

SPRING Working Group
Internet Draft
Intended status: Informational
Expires: March 16, 2024

W. Cheng
W. Jiang
China Mobile
C. Lin
Y. Qiu
New H3C Technologies
Y. Wei
Huawei Technologies
Ran. Chen
ZTE
R. Liang
Ruijie Networks
September 13, 2023

SR Policy Group
draft-cheng-spring-sr-policy-group-03

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on March 16 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 (<http://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 Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

Abstract

Segment Routing is a source routing paradigm that explicitly indicates the forwarding path for packets at the ingress node. An SR Policy is associated with one or more candidate paths, and each candidate path is either dynamic, explicit, or composite. This document describes SR policy Group in MPLS and IPV6 environments and illustrates some use cases for parent SR policy and SR Policy Group to provide best practice cases for operators.

Table of Contents

| | | |
|---------------------|--|--------------------|
| 1. | Introduction..... | 2 |
| 2. | Terminology..... | 3 |
| 3. | Parent SR Policy..... | 3 |
| 4. | SR Policy Group..... | 4 |
| 5. | Steering into SR Policy Group..... | 5 |
| 6. | Relationship Sorting..... | 7 |
| 7. | Information Model of SR Policy Group..... | 8 |
| 8. | Use Cases..... | 10 |
| | 8.1. Parent SR Policy in L3VPN over TE Scenarios..... | 10 |
| | 8.2. SR Policy Group in Multi-VPN Tenants Scenarios..... | 11 |
| 9. | IANA Considerations..... | 14 |
| 10. | Security Considerations..... | 14 |
| 11. | References..... | 15 |
| | 11.1. Normative References..... | 15 |
| | 11.2. Informative References..... | 15 |
| 12. | Acknowledgments..... | 15 |
| | Authors' Addresses..... | 16 |

1. Introduction

Segment routing (SR) [[RFC8402](#)] is a source routing paradigm that explicitly indicates the forwarding path for packets at the ingress node. The ingress node steers packets into a specific path according to the Segment Routing Policy (SR Policy) as defined in [[RFC9256](#)].

In order to distribute SR policies to the headend, [I-D.ietf-idr-segment-routing-te-policy] specifies a mechanism by using BGP.

An SR Policy is associated with one or more candidate paths. A composite candidate path acts as a container for grouping SR Policies. As described in [RFC9256], the composite candidate path construct enables combination of SR Policies, each with explicit candidate paths and or dynamic candidate paths with potentially different optimization objectives and constraints, for load-balanced steering of packet flows over its constituent SR Policies. For convenience, the composite candidate path formed by the combination of SR policies is called Parent SR policy.

This document describes SR Policy Group in MPLS and IPv6 environments and illustrates some use cases for parent SR policy and SR Policy Group to provide best practice cases for operators.

2. Terminology

The definitions of the basic terms are identical to those found in Segment Routing Architecture [RFC8402], Segment Routing Policy Architecture [RFC9256].

SR policy As described in [RFC9256], the general concept of SR Policy provides a framework that enables the instantiation of an ordered list of segments on a node for implementing a source routing policy for the steering of traffic for a specific purpose (e.g., for a specific Service Level Agreement (SLA)) from that node. An SR Policy is a forwarding path that meets the specified forwarding requirements.

Parent SR Policy A Parent SR Policy is the composite candidate path that acts as a container for grouping SR policies which meet different service optimization objectives and constraints and have the same destination endpoint.

SR Policy Group: An SR policy Group is an instantiation of a group of constituent Parent SR Policies to different destination endpoints with the same service forwarding model. It establishes the mapping relationship between the flow characteristics and the color value of the SR Policy, and guide the flows with different SLA requirements to the SR Policy with different colors.

3. Parent SR Policy

An SR Policy is associated with one or more candidate paths. According to [section 2.2 of \[RFC9256\]](#) a parent SR policy is the composite candidate path that acts as a container for grouping SR

policies. The parent SR policy is valid when it has at least one valid constituent SR policy.

