

Network Working Group
Request for Comments: 4972
Category: Standards Track

JP. Vasseur, Ed.
Cisco Systems, Inc
JL. Leroux, Ed.
France Telecom
S. Yasukawa
NTT
S. Previdi
P. Psenak
Cisco Systems, Inc
P. Mabbey
Comcast
July 2007

Routing Extensions for Discovery of Multiprotocol (MPLS)
Label Switch Router (LSR) Traffic Engineering (TE) Mesh Membership

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

The setup of a full mesh of Multi-Protocol Label Switching (MPLS) Traffic Engineering (TE) Label Switched Paths (LSP) among a set of Label Switch Routers (LSR) is a common deployment scenario of MPLS Traffic Engineering either for bandwidth optimization, bandwidth guarantees or fast rerouting with MPLS Fast Reroute. Such deployment may require the configuration of a potentially large number of TE LSPs (on the order of the square of the number of LSRs). This document specifies Interior Gateway Protocol (IGP) routing extensions for Intermediate System-to-Intermediate System (IS-IS) and Open Shortest Path First (OSPF) so as to provide an automatic discovery of the set of LSRs members of a mesh in order to automate the creation of such mesh of TE LSPs.

Table of Contents

1. Introduction	2
2. Definitions	3
2.1. Conventions Used in This Document	4
3. Description of a TE Mesh-Group	4
4. TE-MESH-GROUP TLV Formats	4
4.1. OSPF TE-MESH-GROUP TLV Format	4
4.2. IS-IS TE-MESH-GROUP Sub-TLV Format	7
5. Elements of Procedure	9
5.1. OSPF	9
5.2. IS-IS	10
6. Backward Compatibility	11
7. IANA Considerations	11
7.1. OSPF	11
7.2. IS-IS	11
8. Security Considerations	12
9. Acknowledgements	12
10. References	12
10.1. Normative References	12
10.2. Informative References	13

1. Introduction

There are two well-known approaches in deploying MPLS Traffic Engineering:

(1) The so-called "strategic" approach that consists of setting up a full mesh of TE LSPs between a set of LSRs.

(2) The so-called "tactical" approach, where a set of TE LSPs are provisioned on well-identified "hot spots" in order to alleviate a congestion resulting, for instance, from an unexpected traffic growth in some parts of the network.

The setup of a full mesh of TE LSPs among a set of LSRs is a common deployment scenario of MPLS Traffic Engineering either for bandwidth optimization, bandwidth guarantees, or fast rerouting with MPLS Fast Reroute. Setting up a full mesh of TE LSPs between N LSRs requires the configuration of a potentially large number of TE LSPs ($O(N^2)$). Furthermore, the addition of any new LSR in the mesh requires the configuration of N additional TE LSPs on the new LSR and one new TE LSP on every LSR of the existing mesh destined to this new LSR, which gives a total of $2*N$ TE LSPs to be configured. Such an operation is not only time consuming but also risky (prone to misconfiguration) for Service Providers. Hence, an automatic mechanism for setting up TE LSPs meshes is desirable and requires the ability to automatically discover the set of LSRs that belong to the mesh. This document

specifies routing extensions so as to automatically discover the members of a mesh, also referred to as a "TE mesh-group". Note that the mechanism(s) needed for the dynamic creation of TE LSPs is implementation specific and outside the scope of this document.

Routing extensions have been defined in [RFC4970] and [RFC4971] in order to advertise router capabilities. This document specifies IGP (OSPF and IS-IS) TE Mesh Group (Type Length Value) TLVs allowing for the automatic discovery of a TE mesh-group members, to be carried in the OSPF Router Information (Link State Advertisement) LSA [RFC4970] and IS-IS Router Capability TLV [RFC4971]. The routing extensions specified in this document provide the ability to signal multiple TE mesh groups. An LSR may belong to more than one TE mesh-group(s).

