Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 3, 2016

Z. Li
M. Chen
Huawei
G. Mirsky
Ericsson

August 31, 2015

Routing Extensions for Discovery of Role-based MPLS Label Switching Router (MPLS LSR) Traffic Engineering (TE) Mesh Membership draft-li-teas-role-based-automesh-01

Abstract

A Traffic Engineering (TE) mesh-group is defined as a group of Label Switching Routers (LSRs) that are connected by a full mesh of TE LSPs. Routing protocol (OSPF and IS-IS) extensions facilitate discovery of Multiprotocol Label Switching (MPLS) LSR TE mesh membership and automate the creation of a full mesh of TE Label Switched Paths (LSPs).

This document introduces a role-based TE mesh-group that applies to the scenarios where full mesh TE LSPs are not necessary and TE LSPs setup depends on the roles of LSRs in a TE mesh-group. Interior Gateway Protocol (IGP) routing extensions for automatic discovery of role-based TE mesh membership are defined accordingly.

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 March 3, 2016.

Table of Contents

1. Introduction

A Traffic Engineering (TE) mesh-group [RFC4972] is defined as a group of Label Switching Routers (LSRs) that are connected by a full mesh of TE LSPs. [RFC4972] specifies Intermediate System-to-Intermediate System (IS-IS) and Open Shortest Path First (OSPF) extensions to provide an automatic discovery of the set of LSR members of a TE mesh-group in order to automate the creation of such mesh of TE LSPs.

This is called "auto-mesh TE" or "auto-mesh". Therefore auto-mesh TE largely simplifies the configuration required for the deployment of full mesh TE Label Switched Paths (LSPs).

1.1. Motivation and Scope

In some deployment scenarios, auto configuration of TE LSPs among specific nodes is useful but a full mesh may not be needed . An example where a full mesh is not required, but a partial mesh would significantly reduce operational overhead, is deployment and operation of TE LSPs in a mobile backhaul network ([I-D.li-mpls-seamless-mpls-mbb]). In this scenario, auto-mesh of TE LSPs between the Cell Site Gateways (CSGs), and Radio Network Controller (RNC) Site Gateways (RSGs) would be useful, and TE LSPs between CSGs, and TE LSPs between RSGs, would not be necessary. The amount of CSGs in mobile backhaul networks are very large. If using the auto-mesh mechanism as defined in [RFC4972], a full mesh TE LSPs will be established among all the CSGs and RSGs, thus resulting in large amount of unnecessary TE LSPs being established from CSGs-to-CSGs, and RSGs-to-RSGs. Potentially causing scaling issues and wasting network resources.

Therefore, there are clear requirements to optimize the auto-mesh function for setup of TE LSPs, and allowing specific group membership rather than a full TE mesh between all LSRs.

This document introduces a "role-based auto-mesh TE group" or "role-based auto-mesh" where the setup of the TE LSPs are dependent on the roles of the LSRs within a TE mesh-group. The method and procedure for signaling the TE LSPs is out the scope of this document.

1.2. Terminology

RSVP-TE - Resource Reservation Protocol-Traffic Engineering

P2MP - Point-to-Multipoint

IS-IS - Intermediate System-to-Intermediate System

OSPF - Open Shortest Path First

CSG - Cell Site Gateway

RNC - Radio Network Controller

MBH - Mobile Backhaul

MPLS - Multiprotocol Label Switching

LSP - Label Switched Path

TE LSP - Traffic Engineered LSP

2. Role-based TE Mesh Group

A role-based TE mesh-group is a special TE mesh-group where TE LSPs will not be established among all member LSRs. LSRs in a role-based TE mesh-group will have different roles. The TE LSPs setup depends on the roles of the LSRs in a TE mesh-group.

This document introduces the Hub-Spoke LSR TE mesh-group, where an LSR can be a Hub, a Spoke or both Hub and Spoke (Hub-Spoke) LSR in a mesh-group. The rules for Hub-Spoke TE mesh-group are as follows:

- TE LSPs SHOULD only be setup between Spoke and Hub LSRs.

