LISP Working Group
Internet-Draft
Intended status: Experimental
Expires: 12 September 2025

D. Saucez, Ed.
Inria
L. Iannone, Ed.
Huawei
V. Ermagan
Google
D. Farinacci
lispers.net
D. Lewis
ICANN
F. Maino
M. Portoles Comeras
Cisco Systems, Inc.
J. Skriver
Arista
C. White
Logicalelegance, Inc.
A. Lopez
UPC/BarcelonaTech
11 March 2025

NAT traversal for LISP
draft-ermagan-lisp-nat-traversal-20

Abstract

This document describes a mechanism for IPv4 NAT traversal for LISP tunnel routers (xTR) and LISP Mobile Nodes (LISP-MN) behind a Network Address Translator (NAT) device. A LISP device both detects the NAT and initializes its state. Forwarding to the LISP device through a NAT is enabled by the LISP Re-encapsulating Tunnel Router (RTR) network element, which acts as an anchor point in the data plane, forwarding traffic from unmodified LISP devices through the NAT.

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 https://datatracker.ietf.org/drafts/current/.

Copyright Notice

Table of Contents

1. Introduction

The Locator/ID Separation Protocol [RFC9300] [RFC9301] defines a set of functions for encapsulating routers to exchange information used to map from Endpoint IDentifiers (EIDs) to routable Routing LOCators (RLOCs). The assumption that the LISP Tunnel Routers are reachable at their RLOC breaks when a LISP device is behind a Network Address Translator (NAT [RFC3022]). LISP relies on the xTR being able to receive traffic at its RLOC on destination port 4341. However, nodes behind a NAT are only reachable through the NAT's public address and in most cases only after the appropriate mapping state is set up in the NAT. Depending on the type of the NAT device, this mapping state may be address and port dependent. In other words, the mapping state in the NAT device may be associated with the 5-tuple that forms a specific flow, preventing incoming traffic from any LISP router other than the one associated with the 5-tuple. A NAT traversal mechanism is needed to make the LISP device behind a NAT reachable.

This document briefly discusses available NAT traversal options, and then it introduces in detail a NAT traversal mechanism specific for LISP. Two new LISP control messages, namely LISP Info-Request and LISP Info-Reply, are introduced in order to detect whether a LISP device is behind a NAT, and discover the global IP address and global ephemeral port used by the NAT to forward LISP packets sent by the LISP device. The LISP Re-encapsulating Tunnel Router (RTR) [RFC9300], acts as a re-encapsulating LISP tunnel router to pass traffic through the NAT, to and from the LISP device. A modification to how the LISP Map-Register messages are sent allows LISP device to initialize NAT state to use the RTR services. This mechanism addresses the scenario where the LISP device is behind the NAT, but the associated Map-Server [RFC9301] is on the public side of the NAT.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Definition of Terms

This documents assumes that the reader is familiar with LISP and the LISP terminology. For definitions of terms like Map-Request, Map-Reply, Ingress Tunnel Router (ITR), and Egress Tunnel Router (ETR), please consult the LISP specification in [RFC9300] and [RFC9301].

Basic NAT: See [RFC3022] and also Section 4.1.1 of [RFC2663].

Data Plane Encapsulated Control Message (DP-ECM): Is a LISP control message encapsulated with a LISP data plane header and defined in Section 6.4.

Info-Request: A LISP control message sent by a LISP device to its Map-Server to check whether it is behind a NAT and defined in Section 6.1.1.

Info-Reply: A LISP control message sent by a Map Server to a LISP device in response to an Info-Request control message and defined in Section 6.1.2.

NAPT Network Address Port Translation: See [RFC3022] and also Section 4.1.2 of [RFC2663].

Re-encapsulating Tunnel Router (RTR): As defined in [RFC9301].

Site-ID: As defined in [RFC9301].

xTR-ID: As defined in [RFC9301].

In this document the general term NAT is used to refer to both Basic NAT and NAPT.

4. NAT Traversal Basics

There are a variety of NAT devices and a variety of network topologies utilizing NAT devices in deployments. Most NAT devices deployed today are designed primarily around the client/server paradigm, where clients inside a private network do initiate connections to public servers with public IP addresses. As such, any protocol requiring a device or host, in a private network behind a NAT, to receive packets or accept sessions from public servers/hosts, without first initiating a session or sending packets towards those servers/hosts, will be challenged by deployed NAT devices.

NAT devices are loosely classified based on how restrictive they are. These classifications are essentially identifying the type of mapping state that the NAT device is requiring to allow incoming traffic.

For instance, the mapping state may be end-point independent, where once device A inside the private network sends traffic to a destination outside, a mapping state in the NAT is created that only includes information about device A, namely its IP address and perhaps its port number. Once this mapping is established in the NAT device, any external device with any IP address could send packets to device A. More restrictive NAT devices could include the 5-tuple information of the flow as part of the mapping state, in other words, the mapping state in the NAT is dependent upon source IP and port, as well as destination IP and port (symmetric NAT or endpoint-dependent NAT). Such a NAT only allows traffic from the specified destination IP and port to reach the specified source device on the specified source port. Traffic with a different 5-tuple signature will not be allowed to pass. In general, in the case of less restrictive NATs it may be possible to eventually establish direct peer-to-peer connections, by means of various hole punching techniques and initial rendezvous servers. However, in the case of symmetric NATs or NATs with endpoint-address-and-port-dependent mappings, direct connection may prove impossible. In such cases a relay device is required that is in the public network and can relay packets between the two endpoints.

Various methods have been designed to address NAT traversal challenges, mostly in the context of peer-to-peer applications and protocols. Among these, the Interactive Connectivity Establishment (ICE) [RFC8445] seems the most comprehensive, which defines a protocol that leverages other protocols such as Session Traversal Utilities for NAT (STUN) [RFC8489] and Traversal Using Relays around NAT (TURN) [RFC8656], as well as rendezvous servers to identify and exchange a list of potential transport (IP and port) addresses between the two endpoints. All possible pairs of transport addresses are exhaustively tested to find the best possible option for communication, preferring direct connections to connections using a relay. In the case of most restrictive NATs, ICE leads to use of TURN servers as relay for the traffic. TURN requires a list of allowed peer IP addresses defined as permissions, before allowing a peer to use the relay server to reach a TURN client.

Common NAT traversal techniques such as ICE generally assume bi-directional traffic with the same 5-tuple. LISP, however, requires traffic to use destination UDP port 4341, without specifying the source port. As a result, LISP traffic is generally unidirectional. This means that, in the case of symmetric or endpoint-address-and-port-dependent mapping NATs, even when an outgoing mapping is established, still incoming traffic may not match the established mapping and will not be allowed to pass. As a result, while ICE may be used to traverse less restrictive NATs, use of standard TURN servers as relays to traverse symmetric NATs for LISP protocol is not possible.

The rest of this document specifies a NAT traversal technique for the LISP protocol that enables LISP protocol to traverse multiple types of NATs including symmetric NATs.

