Recommendations

- Huawei Learning Website
  - http://learning.huawei.com/en
- Huawei e-Learning
  - https://ilearningx.huawei.com/portal/#/portal/ebg/51
- Huawei Certification
  - http://support.huawei.com/learning/NavigationAction!createNavi?navId=_31&lang=en
- Find Training
  - http://support.huawei.com/learning/NavigationAction!createNavi?navId=_trainingsearch&lang=en

More Information

- Huawei learning APP

Huawei Certification

HCIP-Routing & Switching-IERS

Implementing Enterprise Routing and Switching Network

Lab Guide

Huawei Technologies Co.,Ltd

Huawei Certification

Implementing Enterprise Routing and Switching Network

Lab Guide

Edition 2.5

Huawei Certification System

Relying on its strong technical and professional training and certification system and in accordance with customers of different ICT technology levels, Huawei certification is committed to providing customers with authentic, professional certification, and addresses the need for the development of quality engineers that are capable of supporting Enterprise networks in the face of an ever changing ICT industry. The Huawei certification portfolio for routing and switching (R&S) is comprised of three levels to support and validate the growth and value of customer skills and knowledge in routing and switching technologies.

The Huawei Certified Network Associate (HCIA) certification level validates the skills and knowledge of IP network engineers to implement and support small to medium-sized enterprise networks. The HCIA certification provides a rich foundation of skills and knowledge for the establishment of such enterprise networks, along with the capability to implement services and features within existing enterprise networks, to effectively support true industry operations.

HCIA certification covers fundamentals skills for TCP/IP, routing, switching and related IP network technologies, together with Huawei data communications products, and skills for versatile routing platform (VRP) operation and management.

The Huawei Certified Network Professional (HCIP-R&S) certification is aimed at enterprise network engineers involved in design and maintenance, as well as professionals who wish to develop an in depth knowledge of routing, switching, network efficiency and optimization technologies. HCIP-R&S consists of three units including Implementing Enterprise Routing and Switching Network (IERS), Improving Enterprise Network Performance (IENP), and Implementing Enterprise Network Engineering Project (IEEP), which includes advanced IPv4 routing and switching technology principles, network security, high availability and QoS, as well as application of the covered technologies in Huawei products.

The Huawei Certified Internet Expert (HCIE-R&S) certification is designed to imbue engineers with a variety of IP network technologies and proficiency in maintenance, for the diagnosis and troubleshooting of Huawei products, to equip engineers with in-depth competency in the planning, design and optimization of large-scale IP networks.

About This Document

Overview

This document is HCIP-Implementing Enterprise Routing and Switching (HCIP-IERS) certification training material. It is intended for those who are preparing for the HCIP-IERS exam and those who want to master common routing protocol principles and Huawei Versatile Routing Platform (VRP) implementation.

Chapters 1, 2, and 3 describe working principles, configurations, and implementation of OSPF and IS-IS (two IGPs) as well as BGP (an EGP). These chapters help readers master IPv4 routing protocol knowledge.

Chapter 4 briefly introduces multicast address, IGMP, and PIM-SM, helping readers learn fundamental multicast knowledge, common multicast protocol principles, and multicast applications.

Chapter 5 illustrates how to flexibly use a variety of tools to control route selection. It helps readers flexibly use routing protocols.

Chapter 6 describes the VLAN principles and implementation, including VLAN Layer 2 interconnection and Layer 3 routing, helping readers learn VLAN working principles and configurations on the VRP.

Chapter 7 describes the working principles and implementation of STP protocols, including STP, RSTP, and MSTP.

This document helps readers understand how to implement routing and switching technologies on Huawei products.

Background Knowledge Required

To fully understand this document, readers should:

- Have participated in HCIA training.
- Have passed HCIA exams.
- Familiarize with the TCP/IP protocol stack and IP addressing.

Icons

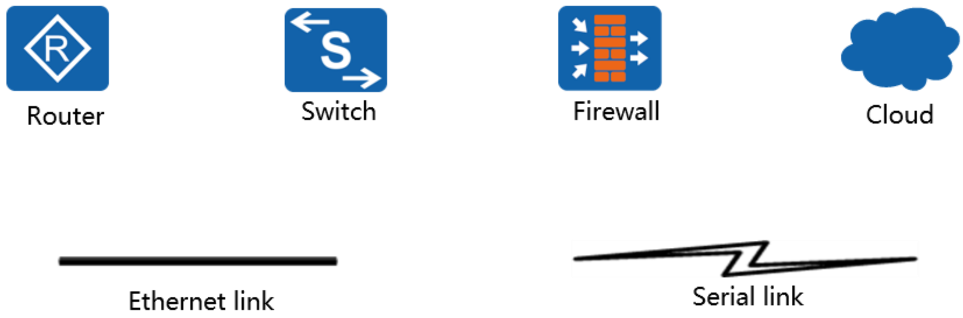

Contents

Lab Environment

Install eNSP

1. Login website of eNSP：

   https://support.huawei.com/enterprise/en/tool/ensp-TL1000000015/23917110

2. Download the latest version of eNSP

Software Type　　Others

Software Name	Size	Publication Date	Downloads	Download
CE.zip	564.58MB	2019/03/08	986	
CX.zip	405.65MB	2019/03/08	802	
NE40E.zip	405.69MB	2019/03/08	886	
NE5000E.zip	405.19MB	2019/03/08	664	
NE9000.zip	405.48MB	2019/03/08	596	
USG6000V.zip	344.93MB	2019/03/08	751	
eNSP V100R003C00SPC100 Setup.zip	542.52MB	2019/03/08	27906	

Download

3. Please refer to the software installation guide below to install eNSP in local PC.

File Name
eNSP V100R003C00SPC100 Software Installation Guide
eNSP V100R003C00SPC100 Release Notes
eNSP V100R003C00SPC100 Virus Scan Report

   Then engineer can practice lab with AR, Router, S57, S37, USG5500, AC, AP .

   If the engineer want to practice lab with USG6000V, CE, NE40, NE5000E, NE9000, CX, please follow Step4

4. Enable USG6000V, CE, NE40, NE5000E, NE9000, CX devices in eNSP:
   1) For example, if you want to enable USG6000V in eNSP, you should download the corresponding mirror file.

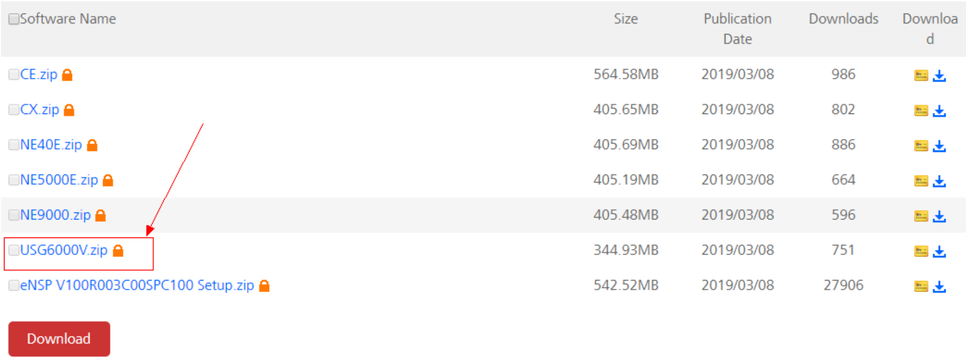

Software Name	Size	Publication Date	Downloads	Download
CE.zip	564.58MB	2019/03/08	986	
CX.zip	405.65MB	2019/03/08	802	
NE40E.zip	405.69MB	2019/03/08	886	
NE5000E.zip	405.19MB	2019/03/08	664	
NE9000.zip	405.48MB	2019/03/08	596	
USG6000V.zip	344.93MB	2019/03/08	751	
eNSP V100R003C00SPC100 Setup.zip	542.52MB	2019/03/08	27906	

Download

2) Select USG6000V into new project of eNSP, then right click "start" of USG6000V :

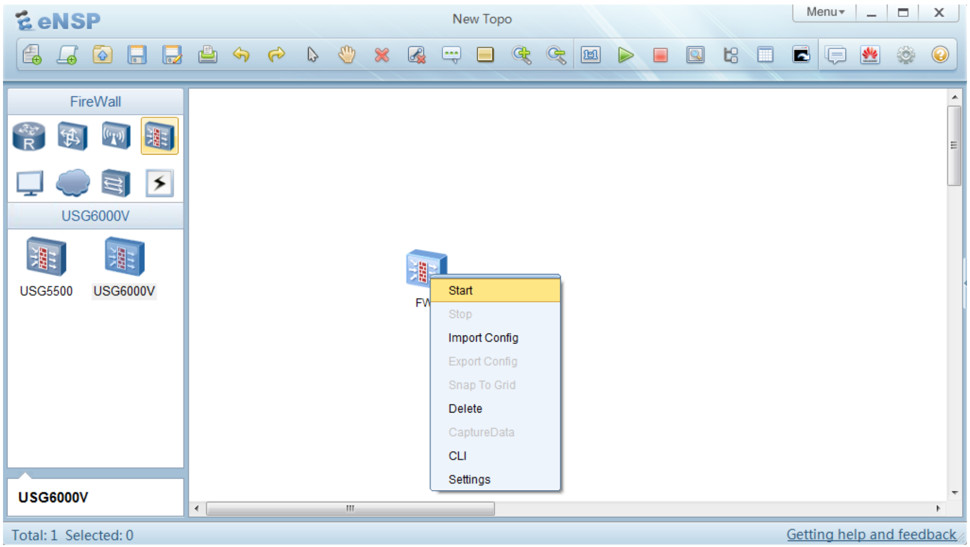

3) The dialog box of "import package" will show up:

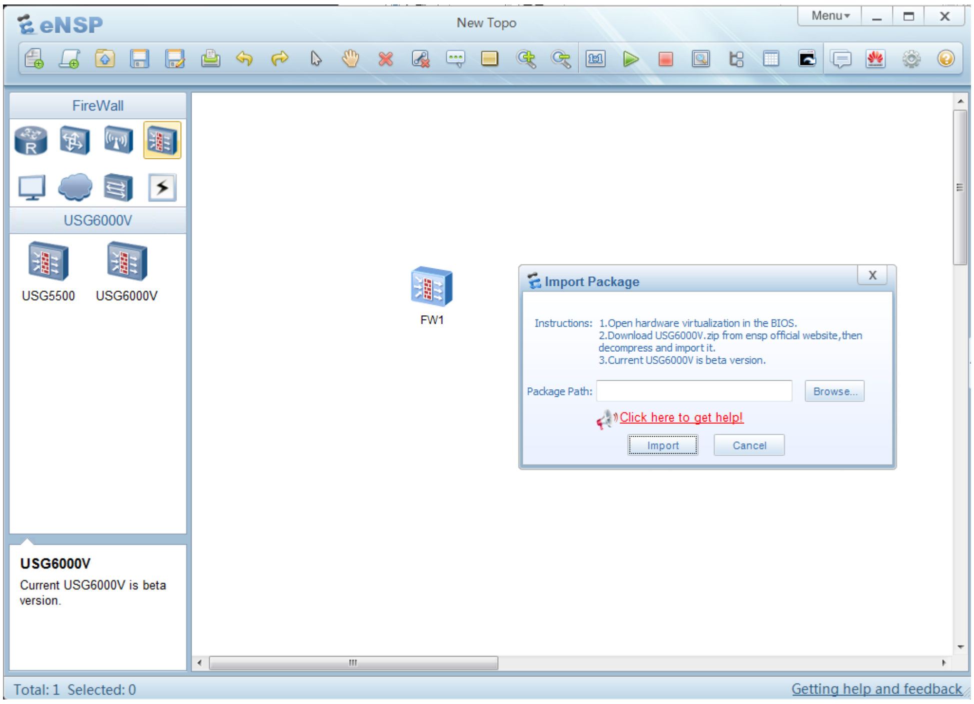

4) Click "Browse" - and import the downloaded mirror files, then engineer can practice lab with USG6000V.

5) If the engineer want to practice CE, NE40, NE5000E, NE9000, CX, please repeat step 4-1) --- 4).

Chapter 1 OSPF Features and Configurations

Lab 1-1 Single-Area OSPF

Learning Objectives

The objectives of this lab are to learn and understand how to perform the following operations:

- How to configure single-area OSPF
- How to configure OSPF authentication
- How to establish neighbor relationships on multi-access networks
- How to use OSPF to advertise the subnet mask of the network to which the loopback interface connects
- How to change cost values for OSPF interfaces
- How to configure an interface as a silent interface
- How to view OSPF status using the **display** command
- How to view OSPF neighbor relationships and troubleshoot faults using the **debug** command

Topology

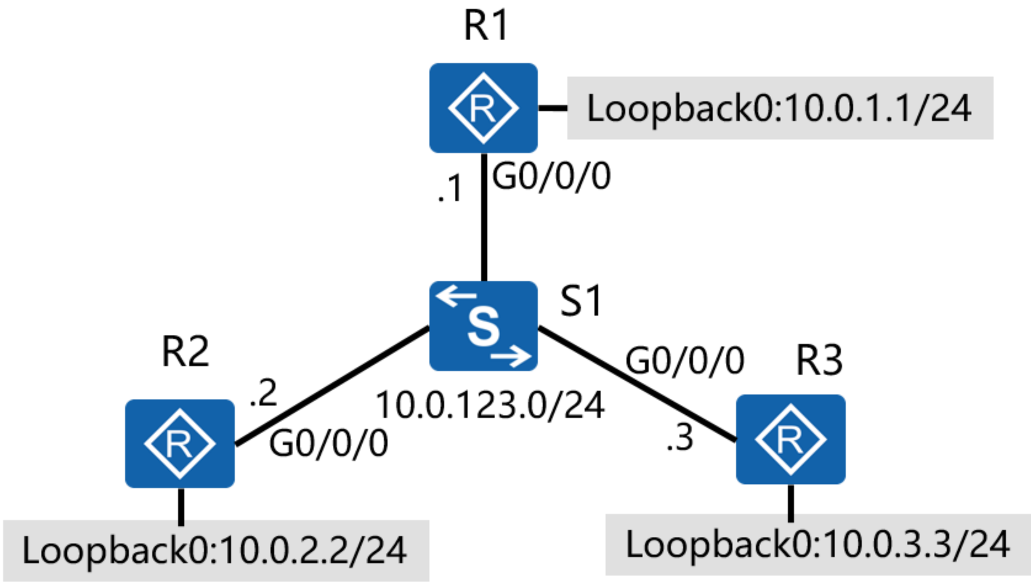

Figure 1-1 Single-area OSPF

Scenario

Assume that you are a network administrator of a company that has three ARG3 routers. These routers are interconnected over the Ethernet. A broadcast multi-access network, such as Ethernet, has security threats. Therefore, OSPF area authentication is required to prevent malicious route attacks. A network connectivity failure occurs during network deployment. You can run the **display** and **debug** commands for fault location.

Tasks

Step 1 **Perform basic configurations and configure IP addresses.**

Configure IP addresses and masks for R1, R2, and R3. Set a 24-bit mask for loopback interfaces to simulate an independent network segment.

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
```

```
[R1]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]ip address 10.0.123.1 24
[R1-GigabitEthernet0/0/0]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 24
[R1-LoopBack0]quit

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.123.2 24
[R2-GigabitEthernet0/0/0]quit
[R2]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24
[R2-LoopBack0]quit

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]ip address 10.0.123.3 24
[R3-GigabitEthernet0/0/0]quit
[R3]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 24
[R3-LoopBack0]quit
```

Verify the connectivity between routers.

```
[R1]ping -c 1 10.0.123.2
  PING 10.0.123.2: 56   data bytes, press CTRL_C to break
```

```
    Reply from 10.0.123.2: bytes=56 Sequence=1 ttl=255 time=2 ms

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

[R1]ping -c 1 10.0.123.3
  PING 10.0.123.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.123.3: bytes=56 Sequence=1 ttl=255 time=2 ms

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

[R2]ping -c 1 10.0.123.3
  PING 10.0.123.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.123.3: bytes=56 Sequence=1 ttl=255 time=2 ms

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

Step 2 **Configure single-area OSPF.**

Configure single-area OSPF and deploy all routers in area 0. Configure OSPF process 1. In addition, configure area authentication and set the password to **huawei**. In an OSPF area, Huawei devices support plain text or MD5 authentication. Plain text authentication is used for this step.

Set the wildcard subnet mask to 0.0.0.0 when you use the **network** command. To ensure the stability of Router IDs, they are usually specified manually as.

```
[R1]ospf 1 router-id 10.0.1.1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.0.123.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]authentication-mode simple plain huawei
[R1-ospf-1-area-0.0.0.0]quit
[R1-ospf-1]quit

[R2]ospf 1 router-id 10.0.2.2
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.123.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]authentication-mode simple plain huawei
[R2-ospf-1-area-0.0.0.0]quit
[R2-ospf-1]quit

[R3]ospf 1 router-id 10.0.3.3
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]network 10.0.123.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]authentication-mode simple plain huawei
```

[R3-ospf-1-area-0.0.0.0]quit

[R3-ospf-1]quit

View the routing tables and test the connectivity of the entire network.

View the routing table of R1.

```
 [R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 12        Routes : 12

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        10.0.1.0/24  Direct  0    0           D   10.0.1.1        LoopBack0
        10.0.1.1/32  Direct  0    0           D   127.0.0.1       LoopBack0
      10.0.1.255/32  Direct  0    0           D   127.0.0.1       LoopBack0
        10.0.2.2/32  OSPF    10   1           D   10.0.123.2      GigabitEthernet0/0/0
        10.0.3.3/32  OSPF    10   1           D   10.0.123.3      GigabitEthernet0/0/0
      10.0.123.0/24  Direct  0    0           D   10.0.123.1      GigabitEthernet0/0/0
      10.0.123.1/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
    10.0.123.255/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
       127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
       127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct  0    0           D   127.0.0.1       InLoopBack0
```

The command output shows that R1 learns two routes from OSPF: 10.0.2.2/32 and 10.0.3.3/32. The next hops of the two routes are 10.0.123.2 and 10.0.123.3 respectively.

Verify the connectivity from R1 to loopback interface addresses of R2 and R3.

```
[R1]ping -c 1 10.0.2.2
  PING 10.0.2.2: 56   data bytes, press CTRL_C to break
    Reply from 10.0.2.2: bytes=56 Sequence=1 ttl=255 time=3 ms

  --- 10.0.2.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 3/3/3 ms

[R1]ping -c 1 10.0.3.3
  PING 10.0.3.3: 56   data bytes, press CTRL_C to break
    Reply from 10.0.3.3: bytes=56 Sequence=1 ttl=255 time=2 ms

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

Run the **display ospf brief** command to view basic OSPF information on R1.

```
[R1]display ospf brief

         OSPF Process 1 with Router ID 10.0.1.1
```

OSPF Protocol Information

RouterID: 10.0.1.1 Border Router:

Multi-VPN-Instance is not enabled

Global DS-TE Mode: Non-Standard IETF Mode

Graceful-restart capability: disabled

Helper support capability : not configured

Applications Supported: MPLS Traffic-Engineering

Spf-schedule-interval: max 10000ms, start 500ms, hold 1000ms

Default ASE parameters: Metric: 1 Tag: 1 Type: 2

Route Preference: 10

ASE Route Preference: 150

SPF Computation Count: 9

RFC 1583 Compatible

Retransmission limitation is disabled

Area Count: 1 Nssa Area Count: 0

ExChange/Loading Neighbors: 0

Process total up interface count: 2

Process valid up interface count: 1

Area: 0.0.0.0 (MPLS TE not enabled)

Authtype: Simple Area flag: Normal

SPF scheduled Count: 9

ExChange/Loading Neighbors: 0

Router ID conflict state: Normal

Area interface up count: 2

Interface: 10.0.1.1 (LoopBack0)

Cost: 0 State: P-2-P Type: P2P MTU: 1500

```
Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1
Interface: 10.0.123.1 (GigabitEthernet0/0/0)
Cost: 1       State: DR        Type: Broadcast    MTU: 1500
Priority: 1
Designated Router: 10.0.123.1
Backup Designated Router: 10.0.123.2
Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1
```

The preceding command output **Authtype: Simple** shows that plaintext authentication is enabled in Area 0. OSPF runs on two interfaces: GigabitEthernet0/0/0 and Loopback0. The network type of GigabitEthernet0/0/0 is broadcast, cost is 1, and priority is 1. DR is R1, and BDR is 10.0.123.2. The network type of another OSPF-enabled Loopback 0 is P2P.

Run the **display ospf peer brief** command on R1 to check information about OSPF neighbor relationships between the routers.

```
[R1]display ospf peer brief

     OSPF Process 1 with Router ID 10.0.1.1
          Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                       Neighbor id      State
 0.0.0.0          GigabitEthernet0/0/0            10.0.2.2         Full
 0.0.0.0          GigabitEthernet0/0/0            10.0.3.3         Full
 ----------------------------------------------------------------------------
 Total Peer(s):       2
```

The preceding command output shows that R1 has two neighbors in Area 0.0.0.0, their router IDs are 10.0.2.2 and 10.0.3.3 respectively, and their OSPF neighbor relationships are in Full state.

Run the **display ospf lsdb** command on R1 to check OSPF LSDB information.

```
[R1]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.1.1
                 Link State Database

                         Area: 0.0.0.0
 Type      LinkState ID    AdvRouter          Age  Len     Sequence    Metric
 Router    10.0.3.3    10.0.3.3           1569   48     80000005         0
 Router    10.0.2.2    10.0.2.2           1568   48     80000006         0
 Router    10.0.1.1    10.0.1.1           1567   48     80000008         0
 Network10.0.123.110.0.1.1            1567   36     80000004         0
```

The preceding command output shows that the LSDB contains four LSAs, the first three of which are Type 1 LSAs generated by R1, R2, and R3 respectively. You can check the **AdvRouter** field to determine which router generates an LSA. The fourth LSA is a Type 2 LSA, which is generated by a DR of a network segment. Because R1 is the DR of the network segment 10.0.123.0/24, you can see that the **AdvRouter** field of this LSA is 10.0.1.1.

```
[R1]display ospf lsdb router self-originate

         OSPF Process 1 with Router ID 10.0.1.1
                         Area: 0.0.0.0
                 Link State Database

  Type         : Router
  Ls id        : 10.0.1.1
  Adv rtr : 10.0.1.1
  Ls age       : 430
```

```
Len          : 48
Options      :  E
seq#         : 80000009
chksum            : 0x8188
Link count: 2
 * Link ID: 10.0.1.1
   Data     : 255.255.255.255
   Link Type: StubNet
   Metric : 0
   Priority : Medium
 * Link ID    : 10.0.123.1
   Data       : 10.0.123.1
   Link Type: TransNet
   Metric     : 1
```

The preceding command output shows detailed information about the Router LSA generated by R1. This LSA describes two networks. The first network is the network segment where the loopback interface resides. The **Link Type** field displays StubNet, and **Link ID** and **Data** fields indicate the IP address and mask of this stub network segment. The second network is the network segment that connects the three routers. The **Link Type** displays TransNet, the **Link ID** field displays 10.0.123.1, which is the interface address of the DR, and the **Data** field displays 10.0.123.1, which is the local interface address on the network segment.

```
[R1]display ospf lsdb network self-originate

         OSPF Process 1 with Router ID 10.0.1.1
                         Area: 0.0.0.0
                   Link State Database
```

```
Type          : Network
Ls id         : 10.0.123.1
Adv rtr : 10.0.1.1
Ls age        : 1662
Len           : 36
Options       :   E
seq#          : 80000005
chksum             : 0x3d58
Net mask      : 255.255.255.0
Priority : Low
   Attached Router     10.0.1.1
   Attached Router     10.0.2.2
   Attached Router     10.0.3.3
```

The preceding command output shows detailed information about the Network LSA generated by R1. This Type 2 LSA describes neighbor information on the network segment where the DR resides.

Step 3 **Observe the OSPF neighbor relationship establishment process on the routers.**

Check DR and BDR election on the network segment 10.0.123.0/24 and analyze whether the results of tests performed by different candidates are the same. According to the following command output, the interface IP address of the DR on this network segment is 10.0.123.1, and that of the BDR on this network segment is 10.0.123.2.

```
[R1]display ospf peer

         OSPF Process 1 with Router ID 10.0.1.1
```

```
                    Neighbors

 Area 0.0.0.0 interface 10.0.123.1(GigabitEthernet0/0/0)'s neighbors
 Router ID: 10.0.2.2          Address: 10.0.123.2
   State: Full   Mode:Nbr is   Master   Priority: 1
   DR: 10.0.123.1   BDR: 10.0.123.2   MTU: 0
   Dead timer due in 40   sec
   Retrans timer interval: 5
   Neighbor is up for 01:03:35
   Authentication Sequence: [ 0 ]

 Router ID: 10.0.3.3          Address: 10.0.123.3
   State: Full   Mode:Nbr is   Master   Priority: 1
   DR: 10.0.123.1   BDR: 10.0.123.2   MTU: 0
   Dead timer due in 33   sec
   Retrans timer interval: 5
   Neighbor is up for 01:02:27
   Authentication Sequence: [ 0 ]
```

The results of tests performed by different candidates may be different. This is because DR election of OSPF is not preempted. That is, when there is a DR or BDR on a network, the router that newly joins the network cannot preempt to be the DR or BDR. On this network, the router whose OSPF process starts first or that connects to this network first becomes the DR, and other routers are the BDR and DR others.

After the DR fails, the BDR becomes the new DR. You can reset an OSPF process to observe the DR role change. The following example resets the OSPF process of R1.

```
<R1>reset ospf process
Warning: The OSPF process will be reset. Continue? [Y/N]:y
[R1]display ospf peer
```

```
          OSPF Process 1 with Router ID 10.0.1.1
                  Neighbors

 Area 0.0.0.0 interface 10.0.123.1(GigabitEthernet0/0/0)'s neighbors
 Router ID: 10.0.2.2              Address: 10.0.123.2
   State: Full   Mode:Nbr is   Master   Priority: 1
   DR: 10.0.123.2   BDR: 10.0.123.3   MTU: 0
   Dead timer due in 34   sec
   Retrans timer interval: 0
   Neighbor is up for 00:00:19
   Authentication Sequence: [ 0 ]

 Router ID: 10.0.3.3              Address: 10.0.123.3
   State: Full   Mode:Nbr is   Master   Priority: 1
   DR: 10.0.123.2   BDR: 10.0.123.3   MTU: 0
   Dead timer due in 39   sec
   Retrans timer interval: 5
   Neighbor is up for 00:00:19
   Authentication Sequence: [ 0 ]
```

After the OSPF process of R1 is reset, the BDR 10.0.123.2 becomes the new DR, and the DR other 10.0.123.3 becomes the new BDR.

Shut down G0/0/0 of R1, R2, and R3 and run the **debugging ospf 1 event** command to check the OSPF neighbor relationship establishment process. Undoshutdown G0/0/0 of R1, R2, and R3 simultaneously, and observe neighbor status change and DR and BDR election on the broadcast multi-access network.

```
<R1>debugging ospf 1 event
```

```
<R1>terminal debugging
[R1]interface  GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]shut
[R1-GigabitEthernet0/0/0]undo shut

<R2>debugging ospf 1 event
<R2>terminal debugging
[R2]interface  GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]shut
[R2-GigabitEthernet0/0/0]undo shut

<R3>debugging ospf 1 event
<R3>terminal debugging
[R3]interface  GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]shutdown
[R3-GigabitEthernet0/0/0]undo shutdown
```

Perform the same operations on R2 and R3 and check debugging information on R3. The default interface priority of all routers is 1. Therefore, router IDs of the three routers are compared during DR election. Among the three routers, R3 has the largest router ID and becomes the DR of the network segment.

```
[R3-GigabitEthernet0/0/0]
Oct 12 2016 11:54:59.220.1+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802c Line: 1326 Level: 0x20
  OSPF 1: Intf 10.0.123.3 Rcv InterfaceUp State Down -> Waiting.
[R3-GigabitEthernet0/0/0]
Oct 12 2016 11:54:59.230.1+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802c Line: 1440 Level: 0x20
  OSPF 1 Send Hello Interface Up on 10.0.123.3
```

```
[R3-GigabitEthernet0/0/0]
Oct 12 2016 11:55:08.550.2+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 1200 Level: 0x20
  OSPF 1: Nbr 10.0.123.1 Rcv HelloReceived State Down -> Init.
[R3-GigabitEthernet0/0/0]
Oct 12 2016 11:55:09.530.2+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 1200 Level: 0x20
  OSPF 1: Nbr 10.0.123.2 Rcv HelloReceived State Down -> Init.
[R3-GigabitEthernet0/0/0]
Oct 12 2016 11:55:18.540.2+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 1796 Level: 0x20
  OSPF 1: Nbr 10.0.123.1 Rcv 2WayReceived State Init -> 2Way.
[R3-GigabitEthernet0/0/0]
Oct 12 2016 11:55:19.570.2+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 1796 Level: 0x20
  OSPF 1: Nbr 10.0.123.2 Rcv 2WayReceived State Init -> 2Way.
[R3-GigabitEthernet0/0/0]
Oct 12 2016 11:55:39.370.1+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 1796 Level: 0x20
  OSPF 1: Nbr 10.0.123.1 Rcv AdjOk? State 2Way -> ExStart.
[R3-GigabitEthernet0/0/0]
Oct 12 2016 11:55:39.370.2+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 1796 Level: 0x20
  OSPF 1: Nbr 10.0.123.2 Rcv AdjOk? State 2Way -> ExStart.
[R3-GigabitEthernet0/0/0]
Oct 12 2016 11:55:39.370.3+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802c Line: 2127 Level: 0x20
  OSPF 1 Send Hello Interface State Changed on 10.0.123.3
[R3-GigabitEthernet0/0/0]
```

```
Oct 12 2016 11:55:39.370.4+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802c Line: 2138 Level: 0x20
  OSPF 1: Intf 10.0.123.3 Rcv WaitTimer State Waiting -> DR.
[R3-GigabitEthernet0/0/0]
Oct 12 2016 11:55:39.390.1+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 1909 Level: 0x20
  OSPF 1: Nbr 10.0.123.1 Rcv NegotiationDone State ExStart -> Exchange.
[R3-GigabitEthernet0/0/0]
Oct 12 2016 11:55:39.390.2+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 1909 Level: 0x20
  OSPF 1: Nbr 10.0.123.2 Rcv NegotiationDone State ExStart -> Exchange.
[R3-GigabitEthernet0/0/0]
Oct 12 2016 11:55:39.400.1+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 2021 Level: 0x20
  OSPF 1: Nbr 10.0.123.1 Rcv ExchangeDone State Exchange -> Loading.
[R3-GigabitEthernet0/0/0]
Oct 12 2016 11:55:39.400.2+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 2423 Level: 0x20
  OSPF 1: Nbr 10.0.123.1 Rcv LoadingDone State Loading -> Full.
[R3-GigabitEthernet0/0/0]
Oct 12 2016 11:55:39.400.3+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 2021 Level: 0x20
  OSPF 1: Nbr 10.0.123.2 Rcv ExchangeDone State Exchange -> Loading.
[R3-GigabitEthernet0/0/0]
Oct 12 2016 11:55:39.400.4+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 2423 Level: 0x20
  OSPF 1: Nbr 10.0.123.2 Rcv LoadingDone State Loading -> Full.
<R1>undo debugging all
<R2>undo debugging all
```

```
<R3>undo debugging all
```

When G0/0/0 is just enabled, the interface state changes from Down to Waiting. Then routers start exchanging Hello packets. After 40 seconds, the status of G0/0/0 on R3 changes from Waiting to DR.

Step 4 **Set the network type for loopback interfaces.**

Check the IP routing table of R1 and focus on the two routes 10.0.2.2/32 and 10.0.3.3/32.

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 12        Routes : 12

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.0/24  Direct  0    0           D   10.0.1.1        LoopBack0
       10.0.1.1/32  Direct  0    0           D   127.0.0.1       LoopBack0
     10.0.1.255/32  Direct  0    0           D   127.0.0.1       LoopBack0
       10.0.2.2/32  OSPF    10   1           D   10.0.123.2      GigabitEthernet0/0/0
       10.0.3.3/32  OSPF    10   1           D   10.0.123.3      GigabitEthernet0/0/0
     10.0.123.0/24  Direct  0    0           D   10.0.123.1      GigabitEthernet0/0/0
     10.0.123.1/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
   10.0.123.255/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32Direct    0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct    0    0           D   127.0.0.1       InLoopBack0
```

During the configuration of IP addresses for loopback interfaces of R2 and R3, the 24-bit mask is used. Analyze why the IP routing table of R1 displays routes with the 32-bit mask.

Run the **display ospf interface LoopBack 0 verbose** command to check the OSPF running status of Loopback0.

```
[R1]display ospf interface LoopBack 0 verbose

         OSPF Process 1 with Router ID 10.0.1.1
                 Interfaces

 Interface: 10.0.1.1 (LoopBack0)
 Cost: 0        State: P-2-P       Type: P2P        MTU: 1500
 Timers: Hello 10 , Dead 40 , Poll   120 , Retransmit 5 , Transmit Delay 1
  IO Statistics
             Type          Input      Output
            Hello        0            0
   DB Description        0            0
   Link-State Req        0            0
Link-State Update        0            0
   Link-State Ack        0            0
 ALLSPF GROUP
 OpaqueId: 0    PrevState: Down
```

OSPF knows that the network segment where Loopback0 resides can have only one IP address. Therefore the subnet mask of the advertised route is 32 bits.

Change the network type of Loopback0 on R2 to broadcast. When OSPF advertises network information of this interface, it will use a 24-bit mask.

```
[R2]interface LoopBack 0

[R2-LoopBack0]ospf network-type broadcast
```

You can see that the subnet mask of the route to Loopback0 address advertised by R2 is 24 bits.

```
[R1]display ip routing-table

Route Flags: R - relay, D - download to fib

------------------------------------------------------------------------------

Routing Tables: Public

         Destinations : 12        Routes : 12

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        10.0.1.0/24  Direct  0    0           D   10.0.1.1        LoopBack0
        10.0.1.1/32  Direct  0    0           D   127.0.0.1       LoopBack0
      10.0.1.255/32  Direct  0    0           D   127.0.0.1       LoopBack0
        10.0.2.2/24  OSPF    10   1           D   10.0.123.2      GigabitEthernet0/0/0
        10.0.3.3/32  OSPF    10   1           D   10.0.123.3      GigabitEthernet0/0/0
      10.0.123.0/24  Direct  0    0           D   10.0.123.1      GigabitEthernet0/0/0
      10.0.123.1/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
    10.0.123.255/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
       127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
       127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct  0    0           D   127.0.0.1       InLoopBack0
```

Run the **display ospf interface LoopBack 0 verbose** command to check the OSPF running status of Loopback0. The command output shows that the network type of Loopback0 is broadcast.

```
[R2]display ospf interface LoopBack 0 verbose

          OSPF Process 1 with Router ID 10.0.2.2
                Interfaces

 Interface: 10.0.2.2 (LoopBack0)
 Cost: 0          State: DR          Type: Broadcast       MTU: 1500
 Priority: 1
 Designated Router: 10.0.2.2
 Backup Designated Router: 0.0.0.0
 Timers: Hello 10 , Dead 40 , Poll   120 , Retransmit 5 , Transmit Delay 1
  IO Statistics
              Type          Input       Output
             Hello              0            0
    DB Description        0            0
      Link-State Req        0            0
Link-State Update        0            0
      Link-State Ack        0            0
 ALLSPF GROUP
 ALLDR GROUP
 OpaqueId: 0     PrevState: Waiting
```

Step 5 **Change the OSPF interface cost.**

Check the cost of the route from R1 to Loopback0 of R3. You can see that the cost of the route to 10.0.3.3/32 is 1.

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 12        Routes : 12

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.0/24  Direct  0    0           D   10.0.1.1        LoopBack0
       10.0.1.1/32  Direct  0    0           D   127.0.0.1       LoopBack0
     10.0.1.255/32  Direct  0    0           D   127.0.0.1       LoopBack0
       10.0.2.2/24  OSPF    10   1           D   10.0.123.2      GigabitEthernet0/0/0
       10.0.3.3/32  OSPF    10   1           D   10.0.123.3      GigabitEthernet0/0/0
     10.0.123.0/24  Direct  0    0           D   10.0.123.1      GigabitEthernet0/0/0
     10.0.123.1/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
   10.0.123.255/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

Change the cost of G0/0/0 on R1 to 20 and that on R3 to 10.

```
[R1]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]ospf cost 20
[R1-GigabitEthernet0/0/0]quit

[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]ospf cost 10
```

[R3-GigabitEthernet0/0/0]quit

Check the cost of the route from R1 to Loopback0 of R3 again. You can see that the cost of the route to 10.0.3.3/32 is 20.

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 12        Routes : 12

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.0/24  Direct  0    0           D   10.0.1.1        LoopBack0
       10.0.1.1/32  Direct  0    0           D   127.0.0.1       LoopBack0
     10.0.1.255/32  Direct  0    0           D   127.0.0.1       LoopBack0
       10.0.2.2/24  OSPF    10   1           D   10.0.123.2      GigabitEthernet0/0/0
       10.0.3.3/32  OSPF    10   20          D   10.0.123.3      GigabitEthernet0/0/0
     10.0.123.0/24  Direct  0    0           D   10.0.123.1      GigabitEthernet0/0/0
     10.0.123.1/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
   10.0.123.255/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

On R3, check the cost of the route to 10.0.1.1/32. You can see that the cost is 10.

```
[R3]display ip routing-table
Route Flags: R - relay, D - download to fib
```

```
------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 12        Routes : 12

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.1/32  OSPF    10   10          D   10.0.123.1      GigabitEthernet0/0/0
       10.0.2.0/24  OSPF    10   10          D   10.0.123.2      GigabitEthernet0/0/0
       10.0.3.0/24  Direct  0    0           D   10.0.3.3        LoopBack0
       10.0.3.3/32  Direct  0    0           D   127.0.0.1       LoopBack0
     10.0.3.255/32  Direct  0    0           D   127.0.0.1       LoopBack0
     10.0.123.0/24  Direct  0    0           D   10.0.123.3      GigabitEthernet0/0/0
     10.0.123.3/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
   10.0.123.255/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

Step 6 **Configure OSPF interfaces as silent interfaces.**

Configure G0/0/0 of R1 as a silent interface.

```
[R1]ospf 1
[R1-ospf-1]silent-interface GigabitEthernet 0/0/0
[R1-ospf-1]quit
```

Run the **display ip routing-table** on R1 to check OSPF neighbor relationship establishment and routing entry learning on R1. The command output shows that the route learned from OSPF disappears in the IP routing table.

