

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 28, 2017

M. Boucadair
C. Jacquenet
Orange
April 26, 2017

LISP Subscription
draft-boucadair-lisp-subscribe-05

Abstract

Mapping Services in Locator/ID Separation Protocol (LISP) networks are key to proper LISP forwarding operation. When considering the deployment of LISP at large scale, these Mapping Services are even more crucial for the sake of the LISP forwarding efficiency. This document introduces two additional LISP messages that are meant to facilitate the dynamic discovery of Mapping Systems, improve Ingress Tunnel Routers (ITR) recovery times and optimize the solicitation of the LISP Mapping System as a function of the ITR location, in particular.

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 October 28, 2017.

Copyright Notice

Copyright (c) 2017 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
1.1.	Issues	3
1.2.	Assumptions	3
1.3.	Improving LISP Mapping Services	4
2.	Map-Subscribe Message Format	5
3.	Map-Subscribe-Ack Message Format	6
4.	Generating a Map-Subscribe Message	9
5.	Processing a Map-Subscribe Message	10
6.	Processing a Map-Subscribe-Ack Message	12
7.	Subscription to Multiple Resolvers	13
8.	Sample Examples	13
8.1.	Map-Resolver Redirect	13
8.2.	Mapping Cache Retrieval	14
8.3.	Unsolicited Map-Reply	16
9.	Security Considerations	17
10.	IANA Considerations	17
11.	Acknowledgments	17
12.	References	17
12.1.	Normative references	18
12.2.	Informative references	18
	Authors' Addresses	18

[1.](#) Introduction

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. The ability of dynamically discovering the Map-Resolver and Map-Server entities that provide such mapping services is meant to facilitate global LISP operation. In particular, the ability to inform Ingress Tunnel Routers (ITR) of a LISP network about the availability and the status

of several Mapping Services is likely to improve the overall LISP forwarding serviceability.

1.1. Issues

This section lists a set of issues that need further investigation:

Discover ITRs: Current LISP design does not allow to automatically discover active ITRs of a LISP domain (nor the mapping system of a given domain is aware of ITRs of the same domain that can solicit its services, let alone ITRs of other domains). The solution proposed in this document allows to fill that gap.

Optimize EID-ROLC resolution time: Leaf LISP networks can be better serviced, for example by avoiding the cascading of Map-Resolvers, or by avoiding the solicitation of a Map-Resolver that is located an ocean away, etc. Policies should be taken into account when configuring Map-Resolver information on an ITR for improving EID-to-ROLC resolution. These policies should be modified and adjusted according to various events (e.g., installation of an additional Map-Resolver).

Accomodate Map-Resolvers constraints: Allows for the protocol to redirect a requesting ITR to another Map-Resolver when some events occur, such as an overload of the initially targeted Map-Resolver or when Map-Resolvers are optimized to service a set of destination EIDs, etc.

Faster Recovery of mapping entries: Whenever an ITR fails for some reason, the faulty ITR needs to recover at least the list of mappings for the most popular prefixes in a timely manner, in particular. Policies for mapping entries (to be recovered) are deployment-specific.

Receive push notifications: For LISP leaf networks that would need to maintain an up-to-date mapping table for a set of destination EIDs (or even the global mapping table) to avoid issues such as the loss of a first packet or to optimize LISP forwarding delay (and therefore the overall forwarding efficiency), a dynamic push mechanism would be helpful.

1.2. Assumptions

This document makes the following assumptions:

- o Various LISP players (network operators, service providers, etc.) are likely to deploy and operate different LISP Mapping Systems. Multiple Mapping Systems will therefore coexist for various

reasons, e.g., avoid country-centric governance, allow for distinct technologies to implement the mapping service, business opportunities, service innovation, etc.

- o Interconnection between these Mapping Systems is required for the sake of global connectivity and also to minimize the risk of fragmenting the Internet.
- o Mapping Services are supposed to be dimensioned to maintain a global mapping database for the entire LISP-enabled Internet.
- o Mapping Service providers may offer advanced services to their customers such as maintain local caches (a la CDN), or update ITR mapping entries that match some criteria requested by a leaf LISP network, redirect ITR requests to the closest Map-Resolvers, structure the mapping resolution service so that the resolution time is optimized, etc.
- o The entries of the mapping tables are exchanged between these Mapping systems so that Map-Request messages can be processed as close to the LISP leaf networks as possible.
- o A leaf LISP-enabled network subscribes to the Mapping Service provided by one or several Mapping Service operators.
- o The contribution of each player involved in the provisioning and the operation of a LISP-based connectivity forwarding service should be rationalized so that clear interfaces are defined and adequate mechanisms for troubleshooting, diagnosis and repair purposes can be easily implemented and adopted. The inability of identifying what is at the origin of the degradation of a LISP connectivity service is seen as one of the hurdles that are likely to jeopardize LISP deployments at large scale.

