

Network Working Group
Internet Draft
Intended status: Informational
Expires: June 22, 2026

L. Gong
W. Cheng
China Mobile
C. Lin
M. Chen
New H3C Technologies
J. Dong
Huawei Technologies
R. Chen
ZTE Corporation
Y. Liang
Ruijie Networks Co., Ltd.
December 22, 2025

Segment Routing based Solution for Hierarchical IETF Network Slices
draft-gong-spring-hierarchical-slice-solution-03

Abstract

This document describes a Segment Routing based solution for two-level hierarchical IETF network slices. Level-1 network slice is realized by associating Flex-Algo with dedicated sub-interfaces, and level-2 network slice is realized by using SR Policy with additional NRP-ID on data plane.

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 June 22, 2026.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction.....	2
1.1. Requirements Language.....	4
2. Solution based on Segment Routing.....	4
3. Example.....	7
4. Security Considerations.....	11
5. IANA Considerations.....	11
6. References.....	11
6.1. Normative References.....	11
6.2. Informative References.....	12
Authors' Addresses.....	14

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. A Network Resource Partition (NRP) is a collection of resources in the underlay network. Each NRP is used as the underlay network construct to support one or a group of IETF network slice services.

Hierarchical composition of IETF Network Slice means that a network slice can be further sliced into other network slices. Figure 1 shows the architecture of two-level hierarchical IETF network slices. Network resources are partitioned in a hierarchical manner. Network resources of the underlay network are partitioned into multiple level-1 network slices. Then network resources of a level-1 network slice are further partitioned into multiple level-2 network slices.




Figure 1: Architecture of Two-level Hierarchical IETF Network Slices

[I-D.dong-teas-hierarchical-ietf-network-slice] describes several possible scenarios of hierarchical IETF network slices. For example, level-1 can be industry slices which are used to deliver services for different vertical industries, and level-2 can be customer slices which are created to meet specific requirements of some or all of the customers within the corresponding industry of level-1.

For the two-level hierarchical IETF network slices discussed in this document, the level-1 and level-2 network slices are both created and managed by the same operator, and they are used to provide services at different granularity.

Segment Routing (SR) [RFC8402] is a source routing paradigm that explicitly indicates the forwarding path for packets at the ingress node. IETF network slices may be realized by using Segment Routing technologies.

This document proposes a Segment Routing based solution for two-level hierarchical IETF network slices. Level-1 network slice is realized by associating Flex-Algo with dedicated sub-interfaces, and

level-2 network slice is realized by using SR Policy with additional NRP-ID on data plane.

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. Solution based on Segment Routing

Flex-Algo is a mechanism that allows IGP to compute the best paths along the constrained topology in a distributed manner. [I-D.ietf-lsr-flex-algo] specifies the way of using Segment Routing (SR) Prefix-SIDs and SRv6 locators to steer packets for Flex-Algo.

As shown in Figure 2, each NRP for level-1 network slices is associated with a Flex-Algo. All the nodes belong to the level-1 NRP participate in the associated Flex-Algo. All the links belong to the level-1 NRP are included by the Admin Group rules of the associated Flex-Algo. Traffics of the level-1 network slices are steered into the Flex-Algo paths by using Prefix-SIDs or SRv6 locators, so that the corresponding level-1 NRPs will be used for forwarding.

Segment Routing Policy (SR Policy) is an ordered list of segments that represent a source-routed policy [I-D.ietf-spring-segment-routing-policy]. The packets steered into an SR Policy carry an ordered list of segments associated with that SR Policy.

In each NRP for level-2 network slices, the connectivity among PEs is achieved by SR Policies. The segment lists of these SR Policies composed with segments associated with the corresponding Flex-Algo of the level-1 NRP. So, the level-2 forwarding paths are restricted in the level-1 topology. Traffics of the level-2 network slice are steered into the SR Policies, so that the corresponding level-2 NRPs will be used for forwarding.

