Network Working Group
Request for Comments: 4205
Updates: 3784
Category: Informational

K. Kompella, Ed.
Y. Rekhter, Ed.
Juniper Networks

October 2005

Intermediate System to Intermediate System (IS-IS) Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)

Status of This Memo

Copyright Notice

Abstract

This document specifies encoding of extensions to the IS-IS routing protocol in support of Generalized Multi-Protocol Label Switching (GMPLS).

1. Introduction

This document specifies extensions to the IS-IS routing protocol in support of carrying link state information for Generalized Multi-Protocol Label Switching (GMPLS). The set of required enhancements to IS-IS are outlined in [GMPLS-ROUTING]. Support for unnumbered interfaces assumes support for the "Point-to-Point Three-Way Adjacency" IS-IS Option type [ISIS-3way].

In this section we define the enhancements to the Traffic Engineering (TE) properties of GMPLS TE links that can be announced in IS-IS Link State Protocol Data Units.

In this document, we enhance the sub-TLVs for the extended IS reachability TLV (see [ISIS-TE]) in support of GMPLS. Specifically, we add the following sub-TLVs:

Sub-TLV Type	Length	Name
4	8	Link Local/Remote Identifiers
20	2	Link Protection Type
21	variable	Interface Switching Capability Descriptor

We further add one new TLV to the TE TLVs:

TLV Type	Length	Name
138	variable	Shared Risk Link Group

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].

1.1. Link Local/Remote Identifiers

A Link Local Interface Identifiers is a sub-TLV of the extended IS reachability TLV. The type of this sub-TLV is 4, and length is eight octets. The value field of this sub-TLV contains four octets of Link Local Identifier followed by four octets of Link Remote Identifier (see Section "Support for unnumbered links" of [GMPLS-ROUTING]). If the Link Remote Identifier is unknown, it is set to 0.

The following illustrates encoding of the Value field of the Link Local/Remote Identifiers sub-TLV.

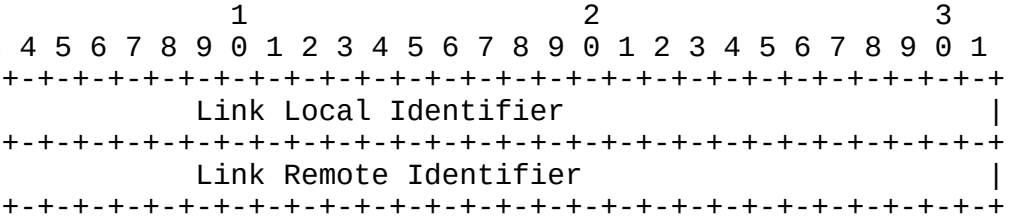

The Link Local/Remote Identifiers sub-TLV MUST NOT occur more than once within the extended IS reachability TLV. If the Link Local/Remote Identifiers sub-TLV occurs more than once within the extended IS reachability TLV, the receiver SHOULD ignore all these sub-TLVs.

1.2. Link Protection Type

The Link Protection Type is a sub-TLV (of type 20) of the extended IS reachability TLV, with length two octets.

The following illustrates encoding of the Value field of the Link Protection Type sub-TLV.

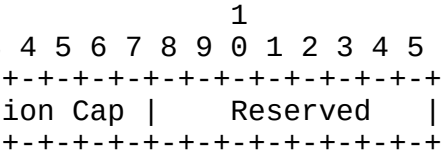

The first octet is a bit vector describing the protection capabilities of the link (see Section "Link Protection Type" of [GMPLS-ROUTING]). They are:

0x01 Extra Traffic

0x02 Unprotected

0x04 Shared

0x08 Dedicated 1:1

0x10 Dedicated 1+1

0x20 Enhanced

0x40 Reserved

0x80 Reserved

The second octet SHOULD be set to zero by the sender, and SHOULD be ignored by the receiver.

The Link Protection Type sub-TLV MUST NOT occur more than once within the extended IS reachability TLV. If the Link Protection Type sub-TLV occurs more than once within the extended IS reachability TLV, the receiver SHOULD ignore all these sub-TLVs.

