IDR
Internet-Draft
Intended status: Standards Track
Expires: November 11, 2021

S. Peng
Y. Liu
ZTE Corporation

May 10, 2021

BGP Tunnel Encapsulation Attribute Extensions for Network Slicing draft-peng-idr-bgp-tea-extensions-network-slicing-00

Abstract

This document defines extension to BGP Tunnel Encapsulation attribute to provide network slicing information needed to create tunnels and their corresponding encapsulation headers.

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 November 11, 2021.

Table of Contents

1. Introduction 2
2. Requirements Language 3
3. Gap Analysis of Tunnel Encapsulation Attribute 3
3.1. Additional Rules to Filter Traffic 3
3.2. Additional Tunnel Information for Network Slice 4
4. BGP Flow Specification Considerations 5
5. TEA Extensions 5
5.1. Flow Classification Sub-TLV (Type Code TBA1) 5
5.1.1. IP Differentiated Service sub-sub-TLV 5
5.1.2. IP Source Address Range sub-sub-TLV 6
5.1.3. IP Destination Address Range sub-sub-TLV 7
5.1.4. IP Protocol Number Range sub-sub-TLV 7
5.1.5. Transport Source Port Range sub-sub-TLV 7
5.1.6. Transport Destination Port Range sub-sub-TLV 8
5.1.7. Ethernet Frame related sub-sub-TLVs 8
5.2. Virtual Network Sub-TLV (Type Code TBA2) 8
5.3. SR-BE Encapsulation Sub-TLV 10
6. Examples 10
6.1. IP Flex-algo Examples 10
6.2. SR Flex-algo Examples 12
6.3. Specifying Network Slice Identifier Examples 13
7. IANA Considerations 14
7.1. BGP Tunnel Encapsulation Attribute Sub-TLVs 14
7.2. sub-sub-Types of Flow Classification Sub-TLVs 14
7.3. BGP Tunnel Encapsulation Attribute Tunnel Types 15
8. Security Considerations 15
9. Acknowledgements 15
10. Normative References 15
Authors' Addresses 17

1. Introduction

[I-D.ietf-teas-ietf-network-slices] describes network slicing in the context of networks built from IETF technologies. In order to provide network slicing in the operators network for different scenarios, some existing control plane technologies are utilized, and also some new technologies are developed. A network slice can be a virtual network or a traffic engineering path with guaranteed resources. For example, it could be implemented as an IGP Multi-Topology (see [RFC5120], [RFC4915], [RFC5340]), or an IGP Flexible Algorithm (see [I-D.ietf-lsr-flex-algo]), or a Slice Aggregate which comprises of multiple IETF network slice traffic streams(see [I-D.bestbar-teas-ns-packet]), or a simple SR policy(see [I-D.ietf-spring-segment-routing-policy]).

Once a network slice is created, it is necessary to configure the corresponding traffic mapping policy on the entry node of the slice to steer the traffic to the that slice. For example, ACL rules can be configured on the entry node of the slice to filter the traffic based on 5-tuple fields or other fields (such as Differentiated Services) of the packets, then steer the matched traffic to that slice; It is also possible to firstly set a Color for the matched traffic, then steer the traffic to an SR policy. However, such configuration is inflexible, especially for the case where the slice entry node is not the endpoint of overlay service, in this case it is not recommended to configure a large number of service-related policies on the slice entry node. From the point of view of simplifying operation and maintenance, automatic slice steering is necessary.

[RFC9012] defines a BGP path attribute known as the "Tunnel Encapsulation attribute", which can be used with BGP UPDATEs of various Subsequent Address Family Identifiers (SAFIs) to provide information needed to create tunnels and their corresponding encapsulation headers. This document describes extensions to Tunnel Encapsulation attribute to specify forwarding path within particular network slice.

2. 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.

3. Gap Analysis of Tunnel Encapsulation Attribute

3.1. Additional Rules to Filter Traffic

[RFC9012] defines two Sub-TLVs, i.e., Protocol Type Sub-TLV (Type Code 2) and Color Sub-TLV (Type Code 4), for aiding tunnel selection.

