

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 28, 2022

Z. Li
Z. Hu
J. Dong
Huawei Technologies
October 25, 2021

Intent-based Routing
draft-li-teas-intent-based-routing-00

Abstract

This document defines the intent-based routing mechanism through which the packet can carry the intent information and the network node can enforce the policy according to the intent information (typically steering the packet into the SR policy or the underlay slice which can meet the intent). The intent-based routing mechanism provides a simple and scalable solution to meet the different service requirements for the inter-domain routing.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

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 April 28, 2022.

Copyright Notice

Copyright (c) 2021 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 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.

Table of Contents

1. Introduction	2
2. Terminologies	3
3. Intent-based Routing	3
4. Illustration	5
5. IPv6 Encapsulation	8
6. Security Considerations	9
7. IANA Considerations	9
8. References	9
8.1. Normative References	9
8.2. Informative References	10
Authors' Addresses	11

1. Introduction

[I-D.hegde-spring-mpls-seamless-sr] describes the requirements for end-to-end intent-based paths spanning multi-domain networks. [I-D.kaliraj-idr-bgp-classful-transport-planes] specifies the BGP based mechanisms to signal the packet paths which span multiple domains and provide different SLA characteristics. Since these SR paths need to setup according to the pair <color, endpoint>, it means more SR paths are introduced and this will cause more challenges on scalability.

In order to reduce the challenge of scalability introduced by the inter-domain routing with different service requirements, this document proposes the intent-based routing mechanism through which the packet can carry the intent information and the network node can steer the packet into the SR policy to satisfy the service requirement (that is, meet the specific intent). With the intent-based routing mechanism, network nodes do not need to maintain the fine-granularity connection state for each destination in the control plane, which can improve the scalability of the end- to-end routing significantly.

Besides steering the packet into the SR policy, the intent-based routing mechanism can also be used to steer the traffic into the

underlay network slice to meet the specific intent or enforce policy for other intents such as network measurement, security, etc. Since the same intent can be satisfied by different solutions in the different network domain, the intent-based routing also improve the flexibility to satisfying the service requirement through the combined solutions for the same intent.

2. Terminologies

The following terminologies are used in this document.

SR: Segment Routing

SRv6: Segment Routing over IPv6

3. Intent-based Routing

The Intent-based routing mechanism introduces the concept of intent as the information carried in the data plane to represent the specific service requirement for the destination on the network. The intent can be associated with a series of service attributes, such as low latency and high bandwidth. The value can be allocated by the administrator. The allocation of values of the intent in the multiple domain must be consistent.

[I-D.ietf-spring-segment-routing-policy] defines the color used for the SR policy. The color is a 32-bit numerical value that associates the SR Policy with an intent (e.g. low-latency). There can be the mapping as follows between the color and the intent. If the intent and the color can be designed and allocated consistently, the value of the color can be the same as that of the intent and the mapping between the color and the intent can be saved in the data plane.




Figure 1 Mapping between Intent and Color

Figure 1: Figure 1: Reference Topology

In the scenario of the inter-domain routing, the SR policy group for a specific Endpoint shown in the Figure 2 can be set up in the data plane in the local network domain. That is, it is not necessary to advertise the pair <color, endpoint> to set up the end-to-end SR path. When the packet carrying the intent information arrives at the

edge node of the network domain, the edge node can search the SR policy group according to the destination, then steer the packet into the corresponding SR policy according to the mapping between the color and the intent and the mapping between the color and the SR policy.




Figure 2: Figure 2: SR Policy Group

In the scenario of the inter-domain network slicing, the following mapping between the color and the local underlay network slice can be set up in the data plane in the local network domain. When the packet carrying the intent information arrives at the edge node of the network domain, the edge node can steer the packet into the local underlay network slice according to the mapping between the color and the intent and the mapping between the color and the local underlay network slice.




Figure 3: Figure 3: Mapping between Color and Underlay Network Slice

Since the same Intent may be satisfied by the SR policy or the underlay network slice, the local network domain can choose the different solutions flexibly without the need of coordination with other network domains. This can also improve the flexibility of the inter-domain routing.