1.3. Interface Switching Capability Descriptor

The Interface Switching Capability Descriptor is a sub-TLV (of type 21) of the extended IS reachability TLV. The length is the length of value field in octets. The following illustrates encoding of the Value field of the Interface Switching Capability Descriptor sub-TLV.

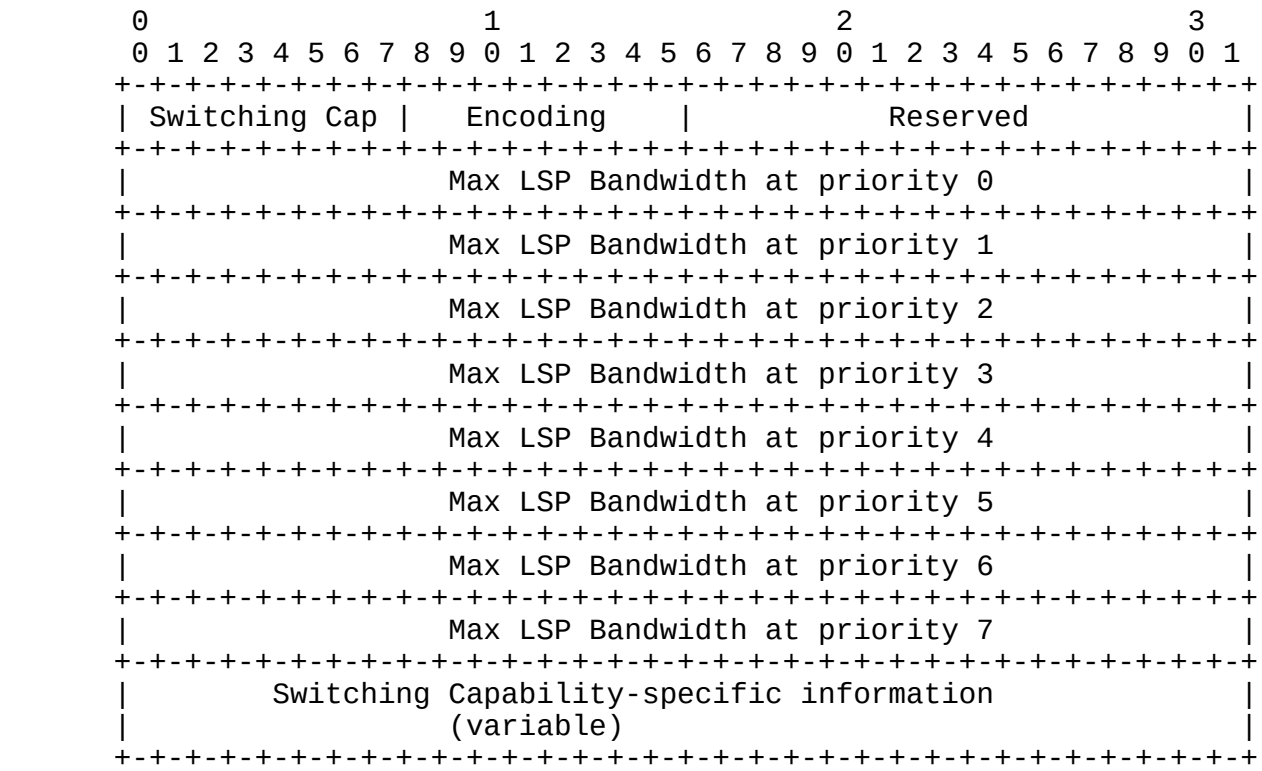

The Switching Capability (Switching Cap) field contains one of the following values:

1	Packet-Switch Capable-1 (PSC-1)
2	Packet-Switch Capable-2 (PSC-2)
3	Packet-Switch Capable-3 (PSC-3)
4	Packet-Switch Capable-4 (PSC-4)
51	Layer-2 Switch Capable (L2SC)
100	Time-Division-Multiplex Capable (TDM)
150	Lambda-Switch Capable (LSC)
200	Fiber-Switch Capable (FSC)

The Encoding field contains one of the values specified in Section 3.1.1 of [GMPLS-SIG].

