

LSR Working Group
Internet Draft
Intended status: Informational
Expires: January 7, 2026

W. Cheng
China Moblie
C. Lin
New H3C Technologies

L. Gong
China Moblie
July 7, 2025

Advertise NRP Group extensions for IGP
draft-cheng-lsr-advertise-nrp-group-extensions-04

Abstract

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. 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. This document describes an IGP mechanism that is used to advertise a large number of NRPs into a smaller number of NRP groups.

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 January 7, 2026.

Copyright Notice

Copyright (c) 2023 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents

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

Table of Contents

1. Introduction.....	3
1.1. Requirements Language.....	3
2. Motivation.....	4
2.1. Performance of NRP group distribution.....	4
2.2. IGP Announcement efficiency.....	4
3. Extend for IS-IS.....	5
3.1. NRP Group TLV.....	5
3.2. NRP Link Group Info Sub-TLV in NBR TLV.....	6
3.3. NRP Group Offset for SRv6 END.X.....	6
3.4. NRP Group Offset for SRv6 END.X for SRv6 LAN END.X.....	8
4. Extend for OSPF & OSPFv3.....	9
4.1. NRP Group TLV.....	9
4.2. Link NRP Group Info Sub-TLV in NBR TLV.....	10
4.3. NRP Group Offset for SRv6 SIDs Associated with Adjacencies	11
5. Example.....	11
6. Security Considerations.....	12
7. IANA Considerations.....	12
7.1. ISIS NRP Group TLV.....	12
7.2. ISIS Link NRP Group Info Sub-TLV.....	13
7.3. OSPF and OSPFv3 NRP Group TLV.....	13
7.4. OSPF Link NRP Group Info Sub-TLV.....	14
7.5. OSPFv3 Link NRP Group Sub-TLVs.....	14
8. References.....	14
8.1. Normative References.....	14
Authors' Addresses.....	15

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.

To allow the network controller and network nodes to perform NRP-specific explicit path computation or shortest path computation, the group of resource-aware SIDs allocated by network nodes to each NRP and the associated topology and resource attributes need to be distributed using the control plane. [I-D.ietf-teas-nrp-scalability] analyzes the scalability requirements and the control plane and data plane scalability considerations of NRP. In order to support a relatively large number of NRPs in the network, one proposed approach is to separate the topology and resource attributes of the NRP in control plane, so that the advertisement and processing of each type of attribute could be decoupled. Multiple NRPs may share the same topology, and multiple NRPs may share the same set of network resources on some network segments, while the difference in either the topology or resource attributes makes them different NRPs. This allows flexible combination of network topology and network resource attributes to build a large number of NRPs with a relatively small number of logical topologies.

When the IGP protocol advertises a large amount of NRPs, there will be some problems. Please refer to chapter 2 for a detailed description.

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. Motivation

2.1. Performance of NRP group distribution

When controller distribute NRPs to each device on the middle path, which occupies resources. When there are a large number of NRPs, the process of resource occupation on the device takes too long.

To address this issue, the preemption method is used. The slice configuration is pre-configured on the devices in the middle path, and the devices can prepare resources in advance without actually occupying them. Only when the slice traffic actually passes through these devices will they actually occupy the corresponding bandwidth resources. This approach can effectively avoid the problem of excessive resource occupation time, and improve network performance and efficiency.

When the controller selects the required bandwidth resources based on user information, it then sends a forwarding policy to the headend. This policy is associated with the corresponding NRP and guides forwarding via Segment-Routing Policy at the headend. When intermediate devices perform forwarding, they will actually consume bandwidth resources based on the previously pre-configured NRP information.

Examples related to pre-configuration are described in chapter 5.

2.2. IGP Announcement efficiency

The IGP protocol needs to advertise the NRP joined on each interface. When interfaces join a large amount of NRPs resources, This will cause the IGP protocol to have overly large packets when advertising adjacency status.

To solve this problem, we can split multiple NRPs into a group and allocate them uniformly according to the group, called NRP Group. When announcing via the IGP protocol, we only need to announce the global NRP grouping status. When announcing the NRP joined on the interface, we only need to notify limited NRP group information.

So we need to extend the IGP protocol to advertise the composition of the global NRP groups and which NRP group is added on which interface. The specific extension can be referred to in Chapter 3 and Chapter 4.

Note that BGP-LS advertisement about global NRP Group TLV and NRP Link Group Info will be defined in a later document and is not covered in this document.

