

IDR Working Group
Internet-Draft
Updates: 8669 (if approved)
Intended status: Standards Track
Expires: 19 September 2026

W. Jiang
Y. Liu
China Mobile
S. Zhuang
Huawei Technologies
G. Mishra
Verizon Communications Inc.
S. Chen
Huawei Technologies
18 March 2026

Traffic Steering using BGP FlowSpec with SR Policy
draft-ietf-idr-ts-flowspec-srv6-policy-10

Abstract

BGP Flow Specification version 1 (FSv1), defined in [RFC8955], [RFC8956] and [RFC9117] has been proposed to distribute BGP [RFC4271] FlowSpec NLRI to FlowSpec clients to mitigate (distributed) denial-of-service attacks, and to provide traffic filtering in the context of a BGP/MPLS VPN service. Recently, traffic steering applications in the context of SR-MPLS and SRv6 using FlowSpec are being used in networks. This document introduces the usage of BGP FSv1 to steer packets into an SR Policy.

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 19 September 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. Definitions and Acronyms	3
3. Operations	4
4. SR-MPLS Application Examples	5
5. SRv6 Application Examples	6
6. Error Handling	8
7. Running Code	9
7.1. Interop-test Status	10
7.2. Deployment Status	11
8. IANA Considerations	11
9. Security Considerations	11
10. Contributors	11
11. Acknowledgements	12
12. References	12
12.1. Normative References	12
12.2. Informative References	14
Authors' Addresses	14

1. Introduction

SR-MPLS [RFC8660] forwards data packets using the source routing model. The core idea of SR-MPLS is to divide a packet forwarding path into different segments, allocate segment identifiers (SIDs) to the segments, and encapsulate segment information into packets at the ingress of the path to guide packet forwarding.

Segment Routing IPv6 (SRv6) is a protocol designed to forward IPv6 data packets on a network using the source routing model. SRv6 enables the ingress network device to add a segment routing header (SRH) [RFC8754] to an IPv6 packet and push an explicit IPv6 address stack into the SRH. After receiving the packet, each transit node updates the IPv6 destination IP address in the packet and segment list to implement hop-by-hop forwarding.

SR Policy (includes SR-MPLS and SRv6 Policy) [RFC9256] is a tunneling technology based on SR-MPLS and SRv6. An SR Policy is a set of candidate paths consisting of one or more segment lists, that is, segment ID (SID) lists. Each SID list identifies an end-to-end path from the source node to the destination node, instructing a network device to forward traffic through the path rather than the shortest path computed using an IGP. The header of a packet steered into an SR Policy is augmented with an ordered list of segments associated with that SR Policy, so that other devices on the network can execute the instructions encapsulated into the list.

The headend of an SR Policy may learn multiple candidate paths for an SR Policy. Candidate paths may be learned via a number of different mechanisms, e.g., CLI, NetConf, PCEP [RFC9862], or BGP [RFC9830].

[RFC8955], [RFC8956] and [RFC9117] define the BGP Flow Specification version 1 (FSv1) that allows conveying Flow Specifications and traffic Action/Rules associated (rate-limiting, redirect, remark ...). BGP Flow Specifications are encoded within the MP_REACH_NLRI and MP_UNREACH_NLRI attributes [RFC4760]. Rules (Actions associated) are encoded in Extended Community attribute [RFC4360]. The BGP Flow Specification function allows BGP Flow Specification routes that carry traffic policies to be transmitted to BGP Flow Specification peers to steer traffic.

This document proposes a new usage for BGP FSv1 in order to steer traffic into an SR Policy. This work is helpful for promoting the deployment of SR-MPLS and SRv6 networks. BGP Flow Specification version 2 (FSv2) defined in [I-D.ietf-idr-fsv2-ip-basic], how to use FSv2 is outside the scope of this document and will be defined in another document.

