SPRING Working Group
Internet Draft
Intended status: Standards Track
Expires: December 09, 2025

Y. Liu
China Mobile
C. Lin
H. Li
New H3C Technologies
L. Gong
China Mobile

June 07, 2025

NRP ID in SRv6 segment
draft-liu-spring-nrp-id-in-srv6-segment-06

Abstract

This document proposes a method to carry the NRP-ID with the packet in the SRv6 segment.

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 December 09 2025.

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 (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.

Table of Contents

1. Introduction

Network slicing provides the ability to partition a physical network into multiple isolated logical networks of varying sizes, structures, and functions so that each slice can be dedicated to specific services or customers. [I-D.ietf-teas-ietf-network-slices] defines the term "IETF Network Slice" and establishes the general principles of network slicing in the IETF context. [I-D.cheng-teas-

network-slice-usecase] describes several use cases of IETF Network Slice. [I-D.ietf-teas-ns-ip-mpls] proposes a solution to realize network slicing in IP/MPLS networks. Network nodes need to identify a packet belonging to a network slice before it can apply the proper forwarding treatment, so a slice ID must be carried in each packet.

Packets belong to a network slice need to be forwarded using the specific network resources. [I-D.draft-ietf-teas-ietf-network-slices] defines the network resource mapped to the network slice as NRP, that is, the Network Resource Partition, and defines the nrp-id to identify the NRP used in the forwarding process.

In a network that provides slicing services, the NRP-ID can be carried in the packet. In the process of packet forwarding, the routers on the forwarding path can extract NRP-ID from the packet, determine the NRP to which the packet belongs, and then forward the packet using the resources associated with the NRP.

Segment Routing (SR) allows a headend node to steer a packet flow along any path. Per-path states of Intermediate nodes are eliminated thanks to source routing. The headend node steers a flow into an SR Policy. The packets steered into an SR Policy carry an ordered list of segments associated with that SR Policy.

When SRv6 network provides network slicing service, it is also necessary to consider how to carry NRP-ID with packet. This document proposes a method to carry the NRP-ID in the SRv6 network. By setting the NRP-ID in the SRv6 segment, the SRv6 endpoint or transit node can be aware the NRP to which the packet belongs and carry out relevant forwarding processing.

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. Description of Applicability

2.1. Applicable scenario

The method proposed in this document is applicable to slice forwarding scenarios based on SRv6 TE mode. The strict hop-by-hop forwarding path is specified to meet the SLA requirements of customer traffic through SRv6 policy.

An example is the SRv6 based hierarchical slice architecture described in [draft-gong-teas-hierarchical-slice-solution]. As per the document, the level-1 slice defines the topology through flex-algo, and ensures the bandwidth through the sub-interface with a dedicated queue. The network topology of the level-2 slice is defined by a set of SRv6 policies on different ingress nodes, and the bandwidth is guaranteed through virtual data channel with a dedicated queue under the sub-interface. Therefore, when the traffic is forwarded in the level-2 slice, the forwarding path needs to be specified through the SRv6 policy. At the same time, the NRP-ID needs to be carried with the message, which is used to associate with the virtual data channel under the outgoing interface during forwarding, so as to obtain bandwidth guarantee.

Through the method proposed in this document, the NRP-ID can be carried in the segment list encoded in the SRH, and there is no need to add an IPv6 extension header.

This document also considers the scenario of non-strict hop-by-hop forwarding paths, Please refer to Section 7 for relevant analysis.

2.2. Global and local NRP-ID

As per [draft-ietf-teas-ietf-network-slices], NRP is a collection of resources (bufferage, queuing, scheduling, etc.) in the underlay network.

The custom traffic belonging to a certain network slice will be forwarded in the corresponding NRP.

For an NRP, each router in the slice domain can use the same NRP-ID to map local resources. This document calls such NRP-IDs global NRP-IDs. NRP-ID needs to be unique in the slice domain.

On the other hand, different NRP-IDs are used on each router to map local resources such as queues. This NRP-ID can be called local NRP-ID. Local NRP-ID only needs to be unique within the router. The local NRP-ID can map any local resource, not limited to links and queues, not even limited to the network slicing.

This document describes how to carry NRP-ID through the argument field of segment. This method can meet the scenarios of both global and local NRP-ID.

3. Encoding NRP-ID in SRv6 segment

The structure of SRv6 segment is defined in [RFC8986]. An SRv6 segment consists of three parts, LOC:FUNCT:ARG.

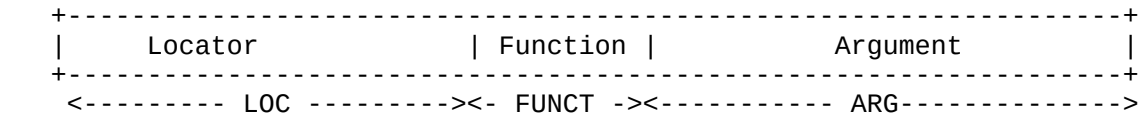

Figure 1: structure of segment

After the packet enters the SRv6 domain, the ingress node (headend) adds SRv6 Encapsulation to packet. In SRv6 TE (traffic engineer) mode, the headend node encapsulates an IPv6 header and an SRH header at the same time. A group of SRv6 segments is encapsulated in the SRH header to indicate the forwarding path. NRP-ID can be carried in segment to identify the NRP to which the packet belongs.

