Huawei HCIP Certification Training

HCIP-Datacom-WAN Planning and Deployment Lab Guide

ISSUE: 1.0

HUAWEI TECHNOLOGIES CO., LTD

Copyright © Huawei Technologies Co., Ltd. 2021. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

Trademarks and Permissions

 and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base Bantian, Longgang Shenzhen 518129 People's Republic of China

Website: https://e.huawei.com

Huawei Certification System

Huawei Certification is an integral part of the company's "Platform + Ecosystem" strategy, and it supports the ICT infrastructure featuring "Cloud-Pipe-Device". It evolves to reflect the latest trends of ICT development. Huawei Certification consists of two categories: ICT Infrastructure Certification, and Cloud Service & Platform Certification, making it the most extensive technical certification program in the industry.

Huawei offers three levels of certification: Huawei Certified ICT Associate (HCIA), Huawei Certified ICT Professional (HCIP), and Huawei Certified ICT Expert (HCIE).

Huawei Certification covers all ICT fields and adapts to the industry trend of ICT convergence. With its leading talent development system and certification standards, it is committed to fostering new ICT talent in the digital era, and building a sound ICT talent ecosystem.

HCIP-Datacom-WAN Planning and Deployment V1.0 aims to train and certificate senior engineers with professional knowledge and skills regarding bearer WAN scenarios in the data communication network field.

Passing HCIP-Datacom-WAN Planning and Deployment V1.0 certification will indicate that you:

1) Understand and master the enterprise bearer WAN solution, enterprise bearer WAN architecture and key technologies, WAN VPN technologies, MPLS TE fundamentals and configuration, SR, SRv6, Huawei CloudWAN solution architecture and fundamentals, Huawei CloudWAN solution O&M and troubleshooting, and Huawei CloudWAN solution design practice (financial scenario).

2) Are qualified for enterprise bearer WAN engineer positions (account manager, project manager, pre-sales engineer, post-sales engineer, and O&M engineer) and are capable of using Huawei datacom devices to design, deploy, and maintain enterprise Bearer WANs. You will be able to use Huawei datacom devices to plan, deploy, and maintain enterprise bearer WANs. Additionally, you will qualify for the corresponding senior engineer positions (such as the customer manager, project manager, pre-sales engineer, post-sales engineer, and O&M engineer).

The Huawei certification system introduces the industry, fosters innovation, and imparts cutting-edge datacom knowledge.


Huawei Certification
Cloud Service & Platform
ICT Infrastructure
HarmonyOS (To Be)
IoT
GaussDB
openGauss (To Be)
openEuler
Cloud(Cloud Computing + Cloud Service)
Intelligent Driving (To Be)
Big Data
AI
Kunpeng Application Developer
Storage
Datacom
Transmission
Intelligent Computing
WLAN
Access
Enterprise Communication
Security
LTE
Intelligent Vision
SDN
5G
Data Center
HCIE
Huawei Certified ICT Expert
HCIP
Huawei Certified ICT Professional
HCIA
Huawei Certified ICT Associate

About This Document

Overview

This textbook accompanies the training courses for HCIP-Datacom-WAN Network Planning and Deployment certification. It is applicable to candidates for the corresponding exam and people who want to understand the bearer WAN and related solutions, bearer WAN architecture and typical technology applications, WAN VPN technology, MPLS TE technology, SR technology, SRv6 technology, WAN controller-based network management and analysis, WAN controller-based network traffic control, bearer WAN O&M and troubleshooting, and bearer WAN design.

Description

This experiment guide introduces three experiments. The first two are traditional CLI-based experiments, whereas the last is performed using both the CLI and controller UI. The experiments are as follows:

1. SR-MPLS experiment
2. SRv6 experiment
3. iMaster NCE-IP experiment

Background Knowledge Required

This textbook is for Huawei's basic certification. To understand the content better, ensure that you meet the following requirements:

1. Have basic computer skills.
2. Have taken the HCIP-Datacom course.
3. Have passed the HCIP-Datacom exam.
4. Have a good understanding of TCP/IP protocol stack fundamentals.
5. Have a good understanding of Ethernet switch and router fundamentals.

Symbol Conventions

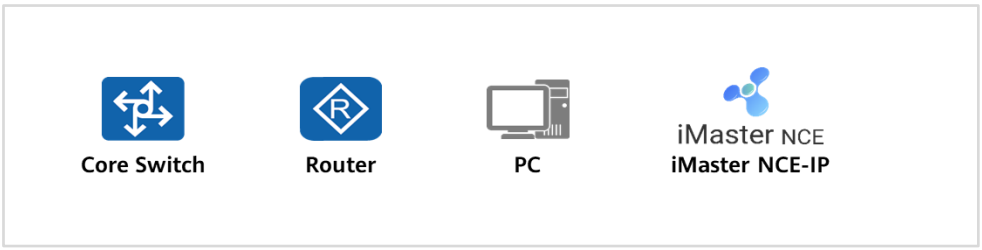

Lab Environment

Networking Introduction

This experiment environment is intended for datacom engineers who are preparing for the HCIP-Datacom-WAN exam. This lab environment includes six routers and several servers.

Device Introduction

The following table lists devices recommended for HCIP-Datacom-WAN experiments and the mappings between the device name, model, and software version.

Device Name	Model	Software Version
Router	NetEngine 8000 M6	V800R012C10 and later

Contents

1 SR-MPLS Experiment

1.1 L3VPNv4 over SR-MPLS BE Experiment

1.1.1 Introduction

1.1.1.1 Objectives

Upon completion of this task, you will be able to:

- Configure IS-IS to ensure that PEs are routable to each other.
- Configure SR-MPLS to establish SR LSPs.
- Recurse L3VPN tunnels used for communication between CEs to SR-MPLS BE tunnels.
- Observe the SR-MPLS tunnel status.
- Observe label changes in packets forwarded through an SR-MPLS BE tunnel.

1.1.1.2 Networking Description

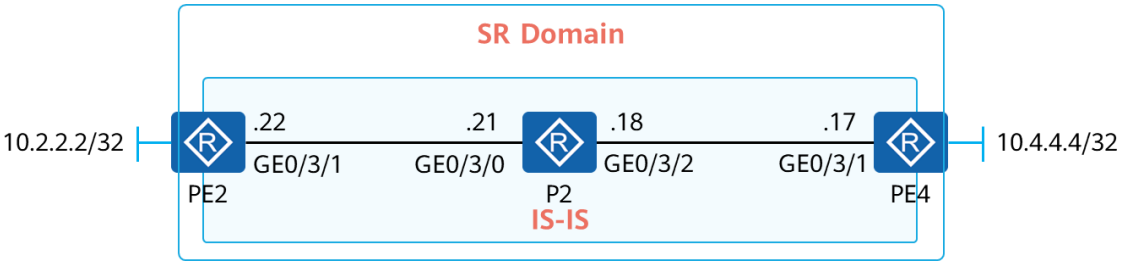

Figure 1-1 L3VPNv4 over SR-MPLS BE experiment topology

The figure shows the device connection and IP address planning. The interface interconnection addresses are in the format of 10.0.0.*Y*/30, and the values represented by *Y* are shown in the figure. Loopback0 is created on all devices. The Loopback0 address is used as the MPLS LSR ID of each device in the SR domain.

Loopback1 is created on PE2 and PE4 to simulate user access. The Loopback1 addresses on PE2 and PE4 are 10.2.2.2/32 and 10.4.4.4/32, respectively, as shown in the preceding figure.

1.1.2 Experiment Task

1.1.2.1 Configuration Roadmap

1. Configure IP addresses for devices.

2. Configure IS-IS in the SR domain. Specifically, enable IS-IS on interconnection and Loopback0 interfaces for communication in the SR domain.
3. Configure MPLS. Specifically, enable MPLS and set MPLS LSR IDs on devices.
4. Configure SR. Specifically, enable SR globally, enable IS-IS extensions for SR capabilities, and configure node SIDs.
5. Establish an MP-IBGP peer relationship between PE2 and P2 and between PE4 and P2. P2 functions as an RR to reflect VPNv4 routes from PE2 and PE4.
6. Create a VPN instance named **vpna**, add Loopback1 to the VPN instance on PE2 and PE4, and import direct routes to the BGP instance.

1.1.2.2 Configuration Procedure

Step 1 Configure IP addresses for interconnection and loopback interfaces.

Configure the configuration validation mode as immediate validation and configure IP addresses for interconnection and Loopback0 interfaces. Loopback0 addresses must be configured as planned in the following table.

Table 1-1 Loopback0 interface IP addresses

Device Number	Loopback0 IP Address
PE2	1.0.0.2
P2	1.0.0.6
PE4	1.0.0.4

Name the devices.

Omitted

Configure the configuration validation mode as immediate validation.

```
<PE2>system-view immediately
```

```
<P2>system-view immediately
```

```
<PE4>system-view immediately
```

Configure IP addresses for GE0/3/1 and Loopback0 on PE2.

```
[PE2] interface LoopBack 0
[PE2-LoopBack0] ip address 1.0.0.2 255.255.255.255
[PE2-LoopBack0] quit
```

```
[PE2]interface GigabitEthernet0/3/1
[PE2-GigabitEthernet0/3/1] ip address 10.0.0.22 255.255.255.252
```

Configure IP addresses for GE0/3/0, GE0/3/2, and Loopback0 on P2.

```
[P2] interface LoopBack 0
[P2-LoopBack0] ip address 1.0.0.6 255.255.255.255
[P2-LoopBack0] quit

[P2]interface GigabitEthernet0/3/0
[P2-GigabitEthernet0/3/0] ip address 10.0.0.21 255.255.255.252
[P2-GigabitEthernet0/3/0] quit
[P2]interface GigabitEthernet0/3/2
[P2-GigabitEthernet0/3/2] ip address 10.0.0.18 255.255.255.252
```

Configure IP addresses for GE0/3/1 and Loopback0 on PE4.

```
[PE4] interface LoopBack 0
[PE4-LoopBack0] ip address 1.0.0.4 255.255.255.255
[PE4-LoopBack0] quit

[PE4]interface GigabitEthernet0/3/1
[PE4-GigabitEthernet0/3/1] ip address 10.0.0.17 255.255.255.252
```

Test interconnection interface connectivity on P2.

```
[P2]ping -c 1 10.0.0.22
  PING 10.0.0.22: 56  data bytes, press CTRL_C to break
    Reply from 10.0.0.22: bytes=56 Sequence=1 ttl=255 time=1 ms

  --- 10.0.0.22 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 1/1/1 ms

[P2]ping -c 1 10.0.0.17
  PING 10.0.0.17: 56  data bytes, press CTRL_C to break
    Reply from 10.0.0.17: bytes=56 Sequence=1 ttl=255 time=1 ms

  --- 10.0.0.17 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 1/1/1 ms
```

Step 2 Configure IS-IS in the SR domain.

Ensure that the IS-IS area ID is 49.0001, the IS-IS process ID is 1, all devices are Level-2 devices, and the NET is converted from the Loopback0 IP address (for example, PE2's NET is 49.0001.0010.0000.0002.00). Then enable IS-IS on Loopback0 and interconnection interfaces.

In this case, you need to set **cost-style** to **wide** to support IS-IS extensions.

Configure PE2.

```
[PE2]isis 1
[PE2-isis-1] is-level level-2
[PE2-isis-1] cost-style wide
[PE2-isis-1] network-entity 49.0001.0010.0000.0002.00
[PE2-isis-1] is-name PE2
```

Configure P2.

```
[P2]isis 1
[P2-isis-1] is-level level-2
[P2-isis-1] cost-style wide
[P2-isis-1] network-entity 49.0001.0010.0000.0006.00
[P2-isis-1] is-name P2
```

Configure PE4.

```
[PE4]isis 1
[PE4-isis-1] is-level level-2
[PE4-isis-1] cost-style wide
[PE4-isis-1] network-entity 49.0001.0010.0000.0004.00
[PE4-isis-1] is-name PE4
```

Enable IS-IS on interfaces.

```
[PE2]interface LoopBack0
[PE2-LoopBack0] isis enable 1
[PE2-LoopBack0] quit
[PE2]interface GigabitEthernet0/3/1
[PE2-GigabitEthernet0/3/1] isis enable 1
[PE2-GigabitEthernet0/3/1] isis circuit-type p2p
```

```
[P2]interface LoopBack0
[P2-LoopBack0] isis enable 1
[P2-LoopBack0] quit
[P2]interface GigabitEthernet0/3/0
[P2-GigabitEthernet0/3/0] isis enable 1
[P2-GigabitEthernet0/3/0] isis circuit-type p2p
[P2-GigabitEthernet0/3/0] quit
[P2]interface GigabitEthernet0/3/2
[P2-GigabitEthernet0/3/2] isis enable 1
[P2-GigabitEthernet0/3/2] isis circuit-type p2p
```

```
[PE4]interface LoopBack0
[PE4-LoopBack0] isis enable 1
[PE4-LoopBack0] quit
[PE4]interface GigabitEthernet0/3/1
[PE4-GigabitEthernet0/3/1] isis enable 1
```

```
[PE4-GigabitEthernet0/3/1] isis circuit-type p2p
```

Check IS-IS neighbor relationships on P2.

```
[P2]display isis peer

                    Peerinformation for ISIS(1)

 System Id     Interface          Circuit Id         State  HoldTime Type     PRI
------------------------------------------------------------------------------
PE2*           GE0/3/0            0000000007         Up     26s      L2       --
PE4*           GE0/3/2            0000000007         Up     23s      L2       --
```

IS-IS neighbor relationships with PE2 and PE4 have been established.

Check IS-IS routes on P2.

```
[P2]display  isis route

                        Route  information for ISIS(1)
                        ----------------------------

                       ISIS(1)  Level-2 Forwarding Table
                       --------------------------------

IPV4Destination    IntCost     ExtCost    ExitInterface      NextHop          Flags
------------------------------------------------------------------------------
1.0.0.2/32         10          NULL       GE0/3/0            10.0.0.22        A/-/-/-
1.0.0.4/32         10          NULL       GE0/3/2            10.0.0.17        A/-/-/-
10.0.0.20/30       10          NULL       GE0/3/0            Direct           D/-/L/-
10.0.0.16/30       10          NULL       GE0/3/2            Direct           D/-/L/-
```

P2 has learned the IS-IS routes generated by Loopback0 on PE2 and PE4.

Step 3 Configure MPLS.

Enable MPLS on the three devices and configure MPLS LSR IDs. MPLS does not need to be enabled on interfaces.

```
[PE2]mpls lsr-id 1.0.0.2
[PE2]mpls
```

```
[P2]mpls lsr-id 1.0.0.6
[P2]mpls
```

```
[PE4]mpls lsr-id 1.0.0.4
[PE4]mpls
```

Step 4 Configure SR capabilities on devices.

Enable SR-MPLS globally, enable IS-IS extensions for SR capabilities, configure an SRGB for IS-IS, and set the SRGB range to 16000 to 17000 on all devices.

Configure a SID for Loopback0 and use an index as the relative label value. The relative label value must be consistent with the planned loopback address. For example, if the IP address of Loopback0 is 1.0.0.2, set the index to 2.

Enable SR-MPLS globally.

```
[PE2]segment-routing
```

```
[P2]segment-routing
```

```
[PE4]segment-routing
```

Enable IS-IS extensions for SR capabilities and configure an SRGB for IS-IS.

```
[PE2]isis 1
[PE2-isis-1]segment-routing mpls
[PE2-isis-1]segment-routing global-block 16000 17000
```

```
[P2]isis 1
[P2-isis-1]segment-routing mpls
[P2-isis-1]segment-routing global-block 16000 17000
```

```
[PE4]isis 1
[PE4-isis-1]segment-routing mpls
[PE4-isis-1]segment-routing global-block 16000 17000
```

Configure a node SID for devices.

```
[PE2]interface LoopBack 0
[PE2-LoopBack0]isis prefix-sid index 2
```

```
[P2]interface LoopBack 0
[P2-LoopBack0]isis prefix-sid index 6
```

```
[PE4]interface LoopBack 0
[PE4-LoopBack0]isis prefix-sid index 4
```

Run the **display tunnel-info all** command on PE2 to check SR LSP establishment.

```
[PE2]display tunnel-info all
Tunnel ID                Type             Destination          Status
------------------------------------------------------------------------------
0x000000002900000002 srbe-lsp         1.0.0.6              UP
0x000000002900000005 srbe-lsp         1.0.0.4              UP
```

SR LSPs to P2 and PE4 have been established.

Check the SR label forwarding table.

```
[PE2]display  segment-routing prefix  mpls  forwarding

                  Segment Routing Prefix MPLS Forwarding Information
              ------------------------------------------------------------
              Role: I-Ingress, T-Transit, E-Egress, I&T-Ingress And Transit

Prefix          Label   OutLabel   Interface  NextHop      Role  MPLSMtu  Mtu     State
------------------------------------------------------------------------------------------------------
1.0.0.2/32      16002   NULL       Loop0      127.0.0.1    E     ---      1500    Active
1.0.0.4/32      16004   16004      GE0/3/1    10.0.0.21    I&T   ---      1500    Active
1.0.0.6/32      16006   3          GE0/3/1    10.0.0.21    I&T   ---      1500    Active

Total information(s): 3
```

The out label of the route from PE2 to P2 (1.0.0.6) is 3, and the out label of the route from PE2 to PE4 (1.0.0.4) is 16004.

Check the connectivity of the CR-LSP from PE2 to PE4.

```
[PE2]ping lsp segment-routing ip 1.0.0.4 32 version draft2
  LSP PING FEC: SEGMENT ROUTING IPV4 PREFIX 1.0.0.4/32 : 100   data bytes, press CTRL_C to
break
    Reply from 1.0.0.4: bytes=100 Sequence=1 time=12 ms
    Reply from 1.0.0.4: bytes=100 Sequence=2 time=3 ms
    Reply from 1.0.0.4: bytes=100 Sequence=3 time=7 ms
    Reply from 1.0.0.4: bytes=100 Sequence=4 time=2 ms
    Reply from 1.0.0.4: bytes=100 Sequence=5 time=2 ms

  --- FEC: SEGMENT ROUTING IPV4 PREFIX 1.0.0.4/32 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 2/5/12 ms
```

The connectivity is normal.

Step 5 Configure an L3VPN.

Create a VPN instance named **vpna** on PE2 and PE4, add Loopback1 to the VPN instance, and establish an MP-BGP VPNv4 peer relationship between PE2 and P2 and between PE4 and P2 (the AS number is 65001). P2 functions as the RR, and PE2 and PE4 function as the RR clients and advertise VPNv4 routes through P2.

Create a VPN instance named **vpna**.

```
[PE2]ip vpn-instance vpna
[PE2-vpn-instance-vpna] ipv4-family
[PE2-vpn-instance-vpna-af-ipv4] route-distinguisher 100:20
[PE2-vpn-instance-vpna-af-ipv4] vpn-target 100:1020 both
```

```
[PE4]ip vpn-instance vpna
[PE4-vpn-instance-vpna] ipv4-family
[PE4-vpn-instance-vpna-af-ipv4] route-distinguisher 100:40
[PE4-vpn-instance-vpna-af-ipv4] vpn-target 100:1020 both
```

Create Loopback1, associate it with the VPN instance, and configure an IP address for the interface.

```
[PE2]interface LoopBack 1
[PE2-LoopBack1]ip binding vpn-instance vpna
Info: All IPv4 and IPv6 related configurations on this interface are removed.
[PE2-LoopBack1]ip address 10.2.2.2 32
```

```
[PE4]interface LoopBack 1
[PE4-LoopBack1]ip binding vpn-instance vpna
Info: All IPv4 and IPv6 related configurations on this interface are removed.
[PE4-LoopBack1]ip address 10.4.4.4 32
```

Note that you need to associate the interface with the VPN instance before configuring an IP address for the interface.

Use Loopback0 to configure the MP-BGP VPNv4 peer relationship and use the Loopback0 address as the router ID.

```
[PE2]bgp 65001
[PE2-bgp] router-id 1.0.0.2
[PE2-bgp] peer 1.0.0.6 as-number 65001
[PE2-bgp] peer 1.0.0.6 connect-interface LoopBack0
[PE2-bgp]ipv4-family vpnv4
[PE2-bgp-af-vpnv4] peer 1.0.0.6 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
```

```
[PE4]bgp 65001
[PE4-bgp] router-id 1.0.0.4
[PE4-bgp] peer 1.0.0.6 as-number 65001
[PE4-bgp] peer 1.0.0.6 connect-interface LoopBack0
[PE4-bgp]ipv4-family vpnv4
[PE4-bgp-af-vpnv4] peer 1.0.0.6 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
```

```
[P2-bgp]bgp 65001
[P2-bgp] router-id 1.0.0.6
[P2-bgp] peer 1.0.0.2 as-number 65001
[P2-bgp] peer 1.0.0.2 connect-interface LoopBack0
[P2-bgp] peer 1.0.0.4 as-number 65001
[P2-bgp] peer 1.0.0.4 connect-interface LoopBack0
[P2-bgp] ipv4-family vpnv4
[P2-bgp-af-vpnv4] undo policy vpn-target
[P2-bgp-af-vpnv4] peer 1.0.0.2 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
[P2-bgp-af-vpnv4] peer 1.0.0.2 reflect-client
[P2-bgp-af-vpnv4] peer 1.0.0.4 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
[P2-bgp-af-vpnv4] peer 1.0.0.4 reflect-client
```

When configuring an RR, disable the RT check on VPNv4 routes.

Check the VPNv4 peer relationship status.

```
[P2]display bgp vpnv4 all peer

BGP local router ID : 1.0.0.6
Local AS number : 65001
Total number of peers: 2                    Peers in established state : 2

 Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State   PrefRcv
 1.0.0.2         4       65001       58      106     0  00:01:00 Established         1
 1.0.0.4         4       65001       52      154     0  00:03:03 Established         0
```

P2 has established MP-BGP VBNv4 peer relationships with PE2 and PE4.

Import the direct routes of Loopback1 to BGP.

```
[PE2]bgp 65001
[PE2-bgp]ipv4-family vpn-instance vpna
[PE2-bgp-vpna] import-route direct
```

```
[PE4]bgp 65001
[PE4-bgp]ipv4-family vpn-instance vpna
[PE4-bgp-vpna] import-route direct
```

Check VPNv4 routes on PE2.

```
[PE2]display bgp vpnv4 all routing-table | include 10.4.4.4
BGP Local router ID is 1.0.0.2
Status codes: *-valid, >-best, d-damped, x-best external, a-add path,
              h-history, i-internal, s-suppressed, S-Stale
              Origin: i-IGP, e-EGP, ?-incomplete
RPKI validation codes: V-valid, I-invalid, N-not-found
Total number of routes from all PE: 3
Route Distinguisher: 100:20
```

```
         Network          NextHop          MED        LocPrf    PrefVal Path/Ogn

Route Distinguisher:  100:40

         Network          NextHop          MED        LocPrf    PrefVal Path/Ogn
*>i      10.4.4.4/32      1.0.0.4           0          100          0        ?

VPN-Instance vpna,Router ID 1.0.0.2:

Total Number of Routes: 3
         Network          NextHop          MED        LocPrf    PrefVal Path/Ogn
*>i      10.4.4.4/32      1.0.0.4           0          100          0        ?
```

PE2 has learned the VPNv4 route from PE4 through MP-BGP.

Check the IP routing table on PE2.

```
[PE2]display ip routing-table vpn-instance vpna
Route Flags: R -relay,D -downloadtofib,T -tovpn-instance, B -blackholeroute
------------------------------------------------------------------------------
RoutingTable: vpna
         Destinations : 4        Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
        10.2.2.2/32 Direct  0    0           D   127.0.0.1       LoopBack1
        10.4.4.4/32 IBGP    255  0           RD  1.0.0.4         GigabitEthernet0/3/1
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

The route to the network segment of the remote CE has been loaded to the VPN instance routing table on PE2.

Check route details.

```
[PE2]display ip routing-table vpn-instance vpna 10.4.4.4 verbose
Route Flags: R -relay,D -downloadtofib,T -tovpn-instance, B -blackholeroute
------------------------------------------------------------------------------
RoutingTable: vpna
Summary Count : 1

Destination:      10.4.4.4/32
    Protocol:     IBGP                Process ID   : 0
  Preference:     255                 Cost         : 0
     NextHop:     1.0.0.4             Neighbour    : 1.0.0.6
       State:     Active Adv Relied   Age          : 00h03m54s
         Tag:     0                   Priority     : low
       Label:     48155               QoSInfo      : 0x0
  IndirectID:     0x1000349           Instance     :
RelayNextHop:     10.0.0.21           Interface    : GigabitEthernet0/3/1
    TunnelID:     0x000000002900000005 Flags       : RD
```

The tunnel ID can be found. We can determine based on previous information that the tunnel is an SR-MPLS BE tunnel.

Check the connectivity between Loopback1 interfaces on PE2 and PE4.

```
[PE2]ping -vpn-instance vpna -a 10.2.2.2 10.4.4.4
  PING 10.4.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.4.4.4: bytes=56 Sequence=1 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=2 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=3 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=4 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=5 ttl=254 time=1 ms

  --- 10.4.4.4 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 1/1/1 ms
```

The connectivity is normal.

Step 6 Observe the forwarding process.

Capture the headers of incoming packets on GE0/3/0 of P2 and check the labels encapsulated into these packets during communication between 10.2.2.2 and 10.4.4.4.

#On PE4, check the label allocated by PE4 to route 10.4.4.4.

```
[PE2]display bgp vpnv4 all routing-table label

BGP Local router ID is 1.0.0.2
Status codes: *-valid, > -best, d -damped, x-best external, a -add path,
              h -history, i -internal, s-suppressed, S -Stale
              Origin: i -IGP, e -EGP, ?-incomplete
RPKI validation codes: V -valid, I -invalid, N -not-found


Total number of routes from all PE: 1
Route Distinguisher: 100:40


      Network      NextHop                  In/OutLabel
*>i   10.4.4.4     1.0.0.4                  NULL/48155

VPN-Instance vpna, Router ID 1.0.0.2:

Total Number of Routes: 1
      Network      NextHop                  In/OutLabel
*>i   10.4.4.4     1.0.0.4                  NULL/48155
```

On PE4, MP-BGP assigns label 48155 to the VPNv4 route 10.4.4.4. This label is the inner label (VPN label) of packets destined for 10.4.4.4.

In SR-MPLS BE mode, the outer label is identical with the node SID. The node SID of PE4 is 16004, which is the outer label (public network label) of VPN packets from PE2 to 10.4.4.4.

Create ACL 10000 on P2 to match the traffic from 10.2.2.2 to 10.4.4.4.

```
[P2]acl number 10000
[P2-acl-mpls-10000] rule permit label 16004 48155
```

Use ACL rules to match packets with the outer label being 16004 and inner label being 48155.

Run the **capture-packet** command on P2 to capture packet headers on GE0/3/0.

```
[P2]capture-packet forwarding interface GigabitEthernet 0/3/0 inbound acl 10000 packet-num 5
packet-len 64 overwrite file SRMPLSBE.cap
Info: Capture-packet data will be saved to cfcard:/logfile/SRMPLSBE.cap.
```

On PE2, ping 10.4.4.4 from 10.2.2.2.

```
[PE2]ping -vpn-instance vpna -a 10.2.2.2 10.4.4.4
PING 10.4.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.4.4.4: bytes=56 Sequence=1 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=2 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=3 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=4 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=5 ttl=254 time=1 ms

  --- 10.4.4.4 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 1/1/1 ms
```

Information about the captured packet headers is saved in the **/logfile** directory of the device. You can download the file through FTP or SFTP. For details about how to enable FTP or SFTP on the device, see the related product documentation (for example: https://support.huawei.com/hedex/hdx.do?docid=EDOC1100168795&lang=en).

Check captured packet headers.

```
Frame 1: 106 bytes on wire (848 bits), 64 bytes captured (512 bits)
Ethernet II, Src: HuaweiTe_7a:c2:8a (dc:99:14:7a:c2:8a), Dst: HuaweiTe_7a:c3:f1 (dc:99:14:7a:c3:f1)
MultiProtocol Label Switching Header, Label: 16004, Exp: 0, S: 0, TTL: 255
MultiProtocol Label Switching Header, Label: 48155, Exp: 0, S: 1, TTL: 255
Internet Protocol Version 4, Src: 10.2.2.2, Dst: 10.4.4.4
Internet Control Message Protocol
```

The outer label (public network label) of these packets is 16004, which is identical with the node SID of PE4. The inner label is the label allocated by MP-BGP to VPNv4 routes on PE4.

1.1.3 Quiz

In an L3VPNv4 over SR-MPLS BE scenario, will the outer label change during packet forwarding?

1.2 EVPN L3VPNv4 over SR-MPLS TE Experiment

1.2.1 Introduction

1.2.1.1 Objectives

Upon completion of this task, you will be able to:

- Configure IS-IS to ensure that PEs are routable to each other.
- Manually configure SR-MPLS TE tunnels.
- Recurse L3VPN tunnels used for communication between CEs to SR-MPLS TE tunnels.
- Observe label changes in packets forwarded through an SR-MPLS TE tunnel.

1.2.1.2 Networking Description

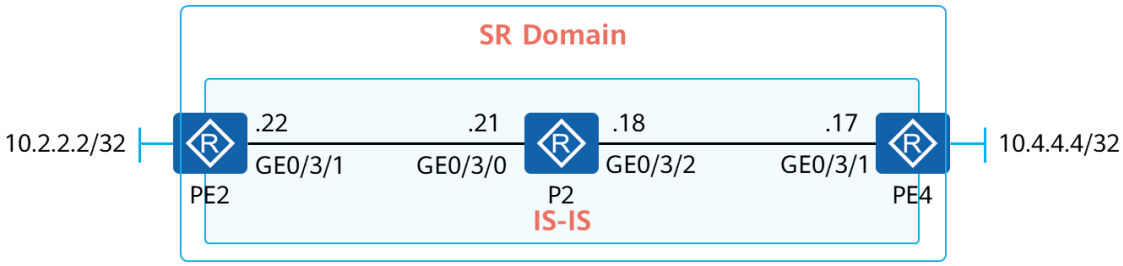

Figure 1-2 EVPN L3VPNv4 over SR-MPLS TE experiment

The figure shows the device connection and IP address planning. The interface interconnection addresses are in the format of 10.0.0.*Y*/30, and the values represented by *Y* are shown in the figure. Loopback0 is created on all devices. The Loopback0 address is used as the MPLS LSR ID of each device in the SR domain.

Loopback1 is created on PE2 and PE4 to simulate user access. The Loopback1 addresses on PE2 and PE4 are 10.2.2.2/32 and 10.4.4.4/32, respectively, as shown in the preceding figure.

1.2.2 Experiment Task

1.2.2.1 Configuration Roadmap

1. Configure IP addresses for devices.
2. Configure IS-IS in the SR domain. Specifically, enable IS-IS on interconnection and Loopback0 interfaces for communication in the SR domain.
3. Configure MPLS. Specifically, enable MPLS and MPLS TE and set MPLS LSR IDs on devices.
4. Configure SR. Specifically, enable SR globally, enable IS-IS extensions for SR capabilities, and configure node SIDs.
5. Configure explicit paths and TE tunnel interfaces on PE2 and PE4.
6. Configure an EVPN instance, add Loopback1 on PE2 and PE4 to the instance, and establish EVPN peer relationships between PE2 and P2 and between PE4 and P2.
7. Configure a tunnel selection policy to recurse EVPN traffic to TE tunnels.

1.2.2.2 Configuration Procedure

Step 1 Configure IP addresses for interconnection and loopback interfaces.

Configure the configuration validation mode as immediate validation and configure IP addresses for interconnection and Loopback0 interfaces. Loopback0 addresses must be configured as planned in the following table.

Table 1-2 Loopback0 IP addresses

Device Number	Loopback0 IP Address
PE2	1.0.0.2
P2	1.0.0.6
PE4	1.0.0.4

Name the devices.

Omitted

Configure the configuration validation mode as immediate validation.

```
<PE2>system-view immediately

<P2>system-view immediately

<PE4>system-view immediately
```

Configure IP addresses for GE0/3/1 and Loopback0 on PE2.

```
[PE2] interface LoopBack 0
[PE2-LoopBack0] ip address 1.0.0.2 255.255.255.255
[PE2-LoopBack0] quit

[PE2]interface GigabitEthernet0/3/1
[PE2-GigabitEthernet0/3/1] ip address 10.0.0.22 255.255.255.252
```

Configure IP addresses for GE0/3/0, GE0/3/2, and Loopback0 on P2.

```
[P2] interface LoopBack 0
[P2-LoopBack0] ip address 1.0.0.6 255.255.255.255
[P2-LoopBack0] quit

[P2]interface GigabitEthernet0/3/0
[P2-GigabitEthernet0/3/0] ip address 10.0.0.21 255.255.255.252
[P2-GigabitEthernet0/3/0] quit
[P2]interface GigabitEthernet0/3/2
[P2-GigabitEthernet0/3/2] ip address 10.0.0.18 255.255.255.252
```

Configure IP addresses for GE0/3/1 and Loopback0 on PE4.

```
[PE4] interface LoopBack 0
[PE4-LoopBack0] ip address 1.0.0.4 255.255.255.255
[PE4-LoopBack0] quit

[PE4]interface GigabitEthernet0/3/1
[PE4-GigabitEthernet0/3/1] ip address 10.0.0.17 255.255.255.252
```

Test interconnection interface connectivity on P2.

```
[P2]ping -c 1 10.0.0.22
  PING 10.0.0.22: 56  data bytes, press CTRL_C to break
    Reply from 10.0.0.22: bytes=56 Sequence=1 ttl=255 time=1 ms

  --- 10.0.0.22 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 1/1/1 ms
```

```
[P2]ping -c 1 10.0.0.17
  PING 10.0.0.17: 56  data bytes, press CTRL_C to break
    Reply from 10.0.0.17: bytes=56 Sequence=1 ttl=255 time=1 ms

  --- 10.0.0.17 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 1/1/1 ms
```

Step 2 Configure IS-IS in the SR domain.

Ensure that the IS-IS area ID is 49.0001, the IS-IS process ID is 1, all devices are Level-2 devices, and the NET is converted from the Loopback0 IP address (for example, PE2's NET is 49.0001.0010.0000.0002.00). Then enable IS-IS on Loopback0 and interconnection interfaces.

In this case, you need to set **cost-style** to **wide** to support IS-IS extensions.

Configure PE2.

```
[PE2]isis 1
[PE2-isis-1] is-level level-2
[PE2-isis-1] cost-style wide
[PE2-isis-1] network-entity 49.0001.0010.0000.0002.00
[PE2-isis-1] is-name PE2
```

Configure P2.

```
[P2]isis 1
[P2-isis-1] is-level level-2
[P2-isis-1] cost-style wide
[P2-isis-1] network-entity 49.0001.0010.0000.0006.00
```

```
[P2-isis-1] is-name P2
```

Configure PE4.

```
[PE4]isis 1
[PE4-isis-1] is-level level-2
[PE4-isis-1] cost-style wide
[PE4-isis-1] network-entity 49.0001.0010.0000.0004.00
[PE4-isis-1] is-name PE4
```

Enable IS-IS on interfaces.

```
[PE2]interface LoopBack0
[PE2-LoopBack0] isis enable 1
[PE2-LoopBack0] quit
[PE2]interface GigabitEthernet0/3/1
[PE2-GigabitEthernet0/3/1] isis enable 1
[PE2-GigabitEthernet0/3/1] isis circuit-type p2p
```

```
[P2]interface LoopBack0
[P2-LoopBack0] isis enable 1
[P2-LoopBack0] quit
[P2]interface GigabitEthernet0/3/0
[P2-GigabitEthernet0/3/0] isis enable 1
[P2-GigabitEthernet0/3/0] isis circuit-type p2p
[P2-GigabitEthernet0/3/0] quit
[P2]interface GigabitEthernet0/3/2
[P2-GigabitEthernet0/3/2] isis enable 1
[P2-GigabitEthernet0/3/2] isis circuit-type p2p

[PE4]interface LoopBack0
[PE4-LoopBack0] isis enable 1
[PE4-LoopBack0] quit
[PE4]interface GigabitEthernet0/3/1
[PE4-GigabitEthernet0/3/1] isis enable 1
[PE4-GigabitEthernet0/3/1] isis circuit-type p2p
```

Check IS-IS neighbor relationships on P2.

```
[P2]display isis peer

                         Peerinformation for ISIS(1)

 System Id     Interface          Circuit Id         State  HoldTime Type       PRI
------------------------------------------------------------------------------
PE2*           GE0/3/0            0000000007         Up     26s      L2         --
PE4*           GE0/3/2            0000000007         Up     23s      L2         --
```

IS-IS neighbor relationships with PE2 and PE4 have been established.

Check IS-IS routes on P2.

```
[P2]display  isis route

                         Route information for ISIS(1)
                         ----------------------------

                        ISIS(1)  Level-2 Forwarding Table
                        ---------------------------------

IPV4Destination    IntCost     ExtCost    ExitInterface      NextHop          Flags
-------------------------------------------------------------------------------
1.0.0.2/32         10          NULL       GE0/3/0            10.0.0.22        A/-/-/-
1.0.0.4/32         10          NULL       GE0/3/2            10.0.0.17        A/-/-/-
10.0.0.20/30       10          NULL       GE0/3/0            Direct           D/-/L/-
10.0.0.16/30       10          NULL       GE0/3/2            Direct           D/-/L/-
```

P2 has learned the IS-IS routes generated by Loopback0 on PE2 and PE4.

Step 3 Configure MPLS.

Enable MPLS on the three devices and configure MPLS LSR IDs. MPLS does not need to be enabled on interfaces.

```
[PE2]Mpls lsr-id 1.0.0.2
[PE2]Mpls

[P2]Mpls lsr-id 1.0.0.6
[P2]Mpls

[PE4]Mpls lsr-id 1.0.0.4
[PE4]Mpls
```

Step 4 Configure SR capabilities on devices.

Enable SR-MPLS globally, enable IS-IS extensions for SR capabilities, configure an SRGB for IS-IS, and set the SRGB range to 16000 to 17000 on all devices.

Configure a SID for Loopback0 and use an index as the relative label value. The relative label value must be the same as the planned loopback address. For example, if the IP address of Loopback0 is 1.0.0.2, set the index to 2.

Enable SR-MPLS globally.

```
[PE2]segment-routing

[P2]segment-routing

[PE4]segment-routing
```

Enable IS-IS extensions for SR capabilities and configure an SRGB for IS-IS.

```
[PE2]isis 1
[PE2-isis-1]segment-routing mpls
[PE2-isis-1]segment-routing global-block 16000 17000

[P2]isis 1
```

```
[P2-isis-1]segment-routing mpls
[P2-isis-1]segment-routing global-block 16000 17000

[PE4]isis 1
[PE4-isis-1]segment-routing mpls
[PE4-isis-1]segment-routing global-block 16000 17000
```

Configure a node SID for devices.

```
[PE2]interface LoopBack 0
[PE2-LoopBack0]isis prefix-sid index 2

[P2]interface LoopBack 0
[P2-LoopBack0]isis prefix-sid index 6

[PE4]interface LoopBack 0
[PE4-LoopBack0]isis prefix-sid index 4
```

Manually configure adjacency SIDs on P2.

```
[P2]segment-routing
[P2-segment-routing] ipv4 adjacency local-ip-addr 10.0.0.21 remote-ip-addr 10.0.0.22 sid 321536
[P2-segment-routing] ipv4 adjacency local-ip-addr 10.0.0.18 remote-ip-addr 10.0.0.17 sid 321537
```

To ensure that the adjacency SIDs specified during explicit path configuration remain unchanged, you are advised to configure static adjacency SIDs. Then, the SIDs remain unchanged after the device restarts.

Step 5 Configure SR-MPLS TE explicit paths and TE tunnel interfaces.

Configure explicit paths on PE2 and PE4, specify the nodes that the paths must pass through by specifying node SIDs, create TE tunnel interfaces on PE2 and PE4, and associate the interfaces with configured explicit paths.

This experiment is implemented through the CLI and does not involve the controller. In normal scenarios where the controller is used, the paths are computed by the controller.

Create explicit paths.

```
[PE2]explicit-path PE2_PE4_Manual
[PE2-explicit-path-PE2_PE4_Manual] next sid label 16006 type prefix
[PE2-explicit-path-PE2_PE4_Manual] next sid label 321537 type adjacency
```

Configure an explicit path named **PE2_PE4_Manual** on PE2 and forcibly specify the path to pass through P2 and GE0/3/2 on P2.

```
[PE4]explicit-path PE4_PE2_Manual
[PE4-explicit-path-PE4_PE2_Manual] next sid label 16006 type prefix
[PE4-explicit-path-PE4_PE2_Manual] next sid label 321536 type adjacency
```

Configure an explicit path named **PE4_PE2_Manual** on PE4 and forcibly specify the path to pass through GE0/3/0 on P2.

Create TE tunnel interfaces.

```
[PE2]interface Tunnel10
[PE2-Tunnel10] ip address unnumbered interface LoopBack0
[PE2-Tunnel10] tunnel-protocol mpls te
[PE2-Tunnel10] destination 1.0.0.4
[PE2-Tunnel10] mpls te signal-protocol segment-routing
[PE2-Tunnel10] mpls te tunnel-id 10
[PE2-Tunnel10] mpls te path explicit-path PE2_PE4_Manual
```

Create tunnel interface 10 on PE2, configure PE2 to borrow the Loopback0 IP address, set the destination address to 1.0.0.4 (Loopback0 address of PE4), and associate tunnel interface 10 with the explicit path **PE2_PE4_Manual**.

```
[PE4]interface Tunnel10
[PE4-Tunnel10] ip address unnumbered interface LoopBack0
[PE4-Tunnel10] tunnel-protocol mpls te
[PE4-Tunnel10] destination 1.0.0.2
[PE4-Tunnel10] mpls te signal-protocol segment-routing
[PE4-Tunnel10] mpls te tunnel-id 10
[PE4-Tunnel10] mpls te path explicit-path PE4_PE2_Manual
```

Create tunnel interface 10 on PE4, configure PE4 to borrow the Loopback0 IP address, set the destination address to 1.0.0.2 (Loopback0 IP address of PE2), and associate tunnel interface 10 with the explicit path **PE4_PE2_Manual**.

Check the SR-MPLS TE tunnel status.

```
[PE2]display  tunnel-info  all
Tunnel ID                Type                Destination         Status
--------------------------------------------------------------------------------
0x000000000300000001 sr-te               1.0.0.4             UP
0x000000002900000002 srbe-lsp            1.0.0.6             UP
0x000000002900000005 srbe-lsp            1.0.0.4             UP
```

```
[PE4]display  tunnel-info  all
Tunnel ID                Type                Destination         Status
--------------------------------------------------------------------------------
0x000000000300000001 sr-te               1.0.0.2             UP
0x000000002900000008 srbe-lsp            1.0.0.6             UP
0x00000000290000000e srbe-lsp            1.0.0.2             UP
```

The tunnel status on PE2 and PE4 is normal.

Check the tunnel connectivity on PE2 and PE4.

```
[PE2]ping lsp segment-routing te Tunnel 10
  LSP PING FEC: SEGMENT ROUTING TE TUNNEL IPV4 SESSION QUERY Tunnel10 : 100  data bytes,
press CTRL_C to break
    Reply from 1.0.0.4: bytes=100 Sequence=1 time=11 ms
    Reply from 1.0.0.4: bytes=100 Sequence=2 time=3 ms
    Reply from 1.0.0.4: bytes=100 Sequence=3 time=3 ms
    Reply from 1.0.0.4: bytes=100 Sequence=4 time=3 ms
    Reply from 1.0.0.4: bytes=100 Sequence=5 time=3 ms
```

```
  --- FEC: SEGMENT ROUTING TE TUNNEL IPV4 SESSION QUERY Tunnel10 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 3/4/11 ms
```

```
[PE4]ping lsp segment-routing te Tunnel 10
  LSP PING FEC: SEGMENT ROUTING TE TUNNEL IPV4 SESSION QUERY Tunnel10 : 100   data bytes,
press CTRL_C to break
    Reply from 1.0.0.2: bytes=100 Sequence=1 time=12 ms
    Reply from 1.0.0.2: bytes=100 Sequence=2 time=2 ms
    Reply from 1.0.0.2: bytes=100 Sequence=3 time=3 ms
    Reply from 1.0.0.2: bytes=100 Sequence=4 time=2 ms
    Reply from 1.0.0.2: bytes=100 Sequence=5 time=4 ms

  --- FEC: SEGMENT ROUTING TE TUNNEL IPV4 SESSION QUERY Tunnel10 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 2/4/12 ms
```

The TE tunnel connectivity is normal.

Step 6 Configure an EVPN.

Create a VPN instance named **vpna** on PE2 and PE4, add Loopback1 to the VPN instance, and establish an MP-BGP EVPN peer relationship between PE2 and P2 and between PE4 and P2 (the AS number is 65001). P2 functions as the RR, and PE2 and PE4 function as the RR clients and advertise EVPN routes through P2.

Create a VPN instance named **vpna**.

```
[PE2]ip vpn-instance vpna
[PE2-vpn-instance-vpna] ipv4-family
[PE2-vpn-instance-vpna-af-ipv4] route-distinguisher 100:20
[PE2-vpn-instance-vpna-af-ipv4] vpn-target 100:1020 evpn
[PE2-vpn-instance-vpna-af-ipv4] evpn mpls routing-enable
```

```
[PE4]ip vpn-instance vpna
[PE4-vpn-instance-vpna] ipv4-family
[PE4-vpn-instance-vpna-af-ipv4] route-distinguisher 100:40
[PE4-vpn-instance-vpna-af-ipv4] vpn-target 100:1020 evpn
[PE4-vpn-instance-vpna-af-ipv4] evpn mpls routing-enable
```

You only need to configure EVPN RTs. Meanwhile, enable EVPN to generate and advertise IP prefix routes and IRB routes.

Create Loopback1, associate it with the VPN instance, and configure an IP address for the interface.

```
[PE2]interface LoopBack 1
[PE2-LoopBack1]ip binding vpn-instance vpna
Info: All IPv4 and IPv6 related configurations on this interface are removed.
[PE2-LoopBack1]ip address 10.2.2.2 32
```

```
[PE4]interface LoopBack 1
[PE4-LoopBack1]ip binding vpn-instance vpna
Info: All IPv4 and IPv6 related configurations on this interface are removed.
[PE4-LoopBack1]ip address 10.4.4.4 32
```

Note that you need to associate the interface with the VPN instance before configuring an IP address for the interface.

Configure EVPN peer relationships. Use Loopback0 to set up peer relationships and Loopback0 addresses as router IDs.

```
[PE2]bgp 65001
[PE2-bgp] router-id 1.0.0.2
[PE2-bgp] peer 1.0.0.6 as-number 65001
[PE2-bgp] peer 1.0.0.6 connect-interface LoopBack0
[PE2-bgp] l2vpn-family evpn
[PE2-bgp-af-vpnv4] peer 1.0.0.6 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
```

```
[PE4]bgp 65001
[PE4-bgp] router-id 1.0.0.4
[PE4-bgp] peer 1.0.0.6 as-number 65001
[PE4-bgp] peer 1.0.0.6 connect-interface LoopBack0
[PE4-bgp]l2vpn-family evpn
[PE4-bgp-af-vpnv4] peer 1.0.0.6 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
```

```
[P2-bgp]bgp 65001
[P2-bgp] router-id 1.0.0.6
[P2-bgp] peer 1.0.0.2 as-number 65001
[P2-bgp] peer 1.0.0.2 connect-interface LoopBack0
[P2-bgp] peer 1.0.0.4 as-number 65001
[P2-bgp] peer 1.0.0.4 connect-interface LoopBack0
[P2-bgp]l2vpn-family evpn
[P2-bgp-af-vpnv4] undo policy vpn-target
[P2-bgp-af-vpnv4] peer 1.0.0.2 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
[P2-bgp-af-vpnv4] peer 1.0.0.2 reflect-client
[P2-bgp-af-vpnv4] peer 1.0.0.4 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
[P2-bgp-af-vpnv4] peer 1.0.0.4 reflect-client
```

When configuring an RR, disable the RT check on VPNv4 routes.

Check EVPN peer relationships.

```
[P2]display bgp evpn peer

BGP local router ID : 1.0.0.6
Local AS number : 65001
Total number of peers: 2                    Peers in established state : 2

 Peer            V    AS  MsgRcvd  MsgSent  OutQ  Up/Down       State    PrefRcv
 1.0.0.2         4 65001       121      276     0  00:53:06 Established          0
 1.0.0.4         4 65001        51      144     0  00:00:14 Established          0
```

P2 has established MP-BGP EVPN peer relationships with PE2 and PE4.

Import the direct routes of Loopback1 to BGP.

```
[PE2]bgp 65001
[PE2-bgp]ipv4-family vpn-instance vpna
[PE2-bgp-vpna] import-route direct
[PE2-bgp-vpna] advertise l2vpn evpn

[PE4]bgp 65001
[PE4-bgp]ipv4-family vpn-instance vpna
[PE4-bgp-vpna] import-route direct
[PE4-bgp-vpna] advertise l2vpn evpn
```

Note that you need to enable the VPN instance to advertise IP routes to the EVPN instance.

Check EVPN routes on PE2 and PE4.

```
[PE2]display bgp evpn all routing-table

Local AS number : 65001
BGP Local router ID is 1.0.0.2
Status codes: *-valid, >-best, d-damped, x-best external, a-add path,
              h-history, i-internal, s-suppressed, S-Stale
              Origin: i-IGP, e-EGP, ?-incomplete

EVPN address family:
Number of Ip Prefix Routes: 2
Route Distinguisher: 100:20
     Network(EthTagId/IpPrefix/IpPrefixLen)                NextHop
*>   0:10.2.2.2:32                                         0.0.0.0

Route Distinguisher: 100:40
     Network(EthTagId/IpPrefix/IpPrefixLen)                NextHop
*>i  0:10.4.4.4:32                                         1.0.0.4
```

```
[PE4]display bgp evpn all routing-table
Local AS number : 65001

BGP Local router ID is 1.0.0.4
```

```
Status codes:  *-valid,> -best, d -damped,x-bestexternal,a -add path,
                h -history,  i -internal, s-suppressed, S -Stale
                Origin: i -IGP, e -EGP, ?-incomplete


EVPN addressfamily:
Number ofIp Prefix Routes: 2
Route Distinguisher:  100:20
      Network(EthTagId/IpPrefix/IpPrefixLen)                NextHop
*>i   0:10.2.2.2:32                                         1.0.0.2

Route Distinguisher:  100:40
      Network(EthTagId/IpPrefix/IpPrefixLen)                NextHop
*>    0:10.4.4.4:32                                         0.0.0.0
```

The Loopback1 EVPN route from the peer end has been learned.

Check the VPN instance routing table on PE2 and PE4.

```
[PE2]display ip routing-table vpn-instance vpna
Route Flags:  R -relay,D -downloadtofib,T -tovpn-instance, B -blackholeroute
------------------------------------------------------------------------------
RoutingTable: vpna
         Destinations : 4        Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
       10.2.2.2/32  Direct  0    0           D   127.0.0.1       LoopBack1
       10.4.4.4/32  IBGP    255  0           RD  1.0.0.4         GigabitEthernet0/3/1
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

```
[PE4]display ip routing-table vpn-instance vpna
Route Flags:  R -relay,D -downloadtofib,T -tovpn-instance, B -blackholeroute
------------------------------------------------------------------------------
RoutingTable: vpna
         Destinations : 4        Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
       10.2.2.2/32  IBGP    255  0           RD  1.0.0.2         GigabitEthernet0/3/1
       10.4.4.4/32  Direct  0    0           D   127.0.0.1       LoopBack1
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

The route advertised by the peer end has been learned.

Step 7 Configure tunnel selection policies.

Configure tunnel selection policies to preferentially select SR-MPLS TE tunnels and associate these tunnel policies with the VPN instance.

Configure PE2.

```
[PE2]tunnel-policy p1
```

```
Info: New tunnel-policy is configured.
[PE2-tunnel-policy-p1]tunnel select-seq sr-te load-balance-number 1
[PE2-tunnel-policy-p1]quit
```

Configure PE4.

```
[PE4] tunnel-policy p1
Info: New tunnel-policy is configured.
[PE4-tunnel-policy-p1]tunnel select-seq sr-te load-balance-number 1
[PE4-tunnel-policy-p1]quit
```

Associate tunnel policies with the VPN instance.

```
[PE2]ip vpn-instance vpna
[PE2-vpn-instance-vpna]ipv4-family
[PE2-vpn-instance-vpna-af-ipv4]tnl-policy p1 evpn
```

```
[PE4]ip vpn-instance vpna
[PE4-vpn-instance-vpna]ipv4-family
[PE4-vpn-instance-vpna-af-ipv4]tnl-policy p1 evpn
```

Note that the **evpn** parameter needs to be added during the association, so that the routes learned through EVPN can recurse to SR-MPLS TE tunnels.

Check the VPN instance routing table on PE2 and PE4.

```
[PE2]display ip  routing-table  vpn-instance vpna
Route Flags:  R -relay,D -downloadtofib,T -tovpn-instance, B -blackholeroute
------------------------------------------------------------------------------
RoutingTable: vpna
         Destinations : 4        Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        10.2.2.2/32  Direct  0    0           D   127.0.0.1       LoopBack1
        10.4.4.4/32  IBGP    255  0           RD  1.0.0.4         Tunnel10
       127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

```
[PE4]display ip  routing-table vpn-instance vpna
Route Flags:  R -relay,D -downloadtofib,T -tovpn-instance, B -blackholeroute
------------------------------------------------------------------------------
RoutingTable: vpna
         Destinations : 4        Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        10.2.2.2/32  IBGP    255  0           RD  1.0.0.2         Tunnel10
        10.4.4.4/32  Direct  0    0           D   127.0.0.1       LoopBack1
       127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
```

```
255.255.255.255/32  Direct  0      0                D    127.0.0.1       InLoopBack0
```

The routes from PE2 and PE4 to the network segment of the remote CE have recursed to SR-MPLS TE tunnels.

Step 8 Verify connectivity and check labels.

Verify connectivity between the Loopback1 interfaces used by PE2 and PE4 to simulate CEs, and capture packet headers to check SR-MPLS labels.

On PE2, check the connectivity between Loopback1 on PE2 and Loopback1 on PE4.

```
[PE2]ping -vpn-instance vpna -a 10.2.2.2 10.4.4.4
  PING 10.4.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.4.4.4: bytes=56 Sequence=1 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=2 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=3 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=4 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=5 ttl=254 time=1 ms

  --- 10.4.4.4 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 1/1/1 ms
```

The connectivity is normal.

#On PE2, check the label of the EVPN route 10.4.4.4.

```
[PE2]display bgp evpn all routing-table prefix-route 0:10.4.4.4:32
 BGP local router ID : 1.0.0.2
 Local AS number : 65001
 Total routes of Route Distinguisher(100:40): 1
 BGP routing table entry information of 0:10.4.4.4:32:
 Label information (Received/Applied): 48156/NULL
 From: 1.0.0.6 (1.0.0.6)
 Route Duration: 0d00h13m59s
 Relay IP Nexthop: 10.0.0.21
 Relay IP Out-Interface: GigabitEthernet0/3/1
 Relay Tunnel Out-Interface: GigabitEthernet0/3/1
 Original nexthop: 1.0.0.4
 Qos information : 0x0
 Ext-Community: RT <100 : 1020>
 AS-path Nil, origin incomplete, MED 0, localpref 100, pref-val 0, valid, internal, best, select, pre 255,
IGP cost 20
 Originator: 1.0.0.4
 Cluster list: 1.0.0.6
 Route Type: 5 (Ip Prefix Route)
 Ethernet Tag ID: 0, IP Prefix/Len: 10.4.4.4/32, ESI: 0000.0000.0000.0000.0000, GW IP Address: 0.0.0.0
 Not advertised to any peer yet
```

PE4 allocates label 48156 to route 10.4.4.4. This label is the inner label (VPN label) carried in packets destined for 10.4.4.4.

The outer labels are used to strictly specify an explicit path. The labels of the outermost and second outermost layers are 16006 and 321537, respectively. In other words, when PE2 sends a packet to Loopback1 on PE4, the labels encapsulated into the packet are 48156, 321537, and 16006 from the innermost layer to the outermost layer.

Create ACL 10000 on P2 to match packets from 10.2.2.2 to 10.4.4.4.

```
[P2]acl number 10000
[P2-acl-mpls-10000] rule permit label 321537 48156
```

Note that the label sequence configured in the ACL is from the outermost layer to the innermost layer. In this case, when the packet reaches P2, the outermost label 16006 is already removed.

Run the **capture-packet** command on P2 to capture packet headers on GE0/3/0.

```
[P2]capture-packet forwarding interface GigabitEthernet 0/3/0 inbound acl 10000 packet-num 5
packet-len 64 overwrite file SRMPLSTE.cap
Info: Capture-packet data will be saved to cfcard:/logfile/SRMPLSTE.cap.
```

On PE2, ping 10.4.4.4 from 10.2.2.2.

```
[PE2]ping -vpn-instance vpna -a 10.2.2.2 10.4.4.4
  PING 10.4.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.4.4.4: bytes=56 Sequence=1 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=2 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=3 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=4 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=5 ttl=254 time=1 ms

  --- 10.4.4.4 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 1/1/1 ms
```

Information about the captured packet headers is saved in the **/logfile** directory of the device. You can download the file through FTP or SFTP.

Check captured packet headers.

```
Ethernet II, Src: HuaweiTe_7a:c2:8a (dc:99:14:7a:c2:8a), Dst: HuaweiTe_7a:c3:f1 (dc:99:14:7a:c3:f1)
MultiProtocol Label Switching Header, Label: 321537, Exp: 0, S: 0, TTL: 255
MultiProtocol Label Switching Header, Label: 48156, Exp: 0, S: 1, TTL: 255
Internet Protocol Version 4, Src: 10.2.2.2, Dst: 10.4.4.4
Internet Control Message Protocol
```

The labels of these packets are as expected.

1.2.3 Quiz

In an SR-MPLS TE scenario, how can we forcibly forward packets through a specific interface on a specific device?

1.3 L3VPNv4 over Static SR-MPLS Policy Experiment

1.3.1 Introduction

1.3.1.1 Objectives

Upon completion of this task, you will be able to:

- Configure IS-IS to ensure that PEs are routable to each other.
- Manually configure TE tunnels.
- Recurse L3VPN tunnels used for communication between CEs to static SR-MPLS Policies based on DSCP values.

1.3.1.2 Networking Description

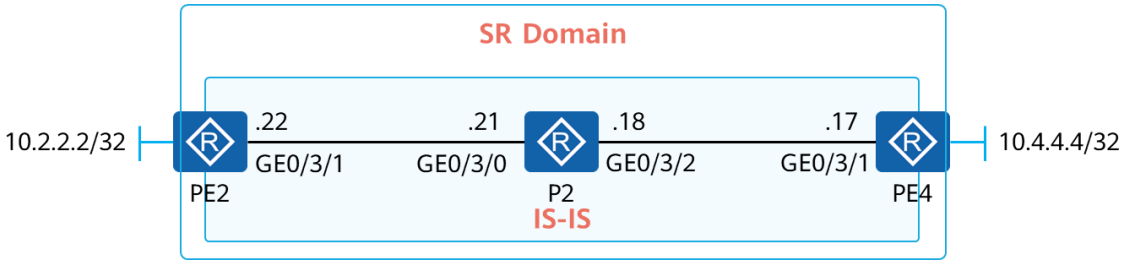

Figure 1-3 L3VPNv4 over static SR-MPLS Policy experiment

The figure shows the device connection and IP address planning. The interface interconnection addresses are in the format of 10.0.0.*Y*/30, and the values represented by *Y* are shown in the figure. Loopback0 is created on all devices. The Loopback0 address is used as the MPLS LSR ID of each device in the SR domain.

Loopback1 is created on PE2 and PE4 to simulate user access. The Loopback1 addresses on PE2 and PE4 are 10.2.2.2/32 and 10.4.4.4/32, respectively, as shown in the figure.

1.3.2 Experiment Task

1.3.2.1 Configuration Roadmap

1. Configure IP addresses for devices.
2. Configure IS-IS in the SR domain. Specifically, enable IS-IS on interconnection and Loopback0 interfaces for communication in the SR domain.
3. Configure MPLS. Specifically, enable MPLS and MPLS TE and set MPLS LSR IDs on devices.
4. Configure SR. Specifically, enable SR globally, enable IS-IS extensions for SR capabilities and TE, and configure node SIDs.
5. Configure explicit paths and TE tunnel interfaces on PE2 and PE4.
6. Configure a VPN instance, add Loopback1 on PE2 and PE4 to the instance, and establish VPNv4 peer relationships between PE2 and P2 and between PE4 and P2.
7. Configure a tunnel selection policy to recurse L3VPN traffic to an SR-MPLS Policy.

1.3.2.2 Configuration Procedure

Step 1 Configure IP addresses for interconnection and loopback interfaces.

Configure the configuration validation mode as immediate validation and configure IP addresses for interconnection and Loopback0 interfaces. Loopback0 addresses must be configured as planned in the following table.

Table 1-3 Loopback0 interface IP addresses

Device Number	Loopback0 IP Address
PE2	1.0.0.2
P2	1.0.0.6
PE4	1.0.0.4

Name the devices.

Omitted

Configure the configuration validation mode as immediate validation.

```
<PE2>system-view immediately

<P2>system-view immediately

<PE4>system-view immediately
```

Configure IP addresses for GE0/3/1 and Loopback0 on PE2.

```
[PE2] interface LoopBack 0
[PE2-LoopBack0] ip address 1.0.0.2 255.255.255.255
[PE2-LoopBack0] quit

[PE2]interface GigabitEthernet0/3/1
[PE2-GigabitEthernet0/3/1] ip address 10.0.0.22 255.255.255.252
```

Configure IP addresses for GE0/3/0, GE0/3/2, and Loopback0 on P2.

```
[P2] interface LoopBack 0
[P2-LoopBack0] ip address 1.0.0.6 255.255.255.255
[P2-LoopBack0] quit

[P2]interface GigabitEthernet0/3/0
[P2-GigabitEthernet0/3/0] ip address 10.0.0.21 255.255.255.252
[P2-GigabitEthernet0/3/0] quit
[P2]interface GigabitEthernet0/3/2
[P2-GigabitEthernet0/3/2] ip address 10.0.0.18 255.255.255.252
```

Configure IP addresses for GE0/3/1 and Loopback0 on PE4.

```
[PE4] interface LoopBack 0
[PE4-LoopBack0] ip address 1.0.0.4 255.255.255.255
[PE4-LoopBack0] quit

[PE4]interface GigabitEthernet0/3/1
[PE4-GigabitEthernet0/3/1] ip address 10.0.0.17 255.255.255.252
```

Test interconnection interface connectivity on P2.

```
[P2]ping -c 1 10.0.0.22
  PING 10.0.0.22: 56  data bytes, press CTRL_C to break
    Reply from 10.0.0.22: bytes=56 Sequence=1 ttl=255 time=1 ms

  --- 10.0.0.22 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 1/1/1 ms

[P2]ping -c 1 10.0.0.17
  PING 10.0.0.17: 56  data bytes, press CTRL_C to break
    Reply from 10.0.0.17: bytes=56 Sequence=1 ttl=255 time=1 ms

  --- 10.0.0.17 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 1/1/1 ms
```

Step 2 Configure IS-IS in the SR domain.

Ensure that the IS-IS area ID is 49.0001, the IS-IS process ID is 1, all devices are Level-2 devices, and the NET is converted from the Loopback0 IP address (for example, PE2's NET is 49.0001.0010.0000.0002.00). Then enable IS-IS on Loopback0 and interconnection interfaces.

In this case, you need to set **cost-style** to **wide** to support IS-IS extensions.

Configure PE2.

```
[PE2]isis 1
[PE2-isis-1] is-level level-2
[PE2-isis-1] cost-style wide
[PE2-isis-1] network-entity 49.0001.0010.0000.0002.00
[PE2-isis-1] is-name PE2
```

Configure P2.

```
[P2]isis 1
[P2-isis-1] is-level level-2
[P2-isis-1] cost-style wide
[P2-isis-1] network-entity 49.0001.0010.0000.0006.00
[P2-isis-1] is-name P2
```

Configure PE4.

```
[PE4]isis 1
[PE4-isis-1] is-level level-2
[PE4-isis-1] cost-style wide
[PE4-isis-1] network-entity 49.0001.0010.0000.0004.00
[PE4-isis-1] is-name PE4
```

Enable IS-IS on interfaces.

```
[PE2]interface LoopBack0
[PE2-LoopBack0] isis enable 1
[PE2-LoopBack0] quit
[PE2]interface GigabitEthernet0/3/1
[PE2-GigabitEthernet0/3/1] isis enable 1
[PE2-GigabitEthernet0/3/1] isis circuit-type p2p
```

```
[P2]interface LoopBack0
[P2-LoopBack0] isis enable 1
[P2-LoopBack0] quit
[P2]interface GigabitEthernet0/3/0
[P2-GigabitEthernet0/3/0] isis enable 1
[P2-GigabitEthernet0/3/0] isis circuit-type p2p
[P2-GigabitEthernet0/3/0] quit
[P2]interface GigabitEthernet0/3/2
[P2-GigabitEthernet0/3/2] isis enable 1
[P2-GigabitEthernet0/3/2] isis circuit-type p2p

[PE4]interface LoopBack0
[PE4-LoopBack0] isis enable 1
[PE4-LoopBack0] quit
[PE4]interface GigabitEthernet0/3/1
[PE4-GigabitEthernet0/3/1] isis enable 1
[PE4-GigabitEthernet0/3/1] isis circuit-type p2p
```

Check IS-IS neighbor relationships on P2.

```
[P2]display isis peer

                         Peerinformation for ISIS(1)

 System Id     Interface          Circuit Id         State  HoldTime Type       PRI
------------------------------------------------------------------------------
PE2*           GE0/3/0            0000000007         Up     26s      L2         --
PE4*           GE0/3/2            0000000007         Up     23s      L2         --
```

IS-IS neighbor relationships with PE2 and PE4 have been established.

Check IS-IS routes on P2.

```
[P2]display  isis route
```

```
                         Route information for ISIS(1)
                         ------------------------------

                        ISIS(1)  Level-2 Forwarding Table
                        ---------------------------------

IPV4Destination    IntCost     ExtCost    ExitInterface      NextHop          Flags
--------------------------------------------------------------------------------
1.0.0.2/32         10          NULL       GE0/3/0            10.0.0.22        A/-/-/-
1.0.0.4/32         10          NULL       GE0/3/2            10.0.0.17        A/-/-/-
10.0.0.20/30       10          NULL       GE0/3/0            Direct           D/-/L/-
10.0.0.16/30       10          NULL       GE0/3/2            Direct           D/-/L/-
```

P2 has learned the IS-IS routes generated by Loopback0 on PE2 and PE4.

Step 3 Configure MPLS.

Enable MPLS on the three devices and configure MPLS LSR IDs. MPLS does not need to be enabled on interfaces.

```
[PE2]Mpls lsr-id 1.0.0.2
[PE2]Mpls

[P2]Mpls lsr-id 1.0.0.6
[P2]Mpls

[PE4]Mpls lsr-id 1.0.0.4
[PE4]Mpls
```

Step 4 Configure SR capabilities on devices.

Enable SR-MPLS globally, enable IS-IS extensions for SR capabilities, configure an SRGB for IS-IS, and set the SRGB range to 16000 to 17000 on all devices.

Configure a SID for Loopback0 and use an index as the relative label value. The relative label value must be the same as the planned loopback address. For example, if the IP address of Loopback0 is 1.0.0.2, set the index to 2.

Enable SR-MPLS globally.

```
[PE2]segment-routing

[P2]segment-routing

[PE4]segment-routing
```

Enable IS-IS extensions for SR capabilities and configure an SRGB for IS-IS.

```
[PE2]isis 1
[PE2-isis-1]segment-routing mpls
[PE2-isis-1]segment-routing global-block 16000 17000

[P2]isis 1
[P2-isis-1]segment-routing mpls
[P2-isis-1]segment-routing global-block 16000 17000
```

```
[PE4]isis 1
[PE4-isis-1]segment-routing mpls
[PE4-isis-1]segment-routing global-block 16000 17000
```

Configure a node SID for devices.

```
[PE2]interface LoopBack 0
[PE2-LoopBack0]isis prefix-sid index 2

[P2]interface LoopBack 0
[P2-LoopBack0]isis prefix-sid index 6

[PE4]interface LoopBack 0
[PE4-LoopBack0]isis prefix-sid index 4
```

Manually configure adjacency SIDs on P2.

```
[P2]segment-routing
[P2-segment-routing] ipv4 adjacency local-ip-addr 10.0.0.21 remote-ip-addr 10.0.0.22 sid 321536
[P2-segment-routing] ipv4 adjacency local-ip-addr 10.0.0.18 remote-ip-addr 10.0.0.17 sid 321537
```

To ensure that the adjacency SIDs specified during explicit path configuration remain unchanged, you are advised to configure static adjacency SIDs. Then, the SIDs remain unchanged after the device restarts.

Step 5 Configure an SR-MPLS Policy.

Configure candidate paths and associate these candidate paths with the SR-MPLS Policy. Then configure an SR-MPLS Policy group and associate colors with DSCP values in the SR-MPLS policy group.

Configure candidate paths and associate them with an SR-MPLS Policy.

```
[PE2]segment-routing
[PE2-segment-routing] segment-list PE2_PE4
[PE2-segment-routing-segment-list-PE2_PE4] index 10 sid label 16006
[PE2-segment-routing-segment-list-PE2_PE4] index 20 sid label 321537
[PE2-segment-routing-segment-list-PE2_PE4] sr-te policy p1 endpoint 1.0.0.4 color 100
[PE2-segment-routing-te-policy-p1] candidate-path preference 100
[PE2-segment-routing-te-policy-p1-path] segment-list PE2_PE4
```

```
[PE4]segment-routing
[PE4-segment-routing] segment-list PE4_PE2
[PE4-segment-routing-segment-list-PE4_PE2] index 10 sid label 16006
[PE4-segment-routing-segment-list-PE4_PE2] index 20 sid label 321536
[PE4-segment-routing-segment-list-PE4_PE2] sr-te policy p1 endpoint 1.0.0.2 color 100
[PE4-segment-routing-te-policy-p1] candidate-path preference 100
[PE4-segment-routing-te-policy-p1-path] segment-list PE4_PE2
```

Associate colors with DSCP values to implement DSCP-based differentiated services. Packets from the same VPN instance carry different DSCP values and are mapped to different colors. The DSCP value -> color -> SR-MPLS TE Policy mapping is then formed (colors are already associated with SR-MPLS TE Policies). This allows packets destined for the same address to recurse to different SR-MPLS TE Policies based on DSCP values carried in these packets, implementing differentiated services.

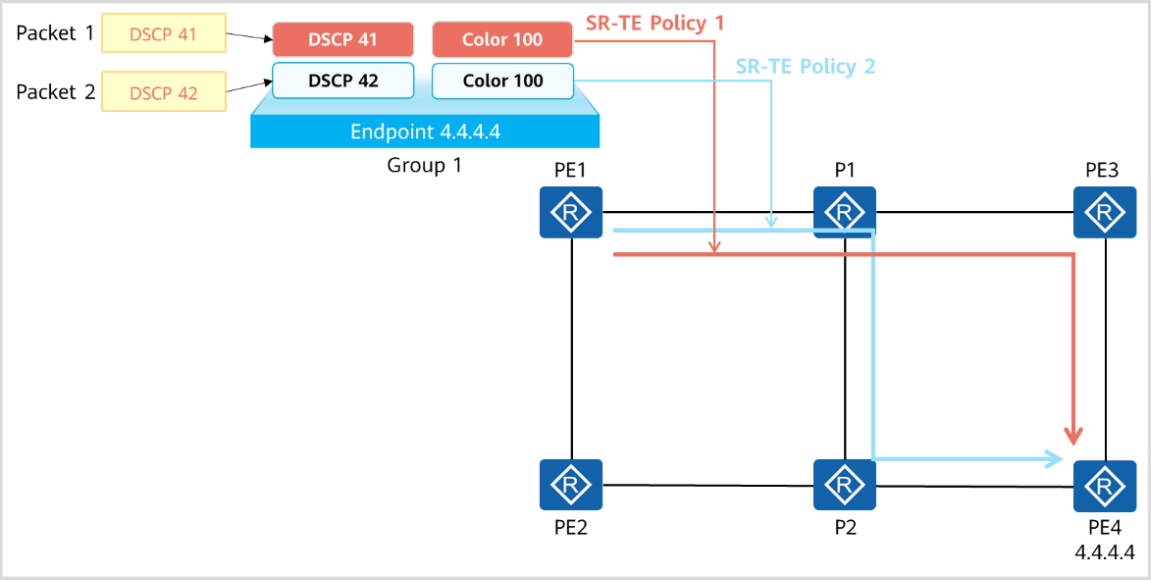

Figure 1-4 SR-MPLS Policy group traffic mapping

Create an SR-MPLS Policy group.

```
[PE2] segment-routing
[PE2-segment-routing] sr-te-policy group 1
[PE2-segment-routing-te-policy-group-1] endpoint 1.0.0.4
[PE2-segment-routing-te-policy-group-1] color 100 match dscp ipv4 41
```

```
[PE4] segment-routing
[PE4-segment-routing] sr-te-policy group 1
[PE4-segment-routing-te-policy-group-1] endpoint 1.0.0.2
[PE4-segment-routing-te-policy-group-1] color 100 match dscp ipv4 41
```

Check the tunnel status on PE2.

```
[PE2]display tunnel-info all
Tunnel ID                Type                 Destination      Status
--------------------------------------------------------------------------------
0x000000002900000002 srbe-lsp             1.0.0.6          UP
0x000000002900000005 srbe-lsp             1.0.0.4          UP
0x00000000330001c002 srtepolicygroup      1.0.0.4          UP
```

A tunnel of the srtepolicygroup type can be found.

Test tunnel connectivity.

```
[PE2]ping lsp sr-te policy policy-name p1
  LSP PING FEC: Nil FEC : 100   data bytes, press CTRL_C to break
  sr-te policy's segment list:
  Preference: 100; Path Type: primary; Protocol-Origin: local; Originator: 0, 0.0.0.0; Discriminator: 100;
Segment-List ID: 106497; Xcindex: 2106497
    Reply from 1.0.0.4: bytes=100 Sequence=1 time=11 ms
    Reply from 1.0.0.4: bytes=100 Sequence=2 time=2 ms
    Reply from 1.0.0.4: bytes=100 Sequence=3 time=2 ms
    Reply from 1.0.0.4: bytes=100 Sequence=4 time=2 ms
    Reply from 1.0.0.4: bytes=100 Sequence=5 time=2 ms

  --- FEC: Nil FEC ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 2/3/11 ms
```

The tunnel connectivity is normal.

Step 6 Configure an L3VPN.

Create a VPN instance named **vpnb** on PE2 and PE4, add Loopback1 to the VPN instance, and establish an MP-BGP VPNv4 peer relationship between PE2 and P2 and between PE4 and P2 (the AS number is 65001). P2 functions as the RR, and PE2 and PE4 function as the RR clients and advertise VPNv4 routes through P2.

Create a VPN instance named **vpnb**.

```
[PE2]ip vpn-instance vpnb
[PE2-vpn-instance-vpnb] ipv4-family
[PE2-vpn-instance-vpnb-af-ipv4] route-distinguisher 100:20
[PE2-vpn-instance-vpnb-af-ipv4] vpn-target 100:1020 both

[PE4]ip vpn-instance vpnb
[PE4-vpn-instance-vpnb] ipv4-family
[PE4-vpn-instance-vpnb-af-ipv4] route-distinguisher 100:40
[PE4-vpn-instance-vpnb-af-ipv4] vpn-target 100:1020 both
```

Create Loopback1, associate it with the VPN instance, and configure an IP address for the interface.

```
[PE2]interface LoopBack 1
[PE2-LoopBack1]ip binding vpn-instance vpnb
Info: All IPv4 and IPv6 related configurations on this interface are removed.
[PE2-LoopBack1]ip address 10.2.2.2 32
```

```
[PE4]interface LoopBack 1
[PE4-LoopBack1]ip binding vpn-instance vpnb
Info: All IPv4 and IPv6 related configurations on this interface are removed.
[PE4-LoopBack1]ip address 10.4.4.4 32
```

Note that you need to associate the interface with the VPN instance before configuring an IP address for the interface.

Use Loopback0 to configure the MP-BGP VPNv4 peer relationship and use the Loopback0 address as the router ID.

```
[PE2]bgp 65001
[PE2-bgp] router-id 1.0.0.2
[PE2-bgp] peer 1.0.0.6 as-number 65001
[PE2-bgp] peer 1.0.0.6 connect-interface LoopBack0
[PE2-bgp]ipv4-family vpnv4
[PE2-bgp-af-vpnv4] peer 1.0.0.6 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
```

```
[PE4]bgp 65001
[PE4-bgp] router-id 1.0.0.11
[PE4-bgp] peer 1.0.0.6 as-number 65001
[PE4-bgp] peer 1.0.0.6 connect-interface LoopBack0
[PE4-bgp]ipv4-family vpnv4
[PE4-bgp-af-vpnv4] peer 1.0.0.6 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
```

```
[P2-bgp]bgp 65001
[P2-bgp] router-id 1.0.0.6
[P2-bgp] peer 1.0.0.2 as-number 65001
[P2-bgp] peer 1.0.0.2 connect-interface LoopBack0
[P2-bgp] peer 1.0.0.4 as-number 65001
[P2-bgp] peer 1.0.0.4 connect-interface LoopBack0
[P2-bgp] ipv4-family vpnv4
[P2-bgp-af-vpnv4] undo policy vpn-target
[P2-bgp-af-vpnv4] peer 1.0.0.2 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
[P2-bgp-af-vpnv4] peer 1.0.0.2 reflect-client
[P2-bgp-af-vpnv4] peer 1.0.0.4 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
[P2-bgp-af-vpnv4] peer 1.0.0.4 reflect-client
```

When configuring an RR, disable the RT check on VPNv4 routes.

