

Network Working Group
Request for Comments: 5329
Category: Standards Track

K. Ishiguro
V. Manral
IP Infusion, Inc
A. Davey
Data Connection Limited
A. Lindem, Ed.
Redback Networks
September 2008

Traffic Engineering Extensions to OSPF Version 3

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Abstract

This document describes extensions to OSPFv3 to support intra-area Traffic Engineering (TE). This document extends OSPFv2 TE to handle IPv6 networks. A new TLV and several new sub-TLVs are defined to support IPv6 networks.

Table of Contents

1. Introduction	2
1.1. Requirements Notation	2
2. Intra-Area-TE-LSA	3
2.1. Intra-Area-TE-LSA Payload	4
3. Router IPv6 Address TLV	4
4. Link TLV	5
4.1. Link ID Sub-TLV	6
4.2. Neighbor ID Sub-TLV	6
4.3. Local Interface IPv6 Address Sub-TLV	6
4.4. Remote Interface IPv6 Address Sub-TLV	7
5. Security Considerations	8
6. Management Considerations	8
7. IANA Considerations	9
8. References	9
8.1. Normative References	9
8.2. Informative References	10
Acknowledgments	10

1. Introduction

OSPFv3 has a very flexible mechanism for adding new LS types. Unknown LS types are flooded properly based on the flooding scope bits in the LS type [OSPFV3]. This document defines the Intra-Area-TE-LSA to OSPFv3.

For Traffic Engineering, this document uses "Traffic Engineering Extensions to OSPF" [TE] as a base for TLV definitions. New TLVs and sub-TLVs are added to [TE] to extend TE capabilities to IPv6 networks. Some existing TLVs and sub-TLVs require clarification for OSPFv3 applicability.

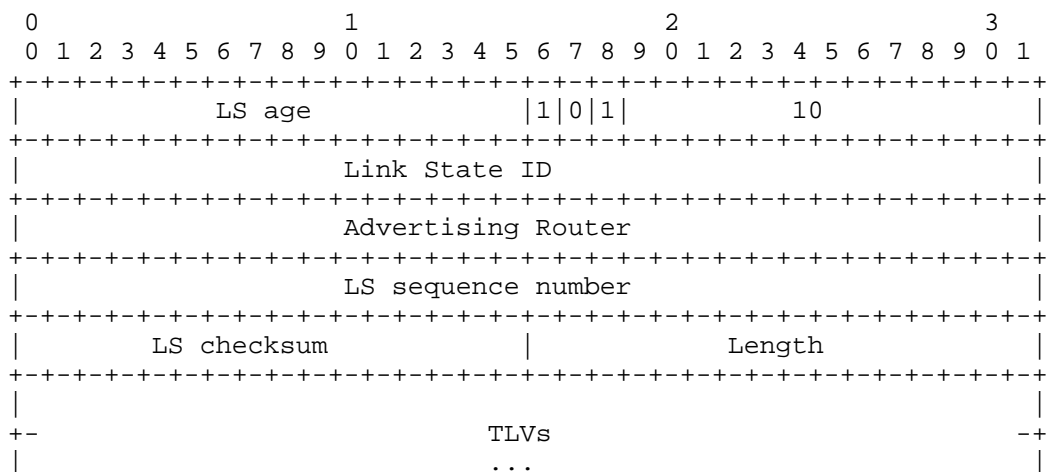
GMPLS [GMPLS] and the Diff-Serv MPLS extensions [TE-DIFF] are based on [TE]. These functions can also be extended to OSPFv3 by utilizing the TLVs and sub-TLVs described in this document.

1.1. Requirements Notation

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 [RFC-KEYWORDS].

2. Intra-Area-TE-LSA

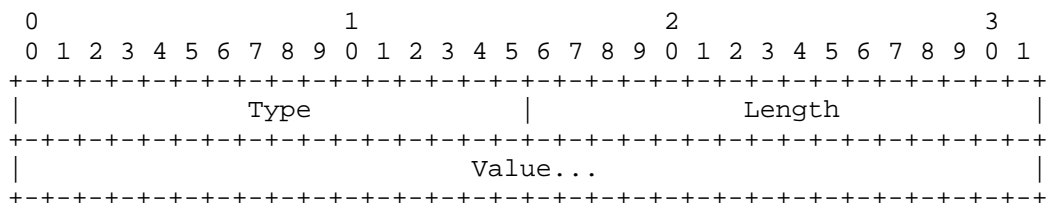
A new LS type is defined for the Intra-Area-TE-LSA. This is different from OSPFv2 Traffic Engineering [TE] where opaque LSAs are used to advertise TE information [OPAQUE]. The LSA function code is 10, the U-bit is set, and the scope is set to 01 for area-scoping. When the U-bit is set to 1, an OSPFv3 router must flood the LSA at its defined flooding scope even if it does not recognize the LS type [OSPFV3].