```
[R1]display ip routing-table
```

```
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 12        Routes : 12

Destination/Mask    Proto   Pre Cost      Flags NextHop         Interface

       10.0.1.0/24  Direct  0    0           D   10.0.1.1         LoopBack0
       10.0.1.1/32  Direct  0    0           D   127.0.0.1        LoopBack0
     10.0.1.255/32  Direct  0    0           D   127.0.0.1        LoopBack0
     10.0.123.0/24  Direct  0    0           D   10.0.123.1       GigabitEthernet0/0/0
     10.0.123.1/32  Direct  0    0           D   127.0.0.1        GigabitEthernet0/0/0
   10.0.123.255/32  Direct  0    0           D   127.0.0.1        GigabitEthernet0/0/0
      127.0.0.0/8   Direct  0    0           D   127.0.0.1        InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1        InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1        InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1        InLoopBack0
```

Check the neighbor list of R1. You can see that OSPF neighbor relationships between R1 and R2 and between R1 and R3 disapear. After a RIP interface is configured as a silent interface, this interface does not send RIP updates. In OSPF, routers can exchange routing information only after they establish an OSPF neighbor relationship. After an OSPF interface is configured as a silent interface, this interface does not receive or send Hello packets. As a result, this interface cannot establish OSPF neighbor relationships with interfaces of other routers.

```
[R1]display ospf interface GigabitEthernet 0/0/0

         OSPF Process 1 with Router ID 10.0.1.1
                 Interfaces
```

```
Interface: 10.0.123.1 (GigabitEthernet0/0/0)
Cost: 20      State: DR       Type: Broadcast    MTU: 1500
Priority: 1
Designated Router: 10.0.123.1
Backup Designated Router: 0.0.0.0
Timers: Hello 10 , Dead 40 , Poll   120 , Retransmit 5 , Transmit Delay 1
Silent interface, No hellos
```

Restore G0/0/0 of R1 to the default state and configure Loopback0 of the three routes as silent interfaces.

```
[R1]ospf 1
[R1-ospf-1]undo silent-interface GigabitEthernet0/0/0
[R1-ospf-1]silent-interface LoopBack 0
[R1-ospf-1]quit

[R2]ospf 1
[R2-ospf-1]silent-interface LoopBack 0
[R1-ospf-1]quit

[R3]ospf 1
[R3-ospf-1]silent-interface LoopBack 0
[R1-ospf-1]quit
```

Check the IP routing table of R1. The command output shows that configuring Loopback0 as a silent interface does not affect its route advertisement.

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
```

```
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 12        Routes : 12

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.0/24  Direct  0    0           D   10.0.1.1        LoopBack0
       10.0.1.1/32  Direct  0    0           D   127.0.0.1       LoopBack0
     10.0.1.255/32  Direct  0    0           D   127.0.0.1       LoopBack0
       10.0.2.0/24  OSPF    10   20          D   10.0.123.2      GigabitEthernet0/0/0
       10.0.3.3/32  OSPF    10   20          D   10.0.123.3      GigabitEthernet0/0/0
     10.0.123.0/24  Direct  0    0           D   10.0.123.1      GigabitEthernet0/0/0
     10.0.123.1/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
   10.0.123.255/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

----End

Additional Exercises: Analysis and Verification

Analyze why the wildcard mask 0.0.0.0 is used in OSPF configuration? The wildcard mask 0.0.0.255 can also be used in actual configuration, what are the differences of the two wildcard masks?

Analyze which types of interfaces should be configured as silent interfaces in real-world networks.

Device Configurations

```
<R1>display current-configuration
```

```
[V200R007C00SPC600]
#
 sysname R1
#
interface GigabitEthernet0/0/0
 ip address 10.0.123.1 255.255.255.0
 ospf cost 20
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.0
#
ospf 1 router-id 10.0.1.1
 silent-interface LoopBack0
 area 0.0.0.0
  authentication-mode simple plain huawei
  network 10.0.1.1 0.0.0.0
  network 10.0.123.1 0.0.0.0
#
return

<R2>display current-configuration
[V200R007C00SPC600]
#
 sysname R2
#
interface GigabitEthernet0/0/0
 ip address 10.0.123.2 255.255.255.0
#
interface LoopBack0
```

```
 ip address 10.0.2.2 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.2.2
 silent-interface LoopBack0
 area 0.0.0.0
  authentication-mode simple plain huawei
  network 10.0.2.2 0.0.0.0
  network 10.0.123.2 0.0.0.0
#
return

<R3>display current-configuration
[V200R007C00SPC600]
#
 sysname R3
#
interface GigabitEthernet0/0/0
 ip address 10.0.123.3 255.255.255.0
 ospf cost 10
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.0
#
ospf 1 router-id 10.0.3.3
 silent-interface LoopBack0
 area 0.0.0.0
  authentication-mode simple plain huawei
  network 10.0.3.3 0.0.0.0
```

```
 network 10.0.123.3 0.0.0.0
#
return
```

Lab 1-2 Multiple OSPF Areas

Learning Objectives

The objectives of this lab are to learn and understand:

- How to configure a router ID for an OSPF router
- How to configure multiple OSPF areas
- How to configure route summarization between OSPF areas
- How to set the reference bandwidth
- How to configure OSPF to import external routes
- How to summarize routes when OSPF imports external routes
- How to import default routes into OSPF
- How to change the priorities of OSPF routes

Topology

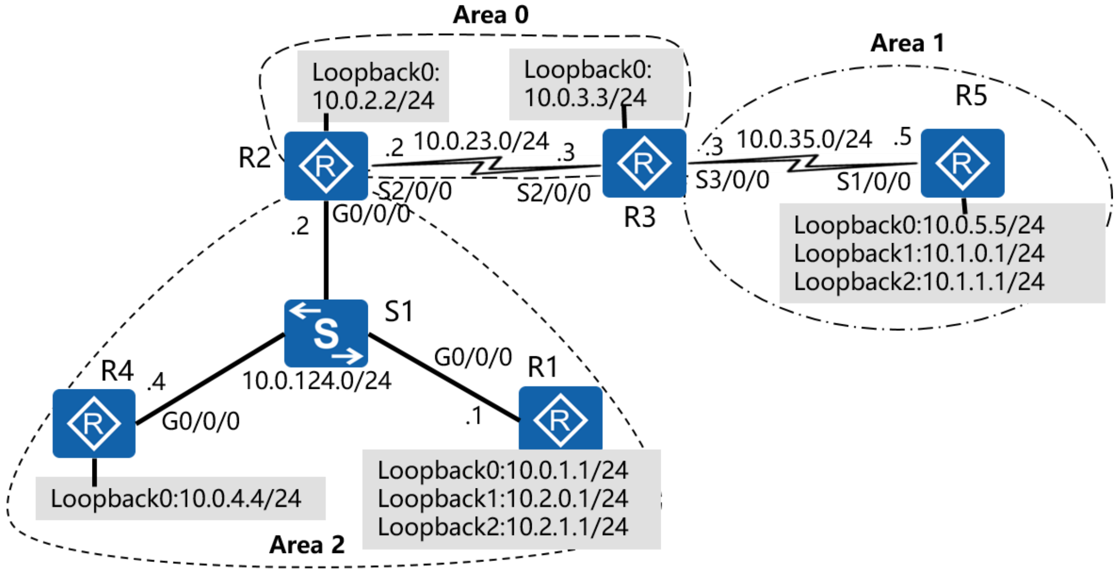

Figure 1-2 Multiple OSPF areas

Scenario

You are a network administrator of a company. There are five AR G3 routers in the network. R1, R2, and R4 are deployed in the headquarters and connected through an Ethernet. R3 and R5 are deployed in the branch. R3 is connected to R2 in the headquarters through a leased line, and R5 is connected to R3 through a leased line. Because of the large network scale, to control the flooding of LSAs, you design multiple OSPF areas for interconnection.

Loopback0 and interconnected interfaces of R2 and R3 belong to Area 0. The interconnected network segment between R3 and R5 and Loopback0/1/2 of R5 belong to Area 1. The interconnected network segment between R1, R2, and R4 and Loopback0 of R1 and R4 belong to Area 2.

To specify router IDs for the routers, configure the routers to use fixed addresses as their router IDs.

To improve routing forwarding efficiency for routers, you configure automatic summarization on the borders between areas.

R1 is connected to an external network of the company. You configure R1 to import routes outside the areas into these areas.

R4 is connected to the Internet. You need to configure a default route on R4 and import it into the areas so that all the routers in these areas know how to access the Internet.

OSPF routes are classified into internal and external routes. You change the priorities of OSPF routes to avoid risks.

In OSPF, the cost of a specific route is the sum of the costs of all the links through which a route reaches a destination network. The link cost is obtained through comparison between the interface bandwidth and reference bandwidth. The reference bandwidth is 100 Mbps, but actual interface bandwidth may be 1000 Mbps. Because the cost is an integer, the OSPF cost of both fast Ethernet (FE) interfaces and gigabit Ethernet (GE) interfaces is 1. To differentiate these links, you can define the reference bandwidth as 10 Gbps.

Some network faults occur during device configuration, you can run the display and debugging commands to rectify these faults.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

Configure IP addresses and masks for all the routers. Set a 24-bit mask for all loopback interfaces to simulate an independent network segment.

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]ip address 10.0.124.1 24
[R1-GigabitEthernet0/0/0]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 24
[R1-LoopBack0]quit
[R1]interface LoopBack 1
[R1-LoopBack1]ip address 10.2.0.1 24
[R1-LoopBack1]quit
[R1]interface LoopBack 2
[R1-LoopBack2]ip address 10.2.1.1 24
[R1-LoopBack2]quit

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.124.2 24
[R2-GigabitEthernet0/0/0]quit
[R2]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 24
```

```
[R2-Serial2/0/0]quit
[R2]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24
[R2-LoopBack0]quit

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 24
[R3-Serial2/0/0]quit
[R3]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 24
[R3-Serial3/0/0]quit
[R3]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 24

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ip address 10.0.124.4 24
[R4-GigabitEthernet0/0/0]quit
[R4]interface LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 24
[R4-LoopBack0]quit

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface Serial 1/0/0
[R5-Serial1/0/0]ip address 10.0.35.5 24
```

```
[R5-Serial1/0/0]quit
[R5]interface LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 24
[R5-LoopBack0]quit
[R5]interface LoopBack 1
[R5-LoopBack1]ip address 10.1.0.1 24
[R5-LoopBack1]quit
[R5]interface LoopBack 2
[R5-LoopBack2]ip address 10.1.1.1 24
[R5-LoopBack2]quit
```

After the configurations are complete, test direct link connectivity.

```
[R2]ping -c 1 10.0.124.1
  PING 10.0.124.1: 56   data bytes, press CTRL_C to break
    Reply from 10.0.124.1: bytes=56 Sequence=1 ttl=255 time=5 ms

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

[R2]ping -c 1 10.0.124.4
  PING 10.0.124.4: 56   data bytes, press CTRL_C to break
    Reply from 10.0.124.4: bytes=56 Sequence=1 ttl=255 time=14 ms

  --- 10.0.124.4 ping statistics ---
    1 packet(s) transmitted
```

```
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 14/14/14 ms

[R2]ping -c 1 10.0.23.3
  PING 10.0.23.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.3: bytes=56 Sequence=1 ttl=255 time=41 ms

  --- 10.0.23.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 41/41/41 ms

[R3]ping -c 1 10.0.35.5
  PING 10.0.35.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.35.5: bytes=56 Sequence=1 ttl=255 time=38 ms

  --- 10.0.35.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 38/38/38 ms
```

Step 2 **Configure multiple OSPF areas.**

To ensure stable router IDs, manually specify router IDs for routers. There are two methods to manually specify a router ID for a router. The first one is to run the **router id** command in the system view.

```
[R1]router id 10.0.1.1
```

The second one is to specify the **router-id** parameter when starting an OSPF process.

```
[R1]ospf 1 router-id 10.0.1.1
```

When both methods are used on a router to specify a router ID, only the router ID configured using the second method takes effect on the router. If multiple OSPF processes need to be started on a router and these processes must use different router IDs, you can only use the second method to specify router IDs for these processes.

On R1, configure Loopback0 and GigabitEthernet0/0/0 to belong to Area 2. To enable OSPF to advertise real masks of loopback interfaces, change the OSPF network type of loopback interfaces in all the areas to broadcast.

```
[R1]ospf 1 router-id 10.0.1.1
[R1-ospf-1]area 2
[R1-ospf-1-area-0.0.0.2]network 10.0.124.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]network 10.0.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]quit
[R1-ospf-1]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ospf network-type broadcast
[R1-LoopBack0]quit
```

On R2, configure Loopback0 and Serial2/0/0 to belong to Area 0 and GigabitEthernet0/0/0 to belong to Area 2.

```
[R2]ospf 1 router-id 10.0.2.2
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.23.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
```

[R2-ospf-1-area-0.0.0.0]quit

[R2-ospf-1]area 2

[R2-ospf-1-area-0.0.0.2]network 10.0.124.2 0.0.0.0

[R2-ospf-1-area-0.0.0.2]quit

[R2-ospf-1]quit

[R2]interface LoopBack 0

[R2-LoopBack0]ospf network-type broadcast

[R2-LoopBack0]quit

On R3, configure Loopback0 and Serial2/0/0 to belong to Area 0 and Serial3/0/0 to belong to Area 1.

[R3]ospf 1 router-id 10.0.3.3

[R3-ospf-1]area 0

[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0

[R3-ospf-1-area-0.0.0.0]network 10.0.23.3 0.0.0.0

[R3-ospf-1-area-0.0.0.0]quit

[R3-ospf-1]area 1

[R3-ospf-1-area-0.0.0.1]network 10.0.35.3 0.0.0.0

[R3-ospf-1-area-0.0.0.1]quit

[R3-ospf-1]quit

[R3]interface LoopBack 0

[R3-LoopBack0]ospf network-type broadcast

[R3-LoopBack0]quit

On R4, configure Loopback0 and GigabitEthernet0/0/0 to belong to Area 2.

[R4]ospf 1 router-id 10.0.4.4

[R4-ospf-1]area 2

[R4-ospf-1-area-0.0.0.2]network 10.0.4.4 0.0.0.0

[R4-ospf-1-area-0.0.0.2]network 10.0.124.4 0.0.0.0

[R4-ospf-1-area-0.0.0.2]quit

[R4-ospf-1]quit

[R4]interface LoopBack 0

[R4-LoopBack0]ospf network-type broadcast

[R4-LoopBack0]quit

On R5, configure loopback interfaces and Serial1/0/0 to belong to Area 1.

[R5]ospf 1 router-id 10.0.5.5

[R5-ospf-1]area 1

[R5-ospf-1-area-0.0.0.1]network 10.0.5.5 0.0.0.0

[R5-ospf-1-area-0.0.0.1]network 10.1.0.1 0.0.0.0

[R5-ospf-1-area-0.0.0.1]network 10.1.1.1 0.0.0.0

[R5-ospf-1-area-0.0.0.1]network 10.0.35.5 0.0.0.0

[R5-ospf-1-area-0.0.0.1]quit

[R5-ospf-1]quit

[R5]interface LoopBack 0

[R5-LoopBack0]ospf network-type broadcast

[R5-LoopBack0]quit

[R5]interface LoopBack 1

[R5-LoopBack1]ospf network-type broadcast

[R5-LoopBack1]quit

[R5]interface LoopBack 2

[R5-LoopBack2]ospf network-type broadcast

[R5-LoopBack2]quit

After the configurations are complete, check the IP routing table of R1.

[R1]display ip routing-table

Route Flags: R - relay, D - download to fib

```
------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 24        Routes : 24

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.0/24  Direct  0    0           D   10.0.1.1        LoopBack0
       10.0.1.1/32  Direct  0    0           D   127.0.0.1       LoopBack0
     10.0.1.255/32  Direct  0    0           D   127.0.0.1       LoopBack0
       10.0.2.0/24  OSPF    10   1           D   10.0.124.2      GigabitEthernet0/0/0
       10.0.3.0/24  OSPF    10   1563        D   10.0.124.2      GigabitEthernet0/0/0
       10.0.4.0/24  OSPF    10   1           D   10.0.124.4      GigabitEthernet0/0/0
       10.0.5.0/24  OSPF    10   3125        D   10.0.124.2      GigabitEthernet0/0/0
      10.0.23.0/24  OSPF    10   1563        D   10.0.124.2      GigabitEthernet0/0/0
      10.0.35.0/24  OSPF    10   3125        D   10.0.124.2      GigabitEthernet0/0/0
     10.0.124.0/24  Direct  0    0           D   10.0.124.1      GigabitEthernet0/0/0
     10.0.124.1/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
   10.0.124.255/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
       10.1.0.0/24  OSPF    10   3125        D   10.0.124.2      GigabitEthernet0/0/0
       10.1.1.0/24  OSPF    10   3125        D   10.0.124.2      GigabitEthernet0/0/0
     10.2.0.0/24    Direct  0    0           D   10.2.0.1        LoopBack1
     10.2.0.1/32    Direct  0    0           D   127.0.0.1       LoopBack1
     10.2.0.255/32  Direct  0    0           D   127.0.0.1       LoopBack1
       10.2.1.0/24  Direct  0    0           D   10.2.1.1        LoopBack2
       10.2.1.1/32  Direct  0    0           D   127.0.0.1       LoopBack2
     10.2.1.255/32  Direct  0    0           D   127.0.0.1       LoopBack2
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32Direct    0    0           D   127.0.0.1       InLoopBack0
```

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

R1 has all routing entries of the network.

On R1, test the connectivity to loopback interfaces of other routers.

```
[R1]ping -c 1 10.0.2.2
  PING 10.0.2.2: 56   data bytes, press CTRL_C to break
    Reply from 10.0.2.2: bytes=56 Sequence=1 ttl=255 time=3 ms

  --- 10.0.2.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 3/3/3 ms

[R1]ping -c 1 10.0.5.5
  PING 10.0.5.5: 56   data bytes, press CTRL_C to break
    Reply from 10.0.5.5: bytes=56 Sequence=1 ttl=253 time=88 ms

  --- 10.0.5.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 88/88/88 ms

[R1]ping -c 1 10.0.4.4
  PING 10.0.4.4: 56   data bytes, press CTRL_C to break
    Reply from 10.0.4.4: bytes=56 Sequence=1 ttl=255 time=3 ms
```

```
  --- 10.0.4.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 3/3/3 ms
```

Run the **display ospf brief** command on R2 to check basic OSPF information.

```
[R2]display ospf brief

         OSPF Process 1 with Router ID 10.0.2.2
                 OSPF Protocol Information

 RouterID: 10.0.2.2           Border Router:   AREA
 Multi-VPN-Instance is not enabled
 Global DS-TE Mode: Non-Standard IETF Mode
 Graceful-restart capability: disabled
 Helper support capability   : not configured
 Spf-schedule-interval: max 10000ms, start 500ms, hold 1000ms
 Default ASE parameters: Metric: 1 Tag: 1 Type: 2
 Route Preference: 10
 ASE Route Preference: 150
 SPF Computation Count: 19
 RFC 1583 Compatible
 Retransmission limitation is disabled
 Area Count: 2    Nssa Area Count: 0
 ExChange/Loading Neighbors: 0

Area: 0.0.0.0          (MPLS TE not enabled)
```

```
 Authtype: None    Area flag: Normal
 SPF scheduled Count: 18
 ExChange/Loading Neighbors: 0
 Router ID conflict state: Normal
 Area interface up count: 2

 Interface: 10.0.2.2 (LoopBack0)
 Cost: 0          State: DR           Type: Broadcast      MTU: 1500
 Priority: 1
 Designated Router: 10.0.2.2
 Backup Designated Router: 0.0.0.0
 Timers: Hello 10 , Dead 40 , Poll   120 , Retransmit 5 , Transmit Delay 1

 Interface: 10.0.23.2 (Serial2/0/0) --> 10.0.23.3
 Cost: 1562      State: P-2-P       Type: P2P          MTU: 1500
 Timers: Hello 10 , Dead 40 , Poll   120 , Retransmit 5 , Transmit Delay 1

Area: 0.0.0.2              (MPLS TE not enabled)
 Authtype: None    Area flag: Normal
 SPF scheduled Count: 16
 ExChange/Loading Neighbors: 0
 Router ID conflict state: Normal
 Area interface up count: 1

 Interface: 10.0.124.2 (GigabitEthernet0/0/0)
 Cost: 1          State: BDR          Type: Broadcast      MTU: 1500
 Priority: 1
 Designated Router: 10.0.124.1
 Backup Designated Router: 10.0.124.2
```

```
Timers: Hello 10 , Dead 40 , Poll    120 , Retransmit 5 , Transmit Delay 1
```

In the preceding command output, "Border Router: AREA" indicates that R2 is an ABR. If R2 is an intra-area router, the **Border Router** field is empty. If it is an ASBR, this field displays AS.

R2 has three interfaces that participate in OSPF route calculation. You have changed the network type of Loopback0 to broadcast. The encapsulation type of Serial2/0/0 is PPP. Therefore, the default network type is point-to-point (P2P). GigabitEthernet 0/0/0 is connected to Area 2 and its network type is broadcast.

Run the **display ospf peer brief** command on R2 to check information about OSPF neighbor relationships between the routers. The command output shows that in Area 0, R2 has a neighbor 10.0.3.3; in Area 2, R2 has two neighbors 10.0.1.1 and 10.0.4.4 and has established neighbor relationships with the two neighbors (in Full state).

```
[R2]display ospf peer brief

         OSPF Process 1 with Router ID 10.0.2.2
                  Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                        Neighbor id      State
 0.0.0.0          Serial2/0/0                        10.0.3.3       Full
 0.0.0.2          GigabitEthernet0/0/0               10.0.1.1       Full
 0.0.0.2          GigabitEthernet0/0/0               10.0.4.4       Full
 ----------------------------------------------------------------------------
```

Run the **display ospf lsdb** command on R2 to check OSPF LSDB information. The command output shows that R2 functioning as an ABR maintains two LSDBs, which describe routes of Area 0 and Area 2 respectively.

```
[R2]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.2.2
                 Link State Database

                         Area: 0.0.0.0
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router      10.0.3.3        10.0.3.3          788  60    80000008       0
 Router      10.0.2.2        10.0.2.2          869  60    80000008       0
 Sum-Net     10.0.35.0       10.0.3.3          846  28    80000002    1562
 Sum-Net     10.0.124.0      10.0.2.2         1259  28    80000002       1
Sum-Net      10.0.1.0        10.0.2.2          143  28    80000001       1
Sum-Net      10.1.1.0        10.0.3.3         1565  28    80000001    1562
 Sum-Net     10.0.5.0        10.0.3.3         1594  28    80000001    1562
 Sum-Net     10.1.0.0        10.0.3.3         1584  28    80000001    1562
 Sum-Net     10.0.4.0        10.0.2.2          538  28    80000002       1

                         Area: 0.0.0.2
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router      10.0.4.4        10.0.4.4          504  48    80000008       1
 Router      10.0.2.2        10.0.2.2          558  36    80000006       1
 Router      10.0.1.1        10.0.1.1          568  60    80000011       1
 Network     10.0.124.1      10.0.1.1          559  36    80000005       0
 Sum-Net     10.0.35.0       10.0.2.2          846  28    80000002    3124
 Sum-Net     10.0.3.0        10.0.2.2          830  28    80000002    1562
 Sum-Net     10.0.2.0        10.0.2.2         1249  28    80000002       0
 Sum-Net     10.1.1.0        10.0.2.2         1565  28    80000001    3124
 Sum-Net     10.0.5.0        10.0.2.2         1595  28    80000001    3124
 Sum-Net     10.1.0.0        10.0.2.2         1584  28    80000001    3124
```

```
Sum-Net      10.0.23.0      10.0.2.2          1261   28     80000002      1562
```

Step 3 **Configure inter-area route summarization.**

Check the OSPF routing tables of R2 and R3.

```
[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 7          Routes : 7

OSPF routing table status : <Active>
         Destinations : 7          Routes : 7

Destination/Mask    Proto   Pre  Cost     Flags NextHop         Interface

       10.0.1.0/24  OSPF    10   1          D   10.0.124.1      GigabitEthernet0/0/0
       10.0.3.0/24  OSPF    10   1562       D   10.0.23.3         Serial2/0/0
       10.0.4.0/24  OSPF    10   1          D   10.0.124.4      GigabitEthernet0/0/0
       10.0.5.0/24  OSPF    10   3124       D   10.0.23.3         Serial2/0/0
      10.0.35.0/24  OSPF    10   3124       D   10.0.23.3         Serial2/0/0
       10.1.0.0/24  OSPF    10   3124       D   10.0.23.3         Serial2/0/0
       10.1.1.0/24  OSPF    10   3124       D   10.0.23.3         Serial2/0/0

OSPF routing table status : <Inactive>
         Destinations : 0          Routes : 0

[R3]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
```

```
------------------------------------------------------------------------

Public routing table : OSPF
         Destinations : 7           Routes : 7

OSPF routing table status : <Active>
         Destinations : 7           Routes : 7

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.0/24  OSPF    10   1563        D   10.0.23.2       Serial2/0/0
       10.0.2.0/24  OSPF    10   1562        D   10.0.23.2       Serial2/0/0
       10.0.4.0/24  OSPF    10   1563        D   10.0.23.2       Serial2/0/0
       10.0.5.0/24  OSPF    10   1562        D   10.0.35.5       Serial3/0/0
     10.0.124.0/24  OSPF    10   1563        D   10.0.23.2       Serial2/0/0
       10.1.0.0/24  OSPF    10   1562        D   10.0.35.5       Serial3/0/0
       10.1.1.0/24  OSPF    10   1562        D   10.0.35.5       Serial3/0/0

OSPF routing table status : <Inactive>
         Destinations : 0           Routes : 0
```

Routing information of 10.1.0.0/24 and 10.1.1.0/24 is displayed as specific routes.

These specific routes can be summarized and then advertised to other areas. Route summarization can reduce the routing entries advertised to other areas and reduces route flappings. Run the **abr-summary** command on R3 to summarize the network segment of Loopback1 and Loopback2 of R5 for advertisement.

```
[R3]ospf 1
[R3-ospf-1]area 1
[R3-ospf-1-area-0.0.0.1]abr-summary 10.1.0.0 255.255.254.0
```

[R3-ospf-1-area-0.0.0.1]quit

[R3-ospf-1]quit

After the configurations are complete, check summarized route information on R3 and R2.

```
[R3]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 7        Routes : 7

OSPF routing table status : <Active>
         Destinations : 7        Routes : 7

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.0/24  OSPF    10   1563        D   10.0.23.2       Serial2/0/0
       10.0.2.0/24  OSPF    10   1562        D   10.0.23.2       Serial2/0/0
       10.0.4.0/24  OSPF    10   1563        D   10.0.23.2       Serial2/0/0
       10.0.5.0/24  OSPF    10   1562        D   10.0.35.5       Serial3/0/0
     10.0.124.0/24  OSPF    10   1563        D   10.0.23.2       Serial2/0/0
       10.1.0.0/24  OSPF    10   1562        D   10.0.35.5       Serial3/0/0
       10.1.1.0/24  OSPF    10   1562        D   10.0.35.5       Serial3/0/0

OSPF routing table status : <Inactive>
         Destinations : 0        Routes : 0

[R2]display ip routing-table protocol ospf
```

```
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 6          Routes : 6

OSPF routing table status : <Active>
         Destinations : 6          Routes : 6

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.0/24  OSPF    10   1           D   10.0.124.1      GigabitEthernet0/0/0
       10.0.3.0/24  OSPF    10   1562        D   10.0.23.3       Serial2/0/0
       10.0.4.0/24  OSPF    10   1           D   10.0.124.4      GigabitEthernet0/0/0
       10.0.5.0/24  OSPF    10   3124        D   10.0.23.3       Serial2/0/0
      10.0.35.0/24  OSPF    10   3124        D   10.0.23.3       Serial2/0/0
       10.1.0.0/23  OSPF    10   3124        D   10.0.23.3       Serial2/0/0

OSPF routing table status : <Inactive>
         Destinations : 0          Routes : 0
```

The preceding command output shows that in the OSPF routing table of R3, routes 10.1.0.0/24 and 10.1.1.0/24 are still displayed as specific routes; in the OSPF routing table of R2, only the summarized route 10.1.0.0/23 exists.

After the configurations are complete, test the connectivity between other routers and network segments 10.1.0.0/24 and 10.1.1.0/24.

```
[R1]ping -c 1 10.1.0.1
  PING 10.1.0.1: 56   data bytes, press CTRL_C to break
    Reply from 10.1.0.1: bytes=56 Sequence=1 ttl=253 time=66 ms
```

```
  --- 10.1.0.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 66/66/66 ms

[R1]ping -c 1 10.1.1.1
  PING 10.1.1.1: 56  data bytes, press CTRL_C to break
    Reply from 10.1.1.1: bytes=56 Sequence=1 ttl=253 time=66 ms

  --- 10.1.1.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 66/66/66 ms

[R2]ping -c 1 10.1.0.1
  PING 10.1.0.1: 56  data bytes, press CTRL_C to break
    Reply from 10.1.0.1: bytes=56 Sequence=1 ttl=254 time=69 ms

  --- 10.1.0.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 69/69/69 ms

[R3]ping -c 1 10.1.0.1
  PING 10.1.0.1: 56  data bytes, press CTRL_C to break
```

```
  Reply from 10.1.0.1: bytes=56 Sequence=1 ttl=255 time=29 ms

 --- 10.1.0.1 ping statistics ---
  1 packet(s) transmitted
  1 packet(s) received
  0.00% packet loss
  round-trip min/avg/max = 29/29/29 ms
```

Step 4 **Change the OSPF reference bandwidth.**

In real-world networks, you may use 1000M Ethernet and even 10G Ethernet. The default OSPF reference bandwidth is 100 Mbps and the interface cost is an integer. Therefore, OSPF cannot differentiate 100M Ethernet and 1000M Ethernet based on the bandwidth.

Run the **bandwidth-reference** command on R2 to change the OSPF reference bandwidth to 10 Gbps.

```
[R2-ospf-1]bandwidth-reference 10000
```

Check the OSPF routing table of R2 to learn OSPF neighbor relationships and routing information learning. In the OSPF routing table, the cost has changed.

```
[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 7           Routes : 7

OSPF routing table status : <Active>
         Destinations : 7           Routes : 7
```

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
10.0.3.0/24	OSPF	10	65535	D	10.0.23.3	Serial2/0/0
10.0.4.0/24	OSPF	10	10	D	10.0.124.4	GigabitEthernet0/0/0
10.0.5.0/24	OSPF	10	67097	D	10.0.23.3	Serial2/0/0
10.0.35.0/24	OSPF	10	67097	D	10.0.23.3	Serial2/0/0
10.1.0.0/23	OSPF	10	67097	D	10.0.23.3	Serial2/0/0

OSPF routing table status : <Inactive>

Destinations : 0 Routes : 0

In multiple OSPF areas, the OSPF reference bandwidth must be consistent. Otherwise, OSPF cannot run normally. Change the OSPF reference bandwidth of all routers to 10 Gbps.

```
[R1]ospf 1
[R1-ospf-1]bandwidth-reference 10000
[R1-ospf-1]quit

[R2]ospf 1
[R2-ospf-1]bandwidth-reference 10000
[R2-ospf-1]quit

[R3]ospf 1
[R3-ospf-1]bandwidth-reference 10000
[R3-ospf-1]quit

[R4]ospf 1
[R4-ospf-1]bandwidth-reference 10000
```

```
[R4-ospf-1]quit


[R5]ospf 1
[R5-ospf-1]bandwidth-reference 10000
[R5-ospf-1]quit
```

Check the neighbor list and OSPF routing table of R2 to determine whether OSPF neighbor relationships and routing information are normal.

```
[R2]display ospf peer brief

     OSPF Process 1 with Router ID 10.0.2.2
           Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                        Neighbor id       State
 0.0.0.0          Serial2/0/0                      10.0.3.3          Full
 0.0.0.2          GigabitEthernet0/0/0             10.0.1.1          Full
 0.0.0.2          GigabitEthernet0/0/0             10.0.4.4          Full
 ----------------------------------------------------------------------------

[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
----------------------------------------------------------------------------
Public routing table : OSPF
     Destinations : 6          Routes : 6

OSPF routing table status : <Active>
     Destinations : 6          Routes : 6
```

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

      10.0.1.0/24  OSPF    10   100          D   10.0.124.1    GigabitEthernet0/0/0
      10.0.3.0/24  OSPF    10   65535        D   10.0.23.3       Serial2/0/0
      10.0.4.0/24  OSPF    10   100          D   10.0.124.4    GigabitEthernet0/0/0
      10.0.5.0/24  OSPF    10   131070       D   10.0.23.3       Serial2/0/0
     10.0.35.0/24  OSPF    10   131070       D   10.0.23.3       Serial2/0/0
      10.1.0.0/23  OSPF    10   131070       D   10.0.23.3       Serial2/0/0

OSPF routing table status : <Inactive>
         Destinations : 0           Routes : 0
```

The preceding command output shows that routing information is normal. You can test network connectivity.

Step 5 **Summarize direct routes and import summarized routes into OSPF areas.**

Loopback1 and Loopback2 of R1 do not belong to an OSPF area. Import the network segments where the two loopback interfaces reside into an OSPF area and summarize the routes on R1.

```
[R1]ospf 1
[R1-ospf-1]import-route direct
[R1-ospf-1]asbr-summary 10.2.0.0 255.255.254.0
[R1-ospf-1]quit
```

Check external routing information on R1.

```
[R1]display ospf lsdb ase 10.2.0.0
```

OSPF Process 1 with Router ID 10.0.1.1

Link State Database

Type : External

Ls id : 10.2.0.0

Adv rtr : 10.0.1.1

Ls age : 293

Len : 36

Options : E

seq# : 80000001

chksum : 0x2b6

Net mask : 255.255.254.0

TOS 0 Metric: 2

E type : 2

Forwarding Address : 0.0.0.0

Tag : 1

Priority : Low

R1 uses a Type 5 LSA to advertise the network segment 10.2.0.0 to other routers. The subnet mask is 255.255.254.0.

Check summarized routes on other routers and test network connectivity.

[R2]display ip routing-table protocol ospf

Route Flags: R - relay, D - download to fib

Public routing table : OSPF

Destinations : 7 Routes : 7

OSPF routing table status : <Active>

```
         Destinations : 7          Routes : 7

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.0/24  O_ASE   150  100         D   10.0.124.1      GigabitEthernet0/0/0
       10.0.3.0/24  OSPF    10   65535       D   10.0.23.3       Serial2/0/0
       10.0.4.0/24  OSPF    10   100         D   10.0.124.4      GigabitEthernet0/0/0
       10.0.5.0/24  OSPF    10   131070      D   10.0.23.3       Serial2/0/0
      10.0.35.0/24  OSPF    10   131070      D   10.0.23.3       Serial2/0/0
       10.1.0.0/23  OSPF    10   131070      D   10.0.23.3       Serial2/0/0
       10.2.0.0/23  O_ASE   150  2           D   10.0.124.1      GigabitEthernet0/0/0

OSPF routing table status : <Inactive>
         Destinations : 0          Routes : 0

[R2]ping -c 1 10.2.0.1
  PING 10.2.0.1: 56  data bytes, press CTRL_C to break
    Reply from 10.2.0.1: bytes=56 Sequence=1 ttl=255 time=2 ms

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

[R2]ping -c 1 10.2.1.1
  PING 10.2.1.1: 56  data bytes, press CTRL_C to break
    Reply from 10.2.1.1: bytes=56 Sequence=1 ttl=255 time=2 ms
```

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

You can see a summarized route with a 23-bit mask on R2.

Delete Loopback2 of R1 and then check the routing entry change on R2. You can see that Loopback2 does not exist but the summarized route still exists.

```
[R1]undo interface LoopBack 2


[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 7        Routes : 7


OSPF routing table status : <Active>
         Destinations : 7        Routes : 7


Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface


       10.0.1.0/24  ospf    150 100          D   10.0.124.1      GigabitEthernet0/0/0
       10.0.3.0/24  OSPF    10   65535       D   10.0.23.3       Serial2/0/0
       10.0.4.0/24  OSPF    10   100         D   10.0.124.4      GigabitEthernet0/0/0
       10.0.5.0/24  OSPF    10   131070      D   10.0.23.3       Serial2/0/0
      10.0.35.0/24  OSPF    10   131070      D   10.0.23.3       Serial2/0/0
       10.1.0.0/23  OSPF    10   131070      D   10.0.23.3       Serial2/0/0
```

```
    10.2.0.0/23  O_ASE   150  2        D   10.0.124.1    GigabitEthernet0/0/0

OSPF routing table status : <Inactive>
        Destinations : 0         Routes : 0
```

On R5, send a tracert packet to the address 10.2.1.1.

```
<R5>tracert 10.2.1.1
 traceroute to  10.2.1.1(10.2.1.1), max hops: 30 ,packet length: 40,press CTRL_C to break
 1 10.0.35.3 62 ms  28 ms  27 ms
 2 10.0.23.2 54 ms  58 ms  57 ms
 3  *  *  *
...
```

Although Loopback2 has been deleted, the packet to this destination address is still forwarded by R2 and R3 until it is discarded by R1.

Step 6 **Configure OSPF to import default routes.**

Loopback0 of R4 is connected to the Internet. Configure a default route on R4 with the next hop pointing to Loopback0.

```
[R4]ip route-static 0.0.0.0 0.0.0.0 LoopBack 0
```

Import this default route into an OSPF area, define its type as Type 1, set its cost to 10, and configure permanent advertisement of this default route.

```
[R4]ospf 1
[R4-ospf-1]default-route-advertise always type 1
[R4-ospf-1]quit
```

Check default route learning on R2. You can see that R2 learns a default route using a Type 5 LSA, and the next hop is the interface address of R4.

```
[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 8          Routes : 8

OSPF routing table status : <Active>
         Destinations : 8          Routes : 8

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        0.0.0.0/0   O_ASE   150  101         D   10.0.124.4      GigabitEthernet0/0/0
       10.0.1.0/24  ospf    10   100         D   10.0.124.1      GigabitEthernet0/0/0
       10.0.3.0/24  OSPF    10   65535       D   10.0.23.3       Serial2/0/0
       10.0.4.0/24  OSPF    10   100         D   10.0.124.4      GigabitEthernet0/0/0
       10.0.5.0/24  OSPF    10   131070      D   10.0.23.3       Serial2/0/0
      10.0.35.0/24  OSPF    10   131070      D   10.0.23.3       Serial2/0/0
       10.1.0.0/23  OSPF    10   131070      D   10.0.23.3       Serial2/0/0
       10.2.0.0/23  O_ASE   150  2           D   10.0.124.1      GigabitEthernet0/0/0

OSPF routing table status : <Inactive>
         Destinations : 0          Routes : 0
```

Test the connectivity between R5 and the address 10.0.4.4.

