

LISP Working Group
Internet-Draft
Intended status: Experimental
Expires: February 3, 2018

A. Rodriguez-Natal
V. Ermagan
J. Leong
F. Maino
Cisco Systems
A. Cabellos-Aparicio
Technical University of Catalonia
S. Barkai
Fermi Serverless
D. Farinacci
lispers.net
August 2, 2017

**Publish-Subscribe mechanism for LISP
draft-rodrigueznatal-lisp-pubsub-00**

Abstract

This document describes simple extensions to the use of Map-Request to enable Publish/Subscribe (PubSub) operation for LISP.

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 February 3, 2018.

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](#)
- [2.](#) Requirements Language [3](#)
- [3.](#) Deployment Assumptions [3](#)
- [4.](#) Map-Request Additions [4](#)
- [5.](#) Mapping Request Subscribe Procedures [5](#)
- [6.](#) Mapping Notification Publish Procedures [6](#)
- [7.](#) Security Considerations [7](#)
- [8.](#) Acknowledgments [7](#)
- [9.](#) IANA Considerations [7](#)
- [10.](#) Normative References [7](#)
- Authors' Addresses [8](#)

1. Introduction

The Locator/ID Separation Protocol (LISP) [[RFC6830](#)] splits current IP addresses in two different namespaces, Endpoint Identifiers (EIDs) and Routing Locators (RLOCs). LISP uses a map-and-encap approach that relies on a Mapping System (basically a distributed database) that stores and disseminates EID-RLOC mappings and on LISP tunnel routers (xTRs) that encapsulate and decapsulate packets based on those mappings.

ITRs/RTRs/PITRs pull EID-to-RLOC mapping information from the Mapping System by means of explicitly requesting mappings. [[RFC6830](#)] indicates how ETRs can tell ITRs/RTRs/PITRs about mapping changes. This document presents a Publish/Subscribe (PubSub) architecture in which the Mapping System can notify ITRs/RTRs/PITRs about mapping changes. When this mechanism is used, mapping changes can be notified faster and can be managed in the Mapping System versus the LISP sites.

In general, when an ITR/RTR/PITR wants to be notified for mapping changes, the following flow occurs:

- (1) The ITR/RTR/PITR sends a Map-Request for the EID-prefix.
- (2) The ITR/RTR/PITR sets the subscribe bit on the Map-Request and includes its xTR-ID.

- (3) The Map-Request flows to one of the Map-Servers that the EID-prefix is registered to.
- (4) The Map-Server creates subscription state for the ITR/RTR/PITR on the EID-prefix.
- (5) The Map-Server sends a Map-Notify to the ITR/RTR/PITR to acknowledge the successful subscription.
- (6) When there is an RLOC-set change for the EID-prefix, the Map-Server sends a Map-Notify message to each ITR/RTR/PITR in the subscription list.
- (7) Each ITR/RTR/PITR sends a Map-Notify-Ack to acknowledge the received Map-Notify.

2. 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 [[RFC2119](#)].

