

Network Working Group  
Internet-Draft  
Intended status: Experimental  
Expires: March 20, 2016

M. Boucadair  
C. Jacquenet  
France Telecom  
September 17, 2015

Triggered LISP Map-Request for Inter-Domain LISP Deployments  
draft-boucadair-lisp-triggered-map-request-00

## Abstract

It is commonly acknowledged that one of the issues raised by the current Locator/ID Separation Protocol (LISP) design is that some packets are likely to be lost in specific situations such as the absence of a mapping entry in the mapping cache maintained by an ITR. This issue is usually referred to as the "first packet loss" problem.

This document specifies a new LISP capability called Triggered Map-Request which aims at addressing the "first packet loss" issue. Also, it proposes to extend the LISP mapping entries with names instead of achieving RLOC resolution based upon EID prefixes only.

## Requirements Language

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 [RFC 2119](#) [[RFC2119](#)].

## 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 March 20, 2016.

Internet-Draft

Triggered Map-Request

September 2015

## Copyright Notice

Copyright (c) 2015 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

<a href="#">1.</a>	Problem Statement . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Sample LISP Flow Example (Focus on the Source Side) . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Triggered Map-Request . . . . .	<a href="#">4</a>
<a href="#">4.</a>	Name as an EID: Updated Map-Request Message Format . . . . .	<a href="#">5</a>
<a href="#">5.</a>	Operation . . . . .	<a href="#">7</a>
<a href="#">6.</a>	Security Considerations . . . . .	<a href="#">13</a>
<a href="#">7.</a>	IANA Considerations . . . . .	<a href="#">13</a>
<a href="#">8.</a>	Acknowledgments . . . . .	<a href="#">13</a>
<a href="#">9.</a>	References . . . . .	<a href="#">14</a>
<a href="#">9.1.</a>	Normative references . . . . .	<a href="#">14</a>
<a href="#">9.2.</a>	Informative references . . . . .	<a href="#">14</a>
	Authors' Addresses . . . . .	<a href="#">15</a>

[1.](#) Problem Statement

Locator/ID Separation Protocol (LISP, [[RFC6830](#)] ) operation relies upon a mapping mechanism that is used by ingress/egress Tunnel Routers (xTR) to forward traffic over the LISP network. It is commonly acknowledged that some packets are likely to be lost in some specific situations, such as the absence of a mapping entry in the mapping cache maintained by an ITR. This problem is usually denoted as the "first packet loss" issue.

Deploying LISP at the scale of the Internet heavily relies upon the reliability of the LISP Mapping service. In particular, LISP deployments at large scale must not degrade the level of quality as

currently provided by several decades of inter-domain routing practices.

This document describes a solution to prepare the local mapping service by anticipating the process of retrieving appropriate mapping

Internet-Draft

Triggered Map-Request

September 2015

entries by ITRs of a LISP-enabled domain before a packet is actually received by one of these ITRs. The LISP resolution result may remain valid until a Map-Request reaches the ultimate ETR.

In addition to better accommodating the risk raised by the "first packet loss" issue, this proposal reduces the delivery time that is likely to be exacerbated by the two indirection levels (DNS and LISP) together with the delay between the DNS resolution and the LISP resolution. An example is shown in [Section 2](#) for illustration purposes.

This document focuses on the sole LISP inter-domain use case. As such, the applicability of the proposed solution to other LISP uses cases is out of the scope of this document.

In addition to the terms defined in [[RFC6830](#)] and [[RFC6833](#)], this document makes uses of the following terms:

- o Authoritative server: A DNS server that can answer authoritatively for a given DNS query.
- o Stub resolver: A resolver with minimum functionality, typically used by endpoints that depend on a recursive resolver.
- o Recursive resolver: A DNS server that accepts requests from one resolver, and asks another resolver for the answer on behalf of the first resolver.

Within this document, "Triggered Map-Request" is used to refer to a Map-Request that is issued by an ITR based on some other events than presenting a packet to the ITR forwarding engine.