This document proposes to use the ARG part of the segment to carry the NRP-ID.

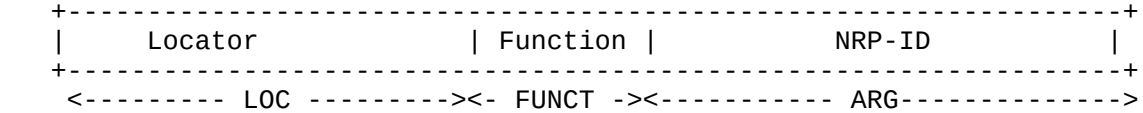

Figure 2: Encoding NRP-ID in segment

4. Deployment consideration of NRP-ID In segment

In the SRv6 TE mode, multiple segments are encoded in the SRH. The last segment in the SRH is usually the service or End SID of the tailend node and does not need to carry the NRP-ID.

Other segments in SRH are usually End or End.X segments, which are used to guide intermediate endpoint nodes to forward packets. The arguments of these SIDs are not defined in [RFC8986] and can carry NRP-IDs.

Different segments can carry the same or different NRP-ID, which is arranged by the controller or operator by CLI according to the actual requirement.

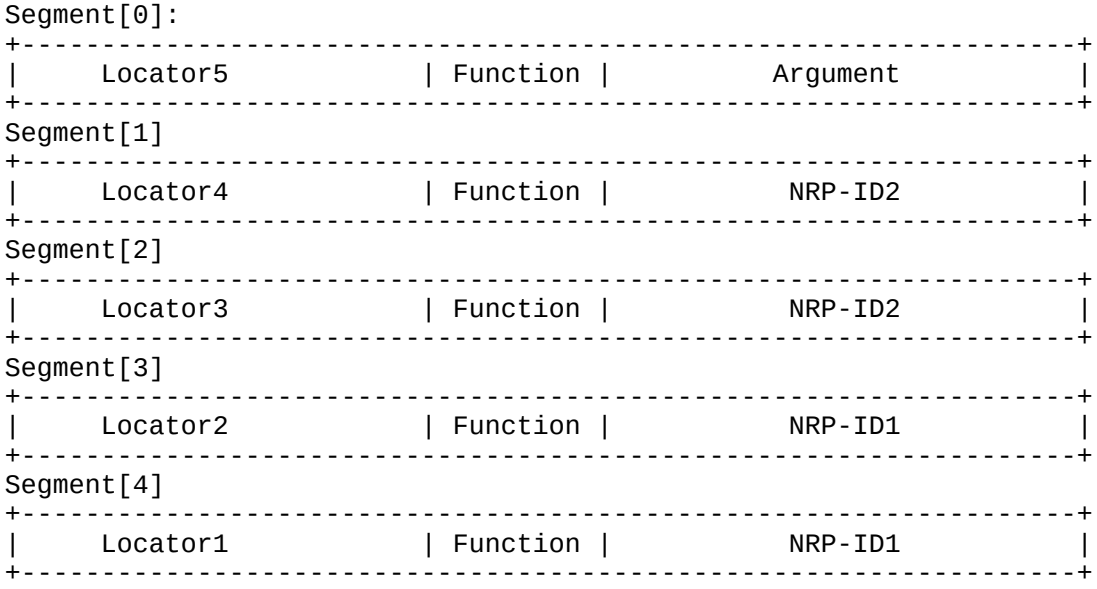

Figure 3: Encoding NRP-ID in segment list

4.1. Carrying Local NRP-ID

Local NRP-ID can be applied to non-network slicing, such as Detnet. One of the goals of DetNet is to provide bounded end-to-end latency for critical flows. [draft-chen-detnet-sr-based-bounded-latency] describes a method to implement bound latency, called Cycle Specified Queuing and Forwarding (CSQF).

By specifying the sending cycle of a packet at a node and making sure that the packet will be transmitted in that cycle, CSQF can achieve bounded latency within the node. By specifying the sending cycle at every node along a path, the end-to-end bounded latency can be achieved.

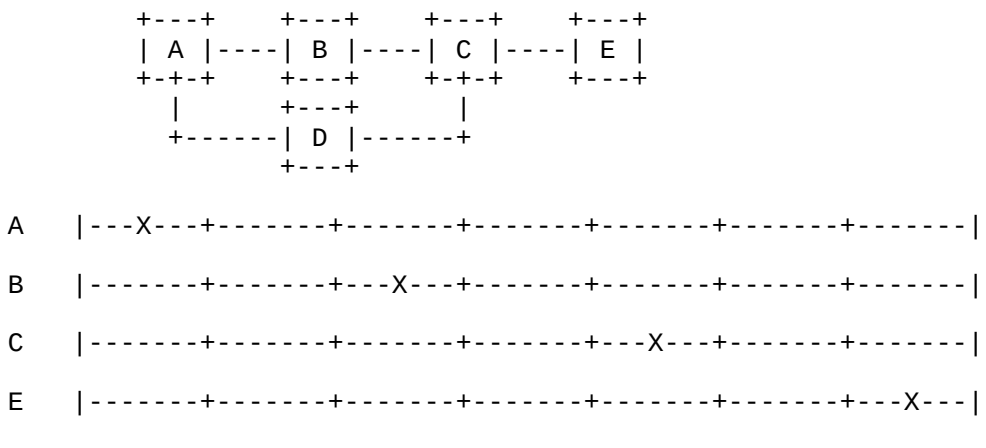

Figure 4: Example for CSQF

The sending cycle can be taken as a local NRP-ID and carried by SRH. After the packet reaches each endpoint, the local NRP-ID (cycle id) carried in the SRH can be used to specify the cycle in which the packet is expected to be sent, so as to achieve the end to end bound latency.