As defined in [[RFC9256](#)], the endpoints of the constituent SR Policies and the parent SR policy MUST be identical, and the colors of each of the constituent SR Policies and the parent SR policy MUST be different. Parent SR policy and its constituent SR Policies follow the same criteria:

- * The endpoints of the constituent SR Policies and its parent SR policy MUST be identical.
- * The colors of each of the constituent SR Policies and its parent SR policy MUST be different.
- * The constituent SR Policies MUST NOT contain parent SR policy.

As a special SR policy, parent SR policy also has color attribute, which is identified by <color, endpoint> on the headend.

If a parent SR policy has at least one valid constituent SR Policy of specified color, flow load-balance steer over its valid constituent SR policies with the same color. When all constituent SR policies of specified color are invalid, packets can be forwarded based on a preconfigured default forwarding path (such as a BE path, an SR policy path, or a composite SR policy path of another color).

The parent SR Policy can be well used for the following scenarios

- * Point-to-point leased line customers have multiple services.
- * These services need to be forwarded through different SR policy paths with different SLAs.

4. SR Policy Group

An SR policy Group is an instantiation of a group of constituent Parent SR Policies to different destination endpoints with the same service forwarding model and represents a composite candidate path defined in [[RFC9256](#)]. It implements a source routing policy to steer the service traffic from different source endpoints for a specific purpose (e.g., for a specific SLA). Based on the Parent SR policy described in [Section 3](#), a SR Policy Group can be built.

According to the service quality requirements, a unified service forwarding model is planned for nodes to determine the forwarding path of service flow. The traffic with the same service forwarding

model from different source endpoints to different destination endpoints uses the same SR Policy Group.

An SR Policy Group is identified by <color, endpoint> on the headend, same with an SR Policy. The color is an unsigned non-zero 32-bit integer value that associates the SR Policy Group with a service forwarding model (e.g., A set of SLA attributes). Different service qualities use different Color values.

The color value identifying the SR Policy Group corresponds to the Color attribute of the BGP route published by the endpoint. The destination endpoint publishes the BGP route and indicates which SR Policy Group path the headend should use to send packets to it through the Color attribute in the route.

In the Policy Group, establish the mapping relationship between the flow characteristics and the color value of the SR Policy path, and guide the service flows with different SLA requirements to the SR Policy path with different colors.

In a multi-site VPN scenario, there are multiple service types of traffic between each two nodes, and these traffics need to be forwarded on SR policy paths with different SLA. Through SR Policy Group, deployment can be simplified and the problem of complex configuration in multi-point scenarios can be effectively solved.

5. Steering into SR Policy Group

A headend can steer a packet flow into a SR policy group in various ways:

- * Per-flow Steering: Specify the mapping relationship between color and flow characteristics (such as DSCP) for SR policy group, and create a parent SR Policy that binds a specified destination endpoint address on the ingress node according the SR policy group. Upon receiving a packet with the specified destination address, the ingress node searches for the SR policy containing the color value mapped to the flow characteristics of the packet in the parent SR policy. The ingress node will use the matching SR policy to forward the packet.

The ingress node obtains a parent SR policy for traffic steering as follows:

- The destination endpoint publishes a BGP route with the specified Color extended community attribute.

- Get the color extended community attribute in the received BGP route.
- Match the color attribute value of the received BGP route with the SR Policy Group.
- Searches for a SR Policy Group with color matching the color extended community attribute.
- Searches for a Parent SR Policy with endpoint address matching the next hop in the BGP route, and recurses the BGP route to the parent SR policy.

The Ingress node can match flow characteristics in its ingress interfaces (upon any field such as Ethernet destination/source/VLAN/TOS or IP destination/source/DSCP or transport ports or application attribute etc.) and color them with an internal per-packet forwarding-class variable. According to the forwarding-class variable the ingress node selects the matching SR policy in the parent SR policy.

- * Policy-based Steering: incoming packets match a routing policy that directs them on a parent SR policy. Parse the flow characteristics (such as DSCP/802.1p value) from the packet header, find its corresponding color, and then match it to an SR policy in the parent SR policy, forward the incoming packets through the matched SR policy.

