

IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: 3 September 2026

J. Dong
Huawei Technologies
R. Chen
ZTE Corporation
S. Wang
China Telecom
W. Jiang
China Mobile
2 March 2026

BGP Flowspec for IETF Network Slice Traffic Steering
draft-ietf-idr-flowspec-network-slice-ts-05

Abstract

BGP Flow Specification (Flowspec) provides a mechanism to distribute traffic Flow Specifications and the forwarding actions to be performed to the specific traffic flows. A set of Flowspec components are defined to specify the matching criteria that can be applied to the packet, and a set of BGP extended communities are defined to encode the actions a routing system can take on a packet which matches the flow specification.

An IETF Network Slice enables connectivity between a set of Service Demarcation Points (SDPs) with specific Service Level Objectives (SLOs) and Service Level Expectations (SLEs) over a common underlay network. To meet the connectivity and performance requirements of network slice services, network slice service traffic may need to be mapped to a corresponding Network Resource Partition (NRP). The edge nodes of the NRP needs to identify the traffic flows of specific connectivity constructs of network slices, and steer the matched traffic into the corresponding NRP, or a specific path within the corresponding NRP.

BGP Flowspec can be used to distribute the matching criteria and the forwarding actions to be performed on network slice service traffic. The existing Flowspec components can be reused for the matching of network slice services flows at the edge of an NRP. New components and traffic action may need to be defined for steering network slice service flows into the corresponding NRP. This document defines the extensions to BGP Flowspec for IETF network slice traffic steering (NS-TS).

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 3 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	3
1.1. Requirements Language	3
2. Matching Rules for Network Slice Traffic	4
2.1. NRP ID Component	4
3. Network Slice Traffic Steering Actions	5
3.1. Traffic Steering to NRP specific BE Path	5
3.1.1. Redirect to NRP specific Resource-aware Segment	5
3.1.2. Encapsulate-NRP-ID Action	5
3.2. Traffic Steering to NRP specific TE Path	6
4. Security Considerations	7
5. IANA Considerations	7
6. Acknowledgments	7
7. References	7
7.1. Normative References	7
7.2. Informative References	8
Authors' Addresses	10

1. Introduction

BGP Flow Specification (Flowspec) [RFC8955] [RFC8956] and BGP Flow Specification Version 2 [I-D.ietf-idr-fsv2-ip-basic] provide the BGP based mechanism to distribute traffic Flow Specifications and the forwarding actions to be performed to the matched traffic flows. A set of Flowspec components are defined to specify the matching criteria that is applied to the packet, and a set of Traffic Filtering Action are defined to encode the actions a routing system can take on a packet which matches the flow specification.

[RFC9543] defines the term "IETF Network Slice" and discusses the general framework for requesting and operating IETF Network Slices, their characteristics, and the necessary system components and interfaces. As described in [RFC9543], an IETF Network Slice enables connectivity between a set of Service Demarcation Points (SDPs) with specific Service Level Objectives (SLOs) and Service Level Expectations (SLEs) over a common underlay network. To meet the connectivity and performance requirements, network slice services may need to be mapped to a Network Resource Partition (NRP). An NRP is a collection of resources (bufferage, queuing, scheduling, etc.) in the underlay network. Each NRP can be identified using a unique NRP ID in control plane and management plane. The NRP ID may also be encapsulated in data packet as an NRP Selector ID to guide the NRP-specific packet forwarding. The edge nodes of an NRP needs to identify the traffic flows of specific connectivity constructs of network slices, and steer the matched packets into the corresponding NRP, so that the packet can be forwarded via either a shortest path or a Traffic Engineering (TE) path within the NRP.

BGP Flowspec can be used to distribute the matching criteria and the forwarding actions to network nodes for steering the matched service traffic into an NRP. The existing Flowspec components can be reused for the matching of network slice service flows. New components and traffic actions may need to be defined for steering network slice service flows into the corresponding NRP. This document defines the extensions to BGP Flowspec for IETF Network Slice Traffic Steering (NS-TS).

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. Matching Rules for Network Slice Traffic

A set of traffic matching rules can be used as the criteria to match the traffic flows of specific connectivity constructs of IETF network slice. The BGP Flowspec components as defined in [RFC8955] [RFC8956] can be used to specify the matching rules for network slice service packets.

In some cases, such as for multi-domain network slices, data packets of a network slice are encapsulated with data plane NRP ID in an upstream network domain using the mechanisms as described in [I-D.ietf-6man-enhanced-vpn-vtn-id]. Then the ingress edge node of the downstream network domain may perform traffic matching based on the NRP ID in the packets and the corresponding network slice matching rules, so that the packets can be steered into a corresponding NRP in the local domain. A new Flow component called NRP ID component is defined for this purpose.