1.3. Improving LISP Mapping Services

This document specifies a couple of additional LISP messages that are meant to improve the subscription to Mapping Services, let alone their serviceability. They are described in the following sections.

A simple method to redirect ITR-originated mapping requests towards another Map-Resolver when some conditions are met (e.g., overload of a Map-Resolver, enforcement of traffic engineering policies, etc.) is defined in [Section 2](#) and [Section 3](#).

2. Map-Subscribe Message Format

The format of the Map-Subscribe message is shown in Figure 1.

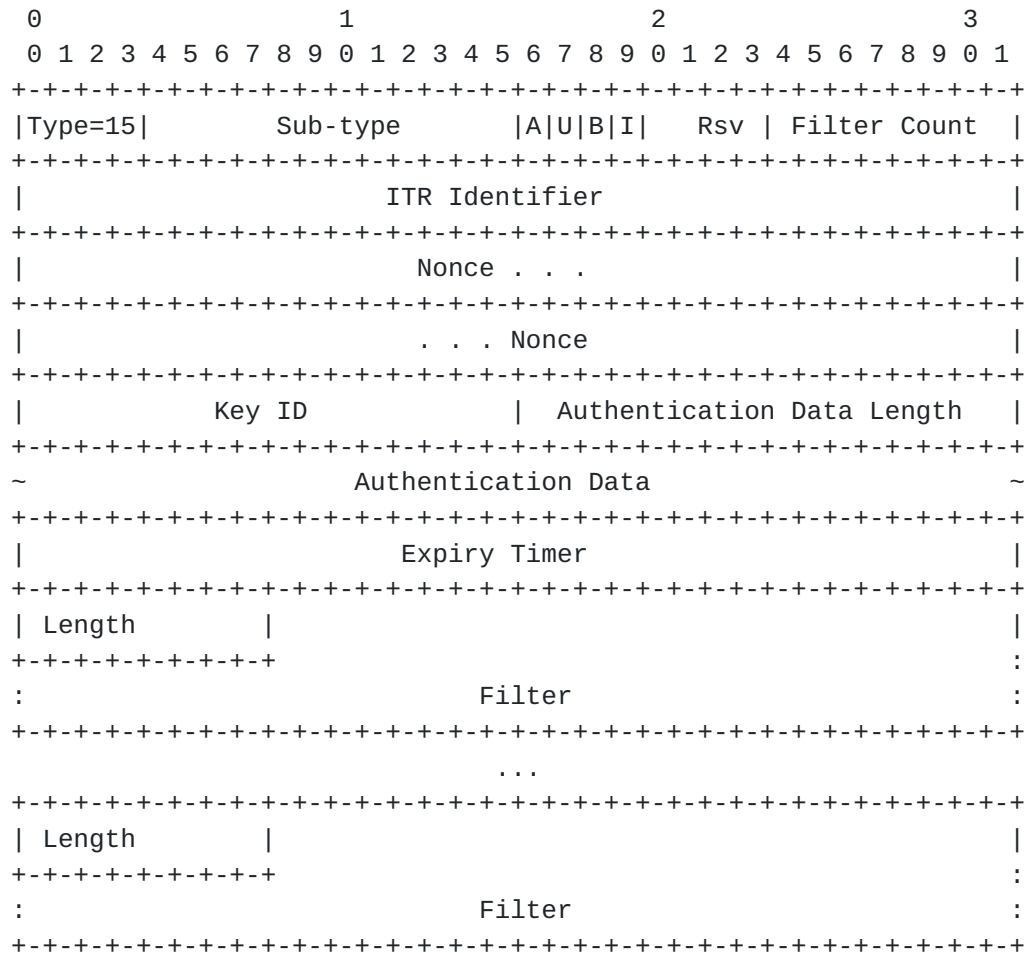


Figure 1: Map-Subscribe Message Format

The description of the fields is as follows:

- o Type MUST be set to 15 [[RFC8113](#)].
- o Sub-type: MUST be set to 1024.
- o A (Ack-bit): this bit MUST be set to 0 for Map-Subscribe requests.
- o U (unsolicited-map-reply bit): When set, this flag indicates that the originating ITR is ready to receive implicit Map-Reply messages.

