

Workgroup: LSR Working Group

Internet-Draft:

draft-wang-lsr-stub-link-attributes-08

Published: 24 November 2023

Intended Status: Standards Track

Expires: 27 May 2024

Authors: A. Wang Z. Hu G. Mishra
 China Telecom Huawei Technologies Verizon Inc.
 J. Sun
 ZTE Corporation

Advertisement of Stub Link Attributes

Abstract

This document describes the mechanism that can be used to advertise the stub link attributes within the IS-IS or OSPF domain.

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 27 May 2024.

Copyright Notice

Copyright (c) 2023 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. Conventions used in this document](#)
- [3. Consideration for Identifying Stub Link](#)
- [4. OSPF Protocol Extension for Stub Link Attributes](#)
 - [4.1. OSPF Stub-Link TLV](#)
 - [4.2. OSPF Stub Link IPv4 Prefix Sub-TLV](#)
 - [4.3. OSPF Stub Link IPv6 Prefix Sub-TLV](#)
- [5. IS-IS Protocol Extension for Stub Link Attributes](#)
 - [5.1. IS-IS Stub-link TLV](#)
 - [5.2. IS-IS Stub Link IPv4 Prefix Sub-TLV](#)
 - [5.3. IS-IS Stub Link IPv6 Prefix Sub-TLV](#)
- [6. Security Considerations](#)
- [7. IANA Considerations](#)
- [8. Acknowledgement](#)
- [9. References](#)
 - [9.1. Normative References](#)
 - [9.2. Informative References](#)
- [Appendix A. Applied Scenarios](#)
 - [A.1. Inter-AS topology recovery](#)
 - [A.2. Egress Engineering for Anycast Servers](#)
 - [A.3. Egress Engineering for the path to BGP Next-hop](#)
- [Authors' Addresses](#)

1. Introduction

Stub links are used commonly within enterprise or service provider networks. One common use case is the inter-AS routing scenario where there are no IGP adjacencies between the adjacent BGP domains, another use case is at the network boundary that the interfaces are used to connect to the application servers.

For operators that have multiple ASes interconnecting with each other via the stub links, there is a requirement to obtain the inter-AS topology information as described in [\[I-D.ietf-idr-bgp-ls-inter-as-topology-ext\]](#). To achieve such goal, it is required that the BGP-LS to be enabled on every router that has the stub links, which is challenging for the network operation. It is desirable to advertise the stub link info into the IGP to ease the deployment of BGP-LS on any router in the IGP domain.

For stub links that are used to connect the servers, knowing the status of these stub links can facilitate the routers within the IGP to accomplish TE tasks in some scenarios.

But OSPF and IS-IS have no capability to identify such stub links and their associated attributes now.

This document defines the protocol extension for OSPFv2/v3 and IS-IS to indicate the stub links and their associated attributes.

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 [[RFC2119](#)] .

3. Consideration for Identifying Stub Link

OSPF[[RFC5392](#)] defines the Inter-AS-TE-v2 LSA and Inter-AS-TE-v3 LSA to carry the TE information about inter-AS links. IS-IS[[RFC9346](#)] defines the Inter-AS Reachability TLV to carry the TE information about inter-AS links. But they are normally being used under RSVP-TE, especially inter-domain RSVP-TE scenarios. As illustrated in the potential scenarios that described in [Appendix A](#), there is still the need for a generic solution which also covers non inter-AS stub links.

Then, to solve the problems that described in the applied scenarios, this document defines the Stub-Link TLV and some additional sub-TLVs to identify the stub link and transmit the associated attributes for OSPF and IS-IS respectively.

4. OSPF Protocol Extension for Stub Link Attributes

The following sections define the protocol extension to indicate the stub link and its associated attributes in OSPFv2/v3.