For Protocol Type Sub-TLV, the value fileld contains a 2-octet value from IANA's "ETHER TYPES" registry, such as IPv4(0x0800), IPv6(0x86dd), MPLS(0x8847), to indicate the type of the payload packets that are allowed to be encapsulated with the tunnel parameters that are being signaled in the parent TLV. However, for network slicing case, it is not enough that the traffic of different slices is distinguished based on the coarse-grained ethernet types, more fine-grained differentiation is needed, such as Differentiated Services field or 5-tuple fields of the IP packet, or Source/ Destination MAC or VLAN ID/PCP of the Ethernet frame.

For Color Sub-TLV, the value field contains a Color Extended Community to "color" the corresponding Tunnel TLV. In more detailed, as section "8. Recursive Next-Hop Resolution" of [RFC9012] described, an overlay route (U1) with a Color Extended Community that is identified in one of Color sub-TLVs of the underlay route (U2) can use tunnel identified in the Tunnel Encapsulation attribute of the underlay route (U2). That is, an overlay route with specific Color Extended Community indicates the expected TE purpose during recursive next-hop resolution, while an underlay route with specific Color sub-TLV of Tunnel Encapsulation attribute indicates the "color" of the corresponding Tunnel. In fact, the corresponding Tunnel itself may not have a color attribute, and the color is temporarily set by the underlay route. It is possible that different colors may be set by different underlay routes. Another example that corresponding TE path itself has a color attribute is SR policy defined in [I-D.ietf-spring-segment-routing-policy]. For network slicing case, although a Color Extended Community can be contained in a route UPDATE to indicate TE purpose within the specific network slice, that means any packets matched that route will have the same forwarding behavior. In some cases these packets may need different treatment. It is possible to advertise multiple Color Extended Community for the same route, however, ACL rules is necessary at entry nodes of the slice to firstly set color for packets and then quest TE purpose, that is inflexible.

3.2. Additional Tunnel Information for Network Slice

[RFC9012] does not define how to specify the tunnel within a network slice. For example, Tunnel Encapsulation Attribute with IP-in-IP Tunnel type will only let the packet be encapsulated in IP-tunnel created in physical network. It is useful to combine the existing Tunnel Egress Endpoint Sub-TLV or BGP next-hop with a Network Slice Identifier to select tunnel within expected network slice. Note that [RFC9012] defines that some of the tunnel types (for example, VXLAN and NVGRE) that can be specified in the Tunnel Encapsulation attribute have an encapsulation header containing a virtual network identifier of some sort, however they are different from the semantic and function of network slice.

In some network slicing techniques, SID allocation has distinguished different network slices, so SR-BE (Best Effort) can be specified directly in Tunnel Encapsulation Attribute. Although, [RFC9012] defines Prefix-SID sub-TLV, it is the Prefix-SID that the tunnel's egress endpoint uses to represent the prefix appearing in the NLRI field of the BGP UPDATE to which the Tunnel Encapsulation attribute is attached, it is not the Prefix-SID of the outer SR-BE tunnel.

4. BGP Flow Specification Considerations

[RFC8955] defines Flow Specification NLRI and also specifies BGP Extended Community encoding formats used to propagate Traffic Filtering Actions along with the Flow Specification NLRI. The corresponding flow routes can be installed on the entry nodes of the network slice, and that is similar with manually configured policy based routes. However, both of them are inflexible, and some implementations may not be easy to route packets in combination with flow routes and unicast routes. Instead, prefix NLRI containing Tunnel Encapsulation Attribute is different with Flow Specification NLRI, it does not need to introduce flow route state on the entry node of the network slice, and all forwarding behaviors are based on traditional unicast route with its TEA attributes.