5. LISP NAT Traversal Overview

There are two attributes of a LISP device behind a typical NAT that requires special consideration in LISP protocol behavior in order to make the device reachable. First, the RLOC assigned to the device is typically not globally unique nor globally routable. Hence, for NAT traversal, outbound packets are required, so to create state before the NAT accepts inbound packets. Second, LISP protocol requires an xTR to receive traffic on the specific UDP port 4341, so the random UDP port allocated by the NAT on its public side to associate with a xTR behind the NAT cannot be used by other xTRs to send LISP traffic. This section provides an overview of the LISP NAT traversal mechanism which deals with these conditions, while the rest of the document specifies the mechanism in details.

When a LISP device needs to register an RLOC in the mapping system, it needs to first discover whether the RLOC is behind a NAT. To do this, the ETR queries its Map-Server to discover the ETR's translated global RLOC and port, via two new LISP messages, namely Info-Request (see Section 6.1.1) and Info-Reply (see Section 6.1.1). Once the ETR detects that it is behind a NAT, it uses a Re-encapsulating Tunnel Router (RTR) as an anchor point for sending and receiving data plane traffic through the NAT device. The ETR registers the RTR RLOC(s) to its Map-Server using the RTR as a proxy for the Map-Register message. The ETR encapsulates the Map-Register message in a LISP Encapsulated Control Message (ECM) destined to the RTR's RLOC using 4341 as source port. The RTR strips the LISP ECM header and sends it to the Map-Server. This initializes state in the NAT device so that the ETR can receive traffic on port 4341 from the RTR. The ETR also registers the RTR RLOC as the RLOC where the ETR EID prefix is reachable. As a result, all packets destined to the ETR's EID will go to the registered RTR. The RTR will then re-encapsulate and forward the traffic, thanks to the existing NAT state, to the ETR.

Outbound LISP data traffic from the ITR can be sent directly to the external destinations, however this will create state on the NAT. To avoid excessive state on the NAT device, the ITR can encapsulate outgoing traffic to the RTR, where the RTR de-capsulates the LISP packets, and then re-encapsulates them or forwards them natively depending on their destination.

A complete and detailed example of LISP NAT traversal can be found in Appendix A.

6. LISP Messages Details

The main modifications in the LISP protocol to enable LISP NAT traversal via an RTR include:

1. two new messages used for NAT discovery, namely Info-Request and Info-Reply;

2. the encapsulation of the Map-Register, between the xTR and the RTR, in an ECM header;

3. the Data Plane Encapsulation of the ECM-encapsulated Map-Notify, between the RTR and the xTR;

This section describes the message formats and details of the Info-Request (Section 6.1.1), Info-Reply (Section 6.1.2), and DP-ECM Map-Notify (Section 6.3) messages, as well as minor changes to Map-Register and Map-Notify messages.

6.1. Info-Request/Info-Reply Messages Format

An ETR sends an Info-Request message to its configured Map-Server, to trigger an Info-Reply message, in order to detect whether there is a NAT device on the path to its Map-Server and to obtain a list of RTR RLOCs that can be used for LISP data plane NAT traversal.

The Info-Request and Info-Reply are actually one single LISP control message, which are distinguished by a bit that indicates whether the message is a request or a reply and the presence of a NAT LCAF [RFC8060] in the latter. The rest of the message is the same and has the format depicted in Figure 1.

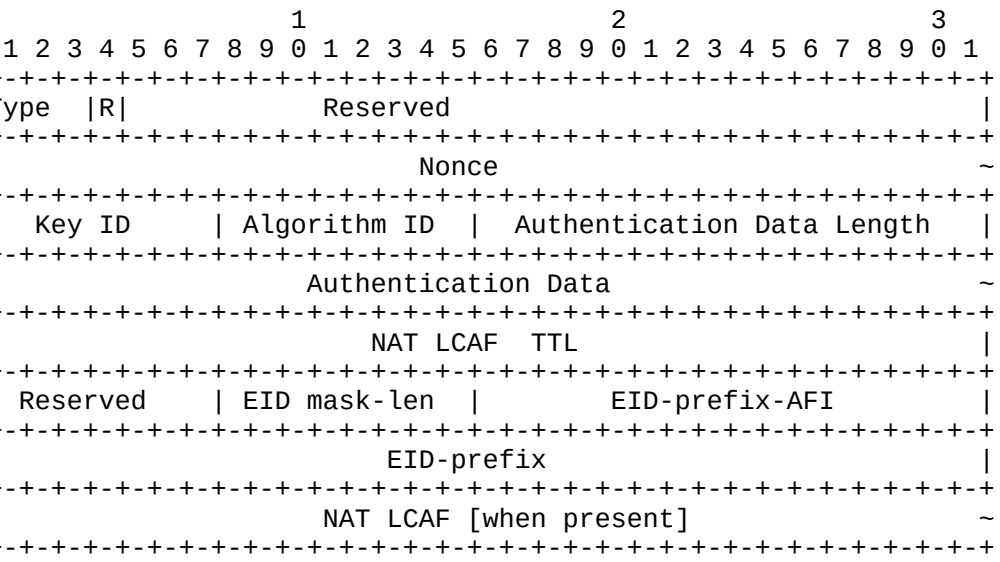

Figure 1: Generic Info-Request/Info-Reply Message Format.

Where:

Type: TBD (Info-Request/Info-Reply)

Reply (R): R bit indicates this is a Info-Reply to an Info-Request. R bit is set to 0 in an Info-Request. R bit is set to 1 in an Info-Reply.

Reserved: MUST be set to 0 on transmit and MUST be ignored on receipt.

NAT LCAF TTL: For Info-Request (R=0) this field MUST be set to 0 on

transmission and MUST be ignored on reception. For Info-Replies (R=1) this field expresses the time, in minutes, the recipient of the Info-Reply CAN store the RTR Information.

Nonce: As defined in Section 5.6 of [RFC9301].

Key ID: As defined in Section 5.6 of [RFC9301].

Algorithm ID: As defined in Section 5.6 of [RFC9301].

Authentication Data Length: As defined in Section 5.6 of [RFC9301].

Authentication Data: As defined in Section 5.6 of [RFC9301].

NAT LCAF: As defined in Section 4.4 of [RFC8060]. This field is only present in Info-Reply messages (R=1) indicating that the RLOC of the Info-Request is behind a NAT device (N=1).

The following sections describe in details the Info-Request and Info-Reply variants.

6.1.1. Info-Request Message

An Info-Request message is a LISP control message, its source port is chosen by the xTR and its destination port is set to the reserved LISP Control Packet port number 4342.