4.1. OSPF Stub-Link TLV

This document defines the Stub-Link TLV to describe stub link of a single router. This Stub-Link TLV is only applicable to the Inter-AS-TE-v2 LSA and Inter-AS-TE-v3 LSA [[RFC5392](#)]. Inclusion in other LSAs MUST be ignored.

The OSPF Stub-Link TLV which is under the IANA codepoint "Top Level Types in TE LSAs" has the following format:

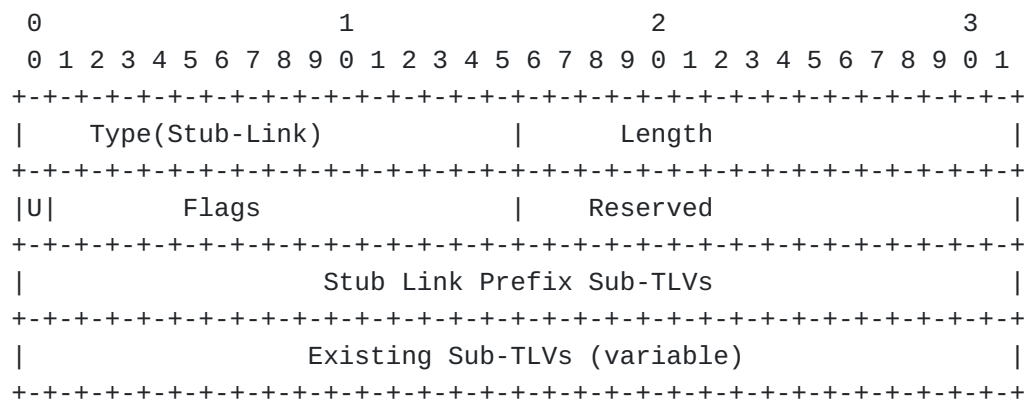


Figure 1: OSPF Stub-Link TLV

Type: The TLV type. The value is 7(TBD) for OSPF Stub-Link.

Length: Variable, dependent on sub-TLVs

Flags: Define the type of the stub-link:

*U bit(bit 0): Identify the unnumbered stub link if this bit is set.

*bit 1-15: Reserved

Stub Link Prefix Sub-TLV: The prefix of the stub-link. It's format is defined in [Section 4.2](#) and [Section 4.3](#).

Existing Sub-TLVs: Sub-TLV that defined within "Open Shortest Path First (OSPF) Traffic Engineering TLVs" for TE Link TLV(Value 2) can be included if necessary.

If the stub-link is identified as unnumbered stub link (U bit is set), then the "Remote IPv4 Address Sub-TLV" or "Remote Interface IPv6 Address Sub-TLV", which should be set to the identifier value of remote router, SHOULD be included to facilitate the pairing of inter-AS link.

If this TLV is advertised multiple times in the same Inter-AS-TE-v2/v3 LSA, only the first instance of the TLV is used by receiving OSPFv2/v3 routers. This situation SHOULD be logged as an error.

If this TLV is advertised multiple times for the same link in different Inter-AS-TE-v2/v3 LSA originated by the same OSPFrouter, the OSPFStub-Link TLV in these LSAs with the smallest Opaque ID is used by receiving OSPF routers. This situation may be logged as a warning.

It is RECOMMENDED that OSPF routers advertising OSPF Stub-Link TLVs in different OSPF Inter-AS-TE v2/v3 LSAs re-originate these LSAs in ascending order of Opaque ID to minimize the disruption.

This document creates a registry for Stub-Link attributes in [Section 7](#).