- o B (bulk-support bit): When set, this flag indicates that the originating ITR supports mapping bulk retrieval method (e.g., [[I-D.boucadair-lisp-bulk](#)]).
- o I (immediate-retrieval bit): When set, this flag indicates that the originating ITR requests immediate retrieval of the portion of the mapping table that matches the filters included in the request.
- o Rsv: reserved bits, MUST be sent as 0 and MUST be ignored when received.
- o Filter Count: This field indicates the number of the filters included in the message.
- o Nonce, Key ID, Authentication Data Length, and Authentication Data are similar to those of a LISP Map-Register message ([[RFC6830](#)]).
- o Expiry Timer: This field indicates, in seconds, the validity timer for the subscription.
- o Length: This field indicates, in octets, the length of the filter that is encoded in the "Filter" field.
- o Filter: This field carries a destination EID (or a set thereof) that is encoded as an UTF-8 string. This specification allows to convey IP prefix literals, Names and/or AS numbers. One or multiple filters may be present in a request. IPv4 prefixes are encoded as IPv4-mapped IPv6 prefixes [[RFC4291](#)] (i.e., starting with ::ffff:0:0/96). A mix of names, IP prefixes and AS numbers may be enclosed in the same request. The value 0 is used to delete existing filters. Filters MUST be applied in the order they appear in the request.

3. Map-Subscribe-Ack Message Format

The format of the Map-Subscribe-Ack message is shown in Figure 2.