```
[R5]ping -c 1 10.0.4.4
  PING 10.0.4.4: 56  data bytes, press CTRL_C to break
```

```
Reply from 10.0.4.4: bytes=56 Sequence=1 ttl=253 time=78 ms


 --- 10.0.4.4 ping statistics ---
  1 packet(s) transmitted
  1 packet(s) received
  0.00% packet loss
  round-trip min/avg/max = 78/78/78 ms
```

Step 7 **Change the priorities of internal and external routes.**

Check the OSPF routing table of R1 and focus on priorities of different types of routes.

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 8        Routes : 8


OSPF routing table status : <Active>
         Destinations : 8        Routes : 8


Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface


        0.0.0.0/0   O_ASE   150  101          D   10.0.124.4      GigabitEthernet0/0/0
      10.0.2.0/24   OSPF    10   100          D   10.0.124.2      GigabitEthernet0/0/0
      10.0.3.0/24   OSPF    10   65635        D   10.0.124.2      GigabitEthernet0/0/0
      10.0.4.0/24   OSPF    10   100          D   10.0.124.4      GigabitEthernet0/0/0
      10.0.5.0/24   OSPF    10   131170       D   10.0.124.2      GigabitEthernet0/0/0
     10.0.23.0/24   OSPF    10   65635        D   10.0.124.2      GigabitEthernet0/0/0
```

```
10.0.35.0/24   OSPF    10    131170  D    10.0.124.2     GigabitEthernet0/0/0
 10.1.0.0/23   OSPF    10    131170  D    10.0.124.2     GigabitEthernet0/0/0


OSPF routing table status : <Inactive>
         Destinations : 0          Routes : 0
```

By default, the priorities of OSPF intra-area and inter-area routes are 10. The priorities of OSPF external routes are 150.

On R1 and R4, change the priorities of OSPF intra-area and inter-area routes to 20 and those of OSPF external routes to 50.

```
[R1]ospf 1
[R1-ospf-1]preference 20
[R1-ospf-1]preference ase 50
[R1-ospf-1]quit

[R4]ospf 1
[R4-ospf-1]preference 20
[R4-ospf-1]preference ase 50
[R4-ospf-1]quit
```

Check the priorities of OSPF internal and external routes in the OSPF routing table of R1. The following command output shows that their priorities have been changed successfully.

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 8          Routes : 8
```

```
OSPF routing table status : <Active>
         Destinations : 8          Routes : 8

Destination/Mask    Proto    Pre   Cost      Flags NextHop          Interface

       0.0.0.0/0    O_ASE    50    101         D   10.0.124.4       GigabitEthernet0/0/0
      10.0.2.0/24   OSPF     20    100         D   10.0.124.2       GigabitEthernet0/0/0
      10.0.3.0/24   OSPF     20    65545       D   10.0.124.2       GigabitEthernet0/0/0
      10.0.4.0/24   OSPF     20    100         D   10.0.124.4       GigabitEthernet0/0/0
      10.0.5.0/24   OSPF     20    131170      D   10.0.124.2       GigabitEthernet0/0/0
     10.0.23.0/24   OSPF     20    65635       D   10.0.124.2       GigabitEthernet0/0/0
     10.0.35.0/24   OSPF     20    131170      D   10.0.124.2       GigabitEthernet0/0/0
      10.1.0.0/23   OSPF     20    131170      D   10.0.124.2       GigabitEthernet0/0/0

OSPF routing table status : <Inactive>
         Destinations : 0          Routes : 0
```

Route priorities take effect only on a router to help select the optimal route among multiple routes learned using multiple methods. If route priorities on different routers within the same area are different, these routers can still work normally.

----End

Additional Exercises: Analysis and Verification

Analyze what is the function of configuring permanent advertisement of default routes in step 6 and what are the advantages and disadvantages ot this function.

Route summarization has advantages and disadvantages. Analyze how to avoid these disadvantages.

Device Configurations

```
<R1>display current-configuration
[V200R007C00SPC600]
#
 sysname R1
#
interface GigabitEthernet0/0/0
 ip address 10.0.124.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.0
 ospf network-type broadcast
#
interface LoopBack1
 ip address 10.2.0.1 255.255.255.0
#
ospf 1 router-id 10.0.1.1
 asbr-summary 10.2.0.0 255.255.254.0
 import-route direct
 preference 20
 preference ase 50
 bandwidth-reference 10000
 area 0.0.0.2
  network 10.0.1.1 0.0.0.0
  network 10.0.124.1 0.0.0.0
#
return
```

```
<R2>display current-configuration
```

```
[V200R007C00SPC600]
#
 sysname R2
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.124.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.2.2
 bandwidth-reference 10000
 area 0.0.0.0
    network 10.0.2.2 0.0.0.0
network 10.0.23.2 0.0.0.0
area 0.0.0.2
  network 10.0.124.2 0.0.0.0
#
return

<R3>display current-configuration
[V200R007C00SPC600]
#
 sysname R3
```

```
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.3.3
 bandwidth-reference 10000
 area 0.0.0.0
  network 10.0.3.3 0.0.0.0
  network 10.0.23.3 0.0.0.0
 area 0.0.0.1
  abr-summary 10.1.0.0 255.255.254.0
  network 10.0.35.3 0.0.0.0
#
return

<R4>display current-configuration
[V200R007C00SPC600]
#
 sysname R4
#
```

```
interface GigabitEthernet0/0/0
 ip address 10.0.124.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.4.4
 default-route-advertise always type 1
 preference 20
 preference ase 50
 bandwidth-reference 10000
 area 0.0.0.2
  network 10.0.4.4 0.0.0.0
  network 10.0.124.4 0.0.0.0
#
 ip route-static 0.0.0.0 0.0.0.0 LoopBack0
#
return

<R5>display current-configuration
[V200R007C00SPC600]
#
 sysname R5
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
#
```

```
interface LoopBack0
 ip address 10.0.5.5 255.255.255.0
 ospf network-type broadcast
#
interface LoopBack1
 ip address 10.1.0.1 255.255.255.0
 ospf network-type broadcast
#
interface LoopBack2
 ip address 10.1.1.1 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.5.5
 bandwidth-reference 10000
 area 0.0.0.1
  network 10.0.5.5 0.0.0.0
  network 10.1.0.1 0.0.0.0
  network 10.1.1.1 0.0.0.0
  network 10.0.35.5 0.0.0.0
#
return
```

Lab 1-3 OSPF Neighbor Relationship and LSA

Learning Objectives

The objectives of this lab are to learn and understand:

- Procedure for establishing OSPF neighbor relationships on an Ethernet
- How to affect DR election
- What are the content and functions of five types of LSAs
- Transmission of OSPF Link State Request (LSR), Link State Update (LSU), and Link State Acknowledgement (LSAck) packets

Topology

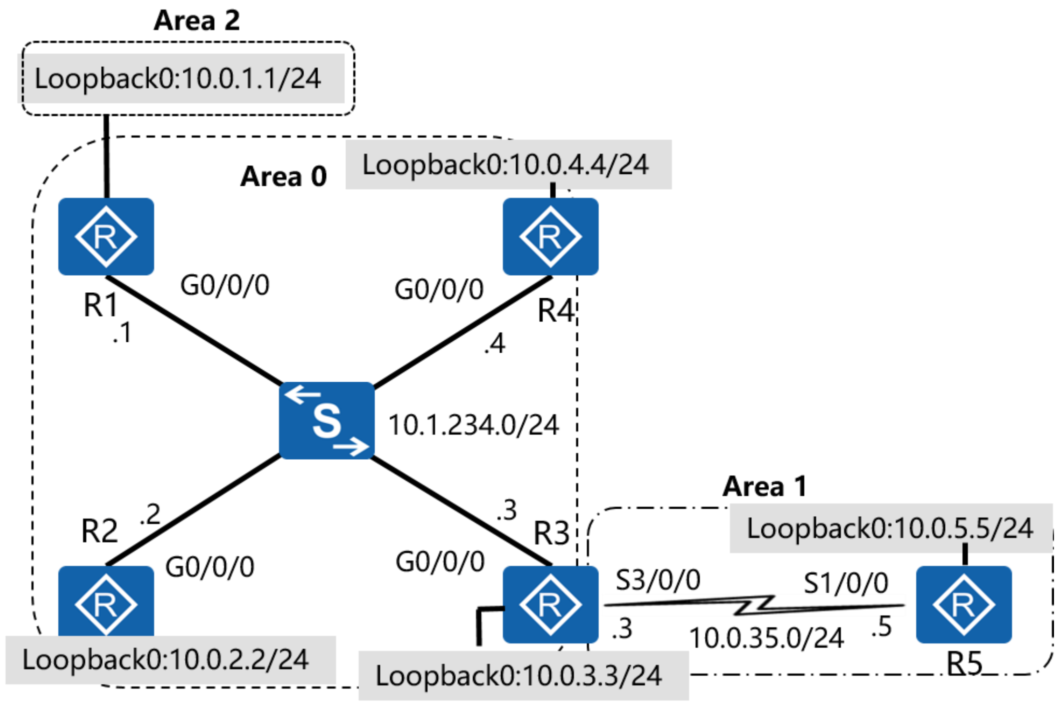

Figure 1-3 OSPF neighbor relationship and LSA

Scenario

You are a network administrator of a company. There are five AR G3 routers in the network. R1, R2, R3, and R4 are deployed in the headquarters and connected through an Ethernet. R5 is deployed in the branch and is connected to R3 in the headquarters through a leased line. Because of the large network scale, to control the flooding of LSAs, you design multiple OSPF areas for interconnection.

Loopback0 of R1 belongs to Area 2. Loopback0 of R2, R3, and R4 and the network segment 10.1.234.0/24 belong to Area 0. The interconnected network segment between R3 and R5 belongs to Area 1. Loopback0 of R5 belongs to an OSPF external network.

To specify router IDs for the routers, configure the routers to use fixed addresses as their router IDs.

You need to affect DR election and BDR election on the interconnected network between R1, R2, R3, and R4. That is, you need to configure R3 as the DR, R2 as the BDR, and R4 as the DR other.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

Configure IP addresses and masks for all the routers. Set a 24-bit mask for all loopback interfaces to simulate an independent network segment.

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]ip address 10.1.234.1 24
[R1-GigabitEthernet0/0/0]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 24
[R1-LoopBack0]quit
```

```
<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.1.234.2 24
[R2-GigabitEthernet0/0/0]quit
[R2]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24
[R2-LoopBack0]quit

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]ip address 10.1.234.3 24
[R3-GigabitEthernet0/0/0]quit
[R3]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 24
[R3-Serial3/0/0]quit
[R3]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 24
[R3-LoopBack0]quit

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ip address 10.1.234.4 24
[R4-GigabitEthernet0/0/0]quit
[R4]interface LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 24
```

```
[R4-LoopBack0]quit

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface Serial 1/0/0
[R5-Serial1/0/0]ip address 10.0.35.5 24
[R5-Serial1/0/0]quit
[R5]interface LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 24
[R5-LoopBack0]quit
```

After the configurations are complete, test direct link connectivity.

```
[R1]ping -c 1 10.1.234.2
  PING 10.1.234.2: 56  data bytes, press CTRL_C to break
    Reply from 10.1.234.2: bytes=56 Sequence=1 ttl=255 time=13 ms

  --- 10.1.234.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 13/13/13 ms

[R1]ping -c 1 10.1.234.4
  PING 10.1.234.4: 56  data bytes, press CTRL_C to break
    Reply from 10.1.234.4: bytes=56 Sequence=1 ttl=255 time=6 ms

  --- 10.1.234.4 ping statistics ---
    1 packet(s) transmitted
```

```
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 6/6/6 ms

[R3]ping -c 1 10.1.234.1
  PING 10.1.234.1: 56  data bytes, press CTRL_C to break
    Reply from 10.1.234.1: bytes=56 Sequence=1 ttl=255 time=13 ms

  --- 10.1.234.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 13/13/13 ms

[R3]ping -c 1 10.0.35.5
  PING 10.0.35.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.35.5: bytes=56 Sequence=1 ttl=255 time=32 ms

  --- 10.0.35.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 32/32/32 ms
```

Step 2 **Configure multiple OSPF areas.**

On R1, configure GigabitEthernet0/0/0 to belong to Area 0 and Looback0 to belong to Area 2. To enable OSPF to advertise real masks of loopback interfaces, change the OSPF network type of loopback interfaces in all the areas to broadcast.

```
[R1]ospf 1 router-id 10.0.1.1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.1.234.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]quit
[R1-ospf-1]area 2
[R1-ospf-1-area-0.0.0.2]network 10.0.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]quit
[R1-ospf-1]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ospf network-type broadcast
[R1-LoopBack0]quit
```

Configure all the interfaces of R2 and R4 to belong to Area 0.

```
[R2]ospf 1 router-id 10.0.2.2
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.1.234.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]quit
[R2-ospf-1]quit
[R2-]interface LoopBack 0
[R2-LoopBack0]ospf network-type broadcast
[R2-LoopBack0]quit

[R4]ospf 1 router-id 10.0.4.4
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]network 10.1.234.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]network 10.0.4.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]quit
```

[R4-ospf-1]quit

[R4-]interface LoopBack 0

[R4-LoopBack0]ospf network-type broadcast

[R4-LoopBack0]quit

On R3, configure Loopback0 and GigabitEthernet0/0/0 to belong to Area 0 and Serial3/0/0 to belong to Area 2.

[R3]ospf 1 router-id 10.0.3.3

[R3-ospf-1]area 0

[R3-ospf-1-area-0.0.0.0]network 10.1.234.3 0.0.0.0

[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0

[R3-ospf-1-area-0.0.0.0]quit

[R3-ospf-1]area 1

[R3-ospf-1-area-0.0.0.1]network 10.0.35.3 0.0.0.0

[R3-ospf-1-area-0.0.0.1]quit

[R3-ospf-1]quit

[R3]interface LoopBack 0

[R3-LoopBack0]ospf network-type broadcast

[R3-LoopBack0]quit

On R5, configure Serial1/0/0 to belong to Area 1 and configure Looback0 not to belong to any area.

[R5]osp 1 router-id 10.0.5.5

[R5-ospf-1]area 1

[R5-ospf-1-area-0.0.0.1]network 10.0.35.5 0.0.0.0

[R5-ospf-1-area-0.0.0.1]quit

[R5-ospf-1]quit

After the configurations are complete, check the IP routing table of R1.

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 14        Routes : 14

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.0/24  Direct  0    0           D   10.0.1.1        LoopBack0
       10.0.1.1/32  Direct  0    0           D   127.0.0.1       LoopBack0
     10.0.1.255/32  Direct  0    0           D   127.0.0.1       LoopBack0
       10.0.2.0/24  OSPF    10   1           D   10.1.234.2      GigabitEthernet0/0/0
       10.0.3.0/24  OSPF    10   1           D   10.1.234.3      GigabitEthernet0/0/0
       10.0.4.0/24  OSPF    10   1           D   10.1.234.4      GigabitEthernet0/0/0
      10.0.35.0/24  OSPF    10   1563        D   10.1.234.3      GigabitEthernet0/0/0
     10.1.234.0/24  Direct  0    0           D   10.1.234.1      GigabitEthernet0/0/0
     10.1.234.1/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
   10.1.234.255/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct  0    0           D   127.0.0.1       InLoopBack0
```

The preceding command output shows that R1 has the routes of the entire network except the network segment 10.0.5.5/24 that is not advertised into OSPF.

Test network connectivity.

```
[R1]ping -c 1 10.0.2.2
```

```
  PING 10.0.2.2: 56   data bytes, press CTRL_C to break
    Reply from 10.0.2.2: bytes=56 Sequence=1 ttl=255 time=2 ms

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

[R1]ping -c 1 10.0.4.4
  PING 10.0.4.4: 56   data bytes, press CTRL_C to break
    Reply from 10.0.4.4: bytes=56 Sequence=1 ttl=255 time=3 ms

  --- 10.0.4.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 3/3/3 ms

[R3]ping -c 1 10.0.1.1
  PING 10.0.1.1: 56   data bytes, press CTRL_C to break
    Reply from 10.0.1.1: bytes=56 Sequence=1 ttl=255 time=3 ms

  --- 10.0.1.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 3/3/3 ms
```

Run the **display ospf brief** command on R1 to check basic OSPF information running on the routers. The command output shows that R1 has become an ABR because its Loopback0 belongs to Area 2. The network segment to which GigabitEthernet0/0/0 of R1 is connected is a broadcast network, and R1 is the DR of this network segment.

```
[R1]display ospf brief

         OSPF Process 1 with Router ID 10.0.1.1
                 OSPF Protocol Information

 RouterID: 10.0.1.1          Border Router:   AREA
 Multi-VPN-Instance is not enabled
 Global DS-TE Mode: Non-Standard IETF Mode
 Graceful-restart capability: disabled
 Helper support capability   : not configured
 Applications Supported: MPLS Traffic-Engineering
 Spf-schedule-interval: max 10000ms, start 500ms, hold 1000ms
 Default ASE parameters: Metric: 1 Tag: 1 Type: 2
 Route Preference: 10
 ASE Route Preference: 150
 SPF Computation Count: 22
 RFC 1583 Compatible
 Retransmission limitation is disabled
 Area Count: 2    Nssa Area Count: 0
 ExChange/Loading Neighbors: 0
 Process total up interface count: 2
 Process valid up interface count: 1

 Area: 0.0.0.0           (MPLS TE not enabled)
```

Authtype: None Area flag: Normal

SPF scheduled Count: 22

ExChange/Loading Neighbors: 0

Router ID conflict state: Normal

Area interface up count: 1

Interface: 10.1.234.1 (GigabitEthernet0/0/0)

Cost: 1 State: DR Type: Broadcast MTU: 1500

Priority: 1

Designated Router: 10.1.234.1

Backup Designated Router: 10.1.234.2

Timers: Hello 10 , Dead 40 , Poll 120 , Retransmit 5 , Transmit Delay 1

Area: 0.0.0.2 (MPLS TE not enabled)

Authtype: None Area flag: Normal

SPF scheduled Count: 20

ExChange/Loading Neighbors: 0

Router ID conflict state: Normal

Area interface up count: 1

Interface: 10.0.1.1 (LoopBack0)

Cost: 0 State: DR Type: Broadcast MTU: 1500

Priority: 1

Designated Router: 10.0.1.1

Backup Designated Router: 0.0.0.0

Timers: Hello 10 , Dead 40 , Poll 120 , Retransmit 5 , Transmit Delay 1

Run the **display ospf peer brief** command on R1 to check information about OSPF neighbor relationships between the routers. Because R1 is the DR, it has established OSPF neighbor relationships with all the routers on this network segment. Run the **display ospf peer brief** command on R3. The command output shows that R3 and R4 establish an OSPF neighbor relationship instead of an adjacency.

```
[R1]display ospf peer brief

          OSPF Process 1 with Router ID 10.0.1.1
                    Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                      Neighbor id      State
 0.0.0.0          GigabitEthernet0/0/0           10.0.2.2         Full
 0.0.0.0          GigabitEthernet0/0/0           10.0.3.3         Full
 0.0.0.0          GigabitEthernet0/0/0           10.0.4.4         Full
 ----------------------------------------------------------------------------

[R3]display ospf peer brief

          OSPF Process 1 with Router ID 10.0.3.3
                    Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                      Neighbor id      State
 0.0.0.0          GigabitEthernet0/0/0            10.0.1.1        Full
 0.0.0.0          GigabitEthernet0/0/0            10.0.2.2        Full
 0.0.0.0          GigabitEthernet0/0/0            10.0.4.4        2-Way
 0.0.0.1          Serial3/0/0                     10.0.5.5        Full
 ----------------------------------------------------------------------------
```

Run the **display ospf lsdb** command on R5 to check OSPF LSDB information.

```
[R5]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.5.5
                 Link State Database

                         Area: 0.0.0.1
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.5.5           10.0.5.5          1182  48    80000002    1562
 Router    10.0.3.3           10.0.3.3          1183  48    80000002    1562
 Sum-Net    10.0.3.0          10.0.3.3          1429  28    80000001       0
 Sum-Net    10.0.2.0          10.0.3.3          1429  28    80000001       1
 Sum-Net    10.0.1.0          10.0.3.3          1429  28    80000001       1
 Sum-Net    10.1.234.0        10.0.3.3          1429  28    80000001       1
 Sum-Net    10.0.4.0          10.0.3.3          1430  28    80000001       1
```

The preceding command output shows that Area 1 has only two routers. Therefore, in the LSDB of R5, there are only two Type 1 LSAs, and the remaining Type 3 LSAs describe inter-area routes that are advertised from R3 to R5.

Run the **display ospf lsdb** command on R2 to check OSPF LSDB information.

```
[R2]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.2.2
                 Link State Database

                         Area: 0.0.0.0
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.3.3           10.0.3.3             4  48    80000009       1
```

```
Router    10.0.4.4       10.0.4.4        150   48    80000009      1
Router    10.0.2.2       10.0.2.2        149   48    8000000C      1
Router    10.0.1.1       10.0.1.1        149   36    8000000B      1
Network   10.1.234.1     10.0.1.1        149   40    80000007      0
Sum-Net   10.0.35.0      10.0.3.3       1790   28    80000001   1562
Sum-Net   10.0.1.0       10.0.1.1        817   28    80000002      0
```

The preceding command output shows that R2 still has one Type 2 LSA in addition to four Type 1 LSAs. GigabitEthernet0/0/0 of R2 is connected to a broadcast network, and the DR on this network will generate a Type 2 LSA to describe all neighbors. The **AdvRouter** field indicates that R1 generates this LSA. That is, the DR of this network segment generates this LSA, which meets the requirements.

Run the **display ospf lsdb** command on R1 to check OSPF LSDB information.

```
[R1]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.1.1
                 Link State Database

                      Area: 0.0.0.0
Type      LinkState ID   AdvRouter       Age  Len    Sequence   Metric
Router    10.0.3.3       10.0.3.3        447   48    80000009      1
Router    10.0.4.4       10.0.4.4        592   48    80000009      1
Router    10.0.2.2       10.0.2.2        592   48    8000000C      1
Router    10.0.1.1       10.0.1.1        591   36    8000000B      1
Network   10.1.234.1     10.0.1.1        591   40    80000007      0
Sum-Net   10.0.35.0      10.0.3.3        434   28    80000002   1562
Sum-Net   10.0.1.0       10.0.1.1       1259   28    80000002      0
```

```
                    Area: 0.0.0.2
Type       LinkState ID    AdvRouter        Age  Len   Sequence   Metric
Router     10.0.1.1         10.0.1.1        1223  36    80000004      0
Sum-Net    10.0.35.0        10.0.1.1         433  28    80000002   1563
Sum-Net    10.0.3.0         10.0.1.1         541  28    80000002      1
Sum-Net    10.0.2.0         10.0.1.1         909  28    80000002      1
Sum-Net    10.1.234.0       10.0.1.1        1269  28    80000002      1
Sum-Net    10.0.4.0         10.0.1.1         711  28    80000002      1;
```

Loopback0 of R1 belongs to Area 2. Therefore, R1 has LSDBs of two areas: Area 0 and Area 2.

Run the **display ospf lsdb** command on R4 to check OSPF LSDB information.

```
[R4]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.4.4
                 Link State Database

                    Area: 0.0.0.0
Type       LinkState ID    AdvRouter        Age  Len   Sequence   Metric
Router     10.0.3.3         10.0.3.3         745  48    80000009      1
Router     10.0.4.4         10.0.4.4         888  48    80000009      1
Router     10.0.2.2         10.0.2.2         889  48    8000000C      1
Router     10.0.1.1         10.0.1.1         889  36    8000000B      1
Network    10.1.234.1       10.0.1.1         889  40    80000007      0
Sum-Net    10.0.35.0        10.0.3.3         732  28    80000002   1562
Sum-Net    10.0.1.0         10.0.1.1        1556  28    80000002      0
```

LSDB information will vary depending on OSPF router roles. Analyze differences in LSDBs of R5, R2, R1, and R4.

Step 3 **Change OSPF interface priorities of routers to affect DR election.**

Set the priority of G0/0/0 on R3 to 255 to ensure that R3 becomes the DR of the network segment 10.1.234.0/24. Set the priority of G0/0/0 on R2 to 254 to ensure that R2 becomes the BDR of the network segment 10.1.234.0/24. Set the priority of G0/0/0 on R4 to 0 to ensure that R4 does not participate in DR/ BDR election and becomes the DR other of the network segment 10.1.234.0/24.

```
[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]ospf dr-priority 255
[R3-GigabitEthernet0/0/0]quit

[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ospf dr-priority 254
[R2-GigabitEthernet0/0/0]quit

[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ospf dr-priority 0
[R4-GigabitEthernet0/0/0]quit
```

After the configurations are complete, the DR and BDR have been elected and cannot be preempted. Therefore, G0/0/0 of R1, R2, R3, and R4 must be shut down and G0/0/0 of R3, R2, R1, and R4 must be enabled in sequence.

```
[R1]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]shutdown

[R2]interface GigabitEthernet 0/0/0
```

```
[R2-GigabitEthernet0/0/0]shutdown

[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]shutdown

[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]shutdown

[R1-GigabitEthernet0/0/0]undo shutdown
[R1-GigabitEthernet0/0/0]quit

[R2-GigabitEthernet0/0/0]undo shutdown
[R2-GigabitEthernet0/0/0]quit

[R3-GigabitEthernet0/0/0]undo shutdown
[R3-GigabitEthernet0/0/0]quit

[R4-GigabitEthernet0/0/0]undo shutdown
[R4-GigabitEthernet0/0/0]quit
```

Check DR/BDR election on the network segment 10.1.234.0/24.

```
[R3]display ospf peer

         OSPF Process 1 with Router ID 10.0.3.3
                 Neighbors

 Area 0.0.0.0 interface 10.1.234.3(GigabitEthernet0/0/0)'s neighbors
 Router ID: 10.0.1.1          Address: 10.1.234.1
```

```
  State: Full   Mode:Nbr is   Slave   Priority: 1
  DR: 10.1.234.3   BDR: 10.1.234.2   MTU: 0
  Dead timer due in 29   sec
  Retrans timer interval: 3
  Neighbor is up for 00:02:17
  Authentication Sequence: [ 0 ]

Router ID: 10.0.2.2          Address: 10.1.234.2
  State: Full   Mode:Nbr is   Slave   Priority: 254
  DR: 10.1.234.3   BDR: 10.1.234.2   MTU: 0
  Dead timer due in 35   sec
  Retrans timer interval: 6
  Neighbor is up for 00:01:14
  Authentication Sequence: [ 0 ]

Router ID: 10.0.4.4          Address: 10.1.234.4
  State: Full   Mode:Nbr is   Master   Priority: 0
  DR: 10.1.234.3   BDR: 10.1.234.2   MTU: 0
  Dead timer due in 32   sec
  Retrans timer interval: 3
  Neighbor is up for 00:01:26
  Authentication Sequence: [ 0 ]

                  Neighbors

Area 0.0.0.1 interface 10.0.35.3(Serial3/0/0)'s neighbors
Router ID: 10.0.5.5          Address: 10.0.35.5
  State: Full   Mode:Nbr is   Master   Priority: 1
  DR: None    BDR: None    MTU: 0
```

Dead timer due in 27 sec

Retrans timer interval: 4

Neighbor is up for 00:53:37

Authentication Sequence: [0]

After their interfaces are restarted, R3 becomes the DR and R2 becomes the BDR of the network segment 10.1.234.0/24.

Check the neighbor relationship between R4 and R1.

```
[R4]display ospf peer 10.0.1.1

         OSPF Process 1 with Router ID 10.0.4.4
                 Neighbors

 Area 0.0.0.0 interface 10.1.234.4(GigabitEthernet0/0/0)'s neighbors
 Router ID: 10.0.1.1          Address: 10.1.234.1
   State: 2-Way   Mode:Nbr is   Slave   Priority: 1
   DR: 10.1.234.3   BDR: 10.1.234.2   MTU: 0
   Dead timer due in 30   sec
   Retrans timer interval: 0
   Neighbor is up for 00:00:00
   Authentication Sequence: [ 0 ]
```

After their neighbor relationship becomes stable, because R1 and R4 are both DR others, they only establish a neighbor relationship in 2-way state.

Step 4 **Summarize direct routes and import summarized routes into OSPF areas.**

Loopback0 of R5 does not belong to any OSPF area. Import the network segment where Loopback0 resides into an OSPF area.

```
[R5]ospf 1
[R5-ospf-1]import-route direct
[R5-ospf-1]quit
```

Check the imported external route on R1 and R3.

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 6        Routes : 6

OSPF routing table status : <Active>
         Destinations : 6        Routes : 6

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        10.0.2.0/24  OSPF    10   1           D   10.1.234.2      GigabitEthernet0/0/0
        10.0.3.0/24  OSPF    10   1           D   10.1.234.3      GigabitEthernet0/0/0
        10.0.4.0/24  OSPF    10   1           D   10.1.234.4      GigabitEthernet0/0/0
        10.0.5.0/24  O_ASE   150  1           D   10.1.234.3      GigabitEthernet0/0/0
       10.0.35.0/24  OSPF    10   1563        D   10.1.234.3      GigabitEthernet0/0/0
       10.0.35.3/32  O_ASE   150  1           D   10.1.234.3      GigabitEthernet0/0/0

OSPF routing table status : <Inactive>
         Destinations : 0        Routes : 0
```

```
[R3]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 5          Routes : 5
OSPF routing table status : <Active>
         Destinations : 4          Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.0/24  OSPF    10   1           D   10.1.234.1      GigabitEthernet0/0/0
       10.0.2.0/24  OSPF    10   1           D   10.1.234.2      GigabitEthernet0/0/0
       10.0.4.0/24  OSPF    10   1           D   10.1.234.4      GigabitEthernet0/0/0
       10.0.5.0/24  O_ASE   150  1           D   10.0.35.5         Serial3/0/0

OSPF routing table status : <Inactive>
         Destinations : 1          Routes : 1

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.35.3/32  O_ASE   150  1               10.0.35.5       Serial3/0/0
```

You can see two external routes on both R1 and R3: 10.0.5.0/24 and 10.0.35.3/32. 10.0.5.0/24 is the loopback interface address of R5. Why there is still a route to 10.0.35.3/32?

Check the IP routing table of R5. PPP encapsulation is used between R3 and R5, therefore, the Serial3/0/0 address of R3 is displayed in the IP routing table of R5 as a direct route. After the **import-route direct** command is run on R5, this direct route is also advertised. Other routing entries are not displayed in the following command output.

```
[R5]display   ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 16          Routes : 16

Destination/Mask     Proto    Pre   Cost       Flags NextHop          Interface
      10.0.35.0/24  Direct   0      0             D     10.0.35.5        Serial1/0/0
      10.0.35.3/32  Direct   0      0             D     10.0.35.3        Serial1/0/0
      10.0.35.5/32  Direct   0      0             D     127.0.0.1        InLoopBack0
    10.0.35.255/32  Direct   0      0             D     127.0.0.1        InLoopBack0
```

Test network connectivity.

```
[R1]ping -c 1 10.0.5.5
  PING 10.0.5.5: 56   data bytes, press CTRL_C to break
    Reply from 10.0.5.5: bytes=56 Sequence=1 ttl=254 time=41 ms

  --- 10.0.5.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 41/41/41 ms
```

Check OSPF external routes in the LSDB of R1. You can see that the LSDB contains three external routes: 10.0.5.0/24, 10.0.35.0/24, and 10.0.35.3/32.

The routing table of R1 has only two external routes.

```
[R1]display ospf lsdb ase

         OSPF Process 1 with Router ID 10.0.1.1
                  Link State Database


 Type       : External
 Ls id      : 10.0.5.0
 Adv rtr    : 10.0.5.5
 Ls age     : 834
 Len        : 36
 Options    :  E
 seq#       : 80000001
 chksum     : 0xa904
 Net mask   : 255.255.255.0
 TOS 0  Metric: 1
 E type     : 2
 Forwarding Address : 0.0.0.0
 Tag        : 1
 Priority   : Low

 Type       : External
 Ls id      : 10.0.35.0
 Adv rtr    : 10.0.5.5
 Ls age     : 1342
 Len        : 36
 Options    :  E
 seq#       : 80000001
 chksum     : 0x5e31
 Net mask   : 255.255.255.0
 TOS 0  Metric: 1
```

```
E type      : 2
Forwarding Address : 0.0.0.0
Tag          : 1
Priority   : Low

Type         : External
Ls id        : 10.0.35.3
Adv rtr      : 10.0.5.5
Ls age       : 1344
Len          : 36
Options      :   E
seq#         : 80000001
chksum       : 0x404c
Net mask     : 255.255.255.255
TOS 0   Metric: 1
E type       : 2
Forwarding Address : 0.0.0.0
Tag          : 1
Priority   : Medium
```

After comparison, you will find that the route 10.0.35.0/24 is displayed as an internal route in the routing table of R1.

Check Type 3 LSAs in the LSDB of R1, and you can see the route 10.0.35.0/24.

```
[R1]display ospf lsdb summary 10.0.35.0

         OSPF Process 1 with Router ID 10.0.1.1
                          Area: 0.0.0.0
                  Link State Database
```

Type : Sum-Net

Ls id : 10.0.35.0

Adv rtr : 10.0.3.3

Ls age : 236

Len : 28

Options : E

seq# : 80000007

chksum : 0x14e5

Net mask : 255.255.255.0

Tos 0 metric: 1562

Priority : Low

Area: 0.0.0.2

Link State Database

Type : Sum-Net

Ls id : 10.0.35.0

Adv rtr : 10.0.1.1

Ls age : 1637

Len : 28

Options : E

seq# : 80000002

chksum : 0x42bf

Net mask : 255.255.255.0

Tos 0 metric: 1563

Priority : Low

When the network bits and mask of the routes advertised by Type 3 and Type 5 LSAs are the same, OSPF prefers and adds the route advertised by a Type 3 LSA into its routing table.

Step 5 **Check various types of LSAs.**

On R1, check detailed information about Type 1 LSA 10.0.1.0 in Area 0 and Area 2.

```
[R1]display ospf lsdb router 10.0.1.1

         OSPF Process 1 with Router ID 10.0.1.1
                         Area: 0.0.0.0
                 Link State Database

  Type        : Router
  Ls id       : 10.0.1.1
  Adv rtr : 10.0.1.1
  Ls age      : 591
  Len         : 36
  Options     :   ABR   E
  seq#        : 8000001e
  chksum          : 0xbc70
  Link count: 1
   * Link ID: 10.1.234.3
     Data     : 10.1.234.1
     Link Type: TransNet
     Metric : 1
                         Area: 0.0.0.2
                 Link State Database

  Type        : Router
```

```
Ls id       : 10.0.1.1
Adv rtr : 10.0.1.1
Ls age      : 627
Len         : 36
Options     :   ABR   E
seq#        : 80000008
chksum          : 0x1018
Link count: 1
 * Link ID: 10.0.1.0
   Date     : 255.255.255.0
   Link Type: StubNet
   Metric   : 0
   Priority: Low
```

For a Type 1 LSA, the **Ls id** field indicates the router ID of the router that generates this LSA.

R1 generates two Type 1 LSAs and floods one within Area 0. In Area 0, R1 is connected to a transit network segment. Therefore, the **Link Type** field displays **TransNet**. For TransNet, the **Link ID** field indicates the interface IP address of the DR on this network segment, and the **Data** field indicates the local interface IP address.

R1 floods the second Type 1 LSA within Area 2 and is connected to Area 2 through loopback interfaces. For a loopback interface, the **Link Type** field displays **StubNet**. The **Link ID** field indicates the IP network address of this stub network segment, and the **Data** field indicates the network mask of this stub network segment.

On R2, R3, and R4, check detailed information about Type 2 LSA 10.1.234.0 in Area 0.

```
[R2]display ospf lsdb network 10.1.234.3

         OSPF Process 1 with Router ID 10.0.2.2
```

```
                    Area: 0.0.0.0
          Link State Database


Type        : Network
Ls id       : 10.1.234.3
Adv rtr     : 10.0.3.3
Ls age      : 115
Len         : 40
Options     :  E
seq#        : 8000000f
chksum      : 0x807e
Net mask    : 255.255.255.0
Priority    : Low
   Attached Router    10.0.3.3
   Attached Router    10.0.1.1
   Attached Router    10.0.2.2
   Attached Router    10.0.4.4
```

You can see that this LSA is the same on R2, R3, and R4.

You can also know that this LSA is generated by R3 according to the **Adv rtr** field. The **Ls id** field of a Type 2 LSA indicates the interface IP address of the DR on this network segment, and the **Attached Router** field indicates the router IDs of all the routers on this network segment.

On R1 and R3, check detailed information about Type 3 LSA 10.0.35.0/24 in Area 0.

```
[R3]display ospf lsdb summary 10.0.35.0


       OSPF Process 1 with Router ID 10.0.3.3
                    Area: 0.0.0.0
```

```
                Link State Database

Type         : Sum-Net
Ls id        : 10.0.35.0
Adv rtr      : 10.0.3.3
Ls age       : 591
Len          : 28
Options      :   E
seq#         : 8000000a
chksum       : 0xee8
Net mask     : 255.255.255.0
Tos 0    metric: 1562
Priority     : Low
```

The preceding command output shows that this route is advertised by R3 within Area 0. The **Ls id** field indicates the network address of the advertised destination network segment, and the **Net mask** field indicates the mask of the destination network segment.

```
[R1]display ospf lsdb summary 10.0.35.0

         OSPF Process 1 with Router ID 10.0.1.1
                        Area: 0.0.0.0
                Link State Database

Type         : Sum-Net
Ls id        : 10.0.35.0
Adv rtr      : 10.0.3.3
Ls age       : 136
Len          : 28
```

```
Options    :  E
seq#       : 80000004
chksum      : 0x1ae2
Net mask   : 255.255.255.0
Tos 0   metric: 1562
Priority   : Low
                 Area: 0.0.0.2
            Link State Database


Type        : Sum-Net
Ls id      : 10.0.35.0
Adv rtr    : 10.0.1.1
Ls age     : 382
Len        : 28
Options    :  E
seq#       : 80000002
chksum      : 0x42bf
Net mask   : 255.255.255.0
Tos 0   metric: 1563
Priority   : Low
```

R1 has a total of two Type 3 LSAs 10.0.35.0/24. The **Adv rtr** field indicates that this LSA in Area 0 is generated by R3. R1 is an ABR, so it generates another LSA after receiving this LSA and advertises it within Area 2.