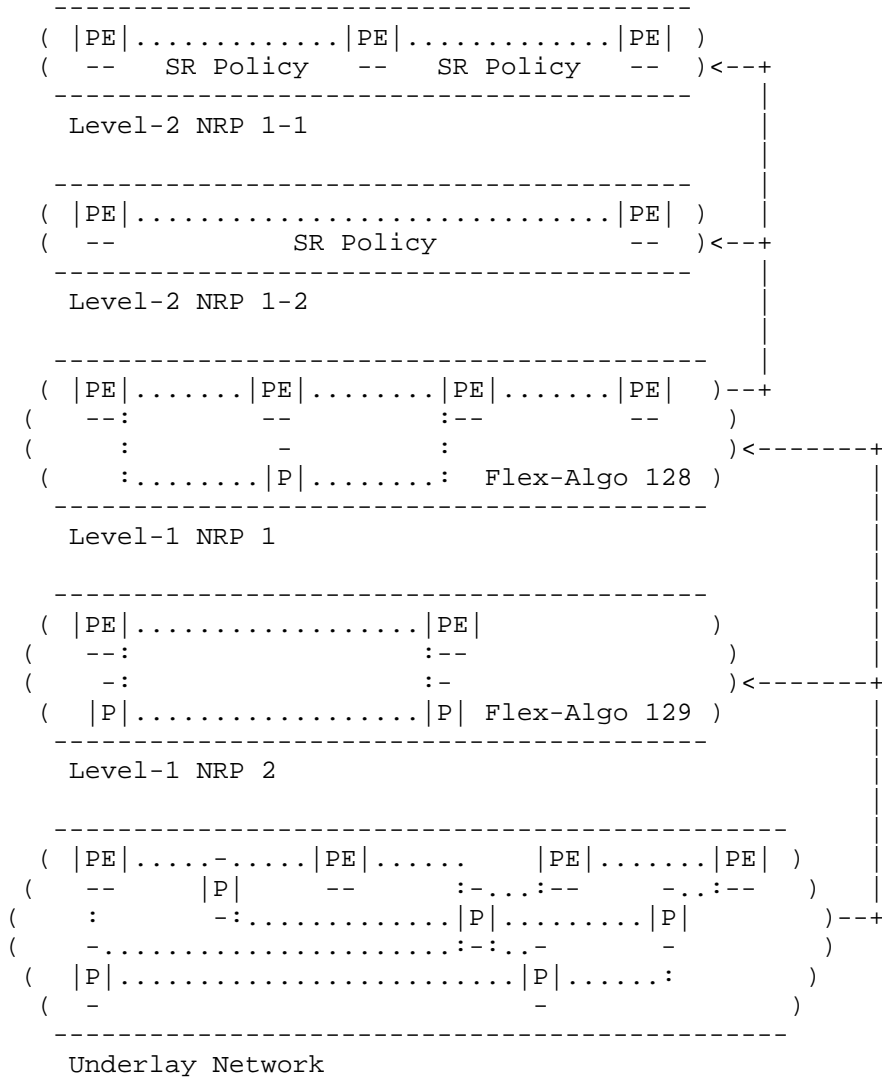


Figure 2: Framework of Solution

The network resources for the two-level network slices are also partitioned in a hierarchical manner.

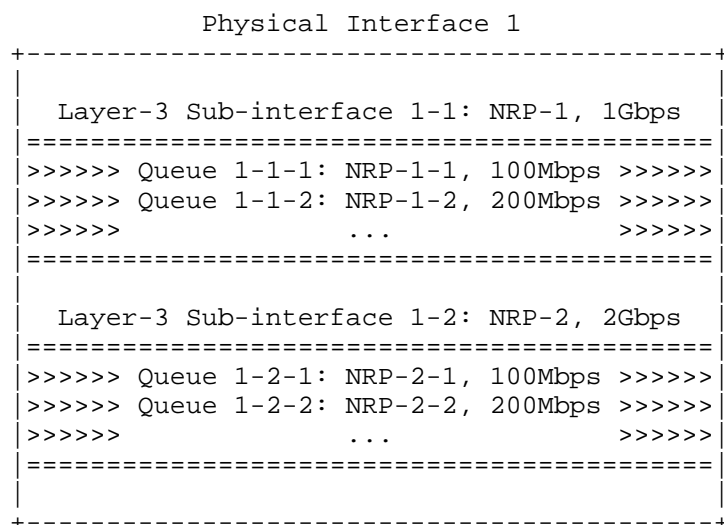


Figure 3: Hierarchical Network Resource Partition

As shown in Figure 3, the bandwidth resource of a physical interface is partitioned in two levels.

