

Internet Engineering Task Force
INTERNET-DRAFT
Expires March 2011
Intended status: Proposed Standard

J. Harrison
J. Berger
M. Bartlett
Metaswitch Networks
September 30, 2010

IPv6 Traffic Engineering in IS-IS
<[draft-ietf-isis-ipv6-te-08.txt](#)>

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [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/1id-abstracts.txt>.

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

This Internet-Draft will expire on March 30, 2011.

Abstract

This document specifies a method for exchanging IPv6 Traffic Engineering information using the IS-IS routing protocol. This information enables routers in an IS-IS network to calculate traffic engineered routes using IPv6 addresses.

Expires March 2011

[Page 1]

1. Overview

The IS-IS routing protocol is defined in [[IS-IS](#)]. Each router generates a Link State Protocol Data Unit (LSP) that contains information describing the router and the links from the router. The information in the LSP is encoded in a variable length data structure consisting of a Type, Length and Value. Such a data structure is referred to as a TLV.

[TE] and [[GMPLS](#)] define a number of TLVs and sub-TLVs that allow Traffic Engineering (TE) information to be disseminated by the IS-IS protocol [[IS-IS](#)]. The addressing information passed in these TLVs is IPv4 specific.

[IPv6] describes how the IS-IS protocol can be used to carry out Shortest Path First (SPF) routing for IPv6. It does this by defining IPv6 specific TLVs that are analogous to the TLVs used by IS-IS for carrying IPv4 addressing information.

Multi-Protocol Label Switching (MPLS) Traffic Engineering is very successful and, as the use of IPv6 grows, there is a need to be able to support Traffic Engineering in IPv6 networks.

This document defines the TLVs that allow Traffic Engineering information (including Generalized-MPLS (GMPLS) TE information) to be carried in IPv6 IS-IS networks.

2. Requirements Language

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 [[KEYWORDS](#)].

Expires March 2011

[Page 2]

3. Summary of operation

3.1 Identifying IS-IS links using IPv6 addresses

Each IS-IS link has certain properties - bandwidth, shared risk link groups (SRLGs), switching capabilities and so on. The IS-IS extensions defined in [TE] and [GMPLS] describe how to associate these traffic engineering parameters with IS-IS links. These TLVs use IPv4 addresses to identify the link (or local/remote link identifiers on unnumbered links).

When IPv6 is used, a numbered link may be identified by IPv4 and/or IPv6 interface addresses. The type of identifier used does not affect the properties of the link - it still has the same bandwidth, SRLGs, switching capabilities.

This document describes an approach for supporting IPv6 traffic engineering, by defining TLV extensions that allow TE links and nodes to be identified by IPv6 addresses.

3.1.1 IPv6 address types

An IPv6 address can have global, unique-local or link-local scope.

- A link-local IPv6 address is valid only within the scope of a single link, and may only be referenced on that link.
- A unique-local IPv6 address is globally unique, but is intended for local communication.
- A global IPv6 address is valid within the scope of the Internet.

Because the IPv6 traffic engineering TLVs present in LSPs are propagated across networks, they MUST NOT use link-local addresses.

IS-IS does not need to differentiate between global and unique-local addresses.

3.2 IP addresses used in Traffic Engineering TLVs

This section lists the IP addresses used in the TLVs defined in [TE] and [GMPLS], and gives an overview of the required IPv6 equivalents.

Expires March 2011

[Page 3]

3.2.1 TE Router ID TLV

The TE Router ID TLV contains a stable IPv4 address that is routable, regardless of the state of each interface.

Similarly, for IPv6, it is useful to have a stable IPv6 address identifying a TE node. The IPv6 TE Router ID TLV is defined in [section 4.1](#).

3.2.2 IPv4 Interface Address sub-TLV

This sub-TLV of the Extended IS Reachability TLV contains an IPv4 address for the local end of a link. The equivalent IPv6 Interface Address sub-TLV is defined in [section 4.2](#).

3.2.3 IPv4 Neighbor Address sub-TLV

This sub-TLV of the Extended IS Reachability TLV is used for point-to-point links, and contains an IPv4 address for the neighbor's end of a link. The equivalent IPv6 Neighbor Address sub-TLV is defined in [section 4.3](#).

A router constructs the IPv4 TLV Neighbor Address TLV using one of the IPv4 addresses received in the IS-IS Hello (IIH) PDU from the neighbor on the link.