## [2.](#) Sample LISP Flow Example (Focus on the Source Side)

In order to further illustrate the issue related to the processing of the first packet, let's consider this example in which Host1 wants to communicate with a remote Host2, identified with "host2.xyz.example.com". To do so, the following steps need to be

followed:

- 1 Host1 does a DNS lookup on host2.xyz.example.com. DNS queries (A and/or AAAA) are issued by the local stub-resolver of Host1 and forwarded to a pre-configured recursive resolver.
- 2 If the recursive resolver is the authoritative server for this record, corresponding records are returned to the requesting stub resolver, otherwise the request is forwarded upstream following the normal DNS resolution procedure.
- 3 Once the recursive resolver receives a response from the DNS infrastructure, it will relay it to the requesting resolver. As a

result of this procedure, A and/or AAAA records are returned to the requesting host (i.e., Host1).

- 4 One of returned IPv4 or IPv6 addresses will be used by Host1 as the destination EID. The locally assigned address of host1.abc.example.com that belongs to the same address family is used as the source EID. An IPv4 or IPv6 packet is then built and forwarded through the LISP site as a normal IP packet until it reaches an ITR.
- 5 Upon receipt of this packet by an ITR, because no mapping entry is present for the destination EID, the ITR must invoke the LISP Mapping Service to retrieve the appropriate mapping entry to forward the packet outside the LISP leaf domain. According to [\[RFC6830\]](#):
  - 5.1 When an alternate mapping system is not in use, the Map-Request packet is routed through the underlying routing system. Otherwise, the Map-Request packet is routed on an alternate logical topology, for example, the [\[RFC6836\]](#) database mapping system. In either case, when the Map-Request arrives at one of the ETRs at the destination site, it will process the packet as a control message.
  - 5.2 The ETR looks at the destination EID of the Map-Request and matches it against the prefixes stored in the EID-to-RLLOC mapping database maintained by the ETR. This is the list of the EID-Prefixes the ETR is aware of, and which have been assigned to the ETR is connected to. If there is no match, the Map-Request is dropped. Otherwise, a LISP Map-Reply is returned to the ITR.
  - 5.3 The ITR receives the Map-Reply message, parses the message

(to check for format validity), and extracts the mapping information from the packet. This information is stored in the ITR's EID-to-RLOC mapping cache. Note that the map-cache is an on-demand cache. Map-cache management by the ITR is optimized to accommodate the ITR's resource constraints.

### 3. Triggered Map-Request

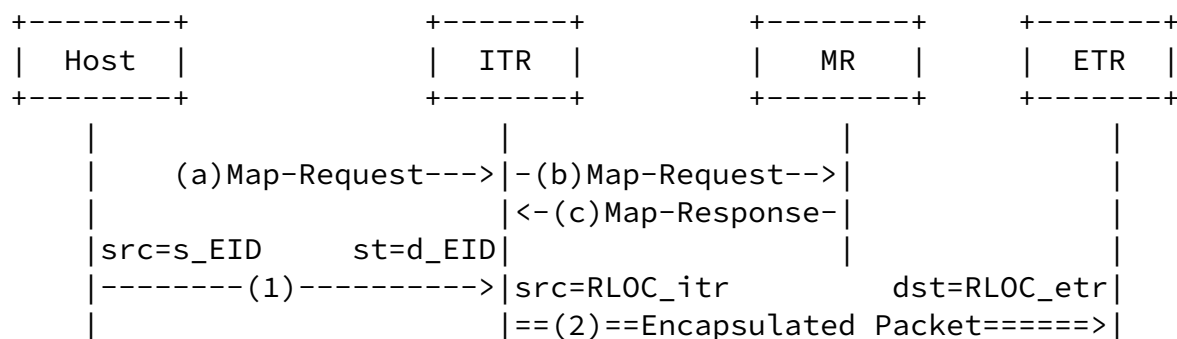
The rationale adopted by this document is that, instead of waiting for a packet to be received by an ITR for issuing a Map-Request message, the request can be triggered by other events so that the local mapping cache is ready to process a packet that needs to be forwarded outside a LISP leaf domain. This mode is called Triggered Map-Request.

Triggered Map-Request has the same message format as a normal Map-Request: that is, an external entity receiving a triggered Map-Request or a normal Map-Request won't be able to make the difference between the two messages. Whether the Map-Request is triggered by an

external entity or carried by a packet that needs to be forwarded outside a LISP leaf domain reflects a context that is local to the LISP domain that originates the Map-Request message.

Triggered Map-Request is meant to anticipate the receipt of a packet that would have to be forwarded outside so that the ITR installs the required state and proceed with the forwarding of the packet over a LISP infrastructure accordingly.

