Internet Engineering Task Force
Internet-Draft
Expires: January 6, 2011

D. Klein
M. Hartmann
M. Menth
University of Wuerzburg

July 5, 2010

NAT Traversal for LISP Mobile Node
draft-klein-lisp-mn-nat-traversal-00

Abstract

The Locator/Identifier Separation Protocol (LISP) is a new naming and addressing architecture to solve the Internet's routing scaling problem. It separates global routing in the Internet from local routing and naming in end-user networks. The basic LISP architecture does not support mobility. The mobility extension LISP Mobile Node (LISP-MN) describes a mechanism that extends LISP to support mobile nodes and enables them to roam into LISP and non-LISP networks while being reachable under the same address. Currently, LISP-MN does not support networks that use network address translation (NAT). This document presents an extension for LISP-MN that makes LISP mobile nodes behind a NAT globally reachable.

Status of this Memo

Copyright Notice

Table of Contents

1. Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Introduction

The Locator/Identifier Separation Protocol (LISP) [LISP] separates naming and local routing in edge networks from global routing in the Internet. Special endpoint identifiers (EIDs), which are independent of the global routing, are used to distinguish end-hosts on a global scale.

The basic LISP architecture does not support mobility of end-hosts. The extension LISP Mobile Node (LISP-MN) [LISP-MN] describes a mechanism that enables MNs to roam into LISP sites and non-LISP networks while being reachable under the same EID. The operation of LISP-MN is illustrated and analyzed for different networking scenarios in [MEKL10]. When a MN roams into a network, it receives a new address from the network, e.g., from DHCP. To be reachable as a LISP node by its EID, it registers this address in the global LISP mapping system.

In a non-LISP network without NAT, the assigned address serves as globally reachable routing locator (RLOC). When packets are sent to the EID of the MN, the RLOC of the MN is requested from the mapping system and the packets are encapsulated and tunneled to this RLOC. The MN acts like an ETR, decapsulates the traffic, and receives the actual packets. When the MN wants to send traffic to another node, it acts like an ITR. If the other node is in a LISP site, it encapsulates the traffic towards the RLOC of this LISP site. If the other node has a globally reachable IP address, the MN encapsulates the traffic towards its configured proxy ETR. This proxy ETR decapsulates the traffic and forwards it to the other node. If the MN roams into a LISP network, the operation is more complex, but the details are not relevant in this draft.

If the MN roams into a network using network address translation (NAT), the MN is assigned a private address which is not routable in the Internet. Thus, packets tunneled to that address from the Internet cannot reach the MN. Therefore, LISP-MN does not work behind NAT boxes.

In this document we present an extension to LISP-MN which allows NAT traversal for LISP-MNs by utilizing special NAT traversal routers (NTRs) whose functionality may be integrated in a MN's MS. In the

following, we assume that a NAT box not only translates IP addresses but also ports (NAPT). Section 3 introduces the most important terms and definitions used in this document. Section 4 gives a short overview of the NAT traversal technique and Section 5 describes the NAT traversal mechanism in detail. Section 6 discusses some security issues which arise due to the NAT traversal mechanism and finally, Section 7 gives a short conclusion. A paper version of this draft is provided in [KLHA10].

3. Terminology

This section lists the most important terms and definitions used throughout this document.

Endpoint Identifier (EID): IPv4 or IPv6 address of an end-host that is used to identify the end-host on a global scale. EIDs are only routable within a LISP site. Transport connections between end-hosts are bound to EIDs. Therefore, EIDs must not change due to a roaming event because otherwise, existing transport connections would fail.

Routing Locator (RLOC): Globally routable IPv4 or IPv6 address which is used to reach LISP end-hosts in another LISP site.

Ingress Tunnel Router (ITR): An ITR is the gateway of a LISP site and receives outgoing packets from LISP nodes within its LISP site destined to nodes in other LISP or non-LISP sites. The (inner) header (IH) of outgoing packets remains unchanged and the ITR adds an additional outer header (OH) that contains RLOC addresses to make the packet globally routable. The ITR uses its own RLOC as source address in the OH and for the destination address, it obtains an RLOC for the destination EID from the mapping system. The ITR also adds a special UDP LISP-header between the outer and inner IP header. UDP port 4341 is used as destination port for data packets and UDP port 4342 is used as destination port for signaling packets. The source port for both packet types is randomly chosen and has no special purpose.

Egress Tunnel Router (ETR): An ETR of a LISP site receives LISP-encapsulated IP packets from the Internet which are addressed to one of its own RLOCs. The ETR decapsulates the packet, removes the additional UDP header, and forwards the packet to the destination node within its own LISP site according to the EID in the inner header.