Check the VPNv4 peer relationship status.

```
[P2]display  bgp vpnv4 all peer

BGP  local   router  ID : 1.0.0.6
Local AS number : 65001
Total number  of peers: 2                    Peers in established   state : 2

 Peer          V        AS      MsgRcvd   MsgSent  OutQ   Up/Down       State       PrefRcv
 1.0.0.2       4        65001        58       106     0     00:01:00    Established         1
 1.0.0.4       4        65001        52       154     0     00:03:03    Established         0
```

P2 has established MP-BGP VBNv4 peer relationships with PE2 and PE4.

Import the direct routes of Loopback1 to BGP.

```
[PE2]bgp 65001
[PE2-bgp]ipv4-family vpn-instance vpnb
[PE2-bgp-vpnb]import-route direct

[PE4]bgp 65001
[PE4-bgp]ipv4-family vpn-instance vpnb
[PE4-bgp-vpnb]import-route direct
```

Check VPNv4 routes on PE2.

```
[PE2]display bgp vpnv4 all routing-table | include 10.4.4.4
BGP Local router ID is 1.0.0.2
Status codes: *-valid, > -best, d -damped, x-best external, a -add path,
              h -history, i -internal, s-suppressed, S -Stale
              Origin : i -IGP, e -EGP, ?-incomplete
RPKI validation codes: V -valid, I -invalid, N -not-found


Total number of routes from all PE: 3
Route Distinguisher: 100:20


         Network          NextHop                MED         LocPrf    PrefVal Path/Ogn

Route Distinguisher: 100:40


         Network          NextHop                MED         LocPrf    PrefVal Path/Ogn
*>i      10.4.4.4/32      1.0.0.4                 0          100          0          ?

VPN-Instance vpnb, Router ID 1.0.0.2:

Total Number of Routes: 3
         Network          NextHop                MED         LocPrf    PrefVal Path/Ogn
*>i      10.4.4.4/32      1.0.0.4                             100          0          ?
```

PE2 has learned the VPNv4 route from PE4 through MP-BGP.

Check the IP routing table on PE2.

```
[PE2-bgp]display ip routing-table vpn-instance vpnb
Route Flags: R -relay, D -download to fib, T -to vpn-instance, B -black hole route
------------------------------------------------------------------------------
Routing Table: vpnb
         Destinations : 4        Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        10.2.2.2/32  Direct  0    0           D   127.0.0.1       LoopBack1
        10.4.4.4/32  IBGP    255  0           RD  1.0.0.4         GigabitEthernet0/3/1
      127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
```

```
255.255.255.255/32  Direct  0      0                D    127.0.0.1       InLoopBack0
```

The route to the remote CE has been loaded to the VPN instance routing table on PE2.

Check route details.

```
[PE2]display  ip routing-table vpn-instance vpnb 10.4.4.4  verbose
Route Flags:  R -relay,D -downloadtofib,T -tovpn-instance, B -blackholeroute
------------------------------------------------------------------------------
RoutingTable: vpnb
Summary Count : 1

Destination:  10.4.4.4/32
     Protocol:  IBGP                 Process ID  : 0
   Preference:  255                        Cost  : 0
      NextHop:  1.0.0.4              Neighbour  : 1.0.0.6
        State : Active Adv Relied          Age  : 00h03m17s
          Tag:  0                     Priority  : low
        Label:  48157                 QoSInfo : 0x0
    IndirectID:  0x1000347            Instance  :
RelayNextHop:  10.0.0.21              Interface : GigabitEthernet0/3/1
     TunnelID : 0x000000002900000005  Flags : RD
```

In this case, the tunnel ID is 000000002900000005.

Check tunnel status.

```
[PE2]display tunnel-info all | include 000000002900000005
Tunnel ID                Type             Destination          Status
------------------------------------------------------------------------------
0x000000002900000005 srbe-lsp             1.0.0.4              UP
```

Judging from the tunnel ID information, the route from PE2 to Loopback1 on PE4 still recurses to the SR-MPLS BE tunnel.

Step 7 Configure a tunnel binding policy.

Configure a tunnel binding policy to preferentially select the SR-MPLS Policy group and associate the tunnel policy with the VPN instance.

Configure a tunnel binding policy.

```
[PE2]tunnel-policy p1
[PE2-tunnel-policy-p1]tunnel binding destination 1.0.0.4 sr-te-policy group 1
```

```
[PE4]tunnel-policy p1
[PE4-tunnel-policy-p1] tunnel binding destination 1.0.0.2 sr-te-policy group 1
```

Associate the tunnel policy with the VPN instance.

```
[PE2]ip vpn-instance vpnb
[PE2-vpn-instance-vpnb] ipv4-family
[PE2-vpn-instance-vpnb-af-ipv4] tnl-policy p1
```

```
[PE4]ip vpn-instance vpnb
[PE4-vpn-instance-vpnb] ipv4-family
[PE4-vpn-instance-vpnb-af-ipv4] tnl-policy p1
```

Check the VPN instance routing table again.

```
[PE2]display ip routing-table  vpn-instance vpnb
Route Flags:  R -relay,D -downloadtofib,T -tovpn-instance, B -blackholeroute
------------------------------------------------------------------------------
RoutingTable: vpnb
         Destinations : 4        Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        10.2.2.2/32  Direct  0    0           D   127.0.0.1       LoopBack1
        10.4.4.4/32  IBGP    255  0           RD  1.0.0.4         SR-TE Policy Group
      127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0

[PE4]display ip routing-table vpn-instance vpnb
Route Flags:  R -relay,D -downloadtofib,T -tovpn-instance, B -blackholeroute
------------------------------------------------------------------------------
RoutingTable: vpnb
         Destinations : 4        Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        10.2.2.2/32  IBGP    255  0           RD  1.0.0.2         SR-TE Policy Group
        10.4.4.4/32  Direct  0    0           D   127.0.0.1       LoopBack1
      127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

The next hop of the VPNv4 route has changed to the SR-TE Policy group, and the route has recursed to an SR-MPLS Policy.

Test the connectivity between CEs.

```
[PE2]ping -vpn-instance vpnb -a 10.2.2.2 10.4.4.4
  PING 10.4.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.4.4.4: bytes=56 Sequence=1 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=2 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=3 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=4 ttl=254 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=5 ttl=254 time=1 ms

  --- 10.4.4.4 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 1/1/1 ms
```

The connectivity is normal.

1.3.3 Quiz

What is the 3-tuple used to uniquely identify an SR-MPLS Policy?

2 SRv6 Experiment

2.1 L3VPNv4 over SRv6 BE Experiment

2.1.1 Introduction

2.1.1.1 Objectives

Upon completion of this task, you will be able to:

- Configure SRv6 locators for automatic allocation of SIDs to local VPN routes.
- Recurse L3VPN tunnels used for communication between CEs to SRv6 BE tunnels.
- Observe packet forwarding over the SRv6 BE tunnel.

2.1.1.2 Networking Description

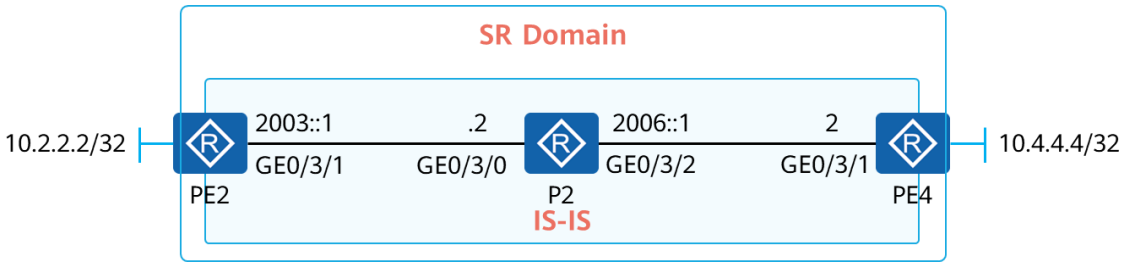

Figure 2-1 L3VPNv4 over SRv6 BE experiment topology

The figure shows the device connection and IP address planning. Loopback0 is created for all devices, and the Loopback0 IP address is FC01::*X*. For details about the value of *X*, see the following table.

Loopback1 is created on PE2 and PE4 to simulate user access. The Loopback1 addresses on PE2 and PE4 are 10.2.2.2/32 and 10.4.4.4/32, respectively, as shown in the figure.

2.1.2 Experiment Task

2.1.2.1 Configuration Roadmap

1. Configure IPv6 addresses for devices.
2. Configure IS-IS in the SR domain. Specifically, enable IS-IS on interconnection and Loopback0 interfaces for communication in the SR domain.
3. Create a VPN instance named **vpna**, add Loopback1 to the VPN instance on PE2 and PE4, and import direct routes to the BGP instance.
4. Establish an MP-IBGP peer relationship between PE2 and P2 and between PE4 and P2. P2 functions as an RR to reflect VPNv4 routes from PE2 and PE4.

5. Configure SRv6. Specifically, enable SRv6 globally, enable IS-IS extensions for SR capabilities, configure the source address for SRv6 encapsulation and locator, and enable the function to assign SIDs to VPN instance routes as well as the function to add SIDs to routes to be advertised to BGP peers.

2.1.2.2 Configuration Procedure

Step 1 Configure IPv6 addresses for interconnection and loopback interfaces.

Configure the configuration validation mode as immediate validation and configure IP addresses for interconnection and Loopback0 interfaces. Loopback0 addresses must be configured as planned in the following table.

Table 2-1 Loopback0 IP addresses

Device Number	Loopback0 IP Address
PE2	FC01::2
P2	FC01::6
PE4	FC01::4

Name the devices.

Omitted

Configure the configuration validation mode as immediate validation.

```
<PE2>system-view immediately

<P2>system-view immediately

<PE4>system-view immediately
```

Configure IP addresses for GE0/3/1 and Loopback0 on PE2.

```
[PE2]interface LoopBack0
[PE2-LoopBack0] ipv6 enable
[PE2-LoopBack0] ipv6 address FC01::2/128
[PE2-LoopBack0] quit
[PE2]interface GigabitEthernet0/3/1
[PE2-GigabitEthernet0/3/1] ipv6 enable
[PE2-GigabitEthernet0/3/1] ipv6 address 2003::1/64
```

Configure IP addresses for GE0/3/0, GE0/3/2, and Loopback0 on P2.

```
[P2]interface LoopBack0
[P2-LoopBack0] ipv6 enable
[P2-LoopBack0] ipv6 address FC01::6/128
[P2-LoopBack0] quit
[P2]interface GigabitEthernet0/3/0
[P2-GigabitEthernet0/3/0] ipv6 enable
[P2-GigabitEthernet0/3/0] ipv6 address 2003::2/64
```

```
[P2-GigabitEthernet0/3/0] quit
[P2]interface GigabitEthernet0/3/2
[P2-GigabitEthernet0/3/2] ipv6 enable
[P2-GigabitEthernet0/3/2] ipv6 address 2006::1/64
```

Configure IP addresses for GE0/3/1 and Loopback0 on PE4.

```
[PE4] interface LoopBack 0
[PE4-LoopBack0] ipv6 enable
[PE4-LoopBack0] ipv6 address FC01::4/128
[PE4-LoopBack0] quit

[PE4]interface GigabitEthernet0/3/1
[PE4-GigabitEthernet0/3/1] ipv6 enable
[PE4-GigabitEthernet0/3/1] ipv6 address 2006::2/64
```

Test interconnection interface connectivity on P2.

```
[P2]ping ipv6 2003::1
  PING 2003::1 : 56  data bytes, press CTRL_C to break
    Reply from 2003::1
    bytes=56 Sequence=1 hop limit=64 time=1 ms
    Reply from 2003::1
    bytes=56 Sequence=2 hop limit=64 time=1 ms
    Reply from 2003::1
    bytes=56 Sequence=3 hop limit=64 time=1 ms
    Reply from 2003::1
    bytes=56 Sequence=4 hop limit=64 time=1 ms
    Reply from 2003::1
    bytes=56 Sequence=5 hop limit=64 time=1 ms

  --- 2003::1 ping statistics---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max=1/1/1 ms

[P2]ping ipv6 2006::2
  PING 2006::2 : 56  data bytes, press CTRL_C to break
    Reply from 2006::2
    bytes=56 Sequence=1 hop limit=64 time=1 ms
    Reply from 2006::2
    bytes=56 Sequence=2 hop limit=64 time=1 ms
    Reply from 2006::2
    bytes=56 Sequence=3 hop limit=64 time=1 ms
    Reply from 2006::2
    bytes=56 Sequence=4 hop limit=64 time=1 ms
    Reply from 2006::2
    bytes=56 Sequence=5 hop limit=64 time=1 ms

  --- 2006::2 ping statistics---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
```

```
round-trip min/avg/max=1/1/1 ms
```

Step 2 Configure IS-IS in the SR domain.

Ensure that the IS-IS area ID is 49.0001, the IS-IS process ID is 1, all devices are Level-2 devices, and the NET is converted from the Loopback0 IP address. For example, the NET of PE2 is 49.0001.0010.0000.0002.00. Enable IPv6 for the IS-IS process and enable IS-IS IPv6 on Loopback0 and interconnection interfaces.

In this case, you need to set **cost-style** to **wide** to support IS-IS extensions.

Configure PE2.

```
[PE2]isis 1
[PE2-isis-1] is-level level-2
[PE2-isis-1] cost-style wide
[PE2-isis-1] network-entity 49.0001.0010.0000.0002.00
[PE2-isis-1] is-name PE2
[PE2-isis-1] ipv6 enable topology ipv6
```

Configure P2.

```
[P2]isis 1
[P2-isis-1] is-level level-2
[P2-isis-1] cost-style wide
[P2-isis-1] network-entity 49.0001.0010.0000.0006.00
[P2-isis-1] is-name P2
[P2-isis-1] ipv6 enable topology ipv6
```

Configure PE4.

```
[PE4]isis 1
[PE4-isis-1] is-level level-2
[PE4-isis-1] cost-style wide
[PE4-isis-1] network-entity 49.0001.0010.0000.0004.00
[PE4-isis-1] is-name PE4
[PE4-isis-1] ipv6 enable topology ipv6
```

Enable IS-IS IPv6 on interfaces.

```
[PE2]interface LoopBack0
[PE2-LoopBack0] isis ipv6 enable 1
[PE2-LoopBack0] quit
[PE2]interface GigabitEthernet0/3/1
[PE2-GigabitEthernet0/3/1] isis ipv6 enable 1
[PE2-GigabitEthernet0/3/1] isis circuit-type p2p
```

```
[P2]interface LoopBack0
[P2-LoopBack0] isis ipv6 enable 1
[P2-LoopBack0] quit
[P2]interface GigabitEthernet0/3/0
[P2-GigabitEthernet0/3/0] isis ipv6 enable 1
```

```
[P2-GigabitEthernet0/3/0] isis circuit-type p2p
[P2-GigabitEthernet0/3/0] quit
[P2]interface GigabitEthernet0/3/2
[P2-GigabitEthernet0/3/2] isis ipv6 enable 1
[P2-GigabitEthernet0/3/2] isis circuit-type p2p
[P2-GigabitEthernet0/3/2] quit
```

```
[PE4]interface LoopBack0
[PE4-LoopBack0] isis ipv6 enable 1
[PE4-LoopBack0] quit
[PE4]interface GigabitEthernet0/3/1
[PE4-GigabitEthernet0/3/1] isis ipv6 enable 1
[PE4-GigabitEthernet0/3/1] isis circuit-type p2p
[PE4-GigabitEthernet0/3/1] quit
```

Check IS-IS neighbor relationships on P2.

```
[P2]display isis  peer

                        Peerinformation for ISIS(1)

 System Id     Interface          CircuitId         State  HoldTime Type     PRI
------------------------------------------------------------------------------
PE2*           GE0/3/0            0000000007        Up     22s      L2       --
PE4*           GE0/3/2            0000000007        Up     25s      L2       --
```

IS-IS neighbor relationships with PE2 and PE4 have been established.

Check IS-IS IPv6 routes on P2.

```
[P2]display isis  route ipv6

                         Route information for ISIS(1)
                         -----------------------------

                        ISIS(1)  Level-2 Forwarding Table
                        ---------------------------------

IPV6  Dest.      ExitInterface       NextHop                       Cost      Flags
------------------------------------------------------------------------------
2003::/64        GE0/3/0             Direct                        10        D/-/-
2006::/64        GE0/3/2             Direct                        10        D/-/L/-
FC01::2/128      GE0/3/0             FE80::DE99:14FF:FE7A:C28A     10        A/-/-/-
FC01::4/128      GE0/3/2             FE80::DE99:14FF:FE7A:C1D6     10        A/-/-/-
FC01::6/128      Loop0               Direct                        0         D/-/L/-
     Flags:  D-Direct,  A-AddedtoURT, L-Advertised in LSPs, S-IGPShortcut,
             U-Up/DownBit Set, LP-Local Prefix-Sid
     Protect Type:  L-Link Protect,N-Node  Protect
```

P2 has learned the IS-IS IPv6 routes generated by Loopback0 on PE2 and PE4.

Step 3 Configure an L3VPN.

Create a VPN instance named **vpna** on PE2 and PE4, add Loopback1 to the VPN instance, and establish an MP-BGP VPNv4 peer relationship between PE2 and P2 and between PE4 and P2 (the AS number is 65001). P2 functions as the RR, and PE2 and PE4 function as the RR clients and advertise VPNv4 routes through P2.

Create a VPN instance named **vpna**.

```
[PE2]ip vpn-instance vpna
[PE2-vpn-instance-vpna] ipv4-family
[PE2-vpn-instance-vpna-af-ipv4] route-distinguisher 100:20
[PE2-vpn-instance-vpna-af-ipv4] vpn-target 100:1020 both

[PE4]ip vpn-instance vpna
[PE4-vpn-instance-vpna] ipv4-family
[PE4-vpn-instance-vpna-af-ipv4] route-distinguisher 100:40
[PE4-vpn-instance-vpna-af-ipv4] vpn-target 100:1020 both
```

Create Loopback1, associate it with the VPN instance, and configure an IP address for the interface.

```
[PE2]interface LoopBack 1
[PE2-LoopBack1] ip binding vpn-instance vpna
Info: All IPv4 and IPv6 related configurations on this interface are removed.
[PE2-LoopBack1] ip address 10.2.2.2 32
```

```
[PE4]interface LoopBack 1
[PE4-LoopBack1] ip binding vpn-instance vpna
Info: All IPv4 and IPv6 related configurations on this interface are removed.
[PE4-LoopBack1] ip address 10.4.4.4 32
```

Note that you need to associate the interface with the VPN instance before configuring an IP address for the interface.

Use Loopback0 to establish MP-BGP VPNv4 peer relationships. Configure router IDs as planned in the following table.

Table 2-2 Loopback0 IP addresses

Device Number	Router ID
PE2	1.0.0.2
P2	1.0.0.6
PE4	1.0.0.4

```
[PE2]bgp 65001
[PE2-bgp] router-id 1.0.0.2
[PE2-bgp] peer FC01::6 as-number 65001
[PE2-bgp] peer FC01::6 connect-interface LoopBack0
```

```
[PE2-bgp]ipv4-family vpnv4
[PE2-bgp-af-vpnv4] peer FC01::6 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
```

```
[P2]bgp 65001
[P2-bgp] router-id 1.0.0.6
[P2-bgp] peer FC01::2 as-number 65001
[P2-bgp] peer FC01::2 connect-interface LoopBack0
[P2-bgp] peer FC01::4 as-number 65001
[P2-bgp] peer FC01::4 connect-interface LoopBack0
[P2-bgp] ipv4-family vpnv4
[P2-bgp-af-vpnv4] undo policy vpn-target
[P2-bgp-af-vpnv4] peer FC01::2 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
[P2-bgp-af-vpnv4] peer FC01::2 reflect-client
[P2-bgp-af-vpnv4] peer FC01::4 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
[P2-bgp-af-vpnv4] peer FC01::4 reflect-client
```

```
[PE4]bgp 65001
[PE4-bgp] router-id 1.0.0.4
[PE4-bgp] peer FC01::6 as-number 65001
[PE4-bgp] peer FC01::6 connect-interface LoopBack0
[PE4-bgp]ipv4-family vpnv4
[PE4-bgp-af-vpnv4] peer FC01::6 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
```

When configuring an RR, disable the RT check on VPNv4 routes.

Check the VPNv4 peer relationship status.

```
[P2]display  bgp vpnv4 all peer

BGP  local   router  ID : 1.0.0.6
Local AS number : 65001
Total number  of  peers: 2                    Peers in established   state : 2

 Peer         V        AS      MsgRcvd   MsgSent  OutQ   Up/Down  State          PrefRcv
 FC01::2      4        65001     88303     89543     0    18m22s     Established        0
 FC01::4      4        65001     88381     89565     0    17m17s     Established        1
```

P2 has established VPNv4 peer relationships with PE2 and PE4.

Import the direct routes of Loopback1 to BGP.

```
[PE2]bgp 65001
[PE2-bgp]ipv4-family vpn-instance vpna
[PE2-bgp-vpna]import-route direct
[PE4]bgp 65001
[PE4-bgp]ipv4-family vpn-instance vpna
[PE4-bgp-vpna]import-route direct
```

Check VPNv4 routes on PE2.

```
[PE2]display bgp vpnv4 all routing-table
BGP Local router ID is 1.0.0.2
Status codes: *-valid,> -best, d -damped,x-best external,a -add path,
              h -history, i -internal, s-suppressed, S -Stale
              Origin : i -IGP, e -EGP, ?-incomplete
RPKI validation codes: V -valid, I -invalid, N -not-found

Total number of routes from all PE: 3
Route Distinguisher: 100:20

         Network          NextHop          MED        LocPrf     PrefVal Path/Ogn

*>       10.2.2.2/32      0.0.0.0          0          0          ?
*>       127.0.0.0/8      0.0.0.0          0          0          ?
Route Distinguisher: 100:40

         Network          NextHop          MED        LocPrf     PrefVal Path/Ogn

*>i      10.4.4.4/32      1.0.0.4           0         100          0         ?

VPN-Instance vpna, Router ID 1.0.0.2:

Total Number of Routes: 3
         Network          NextHop          MED        LocPrf     PrefVal Path/Ogn

*>       10.2.2.2/32      0.0.0.0           0                     0         ?
*>i      10.4.4.4/32      1.0.0.4           0         100         0         ?
*>       127.0.0.0/8      0.0.0.0           0                     0         ?
```

PE2 has learned the VPNv4 route from PE4 through MP-BGP.

Step 4 Configure SRv6.

On PE2 and PE4, enable SRv6 globally, configure the Loopback0 IPv6 addresses as the source addresses for SRv6 encapsulation, configure locators, enable automatic SRv6 SID allocation for VPN routes in the BGP VPN instance, enable the function to add SRv6 SIDs to VPN routes to be advertised in the BGP VPNv4 view, and enable the function to advertise SRv6 locators through IS-IS.

Configure SRv6 locators as planned in the following table.

Table 2-3 SRv6 locator planning

Device	IPv6 Prefix	MASK	Static Segment Length
PE2	FC00:2::	96	16
PE4	FC00:4::	96	16

Enable SR globally and configure the source address for SR encapsulation and locator.

```
[PE2]segment-routing ipv6
[PE2-segment-routing-ipv6] encapsulation source-address FC01::2
[PE2-segment-routing-ipv6] locator SRv6 ipv6-prefix FC00:2:: 96 static 16
```

```
[PE4]segment-routing ipv6
[PE4-segment-routing-ipv6] encapsulation source-address FC01::4
[PE4-segment-routing-ipv6] locator SRv6 ipv6-prefix FC00:4:: 96 static 16
```

Enable the function to add SIDs to VPN routes to be advertised to BGP peers.

```
[PE2]bgp 65001
[PE2-bgp] ipv4-family vpnv4
[PE2-bgp-af-vpnv4] peer FC01::6 prefix-sid
```

```
[P2]bgp 65001
[P2-bgp] ipv4-family vpnv4
[P2-bgp-af-vpnv4] peer FC01::4 prefix-sid
[P2-bgp-af-vpnv4] peer FC01::2 prefix-sid
```

```
[PE4-bgp] 65001
[PE4-bgp] ipv4-family vpnv4
[PE4-bgp-af-vpnv4] peer FC01::6 prefix-sid
```

Enable the function to add SIDs to VPN routes in the BGP VPN instance and specify the previously created SRv6 locator as the locator for allocated SIDs.

```
[PE2]bgp 65001
[PE2-bgp] ipv4-family vpn-instance vpna
[PE2-bgp-vpna] segment-routing ipv6 best-effort
[PE2-bgp-vpna] segment-routing ipv6 locator SRv6
```

```
[PE4]bgp 65001
[PE4-bgp] ipv4-family vpn-instance vpna
[PE4-bgp-vpna] segment-routing ipv6 best-effort
[PE4-bgp-vpna] segment-routing ipv6 locator SRv6
```

Enable IS-IS to advertise SRv6 locators.

```
[PE2]isis 1
[PE2-isis-1]segment-routing ipv6 locator SRv6
```

```
[PE4]isis 1
```

```
[PE4-isis-1] segment-routing ipv6 locator SRv6
```

Check IS-IS IPv6 routes on P2.

```
[P2]display isis route ipv6 FC00:2::

                Route information for ISIS(1)
                ----------------------------

               ISIS(1) Level-2 Forwarding Table
               --------------------------------

IPV6 Dest.      ExitInterface        NextHop                       Cost      Flags
-------------------------------------------------------------------------------
FC00:2::/96     GE0/3/0              FE80::DE99:14FF:FE7A:C28A     10        A/-/-/-
     Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
            U-Up/Down Bit Set, LP-Local Prefix-Sid
     Protect Type: L-Link Protect, N-Node Protect

[P2]display isis route ipv6 FC00:4::

                Route information for ISIS(1)
                ----------------------------

               ISIS(1) Level-2 Forwarding Table
               --------------------------------

 IPV6 Dest.      ExitInterface        NextHop                       Cost      Flags
-------------------------------------------------------------------------------
FC00:4::/96     GE0/3/2              FE80::DE99:14FF:FE7A:C1D6     10        A/-/-/-
     Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
            U-Up/Down Bit Set, LP-Local Prefix-Sid
     Protect Type: L-Link Protect, N-Node Protect
```

In the IS-IS IPv6 routing table, we can find routes generated based on locators on PE2 and PE4. Reachability to SIDs generated based on these locators is implemented through these routes.

Check the SIDs (VPN labels) generated by SRv6 for VPN routes.

```
[PE2]display segment-routing ipv6 local-sid end-dt4 forwarding

                My Local-SID End.DT4 Forwarding Table
                -------------------------------------

SID          : FC00:2::1:9D/128                    FuncType : End.DT4
VPN Name     : vpna                                VPNID    : 14
LocatorName  : SRv6                                LocatorID: 2

Total SID(s): 1
```

PE2 generates the SID FC00:2::1:9D for VPN routes in **vpna** and sends the SID to PE4 through a BGP Update message.

On PE4, check detailed information about the VPNv4 route (10.2.2.2) from PE2.

```
[PE4]display bgp vpnv4 all routing-table 10.2.2.2

  BGP local router ID : 1.0.0.4
  Local AS number : 65001

  Total routes of Route Distinguisher(100:20): 1
  BGP routing table entry information of 10.2.2.2/32:
  Label information (Received/Applied): 3/NULL
  From: FC01::6 (1.0.0.6)
  Route Duration: 0d00h01m34s
  Relay IP Nexthop: FE80::DE99:14FF:FE7A:C3F3
  Relay IP Out-Interface: GigabitEthernet0/3/1
  Relay Tunnel Out-Interface:
  Original nexthop: FC01::2
  Qos information : 0x0
  Ext-Community: RT <100 : 1020>
  Prefix-sid: FC00:2::1:9D
  AS-path Nil, origin incomplete, MED 0, localpref 100, pref-val 0, valid, internal, best, select, pre 255,
IGP cost 20
  Originator: 1.0.0.2
  Cluster list: 1.0.0.6
  Not advertised to any peer yet

  VPN-Instance vpna, Router ID 1.0.0.4:

  Total Number of Routes: 1
  BGP routing table entry information of 10.2.2.2/32:
  Route Distinguisher: 100:20
  Remote-Cross route
  Label information (Received/Applied): 3/NULL
  From: FC01::6 (1.0.0.6)
  Route Duration: 0d00h01m34s
  Relay IP Nexthop: FE80::DE99:14FF:FE7A:C3F3
  Relay IP Out-Interface: GigabitEthernet0/3/1
  Relay Tunnel Out-Interface:
  Original nexthop: FC01::2
  Qos information : 0x0
  Ext-Community: RT <100 : 1020>
  Prefix-sid: FC00:2::1:9D
  AS-path Nil, origin incomplete, MED 0, localpref 100, pref-val 0, valid, internal, best, select, pre 255,
IGP cost 20
  Originator: 1.0.0.2
  Cluster list: 1.0.0.6
  Not advertised to any peer yet
```

The BGP VPNv4 route carries the SID.

Check the VPN instance routing table on PE4.

```
[PE4]display ip routing-table vpn-instance vpna
Route Flags: R -relay,D -downloadtofib,T -tovpn-instance, B -blackholeroute
```

```
------------------------------------------------------------------------------
RoutingTable: vpna
         Destinations : 4        Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        10.2.2.2/32  IBGP   255  0           RD  FC00:2::1:9D    SRv6 BE
        10.4.4.4/32  Direct 0    0           D   127.0.0.1       LoopBack1
       127.0.0.0/8   Direct 0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct 0    0           D   127.0.0.1       InLoopBack0
```

The next hop of the route from PE4 to 10.2.2.2 is FC00:2::1:9D, that is, the SID assigned by PE2 to VPN routes in the VPN instance.

When the CE attached to PE4 accesses the CE attached to PE2, the destination IPv6 address carried in the outer packet header is this address. After receiving the packet, PE2 can determine to which CE the inner packet should be sent according to the destination IPv6 address.

Step 5 Observe the forwarding process.

Capture the headers of outgoing packets on GE0/3/0 of P2 and check the packet encapsulation during communication between 10.4.4.4 and 10.2.2.2.

On P2, create IPv6 ACL 3000 to match the outer headers of packets from 10.4.4.4 to 10.2.2.2.

```
[P2]acl ipv6 3000
[P2-acl6-advance-3000]rule permit ipv6 destination FC00:2::1:9D 128 source FC01::4 128
```

Use the ACL rule to match packets with the destination IPv6 address of FC00:2::1:9D and source IPv6 address of FC01::4.

Run the **capture-packet** command on P2 to capture packet headers on GE0/3/0.

```
[P2]capture-packet forwarding interface GigabitEthernet 0/3/0 outbound ipv6 acl 3000 packet-num 5
packet-len 64 overwrite file SRv6BE.cap
Info: Capture-packet data will be saved to cfcard:/logfile/SRv6BE.cap.
```

On PE4, ping 10.2.2.2 from 10.4.4.4.

```
[PE4]ping -vpn-instance vpna -a 10.4.4.4 10.2.2.2
  PING 10.2.2.2: 56  data bytes, press CTRL_C to break
    Reply from 10.2.2.2: bytes=56 Sequence=1 ttl=255 time=1 ms
    Reply from 10.2.2.2: bytes=56 Sequence=2 ttl=255 time=1 ms
    Reply from 10.2.2.2: bytes=56 Sequence=3 ttl=255 time=1 ms
    Reply from 10.2.2.2: bytes=56 Sequence=4 ttl=255 time=1 ms
    Reply from 10.2.2.2: bytes=56 Sequence=5 ttl=255 time=1 ms

  --- 10.2.2.2 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 1/1/1 ms
```

Information about the captured packet headers is saved in the **/logfile** directory of the device. We can download the file through FTP or SFTP. For details about how to enable FTP or SFTP on the device, see the related product documentation (for example: https://support.huawei.com/hedex/hdx.do?docid=EDOC1100168795&lang=en).

Check captured packet headers.

```
Frame 2: 138 bytes on wire (1104 bits), 64 bytes captured (512 bits)
Ethernet II, Src: HuaweiTe_7a:c3:f1 (dc:99:14:7a:c3:f1), Dst: HuaweiTe_7a:c2:8a (dc:99:14:7a:c2:8a)
Internet Protocol Version 6, Src: fc01::4, Dst: fc00:2::1:9D
Internet Protocol Version 4
```

The destination IPv6 address in the outer IPv6 packet header is fc00:2::19D. The IPv6 packet header is directly followed by the inner IPv4 header. In SRv6 BE, only one SID layer needs to be used for forwarding over public network routes and differentiation of VPN instances to which inner packets belong.

2.1.3 Quiz

In an L3VPNv6 over SRv6 BE scenario, which type of SID does BGP routes in a VPN instance carry?

2.2 L3VPNv4 over SRv6 Policy Experiment

2.2.1 Introduction

2.2.1.1 Objectives

Upon completion of this task, you will be able to:

- Manually allocate SIDs to VPN routes.
- Recurse L3VPN tunnels used for communication between CEs to SRv6 Policies.
- Observe packet forwarding over SRv6 Policies.

2.2.1.2 Networking Description

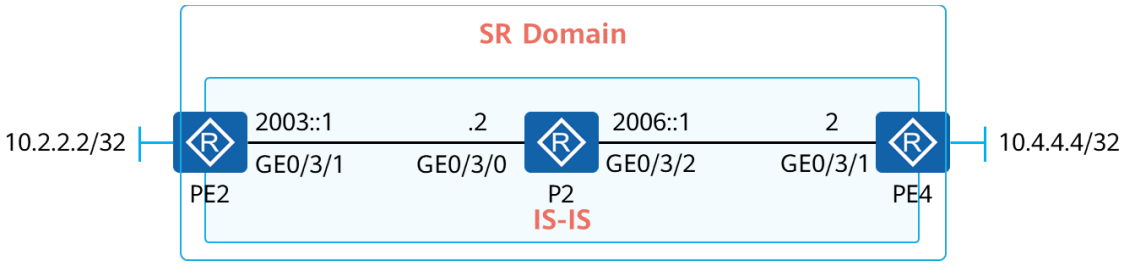

Figure 2-2 L3VPNv4 over SRv6 Policy experiment topology

The figure shows the device connection and IP address planning. Loopback0 is created for all devices, and the Loopback0 IP address is FC01::*X*. For details about the value of *X*, see the following table.

Loopback1 is created on PE2 and PE4 to simulate user access. The Loopback1 addresses on PE2 and PE4 are 10.2.2.2/32 and 10.4.4.4/32, respectively, as shown in the figure.

2.2.2 Experiment Task

2.2.2.1 Configuration Roadmap

1. Configure IP addresses for devices.
2. Configure IS-IS in the SR domain. Specifically, enable IS-IS on interconnection and Loopback0 interfaces for communication in the SR domain.
3. Create a VPN instance named **vpna**, add Loopback1 to the VPN instance on PE2 and PE4, and import direct routes to the BGP instance.
4. Establish an IBGP peer relationship between PE2 and P2 and between PE4 and P2. P2 functions as an RR to reflect VPNv4 routes from PE2 and PE4.
5. Configure a route-policy to allow PE2 and PE4 to add a color to VPNv4 routes to be advertised to each other.
6. Configure SRv6. Specifically, enable SRv6 globally, enable IS-IS extensions for SR capabilities, configure the source address for SRv6 encapsulation and locator, manually assign SIDs to VPN instance routes, manually assign SIDs used for device identification to devices, and enable the function to add SIDs to routes to be advertised to BGP peers.
7. Configure a tunnel policy to recurse VPN routes to SRv6 TE Policies.

2.2.2.2 Configuration Procedure

Step 1 Configure IP addresses for interconnection and loopback interfaces.

Configure the configuration validation mode as immediate validation and configure IP addresses for interconnection and Loopback0 interfaces. Loopback0 addresses must be configured as planned in the following table.

Table 2-4 Loopback0 IP addresses

Device Number	Loopback0 IP Address
PE2	FC01::2
P2	FC01::6
PE4	FC01::4

Name the devices.

Omitted

Configure the configuration validation mode as immediate validation.

```
<PE2>system-view immediately

<P2>system-view immediately
```

```
<PE4>system-view immediately
```

Configure IP addresses for GE0/3/1 and Loopback0 on PE2.

```
[PE2]interface LoopBack0
[PE2-LoopBack0] ipv6 enable
[PE2-LoopBack0] ipv6 address FC01::2/128
[PE2-LoopBack0] quit
[PE2]interface GigabitEthernet0/3/1
[PE2-GigabitEthernet0/3/1] ipv6 enable
[PE2-GigabitEthernet0/3/1] ipv6 address 2003::1/64
```

Configure IP addresses for GE0/3/0, GE0/3/2, and Loopback0 on P2.

```
[P2]interface LoopBack0
[P2-LoopBack0] ipv6 enable
[P2-LoopBack0] ipv6 address FC01::6/128
[P2-LoopBack0] quit
[P2]interface GigabitEthernet0/3/0
[P2-GigabitEthernet0/3/0] ipv6 enable
[P2-GigabitEthernet0/3/0] ipv6 address 2003::2/64
[P2-GigabitEthernet0/3/0] quit
[P2]interface GigabitEthernet0/3/2
[P2-GigabitEthernet0/3/2] ipv6 enable
[P2-GigabitEthernet0/3/2] ipv6 address 2006::1/64
```

Configure IP addresses for GE0/3/1 and Loopback0 on PE4.

```
[PE4] interface LoopBack 0
[PE4-LoopBack0] ipv6 enable
[PE4-LoopBack0] ipv6 address FC01::4/128
[PE4-LoopBack0] quit
[PE4]interface GigabitEthernet0/3/1
[PE4-GigabitEthernet0/3/1] ipv6 enable
[PE4-GigabitEthernet0/3/1] ipv6 address 2006::2/64
```

Test interconnection interface connectivity on P2.

```
[P2]ping ipv6 2003::1
  PING 2003::1 : 56  data bytes, press CTRL_C to break
    Reply from 2003::1
    bytes=56 Sequence=1 hop limit=64 time=1 ms
    Reply from 2003::1
    bytes=56 Sequence=2 hop limit=64 time=1 ms
    Reply from 2003::1
    bytes=56 Sequence=3 hop limit=64 time=1 ms
    Reply from 2003::1
    bytes=56 Sequence=4 hop limit=64 time=1 ms
    Reply from 2003::1
    bytes=56 Sequence=5 hop limit=64 time=1 ms

  --- 2003::1 ping statistics---
    5 packet(s) transmitted
    5 packet(s) received
```

```
    0.00% packet loss
    round-trip min/avg/max=1/1/1 ms
```

```
[P2]ping ipv6 2006::2
  PING 2006::2 : 56  data bytes, press CTRL_C to break
    Reply from 2006::2
    bytes=56 Sequence=1 hop limit=64 time=1 ms
    Reply from 2006::2
    bytes=56 Sequence=2 hop limit=64 time=1 ms
    Reply from 2006::2
    bytes=56 Sequence=3 hop limit=64 time=1 ms
    Reply from 2006::2
    bytes=56 Sequence=4 hop limit=64 time=1 ms
    Reply from 2006::2
    bytes=56 Sequence=5 hop limit=64 time=1 ms

  --- 2006::2 ping statistics---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
round-trip min/avg/max=1/1/1 ms
```

Step 2 Configure IS-IS in the SR domain.

Ensure that the IS-IS area ID is 49.0001, the IS-IS process ID is 1, all devices are Level-2 devices, and the NET is converted from the Loopback0 IP address. For example, the NET of PE2 is 49.0001.0010.0000.0002.00. Enable IPv6 for the IS-IS process and enable IS-IS IPv6 on Loopback0 and interconnection interfaces.

In this case, you need to set **cost-style** to **wide** to support IS-IS extensions.

Configure PE2.

```
[PE2]isis 1
[PE2-isis-1] is-level level-2
[PE2-isis-1] cost-style wide
[PE2-isis-1] network-entity 49.0001.0010.0000.0002.00
[PE2-isis-1] is-name PE2
[PE2-isis-1] ipv6 enable topology ipv6
```

Configure P2.

```
[P2]isis 1
[P2-isis-1] is-level level-2
[P2-isis-1] cost-style wide
[P2-isis-1] network-entity 49.0001.0010.0000.0006.00
[P2-isis-1] is-name P2
[P2-isis-1] ipv6 enable topology ipv6
```

Configure PE4.

```
[PE4]isis 1
[PE4-isis-1] is-level level-2
```

```
[PE4-isis-1] cost-style wide
[PE4-isis-1] network-entity 49.0001.0010.0000.0004.00
[PE4-isis-1] is-name PE4
[PE4-isis-1] ipv6 enable topology ipv6
```

Enable IS-IS IPv6 on interfaces.

```
[PE2]interface LoopBack0
[PE2-LoopBack0] isis ipv6 enable 1
[PE2-LoopBack0] quit
[PE2]interface GigabitEthernet0/3/1
[PE2-GigabitEthernet0/3/1] isis ipv6 enable 1
[PE2-GigabitEthernet0/3/1] isis circuit-type p2p
```

```
[P2]interface LoopBack0
[P2-LoopBack0] isis ipv6 enable 1
[P2-LoopBack0] quit
[P2]interface GigabitEthernet0/3/0
[P2-GigabitEthernet0/3/0] isis ipv6 enable 1
[P2-GigabitEthernet0/3/0] isis circuit-type p2p
[P2-GigabitEthernet0/3/0] quit
[P2]interface GigabitEthernet0/3/2
[P2-GigabitEthernet0/3/2] isis ipv6 enable 1
[P2-GigabitEthernet0/3/2] isis circuit-type p2p
[P2-GigabitEthernet0/3/2] quit
```

```
[PE4]interface LoopBack0
[PE4-LoopBack0] isis ipv6 enable 1
[PE4-LoopBack0] quit
[PE4]interface GigabitEthernet0/3/1
[PE4-GigabitEthernet0/3/1] isis ipv6 enable 1
[PE4-GigabitEthernet0/3/1] isis circuit-type p2p
[PE4-GigabitEthernet0/3/1] quit
```

Check IS-IS neighbor relationships on P2.

```
[P2]display  isis peer

                        Peer information for ISIS(1)

 System Id      Interface          Circuit Id         State  HoldTime Type       PRI
------------------------------------------------------------------------------
PE2*            GE0/3/0            0000000007         Up     22s      L2         --
PE4*            GE0/3/2            0000000007         Up     25s      L2         --
```

IS-IS neighbor relationships with PE2 and PE4 have been established.

Check IS-IS IPv6 routes on P2.

```
[P2]display isis  route  ipv6
```

```
                          Route information for ISIS(1)
                          ------------------------------

                         ISIS(1)  Level-2 Forwarding Table
                         ---------------------------------

IPV6  Dest.       ExitInterface        NextHop                          Cost      Flags
-------------------------------------------------------------------------------
2003::/64          GE0/3/0              Direct                           10         D/-/-
2006::/64          GE0/3/2              Direct                           10         D/-/L/-
FC01::2/128        GE0/3/0              FE80::DE99:14FF:FE7A:C28A        10         A/-/-/-
FC01::4/128        GE0/3/2              FE80::DE99:14FF:FE7A:C1D6        10         A/-/-/-
FC01::6/128        Loop0                Direct                           0          D/-/L/-
      Flags:  D-Direct,  A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
              U-Up/Down Bit Set, LP-Local Prefix-Sid
      Protect Type:  L-Link Protect, N-Node  Protect
```

P2 has learned the IS-IS IPv6 routes generated by Loopback0 on PE2 and PE4.

Step 3 Configure an L3VPN.

Create a VPN instance named **vpna** on PE2 and PE4, add Loopback1 to the VPN instance, and establish an MP-BGP VPNv4 peer relationship between PE2 and P2 and between PE4 and P2 (the AS number is 65001). P2 functions as the RR, and PE2 and PE4 function as the RR clients and advertise VPNv4 routes through P2.

Create a VPN instance named **vpna**.

```
[PE2]ip vpn-instance vpna
[PE2-vpn-instance-vpna] ipv4-family
[PE2-vpn-instance-vpna-af-ipv4] route-distinguisher 100:20
[PE2-vpn-instance-vpna-af-ipv4] vpn-target 100:1020 both

[PE4]ip vpn-instance vpna
[PE4-vpn-instance-vpna] ipv4-family
[PE4-vpn-instance-vpna-af-ipv4] route-distinguisher 100:40
[PE4-vpn-instance-vpna-af-ipv4] vpn-target 100:1020 both
```

Create Loopback1, associate it with the VPN instance, and configure an IP address for the interface.

```
[PE2]interface LoopBack 1
[PE2-LoopBack1] ip binding vpn-instance vpna
Info: All IPv4 and IPv6 related configurations on this interface are removed.
[PE2-LoopBack1] ip address 10.2.2.2 32

[PE4]interface LoopBack 1
[PE4-LoopBack1] ip binding vpn-instance vpna
Info: All IPv4 and IPv6 related configurations on this interface are removed.
[PE4-LoopBack1] ip address 10.4.4.4 32
```

Note that you need to associate the interface with the VPN instance before configuring an IP address for the interface.

Use Loopback0 to establish MP-BGP VPNv4 peer relationships. Configure router IDs as planned in the following table.

Table 2-5 Loopback0 IP addresses

Device Number	Router ID
PE2	1.0.0.2
P2	1.0.0.6
PE4	1.0.0.4

```
[PE2]bgp 65001
[PE2-bgp] router-id 1.0.0.2
[PE2-bgp] peer FC01::6 as-number 65001
[PE2-bgp] peer FC01::6 connect-interface LoopBack0
[PE2-bgp]ipv4-family vpnv4
[PE2-bgp-af-vpnv4] peer FC01::6 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y

[P2]bgp 65001
[P2-bgp] router-id 1.0.0.6
[P2-bgp] peer FC01::2 as-number 65001
[P2-bgp] peer FC01::2 connect-interface LoopBack0
[P2-bgp] peer FC01::4 as-number 65001
[P2-bgp] peer FC01::4 connect-interface LoopBack0
[P2-bgp] ipv4-family vpnv4
[P2-bgp-af-vpnv4] undo policy vpn-target
[P2-bgp-af-vpnv4] peer FC01::2 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
[P2-bgp-af-vpnv4] peer FC01::2 reflect-client
[P2-bgp-af-vpnv4] peer FC01::4 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
[P2-bgp-af-vpnv4] peer FC01::4 reflect-client

[PE4]bgp 65001
[PE4-bgp] router-id 1.0.0.4
[PE4-bgp] peer FC01::6 as-number 65001
[PE4-bgp] peer FC01::6 connect-interface LoopBack0
[PE4-bgp]ipv4-family vpnv4
[PE4-bgp-af-vpnv4] peer FC01::6 enable
Error: Please choose 'YES' or 'NO' first before pressing 'Enter'. [Y/N]:y
```

When configuring an RR, disable the RT check on VPNv4 routes.

Check the VPNv4 peer relationship status.

```
[P2]display  bgp vpnv4 all peer
BGP  local   router  ID  : 1.0.0.6
Local AS number  : 65001
Total number  of  peers: 2                    Peers in established   state : 2
```

```
 Peer          V          AS       MsgRcvd    MsgSent  OutQ  Up/Down      State          PrefRcv
 FC01::2       4          65001      88303      89543     0    18m22s      Established        0
 FC01::4       4          65001      88381      89565     0    17m17s      Established        1
```

P2 has established VPNv4 peer relationships with PE2 and PE4.

Import the direct routes of Loopback1 to BGP.

```
[PE2]bgp 65001
[PE2-bgp]ipv4-family vpn-instance vpna
[PE2-bgp-vpna]import-route direct

[PE4]bgp 65001
[PE4-bgp]ipv4-family vpn-instance vpna
[PE4-bgp-vpna]import-route direct
```

Check VPNv4 routes on PE2.

```
[PE2]display bgp vpnv4 all routing-table
BGP Local router ID  is 1.0.0.2
Status codes:  *-valid,> -best, d -damped,x-best external,a -add path,
                h -history,  i -internal, s-suppressed, S -Stale
                Origin : i -IGP, e -EGP, ?-incomplete
RPKI validation codes:  V -valid, I -invalid, N -not-found

Total number of routes from all PE:  3
Route Distinguisher:  100:20


         Network          NextHop          MED          LocPrf     PrefVal Path/Ogn

*>     10.2.2.2/32        0.0.0.0          0            0          ?
*>     127.0.0.0/8        0.0.0.0          0            0          ?
Route Distinguisher:  100:40

         Network          NextHop          MED          LocPrf     PrefVal Path/Ogn

*>i    10.4.4.4/32        1.0.0.4           0           100          0         ?

VPN-Instance vpna, Router ID 1.0.0.2:

Total Number of Routes:  3
         Network          NextHop          MED          LocPrf     PrefVal Path/Ogn

*>     10.2.2.2/32        0.0.0.0           0                       0         ?
*>i    10.4.4.4/32        1.0.0.4           0           100         0         ?
*>     127.0.0.0/8        0.0.0.0           0                       0         ?
```

PE2 has learned the VPNv4 route from PE4 through MP-BGP.

Create a route-policy on PE2 and PE4.

```
[PE2] route-policy Color permit node 10
[PE2-route-policy] apply extcommunity color 0:100
```

```
[PE4] route-policy Color permit node 10
[PE4-route-policy]apply extcommunity color 0:100
```

Apply the route-policy to routes to be advertised to the VPNv4 peer.

```
[PE2]bgp 65001
[PE2-bgp] ipv4-family vpnv4
[PE2-bgp-af-vpnv4] peer FC01::6 route-policy Color export

[PE4]bgp 65001
[PE4-bgp] ipv4-family vpnv4
[PE4-bgp-af-vpnv4] peer FC01::6 route-policy Color export
```

Check the color carried in VPNv4 routes on PE2 and PE4.

```
[PE2]display bgp vpnv4 all routing-table 10.4.4.4 | include Color

 BGP local router ID : 1.0.0.2
 Local AS number : 65001

 Total routes of Route Distinguisher(100:40): 1
 Ext-Community: RT <100 : 1020>, Color <0 : 100>

 VPN-Instance vpna, Router ID 1.0.0.2:

 Total Number of Routes: 1
 Ext-Community: RT <100 : 1020>, Color <0 : 100>

[PE4]display bgp vpnv4 all routing-table 10.2.2.2 | include Color

 BGP local router ID : 1.0.0.4
 Local AS number : 65001

 Total routes of Route Distinguisher(100:20): 1
 Ext-Community: RT <100 : 1020>, Color <0 : 100>

 VPN-Instance vpna, Router ID 1.0.0.4:

 Total Number of Routes: 1
 Ext-Community: RT <100 : 1020>, Color <0 : 100>
```

Step 4 Configure SRv6.

Enable SRv6 globally on PE2 and PE4, configure Loopback0 IPv6 addresses as source addresses for SRv6 encapsulation, configure locators, manually assign SIDs to PE2, P2, and PE4 as node IDs for these devices, and manually assign SIDs to VPN routes in the VPN instance. Enable the function to add SRv6 SIDs to VPN routes to be advertised in the BGP VPNv4 view and the function to advertise SRv6 locators through IS-IS.

Configure SRv6 locators as planned in the following table.

Table 2-6 SRv6 locator planning

Device	IPv6 Prefix	MASK	Static Segment Length

PE2	FC00:2::	96	16
P2	FC00:6::	96	16
PE4	FC00:4::	96	16

Enable SR globally and configure the source address for SR encapsulation and locator.

```
[PE2]segment-routing ipv6
[PE2-segment-routing-ipv6] encapsulation source-address FC01::2
[PE2-segment-routing-ipv6] locator SRv6 ipv6-prefix FC00:2:: 96 static 16
[PE2-segment-routing-ipv6-locator] opcode ::1 end
[PE2-segment-routing-ipv6-locator] opcode ::22 end-dt4 vpn-instance vpna

[P2]segment-routing ipv6
[P2-segment-routing-ipv6] encapsulation source-address FC01::6
[P2-segment-routing-ipv6] locator SRv6 ipv6-prefix FC00:6:: 96 static 16
[P2-segment-routing-ipv6-locator] opcode ::1 end

[PE4]segment-routing ipv6
[PE4-segment-routing-ipv6] encapsulation source-address FC01::4
[PE4-segment-routing-ipv6] locator SRv6 ipv6-prefix FC00:4:: 96 static 16
[PE4-segment-routing-ipv6-locator] opcode ::1 end
[PE4-segment-routing-ipv6-locator] opcode ::44 end-dt4 vpn-instance vpna
```

Configure node SIDs for PE2, P2, and PE4, and manually assign SIDs to the VPN instance on PE2 and PE4.

Enable the function to add SIDs to VPN routes to be advertised to BGP peers.

```
[PE2]bgp 65001
[PE2-bgp] ipv4-family vpnv4
[PE2-bgp-af-vpnv4] peer FC01::6 prefix-sid

[P2]bgp 65001
[P2-bgp] ipv4-family vpnv4
[P2-bgp-af-vpnv4] peer FC01::4 prefix-sid
[P2-bgp-af-vpnv4] peer FC01::2 prefix-sid

[PE4]bgp 65001
[PE4-bgp] ipv4-family vpnv4
[PE4-bgp-af-vpnv4] peer FC01::6 prefix-sid
```

In the BGP VPN instance, enable the function to recurse the service to an SRv6 Policy.

```
[PE2]bgp 65001
[PE2-bgp] ipv4-family vpn-instance vpna
[PE2-bgp-vpna] segment-routing ipv6 traffic-engineer best-effort
[PE2-bgp-vpna] segment-routing ipv6 locator SRv6 auto-sid-disable

[PE4]bgp 65001
[PE4-bgp] ipv4-family vpn-instance vpna
[PE4-bgp-vpna] segment-routing ipv6 traffic-engineer best-effort
```

```
[PE4-bgp-vpna] segment-routing ipv6 locator SRv6 auto-sid-disable
```

Because END.DT4 SIDs are manually assigned to VPN instance routes, you do not need to enable the function to automatically assign SIDs to routes. To disable this function, run the **auto-sid-disable** command.

Enable IS-IS to advertise SRv6 locators.

```
[PE2]isis 1
[PE2-isis-1] segment-routing ipv6 locator SRv6 auto-sid-disable

[PE4]isis 1
[PE4-isis-1] segment-routing ipv6 locator SRv6 auto-sid-disable
```

Because END SIDs are manually allocated in this example, disable automatic SID allocation here.

Check the END SID on PE2, P2, and PE4.

```
[PE2]display segment-routing ipv6 local-sid end forwarding

                        MyLocal-SID End Forwarding Table
                        ---------------------------------

SID             : FC00:2::1/128                                 FuncType : End
Flavor          : PSP
LocatorName : SRv6                                               LocatorID:   2

Total SID(s):  1

[P2]display segment-routing ipv6 local-sid end forwarding

                        MyLocal-SID End Forwarding Table
                        ---------------------------------

SID             : FC00:6::1/128                                 FuncType : End
Flavor          : PSP
LocatorName : SRv6                                               LocatorID:   2

Total SID(s):  1

[PE4]display segment-routing ipv6 local-sid end forwarding

                         My Local-SID End Forwarding Table
                         ---------------------------------

SID             : FC00:4::1/128                                FuncType : End
Flavor          : PSP
LocatorName : SRv6                                              LocatorID: 2

Total SID(s): 1
```

These END SIDs will be used to configure forwarding paths for SRv6 Policies.

Check the END.DT4 SID on PE2 and PE4.

```
[PE2]display  segment-routing ipv6 local-sid  end-dt4 forwarding

                    My Local-SID End.DT4 Forwarding Table
                    -------------------------------------

SID           : FC00:2::22/128                          FuncType : End.DT4
VPN Name      : vpna                                    VPNID    : 14
LocatorName   : SRv6                                    LocatorID:  2

Total SID(s): 1

[PE4]display  segment-routing ipv6 local-sid  end-dt4 forwarding

                    My Local-SID End.DT4 Forwarding Table
                    -------------------------------------

SID           : FC00:4::44/128                          FuncType  : End.DT4
VPN Name      : vpna                                    VPNID     : 20
LocatorName   : SRv6                                    LocatorID : 2

Total SID(s): 1
```

Check the SIDs carried in VPNv4 routes on PE2 and PE4.

```
[PE2]display bgp vpnv4 all routing-table 10.4.4.4 | include Prefix-sid

 BGP local router ID : 1.0.0.2
 Local AS number : 65001

 Total routes of Route Distinguisher(100:40): 1
 Prefix-sid: FC00:4::44

 VPN-Instance vpna, Router ID 1.0.0.2:

 Total Number of Routes: 1
 Prefix-sid: FC00:4::44
```

The SID of the VPNv4 route received by PE2 is FC00:4::44.

```
[PE4]display bgp vpnv4 all routing-table 10.2.2.2 | include Prefix-sid

 BGP local router ID : 1.0.0.11
 Local AS number : 65001

 Total routes of Route Distinguisher(100:20): 1
 Prefix-sid: FC00:2::22

 VPN-Instance vpna, Router ID 1.0.0.11:

 Total Number of Routes: 1
 Prefix-sid: FC00:2::22
```

The SID of the VPNv4 route received by PE4 is FC00:2::22.

Step 5 Configure SRv6 Policies and tunnel policies.

Configure candidate paths on PE2 and PE4 and use these candidate paths for SRv6 Policies.

Configure candidate paths on PE2.

```
[PE2]segment-routing ipv6
[PE2-segment-routing-ipv6] segment-list PE2_PE4
[PE2-segment-routing-ipv6-segment-list-PE2_PE4] index 5 sid ipv6 FC00:6::1
[PE2-segment-routing-ipv6-segment-list-PE2_PE4] index 10 sid ipv6 FC00:4::1
[PE2-segment-routing-ipv6] srv6-te policy p1 endpoint FC01::4 color 100
[PE2-segment-routing-ipv6-policy-p1] candidate-path preference 100
[PE2-segment-routing-ipv6-policy-p1-path] segment-list PE2_PE4
```

Configure candidate paths on PE4.

```
[PE4]segment-routing ipv6
[PE4-segment-routing-ipv6] segment-list PE4_PE2
[PE4-segment-routing-ipv6-segment-list-PE4_PE2] index 5 sid ipv6 FC00:6::1
[PE4-segment-routing-ipv6-segment-list-PE4_PE2] index 10 sid ipv6 FC00:2::1
[PE4-segment-routing-ipv6-segment-list-PE4_PE2] quit
[PE4-segment-routing-ipv6] srv6-te policy p1 endpoint fc01::2 color 100
[PE4-segment-routing-ipv6-policy-p1] candidate-path preference 100
[PE4-segment-routing-ipv6-policy-p1-path] segment-list PE4_PE2
```

Configure a tunnel policy.

```
[PE2]tunnel-policy p1
[PE2-tunnel-policy-p1] tunnel select-seq ipv6 srv6-te-policy load-balance-number 1
```

```
[PE4]tunnel-policy p1
[PE4-tunnel-policy-p1] tunnel select-seq ipv6 srv6-te-policy load-balance-number 1
```

Apply the tunnel policy to the VPN instance.

```
[PE2]ip vpn-instance vpna
[PE2-vpn-instance-vpna]tnl-policy p1
```

```
[PE4]ip vpn-instance vpna
[PE4-vpn-instance-vpna] tnl-policy p1
```

Check the VPN instance routing table on PE2 and PE4.

```
[PE2]display ip routing-table vpn-instance vpna
Route Flags: R -relay,D -downloadtofib,T -tovpn-instance, B -blackholeroute
------------------------------------------------------------------------------
RoutingTable: vpna
         Destinations : 4        Routes : 4
```

```
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        10.2.2.2/32  Direct  0    0           D   127.0.0.1       LoopBack1
        10.4.4.4/32  IBGP    255  0           RD  FC01::4         p1
       127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0

[PE4]display ip routing-table  vpn-instance vpna
Route Flags:  R -relay,D -downloadtofib,T -tovpn-instance, B -blackholeroute
------------------------------------------------------------------------------
RoutingTable: vpna
         Destinations : 4        Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        10.2.2.2/32  IBGP    255  0           RD  FC01::2         p1
        10.4.4.4/32  Direct  0    0           D   127.0.0.1       LoopBack1
       127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

The route has recursed to a logical interface (based on the tunnel policy).

Check detailed information about route 10.4.4.4.

```
[PE2]display ip routing-table vpn-instance vpna  10.4.4.4  verbose
Route Flags:  R -relay,D -downloadtofib,T -tovpn-instance, B -blackholeroute
------------------------------------------------------------------------------
RoutingTable: vpna
Summary Count : 1

Destination:  10.4.4.4/32
     Protocol  : IBGP                 Process ID  : 0
   Preference  : 255                        Cost: 0
      NextHop:  FC01::4               Neighbour  : FC01::6
        State  : Active Adv Relied          Age  : 00h04m19s
          Tag:  0                     Priority   :  low
        Label: 3                      QoSInfo : 0x0
  IndirectID  : 0x1000306             Instance  :
RelayNextHop:  ::                     Interface  : p1
     TunnelID : 0x000000003400002f41  Flags  : RD
```

The route has recursed to a tunnel with the tunnel ID of 0x000000003400002f41.

Check tunnel information on PE2.

```
[PE2]display tunnel-info all | in 0x000000003400002f41
Tunnel ID                  Type                Destination         Status
------------------------------------------------------------------------------
0x000000003400002f41 srv6tepolicy          FC01::4             UP
```

The tunnel is an SRv6 Policy.

Step 6 Observe the forwarding process.

Capture the headers of incoming packets on GE0/3/0 of P2 and check the packet encapsulation during communication between 10.2.2.2 and 10.4.4.4.

On P2, create IPv6 ACL 3000 to match the outer headers of packets from 10.2.2.2 to 10.4.4.4.

```
[P2]acl ipv6 3000
[P2-acl6-advance-3000] rule permit ipv6 source FC01::2 128 destination FC00:6::1 128
```

When these packets arrive at P2, the source IPv6 address is FC01::2 (source address for encapsulation on PE2) and the destination IPv6 address is FC00:6::1 (END SID of P2).

Run the **capture-packet** command on P2 to capture packet headers on GE0/3/0.

```
[P2]capture-packet forwarding interface GigabitEthernet 0/3/0 inbound ipv6 acl 3000 packet-num 5
packet-len 64 overwrite file SRv6TE.cap
Info: Capture-packet data will be saved to cfcard:/logfile/SRv6TE.cap.
```

On PE2, check the connectivity between Loopback1 on PE2 and Loopback1 on PE4.

```
[PE2]ping -vpn-instance vpna -a 10.2.2.2 10.4.4.4
  PING 10.4.4.4: 56   data bytes, press CTRL_C to break
    Reply from 10.4.4.4: bytes=56 Sequence=1 ttl=255 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=2 ttl=255 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=3 ttl=255 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=4 ttl=255 time=1 ms
    Reply from 10.4.4.4: bytes=56 Sequence=5 ttl=255 time=1 ms

  --- 10.4.4.4 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 1/1/1 ms
```

The connectivity is normal.

Check captured packet headers.

```
Frame 1: 194 bytes on wire (1552 bits), 64 bytes captured (512 bits)
Ethernet II, Src: HuaweiTe_7a:c2:8a (dc:99:14:7a:c2:8a), Dst: HuaweiTe_7a:c3:f1 (dc:99:14:7a:c3:f1)
Internet Protocol Version 6, Src: fc01::2, Dst: fc00:6::1
    0110 .... = Version: 6
    .... 0000 0000 .... .... .... .... .... = Traffic Class: 0x00 (DSCP: CS0, ECN: Not-ECT)
    .... .... .... 0000 0000 0000 0000 0000 = Flow Label: 0x00000
    Payload Length: 140
    Next Header: Routing Header for IPv6 (43)
    Hop Limit: 255
    Source Address: fc01::2
    Destination Address: fc00:6::1
    Routing Header for IPv6 (Segment Routing)
        Next Header: IPIP (4)
        Length: 6
        [Length: 56 bytes]
        Type: Segment Routing (4)
        Segments Left: 2
```

```
Last Entry: 2
Flags: 0x00
```

The SRH still carries two SIDs when packets reach P2. Due to limitations on the captured packet header length, the specific SIDs cannot be viewed. However, we know that one SID is PE4's END SID {FC00:4::1}, and the other SID (the last SID) is an END.DT4 SID {FC00:4::44}.

2.2.3 Quiz

What are End SIDs and End.DT4 SIDs used to identify?

3 iMaster NCE-IP Experiment

3.1 SR-MPLS Service Delivery by the Controller

3.1.1 Introduction

3.1.1.1 Objectives

Upon completion of this task, you will be able to:

- Use iMaster NCE-IP to manage devices and configure services.
- Establish BGP-LS and BGP SRv4 Policy address relationships between the controller and devices.
- Use iMaster NCE-IP to deliver L3VPNv4 over SR-MPLS TE configurations.
- Use iMaster NCE-IP to deliver L3VPNv4 over SR-MPLS Policy configurations.

3.1.1.2 Networking Description

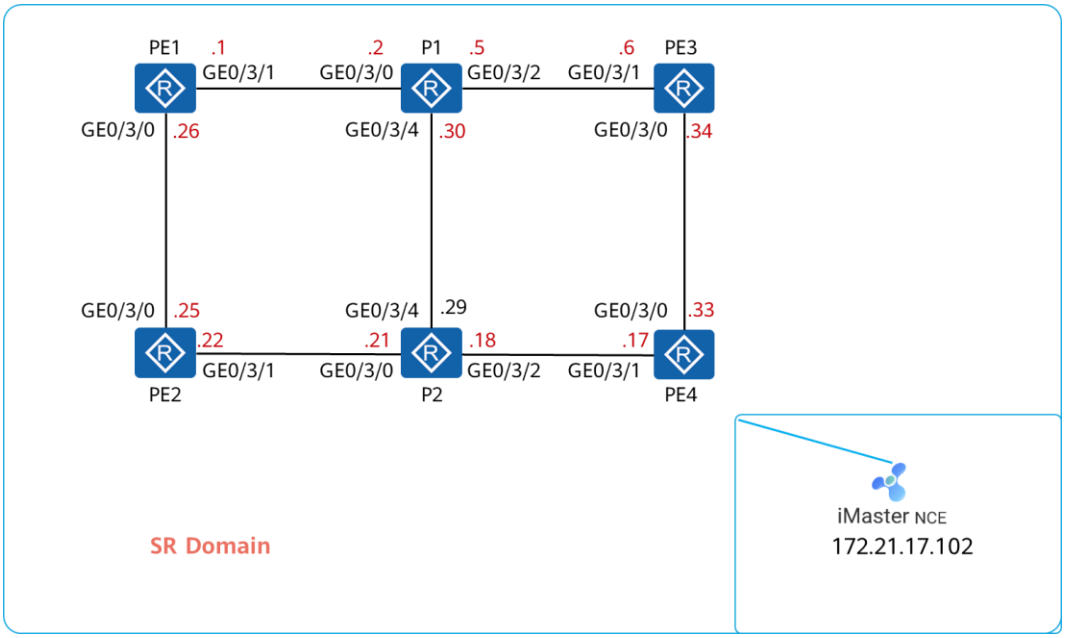

Figure 3-1 Experiment topology for SR-MPLS service delivery through the controller

The figure shows the device connection and IP address planning. The interface interconnection addresses are in the format of 10.0.0.*Y*/30, and the values represented by *Y* are shown in the figure. Loopback0 is created on all devices. The Loopback0 address is used as the MPLS LSR ID of each device in the SR domain.

IS-IS is enabled globally in the entire SR domain, and SR labels are distributed through IS-IS.

BGP runs in the AS. P1 and P2 function as RRs. All PEs establish VPNv4 peer relationships and SR Policy peer relationships with P1 and P2.

All devices connect to iMaster NCE-IP through the management interface (GE0/0/0). The controller address is shown in the figure.

Now we need to use iMaster NCE-IP to manage all devices and then deliver L3VPNv4 over SR-MPLS TE and L3VPNv4 over SR-MPLS Policy configurations.

3.1.2 Experiment Task

3.1.2.1 Configuration Roadmap

1. Perform basic device configurations, such as configuring IP addresses for Loopback0, interconnection, and management interfaces, and configuring SSH and SNMP.
2. Perform IGP, SR, and MPLS configurations, such as enabling SR globally, configuring SRGBs, enabling IS-IS topology to be reported through BGP-LS, and enabling MPLS and MPLS TE globally.
3. Perform BGP configurations, such as configuring VPNv4 and SR Policy peer relationships between devices and BGP-LS and SR Policy peer relationships between the RR and iMaster NCE-IP.
4. Perform controller configurations, such as configuring routes from iMaster NCE-IP to devices (for device management), adding devices to iMaster NCE-IP for management, and configuring BGP-LS and SR Policy peer relationships between iMaster NCE-IP and the RR.
5. Use iMaster NCE-IP to configure and deliver SR-MPLS TE tunnels.
6. Use iMaster NCE-IP to configure L3VPNv4 services, recurse these services to the previously created SR-MPLS TE tunnels, and then deliver these configurations.
7. Use iMaster NCE-IP to configure an SR-MPLS Policy, create an L3VPNv4 service, recurse the L3VPNv4 service to the SR-MPLS Policy, and then deliver these configurations.

3.1.2.2 Device-Side Basic Configurations

To complete the subsequent configuration, you need to create Loopback0 interfaces and configure IP addresses for management, interconnection, and Loopback0 interfaces first.

To enable iMaster NCE-IP to manage devices, enable LLDP and SSH and configure SFTP, NETCONF, and SNMP on all devices.

The management interface IP addresses are in the format of 172.21.17.*X*, and Loopback0 IP addresses are in the format of 1.0.0.*X*. For values represented by *X*, see the following table.

Table 3-1 Address planning

Device Name	Planned X Value
P1	5

P2	6
PE1	1
PE2	2
PE3	3
PE4	4

Step 2 Configure management interface IP addresses.

Configure IP addresses for the management interfaces (GE0/0/0) of all devices. These IP addresses are used for communication between the devices and iMaster NCE-IP.

Configure the configuration validation mode as immediate validation.

```
<PE2>system-view immediately
```

PE2 is used as an example. Repeat this operation for other devices.

#PE1

```
[PE1]interface GigabitEthernet0/0/0
[PE1-GigabitEthernet0/0/0] undo ip binding vpn-instance __LOCAL_OAM_VPN__
[PE1-GigabitEthernet0/0/0] ip address 172.21.17.1 24
[PE1-GigabitEthernet0/0/0] quit
```

#PE2

```
[PE2]interface GigabitEthernet0/0/0
[PE2-GigabitEthernet0/0/0] undo ip binding vpn-instance __LOCAL_OAM_VPN__
[PE2-GigabitEthernet0/0/0] ip address 172.21.17.2 24
[PE2-GigabitEthernet0/0/0] quit
```

#PE3

```
[PE3]interface GigabitEthernet0/0/0
[PE3-GigabitEthernet0/0/0] undo ip binding vpn-instance __LOCAL_OAM_VPN__
[PE3-GigabitEthernet0/0/0] ip address 172.21.17.3 24
[PE3-GigabitEthernet0/0/0] quit
```

#PE4

```
[PE4]interface GigabitEthernet0/0/0
[PE4-GigabitEthernet0/0/0] undo ip binding vpn-instance __LOCAL_OAM_VPN__
[PE4-GigabitEthernet0/0/0] ip address 172.21.17.4 24
[PE4-GigabitEthernet0/0/0] quit

#P1
[P1]interface GigabitEthernet0/0/0
[P1-GigabitEthernet0/0/0] undo ip binding vpn-instance __LOCAL_OAM_VPN__
[P1-GigabitEthernet0/0/0] ip address 172.21.17.5 24
```

```
[P1-GigabitEthernet0/0/0] quit

#P2
[P2]interface GigabitEthernet0/0/0
[P2-GigabitEthernet0/0/0] undo ip binding vpn-instance __LOCAL_OAM_VPN__
[P2-GigabitEthernet0/0/0] ip address 172.21.17.6 24
[P2-GigabitEthernet0/0/0] quit
```

Step 3 Configure Loopback0 IP addresses.

Create Loopback0 on each device and use the Loopback0 IP address as the BGP router ID and MPLS LSR ID of the device.

#PE1

```
[PE1]interface LoopBack0
[PE1-LoopBack0] ip address 1.0.0.1 32
[PE1-LoopBack0] quit
```

#PE2

```
[PE2]interface LoopBack0
[PE2-LoopBack0] ip address 1.0.0.2 32
[PE2-LoopBack0] quit
```

#PE3

```
[PE3]interface LoopBack0
[PE3-LoopBack0] ip address 1.0.0.3 32
[PE3-LoopBack0] quit
```

#PE4

```
[PE4]interface LoopBack0
[PE4-LoopBack0] ip address 1.0.0.4 32
[PE4-LoopBack0] quit
```

#P1

```
[P1]interface LoopBack0
[P1-LoopBack0] ip address 1.0.0.5 32
[P1-LoopBack0] quit
```

#P2

```
[P2]interface LoopBack0
[P2-LoopBack0] ip address 1.0.0.6 32
[P2-LoopBack0] quit
```

Step 4 Configure interconnection interface IP addresses.

Configure the interconnection interface IP address on each device. The interconnection interface address is 10.0.0.X, and the X value is marked in the topology.

By default, DCN is enabled on NE router interfaces. To facilitate the experiment, disable DCN globally on all devices.

Disable DCN globally on each device.

```
[PE1] undo dcn
Warning: This operation will disable DCN function. Continue? [Y/N]:y
```

PE1 is used as an example. Repeat this operation for other devices.

Configure P1 interface addresses.

```
[P1]interface GigabitEthernet0/3/0
[P1-GigabitEthernet0/3/0] ip address 10.0.0.2 30
[P1-GigabitEthernet0/3/0] quit
[P1]interface GigabitEthernet0/3/2
[P1-GigabitEthernet0/3/2] ip address 10.0.0.5 30
[P1-GigabitEthernet0/3/2] quit
[P1]interface GigabitEthernet0/3/4
[P1-GigabitEthernet0/3/4] ip address 10.0.0.30 30
[P1-GigabitEthernet0/3/4] quit
```

Configure P2 interface addresses.

```
[P2]interface GigabitEthernet0/3/0
[P2-GigabitEthernet0/3/0] ip address 10.0.0.21 30
[P2-GigabitEthernet0/3/0] quit
[P2]interface GigabitEthernet0/3/2
[P2-GigabitEthernet0/3/2] ip address 10.0.0.18 30
[P2-GigabitEthernet0/3/2] quit
[P2]interface GigabitEthernet0/3/4
[P2-GigabitEthernet0/3/4] ip address 10.0.0.29 30
[P2-GigabitEthernet0/3/4] quit
```

Configure PE1 interface addresses.

```
[PE1]interface GigabitEthernet0/3/0
[PE1-GigabitEthernet0/3/0] ip address 10.0.0.26 30
[PE1-GigabitEthernet0/3/0] quit
[PE1]interface GigabitEthernet0/3/1
[PE1-GigabitEthernet0/3/1] ip address 10.0.0.1 30
[PE1-GigabitEthernet0/3/1] quit
```

Configure PE2 interface addresses.

```
[PE2]interface GigabitEthernet0/3/0
[PE2-GigabitEthernet0/3/0] ip address 10.0.0.25 30
[PE2-GigabitEthernet0/3/0] quit
[PE2]interface GigabitEthernet0/3/1
[PE2-GigabitEthernet0/3/1] ip address 10.0.0.22 30
[PE2-GigabitEthernet0/3/1] quit
```

Configure PE3 interface addresses.

```
[PE3]interface GigabitEthernet0/3/0
[PE3-GigabitEthernet0/3/0] ip address 10.0.0.34 30
[PE3-GigabitEthernet0/3/0] quit
[PE3]interface GigabitEthernet0/3/1
[PE3-GigabitEthernet0/3/1] ip address 10.0.0.6 30
[PE3-GigabitEthernet0/3/1] quit
```

Configure PE4 interface addresses.

```
[PE4]interface GigabitEthernet0/3/0
[PE4-GigabitEthernet0/3/0] ip address 10.0.0.33 30
[PE4-GigabitEthernet0/3/0] quit
[PE4]interface GigabitEthernet0/3/1
[PE4-GigabitEthernet0/3/1] ip address 10.0.0.17 30
[PE4-GigabitEthernet0/3/1] quit
```

Step 5 Configure SSH.

Enable SSH, create an SSH user named **netconf**, and enable NETCONF on each device.

NETCONF is used by iMaster NCE-IP to deliver configurations to devices.

This step uses PE1 as an example. Repeat the configuration for other devices.

Configure a user interface to allow SSH packets to pass through.

```
[PE1]user-interface vty 0 4
[PE1-ui-vty0-4] authentication-mode aaa
[PE1-ui-vty0-4] protocol inbound all
```

Create a user named **netconf**.

```
[PE1]aaa
[PE1-aaa] local-user netconf password irreversible-cipher Huawei@123
Info: A new user is added.
[PE1-aaa] local-user netconf service-type ftp ssh
[PE1-aaa] local-user netconf level 3
[PE1-aaa] local-user netconf state block fail-times 3 interval 5
[PE1-aaa] local-user netconf user-group manage-ug
```

Create a user named **netconf**, set the user type to FTP and SSH, and add the user to the default user group **manage-ug**. Then, set the maximum number of SSH login attempts allowed to 3 and the lockout interval to 5s.

Enable SSH.

```
[PE1]stelnet server enable
```

Configure NETCONF.

```
[PE1]snetconf server enable
[PE1]ssh user netconf
[PE1]ssh user netconf authentication-type password
[PE1]ssh user netconf service-type all
[PE1]ssh client first-time enable
```

Enable NETCONF, set the password authentication mode for SSH login by NETCONF users, and set the service type to all (including SSH, SFTP, and SNETCONF).

Configure SFTP.

```
[PE1]sftp server enable
Info: Succeeded in starting the SFTP server.
[PE1]sftp client-source -i LoopBack0
Info: Succeeded in setting the source interface of the SFTP client to LoopBack0.
[PE1]ssh user netconf sftp-directory cfcard:
```

Enable SFTP, specify Loopback0 as the source interface for communication, and specify the **netconf** user as the SFTP user.

Step 6 Configure SNMP.

Configure SNMPv3 on each device. Create a user named **snmp** and a user group named **snmp**. Set the authentication algorithm to SHA2-512 and encryption algorithm to AES128.

SNMP is used by iMaster NCE-IP to discover and manage devices.