4.2. OSPF Stub Link IPv4 Prefix Sub-TLV

The OSPF Stub Link IPv4 Prefix Sub-TLV has the following format:

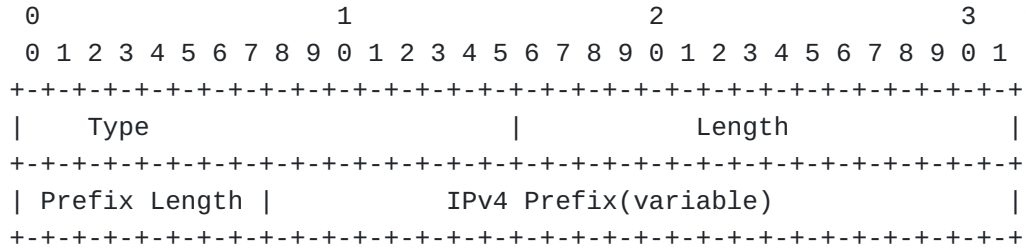


Figure 2: OSPF Stub Link IPv4 Prefix Sub-TLV

Type: IPv4 Prefix Sub-TLV codepoint. Value is 37(TBD) for OSPF(under "Types for sub-TLVs of TE Link TLV (Value 2)")

Length: The length of the value portion in octets.

Prefix Length: the length of the IPv4 Prefix in bits.

IPv4 Prefix: The IPv4 Prefix value of stub link.

4.3. OSPF Stub Link IPv6 Prefix Sub-TLV

The OSPF Stub Link IPv6 Prefix Sub-TLV has the following format:

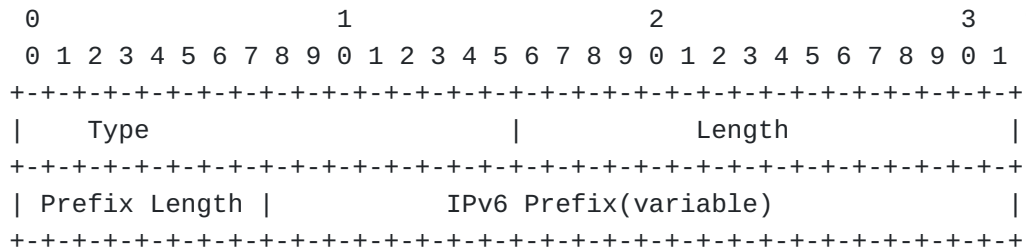


Figure 3: OSPF Stub Link IPv6 Prefix Sub-TLV

Type: IPv6 Prefix Sub-TLV codepoint. Value is Value is 38(TBD) for OSPF(under "Types for sub-TLVs of TE Link TLV (Value 2)")

Length: The length of the value portion in octets.

IPv6 Prefix: The IPv6 Prefix value of stub link.

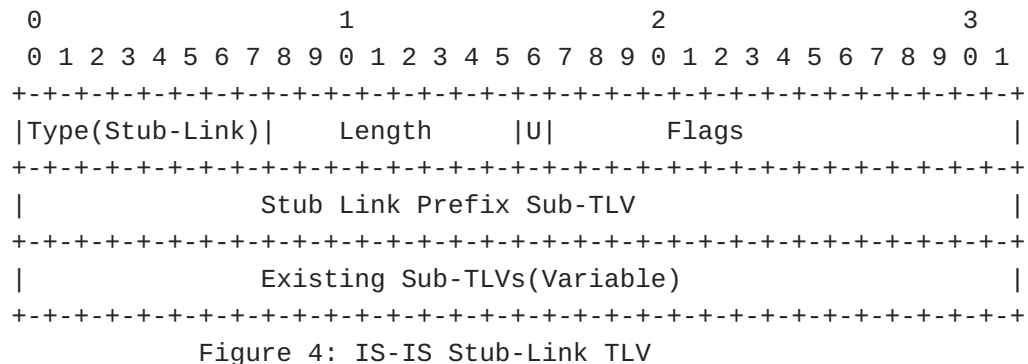
5. IS-IS Protocol Extension for Stub Link Attributes

The following sections define the protocol extension to indicate the stub link and its associated attributes in IS-IS.

5.1. IS-IS Stub-link TLV

This document defines the IS-IS Stub-Link TLV to describes stub link of a single router.