The level-1 NRPs are sliced by layer-3 sub-interfaces with dedicated bandwidth. The Admin Group of layer-3 sub-interface is included by the Flex-Algo which is associated with the level-1 NRP. Meanwhile, it is excluded or not included by irrelevant Flex-Algos. So, the topology of a level-1 network slice consists of a set of layer-3 sub-interfaces with dedicated bandwidth of the relevant level-1 NRP. When the traffics are forwarded according to Prefix-SIDs or SRv6 locators of the associated Flex-Algo, the corresponding bandwidth resources are used.

The level-2 NRPs are sliced by HQoS queues with dedicated bandwidth under the layer-3 sub-interface of level-1 NRP. Since the Flex-Algo associated Prefix-SIDs or SRv6 locators are used as the data plane identifier of level-1 NRP, level-2 NRP needs to be identified by using an extra dimension. On both MPLS-SR and SRv6 data plane, there are several options for realizing level-2 NRP-ID, such as [I-D.ietf-6man-enhanced-vpn-vtn-id], [I-D.cheng-spring-srv6-encoding-network-sliceid], [I-D.decraene-mpls-slid-encoded-entropy-label-id], and [I-D.li-mpls-enhanced-vpn-vtn-id]. As mentioned above, the traffics of level-2 network slice are forwarded according to the segment list of SR Policy. Firstly, the outgoing interface of the Flex-Algo associated segment will be the layer-3 sub-interface of level-1 NRP. Then, the HQoS queue will be selected according to the level-2 NRP-

ID carried in the packets, and the bandwidth resource of level-2 NRP will be used.

Each NRP can be used to support one or a group of network slice. If multiple level-1 network slices need to share the same level-1 NRP, those level-1 network slices should be associated to the same Flex-Algo, while a single level-1 NRP is still mapped to a single Flex-Algo. If multiple level-2 network slices need to share the same level-2 NRP, the SR Policies for those level-2 network slices should be associated to the same level-2 NRP-ID, and those level-2 network slices must belong to the same level-1 network slice, or different level-1 network slices which share the same level-1 NRP.

In the typical per-industrial-per-customer scenario of two-level hierarchical network slices, NRP sharing among different slices may be unnecessary. One-to-one mapping between network slice and NRP may be easier for deployment.

3. Example

The example network in Figure 4 is used for illustration.





Figure 4: Example Network

There are two level-1 network slices to be deployed, slice 1 for education and slice 2 for healthcare. The customers of education access from all PEs. The customers of healthcare access from PE1 and PE2.


Under slice 1, two universities require separate slices for interconnections among branch campuses. University 1 needs interconnection between PE1 and PE2 and interconnection between PE1 and PE3. University 2 needs interconnection between PE1 and PE2. Under slice 2, only one customer requires level-2 network slice.

Assume that the mapping between network slice and NRP is one to one. The topology of NRPs for the above network slices is shown in Figure 5.


Level-1 NRP 1 for Level-1 Network Slice 1:




Level-1 NRP 2 for Level-1 Network Slice 2:



Level-2 NRP 1-1 for Level-2 Network Slice 1-1:



Level-2 NRP 1-2 for Level-2 Network Slice 1-2:



Level-2 NRP 2-1 for Level-2 Network Slice 2-1:




Figure 5: Topology of NRPs

The provider assigns Flex-Algo 128 and 129 respectively for the two level-1 NRPs. All nodes participate in Flex-Algo 128. Only PE1, P1 and PE2 participate in Flex-Algo 129. Layer-3 sub-interfaces are set up for level-1 NRPs. HQoS queues under the layer-3 sub-interfaces are further set up for level-2 NRPs.

Taking PE1 as an example, the network resource partition of link bandwidth is shown in Figure 6.