5. TEA Extensions

5.1. Flow Classification Sub-TLV (Type Code TBA1)

The Flow Classification Sub-TLV MAY be included in a given Tunnel Encapsulation TLV to indicate the flow classification of the payload packets that are allowed to be encapsulated with the tunnel parameters that are being signaled in the parent TLV. It MUST NOT appear more than once in its parent TLV. The Value field of the Flow Classification Sub-TLV comprised of multiple sub-sub-TLVs.

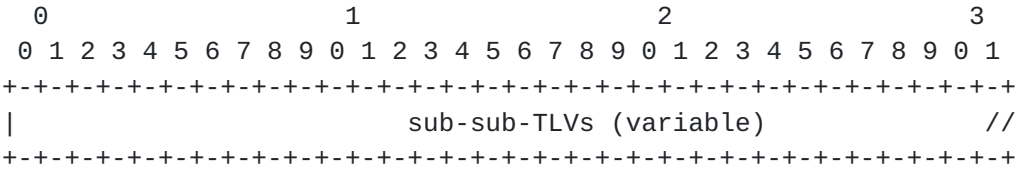

Figure 1: Flow Classification Sub-TLV Value Field

5.1.1. IP Differentiated Service sub-sub-TLV

IP Differentiated Service sub-sub-TLV is used to represent the range of Differentiated Service field (such as the TOS field of the IPv4 header or the TC field of the IPv6 header) that traffic needs to match. It can appear more than once in its parent sub-TLV.

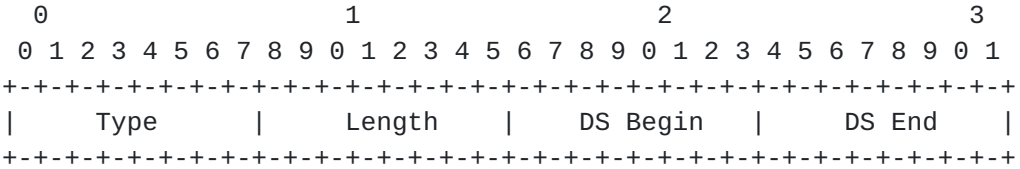

Figure 2: IP Differentiated Service sub-sub-TLV Format

Type: TBD (suggest 1)

Length: 2 octets.

DS Begin: The begin value of the range of Differentiated Service.

DS End: The end value of the range of Differentiated Service. DS Begin and DS End field can be same to specify actions for each DS value.

5.1.2. IP Source Address Range sub-sub-TLV

IP Source Address Range sub-sub-TLV is used to represent the range of Source Address field that traffic needs to match. It can appear more than once in its parent sub-TLV.

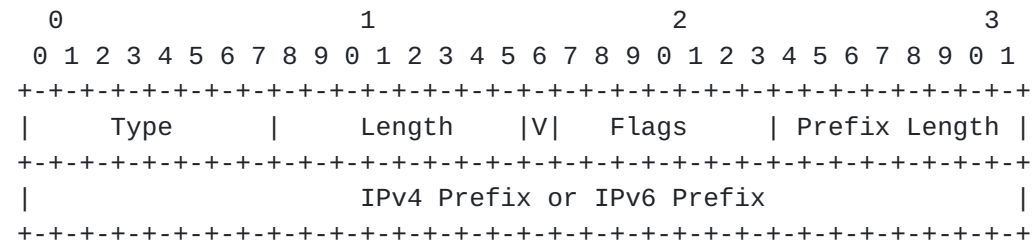

Figure 3: IP Source Address Range sub-sub-TLV Format

Type: TBD (suggest 2)

Length: variable.

Flags: Currently only V-flag is defined. The IP Prefix field contains an IPv4 Prefix if V-Flag is clear and an IPv6 Prefix if V-Flag is set. The remaining bits are reserved for future use. They MUST be set to zero on transmission and MUST be ignored on receipt.

Prefix Length: Number of bits in the Prefix field. MUST be from the range (1 - 32) when V-Flag is clear and (1 - 128) when V-Flag is set. The sub-sub-TLV MUST be ignored if the Prefix Length is outside of this range.