There are relatively tight real-time constraints on the operation of IGPs (such as OSPF and IS-IS). For this reason, some care needs to be applied when proposing to carry additional information in an IGP. The information described in this document is both relatively small in total volume (compared with other information already carried in IGPs), and also relatively stable (i.e., changes are based on configuration changes, but not on dynamic events within the network, or on dynamic triggers, such as the leaking of information from other routing protocols or routing protocol instances).

2. Definitions

Terminology used in this document

IGP: Interior Gateway Protocol

IGP Area: OSPF area or IS-IS level

IS-IS: Intermediate System-to-Intermediate System (IS-IS)

LSR: Label Switch Router

OSPF: Open Shortest Path First

OSPF LSA: OSPF Link State Advertisement

TE LSP: Traffic Engineering Label Switched Path

TE LSP head-end: head/source of the TE LSP

TE LSP tail-end: tail/destination of the TE LSP.

TLV: Type Length Value

2.1. Conventions Used in This Document

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].

3. Description of a TE Mesh-Group

A TE mesh-group is defined as a group of LSRs that are connected by a full mesh of TE LSPs. Routing extensions are specified in this document, allowing for dynamic discovery of the TE mesh-group members. Procedures are also specified for a member to join and leave a TE mesh-group. For each TE mesh-group membership announced by an LSR, the following information is advertised:

- A mesh-group number identifying the TE mesh-group that the LSR belongs to,
- A tail-end address (used as the TE LSP Tail-end address by other LSRs belonging to the same mesh-group),
- A tail-end name: a display string that is allocated to the tail-end used to ease the TE-LSP naming.

4. TE-MESH-GROUP TLV Formats

4.1. OSPF TE-MESH-GROUP TLV Format

The TE-MESH-GROUP TLV is used to advertise the desire of an LSR to join/leave a given TE mesh-group. No sub-TLV is currently defined for the TE-MESH-GROUP TLV.

The OSPF TE-MESH-GROUP TLV (advertised in an OSPF router information LSA defined in [RFC4970]) has the following format:

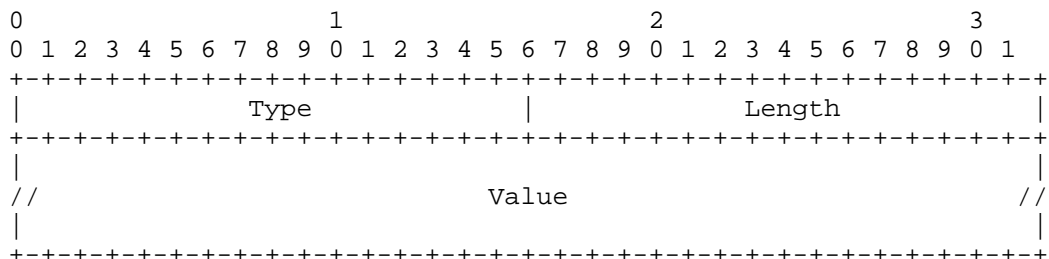


Figure 1 - OSPF TE-MESH-GROUP TLV format

Where

Type: identifies the TLV type

Length: the length of the value field in octets

The format of the OSPF TE-MESH-GROUP TLV is the same as the TLV format used by the Traffic Engineering Extensions to OSPF (see[RFC3630]). The TLV is padded to a four-octet alignment; padding is not included in the length field (so a three-octet value would have a length of three, but the total size of the TLV would be eight octets). Nested TLVs are also 32-bit aligned. Unrecognized types are ignored. All types between 32768 and 65535 are reserved for vendor-specific extensions. All other undefined type codes are reserved for future assignment by IANA.