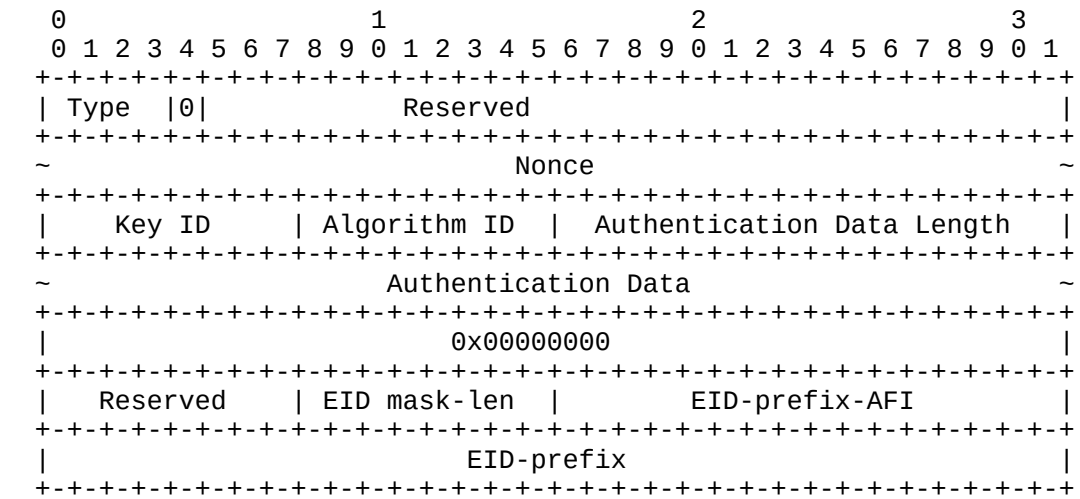

Figure 2: LISP Info-Request Message Format.

Type: TBD (Info-Request/Info-Reply)

Reply (R): MUST be set to 0 (zero).

NAT LCAF TTL: MUST be set to zero (0) on transmit and MUST be ignored on receipt.

The rest of the fields MUST be set in the same way as for a Map-Register message (see Section 5.6 of [RFC9301]) and according to procedures described in Section 7.

6.1.2. Info-Reply Message

The format of the Info-Reply message indicating the presence of a NAT device is depicted in Figure 3, its source port is chosen by the Map-Server and its destination port is set to the reserved LISP Control Packet port number 4342.

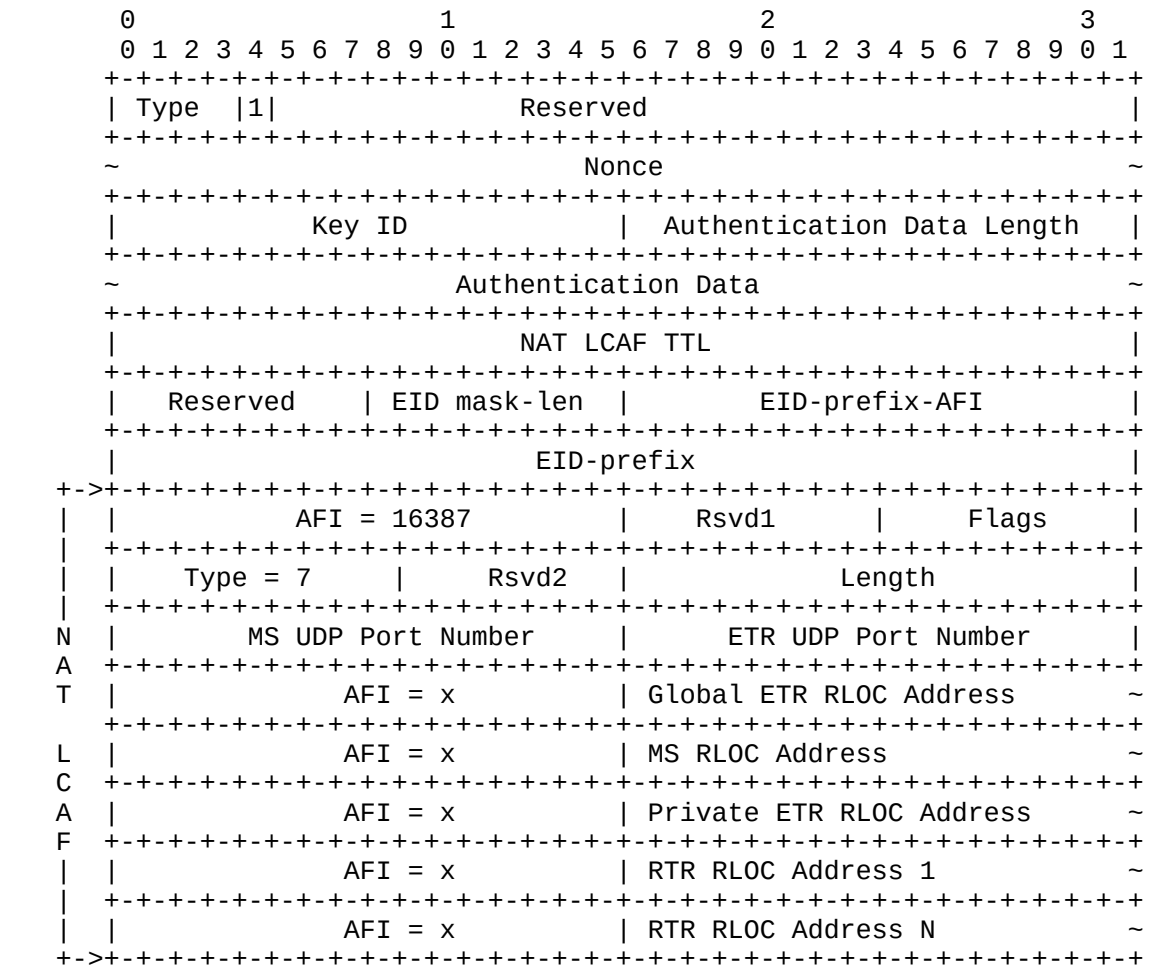

Figure 3: LISP Info-Reply Message Format.

Type: TBD (Info-Request/Info-Reply)

Reply (R): MUST be set to 1 (one).

NAT LCAF TTL: This is the time in minutes the recipient of the Info-Reply can store the NAT LCAF information. MUST be different from 0.

NAT LCAF: NAT LISP Canonical Address Format, as defined in [RFC8060].

The rest of the fields MUST be set in the same way as for a Map-Notify message (see Section 5.7 of [RFC9301]) and according to procedures described in Section 7.

6.2. Map-Register Message

A LISP device that sends a Map-Register to an RTR, MUST encapsulate the Map-Register message using an Encapsulated Control Message (ECM) [RFC9301] and MUST set the "M" bit in the ECM header to indicate that the message is destined to a Map-Server.

The outer header source RLOC of the ECM is set to the LISP device's private RLOC, and the outer header source port is set to 4341. The outer header destination RLOC and port are set to RTR RLOC and 4342 respectively. The inner header source RLOC is set to LISP device's private RLOC, and the inner source port is picked at random. The inner header destination RLOC is set to the xTR's Map-Server RLOC, and inner header destination port is set to 4342.