OSPFv3 Intra-Area-TE-LSA

The Link State ID of an Intra-Area-TE-LSA is an arbitrary value used to maintain multiple Traffic Engineering LSAs. The Link State ID has no topological significance.

The format of the TLVs within the body of an Intra-Area-TE-LSA is the same as the format used by the Traffic Engineering extensions to OSPF [TE]. The LSA payload consists of one or more nested Type/Length/Value (TLV) triplets. The format of each TLV is:



TLV Format

The Length field defines the length of the value portion in octets (thus, a TLV with no value portion would have a length of 0). The TLV is padded to 4-octet alignment; padding is not included in the Length field (so a 3-octet value would have a length of 3, but the total size of the TLV would be 8 octets). Nested TLVs are also 32-bit aligned. For example, a 1-byte value would have the Length field set to 1, and 3 octets of padding would be added to the end of the value portion of the TLV. Unrecognized types are ignored.

2.1. Intra-Area-TE-LSA Payload

An Intra-Area-TE-LSA contains one top-level TLV. There are two applicable top-level TLVs:

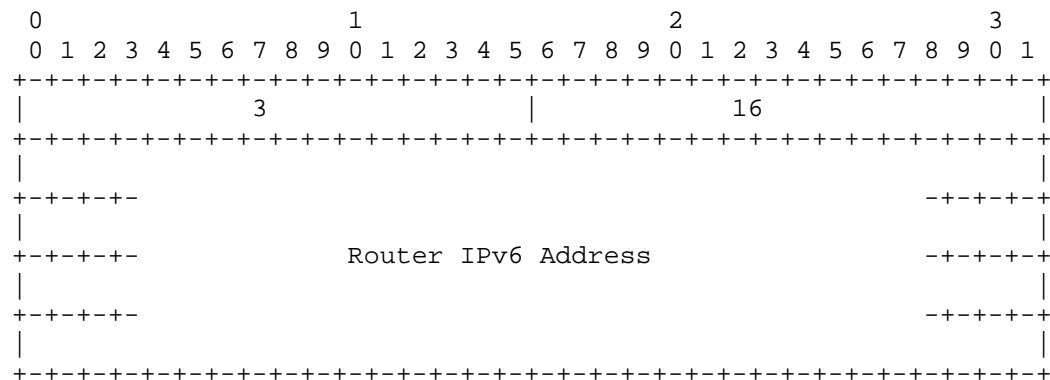
2 - Link TLV

3 - Router IPv6 Address TLV

3. Router IPv6 Address TLV

The Router IPv6 Address TLV advertises a reachable IPv6 address. This is a stable IPv6 address that SHOULD be reachable if there is connectivity to the OSPFv3 router.

The Router IPv6 Address TLV has type 3, length 16, and a value containing a 16-octet local IPv6 address. A link-local address MUST NOT be specified for this TLV. It MUST appear in exactly one Traffic Engineering LSA originated by an OSPFv3 router supporting the TE extensions. The Router IPv6 Address TLV is a top-level TLV as defined in "Traffic Engineering Extensions to OSPF" [TE], and only one top-level TLV may be contained in an LSA.



Type A 16-bit field set to 3.

Length A 16-bit field that indicates the length of the value portion in octets. For this TLV, it is always 16.

Value A stable and routable IPv6 address.

Router IPv6 Address TLV

4. Link TLV

The Link TLV describes a single link and consists of a set of sub-TLVs [TE]. All of the sub-TLVs in [TE] other than the Link ID sub-TLV are applicable to OSPFv3. The Link ID sub-TLV can't be used in OSPFv3 since it is defined to use the OSPFv2 identification for the Designated Router (DR) on multi-access networks. In OSPFv2, neighbors on point-to-point networks and virtual links are identified by their Router IDs while neighbors on broadcast, Non-Broadcast Multi-Access (NBMA), and Point-to-Multipoint links are identified by their IPv4 interface addresses (refer to section 8.2 in [OSPFV2]). The IPv4 interface address is not known to OSPFv3. In contrast to OSPFv2, OSPFv3 always identifies neighboring routers by their Router IDs (refer to section 2.11 in [OSPFV3]).

