IDR
Internet-Draft
Intended status: Standards Track
Expires: 5 January 2026

L. Zhang, Ed.
T. Zhou
Huawei
Z. Li
China Mobile
J. Dong
Huawei

4 July 2025

BGP Flow Specification Extensions for Scheduling
draft-zzd-idr-flowspec-scheduling-03

Abstract

BGP Flow Specification allows conveying flow specifications and traffic Action/Rules associated to perform different actions based on the traffic features. One of the applications is to steer one specific flow into its specific path. However, in some scenarios, the traffic forwarding paths are not constant and change over time.

This document extends BGP Flow Specification with scheduling time information to identify the packets arrived at different time slot. Based on that, the headend can perform different actions at different time for the same traffic.

Table of Contents

1. Introduction

[RFC8955] and [RFC8956] define the BGP [RFC4271] Flow Specification (FlowSpec) that allows conveying flow specifications and traffic Action/Rules associated. 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 existing traffic filter rules and actions in FlowSpec are always effective and will steer specific traffic into one path once been delivered to the headend. However, there are many scenarios that need to schedule routing paths in the network.

[RFC9657] introduces a set of use cases where the topology of the network changes predictably. The topology change may cause some of the paths invalid, and lead to path reselection or even recalculation. However, the reselection or recalculation takes a period of time, which will affect packet forwarding and cause problems such as packet disorder and packet loss. However, on a network with predictable topology changes, if the ingress node knows future topology changes, it can schedule the forwarding paths in advance, and steer flows to different set of paths based on time to prevent packet forwarding from being affected by topology changes.

Another case is in tidal networks, the traffic in network changes periodically(e.g. IP carrier backbone network). In some rush hours(typicall 20-22 o'clock), the increased traffic may lead to network congestion. One possiable solution is to divert some traffic to onther paths or limit the bandwidth of some flows to reduce the congestion in ruch hours.

This document extends BGP Flow Specification with scheduling time information to identify the packets arrived at different time slot. Based on that, the headend can perform different actions at different time for the same traffic.

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. Scheduling Time Information in FlowSpecv1

[RFC8955] defines 12 Components to identify different traffics. Based on [RFC8955], this document defines a new Component to identify the arrival time of packets and perform different actions.

Encoding: <type (1 octet, TBD1), length (1 octet), scheduling time information (variable)+>

Defines the time information that matches the arrival time of packets. This component matches if the arrival time of an IP packet in the scope of scheduling time information.

3. Scheduling time information in FlowSpecv2

[I-D.ietf-idr-flowspec-v2] specifies BGP flow specification v2 to address the issues detected during the deployment of BGP flow specification v1. It defines that the traffic filters are described in the format of sub-TLV and different traffic type have different filter sub-TLVs. This document defines a new sub-TLV for IP filters and L2 filters defined in Section 3 of [I-D.ietf-idr-flowspec-v2] to identify the arrival time of packets and perform different actions. The format of Scheduling Time sub-TLV is shown as follows:

[I-D.ietf-idr-flowspec-v2] specifies BGP flow specification v2(FSv2) to address the issues detected during the deployment of BGP flow specification v1. It defines that the traffic filters are described in the format of sub-TLV and different traffic type have different filter sub-TLVs.

For the IP and VPN IP filters, FSv2 reused the components defined in [RFC8955], [RFC8956], and [I-D.ietf-idr-flowspec-srv6]. Therefore, the new component defined for FlowSpecv1 in Section 3 is also applicable for FlowSpecv2

For L2 Traffic, this document defines a new sub-TLV for L2 filters defined in Section 3 of [I-D.ietf-idr-flowspec-v2] to identify the arrival time of packets and perform different actions. The format of Scheduling Time sub-TLV is shown as follows:

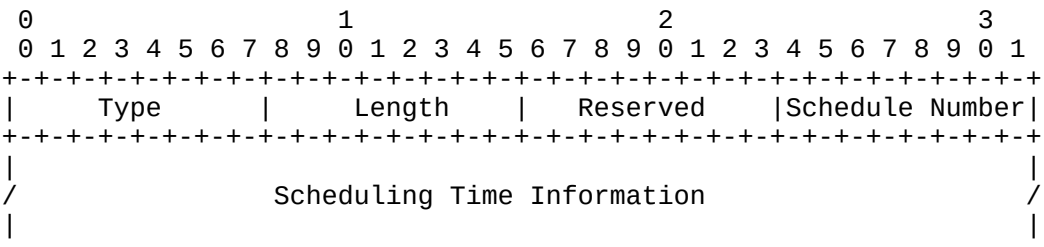

Figure 1: Scheduling Time Sub-TLV

Type: TBD2

Length: the size of the value field in octets, variable.

Schedule Number: indicates the number of schedules.

Schedules Time information: one or more schedules, each schedule indicates when one or more time slots.