- TE LSPs MUST NOT be setup between/among Spoke LSRs.

- TE LSPs MUST NOT be setup between/among Hub LSRs.

A Hub-Spoke LSR has two roles, for a mesh-group, it allows that a Hub-Spoke LSR can connect to any other Hub, Spoke and Hub-Spoke LSRs. This gives a choice to control whether an LSR can connect to any other LSRs through TE LSPs. When an LSR wants to setup TE LSPs with any other LSRs, configure it to Hub-Spoke LSR, otherwise, keep it as pure Hub or Spoke LSR role.

3. IGP Role-based TE Mesh-group Extensions

3.1. OSPF TE-ROLE-MESH-GROUP TLV Format

The OSPF TE-ROLE-MESH-GROUP TLV is used to advertise that an LSR joins/leaves a role-based TE mesh-group and the role of the LSR in the TE mesh-group. The OSPF TE-ROLE-MESH-GROUP TLV format for IPv4 (Figure 2) and IPv6 (Figure 3) is as follows:

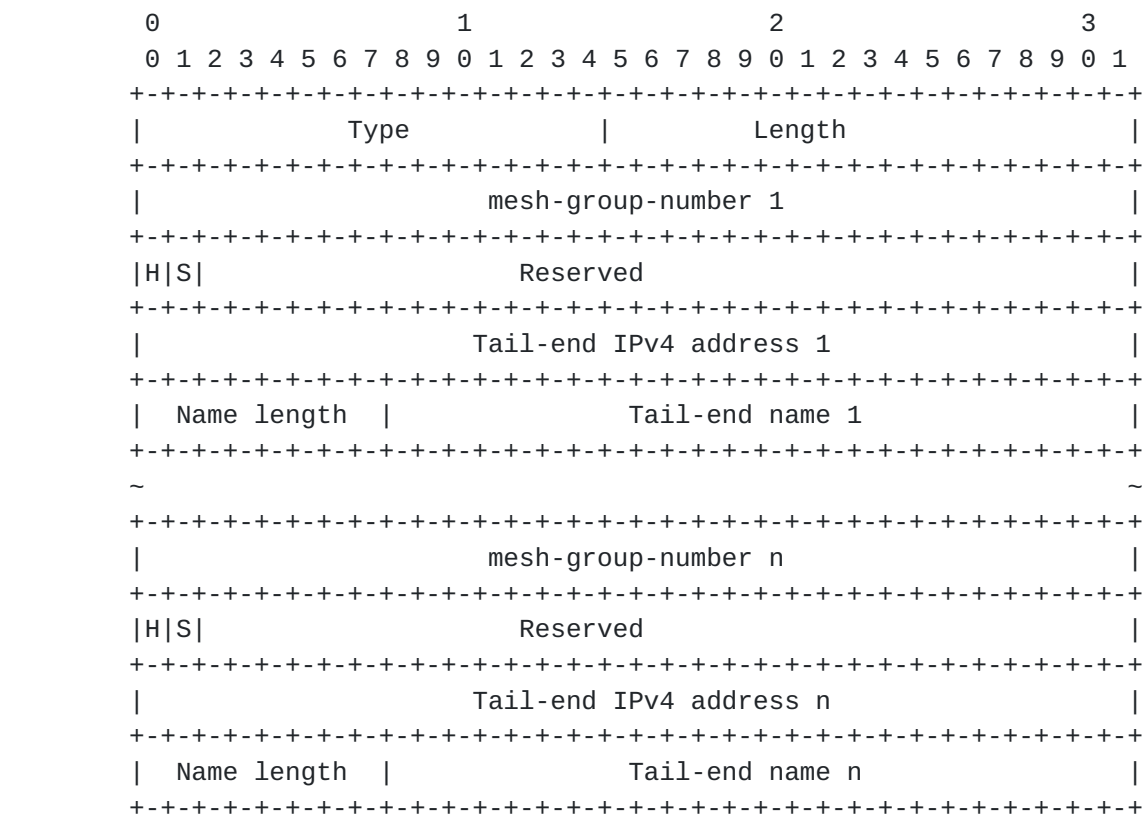

Figure 2 - OSPF TE-ROLE-MESH-GROUP TLV format (IPv4)