3. Extend for IS-IS

3.1. NRP Group TLV

ISIS NRP Group TLV has the following format:

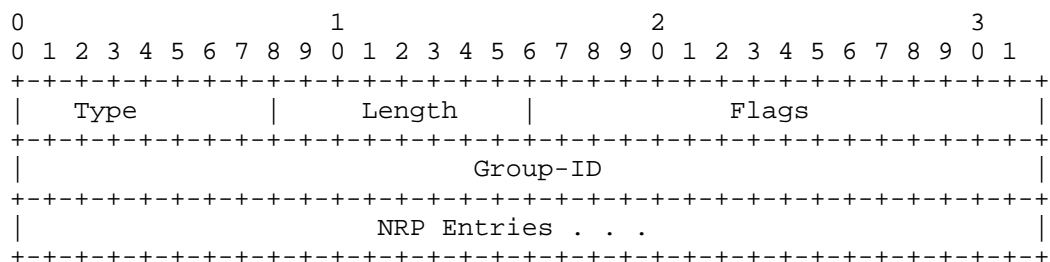


Figure 1: NRP Group TLV

where:

- Type: TBD, 1 octet.
- Length: Single octet, The length value is variable.
- Flags: 2 octets. No flags are currently defined
- Group-ID: NRP Group ID, 4octets.
- NRP Entries: The NRP Entry Sub-TLV

NRP Entry Sub TLV:

The NRP Entry Sub TLV has the following format:

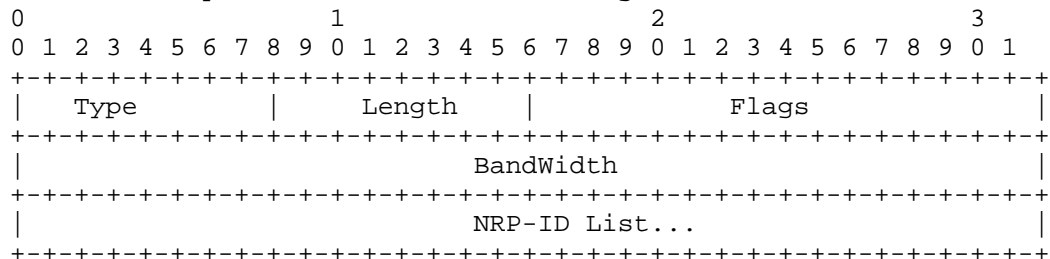


Figure 2: NRP Entry Sub TLV

Type: TBD, 1 octet.
 Length: Single octet, The length value is variable.
 Flags: 2 octets. No flags are currently defined
 BandWidth: 4 octets.
 NRP-ID List: Each NRP-ID is 4 octets, and the list of NRP-IDs is referred to as NRP-ID List.

3.2. NRP Link Group Info Sub-TLV in NBR TLV

This sub-TLV is used to advertise NRP Groups associated with a adjacency. Multiple NRP Group Info sub-TLVs MAY be associated with the same adjacency.

The Link NRP Group Info sub-TLV has the following format:

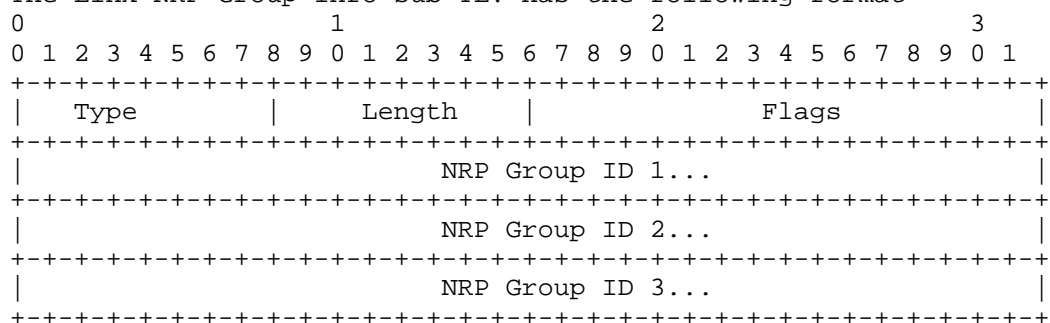


Figure 3: NRP Link Group Info Sub TLV

Type: TBD, 1 octet.
 Length: Single octet, The length value is variable.
 Flags: 2 octets. No flags are currently defined
 NRP Group ID List: Each NRP Group ID is 4 octets, and the list of NRP Group IDs is referred to as NRP-Group-ID List.