An example of Triggered Map-Request is shown in Figure 1. This example does not explicitly identify which entity has triggered the Map-Request in Step (a).






Figure 1: Triggered Map-Request: A Flow Example

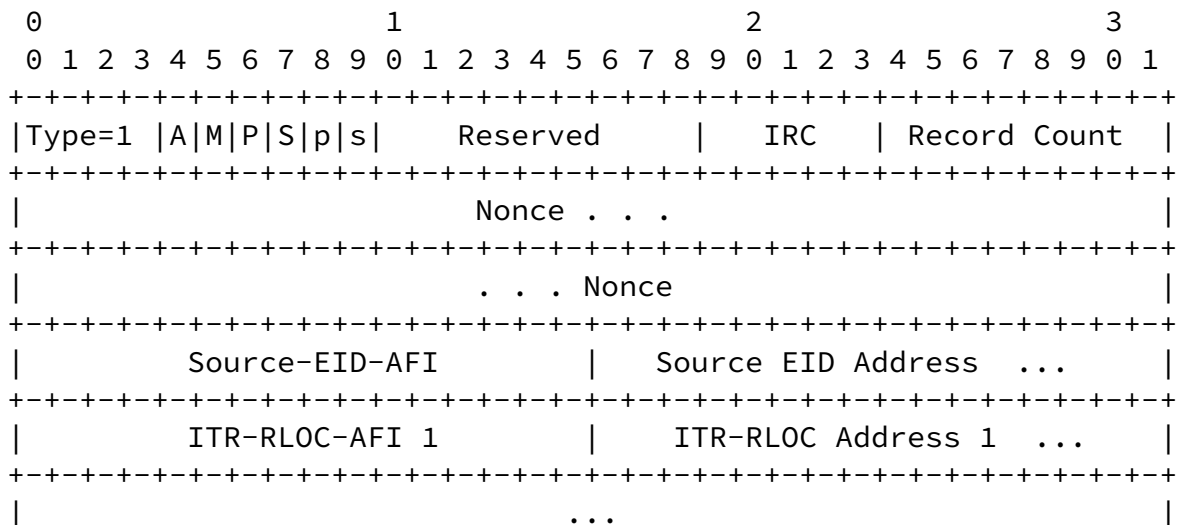
An example of the use of triggered Map-Requests is detailed in [Section 5](#).

#### 4. Name as an EID: Updated Map-Request Message Format

Figure 2 illustrates the changes that are required to the Map-Request message in order to support names as EID identifiers. The design relies upon the definition of one of the reserved bits for this purpose. This bit is called the N-bit. When set (name-as-an-eid bit), this is an indication that the EID-Prefix field must be interpreted as a name.

Note: Another design option is to assign a dedicated value to the "EID-Prefix-AFI" when a name is carried in the message. This design option may offend some purists, since a name is usually not considered as an address family.

OLD:






Figure 2: Name as an EID

The "Reserved" bits MUST each be set to zero and MUST be ignored upon receipt.

When this N-bit is set, the EID-Prefix-AFI MUST be set to zeros and MUST be ignored upon receipt. Also, EID mask-len ( Name-Len) MUST indicate the length of the enclosed "Name". Name is a domain name

(as per [Section 3.1 of \[RFC1035\]](#)) that contains one or more labels. The Name encoding MUST follow the Name Syntax defined in [\[RFC1035\]](#) [\[RFC1123\]](#) [\[RFC2181\]](#) and are represented in ASCII form.

## 5. Operation

The solution relies upon an extension to the DNS resolver (and possibly a management platform) to trigger the sending of Map-Request messages for a given destination EID (that is eventually encoded as a name or an IP address/prefix) by all the ITRs deployed in a given LISP-enabled domain.

This document assumes that in the context of inter-domain LISP deployment use cases, interconnection between Mapping Systems is required for the sake of global reachability. Furthermore, these Mapping Systems are supposed to deploy massive cache systems so that they can service resolution requests as close to the leaf LISP domain as possible, rather than forwarding the Map-Request until it reaches the ultimate ETR. Furthermore, some service innovation can be encouraged by coupling DNS/LISP Mapping services.

The proposed procedure also takes into account CDN environments, at the benefit of relaxing the constraint on the traffic increase on interconnection links, thereby minimizing the need for soliciting inter-domain LISP forwarding.