Stationary Node (SN): A non-mobile end-host which resides in a LISP or non-LISP site and has a relatively fixed IP address. SNs inside a LISP site do not need to be upgraded to communicate via LISP with other LISP nodes within or in a different LISP site.

Mobile Node (MN): A mobile end-host which implements LISP mobile node operations [LISP-MN]. It acts as a lightweight LISP site and has ITR and ETR functionality. It has a fixed EID for transport connections and uses a care-of-address which is dynamically assigned from the hosting domain as locator. Packets to and from mobile nodes are always LISP-encapsulated and carry the current care-of-address in the outer header and the fixed EID in the inner header.

Proxy ITR (PITR): PITRs enable SNs inside LISP sites to be reachable from the non-LISP part of the Internet. PITRs advertise highly aggregated anycast EID-prefixes via BGP in the Internet. IP packets sent from non-LISP sites addressed to EIDs are then routed to the next PITR. The PITR performs ITR functionality on behalf of the non-LISP site and applies the necessary steps to encapsulate and forward a packet to its destination's ETR. The interworking architecture and PITRs are described in [LISP-IW].

Proxy ETR (PETR): PETRs are also part of the interworking architecture [LISP-IW]. They are required by LISP sites to reach non-LISP sites if one of the LISP site's upstream providers performs source address filtering. Normally, ITRs would send IP packets to non-LISP sites without an additional header and with the EID of the sending node as source address. If an upstream provider utilizes source address filtering, it drops packets with an EID source address because EIDs are usually not part of the provider's address range. To circumvent this, ITRs tunnel packets to a pre-configured PETR which acts as ETR on behalf of non-LISP sites.

Map-Server (MS): MSs act as interface for the mapping system and ease the communication between ETRs and different mapping systems [LISP-MS]. A MS learns EID-to-RLOC mappings from authoritative sources like ETRs or MNs via Map-Register messages described in [LISP] and distributes these mappings on behalf of the ETR or MN in the mapping system. For a MN, the MS also acts as proxy-ETR so that non-LISP networks can be reached by the MN.

NAT Traversal Router (NTR): A specific MS which implements the NAT traversal technique proposed in this document. It thereby allows MNs to be reachable behind a NAT-gateway although the MN has only received a non-globally routable private IPv4 [RFC1918] or private IPv6 [RFC4193] address as care-of-address.

4. Architecture Overview

The NAT traversal technique described in this document is implemented inside MSs and requires no new functionality in other entities. In the remainder of this document, we call a modified MS that implements the NAT traversal mechanism a NAT Traversal Router (NTR).

```
                                         ___________________
              __________               ,' INTERNET          `.
            ,' Non-LISP `.            /                       \
           /     site    +---+       |           _________     |
        +--+             |NAT|       |  +-----+ /Traffic  |    |
        |MN|0==========================0| NTR |< for MN   |    |
        +--+             +---+       |  | (MS)| \_________|    |
         ^ \              /          |  +-----+                |
         |  `.__________,'            \    ^                  /
         |                             `.__|________________,'
         |                                 |
         |_________________________________|
             Tunnel between MN and NTR
             used to bypass NAT
```

Figure 1: Architecture overview

When a MN roams into a network, it obtains a care-of-address and registers it as RLOC for its EID with its pre-configured NTR. The Map-Register message also induces context inside the NAT-gateway which allows incoming reply packets from the NTR to the MN to traverse the NAT box. The Map-Register message received by the NTR does not explicitly indicate wheter the MN is behind a NAT, but the NTR is able to determine whether the MN is behind a NAT with the information provided in the Map-Register message. If the MN is behind a NAT, the NTR registers its own IP address as RLOC for the EID of the MN in the mapping system. Thus, when traffic is sent to a MN behind a NAT, (P)ITRs tunnel the traffic to the NTR instead of to the care-of address of the MN. The NTR relays that traffic to the MN and the traffic traverses the NAT due to the context established for the NTR during the registration process. This essentially constitutes a tunnel between the NTR and the MN which is used to bypass the NAT gateway. Figure 1 shows an overview of the architecture and the basic idea of the NAT traversal mechanism.

5. NAT Traversal Mechanism

This section explains the control and data plane operations for NAT traversal by MNs.