On R1, check detailed information about Type 4 LSA 10.0.5.0 in Area 2. A Type 4 LSA describes how to reach an ASBR.

```
[R1]display ospf lsdb asbr 10.0.5.5
```

```
        OSPF Process 1 with Router ID 10.0.1.1
                        Area: 0.0.0.0
                Link State Database

Type         : Sum-Asbr
Ls id        : 10.0.5.5
Adv rtr      : 10.0.3.3
Ls age       : 1119
Len          : 28
Options      :   E
seq#         : 80000008
chksum       : 0x1df3
Tos 0   metric: 1562
                        Area: 0.0.0.2
                Link State Database

Type         : Sum-Asbr
Ls id        : 10.0.5.5
Adv rtr      : 10.0.1.1
Ls age       : 1118
Len          : 28
Options      :   E
seq#         : 80000008
chksum       : 0x41d2
Tos 0   metric: 1563
```

The preceding command output shows that R1 receives a Type 4 LSA from R3. The **Ls id** field indicates the router ID of an ASBR. This LSA cannot be flooded across areas. Therefore, R1 generates another Type 4 LSA and floods it within Area 2.

On R2, R4, and R3, this LSA exists in the LSDB for Area 0. These routers do not belong to the same area as the ASBR (R5), so they need to know the location of this ASBR through a Type 4 LSA.

```
[R2]display ospf lsdb asbr

	 OSPF Process 1 with Router ID 10.0.2.2
		         Area: 0.0.0.0
		     Link State Database

  Type      : Sum-Asbr
  Ls id     : 10.0.5.5
  Adv rtr   : 10.0.3.3
  Ls age    : 1676
  Len       : 28
  Options   :  E
  seq#      : 80000008
  chksum    : 0x1df3
  Tos 0  metric: 1562
```

The preceding command output shows that this Type 4 LSA does not exist in Area 1. Routes in the same area do not need to know the ASBR location through this LSA.

Step 6 **Check LSR, LSU, and LSAck packets.**

Check the transmission of LSU and LSAck packets. Run the **debugging ospf packet update** and **debugging ospf packet ack** commands on R1.

```
<R1>terminal monitor
```

Info: Current terminal monitor is on

<R1>terminal debugging

Info: Current terminal debugging is on

<R1>debugging ospf packet update

<R1>debugging ospf packet ack

By default, when the network is running stably, an OSPF router updates its LSDB at an interval of 30 minutes. To trigger routing information query and update, delete Loopback0 of R3.

[R3]undo interface LoopBack 0

Info: This operation may take a few seconds. Please wait for a moment...succeeded.

[R3]

Oct 25 2016 15:32:27+00:00 R3 %%01IFNET/4/LINK_STATE(l)[58]:The line protocol IP on the interface LoopBack0 has entered the DOWN state

You can see that R1 receives an LSU packet sent from 10.1.234.3. The destination address of the packet is 224.0.0.5 (namely all OSPF routers), describing a network segment (# Links: 1), followed by the Link ID and Link Data of this network segment.

<R1>

Oct 25 2016 15:24:57.790.1+00:00 R1 RM/6/RMDEBUG:

FileID: 0xd0178024 Line: 2271 Level: 0x20

OSPF 1: RECV Packet. Interface: GigabitEthernet0/0/0

<R1>

Oct 25 2016 15:24:57.790.2+00:00 R1 RM/6/RMDEBUG: Source Address: 10.1.234.3

Oct 25 2016 15:24:57.790.3+00:00 R1 RM/6/RMDEBUG: Destination Address: 224.0.0.5

Oct 25 2016 15:24:57.790.4+00:00 R1 RM/6/RMDEBUG: Ver# 2, Type: 4 (Link-State Update)

Oct 25 2016 15:24:57.790.5+00:00 R1 RM/6/RMDEBUG: Length: 64, Router: 10.0.3.3

Oct 25 2016 15:24:57.790.6+00:00 R1 RM/6/RMDEBUG: Area: 0.0.0.0, Chksum: d8ce

Oct 25 2016 15:24:57.790.7+00:00 R1 RM/6/RMDEBUG: AuType: 00

Oct 25 2016 15:24:57.790.8+00:00 R1 RM/6/RMDEBUG: Key(ascii): * * * * * * * *

Oct 25 2016 15:24:57.790.9+00:00 R1 RM/6/RMDEBUG: # LSAS: 1

Oct 25 2016 15:24:57.790.10+00:00 R1 RM/6/RMDEBUG: LSA Type 1

Oct 25 2016 15:24:57.790.11+00:00 R1 RM/6/RMDEBUG: LS ID: 10.0.3.3

Oct 25 2016 15:24:57.790.12+00:00 R1 RM/6/RMDEBUG: Adv Rtr: 10.0.3.3

Oct 25 2016 15:24:57.790.13+00:00 R1 RM/6/RMDEBUG: LSA Age: 1

Oct 25 2016 15:24:57.790.14+00:00 R1 RM/6/RMDEBUG: Options: ExRouting:ON

Oct 25 2016 15:24:57.790.15+00:00 R1 RM/6/RMDEBUG: Length: 36, Seq# 80000020

Oct 25 2016 15:24:57.790.16+00:00 R1 RM/6/RMDEBUG: CheckSum: 9090

Oct 25 2016 15:24:57.790.17+00:00 R1 RM/6/RMDEBUG: NtBit: 0 VBit: 0 EBit: 0 BBit: 1

Oct 25 2016 15:24:57.790.18+00:00 R1 RM/6/RMDEBUG: # Links: 1

Oct 25 2016 15:24:57.790.19+00:00 R1 RM/6/RMDEBUG: LinkID: 10.1.234.3

Oct 25 2016 15:24:57.790.20+00:00 R1 RM/6/RMDEBUG: LinkData: 10.1.234.3

Oct 25 2016 15:24:57.790.21+00:00 R1 RM/6/RMDEBUG: LinkType: 2

Oct 25 2016 15:24:57.790.22+00:00 R1 RM/6/RMDEBUG: TOS# 0 Metric 1

Then check the LSAck packet sent by R1. The source address is GigabitEthernet0/0/0 address of R1 and the destination address is 224.0.0.6. This packet is sent to the DR and BDR. The sequence number of this packet is also 80000020.

<R1>

Oct 25 2016 15:24:58.200.1+00:00 R1 RM/6/RMDEBUG:

FileID: 0xd0178025 Line: 4708 Level: 0x20

OSPF 1: SEND Packet. Interface: GigabitEthernet0/0/0

<R1>

Oct 25 2016 15:24:58.200.2+00:00 R1 RM/6/RMDEBUG: Source Address: 10.1.234.1

Oct 25 2016 15:24:58.200.3+00:00 R1 RM/6/RMDEBUG: Destination Address: 224.0.0.6

Oct 25 2016 15:24:58.200.4+00:00 R1 RM/6/RMDEBUG: Ver# 2, Type: 5 (Link-State Ack)

Oct 25 2016 15:24:58.200.5+00:00 R1 RM/6/RMDEBUG: Length: 44, Router: 10.0.1.1

Oct 25 2016 15:24:58.200.6+00:00 R1 RM/6/RMDEBUG: Area: 0.0.0.0, Chksum: c5ef

Oct 25 2016 15:24:58.200.7+00:00 R1 RM/6/RMDEBUG: AuType: 00

Oct 25 2016 15:24:58.200.8+00:00 R1 RM/6/RMDEBUG: Key(ascii): * * * * * * * *

Oct 25 2016 15:24:58.200.9+00:00 R1 RM/6/RMDEBUG: # LSA Headers: 1

Oct 25 2016 15:24:58.200.10+00:00 R1 RM/6/RMDEBUG: LSA Type 1

Oct 25 2016 15:24:58.200.11+00:00 R1 RM/6/RMDEBUG: LS ID: 10.0.3.3

Oct 25 2016 15:24:58.200.12+00:00 R1 RM/6/RMDEBUG: Adv Rtr: 10.0.3.3

Oct 25 2016 15:24:58.200.13+00:00 R1 RM/6/RMDEBUG: LSA Age: 2

Oct 25 2016 15:24:58.200.14+00:00 R1 RM/6/RMDEBUG: Options: ExRouting:ON

Oct 25 2016 15:24:58.200.15+00:00 R1 RM/6/RMDEBUG: Length: 36, Seq# 80000020

Oct 25 2016 15:24:58.200.16+00:00 R1 RM/6/RMDEBUG: CheckSum: 9090

Restore Loopback0 of R3.

[R3]interface loopback 0

[R3-LoopBack0]ip address 10.0.3.3 24

[R3-LoopBack0]quit

R1 also receives an LSU packet from R3. However, this packet advertises a new network segment. Therefore, # Links displays 2, followed by the network ID and mask of the new network segment.

<R1>

Oct 25 2016 15:51:26.250.1+00:00 R1 RM/6/RMDEBUG:

FileID: 0xd0178024 Line: 2271 Level: 0x20

OSPF 1: RECV Packet. Interface: GigabitEthernet0/0/0

<R1>

Oct 25 2016 15:51:26.250.2+00:00 R1 RM/6/RMDEBUG: Source Address: 10.1.234.3

Oct 25 2016 15:51:26.250.3+00:00 R1 RM/6/RMDEBUG: Destination Address: 224.0.0.5

Oct 25 2016 15:51:26.250.4+00:00 R1 RM/6/RMDEBUG: Ver# 2, Type: 4 (Link-State Update)

Oct 25 2016 15:51:26.250.5+00:00 R1 RM/6/RMDEBUG: Length: 76, Router: 10.0.3.3

Oct 25 2016 15:51:26.250.6+00:00 R1 RM/6/RMDEBUG: Area: 0.0.0.0, Chksum: 2c6f

Oct 25 2016 15:51:26.250.7+00:00 R1 RM/6/RMDEBUG: AuType: 00

Oct 25 2016 15:51:26.250.8+00:00 R1 RM/6/RMDEBUG: Key(ascii): * * * * * * * *

Oct 25 2016 15:51:26.250.9+00:00 R1 RM/6/RMDEBUG: # LSAS: 1

Oct 25 2016 15:51:26.250.10+00:00 R1 RM/6/RMDEBUG: LSA Type 1

Oct 25 2016 15:51:26.250.11+00:00 R1 RM/6/RMDEBUG: LS ID: 10.0.3.3

Oct 25 2016 15:51:26.250.12+00:00 R1 RM/6/RMDEBUG: Adv Rtr: 10.0.3.3

Oct 25 2016 15:51:26.250.13+00:00 R1 RM/6/RMDEBUG: LSA Age: 1

Oct 25 2016 15:51:26.250.14+00:00 R1 RM/6/RMDEBUG: Options: ExRouting:ON

Oct 25 2016 15:51:26.250.15+00:00 R1 RM/6/RMDEBUG: Length: 48, Seq# 8000002a

Oct 25 2016 15:51:26.250.16+00:00 R1 RM/6/RMDEBUG: CheckSum: 2cca

Oct 25 2016 15:51:26.250.17+00:00 R1 RM/6/RMDEBUG: NtBit: 0 VBit: 0 EBit: 0 BBit: 1

Oct 25 2016 15:51:26.250.18+00:00 R1 RM/6/RMDEBUG: # Links: 2

Oct 25 2016 15:51:26.250.19+00:00 R1 RM/6/RMDEBUG: LinkID: 10.1.234.3

Oct 25 2016 15:51:26.250.20+00:00 R1 RM/6/RMDEBUG: LinkData: 10.1.234.3

Oct 25 2016 15:51:26.250.21+00:00 R1 RM/6/RMDEBUG: LinkType: 2

Oct 25 2016 15:51:26.250.22+00:00 R1 RM/6/RMDEBUG: TOS# 0 Metric 1

Oct 25 2016 15:51:26.250.23+00:00 R1 RM/6/RMDEBUG: LinkID: 10.0.3.3

Oct 25 2016 15:51:26.250.24+00:00 R1 RM/6/RMDEBUG: LinkData: 255.255.255.255

Oct 25 2016 15:51:26.250.25+00:00 R1 RM/6/RMDEBUG: LinkType: 3

Oct 25 2016 15:51:26.250.26+00:00 R1 RM/6/RMDEBUG: TOS# 0 Metric 0

R1 first receives the LSAck packet of the BDR.

<R1>

Oct 25 2016 15:51:27.90.1+00:00 R1 RM/6/RMDEBUG:

FileID: 0xd0178024 Line: 2271 Level: 0x20

```
 OSPF 1: RECV Packet. Interface: GigabitEthernet0/0/0
<R1>
Oct 25 2016 15:51:27.90.2+00:00 R1 RM/6/RMDEBUG:   Source Address: 10.1.234.2
Oct 25 2016 15:51:27.90.3+00:00 R1 RM/6/RMDEBUG:   Destination Address: 224.0.0.5
Oct 25 2016 15:51:27.90.4+00:00 R1 RM/6/RMDEBUG:   Ver# 2, Type: 5 (Link-State Ack)
Oct 25 2016 15:51:27.90.5+00:00 R1 RM/6/RMDEBUG:   Length: 44, Router: 10.0.2.2
Oct 25 2016 15:51:27.90.6+00:00 R1 RM/6/RMDEBUG:   Area: 0.0.0.0, Chksum: 289f
Oct 25 2016 15:51:27.90.7+00:00 R1 RM/6/RMDEBUG:   AuType: 00
Oct 25 2016 15:51:27.90.8+00:00 R1 RM/6/RMDEBUG:   Key(ascii): * * * * * * * *
Oct 25 2016 15:51:27.90.9+00:00 R1 RM/6/RMDEBUG:   # LSA Headers: 1
Oct 25 2016 15:51:27.90.10+00:00 R1 RM/6/RMDEBUG:   LSA Type 1
Oct 25 2016 15:51:27.90.11+00:00 R1 RM/6/RMDEBUG:      LS ID: 10.0.3.3
Oct 25 2016 15:51:27.90.12+00:00 R1 RM/6/RMDEBUG:      Adv Rtr: 10.0.3.3
Oct 25 2016 15:51:27.90.13+00:00 R1 RM/6/RMDEBUG:      LSA Age: 2
Oct 25 2016 15:51:27.90.14+00:00 R1 RM/6/RMDEBUG:      Options: ExRouting:ON
Oct 25 2016 15:51:27.90.15+00:00 R1 RM/6/RMDEBUG:      Length: 48, Seq# 8000002a
Oct 25 2016 15:51:27.90.16+00:00 R1 RM/6/RMDEBUG:      CheckSum: 2cca
```

Then check the LSAck packet sent by R1.

```
<R1>
Oct 25 2016 15:51:26.430.1+00:00 R1 RM/6/RMDEBUG:
 FileID: 0xd0178025 Line: 4708 Level: 0x20
 OSPF 1: SEND Packet. Interface: GigabitEthernet0/0/0
<R1>
Oct 25 2016 15:51:26.430.2+00:00 R1 RM/6RMDEBUG:   Source Address: 10.1.234.1
Oct 25 2016 15:51:26.430.3+00:00 R1 RM/6/RMDEBUG:   Destination Address: 224.0.0.6
Oct 25 2016 15:51:26.430.4+00:00 R1 RM/6/RMDEBUG:   Ver# 2, Type: 5 (Link-State Ack)
Oct 25 2016 15:51:26.430.5+00:00 R1 RM/6/RMDEBUG:   Length: 44, Router: 10.0.1.1
```

Oct 25 2016 15:51:26.430.6+00:00 R1 RM/6/RMDEBUG: Area: 0.0.0.0, Chksum: 29a1

Oct 25 2016 15:51:26.430.7+00:00 R1 RM/6/RMDEBUG: AuType: 00

Oct 25 2016 15:51:26.430.8+00:00 R1 RM/6/RMDEBUG: Key(ascii): * * * * * * * *

Oct 25 2016 15:51:26.430.9+00:00 R1 RM/6/RMDEBUG: # LSA Headers: 1

Oct 25 2016 15:51:26.430.10+00:00 R1 RM/6/RMDEBUG: LSA Type 1

Oct 25 2016 15:51:26.430.11+00:00 R1 RM/6/RMDEBUG: LS ID: 10.0.3.3

Oct 25 2016 15:51:26.430.12+00:00 R1 RM/6/RMDEBUG: Adv Rtr: 10.0.3.3

Oct 25 2016 15:51:26.430.13+00:00 R1 RM/6/RMDEBUG: LSA Age: 1

Oct 25 2016 15:51:26.430.14+00:00 R1 RM/6/RMDEBUG: Options: ExRouting:ON

Oct 25 2016 15:51:26.430.15+00:00 R1 RM/6/RMDEBUG: Length: 48, Seq# 8000002a

Oct 25 2016 15:51:26.430.16+00:00 R1 RM/6/RMDEBUG: CheckSum: 2cca

Next, check LSR packets. Normally, routers do not proactively send LSR packets. To check LSR packet transmission, restart the OSPF process of R1. You can see that R1 initiates an LSR packet to R2.

<R1>terminal monitor

Info: Current terminal monitor is on

<R1>terminal debugging

Info: Current terminal debugging is on

<R1>debugging ospf packet update

<R1>debugging ospf packet ack

<R1>debugging ospf packet request

<R1>reset ospf process

Warning: The OSPF process will be reset. Continue? [Y/N]:y

<R1>

Oct 25 2016 16:17:59.750.1+00:00 R1 RM/6/RMDEBUG:

FileID: 0xd0178025 Line: 2993 Level: 0x20

```
 OSPF 1: SEND Packet. Interface: GigabitEthernet0/0/0
<R1>
Oct 25 2016 16:17:59.750.2+00:00 R1 RM/6/RMDEBUG:   Source Address: 10.1.234.1
Oct 25 2016 16:17:59.750.3+00:00 R1 RM/6/RMDEBUG:   Destination Address: 10.1.234.2
Oct 25 2016 16:17:59.750.4+00:00 R1 RM/6/RMDEBUG:   Ver# 2, Type: 3 (Link-State Req)
Oct 25 2016 16:17:59.750.5+00:00 R1 RM/6/RMDEBUG:   Length: 156, Router: 10.0.1.1
Oct 25 2016 16:17:59.750.6+00:00 R1 RM/6/RMDEBUG:   Area: 0.0.0.0, Chksum: 8b05
Oct 25 2016 16:17:59.750.7+00:00 R1 RM/6/RMDEBUG:   AuType: 00
Oct 25 2016 16:17:59.750.8+00:00 R1 RM/6/RMDEBUG:   Key(ascii): * * * * * * * *
Oct 25 2016 16:17:59.750.9+00:00 R1 RM/6/RMDEBUG:   # Requesting LSAs: 11
Oct 25 2016 16:17:59.750.10+00:00 R1 RM/6/RMDEBUG:   LSA Type 1
Oct 25 2016 16:17:59.750.11+00:00 R1 RM/6/RMDEBUG:      LS ID: 10.0.2.2
Oct 25 2016 16:17:59.750.12+00:00 R1 RM/6/RMDEBUG:      Adv Rtr: 10.0.2.2
Oct 25 2016 16:17:59.750.13+00:00 R1 RM/6/RMDEBUG:   LSA Type 1
Oct 25 2016 16:17:59.750.14+00:00 R1 RM/6/RMDEBUG:      LS ID: 10.0.1.1
Oct 25 2016 16:17:59.750.15+00:00 R1 RM/6/RMDEBUG:      Adv Rtr: 10.0.1.1
Oct 25 2016 16:17:59.750.16+00:00 R1 RM/6/RMDEBUG:   LSA Type 1
Oct 25 2016 16:17:59.750.17+00:00 R1 RM/6/RMDEBUG:      LS ID: 10.0.4.4
Oct 25 2016 16:17:59.750.18+00:00 R1 RM/6/RMDEBUG:      Adv Rtr: 10.0.4.4
Oct 25 2016 16:17:59.750.19+00:00 R1 RM/6/RMDEBUG:   LSA Type 1
Oct 25 2016 16:17:59.750.20+00:00 R1 RM/6/RMDEBUG:      LS ID: 10.0.3.3
Oct 25 2016 16:17:59.750.21+00:00 R1 RM/6/RMDEBUG:      Adv Rtr: 10.0.3.3
Oct 25 2016 16:17:59.750.22+00:00 R1 RM/6/RMDEBUG:   LSA Type 2
Oct 25 2016 16:17:59.750.23+00:00 R1 RM/6/RMDEBUG:      LS ID: 10.1.234.3
Oct 25 2016 16:17:59.750.24+00:00 R1 RM/6/RMDEBUG:      Adv Rtr: 10.0.3.3
Oct 25 2016 16:17:59.750.25+00:00 R1 RM/6/RMDEBUG:   LSA Type 3
Oct 25 2016 16:17:59.750.26+00:00 R1 RM/6/RMDEBUG:      LS ID: 10.0.1.0
Oct 25 2016 16:17:59.750.27+00:00 R1 RM/6/RMDEBUG:      Adv Rtr: 10.0.1.1
Oct 25 2016 16:17:59.750.28+00:00 R1 RM/6/RMDEBUG:   LSA Type 3
```

Oct 25 2016 16:17:59.750.29+00:00 R1 RM/6/RMDEBUG: LS ID: 10.0.35.0

Oct 25 2016 16:17:59.750.30+00:00 R1 RM/6/RMDEBUG: Adv Rtr: 10.0.3.3

Oct 25 2016 16:17:59.750.31+00:00 R1 RM/6/RMDEBUG: LSA Type 4

Oct 25 2016 16:17:59.750.32+00:00 R1 RM/6/RMDEBUG: LS ID: 10.0.5.5

Oct 25 2016 16:17:59.750.33+00:00 R1 RM/6/RMDEBUG: Adv Rtr: 10.0.3.3

Oct 25 2016 16:17:59.750.34+00:00 R1 RM/6/RMDEBUG: LSA Type 5

Oct 25 2016 16:17:59.750.35+00:00 R1 RM/6/RMDEBUG: LS ID: 10.0.5.0

Oct 25 2016 16:17:59.750.36+00:00 R1 RM/6/RMDEBUG: Adv Rtr: 10.0.5.5

Oct 25 2016 16:17:59.750.37+00:00 R1 RM/6/RMDEBUG: LSA Type 5

Oct 25 2016 16:17:59.750.38+00:00 R1 RM/6/RMDEBUG: LS ID: 10.0.35.0

Oct 25 2016 16:17:59.750.39+00:00 R1 RM/6/RMDEBUG: Adv Rtr: 10.0.5.5

Oct 25 2016 16:17:59.750.40+00:00 R1 RM/6/RMDEBUG: LSA Type 5

Oct 25 2016 16:17:59.750.41+00:00 R1 RM/6/RMDEBUG: LS ID: 10.0.35.3

Oct 25 2016 16:17:59.750.42+00:00 R1 RM/6/RMDEBUG: Adv Rtr: 10.0.5.5

R1 then receives the LSR packet of R3.

<R1>

Oct 25 2016 16:30:10.80.1+00:00 R1 RM/6/RMDEBUG:

FileID: 0xd0178024 Line: 2271 Level: 0x20

OSPF 1: RECV Packet. Interface: GigabitEthernet0/0/0

<R1>

Oct 25 2016 16:30:10.80.2+00:00 R1 RM/6/RMDEBUG: Source Address: 10.1.234.3

Oct 25 2016 16:30:10.80.3+00:00 R1 RM/6/RMDEBUG: Destination Address: 10.1.234.1

Oct 25 2016 16:30:10.80.4+00:00 R1 RM/6/RMDEBUG: Ver# 2, Type: 3 (Link-State Req)

Oct 25 2016 16:30:10.80.5+00:00 R1 RM/6/RMDEBUG: Length: 48, Router: 10.0.3.3

Oct 25 2016 16:30:10.80.6+00:00 R1 RM/6/RMDEBUG: Area: 0.0.0.0, Chksum: c4c2

Oct 25 2016 16:30:10.80.7+00:00 R1 RM/6/RMDEBUG: AuType: 00

Oct 25 2016 16:30:10.80.8+00:00 R1 RM/6/RMDEBUG: Key(ascii): * * * * * * * *

```
Oct 25 2016 16:30:10.80.9+00:00 R1 RM/6/RMDEBUG:   # Requesting LSAs: 2
Oct 25 2016 16:30:10.80.10+00:00 R1 RM/6/RMDEBUG:   LSA Type 1
Oct 25 2016 16:30:10.80.11+00:00 R1 RM/6/RMDEBUG:      LS ID: 10.0.1.1
Oct 25 2016 16:30:10.80.12+00:00 R1 RM/6/RMDEBUG:      Adv Rtr: 10.0.1.1
Oct 25 2016 16:30:10.80.13+00:00 R1 RM/6/RMDEBUG:   LSA Type 3
Oct 25 2016 16:30:10.80.14+00:00 R1 RM/6/RMDEBUG:      LS ID: 10.0.1.0
Oct 25 2016 16:30:10.80.15+00:00 R1 RM/6/RMDEBUG:      Adv Rtr: 10.0.1.1
```

----End

Additional Exercises: Analysis and Verification

Assume that there is a router R6 in Area 2. What are the differences between the procedure for calculating the routes to the network segment 10.0.5.0/24 on R6 and that on R2 and R3?

When will Type 4 LSAs appear?

If both R1 and R4 are configured as DR others, what are the potential problems?

Device Configurations

```
<R1>display current-configuration
[V200R007C00SPC600]
#
 sysname R1
#
interface GigabitEthernet0/0/0
 ip address 10.1.234.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.0
 ospf network-type broadcast
#
```

```
ospf 1 router-id 10.0.1.1
 area 0.0.0.0
  network 10.1.234.1 0.0.0.0
 area 0.0.0.2
  network 10.0.1.1 0.0.0.0
#
return

<R2>display current-configuration
[V200R007C00SPC600]
#
 sysname R2
#
interface GigabitEthernet0/0/0
 ip address 10.1.234.2 255.255.255.0
 ospf dr-priority 254
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.2.2
 area 0.0.0.0
  network 10.1.234.2 0.0.0.0
  network 10.0.2.2 0.0.0.0
#
return

<R3>display current-configuration
```

```
[V200R007C00SPC600]
#
 sysname R3
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.1.234.3 255.255.255.0
 ospf dr-priority 255
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.3.3
 area 0.0.0.0
  network 10.1.234.3 0.0.0.0
  network 10.0.3.3 0.0.0.0
 area 0.0.0.1
  network 10.0.35.3 0.0.0.0
#
return

<R4>display current-configuration
[V200R007C00SPC600]
#
 sysname R4
```

```
#
interface GigabitEthernet0/0/0
 ip address 10.1.234.4 255.255.255.0
 ospf dr-priority 0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.4.4
 area 0.0.0.0
  network 10.1.234.4 0.0.0.0
  network 10.0.4.4 0.0.0.0
#
return

<R5>display current-configuration
[V200R007C00SPC600]
#
 sysname R5
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.0
#
ospf 1 router-id 10.0.5.5
```

```
 import-route direct
 area 0.0.0.1
  network 10.0.35.5 0.0.0.0
#
return
```

Lab 1-4 OSPF Stub Area and NSSA Area

Learning Objectives

The objectives of this lab are to learn and understand:

- How to configure a stub area
- How to configure an NSSA area
- How to check Type 7 LSAs
- Translation between Type 5 and Type 7 LSAs

Topology

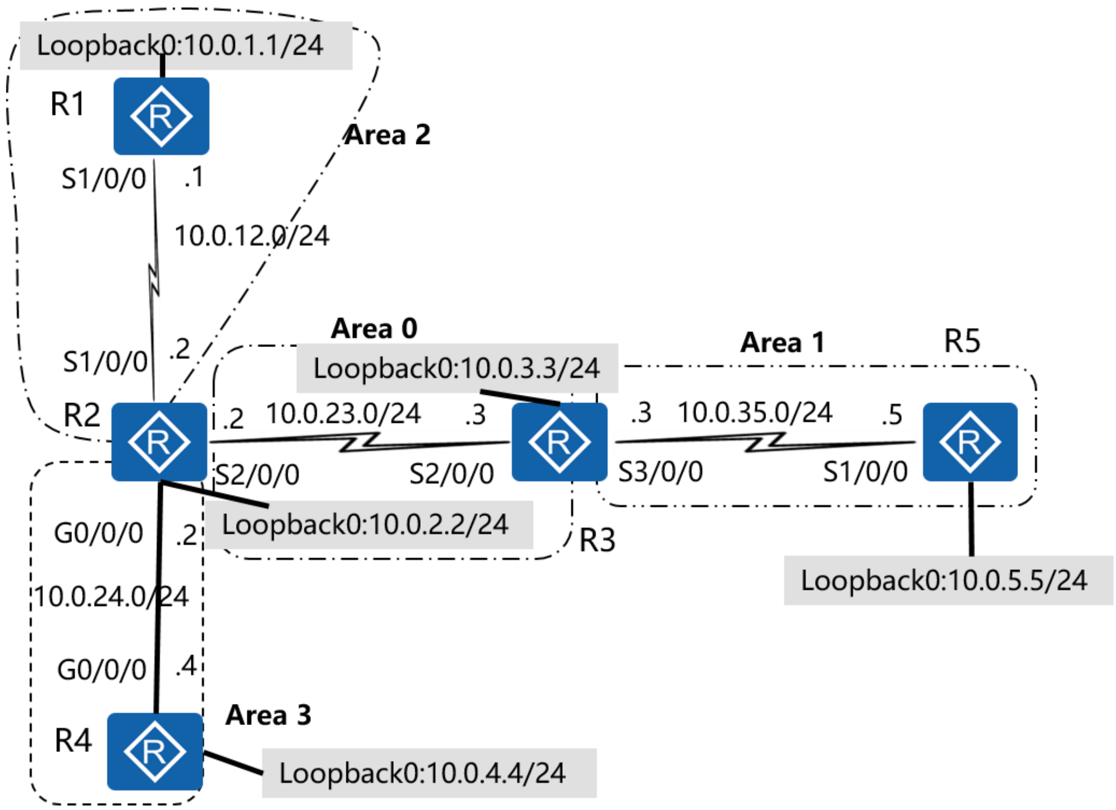

Figure 1-4 OSPF stub area and NSSA area

Scenario

You are a network administrator of a company. The company’s network has five AR G3 routers. R2, R3, and R4 are deployed in the headquarters. R5 is deployed in one branch. R5 is connected to R3 in the headquarters through a leased line. R1 is deployed in the other branch and is connected to R2 in the headquarters through a leased line.

Network segments 10.0.23.0/24, 10.0.2.0/24, and 10.0.3.0/24 belong to Area 0.

The network segment 10.0.35.0/24 belongs to Area 1, which is an NSSA area. Loopback0 of R5 does not belong to any OSPF area.

The network segment 10.0.24.0/24 belongs to Area 3. Loopback0 of R4 is connected to the Internet, requiring a default route to be configured.

Network segments 10.0.12.0/24 and 10.0.1.0/24 belong to Area 2, which is a stub area.

To specify router IDs for the routers, configure the routers to use fixed addresses as their router IDs.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

Configure IP addresses and masks for all the routers. Set a 24-bit mask for all loopback interfaces to simulate an independent network segment.

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 24
[R1-LoopBack0]quit
```

```
<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]quit
[R2]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 24
[R2-Serial2/0/0]quit
[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.24.2 24
[R2-GigabitEthernet0/0/0]quit
[R2]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24
[R2-LoopBack0]quit

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 24
[R3-Serial2/0/0]quit
[R3]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 24
[R3-Serial3/0/0]quit
[R3]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 24
[R3-LoopBack0]quit

<R4>system-view
```

```
Enter system view, return user view with Ctrl+Z.
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ip address 10.0.24.4 24
[R4-GigabitEthernet0/0/0]quit
[R4]interface LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 24
[R4-LoopBack0]quit

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface Serial 1/0/0
[R5-Serial1/0/0]ip address 10.0.35.5 24
[R5-Serial1/0/0]quit
[R5]interface LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 24
[R5-LoopBack0]quit
```

Test direct link connectivity.

```
[R2]ping -c 1 10.0.12.1
  PING 10.0.12.1: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.1: bytes=56 Sequence=1 ttl=255 time=30 ms

  --- 10.0.12.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 30/30/30 ms
```

```
[R2]ping -c 1 10.0.24.4
  PING 10.0.24.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.24.4: bytes=56 Sequence=1 ttl=255 time=6 ms

  --- 10.0.24.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 6/6/6 ms

[R2]ping -c 1 10.0.23.3
  PING 10.0.23.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.3: bytes=56 Sequence=1 ttl=255 time=31 ms

  --- 10.0.23.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 31/31/31 ms

[R3]ping -c 1 10.0.35.5
  PING 10.0.35.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.35.5: bytes=56 Sequence=1 ttl=255 time=38 ms

  --- 10.0.35.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 38/38/38 ms
```

Step 2 **Configure multiple OSPF areas.**

On R1, configure Serial1/0/0 and Loopback0 to belong to Area 2. To enable OSPF to advertise real masks of loopback interfaces, change the OSPF network type of loopback interfaces in all the areas to broadcast. Configure all routers to use IP address of Loopback0 as their router IDs.

```
[R1]ospf 1 router-id 10.0.1.1
[R1-ospf-1]area 2
[R1-ospf-1-area-0.0.0.2]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]network 10.0.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]quit
[R1-ospf-1]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ospf network-type broadcast
[R1-LoopBack0]quit
```

On R2, configure Serial2/0/0 and Loopback0 to belong to Area 0, Serial1/0/0 to belong to Area 2, and GigabitEthernet0/0/0 to belong to Area 3.

```
[R2]ospf 1 router-id 10.0.2.2
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.23.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]quit
[R2-ospf-1]area 2
[R2-ospf-1-area-0.0.0.2]network 10.0.12.2 0.0.0.0
[R2-ospf-1-area-0.0.0.2]quit
[R2-ospf-1]area 3
[R2-ospf-1-area-0.0.0.3]network 10.0.24.2 0.0.0.0
```

[R2-ospf-1-area-0.0.0.3]qui

[R2-ospf-1]quit

[R2]interface LoopBack 0

[R2-LoopBack0]ospf network-type broadcast

[R2-LoopBack0]quit

On R3, configure Serial2/0/0 and Loopback0 to belong to Area 0 and Serial3/0/0 to belong to Area 1.

[R3]ospf 1 router-id 10.0.3.3

[R3-ospf-1]area 0

[R3-ospf-1-area-0.0.0.0]network 10.0.23.3 0.0.0.0

[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0

[R3-ospf-1-area-0.0.0.0]quit

[R3-ospf-1]area 1

[R3-ospf-1-area-0.0.0.1]network 10.0.35.3 0.0.0.0

[R3-ospf-1-area-0.0.0.1]quit

[R3-ospf-1]quit

[R3]interface LoopBack 0

[R3-LoopBack0]ospf network-type broadcast

[R3-LoopBack0]quit

On R4, configure GigabitEthernet0/0/0 to belong to Area 3 and configure Loopback0 not to belong to any area.

[R4]ospf 1 router-id 10.0.4.4

[R4-ospf-1]area 3

[R4-ospf-1-area-0.0.0.3]network 10.0.24.4 0.0.0.0

[R4-ospf-1-area-0.0.0.3]quit

[R4-ospf-1]quit

On R5, configure Serial1/0/0 to belong to Area 1 and configure Looback0 not to belong to any area.

```
[R5]ospf 1 router-id 10.0.5.5

[R5-ospf-1]area 1

[R5-ospf-1-area-0.0.0.1]network 10.0.35.5 0.0.0.0

[R5-ospf-1-area-0.0.0.1]quit

[R5-ospf-1]quit
```

After the configurations are complete, check the IP routing table of R1.

```
[R1]display ip routing-table

Route Flags: R - relay, D - download to fib

------------------------------------------------------------------------------

Routing Tables: Public

         Destinations : 16          Routes : 16

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.0/24  Direct  0    0           D   10.0.1.1        LoopBack0
       10.0.1.1/32  Direct  0    0           D   127.0.0.1       LoopBack0
     10.0.1.255/32  Direct  0    0           D   127.0.0.1       LoopBack0
       10.0.2.0/24  OSPF    10   1562        D   10.0.12.2       Serial1/0/0
       10.0.3.0/24  OSPF    10   3124        D   10.0.12.2       Serial1/0/0
      10.0.12.0/24  Direct  0    0           D   10.0.12.1       Serial1/0/0
      10.0.12.1/32  Direct  0    0           D   127.0.0.1       Serial1/0/0
      10.0.12.2/32  Direct  0    0           D   10.0.12.2       Serial1/0/0
    10.0.12.255/32  Direct  0    0           D   127.0.0.1       Serial1/0/0
      10.0.23.0/24  OSPF    10   3124        D   10.0.12.2       Serial1/0/0
      10.0.24.0/24  OSPF    10   1563        D   10.0.12.2       Serial1/0/0
```

```
      10.0.35.0/24  OSPF    10   4686      D   10.0.12.2       Serial1/0/0
      127.0.0.0/8   Direct  0    0         D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0         D   127.0.0.1       InLoopBack0
127.255.255.255/32Direct    0    0         D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct    0    0         D   127.0.0.1       InLoopBack0
```

Test network connectivity.

```
[R1]ping -c 1 10.0.35.5
  PING 10.0.35.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.35.5: bytes=56 Sequence=1 ttl=253 time=114 ms

  --- 10.0.35.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 114/114/114 ms

[R1]ping -c 1 10.0.3.3
  PING 10.0.3.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.3.3: bytes=56 Sequence=1 ttl=254 time=74 ms

  --- 10.0.3.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 74/74/74 ms

[R1]ping -c 1 10.0.24.4
```

```
PING 10.0.24.4: 56   data bytes, press CTRL_C to break
  Reply from 10.0.24.4: bytes=56 Sequence=1 ttl=254 time=34 ms

--- 10.0.24.4 ping statistics ---
  1 packet(s) transmitted
  1 packet(s) received
  0.00% packet loss
  round-trip min/avg/max = 34/34/34 ms
```

Step 3 **Import external routes into OSPF.**

Import the route to the network segment 10.0.5.0/24 where Loopback0 of R5 resides into an OSPF area. Use the default configuration to import the route.

