

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 2, 2016

M. Boucadair
C. Jacquenet
Orange
May 31, 2016

LISP Mapping Service Functions Discovery (LMSFD) using OSPF
draft-boucadair-lisp-function-discovery-02

Abstract

This document specifies extensions to the Open Shortest Path First (OSPF) protocol for the discovery of Locator/ID Separation Protocol (LISP) Mapping Service functions, especially the Map-Resolver (MR) and Map-Server (MS) LISP components.

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 December 2, 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.	Overview	4
3.	Mapping Service Function Discovery (MSFD) TLV	5
3.1.	MSF-TYPE sub-TLV	6
3.2.	MSF-LOCATOR sub-TLV	6
3.3.	MSF-DESCRIPTION sub-TLV	7
3.4.	MSF-EPOCH sub-TLV	7
3.5.	MSF-UNAVAILABILITY-TIMER sub-TLV	8
3.6.	MSF-REBOOT-TIMER sub-TLV	8
3.7.	MSF-DIAGNOSIS sub-TLV	9
3.8.	MS-STATUS sub-TLV	9
3.9.	MSF-STATUS sub-TLV	9
4.	Mapping Service Reachability Information	10
5.	OSPF Operation	10
6.	Security Considerations	11
7.	IANA Considerations	11
8.	Acknowledgments	11
9.	References	11
9.1.	Normative References	11
9.2.	Informative References	12
	Authors' Addresses	13

[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 (MR) and Map-Server (MS) entities that provide such mapping services is meant to facilitate global LISP operation (automatic discovery of Map-Resolvers and Map-Servers).

The dynamically-acquired information will not only be used by xTR routers but also by management platforms (e.g., a service controller, a network manager, etc.) to forward traffic over the LISP network or to get an up-to-date view of the global LISP network topology, including the location of the resolvers and servers. For example, ETRs will register in all instances that are reachable in a given domain.

The ability to dynamically discover LISP mapping component information and update such information as appropriate is also useful to ease state synchronization among the various Mapping Service Functions that can be solicited in the LISP network, especially whenever a new MS joins the LISP Mapping System. This specification allows the following:

1. Ease the introduction of new MS servers: Additional MS instances may be added to a Mapping Service domain. New MSes need to build an updated mapping database to avoid service disruption. Owing to the mechanism defined in this document,
 - * Peer MSes can be discovered by a new MS, thereby triggering a state synchronisation procedure. How state synchronisation is achieved is out of scope of this document.
 - * xTRs can immediately send registration messages to the new MS.

[RFC7215] indicates that, "for some LISP sites, there is a need for mechanisms to dynamically obtain the address of the Map-Server in the ETR of the AS".

2. Minimize service disruption when multiple MS/MRs are available: this specification allows to disseminate information that will drive the selection process undertaken by an xTR to select an MS/MR and solicit it. For example, MRs with empty databases will be avoided; "ready-to-serve" MRs will be solicited instead. Map-Register requests can thus be sent to multiple MSes whenever needed. When a Mapping Service function loses its state, an explicit message can be generated accordingly so that xTRs (and also management platforms) can trigger appropriate actions.

This document specifies a means to dynamically discover Map-Resolver and Map-Server instances of a LISP network within a domain. Also, it introduces some features to enhance the serviceability of LISP (e.g., detect that a MS lost its state so that a registration procedure is triggered immediately, inform an xTR that a MS/MR instance is about to be unavailable for some reason, indicate the synchronisation status of the mapping database, etc.).

The document exclusively focuses on the dynamic information discovery; how this information is actually used by an xTR to infer its MS/MR selection procedures is out of the scope.

The reader should be familiar with the terms defined in [\[RFC6833\]](#).

2. Overview

This document defines extensions to OSPFv2 [[RFC2328](#)] and OSPFv3 [[RFC5340](#)] so that routers of the OSPF routing domain (a single area or the entire routing domain) can advertise a Mapping Service Function within the domain, along with some other useful information. To do so, a new TLV (named the LISP Mapping Service Function Discovery TLV (LMSFD TLV)) is used to announce LISP MR and MS information. This TLV is carried by the OSPF Router Information LSA [[RFC7770](#)].

The location of each Mapping Service Function is then flooded into an OSPF area or the whole OSPF routing domain (in the case the LSA is AS-scoped). The xTR routers deployed within the OSPF domain must listen to the flooding messages sent by active Mapping Service Function instances. Such messages are referred to "Mapping Service Discovery" messages in this document.

