Network Working Group
IETF Internet Draft

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

Proposed Status: Standard Track
Expires: March 2006

October 2005

Routing extensions for discovery of Traffic Engineering Node Capabilities

draft-vasseur-ccamp-te-node-cap-01.txt

Status of this Memo

Abstract

It is highly desired in several cases, to take into account Traffic Engineering (TE) node capabilities during TE LSP path selection, such as for instance the capability to act as a branch LSR of a P2MP LSP. This requires advertising these capabilities within the IGP.
For that purpose, this document specifies OSPF and IS-IS traffic engineering extensions for the advertisement of control plane and data plane traffic engineering node capabilities.

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

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

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

LSR: Label Switch Router.

TE LSP: Traffic Engineering Label Switched Path.

P2MP TE LSP: A TE LSP that has one unique ingress LSR and one or more egress LSRs.

Branch LSR: An LSR on a P2MP LSP that has more than one directly connected downstream LSRs.

Bud-LSR: An LSR on a P2MP LSP, that is an egress, but also has one or more directly connected downstream LSRs.

3. Introduction

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 (TE) link information used for constraint based routing. Further Generalized MPLS (GMPLS) related routing extensions are defined in [IS-IS-G] and [OSPF-G].

It is desired to complement these routing extensions in order to carry 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, for instance the capability to be a branch LSR or a bud-LSR of a P2MP LSP. These capabilities are then taken into account as constraints when computing TE LSP paths.
It is also useful to advertise control plane TE node capabilities such as for instance 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 signaling feature, or triggering a mechanism to support such nodes. Hence this facilitates backward compatibility.

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

A new TLV is defined for ISIS and OSPF: the TE Node Capability Descriptor TLV, to be carried within:

- the ISIS Capability TLV ([ISIS-CAP]) for ISIS
- the Router Information LSA ([OSPF-CAP]), for OSPF.

4. TE Node Capability Descriptor

4.1. 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 for instance 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 procedure to handle these nodes. In some cases, this may also be useful to ensure backward compatibility.

4.2. Required Information

The TE Node Capability Descriptor contains two variable length sets of bit flags:

-The Data Plane Capabilities: This a variable length set of bit flags where each bit corresponds to a given TE data plane capability.

-The Control Plane Capabilities: This a variable length set of bit flags where each bit corresponds to a given TE control plane capability.

Two Data Plane Capabilities are currently defined:

- -B bit: when set, this flag indicates that the LSR can act as a branch node on a P2MP LSP (see [P2MP-REQ]) and [RSVP-P2MP]).
- -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.

Three Control Plane Capabilities are currently defined:

- -M bit: when set, this flag indicates that the LSR supports MPLS-TE signaling ([RSVP-TE]).
- -G bit: when set this flag indicates that the LSR supports GMPLS signaling ([RSVP-G]).
- -P bit: when set, this flag indicates that the LSR supports P2MP MPLS-TE signaling ([RSVP-P2MP]).

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

Note that bits numbers are under IANA control.

5. OSPF TE extensions

5.1. OSPF TE Node Capability Descriptor TLV format

The OSPF TE Node Capability Descriptor TLV is made of various non ordered sub-TLVs.
The format of the OSPF TE Node Capability Descriptor TLV and its sub-TLVs is the same as the TLV format used by the Traffic Engineering Extensions to OSPF [OSPF-TE]. That is, the TLV is composed of 2 octets for the type, 2 octets specifying the TLV length and a value field.
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 OSPF TE Node Capability Descriptor TLV has the following format:

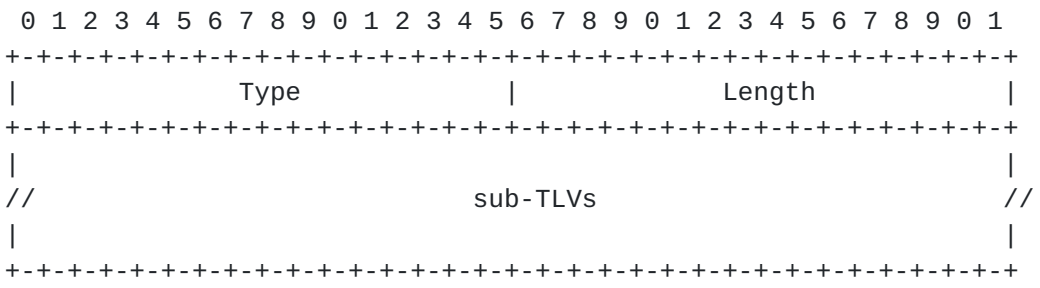

Type To be defined by IANA
Length Variable
Value This comprises one or more sub-TLVs

Currently two sub-TLVs are defined:

Sub-TLV type	Length	Name
1	variable	DATA-PLANE-CAP sub-TLV
2	variable	CONTROL-PLANE-CAP sub-TLV

Any non recognized sub-TLV MUST be silently ignored.
More sub-TLVs could be added in the future to handle new capabilities.

