

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

L. Zhang, Ed.
T. Zhou
J. Dong
Huawei
M. Wang
China Mobile
N. Nzima
MTN
4 July 2025

BGP SR Policy Extensions for Path Scheduling
draft-zzd-idr-sr-policy-scheduling-09

Abstract

Segment Routing (SR) policy enables instantiation of an ordered list of segments with a specific intent for traffic steering. When using SR policy in a time-variant network, delivering the time-variant information associated with paths is necessary in some scenarios.

This document proposes extensions to BGP SR Policy to deliver the schedule time information of candidate path (segment list) and its associated attributes.

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 5 January 2026.

Copyright Notice

Copyright (c) 2025 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
1.1. Requirements Language	3
2. Motivation	3
2.1. Network with Discontinuous Links	3
2.2. Network with Frequent Topology Changes	4
3. Schedule Time Information in SR Policy	4
4. Schedule Time Information Sub-TLV	6
5. Operations	8
5.1. Advertisement of SR Policies with Schedule Time Information	8
5.2. Reception of an SR Policy with Schedule Time Information	9
5.2.1. Validation of Schedule Time Information	9
5.2.2. Enable/Disable Candidate Paths/Segment Lists	10
6. Error Handling and Fault Management	10
7. Manageability Considerations	10
8. Security Considerations	11
9. IANA Considerations	11
10. References	11
10.1. Normative References	11
10.2. Informative References	12
Authors' Addresses	13

1. Introduction

[RFC9657] introduces a set of time-variant network use cases where the topology of the network changes predictably. When the networks uses traditional routing protocols, it takes these topology changes as unexpected events and may cause and packet loss. However, the topology changes of these networks can be predicted in advance, therefore some measures can be taken in advance to prevent the packet loss. With this idea, [I-D.ietf-tvr-requirements] describes the requirements of using the time-variant information in a network. In Section 3.4.1 of [I-D.ietf-tvr-requirements], it describes the centralized routing scenarios with time-variant information, in which the network entities receive the time variable information and traffic forwarding rules directly from a logically centralized

source(an Orchestrator or network controller). The time-variant information is especially essential when there is a risk that a logically centralized source may loses connectivity with the network entities.

Segment Routing (SR) policy [RFC9256] is a set of candidate SR paths consisting of one or more segment lists and necessary path attributes. It describes the traffic forwarding rules for specific flows. [I-D.ietf-idr-sr-policy-safi] introduces how BGP may be used to distribute SR Policy candidate paths. It introduces a BGP SAFI with new NLRI to advertise the candidate paths and specific attributes of a SR Policy from a controller or path computation engine (PCE) to the network entities. However, when using BGP SR Policy in a time-variant network, it can't advertise the schedule time information associated with paths.

This document proposes extensions to BGP SR Policy to carry the schedule time information of candidate paths/segment lists.

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. Motivation

Most of the time-variant network use cases using BGP SR Policy could be benefit from this work. In some cases, carrying the time-variant information with SR Policy is essential.

This section describes the cases that requires extending SR Policy with schedule time information.

2.1. Network with Discontinuous Links

In some time-variant network cases, the links between the network entities and network controller may very weak or intermittent, this is very typical in Resource Preservation and Dynamic Reachability network[RFC9657]. In these cases, Real-time SR policy advertising (before changes occur) may not be timely. For example, when a link of an old path is about to be disconnected, the network controller is going to advertise a new path to the headend. However, the link between the headend and the network controller is not available. As a result, the new path cannot be advertised in time, causing packet loss.

Therefore, in these cases, once the links between the headend and network controller are available, the controller need to advertise the paths with schedule time information for a period in the future to the headend. Then the headend could determine valid paths in the future based on the schedule time information of SR policy.

2.2. Network with Frequent Topology Changes

There are also some time-variant network cases that topology changes frequently. This is very typical when the number of network entities is very large (For example, a Dynamic Reachability network with hundreds or thousands of nodes). In this kind of time-variant network, a path form one network entity to another changes frequently, sometimes it can only be maintained for a few minutes or seconds.

Considering that there are multiple paths in a network that computed by the controller, the SR Policies with candidate paths may be advertised to the headend every few seconds. It poses great changeling to the network controller. However, using schedule time information could advertise several paths once, which greatly mitigate the pressure of network controllers.

