

LISP Working Group
Internet-Draft
Intended status: Experimental
Expires: September 12, 2019

A. Rodriguez-Natal
Cisco Systems
V. Ermagan
Google
J. Leong
F. Maino
Cisco Systems
A. Cabellos-Aparicio
Technical University of Catalonia
S. Barkai
Fermi Serverless
D. Farinacci
lispers.net
M. Boucadair
C. Jacquet
Orange
S. Secci
Cnam
March 11, 2019

Publish/Subscribe Functionality for LISP draft-ietf-lisp-pubsub-03

Abstract

This document specifies an extension 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 <https://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 September 12, 2019.

Copyright Notice

Copyright (c) 2019 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 (<https://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 PubSub Additions	4
5. Mapping Request Subscribe Procedures	6
6. Mapping Notification Publish Procedures	8
7. Security Considerations	9
8. Acknowledgments	9
9. IANA Considerations	9
10. Normative References	10
Authors' Addresses	10

[1. Introduction](#)

The Locator/ID Separation Protocol (LISP) [[I-D.ietf-lisp-rfc6833bis](#)] 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 (1) a Mapping System (basically a distributed database) that stores and disseminates EID-RLOC mappings and on (2) LISP tunnel routers (xTRs) that encapsulate and decapsulate data packets based on the content of those mappings.

Ingress Tunnel Routers (ITRs) / Re-encapsulating Tunnel Routers (RTRs) / Proxy Ingress Tunnel Routers (PITRs) pull EID-to-RLOC mapping information from the Mapping System by means of an explicit request message. [[I-D.ietf-lisp-rfc6833bis](#)] indicates how Egress Tunnel Routers (ETRs) can tell ITRs/RTRs/PITRs about mapping changes. This document presents a Publish/Subscribe (PubSub) extension 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 for a given EID-prefix, the following steps occur:

- (1) The ITR/RTR/PITR sends a Map-Request for that EID-prefix.
- (2) The ITR/RTR/PITR sets the Notification-Requested bit (N-bit) on the Map-Request and includes its xTR-ID and Site-ID.
- (3) The Map-Request is forwarded 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.

This operation is repeated for all EID-prefixes for which ITR/RTR/PITR want to be notified. The ITR/RTR/PITR can set the N-bit for several EID-prefixes within a single Map-Request.

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 (plus a 64-bit Site-ID) identifier is assigned to each xTR.
- (2) Map-Servers are configured in proxy-reply mode, i.e., they are solicited to generate and send Map-Reply messages 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 Site-IDs) as well as the management of security associations are out of the scope of this document and are for further study (see [Section 7](#)).

4. Map-Request PubSub Additions

Figure 1 shows the format of the updated Map-Request to support the PubSub functionality.




Figure 1: Map-Request with I-bit, N-bit, xTR-ID, and Site-ID

The following is added to the Map-Request message defined in Section 5.2 of [[I-D.ietf-lisp-rfc6833bis](#)]:

xTR-ID bit (I-bit): The I-bit of a Map-Request message is set to 1 to indicate that a 128 bit xTR-ID and a 64 bit Site-ID fields are present at the end of the Map-Request message. If an xTR is configured with an xTR-ID or Site-ID, it MUST set the I-bit to 1 and include its xTR-ID and Site-ID in the Map-Request messages it generates. If either the xTR-ID or Site-ID is not configured, an

unspecified value is encoded for whichever ID that is not configured.

Notification-Requested bit (N-bit): The N-bit of an EID-record is set to 1 to specify that the xTR wants to be notified of updates for that mapping record.

xTR-ID field: xTR-ID is a 128 bit field at the end of the Map-Request message, starting after the final Record in the message (or the Map-Reply Record, if present). The xTR-ID is used to uniquely identify the sender of a Map-Request message. The xTR-ID is defined in [[I-D.ietf-lisp-rfc6833bis](#)]

Site-ID field: Site-ID is a 64 bit field at the end of the Map-Request message, following the xTR-ID. Site-ID is used by the Map-Server receiving the Map-Request message to identify which xTRs belong to the same site. The Site-ID is defined in [[I-D.ietf-lisp-rfc6833bis](#)]

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 N-bit set on the EID-Record. The xTR builds a Map-Request according to [[I-D.ietf-lisp-rfc6833bis](#)] but also does the following:

- (1) The xTR MUST set the I-bit to 1 and append its xTR-ID and Site-ID to the Map-Request. The xTR-ID uniquely identifies the xTR.
- (2) The xTR MUST set the N-bit to 1 for each EID-Record to which the xTR wants to subscribe.

