Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 19, 2014

Z. Li
M. Chen
Huawei
G. Mirsky
Ericsson
June 17, 2014

# Routing Extensions for Discovery of Role-based MPLS Label Switching Router (MPLS LSR) Traffic Engineering (TE) Mesh Membership
## draft-li-ccamp-role-based-automesh-02

Abstract

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

This document introduces a role-based TE mesh-group that applies to the scenarios where full mesh TE LSPs is 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 December 19, 2014.

Table of Contents



## 1. Introduction

A TE mesh-group [RFC4972] is defined as a group of 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". The auto-mesh TE significantly simplifies the configuration and deployment of TE LSPs.

In some scenarios, it may not be necessary to establish full mesh TE LSPs among all the LSRs of a TE mesh-group. An example of the use case of non-full mesh of TE LSPs in the mobile backhaul (MBH) networks is presented in ([I-D.li-mpls-seamless-mpls-mbb]). In MBH network TE LSPs are usually setup between the Cell Site Gateways(CSGs) and the Radio Network Controller (RNC) Site Gateways(RSGs). TE LSPs interconnecting CSGs and TE LSPs interconnecting RSGs are not necessary. In most deployments the number of CSGs is very large and there are much more CSGs than RSGs in an MBH domain. With the auto-mesh mechanism defined[RFC4972] full mesh of TE LSPs will be established interconnecting CSGs and RSGs. As result large number of unnecessary TE LSPs will be established interconnecting CSGs and interconnecting RSGs. This likely will not scale well with addition of more CSG devices, would stress control plane with unwarranted RSVP state.

Thus there are requirements to optimize the auto-mesh TE and to reduce the number of unnecessary TE LSPs. This document introduces a "role-based auto-mesh TE" or "role-based auto-mesh" where the setup of the TE LSPs is based on the role of the LSRs within a particular TE mesh-group. Therefore, besides the discovery of the membership of a TE mesh-group, it needs to discover the role of each node in the TE mesh-group.

Another scenario to which the role-based auto-mesh TE can apply is the Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Point-to-Multipoint (P2MP) TE LSP[RFC4875] scenario. For a RSVP-TE P2MP TE LSP, the root LSR has to know all the leaf LSRs before signalling the P2MP TE LSP. The automatic discovery mechanisms defined in this document can be used to discover the leaf LSRs for P2MP TE LSPs.

This document defines IGP routing extensions to automatically discover of the members and their roles of a TE mesh-group.

## 1.1. 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. In a role-based TE mesh-group LSRs will have different roles. TE LSPs setup depends on the roles of the LSRs in a TE mesh-group. This document introduces two types of role-based TE mesh group: Hub-Spoke and Root-Leaf.

For a Hub-Spoke TE mesh-group, an LSR can be a Hub, Spoke or both Hub and Spoke LSR in a 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.

For a Root-Leaf TE mesh-group, an LSR can be a Root, a Leaf or both a Root and Leaf LSR. Once the membership and roles are determined, the root LSRs can signal the P2MP TE LSPs toward all the Leaf LSRs. There may be multiple P2MP TE LSPs within a TE mesh-group.