An SR Policy Group can be instantiated with SR Policies which are associated with different set of network resources (i.e., NRPs). Based on SR Policy Group, it is also a network slice deployment scheme for single user and multiple services. When different services are forwarded through different SR policy paths, different network resources can be used. After associating the SR policy with the network slice, different network slices can be used for forwarding different traffic of the same user.

When all the constituent SR policies in the parent SR policy are not valid, or all the selected paths of the SR policy are unavailable, the service traffic will not be forwarded according to the specified path. At this time, the best-effort forwarding path can be configured for the parent SR policy, and the endpoints through which traffic forwarding must pass can be designed in the best-effort forwarding path.

During network deployment, the best-effort forwarding path can be a SR policy path, an BE forwarding path, or a composite SR policy path of another color. Specify a best-effort forwarding path in the

parent SR policy. When all specified candidate paths are invalid, or the mapping relationship corresponding to their service type is not matched in the parent SR policy, select the default best-effort path forwarding.

6. Relationship Sorting

An SR Policy Group is associated with one or more constituent Parent SR Policies. The hierarchical relationship between SR policy group, Parent SR policy and SR policy is shown in the figure below.

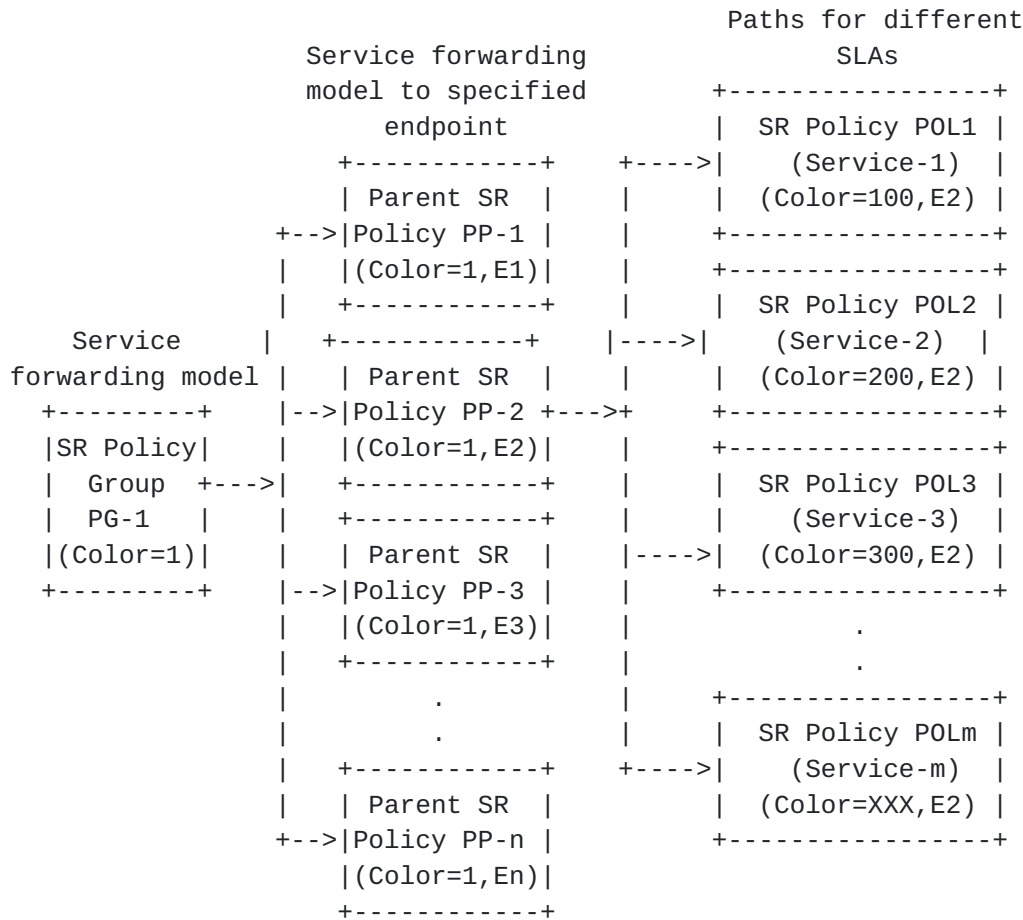


Figure 1 Hierarchical Relationship of SR Policy

The parent SR policy can be generated through static configuration, or dynamically generated when the destination endpoint accesses based on the service forwarding requirements specified by the SR policy group.

