Working Group
Internet Draft
Intended status: Standards Track
Expires: December 2016

Dave Allan, Uma Chunduri
Ericsson

June 2016

OSPF extensions for Computed Multicast applied to MPLS based Segment Routing

draft-allan-ospf-spring-multicast-00

Abstract

This document describes the OSPFv2 extensions required to support multicast for MPLS based Segment Routing. In this approach OSPF speakers compute their role in multicast tree construction based on the information in the OSPF routing information base.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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".

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on December 2016.

Copyright and License Notice

Table of Contents

1. Introduction

[ALLAN-1] describes a solution for multicast for Segment Routing with MPLS data plane in which source specific multicast distribution trees (MDTs) are computed from information distributed via an IGP. Using this approach, both any-source multicast (ASM) and engineered p2mp trees can be supported.
This memo describes three additional TLVs for OSPF to support the segment routing multicast approach described in [ALLAN-1]

1.1. Authors

David Allan, Uma Chunduri

1.2. 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 [RFC2119].

2. Conventions used in this document

2.1. Terminology

Candidate replication point - is a node will potentially need to install state as determined at some intermediate step in MDT computation. It will either resolve to having no role or a role as a replication point once multicast has converged.

Candidate role - refers to any potential combination of roles on a given MDT as determined at some intermediate step in MDT computation. For example, a node with a candidate role may be a leaf and may be a candidate replication point.

Downstream - refers to the direction along the shortest path to one or more leaves for a given multicast distribution tree

Multicast convergence - is when all computation and state installation to ensure the FIB reflects the multicast information in the IGP is complete.

Pinned path - Is a unique shortest path extending from a leaf upstream towards the root for a given MDT. Therefore is a component of an MDT that must be there. It will not necessarily extend from the leaf all the way to the root during intermediate computation steps. A pinned path can result from pruning operations.

Role - refers specifically to a node that is either a root, a leaf or a replication node for a given multicast distribution tree

Unicast convergence- is when all computation and state installation to ensure the FIB reflects the unicast information in the IGP is complete.

Upstream - refers to the direction along the shortest path to the root of a given multicast distribution tree

3. Overview

The role of the IGP in the multicast architecture described in [ALLAN-1] is to synchronize knowledge of the topology, knowledge of

unicast SIDs, knowledge of multicast SIDs, multicast group membership and agreement on the algorithm to use for computation of multicast distribution trees(MDTs) across the set of IGP speakers. This document specifies the TLVs necessary for OSPF version 2 to be the IGP in said multicast architecture.

4. New TLVs

4.1. Additional Router Information TLVs

The following three TLV are in the Router Information Opaque LSA specified in RFC 7770 for SRM compute capabilities and other TLVs in OSPFv2 Extended Prefix Opaque LSA as specified in [RFC7684].

4.1.1. SRM Compute Capability TLV

The presence of this TLV in a router information LSA [RFC7770] indicates both that the originating node supports computed spring multicast, but also indicates the algorithm that is configured to be used. All nodes supporting computed multicast are required to agree on the algorithm for correct operation of the network.

The format of the sub-TLV is:

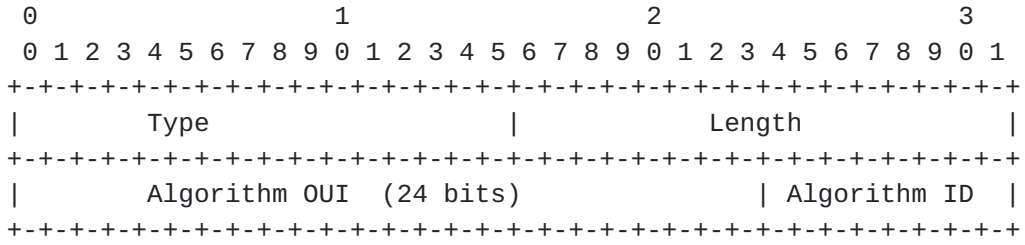

Where:

The upper 24 bits contains an organizationally unique identifier (as per RFC 7042) and the lower 8 bits contains an algorithm identifier.

