Workgroup: LISP Working Group
Internet-Draft: draft-ietf-lisp-pubsub-15
Published: 28 February 2023
Intended Status: Standards Track
Expires: 1 September 2023
Authors: A. Rodriguez-Natal V. Ermagan A. Cabellos
Cisco Google UPC/BarcelonaTech
S. Barkai M. Boucadair
Nexar Orange

Publish/Subscribe Functionality for the Locator/ID Separation Protocol (LISP)

Abstract

This document specifies an extension to the request/reply based Locator/ID Separation Protocol (LISP) control plane to enable Publish/Subscribe (PubSub) operation.

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 1 September 2023.

Table of Contents

1. Introduction

The Locator/ID Separation Protocol (LISP) [RFC9300] [RFC9301] splits IP addresses into two different namespaces: Endpoint Identifiers (EIDs) and Routing Locators (RLOCs). LISP uses a map and encapsulate (a.k.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. Section 6.1 of [RFC9301] 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 main steps occur:

(1) The ITR/RTR/PITR builds a Map-Request for that EID-Prefix with the Notification-Requested bit (N-bit) set and which also includes its xTR-ID and Site-ID.

(2) The Map-Request is forwarded to one of the Map-Servers that the EID-Prefix is registered to.

(3) The Map-Server creates subscription state for the ITR/RTR/PITR on the EID-Prefix.

(4) The Map-Server sends a Map-Notify to the ITR/RTR/PITR to confirm that the subscription has been created and then waits for an acknowledgement of the notification.

(5) The ITR/RTR/PITR sends back a Map-Notify-Ack to acknowledge the successful receipt of the Map-Notify.

(6) When there is a change in the mapping of 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 ITRs/RTRs/PITRs want to be notified. An ITR/RTR/PITR can set the N-bit for several EID-Prefixes within a single Map-Request. Please note that the steps above illustrate only the simplest scenario and that details for this and other scenarios are described later in the document.

The reader may refer to [I-D.boucadair-lisp-pubsub-flow-examples] for sample flows to illustrate the use of the PubSub specification.

1.1. Scope of Applicability

The PubSub procedure specified in this document is intended to be used in contexts with controlled access to the Map-Server. How a deployment controls access to a Map-Server is deployment specific, and therefore out of the scope of this document. However, the Map-Resolvers and Map-Servers need to be configured with the required information to at least ensure the following:

(1) Map-Resolvers MUST verify that an xTR is allowed to (1) set the N-bit to 1 and (2) use the xTR-ID, Site-ID, and ITR-RLOCs included in a Map-Request.

(2) Map-Servers MUST only accept subscription requests from Map-Resolvers that verify Map-Requests as previously described.

2. Terminology and Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

The document uses the terms defined in Section 3 of [RFC9300].

3. Deployment Requirements

In addition to the general assumptions and expectations that [RFC9301] makes for LISP deployments, this document makes the following deployment requirements:

(1) A unique 128-bit xTR-ID (plus a 64-bit Site-ID) identifier is assigned to each xTR.

(2) Map-Servers are configured to proxy Map-Replying (i.e., they are solicited to generate and send Map-Reply messages) for the mappings they are serving.

(3) A security association (e.g., a PubSubKey) is required between the ITRs and the Map-Servers (see Section 7.1).

If a requirement is not met, a subscription cannot be established, and the network will continue operating without this enhancement. The configuration of xTR-IDs and Site-IDs is out of the scope of this document. The reader may refer to [I-D.ietf-lisp-yang] for an example of how these identifiers can be provisioned to LISP nodes.

4. Map-Request PubSub Additions

Figure 1 shows the format of the updated Map-Request to support the PubSub functionality. In particular, this document associates a meaning with one of the reserved bits (see Section 8).