## 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:

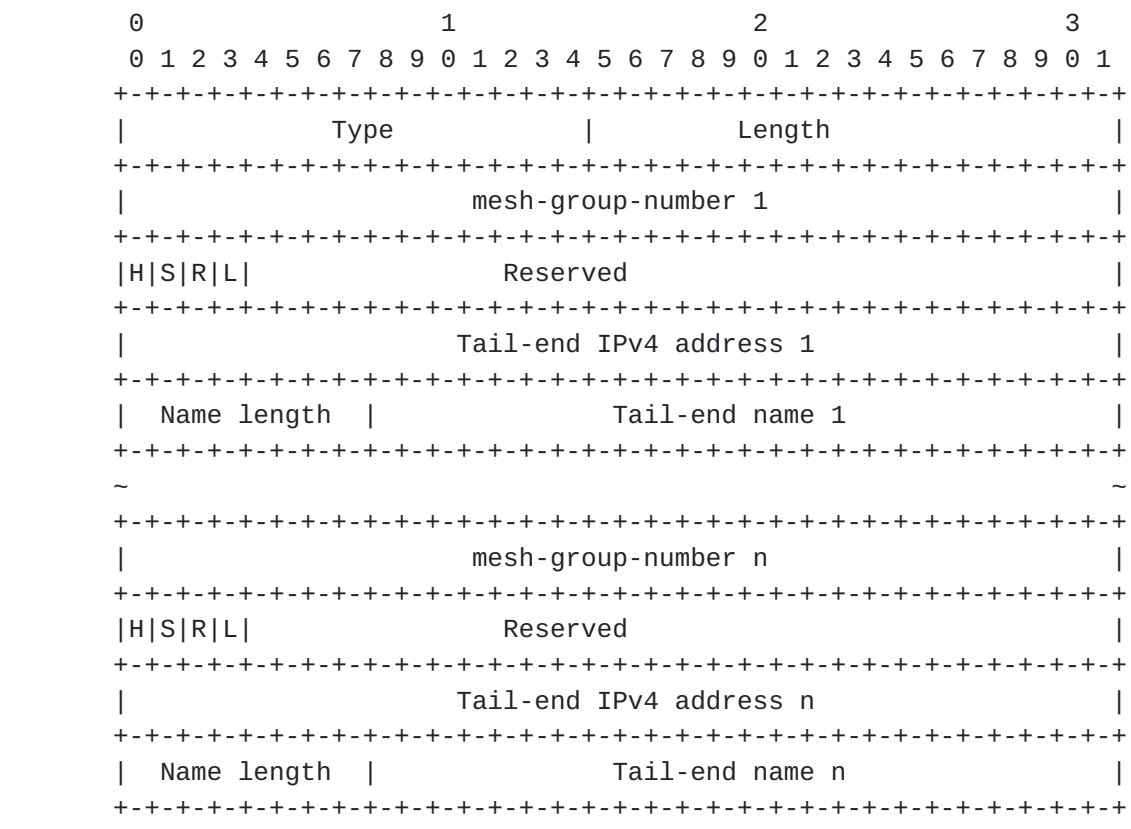


Figure 2 - OSPF TE-ROLE-MESH-GROUP TLV format (IPv4 Address)

