

LISP Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 16, 2014

L. Cheng
M. Sun
ZTE Corporation
July 15, 2013

draft-cheng-lisp-nat-traversal-extension-01
Extension to LISP NAT Traversal Proposal

Abstract

This draft specifies several special NAT traversal scenarios when two or more LISP Sites/MNs which locate behind the same NAT equipment communicate with each other. When these LISP Sites/MNs communicate with each other, it may cause routing latency and will increase re-encapsulation load on Re-encapsulating Tunnel Routers(RTRs) based on existing LISP-NAT strategy.

In this draft, we give detail descriptions of these scenarios. Also we propose some suggestions to solve the problems. According to our strategy, a new kind of message is used for RTRs to send relative information of Corresponding Sites/MNs to xTRs.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 16, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

| | |
|---|---|
| 1. Introduction | 2 |
| 2. Definition of Terms | 3 |
| 3. Issue Statement | 3 |
| 3.1. Basic Scenario-Single Level NAT | 3 |
| 3.2. Extended Scenario I-Multiple Levels NAT | 4 |
| 3.3. Extended Scenario II-Multiple Levels NAT | 6 |
| 4. Solutions | 7 |
| 4.1. RTR Processing | 7 |
| 4.2. ITR Processing | 8 |
| 4.3. ETR Processing | 8 |
| 5. Security Considerations | 8 |
| 6. IANA Considerations | 8 |
| 7. Normative References | 8 |
| Authors' Addresses | 8 |

1. Introduction

Locator/ID Separation Protocol [RFC6830] defines a set of functions for encapsulating routers to exchange information used to map from Endpoint Identifiers (EIDs) to routable Routing Locators (RLOCs). When a LISP site locates behind a NAT, LISP Tunnel Routers (xTRs) are only reachable through the NAT's public address which broke the assumption that xTRs are reachable at their RLOCs.

To make sure LISP devices locate behind a NAT reachable, [LISP-NAT] proposes a NAT traversal mechanism for LISP. To achieve this, an ETR of a LISP site will use its Map-Server to discover whether it is behind a NAT and to get its translated global RLOC and port via two LISP messages: Info-Request and Info-Reply. Once an ETR detects it locates behind a NAT, it use a LISP Re-encapsulating Tunnel Router (RTR) to act as a data plane 'anchor point' to send and receive traffic through the NAT device.

According to RTR proxy strategy introduced in [LISP-NAT], when an ITR behind a NAT needs to encapsulate outbound LISP traffic, it does not send Map-Request for destination EIDs, but just use RTR RLOC as locator for all destination EIDs that it wishes to send data to.

However, in some special scenarios such like two or more LISP Sites/MNs locate behind the same NAT, when these LISP Sites/MNs communicate with each other, this RTR proxy strategy will cause routing latency and extra decapsulation/re-encapsulation cost on RTR. In this draft, we list and describe several special cases. A new kind of message "Info-Notify" is adopted. This message is used by RTR to notify the LISP Site with relative information of its Corresponding Site which locates behind the same NAT with it. Based on this "Info-Notify" message, some feasible solutions are proposed.

2. Definition of Terms

This draft uses terms defined in [RFC6830] and [LISP-NAT]. This section introduces some new terms used in the document.

Corresponding Site/MN When two LISP Sites/MNs communicate with each other, each LISP Sites/MN is considered as Corresponding Site/MN of the other one, regardless of Site's/MN's status, i.e. it is a sender or a receiver.

Info-Notify A message used by RTR to notify LISP Site with relative information of its Corresponding Site.

3. Issue Statement

In this section, several scenarios are specified.

3.1. Basic Scenario-Single Level NAT

As shown in Fig.1, there are two LISP Sites (Site1 and Site2) locate behind the same Nat equipment (NAT1). Furthermore, xTRs of the two sites choose the same RTR as proxy to register mapping information and transfer data traffic.