The following criteria apply for inclusion of constituent Parent SR Policies under a SR Policy Group:

* A SR Policy Group contains one or more Parent SR policies.

- * The colors of SR Policy group and its each constituent Parent SR Policy MUST be identical.
- * The colors of SR Policy group and its each constituent SR Policy of echo constituent Parent SR Policies MUST be different.
- * There can only be one Parent SR Policy with the same source endpoint and the same destination endpoint in the SR Policy group.

7. Information Model of SR Policy Group

In summary, the information model is the following:

SR Policy Group PG-1 <Color = 1>

Parent SR Policy PP-1<Color = 1, Endpoint = E1>

Service Service-1 mapping-to color 100

Service Service-2 mapping-to color 200

Service Service-3 mapping-to color 300

SR Policy POL1 <Headend = H1, Color = 100, Endpoint = E1>
 Candidate Path CP1 <Protocol-Origin = 20, Originator = 64511:192.0.2.1, Discriminator = 1>
 Preference 200
 Priority 10
 Segment List 1 <SID11...SID1i>

SR Policy POL2 <Headend = H1, Color = 200, Endpoint = E1>
 Candidate Path CP1 <Protocol-Origin = 20, Originator = 64511:192.0.2.1, Discriminator = 2>
 Preference 200
 Priority 10
 Segment List 1 <SID21...SID2i>

SR Policy POL3 <Headend = H1, Color = 300, Endpoint = E1>
 Candidate Path CP1 <Protocol-Origin = 20, Originator = 64511:192.0.2.1, Discriminator = 3>
 Preference 200
 Priority 10

Segment List 1 <SID31...SID3i>

Parent SR Policy PP-2<Color = 1, Endpoint = E2>

Service Service-1 mapping-to color 100

Service Service-2 mapping-to color 200

Service Service-3 mapping-to color 300

SR Policy POL4 <Headend = H1, Color = 100, Endpoint = E2>
Candidate Path CP1 <Protocol-Origin = 20, Originator =
64511:192.0.2.1, Discriminator = 4>
Preference 200
Priority 10
Segment List 1 <SID41...SID4i>

SR Policy POL5 <Headend = H1, Color = 200, Endpoint = E2>
Candidate Path CP1 <Protocol-Origin = 20, Originator =
64511:192.0.2.1, Discriminator = 5>
Preference 200
Priority 10
Segment List 1 <SID51...SID5i>

SR Policy POL6 <Headend = H1, Color = 300, Endpoint = E1>
Candidate Path CP1 <Protocol-Origin = 20, Originator =
64511:192.0.2.1, Discriminator = 6>
Preference 200
Priority 10
Segment List 1 <SID61...SID6i>

The SR Policy Group PG-1 is identified by color. It has two constituent Parent SR Policies: PP-1 and PP-2. Each is identified by a tuple <Headend, Color, Endpoint>.

The SR Parent Policy PP-1 is identified by the tuple <Headend = H1, Color = 1, Endpoint = E1>. It has three constituent SR Policies: SR Policy POL1 SR Policy POL2 and SR Policy POL3.

The SR Policy POL1 is identified by the tuple <Headend = H1, Color = 100, Endpoint = E1>. It has one candidate paths: CP1.

The SR Policy POL2 is identified by the tuple <Headend = H1, Color = 200, Endpoint = E1>. It has one candidate paths: CP1.