IP Prefix: The IP Prefix is encoded in the minimal number of octets for the given number of bits. Trailing bits MUST be set to zero and ignored when received.

5.1.3. IP Destination Address Range sub-sub-TLV

IP Destination Address Range sub-sub-TLV is used to represent the range of Destination Address field that traffic needs to match.

It has the same format as IP Source Address Range sub-sub-TLV, suggest Type Code 3.

5.1.4. IP Protocol Number Range sub-sub-TLV

IP Protocol Number Range sub-sub-TLV is used to represent the range of IP Protocol Number field that traffic needs to match. It can appear more than once in its parent sub-TLV.

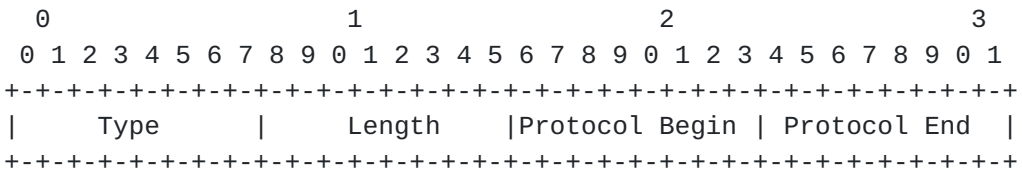

Figure 4: IP Protocol Number Range sub-sub-TLV Format

Type: TBD (suggest 4)

Length: 2 octets.

Protocol Begin: The begin value of the range of IP Protocol Number.

Protocol End: The end value of the range of IP Protocol Number. Protocol Begin and Protocol End field can be same to specify actions for each Protocol value.

5.1.5. Transport Source Port Range sub-sub-TLV

Transport Source Port Range sub-sub-TLV is used to represent the range of Source Port field that traffic needs to match. It can appear more than once in its parent sub-TLV.

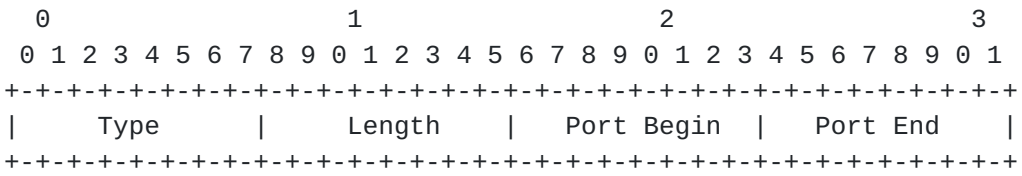

Figure 5: Transport Source Port Range sub-sub-TLV Format

Type: TBD (suggest 5)

Length: 2 octets.

Port Begin: The begin value of the range of Source Port.

Port End: The end value of the range of Source Port. Port Begin and Port End field can be same to specify actions for each Port value.

5.1.6. Transport Destination Port Range sub-sub-TLV

Transport Destination Port Range sub-sub-TLV is used to represent the range of Destination Port field that traffic needs to match. It can appear more than once in its parent sub-TLV.

It has the same format as Transport Source Port Range sub-sub-TLV, suggest Type Code 6.

5.1.7. Ethernet Frame related sub-sub-TLVs

To be defined in future versions.

5.2. Virtual Network Sub-TLV (Type Code TBA2)

The Virtual Network Sub-TLV MAY be included in a given Tunnel Encapsulation TLV to indicate the expected network slice that the traffic is steered to. It MUST NOT appear more than once in its parent TLV. The Value field of the Virtual Network Sub-TLV has the following format:

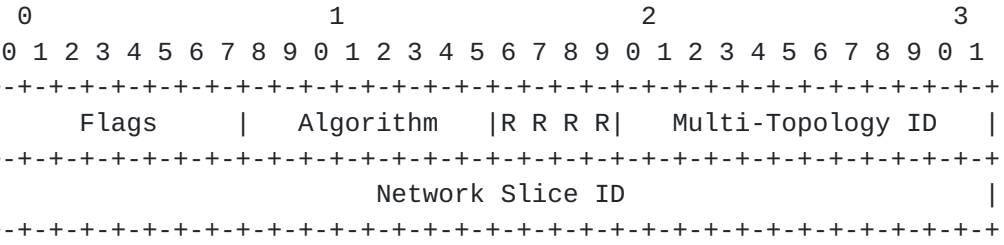

Figure 6: Virtual Network Sub-TLV Value Field

Flags: The following flags are defined:

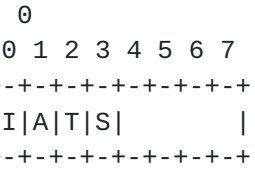

where:

I-Flag: This flag is only valid when S-Flag is set. If I-Flag is set, the entry node of network slice SHOULD insert the value of Network Slice ID field into the outer encapsulated tunnel header, otherwise no necessary.

A-Flag: If A-Flag is set, the Algorithm field contains valid value.

T-Flag: If T-Flag is set, the Multi-Topology ID field contains valid value.

S-Flag: If S-Flag is set, the Network Slice ID field contains valid value.

The remaining bits are reserved for future use. They MUST be set to zero on transmission and MUST be ignored on receipt.

Algorithm: Represents IGP algorithm, see IANA "IGP Algorithm Types" registry, such as 0 (Shortest Path First algorithm based on link metric), 1 (Strict Shortest Path First algorithm based on link metric) defined in [RFC8402], and 128~255 (Flexible Algorithm) defined in [I-D.ietf-lsr-flex-algo].

Multi-Topology ID: Represents IGP Multi-Topology defined in [RFC5120] and [RFC4915].

Network Slice ID: Represents an IETF network Slice defined in [I-D.ietf-teas-ietf-network-slices].

In most cases, it only need to set one of Algorithm, Multi-Topology ID and Network Slice ID. However, it is possible to set them at the same time. When not set, the value is 0.

Virtual Network Sub-TLV is used together with other Sub-TLVs to encapsulate the tunnel information corresponding to the specified tunnel in a specific virtual network. For example, when the Tunnel Type of Tunnel Encapsulation TLV is 7 (representing IP-in-IP), and contains Virtual Network Sub-TLV and Tunnel Egress Endpoint Sub-TLV, it means that the packets is encapsulated in the IP tunnel in the expected virtual network destinated to the Tunnel Egress Endpoint. In detailed, the path of IP tunnel could be determined by <Algorithm, Endpoint>, or <MT-ID, Endpoint>, or <Slice-ID, Endpoint>.

In some network slice control plane schemes, Slice-ID is directly used to identify FIB table, so it is easy to get FIB entry by <Slice-ID, Endpoint>. In other shcemes, Slice-ID need firstly be mapped to an SA-ID, then get FIB entry by <SA-ID, Endpoint>. It is the local matter for the entry node of the network slice to get FIB entry according to the specific deployment mode. In any case, the entry

node of the network slice can, but may not if it has not such capability, insert Slice-ID to the outer tunnel header.

5.3. SR-BE Encapsulation Sub-TLV

This document introduce a new Tunnel Type, termed as SR-BE. For SR-BE tunnel, the structure of the Value field in the Encapsulation sub-TLV (Type Code 1) is shown as the following:

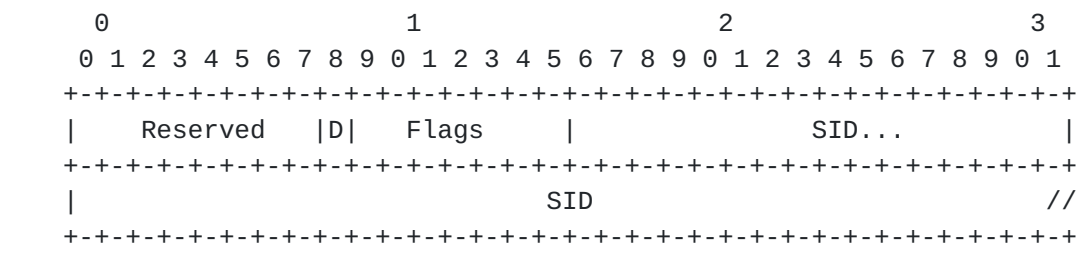

Figure 7: SR-BE Encapsulation Sub-TLV Value Field

Reserved: MUST be set to zero on transmission and disregarded on receipt.