Configure the SNMP version and view.

```
[PE1]snmp-agent sys-info version all
[PE1]snmp-agent mib-view included iso-view iso
[PE1]snmp-agent protocol source-status all-interface
[PE1]snmp-agent group v3 snmp privacy read-view iso-view write-view iso-view notify-view iso-view
```

Set the SNMP version to all and create a view named **iso-view** that corresponds to the iso subtree.

Configure all interfaces to be able to receive and respond to SNMP packets.

Create an SNMP group named **snmp**, and set the read, write, and notification permissions to **iso-view**.

```
[PE1]snmp-agent usm-user v3 snmp
[PE1]snmp-agent usm-user v3 snmp group snmp
[PE1]snmp-agent usm-user v3 snmp authentication-mode sha2-512
Please configure the authentication password (8-255)
Enter Password:
Confirm Password:
[PE1]snmp-agent usm-user v3 snmp privacy-mode aes128
Please configure the privacy password (8-255)
Enter Password:
Confirm Password:
```

Create an SNMPv3 user named **snmp** and add it to the **snmp** group. Set the authentication algorithm to SHA2-512, authentication password to Huawei@123, encryption algorithm to AES128, and authentication password to Huawei@123.

```
[PE1]snmp-agent blacklist ip-block disable
[PE1]snmp-agent protocol source-interface LoopBack0
[PE1]snmp-agent trap enable
```

Disable the SNMP IP address blacklist function, configure Loopback0 as the source interface for communication, and enable the SNMP trap function.

Step 7 Enable LLDP.

Enable LLDP on each device for link discovery between them.

PE1 is used as an example. Repeat the configuration for other devices.

Enable LLDP.

```
[PE1]Lldp enable
Info: Global LLDP is already enabled.
```

3.1.2.3 Device-Side IGP and SR Configurations

Enable SR-MPLS, set the IGP to IS-IS, enable IS-IS to carry link attributes in LSPs, enable IS-IS topology information reporting through BGP-LS, enable IS-IS TE, enable IS-IS to advertise SR-MPLS labels, and configure the same SRGB range (16000 to 17000) on all devices in the SR domain.

Step 1 Configure the IGP.

Ensure that the IS-IS area ID is 49.0001, the IS-IS process ID is 1, all devices are Level-2 devices, and the NET is converted from the Loopback0 IP address (for example, PE2's NET is 49.0001.0010.0000.0002.00). Then enable IS-IS on Loopback0 and interconnection interfaces.

In this case, you need to set **cost-style** to **wide** to support IS-IS extensions.

Enable BFD globally.

```
[PE1]bfd
```

PE1 is used as an example here. Repeat the configuration for other devices.

Description of IS-IS commands:

cost-style wide: The narrow cost type does not support the TE information (such as link bandwidth) required in TE scenarios. Therefore, the wide cost type needs to be configured.

advertise link attributes: This command enables LSPs to carry link attribute TLVs, including interface IP addresses and interface indexes.

bgp-ls enable: This command enables topology information collected by IS-IS to be sent to the controller through BGP-LS. This function only needs to be configured on the RR. That is, only one device in the IGP domain needs to send topology information to the controller through BGP-LS.

traffic-eng: This command enables IS-IS TE, so that link bandwidth information can be sent to the TE module.

set-overload on-startup: This command sets the overload bit, which is used to notify others that the local node cannot forward traffic at this time. The local node is then not

used as a forwarding node during LSP-based path calculation. The command parameters include **on-startup** and **wait-for-bgp**.

metric-delay advertisement enable: This command enables IPv4 delay advertisement. After this function is enabled, IS-IS collects and floods information about the intra-area IPv4 link delay, and BGP-LS reports the information to the controller. The controller can then use the delay information to compute optimal paths on the P2P network.

Configure IS-IS on P1.

```
[P1]isis 1
[P1-isis-1] is-level level-2
[P1-isis-1] cost-style wide
[P1-isis-1] bfd all-interfaces enable
[P1-isis-1] advertise link attributes
[P1-isis-1] bgp-ls enable level-2
[P1-isis-1] network-entity 49.0001.0010.0000.0005.00
[P1-isis-1] is-name P1
[P1-isis-1] traffic-eng level-2
[P1-isis-1] set-overload on-startup
[P1-isis-1] metric-delay advertisement enable level-1-2
```

Configure IS-IS on P2.

```
[P2]isis 1
[P2-isis-1] is-level level-2
[P2-isis-1] cost-style wide
[P2-isis-1] bfd all-interfaces enable
[P2-isis-1] advertise link attributes
[P2-isis-1] bgp-ls enable level-2
[P2-isis-1] network-entity 49.0001.0010.0000.0006.00
[P2-isis-1] is-name P2
[P2-isis-1] traffic-eng level-2
[P2-isis-1] set-overload on-startup
[P2-isis-1] metric-delay advertisement enable level-1-2
```

Configure IS-IS on PE1.

```
[PE1]isis 1
[PE1-isis-1] is-level level-2
[PE1-isis-1] cost-style wide
[PE1-isis-1] bfd all-interfaces enable
[PE1-isis-1] advertise link attributes
[PE1-isis-1] bgp-ls enable level-2
[PE1-isis-1] network-entity 49.0001.0010.0000.0001.00
[PE1-isis-1] is-name PE1
[PE1-isis-1] traffic-eng level-2
[PE1-isis-1] set-overload on-startup
[PE1-isis-1] metric-delay advertisement enable level-1-2
```

Configure IS-IS on PE2.

```
[PE2]isis 1
[PE2-isis-1] is-level level-2
```

```
[PE2-isis-1] cost-style wide
[PE2-isis-1] bfd all-interfaces enable
[PE2-isis-1] advertise link attributes
[PE2-isis-1] bgp-ls enable level-2
[PE2-isis-1] network-entity 49.0001.0010.0000.0002.00
[PE2-isis-1] is-name PE2
[PE2-isis-1] traffic-eng level-2
[PE2-isis-1] set-overload on-startup
[PE2-isis-1] metric-delay advertisement enable level-1-2
```

Configure IS-IS on PE3.

```
[PE3]isis 1
[PE3-isis-1] is-level level-2
[PE3-isis-1] cost-style wide
[PE3-isis-1] bfd all-interfaces enable
[PE3-isis-1] advertise link attributes
[PE3-isis-1] bgp-ls enable level-2
[PE3-isis-1] network-entity 49.0001.0010.0000.0003.00
[PE3-isis-1] is-name P3
[PE3-isis-1] traffic-eng level-2
[PE3-isis-1] set-overload on-startup
[PE3-isis-1] metric-delay advertisement enable level-1-2
```

Configure IS-IS on PE4.

```
[PE4]isis 1
[PE4-isis-1] is-level level-2
[PE4-isis-1] cost-style wide
[PE4-isis-1] bfd all-interfaces enable
[PE4-isis-1] advertise link attributes
[PE4-isis-1] bgp-ls enable level-2
[PE4-isis-1] network-entity 49.0001.0010.0000.0004.00
[PE4-isis-1] is-name P4
[PE4-isis-1] traffic-eng level-2
[PE4-isis-1] set-overload on-startup
[PE4-isis-1] metric-delay advertisement enable level-1-2
```

Enable IS-IS on the interconnection and Loopback0 interfaces of all devices and set the link type to P2P.

#P1

```
[P1]interface GigabitEthernet0/3/0
[P1-GigabitEthernet0/3/0] isis enable 1
[P1-GigabitEthernet0/3/0]isis circuit-type p2p
[P1-GigabitEthernet0/3/0] quit
[P1]interface GigabitEthernet0/3/2
[P1-GigabitEthernet0/3/2] isis enable 1
[P1-GigabitEthernet0/3/2] isis circuit-type p2p
[P1-GigabitEthernet0/3/2] quit
[P1]interface GigabitEthernet0/3/4
[P1-GigabitEthernet0/3/4] isis enable 1
[P1-GigabitEthernet0/3/4] isis circuit-type p2p
[P1-GigabitEthernet0/3/4] quit
```

```
[P1]interface LoopBack0
[P1-LoopBack0] isis enable 1
```

#P2

```
[P2]interface GigabitEthernet0/3/0
[P2-GigabitEthernet0/3/0] isis enable 1
[P2-GigabitEthernet0/3/0] isis circuit-type p2p
[P2-GigabitEthernet0/3/0] quit
[P2]interface GigabitEthernet0/3/2
[P2-GigabitEthernet0/3/2] isis enable 1
[P2-GigabitEthernet0/3/2] isis circuit-type p2p
[P2-GigabitEthernet0/3/2] quit
[P2]interface GigabitEthernet0/3/4
[P2-GigabitEthernet0/3/4] isis enable 1
[P2-GigabitEthernet0/3/4] isis circuit-type p2p
[P2-GigabitEthernet0/3/4] quit
[P2]interface LoopBack0
[P2-LoopBack0] isis enable 1
```

#PE1

```
[PE1]interface GigabitEthernet0/3/0
[PE1-GigabitEthernet0/3/0] isis enable 1
[PE1-GigabitEthernet0/3/0] isis circuit-type p2p
[PE1-GigabitEthernet0/3/0] quit
[PE1]interface GigabitEthernet0/3/1
[PE1-GigabitEthernet0/3/1] isis enable 1
[PE1-GigabitEthernet0/3/1] isis circuit-type p2p
[PE1-GigabitEthernet0/3/1] quit
[PE1]interface LoopBack0
[PE1-LoopBack0] isis enable 1
```

#PE2

```
[PE2]interface GigabitEthernet0/3/0
[PE2-GigabitEthernet0/3/0] isis enable 1
[PE2-GigabitEthernet0/3/0] isis circuit-type p2p
[PE2-GigabitEthernet0/3/0] quit
[PE2]interface GigabitEthernet0/3/1
[PE2-GigabitEthernet0/3/1] isis enable 1
[PE2-GigabitEthernet0/3/1] isis circuit-type p2p
[PE2-GigabitEthernet0/3/1] quit
[PE2]interface LoopBack0
[PE2-LoopBack0] isis enable 1
```

#PE3

```
[PE3]interface GigabitEthernet0/3/0
[PE3-GigabitEthernet0/3/0] isis enable 1
[PE3-GigabitEthernet0/3/0] isis circuit-type p2p
[PE3-GigabitEthernet0/3/0] quit
[PE3]interface GigabitEthernet0/3/1
[PE3-GigabitEthernet0/3/1] isis enable 1
```

```
[PE3-GigabitEthernet0/3/1] isis circuit-type p2p
[PE3-GigabitEthernet0/3/1] quit
[PE3]interface LoopBack0
[PE3-LoopBack0] isis enable 1
```

#PE4

```
[PE4]interface GigabitEthernet0/3/0
[PE4-GigabitEthernet0/3/0] isis enable 1
[PE4-GigabitEthernet0/3/0] isis circuit-type p2p
[PE4-GigabitEthernet0/3/0] quit
[PE4]interface GigabitEthernet0/3/1
[PE4-GigabitEthernet0/3/1] isis enable 1
[PE4-GigabitEthernet0/3/1] isis circuit-type p2p
[PE4-GigabitEthernet0/3/1] quit
[PE4]interface LoopBack0
[PE4-LoopBack0] isis enable 1
```

Step 2 Configure SR-MPLS.

Enable SR-MPLS globally, enable IS-IS to support SR-MPLS, set the SRGB range to 16000–17000, and enable LFA and TI-LFA.

In the PE SR-MPLS view, enable the function to report TE Policy status through BGP-LS, so that iMaster NCE-IP can monitor the tunnel status.

Configure a prefix SID for the Loopback0 IP address, and use the prefix SID as the node SID to identify the node. Specify a relative label value as the prefix SID, and use the *X* value in Loopback0 IP address 1.0.0.*X* as the offset value. For example, if the Loopback0 IP address on PE1 is 1.0.0.1, then the offset value is 1.

Enable SR globally.

```
[PE1]segment-routing
```

PE1 is used as an example.

Enable the function to report SR-MPLS TE Policy information through BGP-LS.

```
[PE1]segment-routing
[PE1-segment-routing] sr-te-policy bgp-ls enable
```

PE1 is used as an example. Repeat the configuration for other PEs.

Step 3 Enable SR-MPLS and configure the SRGB, LFA, and TI-LFA for IS-IS.

#P1

```
[P1]isis 1
[P1-isis-1] segment-routing mpls
[P1-isis-1] segment-routing global-block 16000 17000
[P1-isis-1]frr
[P1-isis-1-frr] loop-free-alternate level-2
[P1-isis-1-frr] ti-lfa level-2
```

P1 is used as an example here. Repeat the configuration for other devices.

Assign prefix SIDs to loopback interfaces.

#P1

```
[P1]interface LoopBack0
[P1-LoopBack0] isis prefix-sid index 5
```

#P2

```
[P2]interface LoopBack0
[P2-LoopBack0] isis prefix-sid index 6
```

#PE1

```
[PE1]interface LoopBack0
[PE1-LoopBack0] isis prefix-sid index 1
```

#PE2

```
[PE2]interface LoopBack0
[PE2-LoopBack0] isis prefix-sid index 2
```

#PE3

```
[PE3]interface LoopBack0
[PE3-LoopBack0] isis prefix-sid index 3
```

#PE4

```
[PE4]interface LoopBack0
[PE4-LoopBack0] isis prefix-sid index 4
```

3.1.2.4 Device-Side MPLS Configurations

Enable MPLS and MPLS TE globally, use the Loopback0 IP address as the LSR ID, and enable MPLS and MPLS TE on interconnection interfaces.

Configure devices as PCE clients and iMaster NCE-IP as the PCE server. Delegate TE tunnels to the PCE server. Huawei's CloudWAN solution does not use PCEP to deliver TE tunnel configurations and only uses PCEP for tunnel status monitoring.

Step 1 Enable MPLS and MPLS TE and configure LSR IDs on all nodes in the SR domain.

#PE1

```
[PE1]mpls lsr-id 1.0.0.1
[PE1]mpls
[PE1-mpls]mpls te
```

PE1 is used as an example. Configure LSR IDs for other devices as planned.

Step 2 Enable MPLS and MPLS TE on interconnection interfaces.

#P1

```
[P1]interface GigabitEthernet0/3/0
[P1-GigabitEthernet0/3/0] mpls
[P1-GigabitEthernet0/3/0]mpls te
[P1-GigabitEthernet0/3/0] quit
[P1]interface GigabitEthernet0/3/2
[P1-GigabitEthernet0/3/2] mpls
[P1-GigabitEthernet0/3/2] mpls te
[P1-GigabitEthernet0/3/2] quit
[P1]interface GigabitEthernet0/3/4
[P1-GigabitEthernet0/3/4] mpls
[P1-GigabitEthernet0/3/4] mpls te
[P1-GigabitEthernet0/3/4] quit
```

#P2

```
[P2]interface GigabitEthernet0/3/0
[P2-GigabitEthernet0/3/0] mpls
[P2-GigabitEthernet0/3/0] mpls te
[P2-GigabitEthernet0/3/0] quit
[P2]interface GigabitEthernet0/3/2
[P2-GigabitEthernet0/3/2] mpls
[P2-GigabitEthernet0/3/2] mpls te
[P2-GigabitEthernet0/3/2] quit
[P2]interface GigabitEthernet0/3/4
[P2-GigabitEthernet0/3/4] mpls
[P2-GigabitEthernet0/3/4] mpls te
[P2-GigabitEthernet0/3/4] quit
```

#PE1

```
[PE1]interface GigabitEthernet0/3/0
[PE1-GigabitEthernet0/3/0] mpls
[PE1-GigabitEthernet0/3/0] mpls te
[PE1-GigabitEthernet0/3/0] quit
[PE1]interface GigabitEthernet0/3/1
[PE1-GigabitEthernet0/3/1] mpls
[PE1-GigabitEthernet0/3/1] mpls te
[PE1-GigabitEthernet0/3/1] quit
```

#PE2

```
[PE2]interface GigabitEthernet0/3/0
[PE2-GigabitEthernet0/3/0] mpls
[PE2-GigabitEthernet0/3/0] mpls te
[PE2-GigabitEthernet0/3/0] quit
[PE2]interface GigabitEthernet0/3/1
[PE2-GigabitEthernet0/3/1] mpls
[PE2-GigabitEthernet0/3/1] mpls te
[PE2-GigabitEthernet0/3/1] quit
```

#PE3

```
[PE3]interface GigabitEthernet0/3/0
[PE3-GigabitEthernet0/3/0] mpls
[PE3-GigabitEthernet0/3/0] mpls te
[PE3-GigabitEthernet0/3/0] quit
[PE3]interface GigabitEthernet0/3/1
[PE3-GigabitEthernet0/3/1] mpls3
[PE3-GigabitEthernet0/3/1] mpls te
[PE3-GigabitEthernet0/3/1] quit
```

#PE4

```
[PE4]interface GigabitEthernet0/3/0
[PE4-GigabitEthernet0/3/0] mpls
[PE4-GigabitEthernet0/3/0] mpls te
[PE4-GigabitEthernet0/3/0] quit
[PE4]interface GigabitEthernet0/3/1
[PE4-GigabitEthernet0/3/1] mpls
[PE4-GigabitEthernet0/3/1] mpls te
[PE4-GigabitEthernet0/3/1] quit
```

Step 3 Configure devices as PCE clients.

Configure all devices as PCE clients and specify iMaster NCE-IP as the PCE server.

#PE1

```
[PE1]pce-client
[PE1-pce-client] connect-server 172.21.17.102
[PE1-pce-client-connect-172.21.17.102] capability segment-routing
```

Here, PE1 is used as an example to describe how to specify iMaster NCE-IP as the PCE server. Repeat the configuration for other devices.

Delegate TE tunnels to the PCE server.

```
[PE1]mpls
[PE1-mpls] mpls te pce delegate
```

PE1 is used as an example here. Repeat the configuration for other devices.

3.1.2.5 Device-Side BGP Configurations

Configure P1 and P2 as RRs and establish BGP VPNv4 peer relationships between PEs and RRs, so that CEs can communicate with each other through L3VPNv4.

Configure BGP-LS peer relationships between RRs and iMaster NCE-IP, so that link information can be reported to iMaster NCE-IP.

To enable iMaster NCE-IP to monitor the path status of SR-MPLS Policies, establish BGP-LS peer relationships between PEs and RRs and use RRs to report the path status of SR-MPLS Policies to iMaster NCE-IP.

To enable iMaster NCE-IP to deliver SR Policy configurations to PEs through BGP SR Policy routes, establish BGP SR Policy peer relationships between iMaster NCE-IP and RRs, and establish BGP SR Policy peer relationships between all PEs and RRs. The RRs then reflect SR Policy routes received from the controller to PEs.

Step 1 Establish VPNv4 peer relationships.

Establish VPNv4 peer relationships between PEs and RRs. Use the Loopback0 IP address as the BGP router ID.

Configure PE1.

```
[PE1]bgp 65001
[PE1-bgp] router-id 1.0.0.1
[PE1-bgp] undo default ipv4-unicast
[PE1-bgp] peer 1.0.0.5 as-number 65001
[PE1-bgp] peer 1.0.0.5 connect-interface LoopBack0
[PE1-bgp] peer 1.0.0.6 as-number 65001
[PE1-bgp] peer 1.0.0.6 connect-interface LoopBack0
[PE1-bgp] ipv4-family vpnv4
[PE1-bgp-af-vpnv4] peer 1.0.0.5 enable
[PE1-bgp-af-vpnv4] peer 1.0.0.5 advertise-community
[PE1-bgp-af-vpnv4] peer 1.0.0.6 enable
[PE1-bgp-af-vpnv4] peer 1.0.0.6 advertise-community
```

PE1 is used as an example. The configurations of other PEs are similar to the configuration of PE1.

Configure P1.

```
[P1]bgp 65001
[P1-bgp] router-id 1.0.0.5
[P1-bgp] undo default ipv4-unicast
[P1-bgp] group RR internal
[P1-bgp] peer RR connect-interface LoopBack0
[P1-bgp] peer 1.0.0.1 as-number 65001
[P1-bgp] peer 1.0.0.1 group RR
[P1-bgp] peer 1.0.0.2 as-number 65001
[P1-bgp] peer 1.0.0.2 group RR
[P1-bgp] peer 1.0.0.3 as-number 65001
[P1-bgp] peer 1.0.0.3 group RR
[P1-bgp] peer 1.0.0.4 as-number 65001
[P1-bgp] peer 1.0.0.4 group RR
[P1-bgp] ipv4-family vpnv4
[P1-bgp-af-vpnv4] undo policy vpn-target
[P1-bgp-af-vpnv4] peer 1.0.0.1 enable
[P1-bgp-af-vpnv4] peer 1.0.0.1 reflect-client
[P1-bgp-af-vpnv4] peer 1.0.0.1 advertise-community
[P1-bgp-af-vpnv4] peer 1.0.0.2 enable
[P1-bgp-af-vpnv4] peer 1.0.0.2 reflect-client
[P1-bgp-af-vpnv4] peer 1.0.0.2 advertise-community
[P1-bgp-af-vpnv4] peer 1.0.0.3 enable
[P1-bgp-af-vpnv4] peer 1.0.0.3 reflect-client
[P1-bgp-af-vpnv4] peer 1.0.0.3 advertise-community
[P1-bgp-af-vpnv4] peer 1.0.0.4 enable
[P1-bgp-af-vpnv4] peer 1.0.0.4 reflect-client
[P1-bgp-af-vpnv4] peer 1.0.0.4 advertise-community
```

P1 is used as an example. Establish a VPNv4 peer relationship between each PE and the RR. The configuration of P2 is the same as that of P1, except each P must have a unique router ID.

Step 2 Establishment BGP-LS peer relationships.

Establish a BGP-LS peer relationship between each RR and iMaster NCE-IP for redundancy protection.

Establish BGP-LS peer relationships between PEs and RRs, so that RRs can report SR-MPLS Policy path status.

This section describes only device-side configurations. Controller-side configurations are described in the following sections.

#P1

```
[P1]bgp 65001
[P1-bgp] peer 172.21.17.102 as-number 65001
[P1-bgp] peer 172.21.17.102 group RR
[P1-bgp] link-state-family unicast
[P1-bgp-af-ls] domain identifier 1.0.0.56
[P1-bgp-af-ls] peer 172.21.17.102 enable
[P1-bgp-af-ls] peer 172.21.17.102 reflect-client
[P1-bgp-af-ls] peer 1.0.0.1 enable
[P1-bgp-af-ls] peer 1.0.0.1 reflect-client
[P1-bgp-af-ls] peer 1.0.0.2 enable
[P1-bgp-af-ls] peer 1.0.0.2 reflect-client
[P1-bgp-af-ls] peer 1.0.0.3 enable
[P1-bgp-af-ls] peer 1.0.0.3 reflect-client
[P1-bgp-af-ls] peer 1.0.0.4 enable
[P1-bgp-af-ls] peer 1.0.0.4 reflect-client
```

#P2

```
[P2]bgp 65001
[P2-bgp] peer 172.21.17.102 as-number 65001
[P2-bgp] peer 172.21.17.102 group RR
[P2-bgp] link-state-family unicast
[P2-bgp-af-ls] domain identifier 1.0.0.56
[P2-bgp-af-ls] peer 172.21.17.102 enable
[P2-bgp-af-ls] peer 172.21.17.102 reflect-client
[P2-bgp-af-ls] peer 1.0.0.1 enable
[P2-bgp-af-ls] peer 1.0.0.1 reflect-client
[P2-bgp-af-ls] peer 1.0.0.2 enable
[P2-bgp-af-ls] peer 1.0.0.2 reflect-client
[P2-bgp-af-ls] peer 1.0.0.3 enable
[P2-bgp-af-ls] peer 1.0.0.3 reflect-client
[P2-bgp-af-ls] peer 1.0.0.4 enable
[P2-bgp-af-ls] peer 1.0.0.4 reflect-client
```

The two RRs must be configured with the same domain identifier, so that iMaster NCE-IP centrally computes link information received from the two RRs.

#PE1

```
[PE1]bgp 65001
[PE1-bgp] link-state-family unicast
[PE1-bgp-af-ls] peer 1.0.0.5 enable
[PE1-bgp-af-ls] peer 1.0.0.6 enable
```

PE1 is used as an example. The configurations of other PEs are similar to the configuration of PE1.

Check the BGP-LS peer status on PE1 and PE2.

```
[P2]display bgp link-state unicast peer
BGP local router ID : 1.0.0.6
Local AS number : 65001
Total number of peers: 14                 Peers in established state : 8

 Peer            V    AS  MsgRcvd  MsgSent  OutQ  Up/Down       State      PrefRcv
 1.0.0.1         4 65001        63      309     0  00:08:51 Established       91
 1.0.0.2         4 65001      1201     1643     0  16:30:46 Established       92
 1.0.0.3         4 65001        54      196     0  00:04:04 Established       90
 1.0.0.4         4 65001        58      199     0  00:04:04 Established       92
 172.21.17.102   4 65001        34      314     0  00:24:02 Active             0
```

Step 3 Establish BGP SR Policy peer relationships.

Establish a BGP SR Policy peer relationship between each PE and each RR (P1 or P2) and a BGP SR Policy peer relationship between each RR and iMaster NCE-IP, so that iMaster NCE-IP can deliver SR Policy routes to PEs.

#PE1

```
[PE1]bgp 65001
[PE1-bgp] ipv4-family sr-policy
[P4-bgp-af-ipv4-srpolicy] peer 1.0.0.5 enable
[P4-bgp-af-ipv4-srpolicy] peer 1.0.0.6 enable
```

PE1 is used as an example. The configurations of other PEs are similar to the configuration of PE1.

#P1

```
[P1]bgp 65001
[P1-bgp] ipv4-family sr-policy
[P1-bgp-af-ipv4-srpolicy] undo router-id filter
[P1-bgp-af-ipv4-srpolicy] peer 1.0.0.1 enable
[P1-bgp-af-ipv4-srpolicy] peer 1.0.0.1 reflect-client
[P1-bgp-af-ipv4-srpolicy] peer 1.0.0.1 advertise-ext-community
[P1-bgp-af-ipv4-srpolicy] peer 1.0.0.2 enable
[P1-bgp-af-ipv4-srpolicy] peer 1.0.0.2 reflect-client
[P1-bgp-af-ipv4-srpolicy] peer 1.0.0.2 advertise-ext-community
[P1-bgp-af-ipv4-srpolicy] peer 1.0.0.3 enable
[P1-bgp-af-ipv4-srpolicy] peer 1.0.0.3 reflect-client
[P1-bgp-af-ipv4-srpolicy] peer 1.0.0.3 advertise-ext-community
[P1-bgp-af-ipv4-srpolicy] peer 1.0.0.4 enable
[P1-bgp-af-ipv4-srpolicy] peer 1.0.0.4 reflect-client
[P1-bgp-af-ipv4-srpolicy] peer 1.0.0.4 advertise-ext-community
```

```
[P1-bgp-af-ipv4-srpolicy] peer 172.21.17.102 enable
[P1-bgp-af-ipv4-srpolicy] peer 172.21.17.102 reflect-client
```

P1 is used as an example. Establish an SR Policy peer relationship between P1 and iMaster NCE-IP, configure PEs as the RR-clients of P1, and enable the function to send extended community attributes to RR clients on P1. In the configurations delivered by the controller, the tunnel color is carried in the extended community attribute. Therefore, this configuration is mandatory.

Note that router ID filtering must be disabled on each RR. In the scenario where RRs are used to push inbound traffic optimization information, all PEs receive traffic optimization policy routes sent by the RRs. To prevent a device from receiving a large number of traffic optimization policy routes that are irrelevant to the device, each traffic optimization policy route carries an extended community attribute in the format of an IP address to identify the node that needs to receive the route. PEs can then filter out unwanted routes based on the extended community attribute. However, RRs need to receive all traffic optimization policy routes from the controller. Therefore, you need to disable this feature on RRs. Disabling this feature is similar to disabling VPNv4 route RT check on RRs.

On RRs, check BGP SR Policy peer relationships with PEs.

```
[P2]display bgp sr-policy peer

BGP local router ID : 1.0.0.6
Local AS number : 65001
Total number of peers: 5                 Peers in established state : 4

 Peer            V    AS     MsgRcvd  MsgSent  OutQ  Up/Down     State        PrefRcv
 1.0.0.1         4    65001  19       22       0     00:13:37    Established  0
 1.0.0.2         4    65001  50401    51119    0     0727h18m    Established  0
 1.0.0.3         4    65001  48152    48653    0     0694h12m    Established  0
 1.0.0.4         4    65001  19       19       0     00:13:01    Established  0
 172.21.17.102   4    65001  52228    52316    0     0745h18m    Active       0
```

```
[P1]display bgp sr-policy peer

BGP local router ID : 1.0.0.5
Local AS number : 65001
Total number of peers: 5                 Peers in established state : 5

 Peer            V    AS     MsgRcvd  MsgSent  OutQ  Up/Down     State        PrefRcv
 1.0.0.1         4    65001  19       21       0     00:13:22    Established  0
 1.0.0.2         4    65001  19       21       0     00:13:23    Established  0
 1.0.0.3         4    65001  20       21       0     00:13:22    Established  0
 1.0.0.4         4    65001  19       20       0     00:13:20    Established  0
 172.21.17.102   4    65001  18       58       0     00:13:20    Active       0
```

A BGP SR Policy peer relationship is established between each RR and PE. After the BGP configuration is complete on iMaster NCE-IP (172.21.17.102), the BGP SR Policy peer relationship with iMaster NCE-IP is also established.

3.1.2.6 Controller-Side Basic Configurations

On the controller, configure routes to devices' Loopback0 interfaces. The controller can then manage devices based on device Loopback0 addresses and establish BGP-LS and SR Policy peer relationships as well as PCEP sessions with RRs.

3.1.2.6.1 Route Configuration

Log in to the iMaster NCE-IP management plane and configure IPv4 routes to devices.

On the home page, choose **maintain** > **Network Configuration** > **Configure routing** from the main menu.

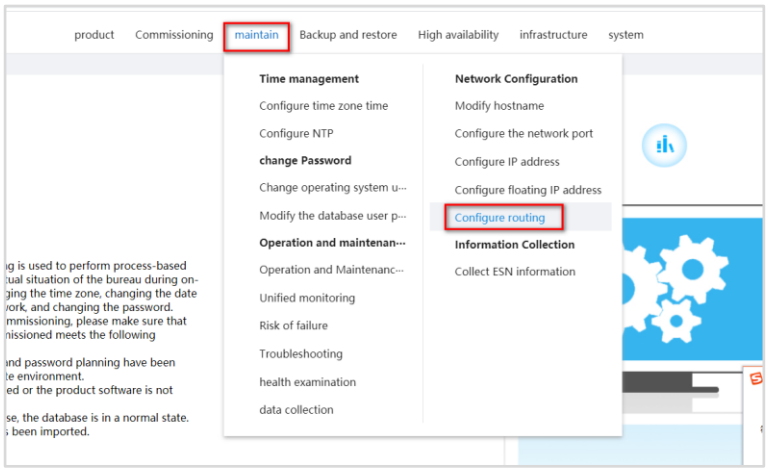

On the **Configure routing** page, select **Analyzer_01**, **Controller_01**, and **NMS_01**, and click **Inquire**.

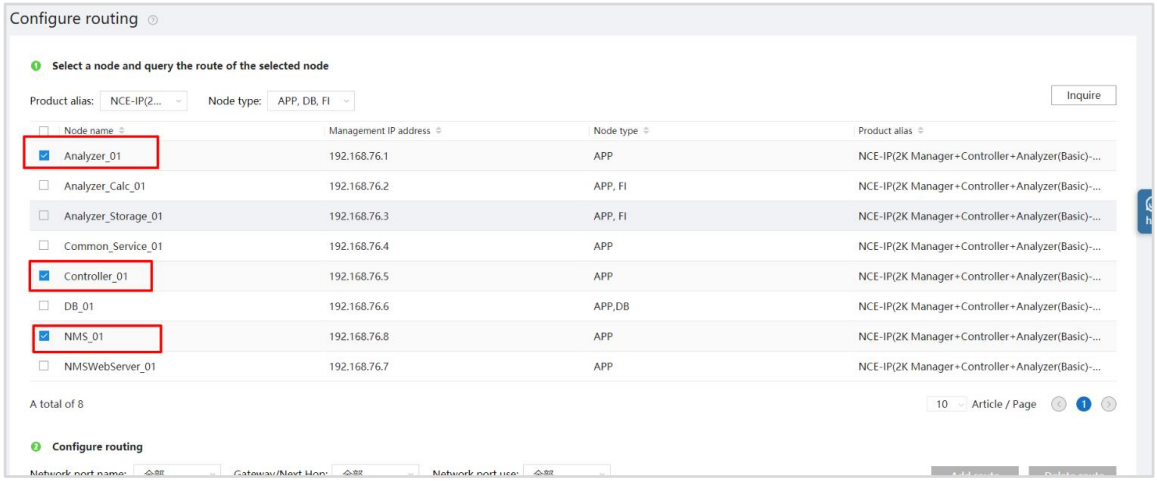

After the query is complete, click **Add route** and select the preceding three nodes.

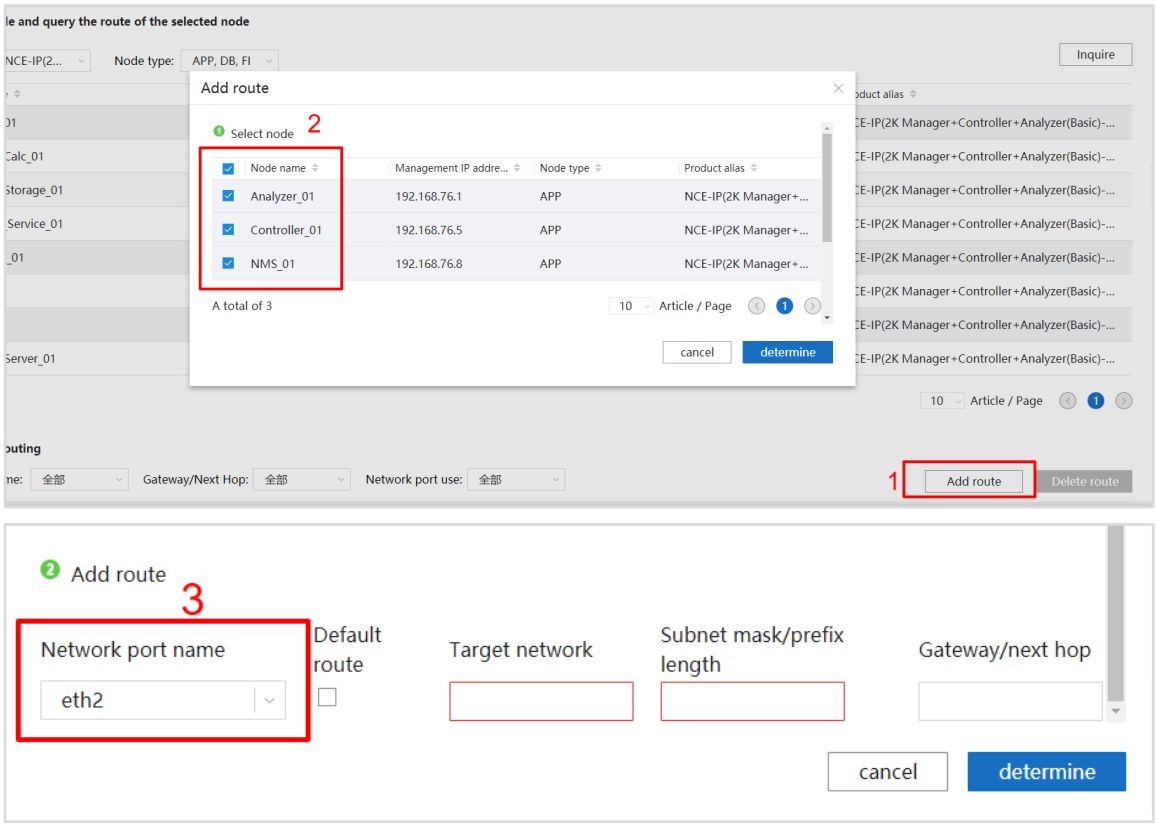

In the **Add route** area, add routes to NEs according to the following table.

Table 3-2 Route configuration table

Destination Device	Destination Network	Subnet Mask/Prefix Length	Gateway/Next Hop
P1	1.0.0.5	255.255.255.255	172.21.17.5
P2	1.0.0.6	255.255.255.255	172.21.17.6
PE1	1.0.0.1	255.255.255.255	172.21.17.1
PE2	1.0.0.2	255.255.255.255	172.21.17.2
PE3	1.0.0.3	255.255.255.255	172.21.17.3
PE4	1.0.0.4	255.255.255.255	172.21.17.4

3.1.2.6.2 NE Addition for Management

Before using iMaster NCE-IP to manage NEs, you need to configure communication templates (SNMP, STelnet, and NETCONF templates) on iMaster NCE-IP. Ensure that the parameters in the templates on iMaster NCE-IP are the same as those configured on NEs.

Step 1 Create SNMP and STelnet parameter templates for to-be-managed NEs.

Log in to the O&M plane of iMaster NCE-IP and open the **Network Management** app on the home page.

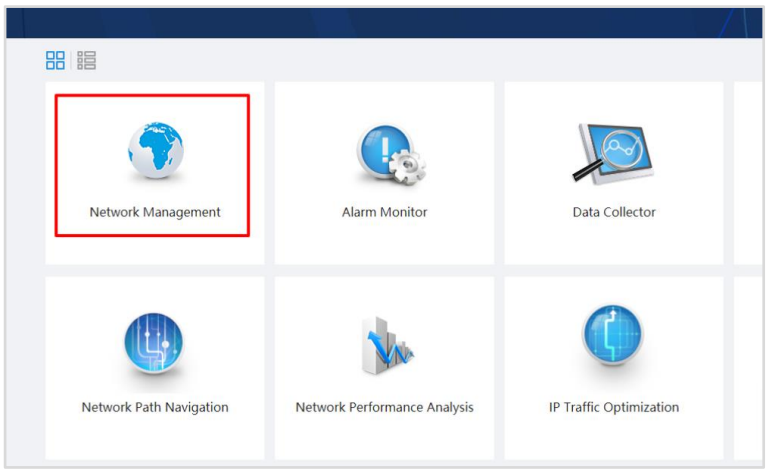

Choose **System** > **NE Communication Parameters** from the main menu.

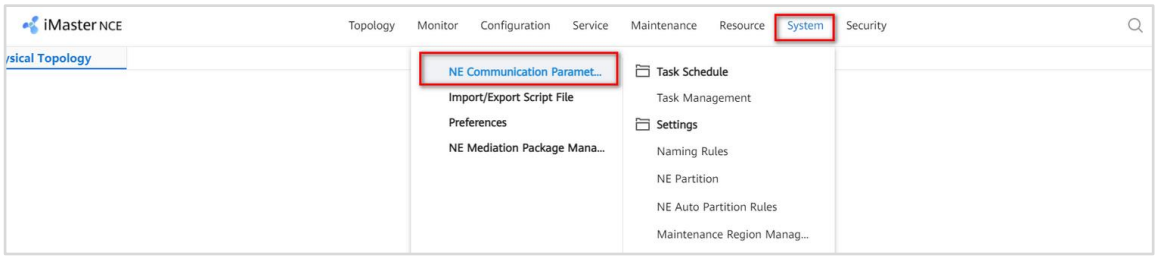

On the page that is displayed, click **Default Access Protocol Parameters**.

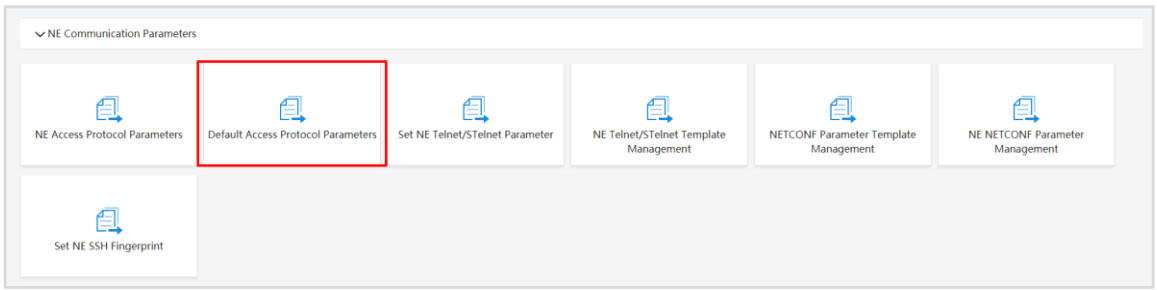

On the **Default Access Protocol Parameters** page, click **Create** to switch to the **SNMP Template** page.

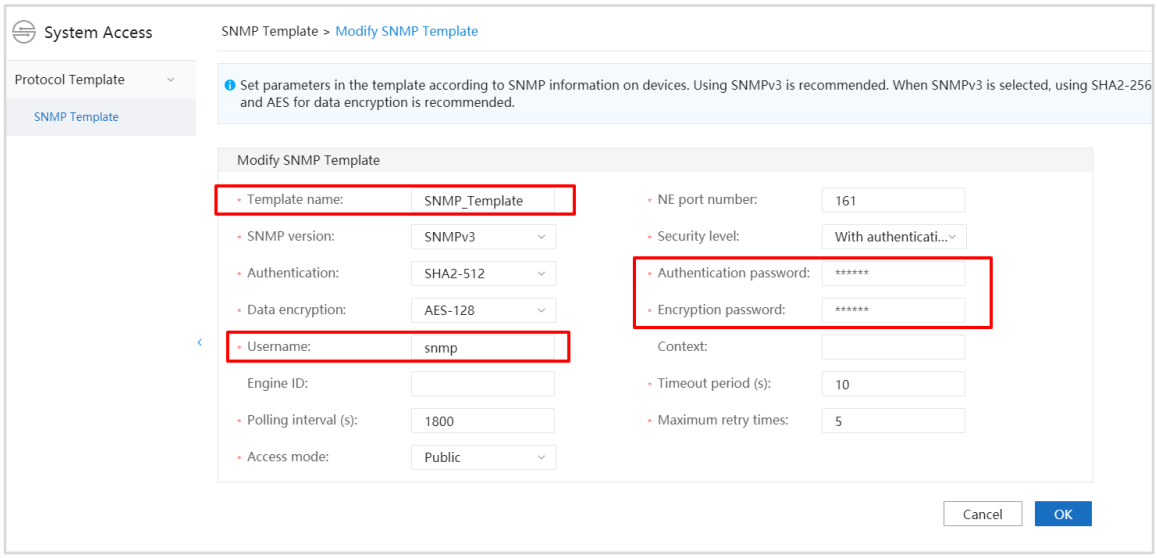

Add SNMP parameters to the template according to those on devices.

On the **NE Communication Parameters** page, click **NE Telnet/STelnet Template Management**.

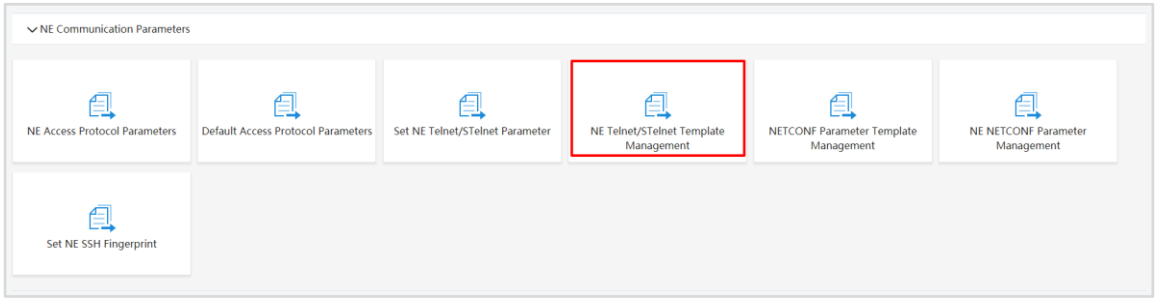

In the dialog box that is displayed, click **Create** at the bottom. In the **Create Template** dialog box, set **Protocol Name** to **STelnet**, **Template Type** to **Private**, and **Template name** to **STelnet_Template**.

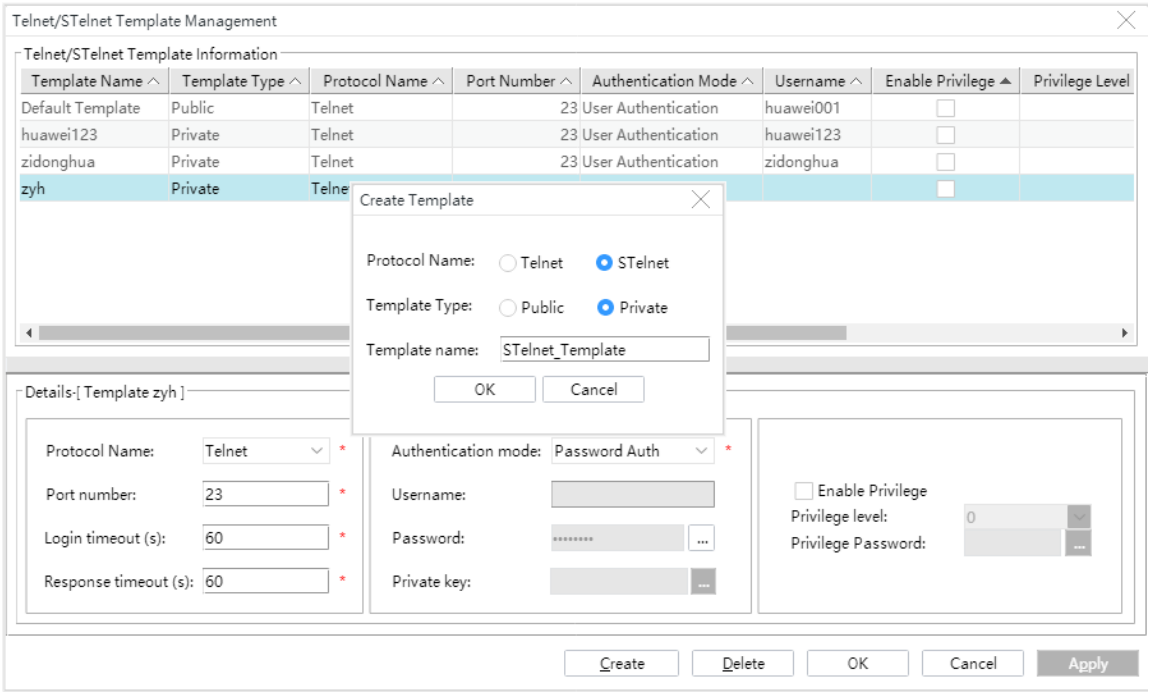

Click the created template, set parameters based on parameter settings on the device, and click **OK** to save the settings.

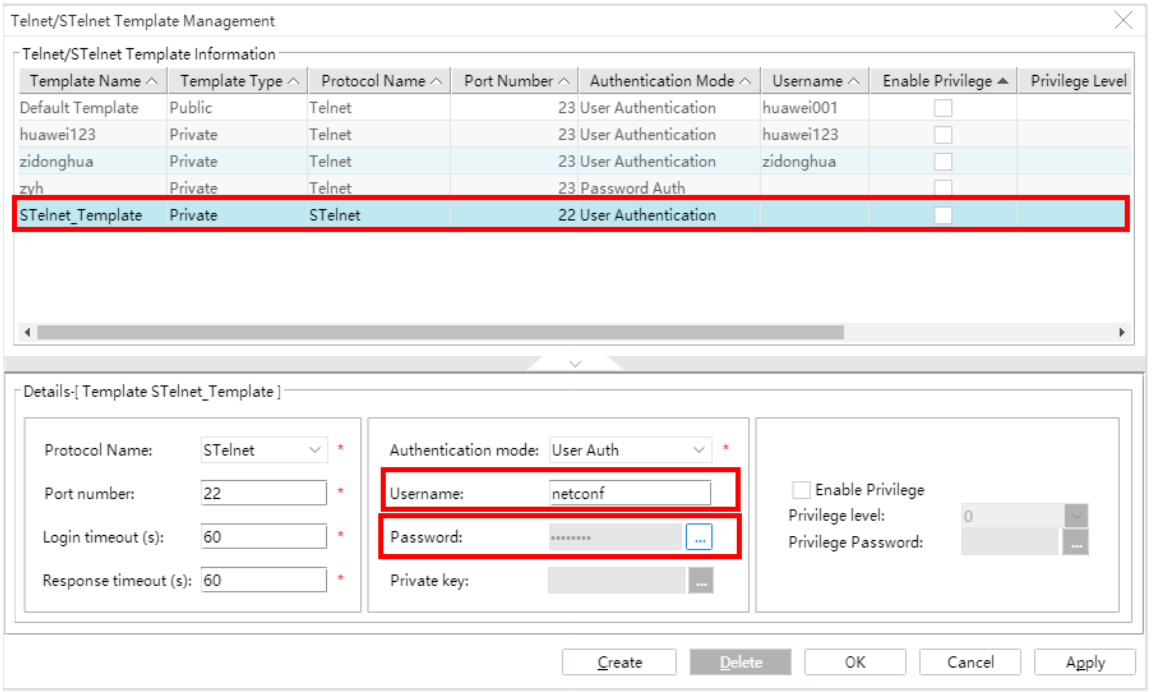

Step 2 Create a NETCONF template for gateway management.

On the **NE Communication Parameters** page, click **NE NETCONF Parameter Template Management** and then click **Create**.

"NETCONF_Template".

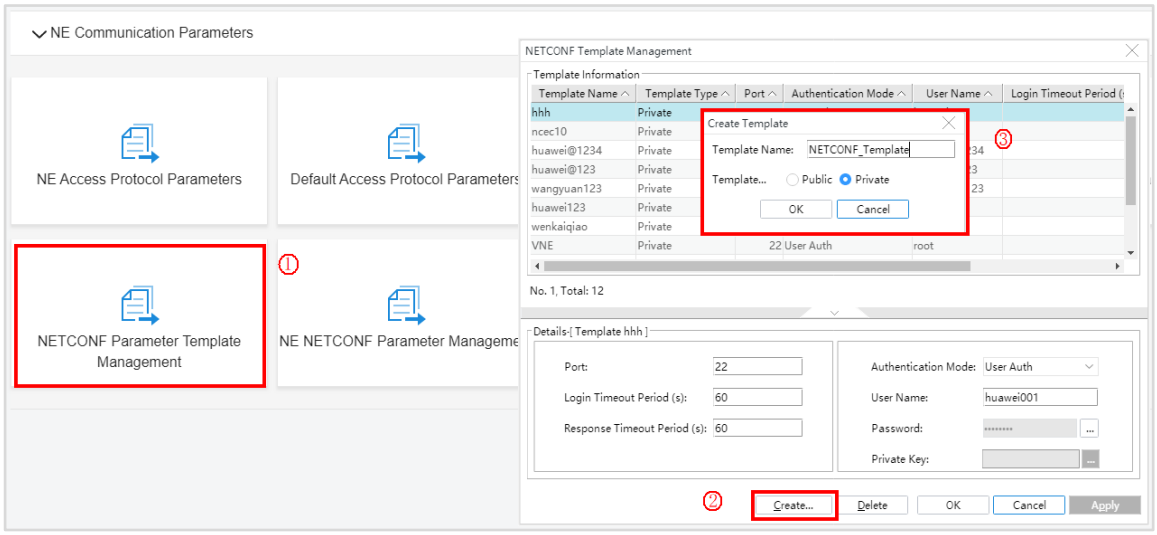

Click the created template, set parameters as planned, and click **OK**.

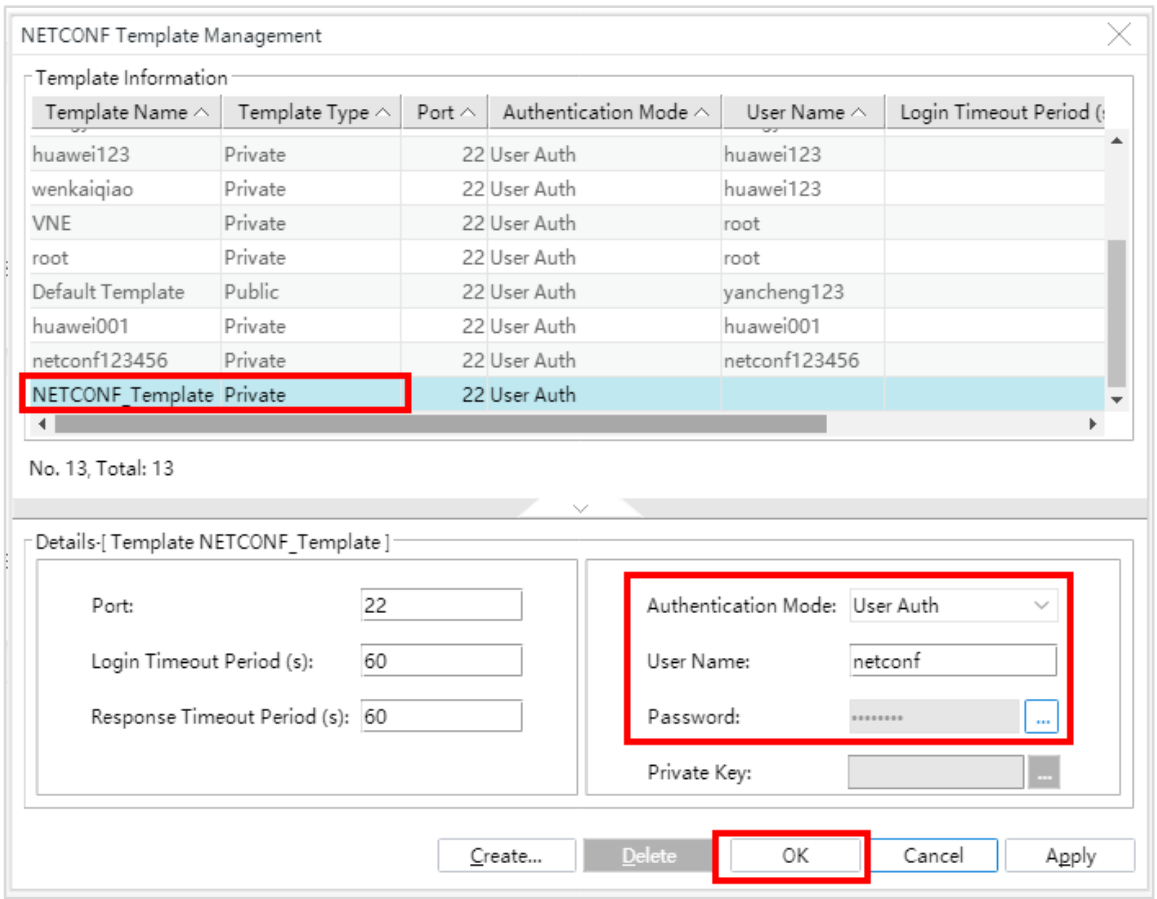

Step 3 Add NEs for management.

Use the Network Management app to add network devices for management.

Open the Network Management app and choose **Topology** > **Discovery** > **NE** from the main menu.

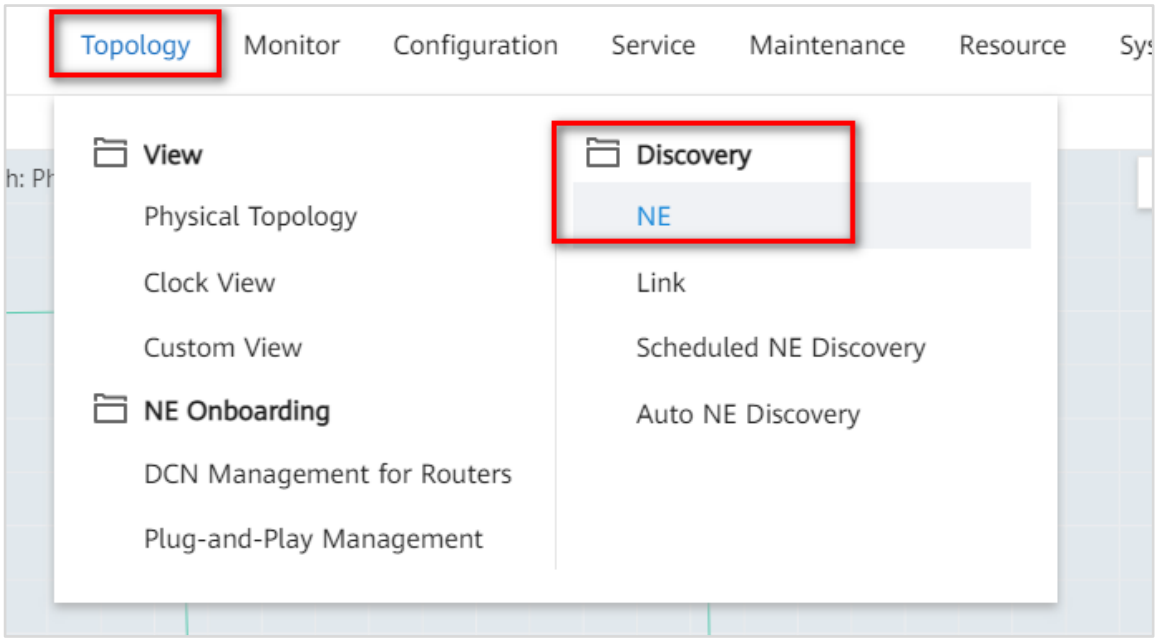

In the **NE Discovery** dialog box, click **Default SNMP Parameters** and select the configured SNMP template.

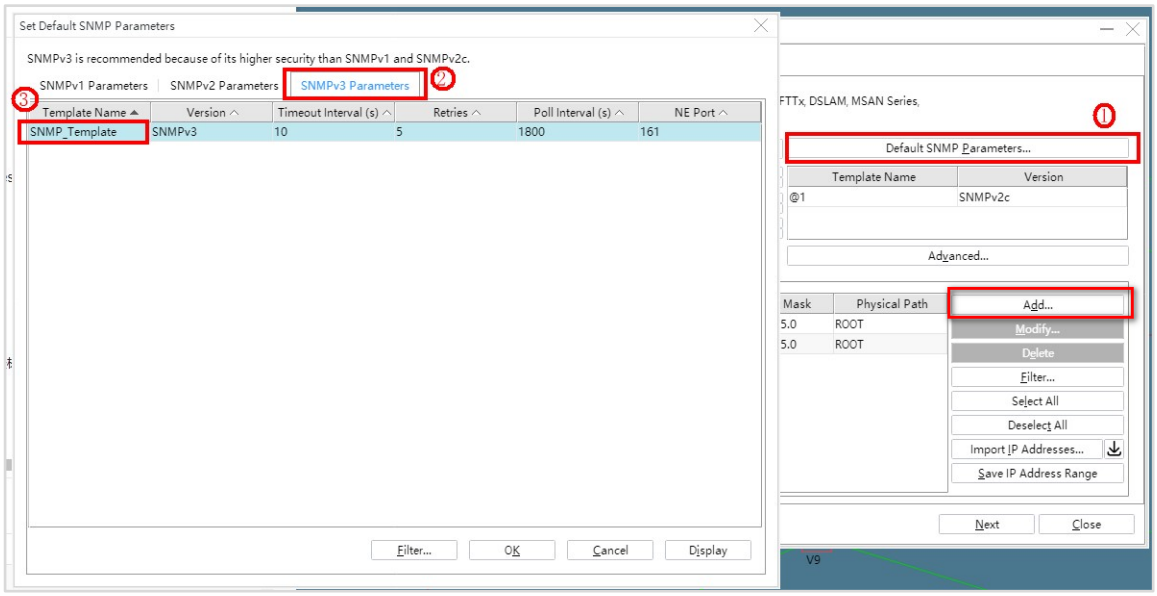

Return to the **NE Discovery** dialog box, click **Advanced**, and select the configured Telnet/STelnet parameter template and NETCONF parameter template.

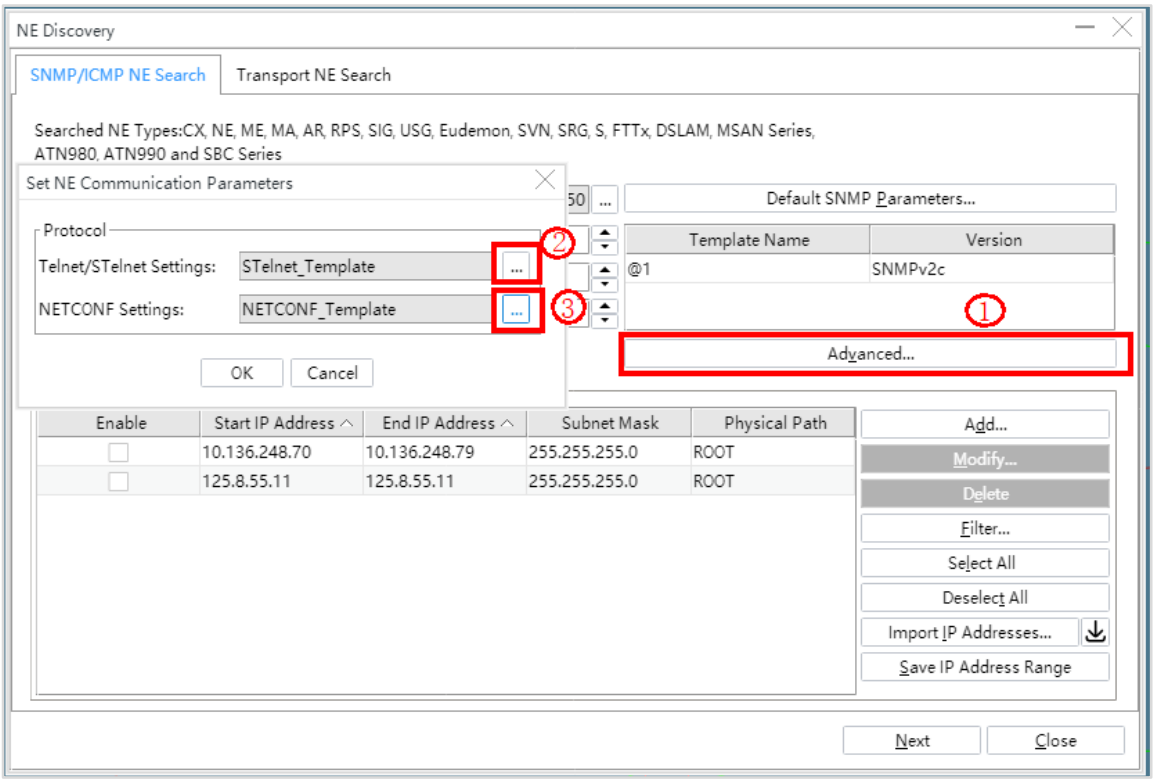

Return to the **NE Discovery** dialog box, click **Add**, and enter the start and end IP addresses of the network segment for management.

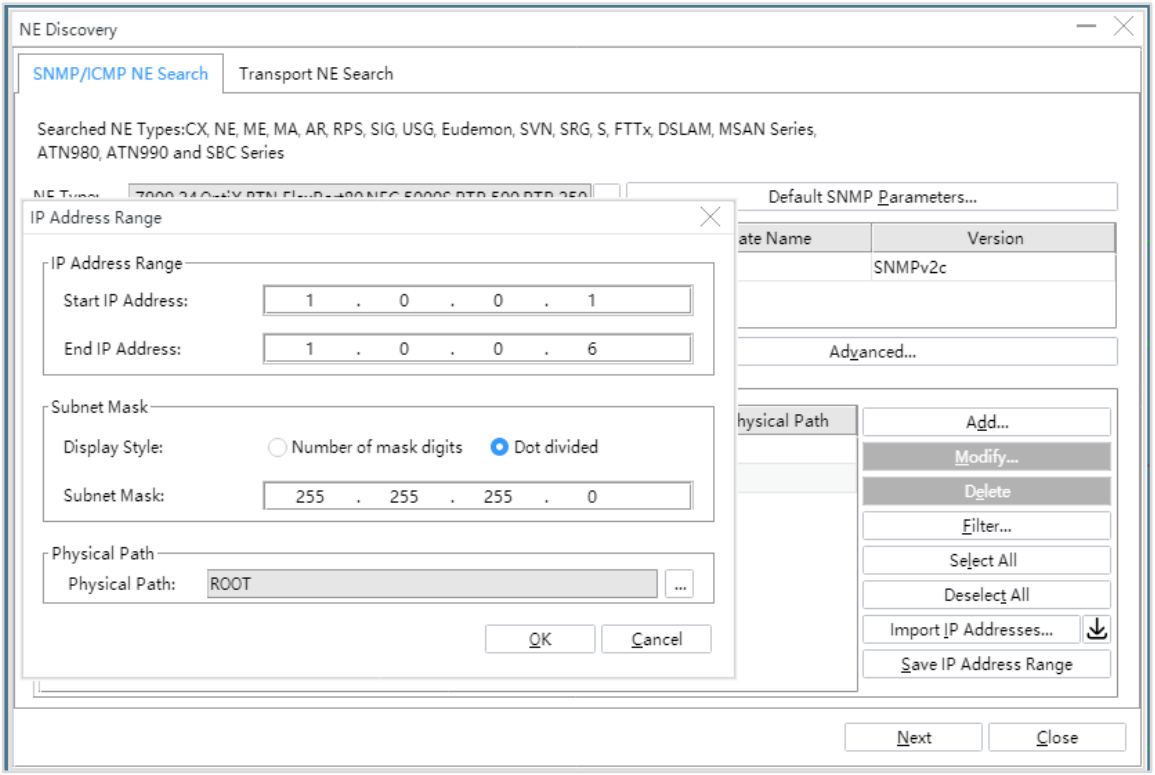

Click **OK**. Then select the created IP address range and click **Next**. In the dialog box that is displayed, iMaster NCE-IP automatically scans for network devices.

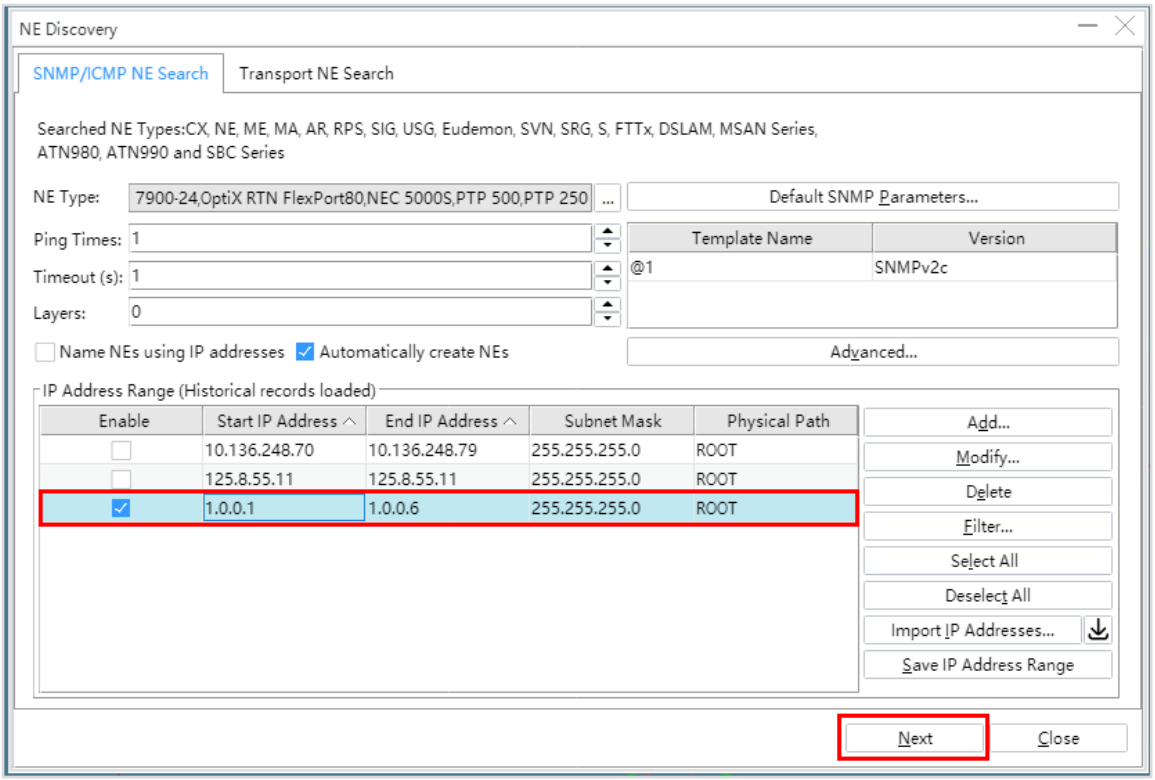

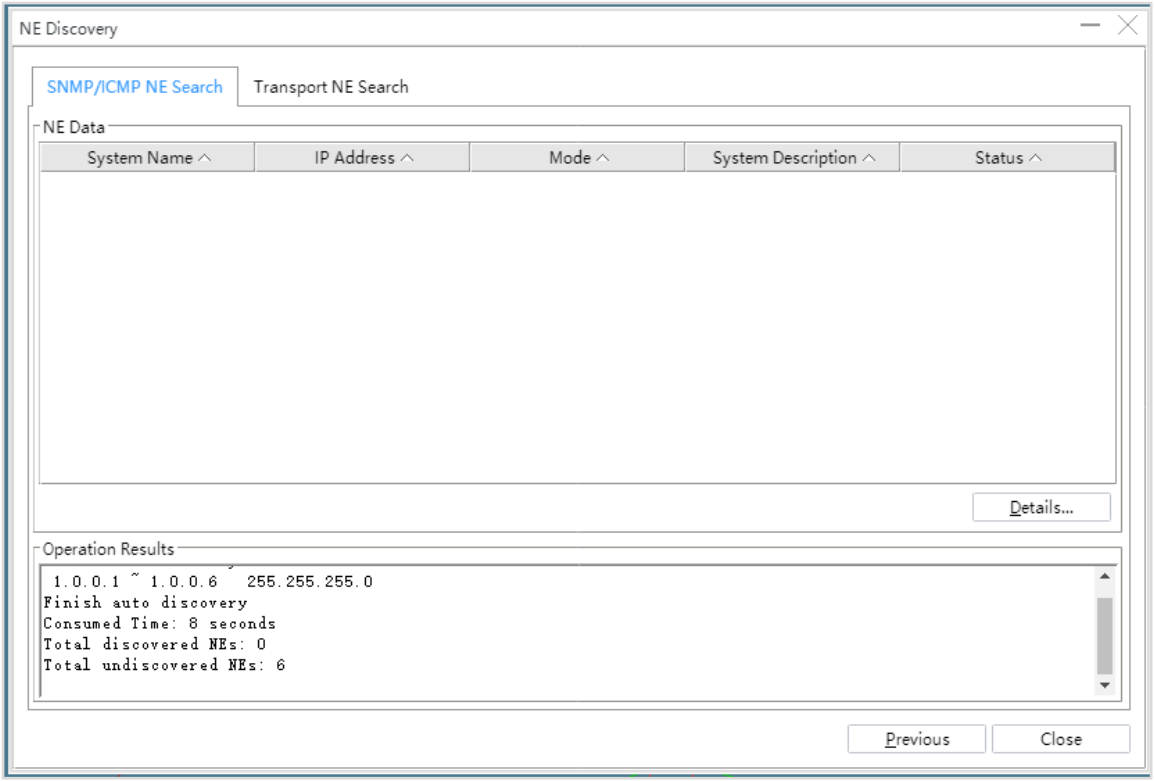

iMaster NCE-IP automatically adds discovered devices to the physical topology.

After NEs are added to iMaster NCE-IP, you can perform the following tests to check whether the SNMP, NETCONF, and STelnet communication between iMaster NCE-IP and NEs is normal.

Test SNMP communication.

Choose **System** > **NE Communication Parameters** from the main menu. Then click **NE Access Protocol Parameters**.

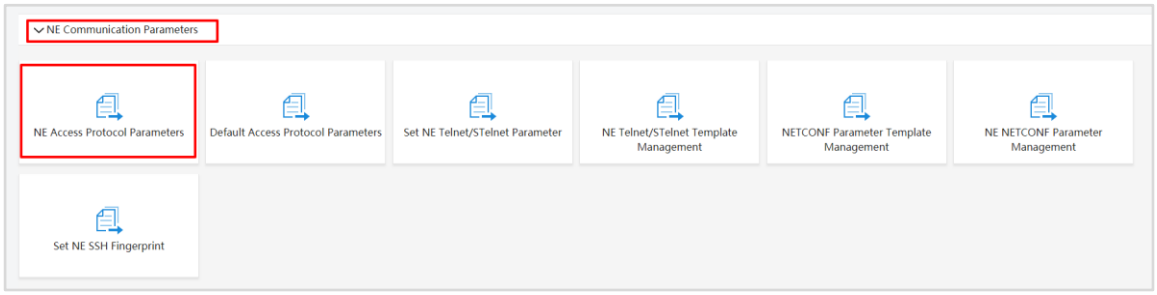

Select all NEs, click **Test** on the right, and wait for the test result.

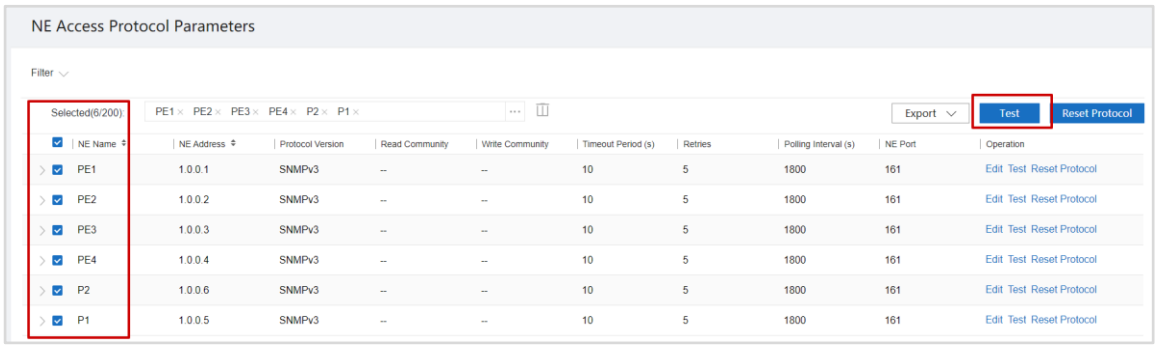

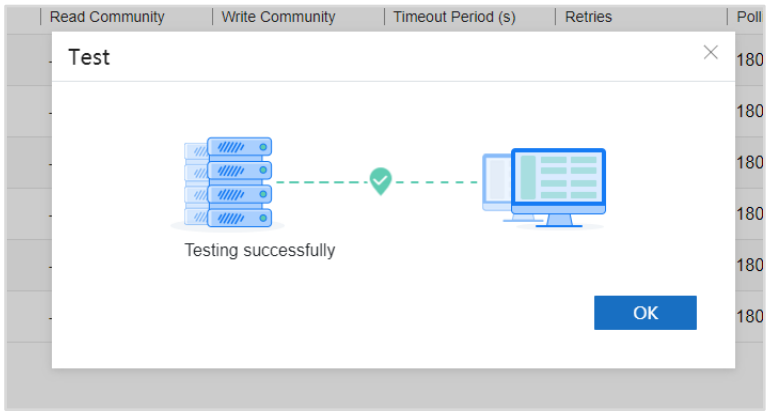

If the SNMP communication with all devices is normal, the system displays a message indicating that the test is successful.

Test STelnet communication.

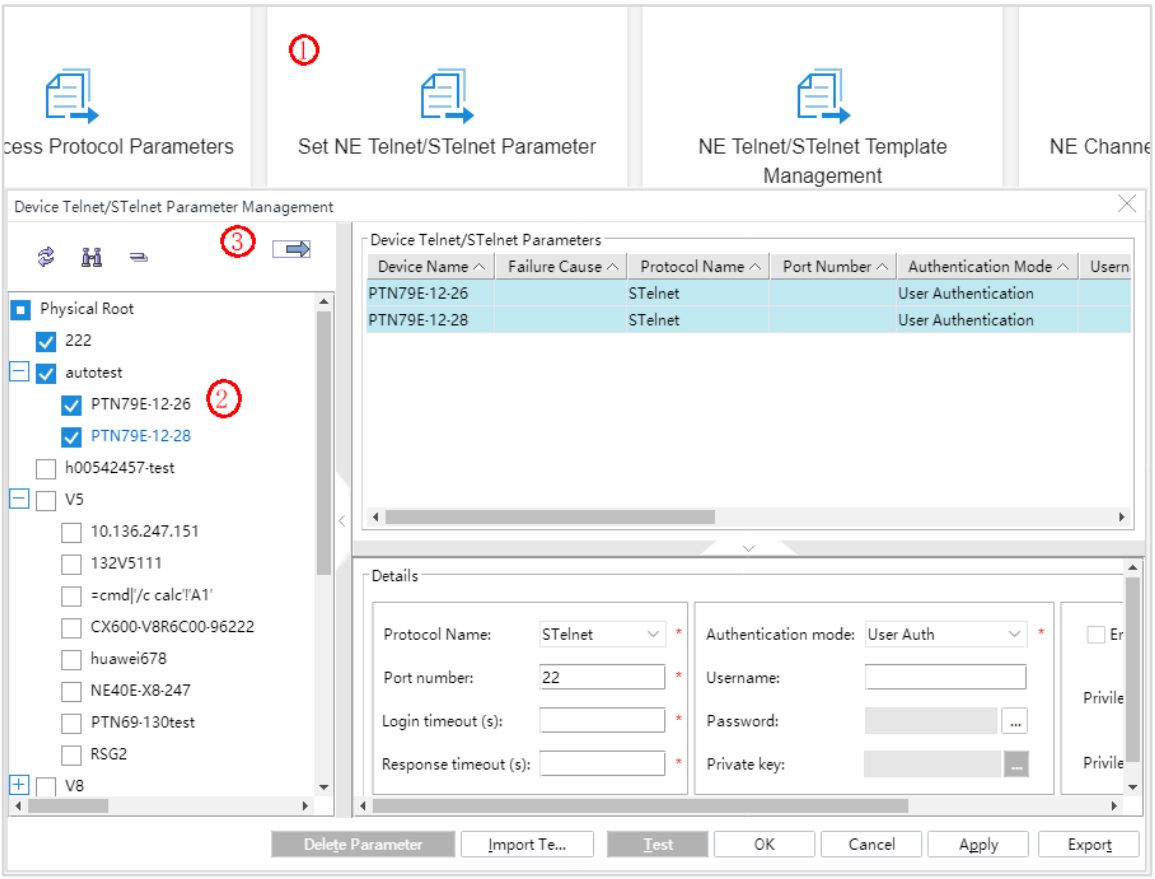

Return to the **NE Communication Parameters** page and click **Set NE Telnet/STelenet Parameter**. In the dialog box that is displayed, select all NEs and click the button in area 3 to add the NEs to the right pane. Then, click **Test**. If a message is displayed indicating that the operation is successful, the STelnet communication between iMaster NCE-IP and NEs is normal.

