

CCAMP Working Group

IETF Internet Draft

JP Vasseur (Ed.)
Cisco System Inc.
JL Le Roux (Ed.)
France Telecom

Proposed Status: Standard

Expires: March 2006

September 2005

Routing extensions for discovery of Multiprotocol (MPLS) Label Switch Router (LSR) Traffic Engineering (TE) mesh membership

[draft-vasseur-ccamp-automesh-02.txt](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of 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/lid-abstracts.txt>.

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

Copyright Notice

Copyright (C) The Internet Society (2005). All Rights Reserved.

Abstract

The set up of a full mesh of MPLS TE LSPs among a set of Label Switch Router (LSR) is common deployment scenario of MPLS Traffic

or fast rerouting with MPLS Fast Reroute. Such deployment requires the configuration of potentially a large number of TE LSPs (on the order of the square of the number LSRs). This document specifies IGP (OSPF and IS-IS) traffic engineering extensions so as to provide an automatic discovery of the set of LSRs members of a mesh, leading to an automatic mechanism to set up TE LSP mesh(es) (also referred to as a mesh-group in this document).

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](#).

Table of Contents

1. Contributors-----	2
2. Terminology-----	3
3. Introduction-----	3
4. TE mesh-roup-----	4
4.1. Description-----	4
4.2. Required Information-----	4
5. TE-MESH-GROUP TLV formats-----	4
5.1. OSPF TE-MESH-GROUP TLV format-----	5
5.2. IS-IS TE-MESH-GROUP TLV format-----	6
6. Elements of procedure-----	7
6.1. OSPF-----	7
6.2. IS-IS-----	7
7. Backward compatibility-----	8
8. Security Considerations-----	8
9. Intellectual Property Statement-----	8
10. Acknowledgment-----	9
11. References-----	9
11.1. Normative references-----	9
11.2. Informative References-----	9
12. Editors' Address-----	10

1. Contributors

This document was the collective work of several. The text and content of this document was contributed by the editors and the

co-authors listed below (the contact information for the editors appears in [section 12](#), and is not repeated below):

Paul Mabey
Qwest Communications
950 17th street
Denver, CO 80202
USA
Email: pmabey@qwest.com

Seisho Yasukawa
NTT
9-11, Midori-Cho 3-Chome
Musashino-Shi, Tokyo 180-8585
JAPAN
Email: yasukawa.seisho@lab.ntt.co.jp

Vasseur, Le Roux et al.

[Page 2]

Internet Draft

[draft-vasseur-ccamp-automesh-02](#)

September 2005

Stefano Previdi
Cisco System, Inc.
Via del Serafico 200
00142 Roma
ITALY
Email: sprevidi@cisco.com

Peter Psenak
Cisco System, Inc.
Pegasus Park
DE Kleetlaan 6A
1831, Diegmen
BELGIUM
Email: ppsenak@cisco.com

[2.](#) Terminology

Terminology used in this document

LSR: Label Switch Router.

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.

IGP Area: OSPF Area or IS-IS level

Link State Advertisement: An OSPF LSA or IS-IS LSP

Intra-area TE LSP: TE LSP whose path does not transit across areas.

Inter-area TE LSP: A TE LSP whose path transits across at least two different IGP areas.

Inter-AS MPLS TE LSP: A TE LSP whose path transits across at least two different ASes or sub-ASes (BGP confederations).

3. Introduction

As of today, there are different approaches in deploying MPLS Traffic Engineering:

- (1) The 'systematic' approach consisting of setting up a full mesh of TE LSPs between a set of LSRs,
- (2) The 'by exception' approach whereby a set of TE LSPs are provisioned on hot spots to alleviate a congestion resulting for instance from an unexpected traffic growth in some part of the network.