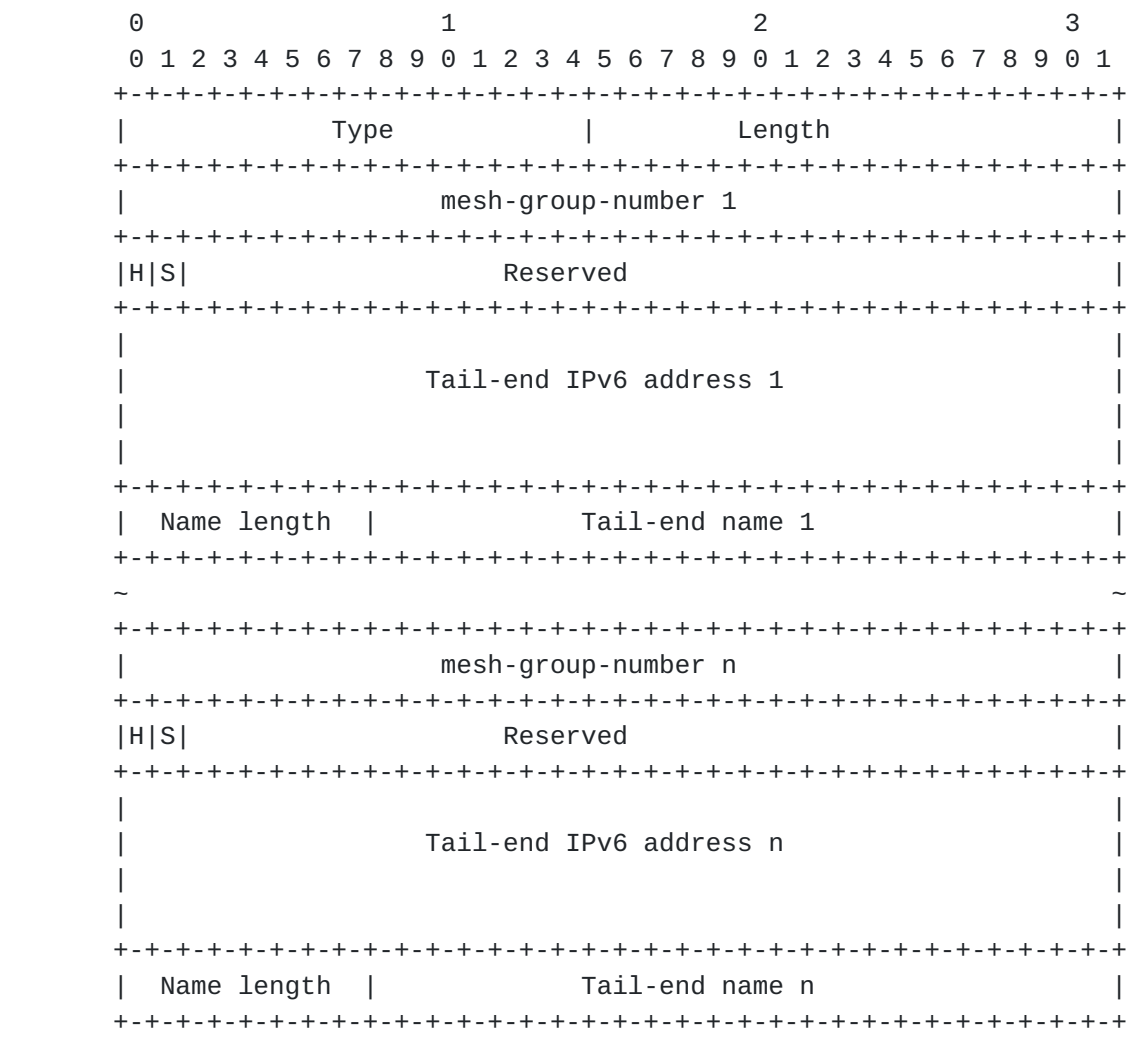

Figure 3 - OSPF TE-ROLE-MESH-GROUP TLV format (IPv6)

The Type of OSPF TE-ROLE-MESH-GROUP TLV for IPv4 is TBD1, the value of the Length is variable.

