Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 9, 2008

M. Bagnulo
Huawei Lab at UC3M
July 8, 2007

Preliminary LISP Threat Analysis
draft-bagnulo-lisp-threat-01

Status of this Memo

Copyright Notice

Abstract

This document performs a preliminary threat analysis on the Locator/ID Separation Protocol (LISP) as defined in draft-farinacci-lisp-01.txt.

Table of Contents

1. Introduction

The Locator/ID Separation Protocol (LISP) is defined in draft-farinacci-lisp-01.txt [1]. The goal of this document is to identify the different threats in the current LISP protocol in order to understand the current level of protection of the LISP protocol and identify additional security mechanisms that are needed to protect it.

As in any ID/loc split approach, the critical operation is the creation of ID to locator binding state in any device that will use it later on. In the particular case of LISP the critical operation is the creation of EID to RLOC mappings in the ITR and the ETR. Such operation is performed in 3 different ways:

Using the information obtained from a LISP data packet (looking into the EID and RLOC information included by the originator of the packet)

Using the information contained in the Map-Reply message

Using a EID-to-RLOC database

This document performs threat analysis for the first two cases. The third case, the one of the EID-to-RLOC database, has a different trust model, and a specific threat analysis needs to be performed for that case.

2. Application Scenario

We will consider the following application scenario.

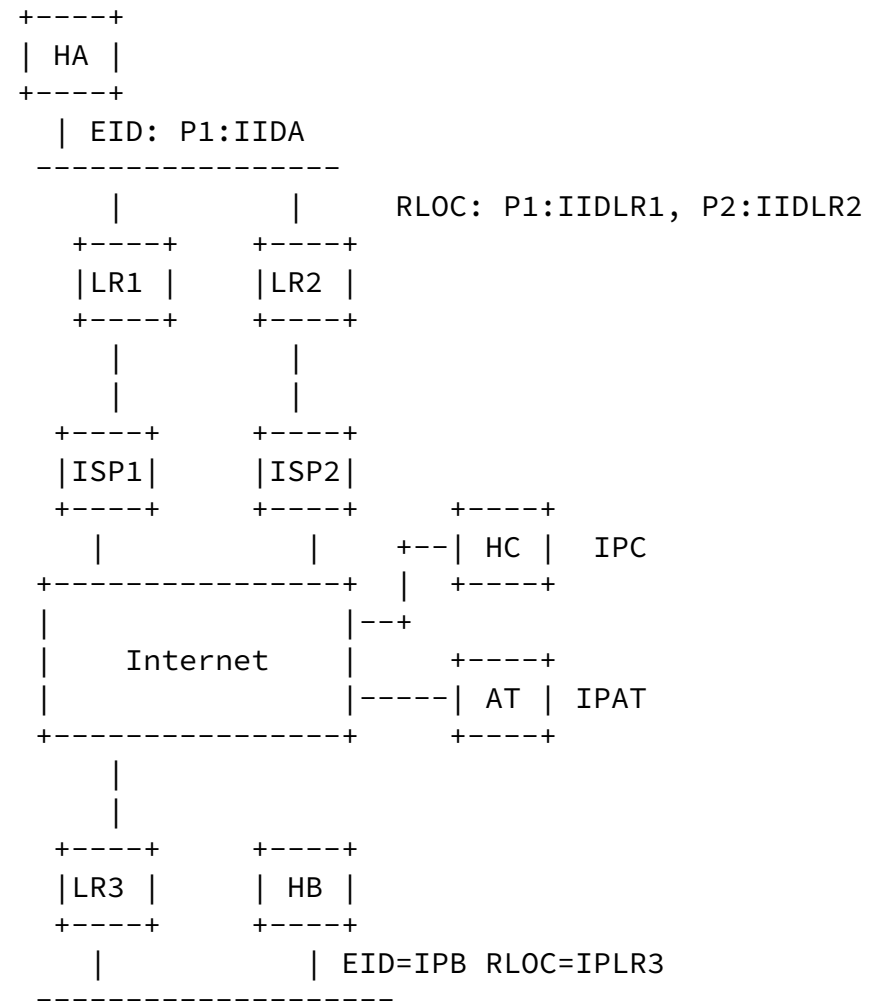
HA
EID: P1:IIDA
RLOC: P1:IIDLR1, P2:IIDLR2
LR1
LR2
ISP1
ISP2
HC
IPC
Internet
AT
IPAT
LR3
HB
EID=IPB RLOC=IPLR3

LR: LISP Router that behaves as both as the ITR and the ETR

In the depicted scenario we have two LISP capable sites. One of the sites, depicted on top of the figure, is multihomed to ISP1 and ISP2. We assume that we are using LISP1, so one of the routable addresses is used as EID, let's say that for host HA P1:IIDA is used as EID. In addition, the locators for that host will be the two addresses of the exit routers that are playing the role of ITR namely LR1 and LR2, so RLOCs are P1:IIDLR1 and P2:IIDLR2.