In case of LISP site having several xTRs, in order to identify the intended recipient xTR for a Map-Notify message the xTR-ID and Site-ID SHOULD be appended to the Map-Register message. If appended, the I-bit in the Map-Register message MUST be set to 1 (see [RFC9301]). An xTR can choose not to append the xTR-ID and Site-ID, in this case the RTR will consider as the intended recipient of the Map-Notify message the xTR identified by the source RLOC of the Map-Register message (i.e. the source address of the inner header of the ECM Map-Register).

6.3. Map-Notify Message

For a full description of all fields in the Map-Notify message refer to Map-Notify section in [RFC9301]. A Map-Server that sends a Map-Notify to an RTR encapsulates the Map-Notify message using an ECM with the "E" bit set to 1 to indicate that the inner message is destined to an ETR.

The outer header source RLOC of the ECM is set to the Map-Server RLOC, and the outer header source port is set to 4342. The outer header destination RLOC and port are set to RTR's RLOC and 4342 respectively. The inner header source RLOC is set to the Map-Server RLOC, and the inner source port is set to 4342. The inner header destination RLOC is set to the LISP device's private RLOC copied from the Map-Register, and inner header destination port is set to 4342.

If the Map-Register carried xTR-ID and Site-ID, the corresponding Map-Notify MUST set the I-bit to 1 and the xTR-ID and Site-ID from the Map-Register appended to the Map-Notify message.

6.4. Data Plane ECM Map-Notify Message

When an RTR receives an ECM Map-Notify message with the E-bit in the ECM header set to 1, it has to relay the Map-Notify to the registering LISP device. After removing the ECM header and processing the Map-Notify message as described in Section 7.3, the RTR encapsulates the Map-Notify first in an ECM and then in a LISP data plane header and sends it to the associated LISP device. This Map-Notify ECM inside a LISP data plane header is referred to as a Data Plane ECM Map-Notify message (or DP-ECM Map-Notify Message).

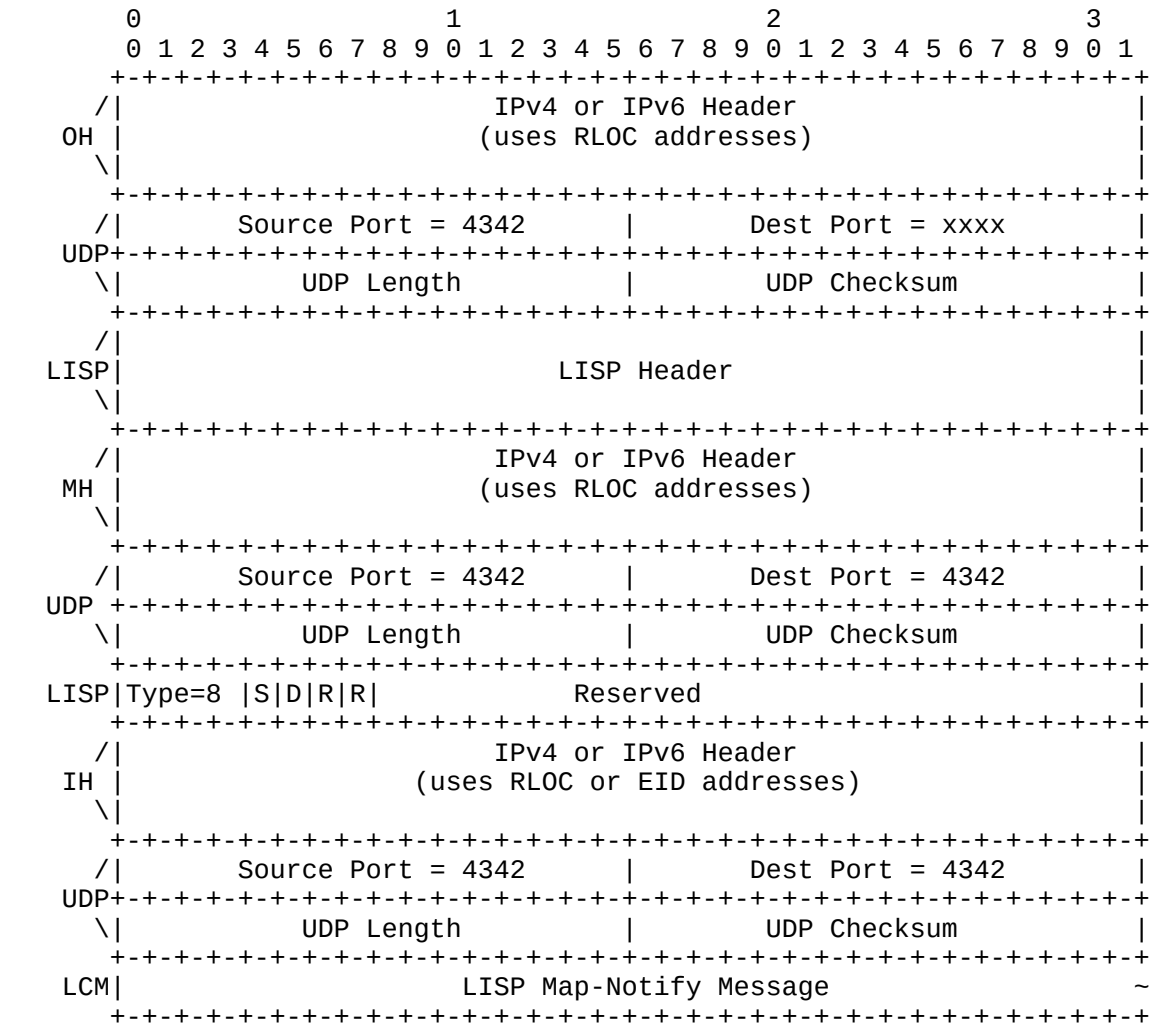

Figure 4: DP-ECM Map-Notify Message.

In a DP-ECM Map-Notify, the outer header source RLOC is set to the RTR's RLOC that was used in the associated ECM Map-Register. This is previously cached by the RTR. The outer header source port is set to 4342. The outer header destination RLOC and port are filled based on the translated global RLOC and port of the registering LISP device previously stored locally at the RTR. The middle ECM header (MH) source RLOC is set to the RTR RLOC, and the source port is set to 4342, the destination RLOC is set to the LISP device's private RLOC copied from the Map-Register, and destination port is set to 4342. The inner Map-Notify is unchanged, just forwarded, hence its header is as set by the Map-Server, namely the source address is the Map-Server's RLOC, the inner header source port is 4342, the destination address is the LISP device's private RLOC, and the destination port is 4342.

7. Protocol Operations

There are two main steps in the NAT traversal procedure. First, the ETR's translated global RLOC must be discovered. Second, the NAT translation table must be updated to accept incoming connections and the RTR must be informed of the ETR's translated global RLOC, including the translated ephemeral port number(s) at which the RTR can reach the LISP device.