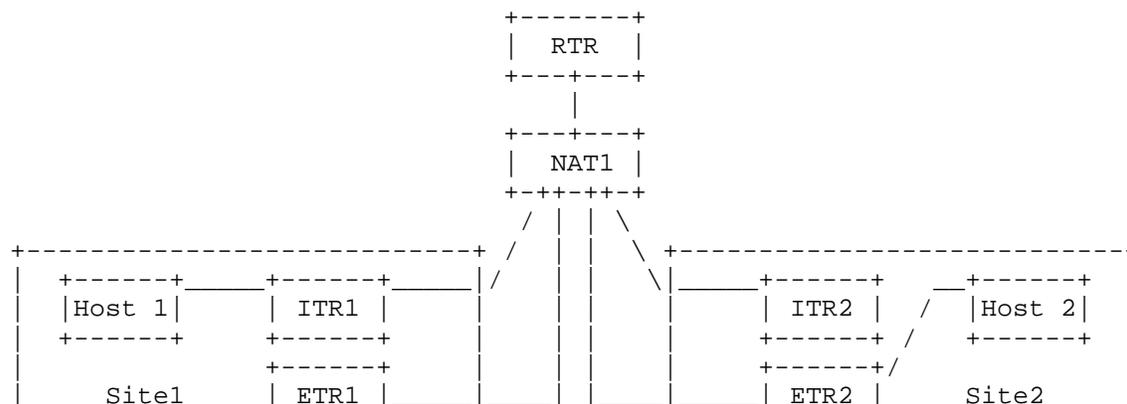




Fig.1 Case 1: Single Level Nat Traversal

Based on RTR proxy strategy proposed in [LISP-NAT], when Host1 in Site1 want to send a packet to Host2 in Site2, following steps will be performed.

1. Host1 sends a packet to ITR1 of Site1, destination of the packet is EID address of Host2.
2. When ITR1 receives the packet, it encapsulates it in a LISP data header with outer header destination set to RTR RLOC and outer header destination port set to 4341. Based on RTR proxy strategy, ITR1 will not send Map-Request for Host2. As a result, ITR1 is not able to get RLOCs of Site2, i.e. ETR2's RLOC.
3. When encapsulated packet passes through NAT, NAT transform the source address and source port in outer header to be global address and port. This may create a state in the NAT device.
4. When RTR receive the encapsulated packet destination to its RLOC, it decapsulates the packet, and look for if there are local mapping information match the destination EID. As Site2 also registers through the RTR, RTR will find mapping information of Site2, as well as global state of ETR2 including the global address and global port in its local cache.
5. RTR uses ETR2's global state information to re-encapsulate the packet, and transfer it to ETR2.

Seen from steps enumerated above, based on existing RTR Proxy strategy, even though Site1 and Site2 locate behind the same NAT, traffic between these Sites need to route to the RTR which locates outside the NAT.

However, as Site1 and Site2 locate behind the same NAT, that's to say, ITR1 and ETR2 locate in the same Intranet, RLOC of ETR2 is reachable to ITR1. As a result, if ITR1 could get mapping information of Site2, it could encapsulate the packet directly to RLOC of ETR2 which could avoid routing latency of packet and could lighten re-encapsulation load on RTR.

3.2. Extended Scenario I-Multiple Levels NAT

routing between NAT1 and NAT2 are invisible to RTR. In RTR's local cache, it will store relative information of Site2, including mapping information, global IP addresses and global ports of xTR2s which are assigned by NAT1, etc.

In this scenario, every packet from Host1 to Host2 will follow the path:

Host1 -> ITR1 -> NAT1 -> RTR -> NAT2 -> ETR2 ->Host2

RTR will perform decapsulation and re-encapsulation process.

Based on existing RTR Proxy strategy, traffic between Site1 and Site2 still needs to route to the RTR which locates outside the NAT.

According to scenario depicted in Fig.2, even if ITR1 could get mapping information of Site2, ETR2's RLOC is not reachable to ITR1.

However, if ITR2 gets global information of Site2, i.e. global addresses and global ports of xTR2s which are used to route outside NAT1, it could encapsulate packets directly to ETR2's global address. If NAT1 supports Hairpin function, the packets destination to ETR2's global address will be route to Site2, and need not to be re-encapsulated by RTR.