Maximum LSP Bandwidth is encoded as a list of eight 4 octet fields in the IEEE floating point format [IEEE], with priority 0 first and priority 7 last. The units are bytes (not bits!) per second.

The content of the Switching Capability specific information field depends on the value of the Switching Capability field.

When the Switching Capability field is PSC-1, PSC-2, PSC-3, or PSC-4, the Switching Capability specific information field includes Minimum LSP Bandwidth and Interface MTU.

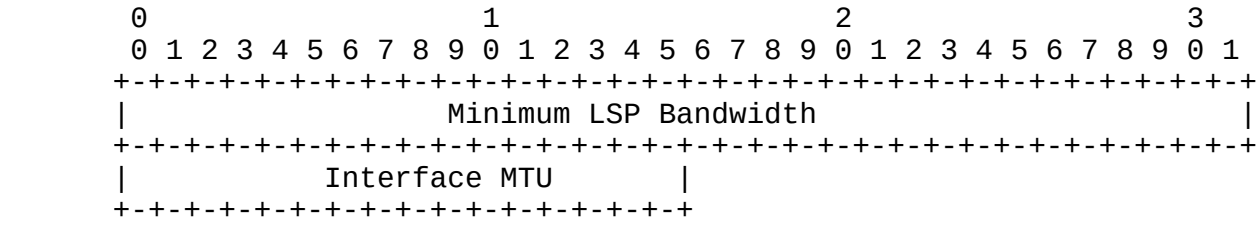

The Minimum LSP Bandwidth is encoded in a 4 octets field in the IEEE floating point format. The units are bytes (not bits!) per second. The Interface MTU is encoded as a 2 octets integer, and carries the MTU value in the units of bytes.

When the Switching Capability field is L2SC, there is no Switching Capability specific information field present.

When the Switching Capability field is TDM, the Switching Capability specific information field includes Minimum LSP Bandwidth and an indication whether the interface supports Standard or Arbitrary SONET/SDH.

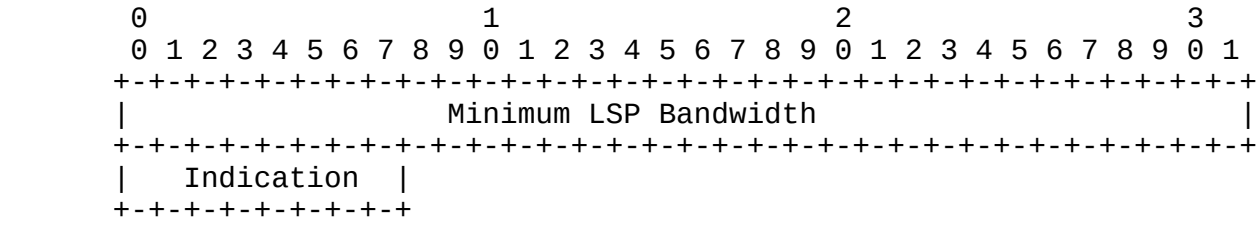

The Minimum LSP Bandwidth is encoded in a 4 octets field in the IEEE floating point format. The units are bytes (not bits!) per second. The indication whether the interface supports Standard or Arbitrary SONET/SDH is encoded as 1 octet. The value of this octet is 0 if the interface supports Standard SONET/SDH, and 1 if the interface supports Arbitrary SONET/SDH.

When the Switching Capability field is LSC, there is no Switching Capability specific information field present.

To support interfaces that have more than one Interface Switching Capability Descriptor (see Section "Interface Switching Capability Descriptor" of [GMPLS-ROUTING]) the Interface Switching Capability Descriptor sub-TLV MAY occur more than once within the extended IS reachability TLV.