Three new sub-TLVs for the Link TLV are defined:

- 18 - Neighbor ID (8 octets)
- 19 - Local Interface IPv6 Address (16N octets, where N is the number of IPv6 addresses)
- 20 - Remote Interface IPv6 Address (16N octets, where N is the number of IPv6 addresses)

The Neighbor ID sub-TLV is mandatory for OSPFv3 Traffic Engineering support. It MUST appear exactly once in a Link TLV. All other sub-TLVs defined in this document SHOULD NOT occur more than once in a Link TLV. If a sub-TLV is specified more than once, instances subsequent to the first are ignored.


4.1. Link ID Sub-TLV

The Link ID sub-TLV is used in OSPFv2 to identify the other end of the link. In OSPFv3, the Neighbor ID sub-TLV MUST be used for link identification. In OSPFv3, the Link ID sub-TLV SHOULD NOT be sent and MUST be ignored upon receipt.

4.2. Neighbor ID Sub-TLV

In OSPFv2, the Link ID is used to identify the other end of a link. In OSPFv3, the combination of Neighbor Interface ID and Neighbor Router ID is used for neighbor link identification. Both are advertised in the Neighbor ID sub-TLV.

Neighbor Interface ID and Neighbor Router ID values are the same as described in RFC 5340 [OSPFV3], A.4.3 Router-LSAs.



Type A 16-bit field set to 18.

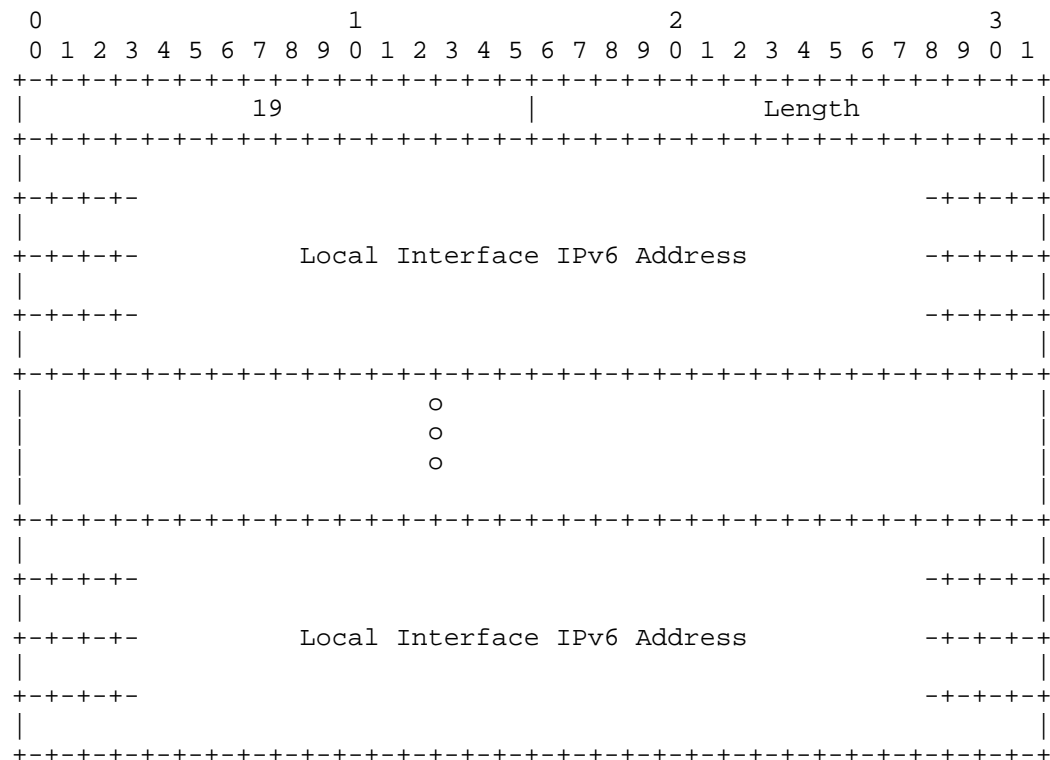
Length A 16-bit field that indicates the length of the value portion in octets. For this sub-TLV, it is always 8.

Value The neighbor's Interface ID and Router ID.

Neighbor ID Sub-TLV

4.3. Local Interface IPv6 Address Sub-TLV

The Local Interface IPv6 Address sub-TLV specifies the IPv6 address(es) of the interface corresponding to this link. If there are multiple local addresses assigned to the link, then they MAY all be listed in this sub-TLV. Link-local addresses MUST NOT be included in this sub-TLV.



Type A 16-bit field set to 19.