The IS-IS Stub-Link TLV has the following format:



Type: IS-IS TLV codepoint. Value is 151 (TBD) for stub-link TLV.

Length: Variable, dependent on sub-TLVs

Flags: Define the type of the stub-link:

*0: U bit(bit 0): Identify the unnumbered stub link if this bit is set.

*bit 1-15: Reserved

Stub Link Prefix Sub-TLV: The prefix of the stub-link. It's format is defined in [Section 5.2](#) and [Section 5.3](#).

Existing Sub-TLVs: Sub-TLVs that defined within "IS-IS Sub-TLVs for TLVs Advertising Neighbor Information" can be included if necessary.

If the stub-link is identified as unnumbered stub link type (U bit is set), then the "IPv4 Remote ASBR ID" or "IPv6 Remote ASBR ID" Sub-TLV SHOULD be included to facilitate the pairing of inter-AS link.

5.2. IS-IS Stub Link IPv4 Prefix Sub-TLV

The IS-IS Stub Link IPv4 Prefix Sub-TLV has the following format:

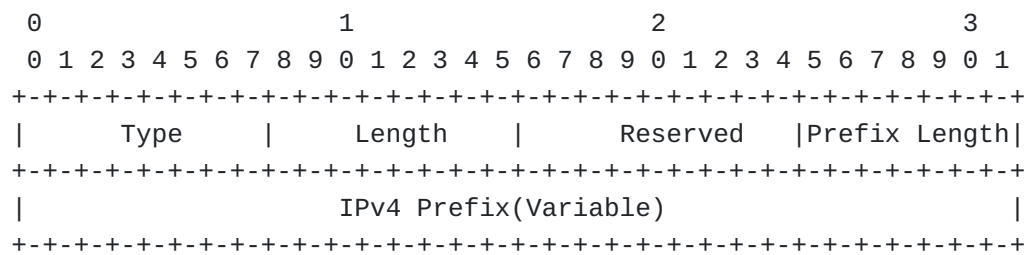


Figure 5: IS-IS Stub Link IPv4 Prefix Sub-TLV

Type: IPv4 Prefix Sub-TLV codepoint. Value is 46(TBD) for IS-IS(under "IS-IS Sub-TLVs for TLVs Advertising Neighbor Information")

Length: Length: The length of the value portion in octets.

Prefix Length: the length of the IPv4 Prefix in bits.

IPv4 Prefix: The IPv4 Prefix value of stub link.

5.3. IS-IS Stub Link IPv6 Prefix Sub-TLV

The IS-IS Stub Link IPv6 Prefix Sub-TLV has the following format:

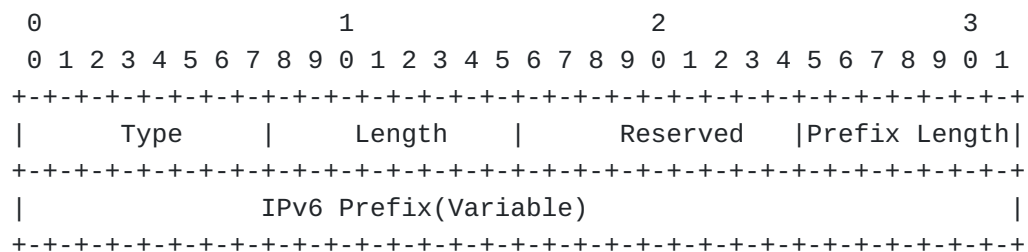


Figure 6: IS-IS Stub Link IPv6 Prefix Sub-TLV

Type: IPv6 Prefix Sub-TLV codepoint. Value is 47(TBD) for IS-IS(under "IS-IS Sub-TLVs for TLVs Advertising Neighbor Information")

Length: Length: The length of the value portion in octets.

Prefix Length: the length of the IPv6 Prefix in bits.

IPv6 Prefix: The IPv6 Prefix value of stub link.