The solution also acknowledges that DNS replies can be policy-based. In particular, this document does not interfere with DNS policies that are enforced on a subnet basis (e.g., [\[I-D.ietf-dnsop-edns-client-subnet\]](#)). When the Mapping System has to undertake a DNS resolution, it will supply the same subnet value as the one that would be indicated by a host connected to the leaf LISP network. Doing so ensures that the address that will be returned to the requesting host during the DNS resolution will match a mapping entry that will be retrieved when Triggered Map-Request operation is enabled.

The detailed procedure to be implemented to minimize the risk of the "first packet loss" issue is specified hereafter:



enable/disable the Triggered-Map-Request procedure. The default value of this parameter is "Disabled".

- 2 All (inter-domain) ITRs MUST subscribe to a well-known multicast group (@MCAST) and listen to port 4342 (default port number).
  - 2.1 The use of multicast transport will help ITRs of the different domains to maintain the same database.
  - 2.2 Also, it does not interfere with the underlying routing and forwarding policies that are configured locally. Whatever the ITR that will be selected when forwarding an outgoing packet, that ITR has issued a triggered Map-Request.
- 3 The DNS resolver is configured with the same @MCAST. If a different port than port number 4342 is used, this port number MUST be explicitly configured on the recursive DNS resolver.
- 4 A recursive DNS resolver within a LISP-enabled domain is updated with one of the following capabilities. The decision about which one to enable is deployment-specific. This decision will help identifying which DNS forwarders will be impacted.
  - 4.1 When receiving a DNS query from a stub-resolver, duplicate that query and forward the duplicate to @MCAST:4342. The original query is forwarded according to normal DNS procedures (see the example shown in Figure 3).

This query is duplicated as close to the stub-resolver as possible so that the LISP resolution process can occur while the DNS resolution is in progress.

- 4.1.1 If the recursive resolver is the authoritative server for this record, or the authoritative server is within the local LISP domain, or a cache is invoked for that name, then corresponding records are returned to the requesting stub resolver following normal DNS procedures. Packets will remain within the same LISP domain. (Inter-domain) ITRs won't be solicited for this communication.
    - 4.1.2 Otherwise, the request is forwarded upstream following the normal DNS resolution procedure. In addition, the DNS recursive resolver MUST duplicate the query and forward it to @MCAST:4342.

4.1.1.3 Upon receipt of the DNS query, an ITR will build a Map-Request with a name ([Section 4](#)). This triggered Map-Request is then forwarded to a Map-Resolver. If the Map-Resolver is updated to support the capability to associate a name with a mapping entry, it can make a query based on the name carried in the Map-Request. If not, the Map-Resolver must proceed first with a DNS resolution locally and then the LISP resolution.