(LR stands for LISP router since it plays both the roles of ITR and ETR).

The other LISP capable site is the one depicted in the lower part of the figure and it has a single ISP and a single ITR/ETR namely, LR3. Host H3 located in this site has IPB as EID and the address of the LR3, LPLR3 as RLOC. Since we are using LISP1, IPB is a routable address

Hosts HC and AT are other hosts in the Internet, with addresses IPC and IPAT respectively.

HA, HB and HC are victims and AT is the attacker.

3. Threat analysis

Full-time off-path attacker assumption

We will limit the considered attacks to those where the attacker is not located along the path used to route packets of the communication being attacked during the whole duration of the attack.

There are then two type of attacks:

- Off-path attacks the attacker is off-path during the whole duration of the attack

- Time shifted attacks: the attacker is located along path during a limited period, but the duration of the attack is significantly longer than the period that the attacker is along the path.

3.1. Identity hijacking

In the attacks considered in this section the attacker will try to hijack the identity of one victim on the eyes of another victim. This means that two parties are deceived, one that thinks that is communicating with the owner of a given identify, but it is communicating with the attacker instead and the party whose identify is being stolen.

As we mentioned earlier, there are two messages an attacker can use to create the state required to launch an attack: the LISP data packets or the Map-Reply message. In this section we will first present the attacks based on using the LISP data packets and then the attacks that use the Map-Reply messages.

3.1.1. Attacks using the LISP data packets to create state

3.1.1.1. Attacker initiated communication (Hijacking a client identity)

In this case, the attacker will initiate a communication with one victim pretending to be someone else.

In the scenario above, the attacker AT will try to initiate a communication with HA pretending to be HC. In order to do that it will send a LISP packet with the following parameters:

- Destination RLOC (outer header destination address): P1:IIDA

- Destination EID (Inner header destination address): P1:IIDA

- Source RLOC (outer header source address): IPAT

- Source EID (inner header source address): IPC

The packet will be received by LR1 who will generate the LISP Map-Reply message back to IPAT and will store the EID to RLOC mapping information for the received data packet. The EID to RLOC mapping information stored at LR1 contains the following information: EID = IPC, RLOC=IPAT

In this case HA will be communicating with the attacker AT but HA

believes that he is communicating with HC. HC identity has been stolen by AT in the eyes of HA.

Basically, in this case the packet that triggers all of this is (sent from AT toward HA (who's EID is P1:11DA)) looks like:

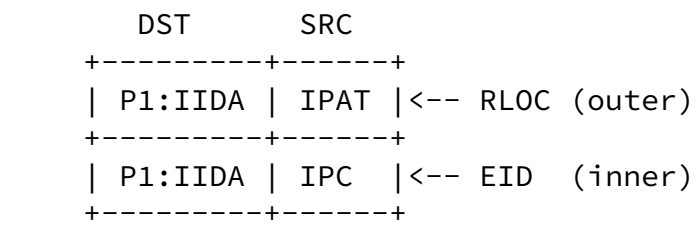

The mapping installed in LR1 is {EID, RLOC} = {IPC, IPAT}, and thus HC's identity is hijacked (at least from HA's perspective) by AT (since it thinks that IPAT is the RLOC associated with EID HC = inaddr(IPC)). As a result, subsequent packets emitted by HA will look like

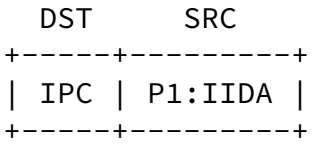

and LR1 will encap as follows (giving the installed mapping):

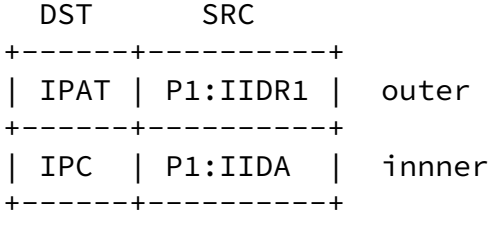

and the hijack is complete.

3.1.1.2. Victim initiated communication (Hijacking a server identity)

In the previous section, the attacker is hijacking the identity of a client, since the attacker is the one that initiates the communication. While this is problematic, a much more ambitious attacks would to hijack the identity of a server, i.e. to hijack the identity of a server when the victim initiates a communication towards the server.