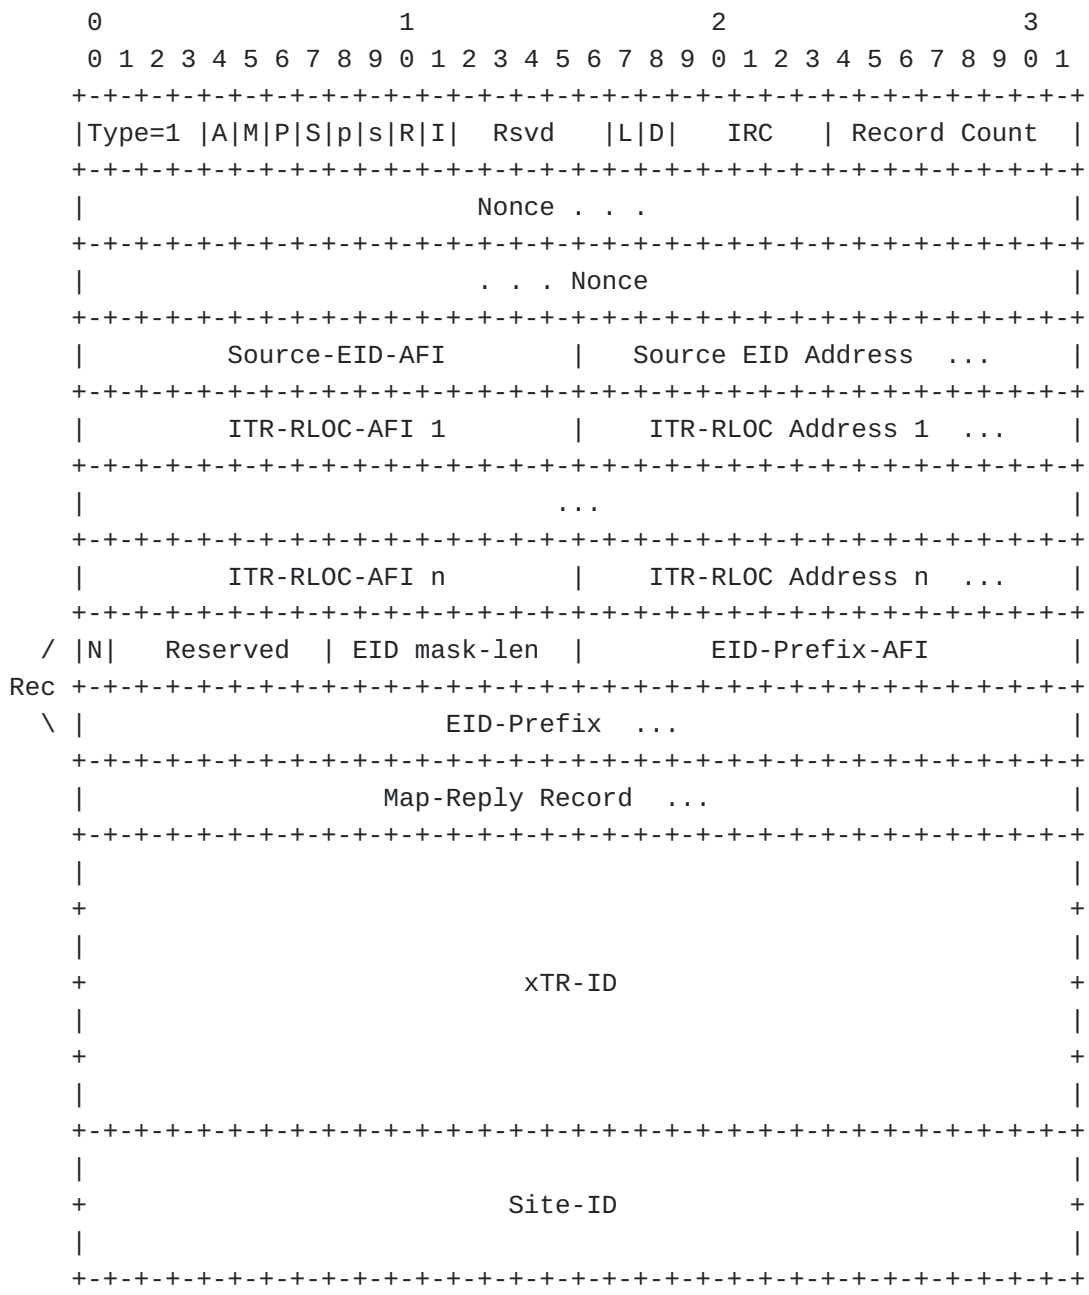

Figure 1: Map-Request with I-bit, N-bit, xTR-ID, and Site-ID

The following is added to the Map-Request message defined in Section 5.2 of [RFC9301]:

xTR-ID bit (I-bit): This bit is set to 1 to indicate that 128-bit xTR-ID and 64-bit Site-ID fields are present in the Map-Request message. For PubSub operation, an xTR MUST be configured with an xTR-ID and Site-ID, and it MUST set the I-bit to 1 and include its xTR-ID and Site-ID in the Map-Request messages it generates. If the I-bit is set, but the Site-ID and/or xTR-ID are not included, a receiver can detect the error because, after processing that last EID-record, there are no bytes left from

processing the message. In this case, the receiver SHOULD log a malformed Map-Request and MUST drop the message.