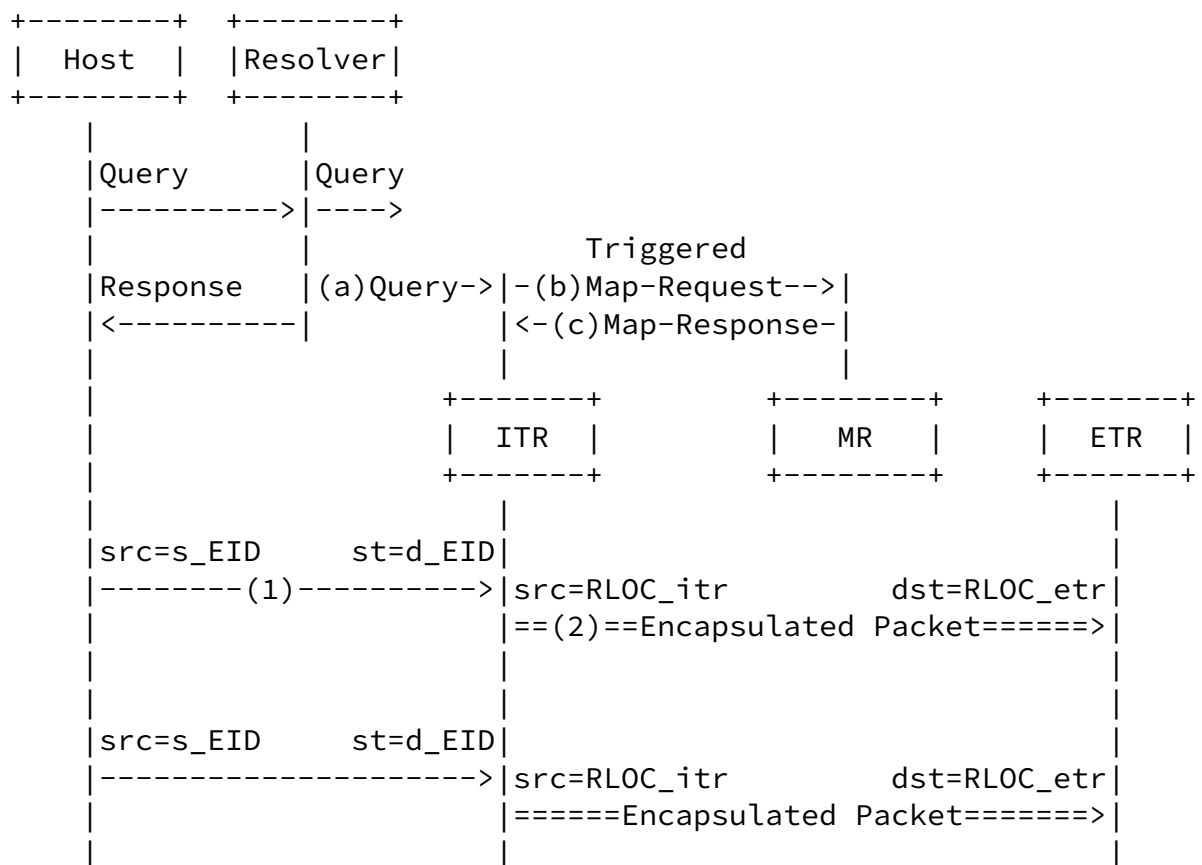


Figure 3: Processing Triggered Map-Request: Close to the Stub-resolver

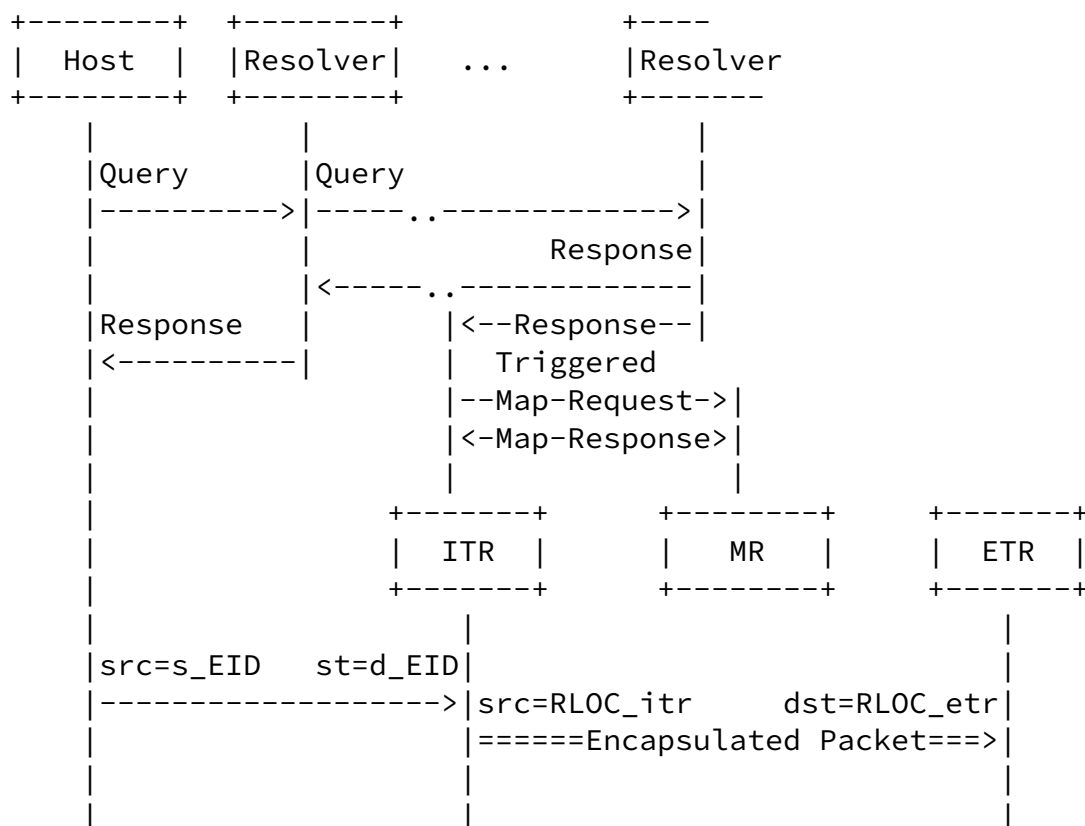
4.2 When forwarding a DNS response to another DNS server, duplicate that response and forward the duplicate to @MCAST:4342. The original response is forwarded according to normal DNS procedures (see the example shown in Figure 4).