The SR Policy POL3 is identified by the tuple <Headend = H1, Color = 300, Endpoint = E1>. It has one candidate paths: CP1.

The SR Parent Policy PP-2 is identified by the tuple <Headend = H1, Color = 1, Endpoint = E2>. It has three constituent SR Policies: SR Policy POL4 SR Policy POL5 and SR Policy POL6.

The SR Policy POL4 is identified by the tuple <Headend = H1, Color = 100, Endpoint = E2>. It has one candidate paths: CP1.

The SR Policy POL2 is identified by the tuple <Headend = H1, Color = 200, Endpoint = E2>. It has one candidate paths: CP1.

The SR Policy POL6 is identified by the tuple <Headend = H1, Color = 300, Endpoint = E2>. It has one candidate paths: CP1.

According to the service forwarding quality requirements, three forwarding paths (Color 100, 200, and 300) are planned.

The service forwarding model of PP-1 is adopted for the destination endpoint E1. According to service characteristics services to E1 are divided into three categories: service-1, service-2, and service-3. The service-1 service is forwarded according to the SR Policy POL1 path of Color 100. The service-2 service is forwarded according to the SR Policy POL2 path of Color 200. The services of service-3 are forwarded according to the SR Policy POL3 path of Color 300.

The destination endpoint E2 also uses the same service forwarding model. The traffic to E2 is differentiated in the same way, and the traffic is sent to E2 according to the SR policy path of Color 100, 200, and 300.

8. Use Cases

8.1. Parent SR Policy in L3VPN over TE Scenarios

In Figure 2, CE1 and CE2 belong to the same L3VPN and access the public network through PE1 and PE2 respectively. There are many

kinds of traffic between CE1 and CE2. When the ordinary traffic is too large, the forwarding of important traffic will be affected.

In order to ensure the forwarding quality of important services, the steering based on Forwarding class can be configured using parent SR policy. After the steering based on forwarding class is configured, the traffic of different service levels will be carried by the specified SR policy tunnel, which can effectively ensure the forwarding quality of important services with high service levels.

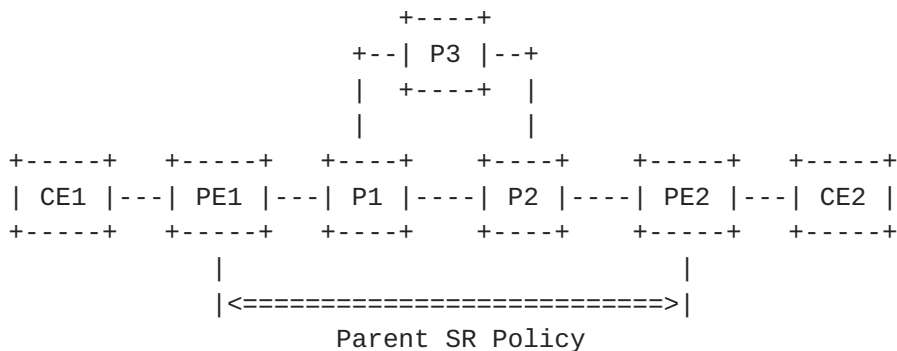


Figure 2 L3VPN over TE Scenario

It is assumed that in this network, the parent SR policy contains three constituent policies: Policy-A, Policy-B and Policy-C. Services with different forwarding class will carry different DSCP values in the packet. Identify the customer's service through DSCP on PE1. The voice traffic of VIP customers is forwarded according to the path of low-delay Policy-A, other traffic of VIP customers is forwarded according to the path of Policy-B, and all businesses of non-VIP customers are carried by Policy-C.

8.2. SR Policy Group in Multi-VPN Tenants Scenarios

In the L3VPN over TE application scenario shown in Figure 3, multi-VPN tenants are connected to the SRv6 network. Controller defines SR Policy Group for each VPN tenant to achieve different forwarding path resource between tenants.