2.1. NRP ID Component

The format of the NRP ID component follows the Flowspec encoding as defined in [I-D.ietf-idr-fsv2-ip-basic], which consists of 1-octet type field, 1-octet length field, and variable value field. The type of NRP ID component is to be assigned by IANA. The format of the value field is shown as below:

```

1             2             3             4
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|g|             Flags             |             Reserved             |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     NRP ID                             |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Where

- * Flags: 2-octet flag field. The first (most significant) bit is defined in this document, the rest of the flag bits SHOULD be set to zero on transmission and MUST be ignored on receipt.
- * Global bit (g): When set, it indicates the NRP ID to be matched is a global unique NRP ID; otherwise the NRP ID is a domain significant NRP ID. The g bit is used for an NRP which span multiple network domains, and a global NRP ID has been coordinated among these domains.
- * Reserved: 2-octet reserved bits. It SHOULD be set to zero on transmission and MUST be ignored on receipt.

* NRP ID: A 4-octet identifier which is used to identify an NRP.

3. Network Slice Traffic Steering Actions

For data packets which match the Flow Specification of a network slice, specific forwarding actions need to be applied. When the network slice service flows are mapped to an NRP in the underlay network, the packets of the flows need to be forwarded in the corresponding NRP using either a shortest (BE) path or a Traffic Engineering (TE) path.

This section describes several actions to be performed on packets which match the Flow Specification of a network slice.

3.1. Traffic Steering to NRP specific BE Path

Packets of a network slice service flow can be steered into an NRP and forwarded to the NRP egress node following the shortest path with the NRP. In this case, the identifier of the NRP needs to be carried in the packet so that the packet forwarding will be performed using the set of resources allocated to the NRP. Depends on the type of the data plane NRP specific identifier, there are two options of this traffic steering.

3.1.1. Redirect to NRP specific Resource-aware Segment

When resource-aware SR segments

[I-D.ietf-spring-resource-aware-segments] are used to represent the network resources allocated to an NRP, packets of a network slice could be steered into an NRP BE path by encapsulating the packets with a resource-aware segment of the egress node in the NRP. For SRv6 data plane, this could be achieved by using the "redirect-to-IP" actions defined in [I-D.ietf-idr-flowspec-redirect-ip]. The mechanism for SR-MPLS data plane will be specified in a future version.

3.1.2. Encapsulate-NRP-ID Action

When a data plane NRP ID (called NRP Selector ID)

[I-D.ietf-teas-nrp-scalability] is used to identify the set of network resources allocated to an NRP, packets of a network slice service flow could be steered into an NRP specific BE path by encapsulating the NRP Selector ID together with the IP address or the SR SID of the egress node in the NRP.

For encapsulating the data plane NRP ID to the matched packets, a new BGP extended community is defined for the "Encapsulate-NRP-ID" action. The format of this extended community is as below:




Figure 1. The format of Encapsulate-NRP-ID action

where:

- * Type: 0x80. It belongs to the Generic Transitive Extended Community Type as defined in [RFC9184].
- * Sub-type: 1 octet to be assigned by IANA.
- * Flags: 2-octet flag field. The first bit is defined in this document. The rest of the flags are unused, which SHOULD be set to zero on transmission and MUST be ignored on receipt.
- * Encapsulate (E) bit: When set, it indicates the NRP ID MUST be encapsulated with an outer header to the packet. Otherwise the NRP ID replaces the NRP ID in the existing header of the packet.
- * NRP ID: A 4-octet identifier which is used to identify an NRP.

If a packet matches the Flow Specification of an IETF network slice, and the traffic actions associated with the flow specification is the Encapsulate-NRP-ID action, then the packet is encapsulated with an corresponding NRP Selector ID in the packet header. The mapping between the control plane NRP ID and the NRP Selector ID of a specific data plane is maintained by the receiving BGP speaker. The Encapsulate-NRP-ID action MAY be used together with the "Redirect-to-IP" action as defined in [I-D.ietf-idr-flowspec-redirect-ip], in that case the destination address of the outer IP header is set to the IP address in the redirect to IP next-hop action. The IPv6 encapsulation of NRP ID is specified in [I-D.ietf-6man-enhanced-vpn-vtn-id]. The encapsulation of NRP-ID in other data plane is for further study and out of the scope of this document.

