Network Working Group
Internet Draft
Expiration Date: November 2002

K. Kompella (Juniper Networks)
Y. Rekhter (Juniper Networks)
A. Banerjee (Calient Networks)
J. Drake (Calient Networks)
G. Bernstein (Ciena)
D. Fedyk (Nortel Networks)
E. Mannie (GTS Network)
D. Saha (Tellium)
V. Sharma (Metanoia, Inc.)

IS-IS Extensions in Support of Generalized MPLS

draft-ietf-isis-gmpls-extensions-12.txt

1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

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/1id-abstracts.txt

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

2. Abstract

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

3. Summary for Sub-IP Area

3.1. Summary

This document specifies encoding of extensions to the IS-IS routing protocol in support of Generalized Multi-Protocol Label Switching (GMPLS). The description of the extensions is specified in [GMPLS-ROUTING].

3.2. Where does it fit in the Picture of the Sub-IP Work

This work fits squarely in either CCAMP or IS-IS boxes.

3.3. Why is it Targeted at this WG

This draft is targeted at either the CCAMP or IS-IS WGs, because this draft specifies the extensions to the IS-IS routing protocols in support of GMPLS, because GMPLS is within the scope of CCAMP WG, and because IS-IS is within the scope of the IS-IS WG.

3.4. Justification

The WG should consider this document as it specifies the extensions to the IS-IS routing protocols in support of GMPLS.

4. Specification of Requirements

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].

5. 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].

6. IS-IS Routing Enhancements

In this section we define the enhancements to the TE properties of GMPLS TE links that can be announced in IS-IS TE LSAs.

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

We further add one new TLV to the TE LSAs.

TLV Type	Length	Name
138	variable	Shared Risk Link Group

6.1. 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 Idenfier (see Section "Support for unnumbered links" of [GMPLS-ROUTING]). If the Link Remote Identifier is unknown, it is set to 0.

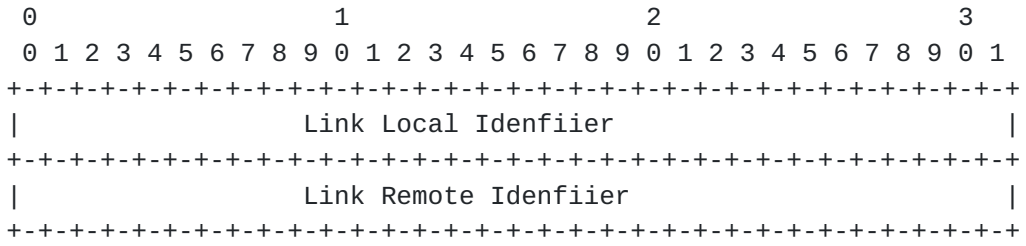

6.2. Link Protection Type

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

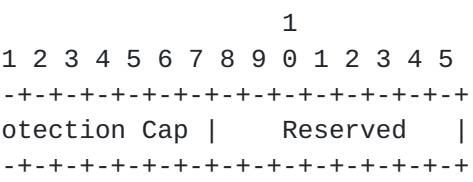

The first octet is a bit vector describing the protection capabilities of the link (see Section "Link Protection Type" of [GMPLS-ROUTING]). They are:

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 may occur at most once within the extended IS reachability TLV.

6.3. 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 format of the value field is as shown below:

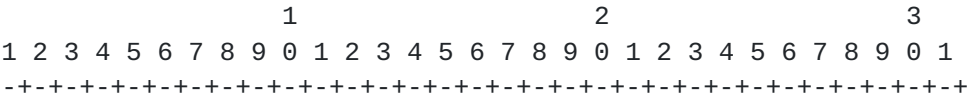

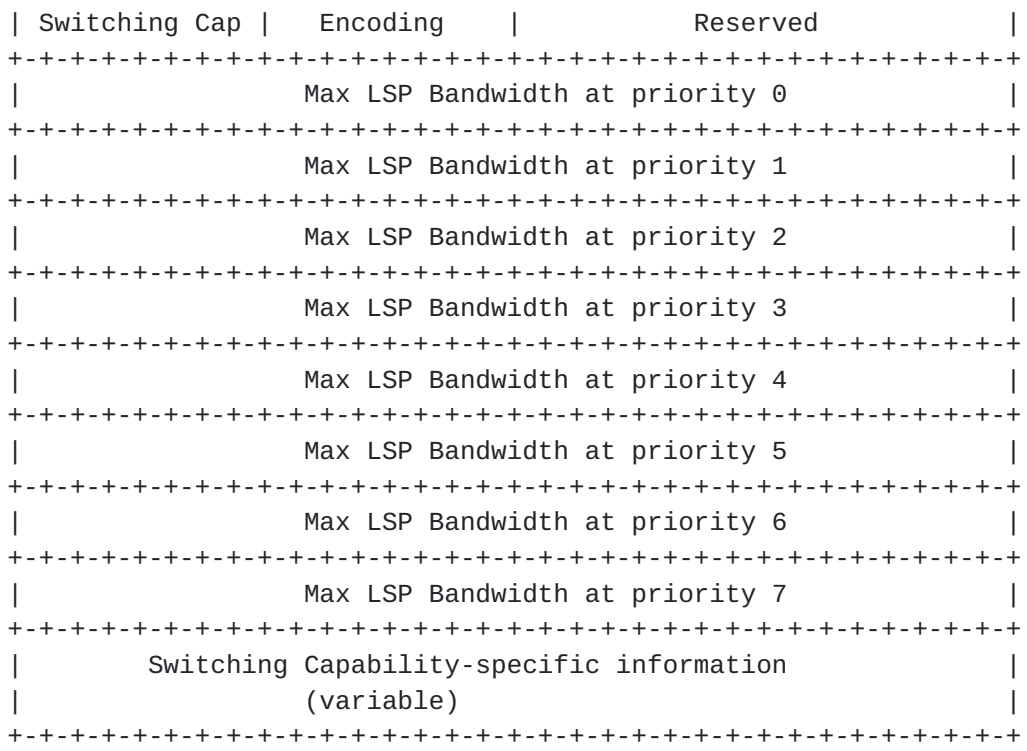

The Switching Capability (Switching Cap) field contains one of the following values:

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, 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.

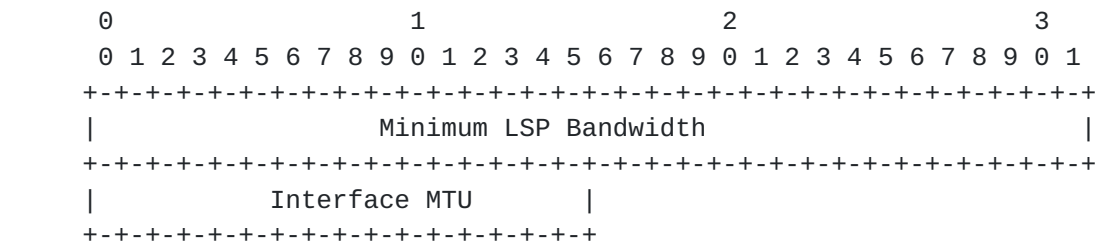

The Minimum LSP Bandwidth is is encoded in a 4 octets field in the IEEE floating point format. The units are bytes (not bits!) per second. The Interface MTU is encoded as a 2 octets integer.

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.

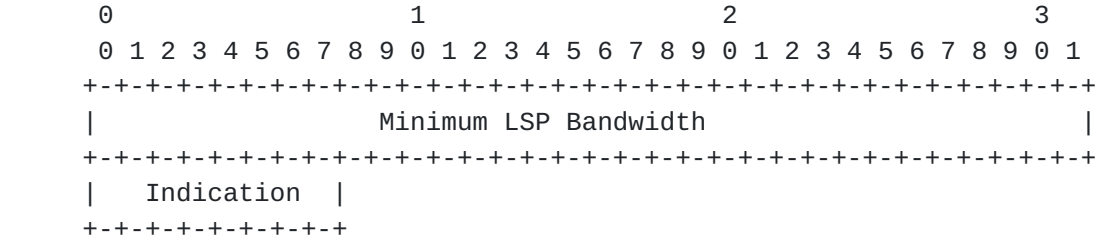

The Minimum LSP Bandwidth is encoded in a 4 octets field in the IEEE floating point format. The units are bytes (not bits!) per second. The indication whether the interface supports Standard or Arbitrary SONET/SDH is encoded as 1 octet. The value of this octet is 0 if the interface supports Standard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

6.4. Shared Risk Link Group TLV

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

7 octets of System ID and 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

and a list of SRLG values, where each element in the list has 4 octets. The length of this TLV is 16 + 4 * (number of SRLG values). 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 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.

