

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: September 10, 2016

M. Boucadair
C. Jacquenet
Orange
March 9, 2016

LISP Mapping Service Discovery at Large
draft-boucadair-lisp-idr-ms-discovery-01

Abstract

Locator/ID Separation Protocol (LISP) 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 (automatic discovery of Map-Resolvers and Map-Servers).

This document specifies a BGP Extended Communities attribute that can be used to dynamically discover LISP Mapping Systems of different domains.

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 September 10, 2016.

Internet-Draft

Mapping Service Discovery

March 2016

Copyright Notice

Copyright (c) 2016 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.	Rationale	5
3.	LISP Mapping System Target BGP Extended Community	5
4.	Security Considerations	6
5.	IANA Considerations	6
6.	Acknowledgments	6
7.	References	7
7.1.	Normative references	7
7.2.	Informative references	7
	Authors' Addresses	7

[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 (automatic discovery of Map-Resolvers and Map-Servers).

Within this document, a Mapping System provides the LISP mapping service [[RFC6833](#)]. Map-Resolvers, Map-Servers, and other components may be part of a Mapping System such as authorization, subscription profiles, etc. These components are considered as black boxes; only the external behavior of the Mapping System is in scope.

Distinct LISP mapping systems may emerge if LISP users or network operators who solicit or manage the Mapping System want to avoid some region-centric systems, for example, or if they want to position themselves as a core provider of the Mapping System. The lack of

clear policies of the management and operation of the LISP Mapping Systems may encourage such practices.

This document assumes a hierarchy in the Mapping System organisation for business, governance, control, and regulatory purposes in particular. In such contexts, the document assumes that a Mapping System may maintain a portion of or a global mapping table.

Because of its experimental nature of LISP and the various platforms LISP operation relies upon (like the platforms used by the LISP mapping systems) should encourage innovation by testing new services that may take advantage of LISP in inter-domain deployment scenarios without requiring the participation of all LISP-enabled domains. Such approach is also meant to avoid any risk of freezing LISP developments.

Because the design and operation of a consistent Mapping System is critical for the adoption of LISP at large scale, this document advocates for means to dynamically discover other Mapping Systems that are open to cooperate in inter-domain LISP deployment scenarios, typically .

Deploying LISP for inter-domain use cases may raise the following issues:

Issue#1: A LISP domain may need to discover available Mapping Systems so that it can rely upon them to extend the reachability scope.

Issue#2: Various Mapping Systems can be deployed over the Internet. These Mapping Systems need to interconnect to extend the reachability scope and avoid pressure on PTR (Proxy Tunnel Router) devices. Also, various Mapping Systems encourage the enforcement of policies that aim at optimizing LISP forwarding: for example, policies that consist in avoiding the solicitation of specific domains/ASes.

Issue#3: Distinct flavors of Mapping Systems may be deployed.

These mappings may not rely on the same database mapping system (e.g., NERD, ALT, CONS, etc.). As such, a clear interface to ease interconnection between these realms is needed. Standard solutions to discover Mapping Systems capabilities are likely to ease the interconnection of Mapping Systems.

Issue#4: Security concerns may arise during the discovery of the available Mapping Systems: for example, a given Mapping System may deny access from another domain, or available Mapping Systems need to make sure that they are entitled to exchange information with

one another or that an xTR of a given LISP network is entitled to solicit a mapping system of another LISP network, etc.

An efficient and scalable deployment of LISP within an inter-domain context for traffic engineering purposes, in particular, relies heavily on the availability of an inter-domain mapping system that spans several domains. From this perspective, the success of a global LISP adoption and deployment will mainly depend on how LISP-enabled domains will graft to existing mapping systems that can guarantee a global reachability scope. To minimize the risk of a fragmented Mapping System that would jeopardize the overall efficiency of an inter-domain LISP routing system, there is a need to encourage and facilitate the coordination of participating Mapping Systems.

This document relies on extended BGP communities [[RFC4360](#)] to advertise that a given domain supports the LISP Mapping Service. A contact IPv4 address and/or IPv6 address are also included in the attribute so that remote LISP Mapping Systems or LISP domains may initiate negotiation cycles for the sake of LISP Mapping System Interconnection or subscription to the Mapping Service offered by that Mapping System.

[Section 3](#) specifies a solution for the discovery of LISP Mapping Systems that are deployed in distinct administrative domains. This BGP-based solution assumes that domains that support a LISP Mapping Service will use the BGP Extended Communities attribute to inform other domains about the support of the service. EIDs that can be serviced with LISP will be tagged accordingly. Note that an EID can

be serviced by multiple Mapping Systems. Remote LISP Mapping Systems will rely upon that BGP-based advertising capability to discover the existence and the status of other Mapping Systems. Once a Mapping System is discovered, a local Mapping System can establish an interconnection agreement with that remote Mapping System. The contact IP address provided as part of the BGP Extended Communities attribute will be used to contact a remote Mapping System to request for further LISP-related capabilities, possibly negotiate an interconnection agreement and, consequently, extend the scope of the networks that can be serviced using LISP. Also, leaf LISP-aware networks can rely upon the information carried in the BGP Extended Communities attribute to discover Mapping Systems that may be solicited to invoke their mapping service. Subscription cycles may then be considered.