```
[R5]ospf 1
[R5-ospf-1]import-route direct
```

After the configurations are complete, check the imported route on R1 and test network connectivity.

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 7            Routes : 7

OSPF routing table status : <Active>
         Destinations : 7            Routes : 7

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

```
   10.0.2.0/24   OSPF     10   1562      D   10.0.12.2      Serial1/0/0
   10.0.3.0/24   OSPF     10   3124      D   10.0.12.2      Serial1/0/0
   10.0.5.0/24   O_ASE    150  1         D   10.0.12.2      Serial1/0/0
  10.0.23.0/24   OSPF     10   3124      D   10.0.12.2      Serial1/0/0
  10.0.24.0/24   OSPF     10   1563      D   10.0.12.2      Serial1/0/0
  10.0.35.0/24   OSPF     10   4686      D   10.0.12.2      Serial1/0/0
  10.0.35.3/32   O_ASE    150  1         D   10.0.12.2      Serial1/0/0

OSPF routing table status : <Inactive>
         Destinations : 0          Routes : 0

[R1]ping -c 1 10.0.5.5
  PING 10.0.5.5: 56   data bytes, press CTRL_C to break
    Reply from 10.0.5.5: bytes=56 Sequence=1 ttl=253 time=111 ms

  --- 10.0.5.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 111/111/111 ms
```

Configure a default route on R4 with the next hop pointing to Loopback0. Import this default route into an OSPF area, define it as a Type 1 route, and set its cost to 20, without using permanent advertisement.

```
[R4]ip route-static 0.0.0.0 0.0.0.0 LoopBack 0
[R4]ospf 1
[R4-ospf-1]default-route-advertise type 1 cost 20
[R4-ospf-1]quit
```

After the configurations are complete, check information about learning this default route on R1, and test network connectivity.

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 8          Routes : 8

OSPF routing table status : <Active>
         Destinations : 8          Routes : 8

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        0.0.0.0/0   O_ASE   150  1583        D   10.0.12.2       Serial1/0/0
      10.0.2.0/24   OSPF    10   1562        D   10.0.12.2       Serial1/0/0
      10.0.3.0/24   OSPF    10   3124        D   10.0.12.2       Serial1/0/0
      10.0.5.0/24   O_ASE   150  1           D   10.0.12.2       Serial1/0/0
     10.0.23.0/24   OSPF    10   3124        D   10.0.12.2       Serial1/0/0
     10.0.24.0/24   OSPF    10   1563        D   10.0.12.2       Serial1/0/0
     10.0.35.0/24   OSPF    10   4686        D   10.0.12.2       Serial1/0/0
     10.0.35.3/32   O_ASE   150  1           D   10.0.12.2       Serial1/0/0

OSPF routing table status : <Inactive>
         Destinations : 0          Routes : 0

[R1]ping -c 1 10.0.4.4
  PING 10.0.4.4: 56  data bytes, press CTRL_C to break
```

```
Reply from 10.0.4.4: bytes=56 Sequence=1 ttl=254 time=39 ms

 --- 10.0.4.4 ping statistics ---
   1 packet(s) transmitted
   1 packet(s) received
   0.00% packet loss
   round-trip min/avg/max = 39/39/39 ms
```

Step 4 **Configure Area 2 as a stub area.**

Check routing information on R1. The default route is an external route (O_ASE), which is learned through the Type 5 LSA advertised by R4.

```
[R1]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.1.1
                 Link State Database

                         Area: 0.0.0.2
Type       LinkState ID    AdvRouter        Age  Len   Sequence   Metric
Router       10.0.2.2        10.0.2.2          12   48    80000003    1562
Router       10.0.1.1        10.0.1.1          11   60    80000003       0
Sum-Net     10.0.35.0       10.0.2.2          33   28    80000001    3124
Sum-Net     10.0.24.0       10.0.2.2          33   28    80000001       1
Sum-Net     10.0.3.0        10.0.2.2          33   28    80000001    1562
Sum-Net     10.0.2.0        10.0.2.2          33   28    80000001       0
Sum-Net     10.0.23.0       10.0.2.2          34   28    80000001    1562
Sum-Asbr   10.0.4.4        10.0.2.2          34   28    80000001       1
Sum-Asbr   10.0.5.5        10.0.2.2          34   28    80000001    3124
```

```
                 AS External Database
 Type      LinkState ID    AdvRouter          Age  Len     Sequence     Metric
 External   0.0.0.0         10.0.4.4          1049  36      80000002       20
 External   10.0.5.0        10.0.5.5          1350  36      80000001        1
 External   10.0.35.0       10.0.5.5          1350  36      80000001        1
 External   10.0.35.3       10.0.5.5          1350  36      80000001        1

[R1]display ospf lsdb ase 0.0.0.0

         OSPF Process 1 with Router ID 10.0.1.1
                 Link State Database

  Type          : External
  Ls id         : 0.0.0.0
  Adv rtr : 10.0.4.4
  Ls age        : 504
  Len           : 36
  Options       :   E
  seq#          : 80000002
  chksum        : 0xa981
  Net mask      : 0.0.0.0
  TOS 0    Metric: 20
  E type        : 1
  Forwarding Address : 0.0.0.0
  Tag           : 1
  Priority   : Low
```

On R1 and R2, configure Area 2 as a stub area.

```
[R1]ospf 1
[R1-ospf-1]area 2
[R1-ospf-1-area-0.0.0.2]stub
[R1-ospf-1-area-0.0.0.2]quit
[R1-ospf-1]quit

[R2]ospf 1
[R2-ospf-1]area 2
[R2-ospf-1-area-0.0.0.2]stub
[R2-ospf-1-area-0.0.0.2]quit
[R2-ospf-1]quit
```

After the configurations are complete, on R1, compare the current IP routing table with the previous one and check routing information learning. You can see that the external route disappears and the default route also becomes an internal route.

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 6          Routes : 6

OSPF routing table status : <Active>
         Destinations : 6          Routes : 6

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        0.0.0.0/0   OSPF    10   1563        D   10.0.12.2       Serial1/0/0
```

```
      10.0.2.0/24  OSPF    10    1562         D    10.0.12.2       Serial1/0/0
      10.0.3.0/24  OSPF    10    3124         D    10.0.12.2       Serial1/0/0
     10.0.23.0/24  OSPF    10    3124         D    10.0.12.2       Serial1/0/0
     10.0.24.0/24  OSPF    10    1563         D    10.0.12.2       Serial1/0/0
     10.0.35.0/24  OSPF    10    4686         D    10.0.12.2       Serial1/0/0

OSPF routing table status : <Inactive>
         Destinations : 0           Routes : 0
```

Check the LSDB of R1. You can see that the LSA describing the external route also disappears, and the default route is learned through a Type 3 LSA.

```
[R1]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.1.1
                 Link State Database

                         Area: 0.0.0.2
 Type       LinkState ID     AdvRouter          Age  Len     Sequence     Metric
 Router     10.0.2.2         10.0.2.2          182   48     80000003     1562
 Router     10.0.1.1         10.0.1.1          182   60     80000004        0
 Sum-Net     0.0.0.0          10.0.2.2          183   28     80000001        1
 Sum-Net     10.0.35.0        10.0.2.2          183   28     80000001     3124
 Sum-Net     10.0.24.0        10.0.2.2          183   28     80000001        1
 Sum-Net     10.0.3.0         10.0.2.2          183   28     80000001     1562
 Sum-Net     10.0.2.0         10.0.2.2          184   28     80000001        0
 Sum-Net     10.0.23.0        10.0.2.2          184   28     80000001     1562
```

Check detailed information about this Type 3 LSA. You can see that the default route described by this LSA is advertised by R2. This proves that after an area is configured as a stub area, an ABR prevents Type 4 and Type 5 LSAs from being sent to this area and uses a Type 3 LSA to flood a default route pointing to itself within this area.

```
[R1]display ospf lsdb summary 0.0.0.0

         OSPF Process 1 with Router ID 10.0.1.1
                         Area: 0.0.0.2
                 Link State Database

  Type        : Sum-Net
  Ls id       : 0.0.0.0
  Adv rtr : 10.0.2.2
  Ls age      : 114
  Len         : 28
  Options     : None
  seq#        : 80000001
  chksum           : 0x1f31
  Net mask    : 0.0.0.0
  Tos 0    metric: 1
  Priority : Low
```

On R2, configure Area 2 as a totally stub area and specify the **no-summary** parameter.

```
[R2]ospf 1
[R2-ospf-1]area 2
[R2-ospf-1-area-0.0.0.2]stub no-summary
[R2-ospf-1-area-0.0.0.2]quit
```

[R2-ospf-1]quit

Check the OSPF routing table of R1. You can see that only one default route is learned through OSPF.

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 1        Routes : 1

OSPF routing table status : <Active>
         Destinations : 1        Routes : 1

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        0.0.0.0/0   OSPF    10   1563        D   10.0.12.2       Serial1/0/0

OSPF routing table status : <Inactive>
         Destinations : 0        Routes : 0
```

Check the LSDB of R1. You can see that the LSDB contains only one Type 3 LSA generated by R2 in addition to the Type 1 LSAs generated by R1 and R2.

This proves that in a totally stub area, an ABR blocks Type 3, Type 4, and Type 5 LSAs and generates a Type 3 LSA to advertise a default route pointing to itself.

```
[R1]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.1.1
                 Link State Database
```

Area: 0.0.0.2

Type	LinkState ID	AdvRouter	Age	Len	Sequence	Metric
Router	10.0.2.2	10.0.2.2	167	48	80000004	1562
Router	10.0.1.1	10.0.1.1	166	60	80000006	0
Sum-Net	0.0.0.0	10.0.2.2	549	28	80000001	1

Step 5 **Configure Area 1 as an NSSA area.**

Check the OSPF routing table of R3. You can see that the network segment 10.0.5.0/24 advertised by R5 is displayed as an external route.

[R3]display ip routing-table protocol ospf

Route Flags: R - relay, D - download to fib

Public routing table : OSPF

Destinations : 7 Routes : 7

OSPF routing table status : <Active>

Destinations : 6 Routes : 6

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
0.0.0.0/0	O_ASE	150	1583	D	10.0.23.2	Serial2/0/0
10.0.1.0/24	OSPF	10	3124	D	10.0.23.2	Serial2/0/0
10.0.2.0/24	OSPF	10	1562	D	10.0.23.2	Serial2/0/0
10.0.5.0/24	O_ASE	150	1	D	10.0.35.5	Serial3/0/0
10.0.12.0/24	OSPF	10	3124	D	10.0.23.2	Serial2/0/0
10.0.24.0/24	OSPF	10	1563	D	10.0.23.2	Serial2/0/0

OSPF routing table status : <Inactive>

Destinations : 1 Routes : 1

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
10.0.35.3/32	O_ASE	150	1		10.0.35.5	Serial3/0/0

Check the OSPF routing table and LSDB of R5. You can see that R5 learns an external route from R4 and that the remaining routes are all internal routes. R5 uses a Type 5 LSA to advertise the network segment 10.0.5.0/24.

[R5]display ip routing-table protocol ospf

Route Flags: R - relay, D - download to fib

Public routing table : OSPF

Destinations : 7 Routes : 7

OSPF routing table status : <Active>

Destinations : 7 Routes : 7

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
0.0.0.0/0	O_ASE	150	3145	D	10.0.35.3	Serial1/0/0
10.0.1.0/24	OSPF	10	4686	D	10.0.35.3	Serial1/0/0
10.0.2.0/24	OSPF	10	3124	D	10.0.35.3	Serial1/0/0
10.0.3.0/24	OSPF	10	1562	D	10.0.35.3	Serial1/0/0
10.0.12.0/24	OSPF	10	4686	D	10.0.35.3	Serial1/0/0
10.0.23.0/24	OSPF	10	3124	D	10.0.35.3	Serial1/0/0
10.0.24.0/24	OSPF	10	3125	D	10.0.35.3	Serial1/0/0

```
OSPF routing table status : <Inactive>
         Destinations : 0          Routes : 0

[R5]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.5.5
                 Link State Database

                         Area: 0.0.0.1
 Type       LinkState ID    AdvRouter          Age  Len    Sequence   Metric
 Router       10.0.5.5        10.0.5.5           882   48     80000004    1562
 Router       10.0.3.3        10.0.3.3          1309   48     80000003    1562
 Sum-Net     10.0.24.0        10.0.3.3            65   28     80000003    1563
 Sum-Net     10.0.12.0        10.0.3.3           819   28     80000001    3124
 Sum-Net     10.0.3.0         10.0.3.3            65   28     80000003       0
 Sum-Net     10.0.2.0         10.0.3.3            65   28     80000003    1562
 Sum-Net     10.0.1.0         10.0.3.3           812   28     80000001    3124
 Sum-Net     10.0.23.0        10.0.3.3            65   28     80000003    1562
 Sum-Asbr   10.0.4.4         10.0.3.3           602   28     80000002    1563


                 AS External Database
 Type       LinkState ID    AdvRouter          Age  Len    Sequence   Metric
 External   10.0.5.0        10.0.5.5         882   36     80000002       1
 External   10.0.35.0       10.0.5.5         883   36     80000002       1
 External   10.0.35.3       10.0.5.5         883   36     80000002       1
 External   0.0.0.0         10.0.4.4         586   36     80000003      20
```

On R3 and R5, configure Area 1 as an NSSA area.

```
[R3]ospf 1
[R3-ospf-1]area 1
[R3-ospf-1-area-0.0.0.1]nssa
[R3-ospf-1-area-0.0.0.1]quit
[R3-ospf-1]quit

[R5]ospf 1
[R5-ospf-1]area 1
[R5-ospf-1-area-0.0.0.1]nssa
[R5-ospf-1-area-0.0.0.1]quit
[R5-ospf-1]quit
```

After a neighbor relationship is established again, check the OSPF routing table of R3.

```
[R3]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 7          Routes : 7

OSPF routing table status : <Active>
         Destinations : 6          Routes : 6

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        0.0.0.0/0   O_ASE   150  1583        D   10.0.23.2       Serial2/0/0
       10.0.1.0/24  OSPF    10   3124        D   10.0.23.2       Serial2/0/0
```

```
   10.0.2.0/24  OSPF      10   1562        D   10.0.23.2       Serial2/0/0
   10.0.5.0/24  O_NSSA    150  1           D   10.0.35.5       Serial3/0/0
  10.0.12.0/24  OSPF      10   3124        D   10.0.23.2       Serial2/0/0
  10.0.24.0/24  OSPF      10   1563        D   10.0.23.2       Serial2/0/0

OSPF routing table status : <Inactive>
         Destinations : 1        Routes : 1

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

  10.0.35.3/32  O_NSSA  150  1               10.0.35.5       Serial3/0/0
```

The preceding command output shows that the external route advertised by R5 is displayed as O_NSSA in the OSPF routing table.

Check the OSPF routing table of R5 again.

```
[R5]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 7        Routes : 7

OSPF routing table status : <Active>
         Destinations : 7        Routes : 7

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

     0.0.0.0/0  O_NSSA  150    1           D   10.0.35.3       Serial1/0/0
   10.0.1.0/24  OSPF    10   4686          D   10.0.35.3       Serial1/0/0
```

```
     10.0.2.0/24   OSPF    10    3124          D    10.0.35.3      Serial1/0/0
     10.0.3.0/24   OSPF    10    1562          D    10.0.35.3      Serial1/0/0
    10.0.12.0/24   OSPF    10    4686          D    10.0.35.3      Serial1/0/0
    10.0.23.0/24   OSPF    10    3124          D    10.0.35.3      Serial1/0/0
    10.0.24.0/24   OSPF    10    3125          D    10.0.35.3      Serial1/0/0

OSPF routing table status : <Inactive>
         Destinations : 0          Routes : 0
```

The default route was previously displayed as an external route (O_ASE) and now becomes an external route (O_NSSA) of an NSSA area.

Check the LSDB of R5.

```
[R5]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.5.5
                 Link State Database

                         Area: 0.0.0.1
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.5.5        10.0.5.5           811   48   80000007     1562
 Router    10.0.3.3        10.0.3.3           811   48   80000007     1562
 Sum-Net   10.0.24.0       10.0.3.3           929   28   80000005     1563
 Sum-Net   10.0.12.0       10.0.3.3           929   28   80000005     3124
 Sum-Net   10.0.3.0        10.0.3.3           929   28   80000005        0
 Sum-Net   10.0.2.0        10.0.3.3           929   28   80000005     1562
 Sum-Net   10.0.1.0        10.0.3.3           930   28   80000005     3124
 Sum-Net   10.0.23.0       10.0.3.3           930   28   80000005     1562
 NSSA      10.0.5.0        10.0.5.5           819   36   80000005        1
```

```
NSSA      10.0.35.0       10.0.5.5       819   36    80000006     1
NSSA      10.0.35.3       10.0.5.5       819   36    80000005     1
NSSA      0.0.0.0         10.0.3.3       930   36    80000005     1
```

You can see that the Type 5 LSA disappears and the external route is advertised using a Type 7 LSA.

Check detailed information about the default route.

```
[R5]display ospf lsdb nssa 0.0.0.0

         OSPF Process 1 with Router ID 10.0.5.5
                      Area: 0.0.0.1
                 Link State Database

  Type        : NSSA
  Ls id       : 0.0.0.0
  Adv rtr : 10.0.3.3
  Ls age      : 1149
  Len         : 36
  Options     : None
  seq#        : 80000005
  chksum          : 0x7745
  Net mask    : 0.0.0.0
  TOS 0   Metric: 1
  E type      : 2
  Forwarding Address : 0.0.0.0
  Tag         : 1
  Priority : Low
```

The previous default route on R5 was advertised by R4, but the current default route is advertised by R3.

This proves that external Type 4 and Type 5 LSAs are prevented from entering an NSSA area, and an ABR uses a Type 7 LSA to advertise a default route within this area. The external route of this area will be advertised by an ASBR as a Type 7 LSA into the NSSA area.

The fundamental difference between an NSSA area and a stub area is that an NSSA area allows importing external routes but a stub area does not.

Step 6 **Observe changes brought by an NSSA area to OSPF.**

Run the **display ospf brief** command to check the role of R3. You can see that the **Border Router** field displays three values: **AREA AS NSSA**. AREA indicates that this router is an ABR; AS indicates that this router is an ASBR; NSSA indicates that this router has at least one interface located in an NSSA area.

```
[R3]display ospf brief


         OSPF Process 1 with Router ID 10.0.3.3
                 OSPF Protocol Information


 RouterID: 10.0.3.3          Border Router:  AREA  AS  NSSA
 Multi-VPN-Instance is not enabled
 Global DS-TE Mode: Non-Standard IETF Mode
 Graceful-restart capability: disabled
 Helper support capability  : not configured
 Applications Supported: MPLS Traffic-Engineering
 Spf-schedule-interval: max 10000ms, start 500ms, hold 1000ms
 Default ASE parameters: Metric: 1 Tag: 1 Type: 2
 Route Preference: 10
 ASE Route Preference: 150
```

SPF Computation Count: 14

RFC 1583 Compatible

Retransmission limitation is disabled

Area Count: 2 Nssa Area Count: 1

ExChange/Loading Neighbors: 0

Process total up interface count: 3

Process valid up interface count: 2

Area: 0.0.0.0 (MPLS TE not enabled)

Authtype: None Area flag: Normal

SPF scheduled Count: 14

ExChange/Loading Neighbors: 0

Router ID conflict state: Normal

Area interface up count: 2

Interface: 10.0.3.3 (LoopBack0)

Cost: 0 State: DR Type: Broadcast MTU: 1500

Priority: 1

Designated Router: 10.0.3.3

Backup Designated Router: 0.0.0.0

Timers: Hello 10 , Dead 40 , Poll 120 , Retransmit 5 , Transmit Delay 1

Interface: 10.0.23.3 (Serial2/0/0) --> 10.0.23.2

Cost: 1562 State: P-2-P Type: P2P MTU: 1500

Timers: Hello 10 , Dead 40 , Poll 120 , Retransmit 5 , Transmit Delay 1

Area: 0.0.0.1 (MPLS TE not enabled)

Authtype: None Area flag: NSSA

SPF scheduled Count: 3

```
ExChange/Loading Neighbors: 0
NSSA Translator State: Elected
Router ID conflict state: Normal
Area interface up count: 1
NSSA LSA count: 0

Interface: 10.0.35.3 (Serial3/0/0) --> 10.0.35.5
Cost: 1562      State: P-2-P       Type: P2P        MTU: 1500
Timers: Hello 10 , Dead 40 , Poll   120 , Retransmit 5 , Transmit Delay 1
```

Type 5 LSAs are not allowed in an NSSA area. Therefore, an ASBR uses a Type 7 LSA to advertise an external route within an NSSA area. However, this Type 7 LSA can only be flooded within an NSSA area. After an ABR of this area receives this LSA, it translates it into a Type 5 LSA and then advertises it to other common areas.

On R3, observe the procedure for translating Type 7 LSAs into Type 5 LSAs. The following example uses the network segment 10.0.5.0/24. For a Type 7 LSA, the **Ls id** field indicates the destination network segment, and the **Net mask** field indicates the mask of the destination network segment. If the **Options** field displays **NP**, this LSA can be translated by an ABR into a Type 5 LSA. If the **Options** field indicates that this LSA cannot be translated into a Type 5 LSA, the Forwarding Address can be set to 0.0.0.0. If the **Options** field indicates that this LSA can be translated into a Type 5 LSA, the Forwarding Address cannot be set to 0.0.0.0.

Here, the next hop of the imported external route is not within an OSPF routing domain, and the Forwarding Address needs to be set as this ASBR's interface IP address of the stub network segment within an OSPF routing domain. The address used here is the address of Serial1/0/0 on R5.

```
[R3]display ospf lsdb nssa 10.0.5.0


         OSPF Process 1 with Router ID 10.0.3.3
```

```
                Area: 0.0.0.0
            Link State Database

                Area: 0.0.0.1
            Link State Database

  Type         : NSSA
  Ls id        : 10.0.5.0
  Adv rtr      : 10.0.5.5
  Ls age       : 836
  Len          : 36
  Options      :   NP
  seq#         : 80000001
  chksum       : 0xb0c2
  Net mask     : 255.255.255.0
  TOS 0    Metric: 1
  E type       : 2
  Forwarding Address : 10.0.35.5
  Tag          : 1
  Priority     : Low
```

Check the Type 5 LSA generated by R3 to describe the network segment 10.0.5.0/24.

```
[R3]display ospf lsdb ase 10.0.5.0

         OSPF Process 1 with Router ID 10.0.3.3
            Link State Database

  Type         : External
```

```
Ls id        : 10.0.5.0
Adv rtr      : 10.0.3.3
Ls age       : 882
Len          : 36
Options      :   E
seq#         : 80000001
chksum       : 0x413e
Net mask     : 255.255.255.0
TOS 0    Metric: 1
E type       : 2
Forwarding Address : 10.0.35.5
Tag          : 1
Priority     : Low
```

The values of the Ls id, Network Mask, and Forwarding Address fields are copied from the previous Type 7 LSA. In this manner, the network segment 10.0.5.0/24 is advertised into other areas.

----End

Additional Exercises: Analysis and Verification

Which scenarios are NSSA areas applicable to?

Analyze why R3 is defined as an ASBR.

Device Configurations

```
<R1>display current-configuration
[V200R007C00SPC600]
#
 sysname R1
#
```

```
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.1.1
 area 0.0.0.2
  network 10.0.12.1 0.0.0.0
  network 10.0.1.1 0.0.0.0
  stub
#
return

<R2>display current-configuration
[V200R007C00SPC600]
#
 sysname R2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
```

```
interface GigabitEthernet0/0/0
 ip address 10.0.24.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.2.2
 area 0.0.0.0
  network 10.0.23.2 0.0.0.0
  network 10.0.2.2 0.0.0.0
 area 0.0.0.2
  network 10.0.12.2 0.0.0.0
  stub no-summary
 area 0.0.0.3
  network 10.0.24.2 0.0.0.0
#
return

<R3>display current-configuration
[V200R007C00SPC600]
#
 sysname R3
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface Serial3/0/0
```

```
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.3.3
 area 0.0.0.0
  network 10.0.23.3 0.0.0.0
  network 10.0.3.3 0.0.0.0
 area 0.0.0.1
  network 10.0.35.3 0.0.0.0
  nssa
#
return

<R4>display current-configuration
[V200R007C00SPC600]
#
 sysname R4
#
interface GigabitEthernet0/0/0
 ip address 10.0.24.4 255.255.255.0
#
interface NULL0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.0
```

```
 ospf network-type broadcast
#
ospf 1 router-id 10.0.4.4
 default-route-advertise cost 20 type 1
 area 0.0.0.3
  network 10.0.24.4 0.0.0.0
#
 ip route-static 0.0.0.0 0.0.0.0 LoopBack0
#
return

<R5>display current-configuration
[V200R007C00SPC600]
#
 sysname R5
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.0
#
ospf 1 router-id 10.0.5.5
 import-route direct
 area 0.0.0.1
  network 10.0.35.5 0.0.0.0
  nssa
#
```

return

Lab 1-5 OSPF Virtual Link and Inter-Area Route Filtering

Learning Objectives

The objectives of this lab are to learn and understand:

- How to configure OSPF virtual links to connect to non-contiguous Area 0
- How to configure OSPF virtual links to connect a non-backbone area to Area 0
- How to filter and control routes between areas

Topology

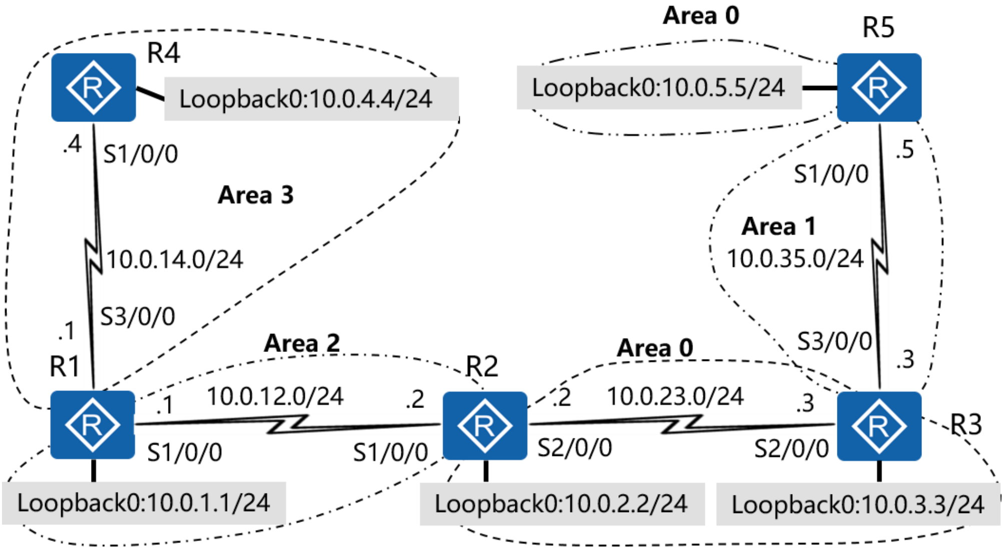

Figure 1-5 OSPF virtual link and inter-area route filtering

Scenario

You are a network administrator of a company. This company recently acquired two small companies, whose routers are R4 and R5 respectively. To combine networks, you decide to configure OSPF virtual links to implement network interconnection, instead of planning the networks again. You find that there is non-contiguous Area 0 and that Area 3 is not directly connected to Area 0. Therefore, you establish a virtual link between R1 and R2 to enable Area 3 to be directly connected to Area 0. Additionally, you establish a virtual link between R3 and R5 to connect non-contiguous Area 0.

To specify router IDs for the routers, configure the routers to use fixed addresses as their router IDs.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

Configure IP addresses and masks for all the routers. Set a 24-bit mask for all loopback interfaces to simulate an independent network segment.

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 24
[R1-Serial3/0/0]quit
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 24
[R1-LoopBack0]quit

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]quit
[R2]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 24
```

```
[R2-Serial2/0/0]quit
[R2]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24
[R2-LoopBack0]quit

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 24
[R3-Serial2/0/0]quit
[R3]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 24
[R3-Serial3/0/0]quit
[R3]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 24
[R3-LoopBack0]quit

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 24
[R4-Serial1/0/0]quit
[R4]interface LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 24
[R4-LoopBack0]quit

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface Serial 1/0/0
```

```
[R5-Serial1/0/0]ip address 10.0.35.5 24
[R5-Serial1/0/0]quit
[R5]interface LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 24
[R5-LoopBack0]quit
```

After the configurations are complete, test direct link connectivity.

```
[R1]ping -c 1 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=39 ms

  --- 10.0.14.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 39/39/39 ms
[R1]ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=29 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 29/29/29 ms

[R3]ping -c 1 10.0.23.2
  PING 10.0.23.2: 56  data bytes, press CTRL_C to break
```

```
    Reply from 10.0.23.2: bytes=56 Sequence=1 ttl=255 time=45 ms

  --- 10.0.23.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 45/45/45 ms

[R3]ping -c 1 10.0.35.5
  PING 10.0.35.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.35.5: bytes=56 Sequence=1 ttl=255 time=32 ms

  --- 10.0.35.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 32/32/32 ms
```

Step 2 **Configure multiple OSPF areas.**

On R1, configure Serial1/0/0 and Loopback0 to belong to Area 2 and Serial3/0/0 to belong to Area 3. To enable OSPF to advertise real masks of loopback interfaces, change the OSPF network type of loopback interfaces in all the areas to broadcast. Configure all routers to use IP address of Loopback0 as their router IDs.

```
[R1]ospf 1 router-id 10.0.1.1
[R1-ospf-1]area 2
[R1-ospf-1-area-0.0.0.2]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]network 10.0.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]quit
```

```
[R1-ospf-1]area 3
[R1-ospf-1-area-0.0.0.3]network 10.0.14.1 0.0.0.0
[R1-ospf-1-area-0.0.0.3]quit
[R1-ospf-1]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ospf network-type broadcast
[R1-LoopBack0]quit
```

On R2, configure Serial2/0/0 and Loopback0 to belong to Area 0 and Serial1/0/0 to belong to Area 2.

```
[R2]ospf 1 router-id 10.0.2.2
[R2-ospf-1]area 2
[R2-ospf-1-area-0.0.0.2]network 10.0.12.2 0.0.0.0
[R2-ospf-1-area-0.0.0.2]quit
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.23.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]quit
[R2-ospf-1]quit
[R2]interface LoopBack 0
[R2-LoopBack0]ospf network-type broadcast
[R2-LoopBack0]quit
```

On R3, configure Serial2/0/0 and Loopback0 to belong to Area 0 and Serial3/0/0 to belong to Area 1.

```
[R3]ospf 1 router-id 10.0.3.3
[R3-ospf-1]area 0
```

[R3-ospf-1-area-0.0.0.0]network 10.0.23.3 0.0.0.0

[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0

[R3-ospf-1-area-0.0.0.0]quit

[R3-ospf-1]area 1

[R3-ospf-1-area-0.0.0.1]network 10.0.35.3 0.0.0.0

[R3-ospf-1-area-0.0.0.1]quit

[R3-ospf-1]quit

[R3]interface LoopBack 0

[R3-LoopBack0]ospf network-type broadcast

[R3-LoopBack0]quit

On R4, configure Serial1/0/0 and Loopback0 to belong to Area 3.

[R4]ospf 1 router-id 10.0.4.4

[R4-ospf-1]area 3

[R4-ospf-1-area-0.0.0.3]network 10.0.14.4 0.0.0.0

[R4-ospf-1-area-0.0.0.3]network 10.0.4.4 0.0.0.0

[R4-ospf-1-area-0.0.0.3]quit

[R4-ospf-1]quit

[R4]interface LoopBack 0

[R4-LoopBack0]ospf network-type broadcast

[R4-LoopBack0]quit

On R5, configure Serial1/0/0 to belong to Area 1 and Looback0 to belong to Area 0.

[R5]ospf 1 router-id 10.0.5.5

[R5-ospf-1]area 0

[R5-ospf-1-area-0.0.0.0]network 10.0.5.5 0.0.0.0

[R5-ospf-1-area-0.0.0.0]quit

[R5-ospf-1]area 1

```
[R5-ospf-1-area-0.0.0.1]network 10.0.35.5 0.0.0.0
[R5-ospf-1-area-0.0.0.3]quit
[R5-ospf-1]quit
[R5]interface LoopBack 0
[R5-LoopBack0]ospf network-type broadcast
[R5-LoopBack0]quit
```

Step 3 **Check the OSPF routing table of each router.**

Check the OSPF routing table of R4. Although R4 establishes a neighbor relationship with R1, it does not learn any OSPF routes.

```
[R4]display ip routing-table protocol ospf
[R4]display ospf peer

         OSPF Process 1 with Router ID 10.0.4.4
                 Neighbors

 Area 0.0.0.3 interface 10.0.14.4(Serial1/0/0)'s neighbors
 Router ID: 10.0.1.1          Address: 10.0.14.1
   State: Full   Mode:Nbr is   Slave   Priority: 1
   DR: None    BDR: None    MTU: 0
   Dead timer due in 39   sec
   Retrans timer interval: 4
   Neighbor is up for 00:21:33
   Authentication Sequence: [ 0 ]
```

Check the LSDB of R4. You can see that there are only Type 1 LSAs. That is, R1 does not advertise routes of other areas into Area 3.

```
[R4]display ospf lsdb
```

```
OSPF Process 1 with Router ID 10.0.4.4
         Link State Database

                  Area: 0.0.0.3
Type      LinkState ID    AdvRouter        Age  Len   Sequence   Metric
Router    10.0.4.4        10.0.4.4         571  60    80000005     0
Router    10.0.1.1        10.0.1.1         616  48    80000003   1562
```

Check the OSPF routing table of R1. The route to 10.0.5.0/24 disappears. After analyzing the LSDB of R3, you will know why this route disappears.

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 5        Routes : 5

OSPF routing table status : <Active>
         Destinations : 5        Routes : 5

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.2.0/24  OSPF    10   1562        D   10.0.12.2       Serial1/0/0
       10.0.3.0/24  OSPF    10   3124        D   10.0.12.2       Serial1/0/0
       10.0.4.0/24  OSPF    10   1562        D   10.0.14.4       Serial3/0/0
      10.0.23.0/24  OSPF    10   3124        D   10.0.12.2       Serial1/0/0
      10.0.35.0/24  OSPF    10   4686        D   10.0.12.2       Serial1/0/0
```

```
OSPF routing table status : <Inactive>
        Destinations : 0        Routes : 0
```

Check the LSDB of R1. To prevent inter-area loops, OSPF does not allow directly advertising routing information between two non-backbone areas. The LSDB shows that an ABR does not forward the Type 3 LSAs received from non-backbone areas.

On R1, the LSDB for Area 2 has four inter-area routes, which are learned from R2 (10.0.2.2). R1 does not forward these LSAs into Area 3. Therefore, R4 cannot learn routes outside its local area.

An ABR does not forward the routes learned from a non-backbone area to another non-backbone area. The routes learned by R1 from R4 will not be advertised as Type 3 LSAs into Area 2. Therefore, R2, R3, and R5 cannot learn routes of Area 3.

```
[R1]display ospf lsdb

        OSPF Process 1 with Router ID 10.0.1.1
                Link State Database

                        Area: 0.0.0.2
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router     10.0.2.2        10.0.2.2          1251  48    80000023    1562
 Router     10.0.1.1        10.0.1.1          1266  60    80000024       0
 Sum-Net    10.0.35.0       10.0.2.2          1178  28    8000001B    3124
 Sum-Net    10.0.3.0        10.0.2.2          1178  28    8000001B    1562
 Sum-Net    10.0.2.0        10.0.2.2          1228  28    80000021       0
 Sum-Net    10.0.23.0       10.0.2.2          1189  28    8000001B    1562

                        Area: 0.0.0.3
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.4.4        10.0.4.4           855   60   80000024       0
```

```
 Router      10.0.1.1          10.0.1.1              898   48     80000022      1562
```

Check the OSPF routing table of R2. Three routes to network segments 10.0.4.0/24, 10.0.5.0/24, and 10.0.14.0/24 respectively disappear from the OSPF routing table of R2.

```
[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 3          Routes : 3

OSPF routing table status : <Active>
         Destinations : 3          Routes : 3

Destination/Mask    Proto    Pre   Cost        Flags NextHop           Interface

       10.0.1.0/24   OSPF     10    1562          D    10.0.12.1         Serial1/0/0
       10.0.3.0/24   OSPF     10    1562          D    10.0.23.3         Serial2/0/0
      10.0.35.0/24   OSPF     10    3124          D    10.0.23.3         Serial2/0/0

OSPF routing table status : <Inactive>
         Destinations : 0          Routes : 0
```

Check the LSDB of R2. You can see that R1 does not advertise routes of Area 3 to R2.

Therefore, R2 does not have routes to network segments 10.0.4.0/24 and 10.0.14.0/24.

In Area 0, R3 does not advertise the route 10.0.5.0 to R2.

```
[R2]display ospf lsdb
```

OSPF Process 1 with Router ID 10.0.2.2

Link State Database

Area: 0.0.0.0

Type	LinkState ID	AdvRouter	Age	Len	Sequence	Metric
Router	10.0.3.3	10.0.3.3	973	60	80000027	0
Router	10.0.2.2	10.0.2.2	972	60	80000028	0
Sum-Net	10.0.35.0	10.0.3.3	984	28	8000001D	1562
Sum-Net	10.0.12.0	10.0.2.2	1035	28	80000022	1562
Sum-Net	10.0.1.0	10.0.2.2	1035	28	80000022	1562

Area: 0.0.0.2

Type	LinkState ID	AdvRouter	Age	Len	Sequence	Metric
Router	10.0.2.2	10.0.2.2	1046	48	80000024	1562
Router	10.0.1.1	10.0.1.1	1063	60	80000025	0
Sum-Net	10.0.35.0	10.0.2.2	973	28	8000001C	3124
Sum-Net	10.0.3.0	10.0.2.2	973	28	8000001C	1562
Sum-Net	10.0.2.0	10.0.2.2	1023	28	80000022	0
Sum-Net	10.0.23.0	10.0.2.2	984	28	8000001C	1562

Check the OSPF routing table of R3. The routes to network segments 10.0.4.0/24, 10.0.5.0/24, and 10.0.14.0/24 disappear from the OSPF routing table.

[R3]display ip routing-table protocol ospf

Route Flags: R - relay, D - download to fib

Public routing table : OSPF

Destinations : 3 Routes : 3

```
OSPF routing table status : <Active>
         Destinations : 3          Routes : 3

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.0/24   OSPF    10   3124        D   10.0.23.2       Serial2/0/0
      10.0.2.0/24   OSPF    10   1562        D   10.0.23.2       Serial2/0/0
     10.0.12.0/24   OSPF    10   3124        D   10.0.23.2       Serial2/0/0

OSPF routing table status : <Inactive>
         Destinations : 0          Routes : 0
```

Check the LSDB of R3. You can see that in Area 1, R3 receives a Type 3 LSA 10.0.5.0 from R5. According to rules, R3 does not forward the Type 3 LSA received from a non-backbone area.