The OSPF TE-MESH-GROUP TLV format for IPv4 (Figure 2) and IPv6 (Figure 3) is as follows:

TYPE: 3
LENGTH: Variable

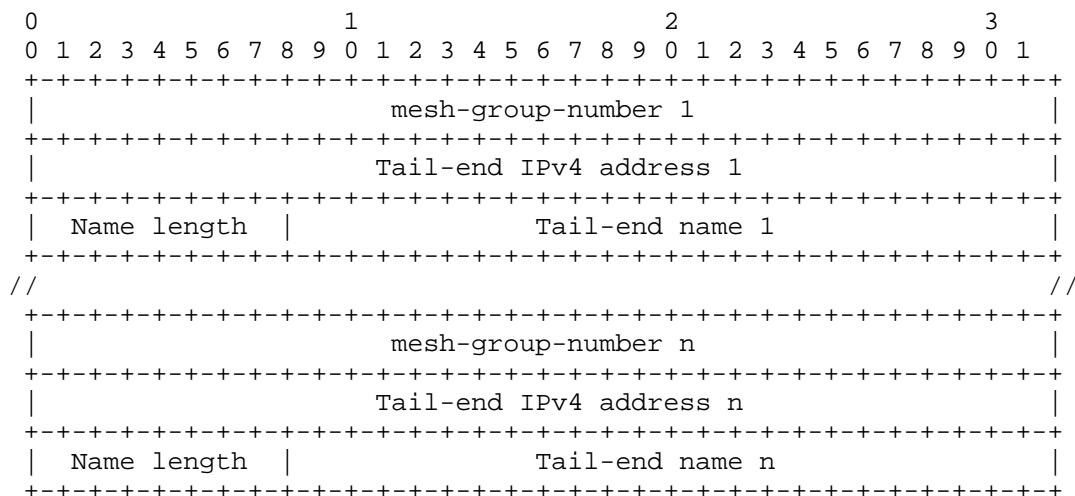


Figure 2 - OSPF TE-MESH-GROUP TLV format (IPv4 Address)

TYPE: 4
 LENGTH: Variable

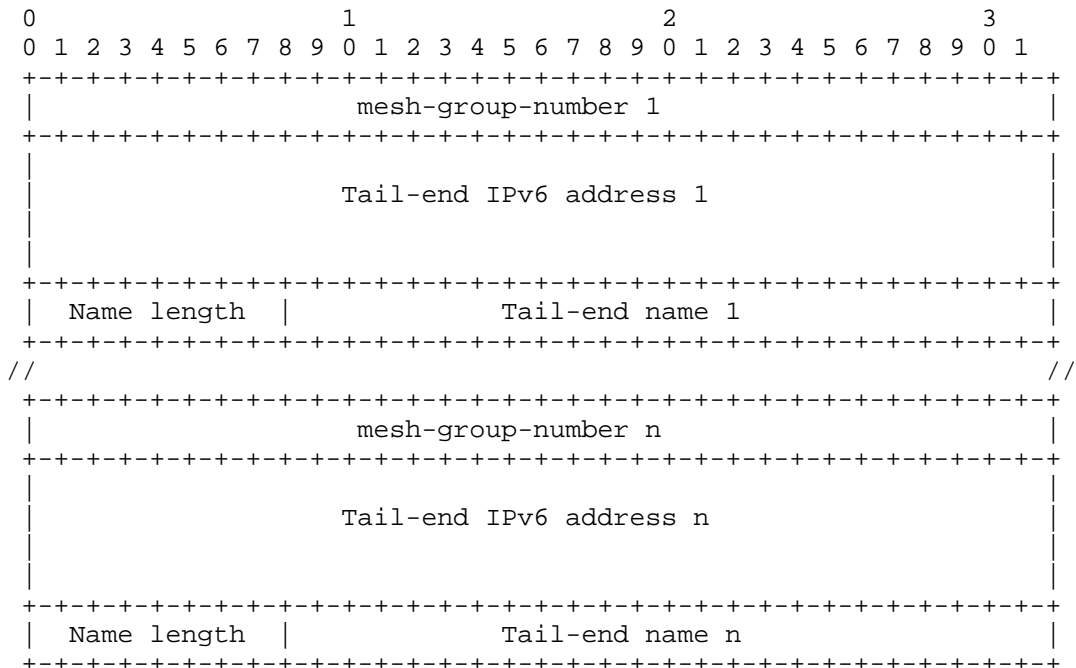


Figure 3 - OSPF TE-MESH-GROUP TLV format (IPv6 Address)

The OSPF TE-MESH-GROUP TLV may contain one or more mesh-group entries, where each entry corresponds to a TE mesh-group and is made of the following fields:

- A mesh-group-number that identifies the mesh-group number.
- A Tail-end address: an IPv4 or IPv6 IP address to be used as a tail-end TE LSP address by other LSRs belonging to the same mesh-group.
- Name length field: An integer, expressed in octets, that indicates the length of the Tail-end name before padding.
- A Tail-end name: A display string that is allocated to the Tail-end. The field is of variable length field and is used to facilitate the TE LSP identification.

