

IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: 16 July 2026

A. Wang
China Telecom
H. Chen
Individual
K. Talaulikar
Cisco Systems
S. Zhuang
Huawei Technologies
C. Lin
New H3C Technologies
12 January 2026

BGP-LS Extension for Inter-AS Topology Retrieval
draft-ietf-idr-bgpls-inter-as-topology-ext-19

Abstract

This document describes the process to distribute Border Gateway Protocol- Link State (BGP-LS) key parameters for inter-domain links between two Autonomous Systems. This document defines a new type within the BGP-LS NLRI for a Stub Link and three new type-length-values (TLVs) for BGP-LS Link descriptor. These additions to BGP-LS let Software Definition Network (SDN) controllers retrieve the network topology automatically under various inter-AS environments.

Such extension and process can enable the network operator to collect the interconnect information between different domains and then calculate the overall network topology automatically based on the information provided by BGP-LS protocol.

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 16 July 2026.

Copyright Notice

Copyright (c) 2026 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Conventions used in this document	3
3. Terminology	3
4. Inter-AS Domain Scenarios	3
5. Stub Link NLRI	4
6. IGP Information for Inter-AS Link	6
7. BGP-LS Extensions for Inter-AS Link	6
7.1. Remote AS Number TLV	7
7.2. IPv4 Remote ASBR ID	8
7.3. IPv6 Remote ASBR ID	8
8. Topology Reconstruction.	9
9. Security Considerations	9
10. IANA Considerations	10
10.1. New BGP-LS NLRI type	10
10.2. New Stub Link Descriptors	10
11. Acknowledgement	11
12. References	11
12.1. Normative References	11
12.2. Informative References	11
Authors' Addresses	11

1. Introduction

BGP-LS [RFC9552] describes the methodology that using BGP protocol to transfer the Link-State information. Such method can enable SDN controller to collect the underlay network topology automatically, but normally it can only get the information within one Interior Gateway Protocol (IGP) domain. If the operator has more than one IGP domain, and these domains interconnect with each other, there is no mechanic within current BGP- LS to transfer the interconnect topology information.

[RFC9086] defines some extensions for exporting BGP peering node topology information (including its peers, interfaces and peering ASs) in a way that is exploitable in order to compute efficient BGP Peering Engineering policies and strategies. Such information can also be used to calculate the interconnection topology among different IGP domains, but it requires every border router to run BGP-LS protocol and report the information to SDN controller. Considering there will be plenty of border routers on the network boundary, such solution restricts its deployment flexibility.

This draft analyzes the situations during which the SDN controller needs to get the interconnected topology information between different AS domains. After describing these situations, this draft defines a new Stub Link type within the BGP-LS NLRI[RFC9552] to describe the Iner-AS link and some new TLVs for that new BGP-LS type. After that, the SDN controller can then deduce the multi-domain topology automatically based on the information from BGP-LS protocol.

2. Conventions used in this document

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 [RFC2119] .

3. Terminology

The following terms are defined in this document:

- * IDCs: Internet Data Centers
- * MAN: Metrio-Area-Network
- * SDN: Software Definition Network

4. Inter-AS Domain Scenarios

Figure 1 illustrates the multi-domain scenarios that this draft discusses. Normally, SDN Controller can get the topology of IGP A and IGP B individually via the BGP-LS protocol, but it can't get the topology connection information between these two IGP domains because there is generally no IGP protocol run on the connected links.

