EID-to-RLOC cache (as described in section
6.3 of [RFC6830]), then this may amplify the attack that caused the
insertion of entries being checked.

3.5. Instance ID

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 the instance ID decides the

forwarding table to use to forward the decapsulated packet in the
LISP site.

An attacker (either a control-plane or data-plane attacker) can use
the instance ID functionality to mount an intrusion attack.

3.6. 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).

As any other system relying on proxies, LISP interworking can be used
by attackers to hide their exact origin in the network.

3.7. Map-Request messages

A control-plane off-path attacker can exploit Map-Request messages to
mount DoS, performance, or amplification attacks. By sending Map-
Request messages at high rate, the attacker can overload nodes
involved in the mapping system. For instance sending Map-Requests at
high rate can considerably increase the state maintained in a Map-
Resolver or consume CPU cycles on ETRs that have to process the Map-
Request packets they receive in their slow path (i.e., performance or
DoS attack). When the Map-Reply packet is larger than the Map-
Request sent by the attacker, that yields to an amplification attack.
The attacker can combine the attack with a spoofing attack to
overload the node to which the spoofed address is actually attached.

It is worth to notice that if the attacker sets the P bit (Probe Bit)
in the Map-Request, it is legitimate the send the Map-Request
directly to the ETR instead of passing through the mapping system.

The SMR bit can be used to mount a variant of these attacks.

For efficiency reasons, Map-Records can be appended to Map-Request
messages. When an xTR receives a Map-Request with appended Map-
Records, it does the same operations as for the other Map-Request
messages and is so subject to the same attacks. However, it also
installs in its EID-to-RLOC cache the Map-Records contained in the
Map-Request. An attacker can then use this vector to force the
installation of mappings in its target xTR. Consequently, the EID-

   to-RLOC cache of the xTR is polluted by potentially forged mappings
   allowing the attacker to mount any of the attacks categorized in
   Section 2.2 (see Section 3.8 for more details). It is worth to
   mention that the attacker does not need to forge the mappings present
   in the Map-Request to achieve a performance or DoS attack. Indeed,
   if the attacker owns a large enough EID prefix it can de-aggregate it
   in many small prefixes, each corresponding to another mapping and it
   installs them in the xTR cache by mean of the Map-Request.

   Moreover, attackers can use Map Resolver and/or Map Server network
   elements to relay its attacks and hide the origin of the attack.
   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.

3.8. Map-Reply messages

   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 an off-path attack is limited given the size of
   the nonce (64 bits). Nevertheless, 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.

   If an attacker manages to send a valid (i.e., in response to a Map-
   Request and with the correct nonce) Map-Reply to an ITR, then it can
   perform any of the attack categorised in Section 2.2 as it can inject
   forged mappings directly in the ITR EID-to-RLOC cache. For instance,
   if the mapping injected to the ITR points to the address of a node
   controlled by the attacker, it can mount replay, packet manipulation,
   packet interception and suppression, or DoS attacks as it will
   receive every packet destined to a destination lying in the EID
   prefix of the injected mapping. In addition, the attacker can inject
   plethora of mappings in the ITR to mount a performance attack by
   filling up the EID-to-RLOC cache of the ITR. If the attacker can
   also mount an amplification attack as soon as the ITR has to send a
   lot of packets to the EIDs matching the injected mapping. In this
   case, the RLOC address associated to the mapping is the address of
   the real target of the attacker and all the traffic of the ITR will
   be sent to the target which means that with one single packet the
   attacker may generate very high traffic towards its final target.

   If the attacker is a valid ETR in the system it can mount a rogue
   attack if it uses prefixes over-claiming. In such a scenario, the
   attacker 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 attacker. For instance if the owned
prefix 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 attacker has authority on. With such
technique, the attacker can mount the attacks presented above as it
can (partially) control the mappings installed on its target ITR. To
force its target ITR to send a Map-Request, nothing prevents the
attacker to initiate some communication with the ITR. This method
can be used by internal attackers that want to control the mappings
installed in their site. To that aim, they simply have to collude
with an external attacker ready to over-claim prefixes on behalf of
the internal attacker.

It is worth to notice that when the Map-Reply is in response to a
Map-Request sent via the mapping system (i.e., not send directly from
the ITR to an ETR), the attacker does not need to use a spoofing
attack to achieve its attack as by design the source IP address of a
Map-Reply is not known in advance by the ITR.

Map-Request and Map-Reply messages are exposed to any type of
attackers, on-path or off-path but also external or internal
attackers. Also, even though they are control message, they can be
leveraged by data-plane attackers. As the decision of removing
mappings is based on the TTL indicated in the mapping, time-shifted
attackers can take benefit of injecting forged mappings as well.

3.9. Map-Register messages

Map-Register messages are sent by ETRs to Map Servers 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 over-claim the prefix it owns in order to
influence the route followed by Map-Requests for EIDs outside the
scope of its legitimate EID prefix (see Section 3.8 for the list of
over-claiming attacks).

A compromised ETR can also de-aggregate its EID prefix in order to
register more EID prefixes than necessary to its Map Servers (see
Section 3.7 for the impact of de-aggregation of prefixes by an
attacker).

Similarly, a compromised Map Server can accept invalid registration

or advertise invalid EID prefix to the mapping system.

3.10. 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.

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.

4. 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].