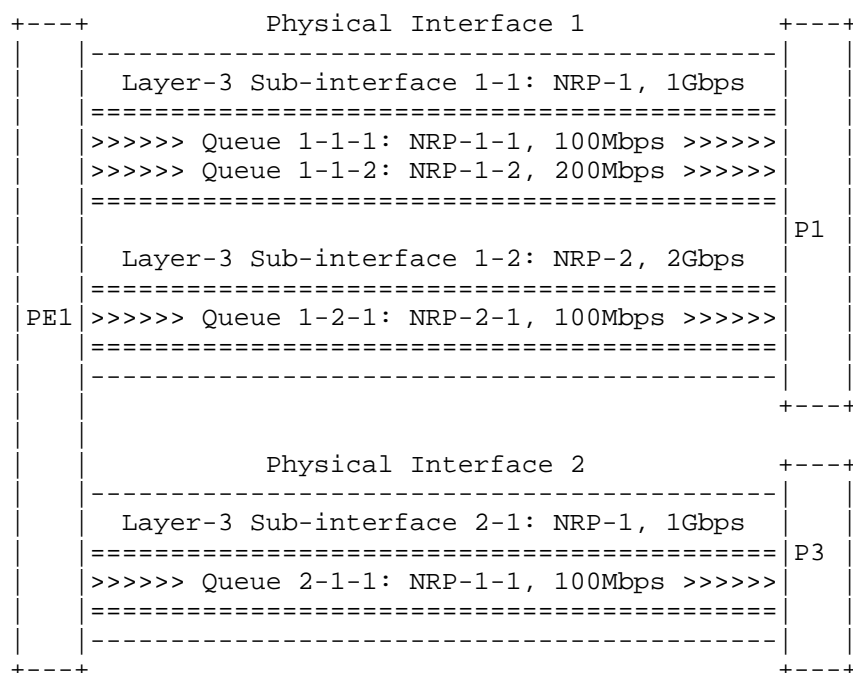


Figure 6: Network Resource Partition on PE1

Physical interface 1 on PE1 corresponds to link PE1-P1, and physical interface 2 corresponds to link PE1-P3.

Under interface 1, there are two layer-3 sub-interfaces 1-1 and 1-2. Sub-interfaces 1-1 is used as NRP-1 with dedicated bandwidth for level-1 network slice 1. Using Admin Group rules, sub-interfaces 1-1 is associated with Flex-Algo 128. Traffics of level-1 network slice 1 are steered into Flex-Algo 128. When the packets are forwarded from PE1 to P1, sub-interfaces 1-1 is selected as the outgoing interface and associated bandwidth resource will be used. Similarly, sub-interface 1-2 is used as NRP-2 for level-1 network slice 2, and associated with Flex-Algo 129.

Under layer-3 sub-interfaces 1-1, two HQoS queues 1-1-1 and 1-1-2 are further used as NRP-1-1 and NRP-1-2, with dedicated bandwidth for level-2 network slice 1-1 and 1-2. These queues are associated with the NRP-ID. When packets are forwarded through sub-interfaces 1-1, level-2 NRP-ID in the packets will be checked. If level-2 NRP-ID exists, the packet will be treated as level-2 network slice traffic, and will be forwarded using the associated queue with dedicated bandwidth for level-2 network slice. Similarly, HQoS queue 1-2-1 is used as NRP-2-1 for level-2 network slice 2-1.

Under interface 2, only layer-3 sub-interface 2-1 for level-1 network slice 1 is configured, along with HQoS queue 2-1-1 for level-2 network slice 1-1. NRPs for other network slices are not necessary, since the link PE1-P3 is not involved in their topologies.

If a packet from university 1 at PE1 needs to be forwarded to university 2 at PE2, the level-1 network slice 1 for education will be used, as shown in Figure 7. PE1 encapsulates the packet with an outer IPv6 header, and the Destination Address in the outer header is End SID for PE2 associated with Flex-Algo 128. Along the path PE1->P1->PE2, the packet is forwarded through layer-3 sub-interface associated with Flex-Algo 128, using dedicated bandwidth for the level-1 network slice 1.

If a packet from a branch campus of university 1 at PE1 needs to be forwarded to another branch campus of the same university at PE2, the level-2 network slice 1-1 for university 1 will be used, as shown in Figure 8. Assume that the level-2 NRP-ID is carried in HBH. PE1 encapsulates the packet with an outer IPv6 header, along with HBH and SRH. The SRH carries the segment-list of SR Policy to PE2, and the SIDs are all associated with Flex-Algo 128. The HBH carries the level-2 NRP-ID associated with level-2 network slice 1-1. Along the path PE1->P1->PE2, the packet is forwarded through the HQoS queue associated with the level-2 NRP-ID, under the layer-3 sub-interface associated with Flex-Algo 128. The dedicated bandwidth for level-2 network slice 1-1 will be used, other than sharing the bandwidth for level-1 network slice 1.