3.3. NRP Group Offset for SRv6 END.X

The SRv6 END.X for NRP Group Offset TLV has the following format:




Figure 4: NRP Group offset for SRv6 END.X TLV

- o Type: TBD.
- o Length: variable.
- o Flags: 1 octet, as defined in [RFC9352].
- o Algorithm: 1 octet, as defined in [RFC9352].
- o NRP Group-ID: NRP Group ID, 4octets.
- o Weight: 1 octet, as defined in [RFC9352].
- o Endpoint Behavior: as defined in [RFC9352].
- o SID: 16 octets. as defined in [RFC9352]..
- o Sub-sub-TLV-length: 1 octet, as defined in [RFC9352]..
- o Optional Sub-sub-TLVs: Supported sub-sub-TLVs are specified, as defined in [RFC9352].

3.4. NRP Group Offset for SRv6 END.X for SRv6 LAN END.X

The SRv6 LAN END.X for NRP Group Offset TLV has the following format:

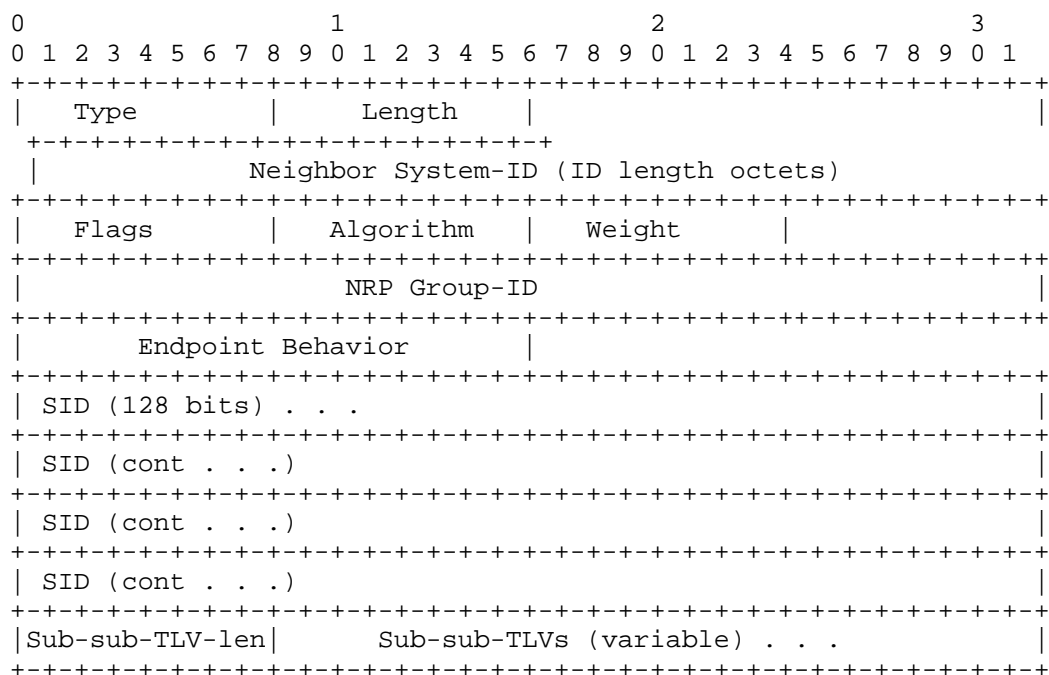


Figure 5: NRP Group offset for SRv6 LAN END.X TLV

- o Type: TBD.
- o Length: variable.
- o Neighbor System-ID: IS-IS System-ID of length "ID Length", as defined in [ISO10589].
- o Flags: 1 octet, as defined in [RFC9352].
- o Algorithm: 1 octet, as defined in [RFC9352].
- o Weight: 1 octet, as defined in [RFC9352].
- o NRP Group-ID: NRP Group ID, 4octets.

- o Endpoint Behavior: as defined in [RFC9352].
- o SID: 16 octets. as defined in [RFC9352]..
- o Sub-sub-TLV-length: 1 octet, as defined in [RFC9352]..
- o Optional Sub-sub-TLVs: Supported sub-sub-TLVs are specified, as defined in [RFC9352].

