INTERNET-DRAFT
Network Working Group
Intended Status: Proposed Standard

Expires: October 2007

J.P. Vasseur (Editor)
Cisco Systems, Inc.
J.L. Le Roux (Editor)
France Telecom

April 2007

IGP Routing Protocol Extensions for Discovery of Traffic Engineering Node Capabilities

draft-ietf-ccamp-te-node-cap-05.txt

Status of this Memo

Abstract

It is highly desired in several cases, to take into account Traffic Engineering (TE) node capabilities during Multi Protocol Label Switching (MPLS) and Generalized MPLS (GMPLS) Traffic Engineered Label Switched Path (TE-LSP) selection, such as for instance the capability to act as a branch Label Switching Router (LSR) of a Point-To-MultiPoint (P2MP) LSP. This requires advertising these capabilities within the Interior Gateway Protocol (IGP). For that purpose, this document specifies Open Shortest Path First (OSPF) and Intermediate System-Intermediate System (IS-IS) traffic engineering extensions for the advertisement of control plane and data plane traffic engineering node capabilities.

Table of Contents

1. Terminology

This document uses terminologies defined in [RFC3031], [RFC3209] and [RFC4461].

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].

2. Introduction

Multi Protocol Label Switching-Traffic Engineering (MPLS-TE) routing ([RFC3784], [RFC3630], [OSPFv3-TE]) relies on extensions to link state Interior Gateway Protocols (IGP) ([IS-IS], [RFC1195], [RFC2328], [RFC2740]) in order to advertise Traffic Engineering (TE) link information used for constraint based routing. Further Generalized MPLS (GMPLS) related routing extensions are defined in [RFC4205] and [RFC4203].

It is desired to complement these routing extensions in order to advertise TE node capabilities, in addition to TE link information. These TE node capabilities will be taken into account as constraints during path selection.

Indeed, it is useful to advertise data plane TE node capabilities, such as the capability for a Label Switching Router (LSR) to be a branch LSR or a bud-LSR of a Point-To-MultiPoint (P2MP) Label Switched Path (LSP). These capabilities can then be taken into account as constraints when computing the route of TE LSPs.

It is also useful to advertise control plane TE node capabilities such as the capability to support GMPLS signaling for a packet LSR, or the capability to support P2MP (Point to Multipoint) TE LSP signaling. This allows selecting a path that avoids nodes that do not support a given control plane feature, or triggering a mechanism to support such nodes on a path. Hence this facilitates backward compatibility.

For that purpose, this document specifies IGP (OSPF and IS-IS) extensions in order to advertise data plane and control plane capabilities of a node.

A new TLV is defined for OSPF, the TE Node Capability Descriptor TLV, to be carried within the Router Information LSA ([OSPF-CAP]).
A new sub-TLV is defined for IS-IS, the TE Node Capability Descriptor sub-TLV, to be carried within the IS-IS Capability TLV ([IS-IS-CAP]).

3. TE Node Capability Descriptor

3.1. Description

LSRs in a network may have distinct control plane and data plane Traffic Engineering capabilities. The TE Node Capability Descriptor information defined in this document describes data and control plane capabilities of an LSR. Such information can be used during path computation so as to avoid nodes that do not support a given TE feature either in the control or data plane, or to trigger procedures to handle these nodes along the path (e.g, trigger LSP hierarchy to support a legacy transit LSR on a P2MP LSP (see [RSVP-P2MP])).

3.2. Required Information

The TE Node Capability Descriptor contains a variable length set of bit flags, where each bit corresponds to a given TE node capability.

Five TE Node Capabilities are defined in this document:

- B bit: when set, this flag indicates that the LSR can act as a branch node on a P2MP LSP (see [RFC4461]);
- E bit: when set, this flag indicates that the LSR can act as a bud LSR on a P2MP LSP, i.e. an LSR that is both transit and egress (see [RFC4461]).
- M bit: when set, this flag indicates that the LSR supports MPLS-TE signaling ([RFC3209]);
- G bit: when set this flag indicates that the LSR supports GMPLS signaling ([RFC3473]);
- P bit: when set, this flag indicates that the LSR supports P2MP MPLS-TE signaling ([RSVP-P2MP]).