Figure 7: Packet Forwarding of Level-1 Network Slice 1




Figure 8: Packet Forwarding of Level-2 Network Slice 1-1

4. Security Considerations

TBD.

5. IANA Considerations

This document has no IANA actions.

6. References

6.1. 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>>.

- [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-teas-ietf-network-slices] Farrel, A., Drake, J., Rokui, R., Homma, S., Makhijani, K., Contreras, L. M., and J. Tantsura, "Framework for IETF Network Slices", Work in Progress, Internet-Draft, draft-ietf-teas-ietf-network-slices-12, 30 June 2022, <<https://www.ietf.org/archive/id/draft-ietf-teas-ietf-network-slices-12.txt>>.
- [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-20 (work in progress), May 2022.
- [I-D.ietf-spring-segment-routing-policy] Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", Work in Progress, Internet-Draft, draft-ietf-spring-segment-routing-policy-22, 22 March 2022, <<http://www.ietf.org/internet-drafts/draft-ietf-spring-segment-routing-policy-22.txt>>.

6.2. Informative References

- [I-D.dong-teas-hierarchical-ietf-network-slice] Dong, J., and Z. Li, "Considerations about Hierarchical IETF Network Slices", Work in Progress, Internet-Draft, draft-dong-teas-hierarchical-ietf-network-slice-01, 7 March 2022, <<http://www.ietf.org/internet-drafts/draft-dong-teas-hierarchical-ietf-network-slice-01.txt>>.
- [I-D.ietf-6man-enhanced-vpn-vtn-id] Dong, J., Li, Z., Xie, C., Ma, C., and G. Mishra, "Carrying Virtual Transport Network (VTN) Identifier in IPv6 Extension Header", Work in Progress, Internet-Draft, draft-ietf-6man-enhanced-vpn-vtn-id-00, 5 March 2022, <<http://www.ietf.org/internet-drafts/draft-ietf-6man-enhanced-vpn-vtn-id-00.txt>>.
- [I-D.cheng-spring-srv6-encoding-network-sliceid] Cheng, W., Lin, C., Gong, L., Zadok, S., and X. Wang, "Encoding Network Slice Identification for SRv6", Work in Progress, Internet-Draft, draft-cheng-spring-srv6-encoding-network-sliceid-04, 8 July 2022, <<http://www.ietf.org/internet-drafts/draft-cheng-spring-srv6-encoding-network-sliceid-04.txt>>.

[I-D.dekraene-mpls-slid-encoded-entropy-label-id] Decraene B., Filsfils, C., Henderickx W., Saad T., Beeram V., "Using Entropy Label for Network Slice Identification in MPLS networks", Work in Progress, Internet-Draft, draft-dekraene-mpls-slid-encoded-entropy-label-id-04, 14 June 2022, <<http://www.ietf.org/internet-drafts/draft-dekraene-mpls-slid-encoded-entropy-label-id-04.txt>>.

[I-D.li-mpls-enhanced-vpn-vtn-id] Li, Z. and J. Dong, "Carrying Virtual Transport Network Identifier in MPLS Packet", Work in Progress, Internet-Draft, draft-li-mpls-enhanced-vpn-vtn-id-02, 7 March 2022, <<http://www.ietf.org/internet-drafts/draft-li-mpls-enhanced-vpn-vtn-id-02.txt>>.

Authors' Addresses

Liyan Gong
China Mobile

Email: gongliyan@chinamobile.com

Weiqiang Cheng
China Mobile

Email: chengweiqiang@chinamobile.com

Changwang Lin
New H3C Technologies

Email: linchangwang.04414@h3c.com

Mengxiao Chen
New H3C Technologies

Email: chen.mengxiao@h3c.com

Jie Dong
Huawei Technologies

Email: jie.dong@huawei.com

Ran Chen
ZTE Corporation

Email: chen.ran@zte.com.cn

Yanrong Liang
Ruijie Networks Co., Ltd.

Email: liangyanrong@ruijie.com.cn