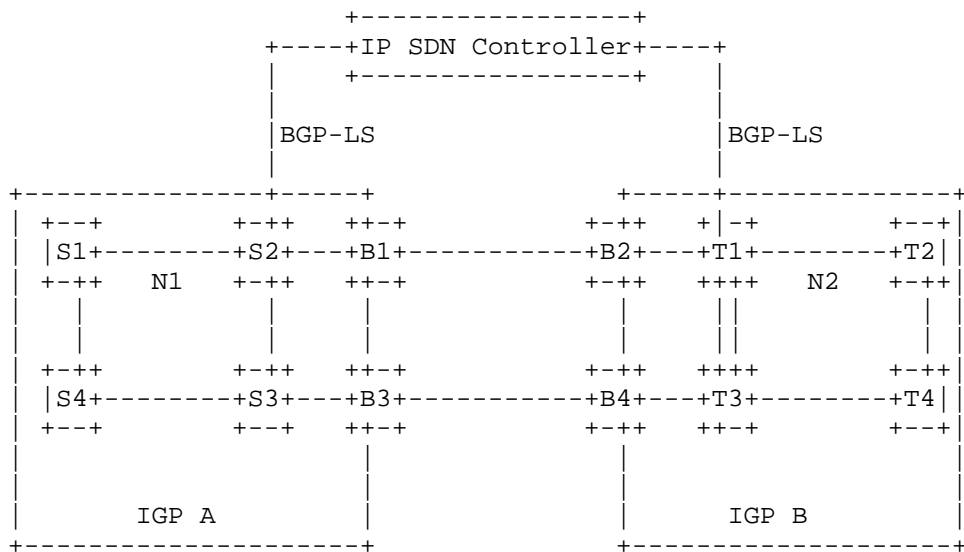


Figure 1: Inter-AS Domain Scenarios

5. Stub Link NLRI

[RFC9552] defines four types within the BGP Link-State NLRI (Node NLRI, Link NLRI, IPv4 Topology Prefix NLRI, and IPv6 Topology Prefix NLRI) to transfer the topology and prefix information. For inter-as link, the two ends of the link exist in different IGP domains, so it is not appropriate to transfer their information within the current defined NLRI types.

This draft defines one new NLRI type 7, see__Section 10) within the BGP Link-State (BGP-LS) NLRI, called the Stub link NLRI. The Stub link NLRI is encoded in the format shown in Figure 2 and explained below:

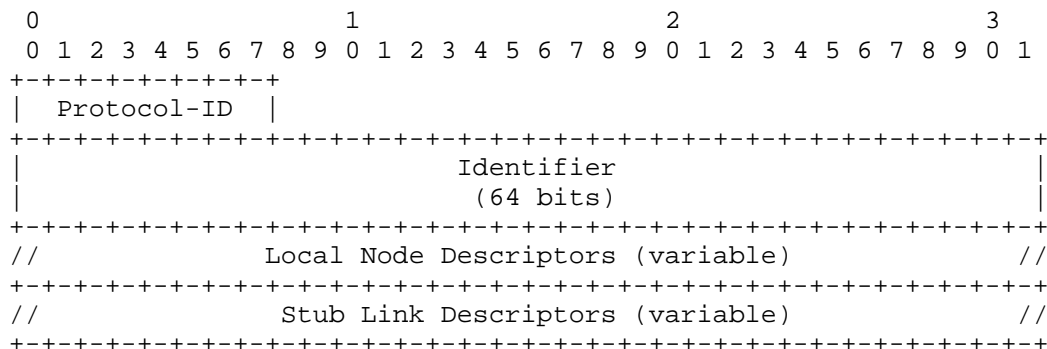


Figure 2: Stub Link NLRI Format

The "Protocol-ID" should be set to the value that indicates the source protocol of the stub link information, as indicated in [RFC9552] in section 5.2.

Local Node descriptors: define the ASBRs that are attached to the Inter-AS stub link, and use the [RFC9552] "Local Node Descriptor" in section 5.2.1.4. The following Node Descriptor sub-TLVs from [RFC9552] are valid for inclusion in the local Node descriptor: AS system, OSPF Area-ID, IGP Router-ID.

Stub Link Descriptors: define the Stub link which has only one end Located in the IGP domain using the [RFC9552] "Link Descriptor definition" in section 5.2.2 with the exceptions noted below.

The Stub link Descriptor supports the inclusion of the following subTLVs:

Link/Local Identifier (TLV 258, [RFC9552]),

IPv4 Interface Address (TLV 259, [RFC9552]),

IPv4 Neighbor Address (TLV 260, [RFC9552]),

IPv6 Interface Address (TLV 261, [RFC9552]),

IPv6 Neighbor Address (TLV 262, [RFC9552]),

Multi-topology identifier (TLV 263, [RFC9552]),

Remote-AS Number (TLV 270, [This document], section Section 7.1),

IPv4 Remote ASBR ID [TLV 271, [This document], section Section 7.2), and

IPv6 Remote ASBR ID [TLV 272, [This document], section Section 7.3).