The Type of OSPF TE-ROLE-MESH-GROUP TLV for IPv6 is TBD2, the value of the Length is variable.

The OSPF TE-ROLE-MESH-GROUP TLV may contain one or more role-based mesh-group entries. Each entry corresponds to a role-based TE mesh-group. The definition of the mesh-group-number, the Tail-end address, the Name length and the Tail-end name in each role-based mesh group entry is the same as that of OSPF TE-MESH-GROUP TLV defined in [RFC4972].

In addition, for each mesh group entry, an four-octet flag field is introduced and four flags are defined in this document. Other bits

are reserved for future use and MUST be set to zero when sent, and MUST be ignored when received.

The H (Hub) bit, when set, it indicates the LSR is a Hub LSR.

The S (Spoke) bit, when set, it indicates the LSR is a Spoke LSR.

The H and S bit are dedicated for Hub-Spoke TE mesh-group and can be both set. When both bits set, it indicates that an LSR has both the Hub and Spoke role in the group. When neither H or S bit set, the element SHOULD be silently ignored.

3.2. IS-IS TE-ROLE-MESH-GROUP Sub-TLV Format

The IS-IS TE-ROLE-MESH-GROUP sub-TLV is used to advertise that an LSR joins/leaves a TE mesh-group and the role of the LSR in the TE mesh-group. The IS-IS TE-ROLE-MESH-GROUP sub-TLV format for IPv4 (Figure 4) and IPv6 (Figure 5) is as follows:

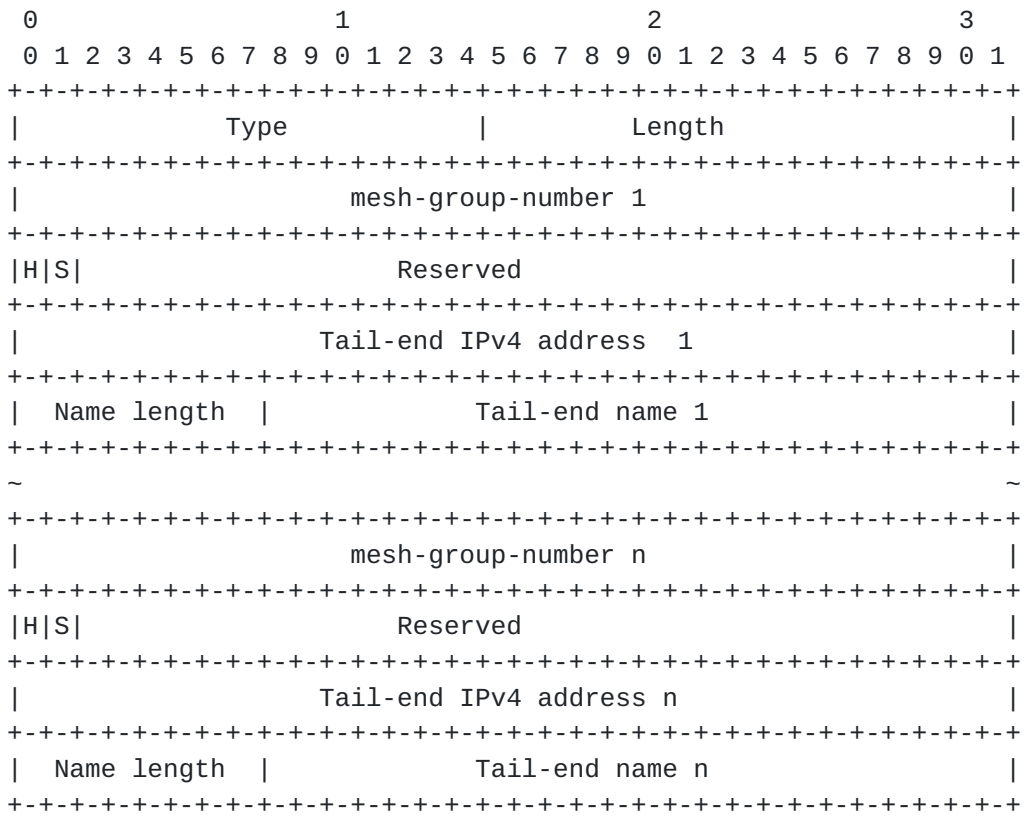

Figure 4 - IS-IS TE-ROLE-MESH-GROUP sub-TLV format (IPv4)