The set up of a full mesh of MPLS 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 ([FRR]). Setting up a full mesh of TE LSPs between a set of LSRs requires the configuration of a potentially large

number of TE LSPs on every head-end LSR. The resulting total number of TE LSP in a full TE mesh of n LSRs is $O(n^2)$. 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 terminating to this new LSR, which gives a total of $2*n$ TE LSPs. Such operation is not only time consuming but also a risky operation for Service Providers. Hence, a more automatic mechanism to setting up one or more full meshes of TE LSPs is desirable and requires the ability to automatically discover the LSRs that belong to the mesh.

MPLS Traffic Engineering (MPLS-TE) routing ([IS-IS-TE], [OSPF-TE]) relies on extensions to link state IGP routing protocols ([OSPF], [IS-IS]) in order to carry Traffic Engineering link information used for constraint based routing. Generalized MPLS (GMPLS) related routing extensions are defined in [IS-IS-G] and [OSPF-G].

Further routing extensions have been defined in [OSPF-CAPS] and [IS-IS-CAPS] so as to advertise router capabilities. This document specifies IGP (OSPF and IS-IS) traffic engineering capability TLVs in order to provide a mechanism to automatically discover the LSR members of a mesh, leading to an automatic mechanism to set up TE LSP mesh (also referred to as a mesh-group in this document) in a network. The routing extensions specified in this document provide the ability to signal multiple TE meshes whereby an LSR can belong to

one or more TE meshes.

4. TE mesh-group

4.1. Description

A TE mesh-group is defined as a group of LSRs that are connected by a full mesh of TE LSPs. It is useful to dynamically advertise the desire of a node to join/leave a particular TE mesh-group. This allows for an automatic provisioning of a full mesh of TE LSPs, and thus drastically reduces the configuration overhead and risk of mis-configuration.

4.2. Required Information


This document specifies a TE-MESH-GROUP TLV that indicates the set of TE mesh-group(s) an LSR belongs to. For each TE mesh group announced by the LSR, the TE-MESH-GROUP TLV carries the following information:

- A mesh-group number identifying the TE mesh-group,
- A Tail-end address (address used as a tail end address by other LSRs belonging to the same mesh-group),
- A Tail-end name: string used to ease the TE-LSP naming (e.g. 'head-name->tail-name').

5. TE-MESH-GROUP TLV formats

5.1. OSPF TE-MESH-GROUP TLV format

The OSPF TE-MESH-GROUP TLV (carried in an OSPF router information LSA as defined in [[OSPF-CAP](#)]) has the following format:



OSPF TE-MESH-GROUP TLV format

Where

Type: identifies the TLV type

Length: length of the value field in octets

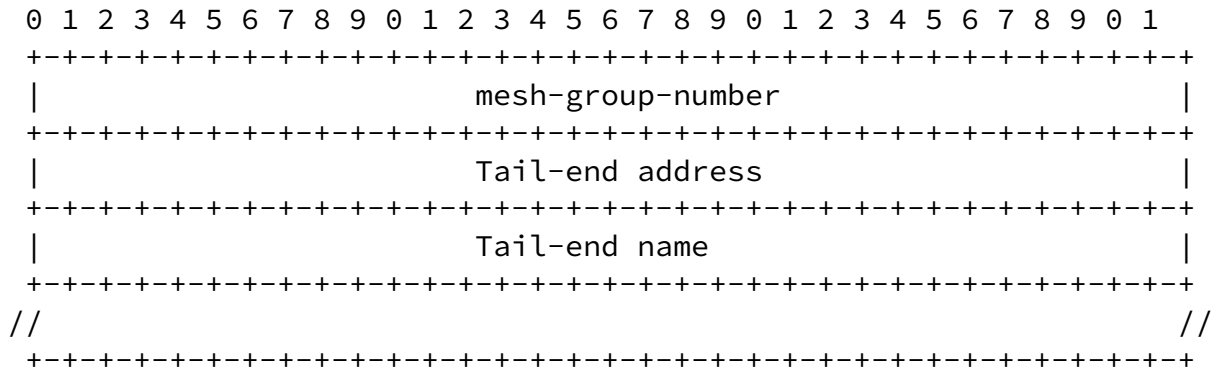
The format of the OSPF TE-MESH-GROUP TLV is the same as the TLV format used by the Traffic Engineering Extensions to OSPF [OSPF-TE]. The TLV is padded to four-octet alignment; padding is not included in the length field (so a three octet value would have a length of three, but the total size of the TLV 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 TE-MESH-GROUP TLV is used to advertise the desire to join/leave a given MPLS TE mesh group. No sub-TLV is currently defined for the TE-mesh-group TLV.