4.2. Carrying Global NRP-ID

4.2.1. Same Global NRP-ID in SRH

In a simple network that uniformly plans NRP, all endpoints share the same global NRP-ID for a NRP, and so the value of the NRP-ID encoded in the SRH is the same.

Referring to the network slicing scenario shown in the figure below, the network slicing is deployed in the entire transmission network, and all network nodes use the same NRP-ID to map the corresponding network slicing. In the example, network slice 1000 and network slice 2000 are created respectively. The corresponding NRP-IDs are 100 and 200, respectively. Therefore, the Global NRP-IDs encapsulated in the SRH are all the same.

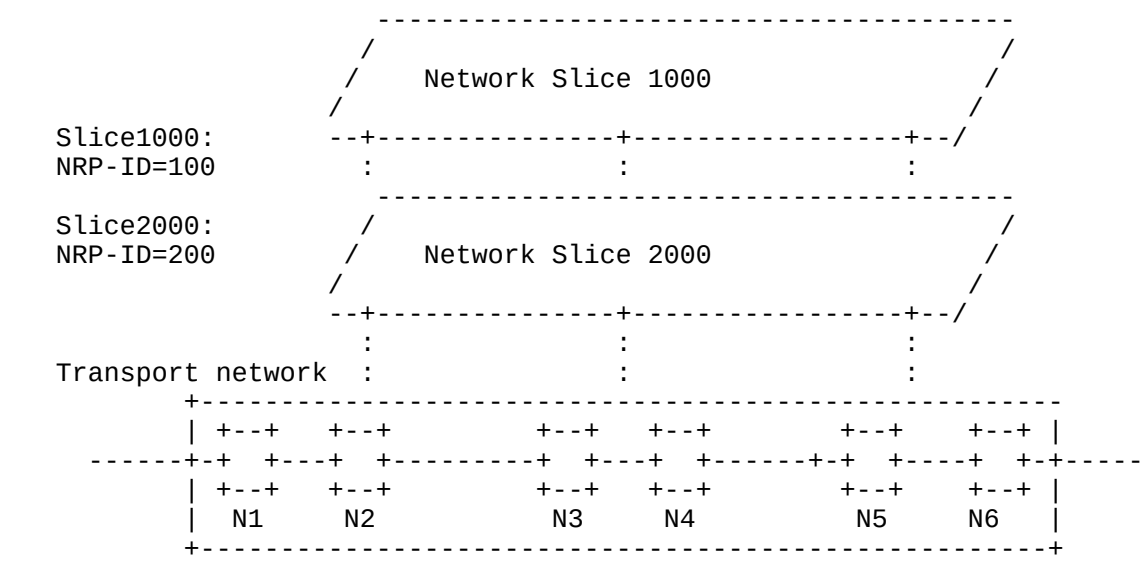

Segment[0]:

Locator5	Function	Argument

Segment[1]

Locator4	Function	Global-NRP-ID1

Segment[2]

Locator3	Function	Global-NRP-ID1

Segment[3]

Locator2	Function	Global-NRP-ID1

Segment[4]

Locator1	Function	Global-NRP-ID1

Figure 5: Example of Same Global NRP-ID

4.2.2. Different Global NRP-ID in SRH

However, in some scenarios, each endpoint will assign different NRP-IDs to the NRP serving the same network slice.

o Scenario1

Customers purchase bandwidth guarantee according to their actual needs. The granularity of bandwidth may vary from several hundred megabytes to several gigabytes. Due to the differences in the capabilities of various endpoint nodes, some endpoints may only provide 1G or 5G bandwidth reservation granularity, while some endpoints may provide 100M bandwidth reservation granularity.

As shown in the following figure, a network slice domain is composed of core routers (C1, C2, C3, and C4) and access routers (A1, A2). The minimum reserved bandwidth that the core router can provide is 1G, and the minimum routing bandwidth that the access router can provide is 200M.

It is assumed that the corresponding bandwidth guarantee is provided for two network slices according to customer requirements

Network slice1 for customer1 Reserved bandwidth 200M

Network slice2 for customer2 Reserved bandwidth 500M

Each endpoint creates dedicated queues on its own interfaces and assigns corresponding NRP-ID to associate with them.

For the core routers:

C1/C2/C3/C4: NRP-ID1 for network slice1

NRP-ID1 for network slice2

For the access routers:

A1/A2: NRP-ID1 for network slice1

NRP-ID2 for network slice2

For core routes, one NRP can be used to serve multiple network slices, so its NRP-ID is the same for multiple network slices. On the contrary, access routes use different NRPS to serve different network slices, so the NRP-IDs are different for multiple network slices.