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 EID-Prefix.

xTR-ID field: If the I-bit is set, this field is added to the Map-Request message as shown in Figure 1, starting right after the final Record in the message (or the Map-Reply Record, if present). The xTR-ID is specified in Section 5.6 of [RFC9301].

Site-ID field: If the I-bit is set, this field is added to the Map-Request message as shown in Figure 1, following the xTR-ID field. The Site-ID is defined in Section 5.6 of [RFC9301].

5. Mapping Request Subscribe Procedures

The xTR subscribes for changes, to 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 Section 5.3 of [RFC9301] 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.

(2) The xTR MUST set the N-bit to 1 for the EID-Record to which the xTR wants to subscribe.

(3) If the xTR has a nonce associated with the EID-Prefix, it MUST use this nonce increased by one in the Map-Request. Otherwise, it generates a nonce following Section 5.2 of [RFC9301]. It is RECOMMENDED that the xTR uses persistent storage to keep nonce state. If the xTR does not have persistent storage and does not have a nonce associated with the EID-Prefix, it MUST reset the nonce by using the procedure described in Section 7.1 to successfully create a new security association with the Map-Server.

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 for notifying 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 Section 8.3 of [RFC9301]. In addition, unless the xTR is using the procedure described in Section 7.1 to create a new security association, the Map-Server MUST verify that the nonce in

the Map-Request is greater than the stored nonce (if any) associated with the xTR-ID (and EID-Prefix, when applicable). Otherwise, the Map-Server MUST silently drop the Map-Request message and SHOULD log the event to record that a replay attack could have occurred. Furthermore, upon processing, for the EID-Record that has the N-bit set to 1, the Map-Server proceeds to add the xTR-ID contained in the Map-Request to the list of xTRs that have requested to be subscribed to that EID-Prefix.

If an xTR-ID is successfully added to the list of subscribers for an EID-Prefix, the Map-Server MUST extract the nonce and ITR-RLOCs present in the Map-Request, and store the association between the EID-Prefix, xTR-ID, ITR-RLOCs, and nonce. Any already present state regarding ITR-RLOCs and/or nonce for the same xTR-ID MUST be overwritten. When the LISP deployment has a single Map-Server, the Map-Server can be configured to keep a single nonce per xTR-ID for all EID-Prefixes (when used, this option MUST be enabled at the Map-Server and all xTRs).

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 builds the Map-Notify according to Sections 5.5 and 5.7 of [RFC9301] 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 (Section 7.1) to sign the authentication data of the Map-Notify. The xTR MUST use the security association to verify the received authentication data.

(3) The Map-Server MUST send the Map-Notify to one of the ITR-RLOCs received in the Map-Request (which one is implementation specific).

As a reminder, the initial transmission and retransmission of Map-Notify messages by a Map-Server follow the procedure specified in Section 5.7 of [RFC9301]. Some state changes may trigger an overload that would impact, e.g., the outbound capacity of a Map-Server. A similar problem may be experienced when a large number of state entries are simultaneously updated. To prevent such phenomena, Map-Servers SHOULD be configured with policies to control the maximum number of subscriptions and also the pace of Map-Notify messages. For example, the Map-Server may be instructed to limit the resources that are dedicated to unsolicited Map-Notify messages to a small fraction (e.g., less than 10%) of its overall processing and forwarding capacity. The exact details to characterize such policies are deployment and implementation specific. Likewise, this document

does not specify which notifications take precedence when these policies are enforced.

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 Section 5.7 of [RFC9301], and MUST use the Map-Notify to populate its Map-Cache with the returned EID-Prefix and RLOC-set. As a reminder, following Section 5.7 of [RFC9301], the xTR has to send a Map-Notify-Ack back to the Map-Server. If the Map-Server does not receive the Map-Notify-Ack after exhausting the Map-Notify retransmissions described in Section 5.7 of [RFC9301], the Map-Server can remove the subscription state. If the Map-Server removes the subscription state, and absent explicit policy, it SHOULD notify the xTR by sending a single Map-Notify with the same nonce but with Loc-Count = 0 (and Loc-AFI = 0), and ACT bits set to 5 "Drop/Auth-Failure". It is OPTIONAL for the xTR to update its map-cache entry for the EID-Prefix (if any) based on this Map-Notify. This message is specifically useful for cases where Map-Notifies are successfully received by an xTR but the corresponding Map-Notify-Acks were lost when forwarded to the Map-Server. xTR implementations can use this signal to try to reinstall their subscription state instead of maintaining stale mappings.

The subscription of an xTR-ID may fail for a number of reasons. For example, it fails because of local configuration policies (such as accept and drop lists of subscribers), because the Map-Server has exhausted the resources to dedicate to the subscription of that EID-Prefix (e.g., the number of subscribers excess the capacity of the Map-Server), or because the xTR tried but was not successful in establishing a new security association (Section 7.1).