Different VPNs use different SR Policy Groups with different colors. The ingress node generates different parent SRv6 policies as required according to the destination endpoint address dynamically. Since user's traffic of different services between two endpoints has different requirements for forwarding quality, identify the service type according to the DSCP of the packet, and steer the flow to the corresponding SR Policy path. The path constituting the SR Policy is calculated by the controller and distributed to the ingress node.

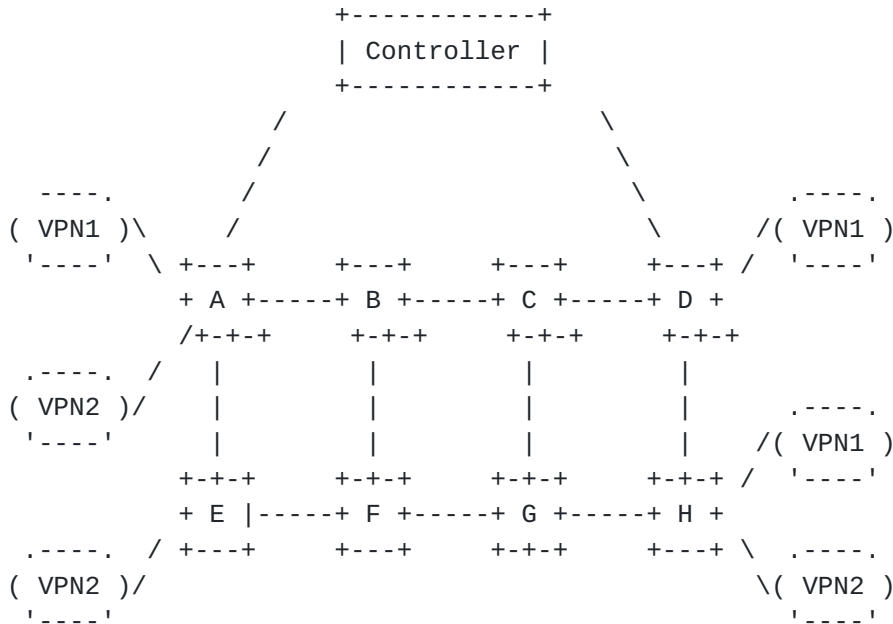


Figure 3 L3VPN over TE application scenario

VPN1 uses SR Policy group 1 identified by Color 100. Plan the forwarding path for VPN1 traffic, and allocate different sets of network resources for forwarding paths as blow:

- * Service-1: Voice service. Low delay forwarding is required, and the DSCP range of the packet is 1~10. The controller calculates the low delay path for the voice traffic, and maps the DSCP 1~10 to Color 500. The voice traffic is forwarded through the constituent SR policy (Color 500) of the Parent SRv6 Policy (Color 100) corresponding to VPN1.
- * Service-2: Video service. The DSCP range of the packet is 11~20, and low delay is not required. However, compared with the traffic of ordinary users, the Video traffic should be forwarded first. The controller calculates the SR Policy path and maps the DSCP 11~20 to Color 501. Video traffic is forwarded along the constituent SR policy (Color 501) of the Parent SRv6 policy (color 100) corresponding to VPN1.
- * Service-3: Non voice and video services. Low latency forwarding and priority forwarding are not required. The controller calculates the SR Policy path and maps all DSCP values outside the range of 1 to 20 to Color 502. Non voice and video traffic is forwarded along the constituent SR policy (Color 502) of the Parent SRv6 policy (color 100) corresponding to VPN1.




Figure 4 SR Policy Group for VPN1

The traffic of VPN1 from A to D of Service-1 will be forwarded based on the SR policy (Headend=A, Color=500, Endpoint=D) of Parent SR policy (Color=100, Endpoint=D) of SR Policy Group1.

Similarly, the traffic of VPN1 from A to H of Service-1 will be forwarded based on the SR policy (Headend=A, Color=500, Endpoint=H) of Parent SR policy (Color=100, Endpoint=H) of SR Policy Group1.