5.1. Control Plane Operations

When a MN roams into a network, it receives a care-of-address, e.g. from the local DHCP service, and sends a Map-Register message to its MS using destination port 4342 without any LISP encapsulation. In contrast to the current behavior in LISP-MN, which uses a randomly chosen source port without special purpose, our NAT traversal proposal requires that source port 4341 is used. The collocated NTR compares the reported care-of-address with the source address of the Map-Register message. If they are the same, the MN is not behind a NAT and the address is registered as RLOC for the EID of the MN in the mapping system.

```
                                             EID->IP:Port   EID->RLOC
                    NAT-TABLE                   Mapping      Mapping
   +-------------+-------------+-----------+   +-------+   +---------+
   | Internal    |  External   |  Peer     |   | EID 1 |   |  EID 1  |
   | IP:Port     |  IP:Port    |  IP:Port  |   |  -->  |   |   -->   |
   |-------------|-------------|-----------|   |1.8.7.2|   |  RLOC N |
   |10.0.0.1:4341|1.8.7.2:20321|RLOC N:4342|   |:20321 |   +---------+
   +-------------+-------------+-----------+   +-------+        *
                                                    *         *
                                                  __*_______*_______
              __________                        ,'  *     *         `.
            ,' Non-LISP `.                     /    *   *    INTERNET \
    +--------+   site  +-------+              |   +------+             |
    |   MN   |         |  NAT  |              |   | NTR  |             |
    |  EID 1 |  -----> |1.8.7.2| ---------------> | (MS) |             |
    |10.0.0.1|    :    +-------+        :     |   |RLOC N|             |
    +--------+    :       /             :     |   +------+             |
            `.____:_____,'              :      \                      /
                  :                     :       `.__________________,'
                  :                     :
                  :                     :
            Src:     Dest:         Src:     Dest:
           +--------+------+      +--------+------+
       OH: |10.0.0.1|RLOC N|      |1.8.7.2 |RLOC N|
           +--------+------+      +--------+------+
      UDP: |4341    |4342  |      |20321   |4342  |
           +--------+------+      +--------+------+
     LISP: |REGISTRATION:  |      |REGISTRATION:  |
           |EID 1->10.0.0.1|      |EID 1->10.0.0.1|
           +--------+------+      +--------+------+
```

Figure 2: Registration process

If the two addresses differ, the MN is behind a NAT and the new NAT traversal concept for MNs behind NATs is used. The NAT traversal concept is explained using the packet flow sequence in Figure 2 as an example. A MN with EID 1 has roamed into a private network and obtained the care-of-address 10.0.0.1. It sends a Map-Register message from source port 4341 containing its care-of-address to the NTR with RLOC N. The intermediate NAT gateway translates the source IP/port combination 10.0.0.1:4341 into 1.8.7.2:20321 and stores this as context for packets to and from destination IP/port RLOC N:4342. The NTR detects that the care-of-address 10.0.0.1 differs from the source address of the Map-Register message (1.8.7.2) and, therefore, it stores its own IP address (RLOC N) as RLOC for EID 1 in the mapping system. In addition, the NTR records the source address and port of the Map-Register message (1.8.7.2:20321) with the EID (EID 1) in an EID-to-IP:Port table. The NTR needs this IP/port combination

to relay packets to the MN behind the NAT. To make the mapping system robust against stale information, an expiration timer is associated with registered EID-to-RLOC mappings. The same may be applied to the EID-to-IP:Port table. The expiration timer should be set to a small value so that the established context in the NAT gateway is also refreshed in time.

5.2. Data Plane Operations

```
                                             EID->IP:Port   EID->RLOC
                    NAT-TABLE                   Mapping      Mapping
   +-------------+-------------+-----------+   +-------+    +--------+
   | Internal    |  External   |  Peer     |   | EID 1 |    | EID 1  |
   | IP:Port     |  IP:Port    |  IP:Port  |   |  -->  |    |  -->   |
   |-------------|-------------|-----------|   |1.8.7.2|    | RLOC N |
   |10.0.0.1:4341|1.8.7.2:20321|RLOC N:4342|   |:20321 |    +--------+
   +-------------+-------------+-----------+  *+-------+   *
                                             *          *
                                        ____*________*
             __________               ,'   *      *  `.
           ,' Non-LISP `.            /    *    *       \   __________
   +--------+   site +-------+      |  +------+         | /LISP site \
   |   MN   |        |  NAT  |      |  | NTR  |     +------+    +-----+
   |  EID 1 | <------|1.8.7.2|<--------| (MS) |<----| ITR  |<---| SN  |
   |10.0.0.1|    :   +-------+   :  |  |RLOC N|  :  |RLOC B|  : |EID 2|
   +--------+    :       /       :  |  +------+  :  +------+  : +-----+
           `.____:_____,'        :   \                 /  \___:______/
                 :               :    `._________:___,'       :
                 :               :               :            :
                 :               :               :      Src:     Dest:
                 :               :               :     +--------+------+
                 :               :               : IH: |EID 2   |EID 1 |
                 :               :               :     +--------+------+
                 :               :               :
         Src:     Dest:     Src:     Dest:      Src:     Dest:
        +--------+------+  +--------+------+   +--------+------+
    OH: |10.0.0.1|RLOC N|  |1.8.7.2 |RLOC N|   |RLOC B  |RLOC N|
        +--------+------+  +--------+------+   +--------+------+
   UDP: |4341    |4342  |  |20321   |4342  |   |30369   |4342  |
        +--------+------+  +--------+------+   +--------+------+
    IH: |EID 2   |EID 1 |  |EID 2   |EID 1 |   |EID 2   |EID 1 |
        +--------+------+  +--------+------+   +--------+------+
```