If the subscription request fails, the Map-Server sends a Map-Reply to the originator of the Map-Request, as described in Section 8.3 of [RFC9301], with the following considerations:

* If the subscription request fails due to policy (e.g. for explicitly configured subscriptions, as described later in this section) the Map-Server MUST respond to the Map-Request with a Negative Map-Reply (Loc-Count = 0 and Loc-AFI = 0) with ACT bits set to 4 "Drop/Policy-Denied".

* If the subscription request fails due to authentication (e.g. when a new security associationg is being established, as described in Section 7.1), the Map-Server MUST respond to the Map-Request with a Negative Map-Reply (Loc-Count = 0 and Loc-AFI = 0) with ACT bits set to 5 "Drop/Auth-Failure".

*If the subscription request fails due to any other reason, the Map-Server MUST follow Section 8.3 of [RFC9301] with no changes.

The xTR processes any (Negative) Map-Reply as specified in Section 8.1 of [RFC9301], with the following considerations: if the xTR receives a Negative Map-Reply with ACT bits set to 4 "Drop/Policy-Denied" or 5 "Drop/Auth-Failure" as a response to a subscription request, it is OPTIONAL for the xTR to update its map-cache entry for the EID-Prefix (if any) based on this Negative Map-Reply. If the subscription request fails (for whichever reason), it is up to the implementation of the xTR to try to subscribe again.

If the Map-Server receives a subscription request for an EID-Prefix not present in the mapping database, it SHOULD follow the same logic described in Section 8.4 of [RFC9301] and create a temporary subscription state for the xTR-ID to the least-specific prefix that both matches the original query and does not match any EID-Prefix known to exist in the LISP-capable infrastructure. Alternatively, the Map-Server can instead determine that such a subscription request fails, and send a Negative Map-Reply following Section 8.3 of [RFC9301]. In both cases, the TTL of the temporary subscription state or the Negative Map-Reply SHOULD be configurable, with a value of 15-minutes being RECOMMENDED.

The subscription state can also be created explicitly by configuration at the Map-Server (possible when a pre-shared security association exists, see Section 7) using a variety of means that are out of scope. If at the time the explicit subscription is configured there is no nonce that can be used for the explicit subscription state (e.g., from a different subscription already established with the same xTR when a single nonce is kept per xTR-ID), then both the xTR and Map-Server MUST be configured with the initial nonce to use. It is RECOMMENDED to have a configuration option to enable (or disable) the xTR to accept publication information for EID-Prefixes the xTR did not explicitly subscribe to. By default, the xTR is allowed to modify explicitly configured subscription state following the procedures described in this section, however this MAY be disabled at the Map-Server via configuration. If the Map-Server is instructed to not allow xTRs to modify explicitly configured subscriptions, and an xTR tries to do so, this triggers a Negative Map-Reply with ACT bits set to 4 "Drop/Policy-Denied" as described earlier in this section.

The following specifies the procedure to remove a subscription: If a valid Map-Request with the N-bit set to 1 only has one ITR-RLOC with AFI = 0 (i.e., Unknown Address), the Map-Server MUST remove the subscription state for that xTR-ID (unless this is disabled via configuration, see previous paragraph). If the subscription state is removed, the Map-Server MUST send a Map-Notify to the source RLOC of

the Map-Request. If the subscription removal fails due to configuration, this triggers a Negative Map-Reply with with ACT bits set to 4 "Drop/Policy-Denied" as described earlier in this section; the Map-Server sends the Negative Map-Reply to the source RLOC of the Map-Request in this case. Removing subscription state at the Map-Server can lead to replay attacks. To soften this, the Map-Server SHOULD keep the last nonce seen per xTR-ID (and EID-Prefix, when applicable). If the Map-Server does not keep last nonces seen, then the Map-Server MUST require the xTRs to subscribe using the procedure described in Section 7.1 to create a new security association with the Map-Server.

If the Map-Server receives a Map-Request asking to remove a subscription for an EID-Prefix without subscription state for that xTR-ID, but covered by a less-specific EID-Prefix for which subscription state exists for the xTR-ID, the Map-Server SHOULD stop publishing updates about this more-specific EID-Prefix to that xTR, until the xTR subscribes to the more-specific EID-Prefix. The same considerations regarding authentication, integrity protection, and nonce checks described in this section and Section 7 for Map-Requests used to update subscription state, apply for Map-Requests used to remove subscription state.