Since the customer network is accessed through the access router, the bandwidth of the customer traffic can be limited through the access router. While on the core router, NRP can be shared by multiple network slices

In this case, for the traffic of customer 2, the forwarding path from A1 to A2 is assumed to be A1, C3, C4 and A2, and the NRP-IDs in the forwarding process are (NRP-ID2, NRP-ID1, NRP-ID1 and NRP-ID2), so the NRP-IDs encoded in the SRH are not the same vale.

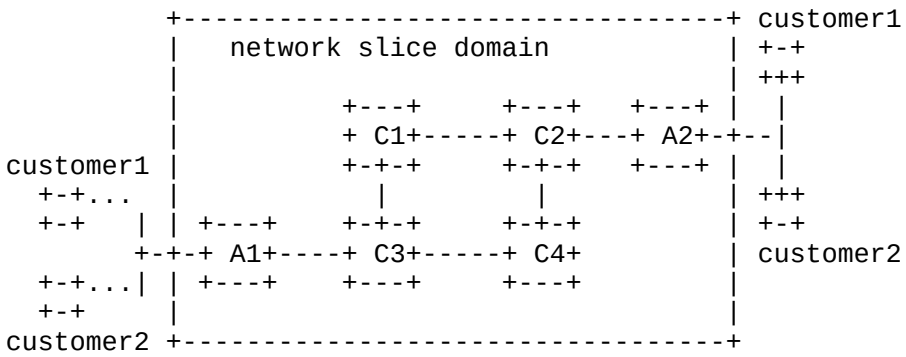

Figure 6: example topology1

o Scenario2

When network slices are deployed in multiple ASes, each AS creates an NRP according to the actual needs and reserves bandwidth resources as needed. Each AS may assign a different NRP-ID to the corresponding NRP. In this case, the NRP-IDs encapsulated in the SRH are different.

As an example, in the following figure, a network slice 2000 is created for a customer. Three ASes create NRPS for the network slice, and the corresponding NRP-IDs are 100,101 and 201 respectively.

In the left to right direction, create the intra-domain SRv6 policy on N1, N10 and N20 respectively, and allocate the corresponding binding SID as BSID1, BSID2 and BSID3

The end-to-end cross domains SRv6 policy is issued on the N1 node of AS100, and its segment list is <BSID1, BSID2, BSID3>

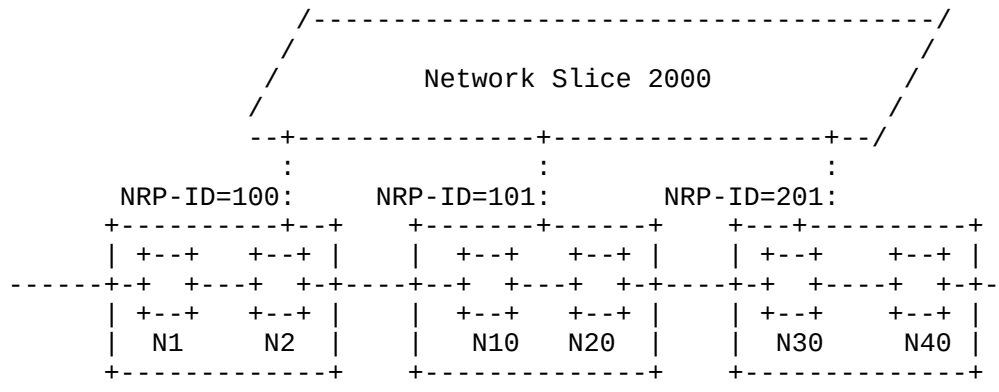

Figure 7: scenario of multi-as

In this case, when the NRP-IDs are carried through the segment, the corresponding NRP-ID needs to be set for each BSID, and its specific values are different. The segment list code of SRH header is as follows:

```
Segment[0]:
+-------------------------------------------------------------------+
|                           Service SID                             |
+-------------------------------------------------------------------+
Segment[1]
+-------------------------------------------------------------------+
|              BSID3               |          201                   |
+-------------------------------------------------------------------+
Segment[2]
+-------------------------------------------------------------------+
|              BSID2               |          101                   |
+-------------------------------------------------------------------+
Segment[3]
+-------------------------------------------------------------------+
|              BSID1               |          100                   |
+-------------------------------------------------------------------+
```

Figure 8: encoding example

Thus, when the packet enters each AS, the ASBR processes the packet according to the corresponding BSID, inherits the NRP-ID carried by the BSID, and sets it in the newly added SRH header to complete the slice forwarding within the AS.

5. NRP-ID position information advertisement

If the network slicing service needs to be supported, when creating a locator, the SRv6 node needs to determine the encoding position of NRP-ID in the segment according to its own role.

The locator and encoding positions of NRP-ID need to be advertised to the controller or headend nodes.

With this information, the controller or the head node will be aware of how to encode the NRP-ID in the segment list when the traffic is steered to the SRv6 policy.

Please note that the description in this section is only for the strict hop-by hop forwarding path scenario.