Test NETCONF communication.

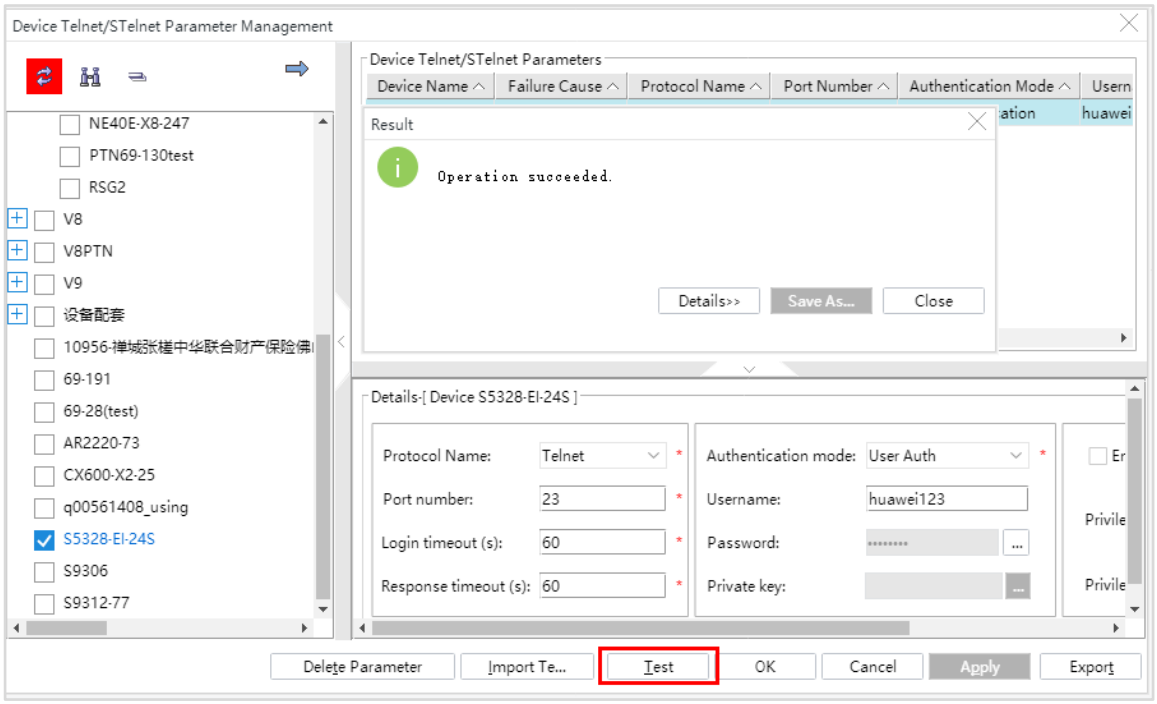

The test process is similar to that of STelnet. On the **NE Communication Parameters** page, click **Set NE NETCONF Parameter**. In the dialog box that is displayed, add devices to the right pane, click **Test**, and wait for the test result.

If all the preceding tests are successful, iMaster NCE-IP can communicate with devices properly.

Step 4 Create a topology.

To facilitate device management, you can create subnets. A subnet, as a logical concept on iMaster NCE-IP, displays topology objects in the same area or with similar attributes.

Open the Network Management app and choose **Topology** > **View** > **Physical Topology** from the main menu.

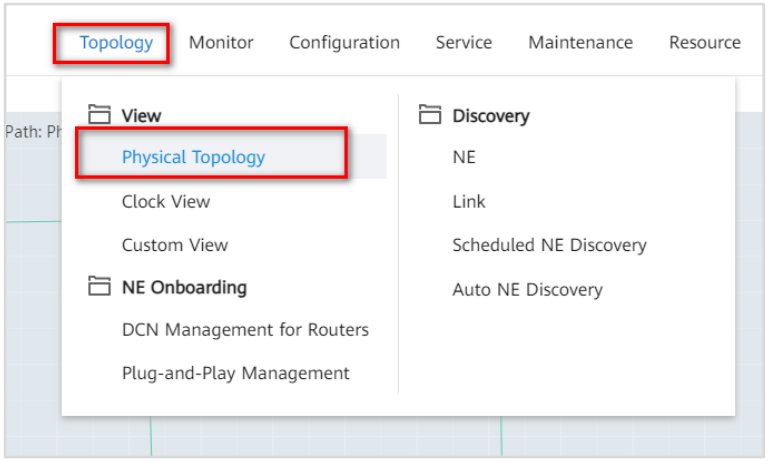

Right-click in the blank area and choose **New** > **Subnet** from the shortcut menu to create a subnet.

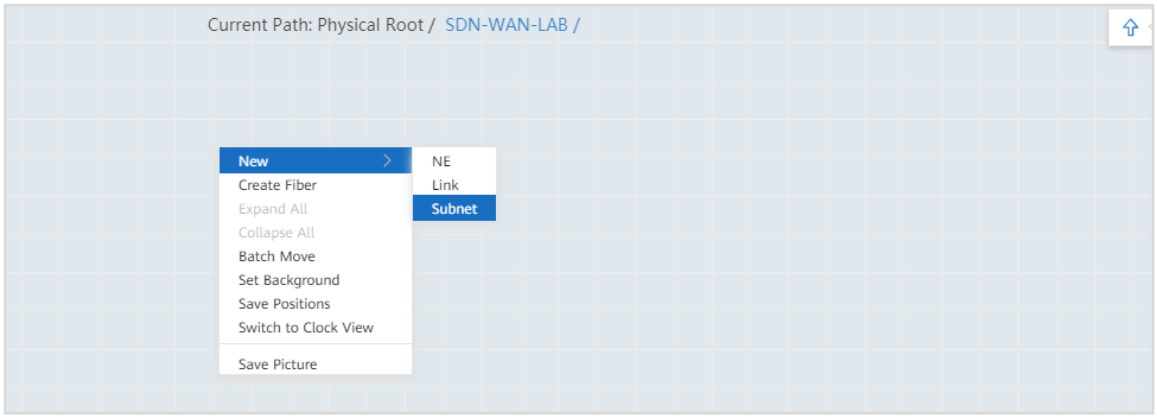

In the dialog box that is displayed, enter the subnet name and click **Select Objects** to add NEs to the subnet.

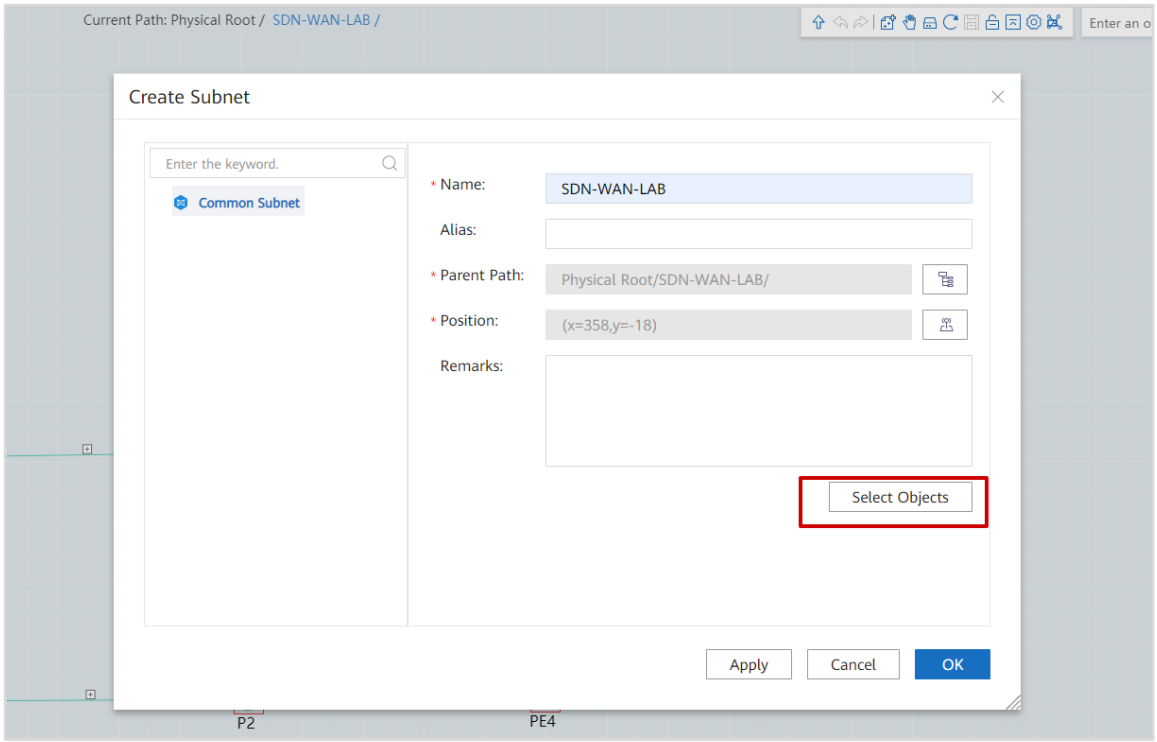

After NEs are added, you can view the added subnet and its subordinate NEs on the left. If the subordinate NEs are displayed in blue, the connections are normal.

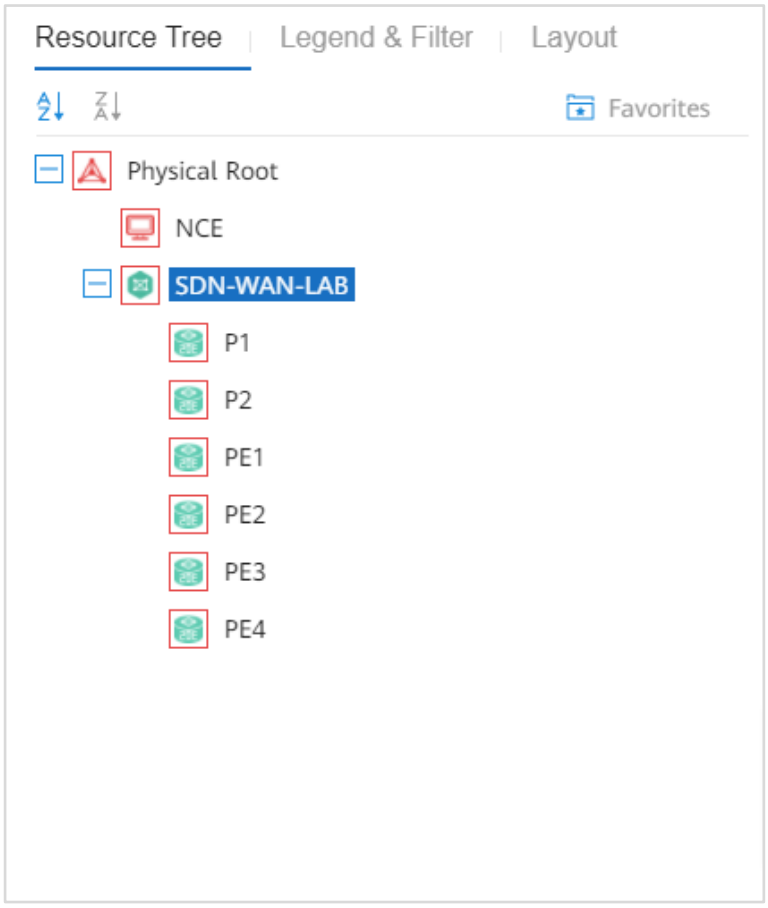

Step 5 Create links.

Open the Network Management app and choose **Topology** > **Discovery** > **Link** from the main menu. On the page that is displayed, select all NEs

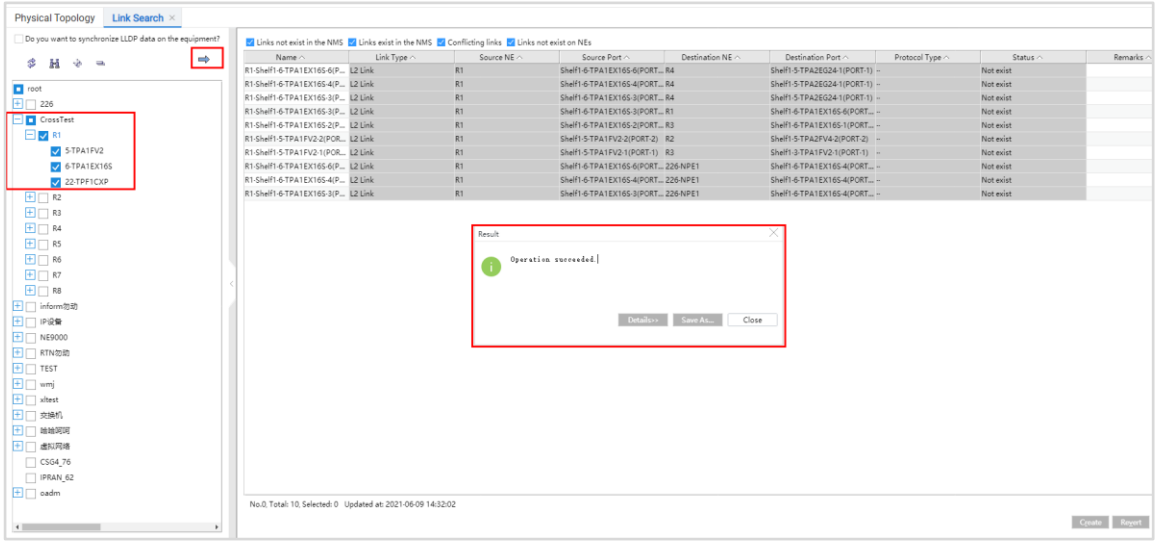

Click the Add icon and wait for the detection result.

Links between devices are detected. If this operation is performed for the first time, the status of all links is **Not exist**.

Hold down **Shift** to select links in the **Not exist** state and click **Create**.

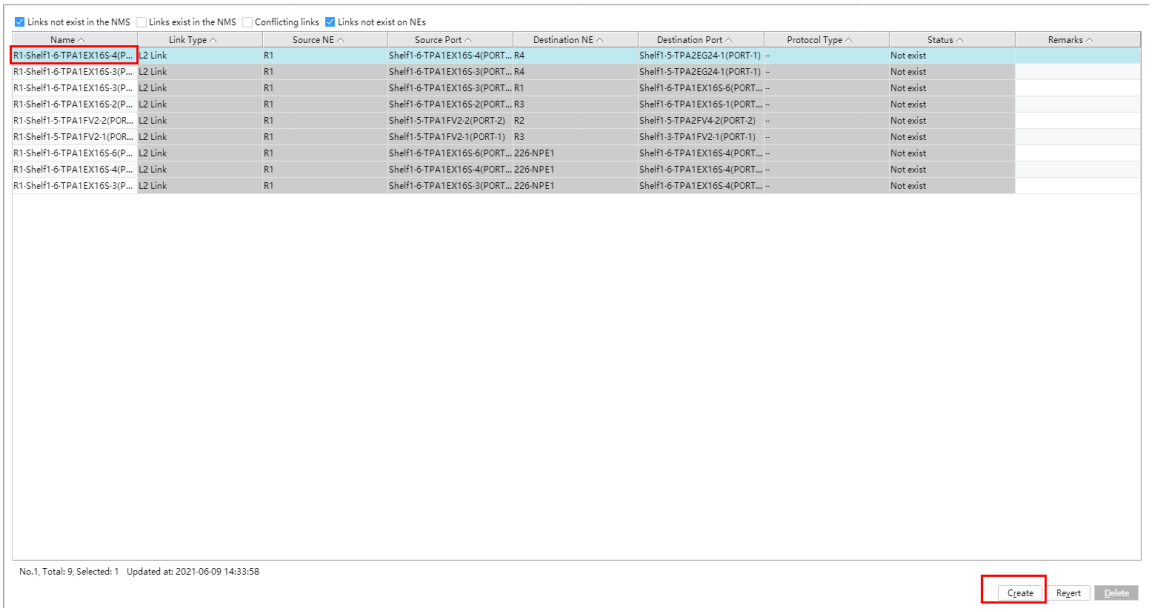

The controller then automatically creates these links and displays a message asking you whether to import these links as fiber links.

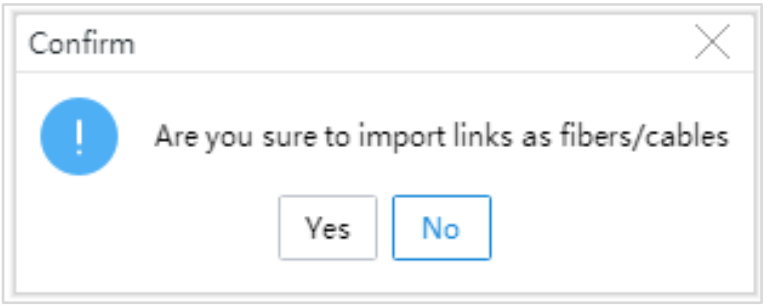

Retain the default value **No**. In this experiment, the links between devices are not fiber links.

Return to the physical topology page and drag the mouse to adjust device locations in the topology.

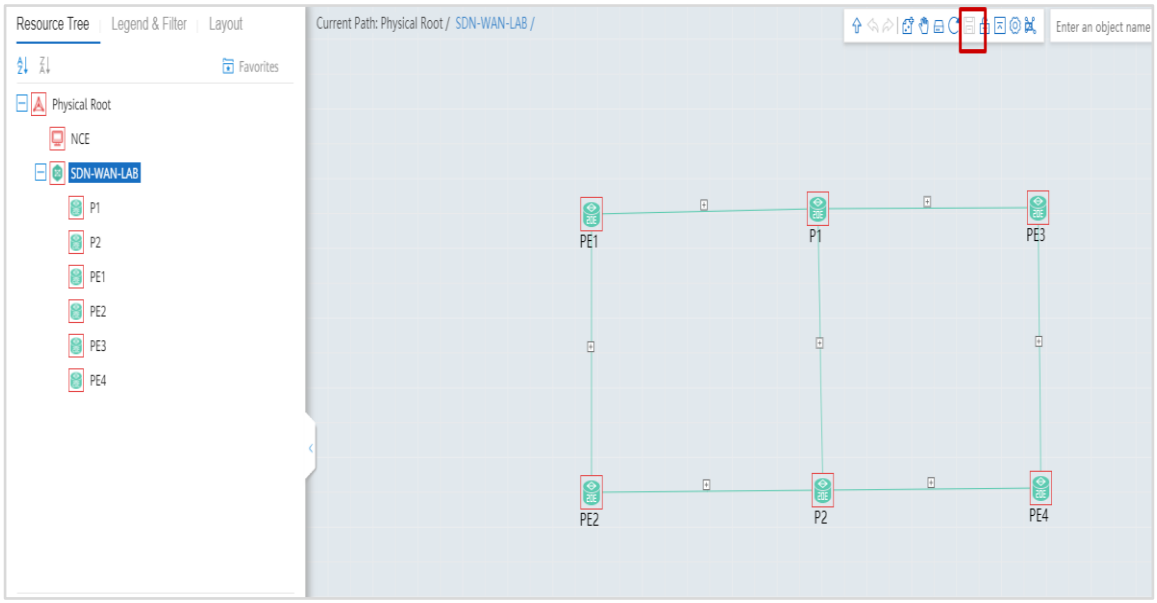

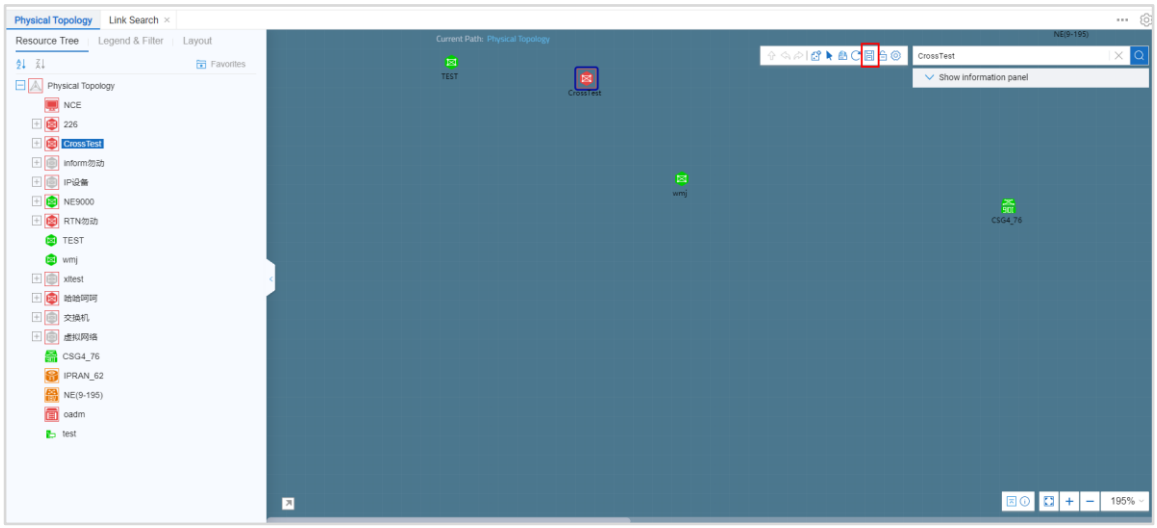

Click the Save icon in the upper right corner to save the current topology layout.

Step 6 Synchronize NE data.

Network maintenance and service configuration can proceed properly only after NE data is synchronized to iMaster NCE-IP and data is consistent between the NE side and controller-side.

Open the Network Management app and choose **Configuration** > **Synchronize NE Configuration Data** from the main menu

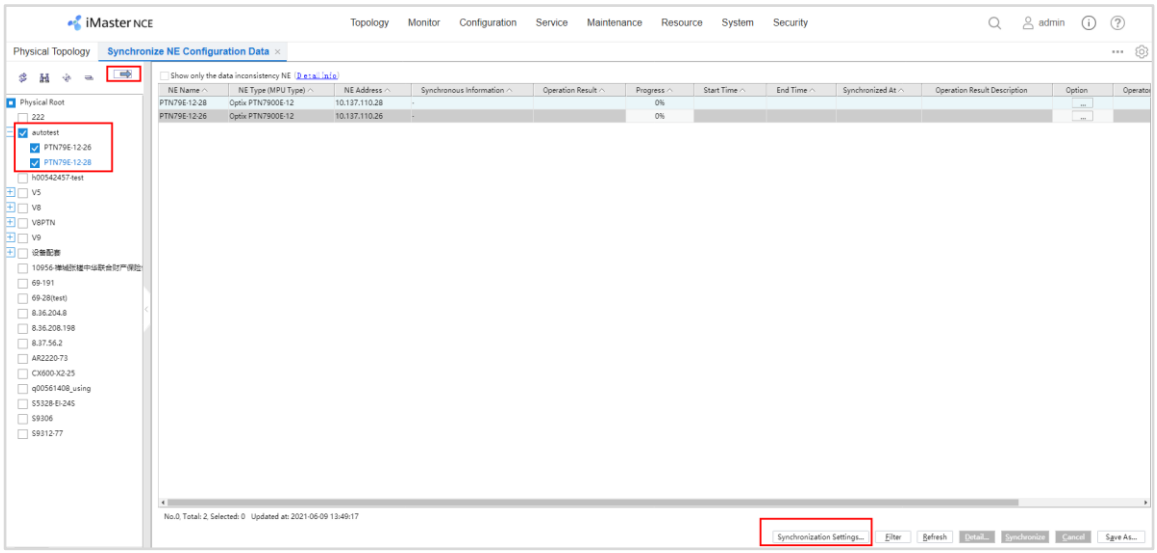

On the page that is displayed, add devices to the right pane, hold down **Shift** to select all devices, and click **Synchronize** in the lower part of the page to synchronize NE configurations.

If the message "Success" is displayed in the **Operation Result** column, iMaster NCE-IP has successfully synchronized device data through SFTP.

3.1.2.7 Controller-Side BGP Configurations

Establish a BGP-LS peer relationship between iMaster NCE-IP and each RR, so that iMaster NCE-IP can receive information, such as link, bandwidth, and TE tunnel status, from NEs.

Establish a BGP SR Policy peer relationship between iMaster NCE-IP and each RR, so that iMaster NCE-IP can deliver SR Policies to NEs through RRs.

Step 1 Configure basic BGP functions.

Open the Network Management app and choose **Configuration** > **Control Unit** > **BGP** from the main menu to configure BGP.

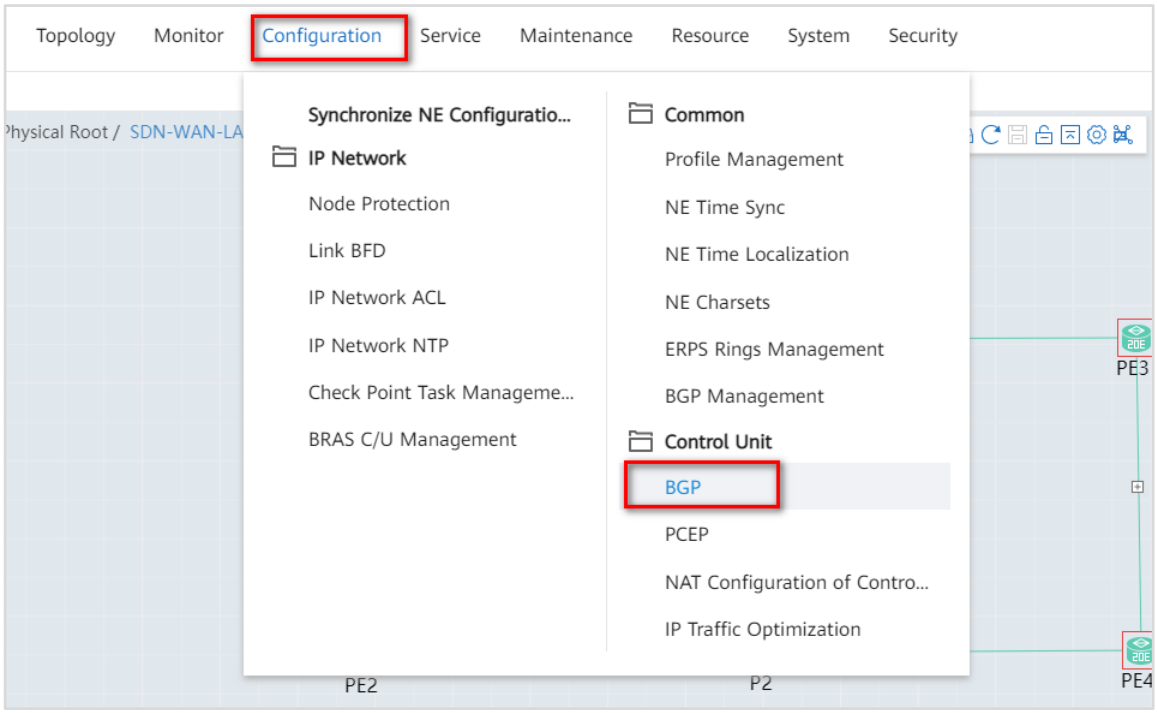

Set the parameters as follows.

Table 3-3 Basic BGP parameters of the controller

Parameter	Value
*Enable BGP	Yes
*Local AS number	65001
Local IPv4 IP	172.21.17.102
*Router ID	172.21.17.102

Set basic BGP parameters according to the parameter planning.

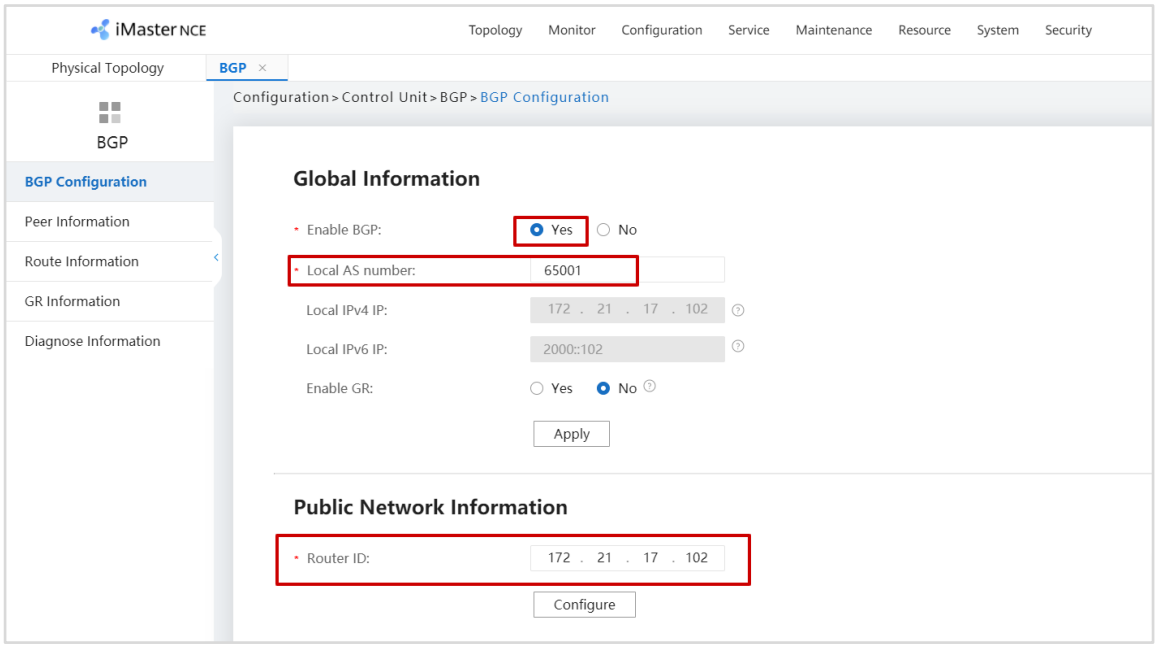

Step 2 Create BGP peers.

Create BGP peers. These BGP peers need to be enabled in specific address families later.

On the **Basic Peer Information** tab page of the BGP configuration page, click **Create Peer**.

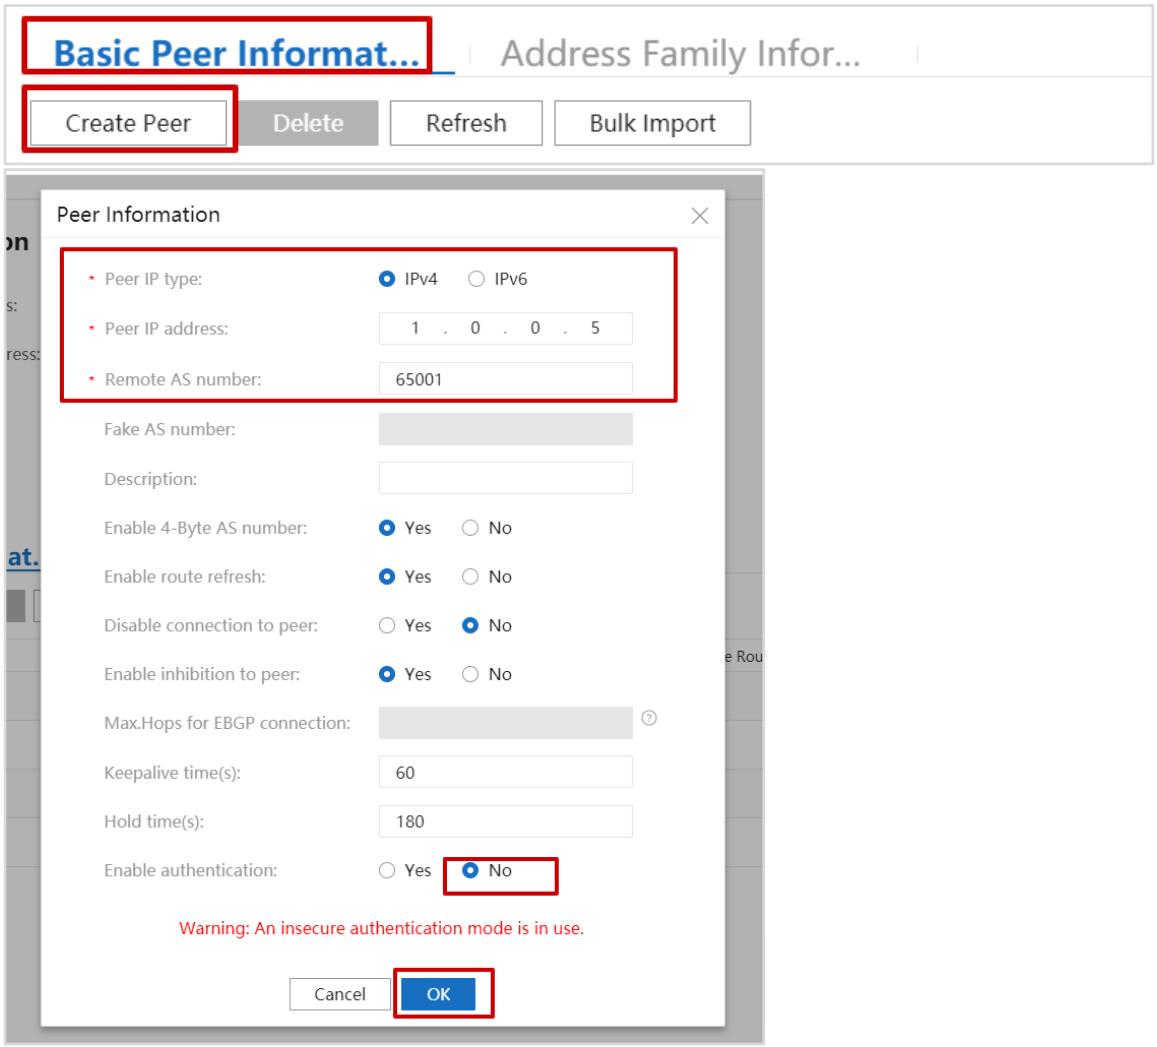

Create IPv4 peers 1.0.0.5 and 1.0.0.6 (corresponding to P1 and P2, which serve as RRs), and disable authentication.

Step 3 Enable IPv4 peers in the BGP-LS address family.

Enable IPv4 peers in the BGP-LS address family, so that iMaster NCE-IP can receive link, bandwidth, and other information from RRs.

On the **Address Family Information** tab page, click **Link-state**. (If **Link-state** is not displayed, click **Create Address Family** to add it.)

Click **Create Peer**. In the dialog box that is displayed, click **Select Peer**.

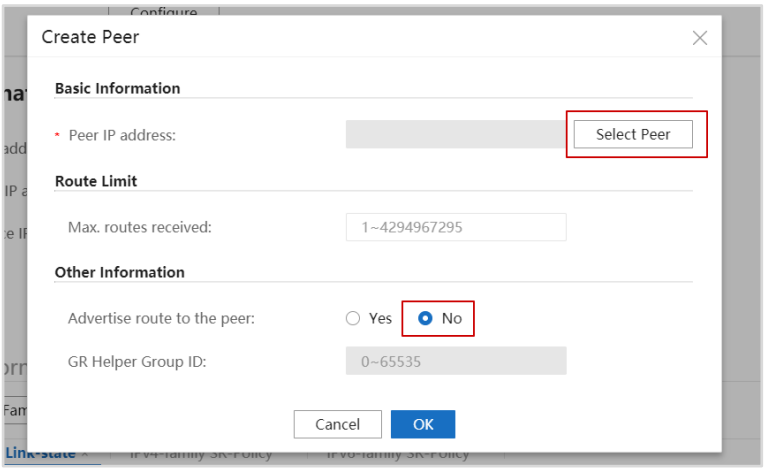

Note that iMaster NCE-IP does not need to send routes to RRs; instead, it only needs to receive routes. Therefore, you need to set **Advertise route to the peer** to **No**.

In the **Select Peer** dialog box, select the previously created peers (1.0.0.5 and 1.0.0.6).

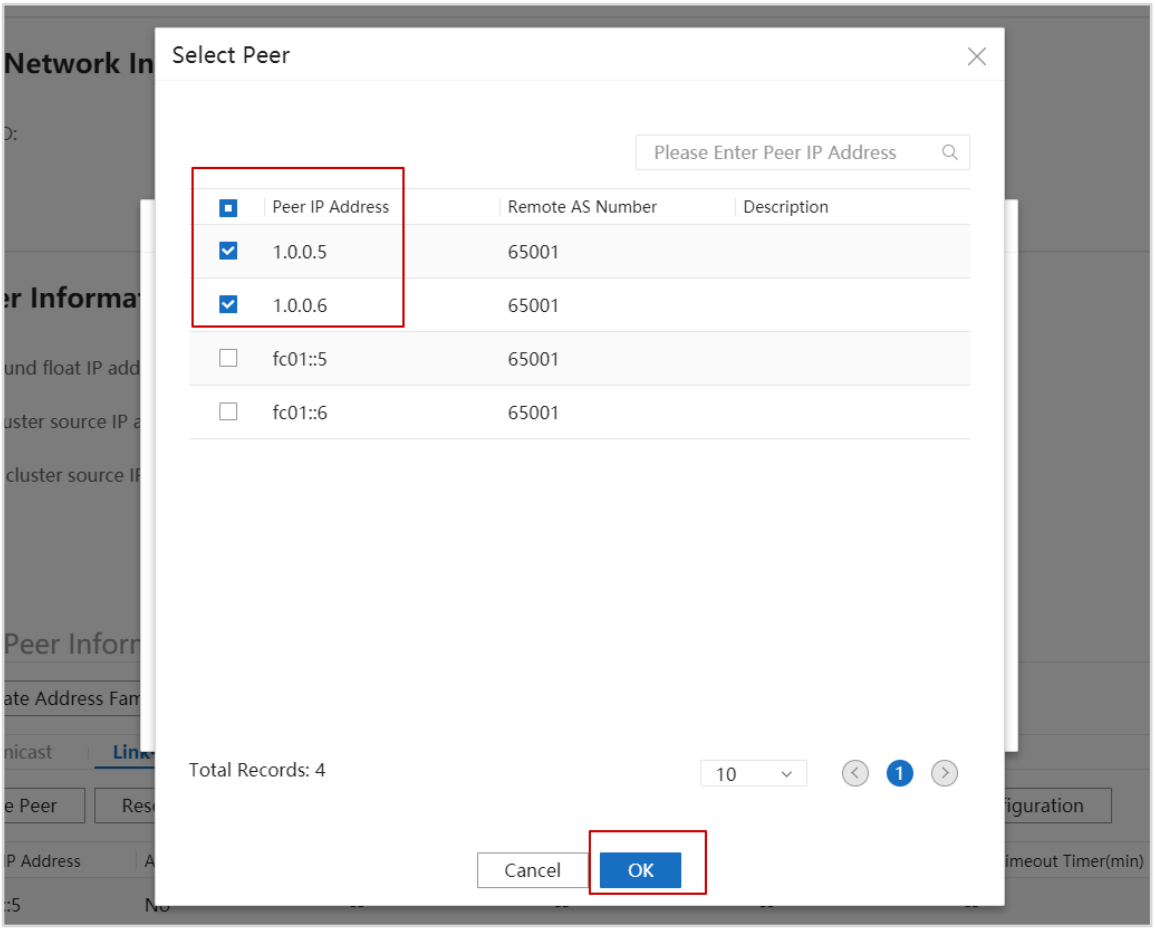

Step 4 Enable IPv4 peers in the BGP SR Policy address family.

Enable IPv4 peers in the BGP SR Policy address family, so that iMaster NCE-IP can deliver SR Policy configurations to NEs.

On the **Address Family Information** tab page, select **IPv4-family SR-Policy**.

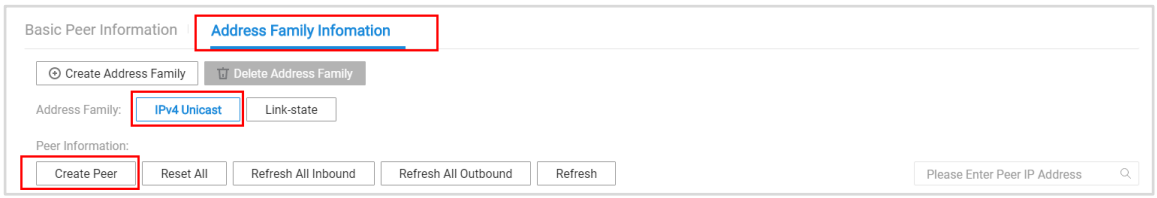

Click **Create Peer**. In the **Create Peer** dialog box, click **Select Peer**.

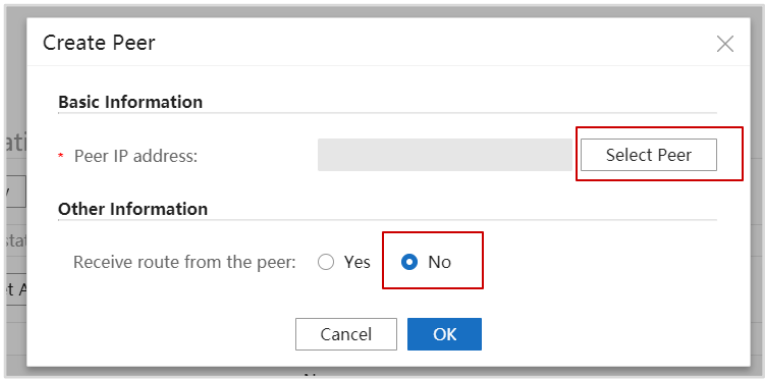

Select the previously created peers (1.0.0.5 and 1.0.0.6).

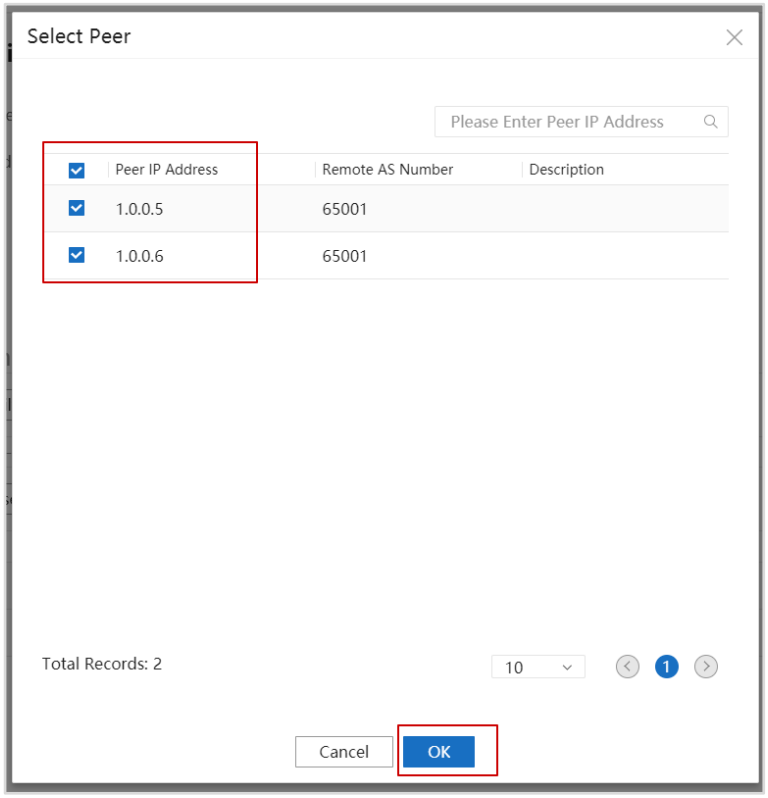

On the **BGP** page, click **Peer Information** to check BGP peer relationships.

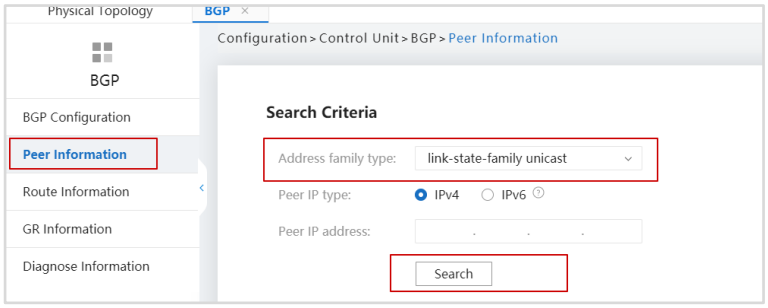

Set **Address family type** to **link-state-family unicast** and click **Search**.

The search results show that iMaster NCE-IP has established BGP-LS peer relationships with P1 and P2 and received route prefixes from P1 and P2.

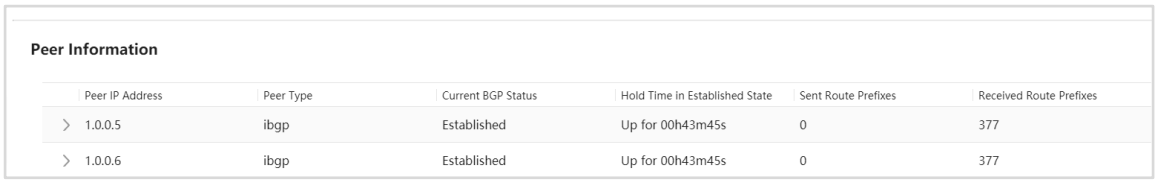

Set **Address family type** to **ipv4-family sr-policy** and click **Search**.

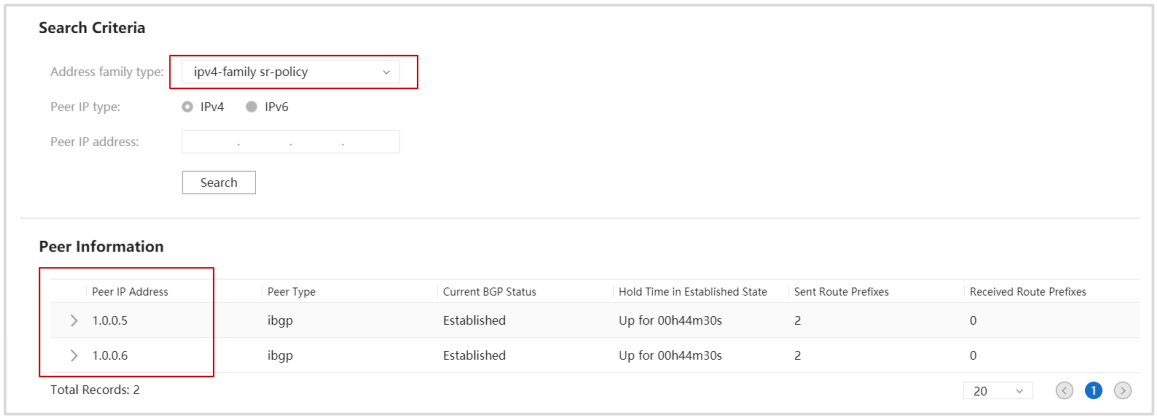

The search results show that BGP SRv4 Policy peer relationships with P1 and P2 have been established.

3.1.2.8 Controller-Side PCE Configuration

Configure iMaster NCE-IP as a PCE server that uses BGP-LS and PCEP to monitor tunnel status.

Step 1 Configure the PCE server.

Open the Network Management app and choose **Configuration** > **Control Unit** > **PCEP** from the main menu to configure the PCE server IP address.

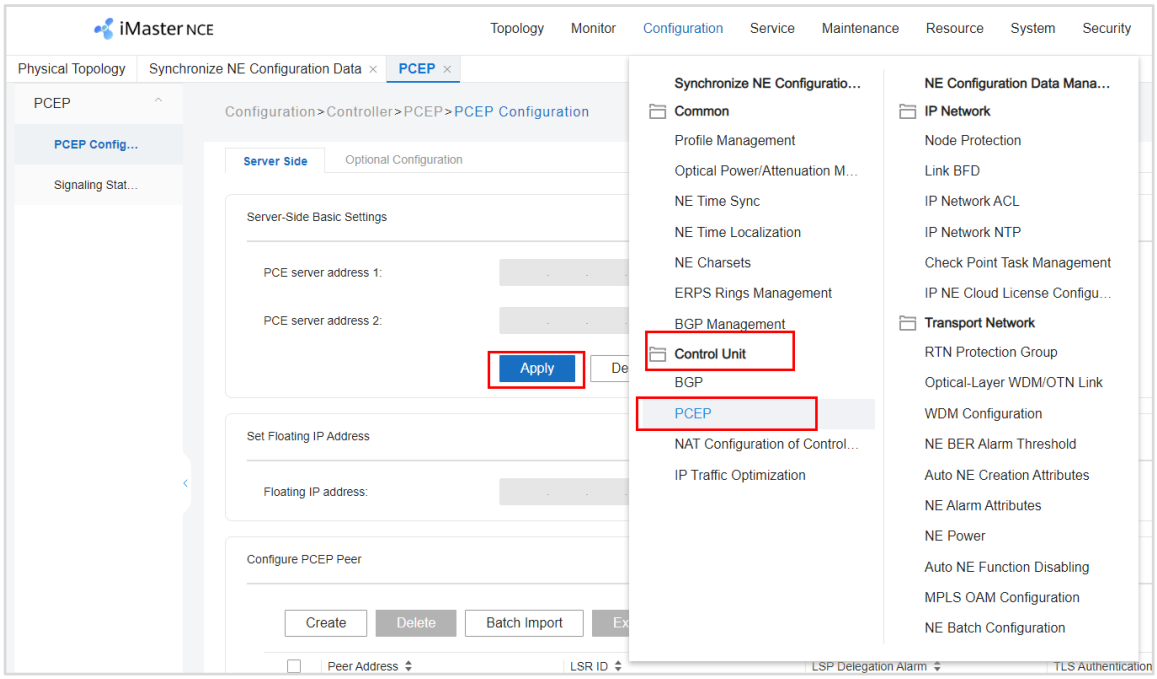

Step 2 Configure PCEP peers.

On the PCEP configuration page, click **Create**. In the dialog box that is displayed, add PCEP peers.

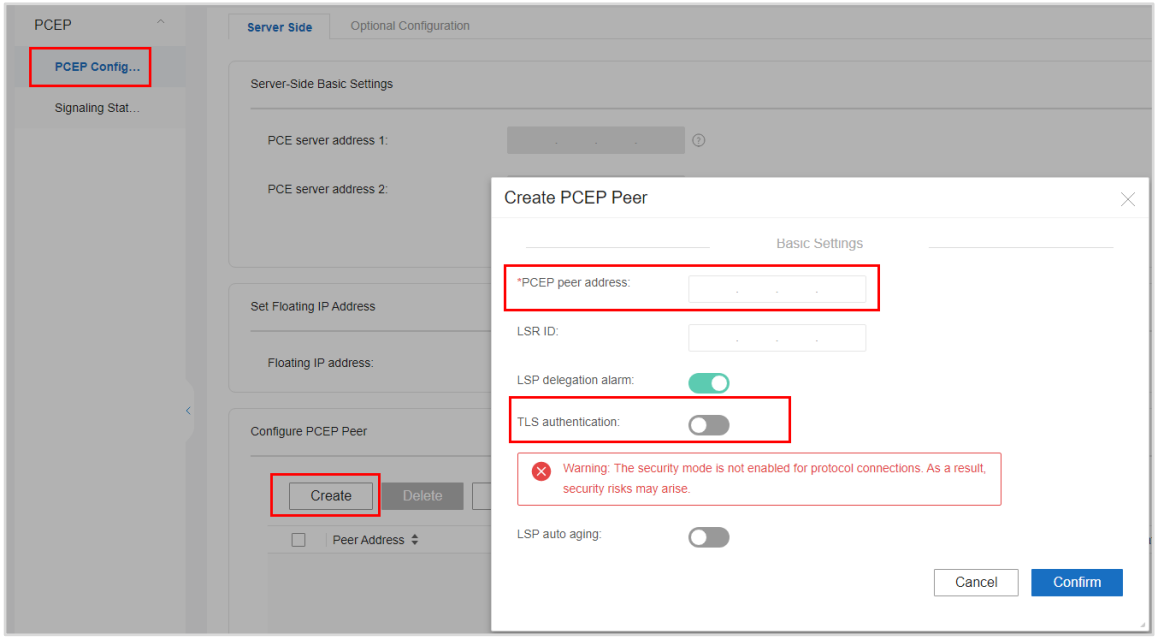

Disable TLS authentication and establish PCEP sessions between iMaster NCE-IP and the four PEs.

Step 3 Check PCEP session status.

Check the PCEP status on iMaster NCE-IP and PEs.

On the PCEP configuration page of iMaster NCE-IP, click **Signaling Statistics** to check PCEP session status.

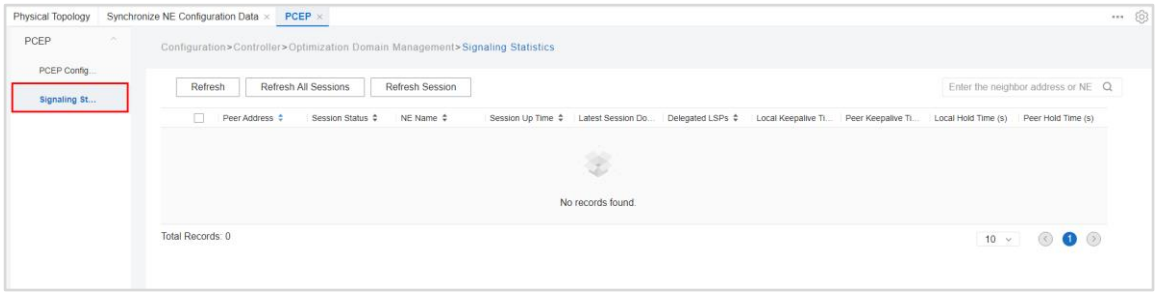

Check the PCEP session status on PEs.

```
[PE1]display pce protocol session

Session IP              State         Session ID
172.21.17.102           UP              2

[PE2]display pce protocol session

Session IP              State         Session ID
172.21.17.102           UP              191

[PE3]display pce protocol session

Session IP              State         Session ID
172.21.17.102           UP              171

[PE4]display pce protocol session

Session IP              State         Session ID
172.21.17.102           UP              139
```

All PEs have established PCEP sessions with iMaster NCE-IP.

3.1.2.9 SR-MPLS TE Tunnel Configuration Delivery by the Controller

iMaster NCE-IP uses NETCONF to deliver tunnel configurations and PCEP to monitor tunnel status.

Before configuring an SR-MPLS TE tunnel, ensure that the PCEP sessions between iMaster NCE-IP and involved NEs are normal, the BGP-LS peer relationships between iMaster NCE-IP and RRs are normal, and the following configurations are ready on NEs.

1. IGP route reachability is available network-wide.
2. MPLS and MPLS TE are enabled both globally and per interface.
3. IGP TE is enabled.
4. SR is enabled globally, and IGP extensions for SR capabilities are enabled.

After the preceding configurations are complete, establish an SR-MPLS TE tunnel between PE1 and PE4.

Open the Network Management app and choose **Service** > **Create** > **Dynamic Tunnel** from the main menu.

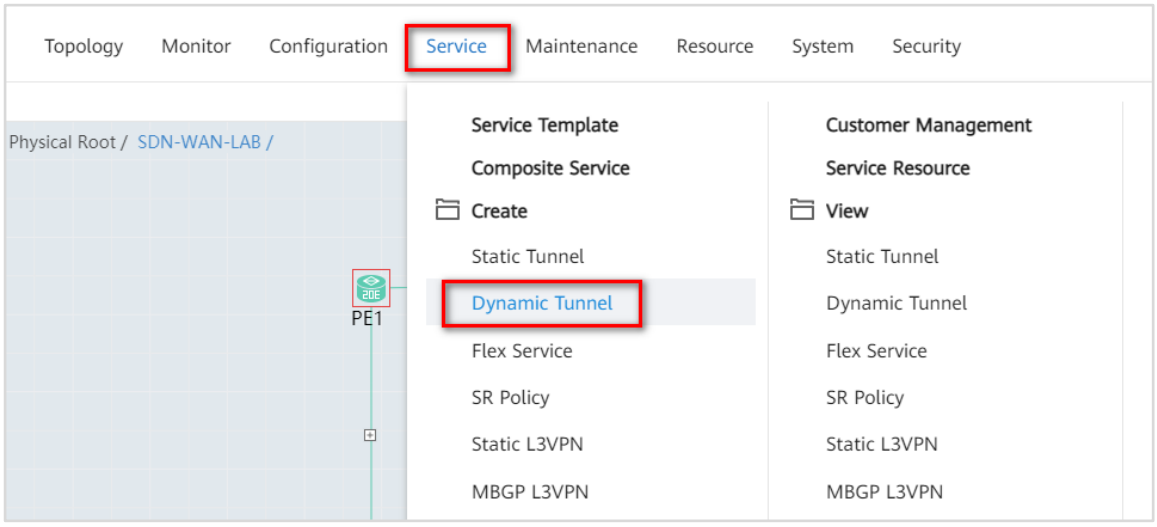

In the dialog box that is displayed, configure the TE tunnel.

Step 1 Configure basic attributes.

In the **Basic Information** area, set **Service template** to **SR_TE_Tunnel** and **Parameter template** to **SR_TE_Tunnel**. Retain the default values for other parameters.

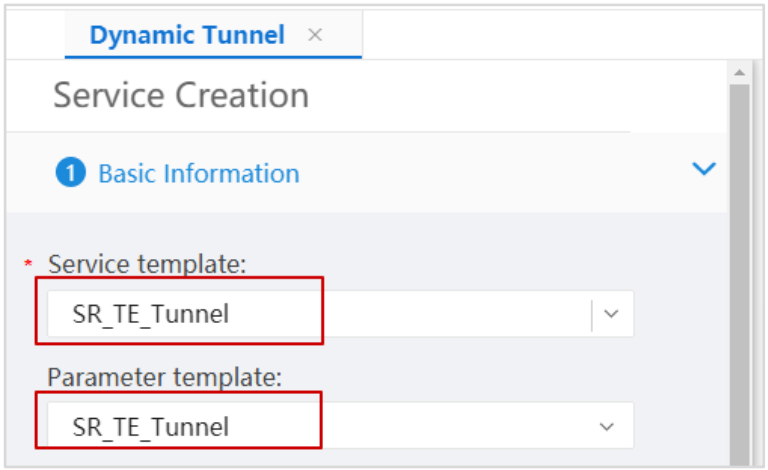

Step 2 Configure the NE list.

In the **NE List** area, set **Source NE** to **PE1** and **Sink NE** to **PE4**.

Step 3 Configure the working tunnel.

In the **Working Tunnel** area, configure **Advanced Information** for the forward and reverse tunnels.

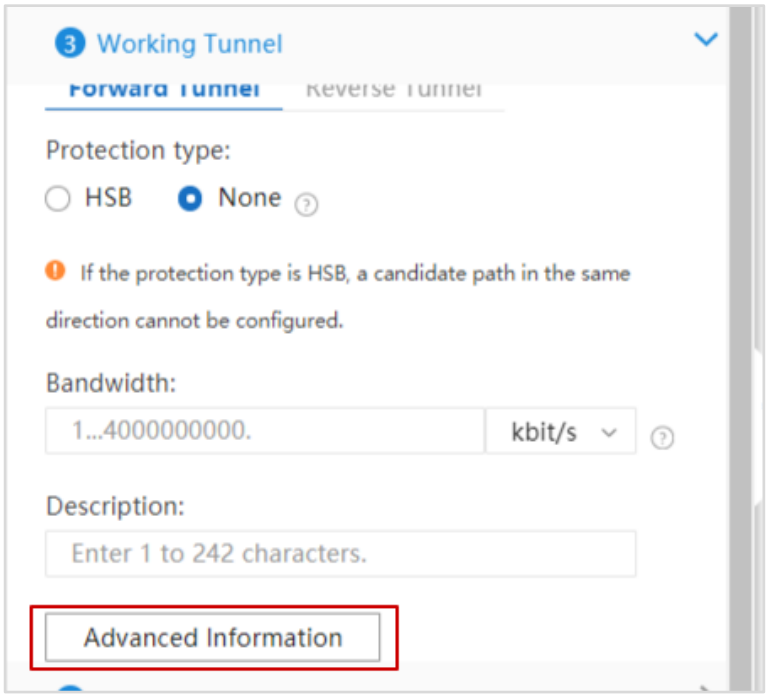

Set **IP address type** to **Unnumbered interface** and **Unnumbered interface** to **LoopBack0**, and retain the default settings for other parameters.

The advanced attribute settings for the forward tunnel are as follows.

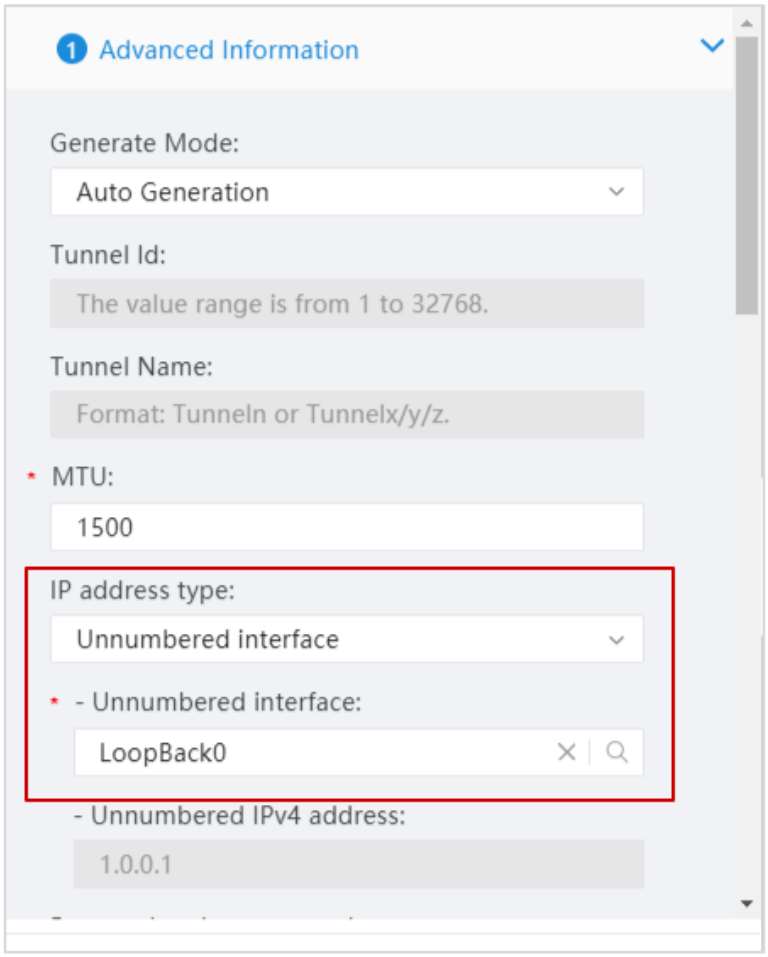

The advanced attribute settings for the reverse tunnel are as follows.

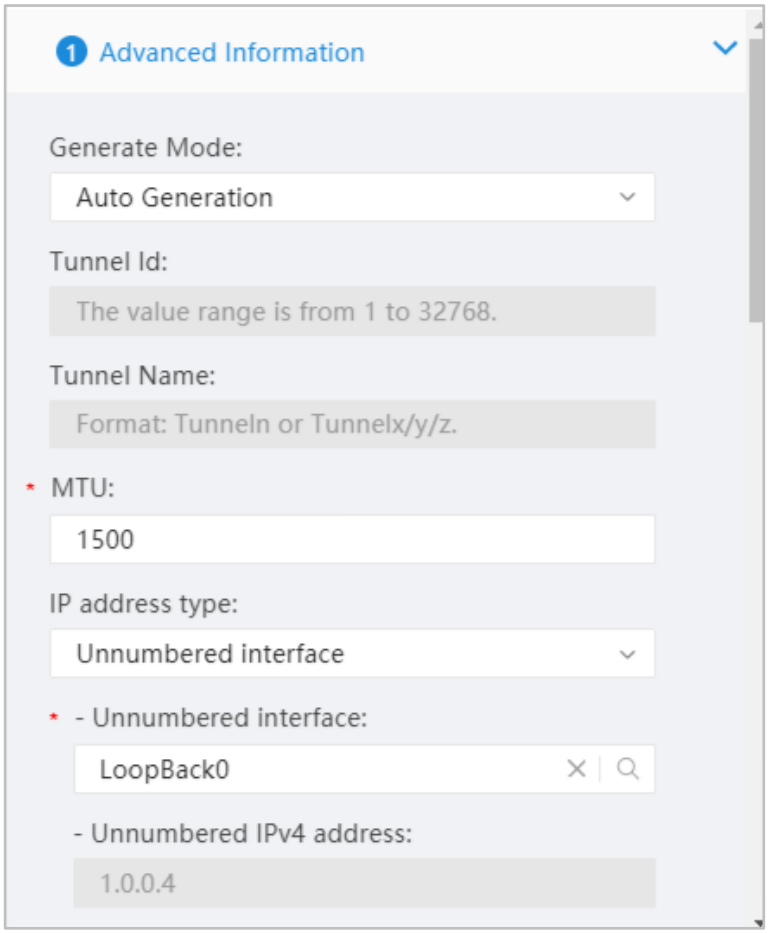

Step 4 Configure route constraints for the working tunnel.

In this template, you can configure path constraints, such as explicit path constraints. The following UIs are for illustration only, and no explicit path is actually configured.

On the **Forward Tunnel** tab page, click the **Explicit Path** value. The **Primary Path Constraints** dialog box is displayed, allowing you to configure explicit path information.

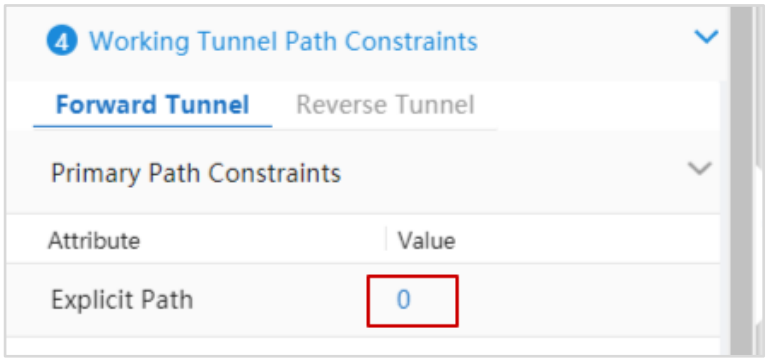

Click **Add** to configure an explicit path.

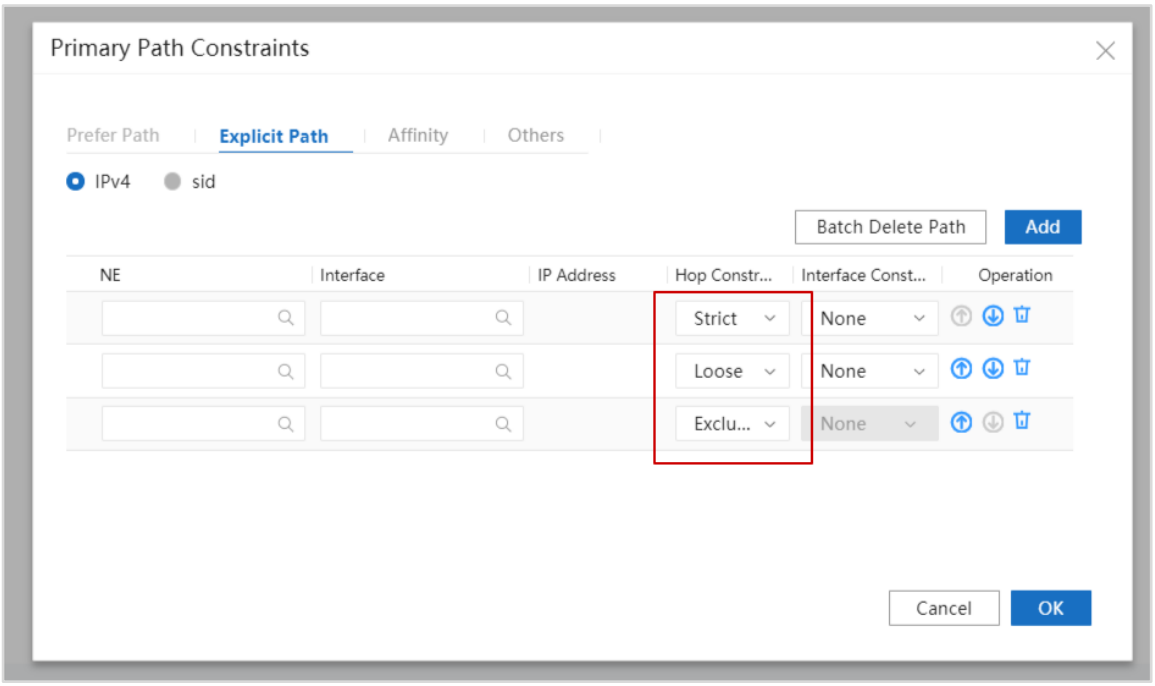

Here, you can configure a strict or loose explicit path or exclude a certain forwarding node. The configuration affects subsequent tunnel path computation result.

Step 5 Compute paths.

Click **Compute Path**. iMaster NCE-IP computes paths based on the least cost, bandwidth balancing, and minimum delay optimization policies.

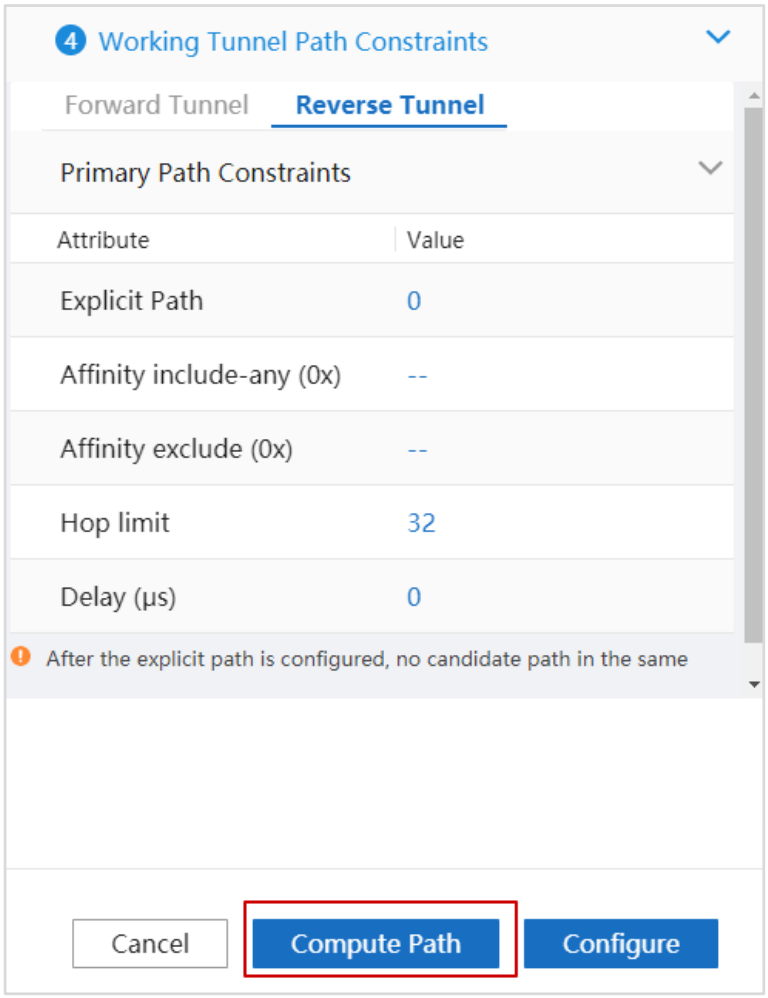

iMaster NCE-IP computes paths based on these three policies.

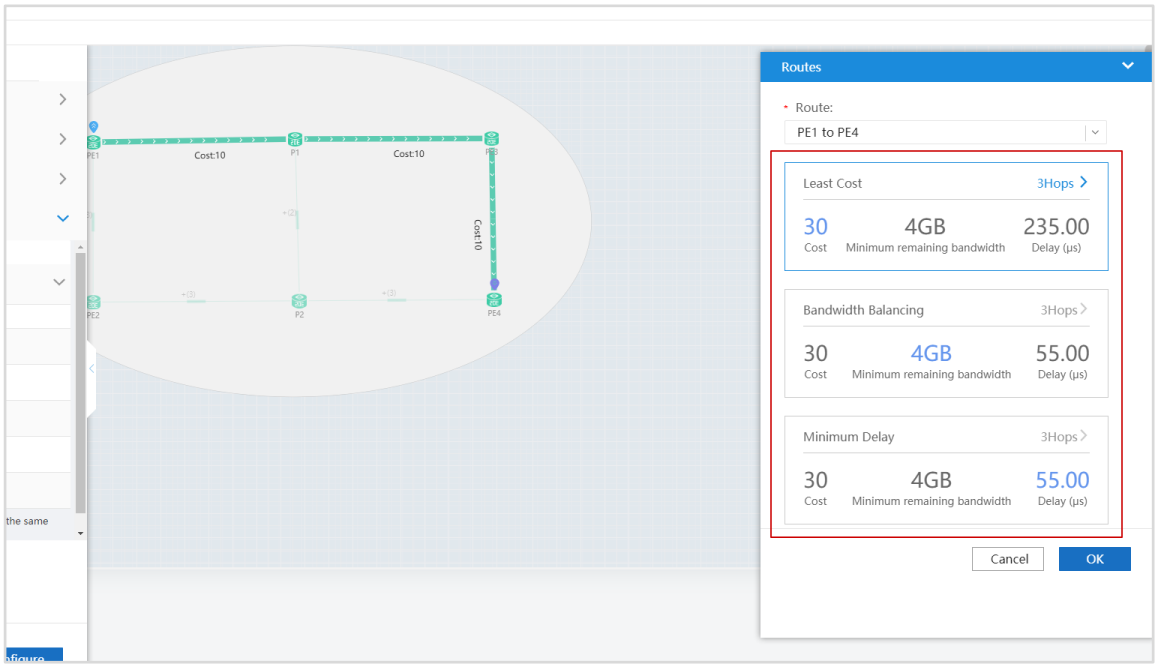

Select the path with the least cost, click **OK**, and then click **Configure**.

In the dialog box that is displayed, click **OK**.

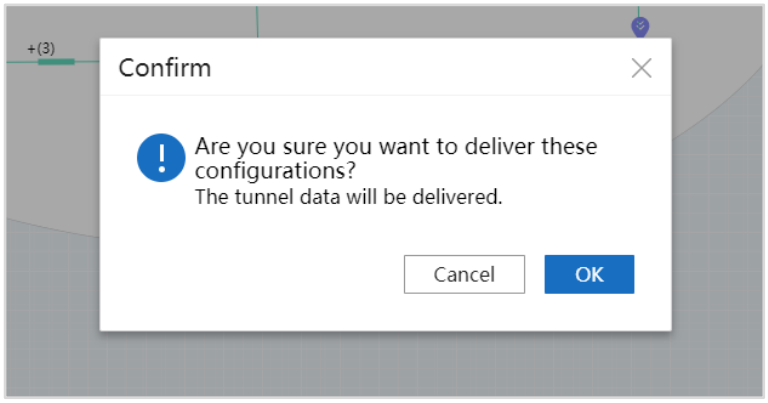

After the tunnel is successfully delivered, a dialog box is displayed, allowing you to view the tunnel status.

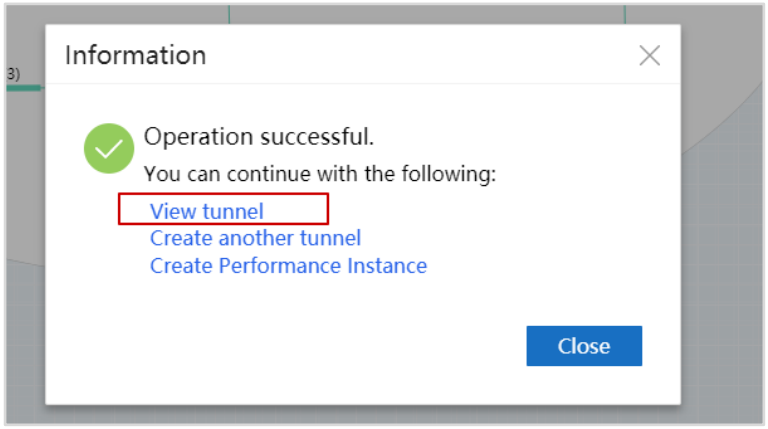

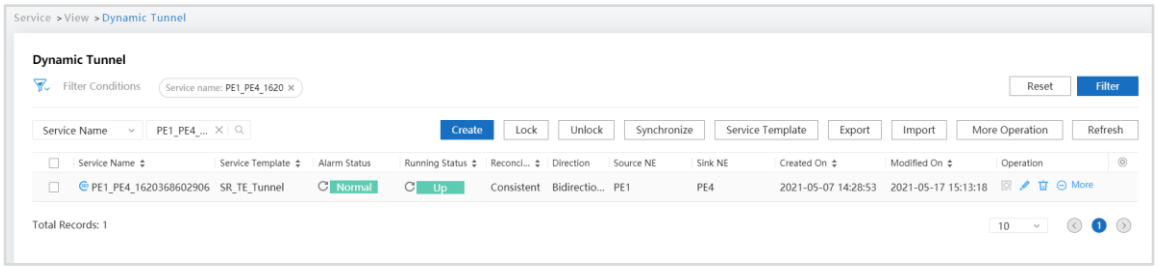

Step 6 Check delivered configurations on a PE.

Check delivered SR-MPLS TE tunnel configurations on PE1.