When an EID-Prefix is removed from the Map-Server (either when explicitly withdrawn or when its TTL expires), the Map-Server notifies its subscribers (if any) 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. If subscription state in the Map-Server exists for a less-specific EID-Prefix and a more-specific EID-Prefix is updated, then the Map-Notify is sent with the more-specific EID-Prefix mapping to the subscribers of the less-specific EID-Prefix mapping. The Map-Notify message sent to each of the subscribers as a result of an update event follows the encoding and logic defined in Section 5.7 of [RFC9301] 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 (which one is implementation specific).

(2) The Map-Server increments the nonce by one every time it sends a Map-Notify as publication to an xTR-ID for a particular EID-Prefix.

(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 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. The xTR MUST keep track of the last nonce seen in a Map-Notify received as a publication from the Map-Server for the EID-Prefix. When the LISP deployment has a single Map-Server, the xTR can be configured to keep track of a single nonce for all EID-Prefix (when used, this option MUST be enabled at the Map-Server and all xTRs). If a Map-Notify received as a publication has a nonce value that is not greater than the saved nonce, the xTR drops the Map-Notify message and logs the fact a replay attack could have occurred. The same considerations discussed in Section 5.6 of [RFC9301] regarding Map-Register nonces apply here for Map-Notify nonces.

The xTR processes the received Map-Notify as specified in Section 5.7 of [RFC9301], with the following considerations: The xTR MUST use its security association with the Map-Server (Section 7.1) 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. If after following Section 5.7 of [RFC9301] regarding retransmission of Map-Notify messages, 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. If the Map-Server tries all the ITR-RLOCs without receiving a response, it may stop trying to send the Map-Notify.

7. Security Considerations

Generic security considerations related to LISP control messages are discussed in Section 9 of [RFC9301].

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.

It is RECOMMENDED to follow guidance from the last paragraph of Section 9 of [RFC9301] to ensure integrity protection of Map-Request messages (e.g., to prevent xTR-ID hijacking).

7.1. Security Association between ITR and Map-Server

Since Map-Notifies from the Map-Server to the ITR need to be authenticated, there is a need for a soft-state or hard-state security association (e.g., a PubSubKey) between the ITRs and the Map-Servers. For some controlled deployments, it might be possible to have a shared PubSubKey (or set of keys) between the ITRs and the Map-Servers. However, if pre-shared keys are not used in the deployment, LISP-SEC [RFC9303] can be used as follows to create a security association between the ITR and the Map-Server.

First, when the ITR is sending a Map-Request with the N-bit set following Section 5, the ITR also performs the steps described in Section 6.4 of [RFC9303]. The ITR can then generate a PubSubKey by deriving a key from the One-Time Key (OTK) and Map-Request's nonce as follows: PubSubKey = KDF(OTK + nonce), where KDF is the Key Derivation Function indicated by the OTK Wrapping ID. If OTK Wrapping ID equals NULL-KEY-WRAP-128, then the PubSubKey is the OTK. Note that as opposed to the pre-shared PubSubKey, this generated PubSubKey is different per EID-Prefix to which an ITR subscribes (since the ITR will use a different OTK per Map-Request).

When the Map-Server receives the Map-Request it follows the procedure specified in Section 5 with the following considerations: The Map-Server MUST verify that the OTK has not been used before. If the Map-Server verifies the OTK and cannot determine that the OTK has not been used before, the subscription request fails due to authentication and this triggers a Negative Map-Reply with ACT bits set to 5 "Drop/Auth-Failure", as described in Section 5. The xTR might try again with a different OTK upon reception of this Negative Map-Reply. Note that a Map-Server implementation might decide to not keep full past OTKs and instead use some form of hash. In that case, hash collisions are handled as if the OTK has been reused. Such an implementation needs to balance the hash length with the rate of collisions expected for the particular deployment; this is implementation specific. If the Map-Server has to reply with a Map-Reply for any other reason (e.g., if PubSub is not supported or a subscription is not accepted), then it follows normal LISP-SEC procedure described in Section 5.7 of [RFC9303]. No PubSubKey, security association, or subscription state is created when the Map-Server responds with any Map-Reply message.

Otherwise, if the Map-Server has to reply with a Map-Notify (e.g., due to subscription accepted) to a received Map-Request, the following extra steps take place:

* The Map-Server extracts the OTK and OTK Wrapping ID from the LISP-SEC ECM Authentication Data.

*The Map-Server generates a PubSubKey by deriving a key from the OTK as described before for the ITR. This is the same PubSubKey derived at the ITR which is used to establish a security association between the ITR and the Map-Server.