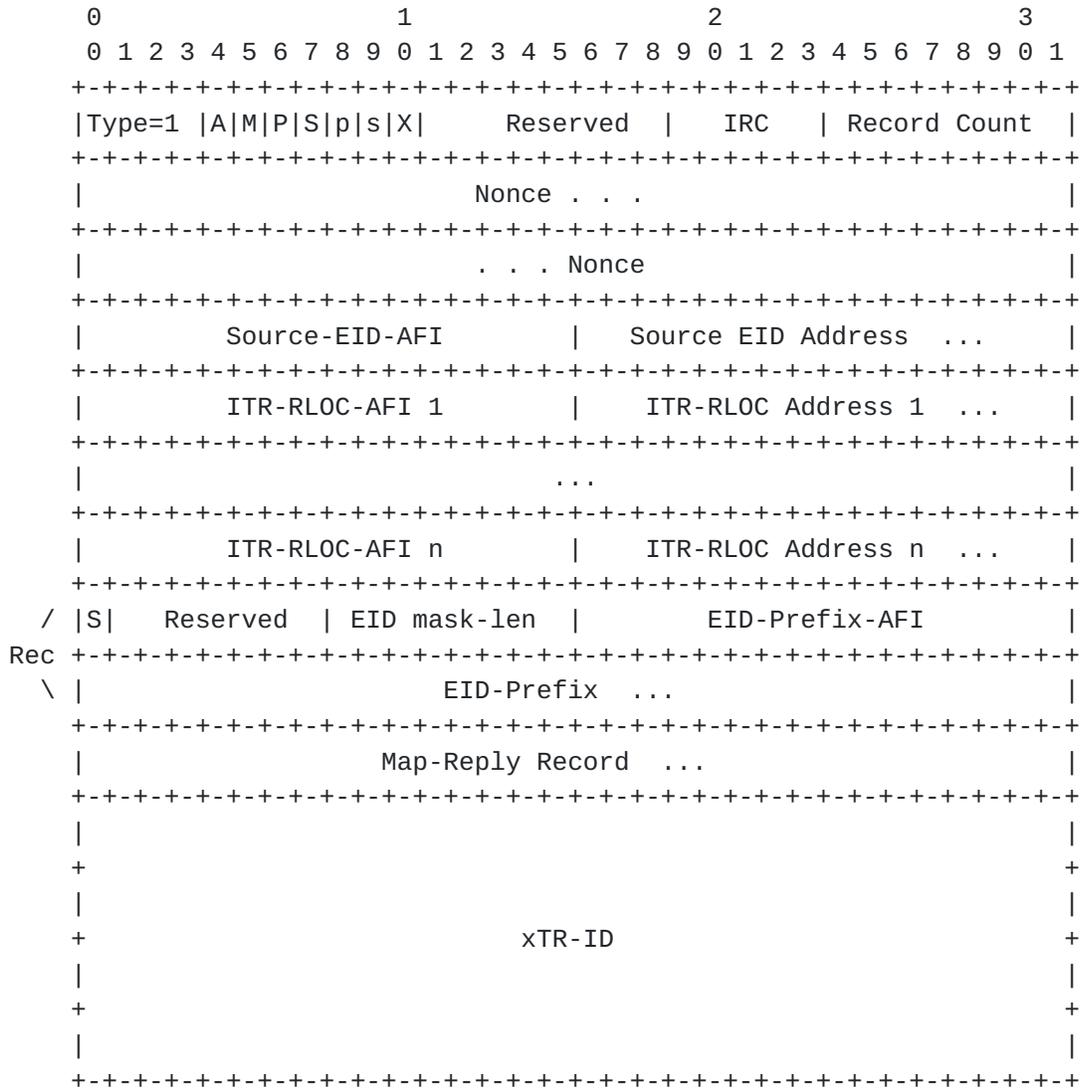
3. Deployment Assumptions

The specification described in this document makes the following deployment assumptions:

- (1) A unique 128-bit xTR-ID identifier is assigned to each xTR.
- (2) Map-Servers are configured in proxy-reply mode i.e. they send Map-Replies for the mappings they are serving.
- (3) There can be either a soft-state or hard-state security association between the xTRs and the Map-Servers.

The distribution of xTR-IDs and the management of security associations are out of the scope of this document.

4. Map-Request Additions



Map-Request with S bit and xTR-ID

The following fields are added to the Map-Request message defined in [RFC6830].

xTR-ID bit (X-bit): the first bit after the s bit in the Map-Request header. The X-bit of a Map-Request message is set to specify that a 128-bit xTR-ID field is appended to the end of the Map-Request, immediately following the last EID-Record (or the Map-Reply Record, if present).

xTR-ID field: The xTR-ID field uniquely identifies each xTR of a given LISP deployment. Provisioning of unique xTR-IDs is out of the scope of this document.

The following field is added to the EID-Record field of a Map-Request message defined in [\[RFC6830\]](#):

Subscribe bit (S-bit): the first bit in the EID-Record section of a Map-Request message. The S-bit of an EID-record is set to 1 to specify that the xTR wants to be notified of updates for that mapping record.

5. Mapping Request Subscribe Procedures

The xTR subscribes for RLOC-set changes for a given EID-prefix by sending a Map-Request to the Mapping System with the S-bit set on the EID-Record. The xTR builds a Map-Request according to [\[RFC6830\]](#) but also does the following:

- (1) The xTR MUST set the S-bit to 1 in each EID-Record to which the xTR wants to subscribe.
- (2) The X-bit of the Map-Request message MUST be set to 1, to specify the presence of an xTR-ID field that uniquely identifies the xTR.

The Map-Request is routed to the appropriate Map-Server through the Mapping System. No Map-Server is pre-assigned to handle the subscription state for a given xTR. The Map-Server that receives the Map-Request will be the Map-Server responsible to notify that specific xTR about future mapping changes for the subscribed mapping records.

Upon reception of the Map-Request, the Map-Server MUST process it as described in [\[RFC6830\]](#). After correct processing, for each EID-Record that has the S-bit set to 1, the Map-Server MUST try to add the xTR-ID contained in the Map-Request to the list of xTR that have requested to be subscribed to that mapping record.

If the xTR-ID is added to the list, the Map-Server MUST send a Map-Notify message back to the xTR to acknowledge the successful subscription. The Map-Server MUST follow the specification on [\[RFC6830\]](#) to build the Map-Notify with the following considerations. The Map-Server MUST use the nonce from the Map-Request as the nonce for the Map-Notify. The Map-Server MUST use its security association with the xTR (see [Section 3](#)) to compute the authentication data of the Map-Notify. The Map-Server MUST send the Map-Notify to one of the ITR-RLOCs received in the Map-Request.