Check tunnel information on PE1.

```
[PE1]display tunnel-info all
Tunnel ID                Type                 Destination          Status
--------------------------------------------------------------------------------
0x000000000300002002 sr-te                1.0.0.4              UP
0x000000002900000003 srbe-lsp             1.0.0.2              UP
0x000000002900000004 srbe-lsp             1.0.0.4              UP
0x000000002900000005 srbe-lsp             1.0.0.6              UP
0x000000002900000008 srbe-lsp             1.0.0.5              UP
0x000000002900000009 srbe-lsp             1.0.0.3              UP
```

Check SR-MPLS TE tunnel information on PE1.

```
[PE1]display tunnel-info 0x000000000300002002
Tunnel ID:          0x000000000300002002
Type:               sr-te
Name:               Tunnel6
Destination:        1.0.0.4
Instance ID:        0
Cost:               0
Status:             UP
Out Interface:      Tunnel6
NextHop:            0.0.0.0
```

iMaster NCE-IP has delivered a tunnel numbered 6.

Check tunnel configurations on PE1.

```
[PE1]display current-configuration interface Tunnel 6
#
interface Tunnel6
 ip address unnumbered interface LoopBack0
 tunnel-protocol mpls te
 destination 1.0.0.4
 mpls te signal-protocol segment-routing
 mpls te tunnel-id 6
 mpls te pce delegate
```

Check the forwarding path of the TE tunnel on PE1.

```
[PE1]display mpls te tunnel path

 Tunnel Interface Name : Tunnel6
 Lsp ID : 1.0.0.1 :6 :19
 Hop Information
   Hop 0     Link label 48091                NAI 10.0.0.1:10.0.0.2
   Hop 1     Link label 48091                NAI 10.0.0.5:10.0.0.6
   Hop 2     Link label 48090                NAI 10.0.0.34:10.0.0.33
```

The command output shows that:

1. The first segment is from PE1 to P1, and the label used to guide packet forwarding is adjacency SID 48091.
2. The second segment is from P1 to PE3, and the label used to guide packet forwarding is adjacency SID 48091.
3. The third segment is from PE3 to PE4, and the label used to guide packet forwarding is adjacency SID 48090.

We can run the **display segment-routing adjacency mpls forwarding** command on each device along the forwarding path to check the outbound interface corresponding to each segment's adjacency SID.

P1 is used as an example. We can see that the TE tunnel is forcibly configured to use GE0/3/2 as the outbound interface when packets are forwarded through P1.

```
[P1]display segment-routing adjacency mpls forwarding

                Segment Routing Adjacency MPLS Forwarding Information

Label       Interface             NextHop             Type      MPLSMtu    Mtu
--------------------------------------------------------------------------------
48090       GE0/3/0               10.0.0.1            ISIS-V4   ---        1500
48091       GE0/3/2               10.0.0.6            ISIS-V4   ---        1500
48092       GE0/3/4               10.0.0.29           ISIS-V4   ---        1500
```

3.1.2.10 L3VPN Service Delivery by the Controller

Use the controller to create an L3VPN service between PE1 and PE4 and recurses the service to an SR-MPLS TE tunnel for traffic forwarding.

Create Loopback1 on PE1 and PE4 to simulate L3VPN access users.

PE1

```
[PE1]interface LoopBack1
[PE1-LoopBack1] ip address 172.16.1.1 32
[PE1-LoopBack1] quit
```

PE4

```
[PE4]interface LoopBack1
[PE4-LoopBack1] ip address 172.16.4.1 32
[PE4-LoopBack1] quit
```

After configuring loopback interface addresses on PE1 and PE4, synchronize NE configurations to iMaster NCE-IP.

Choose **Configuration** > **Synchronize NE Configuration Data** from the main menu. In the **Synchronize NE Configuration Data** dialog box, select PE1 and PE4 and click **Synchronize** to synchronize their configurations to iMaster NCE-IP.

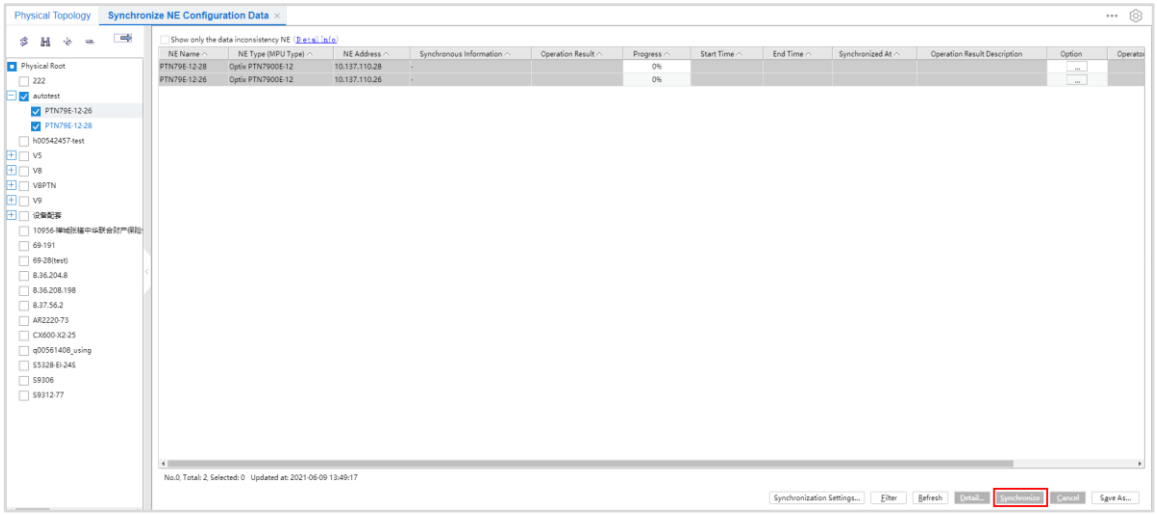

Open the Network Management app and choose **Service** > **Create** > **MBGP L3VPN** from the main menu.

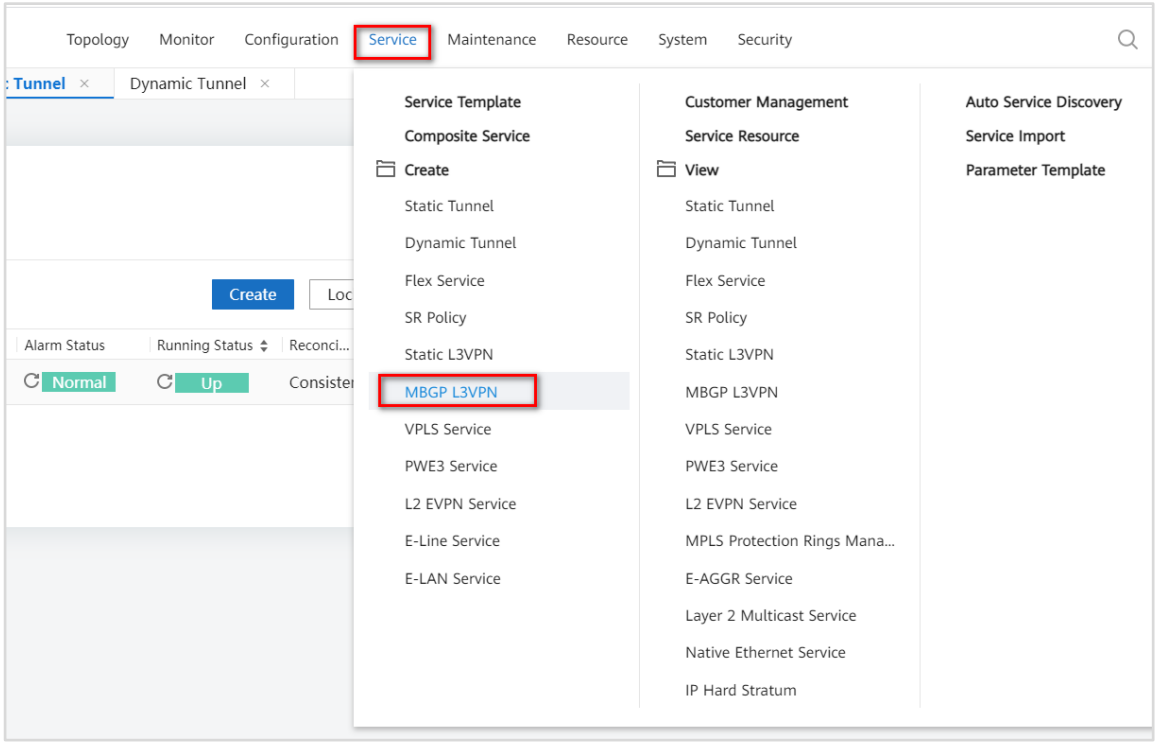

Step 1 Set basic parameters.

In the **Basic Parameter** area, set **Service template** to **L3VPN_Auto_Select** (default value), set **Service name** and **VRF Name**, set **Parameter template** to **MBGP_L3VPN**, and enable **Auto allocate RD and RT: Yes**.

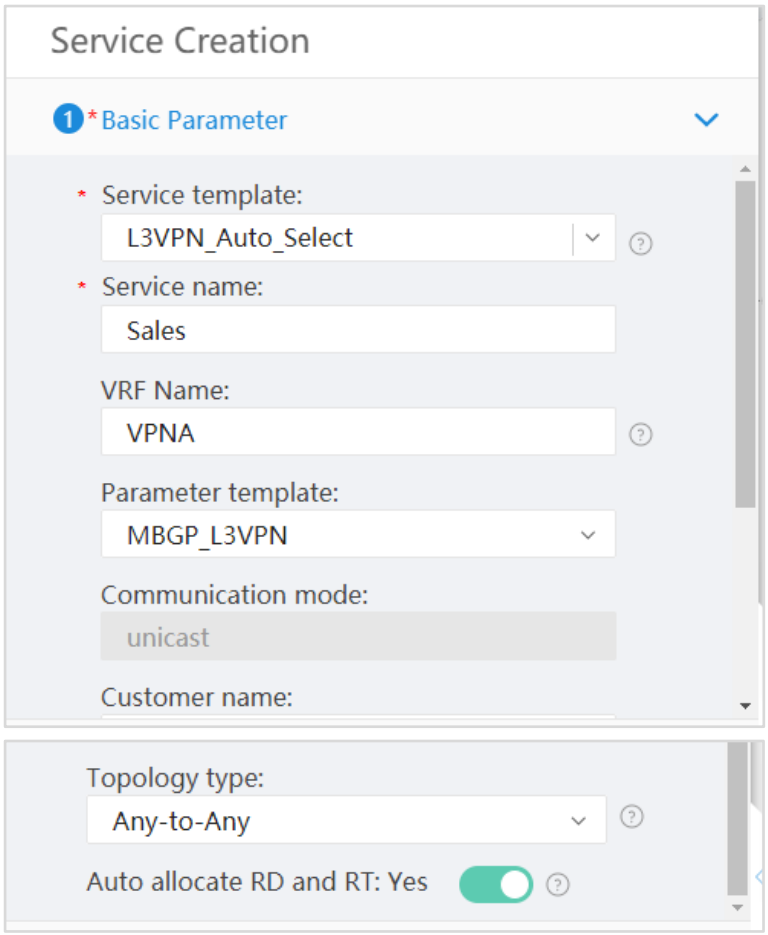

Step 2 Configure service nodes.

Select service nodes PE1 and PE4 and select a tunnel policy.

Click **+** in the **Service Node** area.

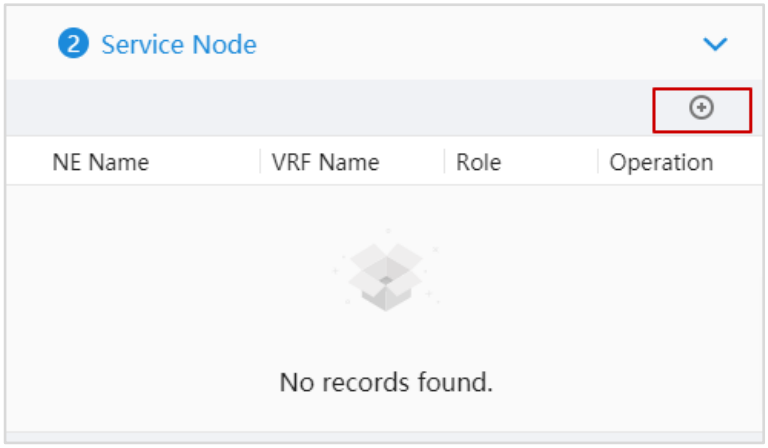

In the dialog box that is displayed, set **NE name** to **PE1** and select **Enable IPv4**.

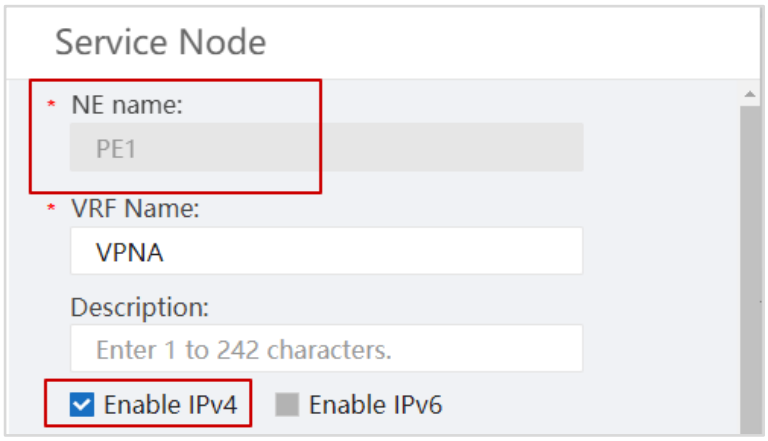

Select **Enable IPv4**.

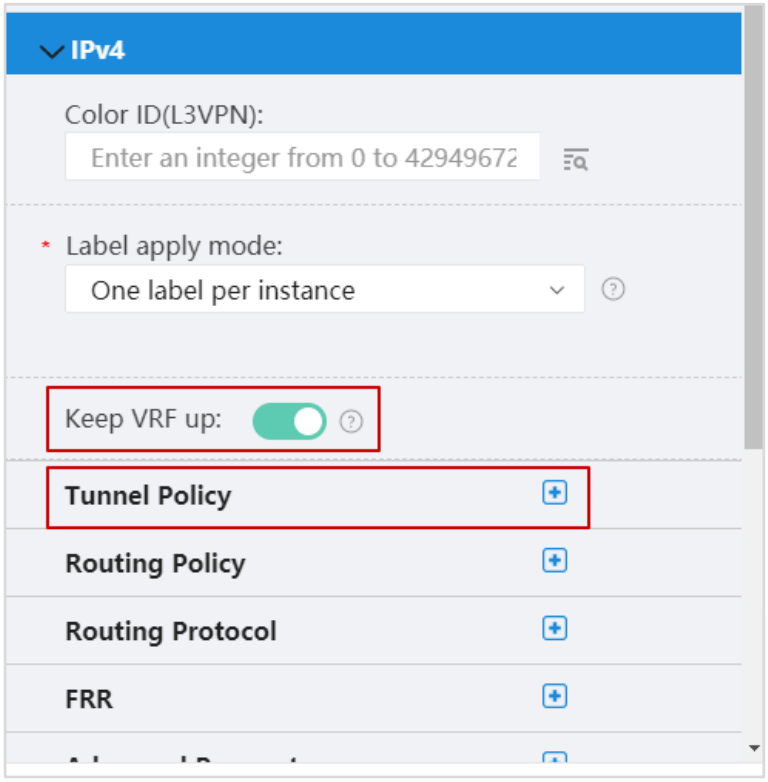

In the dialog box that is displayed, click **+** next to **Tunnel Policy**.

Note that **Keep VRF up** must be selected. In the following figure, **Keep VRF up** is not selected.

In the dialog box that is displayed, set **Policy Type** to **Auto select**, click **+**, and select the SR-MPLS TE tunnel.

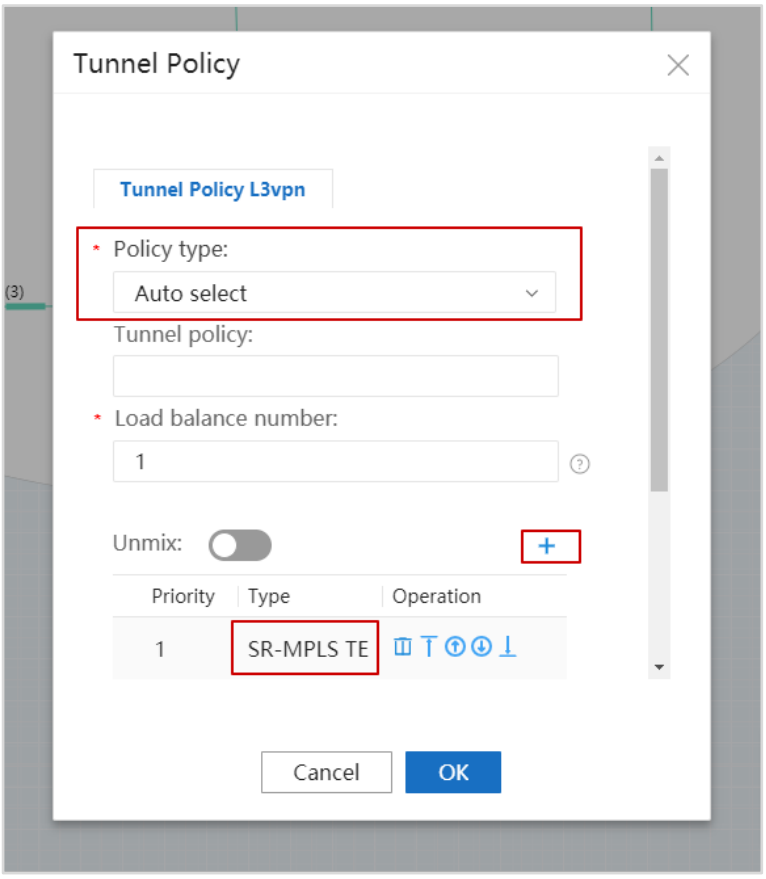

Then, click **OK**.

The service node configuration on PE1 is complete. The service node configuration on PE4 is similar to that on PE1.

Step 3 Configure service access points.

Configure service access points for PE1 and PE4. Here, Loopback1 interfaces are used to simulate user access.

Click **+** in the **Service Access Point** area.

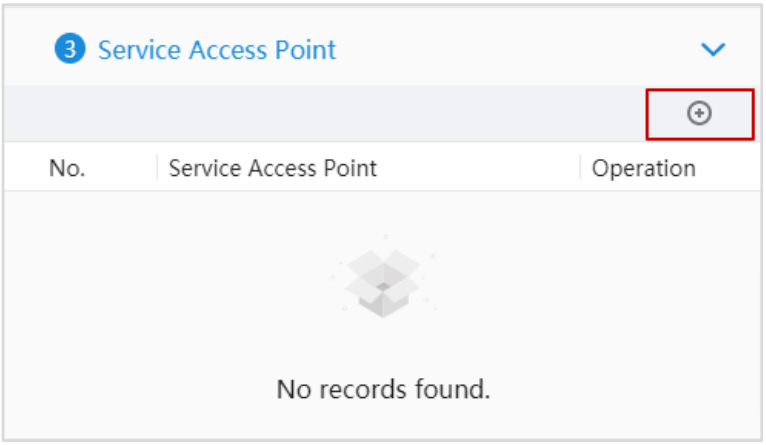

In the **Service Access Point** dialog box, set **Name** to **PE1_Tenant** for PE1, retain the default single-homing access mode, and click **+**.

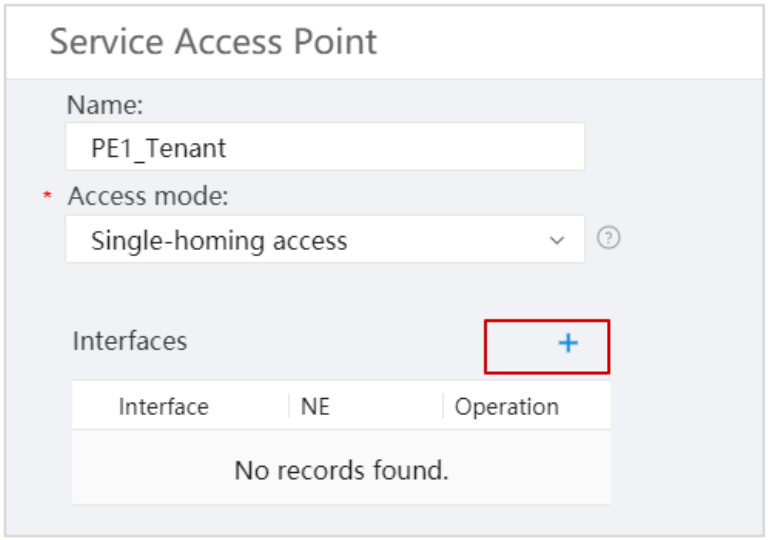

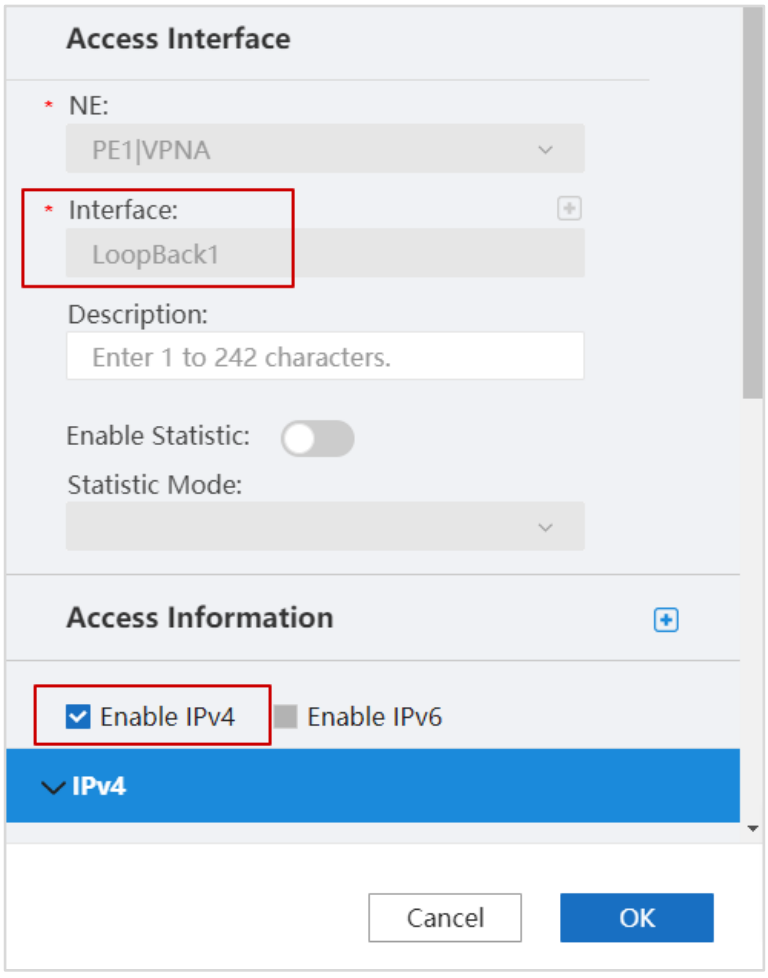

In the dialog box that is displayed, set **NE** to **PE1|VPNA**, **Interface** to **Loopback1**, and select **Enable IPv4**.

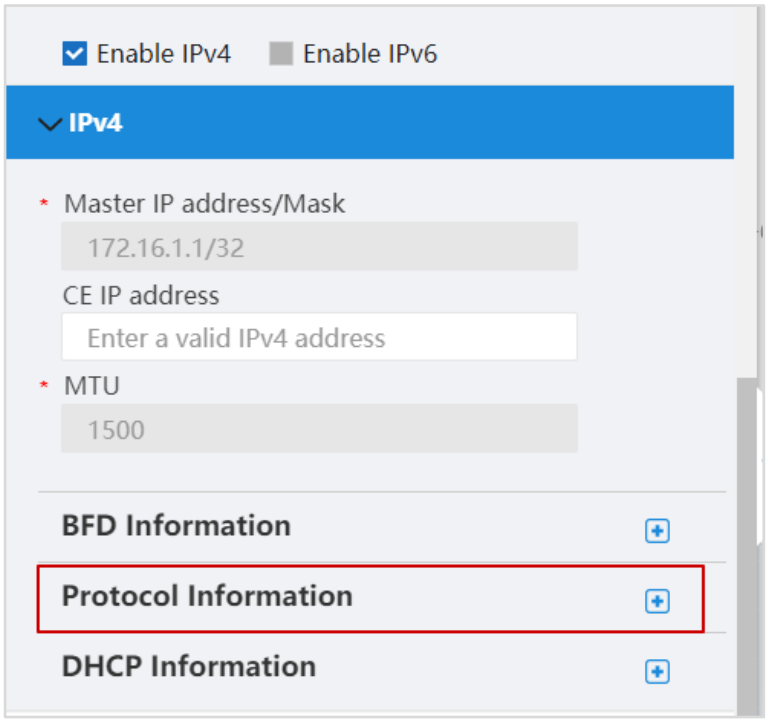

In the new area that is displayed, you can configure protocol information, that is, information about routing protocols between the current device and CE.

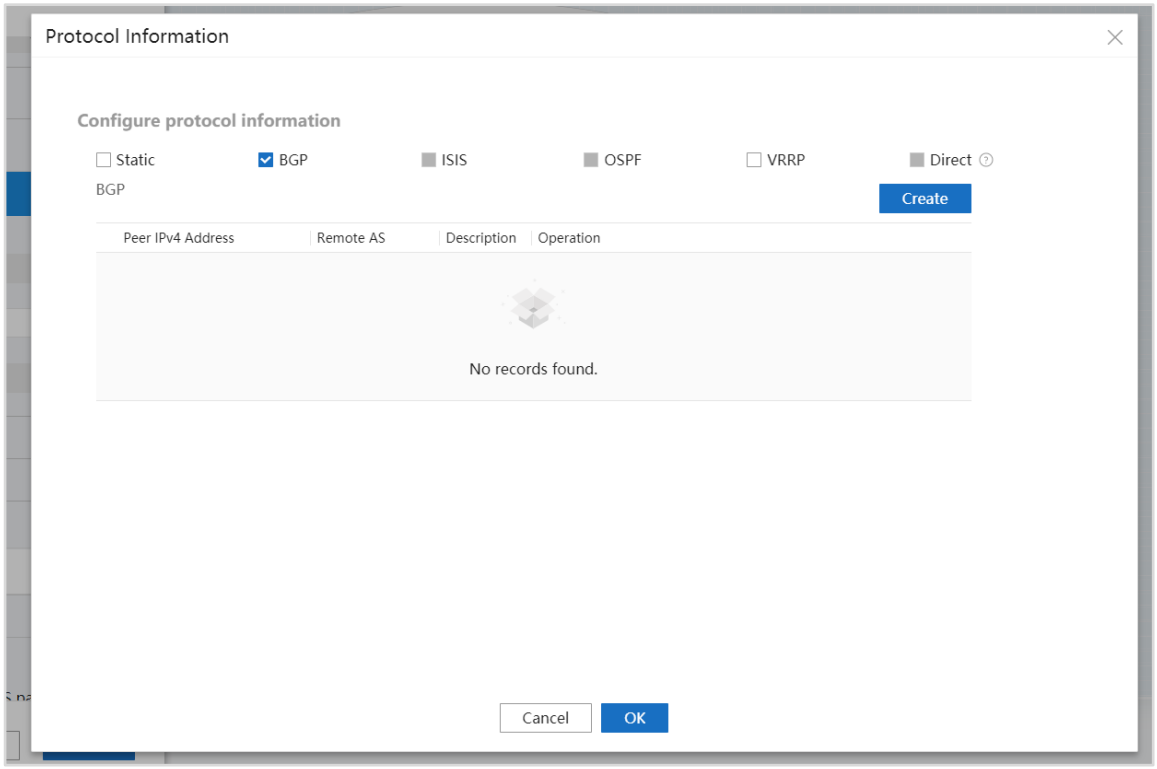

After selecting the corresponding protocol types, you can set protocol parameters. For example, you can configure a BGP peer relationship with the CE if BGP is selected.

In this experiment, Loopback1 interfaces are used to simulate CEs. Therefore, you do not need to configure routing protocols between the current device and CE.

Finally, click **OK**. The service access point PE1_Tenant is configured. The service access point configuration on PE4 is similar to that on PE1.

The service access point configurations (CE-related configurations) on PE1 and PE4 are complete.

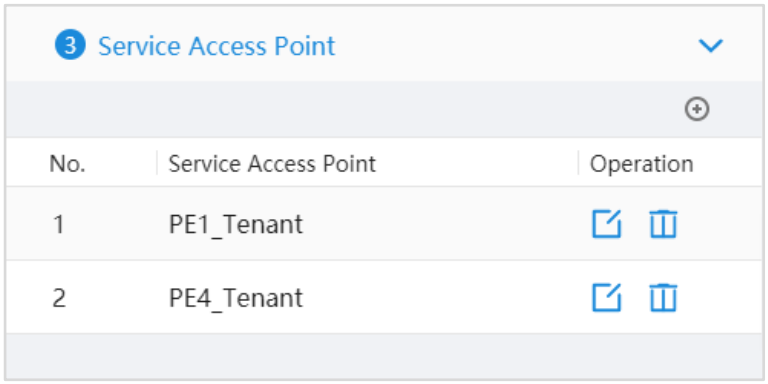

Step 4 Configure tunnels.

Here, associate VRFs with SR-MPLS TE.

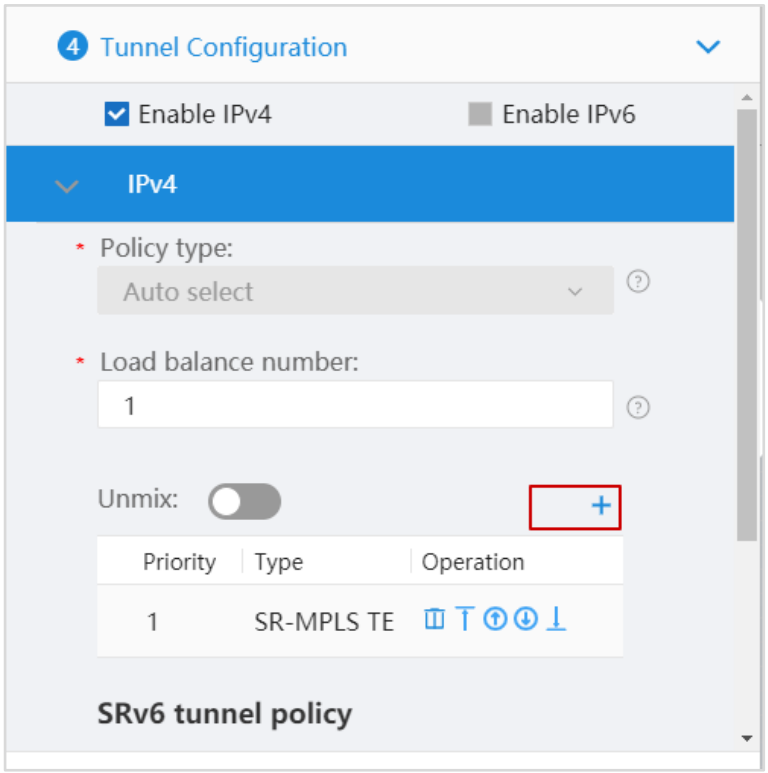

Select **Enable IPv4**, click **+**, and set the tunnel type to **SR-MPLS TE**.

Finally, click **Configure** and wait for the L3VPN service configurations to be delivered to devices.

In the **Success** dialog box, you can click the corresponding hyperlink to check information about the configured L3VPN service.

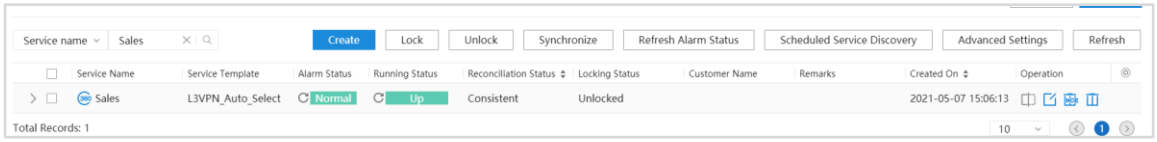

The alarm status and running status of the service are normal. PCEP and BGP-LS are required for the monitoring of the two statuses. Ensure that the PCEP session is normal and iMaster NCE-IP can receive BGP-LS routes from PEs.

In this view, you can click the 360-degree view icon before a service name to access the 360-degree service view. In the 360-degree service view, you can monitor the service information (including service alarms and paths) in real time, diagnose services, and perform active/standby switchovers.

Step 5 Check delivered configurations.

On PE1, check information about the VPN instance delivered by iMaster NCE-IP, binding relationship between the VPN instance and interface, BGP routes, and tunnel policy.

Check VPN instance and Loopback1 configurations.

```
[PE1]display current-configuration configuration vpn-instance VPNA
#
ip vpn-instance VPNA
 ipv4-family
  route-distinguisher 11:18
  tnl-policy NCE-VRF-VPNA
  apply-label per-instance
  transit-vpn
  vpn-target 200:52 export-extcommunity
  vpn-target 200:52 import-extcommunity
```

```
##

[PE1]display current-configuration interface LoopBack 1
#
interface LoopBack1
 ip binding vpn-instance VPNA
 ip address 172.16.1.1 255.255.255.255
#
```

Loopback1 has been bound to VPN instance VPNA, the RD and RTs have been automatically assigned to the VPN instance, and the VPN instance has been associated with tunnel policy NCE-VRF-VPNA.

Check tunnel policy configurations.

```
[PE1]display current-configuration configuration tunnel-policy
#
tunnel-policy NCE-VRF-VPNA
 tunnel select-seq sr-te load-balance-number 1
```

Check the VPN instance route

```
[PE1]display ip  routing-table  vpn-instance VPNA
Route Flags:  R -relay,D -downloadtofib,T -tovpn-instance, B -blackholeroute
------------------------------------------------------------------------------
RoutingTable: VPNA
         Destinations : 4        Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
     172.16.1.1/32  Direct  0    0           D   127.0.0.1       LoopBack1
     172.16.4.1/32  IBGP    255  0           RD  1.0.0.4         Tunnel6
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

An IBGP route exists in the VPN instance routing table on PE1. The outbound interface of the IBGP route is the Tunnel6 interface. This means that the VPN route from PE1 to the peer recurses to an SR-MPLS TE tunnel.

Check the label allocated by PE4 to the VPNv4 route.

```
[PE1]display bgp vpnv4 all routing-table label

BGP Local router ID is 1.0.0.1
Status codes: *-valid,> -best, d -damped,x-bestexternal,a -add path,
              h -history, i -internal, s-suppressed, S -Stale
              Origin : i -IGP, e -EGP, ?-incomplete
RPKI validationcodes: V -valid,I -invalid, N -not-found


Total number ofroutes from all PE: 2
Route Distinguisher: 11:19
```

```
      Network                 NextHop                 In/OutLabel

*>i   172.16.4.1              1.0.0.4                 NULL/48159
*i    172.16.4.1              1.0.0.4                 NULL/48159

Total Number of Routes: 0

VPN-Instance VPNA, Router ID 1.0.0.1:

Total Number of Routes: 2
      Network                 NextHop                 In/OutLabel

*>i   172.16.4.1              1.0.0.4                 NULL/48159
*i    172.16.4.1              1.0.0.4                 NULL/48159
```

The out label of the VPNv4 route 172.16.4.1 on PE1 is 48159, the label allocated to the route by BGP on PE4.

Test L3VPN connectivity on PE1.

```
[PE1]ping -vpn-instance VPNA -a 172.16.1.1 172.16.4.1
  PING 172.16.4.1: 56  data bytes, press CTRL_C to break
    Reply from 172.16.4.1: bytes=56 Sequence=1 ttl=253 time=1 ms
    Reply from 172.16.4.1: bytes=56 Sequence=2 ttl=253 time=1 ms
    Reply from 172.16.4.1: bytes=56 Sequence=3 ttl=253 time=1 ms
    Reply from 172.16.4.1: bytes=56 Sequence=4 ttl=253 time=1 ms
    Reply from 172.16.4.1: bytes=56 Sequence=5 ttl=253 time=1 ms

  --- 172.16.4.1 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 1/1/1 ms
```

The connectivity is normal.

3.1.2.11 SR Policy Delivery by the Controller

In this procedure, we will deploy an SR Policy through iMaster NCE-IP. Create a color for routes advertised by the L3VPN service to carry the color extended community attribute. The L3VPN traffic can then recurse to an SR Policy.

Step 1 Create an SR Policy color.

Open the Network Management app and choose **Configuration** > **Common** > **Profile Management** from the main menu. Then click **SR Policy Color Profile**.

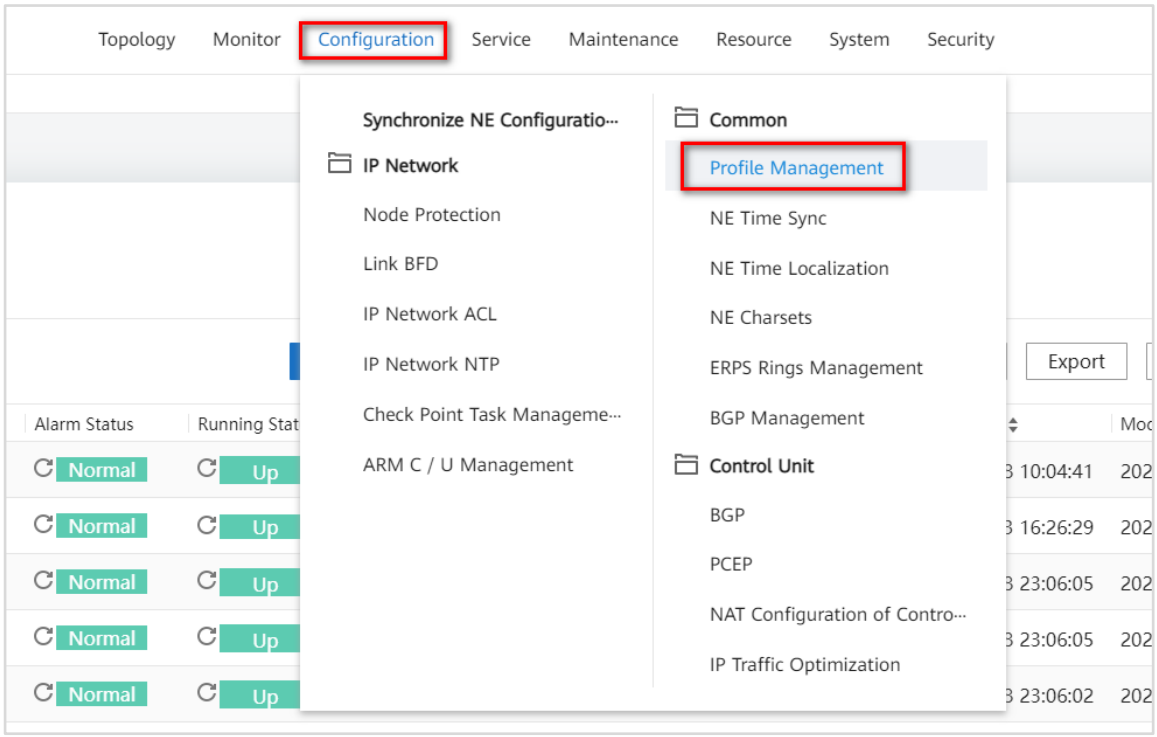
Topology
Monitor
Configuration
Service
Maintenance
Resource
System
Security
Synchronize NE Configuratio...
IP Network
Node Protection
Link BFD
IP Network ACL
IP Network NTP
Check Point Task Manageme...
ARM C / U Management
Common
Profile Management
NE Time Sync
NE Time Localization
NE Charsets
ERPS Rings Management
BGP Management
Control Unit
BGP
PCEP
NAT Configuration of Contro...
IP Traffic Optimization
Export
Alarm Status
Running Stat
Normal
Up
10:04:41
16:26:29
23:06:05
23:06:05
23:06:02

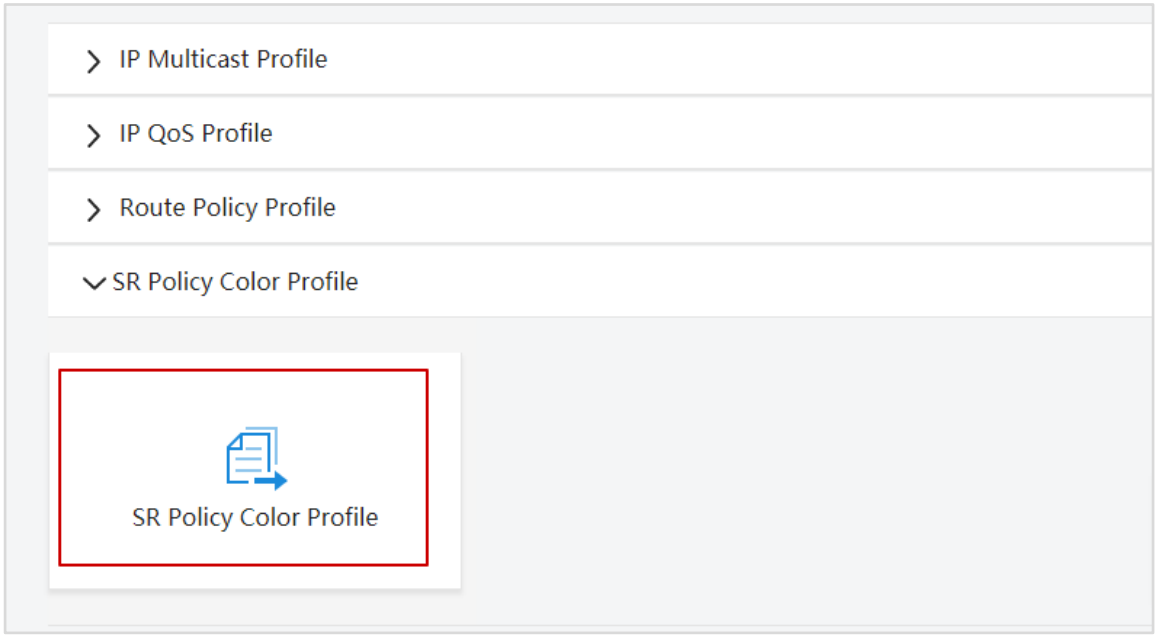
IP Multicast Profile
IP QoS Profile
Route Policy Profile
SR Policy Color Profile
SR Policy Color Profile

Create a color named **PE1_PE4_L3VPN**.

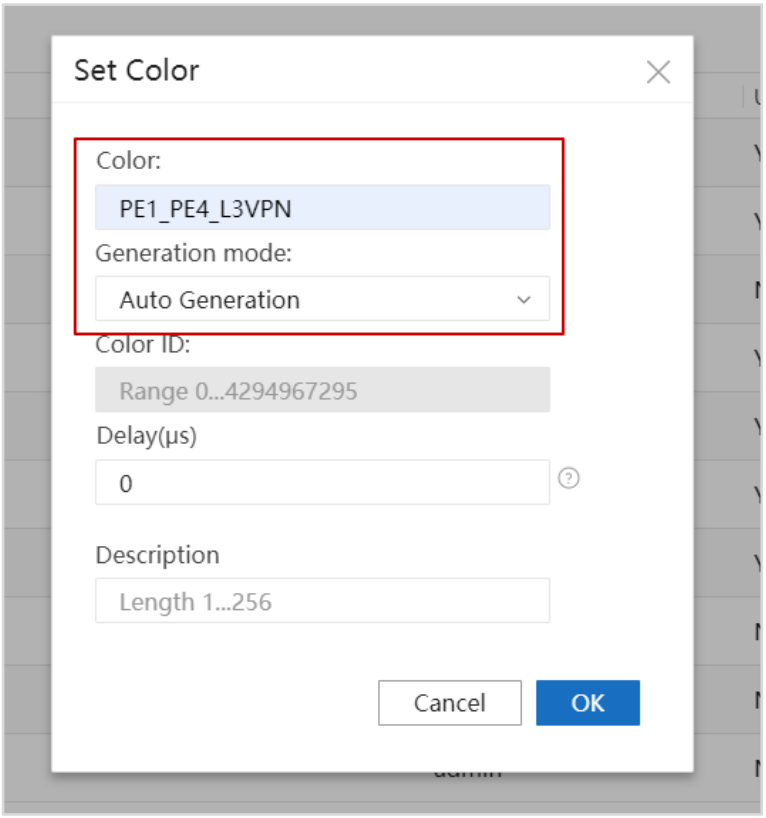

Step 2 Configure an SR Policy.

Before performing this step, ensure that the BGP-LS and BGP SR Policy peer relationships have been established between iMaster NCE-IP and RRs and the following configurations are ready on the NE side:

1. IGP route reachability is available network-wide.
2. MPLS and MPLS TE are enabled both globally and per interface.
3. IGP TE is enabled.
4. SR is enabled globally, and IGP extensions for SR capabilities are enabled.

Open the Network Management app and choose **Service** > **Create** > **SR Policy** from the main menu to create an SR Policy.

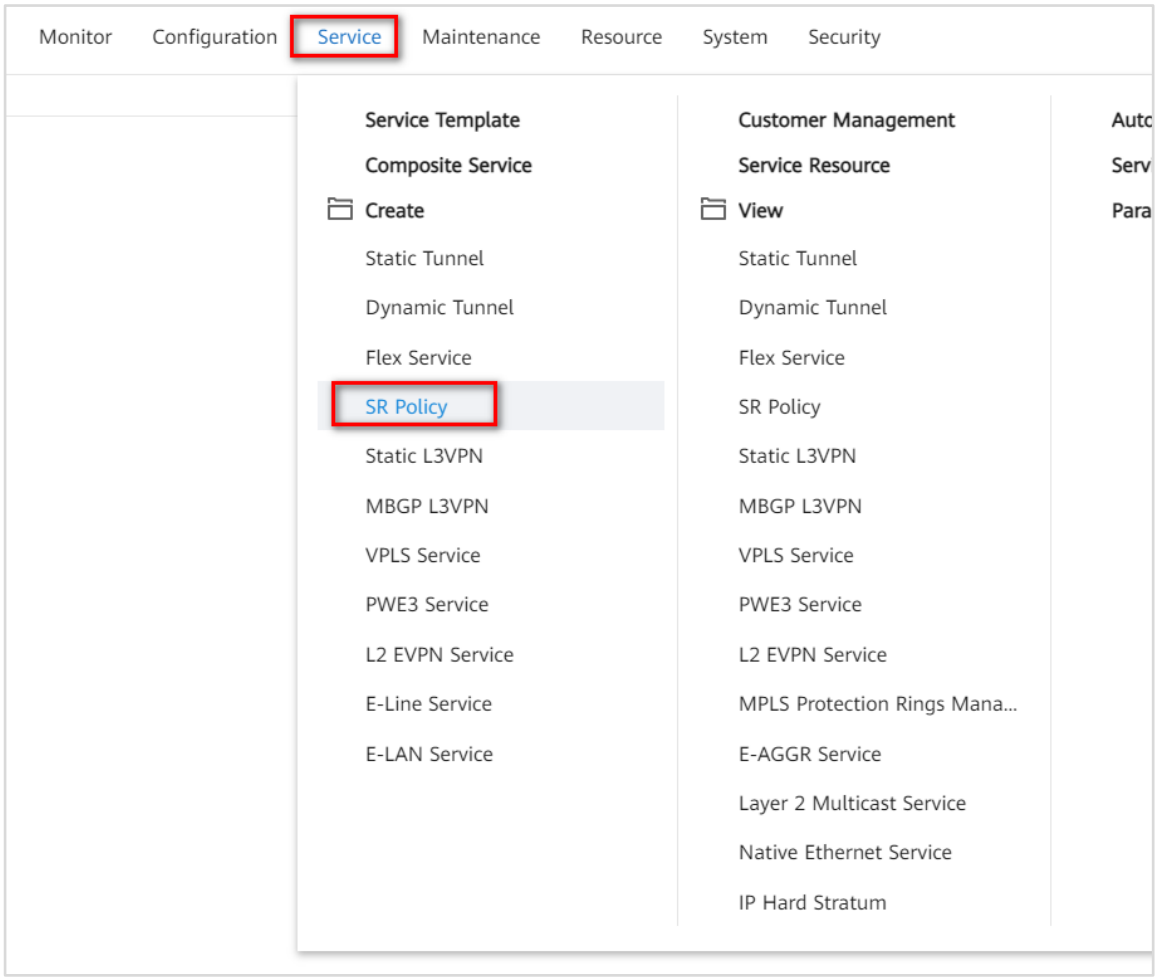

Set **Select a service template** to **SR_TE_Policy**.

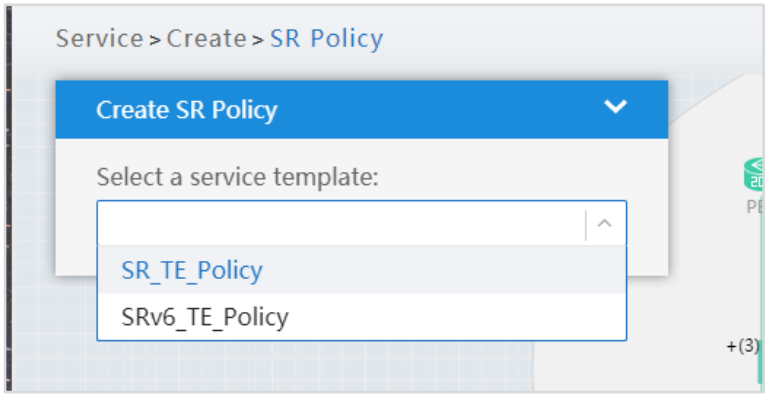

In the **Basic Information** area, set **Parameter template** to **SR_TE_Policy**, **Service name** to **PE1_PE4_L3VPN**, and retain the default values for other parameters.

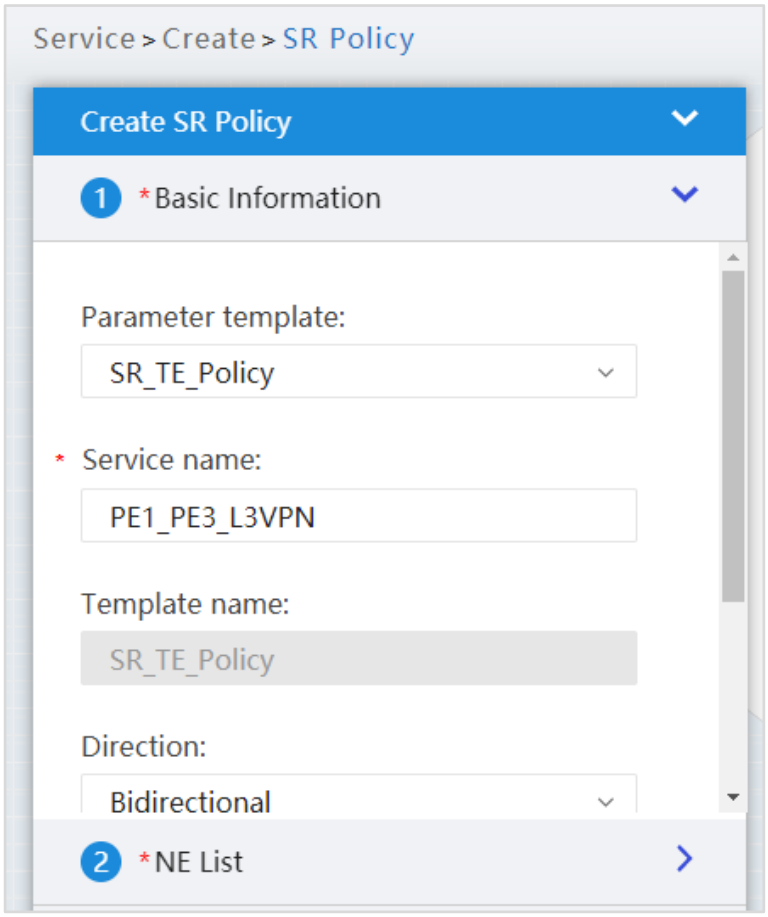

In the **NE List** area, set **Source NE** to **PE1** and **Sink NE** to **PE4**.

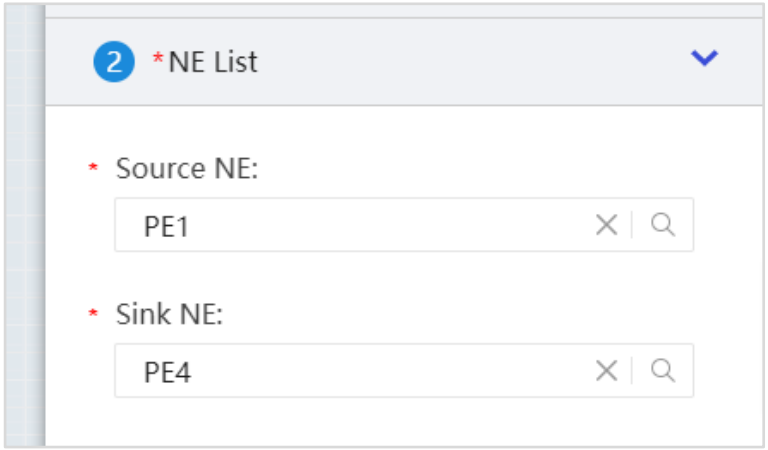

On the **Forward** tab page in the **SR Policy** area, click **Modify** next to **Color**.

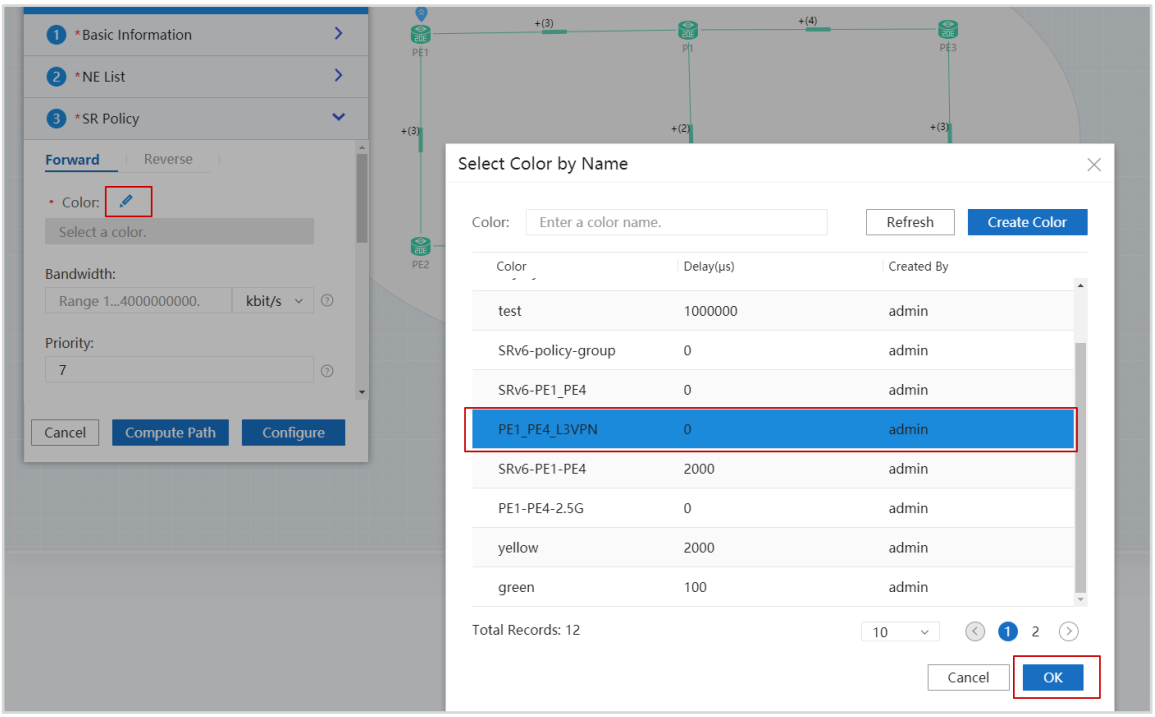

Select the previously created color PE1_PE4_L3VPN.

Configure candidate path preference.

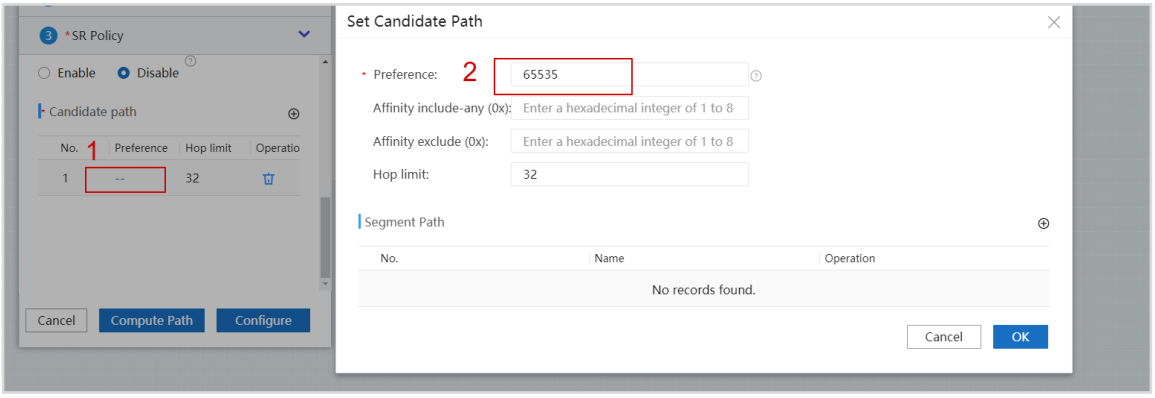

Click the field marked by **1** to access the preference configuration list and set the preference of candidate path 1 to the highest value 65535.

Click **Compute Path**.

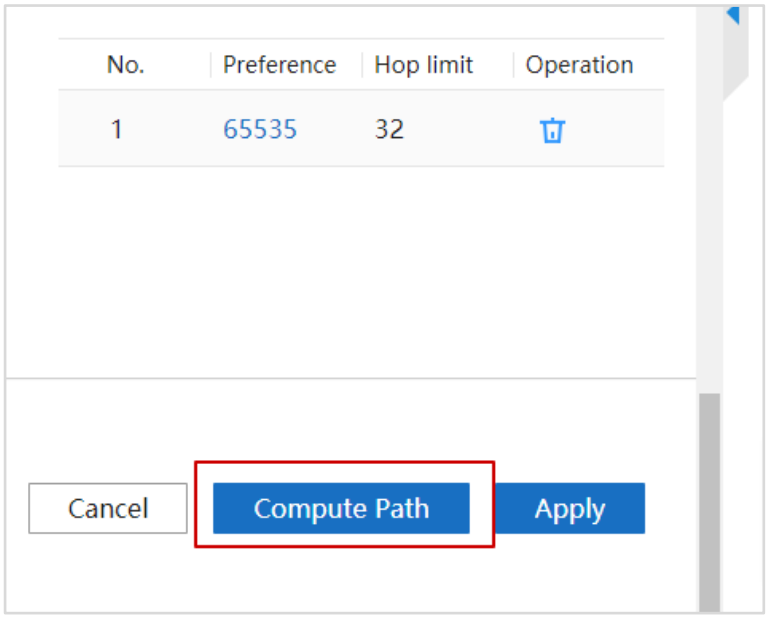

In the computation results, still select the path with the least cost.

Click **OK**, and then click **Apply**.

In the dialog box that is displayed, click **OK**.

In the **Success** dialog box, you can click the corresponding hyperlink to view SR Policy information.

The SR Policy status is normal.

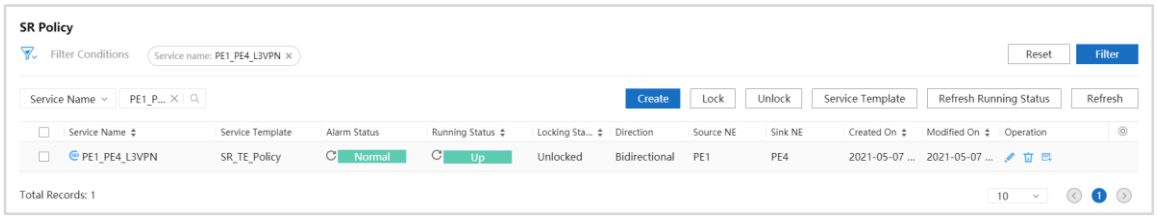

Step 3 Configure route-policies.

To recurse L3VPN traffic to SR Policies, configure a route-policy on PE1 and PE4 to add the color extended community attribute to VPNv4 routes to be advertised, so that these routes can recurse to SR Policies.

Open the Network Management app and choose **Configuration** > **Common** > **Profile Management** from the main menu. Then click **Routing Policy Template**.

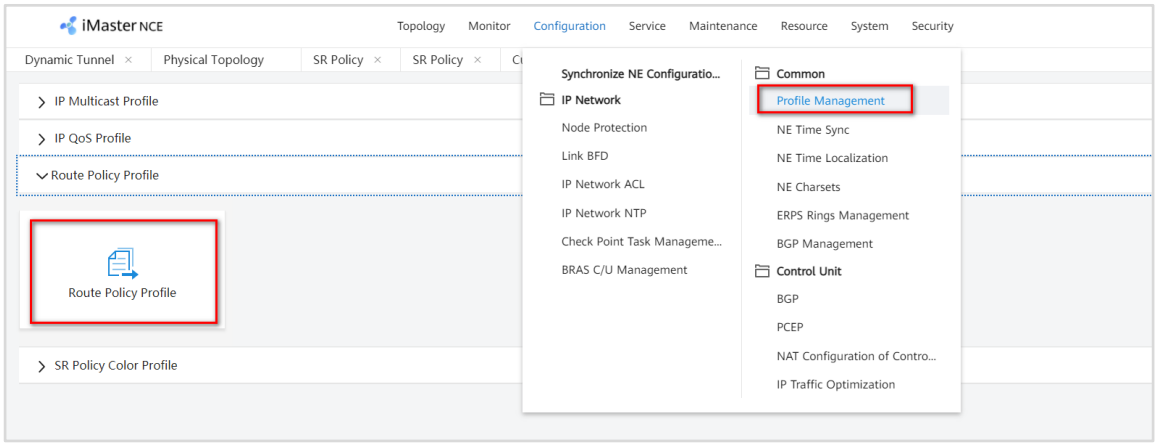

Click **Create** to create a route-policy template.

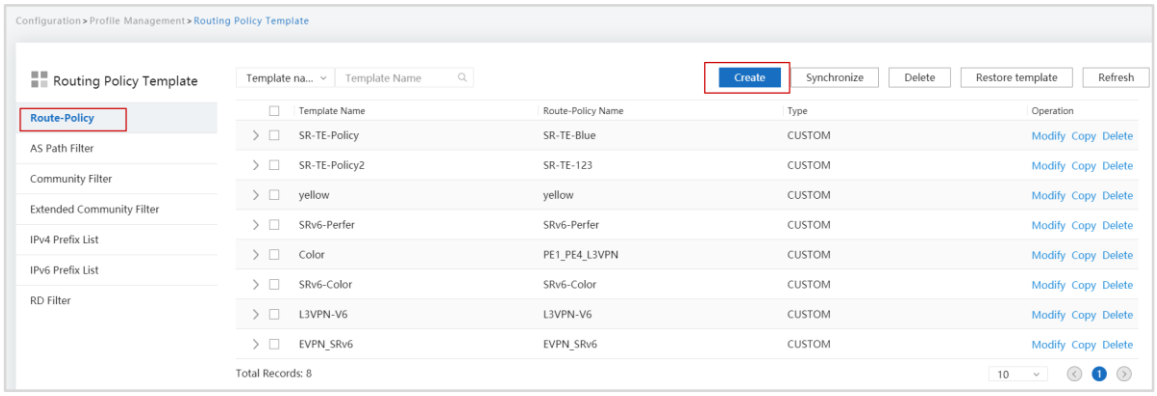

Set the template name to **Color** and route-policy name to **PE1_PE4_L3VPN**.

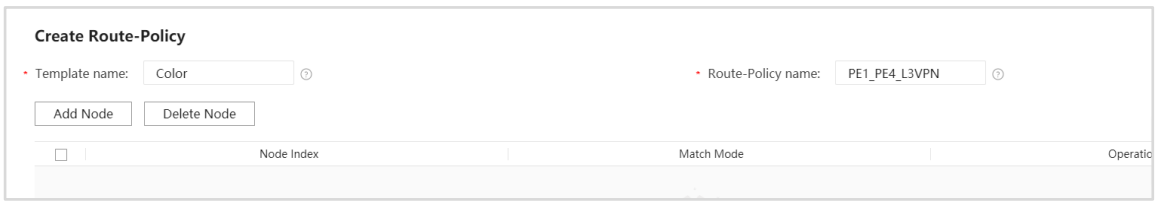

Click **Add Node**.

In the dialog box that is displayed, configure node information. Specifically, set **Node index** to **10**, retain the default value **Permit** for **Match mode**, and leave all match conditions unselected (indicating that all match conditions will be applied).

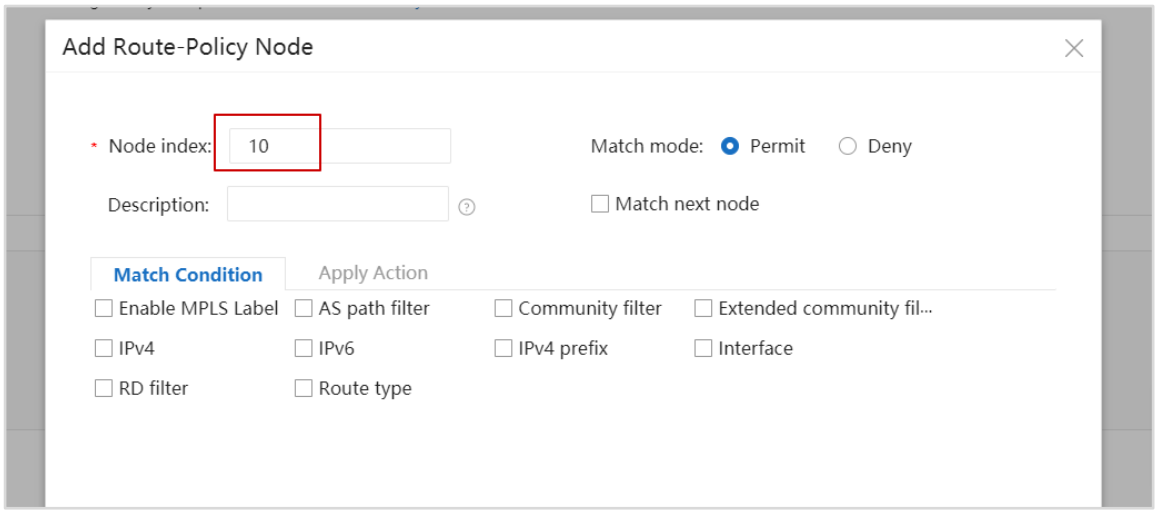

On the **Apply Action** tab page, select **Color**.

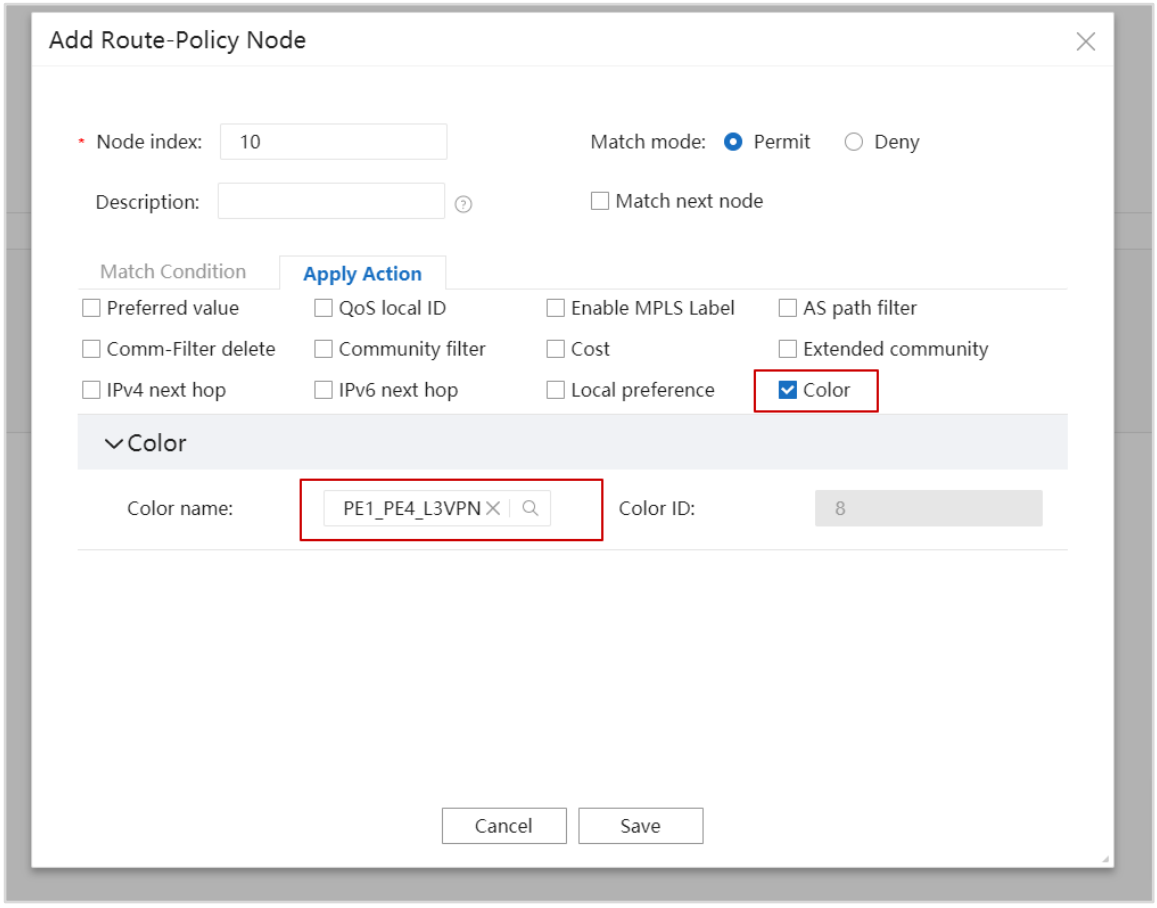

Select the previously created color **PE1_PE4_L3VPN** and click **Save**. Then click **OK**. The template is created.

3.1.2.12 L3VPN Service Delivery by the Controller

Use the controller to deliver a new L3VPN service and recurse the service to an SR Policy for traffic forwarding.

Create Loopback2 on PE1 and PE4 to simulate L3VPN access users.

PE1

```
[PE1]interface LoopBack2
[PE1-LoopBack2] ip address 192.168.1.1 32
[PE1-LoopBack2] quit
```

PE4