R3 does not send this LSA into Area 0 again. This is why R1 and R2 do not have the route 10.0.5.0/24.

```
[R3]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.3.3
                 Link State Database

                         Area: 0.0.0.0
 Type        LinkState ID     AdvRouter          Age  Len    Sequence    Metric
 Router       10.0.3.3         10.0.3.3          111   60    80000028        0
 Router       10.0.2.2         10.0.2.2          112   60    80000029        0
 Sum-Net      10.0.35.0        10.0.3.3          122   28    8000001E     1562
```

```
Sum-Net    10.0.12.0      10.0.2.2       175   28    80000023    1562
Sum-Net    10.0.1.0       10.0.2.2       175   28    80000023    1562

                    Area: 0.0.0.1
Type       LinkState ID   AdvRouter      Age   Len   Sequence    Metric
Router     10.0.5.5       10.0.5.5       117   48    8000001E    1562
Router     10.0.3.3       10.0.3.3       117   48    80000020    1562
Sum-Net    10.0.12.0      10.0.3.3       107   28    8000001D    3124
Sum-Net    10.0.3.0       10.0.3.3       128   28    8000001D       0
Sum-Net    10.0.2.0       10.0.3.3       107   28    8000001D    1562
Sum-Net    10.0.1.0       10.0.3.3       108   28    8000001D    3124
Sum-Net    10.0.5.0       10.0.5.5       128   28    8000001D       0
Sum-Net    10.0.23.0      10.0.3.3       124   28    8000001D    1562
```

The Type 3 LSA 10.0.5.0/24 received from R5 already exists in the LSDB of R3 but does not appear in the routing table of R3.

Check the OSPF routing table of R5.

```
[R5]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 5        Routes : 5

OSPF routing table status : <Active>
         Destinations : 5        Routes : 5

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

```
       10.0.1.0/24   OSPF    10    4686      D    10.0.35.3       Serial1/0/0
       10.0.2.0/24   OSPF    10    3124      D    10.0.35.3       Serial1/0/0
       10.0.3.0/24   OSPF    10    1562      D    10.0.35.3       Serial1/0/0
      10.0.12.0/24   OSPF    10    4686      D    10.0.35.3       Serial1/0/0
      10.0.23.0/24   OSPF    10    3124      D    10.0.35.3       Serial1/0/0

OSPF routing table status : <Inactive>
         Destinations : 0          Routes : 0

 [R5]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.5.5
                 Link State Database

                         Area: 0.0.0.0
 Type       LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router       10.0.5.5        10.0.5.5          820  36    80000002      0
 Sum-Net    10.0.35.0       10.0.5.5          861  28    80000001   1562

                         Area: 0.0.0.1
 Type       LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router       10.0.5.5        10.0.5.5         1096  48    80000003   1562
 Router       10.0.3.3        10.0.3.3         1097  48    80000002   1562
 Sum-Net    10.0.12.0       10.0.3.3         1129  28    80000001   3124
 Sum-Net    10.0.3.0        10.0.3.3         1129  28    80000001      0
 Sum-Net    10.0.2.0        10.0.3.3         1129  28    80000001   1562
 Sum-Net    10.0.1.0        10.0.3.3         1129  28    80000001   3124
 Sum-Net    10.0.5.0        10.0.5.5          861  28    80000001      0
 Sum-Net    10.0.23.0       10.0.3.3         1129  28    80000001   1562
```

R5 does not have routes to network segments 10.0.4.0/24 and 10.0.14.0/24.

R5 has the route to Loopback0 of R3.

R3 has a physical interface connected to Area 0 and can exchange routing information with other routers in Area 0. In this situation, R3 does not add the routes learned through Type 3 LSAs from non-backbone areas to its routing table. Although R5 has an interface belonging to Area 0, this interface is a loopback interface, whose link type is StubNet during OSPF route calculation.

Check the Type 1 LSAs generated by R3. The following command output displays only information about the Type 1 LSAs.

```
[R3]display ospf lsdb router 10.0.3.3

         OSPF Process 1 with Router ID 10.0.3.3
                        Area: 0.0.0.0
                 Link State Database

  Type        : Router
  Ls id       : 10.0.3.3
  Adv rtr     : 10.0.3.3
  Ls age      : 732
  Len         : 60
  Options     :  ABR  E
  seq#        : 80000158
  chksum      : 0xde39
  Link count: 3
   * Link ID: 10.0.3.3
     Data     : 255.255.255.255
     Link Type: StubNet
```

```
    Metric : 0
    Priority : Medium
  * Link ID: 10.0.2.2
    Data     : 10.0.23.3
    Link Type: P-2-P
    Metric : 1562
  * Link ID: 10.0.23.0
    Data     : 255.255.255.0
    Link Type: StubNet
    Metric : 1562
    Priority : Low
```

The preceding command output shows that the type of the link between R3 and R2 is P-2-P. If the link type of an interface is P-2-P, TransNet, or Virtual, a router considers that this interface will exchange routing information with other routers. The router connected to a backbone area through each of the three links does not add the routes learned through Type 3 LSAs from non-backbone areas to its routing table.

```
[R5]display ospf lsdb router 10.0.5.5

          OSPF Process 1 with Router ID 10.0.5.5
                          Area: 0.0.0.0
                  Link State Database

  Type        : Router
  Ls id       : 10.0.5.5
  Adv rtr     : 10.0.5.5
  Ls age      : 583
  Len         : 36
```

```
Options    :  ABR  E
seq#       : 80000040
chksum      : 0x6d69
Link count: 1
 * Link ID: 10.0.5.5
   Data     : 255.255.255.255
   Link Type: StubNet
   Metric : 0
   Priority : Medium
```

R5 has only one Loopback0 belonging to the backbone area. In the LSA describing the route to this interface address, the link type is StubNet, indicating that this interface is not connected to any other router. Then R5 adds the route learned through a Type 3 LSA sent from a non-backbone area to its routing table.

Step 4 **Connect two non-contiguous Areas 0 together.**

Configure a virtual link on R3 and R5 and specify the router ID of the peer ABR in the **vlink-peer** command.

```
[R3]ospf 1
[R3-ospf-1]area 1
[R3-ospf-1-area-0.0.0.1]vlink-peer 10.0.5.5
[R3-ospf-1-area-0.0.0.1]quit
[R3-ospf-1]quit

[R5]ospf
[R5-ospf-1]area 1
[R5-ospf-1-area-0.0.0.1]vlink-peer 10.0.3.3
[R5-ospf-1-area-0.0.0.1]quit
[R5-ospf-1]quit
```

Check whether the neighbor state of the virtual link is Full.

```
[R3]display ospf vlink

          OSPF Process 1 with Router ID 10.0.3.3
                   Virtual Links

 Virtual-link Neighbor-id  -> 10.0.5.5, Neighbor-State: Full

 Interface: 10.0.35.3 (Serial3/0/0)
 Cost: 1562  State: P-2-P  Type: Virtual
 Transit Area: 0.0.0.1
 Timers: Hello 10 , Dead 40 , Retransmit 5 , Transmit Delay 1
 GR State: Normal
```

Observe routing information changes.

```
[R3]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
          Destinations : 4          Routes : 4

OSPF routing table status : <Active>
          Destinations : 4          Routes : 4

Destination/Mask    Proto   Pre  Cost       Flags NextHop         Interface

        10.0.1.0/24  OSPF    10   3124         D    10.0.23.2       Serial2/0/0
```

```
    10.0.2.0/24   OSPF    10    1562        D   10.0.23.2       Serial2/0/0
    10.0.5.0/24   OSPF    10    1562        D   10.0.35.5       Serial3/0/0
   10.0.12.0/24   OSPF    10    3124        D   10.0.23.2       Serial2/0/0

OSPF routing table status : <Inactive>
         Destinations : 0          Routes : 0
```

The preceding command output shows that R3 learns the route to 10.0.5.0/24.

Test network connectivity. You can see that R3 can communicate with the network segment connected to Loopback0 of R5.

```
[R3]ping -c 1 10.0.5.5
  PING 10.0.5.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.5.5: bytes=56 Sequence=1 ttl=255 time=34 ms

  --- 10.0.5.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 34/34/34 ms
```

Check the LSDB of R3.

```
<R3>display ospf lsdb

         OSPF Process 1 with Router ID 10.0.3.3
                 Link State Database

                         Area: 0.0.0.0
 Type       LinkState ID     AdvRouter            Age  Len     Sequence    Metric
```

Router	10.0.5.5	10.0.5.5	1098	48	80000005	0
Router	10.0.3.3	10.0.3.3	1096	72	80000008	0
Router	10.0.2.2	10.0.2.2	920	60	80000006	0
Sum-Net	10.0.35.0	10.0.3.3	830	28	80000002	1562
Sum-Net	10.0.35.0	10.0.5.5	565	28	80000002	1562
Sum-Net	10.0.12.0	10.0.2.2	1124	28	80000002	1562
Sum-Net	10.0.1.0	10.0.2.2	1110	28	80000002	1562

Area: 0.0.0.1

Type	LinkState ID	AdvRouter	Age	Len	Sequence	Metric
Router	10.0.5.5	10.0.5.5	1098	48	80000004	1562
Router	10.0.3.3	10.0.3.3	1096	48	80000003	1562
Sum-Net	10.0.12.0	10.0.3.3	830	28	80000002	3124
Sum-Net	10.0.3.0	10.0.3.3	831	28	80000002	0
Sum-Net	10.0.2.0	10.0.3.3	831	28	80000002	1562
Sum-Net	10.0.1.0	10.0.3.3	831	28	80000002	3124
Sum-Net	10.0.5.0	10.0.5.5	566	28	80000002	0
Sum-Net	10.0.23.0	10.0.3.3	831	28	80000002	1562

R3 receives two Type 1 LSAs from R5. The first Type 1 LSA is received in Area 0, and the virtual link belongs to Area 0. Therefore, this LSA is learned through the virtual link. The second Type 1 LSA is learned in Area 1 and already exists before the virtual link is established. The route to 10.0.5.0/24 is calculated through the LSA learned in Area 0.

Check detailed information about the Type 1 LSA 10.0.5.5 in the LSDB of R3.

```
[R3]display ospf lsdb router 10.0.5.5

          OSPF Process 1 with Router ID 10.0.3.3
```

```
                    Area: 0.0.0.0
             Link State Database

Type        : Router
Ls id       : 10.0.5.5
Adv rtr     : 10.0.5.5
Ls age      : 621
Len         : 48
Options     :   ABR   E
seq#        : 80000005
chksum      : 0x1291
Link count: 2
 * Link ID: 10.0.5.0
   Data     : 255.255.255.0
   Link Type: StubNet
   Metric : 0
   Priority : Low
 * Link ID: 10.0.3.3
   Data     : 10.0.35.5
   Link Type: Virtual
   Metric : 1562
                    Area: 0.0.0.1
             Link State Database

Type        : Router
Ls id       : 10.0.5.5
Adv rtr     : 10.0.5.5
Ls age      : 621
Len         : 48
```

```
Options    :  ABR   VIRTUAL   E
seq#        : 80000004
chksum      : 0x3530
Link count: 2
 * Link ID: 10.0.3.3
   Data     : 10.0.35.5
   Link Type: P-2-P
   Metric : 1562
 * Link ID: 10.0.35.0
   Data     : 255.255.255.0
   Link Type: StubNet
   Metric : 1562
   Priority : Low
```

The preceding command output shows that this LSA describes the network 10.0.5.0/24. Therefore, R3 has the corresponding route. The Type 1 LSA learned in Area 1 describes only the interconnected network segment between R3 and R5.

Check the LSDB of R5.

```
[R5]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.5.5
                 Link State Database

                         Area: 0.0.0.0
Type      LinkState ID     AdvRouter          Age   Len    Sequence    Metric
Router    10.0.5.5         10.0.5.5           577   48     80000005      0
Router    10.0.3.3         10.0.3.3           577   72     80000008      0
Router    10.0.2.2         10.0.2.2           401   60     80000006      0
```

```
Sum-Net    10.0.35.0      10.0.5.5         45   28     80000002     1562
Sum-Net    10.0.35.0      10.0.3.3        312   28     80000002     1562
Sum-Net    10.0.12.0      10.0.2.2        606   28     80000002     1562
Sum-Net    10.0.1.0       10.0.2.2        593   28     80000002     1562

                     Area: 0.0.0.1
Type       LinkState ID   AdvRouter        Age  Len    Sequence     Metric
Router     10.0.5.5       10.0.5.5        578   48     80000004     1562
Router     10.0.3.3       10.0.3.3        578   48     80000003     1562
Sum-Net    10.0.12.0      10.0.3.3        313   28     80000002     3124
Sum-Net    10.0.3.0       10.0.3.3        313   28     80000002        0
Sum-Net    10.0.2.0       10.0.3.3        313   28     80000002     1562
Sum-Net    10.0.1.0       10.0.3.3        313   28     80000002     3124
Sum-Net    10.0.5.0       10.0.5.5         46   28     80000002        0
Sum-Net    10.0.23.0      10.0.3.3        313   28     80000002     1562
```

You can see that the LSDB of R5 is the same as that of R3. After the virtual link is established, R3 and R5 both have interfaces that belong to Area 0. Therefore, their LSDBs are synchronized.

Step 5 **Connect Area 3 to Area 0 through a virtual link.**

Configure a virtual link on R1 and R2.

```
[R1]ospf 1
[R1-ospf-1]area 2
[R1-ospf-1-area-0.0.0.2]vlink-peer 10.0.2.2
[R1-ospf-1-area-0.0.0.2]quit
[R1-ospf-1]quit

[R2]ospf
```

```
[R2-ospf-1]area 2

[R2-ospf-1-area-0.0.0.2]vlink-peer 10.0.1.1

[R2-ospf-1-area-0.0.0.2]quit

[R2-ospf-1]quit
```

Check the OSPF routing table of R4.

```
[R4]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 7          Routes : 7

OSPF routing table status : <Active>
         Destinations : 7          Routes : 7

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.0/24  OSPF    10   1562        D   10.0.14.1       Serial1/0/0
       10.0.2.0/24  OSPF    10   3124        D   10.0.14.1       Serial1/0/0
       10.0.3.0/24  OSPF    10   4686        D   10.0.14.1       Serial1/0/0
       10.0.5.0/24  OSPF    10   6248        D   10.0.14.1       Serial1/0/0
      10.0.12.0/24  OSPF    10   3124        D   10.0.14.1       Serial1/0/0
      10.0.23.0/24  OSPF    10   4686        D   10.0.14.1       Serial1/0/0
      10.0.35.0/24  OSPF    10   6248        D   10.0.14.1       Serial1/0/0

OSPF routing table status : <Inactive>
         Destinations : 0          Routes : 0
```

The preceding command output shows that R4 has network-wide routes.

Test network connectivity.

```
[R4]ping -c 1 10.0.5.5
  PING 10.0.5.5: 56   data bytes, press CTRL_C to break
    Reply from 10.0.5.5: bytes=56 Sequence=1 ttl=252 time=132 ms

  --- 10.0.5.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 132/132/132 ms
```

Check the LSDB of R1.

```
[R1]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.1.1
                 Link State Database

                         Area: 0.0.0.0
 Type       LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router      10.0.5.5        10.0.5.5          419   48    80000006       0
 Router      10.0.3.3        10.0.3.3          418   72    80000009       0
 Router      10.0.2.2        10.0.2.2          232   72    8000000A       0
 Router      10.0.1.1        10.0.1.1          233   36    80000001    1562
 Sum-Net     10.0.35.0       10.0.3.3          151   28    80000003    1562
 Sum-Net     10.0.35.0       10.0.5.5         1687   28    80000002    1562
```

Sum-Net	10.0.14.0	10.0.1.1	291	28	80000001	1562
Sum-Net	10.0.12.0	10.0.1.1	291	28	80000001	1562
Sum-Net	10.0.12.0	10.0.2.2	444	28	80000003	1562
Sum-Net	10.0.1.0	10.0.1.1	291	28	80000001	0
Sum-Net	10.0.1.0	10.0.2.2	430	28	80000003	1562
Sum-Net	10.0.4.0	10.0.1.1	291	28	80000001	1562

Area: 0.0.0.2

Type	LinkState ID	AdvRouter	Age	Len	Sequence	Metric
Router	10.0.2.2	10.0.2.2	235	48	80000005	1562
Router	10.0.1.1	10.0.1.1	234	60	80000009	0
Sum-Net	10.0.35.0	10.0.2.2	151	28	80000003	3124
Sum-Net	10.0.14.0	10.0.1.1	291	28	80000001	1562
Sum-Net	10.0.3.0	10.0.2.2	234	28	80000003	1562
Sum-Net	10.0.2.0	10.0.2.2	443	28	80000003	0
Sum-Net	10.0.5.0	10.0.2.2	402	28	80000002	3124
Sum-Net	10.0.4.0	10.0.1.1	292	28	80000001	1562
Sum-Net	10.0.23.0	10.0.2.2	286	28	80000003	1562

Area: 0.0.0.3

Type	LinkState ID	AdvRouter	Age	Len	Sequence	Metric
Router	10.0.4.4	10.0.4.4	1193	60	80000005	0
Router	10.0.1.1	10.0.1.1	292	48	80000004	1562
Sum-Net	10.0.35.0	10.0.1.1	292	28	80000001	4686
Sum-Net	10.0.12.0	10.0.1.1	294	28	80000001	1562
Sum-Net	10.0.3.0	10.0.1.1	294	28	80000001	3124
Sum-Net	10.0.2.0	10.0.1.1	294	28	80000001	1562
Sum-Net	10.0.1.0	10.0.1.1	294	28	80000001	0
Sum-Net	10.0.5.0	10.0.1.1	294	28	80000001	4686

```
Sum-Net    10.0.23.0       10.0.1.1          294   28     80000001     3124
```

Because a virtual link is created, R1 has LSAs of Area 0. Then Area 0 and Area 3 can exchange routes directly. R1 uses a Type 3 LA to advertise routing information about Area 0 into Area 3.

Check the LSDB of R4.

```
[R4]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.4.4
                 Link State Database

                         Area: 0.0.0.3
Type       LinkState ID    AdvRouter          Age   Len    Sequence    Metric
Router      10.0.4.4        10.0.4.4         1303   60     80000005       0
Router      10.0.1.1        10.0.1.1          404   48     80000004    1562
Sum-Net    10.0.35.0       10.0.1.1          404   28     80000001     4686
Sum-Net    10.0.12.0       10.0.1.1          404   28     80000001     1562
Sum-Net    10.0.3.0        10.0.1.1          404   28     80000001     3124
Sum-Net    10.0.2.0        10.0.1.1          404   28     80000001     1562
Sum-Net    10.0.1.0        10.0.1.1          405   28     80000001        0
Sum-Net    10.0.5.0        10.0.1.1          405   28     80000001     4686
Sum-Net    10.0.23.0       10.0.1.1          405   28     80000001     3124
```

The preceding command output shows that R4 learns the Type 3 LSA advertised by R1.

R4 has routes of other areas.

Step 6 **Configure inter-area route filtering.**

Control advertisement of the route to 10.0.4.0/24 so that R1 can learn this route but R2, R3, and R5 cannot.

Configure an ACL.

```
[R1]acl number 2000
[R1-acl-basic-2000]rule deny source 10.0.4.0 0.0.0.255
[R1-acl-basic-2000]rule permit
[R1-acl-basic-2000]permit
```

Configure Type 3 LSA filtering on R1 when R1 sends routing updates from Area 3 to other areas.

```
[R1]ospf 1
[R1-ospf-1]area 3
[R1-ospf-1-area-0.0.0.3]filter 2000 export
[R1-ospf-1-area-0.0.0.3]quit
[R1-ospf-1]quit
```

Check route filtering on R2.

```
[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 5            Routes : 5

OSPF routing table status : <Active>
         Destinations : 5            Routes : 5

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

```
      10.0.1.0/24   OSPF    10   1562         D   10.0.12.1      Serial1/0/0
      10.0.3.0/24   OSPF    10   1562         D   10.0.23.3      Serial2/0/0
      10.0.5.0/24   OSPF    10   3124         D   10.0.23.3      Serial2/0/0
     10.0.14.0/24   OSPF    10   3124         D   10.0.12.1      Serial1/0/0
     10.0.35.0/24   OSPF    10   3124         D   10.0.23.3      Serial2/0/0

OSPF routing table status : <Inactive>
         Destinations : 0         Routes : 0
```

R2 cannot learn the route 10.0.4.0/24.

R1 still has this route. This is because R1 and R4 belong to the same area and R4 uses a Type 1 LSA to advertise this route to R1.

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 6         Routes : 6

OSPF routing table status : <Active>
         Destinations : 6         Routes : 6

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.2.0/24   OSPF    10   1562         D   10.0.12.2      Serial1/0/0
      10.0.3.0/24   OSPF    10   3124         D   10.0.12.2      Serial1/0/0
      10.0.4.0/24   OSPF    10   1562         D   10.0.14.4      Serial3/0/0
      10.0.5.0/24   OSPF    10   4686         D   10.0.12.2      Serial1/0/0
     10.0.23.0/24   OSPF    10   3124         D   10.0.12.2      Serial1/0/0
```

```
    10.0.35.0/24   OSPF     10    4686          D    10.0.12.2        Serial1/0/0

OSPF routing table status : <Inactive>
         Destinations : 0           Routes : 0
```

----End

Additional Exercises: Analysis and Verification

Why must Areas 0 in OSPF be contiguous? Can Type 1 and Type 2 LSAs be filtered according to the current OSPF design?

Device Configurations

```
<R1>display current-configuration
[V200R007C00SPC600]
#
 sysname R1
#
acl number 2000
 rule 5 deny source 10.0.4.0 0.0.0.255
 rule 10 permit
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
#
```

```
interface LoopBack0
 ip address 10.0.1.1 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.1.1
 area 0.0.0.0
 area 0.0.0.2
  network 10.0.1.1 0.0.0.0
  network 10.0.12.1 0.0.0.0
  vlink-peer 10.0.2.2
 area 0.0.0.3
  filter 2000 export
  network 10.0.14.1 0.0.0.0
#
return

<R2>display current-configuration
[V200R007C00SPC600]
#
 sysname R2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
```

```
interface LoopBack0
 ip address 10.0.2.2 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.2.2
 area 0.0.0.0
  network 10.0.23.2 0.0.0.0
  network 10.0.2.2 0.0.0.0
 area 0.0.0.2
  network 10.0.12.2 0.0.0.0
  vlink-peer 10.0.1.1
#
return

<R3>display current-configuration
[V200R007C00SPC600]
#
 sysname R3
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.0
```

```
 ospf network-type broadcast
#
ospf 1 router-id 10.0.3.3
 area 0.0.0.0
  network 10.0.3.3 0.0.0.0
  network 10.0.23.3 0.0.0.0
 area 0.0.0.1
  network 10.0.35.3 0.0.0.0
  vlink-peer 10.0.5.5
#
return

<R4>display current-configuration
[V200R007C00SPC600]
#
 sysname R4
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.4.4
 area 0.0.0.3
  network 10.0.14.4 0.0.0.0
  network 10.0.4.4 0.0.0.0
```

```
#
return

<R5>display current-configuration
[V200R007C00SPC600]
#
 sysname R5
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.5.5
 area 0.0.0.0
  network 10.0.5.5 0.0.0.0
 area 0.0.0.1
  network 10.0.35.5 0.0.0.0
  vlink-peer 10.0.3.3
#
return
```

Lab 1-6 OSPF Troubleshooting

Learning Objectives

The objectives of this lab are to learn and understand:

- How to troubleshoot inconsistent area IDs in a single OSPF area
- How to troubleshoot unmatched masks in a single OSPF area
- How to troubleshoot inconsistent Hello intervals in a single OSPF area
- How to troubleshoot conflicting router IDs in a single OSPF area
- How to troubleshoot OSPF authentication failures
- How to troubleshoot OSPF route summarization failures
- How to troubleshoot virtual link failures

Topology

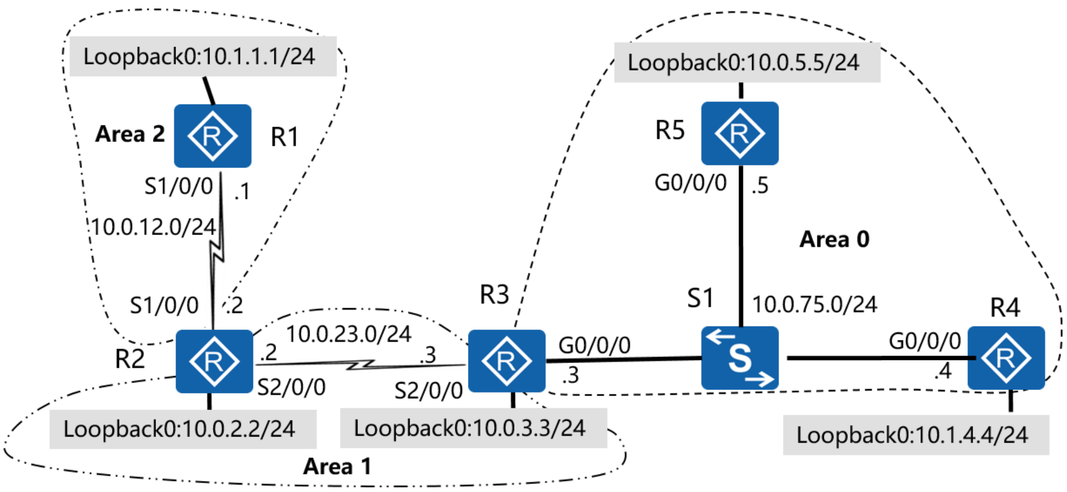

Figure 1-6 OSPF troubleshooting

Scenario

You are a network administrator of a company. The company' s network uses OSPF as the routing protocol. OSPF has powerful functions but also has complex configurations. You use various OSPF features including virtual link in network planning. During network operation, many network communication problems occur.

You use troubleshooting methods to locate and solve these problems, restoring the network.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

Configure IP addresses and masks for all the routers. Set a 24-bit mask for all loopback interfaces to simulate an independent network segment.

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface  Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ip address 10.1.1.1 24
[R1-LoopBack0]quit

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface  Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]quit
[R2]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 24
[R2-Serial2/0/0]quit
[R2]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24
[R2-LoopBack0]quit
```

To simulate failures, configure an IP address 10.0.75.3/25 for G0/0/0 of R3 and configure IP addresses for other interfaces according to the topology.

```
<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 24
[R3-Serial2/0/0]quit
[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]ip address 10.0.75.3 25
[R3-GigabitEthernet0/0/0]quit
[R3]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 24
[R3-LoopBack0]quit

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ip address 10.0.75.4 24
[R4-GigabitEthernet0/0/0]quit
[R4]interface LoopBack 0
[R4-LoopBack0]ip address 10.1.4.4 24
[R4-LoopBack0]quit

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.75.5 24
[R5-GigabitEthernet0/0/0]quit
[R5]interface LoopBack 0
```

```
[R5-LoopBack0]ip address 10.0.5.5 24
[R5-LoopBack0]quit
```

After the configurations are complete, test direct link connectivity.

```
[R3]ping -c 1 10.0.75.4
  PING 10.0.75.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.75.4: bytes=56 Sequence=1 ttl=255 time=5 ms

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

[R3]ping -c 1 10.0.75.5
  PING 10.0.75.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.75.5: bytes=56 Sequence=1 ttl=255 time=5 ms

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

[R3]ping -c 1 10.0.23.2
  PING 10.0.23.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.2: bytes=56 Sequence=1 ttl=255 time=41 ms
```

```
  --- 10.0.23.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 41/41/41 ms

[R1]ping -c 1 10.0.12.2
  PING 10.0.12.2: 56   data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=37 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 37/37/37 ms
```

Step 2 **Configure multiple OSPF areas.**

Configure Serial1/0/0 and Loopback0 of R1 to belong to Area 2 and configure R1 to use the address of Loopback0 as its router ID. To enable OSPF to advertise real masks of loopback interfaces, change the OSPF network type of loopback interfaces in all the areas to broadcast.

```
[R1]ospf 1 router-id 10.1.1.1
[R1-ospf-1]area 2
[R1-ospf-1-area-0.0.0.2]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]network 10.1.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]quit
[R1-ospf-1]quit
[R1]interface LoopBack 0
```

```
[R1-LoopBack0]ospf network-type broadcast
[R1-LoopBack0]quit
```

On R2, configure Serial2/0/0 and Loopback0 to belong to Area 1 and Serial1/0/0 to belong to Area 2. Do not specify a router ID when enabling OSPF.

```
[R2]ospf 1
[R2-ospf-1]area 1
[R2-ospf-1-area-0.0.0.1]network 10.0.23.2 0.0.0.0
[R2-ospf-1-area-0.0.0.1]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.1]quit
[R2-ospf-1]area 2
[R2-ospf-1-area-0.0.0.2]network 10.0.12.2 0.0.0.0
[R2-ospf-1-area-0.0.0.2]quit
[R2-ospf-1]quit
[R2]interface  LoopBack 0
[R2-LoopBack0]ospf network-type broadcast
[R2-LoopBack0]quit
```

On R3, configure Serial2/0/0 and Loopback0 to belong to Area 1 and GigabitEthernet 0/0/0 to belong to Area 0.

```
[R3]ospf 1 router-id 10.0.3.3
[R3-ospf-1]area 1
[R3-ospf-1-area-0.0.0.1]network 10.0.23.3 0.0.0.0
[R3-ospf-1-area-0.0.0.1]network 10.0.3.3 0.0.0.0
[R3-ospf-1-area-0.0.0.1]quit
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]network 10.0.75.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]quit
```

```
[R3-ospf-1]quit
[R3]interface LoopBack 0
[R3-LoopBack0]ospf network-type broadcast
[R3-LoopBack0]quit
```

On R4, configure GigabitEthernet0/0/0 to belong to Area 1 and Loopback0 not to belong to any area. When configuring an OSPF process, run the **ospf 1 router-id** command to specify a router ID 10.0.5.5 for R4.

```
[R4]ospf 1 router-id 10.0.5.5
[R4-ospf-1]area 1
[R4-ospf-1-area-0.0.0.1]network 10.0.75.4 0.0.0.0
[R4-ospf-1-area-0.0.0.1]quit
[R4-ospf-1]quit
```

On R5, configure GigabitEthernet 0/0/0 and Loopback0 to belong to Area 0.

```
[R5]ospf 1 router-id 10.0.5.5
[R5-ospf-1]area 0
[R5-ospf-1-area-0.0.0.0]network 10.0.75.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]network 10.0.5.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]quit
[R5-ospf-1]quit
[R5]interface LoopBack 0
[R5-LoopBack0]ospf network-type broadcast
[R5-LoopBack0]quit
```

Step 3 **Troubleshoot OSPF failures in areas.**

Check the neighbor list of R4. You can see that R4 does not establish neighbor relationships with other routers.

```
[R4]display ospf peer

          OSPF Process 1 with Router ID 10.0.5.5
```

Run the **display ospf error** command on R3, R4, and R5 to check OSPF errors.

```
[R3]display ospf error

          OSPF Process 1 with Router ID 10.0.3.3
                  OSPF error statistics

General packet errors:
 0      : IP: received my own packet   11     : Bad packet
 0      : Bad version                  0      : Bad checksum
 41     : Bad area id                  0      : Drop on unnumbered interface
 0      : Bad virtual link             0      : Bad authentication type
 0      : Bad authentication key       0      : Packet too small
 0      : Packet size > ip length      0      : Transmit error
 2      : Interface down               0      : Unknown neighbor
 0      : Bad net segment              0      : Extern option mismatch
 0      : Router id confusion          0      : Bad authentication sequence number

HELLO packet errors:
 227    : Netmask mismatch                0      : Hello timer mismatch
 0      : Dead timer mismatch             0      : Virtual neighbor unknown
 0      : NBMA neighbor unknown           0      : Invalid Source Address
```

```
[R4]display ospf error

         OSPF Process 1 with Router ID 10.0.5.5
                 OSPF error statistics

General packet errors:
 0      : IP: received my own packet   0      : Bad packet
 0      : Bad version                  0      : Bad checksum
 245    : Bad area id                  0      : Drop on unnumbered interface
 0      : Bad virtual link             0      : Bad authentication type
 0      : Bad authentication key       0      : Packet too small
 0      : Packet size > ip length      0      : Transmit error
 2      : Interface down               0      : Unknown neighbor
 0      : Bad net segment              0      : Extern option mismatch
 235    : Router id confusion          0      : Bad authentication sequence number

[R5]display ospf error

         OSPF Process 1 with Router ID 10.0.5.5
                 OSPF error statistics

General packet errors:
 0      : IP: received my own packet   260    : Bad packet
 0      : Bad version                  0      : Bad checksum
 0      : Bad area id                  0      : Drop on unnumbered interface
 0      : Bad virtual link             0      : Bad authentication type
 0      : Bad authentication key       0      : Packet too small
 0      : Packet size > ip length      0      : Transmit error
 0      : Interface down               0      : Unknown neighbor
```

0 : Bad net segment 0 : Extern option mismatch

286 : Router id confusion 0 : Bad authentication sequence number

HELLO packet errors:

260 : Netmask mismatch 0 : Hello timer mismatch

0 : Dead timer mismatch 0 : Virtual neighbor unknown

0 : NBMA neighbor unknown 0 : Invalid Source Address

The preceding command output shows that five types of errors occur on R3, R4, and R5: Router id confusion (router ID conflict), Netmask mismatch (unmatched subnet mask), Bad area id (incorrect area ID), Bad packet (error packet), and Bad virtual link (incorrect virtual link).

Because no virtual link is configured, the incorrect virtual link indicates an incorrect area ID. If R4 receives an OSPF packet with an area ID 0 on the interface with area ID 1, R4 considers that this packet is sent through a virtual link. No virtual link is configured on R4, this situation indicates that an error occurs.

A subnet mask error also indicates a type of error packet. You can rectify the subnet mask error and then check whether error packets still exist.

First, solve the router ID conflict. Check the router ID of each router in sequence to manually locate the router with the router ID or check system logs to locate the router. Run the **display logbuffer** command to check current system logs.

[R5]display logbuffer

Logging buffer configuration and contents: enabled

Allowed max buffer size: 1024

Actual buffer size: 512

Channel number: 4, Channel name: logbuffer

Dropped messages: 0

Overwritten messages: 0

Current messages: 66

Oct 26 2016 12:34:51+00:00 R5 %%01OSPF/4/CONFLICT_ROUTERID_INTF(l)[12]:OSPF Router id conflict is detected on interface. (ProcessId=1, RouterId=10.0.5.5, AreaId=0.0.0.0, InterfaceName=GigabitEthernet0/0/0, IpAddr=10.0.75.5, PacketSrcIp=10.0.75.4)

The preceding command output of R5 shows that a router ID conflict is detected on the interface with IP address 10.0.75.4. According to the topology, 10.0.75.4 is the interface address of R4. Check the router ID of R4, finding that its router ID is the same as that of R5. Additionally, the area ID configuration of R4 is also incorrect.

```
[R4]display ospf brief

         OSPF Process 1 with Router ID 10.0.5.5
                 OSPF Protocol Information

RouterID: 10.0.5.5           Border Router:
 Multi-VPN-Instance is not enabled
 Global DS-TE Mode: Non-Standard IETF Mode
 Graceful-restart capability: disabled
 Helper support capability   : not configured
 Applications Supported: MPLS Traffic-Engineering
 Spf-schedule-interval: max 10000ms, start 500ms, hold 1000ms
 Default ASE parameters: Metric: 1 Tag: 1 Type: 2
 Route Preference: 10
 ASE Route Preference: 150
 SPF Computation Count: 2
 RFC 1583 Compatible
 Retransmission limitation is disabled
 Area Count: 1    Nssa Area Count: 0
 ExChange/Loading Neighbors: 0
```

Process total up interface count: 1

Process valid up interface count: 1

Area: 0.0.0.1 (MPLS TE not enabled)

Authtype: None Area flag: Normal

SPF scheduled Count: 2

ExChange/Loading Neighbors: 0

Router ID conflict state: Normal

Area interface up count: 1

Interface: 10.0.75.4 (GigabitEthernet0/0/0)

Cost: 1 State: DR Type: Broadcast MTU: 1500

Priority: 1

Designated Router: 10.0.75.4

Backup Designated Router: 0.0.0.0

Timers: Hello 10 , Dead 40 , Poll 120 , Retransmit 5 , Transmit Delay 1

Change the router ID and area ID of R4.

[R4]ospf 1 router-id 10.1.4.4

[R4-ospf-1]area 1

[R4-ospf-1-area-0.0.0.1]undo network 10.0.75.4 0.0.0.0

[R4-ospf-1-area-0.0.0.1]quit

[R4-ospf-1]undo area 1

[R4-ospf-1]area 0

[R4-ospf-1-area-0.0.0.0]network 10.0.75.4 0.0.0.0

[R4-ospf-1-area-0.0.0.0]quit

[R4-ospf-1]quit

<R4>reset ospf process

Warning: The OSPF process will be reset. Continue? [Y/N]:y

Run the **reset ospf counter** command to clear OSPF statistics.

The reset command must be used in the user view.

<R4>reset ospf counters

Wait for a while and then run the **display ospf error** command to check whether the problem of router ID conflict and incorrect area ID is solved.