Besides steering the packet into the SR policy or the underlay network slice, the network node can also enforce the policy for other possible intents such as network measurement, security, etc. This will be defined in the future version of the draft.

4. Illustration

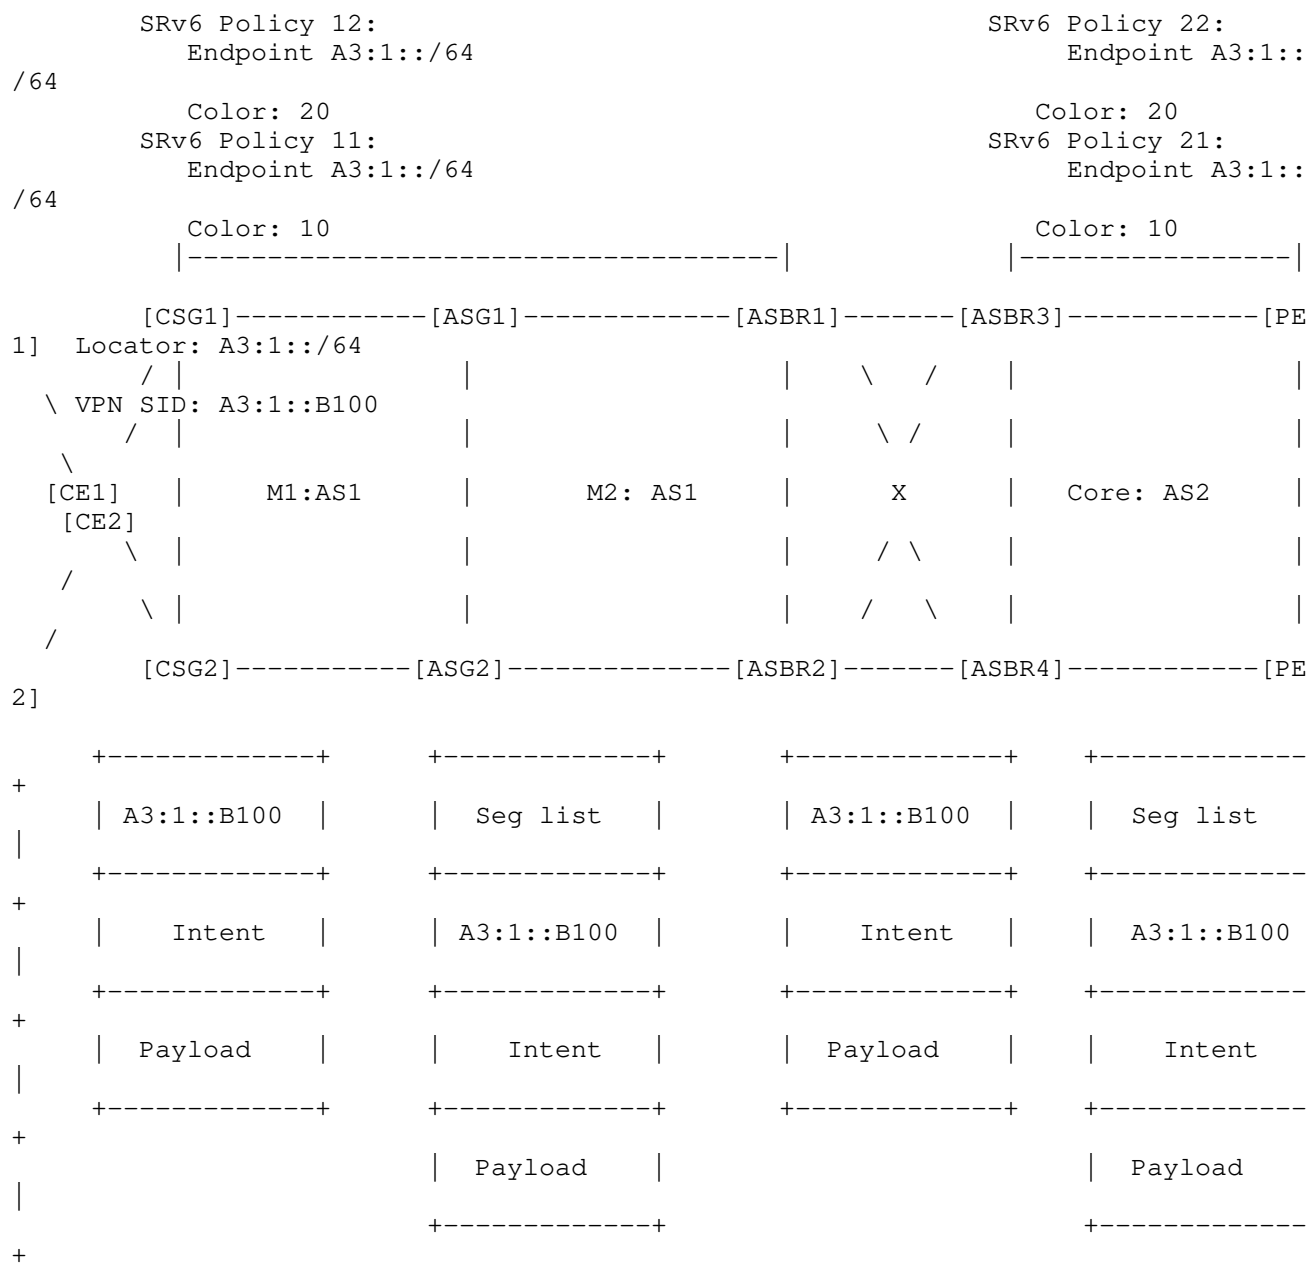


Figure 4: Figure 4: Illustration of Intent-based Inter-domain Routing

Figure 4 shows an example of a service provider network that comprises of two Autonomous systems, AS1 and AS2. The customer requests a leased line that requires bandwidth guarantee from CSG1 to PE1. Assume that the following is applied in the network shown in the Figure 4:

- o Independent ISIS instance in core (C) region.
- o Independent ISIS instance in Metro1 (M1) region.
- o Independent ISIS instance in Metro2 (M2) region.

- o BGP between ASBRs
- o PE1's locator is A3:1::/64, and VPN SID is A3:1::B100.
- o Core's aggregated routes are redistributed from Core to M (M1 and M2).

- o SRv6 policy group is set up in the AS1 between the CSG1 and ASBR1. It includes two SRv6 policies with the same Endpoint A3:1::/64 and color 10 and 20 respectively.
- o SRv6 policy group is set up in the AS2 between the ASBR3 and PE1. It includes two SRv6 policies with the same Endpoint A3:1::/64 and color 10 and 20 respectively.