This newly defined NLRI can be used to describe the link that has only one end located within the IGP domain, as described in the following sections. The Node and Link Descriptor sub-TLVs and Node and Link attributes that are defined in [RFC9552] can be included in the NLRI if necessary. The interface and neighbor address sub-TLVs SHOULD be included in the Local Node Descriptors to differentiate the parallel links between two ASBRs.

6. IGP Information for Inter-AS Link

[RFC9346] and [RFC5392] define IS-IS and OSPF extensions respectively to deal with the reasons for reporting inter-AS link. Three sub-TLVs relating to Inter-Domain Links (Remote AS Number, IPv4 Remote ASBR ID, and IPv6 Remote ASBR ID) are defined in these documents. .

These IGP TLVs are flooded within the IGP domain automatically. This document specifies that these MAY also be carried within the newly defined Stub Link NLRI within the BGP-LS protocol, as the descriptors for the inter-AS stub link. The "Local Node Descriptors" in the Stub Link NLRI within the BGP-LS NLRI should describe the characteristics of ASBRs that are connected across the inter-AS links.

If the SDN controller knows these information via one of the interior router that runs BGP-LS protocol, the SDN controller can rebuild the inter-AS topology correctly according to the procedure described in Section 8

7. BGP-LS Extensions for Inter-AS Link

This section proposes to add three new TLVs to be supported in the Stub Link NLRI in the BGP-LS NLRI. These new TLVs allow BGP-LS to transfer inter-AS information gathered by the SDN controller.

The following Link Descriptor TLVs are added into the BGP-LS protocol :

TLV Code Point	Description	IS-IS/OSPF TLV /Sub-TLV	Reference (RFC/Section)
270	Remote AS Number	24/21	[RFC9346]/3.3.1 [RFC5392]/3.3.1
271	IPv4 Remote ASBR ID	25/22	[RFC9346]/3.3.2 [RFC5392]/3.3.2
272	IPv6 Remote ASBR ID	26/24	[RFC9346]/3.3.3 [RFC5392]/3.3.3

Figure 3: Stub Link Descriptor TLVs

Detail encoding of these TLVs are synchronized with the corresponding parts in [RFC9346] and [RFC5392], which keeps the BGP-LS protocol agnostic to the underly protocol.

7.1. Remote AS Number TLV

A new TLV, the remote AS number TLV, is defined for inclusion in the link descriptor when advertising inter-AS links. The remote AS number TLV specifies the AS number of the neighboring AS to which the advertised link connects.

The remote AS number TLV is TLV type 270 (see Section 10) and is 4 octets in length. The format is as follows:

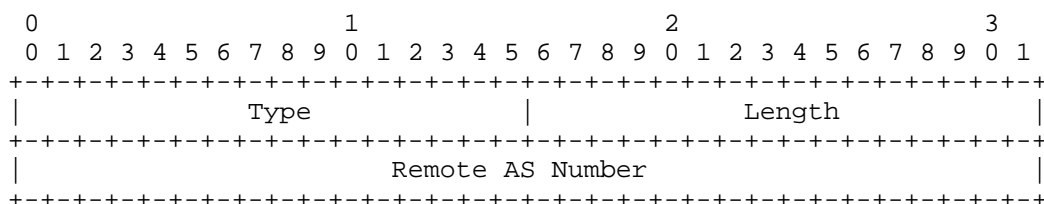


Figure 4: Remote AS Number TLV Format

The Remote AS number field has 4 octets. When only 2 octets are used for the AS number, as in current deployments, the left (high-order) 2 octets MUST be set to 0. The remote AS number TLV MUST be included when a router advertises an inter-AS link.

7.2. IPv4 Remote ASBR ID

A new TLV, which is referred to as the IPv4 remote ASBR ID TLV, is defined for inclusion in the link descriptor when advertising inter-AS links. The IPv4 remote ASBR ID TLV specifies the IPv4 identifier of the remote ASBR to which the advertised inter-AS link connects. This could be any stable and routable IPv4 address of the remote ASBR. Use of the TE Router ID as specified in the Traffic Engineering router ID TLV [RFC9346] is RECOMMENDED.

