

Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: December 1, 2019

K. Talaulikar
P. Psenak
Cisco Systems
J. Tantsura
Apstra
May 30, 2019

**Application Specific Attributes Advertisement with BGP Link-State
draft-ietf-idr-bgp-ls-app-specific-attr-00**

Abstract

Various link attributes have been defined in link-state routing protocols like OSPF and IS-IS in the context of the MPLS Traffic Engineering (TE) and GMPLS. BGP Link-State (BGP-LS) extensions have been defined to distribute these attributes along with other topology information from these link-state routing protocols. Many of these link attributes can be used for applications other than MPLS TE or GMPLS.

Extensions to link-state routing protocols have been defined for such link attributes which enable distribution of their application specific values. This document defines extensions to BGP-LS address-family to enable advertisement of these application specific attributes as a part of the topology information from the network.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

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 <https://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 December 1, 2019.

Copyright Notice

Copyright (c) 2019 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 (<https://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.

Table of Contents

1.	Introduction	2
2.	Application Specific Link Attributes TLV	3
3.	Application Specific Link Attributes	5
4.	Procedures	7
5.	Backward Compatibility	8
6.	IANA Considerations	10
7.	Manageability Considerations	10
	7.1. Operational Considerations	10
	7.2. Management Considerations	10
8.	Security Considerations	10
9.	Acknowledgements	11
10.	References	11
	10.1. Normative References	11
	10.2. Informative References	11
	Authors' Addresses	12

[1.](#) Introduction

Various link attributes have been defined in link-state routing protocols (viz. IS-IS [[RFC1195](#)], OSPFv2 [[RFC2328](#)] and OSPFv3 [[RFC5340](#)]) in the context of the MPLS traffic engineering and GMPLS. All these attributes are distributed by these protocols using TLVs that were originally defined for traditional MPLS Traffic Engineering (i.e. using RSVP-TE [[RFC3209](#)]) or GMPLS [[RFC4202](#)] applications.

In recent years new applications have been introduced which have use cases for many of the link attributes historically used by RSVP-TE and GMPLS. Such applications include Segment Routing (SR) [[RFC8402](#)] and Loop Free Alternates (LFA) [[RFC5286](#)]. This has introduced ambiguity in that if a deployment includes a mix of RSVP-TE support and SR support (for example) it is not possible to unambiguously indicate which advertisements are to be used by RSVP-TE and which advertisements are to be used by SR. If the topologies are fully congruent this may not be an issue, but any incongruence leads to ambiguity. An additional issue arises in cases where both applications are supported on a link but the link attribute values associated with each application differ. Current advertisements do not support advertising application specific values for the same attribute on a specific link. IGP Flexible Algorithm [[I-D.ietf-lsr-flex-algo](#)] is one such application use-case that MAY use application specific link attributes.

[[I-D.ietf-ospf-te-link-attr-reuse](#)] and [[I-D.ietf-isis-te-app](#)] define extensions for OSPF and IS-IS respectively which address these issues. Also, as evolution of use cases for link attributes can be expected to continue in the years to come, these documents define a solution which is easily extensible to the introduction of new applications and new use cases.

BGP Link-State extensions [[RFC7752](#)] have been specified to enable distribution of the link-state topology information from the IGPs to an application like a controller or Path Computation Engine (PCE) via BGP. The controller/PCE gets the end to end topology information across IGP domains so it can perform path computations for use-cases like end to end traffic engineering (TE) using RSVP-TE or SR based mechanisms. A similar challenge to what was describe above is hence also faced by such centralized computation entities.

There is thus a need for BGP-LS extensions to also report link attributes on a per application basis on the same lines as introduced in the link-state routing protocols. This document defines these BGP-LS extensions and also covers the backward compatibility issues related to existing BGP-LS deployments.

2. Application Specific Link Attributes TLV

The BGP-LS [[RFC7752](#)] specifies the Link NLRI for advertisement of links and their attributes using the BGP-LS Attribute. The Application Specific Link Attributes (ASLA) TLV is a new optional top-level BGP-LS Attribute TLV that is introduced for Link NLRIs. It is defined such that it may act as a container for certain existing and future link attributes that require to be defined in an application specific scope.