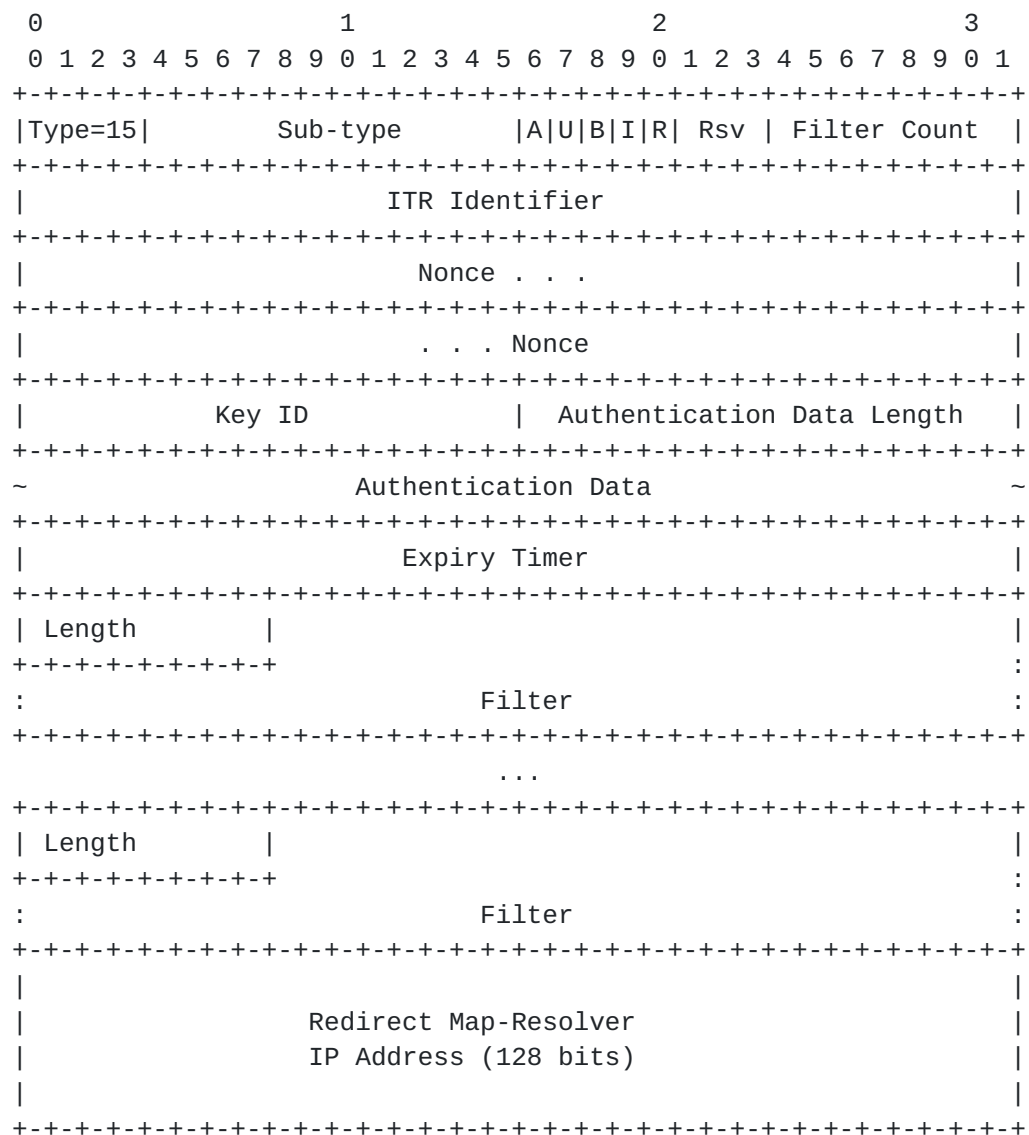


Figure 2: Map-Subscribe-Ack Message Format

The description of the fields is as follows:

- o Type: MUST be set to 15 [[RFC8113](#)].
- o Sub-type: MUST be set to 1024.
- o A (Ack-bit): this bit MUST be set to 1 for Map-Subscribe-Ack responses.
- o U (unsolicited-map-reply bit): When set, this flag indicates that the Map-Resolver can issue implicit Map-Reply messages.

- o B (bulk-support bit): When set, this flag indicates that the Map-Resolver supports mapping bulk retrieval method (e.g., [[I-D.boucadair-lisp-bulk](#)]).
- o I (immediate-retrieval bit): When set, this flag indicates that the Map-Resolver will initiate an immediate push cycle of the portion of the mapping table that matches the filters included in the request.
- o R (Redirect bit): When set, this flag indicates that a redirect Map-Resolver is enclosed in the message.
- o Rsv: reserved bits, MUST be set to 0 and MUST be ignored when received.
- o Result: indicates the result code of the processing of the Map-Subscribe request. The following codes are defined:
 - 0: SUCCESS. This code is used to indicate the request is successfully processed.
 - 1: PARTIAL-FILTERS-INSTALLED-LIMIT. This code is used to indicate a request is successfully processed but some filters were not installed because the number of filters carried in the initial Map-Subscribe message exceeds the filter limit.
 - 2: PARTIAL-FILTERS-INSTALLED-BAD. This code is used to indicate a request is successfully processed but some filters were not installed because they were malformed.
 - 3: PARTIAL-FILTERS-INSTALLED-LOCAL. This code is used to indicate a request is successfully processed but some filters were not installed because of local reasons. The ITR SHOULD, after a certain timer expires, send a Map-Subscribe request message for the set of filters that are not included in the Map-Subscribe-Ack message received by the ITR in response to its initial Map-Subscribe request message.
 - 4: FILTERS-PROHIBITED. This code is used to indicate a request is successfully processed but the installation of filters is prohibited.
- o Filter Count: This field indicates the number of the filters included in the message.
- o Nonce, Key ID, Authentication Data Length, and Authentication Data are similar to those of a LISP Map-Register message ([[RFC6830](#)]).

- o Expiry Timer: This field indicates, in seconds, the validity timer for the subscription.
- o Length: This field indicates, in octets, the length of the filter that is encoded in the "Filter" field.
- o Filter: This field carries the set of filters that were successfully installed.
- o Redirect Map-Resolver IP Address (128 bits): When the R-bit is set, this field carries the IP address of the Map-Resolver where mapping requests should be redirected. An IPv4 address is encoded as an IPv4-mapped IPv6 address.

4. Generating a Map-Subscribe Message

The A-bit of a Map-Subscribe message MUST be set to 0.

An ITR uses the U-bit to inform a Map-Resolver whether it is ready to handle unsolicited Map-Reply messages or not. The ITR MUST set the U-bit when it supports such capability.