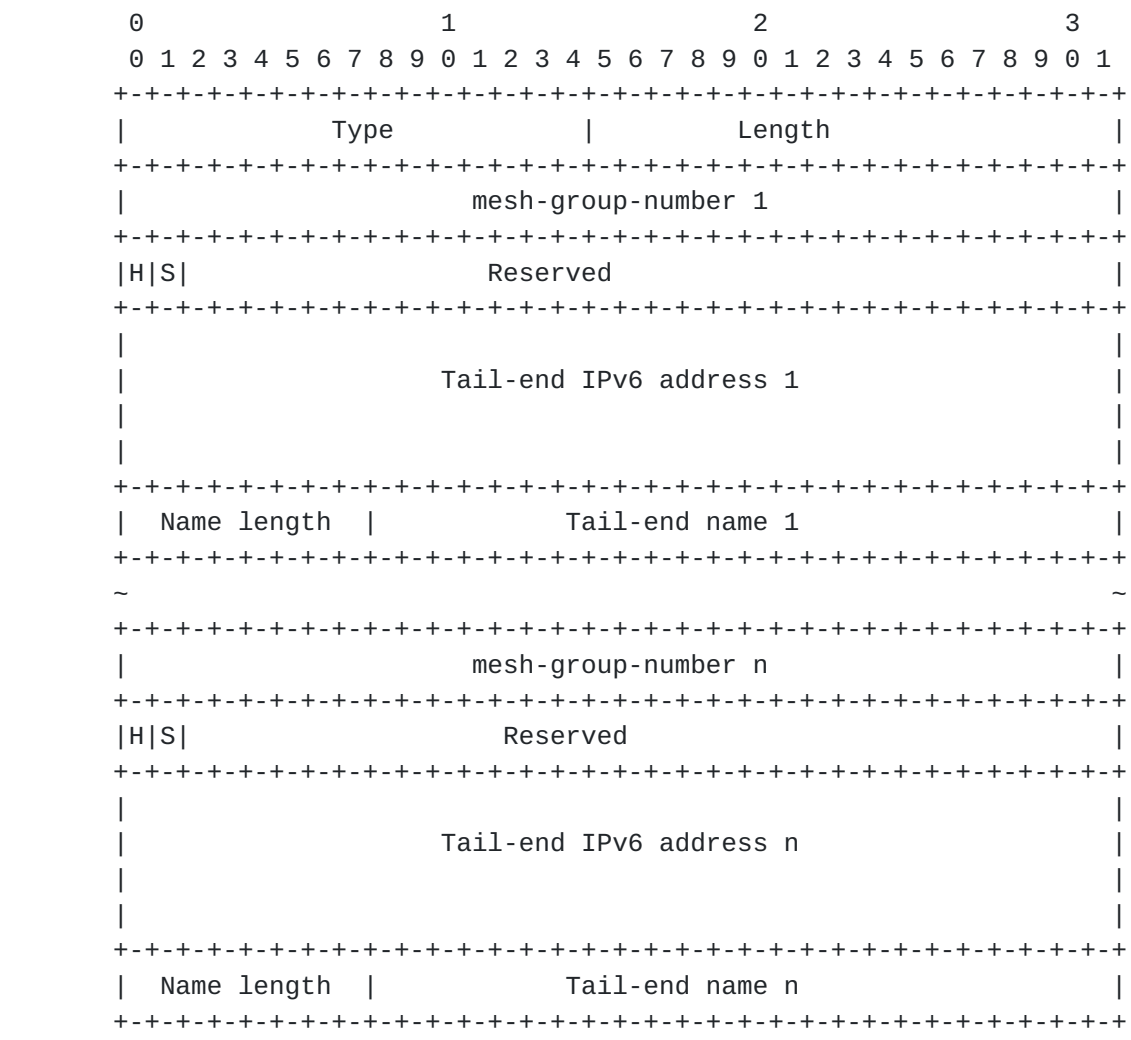

Figure 5 - IS-IS TE-ROLE-MESH-GROUP sub-TLV format (IPv6)

The Type of IS-IS TE-ROLE-MESH-GROUP sub-TLV for IPv4 is TBD3, the value of the Length is variable.