6. Security Considerations

Security concerns for IS-IS are addressed in [[RFC5304](#)] and [[RFC5310](#)]

Security concern for OSPFv3 is addressed in [[RFC4552](#)]

Advertisement of the additional information defined in this document introduces no new security concerns.

7. IANA Considerations

IANA is requested to the allocation in following registries:

Registry	Type	Meaning
	(suggested)	
Top Level Types in TE LSAs	7	OSPF Stub-Link
IS-IS Top-Level TLV	151	IS-IS Stub-Lin
Types for sub-TLVs of TE Link TLV (Value 2)	37	OSPF Stub Link
Types for sub-TLVs of TE Link TLV (Value 2)	38	OSPF Stub Link
IS-IS Sub-TLVs for TLVs		
Advertising Neighbor Information	46	IS-IS Stub Lin
IS-IS Sub-TLVs for TLVs		
Advertising Neighbor Information	47	IS-IS Stub Lin

Figure 7: IANA Allocation for newly defined TLVs and Sub-TLVs

8. Acknowledgement

Thanks Ketan Talaulikar, Acee Lindem, Shunwan Zhang, Peter Psenak, Tony Li, Les Ginsberg, Dhruv Dhody, Jeff Tantsura and Robert Raszuk for their suggestions and comments on this idea.

9. References

9.1. 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>>.
- [RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", RFC 4552, DOI 10.17487/RFC4552, June 2006, <<https://www.rfc-editor.org/info/rfc4552>>.
- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October 2008, <<https://www.rfc-editor.org/info/rfc5304>>.
- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic

Authentication", RFC 5310, DOI 10.17487/RFC5310, February 2009, <<https://www.rfc-editor.org/info/rfc5310>>.

[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>>.

9.2. Informative References

[I-D.ietf-idr-bgppls-inter-as-topology-ext] Wang, A., Chen, H., Talaulikar, K., and S. Zhuang, "BGP-LS Extension for Inter-AS Topology Retrieval", Work in Progress, Internet-Draft, draft-ietf-idr-bgppls-inter-as-topology-ext-13, 3 April 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgppls-inter-as-topology-ext-13>>.

Appendix A. Applied Scenarios

The following sections describe the scenarios that knowing the stub link related attributes information can help solve the corresponding necessity in questions.

A.1. Inter-AS topology recovery

Figure 1 describes the scenario that the necessity of inter-AS topology recovery for Native IP point-to-point stub link scenario.

R10, R11 and R12 are located in AS1. R20, R21, R22 are located in AS2. The controller runs BGP-LS with R10 in AS1 and R20 in AS2 respectively.

There is one BGP session among the border router R11 and R21, which are connected by several stub links (passive interfaces) between them. The situation within the R21 and R22 are the same.

Since the links between the border routers are passive, there will be no IGP neighbors between them. The BGP-LS information carried in each AS will not report these stub links, and the controller can't recover the inter-AS topology automatically.