```
[PE4]interface LoopBack2
[PE4-LoopBack2] ip address 192.168.4.1 32
[PE4-LoopBack2] quit
```

After configuring loopback interface addresses on PE1 and PE4, synchronize NE configurations to iMaster NCE-IP.

Open the Network Management app and choose **Service** > **Create** > **MBGP L3VPN** from the main menu.

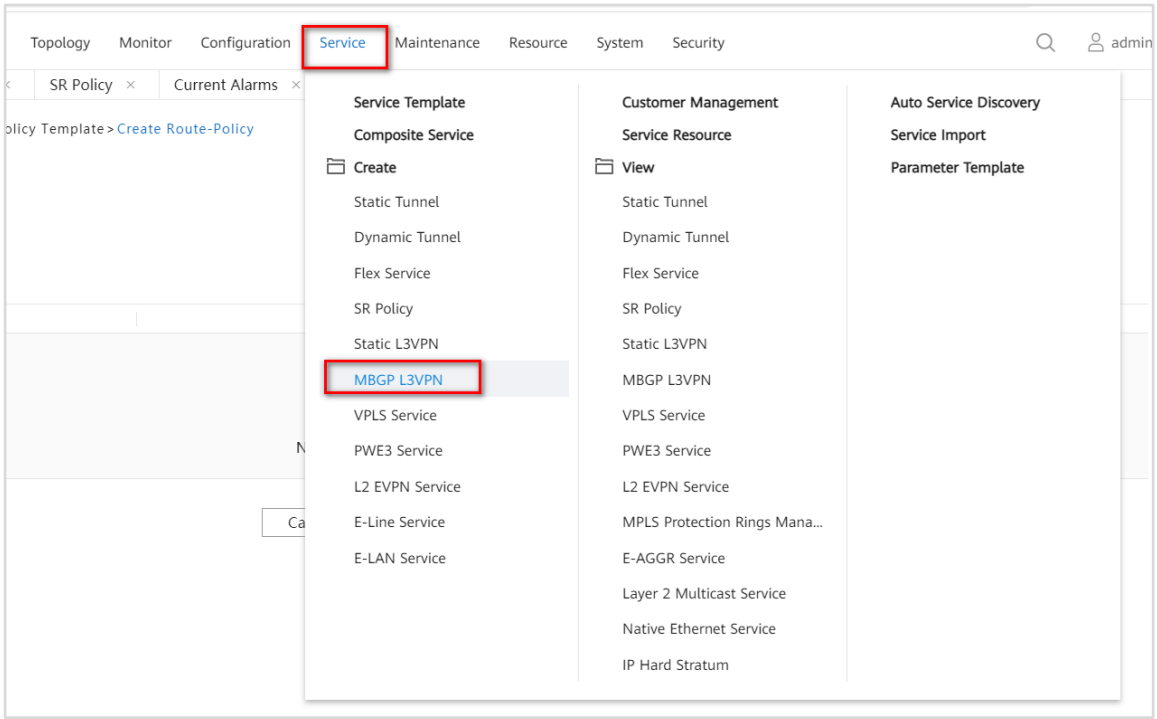

Step 1 Set basic parameters.

In the **Basic Parameter** area, set **Service template** to **L3VPN_Auto_Select** (default value), set **Service name** and **VRF Name**, set **Parameter template** to **MBGP_L3VPN**, and enable **Auto allocate RD and RT: Yes**.

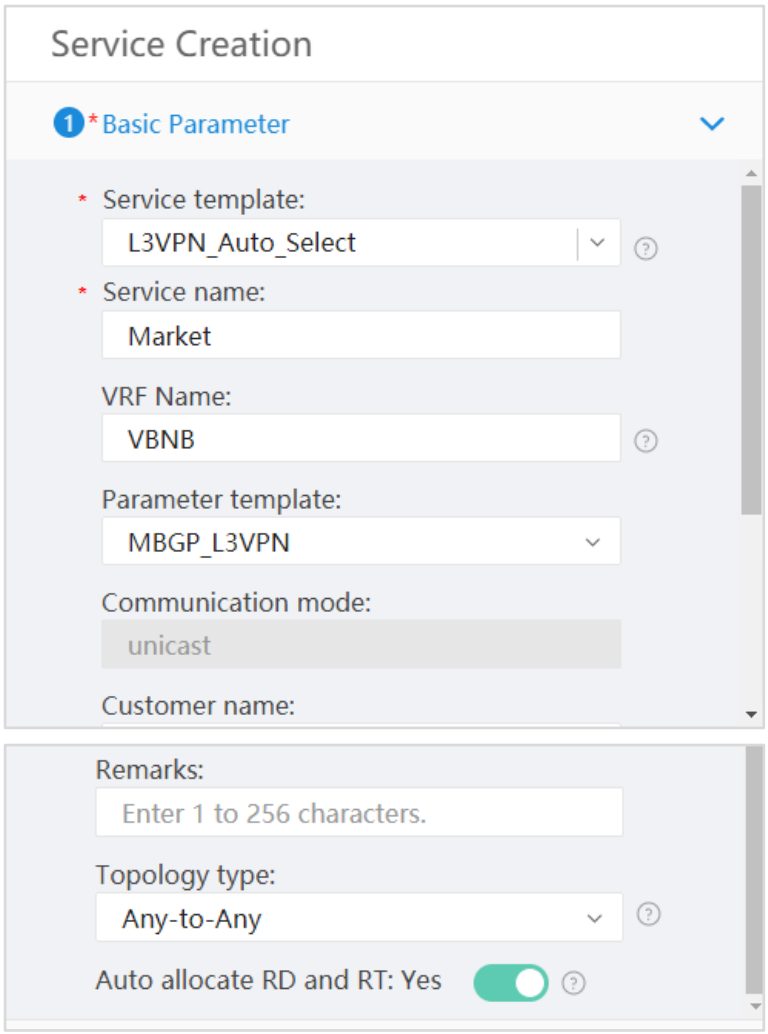

Step 2 Configure service nodes.

In the **Service Node** area, select PE1 and PE4 as service nodes and select a tunnel policy.

Click **+** in the **Service Node** area.

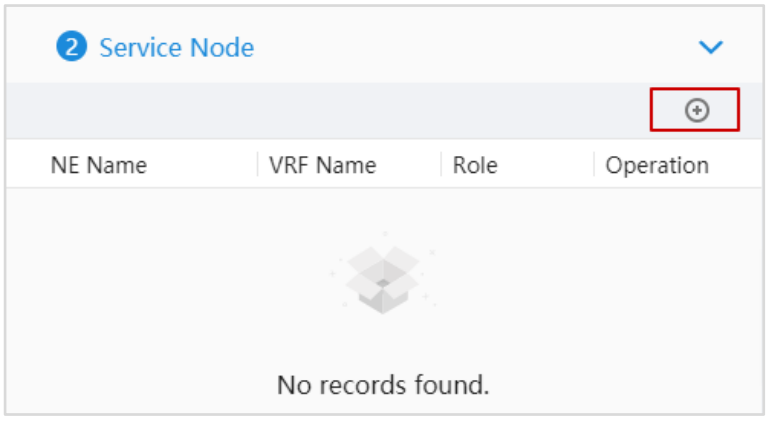

In the dialog box that is displayed, set **NE name** to **PE1** and select **Enable IPv4**.

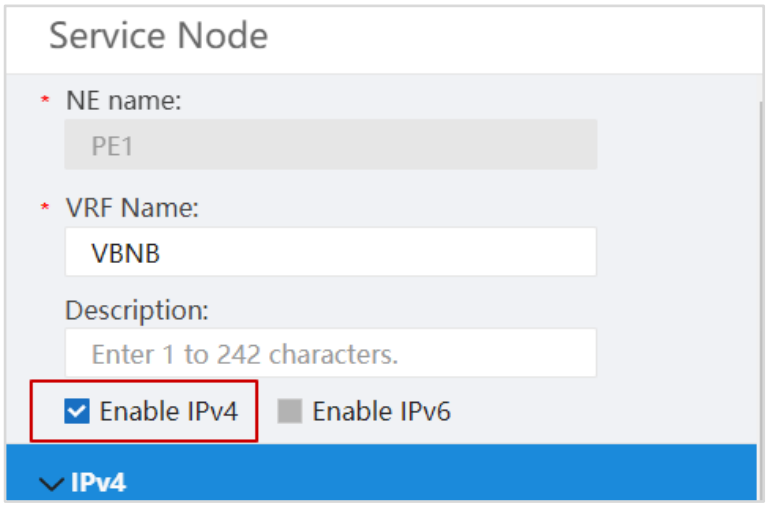

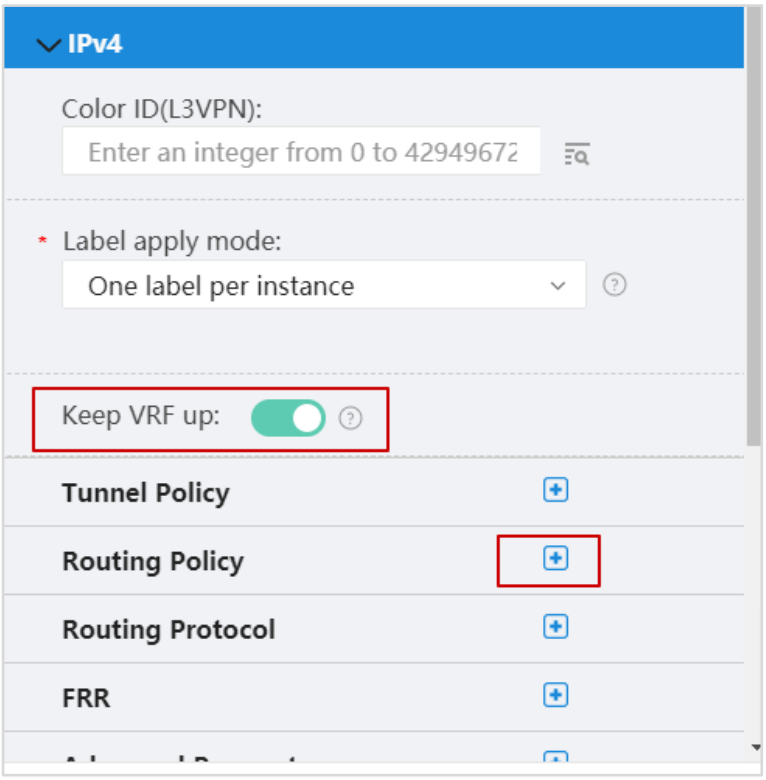

Enable **Keep VRF up**. In the area that is displayed, click **+** next to **Routing Policy**.

In the dialog box that is displayed, set **Direction** to **Out** and **Routing Policy Template** to **Color** for the route-policy, and then click **OK**.

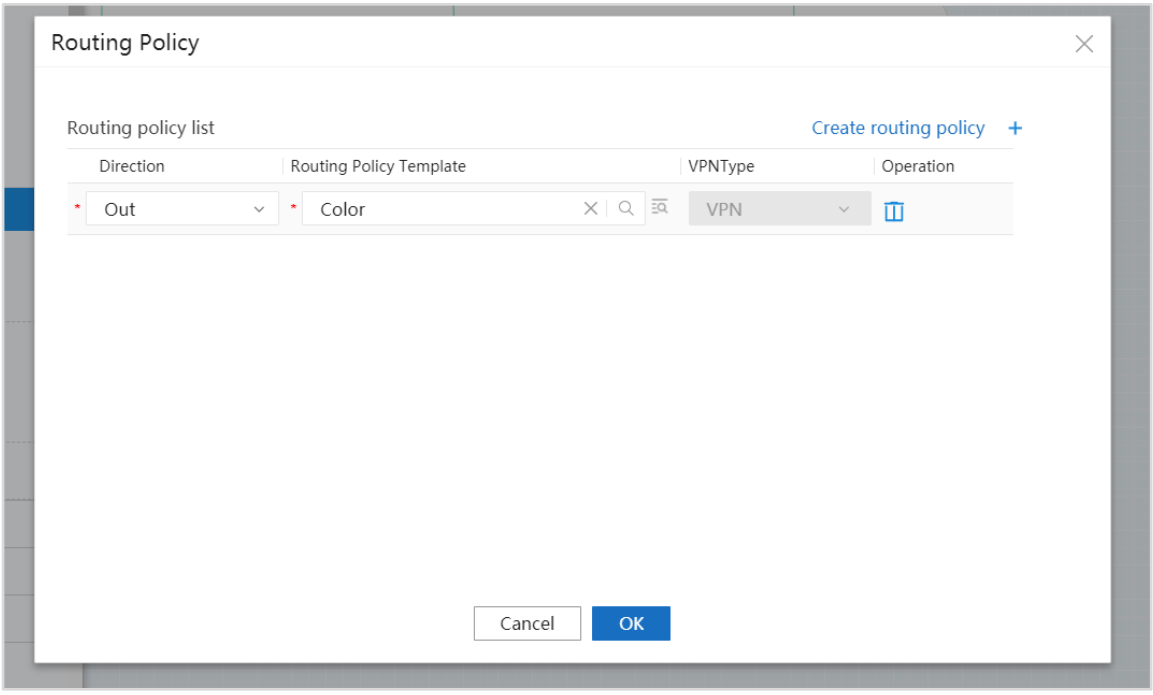

Click **+** next to **Tunnel Policy**. In the dialog box that is displayed, set **Policy type** to **Auto select**, select **Unmix**, and set **Type** to **SR-MPLS TE Policy**.

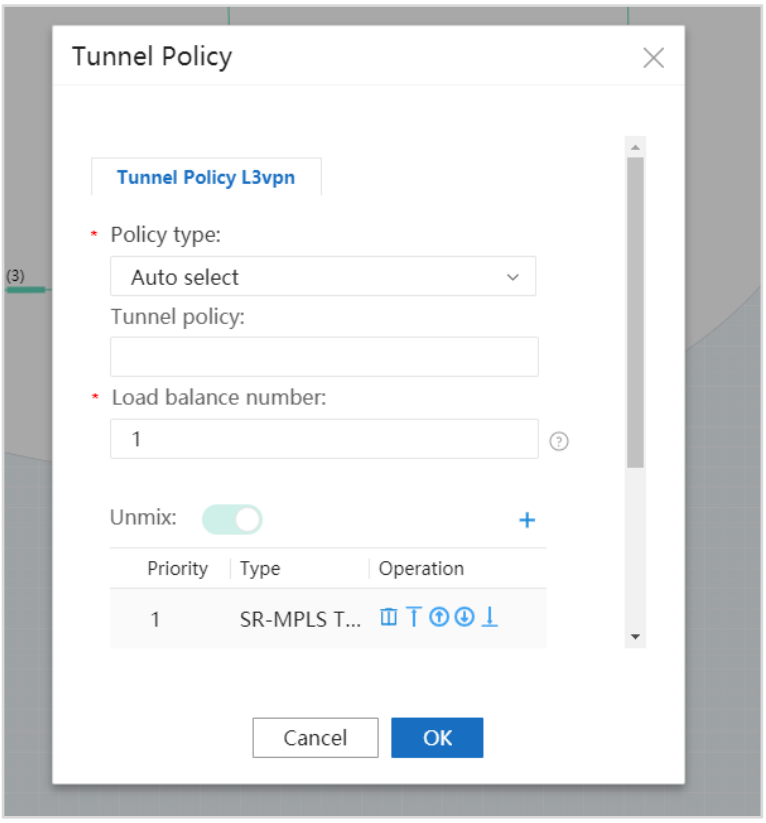

The service node configuration on PE1 is complete. The service node configuration on PE4 is similar to that on PE1.

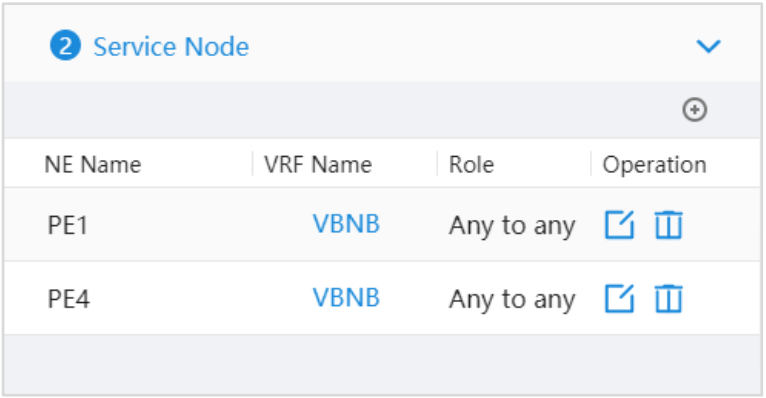

Step 3 Configure service access points.

Configure service access points for PE1 and PE4. Here, Loopback2 interfaces are used to simulate user access.

Click **+** in the **Service Access Point** area.

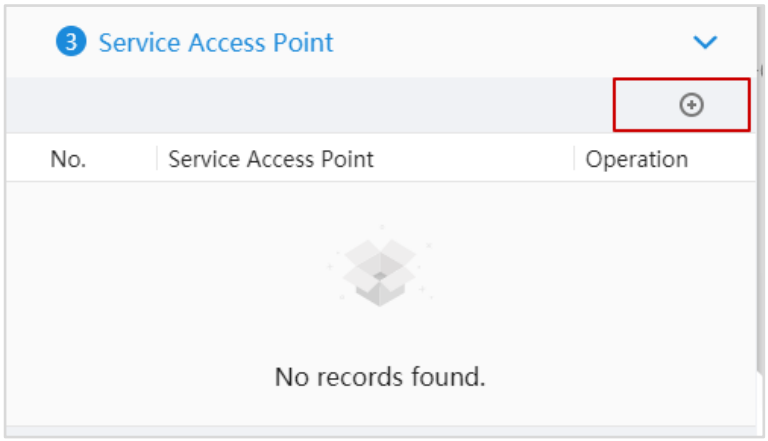

In the **Service Access Point** dialog box, set **Name** to **PE1_Tenant** for PE1, retain the default single-homing access mode, and click **+**.

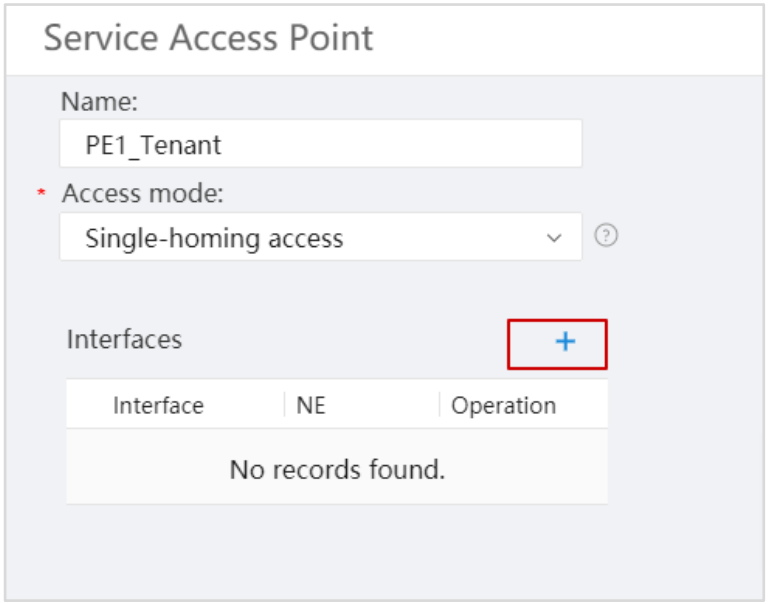

In the dialog box that is displayed, set **NE** to **PE1|VPNB** and **Interface** to **Loopback2**, and select **Enable IPv4**.

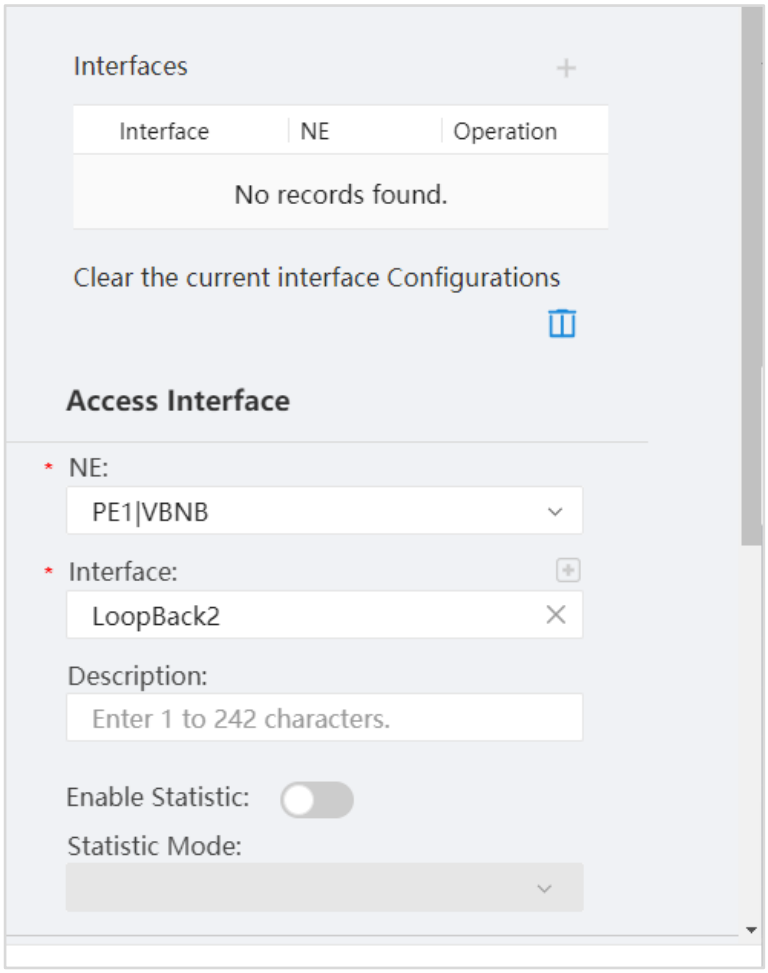
Interfaces
Interface
NE
Operation
No records found.
Clear the current interface Configurations
Access Interface
NE:
PE1|VBNB
Interface:
LoopBack2
Description:
Enter 1 to 242 characters.
Enable Statistic:
Statistic Mode:

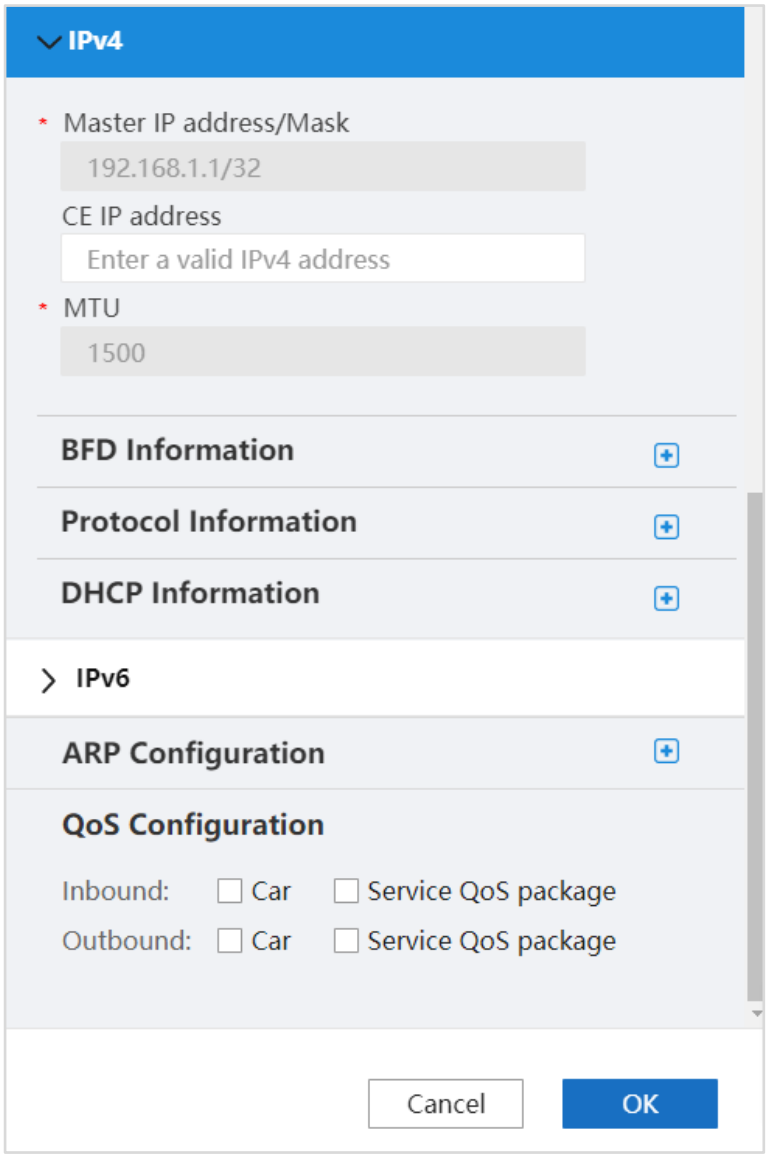

Finally, click **OK**. The service access point PE1_Tenant is configured. The service access point configuration on PE4 is similar to that on PE1.

The service access point configurations (CE-related configurations) on PE1 and PE4 are complete.

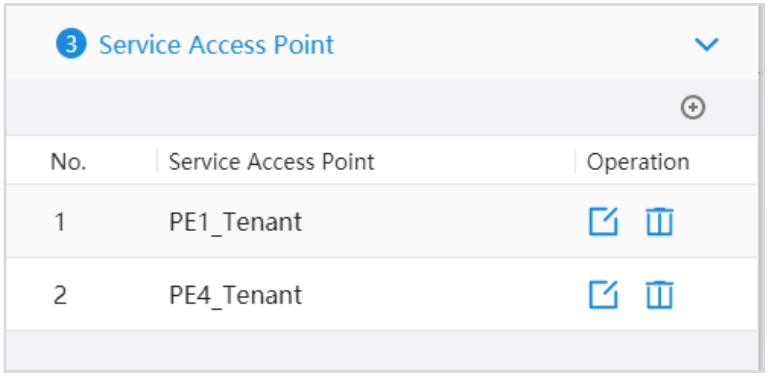

Step 4 Configure tunnels.

Associate the VPN instance with the SR Policy.

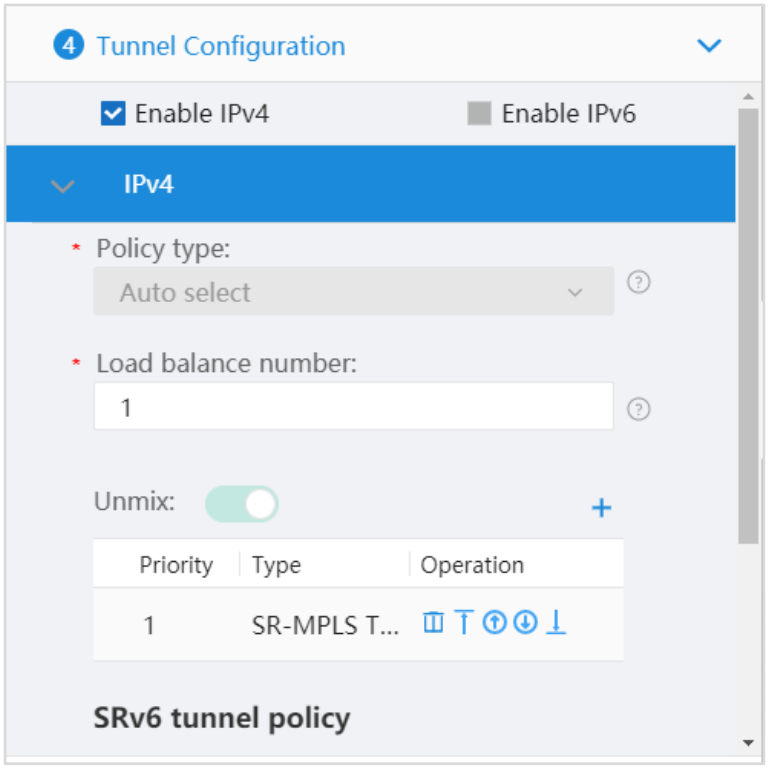

Select **Enable IPv4**, enable **Unmix**, and click **+**. In the dialog box that is displayed, set the tunnel type to **SR-MPLS TE Policy**.

Finally, click **Configure** and wait for the L3VPN service configurations to be delivered to devices.

In the dialog box displayed to indicate service delivery success, click the corresponding hyperlink to view information about the newly created service.

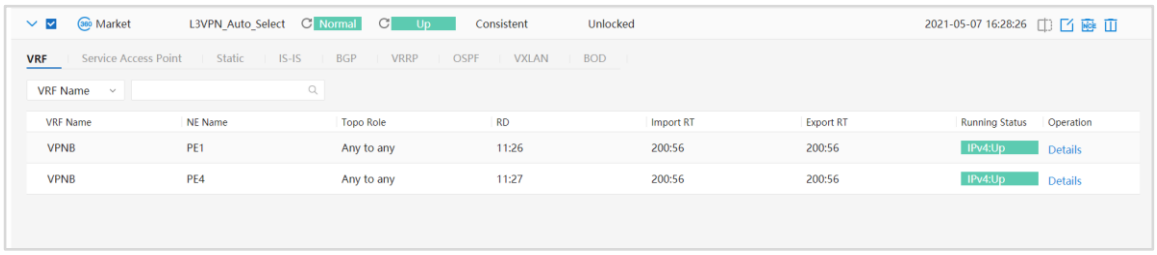

Step 5 Check delivered configurations.

Check VPN instance and Loopback2 configurations.

```
[PE1]display current-configuration configuration vpn-instance VPNB
#
```

```
ip vpn-instance VPNB
 ipv4-family
  route-distinguisher 11:24
  export route-policy PE1_PE4_L3VPN
  tnl-policy NCE-VRF-VPNB
  apply-label per-instance
  transit-vpn
  vpn-target 200:55 export-extcommunity
  vpn-target 200:55 import-extcommunity
#
[PE1]display current-configuration interface LoopBack 2
#
interface LoopBack2
 ip binding vpn-instance VPNB
 ip address 192.168.1.1 255.255.255.255
#
```

Check tunnel information.

```
[PE1]display tunnel-info all
Tunnel ID                Type                 Destination          Status
--------------------------------------------------------------------------------
0x000000000300002002 sr-te                1.0.0.4              UP
0x000000002900000003 srbe-lsp             1.0.0.2              UP
0x000000002900000004 srbe-lsp             1.0.0.4              UP
0x000000002900000005 srbe-lsp             1.0.0.6              UP
0x000000002900000008 srbe-lsp             1.0.0.5              UP
0x000000002900000009 srbe-lsp             1.0.0.3              UP
0x000000003200048001 srtepolicy           1.0.0.4              UP
```

In this case, an SR-TE Policy is delivered.

Check tunnel information.

```
[PE1]display tunnel-info 0x000000003200048001
Tunnel ID:            0x000000003200048001
Type:                 srtepolicy
Name:                 SR-TE Policy
Destination:          1.0.0.4
Instance ID:          0
Cost:                 0
Status:               UP
Color:                8
Group:                0
```

The color whose ID is 8 corresponds to PE1_PE4_L3VPN configured on the controller.

Check the routing table of VPN instance VPNB.

```
[PE1]display ip routing-table vpn-instance VPNB
Route Flags: R - relay, D - download to fib, T - to vpn-instance, B - black hole route
------------------------------------------------------------------------------
Routing Table: VPNB
         Destinations : 4        Routes : 4
```

```
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        127.0.0.0/8 Direct  0    0           D   127.0.0.1       InLoopBack0
      192.168.1.1/32 Direct  0    0           D   127.0.0.1       LoopBack2
      192.168.4.1/32 IBGP    255  0           RD  1.0.0.4         SR-TE Policy
  255.255.255.255/32 Direct  0    0           D   127.0.0.1       InLoopBack0
```

The outbound interface of the route to 192.168.4.1 is an SR-TE Policy, not a specific tunnel interface.

Check SR-TE Policy information.

```
[PE1]display sr-te policy
PolicyName :
Endpoint                : 1.0.0.4                       Color                 : 8
TunnelId                : 1                             TunnelType            : SR-TE Policy
Binding SID             : -                             MTU                   : -
Policy State            : Up                            State Change Time     : 2021-04-02
08:25:58
Admin State             : UP                            Traffic Statistics    : Disable
BFD                     : Disable                       Backup Hot-Standby    : Enable
DiffServ-Mode           : -
Candidate-path Count    : 1

Candidate-path Preference:  65535
PathState               : Active                        Path Type             : Primary
Protocol-Origin         : BGP(20)                       Originator            : 65001,
172.21.17.102
Discriminator           : 87                            Binding SID           : -
GroupId                 : 1                             Policy Name           :
Template ID             : 4294967274
Segment-List Count      : 1
Segment-List            :
  Segment-List ID       : 1                             XcIndex               : 2000001
  List State            : Up                            BFD State             : -
  EXP                   : 0                             TTL                   : 0
  DeleteTimerRemain     : -
  Label: 48091, 48091, 48090
```

The forwarding labels are 48091, 48091, and 48092.

Check BGP SR Policy routes.

```
[PE1]display bgp sr-policy routing-table

BGP Local router ID is 1.0.0.1
Status codes: *-valid, > -best, d -damped, x-best external, a -add path,
              h -history, i -internal, s-suppressed, S -Stale
              Origin : i -IGP, e -EGP, ?-incomplete
RPKI validation codes: V -valid, I -invalid, N -not-found


Total Number of Routes: 2
         Network            Nexthop         MED         LocPrf    PrefVal Path/Ogn
*>i    [15][8][1.0.0.4]     172.21.17.102   4294967286  100       0       ?
```

```
*i                    172.21.17.102    4294967286   100          0          ?
```

Check BGP SR Policy route details.

```
[PE1]display bgp sr-policy routing-table [15][8][1.0.0.4]

 BGP local router ID : 1.0.0.1
 Local AS number : 65001
 Paths:    2 available, 1 best, 1 select, 0 best-external, 0 add-path
 BGP routing table entry information of [15][8][1.0.0.4]:
 From: 1.0.0.5 (1.0.0.5)
 Route Duration: 0d00h54m17s
 Relay IP Nexthop: 172.21.17.102
 Relay IP Out-Interface:GigabitEthernet0/0/0
 Original nexthop: 172.21.17.102
 Qos information : 0x0
 Ext-Community: RT <1.0.0.1 : 0>, SoO <172.21.17.102 : 0>
 AS-path Nil, origin incomplete, MED 4294967283, localpref 100, pref-val 0, valid, internal, best,
select, pre 255
 Originator: 172.21.17.102
 Cluster list: 1.0.0.5
 Tunnel Encaps Attribute (23):
 Tunnel Type: SR Policy (15)
  Preference: 65535
  Segment List
    Weight: 1
    Path MTU: 9600
    Segment: type:1, Label:48091
    Segment: type:1, Label:48091
    Segment: type:1, Label:48090
  Template ID: 4294967274
 Not advertised to any peer yet
```

The forwarding path information of the BGP SR Policy is displayed.

Test L3VPN connectivity on PE1.

```
[PE1]ping -vpn-instance VPNB -a 192.168.1.1 192.168.4.1
  PING 192.168.4.1: 56  data bytes, press CTRL_C to break
    Reply from 192.168.4.1: bytes=56 Sequence=1 ttl=253 time=1 ms
    Reply from 192.168.4.1: bytes=56 Sequence=2 ttl=253 time=1 ms
    Reply from 192.168.4.1: bytes=56 Sequence=3 ttl=253 time=1 ms
    Reply from 192.168.4.1: bytes=56 Sequence=4 ttl=253 time=1 ms
    Reply from 192.168.4.1: bytes=56 Sequence=5 ttl=253 time=1 ms

  --- 192.168.4.1 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 1/1/1 ms
```

The connectivity is normal.

3.1.2.13 Tunnel Optimization

The controller performs tunnel optimization to ensure that delivered tunnels can adjust forwarding paths based on the real-time network delay and bandwidth to dynamically meet service requirements.

Here, the controller is used to perform tunnel optimization based on link delay.

Step 1 Perform basic configurations.

(Optional) Open the Network Path Navigation app and choose **Settings** > **Basic Settings** > **Topology Collection Policy** from the main menu. Then enable **Non-TE link real-time collection**.

This option can be enabled when non-TE links exist in tunnel traffic forwarding.

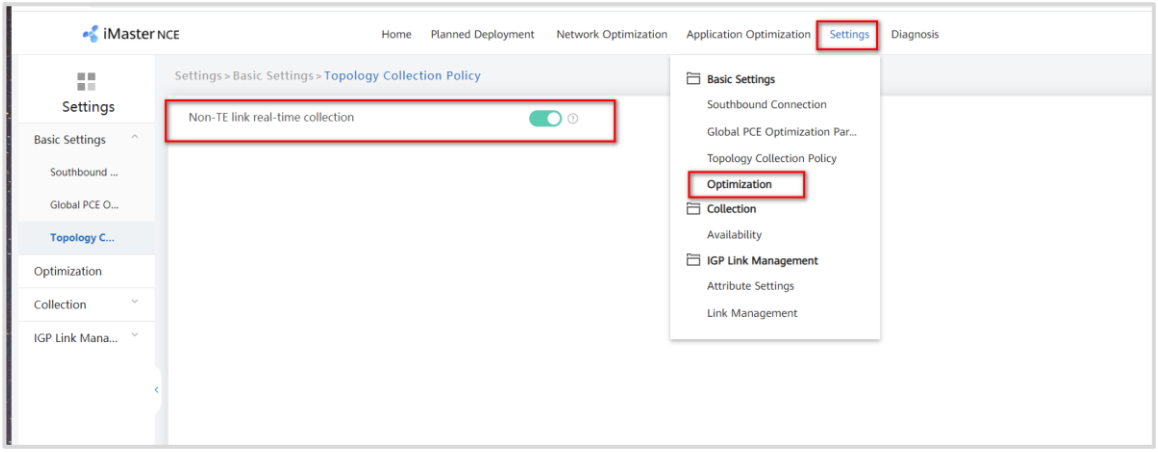

Change the global optimization policy to **Minimum delay**.

Open the Network Path Navigation app and choose **Settings** > **Basic Settings** > **Global Path Computation Parameters** from the main menu. Then change the optimization policy to **Minimum delay**.

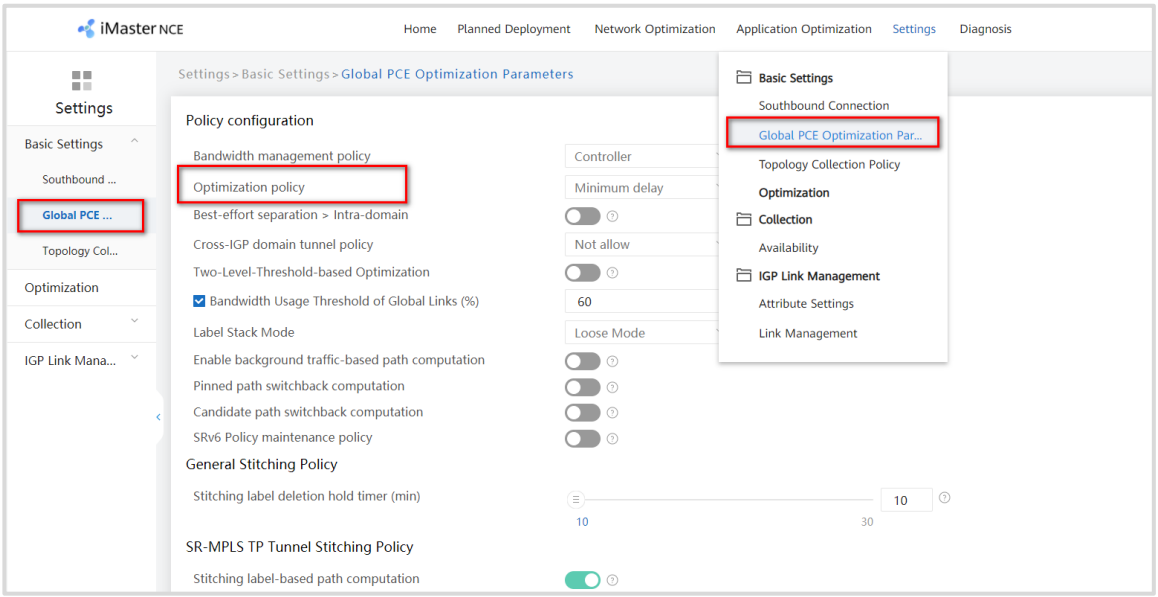

Finally, click **Apply** in the lower part of the page. The controller then uses delay as the criterion for determining tunnel link quality. The controller computes the delay of each forwarding path for the tunnel and selects the path with the minimum delay as the new forwarding path for tunnel optimization.

The controller computes delay in the following ways:

1. Static value: By default, the two-way delay is 200 μs for each interface. You can choose **Settings** > **Attribute Settings** from the main menu of the Network Path Navigation app to change the value.
2. Dynamically measured value: TWAMP is used to measure the two-way delay of a link. The measurement result is updated in real time. You can choose **Settings** > **Test Case Management** from the main menu of the Network Performance Analysis app to create TWAMP test instances and use the test instances to measure the forwarding link delay of a tunnel.

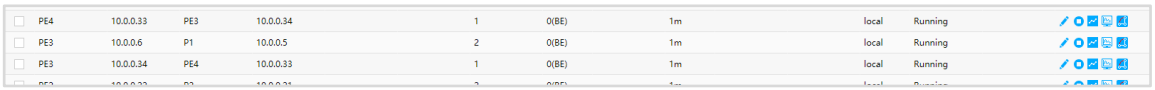

A different TWAMP test instance can be created for each direction of a link to implement two-way delay measurement.

The following describes how to perform manual tunnel optimization by modifying static link delay and how to perform automatic tunnel optimization by configuring TWAMP test instances to modify and monitor link delay.

Step 2 Perform manual optimization.

Manually change link delay. The default delay is 200 μs for each link. Manually change the link delay between P1 and PE3 to 2000 μs, so that the controller switches the L3VPN service between PE1 and PE4 to another path.

Check the delay of the SR-MPLS TE Policy between PE1 and PE4.

Open the Network Path Navigation app and choose **Network Optimization** > **Network Optimization** from the main menu. On the **Tunnel List** tab page, click **Filter** and use **Tunnel Type** as the filter criterion to find the previously created SR-MPLS TE Policy.

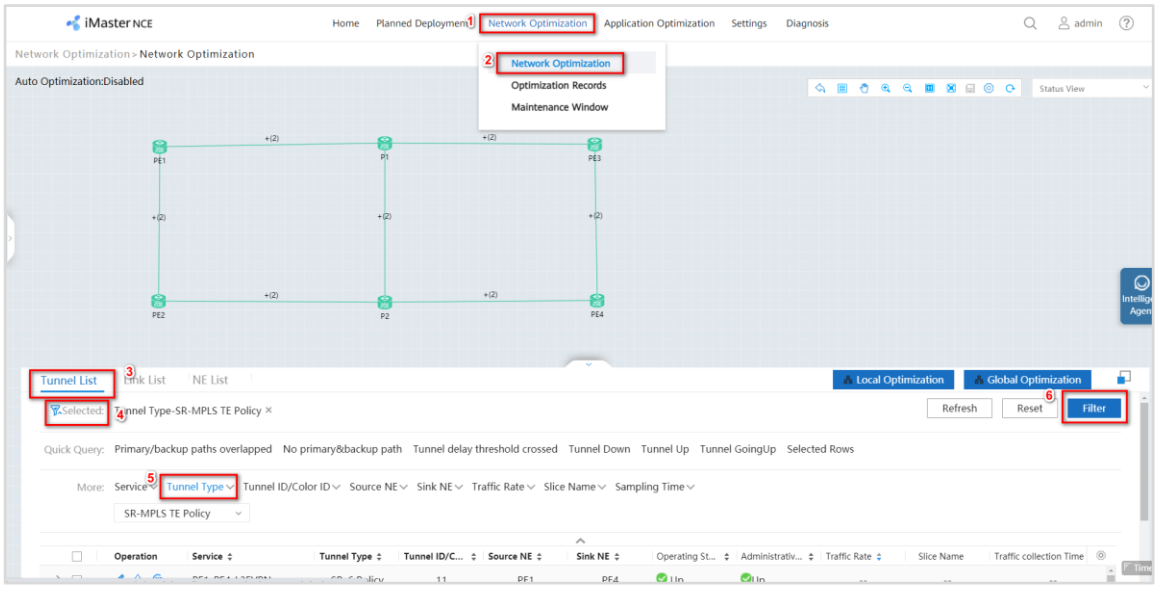

A unidirectional tunnel from PE1 to PE4 and a unidirectional tunnel from PE4 to PE1 can be found, indicating that a bidirectional tunnel is established between PE1 and PE4.

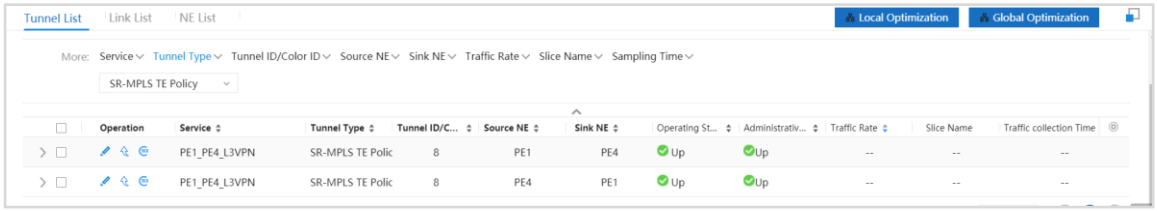

Click > before PE1_PE4_L3VPN to view the current path and accumulated delay.

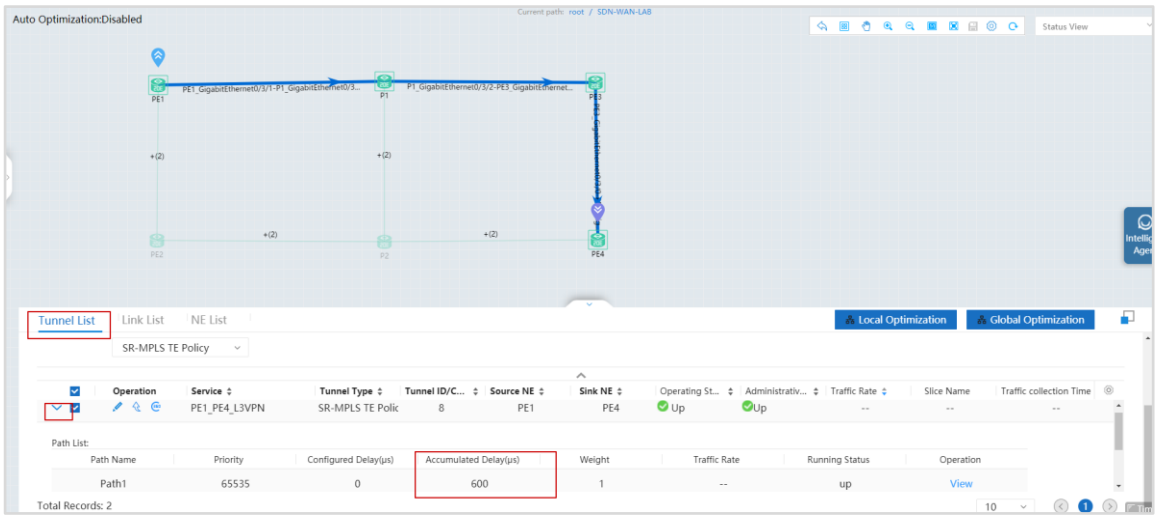

The forwarding path from PE1 to PE4 is PE1 -> P1 -> PE3 -> PE4, and the accumulated delay is 600 μs. The reverse path from PE4 to PE1 is PE4 -> PE3 -> P1 -> PE1, and the accumulated delay is also 600 μs.

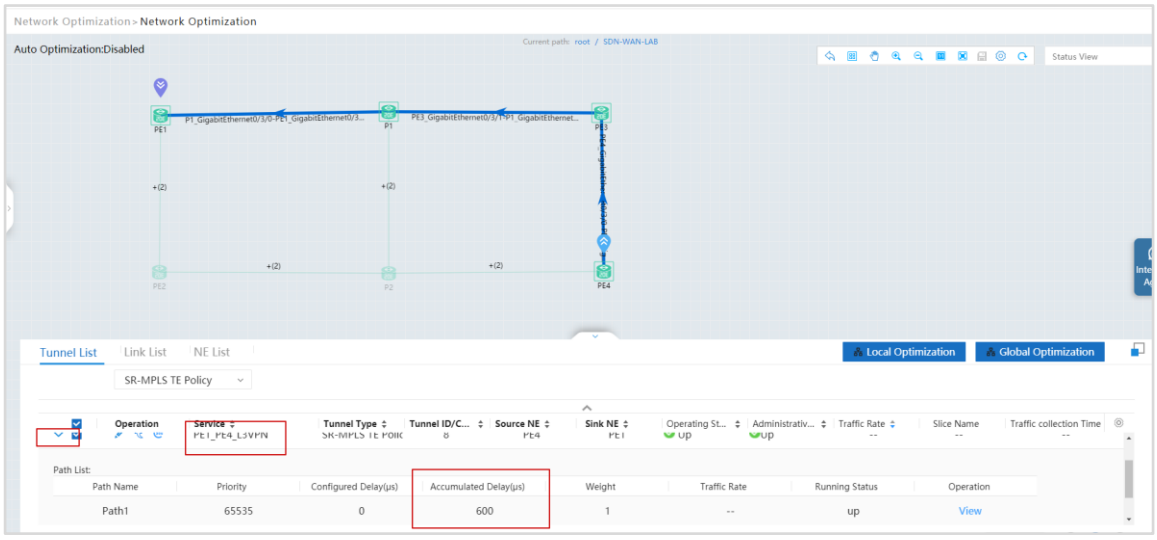

Change the delay of links between P1 and PE3 to 2000 μs.

Choose **Settings** > **IGP Link Management** > **Attribute Settings** from the main menu. In the upper right corner, filter links between P1 and PE3 based on the keyword.

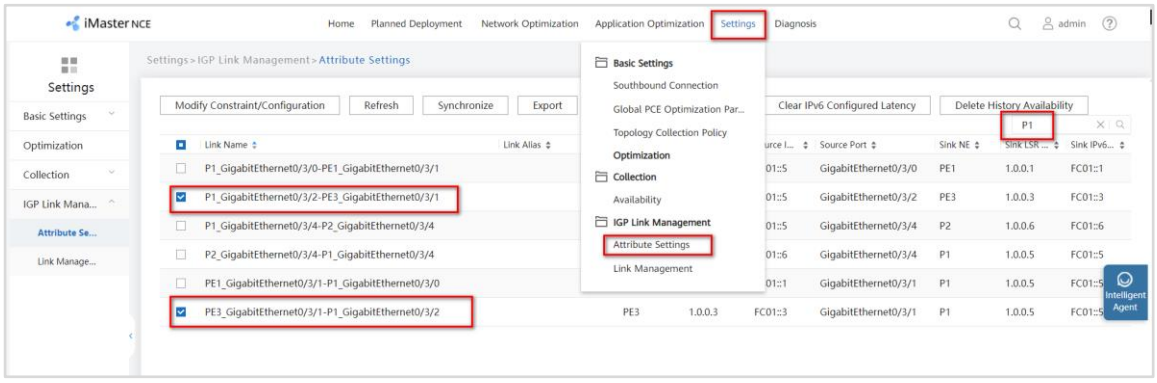

Select the link from P1 to PE3 and that from PE3 to P1, and modify the two-way link delay.

Click **Modify Constraint/Configuration** to modify the delay.

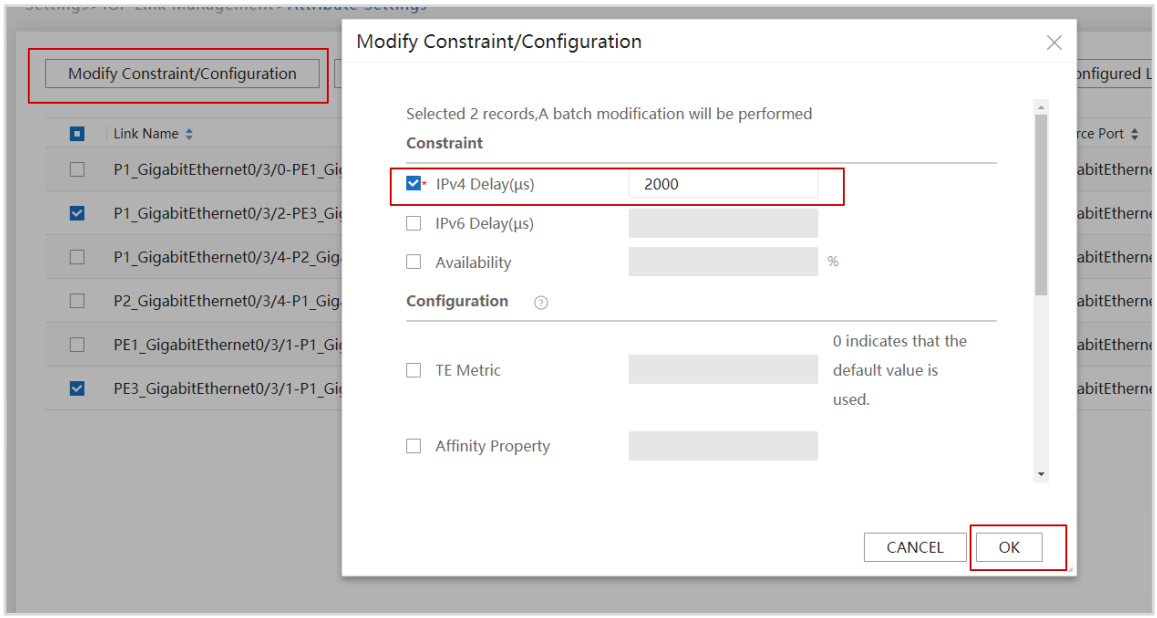

Return to the **Network Optimization** page and check the delay of the SR-MPLS TE Policy between PE1 and PE4.

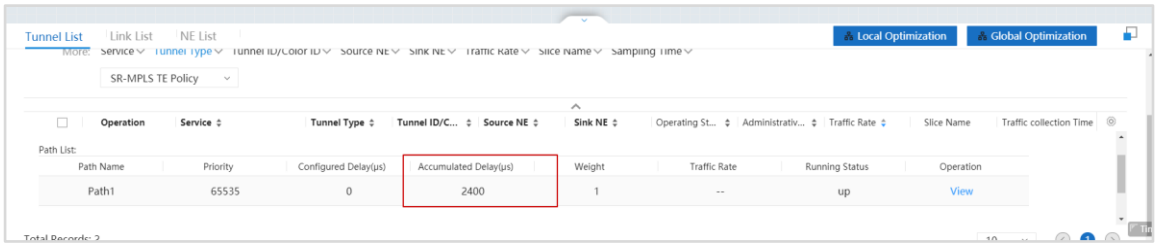

The accumulated delay has changed to 2400 μs.

Perform optimization.

Select the bidirectional tunnel between PE1 and PE4 and click **Local Optimization** in the upper right corner.

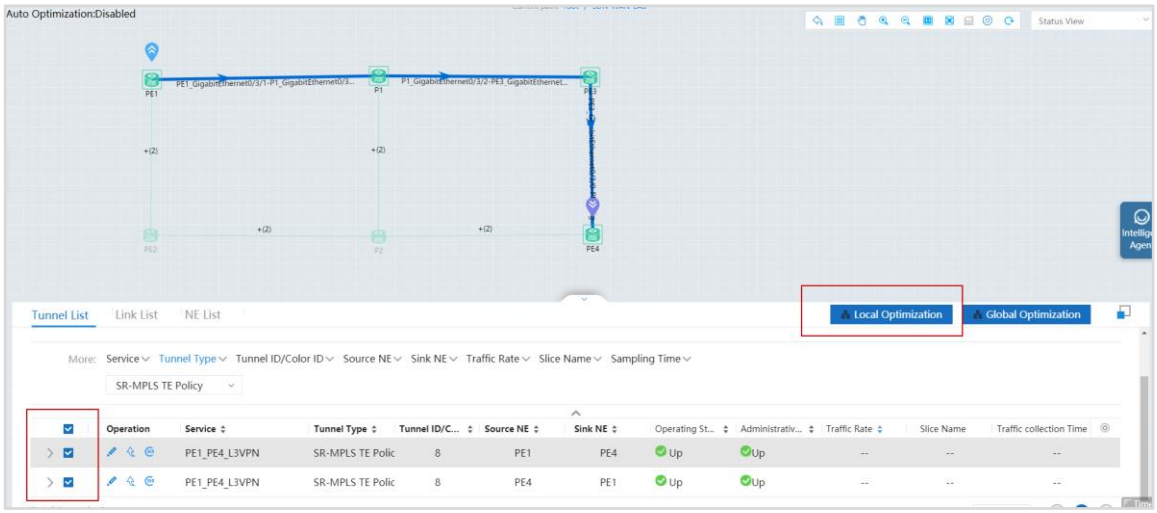

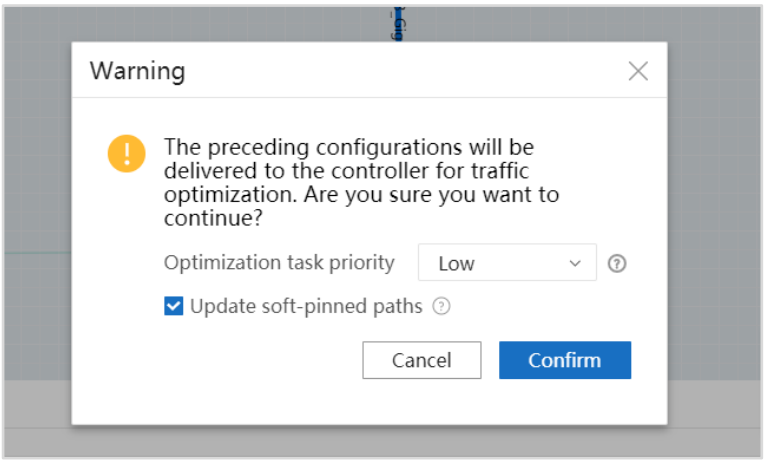

Click **OK** and wait for the result.

Check the optimization result.

After the optimization is complete, iMaster NCE-IP automatically displays a new page.

On the page that is displayed, click the topology to expand it. Then, click the service name to view the paths before and after optimization.

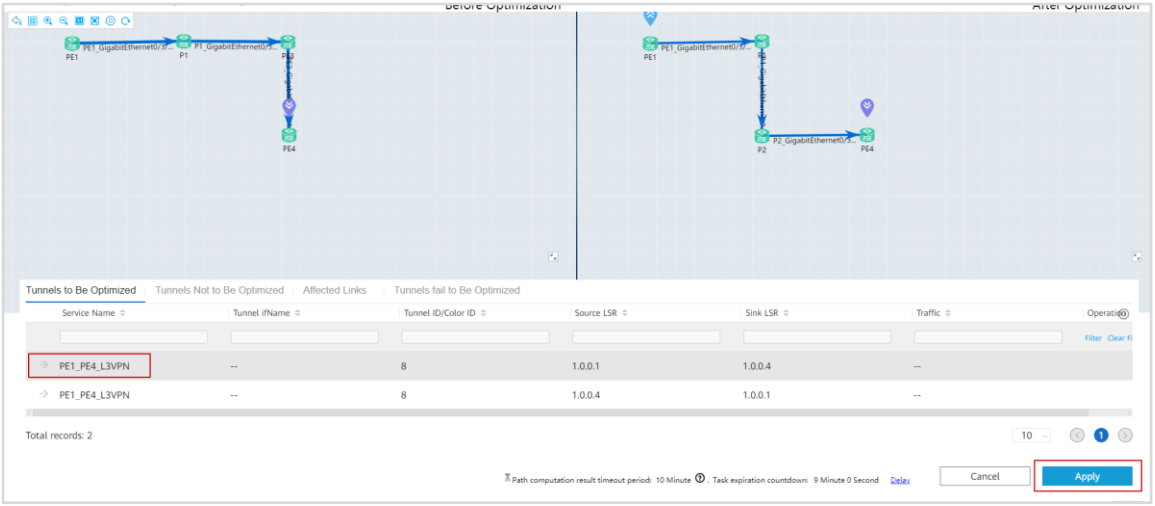

Finally, click **Apply**. The tunnel then switches to the new forwarding path with the minimum delay.

Step 3 Perform automatic optimization.

Use TWAMP to measure network-wide link delay, so that iMaster NCE-IP can perform optimization based on the real-time delay of links.

Before performing automatic optimization, delete the link delay manually modified in the previous step.

Create TWAMP test instances.

Open the Network Performance Analysis app and choose **Settings** > **Test Case Management** from the main menu to create TWAMP test instances.

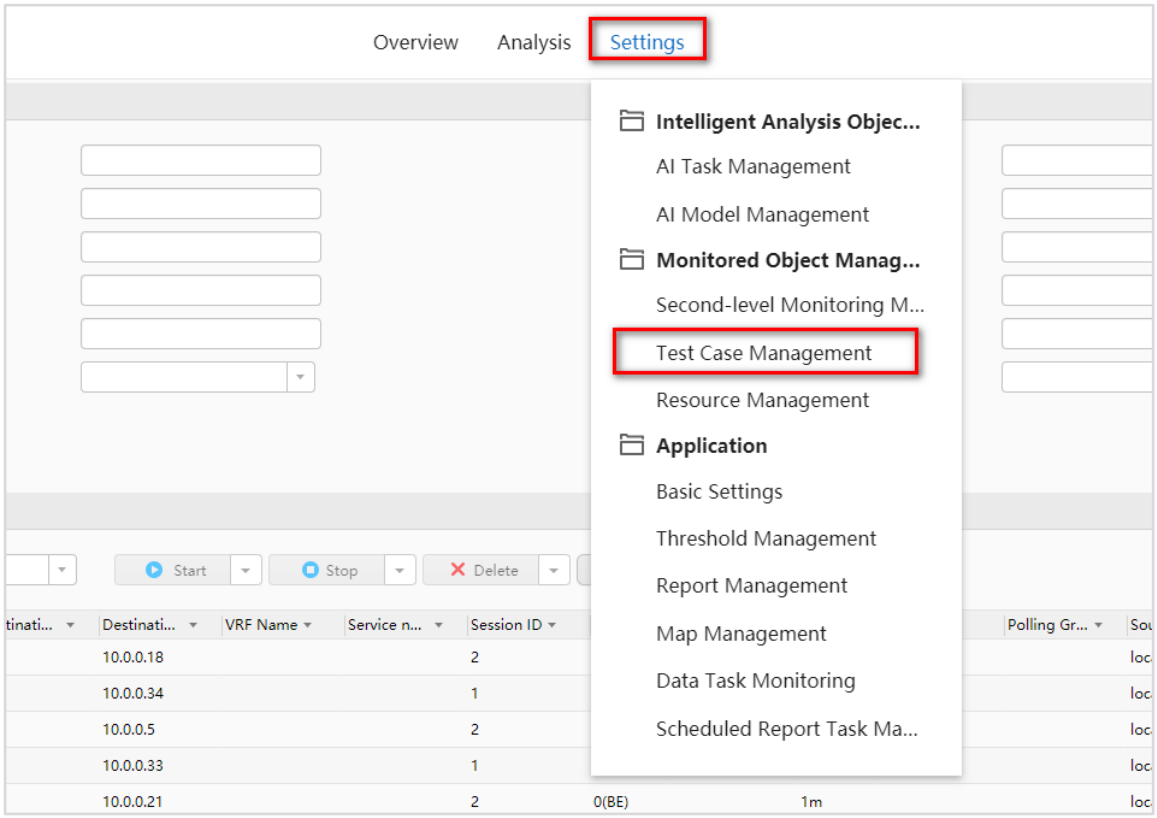

Select **The test source is IGPLink**.

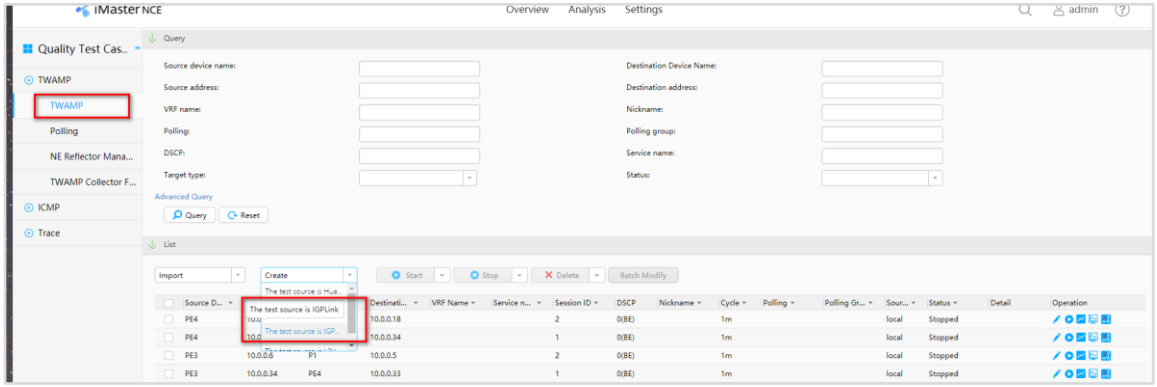

On the page that is displayed, set **Cycle** to **5min**, select all interfaces, and set **Report delay** to **Yes** (the default value is **No**).

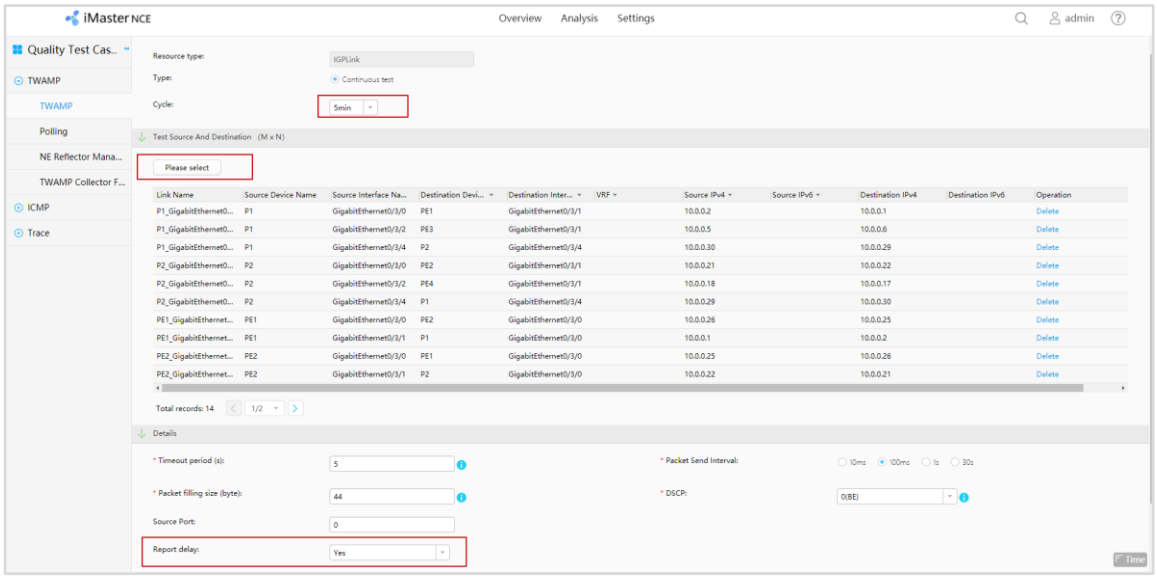

Finally, click **OK**. TWAMP then continuously monitors the bandwidth of all links and provides the bandwidth information for the network optimization module.

Check the real-time link delay of a service tunnel.

Return to the network optimization page and check the real-time delay of the SR-MPLS TE Policy between PE1 and PE4.

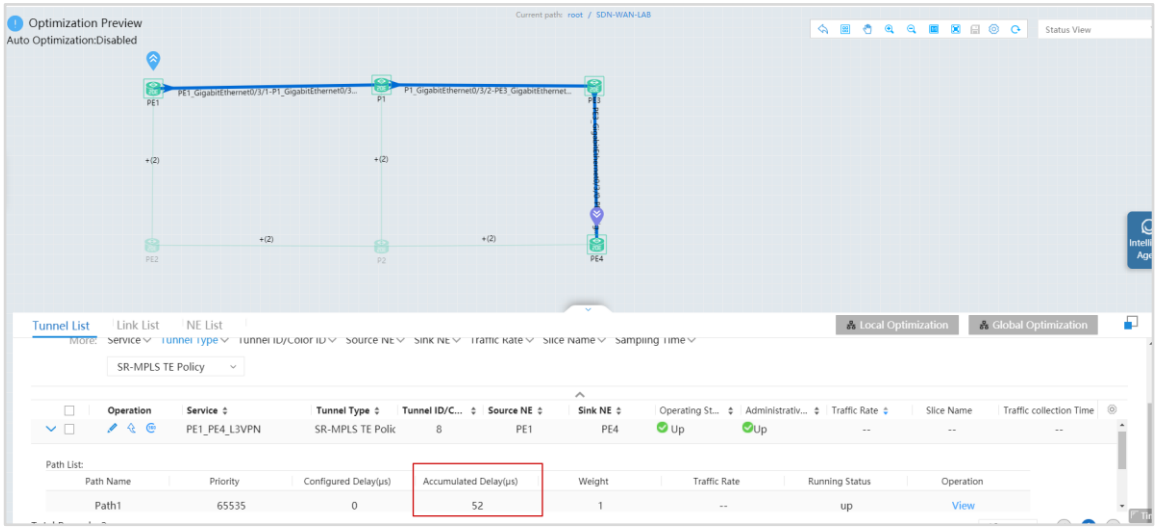

The accumulated delay is 49 μs, which is the real-time measurement result provided by TWAMP.

Enable automatic optimization.

Open the Network Path Navigation app and choose **Settings** > **Optimization** from the main menu. On the page that is displayed, enable automatic optimization.

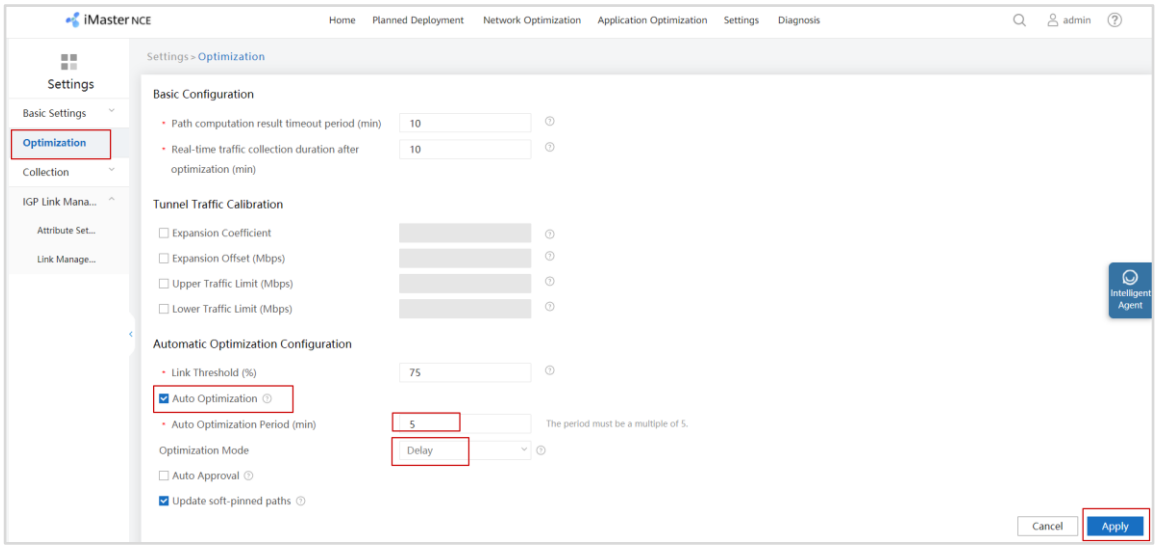

Set **Optimization Mode** to **Delay**, set **Auto Optimization Period (min)** to **5**, and deselect **Auto Approval**.

In delay-based optimization mode, the controller traverses and compares the configured delay of each tunnel with the accumulated delay collected from forwarders, and performs local optimization on tunnels whose accumulated delay exceeds the configured delay.

An SR-MPLS TE tunnel (not the SR-MPLS TE Policy queried during manual optimization) has been created previously. On the network optimization page, find the tunnel through filtering and check its two-way delay.

Check the delay of the SR-MPLS TE tunnel from PE1 to PE4.

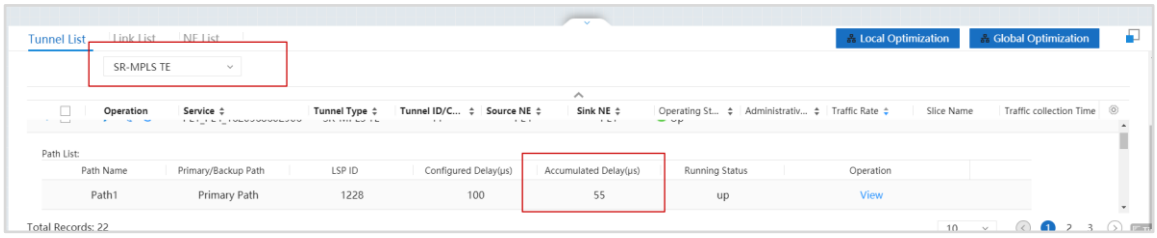

The delay of the tunnel from PE1 to PE4 is 52 μs. (The actual delay depends on the test environment.)

Open the Network Management app and choose **Service** > **Dynamic Tunnel** from the main menu. On the page that is displayed, modify the delay requirement of the SR-MPLS TE tunnel. (No delay requirement has been configured previously.)

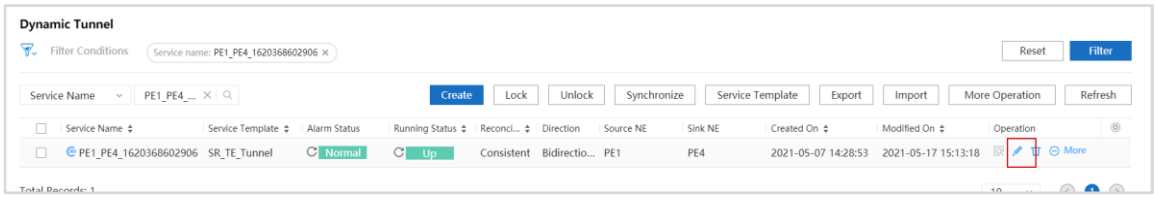

On the **Modify Dynamic Tunnel** page, set the delay constraint to 100 μs for both the forward and reverse paths to restrict tunnel availability.

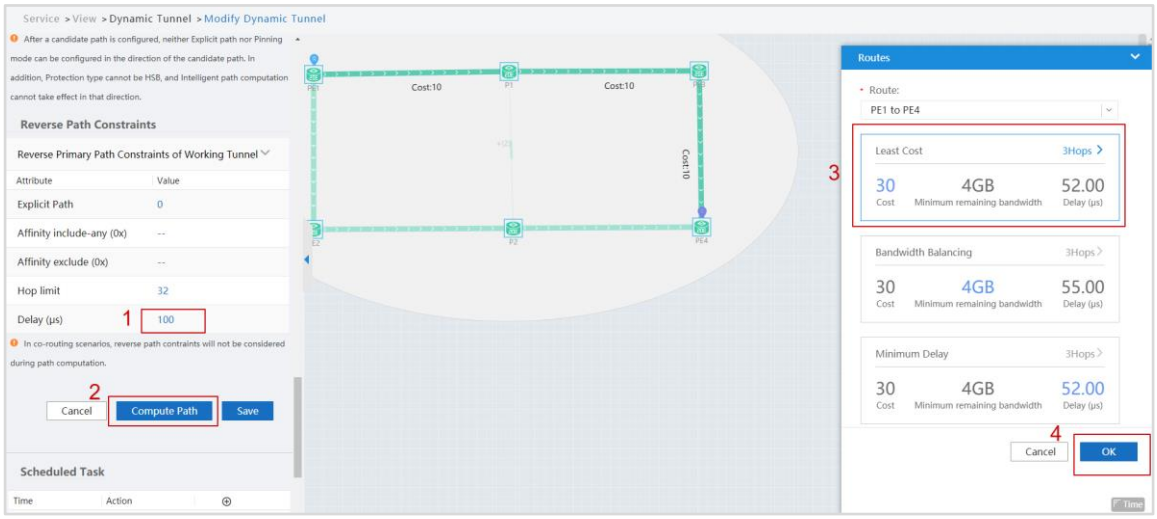

After the tunnel constraints are modified, click **Compute Path**. The controller recomputes tunnel paths. In the computation result area, select the one with the least cost or minimum delay as the optimization path and then click **OK** and **Save** in succession.

In the reconciliation information dialog box that is displayed, click **Next** and **OK** in succession to save the modified SR-MPLS TE configurations.

Then, if the delay of the tunnel between PE1 and PE4 exceeds 100 μs, the controller automatically performs automatic optimization.

Manually interfere in the TWAMP measurement of the path from PE1 to PE4.

In the outbound direction of PE3's GE 0/3/1, filter out the test packets reflected by PE3 to P1, so that P1 cannot detect the delay of the link to PE3's GE 0/3/1.

Check TWAMP session information on P1.

```
[P1]display twamp-light test-session
Total number  : 3
Active number : 3
------------------------------------------------------------------------------
  ID      Sender-IP       Sender-Port  Reflector-IP     Reflector-Port  State
------------------------------------------------------------------------------
   1      10.0.0.2          45001      10.0.0.1           33435        active
   2      10.0.0.5          45002      10.0.0.6           33435        active
   3      10.0.0.30         45003      10.0.0.29          33435        active
```

The command output shows that the packets sent by P1 to PE3 for link delay measurement are UDP packets with the Sender-Port number being 45002 and Reflector-Port number being 33435. In the 5-tuple of the packets reflected by PE3, the source IP address is 10.0.0.6, destination IP address is 10.0.0.5, source port number is 33435, and destination port number is 45002.

Create a traffic policy on PE3 to filter traffic.

```
[PE3]acl 3004
[PE3-acl4-advance-3004] rule 2 permit udp source 10.0.0.6 0 source-port eq 33435 destination
10.0.0.5 0
```

Use an advanced ACL to match the packet reflected by PE3.

Configure a traffic classifier.

```
[PE3]traffic classifier TWAMP
[PE3-classifier-TWAMP] if-match acl 3004
```

Configure a traffic behavior.

```
[PE3]traffic behavior TWAMP
[PE3-behavior-TWAMP]deny
```

Configure a traffic policy and apply the traffic policy to the outbound direction of GE0/3/1.

```
[PE3]traffic policy TWAMP
[PE3-trafficpolicy-TWAMP] classifier TWAMP behavior TWAMP
[PE3-trafficpolicy-TWAMP] quit
[PE3]interface GigabitEthernet0/3/1
[PE3-GigabitEthernet0/3/1] traffic-policy TWAMP outbound
```

On P1, check the TWAMP test result between P1 and PE3.

```
[P1]display twamp-light statistic-type twoway-delay test-session 2
Latest two-way delay statistics(usec):
-------------------------------------------------------------------------------------------
      Index                    Delay(Avg)                    Jitter(Avg)
-------------------------------------------------------------------------------------------
       2714                        -                             -
       2715                        -                             -
       2716                        -                             -
       2717                        -                             -
       2718                        -                             -
       2719                        -                             -
       2720                        -                             -
       2721                        -                             -
       2722                        -                             -
       2723                        -                             -
       2724                        -                             -
       2725                        -                             -
       2726                        -                             -
       2727                        -                             -
```

```
        2728                    -                   -
        2729                    -                   -
        2730                    -                   -
        2731                    -                   -
        2732                    -                   -
        2733                    -                   -
        2734                    -                   -
        2735                    -                   -
        2736                    -                   -
        2737                    -                   -
        2738                    -                   -
        2739                    -                   -
        2740                    -                   -
        2741                    -                   -
        2742                    -                   -
        2743                    -                   -
-------------------------------------------------------------------------------
Average Delay      :        40      Average Jitter     :         0
Maximum Delay      :        41      Maximum Jitter     :         1
Minimum Delay      :        37      Minimum Jitter     :         0
```

The delay information cannot be detected.

On the controller, check the delay of the SR-MPLS TE tunnel from PE1 to PE4.

On the **Network Optimization** page, find the SR-MPLS TE tunnel through filtering and check its two-way delay.

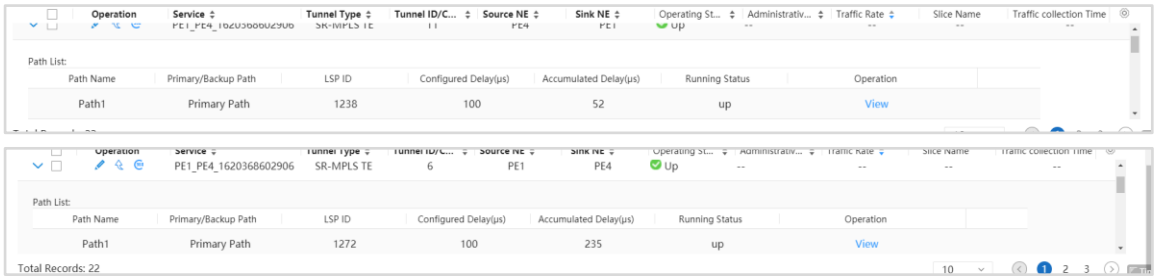

The delay in the direction from PE4 to PE1 is normal, but the delay in the direction from PE1 to PE4 is 233 μs, exceeding the delay limit (100 μs) configured for the tunnel.

In this case, the controller automatically computes a new path that meets the constraint for the tunnel from PE4 to PE1 during the next automatic optimization.

Because **Auto Approval** is not enabled, we need to manually check the optimization result.

Check the optimization result.

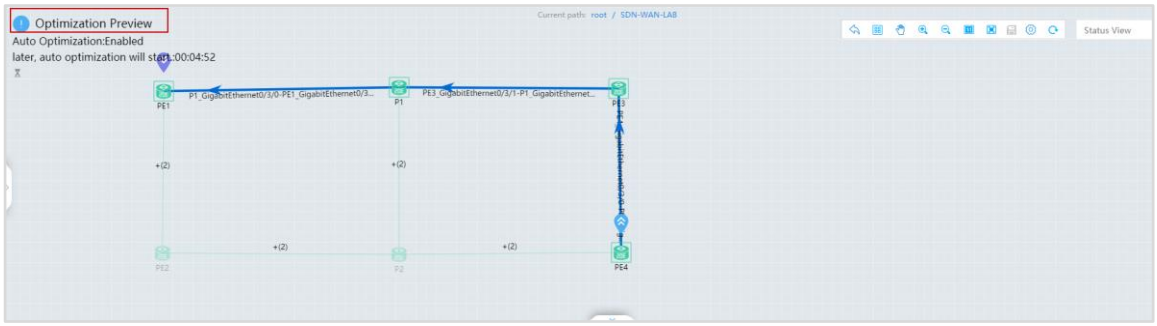

In the upper left corner of the **Network Optimization** page, click **Optimization Preview**.

On the page that is displayed, expand the topology and click the tunnel from PE1 to PE4 (the source LSR is 1.0.0.1).

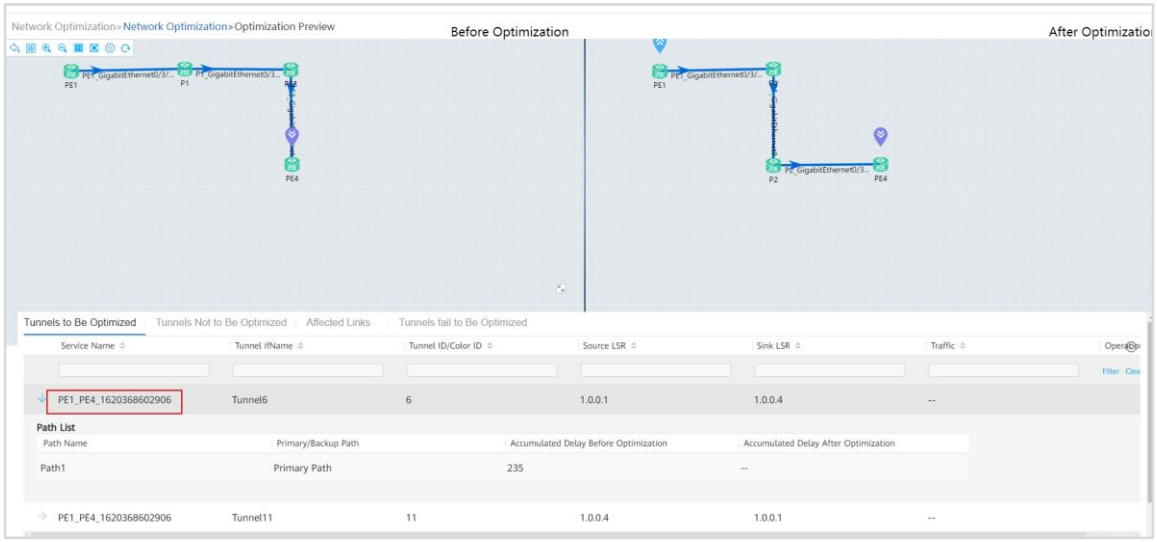

The tunnel has switched to path PE1-> PE2 -> P2 -> PE4 instead of traversing P1 and PE3. Click **Apply** in the lower right corner to deliver the automatic optimization result.

3.1.3 Quiz

Which protocol is used by the controller to deliver SR-MPLS TE configurations?

3.2 SRv6 Service Delivery by the Controller

3.2.1 Introduction

3.2.1.1 Objectives

Upon completion of this task, you will be able to:

- Establish BGP-LS and BGP SRv6 Policy relationships between the controller and devices.
- Deliver EVPN L3VPNv4 over SRv6 Policy configurations through iMaster NCE-IP.

- Observe packet forwarding over the SRv6 TE Policy.

3.2.1.2 Networking Description

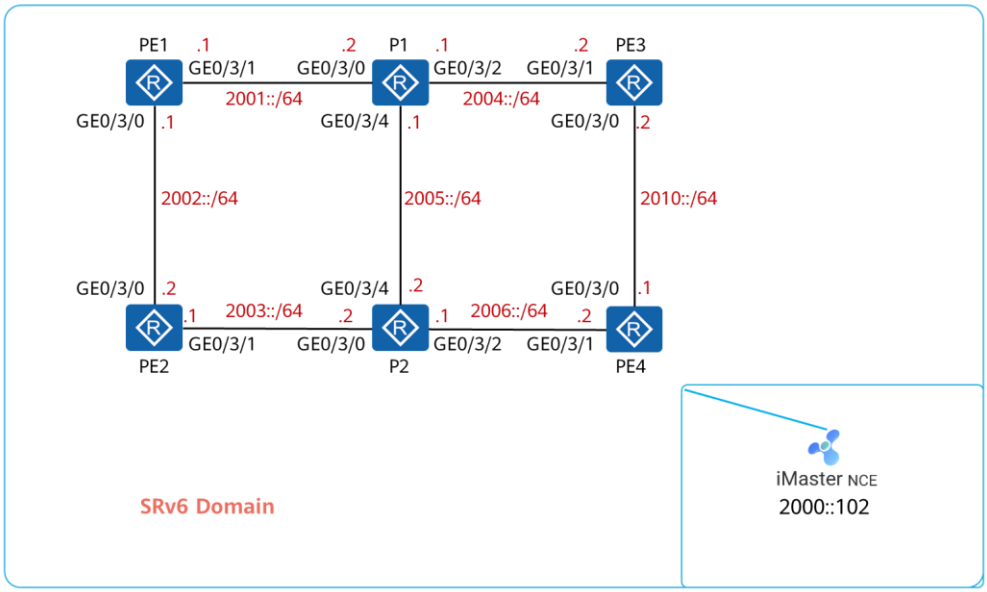

Figure 3-2 Topology for SRv6 service delivery by the controller

The figure shows the device connection and IP address planning.

Loopback0 is created on all devices, and Loopback0 IP addresses are in the format of FC01::*X*, where *X* indicates the device number. For details, see the following address planning table.

IS-IS is enabled globally in the entire SRv6 domain, and SRv6 SIDs are distributed through IS-IS.

BGP runs in the AS. P1 and P2 function as RRs. All PEs establish VPNv4 peer relationships and SR Policy peer relationships with P1 and P2.

All devices connect to iMaster NCE-IP through the management interface (GE0/0/0). The controller address is shown in the figure.

Now we need to use iMaster NCE-IP to manage all devices and then deliver EVPN L3VPNv4 over SRv6 Policy configurations.

3.2.2 Experiment Task

3.2.2.1 Configuration Roadmap

1. Perform basic device configurations, such as configuring IPv6 addresses for Loopback0, interconnection, and management interfaces and configuring SSH.
2. Perform IGP- and SRv6-related device configurations, such as enabling SRv6 globally, configuring SRv6 locators, enabling IS-IS TE, and enabling IS-IS to advertise SRv6 labels.

3. Perform BGP-related device configurations, such as configuring VPNv4 and SR Policy peer relationships between devices and establishing BGP-LS and SR Policy peer relationships between RRs and iMaster NCE-IP.
4. Perform controller configurations, such as configuring routes from iMaster NCE-IP to devices (for device management), adding devices to iMaster NCE-IP for management, and configuring BGP-LS and SR Policy peer relationships between iMaster NCE-IP and the RR.
5. Configure an SRv6 Policy on iMaster NCE-IP and create a color for the SRv6 Policy, so that L3VPN traffic transmitted along EVPN routes carrying the color extended community attribute can recurse to SRv6 Policies.
6. Configure an EVPN L3VPN service on iMaster NCE-IP and recurse the service to the SRv6 Policy for forwarding.

3.2.2.2 Device-Side Basic Configurations

To complete the subsequent configuration, you need to create Loopback0 interfaces and configure IPv6 addresses for Loopback0, management, and interconnection interfaces first.

To enable iMaster NCE-IP to manage devices, enable LLDP and SSH and configure SFTP, NETCONF, and SNMP on all devices.

The IP address is 2000::*X* for the device management interface and FC01::*X* for Loopback0. For values represented by *X*, see the following table.

Table 3-4 Address planning

Device Name	Device Number
PE1	1
PE2	2
PE3	3
PE4	4
P1	5
P2	6

Step 1 Configure IPv6 addresses for management interfaces.

Configure IPv6 addresses for the management interfaces (GE0/0/0) of all devices. These IP addresses are used for communication between the devices and iMaster NCE-IP.

Configure the configuration validation mode as immediate validation.