Figure 3: Incomin flow

When traffic is sent to a MN behind a NAT, a (P)ITR tunnels it to the NTR at which the MN has registered. This is depicted in Figure 3. An NTR relays such traffic as follows. It strips off the LISP and

UDP header, uses the destination EID (EID 1) in the IH of the packet to look up the IP/port combination (1.8.7.2:20321) in the EID-to-IP:Port table, and encapsulates the packets to this IP/port combination using its own IP address and port 4342 as source IP/port combination (RLOC N:4342). The NAT gateway recognizes the destination IP/port and translates it to the address:port of the MN which is 10.0.0.1:4341 in our example. Eventually, the translated packet reaches the MN on the correct port 4341 for incoming LISP-encapsulated traffic. The correct port number is achieved by requiring MNs to send Map-Register messages to the MS using source port 4341. Regarding the behavior of a MN, this constitutes the only difference to the original LISP-MN architecture.

6. Security Considerations

The presented NAT traversal for LISP MN allows other nodes in the Internet to contact MNs behind a NAT gateway which is the intention of the proposal. If the NAT is used as part of a firewall, external nodes can easily circumvent this security feature and contact MNs. However, this is a general concern of all NAT traversal mechanisms. Thus, any type of traffic can reach the MN behind a NAT/firewall. This may be avoided by making the NAT/firewall aware of the NAT traversal mechanism so that deep packet inspection for incoming LISP traffic can be used.

7. Conclusion

NAT traversal for LISP MNs allows MNs that roam into networks behind NATs to be globally reachable. The presented mechanism does not require new architectural components and implements new "NAT Traversal Router" (NTR) functionality only in MSs. The only change to a MN is that it must send Map-Register messages from source port 4341 to its MS.

8. Acknowledgements

The authors thank David Meyer, Darrel Lewis, Dino Farinacci, and Vince Fuller for insightful comments. They also acknowledge the support of G-LAB (support code 01 BK 0800, G-Lab, http://www.german-lab.de/).

9. IANA Considerations

This document makes no request on IANA namespaces [RFC2434].

10. References

10.1. Normative References

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

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

[RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

10.2. Informative References

[KLHA10] Klein, D., Hartmann, M., and M. Menth, "NAT Traversal for LISP Mobile Node", work in progress, April 2010, <http://www.menth.net/Publications/papers/Menth10-Sub-2.pdf>.

[LISP] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "Locator/ID Separation Protocol (LISP)", draft-ietf-lisp-07 (work in progress), April 2010.

[LISP-IW] Lewis, D., Meyer, D., Farinacci, D., and V. Fuller, "Interworking LISP with IPv4 and IPv6", draft-ietf-lisp-interworking-01 (work in progress), March 2010.

[LISP-MN] Farinacci, D., Fuller, V., Lewis, D., and D. Meyer, "LISP Mobile Node", draft-meyer-lisp-mn-01 (work in progress), February 2010.

[LISP-MS] Fuller, V. and D. Farinacci, "LISP Map Server", draft-ietf-lisp-ms-05 (work in progress), April 2010.

[MEKL10] Menth, M., Klein, D., and M. Hartmann, "Improvements to LISP Mobile Node", in proceedings of the 22nd International Teletraffic Congress (ITC), Amsterdam, The Netherlands, Sep 2010. Preprint available at: <http://www.menth.net/Publications/papers/Menth10k.pdf>

Authors' Addresses

Dominik Klein
University of Wuerzburg
Am Hubland
Wuerzburg D-97074
Germany

Phone: +49-931-31-88827
Email: dominik.klein@informatik.uni-wuerzburg.de

Matthias Hartmann
University of Wuerzburg
Am Hubland
Wuerzburg D-97074
Germany

Phone: +49-931-31-83381
Email: hartmann@informatik.uni-wuerzburg.de

Michael Menth
University of Wuerzburg
Am Hubland
Wuerzburg D-97074
Germany

Phone: +49-931-31-86644
Email: menth@informatik.uni-wuerzburg.de