The farthest DNS resolver of a leaf LISP network (i.e., a resolver that forwards DNS queries outside a LISP domain) is updated to fork a query for DNS records that cannot be serviced locally, either because the authoritative server belongs to the local LISP domain or because there is a cache

platform that is enabled in the local LISP domain. This DNS

Query is carried into a Triggered Map-Request message that is forwarded to all the ITRs of that LISP domain. Concretely:

- 4.2.1 If the recursive resolver is the authoritative server for this record, or the authoritative server is within the local domain, or a cache is invoked for that name, then corresponding records are returned to the requesting stub resolver following normal DNS procedures. Packets will remain within the same LISP domain. (Inter-domain) ITRs won't be solicited for this communication.
- 4.2.2 Otherwise, the request is forwarded upstream following the normal DNS resolution procedure. In addition, the DNS recursive resolver MUST duplicate the DNS response and forward it to @MCAST:4342.



```

|src=s_EID   st=d_EID|
|----->|src=RLOC_itr   dst=RLOC_etr|
|               |=====Encapsulated Packet===>|
|               |

```

Figure 4: Processing Triggered Map-Request: Far from to the Stub-Resolver

- 4.3 When forwarding a DNS query to another DNS server, build a corresponding Triggered Map-Request from the contents of the initial DNS query message. The request is then forwarded to @MCAST:4342. The original query is forwarded according to normal DNS procedures (see the example shown in Figure 5).
  - 4.3.1: This procedure is similar to the one described in Bullet 4.1. The only difference is that, instead of forking a DNS message, appropriate Triggered Map-Request messages are generated. The DNS resolvers rely upon the contents of the DNS query to build the Triggered Map-Request message; especially, the destination EID is set to the addresses (IPv4 and/or IPv6) that were included in the DNS response(s). Furthermore, the Map-Request message uses the format defined in [Section 4](#) to set the destination EID.
  - 4.3.2: Upon receipt of the Map-Request that carries a name, an ITR will forward the request to its Map-Resolvers. If the Map-Resolver is updated to support the capability to associate a name with a mapping entry, then it can initiate a query based on the name carried in the Map-Request. If not, the Map-Resolver must proceed first to DNS resolution locally and then a LISP resolution.