[2.](#) Rationale

This document focuses on the discovery of LISP Mapping Systems that are deployed in distinct administrative domains.

The rationale is as follows:

1. **Announce:** Domains that support a LISP Mapping Service will use the BGP Extended Communities attribute ([Section 3](#)) to inform other domains about the support of the service. EIDs that can be serviced with LISP can be tagged accordingly. Note that an EID can be serviced by multiple Mapping Systems.
2. **Discover:** Remote LISP Mapping Systems will rely upon that BGP-based advertising capability ([Section 3](#)) to discover the existence of other Mapping Systems.
3. **Negotiate/Interconnect/Invoke:** The contact IP address provided as part of the BGP Extended Communities attribute ([Section 3](#)) will be used to contact a remote Mapping System to request for further LISP-related capabilities, possibly negotiate an interconnection agreement and, consequently, extend the scope of the networks

that can be serviced using LISP.

4. Negotiate/Subscribe/Invoke: Also, leaf LISP-aware networks can rely upon the information carried in the BGP Extended Communities attribute to discover Mapping Systems that may be solicited to invoke their mapping service. Subscription cycles may then be considered.

Only the first two steps are in scope of this document; the remaining steps can be achieved by other means such as [\[I-D.boucadair-connectivity-provisioning-protocol\]](#).

3. LISP Mapping System Target BGP Extended Community

The LISP Mapping System Target Community identifies one or more Mapping System contact points that can receive mapping system interconnect and/or subscription requests. These contact points are identified with IPv4 and/or IPv6 addresses.

The LISP Mapping System Target Community is of an extended type. As such, the behavior specified in [Section 6 of \[RFC4360\]](#) applies to the LISP Mapping System Target Community.

The presence of this community is an explicit indication that associated networks can be managed by a LISP Mapping System that is reachable at the addresses carried in the attribute.

This document reuses the Transitive IPv4-Address-Specific Extended Community [\[RFC4360\]](#) and Transitive IPv6-Address-Specific Extended Community [\[RFC5701\]](#) for the purpose of this document. Dedicated subtypes are to be allocated (see [Section 5](#)).

The Global Administrator field MUST be set to an IP address of the Mapping System. This address MUST be configured on the originating BGP speaker.

The "Local Administrator" field of the LISP Mapping System Target Community is used to encode an identifier of the Mapping System. Considerations about the assignment of globally unique identifiers to LISP Mapping Systems are out of scope. A configurable parameter may be supported by BGP implementations to provide the value carried in the "Local Administrator" field. If no identifier is configured on

the originating BGP speaker, the "Local Administrator" field MUST be set to 0.

[4.](#) Security Considerations

This document does not introduce any additional security issues other than those discussed in [[RFC4360](#)] and [[RFC5701](#)].

[5.](#) IANA Considerations

According to [[RFC7153](#)], this document requests the assignment of a sub-type in the "0x00-0xbf" range from the Transitive IPv4-Address-Specific Extended Community Sub-Types registry available at <http://www.iana.org/assignments/bgp-extended-communities/bgp-extended-communities.xml#trans-ipv4>:

Type Value	Name	Reference
TBA	LISP Mapping System Target	[This-Document]

Also, this document requests the assignment of a sub-type from the Transitive IPv6-Address-Specific Extended Community Types registry available at <http://www.iana.org/assignments/bgp-extended-communities/bgp-extended-communities.xml#trans-ipv6>:

Type Value	Name	Reference
TBA	LISP Mapping System Target	[This-Document]

[6.](#) Acknowledgments

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

[7.](#) References

[7.1.](#) 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>>.

- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", [RFC 4360](#), DOI 10.17487/RFC4360, February 2006, <<http://www.rfc-editor.org/info/rfc4360>>.
- [RFC5701] Rekhter, Y., "IPv6 Address Specific BGP Extended Community Attribute", [RFC 5701](#), DOI 10.17487/RFC5701, November 2009, <<http://www.rfc-editor.org/info/rfc5701>>.
- [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>>.
- [RFC7153] Rosen, E. and Y. Rekhter, "IANA Registries for BGP Extended Communities", [RFC 7153](#), DOI 10.17487/RFC7153, March 2014, <<http://www.rfc-editor.org/info/rfc7153>>.

[7.2.](#) Informative references

- [I-D.boucadair-connectivity-provisioning-protocol]
Boucadair, M., Jacquenet, C., Zhang, D., and P. Georgatsos, "Connectivity Provisioning Negotiation Protocol (CPNP)", [draft-boucadair-connectivity-provisioning-protocol-10](#) (work in progress), September 2015.
- [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>>.

Authors' Addresses

Mohamed Boucadair
Orange
Rennes 35000
France

EMail: mohamed.boucadair@orange.com

Orange
Rennes 35000
France

E-Mail: christian.jacquet@orange.com