The Map-Request is forwarded to the appropriate Map-Server through the Mapping System. This document does not assume that a 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 receipt of the Map-Request, the Map-Server processes it as described in [[I-D.ietf-lisp-rfc6833bis](#)]. Furthermore, upon processing, for each EID-Record that has the N-bit set to 1, the Map-Server proceeds adding 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 in Section 6.1.7 of [[I-D.ietf-lisp-rfc6833bis](#)] to build the Map-Notify with the following considerations:

- (1) The Map-Server MUST use the nonce from the Map-Request as the nonce for the Map-Notify.
- (2) The Map-Server MUST use its security association with the xTR (see [Section 3](#)) to compute the authentication data of the Map-Notify.
- (3) 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-Request messages sent with an N-bit set, it knows that the Map-Notify is to acknowledge a successful subscription. The xTR processes this Map-Notify as described in [[I-D.ietf-lisp-rfc6833bis](#)] 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 (such as 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 send a Map-Reply to the originator of the Map-Request, as described in [[I-D.ietf-lisp-rfc6833bis](#)]. This is also the case when the Map-Server does not support PubSub operation. The xTR processes the Map-Reply as specified in [[I-D.ietf-lisp-rfc6833bis](#)].

If an xTR-ID is successfully added to the list of subscribers for an EID-Record, the Map-Server MUST extract the nonce and ITR-RLOCs present in the Map-Request, and store the association between the xTR-ID, ITR-RLOCs and nonce. Any already present state regarding ITR-RLOCs and/or nonce 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. When the TTL for the EID-record expires, the EID-prefix is removed from the Map-Server's subscription

cache. On EID-Record removal, the Map-Server notifies the subscribers via a Map-Notify with TTL equal 0.

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 [[I-D.ietf-lisp-rfc6833bis](#)] 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) For this version of the specification, the nonce in the Map-Notify sent as publication is set as follows. If the subscription state at the Map-Server was created by a received Map-Request with the N-bit set, the nonce in the Map-Notify sent as publication MUST be the one used in the Map-Request that created the subscription state. If the subscription state was created by explicit configuration at the Map-Server, the nonce in the Map-Notify sent as publication 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 sent previously in a Map-Request, or with a nonce not present in any list of previously sent nonces but with an EID not local to the xTR, the xTR knows that the Map-Notify has been received to update an entry on its map-cache. Processing of unsolicited Map-Notify messages MUST be explicitly enabled via configuration at the xTR.

The xTR processes the received Map-Notify as specified in [[I-D.ietf-lisp-rfc6833bis](#)], 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

Generic security considerations related to LISP control messages are discussed in [[I-D.ietf-lisp-rfc6833bis](#)].

In the particular case of PubSub, cache poisoning via malicious Map-Notify messages is avoided by the use of nonce and the security association between the ITRs and the Map-Servers. Nevertheless, 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. Such evaluation is one of the goals of this experimental specification.

Misbehaving nodes may send massive subscription requests which may lead to exhaust the resources of Map-Servers. Furthermore, frequently changing the state of a subscription may also be considered as an attack vector. To mitigate such issues, xTRs SHOULD rate-limit Map-Requests and Map-Servers SHOULD rate-limit Map-Notifies. Rate-limiting Map-Requests is discussed in [[I-D.ietf-lisp-rfc6833bis](#)] and the same guidelines apply here. To rate-limit Map-Notifies, a Map-Server MUST NOT send more than one Map-Notify per second to a particular xTR-ID. This parameter MUST be configurable. Note that when the Map-Notify rate-limit threshold is met for a particular xTR-ID, the Map-Server will silently discard additional subscription requests from that xTR-ID. Similarly, for pending mapping updates that need to be notified to that xTR-ID, the Map-Server will combine them into a single Map-Notify (with multiple EID-records) which it will send when the rate-limit mechanism allows it to transmit again Map-Notifies to that xTR-ID.

8. Acknowledgments

This work is partly funded by the ANR LISP-Lab project #ANR-13-INFR-009 (<https://www.lisp-lab.org>).

9. IANA Considerations

This document is requesting bit allocations in the Map-Request message from the "LISP Control Plane Header Bits" registry introduced in [[I-D.ietf-lisp-rfc6833bis](#)]. In particular, this document requests

allocating the following two bits from the sub-registry "Map-Request Header Bits". The position of these two bits in the Map-Request message can be found in Figure 1.

Spec Name	IANA Name	Bit Position	Description
I	map-request-I	11	xTR-ID Bit
N	map-request-N	... + 0	Notification-Requested Bit

Table 1: Additions to the LISP Map-Request Header Bits Sub-Registry

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-24](#) (work in progress), February
 2019.

[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/info/rfc2119>](https://www.rfc-editor.org/info/rfc2119).

Authors' Addresses

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

Email: natal@cisco.com

Vina Ermagan
 Google
 USA

Email: ermagan@gmail.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

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Christian Jacquenet
Orange
Rennes 35000
France

Email: christian.jacquenet@orange.com

Stefano Secci
Cnam
France

Email: stefano.secci@cnam.fr