4. Extend for OSPF & OSPFv3

4.1. NRP Group TLV

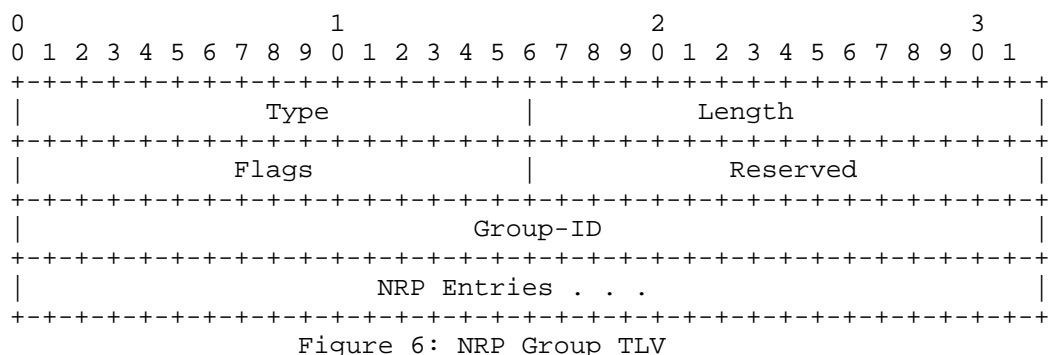


Figure 6: NRP Group TLV

where:

- Type: TBD, 2 octets.
- Length: 2 octets, The length value is variable.
- Flags: 2 octets. No flags are currently defined
- Reserved: 2-octet field. It MUST be set to 0 on transmission and MUST be ignored on receipt.
- Group-ID: NRP Group ID, 4octets.
- NRP Entries: The NRP Entry Sub-TLV

The NRP Entry Sub TLV has the following format:

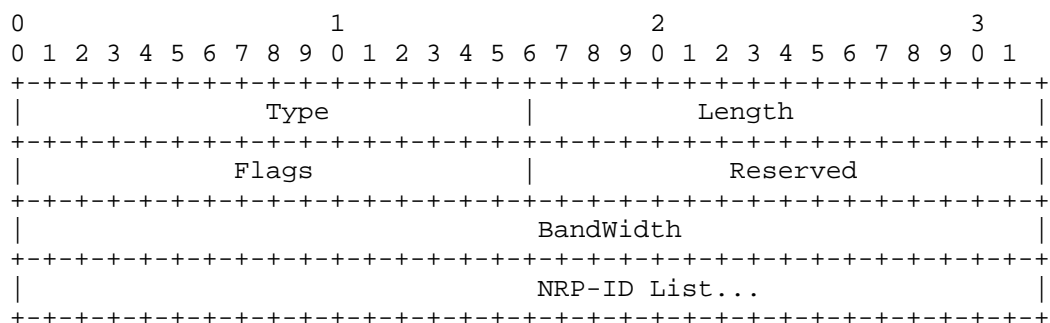


Figure 7: NRP Entry Sub TLV

Type: TBD, 2 octets.
 Length: 2 octets, The length value is variable.
 Flags: 2 octets. No flags are currently defined
 BandWidth: 4 octets.
 NRP-ID List: Each NRP-ID is 4 octets, and the list of NRP-IDs is referred to as NRP-ID List.

4.2. Link NRP Group Info Sub-TLV in NBR TLV

This sub-TLV is used to advertise NRP Groups associated with a adjacency. Multiple NRP Group Info sub-TLVs MAY be associated with the same adjacency.

The Link NRP Group Info sub-TLV has the following format:

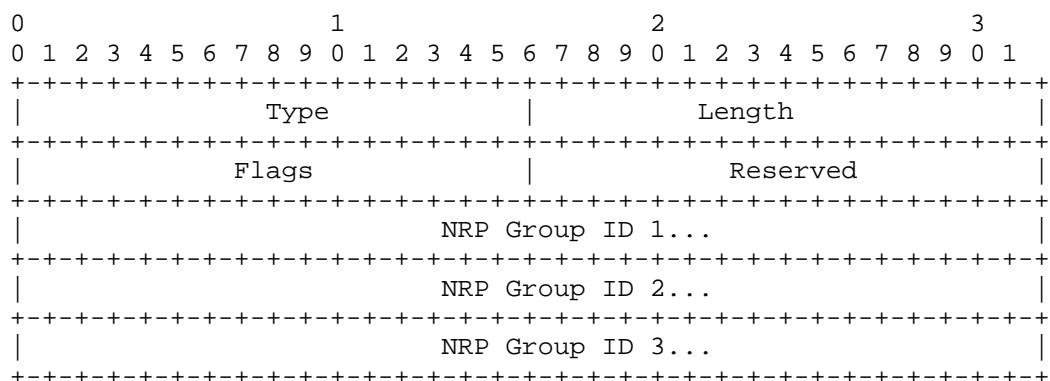


Figure 8: Link NRP Group Info Sub TLV

Type: TBD, 2 octets.