```
<PE2>system-view immediately
```

Here, PE2 is used as an example. Repeat the configuration for other devices.

Name the devices.

Omitted

PE1

```
[PE1]interface GigabitEthernet0/0/0
[PE1-GigabitEthernet0/0/0] undo ip binding vpn-instance __LOCAL_OAM_VPN__
[PE1-GigabitEthernet0/0/0] ipv6 enable
[PE1-GigabitEthernet0/0/0] ipv6 address 2000::1/64
[PE1-GigabitEthernet0/0/0] quit
```

PE2

```
[PE2]interface GigabitEthernet0/0/0
[PE2-GigabitEthernet0/0/0] undo ip binding vpn-instance __LOCAL_OAM_VPN__
[PE2-GigabitEthernet0/0/0] ipv6 enable
[PE2-GigabitEthernet0/0/0] ipv6 address 2000::2/64
[PE2-GigabitEthernet0/0/0] quit
```

PE3

```
[PE3]interface GigabitEthernet0/0/0
[PE3-GigabitEthernet0/0/0] undo ip binding vpn-instance __LOCAL_OAM_VPN__
[PE3-GigabitEthernet0/0/0] ipv6 enable
[PE3-GigabitEthernet0/0/0] ipv6 address 2000::3/64
[PE3-GigabitEthernet0/0/0] quit
```

PE4

```
[PE4]interface GigabitEthernet0/0/0
[PE4-GigabitEthernet0/0/0] undo ip binding vpn-instance __LOCAL_OAM_VPN__
[PE4-GigabitEthernet0/0/0] ipv6 enable
[PE4-GigabitEthernet0/0/0] ipv6 address 2000::4/64
[PE4-GigabitEthernet0/0/0] quit
```

P1

```
[P1]interface GigabitEthernet0/0/0
[P1-GigabitEthernet0/0/0] undo ip binding vpn-instance __LOCAL_OAM_VPN__
[P1-GigabitEthernet0/0/0] ipv6 enable
[P1-GigabitEthernet0/0/0] ipv6 address 2000::5/64
[P1-GigabitEthernet0/0/0] quit
```

P2

```
[P2]interface GigabitEthernet0/0/0
[P2-GigabitEthernet0/0/0] undo ip binding vpn-instance __LOCAL_OAM_VPN__
[P2-GigabitEthernet0/0/0] ipv6 enable
[P2-GigabitEthernet0/0/0] ipv6 address 2000::6/64
[P2-GigabitEthernet0/0/0] quit
```

Step 2 Configure IPv6 addresses for Loopback0 interfaces.

Create Loopback0 on all devices and configure IPv6 addresses for these interfaces.

PE1

```
[PE1]interface LoopBack0
[PE1-LoopBack0] ipv6 enable
[PE1-LoopBack0] ipv6 address FC01::1/128
[PE1-LoopBack0] quit
```

PE2

```
[PE2]interface LoopBack0
[PE2-LoopBack0] ipv6 enable
[PE2-LoopBack0] ipv6 address FC01::2/128
[PE2-LoopBack0] quit
```

PE3

```
[PE3]interface LoopBack0
[PE3-LoopBack0] ipv6 enable
[PE3-LoopBack0] ipv6 address FC01::3/128
[PE3-LoopBack0] quit
```

PE4

```
[PE4]interface LoopBack0
[PE4-LoopBack0] ipv6 enable
[PE4-LoopBack0] ipv6 address FC01::4/128
[PE4-LoopBack0] quit
```

P1

```
[P1]interface LoopBack0
[P1-LoopBack0] ipv6 enable
[P1-LoopBack0] ipv6 address FC01::5/128
[P1-LoopBack0] quit
```

P2

```
[P2]interface LoopBack0
[P2-LoopBack0] ipv6 enable
[P2-LoopBack0] ipv6 address FC01::6/128
[P2-LoopBack0] quit
```

Step 3 Configure IPv6 addresses for interconnection interfaces.

Configure IPv6 addresses for device interconnection interfaces as shown in the topology.

By default, DCN is enabled on NE router interfaces. To facilitate the experiment, disable DCN globally on all devices.

Disable DCN globally on each device.

```
[PE1] undo dcn
Warning: This operation will disable DCN function. Continue? [Y/N]:y
```

Here, PE1 is used as an example. Repeat this operation for other devices.

Configure IPv6 addresses for interconnection interfaces.

P1

```
[P1]interface GigabitEthernet0/3/0
[P1-GigabitEthernet0/3/0] ipv6 enable
[P1-GigabitEthernet0/3/0] ipv6 address 2001::2/64
[P1-GigabitEthernet0/3/0] quit
[P1]interface GigabitEthernet0/3/2
[P1-GigabitEthernet0/3/2] ipv6 enable
[P1-GigabitEthernet0/3/2] ipv6 address 2004::1/64
[P1-GigabitEthernet0/3/2] quit
[P1]interface GigabitEthernet0/3/4
[P1-GigabitEthernet0/3/4] ipv6 enable
[P1-GigabitEthernet0/3/4] ipv6 address 2005::1/64
[P1-GigabitEthernet0/3/4] quit
```

P2

```
[P2]interface GigabitEthernet0/3/0
[P2-GigabitEthernet0/3/0] ipv6 enable
[P2-GigabitEthernet0/3/0] ipv6 address 2003::2/64
[P2-GigabitEthernet0/3/0] quit
[P2]interface GigabitEthernet0/3/2
[P2-GigabitEthernet0/3/2] ipv6 enable
[P2-GigabitEthernet0/3/2] ipv6 address 2006::1/64
[P2-GigabitEthernet0/3/2] quit
[P2]interface GigabitEthernet0/3/4
[P2-GigabitEthernet0/3/4] ipv6 enable
[P2-GigabitEthernet0/3/4] ipv6 address 2005::2/64
[P2-GigabitEthernet0/3/4] quit
```

PE1

```
[PE1]interface GigabitEthernet0/3/0
[PE1-GigabitEthernet0/3/0] ipv6 enable
[PE1-GigabitEthernet0/3/0] ipv6 address 2002::1/64
[PE1-GigabitEthernet0/3/0] quit
[PE1]interface GigabitEthernet0/3/1
[PE1-GigabitEthernet0/3/1] ipv6 enable
[PE1-GigabitEthernet0/3/1] ipv6 address 2001::1/64
[PE1-GigabitEthernet0/3/1] quit
```

PE2

```
[PE2]interface GigabitEthernet0/3/0
[PE2-GigabitEthernet0/3/0] ipv6 enable
[PE2-GigabitEthernet0/3/0] ipv6 address 2002::2/64
[PE2-GigabitEthernet0/3/0] quit
[PE2]interface GigabitEthernet0/3/1
[PE2-GigabitEthernet0/3/1] ipv6 enable
[PE2-GigabitEthernet0/3/1] ipv6 address 2003::1/64
[PE2-GigabitEthernet0/3/1] quit
```

PE3

```
[PE3]interface GigabitEthernet0/3/0
[PE3-GigabitEthernet0/3/0] ipv6 enable
[PE3-GigabitEthernet0/3/0] ipv6 address 2010::2/64
[PE3-GigabitEthernet0/3/0] quit
[PE3]interface GigabitEthernet0/3/1
[PE3-GigabitEthernet0/3/1] ipv6 enable
[PE3-GigabitEthernet0/3/1] ipv6 address 2004::2/64
[PE3-GigabitEthernet0/3/1] quit
```

PE4

```
[PE4]interface GigabitEthernet0/3/0
[PE4-GigabitEthernet0/3/0] ipv6 enable
[PE4-GigabitEthernet0/3/0] ipv6 address 2010::1/64
[PE4-GigabitEthernet0/3/0] quit
[PE4]interface GigabitEthernet0/3/1
[PE4-GigabitEthernet0/3/1] ipv6 enable
[PE4-GigabitEthernet0/3/1] ipv6 address 2006::2/64
[PE4-GigabitEthernet0/3/1] quit
```

Step 4 Configure SSH, SNMP, and LLDP.

The configurations are similar to that in 3.1 SR-MPLS Service Delivery by the Controller.

3.2.2.3 Device-Side IGP and SRv6 Configurations

Enable SRv6 on devices in the entire SRv6 domain, set the IGP to IS-IS, enable IS-IS to carry link attributes in LSPs, enable IS-IS topology information reporting through BGP-LS, enable IS-IS TE, and enable IS-IS to advertise SRv6 labels.

Step 1 Configure the IGP.

The IS-IS area ID is 49.0001, the IS-IS process ID is 1, all devices are Level-2 devices, and the NET is converted from the device number (for example, PE2's NET is 49.0001.0010.0000.0002.00). Enable IS-IS on Loopback0 and interconnection interfaces.

In this case, you need to set **cost-style** to **wide** to support IS-IS extensions.

Enable BFD globally.

```
[PE1]bfd
```

PE1 is used as an example. Repeat the configuration for other devices.

Description of IS-IS commands:

cost-style wide: The narrow cost type does not support the TE information (such as bandwidth) required in TE scenarios. Therefore, the wide cost type needs to be configured.

ipv6 enable topology ipv6: This command enables the IPv6 capability of an IS-IS process.

ipv6 advertise link attributes: This command enables LSPs to carry link attribute TLVs, including interface IPv6 addresses and interface indexes.

ipv6 bgp-ls enable level-2: This command enables topology information collected by IS-IS to be sent to the controller through BGP-LS. This function only needs to be configured on the RR. That is, only one device in the IGP domain needs to send topology information to the controller through BGP-LS.

ipv6 traffic-eng level-2: This command enables IS-IS TE, so that link bandwidth information can be sent to the TE module.

set-overload on-startup: This command sets the overload bit, which is used to notify others that the local node cannot forward traffic at this time. The local node is then not used as a forwarding node during LSP-based path calculation. The command parameters include **on-startup** and **wait-for-bgp**.

ipv6 metric-delay advertisement enable: This command enables IPv6 delay advertisement. After this function is enabled, IS-IS collects and floods information about the intra-area IPv6 link delay, and BGP-LS reports the information to the controller. The controller can then use the delay information to compute optimal paths on a P2P network.

Configure IS-IS on PEs.

```
[PE1]isis 1
[PE1-isis-1] is-level level-2
[PE1-isis-1] network-entity 49.0001.0010.0000.0001.00
[PE1-isis-1] is-name PE1
[PE1-isis-1] set-overload on-startup
[PE1-isis-1] ipv6 enable topology ipv6
[PE1-isis-1] ipv6 advertise link attributes
[PE1-isis-1] ipv6 bfd all-interfaces enable
[PE1-isis-1] ipv6 metric-delay advertisement enable level-1-2
[PE1-isis-1] ipv6 traffic-eng level-2
```

PE1 is used as an example. The configurations of other PEs are similar to the configuration of PE1.

Configure IS-IS on Ps.

```
[P1]isis 1
[P1-isis-1] is-level level-2
[P1-isis-1] cost-style wide
[P1-isis-1] bfd all-interfaces enable
[P1-isis-1] network-entity 49.0001.0010.0000.0005.00
[P1-isis-1] is-name P1
[P1-isis-1] set-overload on-startup
[P1-isis-1] ipv6 enable topology ipv6
[P1-isis-1] ipv6 bgp-ls enable level-2
[P1-isis-1] ipv6 advertise link attributes
[P1-isis-1] ipv6 bfd all-interfaces enable
[P1-isis-1] ipv6 metric-delay advertisement enable level-1-2
[P1-isis-1] ipv6 traffic-eng level-2
```

P1 is used as an example. The configurations of other Ps are similar to the configuration of P1.

Enable FRR on all devices.

```
[PE1]isis 1
[PE1-isis-1] frr
[PE1-isis-1-frr] loop-free-alternate level-2
[PE1-isis-1-frr] ti-lfa level-2
[PE1-isis-1-frr] quit
```

PE1 is used as an example. Repeat the configuration for other devices.

Enable IS-IS on the interconnection and Loopback0 interfaces of all devices and set the link type to P2P.

#PE1

```
[PE1]interface GigabitEthernet0/3/0
[PE1-GigabitEthernet0/3/0] isis ipv6 enable 1
[PE1-GigabitEthernet0/3/0] isis circuit-type p2p
[PE1-GigabitEthernet0/3/0] quit
[PE1]interface GigabitEthernet0/3/1
[PE1-GigabitEthernet0/3/1] isis ipv6 enable 1
[PE1-GigabitEthernet0/3/1] isis circuit-type p2p
[PE1-GigabitEthernet0/3/1] quit
[PE1]interface LoopBack0
[PE1-LoopBack0] isis ipv6 enable 1
```

#PE2

```
[PE2]interface GigabitEthernet0/3/0
[PE2-GigabitEthernet0/3/0] isis ipv6 enable 1
[PE2-GigabitEthernet0/3/0] isis circuit-type p2p
[PE2-GigabitEthernet0/3/0] quit
[PE2]interface GigabitEthernet0/3/1
[PE2-GigabitEthernet0/3/1] isis ipv6 enable 1
[PE2-GigabitEthernet0/3/1] isis circuit-type p2p
[PE2-GigabitEthernet0/3/1] quit
[PE2]interface LoopBack0
[PE2-LoopBack0] isis ipv6 enable 1
```

#PE3

```
[PE3]interface GigabitEthernet0/3/0
[PE3-GigabitEthernet0/3/0] isis ipv6 enable 1
[PE3-GigabitEthernet0/3/0] isis circuit-type p2p
[PE3-GigabitEthernet0/3/0] quit
[PE3]interface GigabitEthernet0/3/1
[PE3-GigabitEthernet0/3/1] isis ipv6 enable 1
[PE3-GigabitEthernet0/3/1] isis circuit-type p2p
[PE3-GigabitEthernet0/3/1] quit
[PE3]interface LoopBack0
[PE3-LoopBack0] isis ipv6 enable 1
```

#PE4

```
[PE4]interface GigabitEthernet0/3/0
[PE4-GigabitEthernet0/3/0] isis ipv6 enable 1
[PE4-GigabitEthernet0/3/0] isis circuit-type p2p
[PE4-GigabitEthernet0/3/0] quit
[PE4]interface GigabitEthernet0/3/1
[PE4-GigabitEthernet0/3/1] isis ipv6 enable 1
[PE4-GigabitEthernet0/3/1] isis circuit-type p2p
[PE4-GigabitEthernet0/3/1] quit
[PE4]interface LoopBack0
[PE4-LoopBack0] isis ipv6 enable 1
```

#P1

```
[P1]interface GigabitEthernet0/3/0
[P1-GigabitEthernet0/3/0] isis ipv6 enable 1
[P1-GigabitEthernet0/3/0] isis circuit-type p2p
[P1-GigabitEthernet0/3/0] quit
[P1]interface GigabitEthernet0/3/2
[P1-GigabitEthernet0/3/2] isis ipv6 enable 1
[P1-GigabitEthernet0/3/2] isis circuit-type p2p
[P1-GigabitEthernet0/3/2] quit
[P1]interface GigabitEthernet0/3/4
[P1-GigabitEthernet0/3/4] isis ipv6 enable 1
[P1-GigabitEthernet0/3/4] isis circuit-type p2p
[P1-GigabitEthernet0/3/4] quit
[P1]interface LoopBack0
[P1-LoopBack0] isis ipv6 enable 1
```

#P2

```
[P2]interface GigabitEthernet0/3/0
[P2-GigabitEthernet0/3/0] isis ipv6 enable 1
[P2-GigabitEthernet0/3/0] isis circuit-type p2p
[P2-GigabitEthernet0/3/0] quit
[P2]interface GigabitEthernet0/3/2
[P2-GigabitEthernet0/3/2] isis ipv6 enable 1
[P2-GigabitEthernet0/3/2] isis circuit-type p2p
[P2-GigabitEthernet0/3/2] quit
[P2]interface GigabitEthernet0/3/4
[P2-GigabitEthernet0/3/4] isis ipv6 enable 1
[P2-GigabitEthernet0/3/4] isis circuit-type p2p
[P2-GigabitEthernet0/3/4] quit
[P2]interface LoopBack0
[P2-LoopBack0] isis ipv6 enable 1
```

Check IS-IS configurations.

Check IS-IS neighbor relationships on P1.

```
[P1]display isis peer
                Peer information for ISIS(1)

System Id Interface         Circuit Id         State  HoldTime  Type     PRI
```

```
--------------------------------------------------------------------------------
PE1*       GE0/3/0       0000000007      Up    22s       L2        --
P3*        GE0/3/2       0000000007      Up    27s       L2        --
P2*        GE0/3/4       0000000010      Up    27s       L2        --

Total Peer(s): 3
```

Check the IS-IS IPv6 routing table on each router. The following example uses the command output on P1.

```
[P1]display isis route ipv6

                         Route information for ISIS(1)
                         -----------------------------

                        ISIS(1) Level-2 Forwarding Table
                        --------------------------------

IPV6 Dest.      ExitInterface        NextHop                       Cost      Flags
--------------------------------------------------------------------------------
2001::/64       GE0/3/0              Direct                        10        D/-/L/-
2002::/64       GE0/3/0              FE80::DE99:14FF:FE7A:C212     20        A/-/-/-
2003::/64       GE0/3/4              FE80::DE99:14FF:FE7A:C3F5     20        A/-/-/-
2004::/64       GE0/3/2              Direct                        10        D/-/L/-
2005::/64       GE0/3/4              Direct                        10        D/-/L/-
2006::/64       GE0/3/4              FE80::DE99:14FF:FE7A:C3F5     20        A/-/-/-
2010::/64       GE0/3/2              FE80::A6BE:2BFF:FEAA:E617     20        A/-/-/-
FC01::1/128     GE0/3/0              FE80::DE99:14FF:FE7A:C212     10        A/-/-/-
FC01::2/128     GE0/3/0              FE80::DE99:14FF:FE7A:C212     20        A/-/-/-
                GE0/3/4              FE80::DE99:14FF:FE7A:C3F5
FC01::3/128     GE0/3/2              FE80::A6BE:2BFF:FEAA:E617     10        A/-/-/-
FC01::4/128     GE0/3/2              FE80::A6BE:2BFF:FEAA:E617     20        A/-/-/-
                GE0/3/4              FE80::DE99:14FF:FE7A:C3F5
FC01::5/128     Loop0                Direct                        0         D/-/L/-
FC01::6/128     GE0/3/4              FE80::DE99:14FF:FE7A:C3F5     10        A/-/-/-
     Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
            U-Up/Down Bit Set, LP-Local Prefix-Sid
     Protect Type: L-Link Protect, N-Node Protect
```

IPv6 routes have been learned through IS-IS.

Step 2 Configure SRv6.

Enable SRv6 globally, configure SRv6 locators, and advertise these locators through IS-IS.

Enable SRv6 globally.

```
[PE1]segment-routing ipv6
[PE1-segment-routing-ipv6] quit
```

PE1 is used as an example.

Globally configure the IPv6 router ID and TE attributes.

Configure the Loopback0 IPv6 address as the global TE IPv6 router ID and globally enable TE on each device.

```
[PE1]te ipv6-router-id FC01::1
[PE1]te attribute enable
```

PE1 is used as an example. Repeat the configuration for other devices.

Command description:

te ipv6-router-id: This command is used to set a global TE IPv6 router ID, which must be configured on the ingress of a tunnel and must be globally unique.

te attribute enable: This command is used to configure interface bandwidth (such as the maximum reservable bandwidth). MPLS TE-related commands can only be used for MPLS TE and SR-MPLS TE. In comparison, this command can be used for SRv6 TE Policy in addition to MPLS TE and SR-MPLS TE.

Configure SRv6 locators.

```
[PE1]segment-routing ipv6
[PE1-segment-routing-ipv6] sr-te frr enable
[PE1-segment-routing-ipv6] encapsulation source-address FC01::1
[PE1-segment-routing-ipv6] locator SRv6 ipv6-prefix FC00:1:: 96 static 16
[PE1-segment-routing-ipv6-locator] opcode ::1 end
[PE1-segment-routing-ipv6-locator] opcode ::F end-op
[PE1-segment-routing-ipv6-locator] quit
[PE1-segment-routing-ipv6-locator] segment-routing ipv6 locator SRv6
```

The following uses PE1 as an example to describe how to configure the source address for encapsulation and SRv6 locator and manually configure End and End.OP SIDs.

Command description:

encapsulation source-address: When traffic enters an SRv6 VPN tunnel, the address configured using this command is used as the source address in the IPv6 packet header. In this experiment, the device's Loopback0 address is used.

locator: This command is used to configure an SRv6 locator. SRv6 SIDs are in the *Locator:Function:Args* format. End SIDs are similar to node SIDs in SR-MPLS and are used to identify destination nodes on a network. End.OP SIDs are used to implement ping and tracert functions in SRv6 scenarios.

```
[PE1]segment-routing ipv6
[PE1-segment-routing-ipv6] sr-te frr enable
[PE1-segment-routing-ipv6] encapsulation source-address FC01::1
[PE1-segment-routing-ipv6] locator SRv6 ipv6-prefix FC00:1:: 96 static 16
[PE1-segment-routing-ipv6-locator] opcode ::1 end
[PE1-segment-routing-ipv6-locator] opcode ::F end-op
[PE1-segment-routing-ipv6-locator] quit
[PE1-segment-routing-ipv6-locator] segment-routing ipv6 locator SRv6
```

PE1 is used as an example to describe how to configure the source address for encapsulation and SRv6 locator and how to manually configure End and End.OP SIDs.

Device locator planning: Each NE uses FC00::*X*:: (*X* indicates the device number, which has been planned during basic configuration). The prefix length is 96, and the static segment length in the Function field is 16. Therefore, the dynamic segment length dynamically allocated by the IGP is 16 (128 – 96 – 16).

Enable the function to report SR-MPLS TE Policy information through BGP-LS.

```
[PE1-segment-routing-ipv6] srv6-te-policy bgp-ls enable
```

The configurations of other PEs are similar to the configuration of PE1.

Configure SRv6 TE bandwidth attributes for interconnection interfaces.

```
[PE1-segment-routing-ipv6] srv6-te-policy bgp-ls enable
```

The configurations of other PEs are similar to the configuration of PE1.

Configure SRv6 TE bandwidth attributes for interconnection interfaces.

```
[PE1]interface GigabitEthernet0/3/0
[PE1-GigabitEthernet0/3/0] te bandwidth max-reservable-bandwidth dynamic 40
[PE1-GigabitEthernet0/3/0] te bandwidth dynamic bc0 100
```

The following uses one interface as an example. Repeat the configuration for other interconnection interfaces.

Verify the configuration.

P1 is used as an example. Check the IS-IS IPv6 routes generated by other NEs based on the SR locator.

```
[P1]display isis route ipv6

                Route information for ISIS(1)
                -----------------------------

               ISIS(1) Level-2 Forwarding Table
               --------------------------------

 IPV6 Dest.     ExitInterface     NextHop                         Cost    Flags
-------------------------------------------------------------------------------
......
FC00:1::/96    GE0/3/0           FE80::DE99:14FF:FE7A:C212       10      A/-/-/-
FC00:2::/96    GE0/3/0           FE80::DE99:14FF:FE7A:C212       20      A/-/-/-
               GE0/3/4           FE80::DE99:14FF:FE7A:C3F5
FC00:3::/96    GE0/3/2           FE80::A6BE:2BFF:FEAA:E617       10      A/-/-/-
FC00:4::/96    GE0/3/2           FE80::A6BE:2BFF:FEAA:E617       20      A/-/-/-
               GE0/3/4           FE80::DE99:14FF:FE7A:C3F5
FC00:5::/96    NULL0               -                             0       A/-/L/-
FC00:6::/96    GE0/3/4           FE80::DE99:14FF:FE7A:C3F5       10      A/-/-/-
.....
```

P1 has learned the IS-IS IPv6 routes generated by other NEs.

Test SRv6 TE BE connectivity.

```
[P1]ping ipv6-sid FC00:1::F
```

```
PING ipv6-sid FC00:1::F : 56   data bytes, press CTRL_C to break
  Reply from FC00:1::F
  bytes=56 Sequence=1 hop limit=64 time=1 ms
  Reply from FC00:1::F
  bytes=56 Sequence=2 hop limit=64 time=1 ms
  Reply from FC00:1::F
  bytes=56 Sequence=3 hop limit=64 time=1 ms
  Reply from FC00:1::F
  bytes=56 Sequence=4 hop limit=64 time=1 ms
  Reply from FC00:1::F
  bytes=56 Sequence=5 hop limit=64 time=1 ms

--- ipv6-sid ping statistics---
  5 packet(s) transmitted
  5 packet(s) received
  0.00% packet loss
  round-trip min/avg/max=1/1/1 ms
```

On P1, ping the End.OP SID of a random NE. The ping operation succeeds.

3.2.2.4 Device-Side BGP Configurations

Configure P1 and P2 as RRs and establish BGP EVPN peer relationships between PEs and RRs, so that CEs can communicate with each other through EVPN L3VPNv4.

Configure BGP-LS peer relationships between RRs and iMaster NCE-IP, so that link information can be reported to iMaster NCE-IP.

To enable iMaster NCE-IP to monitor the path status of SRv6 Policies, establish BGP-LS peer relationships between PEs and RRs and use RRs to report the path status of SRv6 Policies to iMaster NCE-IP.

To enable iMaster NCE-IP to deliver BGP SRv6 Policy routes to PEs, establish BGP SRv6 Policy peer relationships between iMaster NCE-IP and RRs and between PEs and RRs. The RRs then reflect SRv6 Policy routes received from the controller to PEs.

Step 1 Establish EVPN peer relationships.

Establish BGP EVPN peer relationships between PEs and RRs. Set the router ID to 1.0.0.*X* (*X* indicates the device number), and use the Loopback0 address as the source address for initiating a connection.

Configure PEs.

```
[PE1]bgp 65001
[PE1-bgp] router-id 1.0.0.1
[PE1-bgp] undo default ipv4-unicast
[PE1-bgp] peer FC01::5 as-number 65001
[PE1-bgp] peer FC01::5 connect-interface LoopBack0
[PE1-bgp] peer FC01::6 as-number 65001
[PE1-bgp] peer FC01::6 connect-interface LoopBack0
[PE1-bgp] l2vpn-family evpn
[PE1-bgp-af-evpn] policy vpn-target
[PE1-bgp-af-evpn] peer FC01::5 enable
[PE1-bgp-af-evpn] peer FC01::5 advertise-community
[PE1-bgp-af-evpn] peer FC01::5 advertise encap-type srv6
[PE1-bgp-af-evpn] peer FC01::6 enable
```

```
[PE1-bgp-af-evpn] peer FC01::6 advertise-community
[PE1-bgp-af-evpn] peer FC01::6 advertise encap-type srv6
```

PE1 is used as an example. The configurations of other PEs are similar to the configuration of PE1.

By default, EVPN routes advertised by a local device to its peers carry the MPLS encapsulation attribute, which cannot be used for SRv6 forwarding. To enable EVPN routes to recurse to SRv6 tunnels, run the **peer advertise encap-type srv6** command.

Configure RRs.

```
[P1]bgp 65001
[P1-bgp] router-id 1.0.0.5
[P1-bgp] undo default ipv4-unicast
[P1-bgp] group RR-ipv6 internal
[P1-bgp] peer RR connect-interface LoopBack0
[P1-bgp] peer FC01::1 group RR-ipv6
[P1-bgp] peer FC01::2 group RR-ipv6
[P1-bgp] peer FC01::3 group RR-ipv6
[P1-bgp] peer FC01::4 group RR-ipv6
[P1-bgp] peer 2000::102 group RR-ipv6
[P1-bgp] l2vpn-family evpn
[P1-bgp-af-evpn] undo policy vpn-target
[P1-bgp-af-evpn] peer FC01::1 enable
[P1-bgp-af-evpn] peer FC01::1 reflect-client
[P1-bgp-af-evpn] peer FC01::1 advertise-community
[P1-bgp-af-evpn] peer FC01::1 advertise encap-type srv6
[P1-bgp-af-evpn] peer FC01::2 enable
[P1-bgp-af-evpn] peer FC01::2 reflect-client
[P1-bgp-af-evpn] peer FC01::2 advertise-community
[P1-bgp-af-evpn] peer FC01::2 advertise encap-type srv6
[P1-bgp-af-evpn] peer FC01::3 enable
[P1-bgp-af-evpn] peer FC01::3 reflect-client
[P1-bgp-af-evpn] peer FC01::3 advertise-community
[P1-bgp-af-evpn] peer FC01::3 advertise encap-type srv6
[P1-bgp-af-evpn] peer FC01::4 enable
[P1-bgp-af-evpn] peer FC01::4 reflect-client
[P1-bgp-af-evpn] peer FC01::4 advertise-community
[P1-bgp-af-evpn] peer FC01::4 advertise encap-type srv6
```

P1 is used as an example. The configuration of P2 is similar to the configuration of P1.

Check EVPN peer relationships between PEs and RRs.

```
[P1]display bgp evpn peer

BGP local router ID : 1.0.0.5
Local AS number : 65001
Total number of peers: 4                    Peers in established state : 4

  Peer        V     AS  MsgRcvd   MsgSent  OutQ  Up/Down       State       PrefRcv
  FC01::1     4  65001     3291      3338     0   0047h26m     Established     1
  FC01::2     4  65001     7283      7349     0   0105h23m     Established     0
  FC01::3     4  65001     7273      7358     0   0105h23m     Established     1
  FC01::4     4  65001     3284      3317     0   0047h16m     Established     1
```

```
[P2]display bgp evpn peer
BGP local router ID : 1.0.0.6
Local AS number : 65001
Total number of peers: 4                    Peers in established state : 4

  Peer            V    AS  MsgRcvd  MsgSent  OutQ  Up/Down       State       PrefRcv
  FC01::1         4 65001     3291     3338     0  0047h26m Established         1
  FC01::2         4 65001     7283     7349     0  0105h23m Established         0
  FC01::3         4 65001     7273     7358     0  0105h23m Established         1
  FC01::4         4 65001     3284     3317     0  0047h16m Established         1
```

Check whether the peer relationships between PEs and RRs are normal.

Step 2 Establishment BGP-LS peer relationships.

Establish a BGP-LS peer relationship between each RR and iMaster NCE-IP for redundancy protection.

Establish BGP-LS peer relationships between PEs and RRs, so that RRs can report SRv6 Policy path status.

This section describes only device-side configurations. Controller-side configurations are described in the following sections.

Configure RRs.

```
[P1]bgp 65001
[P1-bgp] link-state-family unicast
[P1-bgp-af-ls] domain identifier 1.0.0.56
[P1-bgp-af-ls] peer 2000::102 enable
[P1-bgp-af-ls] peer 2000::102 reflect-client
[P1-bgp-af-ls] peer FC01::1 enable
[P1-bgp-af-ls] peer FC01::1 reflect-client
[P1-bgp-af-ls] peer FC01::2 enable
[P1-bgp-af-ls] peer FC01::2 reflect-client
[P1-bgp-af-ls] peer FC01::3 enable
[P1-bgp-af-ls] peer FC01::3 reflect-client
[P1-bgp-af-ls] peer FC01::4 enable
[P1-bgp-af-ls] peer FC01::4 reflect-client
```

P1 is used as an example.

Configure PEs.

```
[PE1]bgp 65001
[PE1-bgp] link-state-family unicast
[PE1-bgp-af-ls] peer FC01::5 enable
[PE1-bgp-af-ls] peer FC01::6 enable
```

PE1 is used as an example.

Check BGP-LS peer status.

```
[P1] display bgp link-state unicast peer
```

```
BGP local router ID : 1.0.0.5
Local AS number : 65001
Total number of peers: 6                 Peers in established state : 6

 Peer            V    AS  MsgRcvd  MsgSent  OutQ  Up/Down       State   PrefRcv
 2000::102       4    65001       0        0     0  30s           Active        0
 FC01::1         4    65001     283      507     0  03:02:37      Established 136
 FC01::2         4    65001     280      501     0  03:00:28      Established 137
 FC01::3         4    65001     281      498     0  03:00:27      Established 136
 FC01::4         4    65001     281      570     0  03:00:51      Established 136
```

On P1, you can find that the BGP-LS peer relationship between P1 and PE1 is normal. After the controller-side configurations are complete, the BGP-LS peer relationship between P1 and iMaster NCE-IP enters the **Established** state.

Step 3 Establish BGP SR Policy peer relationships.

Establish BGP SRv6 Policy peer relationships between PEs and RRs, between PEs and iMaster NCE-IP, and between RRs and iMaster NCE-IP, so that iMaster NCE-IP can deliver SRv6 Policy configurations to PEs through RRs.

Configure PEs.

```
[PE1]bgp 65001
[PE1-bgp]ipv6-family sr-policy
[PE1-bgp-af-ipv6-srpolicy] undo bestroute nexthop-resolved ip
[PE1-bgp-af-ipv6-srpolicy] peer FC01::5 enable
[PE1-bgp-af-ipv6-srpolicy] peer FC01::6 enable
```

In the IPv6 SR-Policy address family, remove the restriction that routes can be used for route selection when the next hop is iterated to an IP address.

PE1 is used as an example. The configurations of other PEs are similar to the configuration of PE1.

Configure RRs.

```
[P1]bgp 65001
[P1-bgp]ipv6-family sr-policy
[P1-bgp-af-ipv6-srpolicy] undo bestroute nexthop-resolved ip
[P1-bgp-af-ipv6-srpolicy] undo router-id filter
[P1-bgp-af-ipv6-srpolicy] peer 2000::102 enable
[P1-bgp-af-ipv6-srpolicy] peer 2000::102 reflect-client
[P1-bgp-af-ipv6-srpolicy] peer FC01::1 enable
[P1-bgp-af-ipv6-srpolicy] peer FC01::1 reflect-client
[P1-bgp-af-ipv6-srpolicy] peer FC01::1 advertise-ext-community
[P1-bgp-af-ipv6-srpolicy] peer FC01::2 enable
[P1-bgp-af-ipv6-srpolicy] peer FC01::2 reflect-client
[P1-bgp-af-ipv6-srpolicy] peer FC01::2 advertise-ext-community
[P1-bgp-af-ipv6-srpolicy] peer FC01::3 enable
[P1-bgp-af-ipv6-srpolicy] peer FC01::3 reflect-client
[P1-bgp-af-ipv6-srpolicy] peer FC01::3 advertise-ext-community
[P1-bgp-af-ipv6-srpolicy] peer FC01::4 enable
[P1-bgp-af-ipv6-srpolicy] peer FC01::4 reflect-client
[P1-bgp-af-ipv6-srpolicy] peer FC01::4 advertise-ext-community
```

P1 is used as an example. Establish an SR Policy peer relationship between P1 and iMaster NCE-IP, configure PEs as the RR-clients of P1, and enable the function to send extended community attributes to RR clients on P1. In the configurations delivered by the controller, the tunnel color is carried in the extended community attribute. Therefore, this configuration is mandatory.

Note that router ID filtering must be disabled on each RR. In the scenario where RRs are used to push inbound traffic optimization information, all PEs receive traffic optimization policy routes sent by the RRs. To prevent a device from receiving a large number of traffic optimization policy routes that are irrelevant to the device, each traffic optimization policy route carries an extended community attribute in the format of an IP address to identify the node that needs to receive the route. PEs can then filter out unwanted routes based on the extended community attribute. However, RRs need to receive all traffic optimization policy routes from the controller. Therefore, you need to disable this feature on RRs. Disabling this feature is similar to disabling VPNv4 route RT check on RRs.

Check BGP SR Policy peer status.

```
[P1]display bgp sr-policy ipv6 peer

BGP local router ID : 1.0.0.5
Local AS number : 65001
Total number of peers: 5                    Peers in established state : 5

 Peer        V     AS      MsgRcvd   MsgSent  OutQ  Up/Down      State          PrefRcv
 2000::102   4     65001   0         0        0     0106h22m     Active         7
 FC01::1     4     65001   3360      3406     0     0048h25m     Established    0
 FC01::2     4     65001   7352      7417     0     0106h22m     Established    0
 FC01::3     4     65001   7341      7427     0     0106h22m     Established    0
 FC01::4     4     65001   3353      3387     0     0048h15m     Established    0
```

BGP SRv6 Policy peer relationships are established between the RR and PEs. After the BGP configuration is complete on iMaster NCE-IP (2000::102), the BGP SRv6 Policy peer relationship with iMaster NCE-IP is also established.

3.2.2.5 Controller-Side Basic Configurations

Step 1 Configure routes.

On the maintenance page of iMaster NCE-IP, configure routes to NEs.

Log in to the maintenance page and choose **maintain** > **Network Configuration** > **Configure routing** from the main menu.

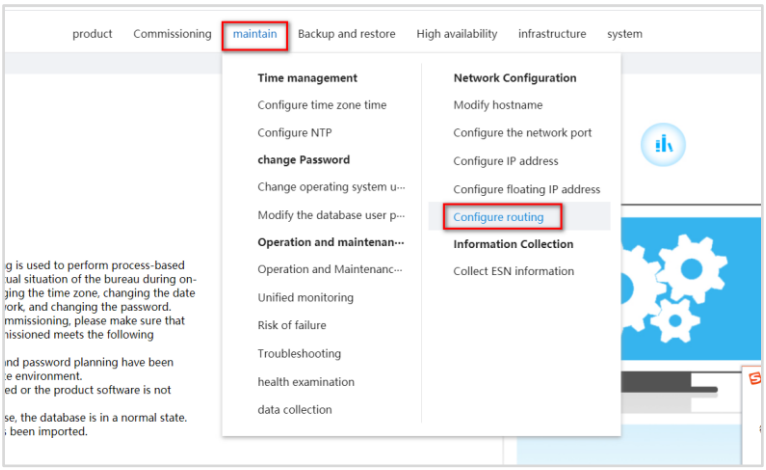

On the **Configure routing** page, select **Analyzer_01**, **Controller_01**, and **NMS_01**, and click **Inquire**.

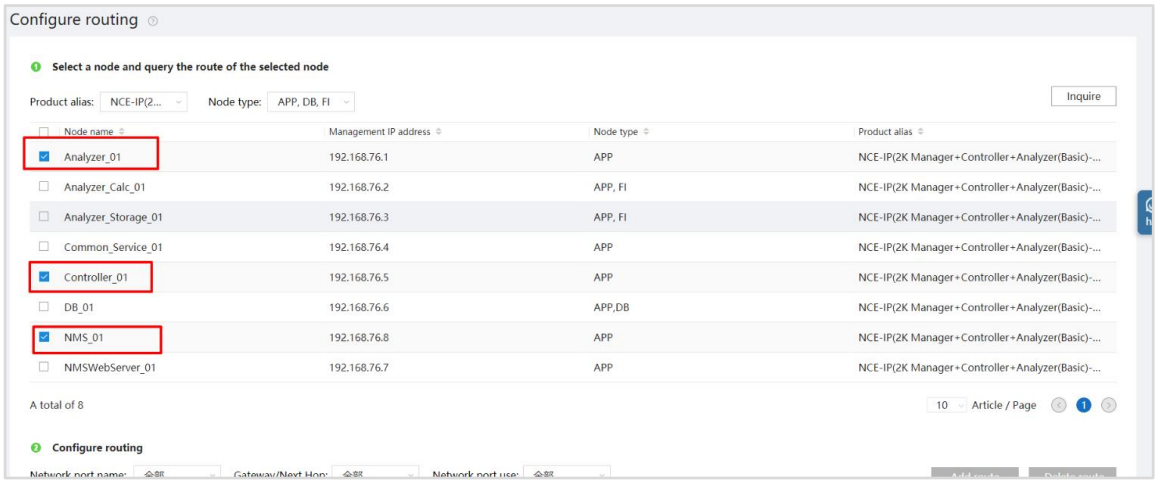

After the query is complete, click **Add Route** at the bottom. In the **Add route** dialog box, select the three nodes previously selected and click **determine**. Then add routes to NEs in the **Add route** area according to the following table.

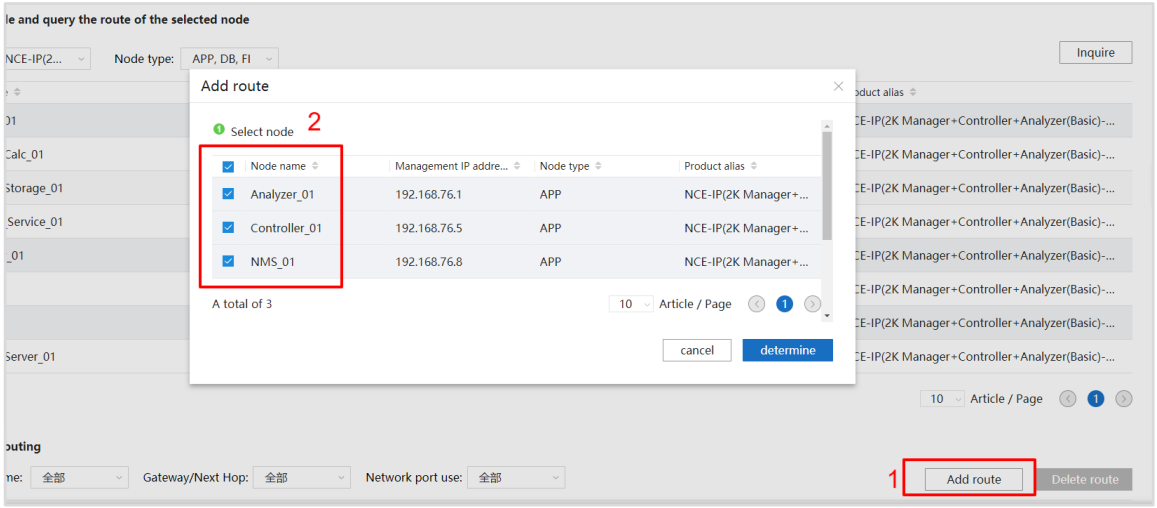

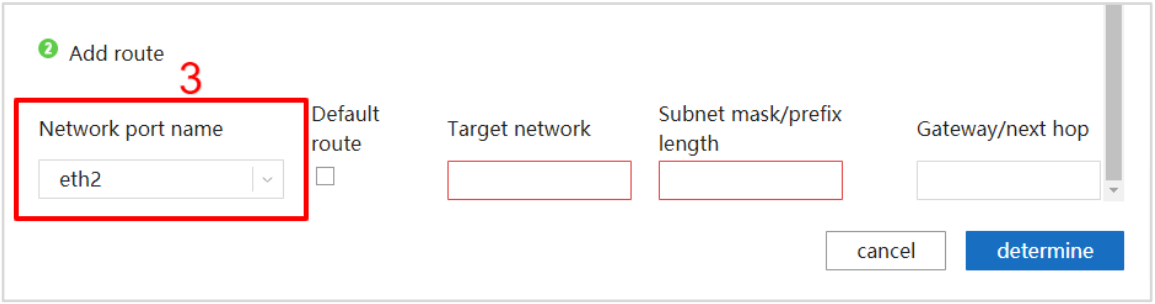

Table 3-5 Route planning

Destination Device	Destination Network	Subnet Mask/Prefix Length	Gateway/Next Hop
PE1	FC01::1	128	2000::1
PE2	FC01::2	128	2000::2
PE3	FC01::3	128	2000::3
PE4	FC01::4	128	2000::4
P1	FC01::5	128	2000::5
P2	FC01::6	128	2000::6

Step 2 Add NEs for management, create links and topologies, and synchronize NE data.

Some operations are similar to those in 3.1 SR-MPLS Service Delivery by the Controller.

Currently, NEs can only be managed based on IPv4 addresses. The IPv4 address configurations are consistent with those in 3.1 SR-MPLS Service Delivery by the Controller.

3.2.2.6 Controller-Side BGP Configurations

To enable iMaster NCE-IP to receive link, bandwidth, and other information from NEs, configure iMaster NCE-IP to establish BGP-LS peer relationships with RRs.

To enable iMaster NCE-IP to deliver SRv6 Policy configurations to NEs, establish BGP SR Policy peer relationship between iMaster NCE-IP and RRs.

Step 1 Configure basic BGP functions.

Open the Network Management app and choose **Configuration** > **Control Unit** > **BGP** from the main menu to configure BGP.

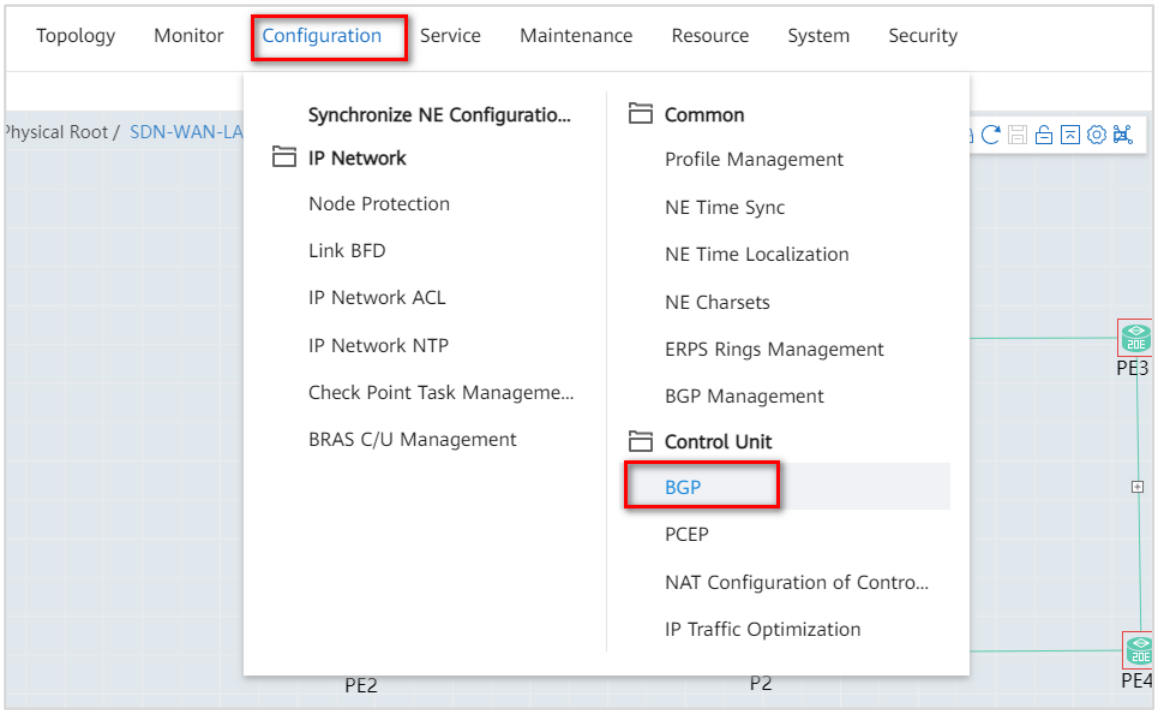

Set the parameters as follows.

Table 3-6 Basic BGP parameters

Parameter	Value
*Enable BGP	Yes
*Local AS number	65001
Local IPv6 address	2000::102
*Router ID	172.21.17.102

Set basic BGP parameters according to the parameter planning.

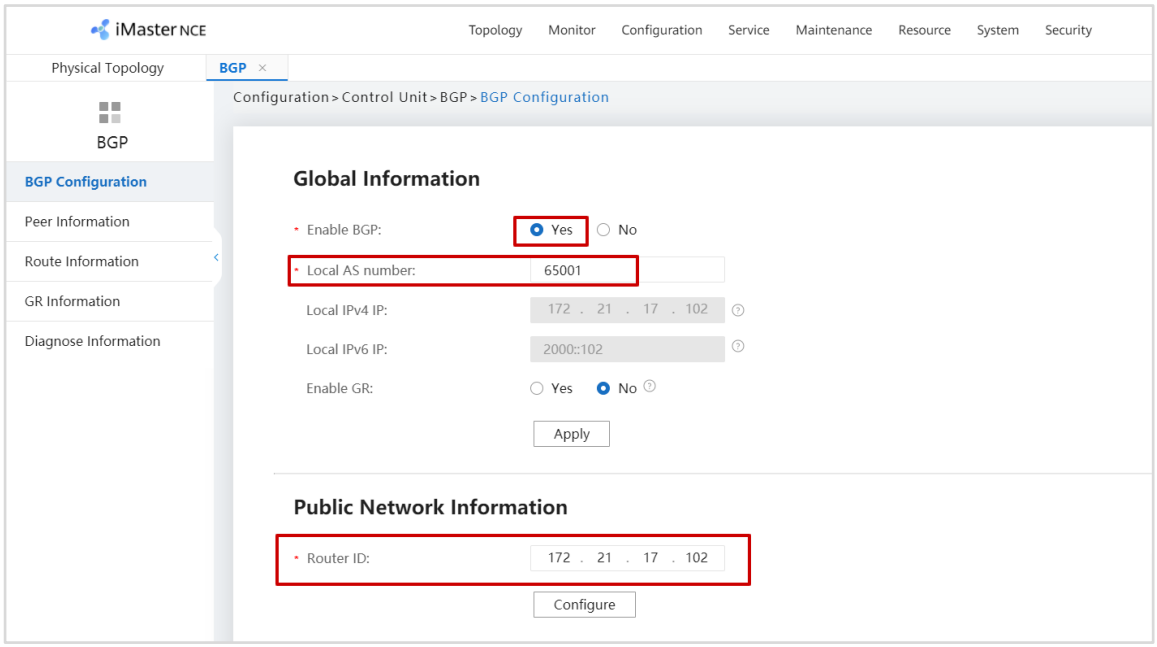

Step 2 Create BGP peers.

Create BGP peers, which will be invoked from specific address families.

On the **Basic Peer Information** tab page of the BGP configuration page, click **Create Peer**.

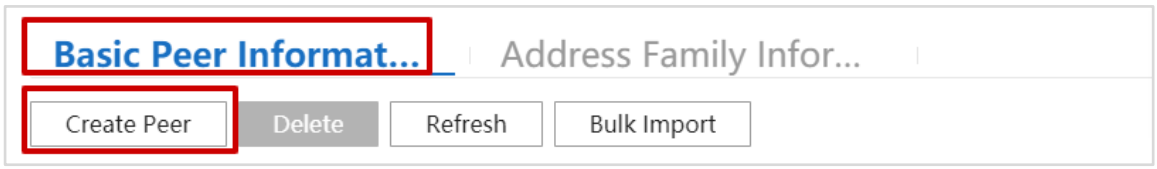

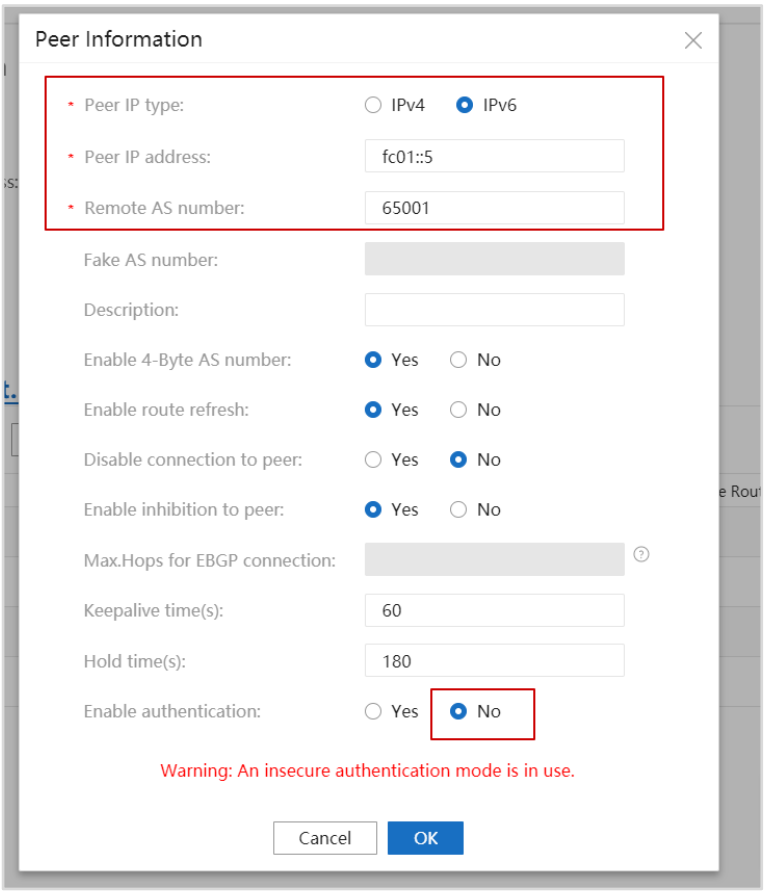

Create IPv6 peers fc01::5 and fc01::6 (corresponding to P1 and P2, which serve as RRs), and disable authentication.

Step 3 Enable IPv6 peers in the BGP-LS address family.

Enable IPv6 peers in the BGP-LS address family, so that iMaster NCE-IP can receive link, bandwidth, and other information from RRs.

On the **Address Family Information** tab page, click **Link-state**. (If **Link-state** is not displayed, click **Create Address Family** to add it.)

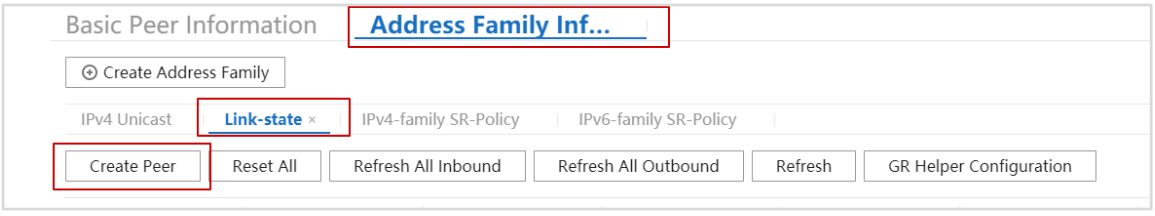

Click **Create Peer**. In the dialog box that is displayed, click **Select Peer**.

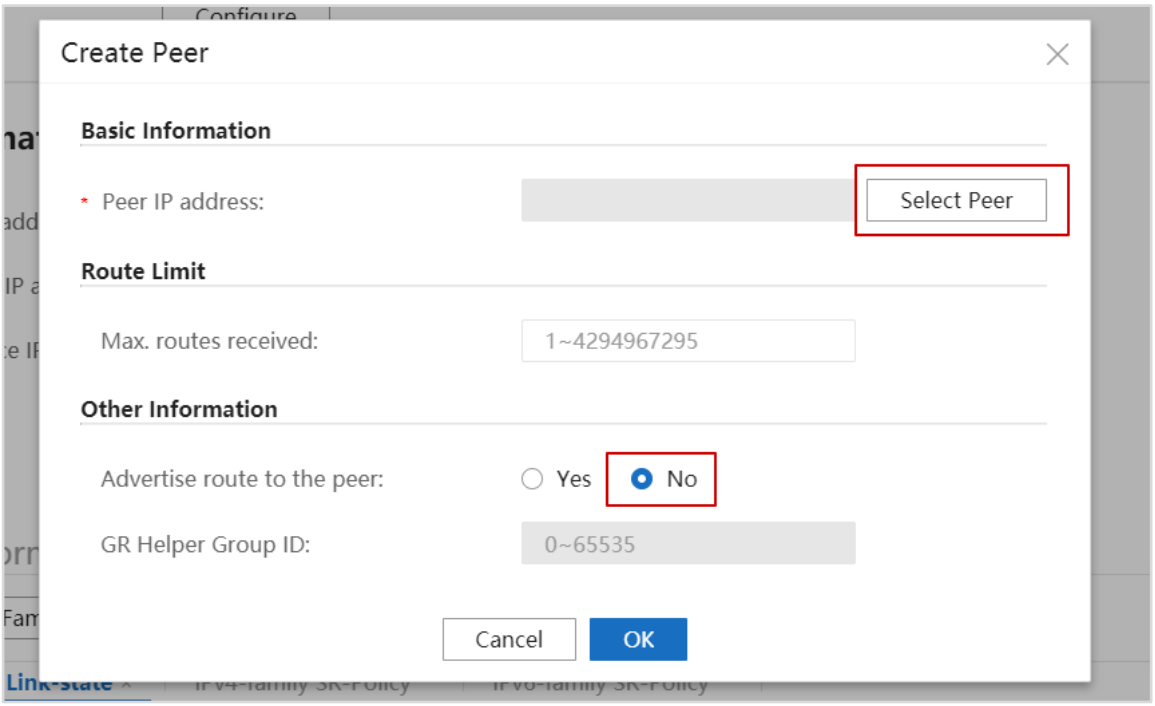

Note that iMaster NCE-IP does not need to send routes to RRs; instead, it only needs to receive routes. Therefore, you need to set **Advertise route to the peer** to **No**.

In the **Select Peer** dialog box, select the previously created peers fc01::5 and fc01::6.

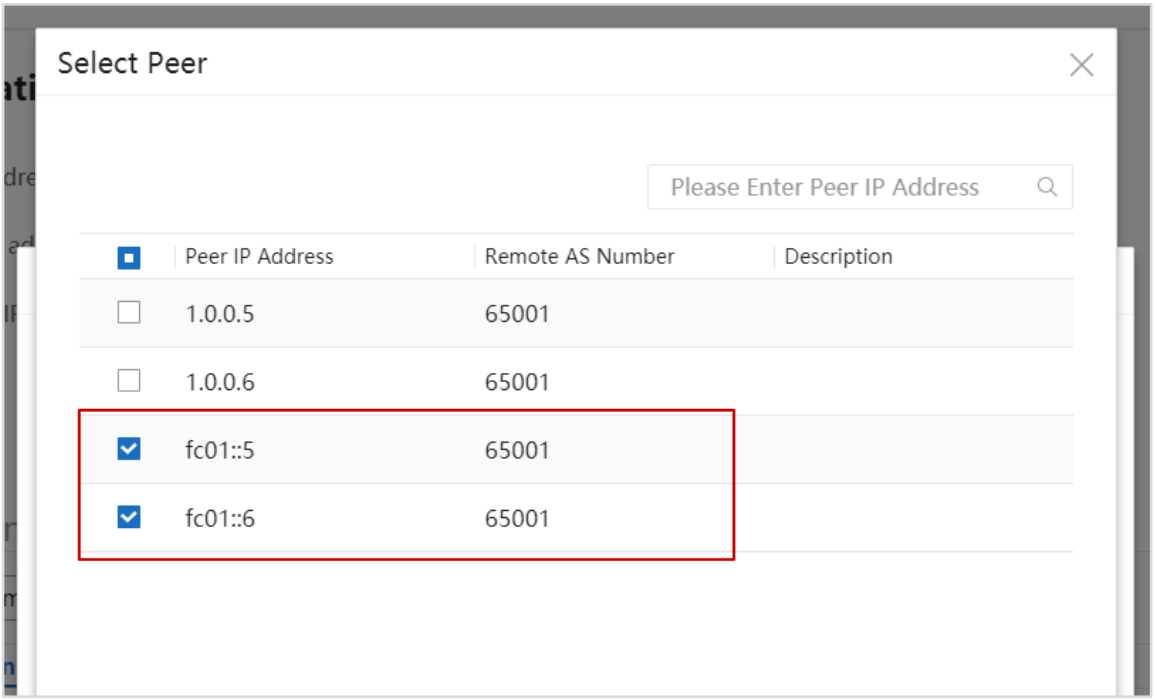

Step 4 Enable IPv6 peers in the BGP SRv6 Policy address family.

Enable IPv6 peers in the BGP SRv6 Policy address family, so that iMaster NCE-IP can deliver SRv6 Policy configurations to NEs.

On the **Address Family Information** tab page, select **IPv6-family SR-Policy**.

Click **Create Peer**. In the **Create Peer** dialog box, click **Select Peer**.

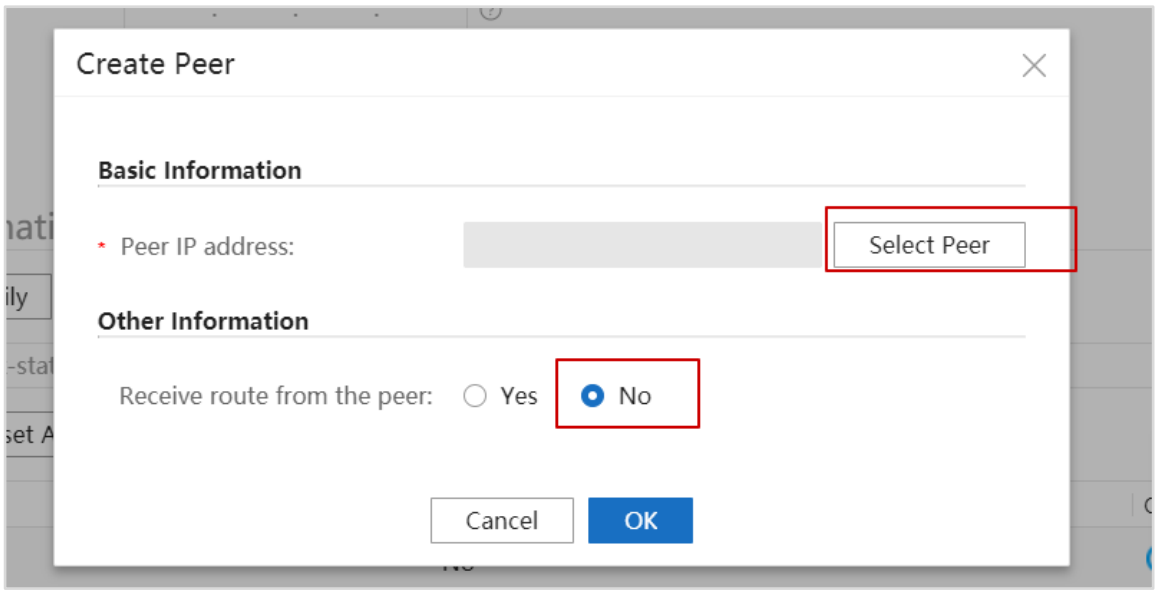

Select the previously created peers (fc01::5 and fc01::6).

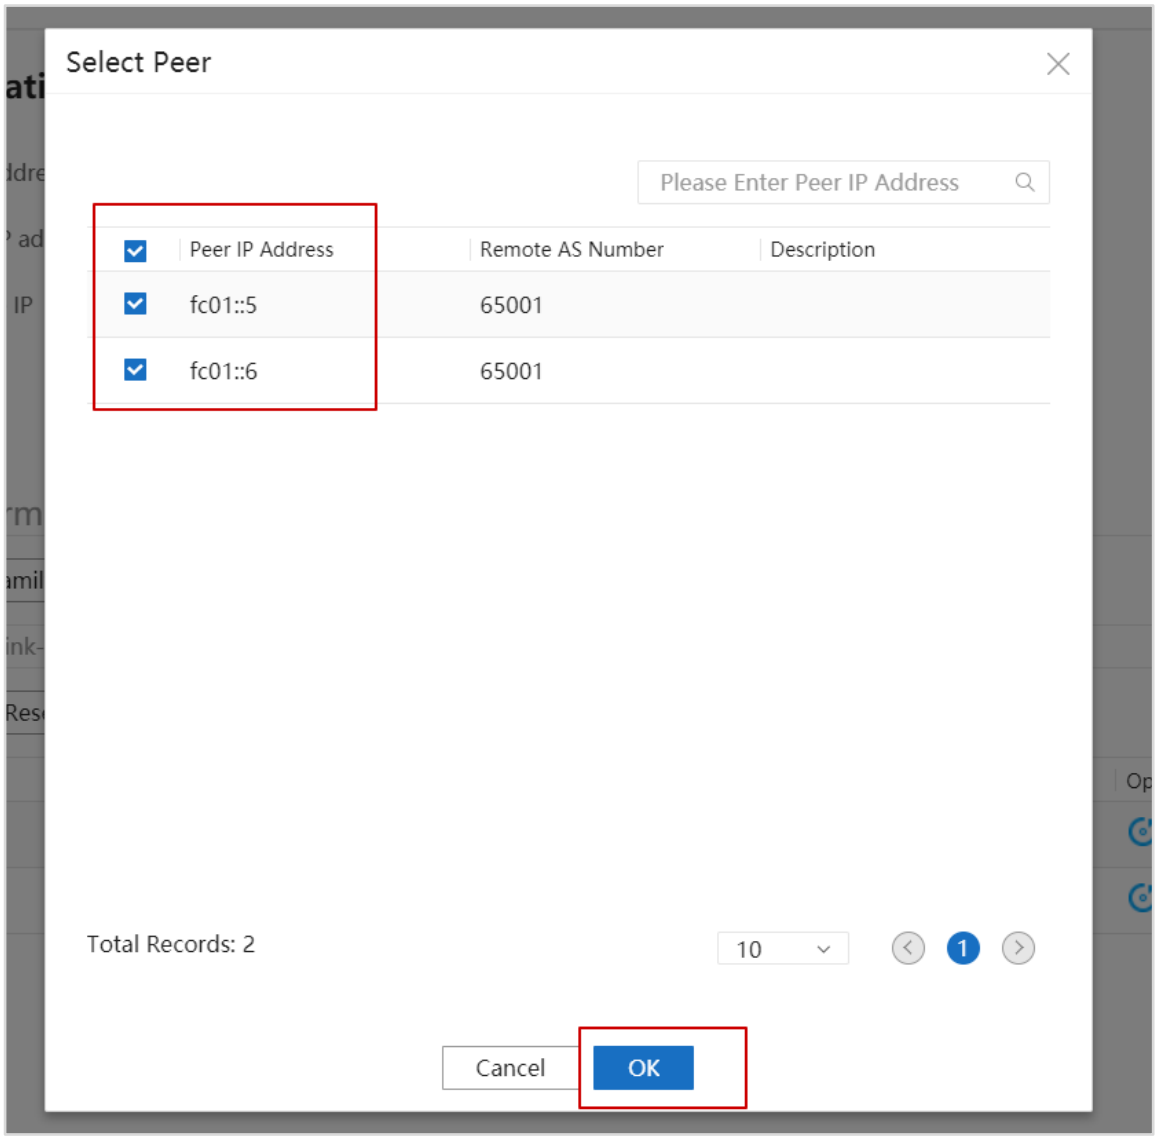

On the **BGP** page, click **Peer Information** to check BGP peer relationships.

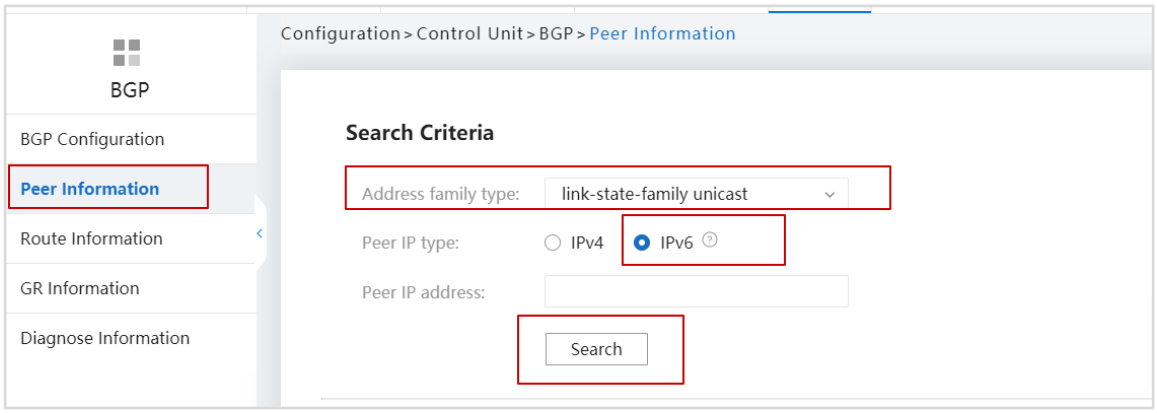

Set **Address family type** to **link-state-family unicast** and click **Search**.

The search results show that iMaster NCE-IP has established BGP-LS peer relationships with P1 and P2 and received route prefixes from P1 and P2.

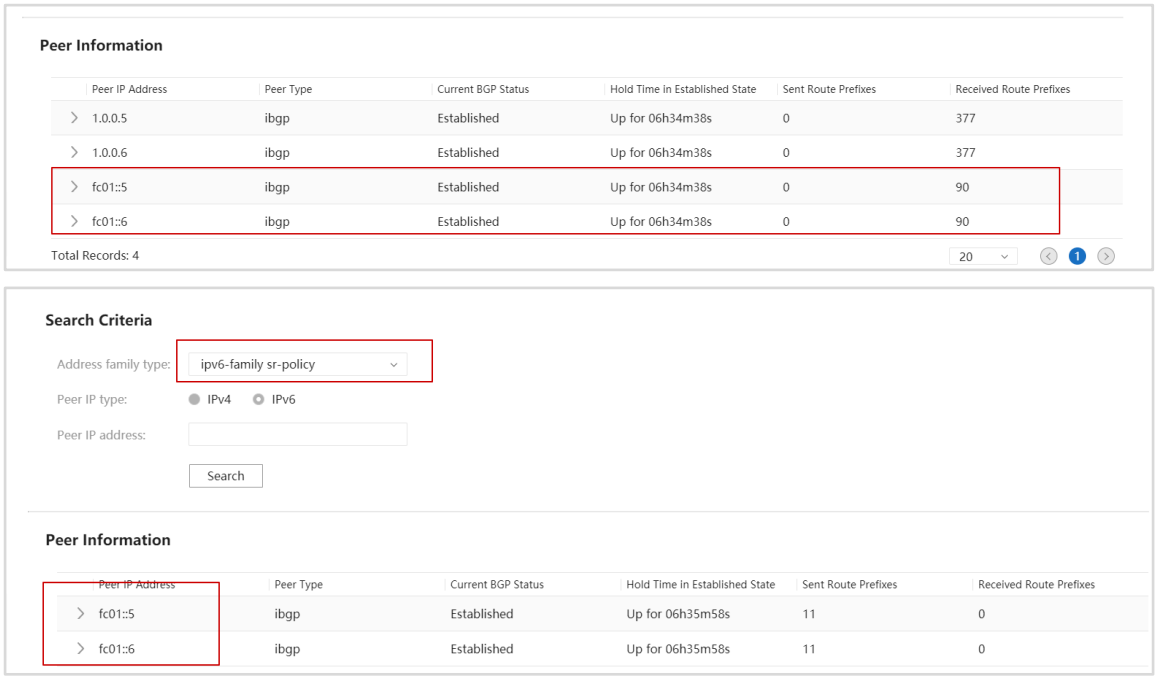

Set **Address family type** to **ipv6-family sr-policy** and click **Search**. The query result shows that BGP SRv6 Policy peer relationships with P1 and P2 have been established.

3.2.2.7 SRv6 Policy Configuration

Use the controller to establish a bidirectional SRv6 Policy between PE1 and PE4.

In this section, we will learn how to configure an SRv6 Policy on iMaster NCE-IP and create a color for the SRv6 Policy, so that L3VPN traffic transmitted along an EVPN route carrying the color extended community attribute can recurse to the SRv6 Policy.

Step 1 Create an SR Policy color.

Open the Network Management app and choose **Configuration** > **Common** > **Profile Management** from the main menu. Then click **SR Policy Color Profile**.

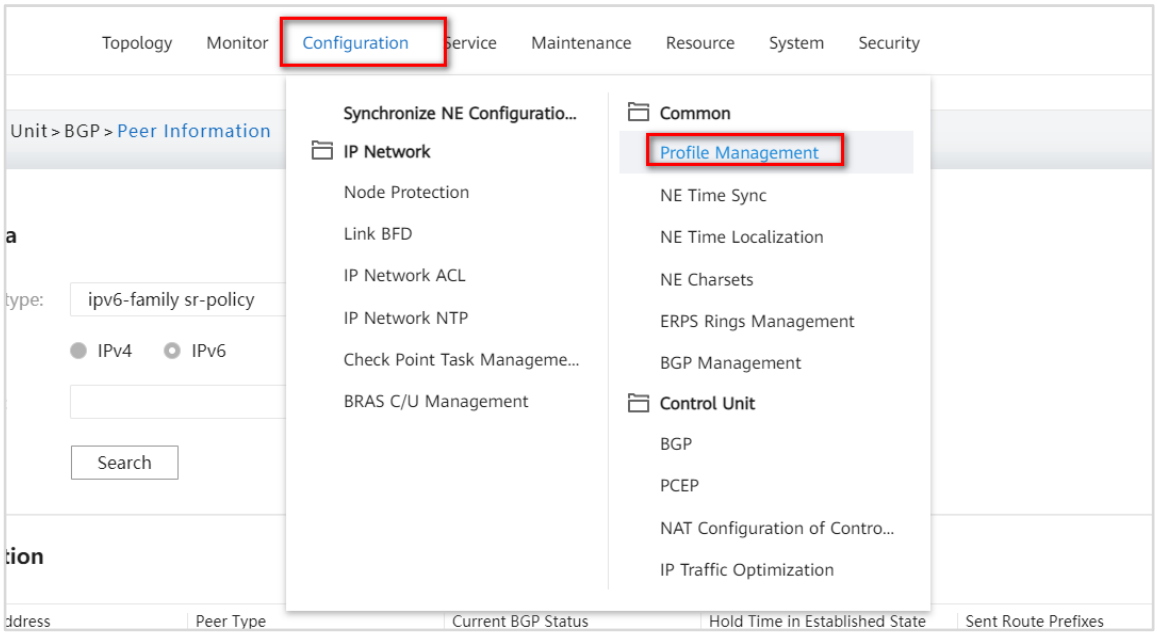

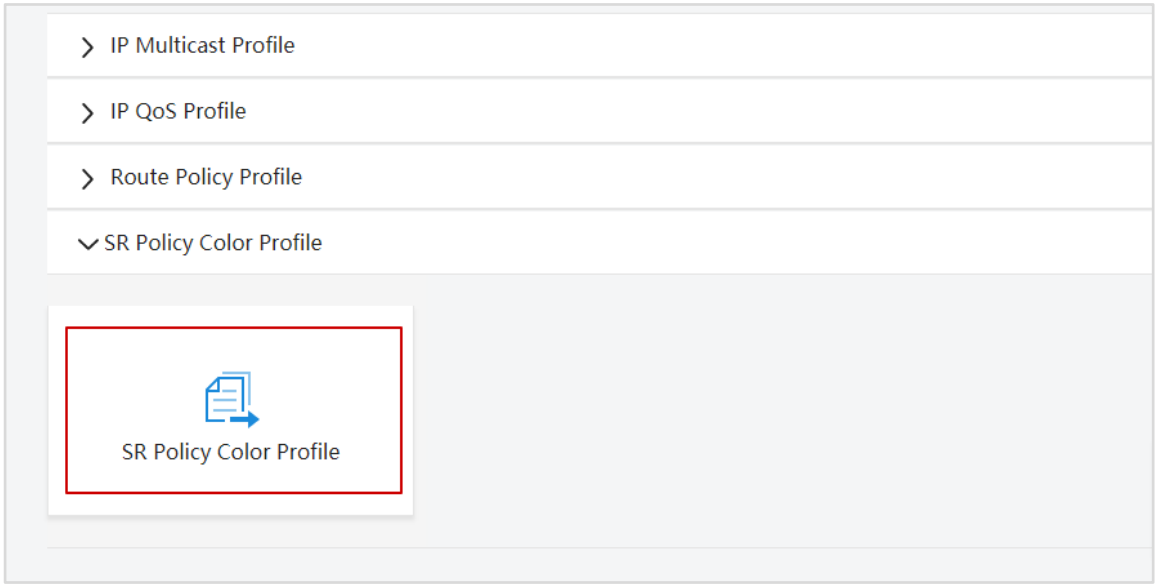

Create a color named **SRv6-PE1_PE4**.

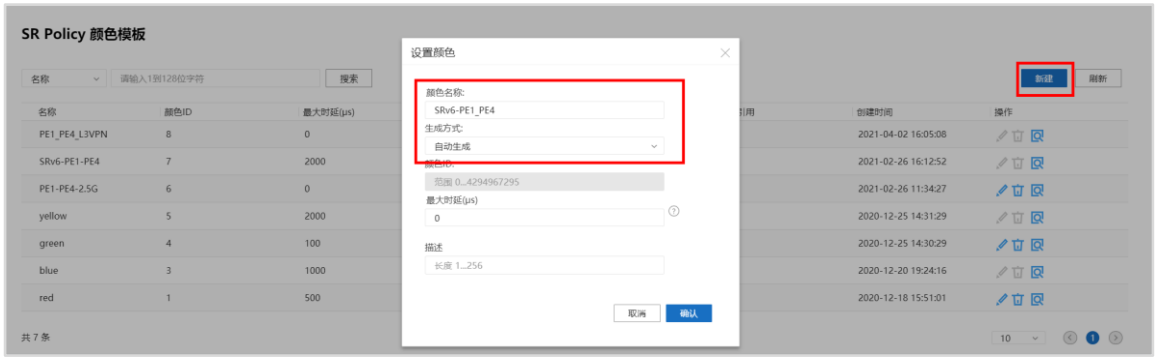

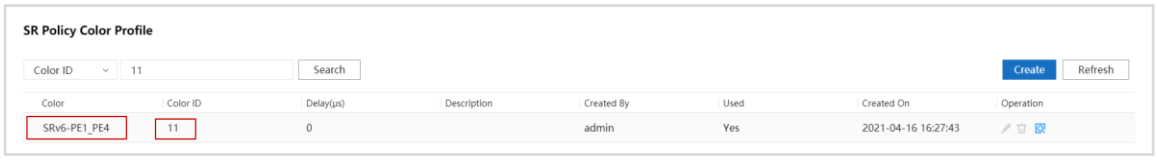

As shown in the figure, a color template with the color ID of 11 is created. After the template is referenced by an SRv6 Policy, color 11 is applied to the tunnel.

Step 2 Configure an SRv6 Policy.

Open the Network Management app and choose **Service** > **Create** > **SR Policy** from the main menu to create an SR Policy.

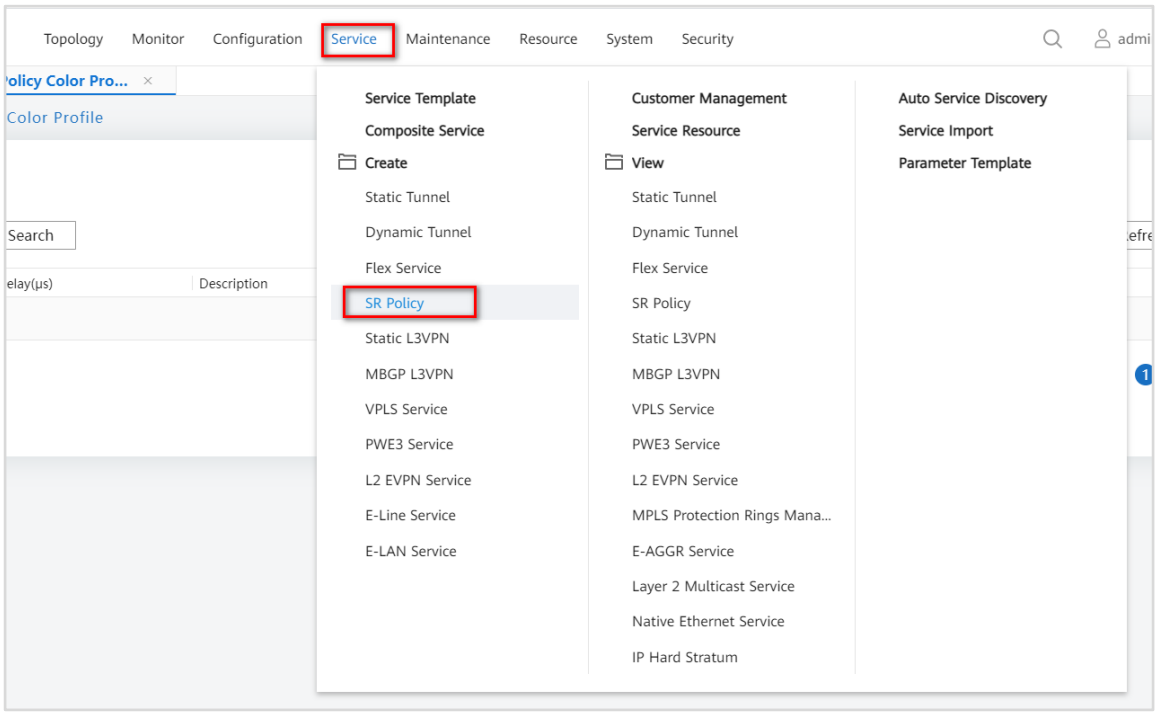

Set **Select a service template** to **SRv6_TE_Policy**.

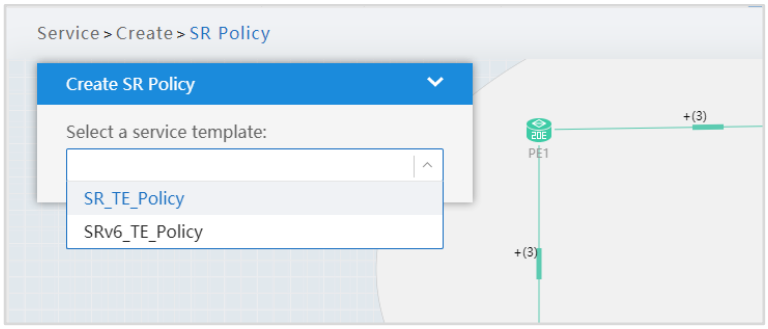

In the **Basic Information** area, set **Parameter template** to **SRv6_TE_Policy**, **Service name** to **PE1_PE4_L3EVPN**, and **Direction** to **Bidirectional**. Retain the default settings for other parameters.

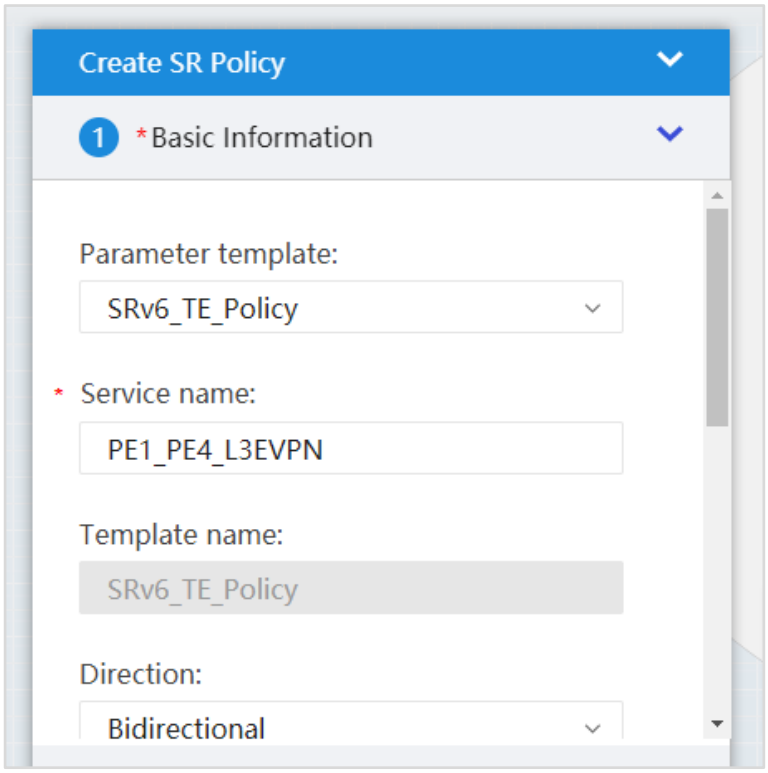

In the **NE List** area, set **Source NE** to **PE1** and **Sink NE** to **PE4**.

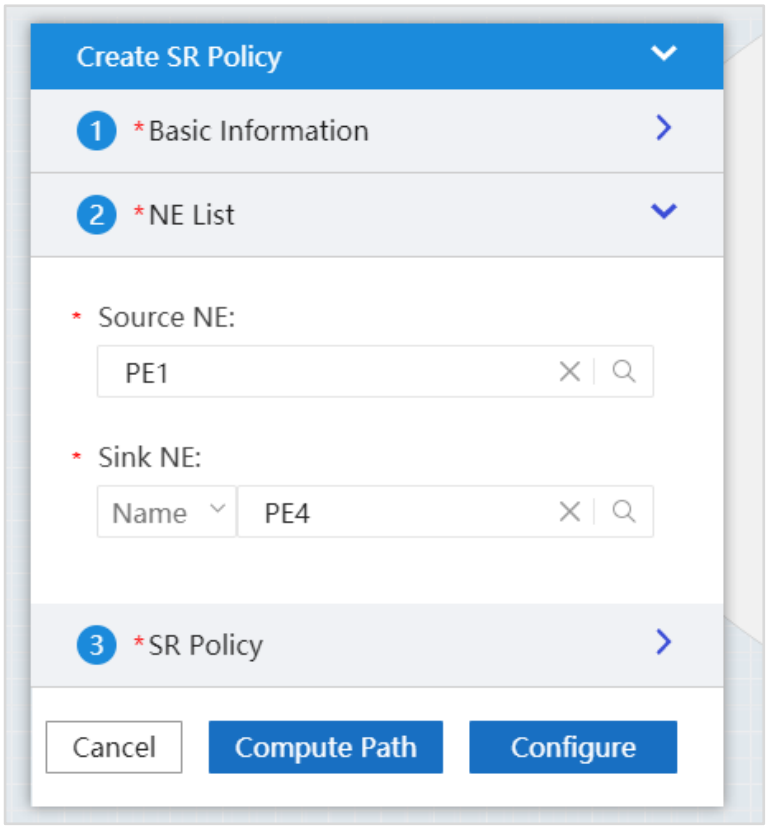

In the **SR Policy** area, configure the forward SR Policy. First, click the modify icon next to **Color** to configure the color.

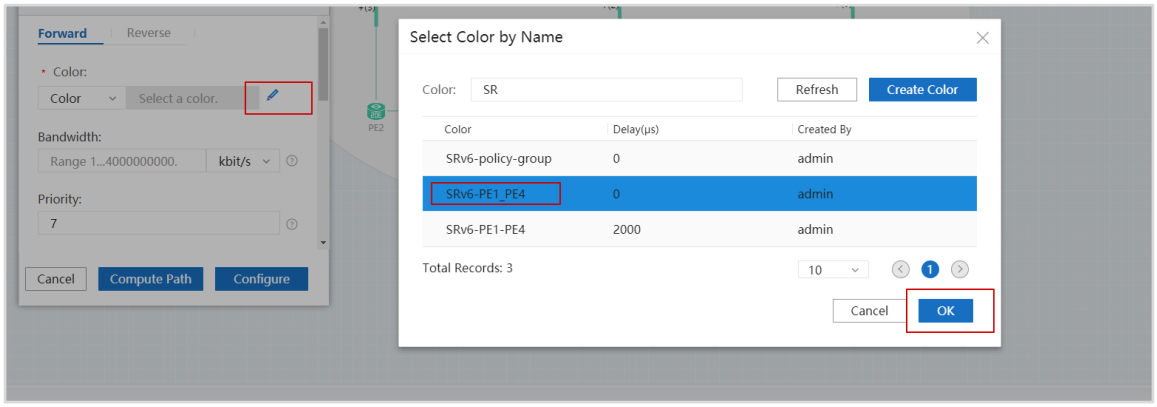

Select the previously created color **SRv6-PE1_PE4**.

Then, select a candidate path and set the candidate path preference to **65535** (the highest preference).

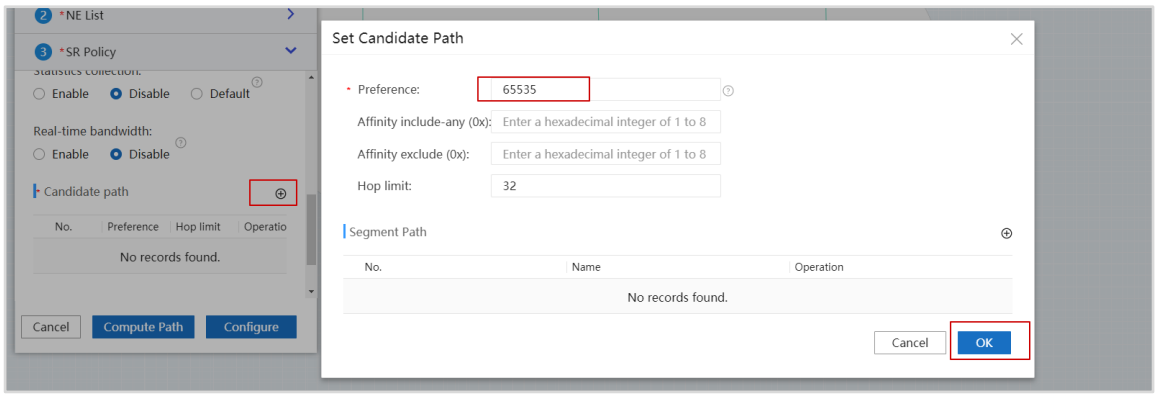

Repeat the preceding steps to configure the reverse SR Policy.

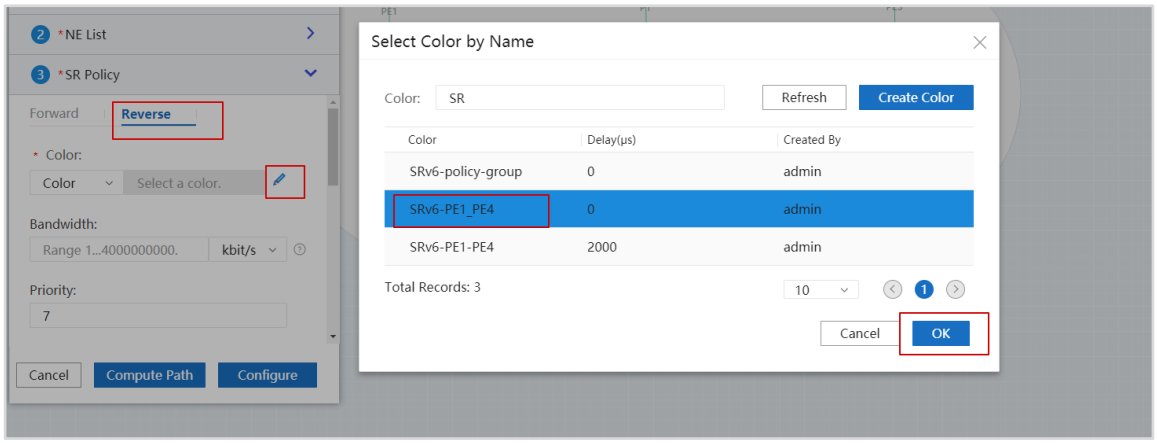

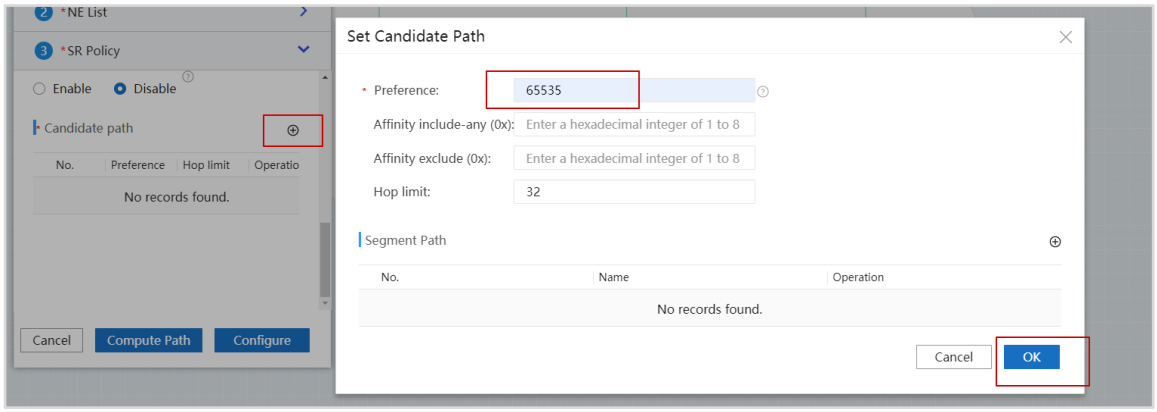

Finally, click **Compute Path**.

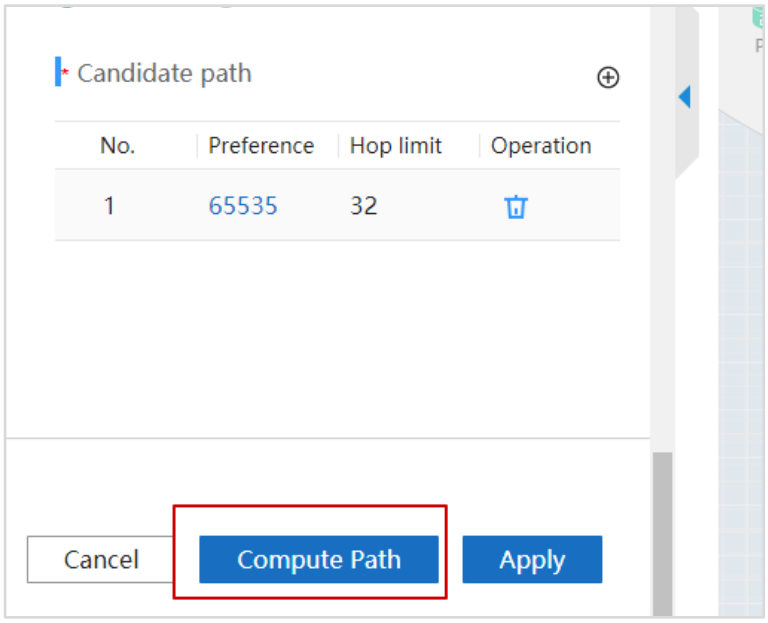

In the computation result area, select the path with the least cost and click **OK**. Then, click **Apply**.

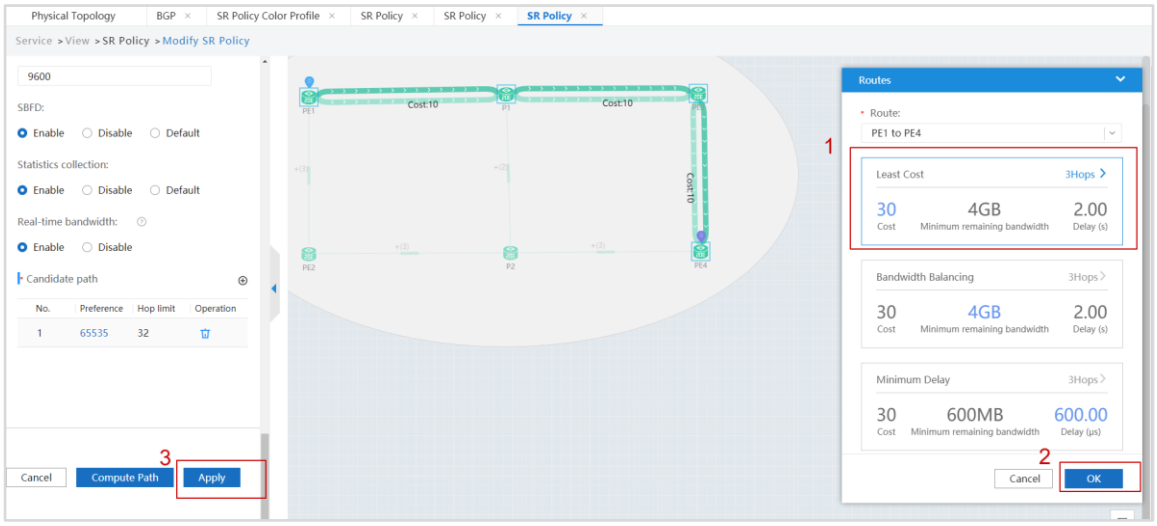

In the dialog box that is displayed, click **OK**.

In the **Success** dialog box, you can click the corresponding hyperlink to view SR Policy information.

Click **View SR Policy** to view the newly created SRv6 Policy.

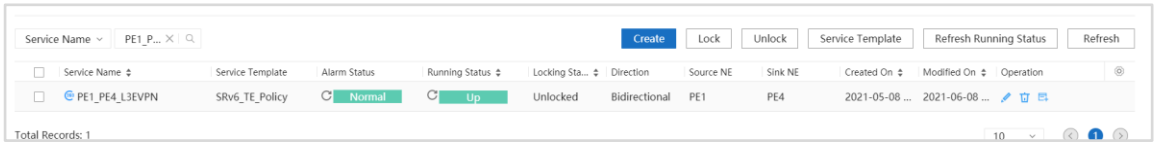

The running status and alarm status are normal.

Step 3 Configure route-policies.

To recurse EVPN L3VPN traffic to SRv6 Policies, configure a route-policy on PE1 and PE4 to add the color extended community attribute to EVPN routes to be advertised, so that these routes can recurse to SRv6 Policies.

Open the Network Management app and choose **Configuration** > **Common** > **Profile Management** from the main menu. Then, click **Route Policy Profile**.

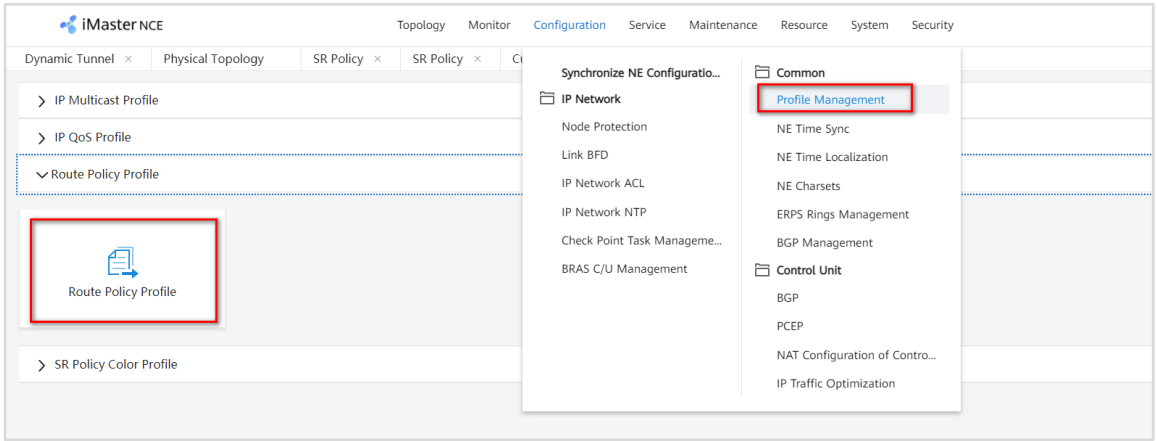

Click **Create** to create a route-policy template.

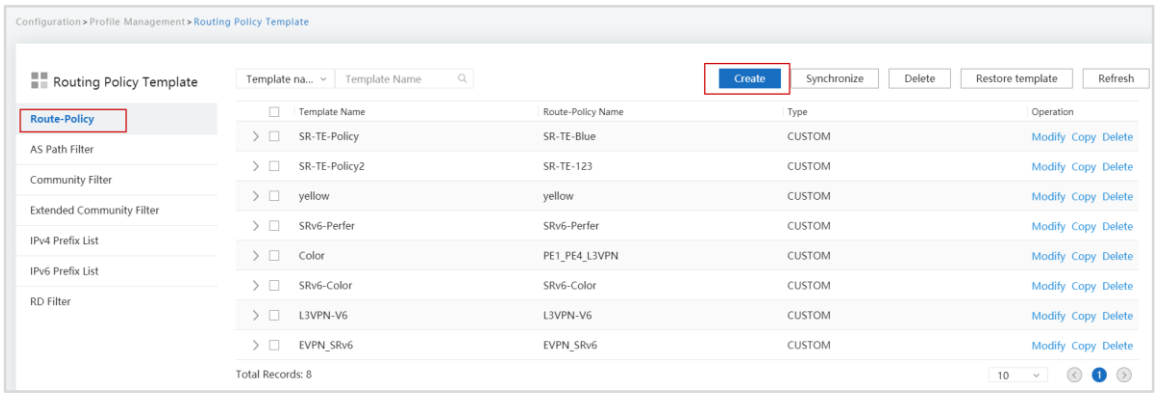

Set **Template name** to **EVPN_SRv6** and **Route-Policy name** to **EVPN_SRv6**.

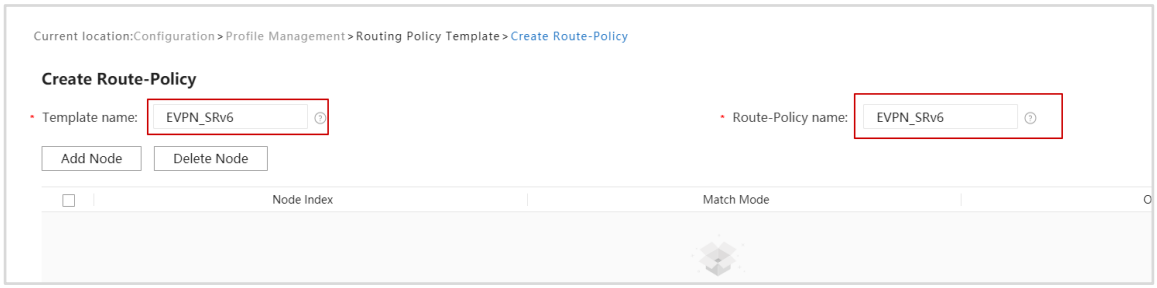

Click **Add Node**.

In the dialog box that is displayed, configure node information. Specifically, set **Node index** to **10**, retain the default value **Permit** for **Match mode**, and leave all match conditions unselected (indicating that all match conditions will be applied).

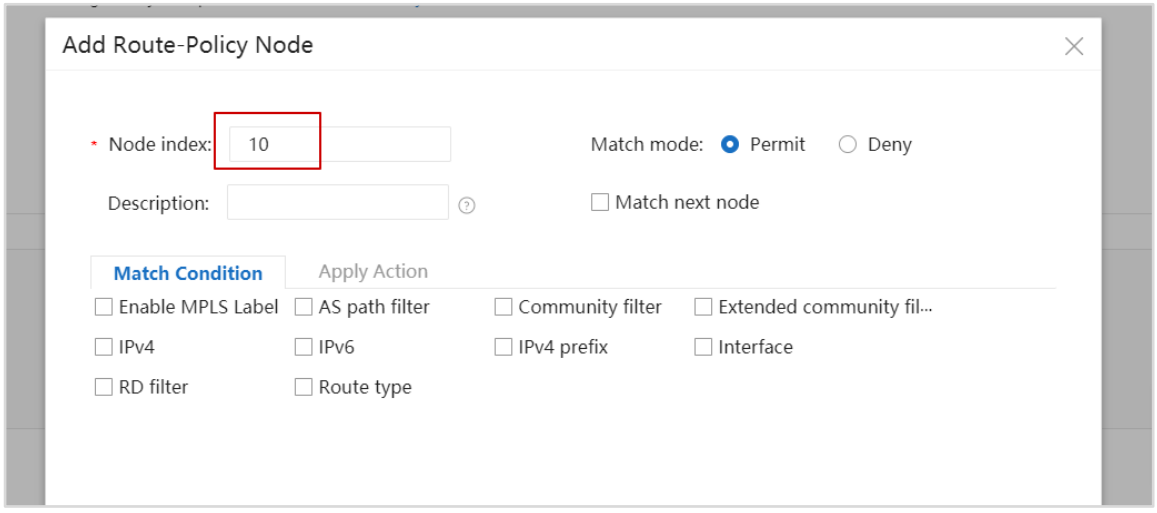

On the **Apply Action** tab page, select **Color**.

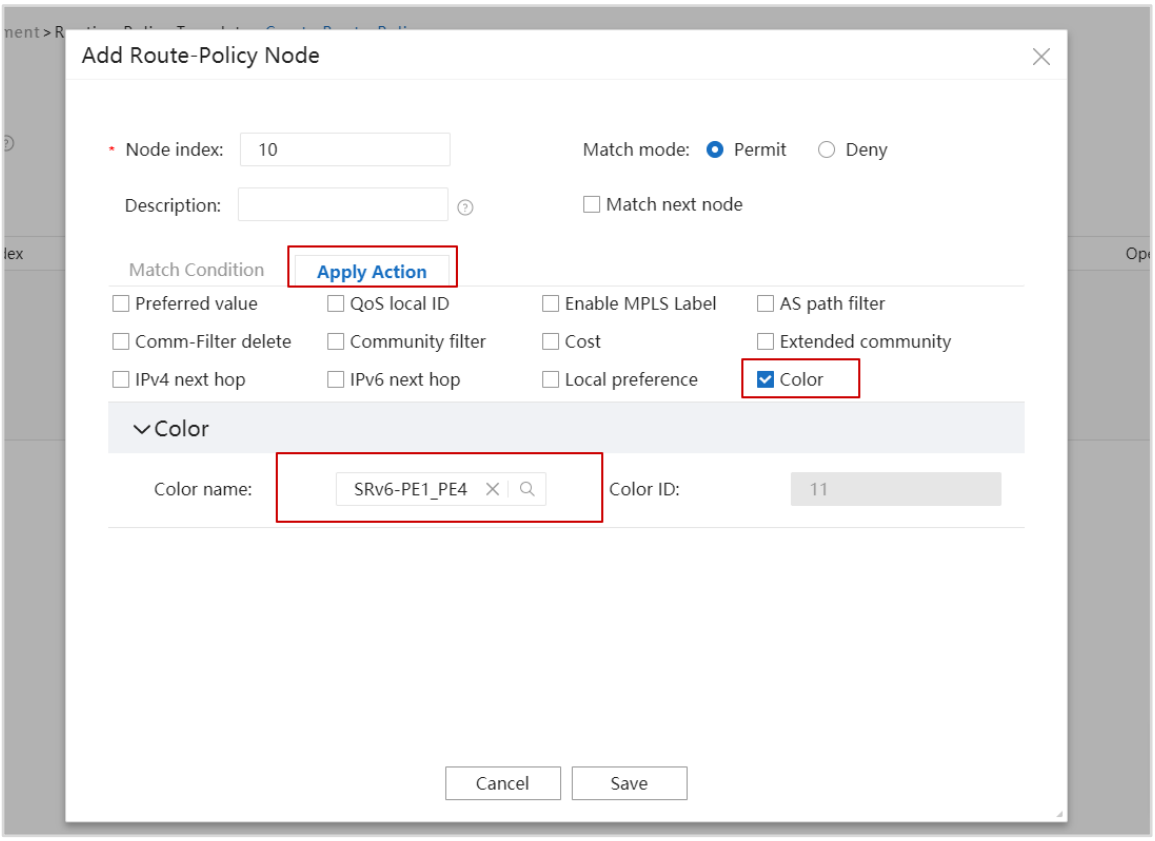

Select the previously created color **SRv6-PE1_PE4** and click **Save**. Then click **OK**. The route-policy template is created.

3.2.2.8 EVPN L3VPN Service Delivery by the Controller

Use the controller to deliver an EVPN L3VPN service and recurse the service to an SR Policy for traffic forwarding.

Create Loopback3 on PE1 and PE4 to simulate EVPN L3VPN access users.

PE1