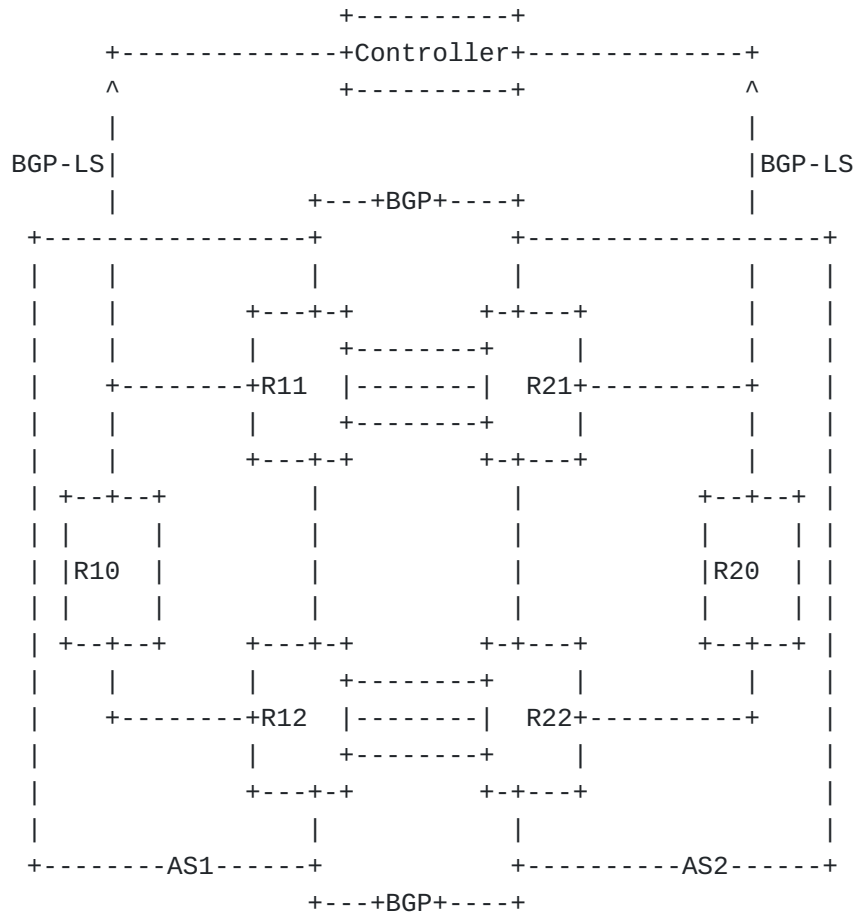


Figure 1: Inter-AS Topology Recovery(P2P Scenario)

Figure 2 describes the similar situation but in LAN environment. The border routers of AS1, AS2 and AS3 are connected via one LAN interfaces(that is to say, the corresponding interfaces on R1, R2 and R3 are on the same subnet). There are three different BGP sessions from the loopback address of the border routers among them respectively. It is necessary to recovery the underlying inter-AS topology automatically.

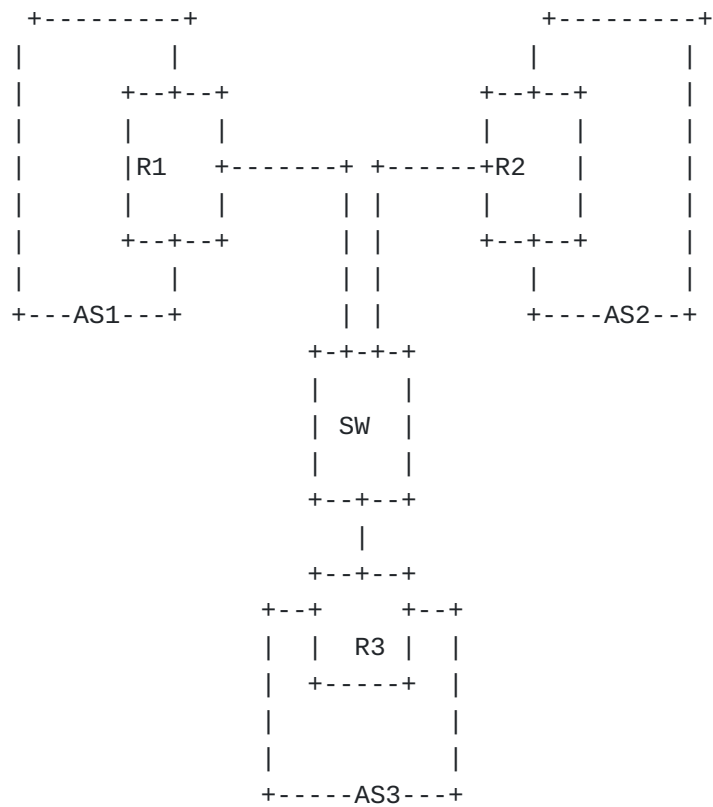


Figure 2: Inter-AS Topology Recovery(LAN Scenario)

A.2. Egress Engineering for Anycast Servers

Figure 3 describes the scenario that the stub link information can be used for egress engineering for Anycast servers that connected to the network. In the example, the R1, R2 and R3 are border routers which are connected directly the server S1, S2 and S3 that have the same IP address IPa. The characteristics of the stub links that connected to these Anycast servers are different. It will be help for the router R0, to know the attributes of the stub links and select the optimal Anycast server to serve the customer's application.