Flags: Currently only D-flag is defined. The SID field contains a 32 bits SR-MPLS SID (index) if D-Flag is clear and a 128 bits SRv6 SID if D-Flag is set. The remaining bits are reserved for future use. They MUST be set to zero on transmission and MUST be ignored on receipt.

SID: 32 bits SR-MPLS SID (index) or 128 bits SRv6 SID.

For SR-MPLS SID (index), the out-label needs to be obtained based on the SRGB of the downstream node and then pushed to the label stack. For SRv6 SID, the SID is filled in the DA field of outer IPv6 header.

6. Examples

6.1. IP Flex-algo Examples

The following example describes a most simple network slice deployment selection, which steers traffic to the corresponding tunnel endpoint according to the traffic class. In the backbone network of some operators, network administrators do not want to deploy a large number of traffic engineering paths in the network, but they also want to automatically select the appropriate path for forwarding according to the traffic characteristics. The network administrator chooses to deploy IGP flex-algo in the backbone network of pure IPv6 (refering to [I-D.ietf-lsr-ip-flexalgo]), and creates an flex-algo plane based on the calculation path of Delay-metric. It is

expected that the traffic with higher traffic class will be forwarded along the flex-algo plane, while the ordinary traffic will continue to be forwarded along the physical network.

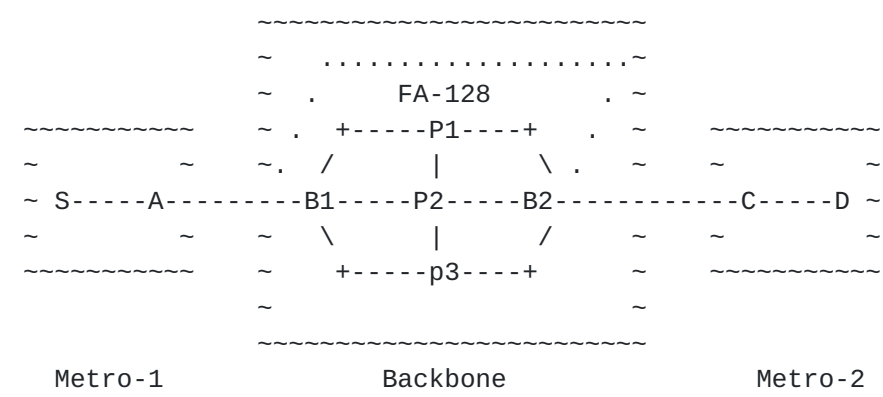

Figure 8: Flex-algo in Backbone

Nodes B1, P1, P2, B2 and their interconnected links join to the flex-algo 128 plane. Suppose that there are two loopback routes on node B2, respectively loopback-B2 and loopback-B200, and loopback-B2 is associated with algorithm 0, loopback-B200 is associated with algorithm 128. On node B1, the route forwarding path to loopback-B2 will be the forwarding path calculated by the minimum IGP metric in the physical network, and the route forwarding path to loopback-B200 will be the forwarding path calculated by the minimum delay metric in flex-algo 128 plane.