The IPv6 Neighbor Address sub-TLV contains a globally unique IPv6 address for the interface from the peer (which can be either a global or unique-local IPv6 address). The IPv6 Interface Address TLV defined in [\[IPv6\]](#) only contains link-local addresses when used in the IIH PDU. Hence a neighbor's IP address from the the IPv6 Interface Address TLV cannot be used when constructing the IPv6 Neighbor Address sub-TLV. Instead, we define an additional TLV, the IPv6 Global Interface Address TLV in [section 4.5](#). The IPv6 Global Interface Address TLV is included in IIH PDUs to provide the globally unique IPv6 address that a neighbor router needs in order to construct the IPv6 Neighbor Address sub-TLV.

3.2.4 IPv4 SRLG TLV

The SRLG TLV (type 138) defined in [\[GMPLS\]](#) contains the set of SRLGs associated with a link. The SRLG TLV identifies the link using either local/remote IPv4 addresses or, (for point-to-point unnumbered links), link local/remote identifiers. The SRLG TLV includes a flags field to indicate which type of identifier is used.

When only IPv6 is used, IPv4 addresses and link local/remote identifiers are not available to identify the link, but IPv6 addresses can be used instead.

Expires March 2011

[Page 3]

There is no back-compatible way to modify the SRLG TLV (type 138) to identify the link by IPv6 addresses, and therefore we need a new TLV.

The IPv6 SRLG TLV is defined in [section 4.4](#).

4. IPv6 TE TLVs

4.1 IPv6 TE Router ID TLV

The IPv6 Traffic Engineering Router ID TLV is TLV type 140.

The IPv6 TE Router ID TLV contains a 16-octet IPv6 address. A stable global IPv6 address **MUST** be used, so that the router ID provides a routable address, regardless of the state of a node's interfaces.

If a router does not implement traffic engineering, it **MAY** include or omit the IPv6 Traffic Engineering Router ID TLV. If a router implements traffic engineering for IPv6, it **MUST** include this TLV in its LSP. This TLV **MUST NOT** be included more than once in an LSP.

An implementation receiving an IPv6 TE Router ID TLV **MUST NOT** consider the router ID as a /128 reachable prefix in the standard SPF calculation, because this can lead to forwarding loops when interacting with systems that do not support this TLV.

4.2 IPv6 Interface Address sub-TLV

The IPv6 Interface Address sub-TLV of the Extended IS Reachability TLV has sub-TLV type 12. It contains a 16-octet IPv6 address for the interface described by the containing Extended IS Reachability TLV. This sub-TLV can occur multiple times.

Implementations **MUST NOT** inject a /128 prefix for the interface address into their routing or forwarding table, because this can lead to forwarding loops when interacting with systems that do not support this sub-TLV.

If a router implements the basic TLV extensions described in [\[TE\]](#), it **MAY** include or omit this sub-TLV. If a router implements IPv6 traffic engineering, it **MUST** include this sub-TLV (except on an unnumbered point-to-point link, in which case the Link Local Interface Identifiers sub-TLV is used instead).

This sub-TLV **MUST NOT** contain an IPv6 link-local address.

Expires March 2011

[Page 4]

[4.3](#) IPv6 Neighbor Address sub-TLV

The IPv6 Neighbor Address sub-TLV of the Extended IS Reachability TLV has sub-TLV type 13. It contains a 16-octet IPv6 address for a neighboring router on the link described by the (main) TLV. This sub-TLV can occur multiple times.

Implementations MUST NOT inject a /128 prefix for the interface address into their routing or forwarding table, because this can lead to forwarding loops when interacting with systems that do not support this sub-TLV.

If a router implements the basic TLV extensions described in [\[TE\]](#), it MAY include or omit this sub-TLV. If a router implements IPv6 traffic engineering, it MUST include this sub-TLV for point-to-point links (except on an unnumbered point-to-point link, in which case the Link Local Interface Identifiers sub-TLV is used instead).

This sub-TLV MUST NOT contain an IPv6 link-local address.