The information to be announced by means of the LMSFD TLV carried in the Router Information LSA during the LISP Mapping Service Function Discovery procedure includes (but is not necessarily limited to):

- o Mapping Service Function type: Indicates whether the MSF acts as Map-Resolver (MR), Map-Server (MS), or both.
- o Mapping Service Function Service locator(s): Includes one or several IPv4 addresses, one or several IPv6 addresses or a combination thereof. This information lists the set of locators that unambiguously indicate where the Mapping Service Function can be reached. The locator information must be included in the Mapping Service Function Discovery messages.
- o Mapping Service Function unavailability timer: Indicates the time when the Mapping Service Function will be unavailable. This parameter can be used for planned maintenance operations, for instance. This parameter does not provide any indication about when the Mapping Service Function instance will be available again. This information is optional and may therefore be included in the Mapping Service Function Discovery messages.
- o Mapping Service Function reboot timer: Operational teams often proceed with a reboot of the devices deployed in the network, within the context of major software upgrade campaigns, for example. This timer is used to indicate that a Mapping Service Function will be unavailable during the reboot of the device that supports the function. Unlike the previous timer, this timer is used to indicate that the Mapping Service Function will be available immediately after the reboot of the device that supports

this function is completed. This information is optional and may therefore be included in the Mapping Service Function Discovery messages.

- o Mapping Service Function Diagnosis: Indicates whether this Mapping Service Function instance supports a diagnostic mechanism. This information is optional and may therefore be included in the Mapping Service Function Discovery messages.
- o Mapping Service Status: Provides information about the status of the mapping database. In particular, it indicates whether the database is empty, synchronized with other MS servers located in the same OSPF domain, etc. This information is optional and may therefore be included in the Mapping Service Function Discovery messages.
- o Mapping Service Function Status: Indicates the status of the Mapping Service Function Instance (enabled, disabled). This information is optional and may therefore be included in the LMSFD messages.
- o Additional capabilities such as the support of mapping bulk retrieval [[I-D.boucadair-lisp-bulk](#)] or notifications [[I-D.boucadair-lisp-subscribe](#)] may be advertised.

3. Mapping Service Function Discovery (MSFD) TLV

The format of the OSPF Mapping Service Function Discovery TLV (LMSFD TLV, Figure 1) and its sub-TLVs use the same TLV format as in the Traffic Engineering Extensions to OSPF [[RFC3630](#)].




Figure 1

The description of the fields is as follows:

- o Type: To be assigned by IANA (see [Section 7](#)).
- o Length: Variable (octets).


- o sub-TLVs: Includes the list of sub-TLVs. The following sub-TLVs are defined in this document:

Sub-TLV type	Length	Name
1	1	MSF-TYPE sub-TLV
2	variable	MSF-LOCATOR sub-TLV
3	variable	MSF-DESCRIPTION sub-TLV
4	4	MSF-EPOCH sub-TLV
5	4	MSF-UNAVAILABILITY-TIMER sub-TLV
6	4	MSF-REBOOT-TIMER sub-TLV
7	1	MSF-DIAGNOSIS sub-TLV
8	4	MS-STATUS sub-TLV
9	4	MSF-STATUS sub-TLV

The MSF-TYPE and MSF-LOCATOR sub-TLVs MUST always be present within the LMSFD TLV. Additional optional sub-TLVs MAY be included.

3.1. MSF-TYPE sub-TLV

The format of MSF-TYPE sub-TLV is shown in Figure 2.



The current type values are defined:

- 0: Map-Server
- 1: Map-Resolver
- 2: Both

Figure 2: MSF-TYPE sub-TLV

3.2. MSF-LOCATOR sub-TLV

The format of MSF-LOCATOR sub-TLV is shown in Figure 3.




Figure 3: MSF-LOCATOR sub-TLV

The MSF-LOCATOR sub-TLV MAY appear twice, especially when the SF can be reached via either an IPv4 or an IPv6 address or both. It MAY also appear more than once for the same address family if the Service Function is assigned several addresses of the same family.

3.3. MSF-DESCRIPTION sub-TLV

The format of MSF-DESCRIPTION sub-TLV is shown in Figure 4.




Figure 4: MSF-DESCRIPTION sub-TLV

When present, the MSF-DESCRIPTION sub-TLV MUST carry UTF-8 encoded [RFC3629] description text. The description text SHOULD NOT be null terminated.

3.4. MSF-EPOCH sub-TLV

The format of MSF-EPOCH sub-TLV is shown in Figure 5.




Figure 5: MSF-EPOCH sub-TLV

When a Mapping Service Function loses its state (e.g., power failure, bug, reset by an administrator, etc.), it may use this sub-TLV with a value set to 0. This value is then incremented by one.

If the value field of the MSF-EPOCH sub-TLV is set to 0, it indicates that the Mapping Service Function instance has been reset or lost its state. This information is particularly important for ETRs so that they can send their registration request immediately.

Receivers may maintain a record of transmitted values to detect anomalies in the Mapping Service Function.

3.5. MSF-UNAVAILABILITY-TIMER sub-TLV

The format of MSF-UNAVAILABILITY-TIMER sub-TLV is shown in Figure 6.




Figure 6: MSF-UNAVAILABILITY-TIMER sub-TLV

The MSF-UNAVAILABILITY-TIMER sub-TLV indicates, in seconds, when the Mapping Service Function instance will be unavailable.

If the value field of the MSF-UNAVAILABILITY-TIMER sub-TLV is set to 0, it indicates the Mapping Service Function instance will be unavailable immediately.

3.6. MSF-REBOOT-TIMER sub-TLV

The format of MSF-REBOOT-TIMER sub-TLV is shown in Figure 7.