The IPv4 remote ASBR ID TLV is TLV type 271 (see Section 10) and is 4 octets in length. The format of the IPv4 remote ASBR ID TLV is as follows:

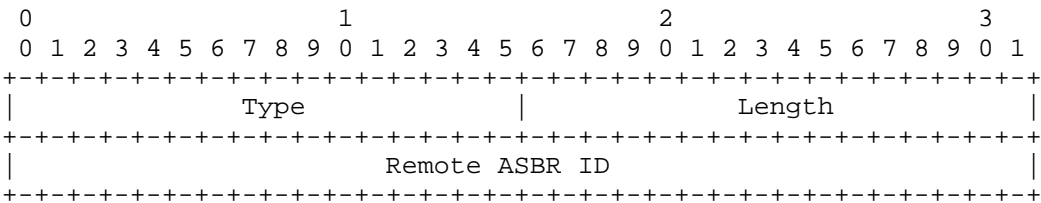


Figure 5: IPv4 Remote ASBR ID TLV Format

The IPv4 remote ASBR ID TLV MUST be included if the neighboring ASBR has an IPv4 address. If the neighboring ASBR does not have an IPv4 address (not even an IPv4 TE Router ID), the IPv6 remote ASBR ID TLV MUST be included instead. An IPv4 remote ASBR ID TLV and IPv6 remote ASBR ID TLV MAY both be present in an inter-AS stub link NLRI.

7.3. IPv6 Remote ASBR ID

A new TLV, which is referred to as the IPv6 remote ASBR ID TLV, is defined for inclusion in the link descriptor when advertising inter-AS links. The IPv6 remote ASBR ID TLV specifies the IPv6 identifier of the remote ASBR to which the advertised inter-AS link connects. This could be any stable and routable IPv6 address of the remote ASBR. Use of the TE Router ID as specified in the IPv6 Traffic Engineering router ID TLV [RFC9346] is RECOMMENDED.

The IPv6 remote ASBR ID TLV is TLV type 272 (see Section 10) and is 16 octets in length. The format of the IPv6 remote ASBR ID TLV is as follows:

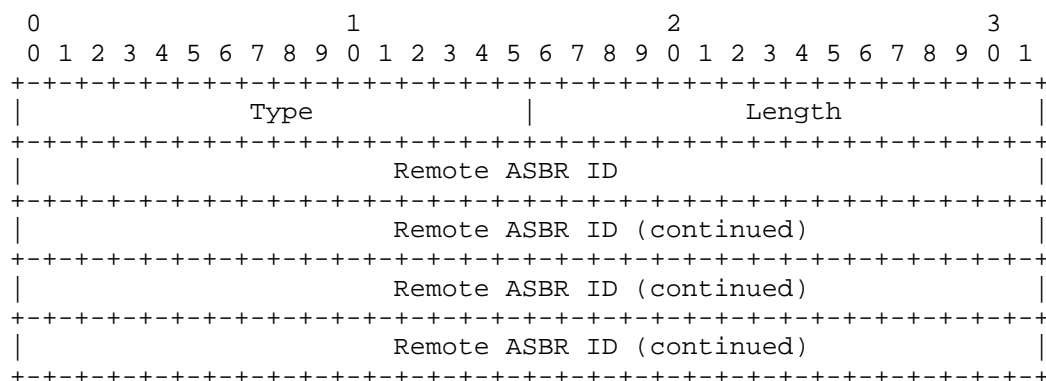


Figure 6: IPv6 Remote ASBR ID TLV Format

The IPv6 remote ASBR ID TLV MUST be included if the neighboring ASBR has an IPv6 address. If the neighboring ASBR does not have an IPv6 address, the IPv4 remote ASBR ID TLV MUST be included instead. An IPv4 remote ASBR ID TLV and IPv6 remote ASBR ID TLV MAY both be present in an inter-AS stub link NLRI.