The Type of IS-IS TE-ROLE-MESH-GROUP sub-TLV for IPv6 is TBD4, the value of the Length is variable.

The IS-IS Role-based TE-ROLE-MESH-GROUP sub-TLV may contain one or more role-based mesh-group entries. Each entry corresponds to a role-based TE mesh-group. The definition of the fields, mesh-group-number, Tail-end address, Name length and Tail-end name in each role-based mesh group entry is the same as that of IS-IS TE-MESH-GROUP sub-TLV defined in [RFC4972].

The H and S bits are defined as in Section 3.1 of this document.

4. Elements of Procedure

4.1. OSPF

The TE-ROLE-MESH-GROUP TLV is advertised within an OSPF Router Information opaque LSA (opaque type of 4, opaque ID of 0) for OSPFv2 [RFC2328] and within a new LSA (Router Information LSA) for OSPFv3 [RFC5340]. The Router Information LSAs for OSPFv2 and OSPFv3 are defined in [RFC4970].

A router MUST originate a new OSPF router information LSA whenever the content of any of the advertised TLV changes or whenever required by the regular OSPF procedure (LSA update (every LSRefreshTime)). If an LSR desires to join or leave a particular role-based TE mesh group or an LSR desires to change its role in a mesh group, it MUST originate a new OSPF Router Information LSA comprising the updated TE-ROLE-MESH-GROUP TLV. In the case of a join, a new entry will be added to the TE-ROLE-MESH-GROUP TLV; if the LSR leaves a mesh-group, the corresponding entry will be removed from the TE-ROLE-MESH-GROUP TLV; if the LSR changes its role in the role-based mesh group, the corresponding entry will be updated in the TE-ROLE-MESH-GROUP TLV. Note that these operations can be performed in the context of a single LSA update. An implementation SHOULD be able to detect any change to a previously received TE-ROLE-MESH-GROUP TLV from a specific LSR.

As defined in [RFC5250] for OSPVv2 and in [RFC5340] for OSPFv3, the flooding scope of the Router Information LSA is determined by the LSA Opaque type for OSPFv2 and the values of the S1/S2 bits for OSPFv3.

If the flooding scope is area local, then the TE-ROLE-MESH-GROUP TLV MUST be carried within an OSPFv2 type 10 router information LSA or an OSPFV3 Router Information LSA with the S1 bit set and the S2 bit clear. If the flooding scope is the entire IGP domain, then the TE-ROLE-MESH-GROUP TLV MUST be carried within an OSPFv2 type 11 Router Information LSA or OSPFv3 Router Information LSA with the S1 bit clear and the S2 bit set.

When the router receives TE-ROLE-MESH-GROUP TLV, it SHOULD setup MPLS TE LSPs according rules which defined in the Section 3.

4.2. IS-IS

The TE-ROLE-MESH-GROUP sub-TLV is advertised within the IS-IS Router CAPABILITY TLV defined in [RFC4971].

An IS-IS router MUST originate a new IS-IS LSP whenever the content of any of the advertised sub-TLV changes or whenever required by

regular IS-IS procedure (LSP updates). If an LSR desires to join or leave a particular role-based TE mesh group or an LSR desires to change its role in a mesh group, it MUST originate a new LSP comprising the refreshed IS-IS Router capability TLV comprising the updated TE-ROLE-MESH-GROUP sub-TLV. In the case of a join, a new entry will be added to the TE-ROLE-MESH-GROUP sub-TLV; if the LSR leaves a mesh-group, the corresponding entry will be deleted from the TE-ROLE-MESH-GROUP sub-TLV; if the LSR changes its role in the role-based mesh group, the corresponding entry will be updated in the TE-ROLE-MESH-GROUP sub-TLV. Note that these operations can be performed in the context of a single update. An implementation SHOULD be able to detect any change to a previously received TE-ROLE-MESH-GROUP sub-TLV from a specific LSR.