3. Schedule Time Information in SR Policy

The NLRI defined in [I-D.ietf-idr-sr-policy-safi] contains the SR Policy candidate path. The content of the SR Policy Candidate Path is encoded in the Tunnel Encapsulation Attribute defined in [RFC9012] using the Tunnel-Type called SR Policy Type with codepoint 15. The SR Policy encoding structure is as follows:

SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>

Attributes:

Tunnel Encapsulation Attribute (23)

Tunnel Type: SR Policy (15)

Binding SID

SRv6 Binding SID

Preference

Priority

Policy Name

Policy Candidate Path Name

Explicit NULL Label Policy (ENLP)

Segment List

Weight

Segment

Segment

...

...

A candidate path includes multiple SR paths, each of which is specified by a segment list. The schedule time information can be applied to a candidate path, indicating the valid time for each candidate path and its associated attributes. The new SR Policy encoding structure is expressed as below:

SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>

Attributes:

Tunnel Encapsulation Attribute (23)

Tunnel Type: SR Policy (15)

Binding SID

SRv6 Binding SID

Preference

Priority

Policy Name

Policy Candidate Path Name

Explicit NULL Label Policy (ENLP)

Schedule Time Information

Segment List

Weight

Segment

Segment

...

...

The schedule time information also can be applied to a segment list to indicate the valid time for a segment list and its associated attributes. The new SR Policy encoding structure is expressed as below:

SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>

Attributes:

Tunnel Encapsulation Attribute (23)

Tunnel Type: SR Policy (15)

Binding SID

SRv6 Binding SID

Preference

Priority

Policy Name

Policy Candidate Path Name

Explicit NULL Label Policy (ENLP)

Segment List

Schedule Time Information

Weight

Segment

Segment

...

...

The Schedule Time Information TLV is either encapsulated at the candidate path level or at the segment list level. It SHOULD NOT be encapsulated at both levels simultaneously.

4. Schedule Time Information Sub-TLV

The Schedule Time Information sub-TLV defined in this document is used in the SR policy Tunnel TLV, it may be extended to other type of Tunnel TLVs, but it is out of scope of this document. The Schedule Time Information sub-TLV indicates one or more valid time slot for one or more paths. It is OPTIONAL and MAY appear more than once in the SR Policy encoding.

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

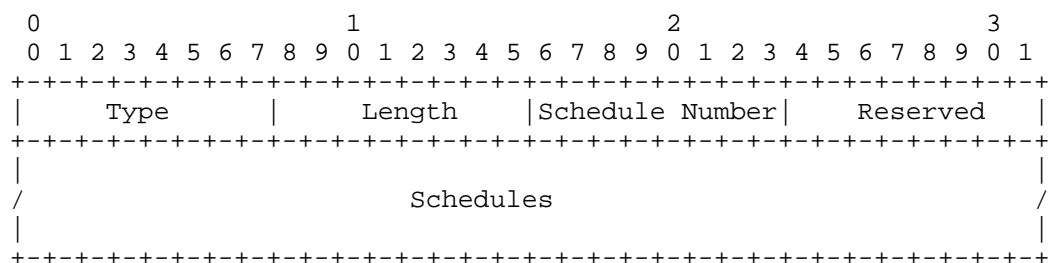


Figure 1: Scheduling Time Information Sub-TLV

Type: TBD1

Length: the size of the value field in octets.

Schedule Number: indicates the number of schedules.

Schedules: one or more schedules, each schedule indicates the duration when the candidate path (segment list) is active. The format of each schedule is shown as follows:

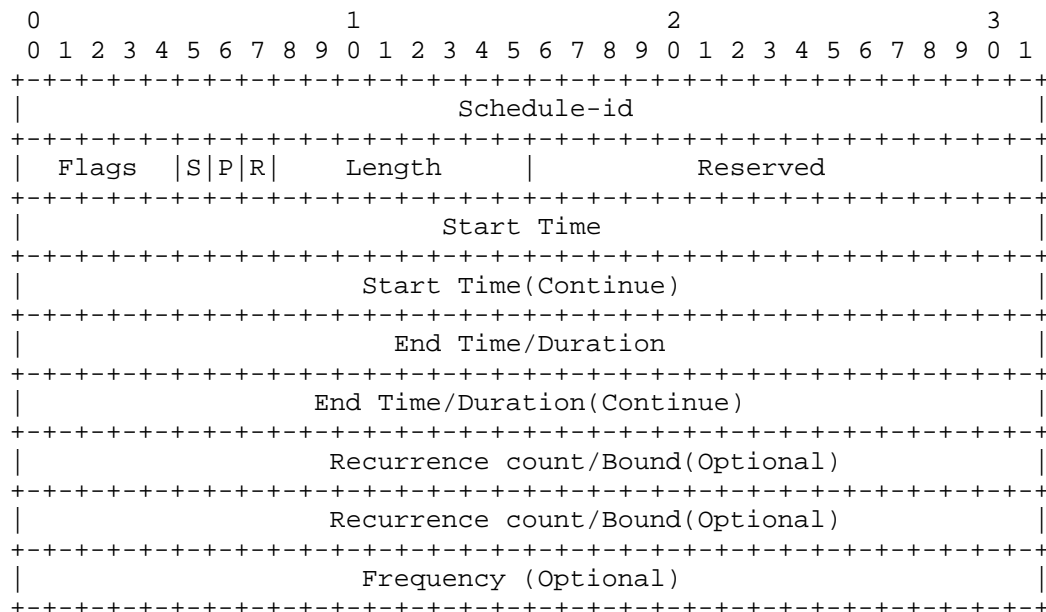


Figure 2: Format of Schedules

Schedule-id: 32-bit value, the unique identifier to distinguish each schedule within a SR Policy, this value is allocated by the SR Policy generator.

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

Length: 8 bits, indicates the length of this schedule in octets.

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 Recurrence Count/Bound, Frequency and Interval field should not be included in the sub-TLV; If S=1, it indicates the schedule has multiple instances, the Recurrence Count/Bound, Frequency and Interval field should be included.

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

R (Recurrence bound type): one-bit flag to indicate the how to determine whether the recurrence is end. if R=1, then the end of recurrence is determined by a detail timepoint; If R = 0, then the end of the recurrence is determined by the number of occurrences.

Start Time: 64-bit value, the number of seconds since the epoch, it indicates when the candidate path (segment list) and its associated attributes 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 candidate path (segment list) and its associated attributes becomes ineffective. If the flag P=0, then it is the number of seconds since the Start Time, it indicates how long the candidate path (segment list) and its associated attributes are effective.

Recurrence Count/Bound(optional): 64-bit value, this field SHOULD be included when the flag P is set to 1. When the flag R=0, then this field indicates the max number of occurrences. For example, if it is set to 2, then the schedule will repeat twice with the specified Frequency and Interval. When the flag R=1, then this field indicates the bounded timepoint of recurrence, it is described by the number of seconds since the epoch.

Frequency(optional): 32-bit value, this field should be included when the flag S is set to 1. 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.

5. Operations

5.1. Advertisement of SR Policies with Schedule Time Information

As described in Section 4.1 of [I-D.ietf-idr-sr-policy-safil], typically, an SR Policy is computed by a controller or a path computation engine (PCE) and originated by a BGP speaker on its behalf. The schedule time information is also computed by a controller or a PCE which can access to all the predicted topology changes. The predicted topology changes could be got from management interfaces or other means.

The controller or PCE SHOULD maintain a time-variant event database as described in Section 6 of [I-D.zdm-tvr-applicability] to store all the predicted information, and compute SR Policies with schedule time information based on the database.

Each Candidate Path or Segment List may have multiple schedules, each schedule is identified by schedule-id, the controller or PCE MUST make the schedule-id unique within a specific SR Policy.

if no schedule is included in the Schedule Time Information sub-TLV, then it is assumed that the candidate path or segment list is not time-varying.

5.2. Reception of an SR Policy with Schedule Time Information

As described in [I-D.ietf-idr-sr-policy-safil], the BGP SR Policy is distinguished by <distinguisher, color, endpoint> tuple. Whenever a headend receives a SR Policy that it has received before, the SR Policy is considered as the fully replacement of the old one.

The headend MUST NOT use the SR Policy with schedule time information when it does not have the capability to deal with the schedule time information.

5.2.1. Validation of Schedule Time Information

On reception of an SR Policy, a BGP speaker first determines if it is valid as described in Section 4.2.1 of [I-D.ietf-idr-sr-policy-safil]. When a headend receives a SR Policy with Schedule Time Information from its neighbors or controller, it SHOULD perform schedule time information validation based on the following rules:

- * When multiple schedules are present within one SR Policy, the schedule-id of each schedule MUST be different. If there are multiple schedules with the same schedule-id, then the first instance of that schedule is used, and the other instances MUST be ignored.
- * If a SR Policy encapsulates Schedule Time Information sub-TLVs at both candidate path level and segment list level simultaneously, then the sub-TLVs at segment list level is used, and the sub-TLVs at candidate path level MUST be ignored.
- * The Start Time of Schedule Time Information sub-TLV MUST be later than the time it is received, if not, the sub-TLV MUST be considered as malformed.
- * If the End Time/Duration field is used to indicate the end time, then it MUST be later than the Start Time, if not, the sub-TLV MUST be considered as malformed.

