IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: 17 August 2026

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

13 February 2026

BGP-LS Extension for Inter-AS Topology Retrieval
draft-ietf-idr-bgpls-inter-as-topology-ext-21

Abstract

This document specifies the procedure for distributing Border Gateway Protocol-Link State (BGP-LS) key parameters for inter-domain links between two Autonomous Systems (ASes). It defines a new type within the BGP-LS Network Layer Reachability Information (NLRI) for a Stub Link, as well as three new type-length-values (TLVs) for the BGP-LS Link descriptor. These BGP-LS extensions enable Software-Defined Networking (SDN) controllers to automatically retrieve network topology across diverse inter-AS environments.

These extensions and procedures allow network operators to collect inter-domain interconnect information and automatically compute the end-to-end network topology using information provided by the 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 17 August 2026.

Table of Contents

1. Introduction

BGP-LS [RFC9552] specifies the methodology for using the BGP protocol to transfer Link-State information. This method enables SDN controllers to automatically collect underlay network topology, but it typically retrieves only information within a single Interior Gateway Protocol (IGP) domain. If an operator manages multiple IGP domains that interconnect with one another, no mechanism exists within the current BGP-LS protocol to transfer inter-domain topology information..

[RFC9086] defines extensions for exporting BGP peering node topology information (including peers, interfaces, and peering ASes) in a manner exploitable for computing efficient BGP Peering Engineering policies and strategies. This information can also be used to compute interconnection topology among different IGP domains, but it requires every border router to run the BGP-LS protocol and report such information to SDN controllers. Given the large number of border routers at the network boundary, this solution limits deployment flexibility.

This document analyzes scenarios in which SDN controllers require inter-domain topology information between different Autonomous Systems (ASes). After describing these scenarios, this document defines a new Stub Link type within the BGP-LS NLRI [RFC9552] to describe inter-AS links and new TLVs for this new BGP-LS type. The SDN controller can then automatically deduce the multi-domain topology using information from the 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 discussed in this document. Typically, the SDN Controller can retrieve the topology of IGP A and IGP B individually via the BGP-LS protocol, but it cannot obtain topology connection information between these two IGP domains, as IGP protocols are generally not run on the connected links.

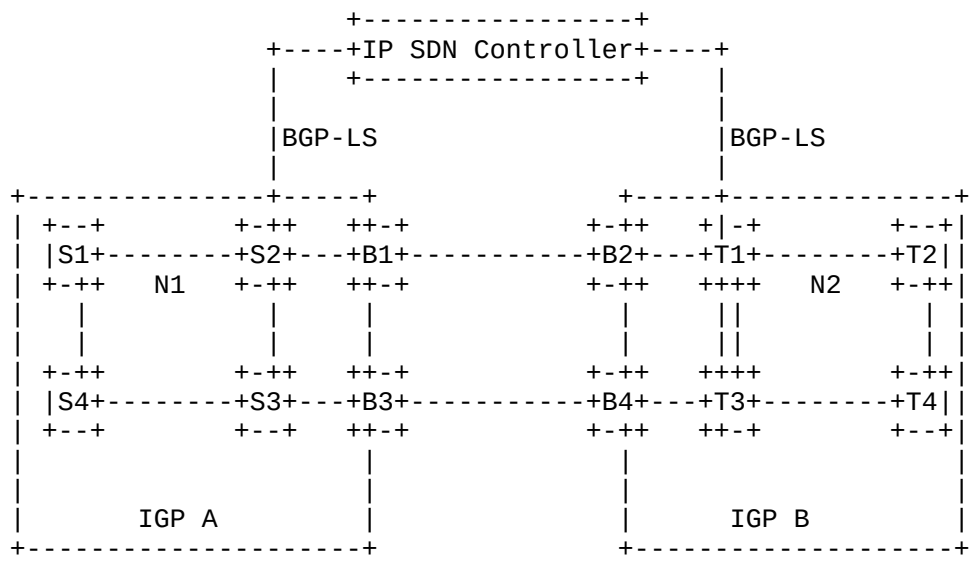

Figure 1: Inter-AS Domain Scenarios

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) for transferring topology and prefix information. For inter-AS links, the two ends of a link reside in different IGP domains; thus, it is not appropriate to transfer their information using the currently defined NLRI types.