An ITR uses the B-bit to inform a Map-Resolver whether it supports the mapping bulk transfer method or not. The ITR MUST set to the B-bit when it supports such method (e.g., [[I-D.boucadair-lisp-bulk](#)]).

An ITR that joins the LISP network is likely to delete all notifications that are bound to its RLOCs. It does so by including a Null filter prior to any filter that it wishes the Map-Resolver to record. Note, an ITR can indicate a Null filter using one of these methods:

1. Send a Map-Subscribe message with a "Filter Count" set to 0, or
2. Include a Filter with a "Filter" field set to zeros.

An ITR that loses its mapping cache for some reason SHOULD generate a Map-Subscribe message towards its Map-Resolver(s) with the I-bit set.

An ITR MAY generate several Map-Subscribe messages to make the Map-Resolver install the required filters. Nevertheless, an ITR MUST expect that the Map-Resolver may limit the number of filters that may be installed. Filters that are not accepted or not processed by the Map-Resolvers are not included in a Map-Subscribe-Ack message.

An ITR that wants to delete one or a set of filters MUST generate a Map-Subscribe message which carries those filters with an Expiry Timer set to 0.

5. Processing a Map-Subscribe Message

A Map-Resolver that does not support the Map-Subscribe message MUST silently ignore any Map-Subscribe message it receives.

Map-Resolvers MUST support a configurable parameter to enable/disable the processing of Map-Subscribe messages. The default value is set to "enabled".

A Map-Resolver SHOULD support a configuration parameter to limit the number of filters per leaf LISP network, per ITR, etc.

If a Map-Resolver receives a Map-Subscribe message and is enabled to process it, a Map-Resolver MUST reply with a Map-Subscribe-Ack message to acknowledge the receipt of the corresponding Map-Subscribe message.

When building a Map-Subscribe-Ack message, the Map-Resolver MUST:

- o Set the A-bit to indicate this is a response to a Map-Subscribe request.
- o Set the U-bit if it supports the unsolicited Map-Reply capability, except if a redirect Map-Resolver is to be returned.
- o Set the B-bit if it supports a method for mapping bulk transfer, except if a redirect Map-Resolver is to be returned.
- o Set the R-bit if it wants to inform the requesting ITR about another Map-Resolver it should contact. The Map-Resolver MAY return a set of filters that are to be applied by the ITR to select the Map-Resolver (i.e., destination EID Map-Resolver address selection).
- o Echo the I-bit if the Map-Resolver accepts to initiate unsolicited mapping retrievals, except if a redirect Map-Resolver is to be returned.
- o If no redirect is enabled and the request includes one or several filters, the Map-Resolver MUST echo the filters it succeeds to install, and in the same order of appearance, in the Map-Subscribe-Ack message.
- o If the Map-Resolver is configured with maximum and minimum values for the expiry timer, the Map-Resolver MUST adjust the Expiry Timer enclosed in the request so that it does not exceed these boundary values.

If the I-bit is set in the Map-Subscribe request and the Map-Resolver supports the unsolicited mapping retrieval capability, the Map-Resolver SHOULD generate unsolicited Map-Reply messages or dedicated bulk transfer messages that carry the EID-RLOC mapping entries that match the filters already present in the Mapping System for that ITR or that match those carried by the Map-Subscribe message.

If filters are included in the request, the Map-Resolver MUST extract those filters and update its mapping system subscription for that ITR accordingly. In particular, the Map-Resolver MUST delete all filters that are active for that ITR if a Null filter is included in the Map-Subscribe request or if the expiry timer is null.

If filters are not allowed for a given ITR, the 'Result' field MUST be set to FILTERS-PROHIBITED.

If all filters are successfully installed, the 'Result' field MUST be set to SUCCESS.

If the Map-Resolver fails to install some of the filters included in a request because the filter limits for that ITR are exceeded, it MUST NOT echo the corresponding filters in the Map-Subscribe-Ack message. The 'Result' field MUST be set to PARTIAL-FILTERS-INSTALLED-LIMIT.

If the Map-Resolver fails to install some of the filters included in a request because these filters were malformed, it MUST NOT echo the corresponding filters in the Map-Subscribe-Ack message; only successfully installed filters MUST be included. The 'Result' field MUST be set to PARTIAL-FILTERS-INSTALLED-BAD.