Node S of Metro-1 needs to send IPv6 packets to node D of Metro-2. Suppose that node D has a local route prefix-D, which is flooded in Metro-2. Node C can advertise prefix-D to node B2 through BGP. When the B2 receives prefix-D from outside the backbone domain, it continues to advertise the routes to the boundary node B1 through BGP. At this time, B2 does not pay attention to the difference of the announced prefixes (i.e. whether it is for prefix-D or other prefix-D'). It just configure a simple local policy to add the Tunnel Encapsulation Attribute to the BGP UPDATE that continues to be advertised, which includes two Tunnel Cncapsulation TLVs:

The first Tunnel Encapsulation TLV:

Tunnel Type: IP in IP

Tunnel Egress Endpoint Sub-TLV: Loopback-B2

Flow Classification Sub-TLV: IP Differentiated Service is [0,3]

The second Tunnel Encapsulation TLV:

Tunnel Type: IP in IP

Tunnel Egress Endpoint Sub-TLV: Loopback-B200

Flow Classification Sub-TLV: IP Differentiated Service is [4,7]

Node B1 will maintain to the prefix-D route entry, which contains the above two Tunnel Encapsulation Attribute options.

Assuming the two data packets, P1 and P2, sent from node S to node D. Suppose that the traffic class in IPv6 header of P1 is 0 and that in IPv6 header of P2 is 7.

When the above packets reache the node B1 of the backbone network, it will match the route entry prefix-D and encapsulate the outer IPv6 tunnel for the packets according to the Tunnel Encapsulation Attribute options. Since the traffic class of P1 is 0, the DA of the encapsulated outer IPv6 header is loopback-B2; while the traffic class of P2 is 7, the DA of the encapsulated outer IPv6 header is loopback-B200.

The above two packets will be forwarded to the destination node B2 along the physical topology and flex-algo 128 plane respectively.

6.2. SR Flex-algo Examples

The network administrator can also deploy IGP flex-algo in the backbone network of SRv6 (refering to [I-D.ietf-lsr-flex-algo]).

In Figure 8, suppose that there are two SRv6 Locators on node B2, respectively LOC-B2 and LOC-B200, and LOC-B2 is associated with algorithm 0, LOC-B200 is associated with algorithm 128. On node B1, the route forwarding path to LOC-B2 will be the forwarding path calculated by the minimum IGP metric in the physical network, and the route forwarding path to LOC-B200 will be the forwarding path calculated by the minimum delay metric in flex-algo 128 plane.

Suppose SID-b2 is allocated in LOC-B2 and SID-B200 is allocated in LOC-B200. These two SIDS may be END SID with USD flavor or END.DT6 SID used to carry global IPv6 packets.

Similar with the above example, B2 sends to B1 about BGP UPDATE with Tunnel Encapsulation Attribute, which includes two Tunnel Cncapsulation TLVs:

The first Tunnel Encapsulation TLV:

Tunnel Type: SR-BE

SR-BE Encapsulation Sub-TLV: D-Flag is set, SID is SID-B2

Flow Classification Sub-TLV: IP Differentiated Service is [0,3]

The second Tunnel Encapsulation TLV:

Tunnel Type: SR-BE

SR-BE Encapsulation Sub-TLV: D-Flag is set, SID is SID-B200

Flow Classification Sub-TLV: IP Differentiated Service is [4,7]

For the above packet P1, the DA of the encapsulated outer IPv6 header is SID-B2; For the above packet P2, the DA of the encapsulated outer IPv6 header is SID-B200. Thus they will be forwarded to the destination node B2 along the physical topology and flex-algo 128 plane respectively.

6.3. Specifying Network Slice Identifier Examples

This example describes steering to a specific network slice according to the traffic level, which requires the entry node to combine the specific Network Slice Identifier with the BGP next-hop of BGP UPDATE or the Tunnel Egress Endpoint Sub-TLV to determine the tunnel encapsulation to be adopted. In this way, the tunnel selection of the entry node is more flexible, and the constructure of BGP UPDATE is more concise.