2. Definitions and Acronyms

- * FlowSpec: Flow Specification
- * FSv1: Flow Specification version 1
- * FSv2: Flow Specification version 2

- * SR: Segment Routing
- * SR-MPLS: SR over the MPLS data plane
- * SRv6: SR over the IPv6 data plane
- * SID: Segment Identifier
- * SRH: Segment Routing Header
- * USD: Ultimate Segment Decapsulation

3. Operations

An SR Policy [RFC9256] is identified through the tuple <Headend, Color, Endpoint>. In the context of a specific Headend, one may identify an SR Policy by the <Color, Endpoint> tuple. The Headend is the node where the SR Policy is instantiated/implemented. The Headend is specified as an IPv4 or IPv6 address and is expected to be unique in the domain. The Endpoint indicates the destination of the SR Policy. The Endpoint is specified as an IPv4 or IPv6 address and is expected to be unique in the domain. The Color is a 32-bit unsigned numerical value that associates with the SR policy, and it defines an application-level network Service Level Agreement (SLA) policy or intent.

The Color and Endpoint are used to automate the steering of BGP service routes on an SR Policy as described in Section 8 of [RFC9256].

[I-D.ietf-idr-flowspec-redirect-ip] defines the redirect to IPv4 and IPv6 Next-hop action. The IPv4 next-hop address in the Flow-spec Redirect to IPv4 Extended Community and the IPv6 next-hop address in the Flow-spec Redirect to IPv6 Extended Community [RFC5701] can be used to specify the endpoint of the SR Policy.

When the packets reach the TailEnd device, some specific function information identifiers can be used to decide how to further process the flows in SRv6 scenarios. Several endpoint functions are already defined, e.g., End.DT6: Endpoint with decapsulation and IPv6 table lookup, and End.DX6: Endpoint with decapsulation and IPv6 cross-connect. The BGP Prefix-SID defined in [RFC8669] is utilized to enable SRv6 VPN services [RFC9252]. SRv6 Services TLVs within the BGP Prefix-SID Attribute can be used to indicate the endpoint functions. This document extends the use of the BGP Prefix-SID attribute [RFC8669] to carry SRv6 SIDs and their associated information with the BGP address families that are defined in [RFC8955] and [RFC8956], where applicable, as described in Section 5.

For SR-MPLS scenarios, this document proposes carrying the Color Extended Community and the Flow-spec Redirect to IPv4/IPv6 Extended Community in the context of a Flowspec NLRI [RFC8955] [RFC8956] to an SR-MPLS Headend to steer traffic into one SR-MPLS Policy.

For SRv6 scenarios, this document proposes carrying the Color Extended Community, the Flow-spec Redirect to IPv6 Extended Community and BGP Prefix-SID Attribute in the context of a Flowspec NLRI [RFC8955] [RFC8956] to an SRv6 Headend to steer traffic into one SRv6 Policy, as well as to indicate specific Tailend functions. The USD (Ultimate Segment Decapsulation) flavor allows the last Segment Endpoint Node to decapsulate the IPv6 header. If the last SID in the SR list is a USD-flavored SRv6 SID, the BGP Prefix-SID Attribute is optional, otherwise, the BGP Prefix-SID attribute MUST be carried to convey a USD-flavored SRv6 SID.

For the case where a FlowSpec route carries multiple Color Extend Communities, the Color Extended community with the highest numerical value will be given higher preference per the description in Section 8.4.1 of [RFC9256].

The method proposed in this document supports load balancing to the tailend device. To achieve that, the Headend device can set up multiple candidate paths in one SR Policy, and use a FlowSpec route to indicate the specific SR Policy.

4. SR-MPLS Application Examples

In the following scenarios shown in Figure 1, BGP FlowSpec Controller signals the filter rules, the Flow-spec Redirect to IPv4/IPv6 action, and the policy color to the SR-MPLS HeadEnd device.

