

Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: November 22, 2021

K. Talaulikar, Ed.
P. Psenak
Cisco Systems
J. Tantsura
Apstra
May 21, 2021

**Application-Specific Attributes Advertisement with BGP Link-State
draft-ietf-idr-bgp-ls-app-specific-attr-06**

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 that 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.

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 November 22, 2021.

Copyright Notice

Copyright (c) 2021 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
1.1.	Requirements Language	3
2.	Application Specific Link Attributes TLV	3
3.	Application-Specific Link Attributes	5
4.	Procedures	7
5.	Deployment Considerations	9
6.	Backward Compatibility	9
7.	IANA Considerations	10
8.	Manageability Considerations	10
9.	Security Considerations	10
10.	Acknowledgements	11
11.	References	11
11.1.	Normative References	11
11.2.	Informative References	11
	Authors' Addresses	13

[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 that have use cases for many of the link attributes historically used by RSVP-TE and GMPLS. Such applications include Segment Routing (SR) Policy [[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 Policy 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 Policy. 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.

[RFC8920] and [RFC8919] define extensions for OSPF and IS-IS respectively that address these issues. Also, as the evolution of use cases for link attributes can be expected to continue in the years to come, these documents define an easily extensible solution for 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 IGP 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 Policy mechanisms. A similar challenge to what was described 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.

1.1. 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.

2. Application Specific Link Attributes TLV

The BGP-LS [RFC7752] specifies the Link NLRI for the 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 application-specific definition.

The format of this TLV is as follows and is similar to the corresponding ASLA sub-TLVs defined for OSPF and IS-IS in [RFC8920] and [RFC8919] respectively.




Figure 1: Application-Specific Link Attributes TLV

where:

- o Type: 1122
- o Length: variable.
- o SABM Length : Standard Application Identifier Bit Mask Length in octets. The values MUST be 0, 4, or 8. If the Standard Application Identifier Bit Mask is not present, the value MUST be set to 0.
- o UDABM Length : User-Defined Application Identifier Bit Mask Length in octets. The values MUST be 0, 4, or 8. If the User-Defined Application Identifier Bit Mask is not present, the value MUST be set to 0.
- o Standard Application Identifier Bit Mask : of size 0, 4, or 8 octets as indicated by the SABML. 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 [[RFC8919](#)].
- o User-Defined Application Identifier Bit Mask : of size 0, 4, or 8 octets as indicated by the UDABML. 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 definition is left to the implementation.

- 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 identifier bit masks) indicates 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 are specified in [Section 4](#).

When the node is not running any of the IGPs 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 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 Policy)
- o those that apply to multiple applications but need to be used only by a 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 that have application-specific semantics. These were originally defined with semantics for RSVP-TE and GMPLS applications.

TLV Code Point	Description	Reference Document
1088	Administrative group (color)	[RFC7752]
1092	TE Default Metric	[RFC7752]
1096	Shared Risk Link Group	[RFC7752]
1114	Unidirectional Link Delay	[RFC8571]
1115	Min/Max Unidirectional Link Delay	[RFC8571]
1116	Unidirectional Delay Variation	[RFC8571]
1117	Unidirectional Link Loss	[RFC8571]
1118	Unidirectional Residual Bandwidth	[RFC8571]
1119	Unidirectional Available Bandwidth	[RFC8571]
1120	Unidirectional Utilized Bandwidth	[RFC8571]
1173	Extended Administrative Group	[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 continue to be advertised at the top-level in the BGP-LS Attribute for carrying attributes specific to RSVP-TE 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, subsequently, 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 that do not have application-specific semantics MUST NOT be advertised within the ASLA TLV. Receivers MUST ignore any non-application-specific attribute sub-TLVs within the ASLA TLV.

When the same application-specific link attributes are advertised both within the ASLA TLV and as top-level TLVs in the BGP-LS Attribute, the attributes advertised within the ASLA TLV take precedence for the applications indicated in the ASLA TLV encoding.

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 are described further in [Section 6](#).

The BGP-LS originator learns of the association of an application-specific attribute to one or more applications from either the underlying IGP protocol LSA/LSPs from which it is advertising 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.

[RFC8920] and [[RFC8919](#)] specify the mechanisms for advertising 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 the respective protocols.