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

5.1.1. The DATA-PLANE-CAP sub-TLV

The DATA-PLANE-CAP sub-TLV is a series of bit flags, where each bit correspond to a data plane TE node capability, and has a variable length.

The format of the DATA-PLANE-CAP sub-TLV is as follows:

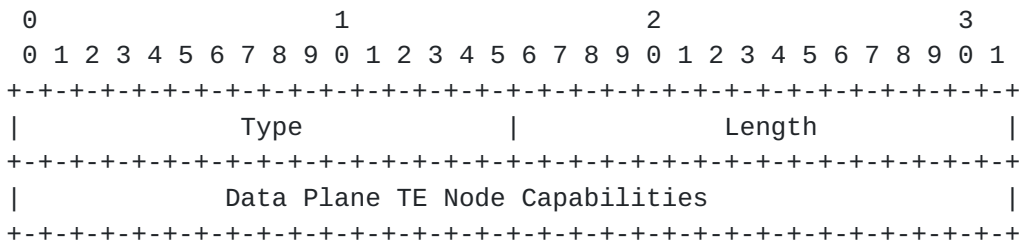

Type To be assigned by IANA (suggested value =1)
Length It is set to N x 4 octets. N starts from 1 and can be increased when there is a need. Each 4 octets are referred to as a capability flag.
Value This comprises one or more capability flags. For each 4 octets, the bits are indexed from the most significant to the least significant, where each bit represents one data plane TE node capability. When the first 32 capabilities are defined, a new capability flag will be used to accommodate the next capability. These bits are under IANA control.

The following bits in the first capability flag are to be assigned by IANA:

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 [P2MP-REQ];
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 [P2MP-REQ];
2-31	Reserved for future assignments by IANA.

5.1.2. The CONTROL-PLANE-CAP sub-TLV

The CONTROL-PLANE-CAP sub-TLV is a series of bit flags, where each bit correspond to a control plane TE node capability, and has a variable length.

The format of the CONTROL-PLANE-CAP sub-TLV is as follows:

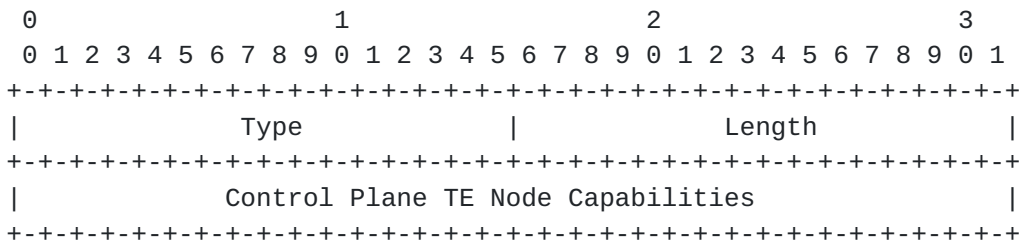

Type	To be assigned by IANA (suggested value = 2)
Length	It is set to N x 4 octets. N starts from 1 and can be increased when there is a need. Each 4 octets are referred to as a capability flag.
Value	This comprises one or more capability flags. For each 4 octets, the bits are indexed from the most Significant to the least significant, where each bit represents one control plane TE node capability. When the first 32 capabilities are defined, a new capability flag will be used to accommodate the next capability. These bits are under IANA control.

The following bits in the first capability flag are to be assigned by IANA:

Bit	Capabilities
0	M bit: If set this indicates that the LSR supports MPLS-TE signaling ([RSVP-TE]).
1	G bit: If set this indicates that the LSR supports GMPLS signaling ([RSVP-G]).
2	P bit: If set this indicates that the LSR supports P2MP MPLS-TE signaling ([RSVP-P2MP]).
3-31	Reserved for future assignments by IANA

5.2. Elements of Procedure

The TE Node Capability Descriptor TLV is carried within an OSPF Router information opaque LSA (opaque type of 4, opaque ID of 0) which is defined in [OSPF-CAP].

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

The TE Node Capability Descriptor TLV advertises capabilities that are taken into account as constraints during path selection. Hence its flooding scope is area-local, and MUST be carried within a type 10 router information LSA.

TE Node Capability Descriptor TLVs are OPTIONAL. 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.

6. IS-IS TE Extensions

6.1. IS-IS TE Node Capability Descriptor TLV format

The IS-IS TE Node Capability Descriptor TLV is made of various non ordered sub-TLVs.

The format of the IS-IS TE Node Capability TLV and its sub-TLVs is the same as the TLV format used by the Traffic Engineering Extensions to IS-IS [ISIS-TE]. 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 TLV has the following format:

TYPE: To be assigned by IANA
LENGTH: Variable, from 3 to 255
VALUE: set of one or more sub-TLVs

Currently two sub-TLVs are defined:

Sub-TLV type	Length	Name
1	variable	DATA-PLANE-CAP sub-TLV
2	variable	CONTROL-PLANE-CAP sub-TLV

Any non recognized sub-TLV MUST be silently ignored.
More sub-TLVs could be added in the future to handle new capabilities.