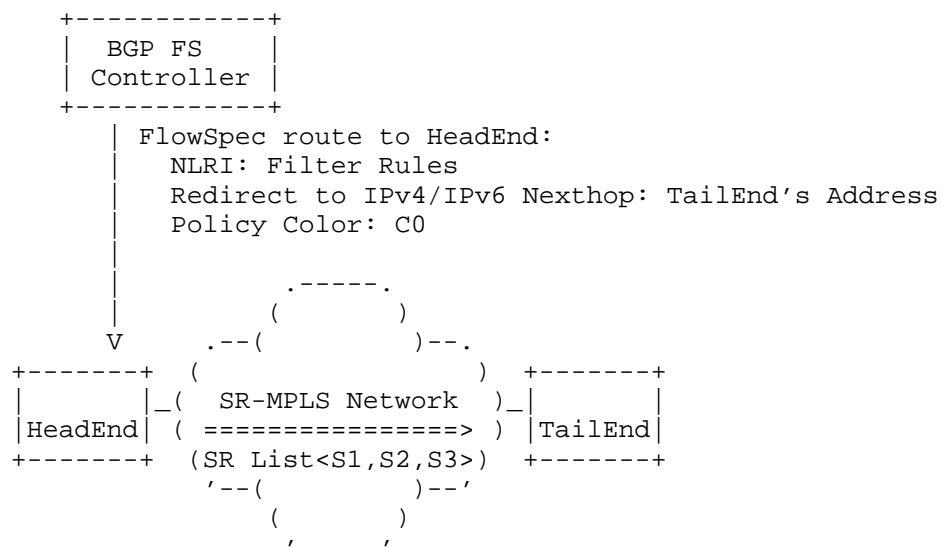


Figure 1: Steering the Traffic Flow into SR-MPLS Policy

When the SR-MPLS HeadEnd device (as a FlowSpec client) receives such instructions from BGP FS Controller, it will steer the traffic flows matching the criteria in the FlowSpec route into the SR-MPLS Policy matching the tuple (Endpoint: TailEnd's Address, Color: C0). And the packets of such traffic flows will be encapsulated with an MPLS stack using the SR List <S1, S2, S3> in the HeadEnd device, then send the packets to the TailEnd device along the path indicated by the SR list.

5. SRv6 Application Examples

In the following scenario shown in Figure 2, BGP FlowSpec Controller signals the filter rules, the redirect to IPv6 Nextthop action, the policy color and the function information (SRv6 SID: Service_id_x) to the HeadEnd device.

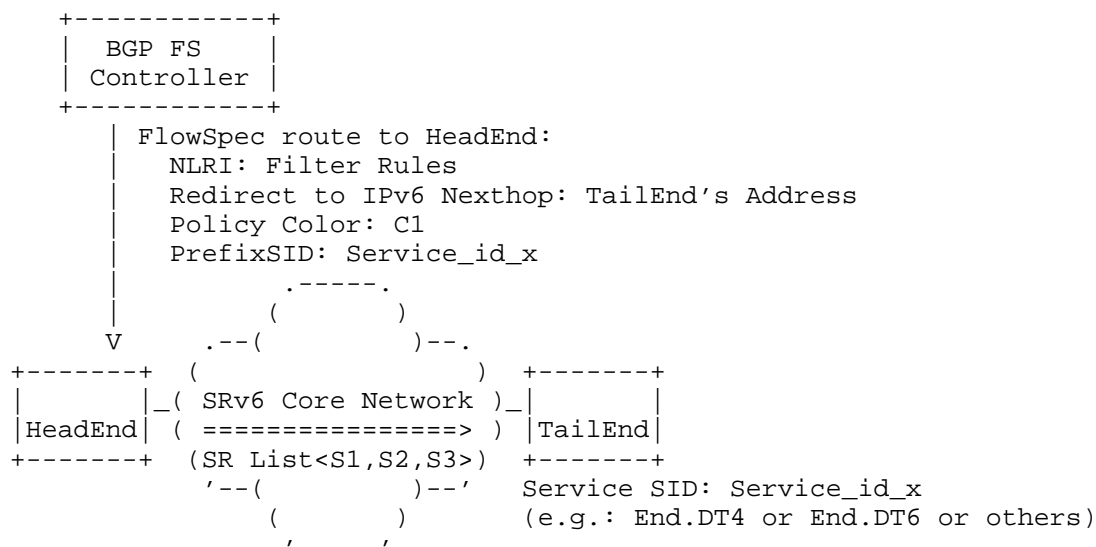
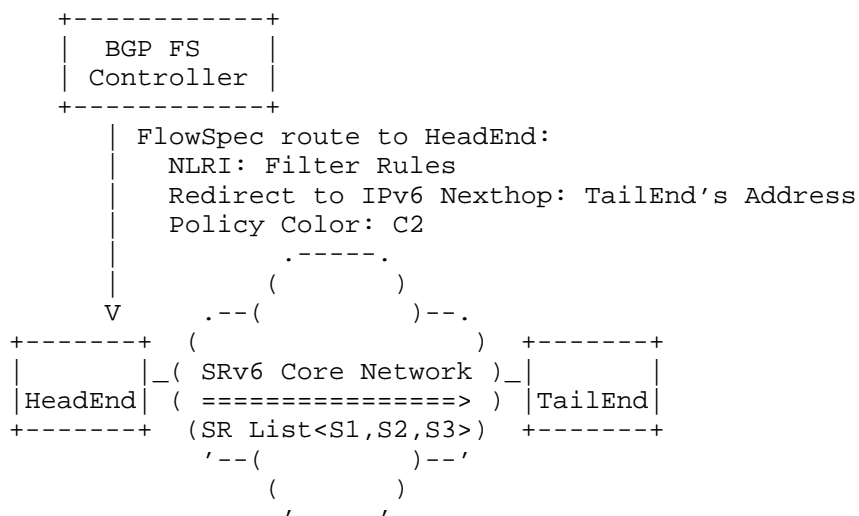