5.1. Static configuration mode

In the static configuration mode, configure the locator encoding information on the controller or headend nodes respectively. For the convenience of description, the locator carrying NRP-ID is named slice prefix in this document.

The aggregation attribute of locator can be used. The following figure is an example. The P nodes only provide End/End.x type segments, and the positions used to encode NRP-ID are usually the same. Therefore, common prefix can be configured to indicate the position of NRP-ID.

If the encoding position of a P node is different from that of most nodes, the slice prefix corresponding to the locator of the P node can be configured separately to specify its encoding position.

Referring to the topology and locators of each node in the Figure 4, the following slice prefix can be configured.

The encoding positions of P1 and P4 nodes are the same, and a slice prefix corresponding to common prefix can be configured to identify the coding position as the low 16 bits. The encoding position of P3 is 96bit to 112bit of segment, so a slice prefix corresponding to its locator is configured separately to explain its coding position

Slice-Prefix1: 2001:1:1::/64 (common prefix)

NRP-ID Position: [112..127]

Slice-Prefix2: 2001:1:1:0:130::/80 (locator for P3)

NRP-ID Position: [96..112] in segment

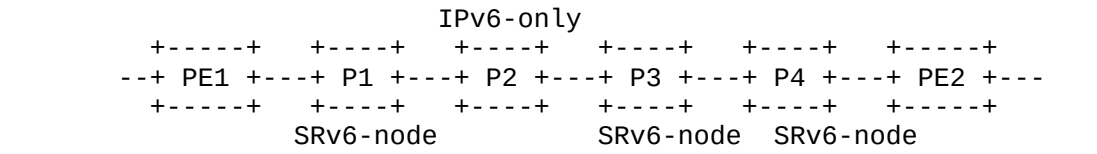

common Prefix: 2001:1:1::/64
locator
P1: 2001:1:1:0:110::/80 for End/End.x
P3: 2001:1:1:0:130::/80 for End/End.x
P4: 2001:1:1:0:140::/80 for End/End.x

Figure 9: example topology for slice prefixes

5.2. Dynamic advertising mode

To simplify the configuration, slice prefix can be advertised to network nodes in the domain through IGP and to controllers through BGP-LS. This reduces the configuration of controllers and SRv6 nodes.

Relevant protocol extensions will be provided in subsequent versions.

6. Behaviors of node

6.1. Behavior of headend

If the network slice function is enabled, the SRv6 headend node determines the network slice to which the customer traffic belongs according to the relevant policies.

The headend node steers the customer traffic to the SRv6 policy and encapsulates the IPv6 header and SRH header for the customer traffic. The headend node encapsulates the segment list of the SRv6 policy in the SRH header.

At the same time, set NRP-ID into segment. These NRP-IDs can be the same or different values according to actual requirement.

Generally, the nodes along the route of the message use the same NRP-ID to identify the NRP associated with the network slice.

Therefore, when the headend node encapsulates the segment list in the SRH, it writes the same NRP-ID into segments except the last segment.

In some special cases, such as cross domain scenarios, different NRP-IDs may be used on the forwarding path. In this case, the controller may need to write the NRP-ID into each segment of the segment list in advance, and then issue the SRv6 policy to the headend node. The headend node only needs to use SRv6 policy to encapsulate customer traffic.

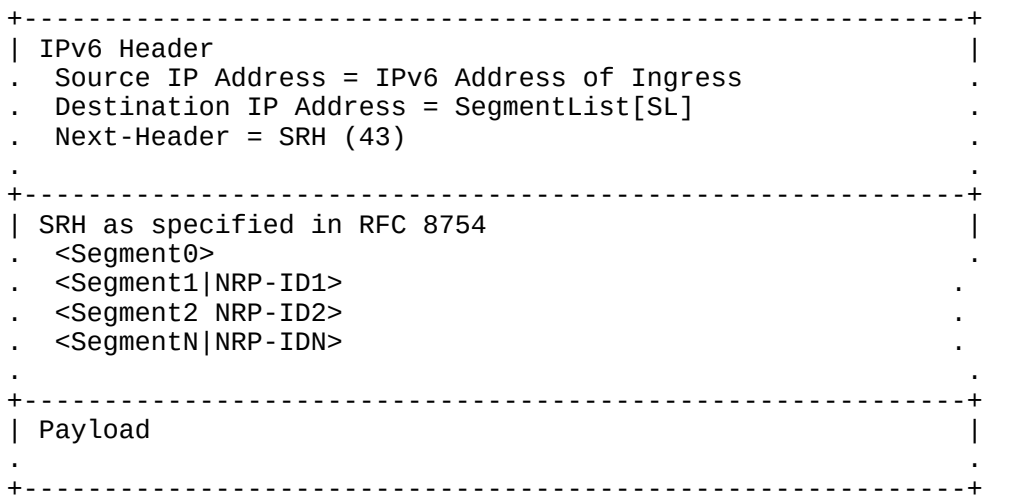

Figure 10: Format of SRv6 TE with slice ID

6.2. Behavior of endpoint

When a SRv6 node receives a packet, the destination address of the packet is the segment instantiated locally by the SRv6 node. At this time, the SRv6 node processes the packet as endpoint node. The endpoint node extracts the NRP-ID from the segment and forwards the packet with the NRP identified by the NRP-ID.