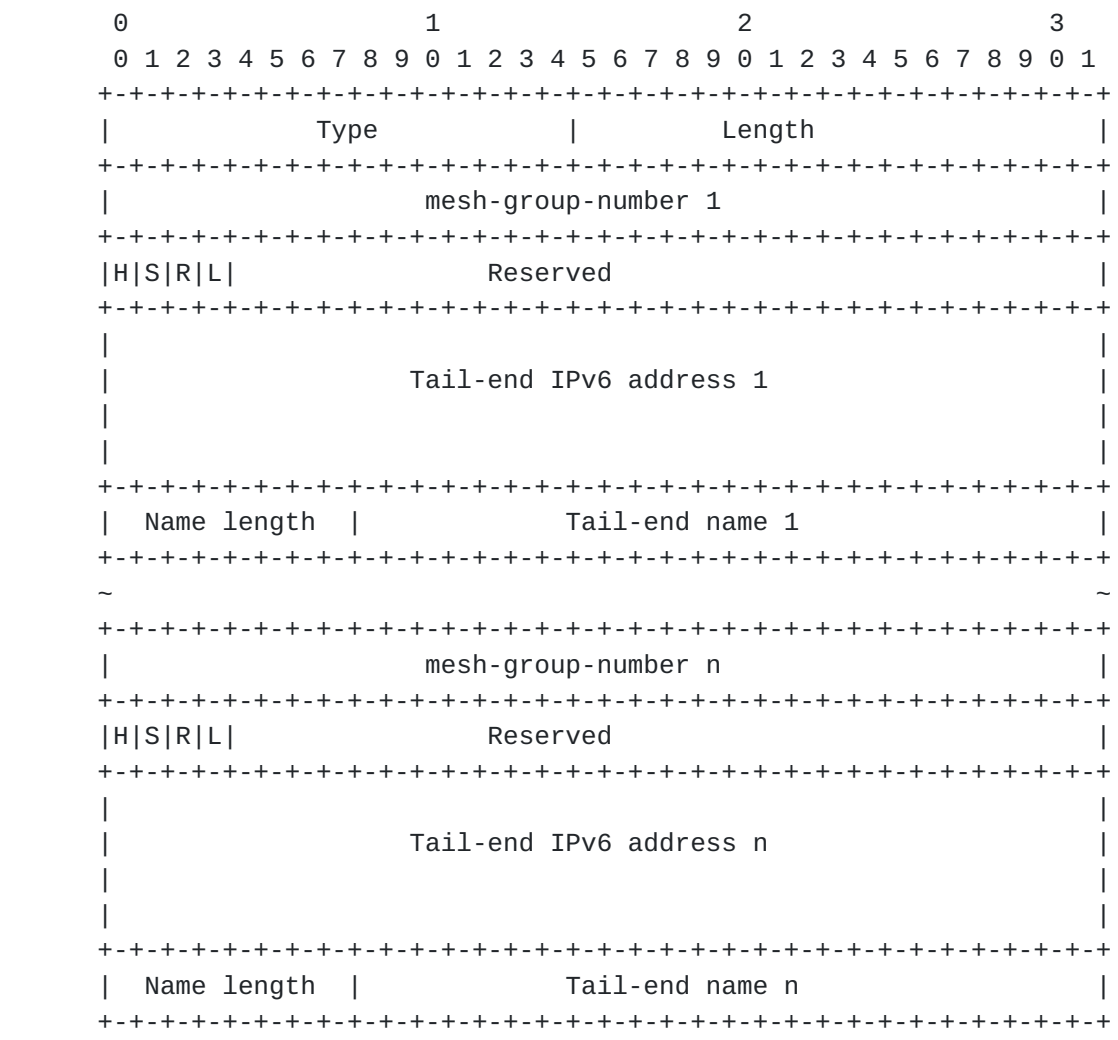


Figure 3 - OSPF TE-ROLE-MESH-GROUP TLV format (IPv6 Address)

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 R (Root) bit, when set, it indicates an LSR is a Root LSR.

The L (Leaf) bit, when set, it indicates an LSR is a Leaf.

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.

The R and Leaf bit can be both set, when both bits set, in indicates an LSR is a Root and Leaf LSR. The R bit and Leaf bit are only used for Root-Leaf TE mesh-group, for other TE mesh-groups, it MUST be set to zero and MUST be ignored when received.