Figure 2: Steering the Traffic Flow into SRv6 Policy (Option 1)

When the HeadEnd device (as a FlowSpec client) receives such instructions from BGP FS Controller, it will steer the traffic flows matching the criteria in the FlowSpec route into the SRv6 Policy matching the tuple (Endpoint: TailEnd's Address, Color: C1). And the packets of such traffic flows will be encapsulated with an SRH (Segment Routing Header) using the SR List <S1, S2, S3, Service_id_x>. When the packets reach to the TailEnd device, they will be further processed per the function denoted by the Service_id_x.

When the HeadEnd device determines (with the help of SRv6 SID Structure) that the Service SID belongs to the same SRv6 Locator as the last SRv6 SID of the TailEnd device in the SRv6 Policy segment list, it MAY exclude that last SRv6 SID when steering the service flow. For example, the effective segment list of the SRv6 Policy associated with SID list <S1, S2, S3> would be replaced with <S1, S2, Service_id_x>.

If the last SRv6 SID (for example, we use S3 here) of the TailEnd device in the SRv6 Policy segment list is USD-flavored, an SRv6 Service SID (e.g., End.DT4 or End.DT6) is not required when a BGP FlowSpec Controller sends the FlowSpec route to the HeadEnd device (as a FlowSpec client). In the following scenario shown in Figure 3, BGP FlowSpec Controller signals the filter rules, the redirect to IPv6 Nexthop action and the policy color to the HeadEnd device.



Note: S3 MUST be a USD-flavored SRv6 SID of the TailEnd

Figure 3: Steering the Traffic Flow into SRv6 Policy (Option 2)

When the HeadEnd device (as a FlowSpec client) receives such instructions from a BGP FS Controller, it will steer the traffic flows matching the criteria in the Flowspec route into the SRv6 Policy matching the tuple (Endpoint: TailEnd's Address, Color: C2). And the packets of such traffic flows will be encapsulated with an SRH (Segment Routing Header) using the SR List <S1, S2, S3>. When the packets reach to the TailEnd device, they will be further processed per the function denoted by the USD-flavored SRv6 SID S3.

For the cases of intra-AS and inter-AS traffic steering using this method, the usages of Flowspec Color Extended Community with BGP prefix SID are the same for both scenarios. The difference lies between the local SRv6 policy configurations. For the inter-domain case, the operator can configure an inter-domain SRv6 policy/path at the Headend device. For the intra-domain case, the operator can configure an intra-domain SRv6 policy/path at the Headend device.

6. Error Handling

The error handling procedures depend on the results of the following:

- * a) The validation procedures of the redirect-to-IP Extended Community as per [I-D.ietf-idr-flowspec-redirect-ip]

- * b) The validation and selection procedures of the Color Extended Community as per [RFC9256], which determine a single color for steering
- * c) The validation procedures of the Prefix-SID attribute per section 6 of [RFC8669] if attached

After the above results are determined, perform the following error-handling procedures:

- * 1) If the Color Extended Community is invalid or is not attached, the actions defined in this document do not apply.
- * 2) If the redirect-to-IP Extended Community is invalid or is not attached, and there are other actions attached, the filter is further processed with those actions.
- * 3) If the redirect-to-IP Extended Community, the Color Extended Community, and the Prefix-SID attribute are attached and valid, the traffic flows are per-destination steered into the corresponding SR-Policy by the [RFC9256] procedures with the Service SID(defined in the Prefix-SID attribute), as illustrated in Sections 3 and 5. The HeadEnd SHOULD load-share the traffic flows across all the corresponding SR-Policies with the redirect-to-IP addresses as their Endpoints if there are multiple valid redirect-to-IP Extended Communities. If the HeadEnd is incapable of doing so, it SHOULD deterministically select one redirect-to-IP address as the Endpoint.
- * 4) If the redirect-to-IP Extended Community and the Color Extended Community are attached and valid, but the Prefix-SID attribute is invalid or is not attached, the traffic flows are per-destination steered into the corresponding SR-Policy by the [RFC9256] procedures without the Service SID, as illustrated in Sections 3, 4, and 5. The HeadEnd SHOULD load-share the traffic flows across all the corresponding SR-Policies with the redirect-to-IP addresses as their Endpoints if there are multiple valid redirect-to-IP Extended Communities. If the HeadEnd is incapable of doing so, it SHOULD deterministically select one redirect-to-IP address as the Endpoint.

7. Running Code

7.1. Interop-test Status

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this document. The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

The Traffic Steering using BGP FlowSpec with SR-MPLS / SRv6 Policy mechanism has been implemented on the following hardware devices, Network Operating System software, and SDN controllers. They have also successfully participated in the series of joint interoperability testing events hosted by China Mobile from July 2021 to October 2021. The following hardware devices and Network Operating System software had successfully passed the interoperability testing (in alphabetical order).

Routers:

Vendors	Device Model	Version
Huawei	NE40-X8A	NE40E V800R021C00SPC091T
New H3C	CR16010H-FA	Version 7.1.075, ESS 8305
Ruijie	RG-N8010-R	N8000-R_RGOS 12.8(1)B08T1
ZTE	M6000-8S Plus	V5.00.10(5.60.5)

Controllers:

Vendors	Device Model	Version
China Unitechs	I-T-E SC	V1.3.6P3
Huawei	NCE-IP	V100R021C00
Ruijie	RG-ONC-AIO-H	RG-ION-WAN-CLOUD_2.00T1
ZTE	ZENIC ONE	R22V16.21.20

7.2. Deployment Status

As of August 2022, this feature has been deployed on the IP backbone network of China Mobile.

China Mobile has fully transitioned to the SRv6 architecture, with no scenarios where SRv6 and SR-MPLSv6 coexist. As a result, when utilizing Color C and IPv6 redirect addresses, traffic steering is executed exclusively via SRv6 policies.

In scenarios where an operator supports the coexistence of SRv6 and SR-MPLSv6, it is recommended to adhere to [RFC9256] to govern policy selection for traffic steering based on Color C and IPv6 endpoints.

8. IANA Considerations

No IANA actions are required for this document. All code points will be used as originally assigned by IANA without any changes.

9. Security Considerations

This document does not change the security properties of SRv6 and BGP.

10. Contributors

The following people made significant contributions to this document:

Yunan Gu
Huawei Technologies
Email: gyunan@huawei.com

Haibo Wang
Huawei Technologies
Email: rainsword.wang@huawei.com

Jie Dong
Huawei Technologies
Email: jie.dong@huawei.com

Xue Yang
China Mobile
Email: yangxuewl@chinamobile.com

11. Acknowledgements

The authors would like to acknowledge the review and inputs from Jeffrey Haas, Susan Hares, Keyur Patel, Weiqiang Cheng, Kaliraj Vairavakkalai, Robin Li, Acee Lindem, Gunter Van De Velde, John Scudder, Rainbow Wu, Linda Dunbar, Gang Yan, Feng Yang, Wim Henderickx, Robert Raszuk, Ketan Talaulikar, Changwang Lin, Aijun Wang, Hao Li, Huaimo Chen, Sheng Fang, Yuanxiang Qiu, Ran Chen, Cheng Li, Zheng Zhang, Xuewei Wang, Yanrong Liang, Xuhui Cai, Haojie Wang, Lili Wang, Nan Geng, Stephane Litkowski and Zhenqiang Li.