PE1 advertises the VPN route with color 10 to CSG1. After CSG1 receive the VPN route, it maps color to the Intent and installs the VPN route with VPN SID A3:1::B100 and the corresponding intent. When CSG1 receives a packet from CE1, assume that CE1 finds the VPN route and the forwarding process is as follows:

1. CE1 encapsulates a new IPv6 header to the packet with the destination IPv6 address set as VPN SID A3:1::B100 and the Intent in the packet.
2. CE1 can search the forwarding entry according to the destination IPv6 address A3:1::B100 and the Intent.
3. After CE1 finds the SRv6 Policy 11 with the color 10, it encapsulates the new IPv6 header with the corresponding segment list to the packet.
4. The packet is forwarded to ASBR1 and the segment list is decapsulated at ASBR1.
5. ASBR1 can send the packet to ASBR3 according to the destination address A3:1::B100 by IPv6 forwarding process.
6. ASBR3 searches the forwarding entry according to the destination IP address A3:1::B100 and the Intent.
7. ASBR3 finds the SRv6 policy 21 with the color 10 and encapsulates the new IPv6 header with the corresponding segment list to the packet.
8. The packet is forwarded to PE1 and the segment list is decapsulated at PE1.
9. The packet is forwarding in the corresponding VPN instance identified by the destination IPv6 address A3:1::B100.

information carried can be read by the destination node along the path.