If, for some other reasons, the Map-Resolver fails to apply the filters included in a request, it MUST NOT echo the said filters in the Map-Subscribe-Ack message; only the successfully installed filters MUST be included. The 'Result' field MUST be set to PARTIAL-FILTERS-INSTALLED-LOCAL.

If a filter that is included in the request is more specific than a filter that is already present in the mapping system for the same ITR, the Map-Resolver MUST NOT add a new filter but MUST include the old filter in the response to the requesting ITR.

If a more specific filter exists in the mapping system for the same ITR, the Map-Resolver MUST replace the old filter (i.e., the one already stored in the system) with the new filter (i.e., the one included in the Map-Subscribe message).

An ITR can replace an existing filter by a more specific one by deleting all filters and install the new ones.

A Map-Resolver that is overloaded or configured by means of a specific policy to redirect requests sent by a set of ITRs to other Map-Resolvers, the Map-Resolver MUST reply with a Map-Subscribe-Ack message with the R-bit set and 'Redirect Map-Resolver IP Address' field set to the redirect Map-Resolver's address. All other bit flags MUST be returned unset. Moreover, the Expiry Timer MUST be set to 0. No Filter MUST be included in the message.

If an event matches one of the filters that have been installed by an ITR, the Map-Resolvers MUST generate the corresponding unsolicited mapping update message (e.g., Map-Reply, mapping bulk method).

Upon expiry of the validity timer associated with a filter, the Map-Resolver MUST delete that filter from its mapping subscription system.

6. Processing a Map-Subscribe-Ack Message

Upon receipt of a Map-Subscribe-Ack message, the ITR MUST check whether the message matches a Map-Subscribe message it sent to the same Map-Resolver. If no matching state is found, the message MUST be silently dropped. If a state is found, in addition to authentication checks, the ITR MUST proceed as follows:

- o If the U-bit is set, it should expect that unsolicited Map-Reply messages will be received from this Map-Resolver. Appropriate security mechanisms (e.g., Access Control Lists) SHOULD be activated to allow the processing of these incoming unsolicited Map-Reply messages.
- o If the B-bit is set, it should expect that (unsolicited) mapping bulk transfer messages are supported by this Map-Resolver. Appropriate security mechanisms (e.g., Access Control Lists) SHOULD be activated to allow the processing of these incoming unsolicited bulk transfer messages.
- o If the R-bit is set and the 'Redirect Map-Resolver IP Address' field embeds a valid IP address, the ITR MUST update its Map-Resolver contact information with the new Map-Resolver's IP address. The ITR MUST use that IP address for subsequent exchanges. Optionally, if filters were included in the reply sent by the Map-Resolver, these filters are used for the destination EID Map-Resolver address selection.

- o If the message includes one or several filters, the ITR MUST check whether the same set of filters as those included in the Map-Subscribe request are carried in the Map-Subscribe-Ack message. Filters that are not returned in the Map-Subscribe-Ack message may not be valid or have exceeded a limit. The "Result" code indicates the reason for not installing these filters. In particular:
 - * An ITR that receives the result code set to PARTIAL-FILTERS-INSTALLED-LIMIT MUST NOT try to install new filters unless it clears all the filters maintained by the Map-Resolver or it removes some of them.
 - * An ITR that receives the result code set to PARTIAL-FILTERS-INSTALLED-BAD MUST NOT resend the same filters that were not returned in the Map-Subscribe-Ack message, in subsequent Map-Subscribe requests.
 - * An ITR that receives the result code set to FILTERS-PROHIBITED MUST NOT include the filters that were not returned in the Map-Subscribe-Ack message, in a Map-Subscribe message sent to that Map-Resolver.
 - * An ITR that receives the result code set to PARTIAL-FILTERS-INSTALLED-LOCAL SHOULD wait for at least 60 seconds before issuing another Map-Subscribe message to install the filters that were not returned in the Map-Subscribe-Ack message.
- o The ITR MUST adjust the Expiry Timer carried in the Map-Subscribe-Ack. Subscription should be renewed before the expiry of that timer.

7. Subscription to Multiple Resolvers

In order to subscribe to multiple Map-Resolvers, an ITR sends Map-Subscribe messages to a list of Map-Resolvers. Each of these requests is handled as specified in [Section 4](#).

8. Sample Examples

This section includes a set of examples to illustrate the usage of the methods defined in [Section 2](#).