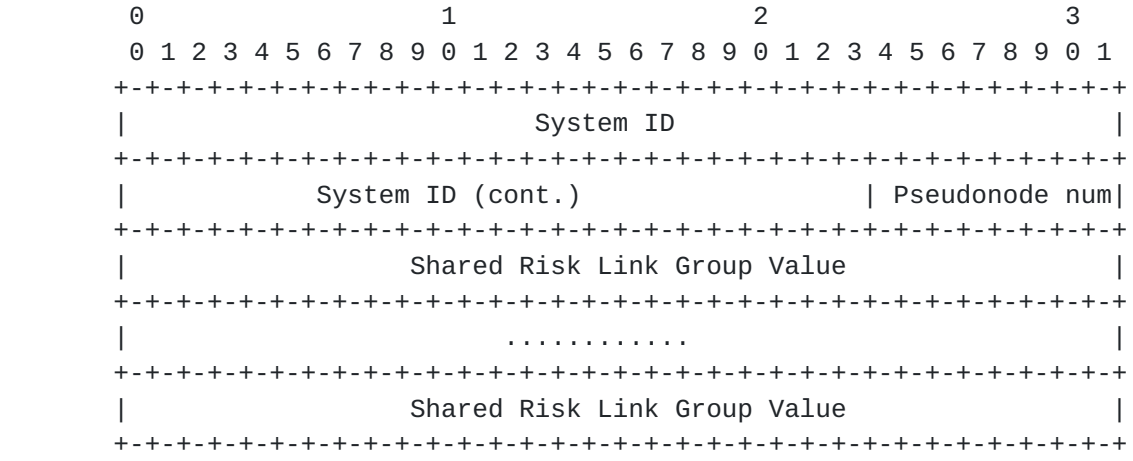

This TLV carries the Shared Risk Link Group information (see Section "Shared Risk Link Group Information" of [GMPLS-ROUTING]).

6.5. Link Identifier for Unnumbered Interfaces

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

7. 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 TE LSAs, these LSAs 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 TE LSAs.

In addition in the case of a planned restart prior to restarting, the restarting node SHOULD originate the TE LSAs 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 advertise the actual unreserved bandwidth on the TE links from the neighbors to that node.

Regular graceful restart should not be aborted if a TE LSA or TE topology changes. TE graceful restart need not be aborted if a TE LSA or TE topology changes.

8. Security Considerations

The extensions proposed in this document does not raise any new security concerns.

9. Acknowledgements

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

10. References

[ISIS-TE] Smit, H., Li, T., "IS-IS Extensions for Traffic Engineering",
draft-ietf-isis-traffic-03.txt (work in progress)

[GMPLS-SIG] Generalized MPLS Group, "Generalized MPLS - Signaling Functional
Description", draft-ietf-mpls-generalized-signaling-04.txt (work in progress)

[GMPLS-ROUTING] "Routing Extensions in Support of Generalized MPLS",
draft-many-ccamp-gmpls-routing-01.txt (work in progress)

[ISIS-3way] "Three-Way Handshake for IS-IS Point-to-Point Adjacencies",
draft-ietf-isis-3way-05.txt (work in progress)

[ISIS-RESTART] "Restart signaling for ISIS", draft-ietf-isis-restart-00.txt
(work in progress)

[GMPLS-RSVP] "Generalized MPLS Signaling - RSVP-TE Extensions",
draft-ietf-mpls-generalized-rsvp-te-06.txt (work in progress)

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

11. Authors' Information

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

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: (408) 972-3720
Email: jdrake@calient.net

Greg Bernstein
Ciena Corporation
10480 Ridgeview Court
Cupertino, CA 94014
Phone: (408) 366-4713
Email: greg@ciena.com

Don Fedyk
Nortel Networks Corp.
600 Technology Park Drive
Billerica, MA 01821
Phone: +1-978-288-4506
Email: dwfedyk@nortelnetworks.com

Eric Mannie
GTS Network Services
RDI Department, Core Network Technology Group
Terhulpsesteenweg, 6A
1560 Hoeilaart, Belgium
Phone: +32-2-658.56.52
E-mail: eric.mannie@gtsgroup.com

Debanjan Saha
Tellium Optical Systems
2 Crescent Place
P.O. Box 901
Ocean Port, NJ 07757
Phone: (732) 923-4264
Email: dsaha@tellium.com

Vishal Sharma
Metanoia, Inc.
335 Elan Village Lane, Unit 203
San Jose, CA 95134-2539
Phone: +1 408-943-1794
Email: v.sharma@ieee.org