1.4. Shared Risk Link Group TLV

The SRLG TLV (of type 138) contains a data structure consisting of:

6 octets of System ID
1 octet of Pseudonode Number
1 octet Flag
4 octets of IPv4 interface address or 4 octets of a Link Local Identifier
4 octets of IPv4 neighbor address or 4 octets of a Link Remote Identifier
(variable) list of SRLG values, where each element in the list has 4 octets.

The following illustrates encoding of the Value field of the SRLG TLV.

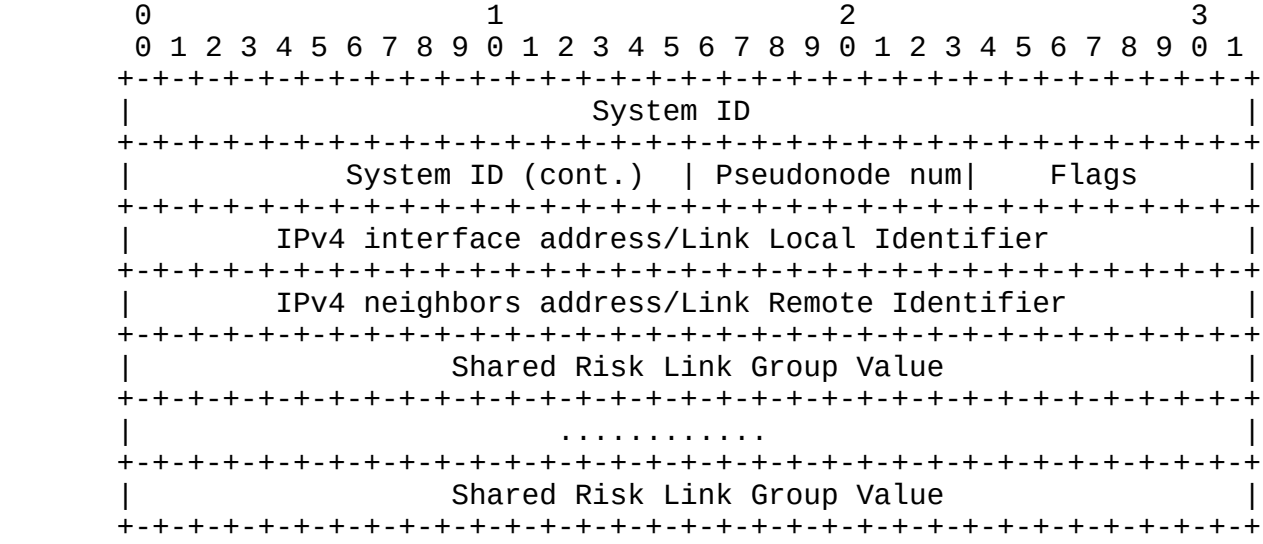

The neighbor is identified by its System Id (6-octets), plus one octet to indicate the pseudonode number if the neighbor is on a LAN interface.

The Least Significant Bit of the Flag octet indicates whether the interface is numbered (set to 1), or unnumbered (set to 0). All other bits are reserved and should be set to 0.

The length of this TLV is 16 + 4 * (number of SRLG values).

This TLV carries the Shared Risk Link Group information (see Section "Shared Risk Link Group Information" of [GMPLS-ROUTING]).

The SRLG TLV MAY occur more than once within the IS-IS Link State Protocol Data Units.

1.5. Link Identifier for Unnumbered Interfaces

Link Identifiers are exchanged in the Extended Local Circuit ID field of the "Point-to-Point Three-Way Adjacency" IS-IS Option type [ISIS-3way].

2. Implications on Graceful Restart

The restarting node SHOULD follow the ISIS restart procedures [ISIS-RESTART], and the RSVP-TE restart procedures [GMPLS-RSVP].

When the restarting node is going to originate its IS-IS Link State Protocol data units for TE links, these Link State Protocol data units SHOULD be originated with 0 unreserved bandwidth, Traffic Engineering Default metric set to 0xffffff, and if the link has LSC or FSC as its Switching Capability then also with 0 as Max LSP Bandwidth, until the node is able to determine the amount of unreserved resources taking into account the resources reserved by the already established LSPs that have been preserved across the restart. Once the restarting node determines the amount of unreserved resources, taking into account the resources reserved by the already established LSPs that have been preserved across the restart, the node SHOULD advertise these resources in its Link State Protocol data units.