Note that new capability bits may be added in the future if required.

4. TE Node Capability Descriptor TLV formats

4.1. OSPF TE Node Capability Descriptor TLV format

The OSPF TE Node Capability Descriptor TLV is a variable length TLV that contains a series of bit flags, where each bit correspond to a TE node capability.

The OSPF TE Node Capability Descriptor TLV is carried within an OSPF Router Information LSA which is defined in [OSPF-CAP].

The format of the OSPF TE Node Capability Descriptor TLV is the same as the TLV format used by the Traffic Engineering Extensions to OSPF [RFC3630]. That is, the TLV is composed of 2 octets for the type, 2 octets specifying the length of the value field and a value field.

The OSPF TE Node Capability Descriptor TLV has the following format:

```
TYPE:    Assigned by IANA - see Section 8.1.
LENGTH:  Variable (multiple of 4).
VALUE:   Array of units of 32 flags numbered from the most
         significant bit as bit zero, where each bit represents
         a TE node capability.
```

The following bits are defined:

Bit	Capabilities
0	B bit: P2MP Branch Node capability: When set this indicates that the LSR can act as a branch node on a P2MP LSP [RFC4461].
1	E bit: P2MP Bud-LSR capability: When set, this indicates that the LSR can act as a bud LSR on a P2MP LSP, i.e. an LSR that is both transit and egress [RFC4461].
2	M bit: If set this indicates that the LSR supports MPLS-TE signaling ([RFC3209]).
3	G bit: If set this indicates that the LSR supports GMPLS signaling ([RFC3473]).
4	P bit: If set this indicates that the LSR supports P2MP MPLS-TE signaling ([RSVP-P2MP]).
5-31	Reserved for future assignments by IANA.

4.2. IS-IS TE Node Capability Descriptor sub-TLV format

The IS-IS TE Node Capability Descriptor sub-TLV is a variable length sub-TLV that contains a series of bit flags, where each bit correspond to a TE node capability.

The IS-IS TE Node Capability Descriptor sub-TLV is carried within an IS-IS CAPABILITY TLV which is defined in [IS-IS-CAP].

The format of the IS-IS TE Node Capability sub-TLV is the same as the TLV format used by the Traffic Engineering Extensions to IS-IS [RFC3784]. That is, the TLV is composed of 1 octet for the type, 1 octet specifying the TLV length and a value field.

The IS-IS TE Node Capability Descriptor sub-TLV has the following format:

TYPE: Assigned by IANA - see Section 8.2.
LENGTH: Variable
VALUE: Array of units of 8 flags numbered from the most significant bit as bit zero, where each bit represents a TE node capability.

The following bits are defined:

Bit	Capabilities
0	B bit: P2MP Branch Node capability: When set this indicates that the LSR can act as a branch node on a P2MP LSP [RFC4461].
1	E bit: P2MP Bud-LSR capability: When set, this indicates that the LSR can act as a bud LSR on a P2MP LSP, i.e. an LSR that is both transit and egress [RFC4461].
2	M bit: If set this indicates that the LSR supports MPLS-TE signaling ([RFC3209]).
3	G bit: If set this indicates that the LSR supports GMPLS signaling ([RFC3473]).
4	P bit: If set this indicates that the LSR supports P2MP MPLS-TE signaling ([RSVP-P2MP]).
5-7	Reserved for future assignments by IANA.

5. Elements of procedure

5.1. OSPF

The TE Node Capability Descriptor TLV is advertised, within an OSPFv2 Router Information LSA (Opaque type of 4 and Opaque ID of 0) or an OSPFv3 Router Information LSA (function code of 12) which are defined in [OSPF-CAP]. As such, elements of procedure are inherited from those defined in [RFC2328], [RFC2740], and [OSPF-CAP].