When N receives a packet whose IPv6 DA is S and S is a local End SID, the pseudo code processed is modified as follows based on RFC[8986]:

S01 - S11.
S12. Decrement IPv6 Hop Limit by 1
S13. Decrement Segments Left by 1
S14. Update IPv6 DA with Segment List[Segments Left]
Insert Extract NRP-ID from destination address.

Modify:
S15. Uses the NRP-ID to select a NRP and apply NRP policies to forward packet
S16. }

When N receives a packet whose IPv6 DA is S and S is a local End.X SID, the pseudo code processed is modified as follows based on RFC[8986]:

S01 - S11.
S12. Decrement IPv6 Hop Limit by 1
S13. Decrement Segments Left by 1
Insert Extract NRP-ID from destination address.
S14. Update IPv6 DA with Segment List[Segments Left]

Modify:
S15. Submit the packet to the IPv6 module for transmission to the new destination via a member of J with using the resource of NRP
S16. }

7. Consideration of transit node

For the network providing network slicing service, there may be some special scenarios due to the consideration of node capacity or business deployment, such as:

- some network nodes only support network slices and do not support srv6
- the SRv6 policy specifies a loose path

As shown in the following figure, B and C are non-SRv6 nodes, and other nodes are SRv6 nodes. Suppose SRv6 policy is created on node A, and the loose path is specified as (D.End, E.End, H.End).

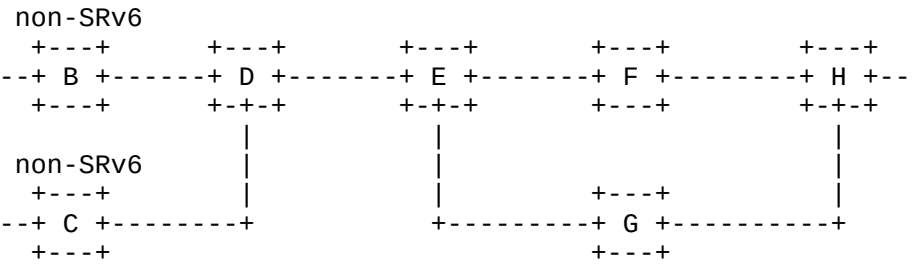

Figure 11: Example of transit node

In this way, during the forwarding process of the packet, either the non-SRv6 node (B or C) or the SRv6 node (F or G) may forward the packet as a transit node. In order to enable these nodes to extract the NRP-ID during forwarding, relevant processing needs to be added

In this scenario, the NRP-ID allocation of each node needs to be restricted accordingly. Take forwarding paths (D.End, E.End, H.End) as example, When node B (or C) receives a packet, its destination address is the segment of node D, then node B (or C) needs to extract the NRP-ID through the argument of the segment of node D and map it to the local NRP. Therefore, the NRP allocated by nodes B and C is required to be consistent with that of node D. The reverse packet is similar, and the NRP-ID allocation of nodes B and C is required to be consistent with that of node A. Therefore, it implies that four nodes A, B, C and D need to have the same NRP-ID allocation result. The same principle applies to nodes E, F, G and H.

Considering the flexibility of the loose path, if the loose path is deployed in the whole network slice domain, all nodes in the whole network slice domain are required to uniformly allocate NRP-IDs, that is, to ensure the consistency of NRP-IDs.

If multiple regions can be divided in the whole network slice domain, and strict paths (through End.X segments) are specified between the regions, the consistency of NRP-ID needs to be ensured in the region, and the NRP-ID of different regions can be different. For example, Node A/B/C/D and Node E/F/G/H are divided into two regions: REG1 and REG2, the segment list can be modified to (D.End.X, H.End). After combining NRP-ID, the segment list in the SRH which encapsulates the customer traffic may become (D.End.X|NRP-

ID1, H.End|NRP-ID2), that is, REG1 uses NRP-ID1, and REG2 uses NRP-ID2 to map the same NRP. Node B, C, F and G as transit nodes can also extract the correct NRP-ID from IPv6 destination address of customer traffic.

The following two sub sections describe the control and forwarding behavior of the transit node.

7.1. Creating LSPT

The transit node needs to extract the NRP-ID from the destination address of the packet. This destination address is not a locally instantiated segment, but a segment of the next endpoint.

Through the method described in Section 5, the slice prefix can be configured on the transit node, or the transit node can learn the slice prefix through protocol extension.

The network node can use these slice prefixes to create a local slice prefix table (LSPT) on the forwarding plane. When forwarding packets as transit node, the network node uses the destination address to lookup the LSPT according to the longest matching principle, and then extracts the NRP-ID from the destination address according to the information of the hit table entry.

7.2. Behavior of transit node

For the transit node, the destination address of the packet is not a local segment, and only IPv6 forwarding is performed for the packet.

The transit node may be a node that supports SRv6 or a node that only supports IPv6.

When processing SRv6 packets, the transit node can use the destination address to lookup the local slice prefix table according to the longest matching principle. If a prefix is hit, the NRP-ID could be extracted according to the configuration information.