*The PubSubKey can now be used to sign and authenticate any Map-Notify between the Map-Server and the ITR for the subscribed EID-Prefix. This includes the Map-Notify sent as a confirmation to the subscription. When the ITR wants to update the security association for that Map-Server and EID-Prefix, it once again follows the procedure described in this section.

Note that if the Map-Server replies with a Map-Notify, none of the regular LISP-SEC steps regarding Map-Reply described in Section 5.7 of [RFC9303] occur.

7.2. DDoS Attack Mitigation

If PubSub is deployed under the scope of applicability defined in Section 1.1 only known nodes can participate on the PubSub deployment. DDoS attacks based on replayed messages by unknown nodes are avoided by the use of nonce and the security association between the ITRs and the Map-Servers. Misbehaving known nodes may send massive subscription requests which may lead to exhausting the resources of a Map-Server. Furthermore, frequently changing the state of a subscription may also be considered as an attack vector. To mitigate such issues, Section 5.3 of [RFC9301] discusses rate-limiting Map-Requests and Section 5.7 of [RFC9301] discusses rate-limiting Map-Notifies. Note that when the Map-Notify rate-limit threshold is met for a particular xTR-ID, the Map-Server will discard additional subscription requests from that xTR-ID and will fall back to [RFC9301] behavior when receiving a Map-Request from that xTR-ID (i.e., the Map-Server will send a Map-Reply).

8. IANA Considerations

This document requests IANA to assign a new bit from the "LISP Control Plane Header Bits: Map-Request" sub-registry under the "Locator/ID Separation Protocol (LISP) Parameters" registry available at [IANA-LISP]. The suggested position of this bit in the Map-Request message can be found in Figure 1.

Spec Name	IANA Name	Bit Position	Description	Reference
I	Map-Request-I	11	xTR-ID Bit	This-Document

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

This document also requests the creation of a new sub-registry entitled "LISP Control Plane Header Bits: Map-Request-Record" under

the "Locator/ID Separation Protocol (LISP) Parameters" registry available at [IANA-LISP].

The initial content of this sub-registry is shown in Table 2:

Spec Name	IANA Name	Bit Position	Description	Reference
N	Map-Request-N	1	Notification-Requested Bit	This-Document

Table 2: Initial Content of LISP Control Plane Header Bits: Map-Request-Record Sub-Registry

The remaining bits (i.e., Bit positions 2-8) are Unassigned.

The policy for allocating new bits from this sub-registry is Specification Required (Section 4.6 of [RFC8126]).

Review requests are evaluated on the advice of one or more designated experts. Criteria that should be applied by the designated experts include determining whether the proposed registration duplicates existing entries and whether the registration description is sufficiently detailed and fits the purpose of this registry. These criteria are considered in addition to those already provided in Section 4.6 of [RFC8126] (e.g., the proposed registration must be documented in a permanent and readily available public specification). The designated experts will either approve or deny the registration request, communicating this decision to IANA. Denials should include an explanation and, if applicable, suggestions as to how to make the request successful.

9. References

9.1. 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, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26,

RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC9300] Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos, Ed., "The Locator/ID Separation Protocol (LISP)", RFC 9300, DOI 10.17487/RFC9300, October 2022, <https://www.rfc-editor.org/info/rfc9300>.

[RFC9301] Farinacci, D., Maino, F., Fuller, V., and A. Cabellos, Ed., "Locator/ID Separation Protocol (LISP) Control Plane", RFC 9301, DOI 10.17487/RFC9301, October 2022, <https://www.rfc-editor.org/info/rfc9301>.

[RFC9303] Maino, F., Ermagan, V., Cabellos, A., and D. Saucez, "Locator/ID Separation Protocol Security (LISP-SEC)", RFC 9303, DOI 10.17487/RFC9303, October 2022, <https://www.rfc-editor.org/info/rfc9303>.

9.2. Informative References

[I-D.boucadair-lisp-pubsub-flow-examples]
Boucadair, M., "LISP PubSub Flow Examples", Work in Progress, Internet-Draft, draft-boucadair-lisp-pubsub-flow-examples-03, 10 February 2023, <https://datatracker.ietf.org/doc/html/draft-boucadair-lisp-pubsub-flow-examples-03>.

[I-D.haindl-lisp-gb-atn] Haindl, B., Lindner, M., Moreno, V., Portoles-Comeras, M., Maino, F., and B. Venkatachalapathy, "Ground-Based LISP for the Aeronautical Telecommunications Network", Work in Progress, Internet-Draft, draft-haindl-lisp-gb-atn-08, 23 September 2022, <https://datatracker.ietf.org/doc/html/draft-haindl-lisp-gb-atn-08>.