The TE Node Capability Descriptor TLV advertises capabilities that may be taken into account as constraints during path selection. Hence its flooding scope is area-local, and it MUST be carried within OSPFv2 type 10 Router Information LSA (as defined in [RFC2370]) or an OSPFv3 Router Information LSA with the S1 bit set and the S2 bit cleared (as defined in [RFC2740]).

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

The TE Node Capability Descriptor TLV is OPTIONAL and MUST NOT appear more than once in an OSPF Router Information LSA. If a TE Node Capability Descriptor TLV appears more than once in an OSPF Router Information LSA, only the first occurrence MUST be processed and other MUST be ignored.

When an OSPF LSA does not contain any TE Node capability Descriptor TLV, this means that the TE Capabilities of that LSR are unknown.

Note that a change in any of these capabilities MAY trigger CSPF computation, but MUST NOT trigger normal SPF computation.

Note also that TE node capabilities are expected to be fairly static. They may change as the result of configuration change, or software upgrade. This is expected not to appear more than once a day.

5.2. IS-IS

The TE Node Capability sub-TLV is carried within an IS-IS CAPABILITY TLV defined in [IS-IS-CAP]. As such, elements of procedure are inherited from those defined in [IS-IS-CAP].

The TE Node Capability Descriptor sub-TLV advertises capabilities that may be taken into account as constraints during path selection. Hence its flooding is area-local, and MUST be carried within an IS-IS CAPABILITY TLV having the S flag cleared.

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

The TE Node Capability Descriptor sub-TLV is OPTIONAL and MUST NOT appear more than once in an ISIS Router Capability TLV.

When an IS-IS LSP does not contain any TE Node capability Descriptor sub-TLV, this means that the TE Capabilities of that LSR are unknown.

Note that a change in any of these capabilities MAY trigger CSPF computation, but MUST NOT trigger normal SPF computation.

Note also that TE node capabilities are expected to be fairly static. They may change as the result of configuration change, or software upgrade. This is expected not to appear more than once a day.

6. Backward Compatibility

The TE Node Capability Descriptor TLVs defined in this document do not introduce any interoperability issue. For OSPF, a router not supporting the TE Node Capability Descriptor TLV will just silently ignore the TLV as specified in [OSPF-CAP]. For IS-IS a router not supporting the TE Node Capability Descriptor sub-TLV will just silently ignore the sub-TLV as specified in [IS-IS-CAP].

When the TE Node capability Descriptor TLV is absent, this means that the TE Capabilities of that LSR are unknown.

The absence of a word of capability flags in OSPF or an octet of capability flags in IS-IS means that these capabilities are unknown.

7. Security Considerations

This document specifies the content of the TE Node Capability Descriptor TLV in ISIS and OSPF, to be used for (G)MPLS-TE path computation. As this TLV is not used for SPF computation or normal routing, the extensions specified here have no direct effect on IP routing. Tampering with this TLV may have an effect on Traffic Engineering computation. Mechanisms defined to secure ISIS Link State PDUs [RFC3567], OSPF LSAs [RFC2154], and their TLVs, can be used to secure this TLV as well.

8. IANA considerations

8.1. OSPF TLV

[OSPF-CAP] defines a new code point registry for TLVs carried in the Router Information LSA defined in [OSPF-CAP].

IANA is requested to make a new codepoint assignment from that registry for the TE Node Capability Descriptor TLV defined in this document and carried within the Router Information LSA. The value 1 is suggested. See Section 4.1 of this document.

8.2. ISIS sub-TLV

[IS-IS-CAP] defines a new code point registry for sub-TLVs carried in the ISIS CAPABILITY TLV defined in [IS-IS-CAP].

IANA is requested to make a new codepoint assignment from that registry for the TE Node Capability Descriptor sub-TLV defined in this document, and carried within the ISIS CAPABILITY TLV. The value 1 is suggested. See Section 4.2 of this document.

8.3. Capability Registry

IANA is requested to create a new registry to manage the space of capability bit flags carried within the OSPF and ISIS TE Node Capability Descriptor.