Therefore, the processing pseudo code can be modified as follows:

```
S1.  If ((Slice forwarding is enabled) && (Destination address hits
network slice prefix)) {
S2      Extracts NRP-ID from destination address by using the
position information of hit prefix
S2.     Uses the NRP-ID to apply NRP policies to forward packet
S3.  }
S4.  Else {
S5.     Forwards the packet without applying any NRP policies
S6. }
```

8. Consideration of compress SRH

[draft-ietf-spring-srv6-srh-compression] describes how to compress the Segment list in SRH through two flavors, namely NEXT-C-SID and REPLACE-C-SID. This section briefly describes how to carry NRP-ID through argument in the case of SRH compression.

8.1. NEXT-C-SID flavor

According to the description in section 4.1 of [draft-ietf-spring-srv6-srh-compression], the structure of the NEXT-C-SID flavor SID is shown in the figure below, and the Argument is used to carry the CSID.

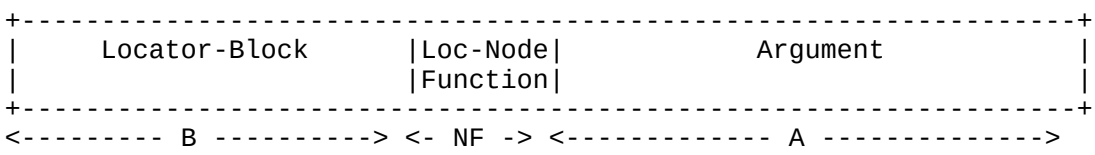

Figure 12: structure of NEXT-C-SID flavor SID

For this type of SID, Argument can be divided into two parts through planning, where Arg.Next is used to carry CSID, and Arg.Extend is used to carry NRP-ID.

In this way, all CSIDs in a container share the same NRP-ID. Different containers can carry different NRP-IDs as needed. This will inevitably lead to a reduction in the number of CSIDs carried by a container, and the corresponding forwarding behavior also needs to be adapted and modified.

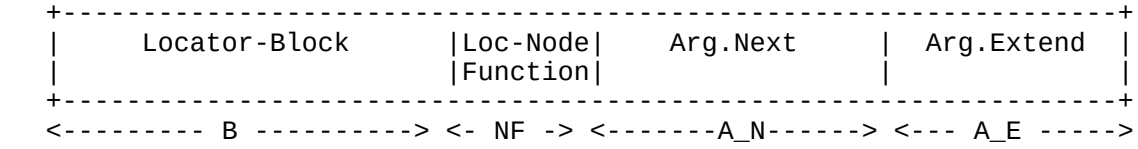

Figure 13: Modified structure of NEXT-C-SID flavor SID

8.2. REPLACE-C-SID flavor

According to the description of [draft-ietf-spring-srv6-srh-compression] in chapter 4.2, the structure of REPLACE-C-SID flavor SID is shown in the figure below. The Argument field is used to identify the number of remaining CSIDs in the container.

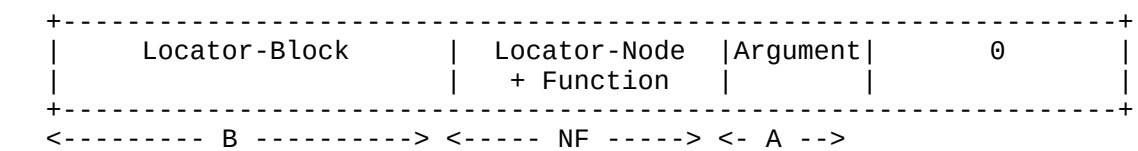

Figure 14: structure of REPLACE-C-SID flavor SID

For this type of SID, the Argument can also be divided into two parts, where Arg.Num is used to identify the number of remaining CSIDs in the container, and Arg.Extend is used to carry the NRP-ID. In this manner, all CSIDs of a compressed path share the same NRP-ID, and different compressed paths may carry different NRP-IDs as required. The corresponding forwarding behavior also needs to be adapted and modified.

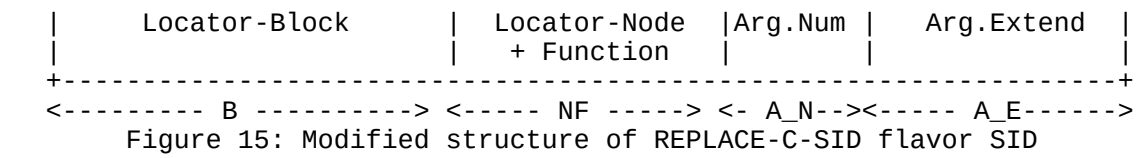

Figure 15: Modified structure of REPLACE-C-SID flavor SID

9. Example

As shown in the following figure, the IP backbone network deploys the network slicing service. The network operator has created two NRPs, NRP1 and NRP2. NRP1 guarantees 100Mbps bandwidth and NRP2 guarantees 200Mbps bandwidth. Set the IDs ID1 and ID2 for the two NRPs respectively.

The IP backbone network provides customers with two network slices. Network slice1 is mapped to NRP1 and network slice2 is mapped to NRP2. SRv6 is used to carry traffic and network slicing services.