8.1. Map-Resolver Redirect

The example shown in Figure 3, illustrates an example of an ITR (ITR1) that is redirected to another Map-Resolver (MR2).

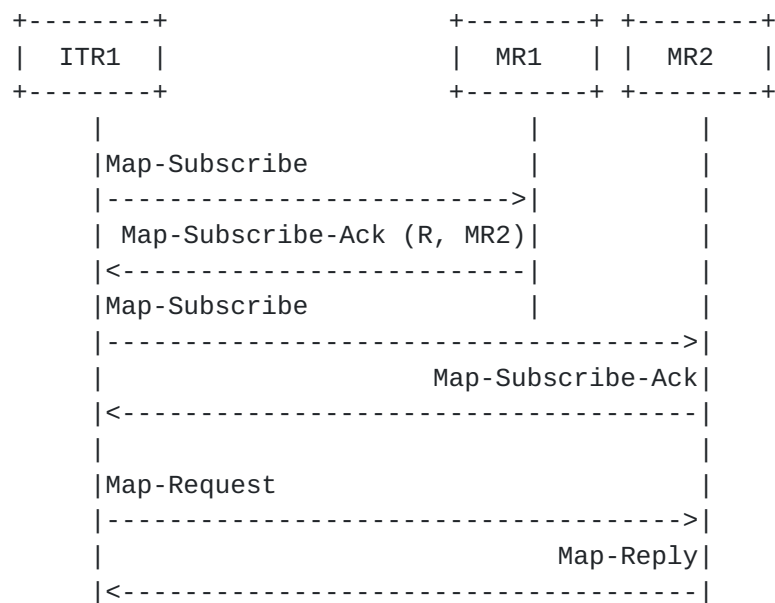


Figure 3: Example of Map-Resolver Redirection

8.2. Mapping Cache Retrieval

The examples shown in Figure 4 and Figure 5, illustrate examples of an ITR (ITR1) that restores its mapping tables upon reboot according to the filters already present in the mapping system.

The example in Figure 4, indicates how an ITR retrieves the mappings that match the filters included in the Map-Subscribe request using distinct Map-Reply messages.

The example in Figure 5, assumes that multiple records bound to distinct EIDs are included in the same Map-Reply message.

This procedure applies to ITRs which do not store the mapping table in a permanent memory storage facility.

Owing to this procedure, the ITR is ready-to-serve as soon as reboot is completed or right after a mapping cache clear event.


```

+-----+
|  ITR1  |
+-----+

+-----+
|  MR    |
+-----+

|
|
| Map-Subscribe(d_EID, d_EID2, |
| ..., d_EIDn)                |
|----->|
|   Map-Subscribe-Ack (d_EID, |
|       d_EID2, ..., d_EIDn) |
|<-----|
|
|   <<ITR Reboot>>
| Map-Subscribe(I)            |
|----->|
|   Map-Subscribe-Ack (I)     |
|<-----|
|       Map-Reply (d_EID)    |
|<-----|
|       Map-Reply (d_EID2)   |
|<-----|
|
|   ....
|       Map-Reply (d_EIDn)   |
|<-----|

```

Figure 4: Example of Mapping Cache Retrieval: Matching the Filters
Installed in the Mapping System

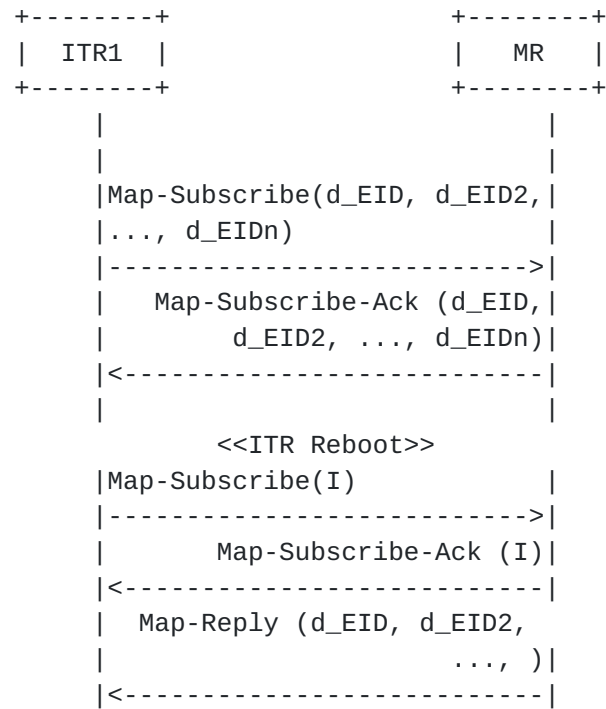


Figure 5: Example of Bulk Mapping Cache Retrieval: Matching the Filters Installed in the Mapping System