This is also possible with LISP. It would work in the following way.

Suppose that HC is a server that HA uses regularly (e.g. a newspaper web site)

Suppose that AT wants that future communication initiated by HA to HC are directed to AT i.e. AT wants to hijack HC identity for all the communications initiated by HA.

In order to do that, AT performs the following actions. It first needs to install false EID-to-RLOC mapping information in LR1. In order to do that, it sends a data packet containing the following information

- Destination RLOC (outer header destination address): P1:IIDA

- Destination EID (Inner header destination address): P1:IIDA

- Source RLOC (outer header source address): IPAT

- Source EID (inner header source address): IPC

The data packet could be a UDP packet that will be discarded upon reception because there is no process listening in the requested port.

LR1 will store that in order to reach IPC it must tunnel the packets to IPAT.

However, there is no actual ongoing communication between HA and HC

at the moment, so the attack has no effect so far. The attacker must then keep the EID to RLOC mapping information alive in LR1 for when HA decides to initiate a communication to HC. The attacker can do that by sending periodic data packets with the same information detailed before.

Suppose that at any point in the future, HA decides to initiate a communication with HC. It will send a data packet with destination address IPC. The data packet will be intercepted by LR1 and tunnelled to the attacker at IPAT, since this is the mapping information available at LR1.

Note that these attacks affect all future communications started by HA and that it affects communication initiated by HA.

3.1.1.3. Intercepting ongoing communications (becoming a MITM)

In the two previous sections, we have considered the case where the attacker wants to hijack a future communication pretending to be one of the involved parties.

In this section we will consider the case where there is an ongoing communication and the attacker wants to hijack the ongoing communication.

Suppose that there is an ongoing communication between HA and HB. In this case, LR1 contains a mapping between EID=IPB and RLOC=IPLR3. LR3 contain a mapping between EID= P1:IIDA and RLOC=P1:IIDLR1, P2: IIDLR2.

Suppose now that AT sends two packets, one to LR1 and another to LR3.

The packet sent to LR1 will contain mapping information of EID=IPB and RLOC=IPAT. The packet sent to LR3 will contain mapping information EID=P1:IIDA and RLOC=IPAT.

(One may think that ingress filtering could help here, but note that the attacker is sending packets from the claimed locator, so ingress filters won't be of any use to prevent this attack)

If the new EI-to-RLOC information overrides the previously available mapping information (this would depend on how the new mapping information is managed, but it seems that in the current version, latest information supersedes older information) , the result is that LR1 will tunnel packets addressed to HB to the attacker at IPAT and LR2 will tunnel packets addressed to HA to the attacker at IPAT. The attacker has now placed himself as a man in the middle in the

communication. It can either modify packets or simply sniff them.

In this case, packets exchanged in this attack would look like this: Suppose HA and HB are communicating. In this case, LR1 has {IPB,IPLR3} and LR3 has {P1:IIDA, {P1:IIDLR1, P2:IIDLR2}} (P1:IIDA has 2 possible RLOCs). Now, the attacker at IPAT sends

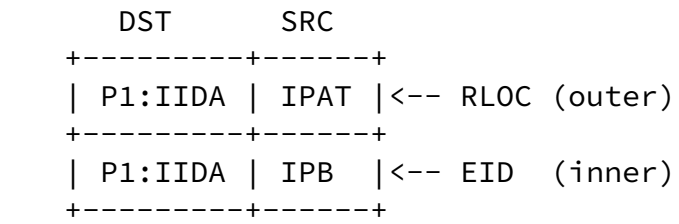

to HA. the packet is intercepted by LR1, which results in the mapping {IPB, IPAT} (so AT has half of the connection). That is, packets from HA -> HB look like

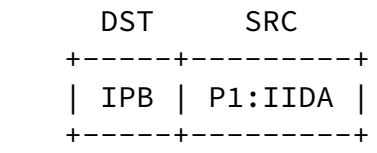

LR1 builds

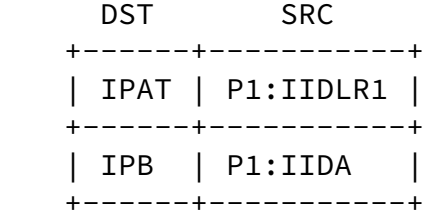

and packets are delivered to the MITM.