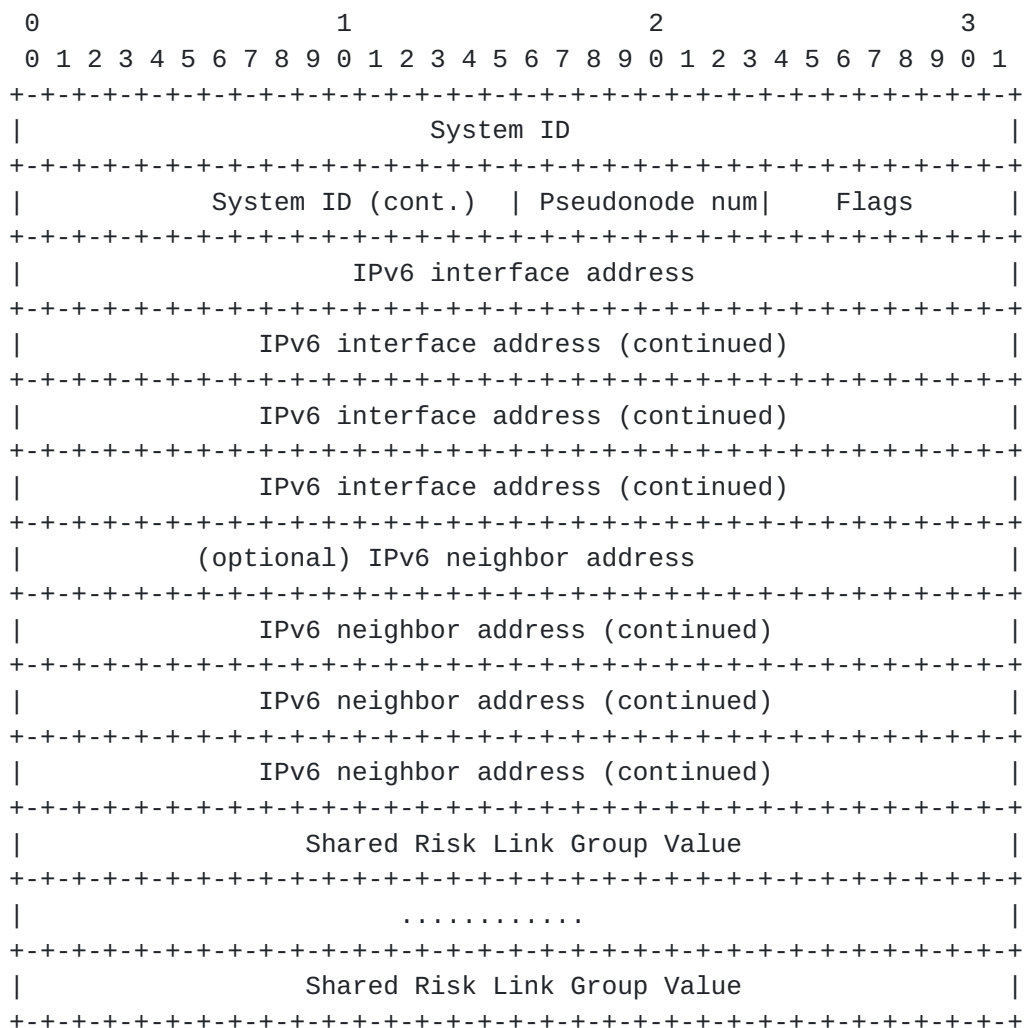
[4.4](#) IPv6 SRLG TLV

The IPv6 SRLG TLV has type 139. The TLV carries the Shared Risk Link Group information (see Section "Shared Risk Link Group Information" of [\[GMPLS-ROUTING\]](#)).

It contains a data structure consisting of:

- 6 octets of System ID
- 1 octet of Pseudonode Number
- 1 octet flags
- 16 octets of IPv6 interface address
- (optional) 16 octets of IPv6 neighbor address
- (variable) list of SRLG values, where each element in the list has 4 octets.


The following illustrates the encoding of the Value field of the IPv6 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 1-octet flags field is interpreted as follows.

Flags (1 octet)



NA - Neighbor Address included.

The flags field currently contains one flag to indicate whether the

IPv6 neighbor address is included (the NA bit is set to 1), or not included (the NA bit is set to 0).

Expires March 2011

[Page 6]

Other bits in the flags field are reserved for future use. Any bits not understood by an implementation MUST be set to zero by the sender. If a router receives an IPv6 SRLG TLV with non-zero values for any bit that it does not understand, it MUST ignore the TLV (in other words, it does not use the TLV locally, but floods the TLV unchanged to neighbors as normal).