Like public deployments of any other control plane protocols, in an
Internet deployment mappings are public and hence provide information
about the infrastructure and reachability of LISP sites (i.e., the
addresses of the edge routers). Depending upon deployment details,
LISP map replies might or might not provide finer grained and more
detailed information than is available with currently deployed
routing and control protocols.

5. Threats Mitigation

Most of threats can be mitigated 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.

The control-plane is the most critical part of LISP from a security
viewpoint and it is worth to notice that the specifications already
offer authentication mechanism for mappings registration ([RFC6833])
and this mechanism combined with LISP-SEC [I-D.ietf-lisp-sec]
strongly mitigates threats in non-trustable environments such as the
Internet. Moreover, LISP specifications define an authentication
data field for Map-Request messages and Encapsulated Control messages
without specifying how to use it [RFC6830]. The presence of this

field in the specifications allows to propose a general
authentication mechanisms for the LISP control-plane while staying
backward compatible. The exact technique still has to be designed
and defined. To maximally mitigate the threats on the mapping
system, authentication must be used, whenever possible, for both Map-
Request and Map-Reply messages and for messages exchanged internally
among elements of the mapping system, such as specified in
[I-D.ietf-lisp-sec] and [I-D.ietf-lisp-ddt].

Systematically applying filters and rate-limitation, as proposed in
[RFC6830], mitigates most of the threats presented in this document.
In order to minimise the risk of overloading the control-plane with
actions triggered from data-plane events, such actions should be rate
limited.

Finally, all information opportunistically learned (e.g., with LSB or
gleaning) should be used with care until they are verified. For
instance, a reachability change learned with LSB should not be used
directly to decide the destination RLOC, but instead should trigger a
rate-limited reachability test. Similarly, a gleaned entry should be
used only for the flow that triggered the gleaning procedure until
the gleaned entry has been verified [Trilogy].

6. Security Considerations

This entire document is dedicated to threat analysis and mitigation
of the Locator/Identifier Separation Protocol, aiming at helping to
understand the security risks at stake, and how to mitigate them,
while deploying LISP in non-trustable environments.

7. IANA Considerations

This document makes no request to IANA.

8. 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, Ross
Callon, 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.

9. References

9.1. Normative References

[RFC6830]   Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The
            Locator/ID Separation Protocol (LISP)", RFC 6830,
            DOI 10.17487/RFC6830, January 2013,
            <http://www.rfc-editor.org/info/rfc6830>.

[RFC6832]   Lewis, D., Meyer, D., Farinacci, D., and V. Fuller,
            "Interworking between Locator/ID Separation Protocol
            (LISP) and Non-LISP Sites", RFC 6832, DOI 10.17487/
            RFC6832, January 2013,
            <http://www.rfc-editor.org/info/rfc6832>.

[RFC6833]   Fuller, V. and D. Farinacci, "Locator/ID Separation
            Protocol (LISP) Map-Server Interface", RFC 6833,
            DOI 10.17487/RFC6833, January 2013,
            <http://www.rfc-editor.org/info/rfc6833>.

[RFC6834]   Iannone, L., Saucez, D., and O. Bonaventure, "Locator/ID
            Separation Protocol (LISP) Map-Versioning", RFC 6834,
            DOI 10.17487/RFC6834, January 2013,
            <http://www.rfc-editor.org/info/rfc6834>.

[RFC6973]   Cooper, A., Tschofenig, H., Aboba, B., Peterson, J.,
            Morris, J., Hansen, M., and R. Smith, "Privacy
            Considerations for Internet Protocols", RFC 6973,
            DOI 10.17487/RFC6973, July 2013,
            <http://www.rfc-editor.org/info/rfc6973>.

9.2. Informative References

[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-03 (work in
              progress), April 2015.

   [I-D.ietf-lisp-sec]
              Maino, F., Ermagan, V., Cabellos-Aparicio, A., and D.
              Saucez, "LISP-Security (LISP-SEC)", draft-ietf-lisp-sec-08
              (work in progress), April 2015.

   [RFC7215] Jakab, L., Cabellos-Aparicio, A., Coras, F., Domingo-
              Pascual, J., and D. Lewis, "Locator/Identifier Separation
              Protocol (LISP) Network Element Deployment
              Considerations", RFC 7215, DOI 10.17487/RFC7215,
              April 2014, <http://www.rfc-editor.org/info/rfc7215>.

   [Trilogy] Saucez, D. and L. Iannone, "How to mitigate the effect of
              scans on mapping systems", Trilogy Future Internet Summer
              School., 2009.

Appendix A. Document Change Log

   o Version 13 Posted August 2015.

      * Keepalive version.

   o Version 12 Posted March 2015.

      * Addressed comments by Ross Callon on the mailing list (http://
         www.ietf.org/mail-archive/web/lisp/current/msg05829.html).

      * Addition of a section discussing mitigation techniques for
         deployments in non-trustable environments.

   o Version 11 Posted December 2014.

      * Editorial polishing. Clarifications added in few points.

   o Version 10 Posted July 2014.

      * Document completely remodeled according to the discussions on
         the mailing list in the thread
         http://www.ietf.org/mail-archive/web/lisp/current/msg05206.html
         and to address comments from Ronald Bonica and Ross Callon.

   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 over-claiming and de-
       aggregation (see Section 3.8).

    * 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.

    * 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: ggx@gigix.net

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

Email: olivier.bonaventure@uclouvain.be