[I-D.ietf-lisp-eid-mobility] Portoles-Comeras, M., Ashtaputre, V., Maino, F., Moreno, V., and D. Farinacci, "LISP L2/L3 EID Mobility Using a Unified Control Plane", Work in Progress, Internet-Draft, draft-ietf-lisp-eid-mobility-11, 10 January 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-lisp-eid-mobility-11>.

[I-D.ietf-lisp-yang]
Ermagan, V., Rodriguez-Natal, A., Coras, F., Moberg, C., Rahman, R., Cabellos-Aparicio, A., and F. Maino, "LISP

YANG Model", Work in Progress, Internet-Draft, draft-ietf-lisp-yang-18, 29 August 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-lisp-yang-18>.

[I-D.moreno-lisp-uberlay] Moreno, V., Farinacci, D., Rodriguez-Natal, A., Portoles-Comeras, M., Maino, F., and S. Hooda, "Uberlay Interconnection of Multiple LISP overlays", Work in Progress, Internet-Draft, draft-moreno-lisp-uberlay-06, 28 September 2022, <https://datatracker.ietf.org/doc/html/draft-moreno-lisp-uberlay-06>.

[IANA-LISP] IANA, "Locator/ID Separation Protocol (LISP) Parameters", <https://www.iana.org/assignments/lisp-parameters/lisp-parameters.xhtml>.

[RFC6835] Farinacci, D. and D. Meyer, "The Locator/ID Separation Protocol Internet Groper (LIG)", RFC 6835, DOI 10.17487/RFC6835, January 2013, <https://www.rfc-editor.org/info/rfc6835>.

Appendix A. Sample PubSub Deployment Experiences

Some LISP production networks have been running different forms of PubSub for some time. The following subsections provide an inventory of some experience lessons from these deployments.

A.1. PubSub as a Monitoring Tool

Some LISP deployments are using PubSub as a way to monitor EID-Prefixes (particularly, EID-to-RLOC mappings). To that aim, some LISP implementations have extended the LISP Internet Groper (lig) [RFC6835] tool to use PubSub. Such an extension is meant to support an interactive mode with lig, and request subscription for the EID of interest. If there are RLOC changes, the Map-Server sends a notification and then the lig client displays that change to the user.

A.2. Mitigating Negative Map-Cache Entries

Section 8.1 of [RFC9301] suggests two TTL values for Negative Map-Replies: either 15-minute (if the EID-Prefix does not exist) or 1-minute (if the prefix exists but has not been registered). While these values are based on the original operational experience of the LISP protocol designers, negative cache entries have two unintended effects that were observed in production.

First, if the xTR keeps receiving traffic for a negative EID destination (i.e., an EID-Prefix with no RLOCs associated with it), it will try to resolve the destination again once the cached state

expires, even if the state has not changed in the Map-Server. It was observed in production that this is happening often in networks that have a significant amount of traffic addressed for outside of the LISP network. This might result on excessive resolution signaling to keep retrieving the same state due to the cache expiring. PubSub is used to relax TTL values and cache negative mapping entries for longer periods of time, avoiding unnecessary refreshes of these forwarding entries, and drastically reducing signaling in these scenarios. In general, a TTL-based schema is a “polling mechanism” that leads to more signaling where PubSub provides an "event triggered mechanism" at the cost of state.

Second, if the state does indeed change in the Map-Server, updates based on TTL timeouts might prevent the cached state at the xTR from being updated until the TTL expires. This behavior was observed during configuration (or reconfiguration) periods on the network, where no-longer-negative EID-Prefixes do not receive the traffic yet due to stale Map-Cache entries present in the network. With the activation of PubSub, stale caches can be updated as soon as the state changes.

A.3. Improved Mobility Latency

An improved convergence time was observed on the presence of mobility events on LISP networks running PubSub as compared with running LISP [RFC9301]. As described in Section 4.1.2.1 of [I-D.ietf-lisp-eid-mobility], LISP can rely on data-driven Solicit-Map-Requests (SMRs) to ensure eventual network converge. Generally, PubSub offers faster convergence due to (1) no need to wait for a data triggered event and (2) less signaling as compared with the SMR-based flow. Note that when a Map-Server running PubSub has to update a large number of subscribers at once (i.e., when a popular mapping is updated) SMR based convergence may be faster for a small subset of the subscribers (those receiving PubSub updates last). Deployment experience reveals that data-driven SMRs and PubSub mechanisms complement each other and provide a fast and resilient network infrastructure in the presence of mobility events.