## 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:

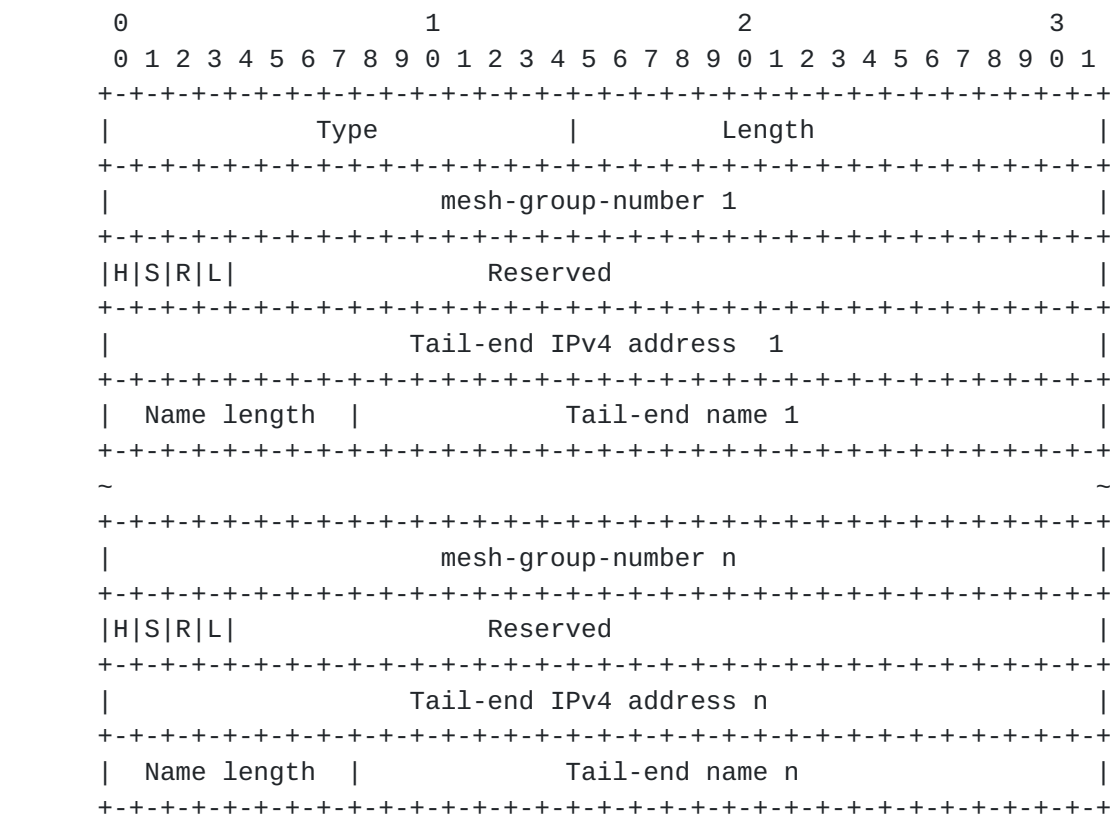


Figure 4 - IS-IS TE-ROLE-MESH-GROUP sub-TLV format (IPv4 Address)