A BGP-LS originator node that is advertising link-state information from the underlying IGP determines the protocol encoding of 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 MUST be encoded using the respective BGP-LS top-level TLVs listed in Table 1.
2. Application-specific link attributes received from an IGP node using ASLA sub-TLV MUST be encoded in the BGP-LS ASLA TLV as sub-TLVs.

3. In the case of IS-IS, the following specific procedures are to be followed:

- * When application-specific link attributes are received from a node with the L bit set in the ASLA sub-TLV AND application bits other than RSVP-TE are set in the application bit masks then the application-specific link attributes advertised in the corresponding legacy IS-IS TLVs/sub-TLVs MUST be encoded within the BGP-LS ASLA TLV as sub-TLVs with the application bits, other than the RSVP-TE bit, copied from the IS-IS ASLA sub-TLV. The link attributes advertised in the legacy IS-IS TLVs/sub-TLVs are also advertised in BGP-LS top-level TLVs listed in Table 1. Note that this is true regardless of whether the RSVP-TE bit was set in the IS-IS ASLA TLV/sub-TLV.
- * When the ASLA sub-TLV has the RSVP-TE application bit set, then the link attributes for the corresponding ASLA sub-TLV MUST be encoded using the respective BGP-LS top-level TLVs listed in Table 1.
- * [\[RFC8919\]](#) allows the advertisement of the Maximum Link Bandwidth within an ASLA sub-TLV even though it is not an application-specific attribute. However, when originating the Maximum Link Bandwidth into BGP-LS, the attribute MUST be encoded only in the top-level BGP-LS Maximum Link Bandwidth TLV (1089) and the receiver MUST ignore them when advertised within the BGP-LS ASLA TLV.
- * [\[RFC8919\]](#) also allows the advertisement of the Maximum Reservable Link Bandwidth and the Unreserved Bandwidth within an ASLA sub-TLV even though these attributes are specific to RSVP-TE application. However, when originating the Maximum Reservable Link Bandwidth and Unreserved Bandwidth into BGP-LS, these attributes MUST be encoded only in the BGP-LS top-level Maximum Reservable Link Bandwidth TLV (1090) and Unreserved Bandwidth TLV (1091) respectively and not within the BGP-LS ASLA TLV.

These rules ensure that a BGP-LS originator performs the advertisement for all application-specific link attributes from the IGP nodes that support or do not support the ASLA extension. Furthermore, it also ensures that the top-level BGP-LS TLVs defined for RSVP-TE and GMPLS applications continue to be used for advertisement of their application-specific attributes.

A BGP-LS consumer node would normally receive all the 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. A BGP-LS consumer that 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.

5. Deployment Considerations

SR Policy and LFA applications have been deployed in some networks using the IGP link attributes defined originally for RSVP-TE as discussed in [[RFC8920](#)] and [[RFC8919](#)]. The corresponding BGP-LS top-level link attribute TLVs originally defined for RSVP-TE have also been similarly used for SR Policy and LFA applications by BGP-LS consumers. Such usage MAY continue without requiring the support of the application-specific link attribute encodings described in this document as long as the following conditions are met:

- o The application is SR Policy or LFA and RSVP-TE is not deployed anywhere in the network
- o The application is SR Policy or LFA, RSVP-TE is deployed in the network, and both the set of links on which SR Policy and/or LFA advertisements are required and the attribute values used by SR Policy and/or LFA on all such links is fully congruent with the links and attribute values used by RSVP-TE

6. Backward Compatibility

The backward compatibility aspects for BGP-LS are associated with the originators (i.e., nodes) and consumers (e.g., PCE, controllers, applications, etc.) of the topology information. BGP-LS implementations have been originating link attributes and consuming them without any application-specific scoping prior to the extensions specified in this document.