8.3. Unsolicited Map-Reply

The example shown in Figure 6, illustrates an ITR (ITR1) that is notified with the new EID-RLLOC mapping that it subscribed to.

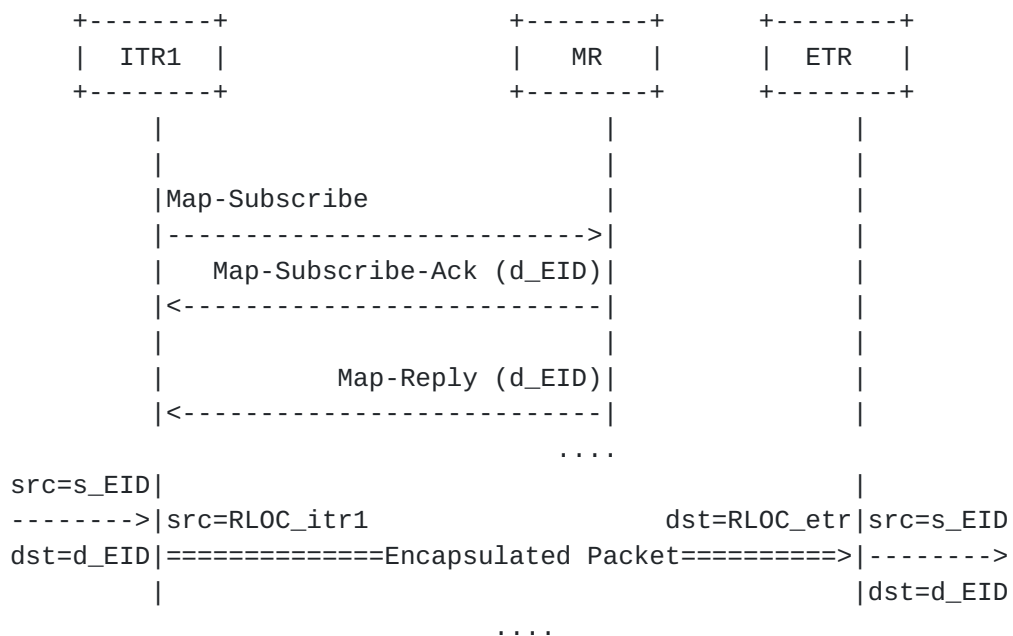


Figure 6: Example of Unsolicited Map-Reply

9. Security Considerations

This document does not introduce any additional security issues other than those discussed in [RFC6830] and [RFC6833].

10. IANA Considerations

This document requests IANA to assign a new code from the LISP Shared Extension Message Type Sub-types ([\[RFC8113\]](#)):

Message	Sub-type	Reference
=====	=====	=====
Map-Subscribe/Map-Subscribe-Ack	1024	[This document]

11. Acknowledgments

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

Many thanks to Stefano Secci and Chi-Dung Phung for their review.

12. References

12.1. 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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), DOI 10.17487/RFC4291, February 2006, <<http://www.rfc-editor.org/info/rfc4291>>.
- [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>>.
- [RFC8113] Boucadair, M. and C. Jacquenet, "Locator/ID Separation Protocol (LISP): Shared Extension Message & IANA Registry for Packet Type Allocations", [RFC 8113](#), DOI 10.17487/RFC8113, March 2017, <<http://www.rfc-editor.org/info/rfc8113>>.

12.2. Informative references

- [I-D.boucadair-lisp-bulk]
Boucadair, M. and C. Jacquenet, "LISP Mapping Bulk Retrieval", [draft-boucadair-lisp-bulk-04](#) (work in progress), February 2017.

Authors' Addresses

Mohamed Boucadair
Orange
Rennes 35000
France

EMail: mohamed.boucadair@orange.com

Christian Jacquenet
Orange
Rennes 35000
France

EMail: christian.jacquenet@orange.com