Furthermore, experience showed that not all LISP entities on the network need to implement PubSub for the network to get the benefits. In scenarios with significant traffic coming from outside of the LISP network, the experience showed that enabling PubSub in the border routers significantly improves mobility latency overall. Even if edge xTRs do not implement PubSub, and traffic is exchanged between EID-Prefixes at the edge, xTRs still converge based on data-driven events and SMR-triggered updates.

A.4. Enhanced Reachability with Dynamic Redistribution of Prefixes

There is a need to interconnect LISP networks with other networks that might or might not run LISP. Some of those scenarios are similar to the ones described in [I-D.haindl-lisp-gb-atn] and [I-D.moreno-lisp-uberlay]. When connecting LISP to other networks, the experience revealed that in many deployments the point of interaction with the other domains is not the Mapping System but rather the border router of the LISP site. For those cases the border router needs to be aware of the LISP prefixes to redistribute them to the other networks. Over the years different solutions have been used.

First, Map-Servers were collocated with the border routers, but this was hard to scale since border routers scale at a different pace than Map-Servers. Second, decoupled Map-Servers and border routers were used with static configuration of LISP entries on the border, which was problematic when modifications were made. Third, a routing protocol (e.g., BGP) can be used to redistributed LISP prefixes from the Map-Servers to a border router, but this comes with some implications, particularly the Map-Servers needs to implement an additional protocol which consumes resources and needs to be properly configured. Therefore, once PubSub was available, deployments started to adapt it to enable border routers to dynamically learn the prefixes they need to redistribute without the need of extra protocols or extra configuration on the network.

In other words, PubSub can be used to discover EID-Prefixes so they can be imported into other routing domains that do not use LISP. Similarly, PubSub can also be used to discover when EID-Prefixes need to be withdrawn from other routing domains. That is, in a typical deployment, a border router will withdraw an EID-Prefix it has been announcing to external routing domains, if it receives a notification that the RLOC-set for that EID-Prefix is now empty.

A.5. Better Serviceability

EID-to-RLOC mappings can have very long TTL, sometimes in the order of several hours. Upon the expiry of that TTL, the xTR checks if these entries are being used and removes any entry that is not being used. The problem with very long Map-Cache TTL is that (in the absence of PubSub) if a mapping changes, but it is not being used, the cache remains but it is stale. This is due to no data traffic being sent to the old location to trigger an SMR based Map-Cache update as described in Section 4.1.2.1 of [I-D.ietf-lisp-eid-mobility]. If the network operator runs a show command on a router to track the state of the Map-Cache, the router will display multiple entries waiting to expire but with stale RLOC information. This might be confusing for operators sometimes,

particularly when they are debugging problems. With PubSub, the Map-Cache is updated with the correct RLOC information, even when it is not being used or waiting to expire, and this helps with debugging.

Acknowledgments

We would like to thank Marc Portoles, Balaji Venkatachalapathy, Bernhard Haindl, Luigi Iannone, and Padma Pillay-Esnault for their great suggestions and help regarding this document.

Many thanks to Alvaro Retana for the careful AD review.

Thanks to Chris M. Lonvick for the security directorate review, Al Morton for the OPS-DIR review, Roni Even for the Gen-ART review, Mike McBride for the rtg-dir review, Magnus Westerlund for the tsv directorate review, and Sheng Jiang for the int-dir review.

Thanks to John Scudder, Erik Kline, Lars Eggert, Warren Kumari, Martin Duke, Murray Kucherawy, Éric Vyncke, Robert Wilton, Zaheduzzaman Sarker, and Roman Danyliw for the IESG review.

This work was partly funded by the ANR LISP-Lab project #ANR-13-INFR-009 (https://anr.fr/Projet-ANR-13-INFR-0009).

Contributors

Dino Farinacci
lispers.net
San Jose, CA
USA

Email: farinacci@gmail.com

Johnson Leong

Email: johnsonleong@gmail.com

Fabio Maino
Cisco
San Jose, CA
USA

Email: fmaino@cisco.com

Christian Jacquenet
Orange
Rennes
France

Email: christian.jacquenet@orange.com

Stefano Secci
Cnam
France

Email: stefano.secci@cnam.fr

Authors' Addresses

Alberto Rodriguez-Natal
Cisco
Barcelona
Spain

Email: natal@cisco.com

Vina Ermagan
Google
United States of America

Email: ermagan@gmail.com

Albert Cabellos
UPC/BarcelonaTech
Barcelona
Spain

Email: acabello@ac.upc.edu

Sharon Barkai
Nexar

Email: sharon.barkai@getnexar.com

Mohamed Boucadair
Orange
Rennes
France

Email: mohamed.boucadair@orange.com