```
PE1
[PE1]interface LoopBack3
[PE1-LoopBack3] ip address 172.20.1.1 255.255.255.255
```

PE4

```
[PE4]interface LoopBack3
[PE4-LoopBack3] ip address 172.20.4.1 255.255.255.255
```

After configuring loopback interface addresses on PE1 and PE4, synchronize NE configurations to iMaster NCE-IP.

Choose **Configuration** > **Synchronize NE Configuration Data** from the main menu. In the **Synchronize NE Configuration Data** dialog box, select PE1 and PE4 and click **Synchronize** to synchronize their configurations to iMaster NCE-IP.

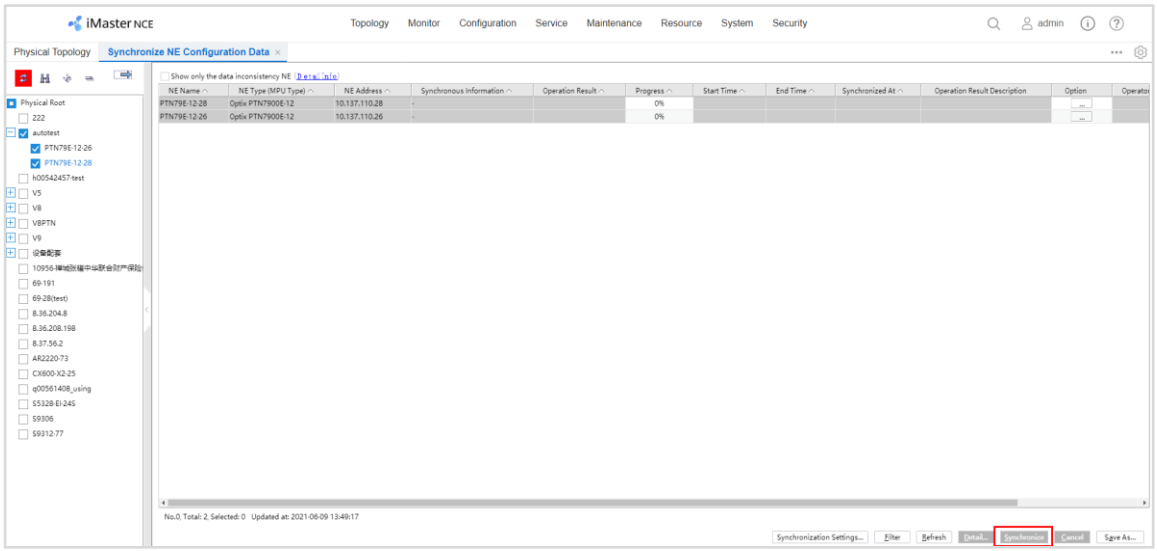

Open the Network Management app and choose **Service** > **MBGP L3VPN** from the main menu to configure EVPN L3VPNv4. EVPN L3VPNv4 configuration mainly consists of the following four aspects:

1. Basic parameters
2. Service nodes
3. Service access points
4. Tunnels

Step 1 Set basic parameters.

In the **Basic Parameter** area, set **Service template** to **L3VPN_EVPN_Auto_Select** (default value), set **Service name** and **VRF Name**, set **Parameter template** to **MBGP_L3VPN**, and set **Topology type** to **Any-to-Any**.

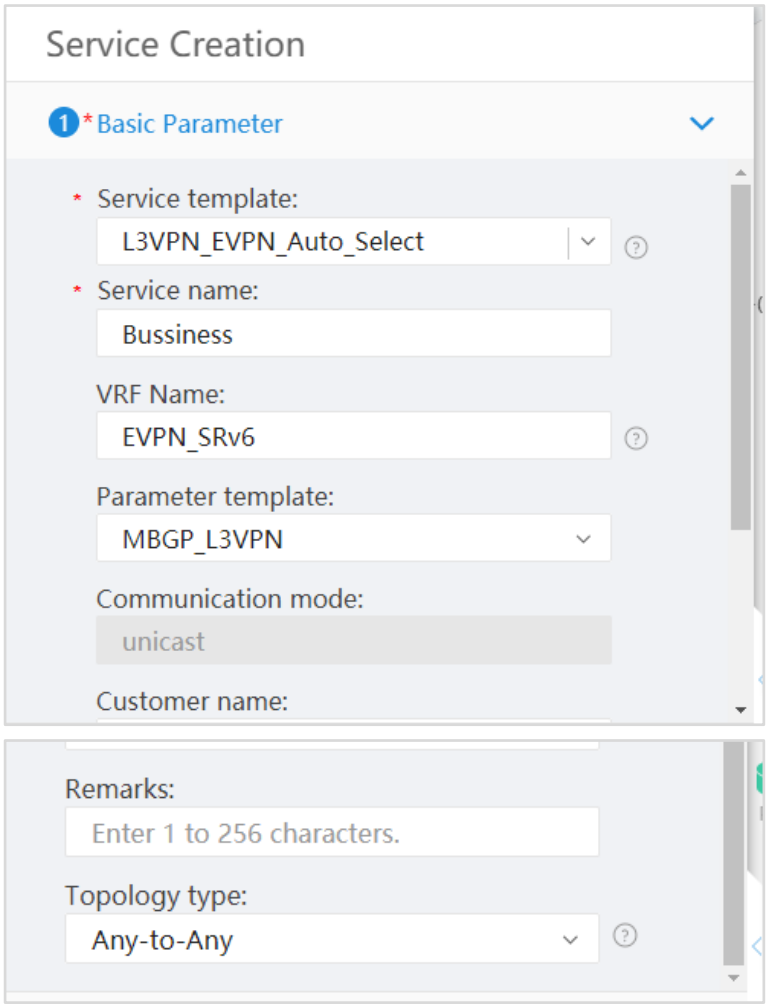

Step 2 Configure service nodes.

Configure service nodes to determine the tunnel ingress and egress. In this example, select PE1 and PE4.

Click **+** in the **Service Node** area.

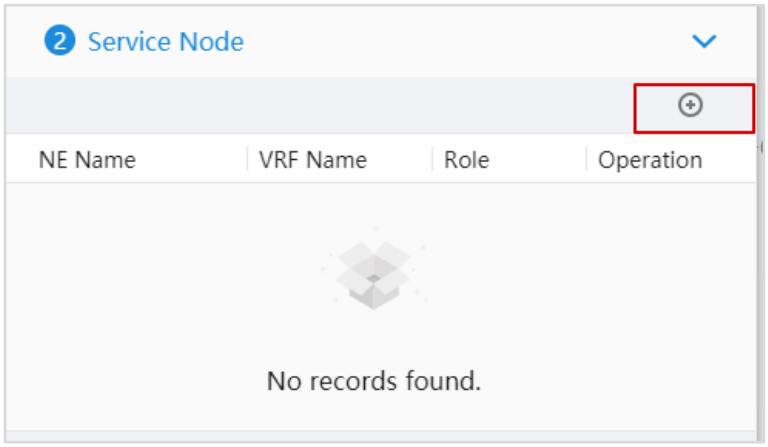

In the dialog box that is displayed, set **NE name** to **PE1** and select **Enable IPv4**.

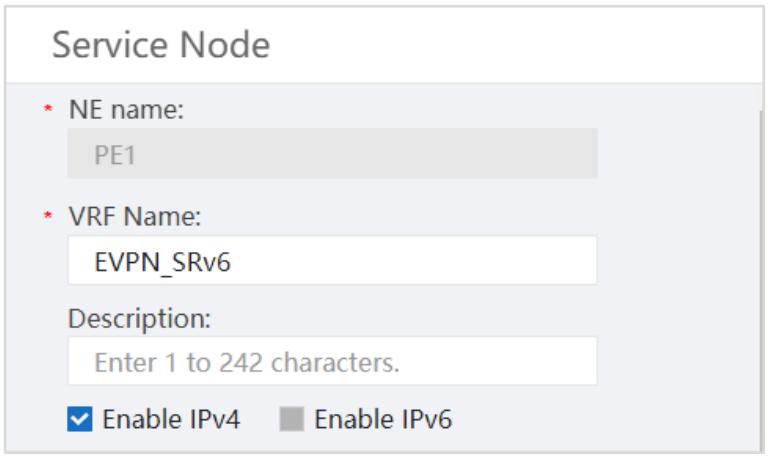

In the new area that is displayed, set the following parameters for the VRF used for service access. You only need to set the EVPN RD and RT values, and do not need to set the L3VPN RD and RT values. This can prevent VPNv4 routes from being imported. If VPNv4 routes are imported, route selection is affected.

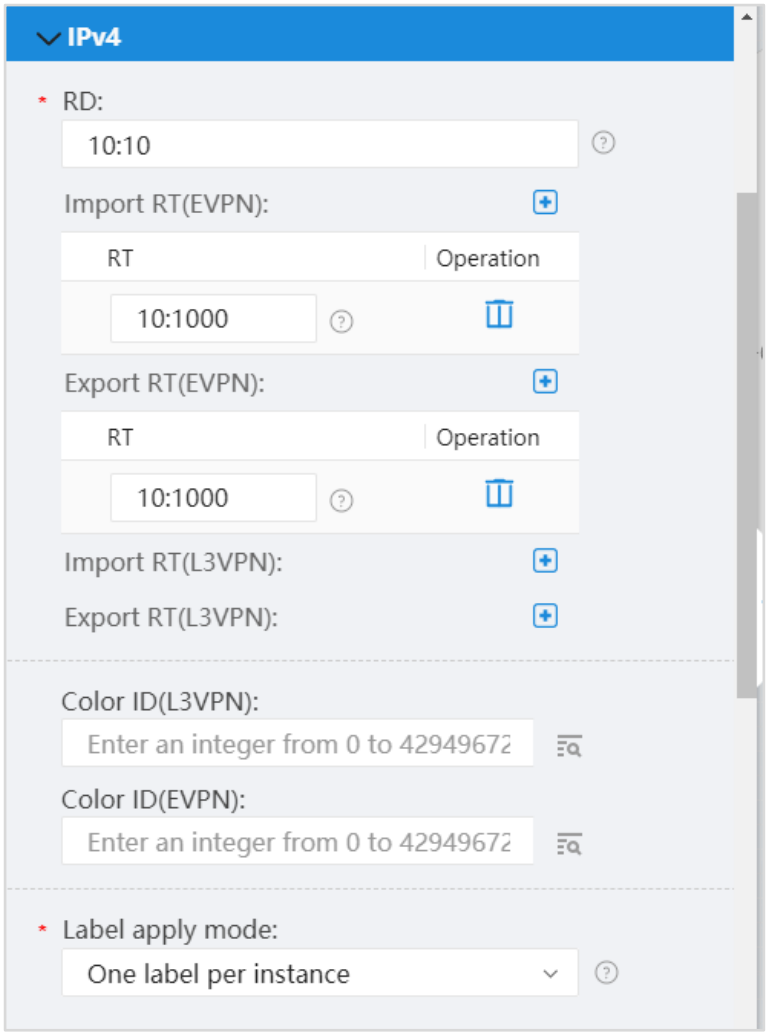

In the **Service Node** area, configure a tunnel policy. Specifically, click **+** next to **Tunnel Policy**. In the **Tunnel Policy** dialog box, click the **Tunnel Policy Evpn** tab and set **Policy type** to **Auto select**. In the **SRv6 tunnel policy** area, click **+** and set **Type** to **SRv6 Policy** for route recursion to SRv6 Policies. EVPN routes then recurse to SRv6 TE Policies based on the color attribute.

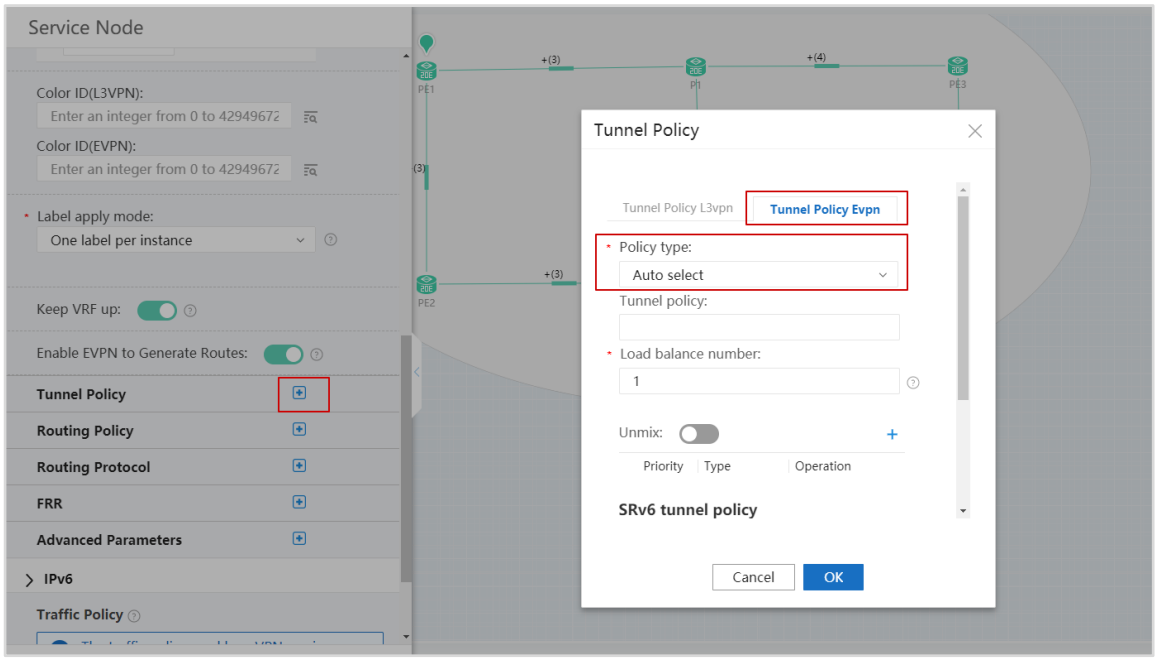

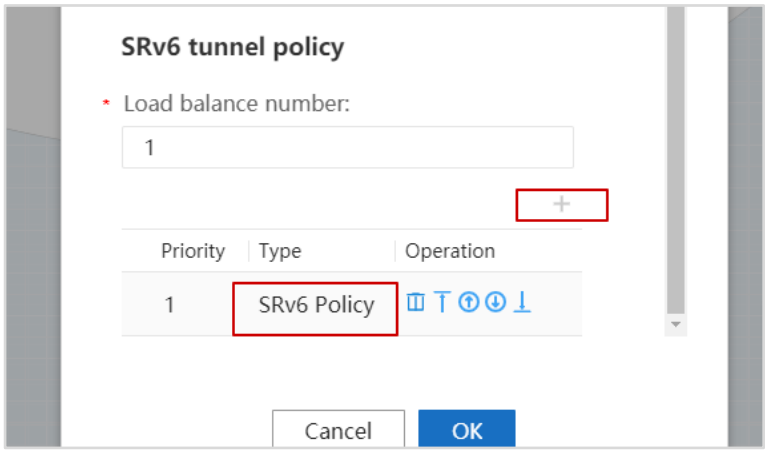

After the parameters are set, click **OK**.

Note that **Keep VRF up** must be selected in the **Service Node** area.

In the **Service Node** area, configure a route-policy. Specifically, click **+** next to **Routing Policy**. In the **Routing Policy** dialog box, set **Direction** to **Out** and click the list button next to the **Routing Policy Template** value. In the dialog box that is displayed, select the previously created route-policy template **EVPN_SRv6**. Then click **OK**.

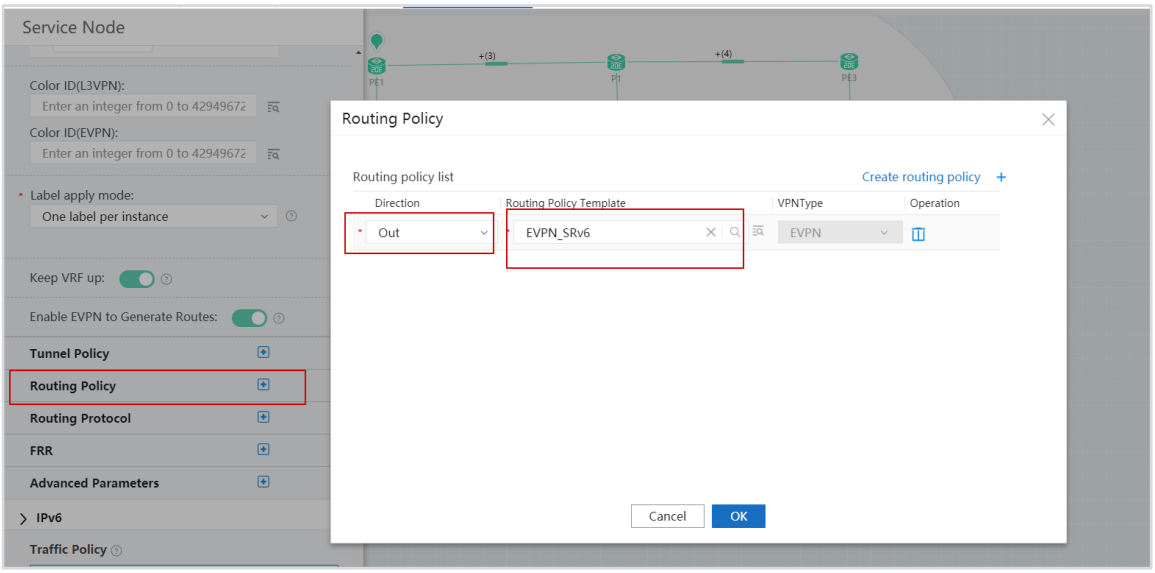

In the **Service Node** area, configure a routing protocol. Specifically, click **+** next to **Routing Protocol**. In the **Routing Policy** dialog box, set **SRv6 VPN Locator** to **SRv6** and turn on **Enable SRv6 BE: Yes** and **Enable SRv6 Policy: Yes**.

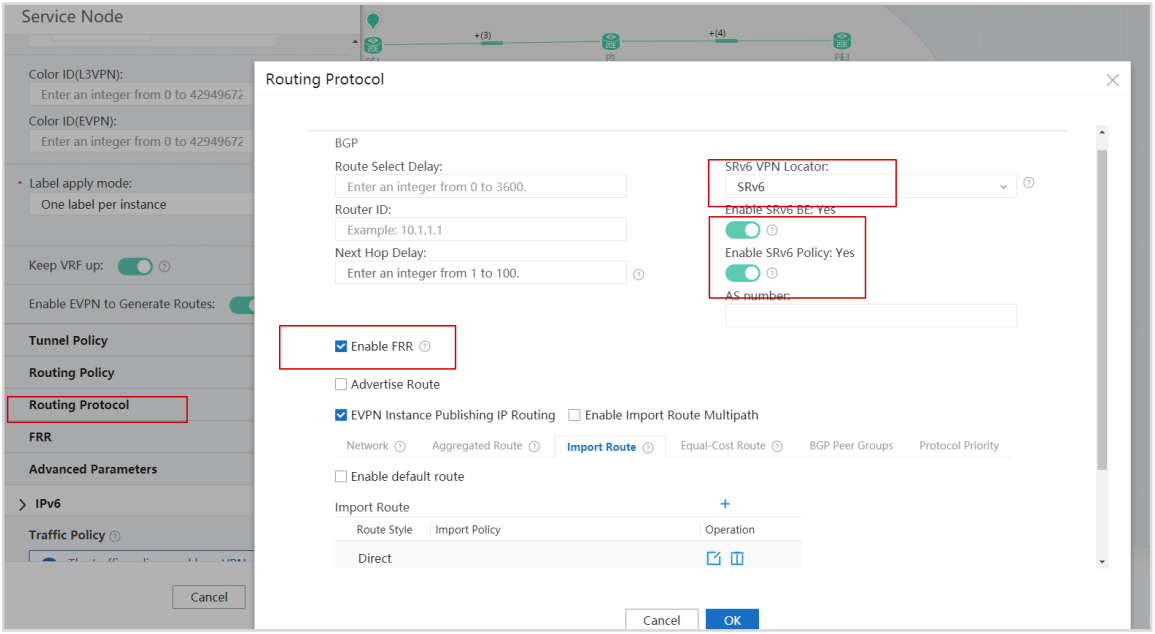

After the configuration is complete, click **OK**.

After the basic VRF parameters, tunnel policy, route-policy, and routing protocol are configured for a service node, click **OK** in the **Service Node** area. The service node configuration is complete.

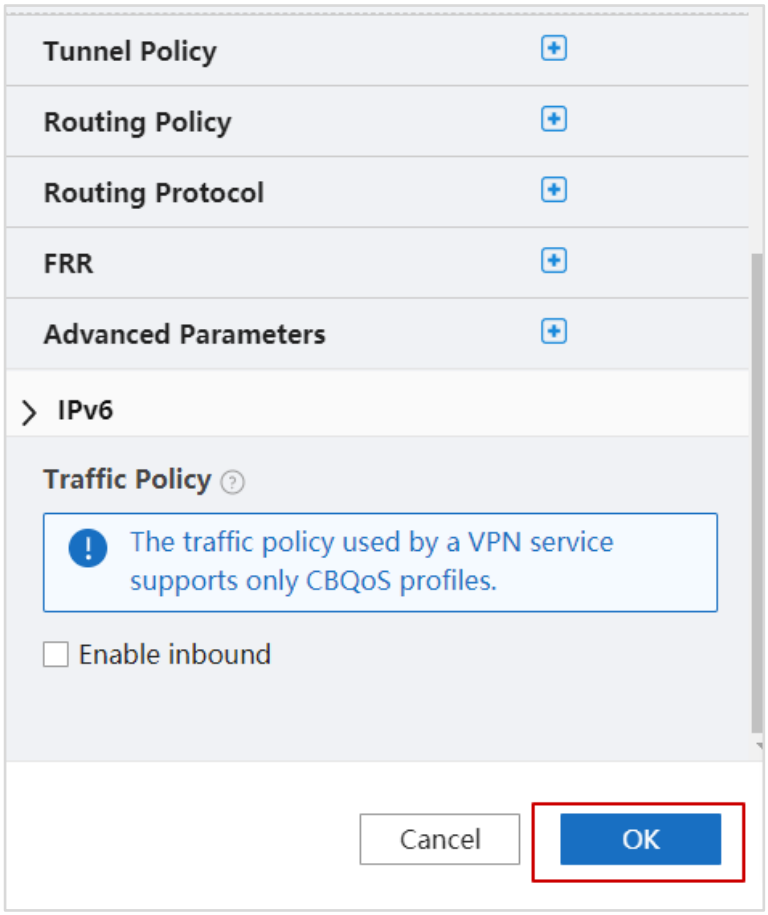

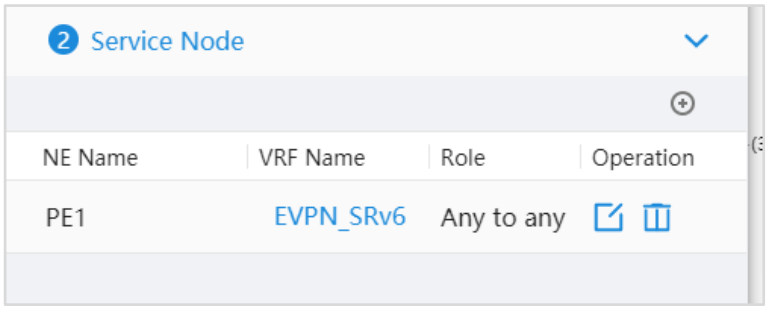

The service node configuration on PE1 is complete. The service node configuration on PE4 is similar to that on PE1.

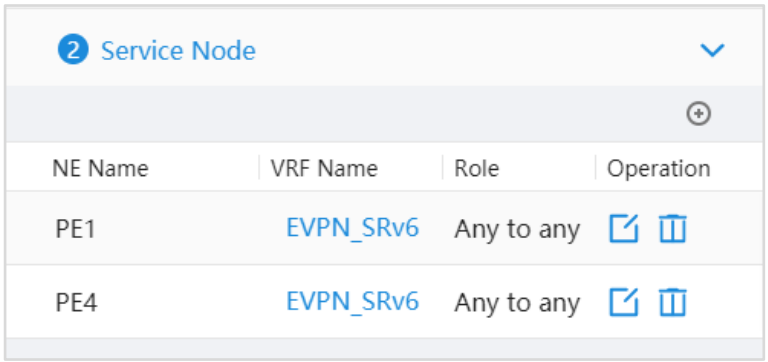

Step 3 Configure service access points.

Configure a service access point to determine the mode in which the user-side CE accesses the PE. In this experiment, Loopback3 is used to simulate user access.

Click **+** in the **Service Access Point** area.

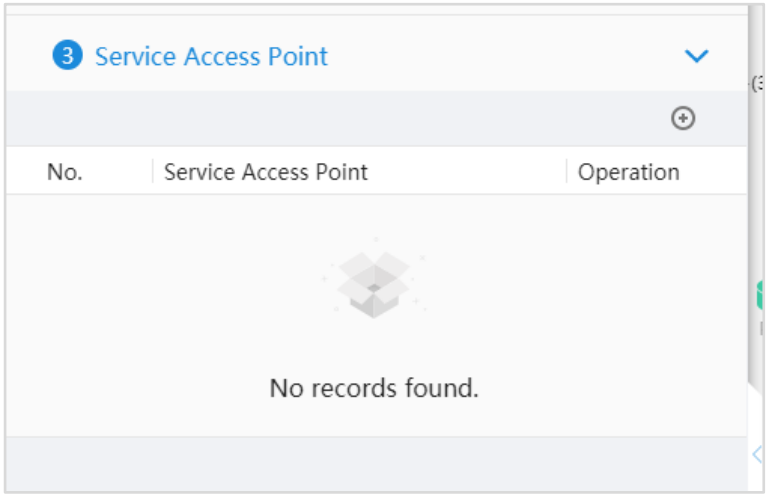

In the **Service Access Point** dialog box, set **Name** to **PE1_Customer** for PE1, retain the default single-homing access mode, and click **+**.

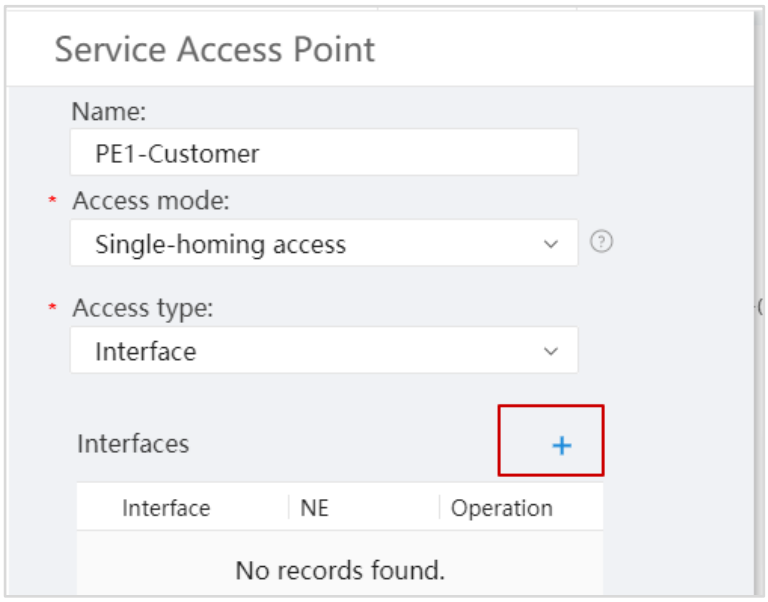

In the dialog box that is displayed, set **NE** to **PE1|EVPN_SRv6** and **Interface** to **Loopback3**, and select **Enable IPv4**.

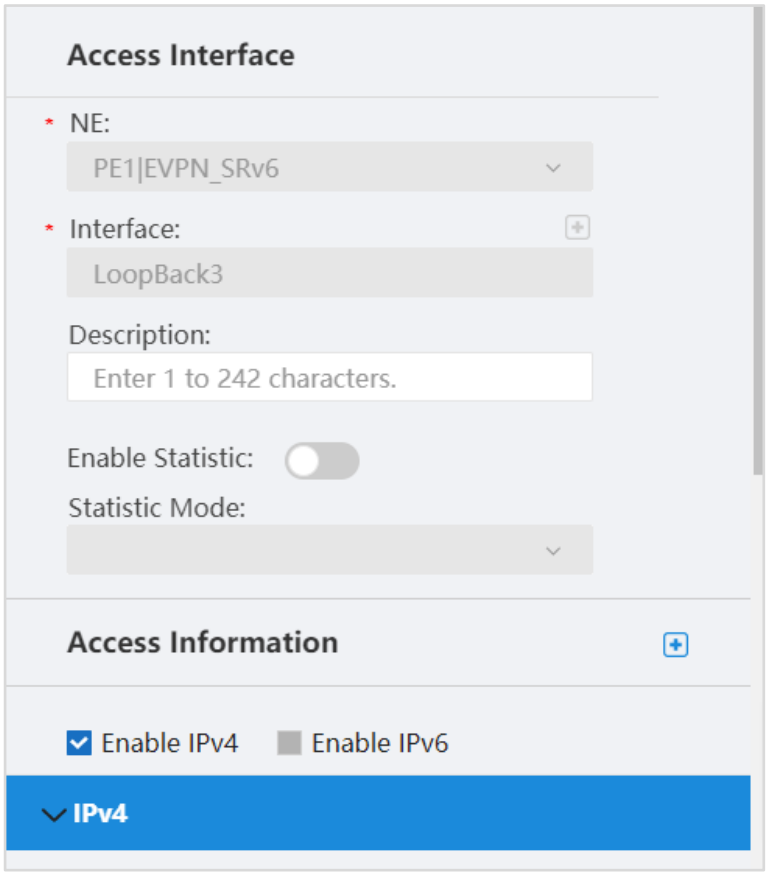
Access Interface
* NE:
PE1|EVPN_SRv6
* Interface:
LoopBack3
Description:
Enter 1 to 242 characters.
Enable Statistic:
Statistic Mode:
Access Information
Enable IPv4
Enable IPv6
IPv4

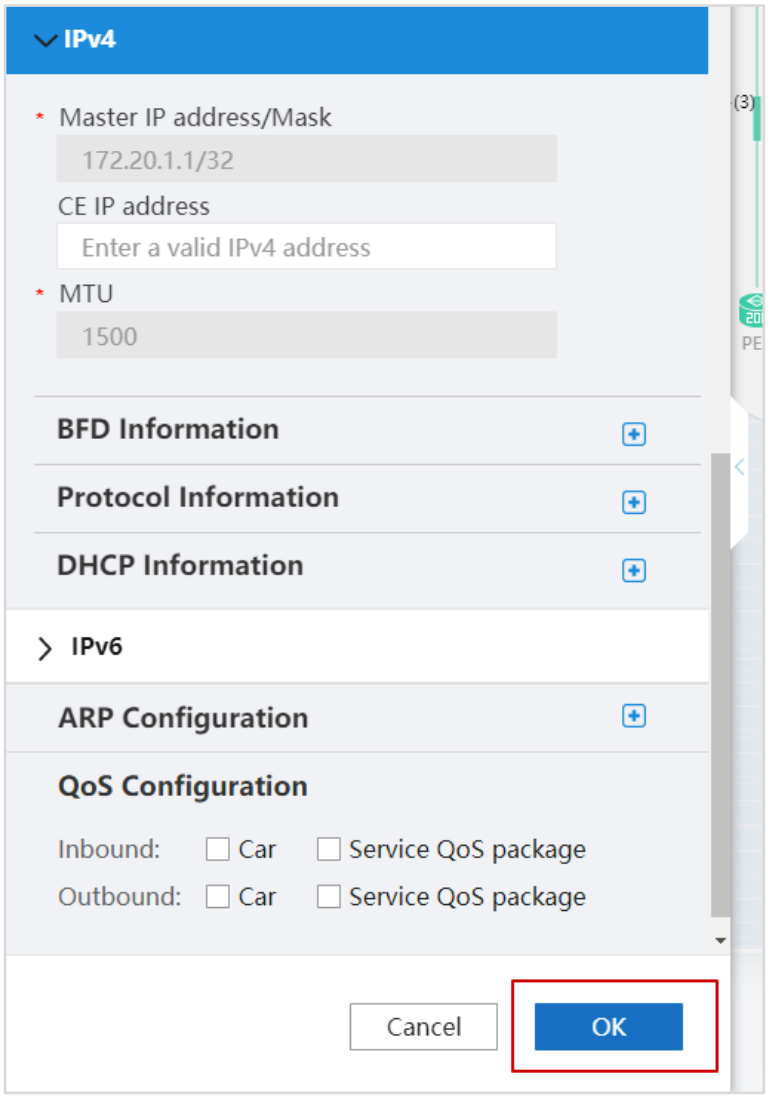

Finally, click **OK**. The service access point PE1_Customer is configured. The service access point configuration on PE4 is similar to that on PE1.

Because loopback interfaces are used to simulate user access, you do not need to configure the interconnection mode between the PE and CE. You can configure the interconnection mode (such as static or BGP) during protocol information configuration in actual scenarios.

The service access point configurations (CE-related configurations) on PE1 and PE4 are complete.

Step 4 Configure tunnels.

Associate the VRF instance with an SRv6 Policy. In the **Tunnel Configuration** area, select **Enable IPv4**.

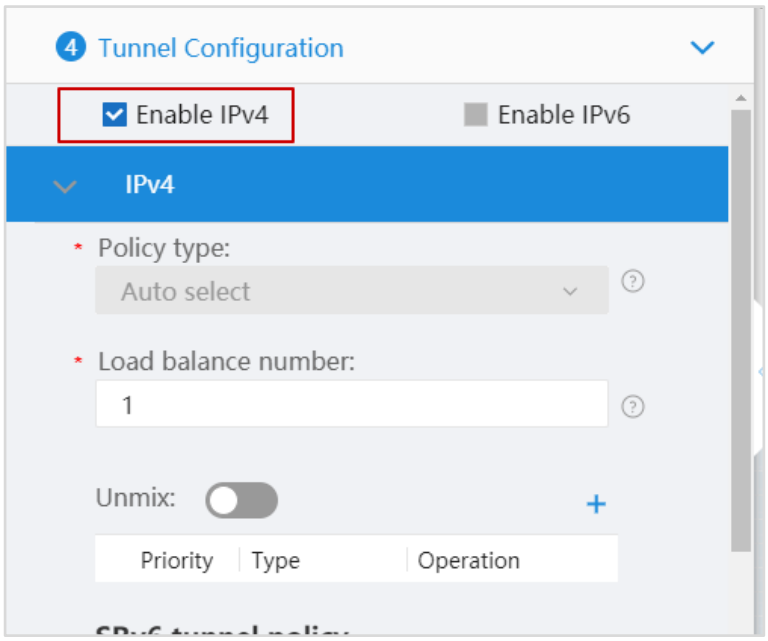

In the **SRv6 Tunnel Policy** area, click **+** and set **Tunnel type** to **SR-MPLS TE Policy**.

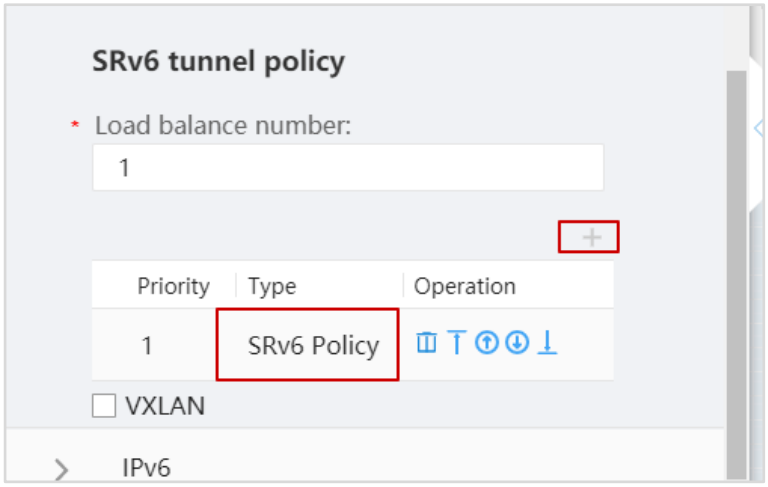

Finally, click **Apply** and wait until the EVPN L3VPN service configurations are delivered to devices.

After the service configurations are successfully delivered, click **View Service**.

The alarm status and running status are normal.

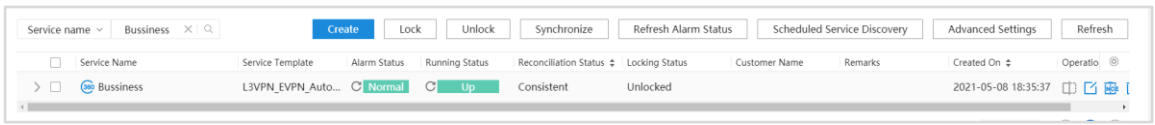

Step 5 Verify delivered configurations.

Check VRF and Loopback3 configurations.

```
[PE1]display current-configuration configuration vpn-instance EVPN_SRV6
#
ip vpn-instance EVPN_SRV6
 ipv4-family
  route-distinguisher 10:10
  tnl-policy NCE-VRF-EVPN_SRV6
  apply-label per-instance
  transit-vpn
  export route-policy EVPN_SRv6 evpn
  vpn-target 10:1000 export-extcommunity evpn
  vpn-target 10:1000 import-extcommunity evpn
  tnl-policy NCE-VRF-E-EVPN_SRV6 evpn
  evpn mpls routing-enable
  default-color 7 evpn
#
[PE1]display current-configuration interface LoopBack 3
#
interface LoopBack3
 ip binding vpn-instance EVPN_SRV6
 ip address 172.20.1.1 255.255.255.255
#
Return
```

According to the configuration information, the tunnel policy applied to the VRF is NCE-VRF-E-EVPN_SRV6.

Check tunnel information.

```
[PE1]display tunnel-info all
Tunnel ID                Type                 Destination          Status
-------------------------------------------------------------------------------
0x000000000300002002 sr-te                1.0.0.4              UP
0x000000002900000003 srbe-lsp             1.0.0.2              UP
0x000000002900000004 srbe-lsp             1.0.0.4              UP
0x000000002900000005 srbe-lsp             1.0.0.6              UP
0x000000002900000008 srbe-lsp             1.0.0.5              UP
0x000000002900000009 srbe-lsp             1.0.0.3              UP
0x00000000320004e002 srtepolicy           1.0.0.4              UP
0x000000003400028001 srv6tepolicy         FC01::4              UP
```

In this case, an SRv6 Policy is delivered.

Check tunnel information.

```
[PE1]display   tunnel-info 0x000000003400028001
Tunnel ID:           0x000000003400028001
Type:                srv6tepolicy
Name:                SRv6-TE Policy
Destination:         FC01::4
Instance ID:          0
Cost:                 0
```

```
Status:              UP
Color:               11
```

The color with ID 11 corresponds to the color in the color template (SRv6-PE1_PE4) applied to the SR Policy configured on the controller.

Check the Color Value of an EVPN Route.

```
[PE1-bgp]display bgp evpn all routing-table prefix-route 0:172.20.4.1:32 | include Color

 BGP local router ID : 1.0.0.1
 Local AS number : 65001
 Ext-Community: RT <10 : 1000>, Color <0 : 11>
```

The command output shows that the route carries the extended community attribute Color 11.

Check VRF EVPN_SRV6's routing table information.

```
[PE1]display ip routing-table vpn-instance  EVPN_SRv6
Route Flags:  R -relay,D -downloadtofib,T -tovpn-instance, B -blackholeroute
------------------------------------------------------------------------------
RoutingTable: EVPN_SRV6
         Destinations : 4        Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
     172.20.1.1/32  Direct  0    0           D   127.0.0.1       LoopBack3
     172.20.4.1/32  IBGP    255  0           RD  FC01::4         SRv6-TE Policy
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

The outbound interface of the route to 172.20.4.1 is an SRv6 Policy, not a specific tunnel interface.

Check SRv6 Policy information.

```
[PE1]display srv6-te policy
PolicyName : -
Color                    : 11                        Endpoint              : FC01::4
TunnelId                 : 163841                    BindingSID            : FC00:1::1:B
TunnelType               : SRv6-TE Policy            DelayTimerRemain      : -
Policy State             : Up                        State ChangeTime      : 2021-05-08
10:45:55
AdminState               : UP                        TrafficStatistics     : Disable
Backup Hot-Standby       : Disable                   BFD                   : Disable

Candidate-pathCount      : 1

Candidate-path Preference: 65535
Path State               : Active                    Path Type             : Primary
```

```
Protocol-Origin          : BGP(20)                  Originator          : 65001,
172.21.17.102
Discriminator            : 5                        Binding SID         : FC00:1::1:B
GroupId                  : 163841                   Policy Name         : -
Template ID              : 4294967278                Path Verification  : Disable
DelayTimerRemain         : -                        Segment-ListCount   : 1
 Segment-List            : -
  Segment-ListID         : 16385                    XcIndex             : 16385
  List State             : Up                       DelayTimerRemain    : -
  Verification State     : -                        SuppressTimeRemain  : -
  PMTU                   : 9600                     Active PMTU         : 9600
  Weight                 : 1                        BFD State           : -
  SID :
        FC00:1::1:22
        FC00:5::1:22
        FC00:3::1:2
```

The tunnel egress is FC01::4 (PE4) and the SIDs of nodes along the tunnel are FC00:1::1:22, FC00:5::1:22, and FC00:3::1:2.

Check BGP SRv6 Policy route information.

```
[PE1]display bgp sr-policy  ipv6 routing-table


BGP Local router ID is 1.0.0.1
Status codes: *-valid,> -best, d -damped,x-best external,a -add path,
              h -history,  i -internal, s-suppressed, S -Stale
              Origin : i -IGP, e -EGP, ?-incomplete
RPKI validation codes: V -valid, I -invalid, N -not-found


Total Number of Routes: 2
         Network            Nexthop         MED         LocPrf    PrefVal Path/Ogn

*>i    [5][11][FC01::4]     2000::102       4294967276  100       0       ?
*i                          2000::102       4294967276  100       0       ?
```

Check BGP SRv6 Policy route details.

```
[PE1]display bgp sr-policy ipv6 routing-table [5][11][FC01::4]

 BGP local router ID : 1.0.0.1
 Local AS number : 65001
 Paths:    2 available, 1 best, 1 select, 0 best-external, 0 add-path
 BGP routing table entry information of [5][11][FC01::4]:
 From: FC01::5 (1.0.0.5)
 Route Duration: 0d00h12m47s
 Relay IP Nexthop: ::
 Relay IP Out-Interface:GigabitEthernet0/0/0
 Original nexthop: 2000::102
 Qos information : 0x0
 Ext-Community: RT <1.0.0.1 : 0>, SoO <172.21.17.102 : 0>
 AS-path Nil, origin incomplete, MED 4294967276, localpref 100, pref-val 0, valid, internal, best,
select, pre 255
```

```
 Originator: 172.21.17.102
 Cluster list: 1.0.0.5
 Tunnel Encaps Attribute (23):
 Tunnel Type: SR Policy (15)
  Preference: 65535
  Binding SID: FC00:1::1:B, s-flag(0), i-flag(0)
  Segment List
    Weight: 1
    Path MTU: 9600
    Segment: type:2, SID: FC00:1::1:22
    Segment: type:2, SID: FC00:5::1:22
    Segment: type:2, SID: FC00:3::1:2
  Template ID: 4294967278
 Not advertised to any peer yet
....
```

In route details, we can see the route color, segment list, and other information.

Test EVPN L3VPN connectivity on PE1.

```
[PE1]ping -vpn-instance EVPN_SRV6 172.20.4.1
  PING 172.20.4.1: 56  data bytes, press CTRL_C to break
    Reply from 172.20.4.1: bytes=56 Sequence=1 ttl=255 time=1 ms
    Reply from 172.20.4.1: bytes=56 Sequence=2 ttl=255 time=1 ms
    Reply from 172.20.4.1: bytes=56 Sequence=3 ttl=255 time=1 ms
    Reply from 172.20.4.1: bytes=56 Sequence=4 ttl=255 time=1 ms
    Reply from 172.20.4.1: bytes=56 Sequence=5 ttl=255 time=1 ms

  --- 172.20.4.1 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 1/1/1 ms
The connectivity is normal.
```

3.2.3 Quiz

What are the three components of an SRv6 SID?

Reference Answers

SR-MPLS Experiment:

1. Unlike in MPLS forwarding, the outer label remains unchanged in a BE scenario.
2. Configure an explicit path and specify first the node SID and then the target adjacency SID of the device for the explicit path.
3. An SR-MPLS Policy is identified by <headend, color, endpoint>.

SRv6 Experiment:

1. END.DT6 SID.
2. An End SID is used to identify a local node, and an End.DT4 SID is used to identify an IPv4 VPN instance on a node.

iMaster NCE-IP Experiment:

1. NETCONF is used to deliver SR-MPLS TE configurations.
2. An SRv6 SID consists of three fields: Locator, Function, and Arguments. SRv6 SIDs are expressed in the *Locator*:*Function*:*Arguments* format.