In addition, in the case of a planned restart prior to restarting, the restarting node SHOULD originate the IS-IS Link State Protocol data units for TE links with 0 as unreserved bandwidth, and if the link has LSC or FSC as its Switching Capability then also with 0 as Max LSP Bandwidth. This would discourage new LSP establishment through the restarting router.

Neighbors of the restarting node SHOULD continue to advertise the actual unreserved bandwidth on the TE links from the neighbors to that node.

3. Contributors

Ayan Banerjee
Calient Networks
5853 Rue Ferrari
San Jose, CA 95138

Phone: +1 408 972 3645
EMail: abanerjee@calient.net

John Drake
Calient Networks
5853 Rue Ferrari
San Jose, CA 95138

Phone: +1 408 972 3720
EMail: jdrake@calient.net

Greg Bernstein
Grotto Networking

EMail: gregb@grotto-networking.com

Don Fedyk
Nortel Networks Corp.
600 Technology Park Drive
Billerica, MA 01821

Phone: +1 978 288 4506
EMail: dwfedyk@nortelnetworks.com

Eric Mannie
Independent Consultant

EMail: eric_mannie@hotmail.com

Debanjan Saha
Tellium Optical Systems
2 Crescent Place
P.O. Box 901
Ocean Port, NJ 07757

Phone: +1 732 923 4264
EMail: dsaha@tellium.com

Vishal Sharma

EMail: v.sharma@ieee.org

4. Acknowledgements

The authors would like to thank Jim Gibson, Suresh Katukam, Jonathan Lang and Quaizar Vohra for their comments on the draft.

5. Security Considerations

This document specifies the contents of GMPLS TE TLVs in ISIS. As these TLVs are not used for SPF computation or normal routing, the extensions specified here have no direct effect on IP routing. Tampering with GMPLS TE TLVs may have an effect on the underlying transport (optical and/or SONET-SDH) network. Mechanisms to secure ISIS Link State PDUs and/or the TE TLVs [ISIS-HMAC] can be used to secure the GMPLS TE TLVs as well.

6. IANA Considerations

This document defines the following new ISIS TLV type that needs to be reflected in the ISIS TLV code-point registry:

Type	Description	IIH	LSP	SNP
138	Shared Risk Link Group	n	y	n

This document also defines the following new sub-TLV types of top-level TLV 22 that need to be reflected in the ISIS sub-TLV registry for TLV 22:

Type	Description	Length
4	Link Local/Remote Identifiers	8
20	Link Protection Type	2
21	Interface Switching Capability Descriptor	variable

References

Normative References

[GMPLS-ROUTING] Kompella, K., Ed., and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, October 2005.

[GMPLS-RSVP] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

[GMPLS-SIG] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

[IEEE] IEEE, "IEEE Standard for Binary Floating-Point Arithmetic", Standard 754-1985, 1985 (ISBN 1-5593-7653-8).

[ISIS-3way] Katz, D. and R. Saluja, "Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Point-to-Point Adjacencies", RFC 3373, September 2002.

[ISIS-RESTART] Shand, M. and L. Ginsberg, "Restart Signaling for Intermediate System to Intermediate System (IS-IS)", RFC 3847, July 2004.

[ISIS-TE] Smit, H. and T. Li, "Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)", RFC 3784, June 2004.

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

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

Authors' Addresses

Kireeti Kompella
Juniper Networks, Inc.
1194 N. Mathilda Ave
Sunnyvale, CA 94089

EMail: kireeti@juniper.net

Yakov Rekhter
Juniper Networks, Inc.
1194 N. Mathilda Ave
Sunnyvale, CA 94089

EMail: yakov@juniper.net

Full Copyright Statement

Intellectual Property

Acknowledgement

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