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

6.1.1. DATA-PLANE-CAP sub-TLV

The DATA-PLANE-CAP sub-TLV is a series of bit flags, where each bit correspond to a data plane TE node capability, and has a variable length. These bits are under IANA control.

The DATA-PLANE-CAP sub-TLV has the following format:

TYPE: To be assigned by IANA (Suggested value =1)
LENGTH: It is set to N. N starts from 1 and can be increased when there is a need. Each octet is referred to as a capability flag.
VALUE: This comprises one or more data plane TE node capability flags.

The following bits in the first capability flag are to be assigned by IANA:

```
0 1 2 3 4 5 6 7
+-+-+-+-+-+-+-+-+
|B|E| Reserved  |
+-+-+-+-+-+-+-+-+
```

B bit: P2MP Branch node capability: When set this indicates that the LSR can act as a branch node on a P2MP LSP [P2MP-REQ]

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 [P2MP-REQ].

Reserved bits are for future assignment by IANA

6.1.2. CONTROL-PLANE-CAP sub-TLV

The CONTROL-PLANE-CAP sub-TLV is a series of bit flags, where each bit correspond to a control plane TE node capability, and has a variable length. These bits are under IANA control.

The CONTROL-PLANE-CAP sub-TLV has the following format:

TYPE: To be assigned by IANA (suggested value = 2)
LENGTH: It is set to N. N starts from 1 and can be increased when there is a need. Each octet is referred to as a capability flag.
VALUE: This comprises one or more control plane TE node capability flags.

The following bits in the first capability flag are to be assigned by IANA.

```
0 1 2 3 4 5 6 7
+-+-+-+-+-+-+-+-+
|M|G|P|Reserved |
+-+-+-+-+-+-+-+-+
```

-M bit: If set this indicates that the LSR supports MPLS-TE signaling ([RSVP-TE]).

-G bit: If set this indicates that the LSR supports GMPLS signaling ([RSVP-G]).

-P bit: If set this indicates that the LSR supports P2MP MPLS-TE signaling ([RSVP-P2MP]).

Reserved bits are for future assignment by IANA.

6.2. Elements of procedure

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

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

The TE Node Capability Descriptor TLV advertises capabilities that are 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.

TE Node Capability Descriptor TLVs are OPTIONAL. When a IS-IS LSP 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.

7. 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 SHOULD just silently ignore the TLV as specified in RFC2370. For IS-IS a router not supporting the TE Node Capability Descriptor TLV SHOULD just silently ignore the TLV.

8. Security Considerations

No new security issues are raised in this document.

9. IANA considerations

9.1. OSPF TLVs

IANA will assign a new codepoint for the TE Node Capability Descriptor TLV defined in this document and carried within the Router Information LSA.

Two sub-TLVs types are defined for this TLV and should be assigned by IANA:

- -CONTROL-PLANE-CAP sub-TLV (suggested value =1)
- -DATA-PLANE-CAP sub-TLV (suggested value =2)

9.2. ISIS TLVs

IANA will assign a new codepoint for the TE Node Capability Descriptor TLV defined in this document, and carried within the ISIS CAPABILITY TLV.

Two sub-TLVs types are defined for this TLV and should be assigned by IANA:

-CONTROL-PLANE-CAP sub-TLV (suggested value =1)
-DATA-PLANE-CAP sub-TLV (suggested value =2)

9.3. Capability bits

IANA is requested to manage the space of control plane and data plane capability bit flags, numbering them in the usual IETF notation starting at zero and continuing at least through 31.
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

Currently two bits are defined in the data plane capability flags. Here are the suggested values:

-0x01: P2MP Branch LSR capability
-0x02: P2MP Bud LSR capability

Currently three bits are defined in the control plane capability flags. Here are the suggested values:

-0x01: MPLS-TE support
-0x02: GMPLS support
-0x04: P2MP RSVP-TE support

10. Acknowledgments

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

We would also like to thank authors of [LSP-ATTRIBUTE] from which some text of this document has been inspired.

11. References

11.1. Normative references

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

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

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

[BCP79] Bradner, S., "Intellectual Property Rights in IETF Technology", RFC 3979, March 2005.

[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.

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

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

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, 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, work in progress.

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

[P2MP-REQ] Yasukawa, S., et. al., "Signaling Requirements for Point to Multipoint Traffic Engineered MPLS LSPs", draft-ietf-mpls-p2mp-sig-requirement, 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-01, work in progress.

[LSP-ATTRIBUTE] Farrel, A., and al., "Encoding of attributes for MPLS LSPs establishment Using RSVP-TE", draft-ietf-mpls-rsvpte-attributes, 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

13. Intellectual Property Statement

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