Special thanks to Nat Kao, who suggested adding SR-MPLS use cases to this document and provided detailed feedback on the error handling section.

Special thanks to Donald E. Eastlake, 3rd, who thoroughly reviewed the entire document and made many useful suggestions for improvement.

12. References

12.1. Normative References

- [I-D.ietf-idr-flowspec-redirect-ip]
Haas, J., Henderickx, W., and A. Simpson, "BGP Flow-Spec Redirect-to-IP Action", Work in Progress, Internet-Draft, draft-ietf-idr-flowspec-redirect-ip-06, 2 March 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-redirect-ip-06>>.
- [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>>.
- [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>>.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <<https://www.rfc-editor.org/info/rfc4360>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.

- [RFC5701] Rekhter, Y., "IPv6 Address Specific BGP Extended Community Attribute", RFC 5701, DOI 10.17487/RFC5701, November 2009, <<https://www.rfc-editor.org/info/rfc5701>>.
- [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>>.
- [RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", RFC 8660, DOI 10.17487/RFC8660, December 2019, <<https://www.rfc-editor.org/info/rfc8660>>.
- [RFC8669] Previdi, S., Filsfils, C., Lindem, A., Ed., Sreekantiah, A., and H. Gredler, "Segment Routing Prefix Segment Identifier Extensions for BGP", RFC 8669, DOI 10.17487/RFC8669, December 2019, <<https://www.rfc-editor.org/info/rfc8669>>.
- [RFC8955] Loibl, C., Hares, S., Raszuk, R., McPherson, D., and M. Bacher, "Dissemination of Flow Specification Rules", RFC 8955, DOI 10.17487/RFC8955, December 2020, <<https://www.rfc-editor.org/info/rfc8955>>.
- [RFC8956] Loibl, C., Ed., Raszuk, R., Ed., and S. Hares, Ed., "Dissemination of Flow Specification Rules for IPv6", RFC 8956, DOI 10.17487/RFC8956, December 2020, <<https://www.rfc-editor.org/info/rfc8956>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.
- [RFC9117] Uttaro, J., Alcaide, J., Filsfils, C., Smith, D., and P. Mohapatra, "Revised Validation Procedure for BGP Flow Specifications", RFC 9117, DOI 10.17487/RFC9117, August 2021, <<https://www.rfc-editor.org/info/rfc9117>>.
- [RFC9252] Dawra, G., Ed., Talaulikar, K., Ed., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "BGP Overlay Services Based on Segment Routing over IPv6 (SRv6)", RFC 9252, DOI 10.17487/RFC9252, July 2022, <<https://www.rfc-editor.org/info/rfc9252>>.

- [RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.
- [RFC9830] Previdi, S., Filsfils, C., Talaulikar, K., Ed., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", RFC 9830, DOI 10.17487/RFC9830, September 2025, <<https://www.rfc-editor.org/info/rfc9830>>.

12.2. Informative References

- [I-D.ietf-idr-fsv2-ip-basic] Hares, S., Eastlake, D., Dong, J., Yadlapalli, C., Maduschke, S., and J. Haas, "BGP Flow Specification Version 2 - for Basic IP", Work in Progress, Internet-Draft, draft-ietf-idr-fsv2-ip-basic-04, 16 March 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-fsv2-ip-basic-04>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [RFC9862] Koldychev, M., Sivabalan, S., Sidor, S., Barth, C., Peng, S., and H. Bidgoli, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing (SR) Policy Candidate Paths", RFC 9862, DOI 10.17487/RFC9862, October 2025, <<https://www.rfc-editor.org/info/rfc9862>>.

Authors' Addresses

Wenying Jiang
China Mobile
No.32 XuanWuMen West Street
Beijing
100053
China
Email: jiangwenying@chinamobile.com

Yisong Liu
China Mobile
No.32 XuanWuMen West Street
Beijing
100053
China
Email: liuyisong@chinamobile.com

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

Gyan Mishra
Verizon Communications Inc.
13101 Columbia Pike
Silver Spring, MD 20904,
United States of America
Email: gyan.s.mishra@verizon.com

Shuanglong Chen
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing
100095
China
Email: chenshuanglong@huawei.com