Figure 7: MSF-REBOOT-TIMER sub-TLV

The MSF-REBOOT-TIMER sub-TLV indicates, in seconds, when the Mapping Service Function instance will start to reboot. The function will be operational right after the reboot is completed.

3.7. MSF-DIAGNOSIS sub-TLV

The format of MSF-DIAGNOSIS sub-TLV is shown in Figure 8.





Figure 8: MSF-DIAGNOSIS sub-TLV

The presence of this sub-TLV indicates that the Mapping Service Function supports diagnostic functions.

3.8. MS-STATUS sub-TLV

The format of MS-STATUS sub-TLV is shown in Figure 9



The current Status values are defined:


- 0: Reset
- 1: Partial
- 2: Synchronized

Figure 9: MS-STATUS sub-TLV

The presence of this sub-TLV indicates the status of the Mapping Service database. This is important for influencing the selection process of Map-Resolvers, in particular.

3.9. MSF-STATUS sub-TLV

The format of MSF-STATUS sub-TLV is shown in Figure 10



The current Status values are defined:

- 0: Enabled
- 1: Disabled

Figure 10: MSF-STATUS sub-TLV

The presence of this sub-TLV indicates the status of Mapping Service Function.

The presence of this sub-TLV is particularly useful to indicate that a given instance is disabled.

4. Mapping Service Reachability Information

This document assumes that Mapping Service Reachability information can be injected into OSPF by a router that embeds a Mapping Service Function instance, or which has been instructed (by means of specific configuration tasks, for example) to advertise such information on behalf of a third party Mapping Service Function.

The mechanism defined in this document may be used to advertise and learn Mapping Service Functions that are available in the same administrative domain than xTRs. Also, it can be used to dynamically advertise related reachability information that is learned using other means when the Mapping Service Functions and xTRs do not belong to the same administrative entity.

Some of the information carried in the LMSFD TLV may be automatically set by an OSPF speaker while other may be explicitly configured.

5. OSPF Operation

The LMSFD TLV is advertised within OSPFv2 or OSPFv3 Router Information LSAs [[RFC7770](#)].

A change in the operational status of a Mapping Service Function instance (e.g., enabled, disabled) MUST trigger the generation of a Router Information LSA that carries the LMSFD TLV with the updated information.

The flooding scope is defined by the Opaque LSA type for OSPFv2 [[RFC5250](#)], and by the S1/S2 bits for OSPFv3 [[RFC5340](#)].

6. Security Considerations

The extensions defined in this document do not introduce any additional security threats than those already documented in the current OSPF protocol specifications.

OSPF does not support any encryption mechanism for protecting the integrity of Mapping Service Function discovery information. Means such as [[RFC2154](#)] may be enabled.

7. IANA Considerations

This document requests IANA to assign a Router Information TLV type for the LISP Mapping Service Discovery Function (LMSFD) TLV.

8. Acknowledgments

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

Thanks to Liu Xiang for the comments.

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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), DOI 10.17487/RFC2328, April 1998, <<http://www.rfc-editor.org/info/rfc2328>>.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, [RFC 3629](#), DOI 10.17487/RFC3629, November 2003, <<http://www.rfc-editor.org/info/rfc3629>>.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), DOI 10.17487/RFC3630, September 2003, <<http://www.rfc-editor.org/info/rfc3630>>.

- [RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", [RFC 5250](#), DOI 10.17487/RFC5250, July 2008, <<http://www.rfc-editor.org/info/rfc5250>>.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", [RFC 5340](#), DOI 10.17487/RFC5340, July 2008, <<http://www.rfc-editor.org/info/rfc5340>>.
- [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>>.
- [RFC7770] Lindem, A., Ed., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", [RFC 7770](#), DOI 10.17487/RFC7770, February 2016, <<http://www.rfc-editor.org/info/rfc7770>>.

9.2. Informative References

- [I-D.boucadair-lisp-bulk] Boucadair, M. and C. Jacquenet, "LISP Mapping Bulk Retrieval", [draft-boucadair-lisp-bulk-01](#) (work in progress), March 2016.
- [I-D.boucadair-lisp-subscribe] Boucadair, M. and C. Jacquenet, "Improving Mapping Services in LISP Networks", [draft-boucadair-lisp-subscribe-02](#) (work in progress), March 2016.
- [RFC2154] Murphy, S., Badger, M., and B. Wellington, "OSPF with Digital Signatures", [RFC 2154](#), DOI 10.17487/RFC2154, June 1997, <<http://www.rfc-editor.org/info/rfc2154>>.
- [RFC7215] Jakab, L., Cabellos-Aparicio, A., Coras, F., Domingo-Pascual, J., and D. Lewis, "Locator/Identifier Separation Protocol (LISP) Network Element Deployment Considerations", [RFC 7215](#), DOI 10.17487/RFC7215, April 2014, <<http://www.rfc-editor.org/info/rfc7215>>.

Authors' Addresses

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Christian Jacquenet
Orange
Rennes 35000
France

Email: christian.jacquenet@orange.com