VPN2 uses SR Policy Group 2 identified by Color 101. Plan the forwarding path for VPN2 traffic. Different sets of network resources are further allocated for the forwarding paths as blow:

- * Service-4: Voice service. Low latency forwarding is required. The DSCP range of the packet is 1 to 10. The controller calculates the low delay path for voice service and maps the DSCP 1~10 to Color 600. The voice traffic is forwarded through the constituent SR policy (Color 600) of the Parent SRv6 Policy (color 101) corresponding to VPN2.
- * Service-5: Non voice services. No special forwarding quality requirements. The controller calculates the SR Policy path and maps all DSCP values outside the range of 1 to 10 to Color 601. Non voice service messages are forwarded along the constituent SR policy (Color 601) of the Parent SR policy (color 101) corresponding to VPN2.

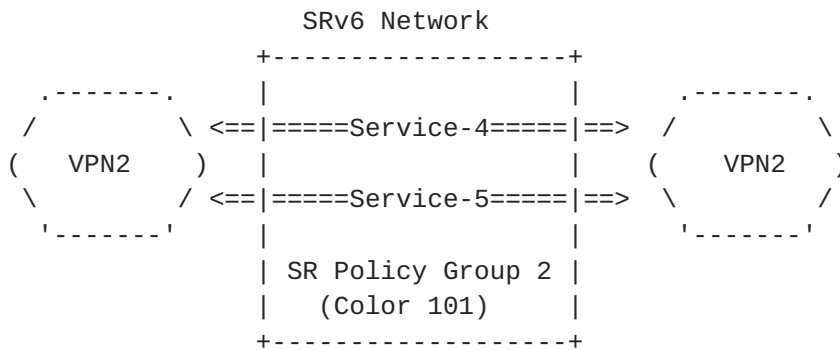


Figure 5 SR Policy Group for VPN2

The traffic of VPN2 from A to H of Service-4 will be forwarded based on the SR policy (Headend=A, Color=600, Endpoint=H) of Parent SR policy (Color=101, Endpoint=H) of SR Policy Group2.

Similarly, the traffic of VPN2 from A to E of Service-1 will be forwarded based on the SR policy (Headend=A, Color=600, Endpoint=E) of Parent SR policy (Color=101, Endpoint=E) of SR Policy Group1.

Because there are many access endpoints and each endpoint may act as an ingress node, compared with the traditional method of distributing service forwarding policies on each ingress node, the above SR Policy Group can greatly simplify the configuration of VPN access endpoints and effectively improve the efficiency of network deployment and operation and maintenance.

For reliability reasons, a default best-effort forwarding path can also be configured for each VPN tenant. If all the SR policy forwarding paths of the SR policy group are invalid, the default path can be used. The endpoints through which traffic forwarding must pass can be designed in the default path.

9. IANA Considerations

This document has no IANA actions.

10. Security Considerations

This document presents use cases to be considered by the deployment of SR Policy. It does not introduce any security considerations.

11. References

11.1. 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>>.
- [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>>.
- [RFC9256] Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", [RFC9256](#), DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [RFC 8402](#), DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [I-D.ietf-idr-segment-routing-te-policy] Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., Rosen, E., Jain, D., and S. Lin, "Advertising Segment Routing Policies in BGP", [draft-ietf-idr-segment-routing-te-policy-20](#) (work in progress), July 2022.

11.2. Informative References

TBD

12. Acknowledgments

The authors would like to thank the following for their valuable contributions of this document:

TBD

Authors' Addresses

Weiqiang Cheng
China Mobile
china

Email: chengweiqiang@chinamobile.com

Wenyong Jiang
China Mobile
China

Email: jiangwenyong@chinamobile.com

Changwang Lin
New H3C Technologies
china

Email: linchangwang.04414@h3c.com

Yuanxiang Qiu
New H3C Technologies
China

Email: qiuyuanxiang@h3c.com

YaWei Zhang
Huawei Technologies
China

Email: zhangyawei@huawei.com

Ran Chen
ZTE Corporation
China

Email: chen.ran@zte.com.cn

Yanrong Liang
Ruijie Networks
China

Email: liangyanrong@ruijie.com.cn