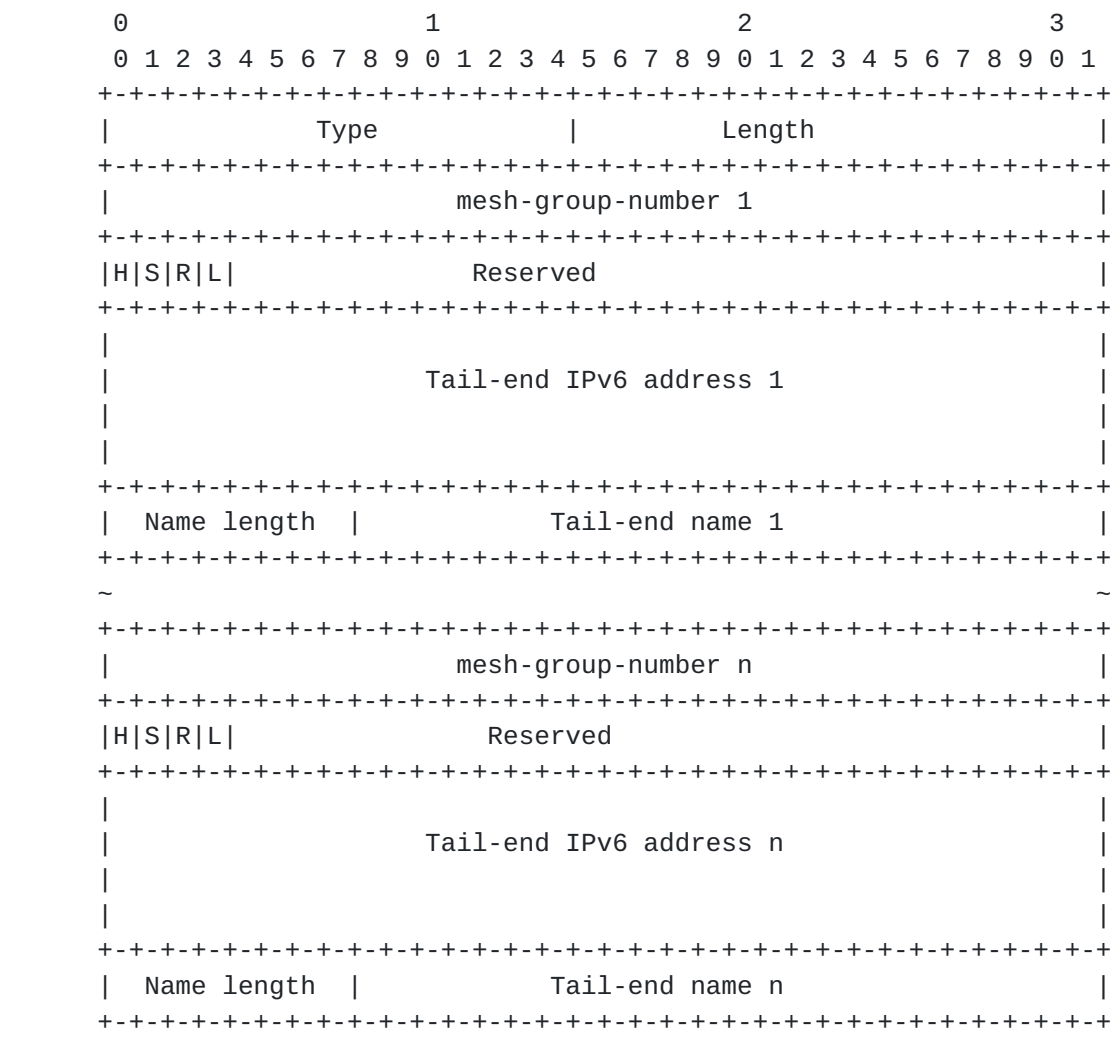


Figure 5 - IS-IS TE-ROLE-MESH-GROUP sub-TLV format (IPv6 Address)

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, S, R and L bits are defined as in Section 3.1 of this document.

## 4. Elements of Procedure

The OSPF TE-ROLE-MESH-GROUP TLV is carried within the OSPF Routing Information LSA, and the IS-IS TE-ROLE-MESH-GROUP sub-TLV is carried within the IS-IS Router capability TLV. As such, elements of procedure are inherited from those defined in [RFC4970] and [RFC4971] for OSPF and IS-IS respectively. Specifically, a router MUST originate a new LSA/LSP whenever the content of this information changes, or whenever required by regular routing procedure (e.g., updates).

The TE-ROLE-MESH-GROUP TLV is OPTIONAL and MUST NOT include more than one of each of the IPv4 instances or the IPv6 instance. If either the IPv4 or the IPv6 OSPF TE-ROLE-MESH-GROUP TLV occurs more than once within the OSPF Router Information LSA, only the first instance is processed, subsequent TLV(s) MUST be ignored. Similarly, if either the IPv4 or the IPv6 IS-IS TE-ROLE-MESH-GROUP sub-TLV occurs more than once within the IS-IS Router capability TLV, only the first instance is processed, subsequent TLV(s) MUST be ignored.

### 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.

For OSPFv2 Router Information opaque LSA:

- Link-local scope: type 9;

- Area-local scope: type 10;

- Routing-domain scope: type 11. In this case, the flooding scope is equivalent to the Type 5 LSA flooding scope.

For OSPFv3 Router Information LSA:

- Link-local scope: OSPFv3 Router Information LSA with the S1 and S2 bits cleared;

- Area-local scope: OSPFv3 Router Information LSA with the S1 bit set and the S2 bit cleared;

- Routing-domain scope: OSPFv3 Router Information LSA with S1 bit cleared and the S2 bit set.

A router may generate multiple OSPF Router Information LSAs with different flooding scopes.

The Role-based TE-MESH-GROUP TLV may be advertised within an Area-local or Routing-domain scope Router Information LSA, depending on the MPLS TE mesh group profile:

- If the MPLS TE mesh-group is contained within a single area (all the LSRs of the mesh-group are contained within a single area), the TE-ROLE-MESH-GROUP TLV MUST be generated within an Area-local Router Information LSA.

- If the MPLS TE mesh-group spans multiple OSPF areas, the TE Role-based mesh- group TLV MUST be generated within a Routing-domain scope router information LSA.

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.

As specified in [RFC4971], a router may generate multiple IS-IS Router CAPABILITY TLVs within an IS-IS LSP with different flooding scopes.

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 assigned a new OSPF TLV code-point for the TE-ROLE-MESH-GROUP TLVs carried within the Router Information LSA.

| 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 assigned a new IS-IS 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, March 1997.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.

[RFC4970] Lindem, A., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", RFC 4970, July 2007.

[RFC4971] Vasseur, JP., Shen, N., and R. Aggarwal, "Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information", RFC 4971, July 2007.

[RFC4972] Vasseur, JP., Leroux, JL., 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, July 2007.

[RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", RFC 5250, July 2008.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.

## 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.

[RFC4875] Aggarwal, R., Papadimitriou, D., and S. Yasukawa, "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, May 2007.

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