8. Topology Reconstruction.

When SDN controller gets such information from BGP-LS protocol, it should find information from the associated router. Based on this information it can create a logical topology that contains the link between these two border routers. Iterating the above procedures for all of the stub links, the SDN controller can automatically retrieve the Inter-AS connection topology.

9. Security Considerations

BGP-LS security is described in [RFC9552]. This addition to BGP-LS focuses on the case when one network operated by a single entity has several IGP domains that are composited by its backbone network and several MANs (Metro-Area- Networks) and Internet Data Centers (IDCs). The configuration of these networks operated by the single administrative entity creates a "walled garden". Within this single Administrative Domain, the network operator needs to monitor and engineer traffic flows that traverse such a network that spans multiple Autonomous Systems. The network operators can obtain this information on inter-as topology via the process described in this draft.

The single Administrative domain consisting of two AS passing information about the Stub-link characteristics does not cause problems within a "walled garden". However, the Stub-link NLRI and characteristics (Link/local identifier, IPv4 Interface Address, IPv4

Neighbor Address, IPv6 interface Address, IPv6 Neighbor Address, Multi-topology Identifier, Remote-AS Number, IPv4 Remote ASBRI ID, and IPv6 Remote ASBR ID) are critical information to a network. As such, operators should handle this critical information in a way that restricts it to the walled garden.

10. IANA Considerations

This document defines:

- * A new BGP NLRI Type: Stub Link NLRI. The codepoint is from the "BGP-LS NLRI Types"
- * Three new Link Descriptors TLV: Remote AS Number TLV, IPv4 Remote ASBR ID, IPv6 Remote ASBR ID. The codepoint are from "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" registry.

10.1. New BGP-LS NLRI type

This document defines a new value in the registry "BGP-LS NLRI Types":

Code Point	Description	Status
7	Stub Link NLRI	Allocation from IANA

Figure 7: Stub Link NLRI Codepoint

10.2. New Stub Link Descriptors

This document defines three new values in the registry "BGP-LS NLRI and Attribute TLVs":

Code Point	Description	Status
270	Remote AS Number	Allocation from IANA
271	IPv4 Remote ASBR ID	Allocation from IANA
272	IPv6 Remote ASBR ID	Allocation from IANA

Figure 8: BGP-LS Link Descriptors TLV

11. Acknowledgement

The author would like to thank Susan Hares, Acee Lindem, Jie Dong, Shaowen Ma, Jeff Tantsura and Dhruv Dhody for their valuable comments and suggestions.

12. References

12.1. 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>>.
- [RFC5392] Chen, M., Zhang, R., and X. Duan, "OSPF Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering", RFC 5392, DOI 10.17487/RFC5392, January 2009, <<https://www.rfc-editor.org/info/rfc5392>>.
- [RFC9346] Chen, M., Ginsberg, L., Previdi, S., and D. Xiaodong, "IS-IS Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering", RFC 9346, DOI 10.17487/RFC9346, February 2023, <<https://www.rfc-editor.org/info/rfc9346>>.
- [RFC9552] Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <<https://www.rfc-editor.org/info/rfc9552>>.

12.2. Informative References

- [RFC9086] Previdi, S., Talaulikar, K., Ed., Filsfils, C., Patel, K., Ray, S., and J. Dong, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing BGP Egress Peer Engineering", RFC 9086, DOI 10.17487/RFC9086, August 2021, <<https://www.rfc-editor.org/info/rfc9086>>.

Authors' Addresses

Aijun Wang
China Telecom
Beiqijia Town, Changping District
Beijing
Beijing, 102209
China
Email: wangaj3@chinatelecom.cn

Huaimo Chen
Individual
Boston, MA
United States of America
Email: hchen.ietf@gmail.com

Ketan Talaulikar
Cisco Systems
India
Email: ketant.ietf@gmail.com

Shunwan Zhuang
Huawei Technologies
Huawei Building, No.156 Beiqing Rd.
Beijing
100095
China
Email: zhuangshunwan@huawei.com

Changwang Lin
New H3C Technologies
Huawei Building, No.156 Beiqing Rd.
Beijing
China
Email: linchangwang.04414@h3c.com