

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 7, 2017

P. Psenak
A. Lindem
Cisco Systems
W. Henderickx
Nokia
J. Tantsura
H. Gredler
Individual
July 6, 2016

OSPFv2 Link Traffic Engineering (TE) Attribute Reuse
draft-ppsenak-ospf-te-link-attr-reuse-02.txt

Abstract

Various link attributes have been defined in OSPFv2 in the context of the MPLS Traffic Engineering (TE) and GMPLS. Many of these link attributes can be used for purposes other than MPLS Traffic Engineering or GMPLS. This document defines how to distribute such attributes in OSPFv2 for applications other than MPLS Traffic Engineering or GMPLS purposes.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on January 7, 2017.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) 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.

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.

Table of Contents

1.	Introduction	2
1.1.	Requirements notation	3
2.	Link attributes examples	3
3.	Advertising Link Attributes	3
3.1.	TE Opaque LSA	3
3.2.	Extended Link Opaque LSA	4
3.3.	Proposed solution	5
4.	Reused TE link attributes	5
4.1.	Remote interface IP address	5
4.2.	Link Local/Remote Identifiers	6
4.3.	Shared Risk Link Group (SRLG)	6
4.4.	Extended Metrics	6
5.	Backward Compatibility	7
6.	Security Considerations	7
7.	IANA Considerations	8
8.	Acknowledgments	8
9.	References	8
9.1.	Normative References	8
9.2.	Informative References	10
	Authors' Addresses	10

[1.](#) Introduction

Various link attributes have been defined in OSPFv2 [[RFC2328](#)] in the context of the MPLS traffic engineering and GMPLS. All these attributes are distributed by OSPFv2 as a sub-TLVs of the Link-TLV advertised in the OSPFv2 TE Opaque LSA [[RFC3630](#)].

Many of these link attributes are useful outside of the traditional MPLS Traffic Engineering or GMPLS. This brings its own set of problems, in particular how to distribute these link attributes in OSPFv2 when MPLS TE or GMPLS are not deployed or are deployed in parallel with other applications that use these link attributes.

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 [[RFC2119](#)].

2. Link attributes examples

This section lists some of the link attributes originally defined for MPLS Traffic Engineering that can be used for other purposes in OSPFv2. The list doesn't necessarily contain all the required attributes.

1. Remote Interface IP address [[RFC3630](#)] - OSPFv2 currently cannot distinguish between parallel set of links between two remote OSPFv2 routers. As a result, the two-way connectivity check performed during SPF may succeed when the two routers disagree on which of the links to use for data traffic.
2. Link Local/Remote Identifiers - [[RFC4203](#)] - Used for the two-way connectivity check for parallel unnumbered links. Also used for identifying adjacencies for unnumbered links in Segment Routing traffic engineering.
3. Shared Risk Link Group (SRLG) [[RFC4203](#)] - In IPFRR, the SRLG is used to compute diverse backup paths [[RFC5714](#)].
4. Unidirectional Link Delay/Loss Metrics [[RFC7471](#)] - Could be used for the shortest path first (SPF) computation using alternate metrics within an OSPF area.

3. Advertising Link Attributes

This section outlines possible approaches for advertising link attributes originally defined for MPLS Traffic Engineering purposes or GMPLS when they are used for other applications.

3.1. TE Opaque LSA

One approach for advertising link attributes is to continue to use TE Opaque LSA ([[RFC3630](#)]). There are several problems with this approach:

1. Whenever the link is advertised in a TE Opaque LSA, the link becomes a part of the TE topology, which may not match IP routed topology. By making the link part of the TE topology, remote nodes may mistakenly believe that the link is available for MPLS TE or GMPLS, when, in fact, MPLS is not enabled on the link.
2. The TE Opaque LSA carries link attributes that are not used or required by MPLS TE or GMPLS. There is no mechanism in TE Opaque LSA to indicate which of the link attributes should be passed to MPLS TE application and which should be used by OSPFv2 and other applications.
3. Link attributes used for non-TE purposes is partitioned across multiple LSAs - the TE Opaque LSA and the Extended Link Opaque LSA. This partitioning will require implementations to lookup multiple LSAs to extract link attributes for a single link, bringing needless complexity to the OSPFv2 implementations.

The advantage of this approach is that there is no additional standardization requirement to advertise the TE/GMPL attributes for other applications. Additionally, link attributes are only advertised once when both OSPF TE and other applications are deployed on the same link. This is not expected to be a common deployment scenario.

3.2. Extended Link Opaque LSA

An alternative approach for advertising link attributes is to use Extended Link Opaque LSAs as defined in [[RFC7684](#)]. This LSA was defined as a generic container for distribution of the extended link attributes. There are several advantages in using Extended Link LSA:

1. Advertisement of the link attributes does not make the link part of the TE topology. It avoids any conflicts and is fully compatible with the [[RFC3630](#)].
2. The TE Opaque LSA remains truly opaque to OSPFv2 as originally defined in [[RFC3630](#)]. Its content is not inspected by OSPFv2 and OSPFv2 acts as a pure transport.
3. There is clear distinction between link attributes used by TE and link attributes used by other OSPFv2 applications.
4. All link attributes that are used by OSPFv2 applications are advertised in a single LSA, the Extended Link Opaque LSA.