If the flooding scope of a TE-ROLE-MESH-GROUP sub-TLV is limited to an IS-IS level/area, the sub-TLV MUST NOT be leaked across level/area and the S flag of the Router CAPABILITY TLV MUST be cleared. Conversely, if the flooding scope of a TE-ROLE-MESH-GROUP sub-TLV is the entire routing domain, the TLV MUST be leaked across IS-IS levels/areas, and the S flag of the Router CAPABILITY TLV MUST be set. In both cases, the flooding rules specified in [RFC4971] apply.

When the router receives TE-ROLE-MESH-GROUP sub-TLV, it SHOULD setup MPLS TE LSPs according rules which defined in the Section 3.

5. Backward Compatibility

For a role-based TE mesh-group, if there are some LSRs only supporting mechanisms defined [RFC4972], all the LSRs of the mesh-group MUST process as defined in [RFC4972]. The operators should avoid to add an LSR that does not support role-based auto-mesh TE to a role-based TE mesh-group.

6. IANA Considerations

6.1. OSPF

The registry for the Router Information LSA is defined in [RFC4970]. IANA is requested to assign two new TLV types from the Standards Action allocation range of the registry "OSPF Router Information (RI) TLVs".

Value	TLV	References
TBD1	TE-ROLE-MESH-GROUP TLV (IPv4)	this document
TBD2	TE-ROLE-MESH-GROUP TLV (IPv6)	this document

6.2. IS-IS

The registry for the Router Capability TLV is defined in [RFC4971]. IANA is requested to assign two new sub-TLV code-point for the TE-ROLE-MESH-GROUP sub-TLVs carried within the IS-IS Router Capability TLV.

Value	Sub-TLV	References
TBD3	TE-ROLE-MESH-GROUP sub-TLV (IPv4)	this document
TBD4	TE-ROLE-MESH-GROUP sub-TLV (IPv6)	this document

7. Security Considerations

The function described in this document does not create any new security issues for the OSPF and IS-IS protocols, the security considerations described in [RFC4972] apply here.

8. Acknowledgements

The authors would like to thank Loa Andersson for his valuable comments.

The authors would also like to thank the authors of [RFC4972] from where we have taken most of the elements procedures.

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>.

[RFC4970] Lindem, A., Ed., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", RFC 4970, DOI 10.17487/RFC4970, July 2007, <http://www.rfc-editor.org/info/rfc4970>.

[RFC4971] Vasseur, JP., Ed., Shen, N., Ed., and R. Aggarwal, Ed., "Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information", RFC 4971, DOI 10.17487/RFC4971, July 2007, <http://www.rfc-editor.org/info/rfc4971>.

[RFC4972] Vasseur, JP., Ed., Leroux, JL., Ed., Yasukawa, S., Previdi, S., Psenak, P., and P. Mabbey, "Routing Extensions for Discovery of Multiprotocol (MPLS) Label Switch Router (LSR) Traffic Engineering (TE) Mesh Membership", RFC 4972, DOI 10.17487/RFC4972, July 2007, <http://www.rfc-editor.org/info/rfc4972>.

[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>.

9.2. Informative References

[I-D.li-mpls-seamless-mpls-mbb]
Li, Z., Li, L., Morillo, M., and T. Yang, "Seamless MPLS for Mobile Backhaul", draft-li-mpls-seamless-mpls-mbb-01 (work in progress), February 2014.

Authors' Addresses

Zhenbin Li
Huawei

Email: lizhenbin@huawei.com

Mach(Guoyi) Chen
Huawei

Email: mach.chen@huawei.com

Greg Mirsky
Ericsson

Email: gregory.mirsky@ericsson.com