The TE-MESH-GROUP TLV has the following format:

CODE: 3

LENGTH: Variable (N*12 octets)



TE-MESH-GROUP TLV format

N is the number of mesh-groups.

For each TE mesh group announced by the LSR, the TE-MESH-GROUP TLV contains:

- A mesh-group-number: identifies the mesh-group number,
- A Tail-end address: user configurable IP address to be used as a tail-end address by other LSRs belonging to the same mesh-group.

- A Tail-end name: 32-bits string which facilitates the TE LSP identification which can be very useful in some environments such as inter-area/AS MPLS TE environments.

5.2. IS-IS TE-MESH-GROUP TLV format

The IS-IS TE-MESH-GROUP TLV 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 TLV is identical to the TLV format used by the Traffic Engineering Extensions to IS-IS [[IS-IS-TE](#)].

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

The TE-MESH-GROUP TLV has the following format:

```

CODE: 2
LENGTH: Variable (N*12 octets)

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     mesh-group-number                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Tail-end address                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Tail-end name                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
//                                                                                               //
+-----+-----+-----+-----+-----+-----+-----+-----+

```

TE-MESH-GROUP TLV format

N is the number of mesh-groups.

For each Mesh-group announced by an LSR, the TLV contains:

- A mesh-group-number: identifies the mesh-group number,
- A Tail-end address: user configurable IP address to be used as a tail-end address by other LSRs belonging to the same mesh-group.
- A Tail-end name: 32-bits string which facilitates the TE LSP identification which can be very useful in inter-area/AS MPLS TE environments.

6. Elements of procedure

The TE-MESH-GROUP TLV is carried in Link State Advertisements (LSA) and Router capability TLV (carried itself within a Link State Packet (LSP)) for OSPF and ISIS respectively. As such, elements of procedures are inherited from those defined in [OSPF-CAPS] and [IS-IS-CAPS]. 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. refresh).

The TE-MESH-GROUP TLV is OPTIONAL.

6.1. OSPF

The TE-MESH-GROUP TLV is carried within an OSPF router information opaque LSA (opaque type of 4, opaque ID of 0) as defined in [OSPF-CAP].

A router MUST originate a new OSPF router information LSA whenever the content of the any of the carried TLV changes or whenever required by the regular OSPF procedure (LSA refresh (every LSRefreshTime)).

As defined in [RFC2370](#), an opaque LSA has a flooding scope determined by its LSA type:

- link-local (type 9),
- area-local (type 10)
- entire OSPF routing domain (type 11). In this case, the flooding scope is equivalent to the Type 5 LSA flooding scope.

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

The TE-MESH-GROUP TLV may be carried within a type 10 or 11 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 have their head-end and tail-end LSR within the same OSPF area), the TE-MESH-GROUP TLV MUST be generated within a Type 10 router information LSA,
- If the MPLS TE mesh-group spans multiple OSPF areas, the TE mesh-group TLV MUST be generated within a Type 11 router information LSA,

6.2. IS-IS

The TE-MESH-GROUP TLV is carried within the IS-IS Router CAPABILITY TLV defined in [\[IS-IS-CAP\]](#).

An IS-IS router MUST originate a new IS-IS LSP whenever the content of the any of the carried sub-TLV changes or whenever required by the regular IS-IS procedure (LSP refresh).