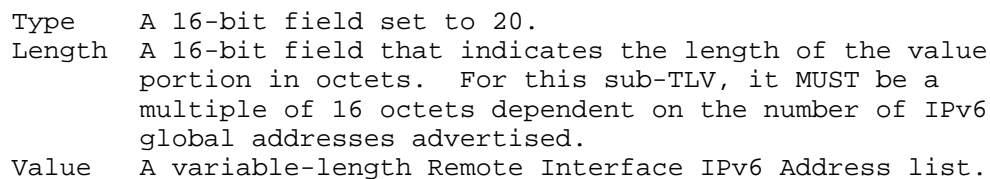
Length A 16-bit field that indicates the length of the value portion in octets. For this sub-TLV, it MUST always be a multiple of 16 octets dependent on the number of IPv6 global addresses advertised.

Value	A list of one or more local IPv6 interface addresses each consuming 16 octets.
-------	--------------------------------------------------------------------------------

Local Interface IPv6 Address Sub-TLV

4.4. Remote Interface IPv6 Address Sub-TLV

The Remote Interface IPv6 Address sub-TLV advertises the IPv6 address(es) associated with the neighbor's interface. This sub-TLV and the Local Interface IPv6 Address sub-TLV are used to discern amongst parallel links between OSPFv3 routers. If the link type is multi-access, the Remote Interface IPv6 Address MAY be set to ::. Alternately, an implementation MAY choose not to send this sub-TLV. Link-local addresses MUST NOT be advertised in this sub-TLV. Neighbor addresses advertised in link-LSAs with a prefix length of 128 and the LA-bit set MAY be advertised.



objects for configuration operations and statistics are defined in [OSPFV3-MIB], and an implementation of the extensions defined in this document SHOULD provide for the appropriate hooks or instrumentation that allow for the MIB objects to be implemented.

The following MIB variables have been added to the OSPFv3 MIB in support of TE:

ospfv3AreaTEEnabled

This TruthValue MIB variable in the ospfv3AreaEntry table entry indicates whether or not OSPFv3 TE advertisement for OSPFv3 interfaces is enabled for the corresponding area. The default value is FALSE.

ospfv3IfTEDisabled

This TruthValue MIB variable in the ospfv3IfEntry table entry indicates whether or not OSPFv3 TE advertisement for OSPFv3 for the corresponding interface is disabled. This MIB variable is only applicable if ospfv3AreaTEEnabled is TRUE for the interface's area. The default value is FALSE.

7. IANA Considerations

The following IANA assignments have been made from existing registries:

1. The OSPFv3 LSA type function code 10 has been assigned to the OSPFv3 Intra-Area-TE-LSA.
2. The Router IPv6 Address TLV type 3 has been assigned from the existing registry for OSPF TE TLVs.
3. The Neighbor ID (18), Local Interface IPv6 Address (19), and Remote Interface IPv6 Address (20) sub-TLVs have been assigned from the existing registry for OSPF TE sub-TLVs.

8. References

8.1. Normative References

- [OSPFV2] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [OSPFV3] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.
- [RFC-KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

- [TE] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.

8.2. Informative References

- [GMPLS] Kompella, K., Ed., and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, October 2005.
- [OPAQUE] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", RFC 5250, July 2008.
- [OSPFV3-MIB] Joyal, D. and V. Manral, "Management Information Base for OSPFv3", Work in Progress, September 2007.
- [TE-DIFF] Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen, P., Krishnan, R., Cheval, P., and J. Heinanen, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", RFC 3270, May 2002.

Acknowledgments

Thanks to Kireeti Kompella, Alex Zinin, Adrian Farrell, and Mach Chen for their comments.

Thanks to Vijay K. Gurbani for providing the General Area Review Team (Gen-ART) review.

Thanks to Rob Austein for providing the Security Directorate (secdir) review.

Thanks to Dan Romascanu for providing the text for the "Management Considerations" section in the context of the IESG review.

Thanks to Dave Ward, Tim Polk, Jari Arkko, and Pasi Eronen for comments and relevant discussion in the context of the IESG review.

The RFC text was produced using Marshall Rose's xml2rfc tool.

Authors' Addresses

Kunihiro Ishiguro
IP Infusion, Inc.
1188 East Arques Avenue,
Sunnyvale, CA 94085
USA

EMail: kunihiro@ipinfusion.com

Vishwas Manral
IP Infusion, Inc
#41, Ground Floor, 5th Cross Road
8th Main Road
Vasanth Nagar, Bangalore 560052
India

EMail: vishwas@ipinfusion.com

Alan Davey
Data Connection Limited
100 Church Street
Enfield
EN2 6BQ
UK

EMail: Alan.Davey@dataconnection.com

Acee Lindem
Redback Networks
102 Carric Bend Court
Cary, NC 27519
USA

EMail: acee@redback.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document 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

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.