3.2. Traffic Steering to NRP specific TE Path

Packets of a network slice can be steered into a TE path within the corresponding NRP. In an SR network, the network slice traffic can be steered into an SR Policy [RFC9256] which is associated with the corresponding NRP.

In SR networks where the NRP is instantiated using NRP specific resource-aware segments [I-D.ietf-spring-resource-aware-segments], the segment list of the SR policy are built with resource-aware SR segments which represents the set of network resources allocated to the NRP on different network segments.

In SR networks where the data plane NRP Selector ID is used to identify the set of network resources allocated to the NRP, the mechanism as defined in [I-D.ietf-idr-sr-policy-nrp] provides the BGP SR Policy extensions to associate an SR Policy candidate path with an NRP-ID.

In both of the above two cases, the mechanism defined in [I-D.ietf-idr-ts-flowspec-srv6-policy] could be used to steer traffic to an SR Policy which is associated with an NRP.

4. Security Considerations

The security considerations of BGP [RFC4271] and BGP Flowspec [RFC8955] [RFC8956] apply to this document.

5. IANA Considerations

IANA is requested to assign a new type code point from "Flow Spec Component Types" registry.

Type Value	IPv4 Name	IPv6 Name	Reference
-----	-----	-----	-----
TBA1	NRP ID	NRP ID	This document

IANA is requested to assign a new sub-type from "Generic Transitive Extended Community Sub-Types" registry.

Value	Description	Reference
-----	-----	-----
TBA2	Flowspec Encapsulate-NRP-ID	This document

6. Acknowledgments

The authors would like to thank Haisheng Wu, Haibo Wang and Shunwan Zhuang for the review and discussion of this document.

7. References

7.1. Normative References

[I-D.ietf-idr-fsv2-ip-basic]

Hares, S., Eastlake, D. E., Dong, J., Yadlapalli, C., and S. Maduschke, "BGP Flow Specification Version 2 - for Basic IP", Work in Progress, Internet-Draft, draft-ietf-idr-fsv2-ip-basic-03, 3 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-fsv2-ip-basic-03>>.

[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>>.

[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>>.

[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>>.

[RFC9184] Loibl, C., "BGP Extended Community Registries Update", RFC 9184, DOI 10.17487/RFC9184, January 2022, <<https://www.rfc-editor.org/info/rfc9184>>.

7.2. Informative References

[I-D.ietf-6man-enhanced-vpn-vtn-id]

Dong, J., Li, Z., Xie, C., Ma, C., and G. S. Mishra, "Carrying Network Resource (NR) related Information in IPv6 Extension Header", Work in Progress, Internet-Draft, draft-ietf-6man-enhanced-vpn-vtn-id-14, 11 February 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-6man-enhanced-vpn-vtn-id-14>>.

[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-05, 13 February 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-redirect-ip-05>>.

[I-D.ietf-idr-sr-policy-nrp]

Dong, J., Hu, Z., and R. Pang, "BGP SR Policy Extensions for Network Resource Partition", Work in Progress, Internet-Draft, draft-ietf-idr-sr-policy-nrp-07, 11 February 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-nrp-07>>.

[I-D.ietf-idr-ts-flowspec-srv6-policy]

Wenying, J., Liu, Y., Zhuang, S., Mishra, G. S., and S. Chen, "Traffic Steering using BGP FlowSpec with SR Policy", Work in Progress, Internet-Draft, draft-ietf-idr-ts-flowspec-srv6-policy-08, 1 February 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-ts-flowspec-srv6-policy-08>>.

[I-D.ietf-spring-resource-aware-segments]

Dong, J., Miyasaka, T., Zhu, Y., Qin, F., and Z. Li, "Introducing Resource Awareness to SR Segments", Work in Progress, Internet-Draft, draft-ietf-spring-resource-aware-segments-17, 19 January 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-resource-aware-segments-17>>.

[I-D.ietf-teas-nrp-scalability]

Dong, J., Li, Z., Gong, L., Yang, G., and G. S. Mishra, "Scalability Considerations for Network Resource Partition", Work in Progress, Internet-Draft, draft-ietf-teas-nrp-scalability-09, 11 February 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-nrp-scalability-09>>.

[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>>.

[RFC9543] Farrel, A., Ed., Drake, J., Ed., Rokui, R., Homma, S., Makhijani, K., Contreras, L., and J. Tantsura, "A Framework for Network Slices in Networks Built from IETF Technologies", RFC 9543, DOI 10.17487/RFC9543, March 2024, <<https://www.rfc-editor.org/info/rfc9543>>.

Authors' Addresses

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

Ran Chen
ZTE Corporation
Email: chen.ran@zte.com.cn

Subin Wang
China Telecom
Email: wangsb6@chinatelecom.cn

Wenying Jiang
China Mobile
Email: jiangwenying@chinamobile.com