If the flooding scope of an MPLS Traffic Engineering capability is limited to an IS-IS level/area, the 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 an MPLS Traffic Engineering capability is the entire routing domain, the TLV MUST be leaked across levels for IS-IS the S flag of the CAPABILITY TLV MUST be set.

In both cases the flooding rules as specified in [[IS-IS-CAP](#)] apply.

As specified in [[IS-IS-CAP](#)], a router may generate multiple IS-IS CAPABILITY TLVs within an IS-IS LSP with different flooding scopes.

[7](#). 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 IS-IS a router not supporting the TE-MESH-GROUP TLV SHOULD just silently ignore the TLV.

[8](#). Security Considerations

No new security issues are raised in this document.

[9](#). Intellectual Property Statement

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.

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".

10. Acknowledgment

We would like to thank Yannick Le Louedec for his useful comments.

11. References

11.1. Normative references

[RFC] Bradner, S., "Key words for use in RFCs to indicate requirements levels", [RFC 2119](#), March 1997.

[RFC3667] Bradner, S., "IETF Rights in Contributions", [BCP 78](#), [RFC 3667](#), February 2004.

[RFC3668] Bradner, S., Ed., "Intellectual Property Rights in IETF Technology", [BCP 79](#), [RFC 3668](#), February 2004.

[OSPF-v2] Moy, J., "OSPF Version 2", [RFC 2328](#), April 1998.

[IS-IS] "Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol " ISO 10589.

[IS-IS-IP] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", [RFC 1195](#), December 1990.

[OSPF-TE] Katz, D., Yeung, D., Kompella, K., "Traffic Engineering Extensions to OSPF Version 2", [RFC 3630](#), September 2003.

[IS-IS-TE] Li, T., Smit, H., "IS-IS extensions for Traffic Engineering", [RFC 3784](#), June 2004.

[OSPF-CAP] Lindem, A., Shen, N., Aggarwal, R., Shaffer, S., Vasseur, J.P., "Extensions to OSPF for advertising Optional Router Capabilities", [draft-ietf-ospf-cap](#), work in progress.

[IS-IS-CAP] Vasseur, J.P. et al., "IS-IS extensions for advertising router information", [draft-ietf-isis-caps](#), work in progress.

11.2. Informative References

[GMPLS-RTG] Kompella, K., Rekhter, Y., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching", [draft-ietf-ccamp-gmpls-routing-09.txt](#) (work in progress)

[OSPF-G] Kompella, K., Rekhter, Y., "OSPF extensions in support of Generalized Multi-protocol Label Switching", [draft-ietf-ccamp-ospf-gmpls-extensions-12.txt](#), work in progress.

Vasseur, Le Roux et al.

[Page 9]

Internet Draft

[draft-vasseur-ccamp-automesh-02](#)

September 2005

[IS-IS-G] Kompella, K., Rekhter, Y., "IS-IS extensions in support of Generalized Multi-protocol Label Switching", [draft-ietf-isis-gmpls-extensions-19.txt](#), work in progress.

[INT-AREA-REQ] Le Roux, J.L., Vasseur, J.P., Boyle, J. et al, "Requirements for inter-area MPLS Traffic Engineering", [RFC4105](#), June 2005.

[INT-AS-REQ] Zhang, R., Vasseur, J.P. et al, "MPLS Inter-AS Traffic Engineering Requirements", [draft-ietf-tewg-interas-mpls-te-req](#), work in progress.

[INT-DOMAIN-FRWK] Farrel, A., Vasseur, J.P., Ayyangar, A., "A Framework for Inter-Domain MPLS Traffic Engineering", [draft-ietf-ccamp-inter-domain-framework](#), work in progress.

12. Editors' Address

Jean-Philippe Vasseur
Cisco Systems, Inc.
300 Beaver Brook Road
Boxborough , MA - 01719
USA
Email: jpv@cisco.com

Jean-Louis Le Roux
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex
FRANCE
Email: jeanlouis.leroux@francetelecom.com

Full Copyright Statement

Copyright (C) The Internet Society (2005).

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 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.