<R4>display ospf error

OSPF Process 1 with Router ID 10.1.4.4

OSPF error statistics

General packet errors:

0	: IP: received my own packet	13	: Bad packet
0	: Bad version	0	: Bad checksum
0	: Bad area id	0	: Drop on unnumbered interface
0	: Bad virtual link	0	: Bad authentication type
0	: Bad authentication key	0	: Packet too small
0	: Packet size > ip length	0	: Transmit error
0	: Interface down	0	: Unknown neighbor
0	: Bad net segment	0	: Extern option mismatch
0	: Router id confusion	0	: Bad authentication sequence number

HELLO packet errors:

13	: Netmask mismatch	0	: Hello timer mismatch
0	: Dead timer mismatch	0	: Virtual neighbor unknown

0 : NBMA neighbor unknown 0 : Invalid Source Address

The preceding command output shows that after the router ID and area ID of R4 are changed, the problem of router ID conflict and incorrect area ID is solved, and there is only the problem of unmatched subnet masks. To locate the router with an incorrect subnet mask, check debugging information on R4.

```
<R4>terminal debugging
Info: Current terminal debugging is on.
<R4>debugging ospf packet hello
Oct 26 2016 14:30:08.350.1+00:00 R4 RM/6/RMDEBUG:
 FileID: 0xd0178024 Line: 2271 Level: 0x20
 OSPF 1: RECV Packet. Interface: GigabitEthernet0/0/0
<R4>
Oct 26 2016 14:30:08.360.1+00:00 R4 RM/6/RMDEBUG:   Source Address: 10.0.75.3
Oct 26 2016 14:30:08.360.2+00:00 R4 RM/6/RMDEBUG:   Destination Address: 224.0.0.5
Oct 26 2016 14:30:08.360.3+00:00 R4 RM/6/RMDEBUG:   Ver# 2, Type: 1 (Hello)
Oct 26 2016 14:30:08.360.4+00:00 R4 RM/6/RMDEBUG:   Length: 44, Router: 10.0.3.3
Oct 26 2016 14:30:08.360.5+00:00 R4 RM/6/RMDEBUG:   Area: 0.0.0.0, Chksum: 9a18
Oct 26 2016 14:30:08.360.6+00:00 R4 RM/6/RMDEBUG:   AuType: 00
Oct 26 2016 14:30:08.360.7+00:00 R4 RM/6/RMDEBUG:   Key(ascii): * * * * * * * *
Oct 26 2016 14:30:08.360.8+00:00 R4 RM/6/RMDEBUG:   Net Mask: 255.255.255.128
Oct 26 2016 14:30:08.360.9+00:00 R4 RM/6/RMDEBUG:   Hello Int: 10, Option: _E_
Oct 26 2016 14:30:08.360.10+00:00 R4 RM/6/RMDEBUG:   Rtr Priority: 1, Dead Int: 40
Oct 26 2016 14:30:08.360.11+00:00 R4 RM/6/RMDEBUG:   DR: 10.0.75.3
Oct 26 2016 14:30:08.360.12+00:00 R4 RM/6/RMDEBUG:   BDR: 0.0.0.0
Oct 26 2016 14:30:08.360.13+00:00 R4 RM/6/RMDEBUG:   # Attached Neighbors: 0
```

The preceding command output shows that the subnet mask in the Hello packet sent from 10.0.75.3 is 255.255.255.128. According to the topology, the interface configuration of R3 is incorrect.

```
[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]display this
[V200R007C00SPC600]
#
interface GigabitEthernet0/0/0
 ip address 10.0.75.3 255.255.255.128
#
return
[R3-GigabitEthernet0/0/0]ip address 10.0.75.3 24
[R3-GigabitEthernet0/0/0]quit
```

Clear OSPF statistics again to check whether OSPF errors still exist.

```
<R3>reset ospf counters

<R3>display ospf error

         OSPF Process 1 with Router ID 10.0.3.3
                 OSPF error statistics

General packet errors:
0      : IP: received my own packet     0      : Bad packet
 0      : Bad version                    0      : Bad checksum
 0      : Bad area id                    0      : Drop on unnumbered interface
 0      : Bad virtual link               0      : Bad authentication type
 0      : Bad authentication key         0      : Packet too small
```

0	: Packet size > ip length	0	: Transmit error
0	: Interface down	0	: Unknown neighbor
0	: Bad net segment	0	: Extern option mismatch
0	: Router id confusion	0	: Bad authentication sequence number

HELLO packet errors:

0	: Netmask mismatch	0	: Hello timer mismatch
0	: Dead timer mismatch	0	: Virtual neighbor unknown
0	: NBMA neighbor unknown	0	: Invalid Source Address

Check the neighbor list of R3. You can see that its neighbor relationships with neighbors are normal.

```
[R3]display ospf peer brief

         OSPF Process 1 with Router ID 10.0.3.3
                  Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                         Neighbor id      State
 0.0.0.0          GigabitEthernet0/0/0                10.1.4.4       Full
 0.0.0.0          GigabitEthernet0/0/0                10.0.5.5       Full
 0.0.0.1          Serial2/0/0                         10.0.2.2       Full
 ----------------------------------------------------------------------------
```

Change the Hello interval of GigabitEthernet0/0/0 on R4 to 5 seconds to observe whether neighbor relationships can be established.

```
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ospf timer hello 5
[R4-GigabitEthernet0/0/0]quit
```

After about 30 seconds, you can see that all neighbor information of R4 disappears.

```
[R4]display ospf peer brief

         OSPF Process 1 with Router ID 10.1.4.4
                  Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                        Neighbor id      State
 ----------------------------------------------------------------------------
```

Clear OSPF statistics of R4 to check whether OSPF errors exist.

```
<R4>reset ospf counters
<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]display ospf error

         OSPF Process 1 with Router ID 10.1.4.4
                 OSPF error statistics

General packet errors:
0      : IP: received my own packet    4      : Bad packet
 0      : Bad version                   0      : Bad checksum
 0      : Bad area id                   0      : Drop on unnumbered interface
 0      : Bad virtual link              0      : Bad authentication type
 0      : Bad authentication key        0      : Packet too small
 0      : Packet size > ip length       0      : Transmit error
 0      : Interface down                0      : Unknown neighbor
 0      : Bad net segment               0      : Extern option mismatch
```

```
0       : Router id confusion          0       : Bad authentication sequence number

HELLO packet errors:
0       : Netmask mismatch             4       : Hello timer mismatch
0       : Dead timer mismatch              0       : Virtual neighbor unknown
0       : NBMA neighbor unknown            0       : Invalid Source Address
```

The preceding command output shows Hello timer mismatch, indicating that Hello intervals of neighbors are inconsistent.

Cancel the Hello interval configuration and then check the neighbor list again.

```
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]undo ospf timer hello
[R4-GigabitEthernet0/0/0]quit

[R4]display ospf peer brief

         OSPF Process 1 with Router ID 10.1.4.4
                  Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                     Neighbor id       State
 0.0.0.0          GigabitEthernet0/0/0          10.0.3.3          Full
 0.0.0.0          GigabitEthernet0/0/0          10.0.5.5          Full
 ----------------------------------------------------------------------------
```

The preceding command output shows that neighbor relationships become normal.

Step 4 **Troubleshoot OSPF authentication failures.**

Configure interface authentication on R1 and R2.

Configure simple authentication on R1 and set the key to 123.

Configure MD5 authentication on R5 and set the key to huawei.

[R1]interface Serial 1/0/0

[R1-Serial1/0/0]ospf authentication-mode simple plain 123

[R1-Serial1/0/0]quit

[R2]interface Serial 1/0/0

[R2-Serial1/0/0]ospf authentication-mode md5 1 plain huawei

[R2-Serial1/0/0]quit

After the configurations are complete, clear OSPF statistics of R1 and then check OSPF errors.

<R1>reset ospf counters

<R1>system-view

Enter system view, return user view with Ctrl+Z.

[R1]display ospf error

OSPF Process 1 with Router ID 10.1.1.1

OSPF error statistics

General packet errors:

0	: IP: received my own packet	3	: Bad packet
0	: Bad version	0	: Bad checksum
0	: Bad area id	0	: Drop on unnumbered interface
0	: Bad virtual link	3	: Bad authentication type
0	: Bad authentication key	0	: Packet too small
0	: Packet size > ip length	0	: Transmit error
0	: Interface down	0	: Unknown neighbor

0 : Bad net segment 0 : Extern option mismatch

0 : Router id confusion 0 : Bad authentication sequence number

Configure MD5 authentication on R1 and then check whether OSPF errors still exist.

[R1]interface Serial 1/0/0

[R1-Serial1/0/0]ospf authentication-mode md5 1 plain 123

[R1-Serial1/0/0]return

<R1>reset ospf counters

<R1>display ospf error

OSPF Process 1 with Router ID 10.1.1.1

OSPF error statistics

General packet errors:

0 : IP: received my own packet 9 : Bad packet

0 : Bad version 0 : Bad checksum

0 : Bad area id 0 : Drop on unnumbered interface

0 : Bad virtual link 0 : Bad authentication type

9 : Bad authentication key 0 : Packet too small

0 : Packet size > ip length 0 : Transmit error

0 : Interface down 0 : Unknown neighbor

0 : Bad net segment 0 : Extern option mismatch

0 : Router id confusion 0 : Bad authentication sequence number

The preceding command output shows that OSPF errors still exist.

Change the key of R1 to huawei and then check neighbor relationships.

[R1]interface Serial 1/0/0

```
[R1-Serial1/0/0]ospf authentication-mode md5 1 plain huawei
[R1-Serial1/0/0]quit
[R1]display ospf peer brief

         OSPF Process 1 with Router ID 10.1.1.1
                  Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                      Neighbor id      State
 0.0.0.2          Serial1/0/0                    10.0.2.2         Full
 ----------------------------------------------------------------------------
```

The preceding command output shows that R1 and R2 have established a neighbor relationship.

Step 5 **Troubleshoot virtual link failures.**

To ensure connectivity between Area 2 and Area 0, create a virtual link between R2 and R3.

```
[R2]ospf 1
[R2-ospf-1]area 1
[R2-ospf-1-area-0.0.0.1]vlink-peer 10.0.3.3
[R2-ospf-1-area-0.0.0.1]quit
[R2-ospf-1]quit

[R3]ospf 1
[R3-ospf-1]area 1
[R3-ospf-1-area-0.0.0.1]vlink-peer 10.0.2.2
[R3-ospf-1-area-0.0.0.1]quit
[R3-ospf-1]quit
```

Check whether the virtual link is established normally and whether R1 learns network-wide routes.

```
[R2]display ospf vlink

          OSPF Process 1 with Router ID 10.0.2.2
                   Virtual Links

 Virtual-link Neighbor-id  -> 10.0.3.3, Neighbor-State: Full

 Interface: 10.0.23.2 (Serial2/0/0)
 Cost: 1562  State: P-2-P  Type: Virtual
 Transit Area: 0.0.0.1
 Timers: Hello 10 , Dead 40 , Retransmit 5 , Transmit Delay 1
 GR State: Normal

[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
          Destinations : 5         Routes : 5

OSPF routing table status : <Active>
          Destinations : 5         Routes : 5

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.2.0/24  OSPF    10   1562        D   10.0.12.2       Serial1/0/0
       10.0.3.0/24  OSPF    10   3124        D   10.0.12.2       Serial1/0/0
       10.0.5.0/24  OSPF    10   3125        D   10.0.12.2       Serial1/0/0
```

```
    10.0.23.0/24   OSPF    10    3124        D    10.0.12.2       Serial1/0/0
    10.0.75.0/24   OSPF    10    3125        D    10.0.12.2       Serial1/0/0

OSPF routing table status : <Inactive>
         Destinations : 0           Routes : 0
```

Test connectivity from R1 to R5. The following command output shows that R1 can reach R5.

```
[R1]ping -c 1 10.0.5.5
  PING 10.0.5.5: 56   data bytes, press CTRL_C to break
    Reply from 10.0.5.5: bytes=56 Sequence=1 ttl=253 time=81 ms

  --- 10.0.5.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 81/81/81 ms
```

Delete Loopback0 of R2 to meet test requirements.

```
[R2]undo interface LoopBack 0
```

R2 is restarted because of an accident. Here, you restart an OSPF process to simulate the restart of R2.

```
<R2>reset ospf process
Warning: The OSPF process will be reset. Continue? [Y/N]:y
```

Users connected to R1 find that they cannot access addresses outside Area 2. An administrator logs in to R1 and finds that R1 cannot communicate with R5 using the loopback interface address.

```
[R1]ping -c 1 10.0.5.5
  PING 10.0.5.5: 56   data bytes, press CTRL_C to break
    Request time out

  --- 10.0.5.5 ping statistics ---
    1 packet(s) transmitted
    0 packet(s) received
100.00% packet loss
```

Check the virtual link between R2 and R3. You can see that the virtual link status is not normal and the router ID of R2 changes.

```
[R2]display ospf vlink

         OSPF Process 1 with Router ID 10.0.23.2
                 Virtual Links

 Virtual-link Neighbor-id   -> 10.0.3.3, Neighbor-State: Down

 Interface: 10.0.23.2 (Serial2/0/0)
 Cost: 1562   State: P-2-P   Type: Virtual
 Transit Area: 0.0.0.1
 Timers: Hello 10 , Dead 40 , Retransmit 5 , Transmit Delay 1
 GR State: Normal
```

A virtual link is established based on the router ID of the peer device. The router ID of R2 changes, so the virtual link fails.

Generally, to prevent a router ID from changing during the operation of a router, you need to specify a router ID for this router when starting an OSPF process.

On R2, set the router ID to 10.0.2.2, add the address of Loopback0, and then restart the OSPF process.

```
[R2]ospf 1 router-id 10.0.2.2
Info: The configuration succeeded. You need to restart the OSPF process to validate the new router ID.
[R2-ospf-1]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24
[R2-LoopBack0]quit
<R2>reset ospf process
Warning: The OSPF process will be reset. Continue? [Y/N]:y
```

Check the virtual link status again.

```
[R2]display ospf vlink

          OSPF Process 1 with Router ID 10.0.2.2
                  Virtual Links

 Virtual-link Neighbor-id  -> 10.0.3.3, Neighbor-State: Full

 Interface: 10.0.23.2 (Serial2/0/0)
 Cost: 1562  State: P-2-P  Type: Virtual
 Transit Area: 0.0.0.1
 Timers: Hello 10 , Dead 40 , Retransmit 5 , Transmit Delay 1
 GR State: Normal
```

The virtual link has recovered.

For security, the administrator uses area authentication in Area 0, enable MD5 encryption to encrypt packets, and set the key to huawei.

```
[R3]ospf 1
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]authentication-mode md5 1 plain huawei
[R3-ospf-1-area-0.0.0.0]quit
[R3-ospf-1]quit

[R4]ospf 1
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]authentication-mode md5 1 plain huawei
[R4-ospf-1-area-0.0.0.0]quit
[R4-ospf-1]quit

[R5]ospf 1
[R5-ospf-1]area 0
[R5-ospf-1-area-0.0.0.0]authentication-mode md5 1 plain huawei
[R5-ospf-1-area-0.0.0.0]quit
[R5-ospf-1]quit
```

The administrator finds that users in Area 2 cannot access networks outside Area 2 and then check the virtual link, finding that the virtual link fails again.

```
[R2]display ospf vlink

         OSPF Process 1 with Router ID 10.0.2.2
                  Virtual Links
```

Virtual-link Neighbor-id -> 10.0.3.3, Neighbor-State: Down

Interface: 10.0.23.2 (Serial2/0/0)

Cost: 1562 State: P-2-P Type: Virtual

Transit Area: 0.0.0.1

Timers: Hello 10 , Dead 40 , Retransmit 5 , Transmit Delay 1

Clear OSPF statistics and then check OSPF errors. The following command output shows that authentication errors occur.

<R2>reset ospf counters

<R2>display ospf error

OSPF Process 1 with Router ID 10.0.2.2

OSPF error statistics

General packet errors:

0	: IP: received my own packet	7	: Bad packet
0	: Bad version	0	: Bad checksum
0	: Bad area id	0	: Drop on unnumbered interface
0	: Bad virtual link	7	: Bad authentication type
9	: Bad authentication key	0	: Packet too small
0	: Packet size > ip length	0	: Transmit error
0	: Interface down	0	: Unknown neighbor
0	: Bad net segment	0	: Extern option mismatch
0	: Router id confusion	0	: Bad authentication sequence number

The virtual link belongs to Area 0. Area authentication is enabled in Area 0, so area authentication also needs to be enabled on the virtual link.

```
[R2]ospf 1
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]authentication-mode md5 1 plain huawei
[R2-ospf-1-area-0.0.0.0]quit
[R2-ospf-1]quit
```

The preceding command output shows that the virtual link status becomes normal and R1 can access other areas normally.

```
[R2]display ospf vlink

          OSPF Process 1 with Router ID 10.0.2.2
                  Virtual Links

 Virtual-link Neighbor-id  -> 10.0.3.3, Neighbor-State: Full

 Interface: 10.0.23.2 (Serial2/0/0)
 Cost: 1562  State: P-2-P  Type: Virtual
 Transit Area: 0.0.0.1
 Timers: Hello 10 , Dead 40 , Retransmit 5 , Transmit Delay 1
 GR State: Normal

[R1]ping -c 1 10.0.5.5
  PING 10.0.5.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.5.5: bytes=56 Sequence=1 ttl=253 time=73 ms

  --- 10.0.5.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
```

0.00% packet loss

round-trip min/avg/max = 73/73/73 ms

Step 6 **Troubleshoot OSPF route summarization failures.**

On R4, import the address of Loopback0 as an external route and configure route summarization using the 16-bit subnet mask.

```
[R4]ospf 1
[R4-ospf-1]import-route direct
[R4-ospf-1]asbr-summary 10.1.0.0 255.255.0.0
[R4-ospf-1]quit
```

After a certain period, the administrator configures inter-area route summarization on R2 and summarizes the network segment connected to Loopback0 of R1 into a route with a 16-bit mask.

```
[R2]ospf 1
[R2-ospf-1]area 2
[R2-ospf-1-area-0.0.0.2]abr-summary 10.1.0.0 255.255.0.0
[R2-ospf-1-area-0.0.0.2]quit
[R2-ospf-1]quit
```

All users on the network except users connected to R4 reflect that they cannot access the loopback interface address 10.1.4.4 of R4.

Check the OSPF routing of R5 that is located in the same area as R4. The following command output shows that to reach 10.1.4.4, the route 10.1.0.0/16 must be used. The next hop of this route is 10.0.75.3.

Why is this incorrect route generated?

```
[R5]display ip routing-table protocol ospf
```

Route Flags: R - relay, D - download to fib

Public routing table : OSPF

Destinations : 5 Routes : 5

OSPF routing table status : <Active>

Destinations : 5 Routes : 5

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
10.0.2.0/24	OSPF	10	1563	D	10.0.75.3	GigabitEthernet0/0/0
10.0.3.0/24	OSPF	10	1	D	10.0.75.3	GigabitEthernet0/0/0
10.0.12.0/24	OSPF	10	3125	D	10.0.75.3	GigabitEthernet0/0/0
10.0.23.0/24	OSPF	10	1563	D	10.0.75.3	GigabitEthernet0/0/0
10.1.0.0/16	OSPF	10	3125	D	10.0.75.3	GigabitEthernet0/0/0

OSPF routing table status : <Inactive>

Destinations : 0 Routes : 0

Check the LSDB of R5.

[R5]display ospf lsdb

OSPF Process 1 with Router ID 10.0.5.5

Link State Database

Area: 0.0.0.0

Type	LinkState ID	AdvRouter	Age	Len	Sequence	Metric
Router	10.0.5.5	10.0.5.5	214	48	80000025	0

Router	10.0.3.3	10.0.3.3	1246	48	80000024	1
Router	10.0.2.2	10.0.2.2	1247	36	80000005	1562
Router	10.1.4.4	10.1.4.4	648	36	8000000D	1
Network	10.0.75.4	10.1.4.4	206	36	80000004	0
Sum-Net	10.0.12.0	10.0.2.2	916	28	80000002	1562
Sum-Net	10.0.3.0	10.0.3.3	893	28	80000008	0
Sum-Net	10.0.3.0	10.0.2.2	916	28	80000002	1562
Sum-Net	10.0.2.0	10.0.3.3	919	28	80000003	1562
Sum-Net	10.0.2.0	10.0.2.2	916	28	80000002	0
Sum-Net	10.1.0.0	10.0.2.2	538	28	80000001	1562
Sum-Net	10.0.23.0	10.0.3.3	893	28	80000008	1562
Sum-Net	10.0.23.0	10.0.2.2	917	28	80000002	1562

AS External Database

Type	LinkState ID	AdvRouter	Age	Len	Sequence	Metric
External	10.0.75.0	10.1.4.4	649	36	80000001	1
External	10.1.0.0	10.1.4.4	620	36	80000001	2

The LSDB of R5 has two LSAs describing the route to 10.1.0.0. Check detailed information about LSAs. The following Type 3 LSA is originated by R2, and the Type 5 LSA is originated by R5. The two LSAs describe the same network segment.

[R5]display ospf lsdb summary 10.1.0.0

OSPF Process 1 with Router ID 10.0.5.5

Area: 0.0.0.0

Link State Database

```
  Type         : Sum-Net
  Ls id        : 10.1.0.0
  Adv rtr      : 10.0.2.2
  Ls age       : 767
  Len          : 28
  Options      :   E
  seq#         : 80000001
  chksum       : 0xa380
  Net mask     : 255.255.0.0
  Tos 0   metric: 1562
  Priority     : Low

[R5]display ospf lsdb ase 10.1.0.0

          OSPF Process 1 with Router ID 10.0.5.5
                  Link State Database

  Type         : External
  Ls id        : 10.1.0.0
  Adv rtr      : 10.1.4.4
  Ls age       : 871
  Len          : 36
  Options      :   E
  seq#         : 80000001
  chksum       : 0xe3cd
  Net mask     : 255.255.0.0
  TOS 0   Metric: 2
  E type       : 2
  Forwarding Address : 0.0.0.0
```

```
Tag         : 1
Priority    : Low
```

In OSPF, Type 3 LSAs are always preferred over Type 5 LSAs. Therefore, in the OSPF routing table of R5, the next hop of the route to 10.1.0.0/16 is R3.

To prevent this problem, cancel external route summarization. This route then will appear in the OSPF routing tables of other routers.

```
[R4]ospf 1
[R4-ospf-1]undo asbr-summary 10.1.0.0 255.255.0.0
[R4-ospf-1]quit

[R5]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 6         Routes : 6

OSPF routing table status : <Active>
         Destinations : 6         Routes : 6

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.2.0/24  OSPF    10   1563        D   10.0.75.3       GigabitEthernet0/0/0
       10.0.3.0/24  OSPF    10   1           D   10.0.75.3       GigabitEthernet0/0/0
      10.0.12.0/24  OSPF    10   3125        D   10.0.75.3       GigabitEthernet0/0/0
      10.0.23.0/24  OSPF    10   1563        D   10.0.75.3       GigabitEthernet0/0/0
       10.1.0.0/16  OSPF    10   3125        D   10.0.75.3       GigabitEthernet0/0/0
       10.1.4.4/24  O_ASE   150  1           D   10.0.75.4       GigabitEthernet0/0/0
```

```
OSPF routing table status : <Inactive>
         Destinations : 0           Routes : 0
```

The preceding command output shows that R5 learns a correct route to 10.1.4.4/24.

Test network connectivity on R1.

```
[R1]ping -c 1 10.1.4.4
  PING 10.1.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.1.4.4: bytes=56 Sequence=1 ttl=253 time=71 ms

  --- 10.1.4.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 71/71/71 ms
```

The preceding command output shows that the network recovers.

----End

Additional Exercises: Analysis and Verification

Can area authentication and interface authentication be enabled in the same area?

Can area IDs of non-backbone areas be the same?

Device Configurations

```
<R1>display current-configuration
[V200R007C00SPC600]
#
 sysname R1
```

```
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
 ospf authentication-mode md5 1 plain huawei
#
interface LoopBack0
 ip address 10.1.1.1 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.1.1.1
 area 0.0.0.2
  network 10.0.12.1 0.0.0.0
  network 10.1.1.1 0.0.0.0
#
return

<R2>display current-configuration
[V200R007C00SPC600]
#
 sysname R2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
 ospf authentication-mode md5 1 plain huawei
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.0
```

```
 ospf network-type broadcast
#
ospf 1 router-id 10.0.2.2
 area 0.0.0.0
  authentication-mode md5 1 plain huawei
 area 0.0.0.1
  network 10.0.23.2 0.0.0.0
  network 10.0.2.2 0.0.0.0
  vlink-peer 10.0.3.3
 area 0.0.0.2
  abr-summary 10.1.0.0 255.255.0.0
  network 10.0.12.2 0.0.0.0
#
return

<R3>display current-configuration
[V200R007C00SPC600]
#
 sysname R3
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.75.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.0
```

```
 ospf network-type broadcast
#
ospf 1 router-id 10.0.3.3
 area 0.0.0.0
  authentication-mode md5 1 plain huawei
  network 10.0.75.3 0.0.0.0
 area 0.0.0.1
  network 10.0.23.3 0.0.0.0
  network 10.0.3.3 0.0.0.0
  vlink-peer 10.0.2.2
#
return

<R4>display current-configuration
[V200R007C00SPC600]
#
 sysname R4
#
interface GigabitEthernet0/0/0
 ip address 10.0.75.4 255.255.255.0
#
interface LoopBack0
 ip address 10.1.4.4 255.255.255.0
#
ospf 1 router-id 10.1.4.4
 import-route direct
 area 0.0.0.0
  authentication-mode md5 1 plain huawei
  network 10.0.75.4 0.0.0.0
```

```
#
return

<R5>display current-configuration
[V200R007C00SPC600]
#
 sysname R5
#
interface GigabitEthernet0/0/0
 ip address 10.0.75.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.5.5
 area 0.0.0.0
  authentication-mode md5 1 plain huawei
  network 10.0.75.5 0.0.0.0
  network 10.0.5.5 0.0.0.0
#
Return
```

Chapter 2 IS-IS Features and Configurations

Lab 2-1 IS-IS Configurations

Learning Objectives

The objectives of this lab are to learn and understand:

- Basic IS-IS configurations
- How to configure the IS-IS designated intermediate system (DIS) priority
- How to configure the IS-IS network type
- How to import external routes into IS-IS
- How to configure the IS-IS interface cost
- How to configure IS-IS route leaking

Topology

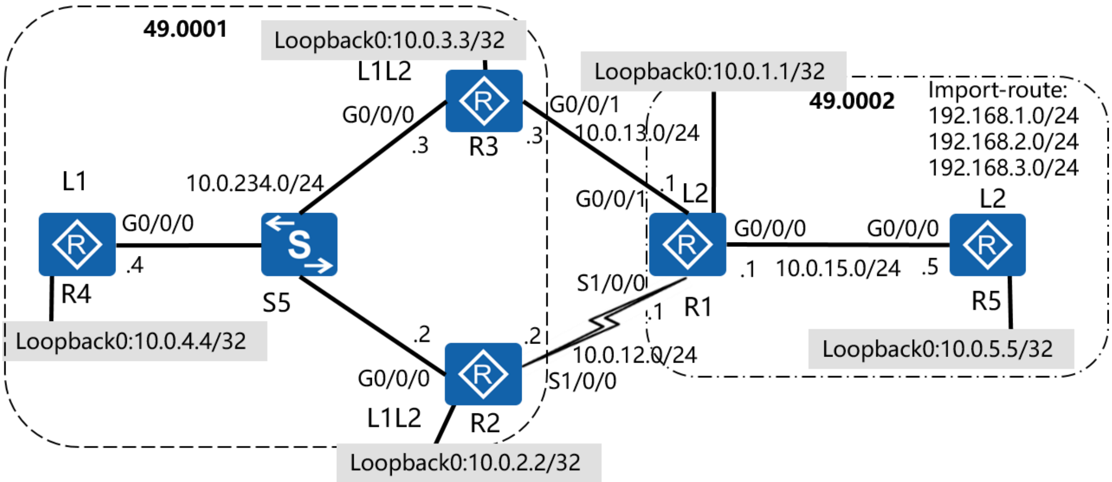

Figure 3-1 IS-IS topology

Scenario

IS-IS runs as an IGP in a network. R1 and R5 run in Area 49.0002 as Level-2 routers. R2, R3, and R4 run in Area 49.0001. R4 is a Level-1 router, while R2 and R3 are Level-1-2 routers. The requirements are as follows: R4 can use the external routes imported by R5. R4 connected to S5 functions as the DIS. There is a P2P link between R1 and R5. Uplink and downlink traffic from R4 to R5 is forwarded through Ethernet interfaces, and route selection is controlled using the cost and route leaking. Switches do not require additional configurations and are only responsible for transparent forwarding.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

Configure IP addresses for all the routers.

```
[R1]interface loopback 0
[R1-LoopBack0]ip address 10.0.1.1 32
[R1-LoopBack0]quit
[R1]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]ip address 10.0.15.1 24
[R1-GigabitEthernet0/0/0]quit
[R1]interface GigabitEthernet 0/0/1
[R1-GigabitEthernet0/0/1]ip address 10.0.13.1 24
[R1-GigabitEthernet0/0/1]quit
[R1]interface interface Serial1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]quit

[R2]interface   LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 32
[R2-LoopBack0]quit
[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.234.2 24
[R2-GigabitEthernet0/0/0]quit
[R2]interface Serial1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]quit

[R3]interface   LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 32
```

```
[R3-LoopBack0]quit
[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]ip address 10.0.234.3 24
[R3-GigabitEthernet0/0/0]quit
[R3]interface GigabitEthernet 0/0/1
[R3-GigabitEthernet0/0/1]ip address 10.0.13.3 24
[R3-GigabitEthernet0/0/1]quit

[R4]interface  LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 32
[R4-LoopBack0]quit
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ip address 10.0.234.4 24
[R4-GigabitEthernet0/0/0]quit

[R5]interface  LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 32
[R5-LoopBack0]quit
[R5]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.15.5 24
[R5-GigabitEthernet0/0/0]quit
```

After the configurations are complete, test connectivity from R1 to directly connected interfaces of R2, R3, and R5.

```
[R1]ping -c 1 10.0.13.3
  PING 10.0.13.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.13.3: bytes=56 Sequence=1 ttl=255 time=1 ms
```

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

[R1]ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=26 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 26/26/26 ms

[R1]ping -c 1 10.0.15.5
  PING 10.0.15.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.15.5: bytes=56 Sequence=1 ttl=255 time=1 ms

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

Test connectivity from R4 to directly connected interfaces of R2 and R3.

```
[R4]ping -c 1 10.0.234.2
```

```
  PING 10.0.234.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.234.2: bytes=56 Sequence=1 ttl=255 time=1 ms

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

[R4]ping -c 1 10.0.234.3
  PING 10.0.234.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.234.3: bytes=56 Sequence=1 ttl=255 time=1 ms

  --- 10.0.234.3 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 **Perform basic IS-IS configurations.**

Configure an IS-IS process 1 for each router according to the topology. The following example configures IS-IS process 1 for R1.

R1 resides in Area 49.0002 and uses network-entity 49.0002.0000.0000.0001.

```
[R1]isis
[R1-isis-1]network-entity 49.0002.0000.0000.0001.00
```

By default, after an IS-IS process is enabled on a router, the router works in Level-1-2 mode. According to the planning, R1 needs to work in Level-2 mode. Therefore, you need to change its IS level.

```
[R1-isis-1]is-level level-2
[R1-isis-1]quit
```

Enable IS-IS on related interfaces, including loopback interfaces. When no IS-IS process ID is specified, by default, IS-IS is enabled in IS-IS process 1.

```
[R1]interface LoopBack 0
[R1-LoopBack0]isis enable
[R1-LoopBack0]quit
[R1]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]isis enable
[R1-GigabitEthernet0/0/0]quit
[R1]interface GigabitEthernet 0/0/1
[R1-GigabitEthernet0/0/1]isis enable
[R1-GigabitEthernet0/0/1]quit
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]isis enable
[R1-Serial1/0/0]quit
```

On R1, check IS-IS enabling information on interfaces.

```
[R1]display isis interface

                    Interface information for ISIS(1)
                    ---------------------------------
 Interface       Id      IPV4.State          IPV6.State      MTU  Type  DIS
 Loop0             001          Up                  Down          1500 L1/L2 --
```

```
GE0/0/0    001    Up    Down    1497 L1/L2 No/No
GE0/0/1    002    Up    Down    1497 L1/L2 No/No
S1/0/0     002    Up    Down    1500 L1/L2 --
```

The preceding command output shows that ISIS(1) has been enabled on a total of four interfaces, whose **IPV4.State** field displays Up.

Similarly, configure other routers. R2 and R3 work in Level-1-2 mode, so you do not need to change their IS levels.

```
[R2]isis 1
[R2-isis-1]network-entity 49.0001.0000.0000.0002.00
[R2-isis-1]quit
[R2]interface LoopBack 0
[R2-LoopBack0]isis enable
[R2-LoopBack0]quit
[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]isis enable
[R2-GigabitEthernet0/0/0]quit
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]isis enable
[R2-Serial1/0/0]quit
```

On R2, check IS-IS enabling information on interfaces.

```
[R2]display isis interface

                    Interface information for ISIS(1)
                    --------------------------------
 Interface      Id      IPV4.State        IPV6.State   MTU  Type  DIS
Loop0           002       Up                Down        1500 L1/L2 --
```

```
GE0/0/0          001         Up              Down          1497 L1/L2 No/No
S1/0/0           001         Up              Down          1500 L1/L2 --
```

Configure IS-IS on R3.

```
[R3]isis 1
[R3-isis-1]network-entity 49.0001.0000.0000.0003.00
[R3-isis-1]quit
[R3]interface LoopBack 0
[R3-LoopBack0]isis enable
[R3-LoopBack0]quit
[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]isis enable
[R3-GigabitEthernet0/0/0]quit
[R3]interface GigabitEthernet 0/0/1
[R3-GigabitEthernet0/0/1]isis enable
[R3-GigabitEthernet0/0/1]quit
```

On R3, check IS-IS enabling information on interfaces.

```
[R3]display isis interface

                      Interface information for ISIS(1)
                      ---------------------------------
 Interface        Id     IPV4.State         IPV6.State      MTU  Type  DIS
 Loop0       001         Up                 Down          1500 L1/L2 --
GE0/0/0      002         Up                 Down          1497 L1/L2 No/No
 GE0/0/1     001         Up                 Down          1497 L1/L2 No/No
```

Configure IS-IS on R4.

[R4]isis 1

[R4-isis-1]network-entity 49.0001.0000.0000.0004.00

[R4-isis-1]is-level level-1

[R4-isis-1]quit

[R4]interface LoopBack 0

[R4-LoopBack0]isis enable

[R4-LoopBack0]quit

[R4]interface GigabitEthernet 0/0/0

[R4-GigabitEthernet0/0/0]isis enable

[R4-GigabitEthernet0/0/0]quit

On R4, check IS-IS enabling information on interfaces.

[R4]display isis interface

Interface information for ISIS(1)

Interface	Id	IPV4.State	IPV6.State	MTU	Type	DIS
Loop0	001	Up	Down	1500	L1/L2	--
GE0/0/0	001	Up	Down	1497	L1/L2	No/No

Configure IS-IS on R5.

[R5]isis 1

[R5-isis-1]network-entity 49.0002.0000.0000.0005.00

[R5-isis-1]is-level level-2

[R5-isis-1]quit

[R5]interface LoopBack 0

[R5-LoopBack0]isis enable

[R5-LoopBack0]quit

[R5]interface GigabitEthernet 0/0/0

[R5-GigabitEthernet0/0/0]isis enable

[R5-GigabitEthernet0/0/0]quit

On R5, check IS-IS enabling information on interfaces.

[R5]display isis interface

```
                        Interface information for ISIS(1)
                        ---------------------------------
Interface    Id     IPV4.State          IPV6.State    MTU   Type   DIS
 Loop0       001        Up                  Down          1500 L1/L2   --
 GE0/0/0     001        Up                  Down          1497 L1/L2 No/No
```

After the configurations are complete, check IS-IS neighbor statuses of routers. The following example displays the IS-IS neighbor status of R1. R1 has three neighbors: R2, R3, and R5.

[R1]display isis peer

```
                          Peer information for ISIS(1)

  System Id        Interface          Circuit Id        State HoldTime Type      PRI
------------------------------------------------------------------------------
0000.0000.0005  GE0/0/0           0000.0000.0005.01 Up     7s        L2        64
0000.0000.0003  GE0/0/1           0000.0000.0001.02 Up     21s       L2        64
0000.0000.0002  S1/0/0            0000000001        Up     28s       L2        --
```

```
Total Peer(s): 3
```

In the preceding command output, the **System Id** field is similar to the **Router Id** field of other routing protocols. You can see that R2, R3, and R5 are in Up state. Their IS-IS neighbor relationships with R1 are normal.

Continue to check IS-IS neighbor statuses of other devices.

```
[R2]display isis peer

                          Peer information for ISIS(1)

  System Id       Interface          Circuit Id         State HoldTime Type        PRI
-------------------------------------------------------------------------------
0000.0000.0001   S1/0/0             0000000001           Up    22s      L2            --
0000.0000.0003   GE0/0/0            0000.0000.0004.01 Up    24s      L1(L1L2)  64
0000.0000.0004   GE0/0/0            0000.0000.0004.01 Up    7s       L1            64
0000.0000.0003   GE0/0/0            0000.0000.0002.01 Up    26s      L2(L1L2)  64

Total Peer(s): 4

[R3]display isis peer

                          Peer information for ISIS(1)

  System Id       Interface          Circuit Id         State HoldTime Type        PRI
-------------------------------------------------------------------------------
0000.0000.0001   GE0/0/1            0000.0000.0001.02 Up    8s       L2            64
0000.0000.0002   GE0/0/0            0000.0000.0004.01 Up    30s      L1(L1L2)      64
0000.0000.0004   GE0/0/0            0000.0000.0004.01 Up    7s       L1            64
```

```
0000.0000.0002   GE0/0/0            0000.0000.0002.01 Up     9s        L2(L1L2)  64

Total Peer(s): 4

[R4]display isis peer

                          Peer information for ISIS(1)

  System Id       Interface           Circuit Id          State HoldTime Type       PRI
--------------------------------------------------------------------------
0000.0000.0003   GE0/0/0            0000.0000.0004.01 Up     29s       L1          64
0000.0000.0002   GE0/0/0            0000.0000.0004.01 Up     23s       L1          64

Total Peer(s): 2

[R5]display isis peer

                          Peer information for ISIS(1)

  System Id       Interface           Circuit Id          State HoldTime Type       PRI
--------------------------------------------------------------------------
0000.0000.0001   GE0/0/0            0000.0000.0005.01 Up     29s       L2          64

Total Peer(s): 1
```

Step 3 **Change the DIS priority.**

R2, R3, and R4 establish IS-IS neighbor relationships in a broadcast network. Therefore, the DIS needs to be elected. By default, the DIS priority is 64. If these routers have the same DIS priority, the router with the highest MAC address becomes the DIS. To ensure that R4 becomes the DIS, change its DIS priority.

```
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]isis dis-priority 120
[R4-GigabitEthernet0/0/0]quit

[R4]display isis interface

                      Interface information for ISIS(1)
                      ---------------------------------
 Interface     Id      IPV4.State        IPV6.State    MTU  Type  DIS
 GE0/0/0        001        Up                Down          1497 L1/L2 Yes/No
 Loop0          001        Up                Down          1500 L1/L2 –
```

Check DIS priorities of neighbors on R2 and R3.