Going the other way, AT sends

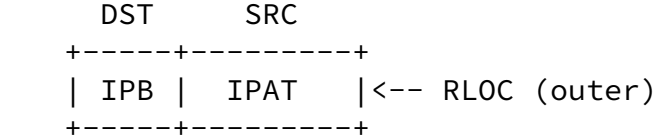

```
| IPB | P1:IIDA |<-- EID  (inner)
+-----+---------+
```


to HB. The packet is intercepted by LR3, which results in the mapping {P1:IIDA, IPAT} (so AT has the other half of the connection), so packets sent to P1:IIDA (HA) get delivered to AT (inaddr(IPAT)).

3.1.2. Attacks using the Map-Reply message to create state

Map-Reply messages are protected by a nonce, which is copied from the LISP data packet that triggered the Map-Reply generation. Such nonce protects from Map-Reply messages generated spontaneously (i.e. not generated as a reply to an actual LISP data packet). While this protection prevents a number of attacks, there are still a few attacks that are possible, which are presented in this section.

3.1.2.1. Less-specific prefix attack

Map-Reply messages are very powerful because they can contain single EIDs but also EID prefixes. Such capability allows any malicious party receiving a LISP data packet to reply with a Map-Reply message that includes a less specific EID prefix that contains more than its own EIDs. Basically, the problem here is that the return routability check only verifies the presence of a single EID and not the whole EID prefix. The attack could be performed in the following way:

For whatever reason, HB decides to start a communication with AT. HB generates a data packet containing:

```
  DST    SRC
+------+-----+
| IPAT | IPB |
+------+-----+
```

and LR3 will encap as follows:

```
  DST      SRC
+------+--------+
| IPAT | IPLR3 |  outer
```

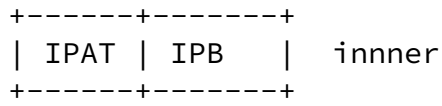

In addition, the encapsulated packet will contain a nonce NB.

AT will receive the packet, will store the state corresponding to HB (EID=IPB, RLOC=IPLR3). In addition, AT can reply with a Map_Reply

message. However, AT can reply with an Map-Reply message that contains amore specific prefix that contains other prefixes than its own. In particular, AT can include the 0/0 prefix in the Map-Reply message. The Map-Reply message is validated by the nonce NB the was included in the received ISP data packet. The Map-Reply message that AT will send back to LR3 will be:

- Nonce: NB

- EID mask-len: 0

- EID prefix: 0

- Locator: IPAT

Upon the reception of such Map-Reply message, LR3 will install the EID-to-RLOC mapping which will map the whole EID space to the attacker locator IPAT. All the following packets routed by LR3 will be forwarded to the attacker AT. Note that this not only affects the packets generated by HB but to all the packets generated by any host behind LR3. Also note that this is the more extreme case, where the whole EID space is hijacked, but it would also be possible to hijack parts of the EID space, which would result in less severe attacks, but probably more difficult to detect.

3.1.2.2. Time-shifted attack

Time-shifted attacks are those where the attacker is located along the path during a short period and launches an attack that persists in time long after the attacker left the on-path position. These attacks have been considered relevant during the design of other protocols for mapping identities and location, such as MIP. In

particular, in order to prevent time-shifted attacks in MIPv6 route optimization, the MIPv6 specification requires the periodic performance of the return routability check. The lifetime of the state created using the validation information obtained using a single return routability check is limited to 7 minutes in the MIPv6 spec. This implies that the maximum time span that a time-shifted attack can be active after the attacker left the on-path position is 7 minutes. For additional information about the MIPv6 security design, the reader is referred to [2].

In the case of LISP, an attacker can launch a time-shifted attacks if he is able to learn a nonce of a LISP data packet generated by the victim. Once the attacker has obtained the nonce, it can then generate a Map-Reply message and hijack any portion of the the EID space, thanks to the aforementioned capabilities of EID aggregation by the means of less-specific EID prefixes. The Map-reply message

would be accepted by the victim because it contains a valid nonce and it will install the ED-to-RLOC mapping. The state will remain in the victim even when the attacker has left the on-path position. The attacker is able to keep the state alive by refreshing the state (is likely that LISP provides some means for this since it should be able for a genuine host to preserve the state in its peers, even when the original path is unreachable)

The attack is then as follows:

The attacker is located along a path sniffing packets and looking for any LISP data packet generated by the victim.

Once the attacker finds such a packet, it learns the nonce in the LISP header.

The attacker generated a Map-Reply packet containing the nonce, the EID prefix 0/0 and its own locator.

The attacker sends the Map-Reply which is processed by the victim's router which installs the state.

From this moment, all the outgoing packets of the victim's site are forwarded to the locator selected by the attacker.

The attacker can leave its position and move to its own locator, but the attack is still active, since the state is still installed in the victim's router.