Internet-Draft

Triggered Map-Request

September 2015

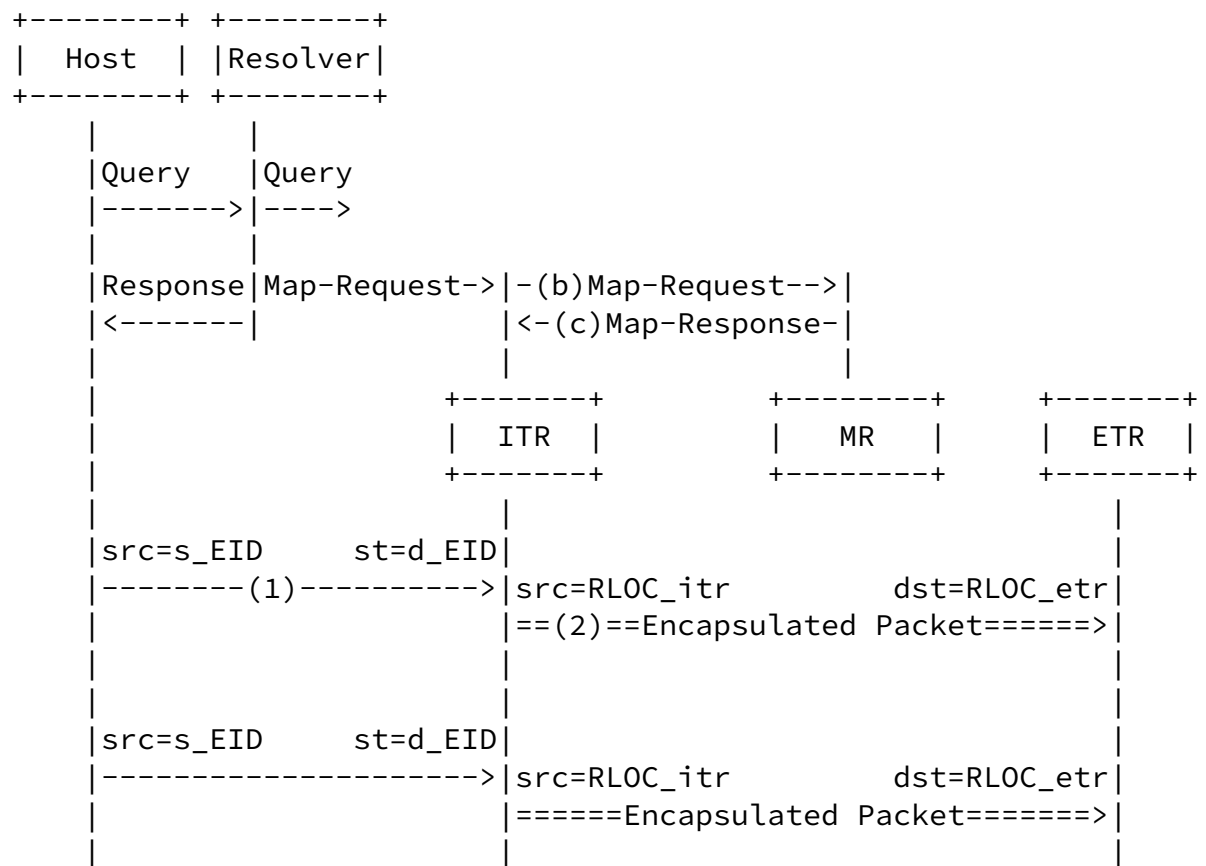
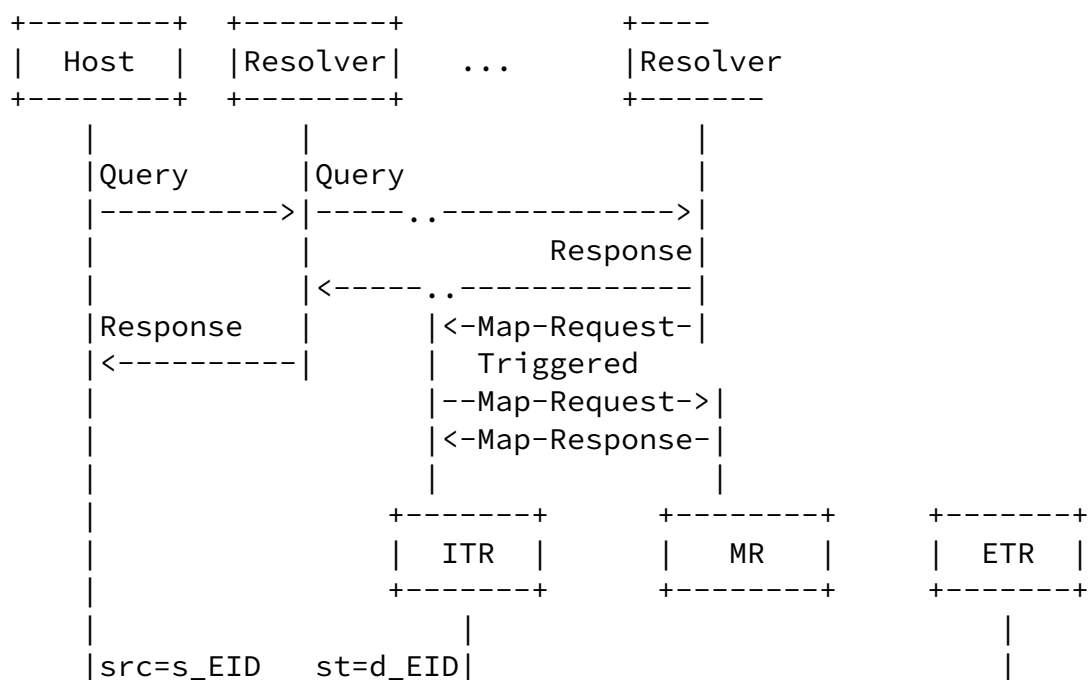


Figure 5: Processing Triggered Map-Request: Close to the Stub-

## Resolver

4.4 When forwarding a DNS response to another DNS server, trigger a corresponding Map-Request that is formed after the contents of the said DNS response. The request is then forwarded to @MCAST:4342. The original response is forwarded according to normal DNS procedures (see the example shown in Figure 6).