4.2. IS-IS TE-MESH-GROUP Sub-TLV Format

The TE-MESH-GROUP sub-TLV is used to advertise the desire of an LSR to join/leave a given TE mesh-group. No sub-TLV is currently defined for the TE-MESH-GROUP sub-TLV.

The IS-IS TE-MESH-GROUP sub-TLV (advertised in the IS-IS CAPABILITY TLV defined in [RFC4971]) is composed of 1 octet for the type, 1 octet specifying the TLV length and a value field. The format of the TE-MESH-GROUP sub-TLV is identical to the TLV format used by the Traffic Engineering Extensions for IS-IS [RFC3784].

The IS-IS TE-MESH-GROUP sub-TLV format for IPv4 (Figure 4) and IPv6 (Figure 5) is as follows:

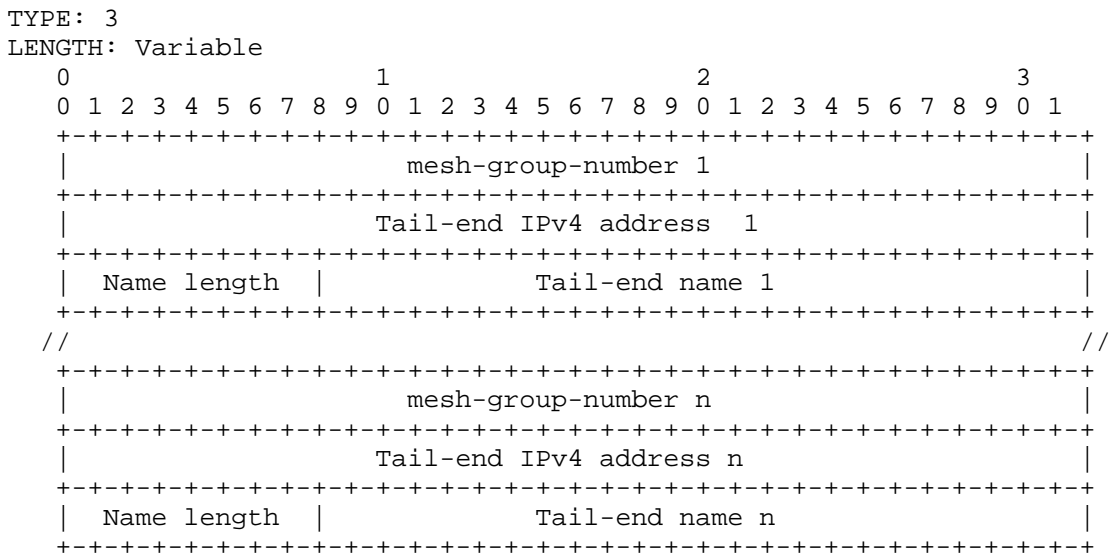


Figure 4 - IS-IS TE-MESH-GROUP sub-TLV format (IPv4 Address)