The default algorithm supported as per [ALLAN-1] is identified by Algorithm OUI = 0x008037 (Ericsson), Algorithm ID = 0x01 (default).

4.1.2. SRM SID Multicast Group Binding Sub-TLV

The SID Multicast Group Binding TLV communicates the binding between the SID specific to the MDT for the multicast group originating at the advertising node and the multicast group address, as well as transmit and receive interest for the advertising node. Note that if the sub-TLV does not have the T bit set, the SID TLV is not included

in the message. This sub-TLV is part of OSPFv2 Extended Prefix TLV as a sub-TLV from OSPFv2 Extended Prefix Opaque LSA as specified in [RFC7684].

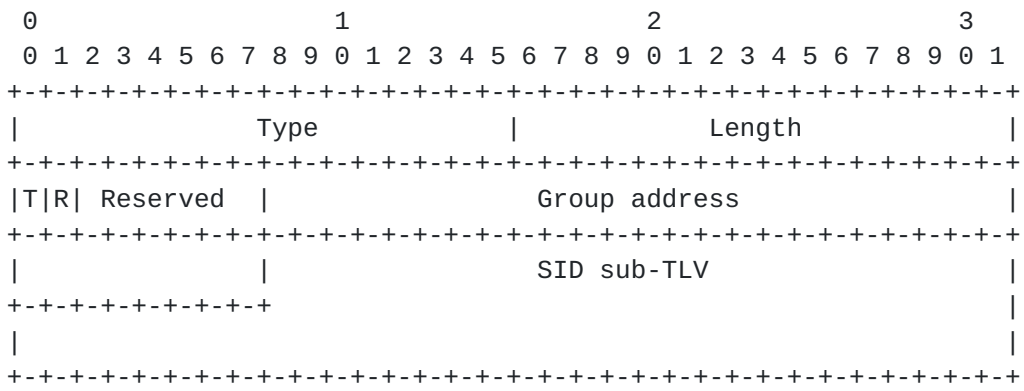

Where:

Type = TBD (IANA assignment from OSPFv2 Extended Prefix TLV Sub-TLVs Registry)

Length = octet size of the T/R, reserved fields, Group ID and the SID sub-tlv information

T-bit indicates that this node is a source for the multicast group specified in the sub-tlv.

R-bit indicates that this node is a receiver for the multicast group specified in the sub-tlv.

Group ID: 4 octet IPv4 multicast address

SID sub-TLV = the segment ID to use for this source/group MDT. The format is as per the SID/label sub-TLV defined in section 2.1 of [SPRING-OSPF].

4.1.3. SRM Pinned Tree Descriptor sub-TLV

The pinned tree descriptor defines all nodes that have a role in a multicast distribution tree, and their relationship to the individual multicast segments that define the tree. This sub-TLV is part of OSPFv2 Extended Prefix TLV as a sub-TLV from OSPFv2 Extended Prefix Opaque LSA as specified in [RFC7684]. The encoding is an unstructured list, where if the tree description exceeds the size, it may simply use more than one sub-TLV.

The encoding of a role descriptor is in the form of upstream_SID/unicast_SID/downstream _SID, where each of the SIDs is encoded as a sub-TLV as per [SPRING-OSPF].

The root of the MDT (and originator of the TLV) will have a NULL upstream SID, transit waypoints will have both a defined upstream and downstream multicast segment SID, and a leaf will have a NULL downstream SID. The unicast SID corresponds to the node for which the entry defines its role.