Besides the Intent option, the intent can also be carried combining with Application-aware Networking ([I-D.li-apn-framework]). [I-D.li-apn-header] and [I-D.li-apn-ipv6-encap] defines that the intent can be carried in the APN header which is encapsulated in the APN option in the IPv6 data plane.

6. Security Considerations

TBD

7. IANA Considerations

TBD

8. References

8.1. 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>>.
- [RFC7356] Ginsberg, L., Previdi, S., and Y. Yang, "IS-IS Flooding Scope Link State PDUs (LSPs)", RFC 7356, DOI 10.17487/RFC7356, September 2014, <<https://www.rfc-editor.org/info/rfc7356>>.
- [RFC7490] Bryant, S., Filsfils, C., Previdi, S., Shand, M., and N. So, "Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)", RFC 7490, DOI 10.17487/RFC7490, April 2015, <<https://www.rfc-editor.org/info/rfc7490>>.
- [RFC8400] Chen, H., Liu, A., Saad, T., Xu, F., and L. Huang, "Extensions to RSVP-TE for Label Switched Path (LSP) Egress Protection", RFC 8400, DOI 10.17487/RFC8400, June 2018, <<https://www.rfc-editor.org/info/rfc8400>>.
- [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>>.

- [RFC8665] Psenak, P., Ed., Previdi, S., Ed., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", RFC 8665, DOI 10.17487/RFC8665, December 2019, <<https://www.rfc-editor.org/info/rfc8665>>.
- [RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", RFC 8667, DOI 10.17487/RFC8667, December 2019, <<https://www.rfc-editor.org/info/rfc8667>>.
- [RFC8679] Shen, Y., Jeganathan, M., Decraene, B., Gredler, H., Michel, C., and H. Chen, "MPLS Egress Protection Framework", RFC 8679, DOI 10.17487/RFC8679, December 2019, <<https://www.rfc-editor.org/info/rfc8679>>.

8.2. Informative References

- [I-D.hegde-spring-mpls-seamless-sr]
Hegde, S., Bowers, C., Xu, X., Gulko, A., Bogdanov, A., Uttaro, J., Jalil, L., Khaddam, M., Alston, A., and L. M. Contreras, "Seamless SR Problem Statement", draft-hegde-spring-mpls-seamless-sr-06 (work in progress), September 2021.
- [I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", draft-ietf-spring-segment-routing-policy-14 (work in progress), October 2021.
- [I-D.kaliraj-idr-bgp-classful-transport-planes]
Vairavakkalai, K., Venkataraman, N., Rajagopalan, B., Mishra, G., Khaddam, M., Xu, X., Szarecki, R. J., and D. J. Gowda, "BGP Classful Transport Planes", draft-kaliraj-idr-bgp-classful-transport-planes-12 (work in progress), August 2021.
- [I-D.li-apn-framework]
Li, Z., Peng, S., Voyer, D., Li, C., Liu, P., Cao, C., Mishra, G., Ebisawa, K., Previdi, S., and J. N. Guichard, "Application-aware Networking (APN) Framework", draft-li-apn-framework-03 (work in progress), May 2021.

[I-D.li-apn-header]

Li, Z. and S. Peng, "Application-aware Networking (APN) Header", draft-li-apn-header-00 (work in progress), October 2021.

[I-D.li-apn-ipv6-encap]

Li, Z. and S. Peng, "Application-aware IPv6 Networking (APN6) Encapsulation", draft-li-apn-ipv6-encap-00 (work in progress), October 2021.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, DOI 10.17487/RFC5226, May 2008, <<https://www.rfc-editor.org/info/rfc5226>>.

[RFC5462] Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field", RFC 5462, DOI 10.17487/RFC5462, February 2009, <<https://www.rfc-editor.org/info/rfc5462>>.

Authors' Addresses

Zhenbin Li
Huawei Technologies
Beijing 100095
China

Email: lizhenbin@huawei.com

Zhibo Hu
Huawei Technologies
Beijing 100095
China

Email: huzhibo@huawei.com

Jie Dong
Huawei Technologies
Beijing 100095
China

Email: jie.dong@huawei.com