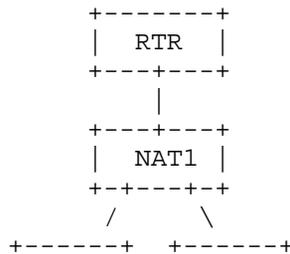
3.3. Extended Scenario II-Multiple Levels NAT

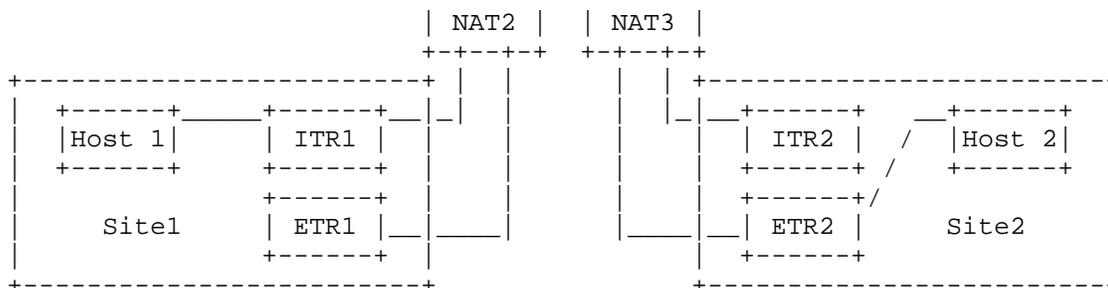
As shown in Fig.3, Site1 and Site2 both locate behind two levels NAT. Site1 and Site2 register through RTR as we described in Section 3.2.

In this scenario, each packet from Host1 to Host2 follows the path:

Host1 -> ITR1 -> NAT2 -> NAT1 -> RTR -> NAT3 -> ETR2 ->Host2

In this situation, if Site1 gets global information of the Corresponding Site, ITR1 could probe ETR2's global IP address. If NAT1 support Hairpin function, probe message will be transfer to Site2, and Site1 need not to encapsulate packet to RTR anymore.





4. Solutions

According to the analysis we described above, in this section we propose our solution strategy. In this

This mechanism enables RTR to send an "Info-Notify" message to the LISP Site with relative information of its Corresponding Site.

4.1. RTR Processing

When RTR receives an encapsulated packet from ITR1 of Site1, following steps will be performed.

1. RTR strips the outer header and lookup in local cache for mapping information of destination EID.
2. In our extension, RTR will judge if Source Site and Destination Site locate behind the same NAT. For example, based on whether the global address of Source Site ITR1 and global address of Destination Site ETR2 are the same, or whether the two global addresses may locate in the same address prefix. Furthermore, if ISP would like to build a relationship between NAT and corresponding RTR, RTR could be informed with public address information of the NAT.
3. If RTR judge that source site and destination site both locate behind the same NAT, it could send "Info-Notify" messages which contain relative information of Corresponding Site to Sender ITR and Receiver ETR respectively. Relative information mentioned above may include mapping information of Site, global addresses and global ports information of Site xTRs.

Note: After RTR sending message to ITR, if RTR receive packet from ITR to the particular ETR, RTR continue to transfer the packet to destination ETR.

4.2. ITR Processing

When ITR receives "Info-Notify" message from RTR,

1. Based on mapping information of Destination Site, ITR could send a RLOC probe message to ETR's RLOC. After receiving Mapping Reply from ETR, ITR could encapsulate packet directly to ETR's RLOC.
2. If Destination Site locate behind multilevel NATs, such like scenarios described in Fig.2 and Fig.3, ETR's RLOC is not reachable to ITR, and ITR will not receive Map Reply form Destination Site when it send RLOC Probe message to ETR's RLOC. In this situation, ITR could then send Probe message to ETR's global address. If ITR could get map reply message, then it could use ETR's global address as destination address of outer header to encapsulate packet.

Note In order to avoid traffic affection, during RLOC Probe phase, ITR could continue to encapsulate data packet to RTR.

4.3. ETR Processing

When ETR receives "Info-Notify" message from RTR, ETR could choose to send a message such like a SMR (Solicit-Map-Request)[RFC6830] to ITR to trigger a mapping request operation. Furthermore, ETR could use ITR's RLOC or global address as destination when it send the message.

5. Security Considerations

TBD

6. IANA Considerations

This document makes no requests to IANA.

7. Normative References

[LISP-NAT]

Ermagan, V., Farinacci, D., Lewis, D., Skriver, J., Maino, F., and C. White, "NAT traversal for LISP", March 2013.

[RFC6830]

Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "Locator/ID Separation Protocol (LISP)", January 2013.

Authors' Addresses

Li Cheng
ZTE Corporation
R&D Building 1, Zijinghua Road No.68
Nanjing, Yuhuatai District 210012
P.R.China

Email: cheng.li2@zte.com.cn

Mo Sun
ZTE Corporation
R&D Building 1, Zijinghua Road No.68
Nanjing, Yuhuatai District 210012
P.R.China

Email: sun.mo@zte.com.cn