The disadvantage of this approach is that in rare cases, the same link attribute is advertised in both the TE Opaque and Extended Link

Attribute LSAs. Additionally, there will be additional standardization effort. However, this could also be viewed as an advantage as the non-TE use cases for the TE link attributes are documented and validated by the OSPF working group.

3.3. Proposed solution

It is RECOMMENDED to use the Extended Link Opaque LSA ([[RFC7684](#)] to advertise any link attributes used for non-TE purposes in OSPFv2, including those that have been originally defined for TE purposes. TE link attributes used for TE purposes continue to use TE Opaque LSA ([[RFC3630](#)]).

It is also RECOMMENDED to keep the format of the link attribute TLVs that have been defined for TE purposes unchanged even when they are used for non-TE purposes.

Finally, it is RECOMMENDED to allocate unique code points for link attribute TLVs that have been defined for TE purposes for the OSPFv2 Extended Link TLV Sub-TLV Registry as defined in [[RFC7684](#)]. For each reused TLV, the code point will be defined in an IETF document along with the expected usecase(s).

4. Reused TE link attributes

This section defines the use case and code points for the OSPFv2 Extended Link TLV Sub-TLV Registry for some of the link attributes that have been originally defined for TE or GMPLS purposes.

4.1. Remote interface IP address

The OSPFv2 description of an IP numbered point-to-point adjacency does not include remote IP address. As described in [Section 2](#), this makes the two-way connectivity check ambiguous in the presence of the parallel point-to-point links between two OSPFv2 routers.

The Remote IP address of the link can also be used for Segment Routing traffic engineering to identify the link in a set of parallel links between two OSPFv2 routers [[I-D.ietf-ospf-segment-routing-extensions](#)]. Similarly, the remote IP address is useful in identifying individual parallel OSPF links advertised in BGP Link-State as described in [[I-D.ietf-idr-ls-distribution](#)].

To advertise the Remote interface IP address in the OSPFv2 Extended Link TLV, the same format of the sub-TLV as defined in [section 2.5.4. of \[RFC3630\]](#) is used and TLV type 4 is used.

4.2. Link Local/Remote Identifiers

The OSPFv2 description of an IP unnumbered point-to-point adjacency does not include remote link identifier. As described in [Section 2](#), this makes the two-way connectivity check ambiguous in the presence of the parallel point-to-point IP unnumbered links between two OSPFv2 routers.

The local and remote link identifiers can also be used for Segment Routing traffic engineering to identify the link in a set of parallel IP unnumbered links between two OSPFv2 routers [[I-D.ietf-ospf-segment-routing-extensions](#)]. Similarly, these identifiers are useful in identifying individual parallel OSPF links advertised in BGP Link-State as described in [[I-D.ietf-idr-ls-distribution](#)].

To advertise the link Local/Remote identifiers in the OSPFv2 Extended Link TLV, the same format of the sub-TLV as defined in [section 1.1. of \[RFC4203\]](#) is used and TLV type 5 is used.

4.3. Shared Risk Link Group (SRLG)

The SRLG of a link can be used in IPFRR to compute a backup path that does not share any SRLG group with the protected link.

To advertise the SRLG of the link in the OSPFv2 Extended Link TLV, the same format of the sub-TLV as defined in [section 1.3. of \[RFC4203\]](#) is used and TLV type 6 is used.

4.4. Extended Metrics

[RFC3630] defines several link bandwidth types. [[RFC7471](#)] defines extended link metrics that are based on link bandwidth, delay and loss characteristics. All these can be used to compute best paths within an OSPF area to satisfy requirements for bandwidth, delay (nominal or worst case) or loss.

To advertise extended link metrics in the OSPFv2 Extended Link TLV, the same format of the sub-TLVs as defined in [[RFC7471](#)] is used with following TLV types:

- 7 - Unidirectional Link Delay
- 8 - Min/Max Unidirectional Link Delay
- 9 - Unidirectional Delay Variation
- 10 - Unidirectional Link Loss

- 11 - Unidirectional Residual Bandwidth
- 12 - Unidirectional Available Bandwidth
- 13 - Unidirectional Utilized Bandwidth

To advertise link maximum bandwidth, maximum reservable bandwidth and unreserved bandwidth in the OSPFv2 Extended Link TLV, the same format of the sub-TLVs as defined in [[RFC3630](#)] is used with following TLV types:

- 7 - Maximum bandwidth
- 8 - Maximum reservable bandwidth
- 9 - Unreserved bandwidth

5. Backward Compatibility

Link attributes may be concurrently advertised in both the TE Opaque LSA [[RFC3630](#)] and the Extended Link Opaque LSA [[RFC7684](#)].

In fact, there is at least one OSPF implementation that utilizes the link attributes advertised in TE Opaque LSAs [[RFC3630](#)] for Non-RSVP TE applications. For example, this implementation of LFA and remote LFA utilizes links attributes such as Shared Risk Link Groups (SRLG) [[RFC4203](#)] and Admin Group [[RFC3630](#)] advertised in TE Opaque LSAs. These applications are described in [[RFC5286](#)], [[RFC7490](#)], [[I-D.ietf-rtgwg-lfa-manageability](#)] and [[I-D.psarkar-rtgwg-rlfa-node-protection](#)].