The format of this TLV is as follows and is similar to the corresponding ASLA sub-TLVs defined for OSPF and IS-IS in [[I-D.ietf-ospf-te-link-attr-reuse](#)] and [[I-D.ietf-isis-te-app](#)] respectively.

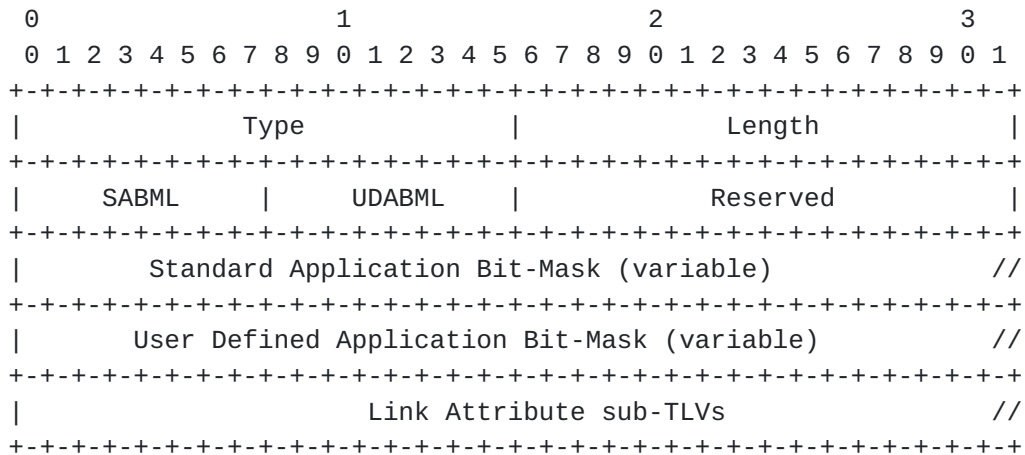


Figure 1: Application Specific Link Attributes TLV

where:

- o Type: TBD (see IANA Considerations [Section 6](#))
- o Length: variable.
- o SABML : 1 octet value carrying the Standard Application Bit-Mask Length in multiples of 4 octets. If the Standard Application Bit-Mask is not present, the SABML MUST be set to 0.
- o UDABML : 1 octet value carrying the User Defined Application Bit-Mask Length in multiples of 4 octets. If the User Defined Application Bit-Mask is not present, the UDABML MUST be set to 0.
- o Standard Application Bit-Mask : variable size in multiple of 4 octets and optional set of bits, where each bit represents a single standard application. The bits are defined in the IANA "IGP Parameters" registries under the "Link Attribute Applications" registry [[I-D.ietf-isis-te-app](#)].
- o User Defined Application Bit-Mask : variable size in multiple of 4 octets and optional set of bits, where each bit represents a single user defined application. The bits are not managed or assigned by IANA or any other standards body and are left to implementation specifics.

- o sub-TLVs : BGP-LS Attribute TLVs corresponding to the Link NLRI that are application specific (as specified in [Section 3](#)) are included as sub-TLVs of the ASLA TLV

An ASLA TLV with both the SABML and UDABML set to 0 (i.e. without any application specific bitmasks) indicate that the link attribute sub-TLVs that it encloses are applicable for all applications.

The ASLA TLV and its sub-TLVs can only be added to the BGP-LS Attribute associated with the Link NLRI of the node that originates the underlying IGP link attribute TLVs/sub-TLVs. The procedures for originating link attributes in the ASLA TLV from underlying IGP is specified in [Section 4](#).

When the node is not running any of the IGP but running a protocol like BGP, then the link attributes for the node's local links MAY be originated as part of the BGP-LS Attribute using the ASLA TLV and its sub-TLVs within the Link NLRI corresponding to the local node.

3. Application Specific Link Attributes

Several BGP-LS Attribute TLVs corresponding to the Link NLRI are defined in BGP-LS and more may be added in the future. The following types of link attributes are required to be considered as application specific.