7.1. xTR Processing

If an ETR is configured to perform NAT discovery and traversal, when it needs to register an RLOC in the mapping system, it has first to detect whether the RLOC is behind a NAT device. For this purpose, the ETR sends an Info-Request message to its Map-Server in order to discover the ETR's translated global RLOC as it is visible to the Map-Server. The ETR uses the private to-be-registered RLOC as the source RLOC of the message. The Map-Server, after authenticating the message, responds with an Info-Reply message. The Map-Server includes the source RLOC and port from the Info-Request message in the Global ETR RLOC Address and ETR UDP Port Number fields of the NAT-Traversal LCAF appended to the Info-Reply. The Map Server also includes the destination RLOC and port number of the Info-Request message in the MS RLOC Address and MS UDP Port Number fields of the NAT LCAF. In addition, the Map-Server provides the list of RTR RLOCs that the ETR may use for NAT traversal services. The source port of the Info-Reply is set to 4342 and the destination port is copied from the source port of the triggering Info-Request message.

Upon receiving and authenticating the Info-Reply message, checks whether the ETR compares the source RLOC and source port used for the Info-Request message with the Global ETR RLOC Address and ETR UDP Port Number fields in the NAT LCAF in the Info-Reply message. If the

two are identical, indicating there is no NAT on the path identified by an info-Request and an Info-Reply, the ETR registers the associated RLOC with its Map-Server as described in [RFC9301]. Otherwise, the ETR concludes that the to-be-registered RLOC is a private address behind a NAT and that it requires an RTR for NAT traversal services in order to be reachable at that RLOC. The Info-Reply will contain a NAT LCAF, whose content the ETR MUST register in its local EID-to-RLOC Database. This entry CAN be stored the number of minutes indicated in the NAT LCAF TTL field of the Info-Reply message. If NAT is detected but the NAT LCAF does not contain any RTR, then NAT traversal cannot be accomplished. The ETR still stores the content in the EID-to-RLOC Database for up to NAT LCAF TTL minutes, but it MUST also create a log message to signal that the RLOC cannot be registered because of an empty RTR set and MUST NOT use the RLOC for LISP operation.

In the case where an xTR has multiple RLOCs, Info-Requests SHOULD be sent per each RLOC, so to perform NAT discovery each RLOC. NAT discovery can be disabled in deployments where it is known to not include NAT devices. Unless the xTR is not willing to receive traffic via all RLOCs, NAT traversal SHOULD accomplished per each RLOC that has been detected as being a private address behind a NAT. This is required to establish the state in the NAT device for that RLOC.

It is worth noting that a STUN MAY also be used to do NAT detection and to discover the NAT-translated public IP address and port number for the ETR behind NAT. If STUN is used, the list of RTR devices that can be used by the xTR for NAT traversal have to be provisioned via other means which are outside the scope of this document.

7.1.1. ETR Processing

Once an ETR has detected its RLOC is behind a NAT, based on local policy, the ETR selects one (or more) RTR(s), from the RTR RLOCs provided in the Info-Reply message, as both control and data plane proxy. Next it needs to initializes state in the NAT in order to receive LISP data traffic on UDP port 4341 from the selected RTR(s). To do so, the ETR sends an ECM-encapsulated Map-Register to the selected RTR(s). The Map-Register message is created as specified in [RFC9301]. More specifically, the source RLOC of the Map-Register is set to ETR’s private RLOC, while the destination RLOC is set to the ETR’s Map-Server RLOC, and destination port is set to 4342. The ETR MUST set the P bit (proxy Map-Reply) and the M bit (want-Map-Notify) in Map-Register to 1, and include the selected RTR RLOC(s) as the locators in the Map-Register message.

The ETR MAY also include its private RLOCs as locators in the Map-Register, including weight and priorities, but MUST set the R bit (reachability bit) in the locator record to 0. This enables the RTR to perform load balancing when forwarding data to an xTR with several RLOCs behind a NAT. The R bit MUST be set to 1 for all RTR locators included in the Map-Register. The ETR MAY also set the I bit in the Map-Register message to 1 and include its xTR-ID and Site-ID in the corresponding fields. In the ECM header of this Map-Register, the source RLOC is set to ETR's private RLOC and the source port is set to 4341, while the destination RLOC is the RTR's RLOC and the destination port is set to LISP control port 4342. The M bit in the ECM header MUST be set to 1, to indicate that this ECM Map-Register is to be forwarded to a Map-Server.

This ECM Map-Register is then sent to the RTR, its processing is described in Section 7.3.

Upon receiving DP-ECM Map-Notify from the RTR, the ETR MUST strip the outer LISP data header, and process the inner ECM Map-Notify message as described in [RFC9301]. When decapsulating, in accordance to [RFC9300], the ETR checks its EID-to-RLOC Database. In case of a DP-ECM, the inner header destination address is not an EID but the same private RLOC as the outer header, yet, since this RLOC is part of a NAT LCAF and the source port is 4342, the inner messages is considered a LISP control message and processed according to [RFC9301].

If ETR did send its xTR-ID and Site-ID in the Map-Register message and receives a DP-ECM Map-Notify with different xTR-ID and/or Site-ID, it MUST log this as an error. The ETR MUST discard such DP-ECM Map-Notify message.

At this point the registration and state initialization is complete and the xTR can use the RTR for both control and data plane proxy. The state created in the NAT is used by the xTR behind the NAT to send and receive LISP control packets to/from the RTR, as well as for receiving LISP data packets form the RTR.

The ETR MUST periodically send ECM Map-Register messages to its RTR in order to both refresh its registration to the RTR and the Map-Server as for [RFC9301], as well as a keep-alive to preserve the state in the NAT device. [RFC2663] points out that the period for sending the keep-alive messages can be set to a default value of two minutes, however since shorter timeouts may exist in some NAT deployments, the interval for sending periodic ECM Map-Registers has to be configured accordingly.

When a Map-Request for a LISP device behind a NAT is received by its Map-Server, the Map-Server responds with a Map-Reply including RTR's RLOC as the locator for the requested EID. As a result, all LISP data traffic destined for the ETR's EID behind the NAT is encapsulated to its RTR. The RTR re-encapsulates the LISP data packets to the ETR's translated global RLOC and port number so the data can pass through the NAT device and reach the ETR. As a result, the ETR receives LISP data traffic with outer header destination port set to 4341 as specified in [RFC9301].

7.1.2. ITR Processing

If ITRs have only private RLOCs, they cannot send Map-Request and receive Map-Reply messages for EID mapping lookups directly, since the Map-Requests (creating state in the NAT) are sent to Map-Resolvers, but Map-Replies come back from ETRs or Map-Servers, for which there is no state in the NAT device, which will drop the packets.