4. Scheduling Time Information

The format of Scheduling time information sub-TLV is shown as follows:

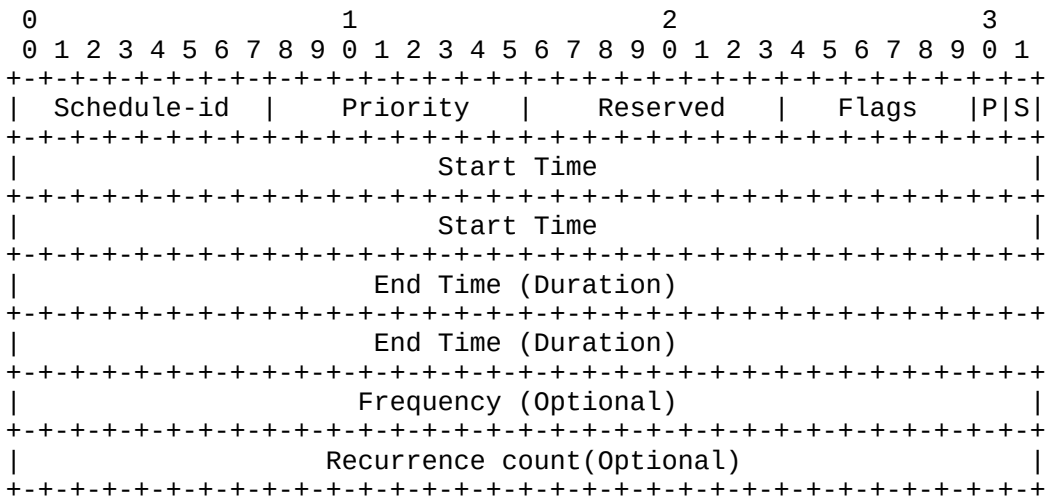

Figure 2: Schedule of SR Policy

Schedule-id: 8-bit value, the unique identifier to distinguish each schedule within a FlowSpec, this value is allocated by the FlowSpec generator.

Priority: 8-bit value, this field is used when there are multiple schedules valid at the same point. The higher value indicates higher priority and the default Preference value is 10.

Flags: 8 bits, currently only 2 bits are used, the other bits are reserved.

P (Period format): one-bit flag to indicate the format of a period. if P=1, then the period is described by a start time filed and an end time field; If P =0, then the period is described by a start time field and a duration time field.

S (Schedule type): one-bit flag to indicate the type of a schedule. If S=0, it indicates the schedule only has one instance, the Frequency and Recurrence count field should not be included in the sub-TLV; If S=1, it indicates the schedule has multiple instances, the Frequency and Recurrence count field should be included.

Start Time: 64-bit value, the number of seconds since the epoch, it indicates when the FlowSpec start to take effect.The epoch is 1 January 1970 at 00:00 UTC.

End Time (Duration): 64-bit value, if the flag P=1, then it is the number of seconds since the epoch, it indicates when the FlowSpec becomes ineffective. If the flag P=0, then it is the number of seconds since the Start Time, it indicates how long the FlowSpec take effect.

Frequency(optional): 32-bit value, it is the numbers of seconds since the Start Time of an instance to the Start Time of next instance. This field indicates the recurrence frequency for all the instance of this schedule. This field should not be included if S=0.

Recurrence Count(optional): 32-bit value, it indicates the number of occurrences. For example, if it is set to 2, then the schedule will repeat twice with the specified Frequency. This field should not be included if P=0.

5. Security Considerations

These extensions to BGP FlowSpec do not add any new security issues to the existing protocol.

6. IANA Considerations

IANA is requested to allocate a new type value for "Scheduling Time Information" Component in "Flow Spec Component Types" registry:

Value	Description	Reference
TBD1	Scheduling Time Information	This document

Table 1

IANA is requested to allocate a new type value for "Scheduling Time sub-TLV" in "L2 Flow Specification Component Types" registry:

Value	Description	Reference
TBD2	Scheduling Time sub-TLV	This document

Table 2

7. References

7.1. Normative References

[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>.

[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>.

[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>.

[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>.

7.2. Informative References

[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>.

[RFC9657] Birrane, III, E., Kuhn, N., Qu, Y., Taylor, R., and L. Zhang, "Time-Variant Routing (TVR) Use Cases", RFC 9657, DOI 10.17487/RFC9657, October 2024, <https://www.rfc-editor.org/info/rfc9657>.

[I-D.ietf-idr-flowspec-v2] Hares, S., Eastlake, D. E., Yadlapalli, C., and S. Maduschke, "BGP Flow Specification Version 2", Work in Progress, Internet-Draft, draft-ietf-idr-flowspec-v2-04, 28 April 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-v2-04>.

[I-D.ietf-idr-flowspec-srv6]
Li, Z., Chen, H., Loibl, C., Mishra, G. S., Fan, Y., Zhu, Y., Liu, L., Liu, X., and S. Zhuang, "BGP Flow Specification for SRv6", Work in Progress, Internet-Draft, draft-ietf-idr-flowspec-srv6-07, 12 May 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-srv6-07>.

Authors' Addresses

Li Zhang (editor)
Huawei
Beiqing Road
Beijing
China
Email: zhangli344@huawei.com

Tianran Zhou
Huawei
Email: zhoutianran@huawei.com

Zhenqiang Li
China Mobile
Email: lizhenqiang@chinamobile.com

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