TYPE: 4
LENGTH: Variable

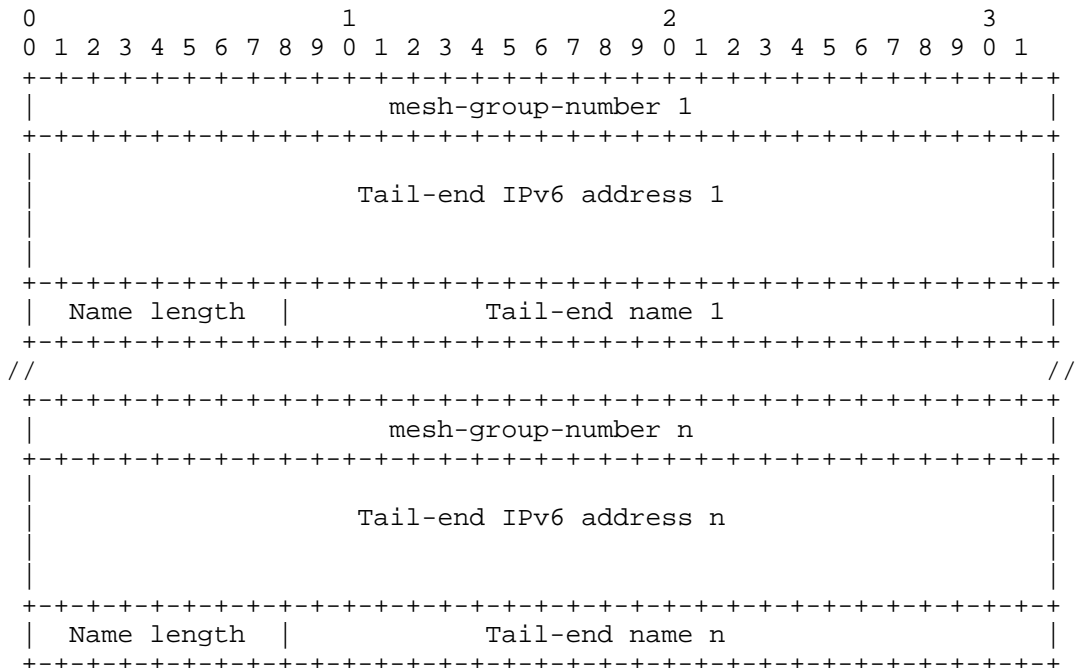


Figure 5 - IS-IS TE-MESH-GROUP sub-TLV format (IPv6 Address)

The IS-IS TE-MESH-GROUP sub-TLV may contain one or more mesh-group entries where each entry correspond to a TE mesh-group and is made of the following fields:

- A mesh-group-number that identifies the mesh-group number.
- A Tail-end address: an IPv4 or IPv6 IP address to be used as a tail-end TE LSP address by other LSRs belonging to the same mesh-group.
- Name length field: An integer, expressed in octets, that indicates the length of the Tail-end name before padding.
- A Tail-end name: A display string that is allocated to the Tail-end. The field is of variable length and is used to facilitate the TE LSP identification.

5. Elements of Procedure

The OSPF TE-MESH-GROUP TLV is carried within the OSPF Routing Information LSA and the IS-IS TE-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-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-MESH-GROUP TLV occurs more than once within the OSPF Router Information LSA, only the first instance is processed, subsequent TLV(s) SHOULD be silently ignored. Similarly, if either the IPv4 or the IPv6 IS-IS TE-MESH-GROUP sub-TLV occurs more than once within the IS-IS Router capability TLV, only the first instance is processed, subsequent TLV(s) SHOULD be silently ignored.

5.1. OSPF

The TE-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 [RFC2740]. 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 TE mesh group, it MUST originate a new OSPF Router Information LSA comprising the updated TE-MESH-GROUP TLV. In the case of a join, a new entry will be added to the TE-MESH-GROUP TLV; conversely, if the LSR leaves, a mesh-group the corresponding entry will be removed from the TE-MESH-GROUP TLV. Note that both 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-MESH-GROUP TLV from a specific LSR.

As defined in [RFC2370] for OSPFv2 and in [RFC2740] 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 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-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 mesh-group TLV MUST be generated within a Routing-domain scope router information LSA.

5.2. IS-IS

The TE-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 TE mesh group, it MUST originate a new LSP comprising the refreshed IS-IS Router capability TLV comprising the updated TE-MESH-GROUP sub-TLV. In the case of a join, a new entry will be added to the TE-MESH-GROUP sub-TLV; conversely, if the LSR leaves a mesh-group, the corresponding entry will be deleted from the TE-MESH-GROUP sub-TLV. Note that both 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-MESH-GROUP sub-TLV from a specific LSR.