Length: 2 octets, The length value is variable.

Flags: 2 octets. No flags are currently defined

NRP Group ID List: Each NRP Group ID is 4 octets, and the list of NRP Group IDs is referred to as NRP-Group-ID List.

4.3. NRP Group Offset for SRv6 SIDs Associated with Adjacencies

TBD.

5. Example

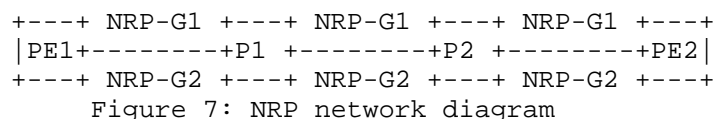


Figure 7: NRP network diagram

The NRP Group configuration on all devices is as follows:

NRP Group 1: (Total 10G)

nrp-1, bandwidth 1G

nrp-2, bandwidth 1G

...

Nrp-10, bandwidth 1G

NRP Group2: (Total 30G)

nrp-101, bandwidth 3G

nrp-102, bandwidth 3G

...

nrp-110, bandwidth 3G

According to pre-configuration based on network planning, initially, all interfaces on the paths of P1, P2, P3, and P4 will be added to both NRP Group1 and NRP Group2.

When users do not have a special service type, traffic forwarding is performed through the BE type.

When a user purchases a specific service type with a bandwidth 10G, the controller selects the corresponding NRP Group1 and sends an SR-TE path to the head node PE1 based on the service type, associating it with the appropriate NRP group ID. The devices on the intermediate paths occupy the corresponding network resources according to the respective NRP group during traffic forwarding, and their priority is higher than that of BE type.

If the user's service type changes to a bandwidth of 30G. the controller revokes the previously associated NRP Group2, issues a new SR-TE path, and associates it with the new NRP group ID.

6. Security Considerations

TBD.

7. IANA Considerations

7.1. ISIS NRP Group TLV

IANA is requested to make a new allocation in the "IS-IS TLV Codepoint Registry" under the registry name "IS-IS TLV Codepoints" as follows:

Type	Name	IIH	LSP	SNP	Purge	reference
TBD	NRP Group	N	Y	N	N	This document

Table 1: Requested ISIS Type Entries

7.2. ISIS Link NRP Group Info Sub-TLV

This document makes the following registrations in the "IS-IS Sub-TLVs for TLVs Advertising Neighbor Information" registry:

Type	Description	22	23	25	141	222	223	Reference
TBD	Link NRP Group Info	y	y	y	y	y	y	RFC 9352, Section 8.1

Table 2: IS-IS Sub-TLVs for TLVs Advertising Neighbor Information Registry

7.3. OSPF and OSPFv3 NRP Group TLV

IANA is requested to register TLV type from the "OSPF Router Information (RI) TLVs" registry as defined by [RFC7770].

Value	Description	Reference
TBD	NRP Group TLV	This document

7.4. OSPF Link NRP Group Info Sub-TLV

IANA is requested to register TLV type from the "OSPF Router Information (RI) TLVs" registry as defined by [RFC7770].

Value	Description	Reference
-----	-----	-----
TBD	Link NRP Group Sub-TLV	This document

7.5. OSPFv3 Link NRP Group Sub-TLVs

IANA is requested the following code points via the early allocation process in the "OSPFv3 Extended-LSA Sub-TLVs" registry under the "OSPFv3 Parameters" registry group for the new Sub-TLVs below that need to be made permanent:

* TBD: Link NRP Group Info Sub-TLV: Add NRP Group Info information to the Neighbor TLV.

8. References

8.1. Normative References

[I-D.ietf-teas-ietf-network-slices] Farrel, A., "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>>.

[RFC9352] P. Psenak, Ed., "IS-IS Extensions to Support Segment Routing over the IPv6 Data Plane", February 2023, <<https://datatracker.ietf.org/doc/rfc9352/>>.

[RFC7770], A. Lindem, Ed., Ed., "Extensions to OSPF for Advertising Optional Router Capabilities", February 2016, <<https://datatracker.ietf.org/doc/rfc7770/>>

Authors' Addresses

Weiqiang Cheng
China Mobile
China

Email: chengweiqiang@chinamobile.com

Changwang Lin
New H3C Technologies
China

Email: linchangwang.04414@h3c.com

Liyan Gong
China Mobile
China

Email: gongliyan@chinamobile.com