The ITR behind a NAT can use its RTR(s) RLOC(s) as locator(s) for all destination EIDs that it wishes to send data to. The ITR encapsulates the LISP traffic in a LISP data header with outer header destination set to RTR RLOC and outer header destination port set to 4341. This creates a secondary state in the NAT device. It is RECOMMENDED that the ITR sets the outer header source port in all egress LISP data packets to a random but static port number in order to avoid creating excessive state in the NAT device. By using the RTR as proxy, the ITR does not need to send Map-Requests for finding EID-to-RLOC mappings. However, if the ITR is multi-homed and has at least one RLOC not behind a NAT, it can choose to send Map-Requests using non-private RLOCs. For this, the ITR specifies in the ITR-RLOC field of the Map-Request the list of RLOCs that are not behind NAT that can receive Map-Reply messages. If all RLOCs of an ITR are behind the NAT and use the same RTR, then the xTR can even map the EID prefix 0/0 to its RTR RLOC(s) in its EID-to-RLOC Map-Cache.

It should be noted that sending packets directly to destination RLOCs through the interface behind NAT will result in creating additional state in the NAT device. Furthermore, outgoing packets use a direct path while the incoming packets are forwarded through the RTR.

Periodic ECM Map-Register and corresponding DP-ECM Map-Notify messages between xTR and RTR, can serve the purpose of RLOC probes as of [RFC9301], except that the LISP device behind a NAT only can probe the RTR's RLOCs.

If the ITR and ETR of a site are not collocated, the RTR RLOC need to be configured in the ITR via an out-of-band mechanism. Other procedures specified here would still apply.

7.2. Map-Server Processing

When a Map-Server receives an Info-Request message, it responds with an Info-Reply message by copying back the same message, but setting the R bit. Furthermore, it appends the NAT LCAF with the mapping between private and public addresses. Map-Server fills the NAT LCAF (LCAF Type = 7) fields according to Section 4.4 of [RFC8060], in particular, it copies the source RLOC and port number of the Info-Request message to the Global ETR RLOC Address and ETR UDP Port Number fields of the NAT LCAF and includes a list of RTR RLOCs that the ETR may use for NAT traversal services for the EID-Prefix communicated in the Info-Request message. In case that there is no list of RTRs in the NAT LCAF, this means that there is no NAT traversal service for the specific EID. The Info-Reply message source port is 4342, and destination port and destination address is taken from the source port and source address of the triggering Info-Request.

Upon receiving an ECM Map-Register message, with the M bit set, the Map-Server processes the inner Map-Register message and generates the resulting Map-Notify as described in [RFC9301]. If the I bit is set in the Map-Register message, the Map-Server also locally stores the xTR-ID and Site-ID from the Map-Register, and sets the I bit in the corresponding Map-Notify message and includes the same xTR-ID and Site-ID in the Map-Notify. The Map-Server then encapsulates the Map-Notify in an ECM header and sets the E bit in the ECM header to 1. This indicates that the ECM Map-Notify is to be processed by an RTR and forwarded to an ETR. The ECM Map-Notify is then sent to the RTR RLOC and port number from which the ECM Map-Register has been received.

7.3. RTR Processing

7.3.1. RTR Control Plane Operations

Upon receiving an ECM Map-Register with the M bit set in the ECM header, the RTR creates an EID-to-RLOC Map-Cache entry for the EID-prefix that is specified in the inner Map-Register message. The EID-to-RLOC Map-Cache entry MUST include the source RLOC, the source port number, which are the translated global RLOC and port number visible to the RTR, the destination RLOC (RTR’s own RLOC) of the outer header, as well as the source RLOC (xTR’s private RLOC), the EID mapping record present in the inner Map-Register message (see [RFC9301] for details), and, if present, the xTR-ID and the Site-ID.

The RTR can later use these fields as for sending DP-ECM Map-Notify back to the ETR. The Nonce field MUST also be stored and used for security purposes and is matched with the Nonce field in the corresponding Map-Notify message. This EID-to-RLOC Map-Cache entry is marked as "pending", until the corresponding Map-Notify message is received and in this state MUST NOT be used to send data packets.

In the case where the xTR has multiple RLOCs behind the NAT, the RLOCs with reachability bit set to zero (R=0) in the record from Map-Register, the RTR CAN store all the private RLOCs in the record from the Map-Register message, but MUST NOT use private RLOCs for which no explicit ECM Map-Register using it as source address in the inner header has been received. This is because without an ECM Map-Register there is no state in the NAT device that allow correct delivery of the returning packets. Furthermore, the EID-to-RLOC Map-Cache entry does not contain the translated global RLOC and port number visible to the RTR. However, an ECM Map-Register message MAY be used to update priority and weight of private RLOCs for which an ECM Map-Register has been already received, even if the message is not originating from that RLOC but originated from the same xTR, identified by the unique xTR-ID.

A Map-Register originating from a unique xTR-ID will always overwrite previously stored information for that xTR-ID, but it does not modify the information indicated by any other xTR-ID serving the same EID prefix. As a result, in the case of a renumbering or xTR reboot, the xTR can use its unique xTR-ID in a Map-Register, overwriting the previously stored information for that xTR. Using this method, the xTR can, for instance, immediately remove any private RLOC from the RTR EID-to-RLOC Map-Cache that is no longer present in the locator record of the ECM Map-Register, for which NAT state is most probably non-existing anymore.

After filling the local EID-to-RLOC Map-Cache entry, the RTR strips the outer header and extracts the Map-Register message, encapsulated in a new ECM header with the E bit set to 0, and M bit set to 1, and sends the ECM Map-Register to destination Map-Server. Map-Server responds with an ECM Map-Notify message to the RTR as described Section 7.2.

Upon receiving an ECM Map-Notify message with E bit set to 1 in the ECM header and the Nonce of a pending EID-to-RLOC Map-Cache entry, the RTR is notified that the Map-Register message was accepted by the Map-Server. The RTR MUST verify that the returned information is the same as seen in the ECM Map-Register, if not the message MUST be silently dropped. If the information is correct, at this point the RTR can change the state of the associated EID-to-RLOC Map-Cache entry to "active" for the duration of the TTL indicated in the Map-Notify message.

The RTR then uses the information in the associated EID-to-RLOC Map-Cache entry to create a DP-ECM Map-Notify message, where the outer header destination RLOC and port number are set to the ETR’s translated global RLOC and port number. If more than one ETR translated RLOC and port exists in the EID-to-RLOC Map-Cache entry for the same EID prefix specified in the Map-Notify, the RTR uses the xTR-ID from the Map-Notify to identify which ETR is the correct destination for the DP-ECM Map-Notify. The RTR sets the LISP data plane outer header source RLOC to RTR’s RLOC from the EID-to-RLOC Map-Cache entry and the outer header source port is set to 4342. Then it set the ECM header using source and destination port 4342, destination address the ETR private RLOC and as source address the RTR RLOC. Finally, the RTR sends the DP-ECM Map-Notify to the ETR.