IGP backward compatibility aspects associated with application-specific link attributes for RSVP-TE, SR Policy, and LFA applications are discussed in the Backward Compatibility sections of [[RFC8920](#)] and [[RFC8919](#)]. While the backward compatibility aspects ensure compatibility of IGP advertisements, they also serve to ensure the backward compatibility of the BGP-LS advertisements used by BGP-LS consumers. In deployments where the BGP-LS originators or consumers do not support the extensions specified in this document, the IGP's need to continue to advertise link attributes intended for use by SR Policy and LFA applications using the RSVP-TE/GMPLS encodings. This allows BGP-LS advertisements to be consistent with the behavior prior to the extensions defined in this document

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

7. IANA Considerations

This document requests assignment of code-points from the registry "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" based on the table below which reflects the values assigned via the early allocation process. 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
1122	Application-Specific Link Attributes TLV	variable

8. Manageability Considerations

The new protocol extensions introduced in this document augment the existing IGP topology information defined in [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 encompasses the BGP-LS TLVs defined in this document.

The extensions specified in this document do not specify any new configuration or monitoring aspects in BGP or BGP-LS. The specification of BGP models is an ongoing work based on [I-D.ietf-idr-bgp-model].

9. Security Considerations

The procedures and protocol extensions defined in this document do not affect the BGP security model. See the "Security Considerations" section of [RFC4271] for a discussion of BGP security. Also, refer to [RFC4272] and [RFC6952] for analyses of security issues for BGP. Security considerations for acquiring and distributing BGP-LS information are discussed in [RFC7752]. The TLVs introduced in this document are used to propagate the application-specific link attributes IGP extensions defined in [RFC8919] and [RFC8920]. It is assumed that the IGP instances originating these TLVs will support all the required security (as described in [RFC8919] and [RFC8920]) in order to prevent any security issues when propagating the TLVs

into BGP-LS. The advertisement of the link attribute information defined in this document presents no significant additional risk beyond that associated with the existing link attribute information already supported in [RFC7752].

10. Acknowledgements

The authors would like to thank Les Ginsberg, Baalajee S, Amalesh Maity, Acee Lindem, Keyur Patel and Paul Wouters for their review and feedback on this document.

11. References

11.1. Normative References

- [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>>.
- [RFC8919] Ginsberg, L., Psenak, P., Previdi, S., Henderickx, W., and J. Drake, "IS-IS Application-Specific Link Attributes", [RFC 8919](#), DOI 10.17487/RFC8919, October 2020, <<https://www.rfc-editor.org/info/rfc8919>>.
- [RFC8920] Psenak, P., Ed., Ginsberg, L., Henderickx, W., Tantsura, J., and J. Drake, "OSPF Application-Specific Link Attributes", [RFC 8920](#), DOI 10.17487/RFC8920, October 2020, <<https://www.rfc-editor.org/info/rfc8920>>.

11.2. Informative References

- [I-D.ietf-idr-bgp-model] Jethanandani, M., Patel, K., Hares, S., and J. Haas, "BGP YANG Model for Service Provider Networks", [draft-ietf-idr-bgp-model-10](#) (work in progress), November 2020.

- [I-D.ietf-idr-eag-distribution]
Tantsura, J., Wang, Z., Wu, Q., and K. Talaulikar,
"Distribution of Traffic Engineering Extended
Administrative Groups using BGP-LS", [draft-ietf-idr-eag-distribution-16](#) (work in progress), April 2021.
- [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>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A
Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#),
DOI 10.17487/RFC4271, January 2006,
<<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis",
[RFC 4272](#), DOI 10.17487/RFC4272, January 2006,
<<https://www.rfc-editor.org/info/rfc4272>>.
- [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>>.
- [RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of
BGP, LDP, PCEP, and MSDP Issues According to the Keying
and Authentication for Routing Protocols (KARP) Design
Guide", [RFC 6952](#), DOI 10.17487/RFC6952, May 2013,
<<https://www.rfc-editor.org/info/rfc6952>>.

[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>>.

[RFC8571] Ginsberg, L., Ed., Previdi, S., Wu, Q., Tantsura, J., and C. Filsfils, "BGP - Link State (BGP-LS) Advertisement of IGP Traffic Engineering Performance Metric Extensions", [RFC 8571](#), DOI 10.17487/RFC8571, March 2019, <<https://www.rfc-editor.org/info/rfc8571>>.

Authors' Addresses

Ketan Talaulikar (editor)
Cisco Systems
India

Email: ketant@cisco.com

Peter Psenak
Cisco Systems
Slovakia

Email: ppsenak@cisco.com

Jeff Tantsura
Apstra

Email: jefftant.ietf@gmail.com