3.2. DoS attacks

In this section, we describe DoS attacks related to LISP.

3.2.1. Flooding a third party

In this case, the attacker AT wants to use HA to flood a victim HC.

In order to do that, AT first initiates a communication with HA and starts a download of a heavy flow. Once that the flow is downloading, it redirects the heavy flow towards HC, flooding it. This is performed in the following way.

AT initiates a communication with HA. In this case, AT uses IPAT as EID and IPAT as RLOC. This mapping information is stored in LR1 since it is contained in the initial LISP data packet as described previously. AT then starts downloading a heavy flow form HA. At some point then, AT redirects the flow towards HC. It can do so by including IPC as a RLOC for the EID IPAT that is being used in the

communication that is downloading the heavy flow. The IPC RLOC could be included since the beginning with a low priority and now simply send a LISP Map-Reply message with a higher priority to IPC. In any case the result is that the flow will flood HC.(Note that it is expected that AT can include additional locators associated to the EID IPAT during the initial period where there is a direct communication between AT and HA)

It should be noted that in some cases, in order to keep the flow going, it is necessary that the receiver sends some ACK packets or similar. In this case, it is possible that the attacker can send such packets, since EID IPAT in LR1 can contain two valid RLOCs i.e. IPAT and IPC. In this case, if IPC has higher priority that IPAT, LR1 will send packets to IPC but will accept packets coming from IPAT as valid packets from the EID IPAT. In this case, the attacker could send ACK packets from its own location, and keep the flooding going towards IPC.

3.2.2. Preventing future communications

This case is similar to the one described in section Victim initiated communication (Hijacking a server identity), but only that instead of the attackers RLOC, a black hole IP address is included as the RLOC for a given EID. The result is that the traffic addressed to the EID is sent to a black hole, resulting in a DoS attacks form communications to that EID.

In addition to that, it is possible to use the Less-specific prefix attack to perform a DoS attack. In this case, a larger number of destinations are affected, since it affects a whole prefix.

Finally, it should be noted that the attacks do not affect the traffic generated by a single machine, but to all the traffic routed by the affected ITR i.e. to the traffic generated by all the hosts behind the ITR.

3.2.3. Interrupt ongoing communication

This case is similar to the one described in the section Intercepting ongoing communications (becoming a MITM) with the only difference that an back hole IP address is included as RLOC for the ongoing communication, terminating it.

3.2.4. DoS attacks against LISP infrastructure

Another type of DoS attacks that must be considered are the DoS attacks against the LISP architecture itself. LISP infrastructure is likely to become a critical part of the network, since EIDs may not

be reachable without the LISP service. This makes the LISP routers a preferred target for attackers.

In particular LISP routers (ITR and ETR) are susceptible to DoS attacks. LISP routers store information about EID-to- RLOC mappings. They learn that information from data packets and from ICMP messages. An attacker could massively generate either tunnelled data packets so that the router cache is overflowed. The result is that routers will not be able to store legitimate EID-to-RLOC mapping information and that LISP features will be annulled. (in the case of using non routable EIDs, all communication would be annulled if LISP

functionality is not available)

Current LISP proposal includes rate-limiting techniques for protecting against DoS attacks. Such technique can be useful to prevent LISP routers from crashing under the attack. However, this technique does not prevent the actual effect of the attack, which that hosts served by the LISP router under attack will not be able to communicate. This is so, because the LISP router rate limiting techniques, will also affect the legitimate request from internal and external hosts, making communication hard if not impossible (depending on the strength of the actual attack)

4. Security considerations

This note outlines security issues of the LISP protocol.

5. Acknowledgments

Pekka Nikander and Dave Meyer reviewed this document and provided comments. In addition, Dave Meyer wrote the text containing the packet description of attacks described in sections 3.1.1.1 and 3.1.1.3. Dino Farinacci provided clarifying comments about how LISP works.

6. Normative References

[1] Farinacci, D., Fuller, V., Oran, D., and D. Meyer, "Locator/ID Separation Protocol (LISP)", draft-farinacci-lisp-01 (work in progress), June 2007.

[2] Nikander, P., Arkko, J., Aura, T., Montenegro, G., and E. Nordmark, "Mobile IP Version 6 Route Optimization Security Design Background", RFC 4225, December 2005.

Author's Address

Marcelo Bagnulo
Huawei Lab at UC3M

Av. Universidad 30
Leganes, Madrid 28911
SPAIN

Phone: 34 91 6249500
Email: marcelo@it.uc3m.es
URI: http://www.it.uc3m.es

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).