Note that this rule for processing the flags octet allows for future extensibility of the IPv6 SRLG TLV. In particular, it allows alternative means of identifying the corresponding link to be added in the future. An implementation that does not understand such an extension will ignore the TLV, rather than attempting to interpret the TLV incorrectly.

The length of this TLV is $24 + 4 * (\text{number of SRLG values}) + 16$ (if the IPv6 neighbor address is included).

To prevent an SRLG TLV and an IPv6 SRLG TLV in the same logical LSP from causing confusion of interpretation, the following rules are applied.

- The IPv6 SRLG TLV MAY occur more than once within the IS-IS logical LSP.
- There MUST NOT be more than one IPv6 SRLG TLV for a given link.
- The IPv6 SRLG TLV (type 139) MUST NOT be used to describe the SRLGs for a given link if it is possible to use the SRLG TLV (type 138).
- If both an SRLG TLV and an IPv6 SRLG are received describing the SRLGs for the same link, the receiver MUST apply the SRLG TLV and ignore the IPv6 SRLG TLV.

In other words, if SRLGs are to be advertised for a link, and if the Extended IS Reachability TLV describing a link contains IPv4 interface/neighbor address sub-TLVs or the link local identifiers sub-TLV, then the SRLGs MUST be advertised in the SRLG TLV (type 138).

4.5 IPv6 Global Interface Address TLV

The IPv6 Global Interface Address TLV is TLV type 233. The TLV structure is identical to that of the IPv6 Interface Address TLV defined in [[IPv6](#)], but the semantics are different. In particular, the TLV is included in IIH PDUs for those interfaces using IPv6 TE extensions. The TLV contains global or unique-local IPv6 addresses assigned to the interface that is sending the Hello.

The IPv6 Global Interface Address TLV is not used in LSPs.

Expires March 2011

[Page 7]

5. Security Considerations

This document raises no new security issues for IS-IS; for general security considerations for IS-IS see [[ISIS-AUTH](#)].

6. IPv4/IPv6 Migration

The IS-IS extensions described in this document allow the routing of GMPLS Label Switched Paths using IPv6 addressing through an IS-IS network. There are no migration issues introduced by the addition of this IPv6 TE routing information into an existing IPv4 GMPLS network. Migration of Label Switched Paths from IPv4 to IPv6 is an issue for GMPLS signaling and is outside the scope of this document.

7. IANA Considerations

This document defines the following new IS-IS TLV types that need to be reflected in the IS-IS TLV code-point registry:

Type	Description	IIH	LSP	SNP
----	-----	---	---	---
139	IPv6 SRLG TLV	n	y	n
140	IPv6 TE Router ID	n	y	n
233	IPv6 Global Interface Address TLV	y	n	n

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

Type	Description	Length
----	-----	-----
12	IPv6 Interface Address	16
13	IPv6 Neighbor Address	16

8. References

8.1 Normative References

- [IS-IS] ISO, "Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", International Standard 10589: 2002, Second Edition, 2002.
- [IPv6] C. Hopps, "Routing IPv6 with IS-IS", [RFC 5308](#),

October 2008.

Expires March 2011

[Page 8]

[TE] H. Smit and T. Li, "IS-IS extensions for Traffic Engineering", [RFC 5305](#), October 2008.

[KEYWORDS]

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

[ISIS-AUTH]

T. Li and R. Atkinson, "IS-IS Cryptographic Authentication", [RFC 5304](#), October 2008.

[GMPLS] K.Kompella and Y.Rekhter, "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching", [RFC 5307](#), October 2008.

[GMPLS-ROUTING]

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

9. Authors' Addresses

Jon Harrison
Metaswitch Networks
100 Church Street
Enfield
EN2 6BQ
U.K.
Phone: +44 20 8366 1177
Email: jon.harrison@metaswitch.com

Jon Berger
Metaswitch Networks
100 Church Street
Enfield
EN2 6BQ
U.K.
Phone: +44 20 8366 1177
Email: jon.berger@metaswitch.com

Mike Bartlett
Metaswitch Networks
100 Church Street
Enfield
EN2 6BQ
U.K.
Phone: +44 20 8366 1177
Email: mike.bartlett@metaswitch.com

10. Full Copyright Statement

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

ALL IETF Documents 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, THE IETF TRUST 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.

Intellectual Property

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Expires March 2011

[Page 10]