4.4.1: This procedure is similar to the one described in Bullet 4.2. The only difference is that, instead of forking a DNS message, appropriate Map-Request messages are generated. The DNS resolver relies upon the content of the DNS response to build the Map-Request message; especially, the destination EID is set to the addresses (IPv4 and/or IPv6) that were included in the DNS response(s).



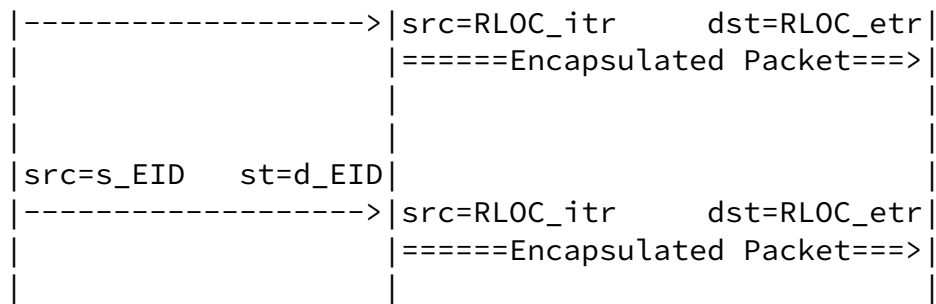


Figure 6: Processing Triggered Map-Request: Far from to the Stub-resolver

- 5 Subsequent packets associated with the same flow are handled locally (i.e., normal LISP operation applies).

## 6. Security Considerations

Security considerations discussed in [RFC6830] and [RFC6833] should be taken into account.

## 7. IANA Considerations

To be completed.

## 8. Acknowledgments

This work is partly funded by ANR LISP-Lab project #ANR-13-INFR-009-X.

## 9. References

### 9.1. Normative references

- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](http://www.rfc-editor.org/info/rfc1035), DOI 10.17487/RFC1035, November 1987, <<http://www.rfc-editor.org/info/rfc1035>>.
- [RFC1123] Braden, R., Ed., "Requirements for Internet Hosts - Application and Support", STD 3, [RFC 1123](http://www.rfc-editor.org/info/rfc1123), DOI 10.17487/RFC1123, October 1989,

<<http://www.rfc-editor.org/info/rfc1123>>.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", [RFC 2181](#), DOI 10.17487/RFC2181, July 1997, <<http://www.rfc-editor.org/info/rfc2181>>.
- [RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", [RFC 6830](#), DOI 10.17487/RFC6830, January 2013, <<http://www.rfc-editor.org/info/rfc6830>>.
- [RFC6833] Fuller, V. and D. Farinacci, "Locator/ID Separation Protocol (LISP) Map-Server Interface", [RFC 6833](#), DOI 10.17487/RFC6833, January 2013, <<http://www.rfc-editor.org/info/rfc6833>>.

## 9.2. Informative references

- [I-D.ietf-dnsop-edns-client-subnet]  
Contavalli, C., Gaast, W., Lawrence, D., and W. Kumari, "Client Subnet in DNS Queries", [draft-ietf-dnsop-edns-client-subnet-03](#) (work in progress), August 2015.
- [RFC6836] Fuller, V., Farinacci, D., Meyer, D., and D. Lewis, "Locator/ID Separation Protocol Alternative Logical Topology (LISP+ALT)", [RFC 6836](#), DOI 10.17487/RFC6836, January 2013, <<http://www.rfc-editor.org/info/rfc6836>>.

### Authors' Addresses

Mohamed Boucadair  
France Telecom



Rennes 35000  
France

EMail: mohamed.boucadair@orange.com

Christian Jacquenet  
France Telecom  
Rennes 35000  
France

EMail: christian.jacquenet@orange.com