This document defines a new NLRI type 7, see__Section 9) within the BGP Link-State (BGP-LS) NLRI, referred to as the Stub Link NLRI. The Stub Link NLRI is encoded in the format shown in Figure 2 and is explained below:

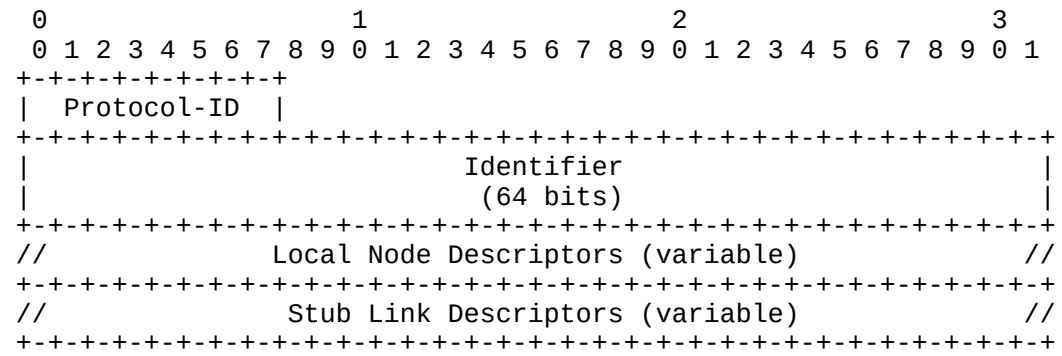

Figure 2: Stub Link NLRI Format

The "Protocol-ID" SHOULD be set to the value indicating the source protocol of the stub link information, as specified in[RFC9552] Section 5.2.

Local Node Descriptors: define the ASBRs attached to the inter-AS stub links, and use the "Local Node Descriptor" specified in[RFC9552] , 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 Links that have only one end located in an IGP domain, using the "Link Descriptor definition" specified in[RFC9552] ,Section 5.2.2 with the exceptions noted below.

The Stub Link Descriptors support the inclusion of the following sub-TLVs:

- 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])
- Remote-AS Number (TLV 270, [This document], section Section 7.1)
- IPv4 Remote ASBR ID (TLV 271, [This document], section Section 7.2)
- IPv6 Remote ASBR ID (TLV 272, [This document], section Section 7.3)

This newly defined NLRI can be used to describe links that have only one end located within an IGP domain, as described in the following sections.

To match the stub link connecting the two bordering Autonomous Systems (ASes), the Remote AS Number TLV, the IPv4 Remote ASBR ID TLV (if the neighboring ASBR has an IPv4 address), and/or the IPv6 Remote ASBR ID TLV (if the neighboring ASBR has an IPv6 address) MUST be included.

Correspondingly, the IPv4 Router-ID of Local Node TLV (1028) and/or the IPv6 Router-ID of Local Node TLV (1029) MUST be included as Link Attribute TLVs of the Stub Link NLRI.

The Node and Link Descriptor sub-TLVs, as well as Node and Link attributes defined in[RFC9552] MAY be included in the NLRI if necessary. The interface and neighbor address sub-TLVs SHOULD be included in the Local Node Descriptors to differentiate parallel links between two ASBRs.

6. IGP Information for Inter-AS Link

[RFC9346] and [RFC5392] define IS-IS and OSPF extensions, respectively, to address the requirements for reporting inter-AS links. Three sub-TLVs related 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 automatically flooded within an IGP domain. This document specifies that these MAY also be carried within the newly defined Stub Link NLRI in the BGP-LS protocol, as descriptors for inter-AS stub links.

7. BGP-LS Extensions for Inter-AS Link

This section proposes adding three new TLVs to be supported within the Stub Link NLRI of the BGP-LS NLRI. These new TLVs enable BGP-LS to transfer inter-AS information collected by the SDN controller.

The following Link Descriptor TLVs are added to 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

The detailed encoding of these TLVs is synchronized with the corresponding sections in[RFC9346] and [RFC5392], which maintains BGP-LS protocol agnosticism to the underlying protocol.

7.1. Remote AS Number TLV

A new TLV, referred to as 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 9) and is 4 octets in length. Its format is as follows:

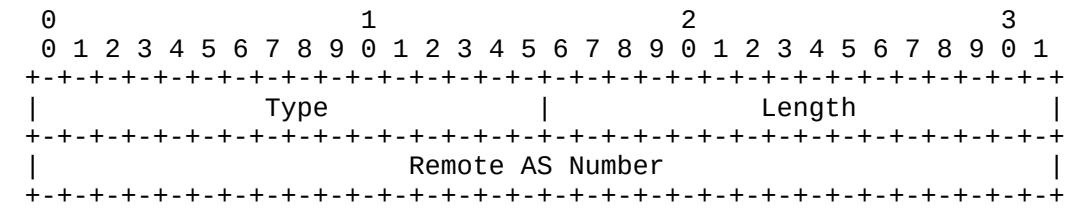

Figure 4: Remote AS Number TLV Format

The Remote AS Number field has 4 octets. When only 2 octets are used for the AS number (for example, when such information is advertised from OSPF, 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, 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 can be any stable, routable IPv4 address of the remote ASBR. The 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 9) and is 4 octets in length. Its format is as follows:

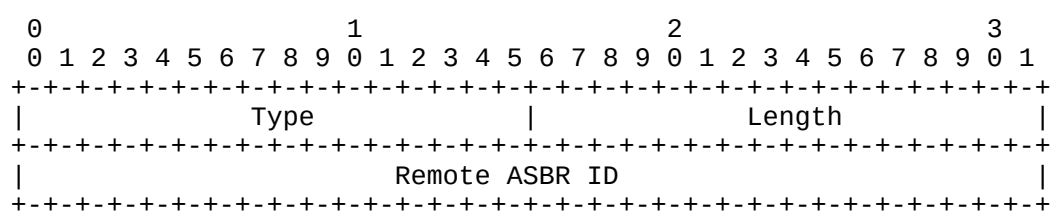

Figure 5: IPv4 Remote ASBR ID TLV Format

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 (including no IPv4 TE Router ID), the IPv6 Remote ASBR ID TLV MUST be included instead. Both an IPv4 Remote ASBR ID TLV and an IPv6 Remote ASBR ID TLV MAY be present in an inter-AS Stub Link NLRI.

7.3. IPv6 Remote ASBR ID

A new TLV, 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 can be any stable, routable IPv6 address of the remote ASBR. The 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 9) and is 16 octets in length. Its format is as follows:

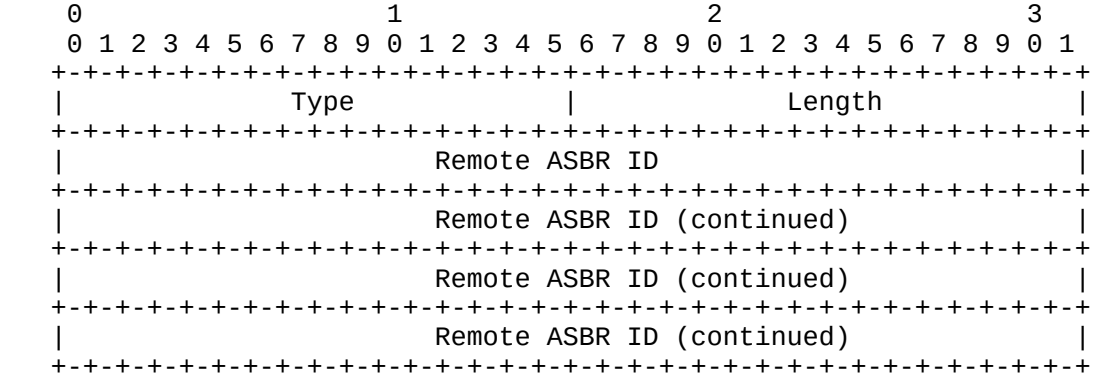

Figure 6: IPv6 Remote ASBR ID TLV Format

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. Both an IPv4 Remote ASBR ID TLV and an IPv6 Remote ASBR ID TLV MAY be present in an inter-AS Stub Link NLRI.

8. Security Considerations

BGP-LS security is specified in [RFC9552]. This extension to BGP-LS focuses on scenarios where a single entity-operated network includes multiple IGP domains composed of its backbone network, several Metro-Area Networks (MANs), and Internet Data Centers (IDCs). The configuration of these networks, operated by a single administrative entity, creates a "walled garden". Within this single administrative domain, the network operator needs to monitor and engineer traffic flows traversing a network that spans multiple Autonomous Systems (ASes). The network operator can obtain this inter-AS topology information via the procedure described in this document.

A single administrative domain consisting of two ASes that passes information about Stub Link characteristics does not cause issues within a "walled garden". However, the Stub Link NLRI and its characteristics (Link/Local Identifier, IPv4 Interface Address, IPv4 Neighbor Address, IPv6 Interface Address, IPv6 Neighbor Address, Multi-Topology Identifier, Remote-AS Number, IPv4 Remote ASBR ID, and IPv6 Remote ASBR ID) constitute critical network information. As such, operators SHOULD handle this critical information in a manner that restricts it to the walled garden.

9. 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.

9.1. 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

9.2. 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

10. 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.

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>.

[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>.

11.2. 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