When the xTR receives a Map-Notify with a nonce that matches one in the list of outstanding Map-Requests sent with a S-bit set, it knows that the Map-Notify is to acknowledge a successful subscription. The

xTR processes this Map-Notify as described in [RFC6830] with the following considerations. The xTR MUST use its security association with the Map-Server (see Section 3) to validate the authentication data on the Map-Notify. The xTR MUST use the Map-Notify to populate its map-cache with the returned EID-prefix and RLOC-set.

The subscription of an xTR-ID to the list of subscribers for the EID-Record may fail for a number of reasons. For example, because of local configuration policies (e.g. white/black lists of subscribers), or because the Map-Server has exhausted the resources to dedicate to the subscription of that EID-Record (e.g. the number of subscribers excess the capacity of the Map-Server).

If the subscription fails, the Map-Server MUST reply to the Map-Request with a Map-Reply as described in [RFC6830]. This is also the case when the Map-Server does not support PubSub operation. The xTR processes the Map-Reply as specified in [RFC6830].

If an xTR-ID is successfully added to the list of subscribers for an EID-Record, the Map-Server MUST extract the ITR-RLOCs present in the Map-Request, and store the association between the xTR-ID and those RLOCs. Any already present state regarding ITR-RLOCs for the same xTR-ID MUST be overwritten.

If the Map-Request only has one ITR-RLOC with AFI = 0 (i.e. Unknown Address), the Map-Server MUST remove the subscription state for that xTR-ID. In this case, the Map-Server MUST send the Map-Notify to the source RLOC of the Map-Request.

6. Mapping Notification Publish Procedures

The publish procedure is implemented via Map-Notify messages that the Map-Server sends to xTRs. The xTRs acknowledge the reception of Map-Notifies via sending Map-Notify-Ack messages back to the Map-Server. The complete mechanism works as follows.

When a mapping stored in a Map-Server is updated (e.g. via a Map-Register from an ETR), the Map-Server MUST notify the subscribers of that mapping via sending Map-Notify messages with the most updated mapping information. The Map-Notify message sent to each of the subscribers as a result of an update event MUST follow the exact encoding and logic defined in [RFC6830] for Map-Notify, except for the following.

- (1) The Map-Notify MUST be sent to one of the ITR-RLOCs associated with the xTR-ID of the subscriber.

- (2) The nonce of the Map-Notify MUST be randomly generated by the Map-Server.
- (3) The Map-Server MUST use its security association with the xTR to compute the authentication data of the Map-Notify.

When the xTR receives a Map-Notify with a nonce not present in any list of previously sent nonces, and an EID not local to the xTR, the xTR knows that the Map-Notify has been received due to an update on the RLOC-set of a cached mapping.

The xTR processes the received Map-Notify as specified in [[RFC6830](#)], with the following considerations. The xTR MUST use its security association with the Map-Server (see [Section 3](#)) to validate the authentication data on the Map-Notify. The xTR MUST use the mapping information carried in the Map-Notify to update its internal map-cache. The xTR MUST acknowledge the Map-Notify by sending back a Map-Notify-Ack (specified in [[I-D.ietf-lisp-rfc6833bis](#)]), with the nonce from the Map-Notify, to the Map-Server. If after a configurable timeout, the Map-Server has not received back the Map-Notify-Ack, it CAN try to send the Map-Notify to a different ITR-RLOC for that xTR-ID.

[7.](#) Security Considerations

The way to provide a security association between the ITRs and the Map-Servers must be evaluated according to the size of the deployment. For small deployments, it is possible to have a shared key (or set of keys) between the ITRs and the Map-Servers. For larger and Internet-scale deployments, scalability is a concern and further study is needed.

[8.](#) Acknowledgments

TBD.

[9.](#) IANA Considerations

This document makes no request to IANA.

[10.](#) Normative References

[[I-D.ietf-lisp-rfc6833bis](#)]
Fuller, V., Farinacci, D., and A. Cabellos-Aparicio,
"Locator/ID Separation Protocol (LISP) Control-Plane",
[draft-ietf-lisp-rfc6833bis-05](#) (work in progress), May
2017.

[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>>.

[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>>.

Authors' Addresses

Alberto Rodriguez-Natal
Cisco Systems
170 Tasman Drive
San Jose, CA
USA

Email: natal@cisco.com

Vina Ermagan
Cisco Systems
170 Tasman Drive
San Jose, CA
USA

Email: vermagan@cisco.com

Johnson Leong
Cisco Systems
170 Tasman Drive
San Jose, CA
USA

Email: joleong@cisco.com

Fabio Maino
Cisco Systems
170 Tasman Drive
San Jose, CA
USA

Email: fmaino@cisco.com

Albert Cabellos-Aparicio
Technical University of Catalonia
Barcelona
Spain

Email: acabello@ac.upc.edu

Sharon Barkai
Fermi Serverless
CA
USA

Email: sharon@fermicloud.io

Dino Farinacci
lispers.net
San Jose, CA
USA

Email: farinacci@gmail.com