LISP defines several mechanisms to signal updated mappings in the data-plane [RFC9300] and in the control plane [RFC9301]. If the Map-Register / Map-Notify modifies and existing EID-to-RLOC Map-Cache entry, the RTR can use these mechanisms to signal the change to device using the RTR to send traffic to the xTRs behind a NAT device.

7.3.2. RTR Data Plane Forwarding

An RTR processes LISP data plane packets according to [RFC9300]. In the case where the destination EID is a previously registered EID behind a NAT device, the RTR strips the LISP data header and re-encapsulate the packet in a new LISP data header. The outer header RLOCs and UDP ports are then filled based on the matching EID-to-RLOC Map-Cache entry for the associated destination EID prefix. The RTR uses the RTR RLOC from the EID-to-RLOC Map-Cache entry as the outer header source RLOC. The outer header source port is set to 4342. The RTR sets the outer header destination RLOC and outer header destination port based on the ETR translated global RLOC and port stored in the Map-Cache entry. Then the RTR forwards the LISP data packet.

8. Security Considerations

These specifications leverage on mechanism defined in [RFC9300], [RFC9301], and [RFC8060], as such security considerations in those documents apply as well here.

To protect origin authentication and integrity of control plane messages the LISP-SEC [RFC9303] MUST be used.

Info-Request and Info-Reply messages have similar header structure of Map-Register and Map-Notify, as such authentication and integrity can be performed in the same way, hence the procedures described in Section 5.6 of [RFC9301] MUST be used with a pre-shared key between the xTR and the Map-Server.

Similarly, Map-Register and Map-Notify MUST be also be verified for authentication and integrity using pre-shared key between the xTR and the Map-Server as for [RFC9301].

The pre-shared key for authentication and integrity verification of the Info-Request/Info-Reply messages SHOULD be different from the pre-shared key used for authentication and integrity verification of the Info-Request/Info-Reply messages.

For the authentication and integrity protection of the RTR, the Encapsulated Control Message LISP-SEC Extensions defined in Section 6.1 of [RFC9303] MUST be used, employing two different pre-shared keys, one between the RTR and the ETR and one between the RTR and the Map-Server. [RFC9303] defines how to protect the ECM-encapsulated Map-Request/Map-reply messages, however, the exact same operations can used to protect ECM-encapsulated Map-Register/Map-Notify messages. In this way, ETR and RTR can authenticate each other and verify messages' integrity. RTR and Map-Server can do exactly the same, authenticating each other and verifying messages' integrity.

Concerning the shared key among the various LISP entities, considerations in Section 7.5 of [RFC9303] apply.

Mapping lookups, through Map-Request/Map-Reply messages exchange, performed by the RTR(s) and the xTR(s) MUST be protected using LISP-SEC [RFC9303].

The RTR MUST re-encapsulate traffic only when the source or the destination are EIDs registered with the procedures described in this document and authenticated using LISP-SEC [RFC9303], which protects against the adverse use of an RTR for EID spoofing.

9. IANA Considerations

IANA is requested to allocate one codepoint from the "LISP Packet Types" registry, to be used to identify Info-Request/Info-Reply messages. The new entry should be defined as in Table 1.

Code	Message	Reference
7 (suggested)	LISP Info-Request/Info-Reply	[This Document]

Table 1

10. Acknowledgments

The authors would like to thank Noel Chiappa, Alberto Rodriguez Natal, Lorand Jakab, Dominik Klein, Matthias Hartmann, and Michael Menth for their previous work, feedback and helpful suggestions.

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/rfc/rfc2119>.

[RFC2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, DOI 10.17487/RFC2663, August 1999, <https://www.rfc-editor.org/rfc/rfc2663>.

[RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, DOI 10.17487/RFC3022, January 2001, <https://www.rfc-editor.org/rfc/rfc3022>.

[RFC8060] Farinacci, D., Meyer, D., and J. Snijders, "LISP Canonical Address Format (LCAF)", RFC 8060, DOI 10.17487/RFC8060, February 2017, <https://www.rfc-editor.org/rfc/rfc8060>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/rfc/rfc8174>.

[RFC9300] Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos, Ed., "The Locator/ID Separation Protocol (LISP)", RFC 9300, DOI 10.17487/RFC9300, October 2022, <https://www.rfc-editor.org/rfc/rfc9300>.

[RFC9301] Farinacci, D., Maino, F., Fuller, V., and A. Cabellos, Ed., "Locator/ID Separation Protocol (LISP) Control Plane", RFC 9301, DOI 10.17487/RFC9301, October 2022, <https://www.rfc-editor.org/rfc/rfc9301>.

[RFC9303] Maino, F., Ermagan, V., Cabellos, A., and D. Saucez, "Locator/ID Separation Protocol Security (LISP-SEC)", RFC 9303, DOI 10.17487/RFC9303, October 2022, <https://www.rfc-editor.org/rfc/rfc9303>.

11.2. Informative References

[RFC8445] Keranen, A., Holmberg, C., and J. Rosenberg, "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal", RFC 8445, DOI 10.17487/RFC8445, July 2018, <https://www.rfc-editor.org/rfc/rfc8445>.

[RFC8489] Petit-Huguenin, M., Salgueiro, G., Rosenberg, J., Wing, D., Mahy, R., and P. Matthews, "Session Traversal Utilities for NAT (STUN)", RFC 8489, DOI 10.17487/RFC8489, February 2020, <https://www.rfc-editor.org/rfc/rfc8489>.

[RFC8656] Reddy, T., Ed., Johnston, A., Ed., Matthews, P., and J. Rosenberg, "Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN)", RFC 8656, DOI 10.17487/RFC8656, February 2020, <https://www.rfc-editor.org/rfc/rfc8656>.

Appendix A. NAT Traversal Example

In what follows there is an example of an ETR initiating a registration of a new RLOC to its Map-Server, when the RLOC is behind a NAT (Appendix A.1). A second example shows how data communication between a LISP site behind a NAT and a second LISP site works (Appendix A.2).

A.1. EID Prefix Registration Example

In this example, the ETR (Site1-ETR) is configured with the local RLOC of 172.16.1.2. The NAT's global (external) addresses are from 192.0.2.1/25 prefix. The Map-Server is at 203.0.113.169. And one potential RTR has an IP address of 203.0.113.1. The Site1-ETR has an EID Prefix of 198.51.100.0/24.

The steps to register the EID-Prefix of Site1-ETR in the MAP Server are as follows:

1. Site1-ETR receives the private IP address, 172.16.1.2 as its RLOC, for instance via DHCP.

2. Site1-ETR sends an Info-Request message with the destination RLOC of the Map-Server, 203.0.113.169, and source RLOC of 172.16.1.2. This packet has the destination port set to 4342 and the source port is set to (for example) 5001.

3. The NAT device translates the source IP from 172.16.1.2 to 192.0.2.1, and source port to (for example) 20001 global ephemeral source port.