If the flooding scope of a TE-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-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.

6. Backward Compatibility

The TE-MESH-GROUP TLVs defined in this document do not introduce any interoperability issue. For OSPF, a router not supporting the TE-MESH-GROUP TLV SHOULD just silently ignore the TLV as specified in [RFC2370]. For an IS-IS, a router not supporting the TE-MESH-GROUP sub-TLV SHOULD just silently ignore the sub-TLV.

7. IANA Considerations

7.1. OSPF

The registry for the Router Information LSA is defined in [RFC4970]. IANA assigned a new OSPF TLV code-point for the TE-MESH-GROUP TLVs carried within the Router Information LSA.

Value	Sub-TLV	References
-----	-----	-----
3	TE-MESH-GROUP TLV (IPv4)	RFC 4972 (this doc)
4	TE-MESH-GROUP TLV (IPv6)	RFC 4972 (this doc)

7.2. 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-MESH-GROUP sub-TLVs carried within the IS-IS Router Capability TLV.

Value	Sub-TLV	References
-----	-----	-----
3	TE-MESH-GROUP TLV (IPv4)	RFC 4972 (this doc)
4	TE-MESH-GROUP TLV (IPv6)	RFC 4972 (this doc)

8. Security Considerations

The function described in this document does not create any new security issues for the OSPF and IS-IS protocols. Security considerations are covered in [RFC2328] and [RFC2740] for the base OSPF protocol and in [RFC1195] for IS-IS. It must be noted that the advertisement of "fake" TE Mesh Group membership(s) by a mis-configured or malicious LSR Y would not have any major impact on the network (other than overloading the IGP), such as triggering the set up of new MPLS TE LSP: indeed, for a new TE LSP originated by another LSR X destined to LSR Y to be set up, the same TE Mesh group membership must be configured on both LSRs. Thus such fake advertisement could not amplify any Denial of Service (DoS) attack.

9. Acknowledgements

We would like to thank Dean Cheng, Adrian Farrel, Yannick Le Louedec, Dave Ward, Les Ginsberg, Stephen Nadas, Acee Lindem, Dimitri Papadimitriou, and Lakshminath Dondeti for their useful comments.

10. References

10.1. Normative References

- [RFC4971] Vasseur, J-P., Ed., Shen, N., Ed., and R. Aggarwal, Ed., "Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information", RFC 4971, July 2007.
- [RFC4970] Lindem, A., Ed., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", RFC 4970, July 2007.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, December 1990.
- [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.
- [RFC2370] Coltun, R., "The OSPF Opaque LSA Option", RFC 2370, July 1998.
- [RFC2740] Coltun, R., Ferguson, D., and J. Moy, "OSPF for IPv6", RFC 2740, December 1999.

10.2. Informative References

- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.
- [RFC3784] Smit, H. and T. Li, "Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)", RFC 3784, June 2004.

Authors' Addresses

JP Vasseur (editor)
Cisco Systems, Inc
1414 Massachusetts Avenue
Boxborough, MA 01719
USA

E-Mail: jpv@cisco.com

JL Le Roux (editor)
France Telecom
2, Avenue Pierre-Marzin
Lanion, 22307
FRANCE

E-Mail: jeanlouis.leroux@orange-ftgroup.com

Seisho Yasukawa
NTT
3-1, Otemachi 2-Chome Chiyoda-ku
Tokyo, 100-8116
JAPAN

E-Mail: s.yasukawa@hco.ntt.co.jp

Stefano Previdi
Cisco Systems, Inc
Via Del Serafico 200
Roma, 00142
Italy

E-Mail: sprevidi@cisco.com

Peter Psenak
Cisco Systems
Mlynske Nivy 43
821 09
Bratislava
Slovakia

EEmail: ppsenak@cisco.com

Paul Mabbey
Comcast Cable
4100 E. Dry Creek Rd
Centennial, CO 80122
USA

EEmail: Paul_Mabey@cable.comcast.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.