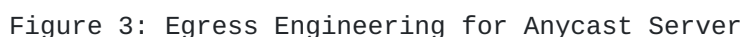


Figure 4 describes the scenario that the stub link information can facilitate the selection of path to BGP next hop. The router R10 and R20 which are located in different AS establish the BGP session directly. The attributes of the stub links among the border routers are vary. It is certainly will be helpful for the router R10 and R20 to select the optimized forwarding path to BGP next hop, that is via the stub links among them, to reach each other.

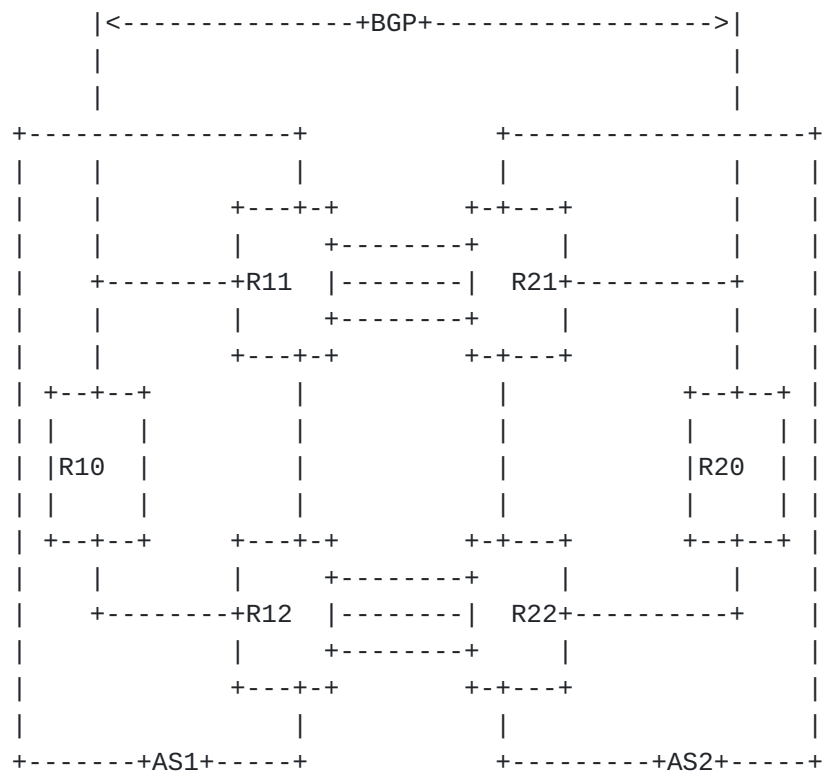


Figure 4: Optimized BGP next hop selection

Authors' Addresses

Aijun Wang
 China Telecom
 Beiqijia Town, Changping District
 Beijing
 102209
 China

Email: wangaj3@chinatelecom.cn

Zhibo Hu
 Huawei Technologies
 Huawei Bld., No.156 Beiqing Rd.
 Beijing
 100095
 China

Email: huzhibo@huawei.com

Gyan S. Mishra
 Verizon Inc.
 13101 Columbia Pike
 Silver Spring, MD 20904
 United States of America

Email: gyan.s.mishra@verizon.com

Jinsong Sun
ZTE Corporation
No. 68, Ziiijnhua Road
Nan Jing
210012
China

Email: sun.jinsong@zte.com.cn