Along with the forwarding path <PE1-P1-P2-PE2>, dedicated queues with guaranteed bandwidth for NRP1 and NRP2 are configured at corresponding interfaces of each router. Taking the interface P1-P2 of router P1 as an example, Queue 1 is configured with NRP-ID1 and guaranteed bandwidth of 100Mbps, and Queue 2 is configured with NRP-ID2 and 200Mbps. When P1 transmits a packet through interface P1-P2, the NRP-ID carried in the destination address is checked.

If ID1 is encapsulated in the destination address, P1 uses queue 1 to transmit the packet. If ID2 is encapsulated in the destination address, P1 uses queue 2 to transmit the packet.

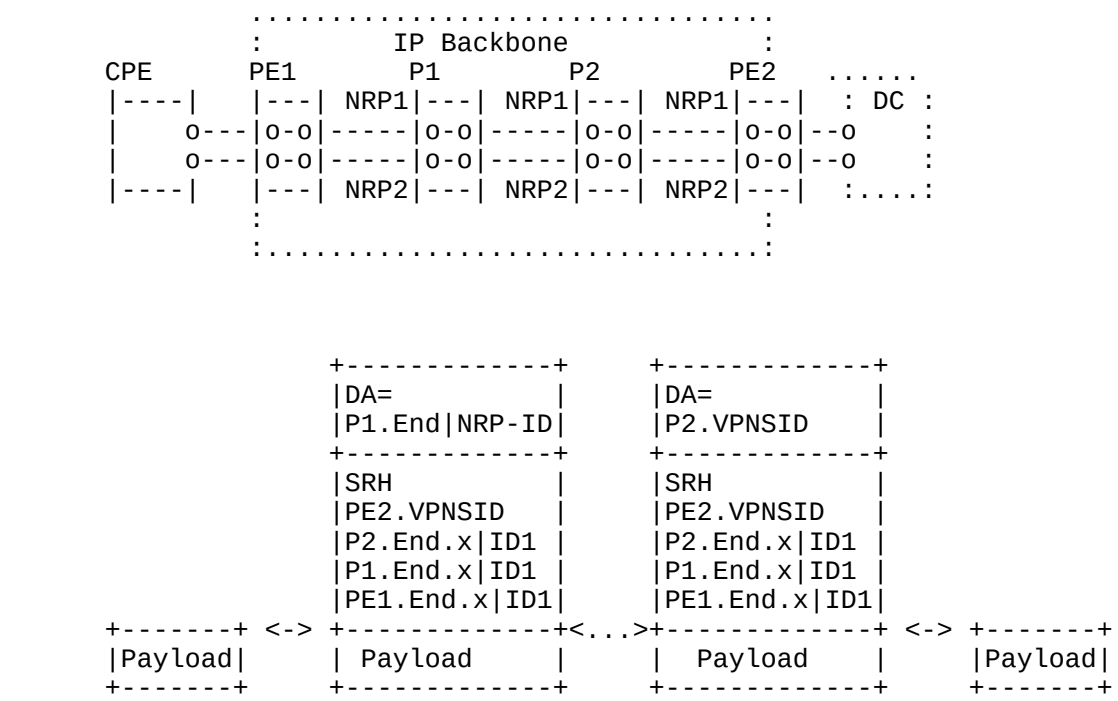

Figure 16: example topology

10. IANA Considerations

This document has no IANA actions.

11. Security Considerations

The security requirements and mechanisms described in [RFC8402] and [RFC8754] also apply to this document.

This document does not introduce any new security consideration.

12. References

12.1. Normative References

[I-D.cheng-teas-network-slice-usecase] Cheng, W., Jiang, W., Chen, R., Gong, L., and S. Peng, "IETF Network Slice use cases", draft-cheng-teas-network-slice-usecase-01 (work in progress), August 2021.

[I-D.ietf-teas-ietf-network-slices] Farrel, A., Gray, E., Drake, J., Rokui, R., Homma, S., Makhijani, K., Contreras, L., and J. Tantsura, "Framework for IETF Network Slices", draft-ietf-teas-ietf-network-slices-19 (work in progress), January 2023.

[I-D.ietf-teas-ns-ip-mpls] Saad, T., Beeram, Dong, J., Wen, B., Ceccarelli, D., Halpern, J., Peng, S., Chen, R., Liu, X., Contreras, L. M., Rokui, R., and L. Jalil, "Realizing Network Slices in IP/MPLSNetworks", Work in Progress, Internet-Draft, draft-ietf-teas-ns-ip-mpls-02, March 2023, <https://www.ietf.org/archive/id/draft-ietf-teas-ns-ip-mpls-02.txt>.

[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>.

[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>.

[RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <https://www.rfc-editor.org/info/rfc8754>.

[RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 0.17487/RFC8986, February 2021, <https://www.rfc-editor.org/info/rfc8986>.

Authors' Addresses

Yisong Liu
China Mobile
China
Email: liuyisong@chinamobile.com

Changwang Lin
New H3C Technologies
China
Email: linchangwang.04414@h3c.com

Hao Li
New H3C Technologies
China
Email: lihao@h3c.com

Liyan Gong
China Mobile
China
Email: gongliyan@chinamobile.com