- * If the Frequency field is present in the Schedule Time Information sub-TLV, then it MUST be greater than the difference between the End Time and Start Time(P=1), or greater than the Duration(P=0), if not, the sub-TLV MUST be considered as malformed.
- * If the Recurrence Count/Bound field is present and used to indicate the boundary time point, then it MUST be greater than the End Time(P=1), or greater than the sum of Start Time and Duration(P=0), if not, the sub-TLV MUST be considered as malformed.

5.2.2. Enable/Disable Candidate Paths/Segment Lists

When a headend receives a SR Policy with schedule time information, it SHOULD parse the SR Policy and save the schedule time information locally. When a data packet arrives, it will steer it to a specific SR Policy by color or other means.

Within a specific SR Policy, the headend dynamically enable/disable different Candidate Paths/Segment Lists based on the schedule time information.

6. Error Handling and Fault Management

The error handling actions defined in Section 5 of [I-D.ietf-idr-sr-policy-safi] are also applicable in this document.

If a BGP speaker receives a SR Policy with Schedule Time Information but it does not have the capability to deal with the schedule time information, then the SR Policy SHOULD NOT be considered usable as described in {Section 4.2.2 of !!I-D.ietf-idr-sr-policy-safi}.

The validation rules defined in Section 5.2.1 also SHOULD be performed by SRPM to determine if the Schedule Time Information sub-TLV is malformed or invalid. If the Schedule Time Information sub-TLV is invalid or malformed, then the implementation SHOULD log errors, and the corresponding candidate path or segment list MUST NOT be used.

7. Manageability Considerations

The specification of BGP models is an ongoing work based on [I-D.ietf-idr-bgp-model] and its future extensions are expected to cover the SR Policy SAFI and its extensions. Existing BGP operational procedures also apply to the extension specified in this document.

The YANG model for the operation and management of SR Policies with Schedule Time Information will be defined in the future.

8. Security Considerations

The security considerations described in the [I-D.ietf-idr-sr-policy-safil] apply to the extensions described in this document as well.

The Schedule Time Information TLV defined in this document could provide details about the network operations thant could be sensitive. Attacker can obtain these sensitive data by snooping BGP messages and select the optimal attack time. Thus suitalbe BGP security mechanisms like TLS is necessary.

Time sychronization is essential for schedules, the attackers can attack the network by generating incorrect time synchronization messages or block the time synchronization process. Therefore, redundant time synchronization mechanisms and secure time synchronization mechanisms(such as NTS (Network Time Security)) are also necessary.

9. IANA Considerations

This document defines a sub-TLV in the registry "BGP Tunnel Encapsulation Attribute sub-TLVs" to be assigned by IANA:

Value	Description	Reference
TBD1	Schedule Time Information (STI)	This document

Table 1

The type value of this sub-TLV is recommended to be token from the range of 64-125(First Come First Served).

10. References

10.1. Normative References

[I-D.ietf-tvr-requirements]
King, D., Contreras, L. M., Sipos, B., and L. Zhang, "TVR (Time-Variant Routing) Requirements", Work in Progress, Internet-Draft, draft-ietf-tvr-requirements-05, 3 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-tvr-requirements-05>>.

- [RFC9256] Filtsils, 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>>.
- [I-D.ietf-idr-sr-policy-safi]
Previdi, S., Filtsils, C., Talaulikar, K., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-sr-policy-safi-13, 6 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-safi-13>>.
- [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>>.
- [I-D.zdm-tvr-applicability]
Zhang, L., Dong, J., and M. Boucadair, "Applicability of TVR YANG Data Models", Work in Progress, Internet-Draft, draft-zdm-tvr-applicability-02, 28 February 2025, <<https://datatracker.ietf.org/doc/html/draft-zdm-tvr-applicability-02>>.

10.2. Informative References

- [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>>.
- [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>>.
- [I-D.ietf-idr-bgp-model]
Jethanandani, M., Patel, K., Hares, S., and J. Haas, "YANG Model for Border Gateway Protocol (BGP-4)", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-model-18, 21 October 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-model-18>>.

Authors' Addresses

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

Tianran Zhou
Huawei
Email: zhoutianran@huawei.com

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

Minxue Wang
China Mobile
Email: wangminxue@chinamobile.com

Nkosinathi Nzima
MTN
Email: Nkosinathi.Nzima@mtn.com