- o those that have different values for different applications (e.g. a different TE metric value used for RSVP-TE than for SR TE)
- o those that are applicable to multiple applications but need to be used only by specific application (e.g. certain SRLG values are configured on a node for LFA but the same do not need to be used for RSVP-TE)

The following table lists the currently defined BGP-LS Attributes TLVs corresponding to Link NLRI which have application specific semantics. They were originally defined with semantics for RSVP-TE and GMPLS applications.

TLV Code Point	Description	Reference Document
1088	Administrative group (color)	[RFC7752]
1090	Max Reservable Bandwidth	[RFC7752]
1091	Unreserved Bandwidth	[RFC7752]
1092	TE Metric	[RFC7752]
1096	SRLG	[RFC7752]
1114	Unidirectional link delay	[I-D.ietf-idr-te-pm-bgp]
1115	Min/Max Unidirectional link delay	[I-D.ietf-idr-te-pm-bgp]
1116	Unidirectional link delay variation	[I-D.ietf-idr-te-pm-bgp]
1117	Unidirectional packet loss	[I-D.ietf-idr-te-pm-bgp]
1118	Unidirectional residual bandwidth	[I-D.ietf-idr-te-pm-bgp]
1119	Unidirectional available bandwidth	[I-D.ietf-idr-te-pm-bgp]
1120	Unidirectional bandwidth utilization	[I-D.ietf-idr-te-pm-bgp]
1173	Extended Administrative group (color)	[I-D.ietf-idr-eag-distribution]

Table 1: BGP-LS Attribute TLVs also used as sub-TLVs of ASLA TLV

All the BGP-LS Attribute TLVs defined in the table above are RECOMMENDED to be continued to be used at the top-level in the BGP-LS Attribute for carrying attributes specific to RSVP-TE/GMPLS application without the use of the ASLA TLV.

When a new link attribute is introduced, it may be thought of as being specific to only a single application. However, down the line, it may be also shared by other applications and/or require application specific values. In such cases, it is RECOMMENDED to err on the side of caution and define such attributes as application specific to ensure flexibility in the future.

BGP-LS Attribute TLVs corresponding to Link NLRI that are defined in the future MUST specify if they are application specific and hence are REQUIRED to be encoded within an ASLA TLV.

Only application specific link attributes need to be advertised within the ASLA TLV. Link attributes which do not have application specific semantics SHOULD NOT be advertised within the ASLA TLV. Receivers SHOULD ignore any non-application specific attribute sub-TLVs within the ASLA TLV.

4. Procedures

The procedures described in this section apply to networks where all BGP-LS originators and consumers support this specification. The backward compatibility aspects and operations in deployments where there are some BGP-LS originators or consumers that do not support this specification is described further in [Section 5](#).

The BGP-LS originator learns of the association of an application specific attribute to one or more set of applications from either the underlying IGP protocol LSA/LSPs from which it is sourcing the topology information or from the local node configuration when advertising attributes for the local node only.

The association of an application specific link attribute with a specific application context when advertising attributes for the local node only (e.g. when running BGP as the only routing protocol) is an implementation specific matter and outside the scope of this document.

[I-D.ietf-ospf-te-link-attr-reuse] and [[I-D.ietf-isis-te-app](#)] specify the mechanisms for flooding of application specific link attributes in OSPFv2/v3 and IS-IS respectively. These IGP specifications also describe the backward compatibility aspects and the existing RSVP-TE/GMPLS specific TLV encoding mechanisms in respective protocols.

A BGP-LS originator node which is sourcing link-state information from the underlying IGP determines the mechanism of flooding application specific link attributes based on the following rules:

1. Application specific link attributes received from an IGP node using existing RSVP-TE/GMPLS encodings only (i.e. without any ASLA sub-TLV) MUST be encoded using the respective BGP-LS top-level TLVs listed in Table 1 (i.e. not within ASLA TLV). When the IGP node is also SR enabled then another copy of application specific link attributes SHOULD be also encoded as ASLA sub-TLVs with the SR application bit for them. Further rules do not apply

for such IGP nodes that do not use ASLA sub-TLVs in their advertisements.