A single registry must be defined for both protocols. It is suggested that a new base registry be created to cover IGP-TE registries that apply to both OSPF and ISIS, and that the new registry requested by this document should be a sub-registry of this new bas registry.

Bits in the new regstry should be numbered in the usual IETF notation starting with the most significant bit as bit zero.

New bit numbers may be allocated only by an IETF Consensus action.

Each bit should be tracked with the following qualities:
- Bit number
- Defining RFC
- Name of bit

IANA is requested to make assignments for the five TE node capabilities defined in this document (see Sections 8.1 and 8.2) using the following suggested values:

Bit No.	Name	Reference
1	B bit: P2MP Branch LSR capability	[This.I-D]
2	E bit: P2MP Bud LSR capability	[This.I-D]
3	M bit: MPLS-TE support	[This.I-D]
4	G bit: GMPLS support	[This.I-D]
5	P bit: P2MP RSVP-TE support	[This.I-D]

9. Acknowledgments

We would like to thank Benoit Fondeviole, Adrian Farrel, Dimitri Papadimitriou, Acee Lindem and David Ward for their useful comments and suggestions.

We would also like to thank authors of [RFC4420] and [OSPF-CAP] by which some text of this document has been inspired.

Adrian Farrel prpeared the final version of this document for submission to the IESG.

10. References

10.1. Normative references

[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", 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.

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[RFC3209] Awduche, D., et. al., "RSVP-TE: Extensions to RSVP for LSP tunnels", RFC 3209, December 2001.

[RFC3473] Berger, L, et. al., "GMPLS Signaling RSVP-TE extensions", RFC 3473, January 2003.

[RFC3630] Katz, D., Yeung, D., Kompella, K., "Traffic Engineering Extensions to OSPF Version 2", RFC 3630, September 2003.

[RFC3784] Li, T., Smit, H., "IS-IS extensions for Traffic Engineering", RFC 3784, June 2004.

[IS-IS] "Intermediate System to Intermediate System Intra-Domain Routeing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", ISO 10589.

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

[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.

[RSVP-P2MP] Aggarwal, Papadimitriou, Yasukawa, et. al. "Extensions to RSVP-TE for point-to-multipoint TE LSPs", draft-ietf-mpls-rsvp-te-p2mp, work in progress.

10.2. Informative References

[RFC2154] Murphy, S., Badger, M., and B. Wellington, "OSPF with Digital Signatures", RFC 2154, June 1997.

[RFC3567] Li, T. and R. Atkinson, "Intermediate System to Intermediate System (IS-IS) Cryptographic Authentication", RFC 3567, July 2003.

[RFC4203] Kompella, K., Rekhter, Y., "OSPF extensions in support of Generalized Multi-protocol Label Switching", RFC4203, October 2005.

[RFC4205] Kompella, K., Rekhter, Y., "IS-IS extensions in support of Generalized Multi-protocol Label Switching", RFC4205, October 2005.

[RFC4420] Farrel, A., and al., "Encoding of attributes for MPLS LSPs establishment Using RSVP-TE", RFC4420, February 2006.

[RFC4461] Yasukawa, S., et. al., "Signaling Requirements for Point to Multipoint Traffic Engineered MPLS LSPs", RFC4461, April 2006.

[OSPFv3-TE] Ishiguro K., Manral V., Davey A., and Lindem A. "Traffic Engineering Extensions to OSPF version 3", draft-ietf-ospf-ospfv3-traffic, work in progress.

11. Editors' Addresses

Jean-Philippe Vasseur
Cisco Systems, Inc.
1414 Massachusetts Avenue
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@orange-ftgroup.com

12. Contributors' Addresses

Seisho Yasukawa
NTT
3-9-11 Midori-cho,
Musashino-shi, Tokyo 180-8585, Japan
Email: s.yasukawa@hco.ntt.co.jp

Stefano Previdi
Cisco Systems, Inc
Via Del Serafico 200
Roma, 00142
Italy
Email: sprevidi@cisco.com

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

Paul Mabbey
Comcast
USA

13. 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.

Disclaimer of Validity

Copyright Statement