When an OSPF routing domain includes routers using link attributes from TE Opaque LSAs for Non-RSVP TE applications such as LFA, OSPF routers in that domain should continue to advertise such TE Opaque LSAs. If there are also OSPF routers using the link attributes described herein for Non-RSVP applications, OSPF routers in the routing domain will also need to advertise these attributes in OSPF Extended Link Attributes LSAs [[RFC7684](#)]. In such a deployment, the advertised attributes SHOULD be the same and Non-RSVP application access to link attributes is a matter of local policy.

6. Security Considerations

Implementations must assure that malformed TLV and Sub-TLV permutations do not result in errors that cause hard OSPFv2 failures.

7. IANA Considerations

This specification updates the OSPFv2 Extended Link TLV sub-TLV registry that is defined in [[RFC7684](#)] with the following TLV types:

- 4 - Remote interface IP address
- 5 - Link Local/Remote Identifiers
- 6 - Shared Risk Link Group
- 7 - Unidirectional Link Delay
- 8 - Min/Max Unidirectional Link Delay
- 9 - Unidirectional Delay Variation
- 10 - Unidirectional Link Loss
- 11 - Unidirectional Residual Bandwidth
- 12 - Unidirectional Available Bandwidth
- 13 - Unidirectional Utilized Bandwidth

8. Acknowledgments

Thanks to Chris Bowers for his review and comments.

9. References

9.1. Normative References

- [I-D.ietf-idr-ls-distribution]
Gredler, H., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and TE Information using BGP", [draft-ietf-idr-ls-distribution-13](#) (work in progress), October 2015.
- [I-D.ietf-ospf-segment-routing-extensions]
Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", [draft-ietf-ospf-segment-routing-extensions-08](#) (work in progress), April 2016.

- [I-D.ietf-rtgwg-lfa-manageability]
Litkowski, S., Decraene, B., Filsfils, C., Raza, K., and M. Horneffer, "Operational management of Loop Free Alternates", [draft-ietf-rtgwg-lfa-manageability-11](#) (work in progress), June 2015.
- [I-D.psarkar-rtgwg-rlfa-node-protection]
psarkar@juniper.net, p., Gredler, H., Hegde, S., Bowers, C., Litkowski, S., and H. Raghuvier, "Remote-LFA Node Protection and Manageability", [draft-psarkar-rtgwg-rlfa-node-protection-05](#) (work in progress), June 2014.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), DOI 10.17487/RFC2328, April 1998, <<http://www.rfc-editor.org/info/rfc2328>>.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), DOI 10.17487/RFC3630, September 2003, <<http://www.rfc-editor.org/info/rfc3630>>.
- [RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", [RFC 5250](#), DOI 10.17487/RFC5250, July 2008, <<http://www.rfc-editor.org/info/rfc5250>>.
- [RFC5286] Atlas, A., Ed. and A. Zinin, Ed., "Basic Specification for IP Fast Reroute: Loop-Free Alternates", [RFC 5286](#), DOI 10.17487/RFC5286, September 2008, <<http://www.rfc-editor.org/info/rfc5286>>.
- [RFC5714] Shand, M. and S. Bryant, "IP Fast Reroute Framework", [RFC 5714](#), DOI 10.17487/RFC5714, January 2010, <<http://www.rfc-editor.org/info/rfc5714>>.
- [RFC7490] Bryant, S., Filsfils, C., Previdi, S., Shand, M., and N. So, "Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)", [RFC 7490](#), DOI 10.17487/RFC7490, April 2015, <<http://www.rfc-editor.org/info/rfc7490>>.
- [RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", [RFC 7684](#), DOI 10.17487/RFC7684, November 2015, <<http://www.rfc-editor.org/info/rfc7684>>.

9.2. Informative References

- [RFC4203] Kompella, K., Ed. and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 4203](#), DOI 10.17487/RFC4203, October 2005, <<http://www.rfc-editor.org/info/rfc4203>>.
- [RFC7471] Giacalone, S., Ward, D., Drake, J., Atlas, A., and S. Previdi, "OSPF Traffic Engineering (TE) Metric Extensions", [RFC 7471](#), DOI 10.17487/RFC7471, March 2015, <<http://www.rfc-editor.org/info/rfc7471>>.

Authors' Addresses

Peter Psenak
Cisco Systems
Apollo Business Center
Mlynske nivy 43
Bratislava, 821 09
Slovakia

Email: ppsenak@cisco.com

Acee Lindem
Cisco Systems
301 Midenhall Way
Cary, NC 27513
USA

Email: acee@cisco.com

Wim Henderickx
Nokia
Copernicuslaan 50
Antwerp, 2018 94089
Belgium

Email: wim.henderickx@nokia.com

Jeff Tantsura
Individual
USA

Email: jefftant.ietf@gmail.com

Hannes Gredler
Individual
Austria

Email: hannes@gredler.at