2. In case of IS-IS, when application specific link attributes are received from a node with the L bit set in the ASLA sub-TLV then the application specific link attributes are picked up from the legacy ISIS TLVs/sub-TLVs and MUST be encoded within the BGP-LS ASLA TLV as sub-TLVs with the application bitmask set as per the IGP ASLA sub-TLV. When the ASLA sub-TLV with the L bit set also has the RSVP-TE application bit set then the link attributes from such an ASLA sub-TLV MUST be also encoded using the respective BGP-LS top-level TLVs listed in Table 1 (i.e. not within ASLA TLV).
3. In case of OSPFv2/v3, when application specific link attributes are received from a node via TE LSAs then the application specific link attributes from those LSAs MUST be encoded using the respective BGP-LS TLVs listed in Table 1 (i.e. not within ASLA TLV).
4. Application specific link attributes received from an IGP node within its ASLA sub-TLV MUST be encoded in the BGP-LS ASLA TLV as sub-TLVs with the application bitmask set as per the IGP advertisement.

These rules ensure that a BGP-LS originator performs the translation for all application specific link attributes from the IGP nodes into the new BGP-LS ASLA TLVs irrespective of the IGP node supporting the ASLA extension. Furthermore, it also ensures that BGP-LS TLVs defined for RSVP-TE and GMPLS applications continue to be used for those respective applications.

A BGP-LS consumer node always gets all application specific link attributes corresponding to RSVP-TE and GMPLS applications as existing top-level BGP-LS TLVs while for other applications they are encoded in ASLA TLV(s) with appropriate applicable bit mask setting.

5. Backward Compatibility

When it comes to BGP-LS, the backward compatibility aspects are associated with the originators (i.e. nodes) and consumers (e.g. PCE, controllers, applications, etc.) of the topology information. The originators of BGP-LS information need to ensure that their encoding of application specific link attributes is done such that consumers running BGP-LS implementations without this specification support can still support existing applications like RSVP-TE and SR. The consumers running BGP-LS implementations that support this

specification should also be able to work with BGP-LS originators that do not support this specification and vice versa.

BGP-LS implementations have been originating link attributes and consuming them without any application specific scoping. While the ASLA TLV can be used without any backward compatibility considerations for any new application (e.g. IGP Flexible Algorithm) specific attribute advertisements, for existing applications like RSVP-TE and SR some backward compatibility aspects need to be taken care of.

This requires the introduction of a "compatibility mode" of operations at originators of BGP-LS information for encoding of information such that older implementations of BGP-LS consumers can still support applications like RSVP-TE and SR. In addition to the rules specified in [Section 4](#), the following rules are to be followed when operating in "compatibility mode" :

- o Application specific link attribute received in IGP ASLA sub-TLVs, corresponding to RSVP-TE or SR applications, MUST be also encoded in their existing top level TLVs (as listed in Table 1) outside of the ASLA TLV in addition to them being also advertised within the ASLA TLV
- o When the same application specific attribute, received in IGP ASLA sub-TLVs, has different values for RSVP-TE and SR applications then the value for RSVP-TE application SHOULD be preferred over the value for SR application for advertisement as the top level TLV (as listed in Table 1). An implementation MAY provide a knob to reverse this preference.

It is RECOMMENDED that implementations operate in "compatibility mode" by default. Implementations SHOULD have a knob for turning the "compatibility mode" on or off. Operators MAY turn the "compatibility mode" off when they are assured that all BGP-LS consumers have been upgraded to support the extensions in this document.

It is RECOMMENDED that the nodes which support this specification are selected as originators of BGP-LS information when sourced from the IGP.

A BGP-LS consumer which does not implement this specification will ignore the ASLA TLV and instead continue to use the attributes from the existing top level TLVs.

A BGP-LS consumer which implements this specification SHOULD prefer the application specific attribute value received via sub-TLVs within the ASLA TLV over the value received via the top level TLVs.