Similar with the above example, B2 sends to B1 about BGP UPDATE with Tunnel Encapsulation Attribute, which includes two Tunnel Cncapsulation TLVs:

The first Tunnel Encapsulation TLV:

Tunnel Type: Any-Encapsulation

Virtual Network Sub-TLV: Algorithm 0, MT-ID 0, Slice-ID 0

Flow Classification Sub-TLV: IP Differentiated Service is [0,3]

The second Tunnel Encapsulation TLV:

Tunnel Type: Any-Encapsulation

Virtual Network Sub-TLV: Algorithm 128, MT-ID 0, Slice-ID 0

Flow Classification Sub-TLV: IP Differentiated Service is [4,7]

Node B1 will maintain to the prefix-D route entry, which contains the tunnel encapsulation attribute information and specifically contains the above two Tunnel Encapsulation options. Because the Tunnel Type is Any-Encapsulation, B1 node needs to select the corresponding tunnel according to the actual deployment modes in the backbone network. For example, if the backbone network is an SR-MPLS network, it needs to find the prefix-SID allocated by the node B2 for the corresponding algorithm in the link-state database, and then obtain the corresponding MPLS SR-BE tunnel; If the backbone network is SRv6 network, it needs to find the END SID assigned by the node B2 for the corresponding algorithm in the link state database, and then obtain the corresponding IPv6 SR-BE tunnel.

7. IANA Considerations

7.1. BGP Tunnel Encapsulation Attribute Sub-TLVs

This document request the following entries to the "BGP Tunnel Encapsulation Attribute Sub-TLVs" registry:

Value	Description	Reference
TBA1	Flow Classification	This Document
TBA2	Virtual Network	This Document

7.2. sub-sub-Types of Flow Classification Sub-TLVs

This document request new registry named as "sub-sub-Types of Flow Classification Sub-TLVs" under the "Border Gateway Protocol (BGP) Tunnel Encapsulation" grouping.

The initial types for this new registry are indicated as the following:

Value	Description	Reference
1	IP Differentiated Service	This Document
2	IP Source Address Range	This Document
3	IP Destination Address Range	This Document
4	IP Protocol Number Range	This Document
5	Transport Source Port Range	This Document
6	Transport Destination Port Range	This Document

7.3. BGP Tunnel Encapsulation Attribute Tunnel Types

This document request the following entries to the "BGP Tunnel Encapsulation Attribute Tunnel Types" registry:

Value	Description	Reference
TBA3	SR-BE	This Document

8. Security Considerations

This document inherits the security consideration from [RFC9012].

9. Acknowledgements

TBD

10. Normative References

[I-D.bestbar-teas-ns-packet]
Saad, T., Beeram, V. P., Wen, B., Ceccarelli, D., Halpern, J., Peng, S., Chen, R., Liu, X., and L. M. Contreras, "Realizing Network Slices in IP/MPLS Networks", draft-bestbar-teas-ns-packet-02 (work in progress), February 2021.

[I-D.ietf-lsr-flex-algo]
Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", draft-ietf-lsr-flex-algo-15 (work in progress), April 2021.

[I-D.ietf-lsr-ip-flexalgo]
Britto, W., Hegde, S., Kaneriya, P., Shetty, R., Bonica, R., and P. Psenak, "IGP Flexible Algorithms (Flex-Algorithm) In IP Networks", draft-ietf-lsr-ip-flexalgo-02 (work in progress), April 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-11 (work in progress), April 2021.

[I-D.ietf-teas-ietf-network-slices]
Farrel, A., Gray, E., Drake, J., Rokui, R., Homma, S., Makhijani, K., Contreras, L. M., and J. Tantsura, "Framework for IETF Network Slices", draft-ietf-teas-ietf-network-slices-00 (work in progress), April 2021.

[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>.

[RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", RFC 4915, DOI 10.17487/RFC4915, June 2007, <https://www.rfc-editor.org/info/rfc4915>.

[RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/RFC5120, February 2008, <https://www.rfc-editor.org/info/rfc5120>.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <https://www.rfc-editor.org/info/rfc5340>.

[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>.

[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>.

[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>.

[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>.

Authors' Addresses

Shaofu Peng
ZTE Corporation

Email: peng.shaofu@zte.com.cn

Yao Liu
ZTE Corporation

Email: liu.yao71@zte.com.cn