4. The Map-Server receives and responds to this Info-Request with an Info-Reply message. This Info-Reply has the destination address set to Site1-ETR's translated address of 192.0.2.1 and the source address is the Map-Server's RLOC, namely 203.0.113.169. The destination port is 20001 and the source port is 4342. Map-Server includes a copy of the source address and port of the Info-Request message (192.0.2.1:20001), and a list of RTR RLOCs including RTR RLOC 203.0.113.1 in the Info-Reply contents.

5. The NAT translates the Info-Reply packet's destination IP from 192.0.2.1 to 172.16.1.2, and translates the destination port from 20001 to 5001, and forwards the Info-Reply to Site1-ETR at 172.16.1.2.

6. Site1-ETR detects that it is behind a NAT by comparing its local RLOC (172.16.1.2) with the Global ETR RLOC Address in the Info-Reply (192.0.2.1) . Then, Site1-ETR picks the RTR 203.0.113.1 from the list of RTR RLOCs in the Info-Reply. Site1-ETR stores the RTR RLOC in a default EID-to-RLOC Map-Cache entry to periodically send ECM Map-Registers to.

7. The ETR sends an ECM encapsulated Map-Register to RTR at 203.0.113.1. The outer header source RLOC of this Map-Register is set to 172.16.1.2 and the outer header source port is set to

4341. The outer header destination RLOC and port are set to RTR RLOC at 203.0.113.1 and 4342 respectively. The M bit in ECM header is set to 1. The inner header destination RLOC is set to ETR's Map-Server 203.0.113.169, and the inner header destination port is set to 4342. The inner header source RLOC is set to ETR's local RLOC 172.16.1.2 and the source port is set to (for example) 5002. In the Map-Register message the RTR RLOC 203.0.113.1 appears as the locator set for the ETR's EID prefix (198.51.100.0/24). In this example ETR also sets the Proxy bit in the Map-Register to 1, and sets I bit to 1, and includes its xTR-ID and Site-ID in the Map-Register.

8. The NAT translates the source RLOC in the ECM header of the Map-Register, by changing it from 172.16.1.2 to 192.0.2.1, and translates the source port in the ECM header from 4341 to (for example) 20002, and forwards the Map-Register to RTR.

9. The RTR receives the Map-Register and creates an EID-to-RLOC Map-Cache entry with the ETR's xTR-ID, EID prefix, and the source RLOC and port of the ECM header of the Map-Register as the locator (198.51.100.0/24 is mapped to 192.0.2.1:20002). RTR also caches the inner header source RLOC of the Map-Register namely 172.16.1.2, and the outer header destination RLOC of the ECM header in the Map-Register (this would be RTR's RLOC 203.0.113.1) to use for sending back a DP-ECM Map-Notify. RTR then removes the outer header, adds a new ECM header with M=0, and E=1, and forwards the Map-Register to the destination Map-Server.

10. The Map-Server receives the ECM Map-Register with N bit set to 1, removes the ECM header, and processes it according to [RFC9301]. After registering the ETR, Map-Server responds with an ECM Map-Notify with the E bit set to 1 in the ECM header. Since the I bit is set in the Map-Register, the Map-Server also sets the I bit in the Map-Notify and copies the xTR-ID and Site-ID from the Map-Register to the Map-Notify. The source address of this Map-Notify is set to 203.0.113.169. The destination is copied from the local source address of the Map Register (172.16.1.2), and both source and destination ports are set to 4342.

11. The RTR receives the ECM Map-Notify. It decapsulated the message stripping the ECM header, then re-encapsulates in a new ECM header and a LISP data plane header and sends the resulting DP-ECM Map-Notify to Site1-ETR with a matching xTR-ID. The outer header source RLOC and port of the DP-ECM Map-Notify are set to 203.0.113.1:4342. The outer header destination RLOC and port are retrieved from previously cached Map-Cache entry in

step 9, namely 192.0.2.1:20002. The ECM header source RLOC and port of are set to 203.0.113.1:4342, while destination RLOC and port are the local RLOC 172.16.1.2 and 4342 respectively. At this point RTR marks ETR's EID prefix as "active" status and forwards the DP-ECM Map-Notify to ETR.

12. The NAT device translates the destination RLOC and port of the Data-Map-Notify to 172.16.1.2:4341 and forwards the packet to ETR.

13. The Site1-ETR receives the packet with a destination port 4341, and processes the packet as a control packet after observing the outer destination address and the inner destination address is a private RLOC belonging to a NAT LCAF. At this point ETR's registration to the RTR is complete.

A.2. Data Communication example

Assume a requesting ITR in a second LISP site (Site2-ITR) has an RLOC of 192.0.2.129/25. The following is an example process of an EID behind Site2-ITR sending a data packet to an EID behind the Site1-ETR.

1. The ITR sends a Map-Request which arrives via the LISP mapping system to the ETR's Map Server.

2. The Map-Server sends a Map-Reply on behalf of the ETR, using the RTR's RLOC (203.0.113.1) in the Map-Reply's Locator Set.

3. The ITR encapsulates a LISP data packet with ITR's local RLOC (192.0.2.129/25) as the source RLOC and the RTR as the destination RLOC (203.0.113.1) in the outer header.

4. The RTR decapsulates the packet, evaluates the inner header against its EID-to-RLOC Map-Cache and then re-encapsulates the packet. The new outer header's source RLOC is the RTR's RLOC 203.0.113.1 and the new outer header's destination RLOC is the Global NAT address 192.0.2.1. The destination port of the packet is set to 20002 (discovered above during the registration phase) and the source port is 4342.

5. The NAT translates the LISP data packet's destination IP from to 192.0.2.1 to 172.16.1.2, and translates the destination port from 20002 to 4341, and forwards the LISP data packet to the ETR at 172.16.1.2.

6. For the reverse path the ITR uses its local EID-to-RLOC Map-Cache entry with the RTR RLOC as the default locator and encapsulates the LISP data packets using RTR RLOC, and 4341 as destination RLOC and port.

Contributors

Albert Cabellos
UPC/BarcelonaTech
Email: acabello@ac.upc.edu

Authors' Addresses

Damien Saucez (editor)
Inria
France
Email: damien.saucez@inria.fr

Luigi Iannone (editor)
Huawei Technologies France S.A.S.U.
18, Quai du Point du Jour
92100 Boulogne-Billancourt
France
Email: luigi.iannone@huawei.com

Vina Ermagan
Google
United States of America
Email: ermagan@gmail.com

Dino Farinacci
lispers.net
Email: farinacci@gmail.com

Darrel Lewis
ICANN
Los Angeles, CA 90292
United States of America
Email: darrel.lewis@icann.org

Fabio Maino
Cisco Systems, Inc.

Email: fmaino@cisco.com

Marc Portoles Comeras
Cisco Systems, Inc.
Email: mportole@cisco.com

Jesper Skriver
Arista
Email: jesper@skriver.dk

Chris White
Logicalelegance, Inc.
Email: chris@logicalelegance.com

Albert Lopez
UPC/BarcelonaTech
Email: alopez@ac.upc.edu