6. IANA Considerations

This document requests assigning code-points from the registry "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" based on table below. The column "IS-IS TLV/Sub-TLV" defined in the registry does not require any value and should be left empty.

Code Point	Description	Length
TBD	Application Specific Link Attributes TLV	variable

7. Manageability Considerations

This section is structured as recommended in [\[RFC5706\]](#).

The new protocol extensions introduced in this document augment the existing IGP topology information that was distributed via [\[RFC7752\]](#). Procedures and protocol extensions defined in this document do not affect the BGP protocol operations and management other than as discussed in the Manageability Considerations section of [\[RFC7752\]](#). Specifically, the malformed NLRIs attribute tests in the Fault Management section of [\[RFC7752\]](#) now encompass the new TLVs for the BGP-LS NLRI in this document.

7.1. Operational Considerations

No additional operation considerations are defined in this document.

7.2. Management Considerations

No additional management considerations are defined in this document.

8. Security Considerations

The new protocol extensions introduced in this document augment the existing IGP topology information that was distributed via [\[RFC7752\]](#). Procedures and protocol extensions defined in this document do not affect the BGP security model other than as discussed in the Security Considerations section of [\[RFC7752\]](#).

9. Acknowledgements

The authors would like to thank Les Ginsberg for his review and contributions to this work.

10. References

10.1. Normative References

- [I-D.ietf-isis-te-app]
Ginsberg, L., Psenak, P., Previdi, S., Henderickx, W., and J. Drake, "IS-IS TE Attributes per application", [draft-ietf-isis-te-app-06](#) (work in progress), April 2019.
- [I-D.ietf-ospf-te-link-attr-reuse]
Psenak, P., Ginsberg, L., Henderickx, W., Tantsura, J., and J. Drake, "OSPF Link Traffic Engineering (TE) Attribute Reuse", [draft-ietf-ospf-te-link-attr-reuse-07](#) (work in progress), April 2019.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", [RFC 7752](#), DOI 10.17487/RFC7752, March 2016, <<https://www.rfc-editor.org/info/rfc7752>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

10.2. Informative References

- [I-D.ietf-idr-eag-distribution]
Wang, Z., Wu, Q., and J. Tantsura, "Distribution of MPLS-TE Extended admin Group Using BGP", [draft-ietf-idr-eag-distribution-08](#) (work in progress), October 2018.
- [I-D.ietf-idr-te-pm-bgp]
Ginsberg, L., Previdi, S., Wu, Q., Tantsura, J., and C. Filsfils, "BGP-LS Advertisement of IGP Traffic Engineering Performance Metric Extensions", [draft-ietf-idr-te-pm-bgp-18](#) (work in progress), December 2018.

[I-D.ietf-lsr-flex-algo]

Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", [draft-ietf-lsr-flex-algo-02](#) (work in progress), May 2019.

[RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", [RFC 1195](#), DOI 10.17487/RFC1195, December 1990, <<https://www.rfc-editor.org/info/rfc1195>>.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), DOI 10.17487/RFC2328, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.

[RFC4202] Kompella, K., Ed. and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 4202](#), DOI 10.17487/RFC4202, October 2005, <<https://www.rfc-editor.org/info/rfc4202>>.

[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, <<https://www.rfc-editor.org/info/rfc5286>>.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", [RFC 5340](#), DOI 10.17487/RFC5340, July 2008, <<https://www.rfc-editor.org/info/rfc5340>>.

[RFC5706] Harrington, D., "Guidelines for Considering Operations and Management of New Protocols and Protocol Extensions", [RFC 5706](#), DOI 10.17487/RFC5706, November 2009, <<https://www.rfc-editor.org/info/rfc5706>>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [RFC 8402](#), DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

Authors' Addresses

Ketan Talaulikar
Cisco Systems

Email: ketant@cisco.com

Peter Psenak
Cisco Systems
Slovakia

Email: ppsenak@cisco.com

Jeff Tantsura
Apstra

Email: jefftant.ietf@gmail.com