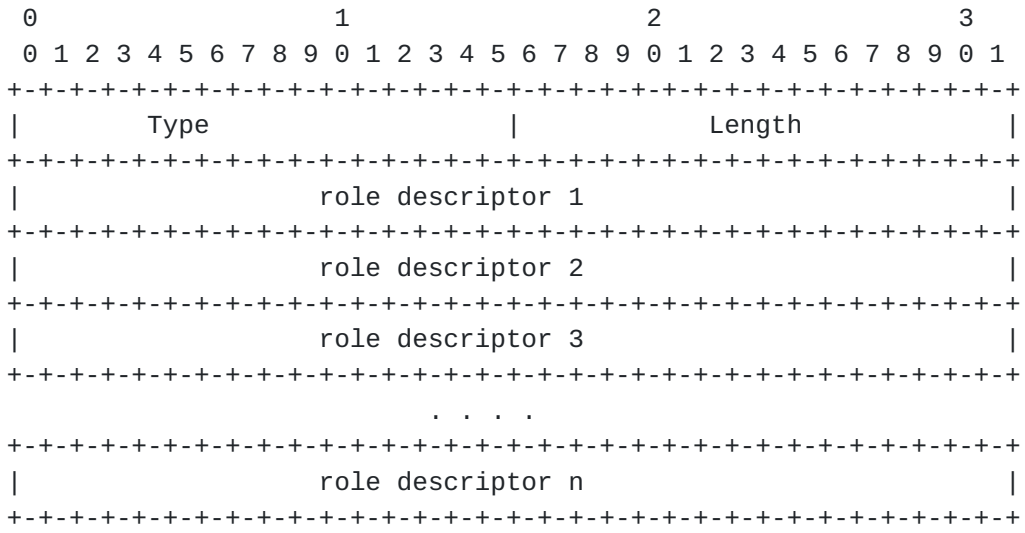

Where:

The encoding of a role descriptor is in the form:

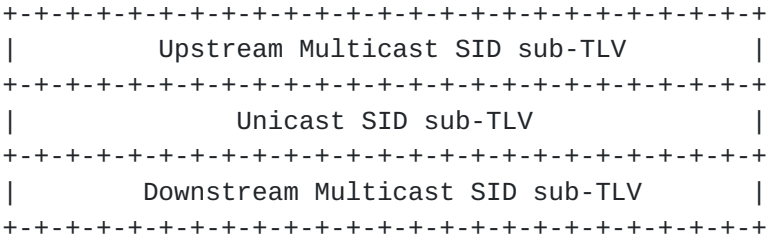

5. Acknowledgements

6. Security Considerations

For a future version of this document.

7. IANA Considerations

This memo requires the assignment of 1 value from the OSPF Router Information (RI) TLV registry:

- SRM Compute Capability TLV

This memo also requires the assignment of the following 2 Sub-TLV values from the OSPFv2 Extended Prefix TLV Sub-TLV registry:

- SRM SID Multicast Group Binding sub-TLV

- SRM Pinned Tree Descriptor sub-TLV

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC7042] Eastlake, D. et. al.,"IANA Considerations and IETF Protocol and Documentation Usage for IEEE 802 Parameters", IETF RFC 7042, October 2013

[RFC7684] Psenak et.al., "OSPFv2 Prefix/Link Attribute Advertisement", IETF RFC 7684, November 2015

[RFC7770] Lindem et.al., "Extensions to OSPF for Advertising Optional Router Capabilities", IETF RFC 7770, February 2016

[SPRING-OSPF] Psenak et.al., "OSPF Extensions for Segment Routing", IETF work in progress, draft-ietf-ospf-segment-routing-extensions-08, April 2016

8.2. Informative References

[ALLAN-1] Allan et.al., "A Framework for Computed Multicast applied to MPLS based Segment Routing", draft-allan-spring-mpls-mcast-framework-01, July 2016

9. Authors' Addresses

Dave Allan (editor)
Ericsson
300 Holger Way
San Jose, CA 95134
USA
Email: david.i.allan@ericsson.com

Uma Chunduri
Ericsson
300 Holger Way
San Jose, CA 95134
USA
Email: uma.chunduri@ericsson.com