```
[R2]display isis peer

                         Peer information for ISIS(1)

  System Id      Interface           Circuit Id        State HoldTime Type       PRI
-------------------------------------------------------------------------------
0000.0000.0001   S1/0/0              0000000001           Up    29s       L2             --
0000.0000.0003   GE0/0/0             0000.0000.0004.01 Up    25s       L1(L1L2)       64
0000.0000.0004   GE0/0/0             0000.0000.0004.01 Up    8s        L1             120
0000.0000.0003   GE0/0/0             0000.0000.0002.01 Up    20s       L2(L1L2)       64

Total Peer(s): 4
```

```
[R3]display isis peer

                          Peer information for ISIS(1)

  System Id      Interface          Circuit Id         State HoldTime Type        PRI
-------------------------------------------------------------------------------
0000.0000.0001   GE0/0/1            0000.0000.0001.02 Up   8s       L2            64
0000.0000.0002   GE0/0/0            0000.0000.0004.01 Up   22s      L1(L1L2)  64
0000.0000.0004   GE0/0/0            0000.0000.0004.01 Up   7s       L1            120
0000.0000.0002   GE0/0/0            0000.0000.0002.01 Up   8s       L2(L1L2)  64

Total Peer(s): 4
```

Step 4 **Configure the IS-IS network type.**

In a broadcast network, by default, an IS-IS router sets the circuit-type of interfaces to broadcast and participates in DIS election. In the topology, the Ethernet between R1 and R5 has only two routers. You can set the circuit-type of interfaces between the two routers to P2P for optimization.

```
[R1]interface  GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]isis circuit-type p2p
[R1-GigabitEthernet0/0/0]quit

[R5]interface  GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]isis circuit-type p2p
[R5-GigabitEthernet0/0/0]quit
```

During the change of the circuit-type, neighbor relationships are established again. Check the configuration. The **Circuit Id** field format changes and the **Circuit Parameters** field displays **p2p**. The following uses the display of R1 as an example.

```
[R1]display isis peer

                          Peer information for ISIS(1)

  System Id       Interface           Circuit Id        State HoldTime Type      PRI
-------------------------------------------------------------------------------
0000.0000.0005   GE0/0/0              0000000002             Up    22s       L2        --
0000.0000.0003   GE0/0/1              0000.0000.0001.02      Up    27s       L2        64
0000.0000.0002   S1/0/0               0000000001             Up    22s       L2        --

[R1]display isis interface GigabitEthernet 0/0/0 verbose

                       Interface information for ISIS(1)
                       ---------------------------------
 Interface         Id        IPV4.State            IPV6.State        MTU   Type   DIS
 GE0/0/0            003          Up                     Down            1497 L1/L2 --
  Circuit MT State                : Standard
  Circuit Parameters              : p2p
  Description                     : HUAWEI, AR Series, GigabitEthernet0/0/0 Interface
  SNPA Address                        : d0d0-4b03-d3fc
  IP Address                      : 10.0.15.1
  IPV6 Link Local Address         :
  IPV6 Global Address(es)         :
  Csnp Timer Value                :   L12     10
  Hello Timer Value                   :           10
  DIS Hello Timer Value           :
```

```
Hello Multiplier Value :        3
Cost                        :  L1      10   L2      10
Ipv6 Cost                   :  L1      10   L2      10
Retransmit Timer Value    :  L12    5
LSP-Throttle Timer        :  L12    50
Bandwidth-Value           :  Low    100000000   High           0
Static Bfd                :  NO
Dynamic Bfd                  :  NO
Fast-Sense Rpr            :  NO
Extended-Circuit-Id Value  :  0000000003
```

Step 5 **Configure IS-IS to import external routes.**

Before importing external routes into IS-IS, check current route learning. The following command output shows that traffic from R1 to R4 is load balanced between GE0/0/1 and S1/0/0.

```
[R1]display isis route

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

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

IPV4 Destination       IntCost      ExtCost ExitInterface     NextHop            Flags
-------------------------------------------------------------------------
10.0.4.4/32              20             NULL    GE0/0/1              10.0.13.3        A/-/-/-
                                                S1/0/0               10.0.12.2
10.0.5.5/32              10             NULL    GE0/0/0              10.0.15.5        A/-/-/-
```

```
10.0.12.0/24     10     NULL   S1/0/0       Direct       D/-/L/-
10.0.13.0/24     10     NULL   GE0/0/1      Direct       D/-/L/-
10.0.234.0/24    20     NULL   S1/0/0       10.0.12.2    A/-/-/-
                               GE0/0/1      10.0.13.3
10.0.15.0/24     10     NULL   GE0/0/0      Direct       D/-/L/-
10.0.1.1/32      0      NULL   Loop0        Direct       D/-/L/-
10.0.2.2/32      10     NULL   S1/0/0       10.0.12.2    A/-/-/-
10.0.3.3/32      10     NULL   GE0/0/1      10.0.13.3    A/-/-/-
     Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
                              U-Up/Down Bit Set

[R1]display ip routing-table protocol isis
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : ISIS
        Destinations : 5        Routes : 7

ISIS routing table status : <Active>
        Destinations : 5        Routes : 7

Destination/Mask    Proto    Pre   Cost   Flags NextHop      Interface

     10.0.2.2/32      ISIS-L2  15    10        D   10.0.12.2  Serial1/0/0
     10.0.3.3/32      ISIS-L2  15    10        D   10.0.13.3  GigabitEthernet0/0/1
     10.0.4.4/32      ISIS-L2  15    20        D   10.0.13.3  GigabitEthernet0/0/1
                      ISIS-L2  15    20        D   10.0.12.2  Serial1/0/0
     10.0.5.5/32      ISIS-L2  15    10        D   10.0.15.5  GigabitEthernet0/0/0
     10.0.234.0/24    ISIS-L2  15    20        D   10.0.12.2  Serial1/0/0
                      ISIS-L2  15    20        D   10.0.13.3  GigabitEthernet0/0/1
```

ISIS routing table status : <Inactive>

Destinations : 0 Routes : 0

R2 is a Level-1-2 router and so generates different routes for Level-1 and Level-2 routers. For Level-1 router, it generates a default route pointing to the null interface. This situation also exists on R3.

[R2]display isis route

Route information for ISIS(1)

ISIS(1) Level-1 Forwarding Table

IPV4 Destination	IntCost	ExtCost	ExitInterface	NextHop	Flags
0.0.0.0/0	10	NULL			
10.0.4.4/32	10	NULL	GE0/0/0	10.0.234.4	A/-/L/-
10.0.12.0/24	10	NULL	S1/0/0	Direct	D/-/L/-
10.0.13.0/24	20	NULL	GE0/0/0	10.0.234.3	A/-/L/-
10.0.234.0/24	10	NULL	GE0/0/0	Direct	D/-/L/-
10.0.2.2/32	0	NULL	Loop0	Direct	D/-/L/-
10.0.3.3/32	10	NULL	GE0/0/0	10.0.234.3	A/-/L/-

Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,

U-Up/Down Bit Set

ISIS(1) Level-2 Forwarding Table

```
                        ------------------------------

IPV4 Destination      IntCost     ExtCost ExitInterface     NextHop          Flags
--------------------------------------------------------------------------
10.0.4.4/32              20          NULL
10.0.5.5/32              20          NULL       S1/0/0              10.0.12.1        A/-/-/-
10.0.12.0/24             10          NULL       S1/0/0              Direct           D/-/L/-
10.0.13.0/24             20          NULL
10.0.234.0/24            10          NULL       GE0/0/0             Direct           D/-/L/-
10.0.15.0/24             20          NULL       S1/0/0              10.0.12.1        A/-/-/-
10.0.1.1/32              10          NULL       S1/0/0              10.0.12.1        A/-/-/-
10.0.2.2/32              0           NULL       Loop0               Direct           D/-/L/-
10.0.3.3/32              10          NULL
     Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
                              U-Up/Down Bit Set

[R2]display ip routing-table protocol isis
Route Flags: R - relay, D - download to fib
--------------------------------------------------------------------------
Public routing table : ISIS
         Destinations : 6          Routes : 6

ISIS routing table status : <Active>
         Destinations : 6          Routes : 6

Destination/Mask     Proto    Pre   Cost        Flags NextHop        Interface

       10.0.1.1/32   ISIS-L2   15    10          D    10.0.12.1       Serial1/0/0
       10.0.3.3/32   ISIS-L1   15    10          D    10.0.234.3      GigabitEthernet0/0/0
```

```
     10.0.4.4/32   ISIS-L1   15    10       D    10.0.234.4      GigabitEthernet0/0/0
     10.0.5.5/32   ISIS-L2   15    20       D    10.0.12.1       Serial1/0/0
    10.0.13.0/24   ISIS-L1   5     20       D    10.0.234.3      GigabitEthernet0/0/0
    10.0.15.0/24   ISIS-L2   15    20       D    10.0.12.1       Serial1/0/0

ISIS routing table status : <Inactive>
         Destinations : 0          Routes : 0

[R3]display isis route

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

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

IPV4 Destination      IntCost     ExtCost ExitInterface    NextHop           Flags
--------------------------------------------------------------------------
0.0.0.0/0            10          NULL
10.0.4.4/32          10          NULL      GE0/0/0           10.0.234.4          A/-/L/-
10.0.12.0/24         20          NULL      GE0/0/0           10.0.234.2          A/-/L/-
10.0.13.0/24         10          NULL      GE0/0/1           Direct              D/-/L/-
10.0.234.0/24        10          NULL      GE0/0/0           Direct              D/-/L/-
10.0.2.2/32          10          NULL      GE0/0/0           10.0.234.2          A/-/L/-
10.0.3.3/32          0           NULL      Loop0             Direct              D/-/L/-
     Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
                              U-Up/Down Bit Set

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

```
                              -------------------------------

IPV4 Destination      IntCost      ExtCost ExitInterface     NextHop            Flags
-------------------------------------------------------------------------
10.0.4.4/32              20            NULL
10.0.5.5/32              20            NULL      GE0/0/1           10.0.13.1         A/-/-/-
10.0.12.0/24             20            NULL
10.0.13.0/24             10            NULL      GE0/0/1           Direct            D/-/L/-
10.0.234.0/24            10            NULL      GE0/0/0           Direct            D/-/L/-
10.0.15.0/24             20            NULL      GE0/0/1           10.0.13.1         A/-/-/-
10.0.1.1/32              10            NULL      GE0/0/1           10.0.13.1         A/-/-/-
10.0.2.2/32              10            NULL
10.0.3.3/32              0             NULL      Loop0             Direct            D/-/L/-
     Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
                              U-Up/Down Bit Set

[R3]display ip routing-table protocol isis
Route Flags: R - relay, D - download to fib
-------------------------------------------------------------------------
Public routing table : ISIS
         Destinations : 6          Routes : 6

ISIS routing table status : <Active>
         Destinations : 6          Routes : 6

Destination/Mask      Proto    Pre   Cost        Flags NextHop          Interface

    10.0.1.1/32   ISIS-L2 15    10          D    10.0.13.1   GigabitEthernet0/0/1
    10.0.2.2/32   ISIS-L1 15    10          D    10.0.234.2 GigabitEthernet0/0/0
```

```
10.0.4.4/32   ISIS-L1 15   10      D   10.0.234.4 GigabitEthernet0/0/0
10.0.5.5/32   ISIS-L2 15   20      D   10.0.13.1  GigabitEthernet0/0/1
10.0.12.0/24  ISIS-L1 15   20      D   10.0.234.2 GigabitEthernet0/0/0
10.0.15.0/24  ISIS-L2 15   20      D   10.0.13.1  GigabitEthernet0/0/1

ISIS routing table status : <Inactive>
         Destinations : 0          Routes : 0
```

R4 is a Level-1 router and can only establish IS-IS neighbor relationships with Level-1 or Level-1-2 routers in the same area. By default, a Level-1 router cannot learn routing information of Level-2 routers and can only access external networks through default routes. The following command output shows that R4 has two default routes pointing to R2 and R3, and the two routes work in load balancing mode.

```
[R4]display isis route

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

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

IPV4 Destination      IntCost     ExtCost ExitInterface    NextHop          Flags
-------------------------------------------------------------------------------
0.0.0.0/0             10          NULL    GE0/0/0          10.0.234.3       A/-/-/-
                                          GE0/0/0          10.0.234.2
10.0.4.4/32           0           NULL    Loop0            Direct           D/-/L/-
10.0.12.0/24          20          NULL    GE0/0/0          10.0.234.2       A/-/-/-
10.0.13.0/24          20          NULL    GE0/0/0          10.0.234.3       A/-/-/-
```

```
10.0.234.0/24     10        NULL    GE0/0/0        Direct           D/-/L/-
10.0.2.2/32       10        NULL    GE0/0/0        10.0.234.2       A/-/-/-
10.0.3.3/3        10        NULL    GE0/0/0        10.0.234.3       A/-/-/-
     Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
                              U-Up/Down Bit Set

[R4]display ip routing protocol isis
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : ISIS
         Destinations : 5          Routes : 6

ISIS routing table status : <Active>
         Destinations : 5          Routes : 6

Destination/Mask    Proto    Pre  Cost      Flags NextHop          Interface

        0.0.0.0/0   ISIS-L1 15    10          D   10.0.234.3       GigabitEthernet0/0/0
                    ISIS-L1 15    10          D   10.0.234.2       GigabitEthernet0/0/0
       10.0.2.2/32  ISIS-L1 15    10          D   10.0.234.2       GigabitEthernet0/0/0
       10.0.3.3/32  ISIS-L1 15    10          D   10.0.234.3       GigabitEthernet0/0/0
      10.0.12.0/24  ISIS-L1 15    20          D   10.0.234.2       GigabitEthernet0/0/0
      10.0.13.0/24  ISIS-L1 15    20          D   10.0.234.3       GigabitEthernet0/0/0

ISIS routing table status : <Inactive>
         Destinations : 0          Routes : 0
```

Before configuring IS-IS to import external routes, check route learning.

[R5]display isis route

Route information for ISIS(1)

ISIS(1) Level-2 Forwarding Table

IPV4 Destination	IntCost	ExtCost	ExitInterface	NextHop	Flags
10.0.4.4/32	30	NULL	GE0/0/0	10.0.15.1	A/-/-/-
10.0.5.5/32	0	NULL	Loop0	Direct	D/-/L/-
10.0.12.0/24	20	NULL	GE0/0/0	10.0.15.1	A/-/-/-
10.0.13.0/24	20	NULL	GE0/0/0	10.0.15.1	A/-/-/-
10.0.234.0/24	30	NULL	GE0/0/0	10.0.15.1	A/-/-/-
10.0.15.0/24	10	NULL	GE0/0/0	Direct	D/-/L/-
10.0.1.1/32	10	NULL	GE0/0/0	10.0.15.1	A/-/-/-
10.0.2.2/32	20	NULL	GE0/0/0	10.0.15.1	A/-/-/-
10.0.3.3/32	20	NULL	GE0/0/0	10.0.15.1	A/-/-/-

Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,

U-Up/Down Bit Set

[R5]display ip routing-table protocol isis

Route Flags: R - relay, D - download to fib

Public routing table : ISIS

Destinations : 7 Routes : 7

ISIS routing table status : <Active>

Destinations : 7 Routes : 7

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

      10.0.1.1/32  ISIS-L2 15   10           D   10.0.15.1       GigabitEthernet0/0/0
      10.0.2.2/32  ISIS-L2 15   20           D   10.0.15.1       GigabitEthernet0/0/0
      10.0.3.3/32  ISIS-L2 15   20           D   10.0.15.1       GigabitEthernet0/0/0
      10.0.4.4/32  ISIS-L2 15   30           D   10.0.15.1       GigabitEthernet0/0/0
     10.0.12.0/24  ISIS-L2 15   20           D   10.0.15.1       GigabitEthernet0/0/0
     10.0.13.0/24  ISIS-L2 15   20           D   10.0.15.1       GigabitEthernet0/0/0
    10.0.234.0/24  ISIS-L2 15   30           D   10.0.15.1       GigabitEthernet0/0/0

ISIS routing table status : <Inactive>
         Destinations : 0        Routes : 0
```

Create new loopback interfaces on R5 and import direct routes into the Level-2 routing table in IS-IS process 1.

```
[R5]interface LoopBack 1
[R5-LoopBack1]ip address 192.168.1.1 24
[R5-LoopBack1]quit
[R5]interface LoopBack 2
[R5-LoopBack2]ip address 192.168.2.1 24
[R5-LoopBack2]quit
[R5]interface LoopBack 3
[R5-LoopBack3]ip address 192.168.3.1 24
[R5-LoopBack3]quit
[R5]isis
[R5-isis-1]import-route direct level-2
[R5-isis-1]quit
```

Check IS-IS routes of R5.

```
[R5]display isis route

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

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

IPV4 Destination     IntCost    ExtCost ExitInterface    NextHop         Flags
-------------------------------------------------------------------------------
10.0.4.4/32             30          NULL     GE0/0/0          10.0.15.1       A/-/-/-
10.0.5.5/32             0           NULL     Loop0             Direct          D/-/L/-
10.0.12.0/24            20          NULL     GE0/0/0          10.0.15.1       A/-/-/-
10.0.13.0/24            20          NULL     GE0/0/0          10.0.15.1       A/-/-/-
10.0.234.0/24           30          NULL     GE0/0/0          10.0.15.1       A/-/-/-
10.0.15.0/24            10          NULL     GE0/0/0          Direct          D/-/L/-
10.0.1.1/32             10          NULL     GE0/0/0          10.0.15.1       A/-/-/-
10.0.2.2/32             20          NULL     GE0/0/0          10.0.15.1       A/-/-/-
10.0.3.3/32             20          NULL     GE0/0/0          10.0.15.1       A/-/-/-
     Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
                               U-Up/Down Bit Set

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

 Type IPV4 Destination     IntCost    ExtCost Tag
-------------------------------------------------------------------------------
 D    192.168.1.0/24       0          0
```

D	192.168.2.0/24	0	0
D	192.168.3.0/24	0	0

Type: D-Direct, I-ISIS, S-Static, O-OSPF, B-BGP, R-RIP, U-UNR

Check the IS-IS routing table of R4 again. No changes are found. This is because Level-2 routes are not leaked into Level-1 routers by default. R4 can access 192.168.1.0/24, 192.168.2.0/24, and 192.168.3.0/24 through default routes.

```
[R4]display ip routing-table protocol isis
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : ISIS
         Destinations : 5          Routes : 6

ISIS routing table status : <Active>
         Destinations : 5          Routes : 6

Destination/Mask    Proto    Pre  Cost      Flags NextHop         Interface

        0.0.0.0/0   ISIS-L1 15   10          D   10.0.234.3   GigabitEthernet0/0/0
                    ISIS-L1 15   10          D   10.0.234.2   GigabitEthernet0/0/0
      10.0.2.2/32   ISIS-L1 15   10          D   10.0.234.2   GigabitEthernet0/0/0
      10.0.3.3/32   ISIS-L1 15   10          D   10.0.234.3   GigabitEthernet0/0/0
     10.0.12.0/24   ISIS-L1 15   20          D   10.0.234.2   GigabitEthernet0/0/0
     10.0.13.0/24   ISIS-L1 15   20          D   10.0.234.3   GigabitEthernet0/0/0

ISIS routing table status : <Inactive>
         Destinations : 0          Routes : 0
```

```
[R4]ping -c 1 192.168.1.1
  PING 192.168.1.1: 56  data bytes, press CTRL_C to break
    Reply from 192.168.1.1: bytes=56 Sequence=1 ttl=253 time=14 ms

  --- 192.168.1.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 14/14/14 ms

[R4]ping -c 1 192.168.2.1
  PING 192.168.2.1: 56  data bytes, press CTRL_C to break
    Reply from 192.168.2.1: bytes=56 Sequence=1 ttl=253 time=13 ms

  --- 192.168.2.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 13/13/13 ms

[R4]ping -c 1 192.168.3.1
  PING 192.168.3.1: 56  data bytes, press CTRL_C to break
    Reply from 192.168.3.1: bytes=56 Sequence=1 ttl=253 time=1 ms

  --- 192.168.3.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
```

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

Step 6 **Change the IS-IS interface cost.**

By default, the IS-IS interface cost is 10, which is not automatically calculated based on the bandwidth. For R1, traffic destined for R4 is load balanced between R2 and R3. R1 and R2 are connected using Serial interfaces, which have low bandwidth and are prone to bandwidth bottlenecks. Therefore, you can change the IS-IS interface cost to control route selection of R1.

Increase the outbound interface cost.

```
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]isis cost 15
[R1-Serial1/0/0]quit

[R1]display isis route

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

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

IPV4 Destination     IntCost    ExtCost ExitInterface     NextHop        Flags
-------------------------------------------------------------------------
10.0.4.4/32            20         NULL       GE0/0/1          10.0.13.3  A/-/-/-
10.0.5.5/32            10         NULL       GE0/0/0          10.0.15.5      A/-/-/-
192.168.1.0/24         10         0          GE0/0/0          10.0.15.5  A/-/-/-
10.0.12.0/24           15         NULL       S1/0/0           Direct         D/-/L/-
192.168.2.0/24         10         0          GE0/0/0          10.0.15.5  A/-/-/-
```

```
10.0.13.0/24     10     NULL    GE0/0/1     Direct      D/-/L/-
192.168.3.0/24   10     0       GE0/0/0     10.0.15.5  A/-/-/-
10.0.234.0/24    20     NULL    GE0/0/1     10.0.13.3  A/-/-/-
10.0.15.0/24     10     NULL    GE0/0/0     Direct      D/-/L/-
10.0.1.1/32      0      NULL    Loop0       Direct      D/-/L/-
10.0.2.2/32      15     NULL    S1/0/0      10.0.12.2  A/-/-/-
10.0.3.3/32      10     NULL    GE0/0/1     10.0.13.3  A/-/-/-
     Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
                              U-Up/Down Bit Set
```

The preceding command output shows that traffic from R1 to R4 is forwarded through Ethernet interfaces.

Step 7 **Configure IS-IS route leaking.**

Observe data forwarding on R4. R4 does not know Level-2 network information and forwards data to R2 and R3 for load balancing. If you do not want R4 to use the link between R2 and R1, configure route leaking to import Level-2 routes into Level-1. According to the longest match principle, R3 forwards packets destined for R5.

Before performing a tracert operation, enable interface unreachable response on R5. By default, interface unreachable response is disabled. If it is disabled, the last hop will expire.

Before interface unreachable response is enabled:

```
[R4]tracert 192.168.1.1
 traceroute to  192.168.1.1(192.168.1.1), max hops: 30 ,packet length: 40,press CTRL_C to break
 1 10.0.234.2 2 ms 10.0.234.3 2 ms 10.0.234.2 2 ms
 2 10.0.13.1 11 ms 10.0.12.1 9 ms 10.0.13.1 11 ms
 3  *  *  *
```

After interface unreachable response is enabled on R5:

```
[R5]icmp port-unreachable send

[R4]tracert 192.168.1.1
 traceroute to  192.168.1.1(192.168.1.1), max hops: 30 ,packet length: 40,press CTRL_C to break
 1 10.0.234.2 2 ms 10.0.234.3 2 ms 10.0.234.2 2 ms
 2 10.0.13.1 2 ms 10.0.12.1 9 ms 10.0.13.1 1 ms
 3 192.168.1.1 8 ms  1 ms  9 ms
```

The preceding command output shows that tracert packets sent each time are load balanced between two next-hop addresses and then reach R5. You can enable route leaking on R3 to make R3 become the preferred next hop.

```
[R3]isis
[R3-isis-1]import-route isis level-2 into level-1
[R3-isis-1]quit

[R4]display isis route

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

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

IPV4 Destination     IntCost    ExtCost ExitInterface    NextHop          Flags
-------------------------------------------------------------------------------
0.0.0.0/0            10         NULL GE0/0/0             10.0.234.3       A/-/-/-
                                     GE0/0/0             10.0.234.2
```

```
10.0.4.4/32        0     NULL   Loop0      Direct       D/-/L/-
10.0.5.5/32        30    NULL   GE0/0/0    10.0.234.3   A/-/-/U
192.168.1.0/24     10    20     GE0/0/0    10.0.234.3   A/-/-/U
10.0.12.0/24       20    NULL   GE0/0/0    10.0.234.2   A/-/-/-
192.168.2.0/24     10    20     GE0/0/0    10.0.234.3   A/-/-/U
10.0.13.0/24       20    NULL   GE0/0/0    10.0.234.3   A/-/-/-
192.168.3.0/24     10    20     GE0/0/0    10.0.234.3   A/-/-/U
10.0.234.0/24      10    NULL   GE0/0/0    Direct       D/-/L/-
10.0.15.0/24       30    NULL   GE0/0/0    10.0.234.3   A/-/-/U
10.0.1.1/32        20    NULL   GE0/0/0    10.0.234.3   A/-/-/U
10.0.2.2/32        10    NULL   GE0/0/0    10.0.234.2   A/-/-/-
10.0.3.3/32        10    NULL   GE0/0/0    10.0.234.3   A/-/-/-
     Flags: D-Direct, A-Added to URT, L-Advertised in LSPs, S-IGP Shortcut,
                          U-Up/Down Bit Set

[R4]display ip routing-table protocol isis
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : ISIS
         Destinations : 11        Routes : 12

ISIS routing table status : <Active>
         Destinations : 11        Routes : 12

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        0.0.0.0/0  ISIS-L1 15   10         D    10.0.234.3     GigabitEthernet0/0/0
                   ISIS-L1 15   10         D    10.0.234.2     GigabitEthernet0/0/0
      10.0.1.1/32  ISIS-L1 15   20         D    10.0.234.3     GigabitEthernet0/0/0
```

```
   10.0.2.2/32  ISIS-L1 15     10          D     10.0.234.2      GigabitEthernet0/0/0
   10.0.3.3/32  ISIS-L1 15     10          D     10.0.234.3      GigabitEthernet0/0/0
   10.0.5.5/32  ISIS-L1 15     30        D     10.0.234.3    GigabitEthernet0/0/0
  10.0.12.0/24  ISIS-L1 15     20        D     10.0.234.2    GigabitEthernet0/0/0
  10.0.13.0/24  ISIS-L1 15     20        D     10.0.234.3    GigabitEthernet0/0/0
  10.0.15.0/24  ISIS-L1 15     30        D     10.0.234.3    GigabitEthernet0/0/0
192.168.1.0/24 ISIS-L1 15     94        D     10.0.234.3    GigabitEthernet0/0/0
192.168.2.0/24 ISIS-L1 15     94        D     10.0.234.3    GigabitEthernet0/0/0
192.168.3.0/24 ISIS-L1 15     94        D     10.0.234.3    GigabitEthernet0/0/0

ISIS routing table status : <Inactive>
         Destinations : 0           Routes : 0

[R4]tracert 192.168.1.1
 traceroute to   192.168.1.1(192.168.1.1), max hops: 30 ,packet length: 40,press CTRL_C to break
 1 10.0.234.3 2 ms   1 ms   1 ms
 2 10.0.13.1 2 ms   2 ms   2 ms
 3 192.168.1.1 1 ms   1 ms   2 ms
```

The preceding tracert operation proves the impact of route leaking on IS-IS route selection.

----End

Device Configurations

```
<R1>display current-configuration
[V200R007C00SPC600]
#
 sysname R1
```

```
#
isis 1
 is-level level-2
 network-entity 49.0002.0000.0000.0001.00
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
 isis enable 1
 isis cost 15
#
interface GigabitEthernet0/0/0
 ip address 10.0.15.1 255.255.255.0
 isis enable 1
 isis circuit-type p2p
#
interface GigabitEthernet0/0/1
 ip address 10.0.13.1 255.255.255.0
 isis enable 1
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
 isis enable 1
#
return

<R2>display current-configuration
[V200R007C00SPC600]
#
```

```
 sysname R2
#
isis 1
 network-entity 49.0001.0000.0000.0002.00
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
 isis enable 1
#
interface GigabitEthernet0/0/0
 ip address 10.0.234.2 255.255.255.0
 isis enable 1
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.255
 isis enable 1
#
return

<R3>display current-configuration
[V200R007C00SPC600]
#
 sysname R3
#
isis 1
 network-entity 49.0001.0000.0000.0003.00
 import-route isis level-2 into level-1
#
```

```
interface GigabitEthernet0/0/0
 ip address 10.0.234.3 255.255.255.0
 isis enable 1
#
interface GigabitEthernet0/0/1
 ip address 10.0.13.3 255.255.255.0
 isis enable 1
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
 isis enable 1
#
return

<R4>display current-configuration
[V200R007C00SPC600]
#
 sysname R4
#
isis 1
 is-level level-1
 network-entity 49.0001.0000.0000.0004.00
#
interface GigabitEthernet0/0/0
 ip address 10.0.234.4 255.255.255.0
 isis enable 1
 isis dis-priority 120
#
interface LoopBack0
```

```
 ip address 10.0.4.4 255.255.255.255
 isis enable 1
#
return

<R5>display current-configuration
[V200R007C00SPC600]
#
 sysname R5
#
 icmp port-unreachable send
#
isis 1
 is-level level-2
 network-entity 49.0002.0000.0000.0005.00
 import-route direct
#
interface GigabitEthernet0/0/0
 ip address 10.0.15.5 255.255.255.0
 isis enable 1
 isis circuit-type p2p
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.255
 isis enable 1
#
interface LoopBack1
 ip address 192.168.1.1 255.255.255.0
#
```

```
interface LoopBack2
 ip address 192.168.2.1 255.255.255.0
#
interface LoopBack3
 ip address 192.168.3.1 255.255.255.0
#
return
```

Chapter 3 BGP Features and Configurations

Lab 3-1 IBGP and EBGP

Learning Objectives

The objectives of this lab are to learn and understand:

- How to configure IBGP
- How to configure BGP in multiple areas
- How to check the BGP neighbor table and database
- How to configure a BGP source address for initiating a connection
- How to configure EBGP multihop
- How to observe next-hop changes in IBGP and EBGP routes
- How to configure a next hop in IBGP
- How to configure the **network** command in BGP

Topology

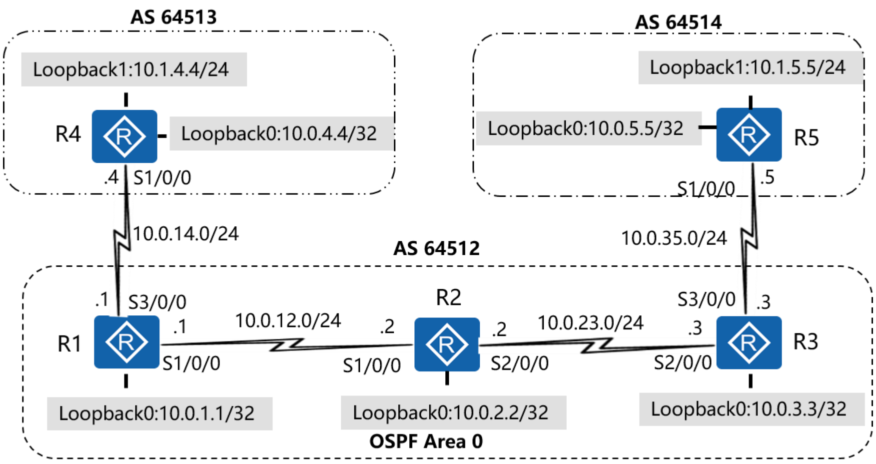

Figure 3-1 IBGP and EBGP

Scenario

You are a network administrator of a company. The company’s network uses BGP as the routing protocol. This network consists of multiple autonomous systems (ASs). Different branches use different AS numbers. You need to build this network. OSPF is used in the headquarters as an IGP. Different branches in the company use private BGP AS numbers. After building the network, you still need to observe BGP routing information transmission.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

Configure IP addresses and masks for all the routers. The mask of IP addresses for Loopback1 of R4 and R5 is 24 bits, which is used to simulate a user network.

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]quit
[R1]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 24
[R1-Serial3/0/0]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 32
[R1-LoopBack0]quit

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]quit
```

```
[R2]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 24
[R2-Serial2/0/0]quit
[R2]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 32

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 24
[R3-Serial2/0/0]quit
[R3]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 24
[R3-Serial3/0/0]quit
[R3]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 32

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 24
[R4-Serial1/0/0]quit
[R3]interface LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 32

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface Serial 1/0/0
[R5-Serial1/0/0]ip address 10.0.35.5 24
```

```
[R5-Serial1/0/0]quit
[R3]interface LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 32
```

Test direct link connectivity.

```
<R1>ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=34 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 34/34/34 ms

<R1>ping -c 1 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=40 ms

  --- 10.0.14.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 40/40/40 ms

<R3>ping -c 1  10.0.23.2
  PING 10.0.23.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.2: bytes=56 Sequence=1 ttl=255 time=33 ms
```

```
   --- 10.0.23.2 ping statistics ---
     1 packet(s) transmitted
     1 packet(s) received
     0.00% packet loss
     round-trip min/avg/max = 33/33/33 ms

<R3>ping -c 1 10.0.35.5
   PING 10.0.35.5: 56   data bytes, press CTRL_C to break
     Reply from 10.0.35.5: bytes=56 Sequence=1 ttl=255 time=35 ms

   --- 10.0.35.5 ping statistics ---
     1 packet(s) transmitted
     1 packet(s) received
     0.00% packet loss
round-trip min/avg/max = 35/35/35 ms
```

The preceding command output shows that direct link connectivity is normal.

Step 2 **Configure an IGP in an AS.**

Use OSPF as an IGP in AS 64512 and advertise the network segment connected to Loopback0 into OSPF. Run OSPF on the network segment connected to S1/0/0 of R1.

```
[R1]router id 10.0.1.1
[R1]ospf 1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]quit
[R1-ospf-1]quit
```

Run OSPF on the network segments connected to S1/0/0 and S2/0/0 of R2.

```
[R2]router id 10.0.2.2

[R2]ospf 1

[R2-ospf-1]area 0

[R2-ospf-1-area-0.0.0.0]network 10.0.12.2 0.0.0.0

[R2-ospf-1-area-0.0.0.0]network 10.0.23.2   0.0.0.0

[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0

[R2-ospf-1-area-0.0.0.0]quit

[R2-ospf-1]quit
```

Run OSPF on the network segment connected to S2/0/0 of R3.

```
[R3]router id 10.0.3.3

[R3]ospf 1

[R3-ospf-1]area 0

[R3-ospf-1-area-0.0.0.0]network 10.0.23.3 0.0.0.0

[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0

[R3-ospf-1-area-0.0.0.0]quit

[R3-ospf-1]quit
```

When configuring the **network** command, use the wildcard mask 0.0.0.0.

After the configurations are complete, check whether OSPF neighbor relationships are established.

```
[R2]display ospf peer


         OSPF Process 1 with Router ID 10.0.2.2

                 Neighbors
```

```
 Area 0.0.0.0 interface 10.0.12.2(Serial1/0/0)'s neighbors
 Router ID: 10.0.1.1          Address: 10.0.12.1
   State: Full   Mode:Nbr is   Slave   Priority: 1
   DR: None    BDR: None    MTU: 0
   Dead timer due in 37   sec
   Retrans timer interval: 5
   Neighbor is up for 00:01:05
   Authentication Sequence: [ 0 ]

                 Neighbors

 Area 0.0.0.0 interface 10.0.23.2(Serial2/0/0)'s neighbors
 Router ID: 10.0.3.3          Address: 10.0.23.3
   State: Full   Mode:Nbr is   Master   Priority: 1
   DR: None    BDR: None    MTU: 0
   Dead timer due in 33   sec
   Retrans timer interval: 5
   Neighbor is up for 00:00:19
   Authentication Sequence: [ 0 ]
```

Check the IP routing table of each router. Check whether these routers can learn routes to the network segments connected to loopback interfaces of the peer device.

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 15         Routes : 15
```

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

        10.0.1.1/32   Direct    0      0              D     127.0.0.1         LoopBack0
        10.0.2.2/32  OSPF     10     1562             D     10.0.12.2         Serial1/0/0
        10.0.3.3/32  OSPF     10     3124             D     10.0.12.2         Serial1/0/0
       10.0.12.0/24  Direct   0      0                D     10.0.12.1         Serial1/0/0
       10.0.12.1/32  Direct   0      0                D     127.0.0.1         Serial1/0/0
       10.0.12.2/32  Direct   0      0                D     10.0.12.2         Serial1/0/0
     10.0.12.255/32  Direct   0      0                D     127.0.0.1         Serial1/0/0
       10.0.14.0/24  Direct   0      0                D     10.0.14.1         Serial3/0/0
       10.0.14.1/32  Direct   0      0                D     127.0.0.1         Serial3/0/0
       10.0.14.4/32  Direct   0      0                D     10.0.14.4         Serial3/0/0
     10.0.14.255/32  Direct   0      0                D     127.0.0.1         Serial3/0/0
       10.0.23.0/24  OSPF    10     3124              D     10.0.12.2         Serial1/0/0
        127.0.0.0/8  Direct   0      0                D     127.0.0.1         InLoopBack0
       127.0.0.1/32  Direct   0      0                D     127.0.0.1          InLoopBack0
127.255.255.255/32Direct   0      0                D     127.0.0.1          InLoopBack0
255.255.255.255/32Direct   0      0                D     127.0.0.1          InLoopBack0

[R2]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 15          Routes : 15

Destination/Mask          Proto    Pre   Cost        Flags NextHop          Interface

10.0.1.1/32     OSPF      10     1562           D     10.0.12.1         Serial1/0/0
```

```
10.0.2.2/32       Direct   0      0             D     127.0.0.1         LoopBack0
10.0.3.3/32     OSPF      10     1562           D     10.0.23.3         Serial2/0/0
10.0.12.0/24    Direct   0      0             D     10.0.12.2         Serial1/0/0
10.0.12.1/32    Direct   0      0             D     10.0.12.1         Serial1/0/0
10.0.12.2/32    Direct   0      0             D     127.0.0.1         Serial1/0/0
10.0.12.255/32Direct   0      0             D     127.0.0.1         Serial1/0/0
10.0.23.0/24    Direct   0      0             D     10.0.23.2         Serial2/0/0
10.0.23.2/32    Direct   0      0             D     127.0.0.1         Serial2/0/0
10.0.23.3/32    Direct   0      0             D     10.0.23.3         Serial2/0/0
10.0.23.255/32Direct   0      0             D     127.0.0.1         Serial2/0/0
127.0.0.0/8     Direct   0      0             D     127.0.0.1         InLoopBack0
127.0.0.1/32    Direct   0      0             D     127.0.0.1         InLoopBack0
127.255.255.255/32       Direct   0      0             D     127.0.0.1         InLoopBack0
255.255.255.255/32       Direct   0      0             D     127.0.0.1         InLoopBack0
```

```
[R3]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 16        Routes : 16

Destination/Mask      Proto    Pre   Cost        Flags NextHop           Interface

       10.0.1.1/32  OSPF      10     3124          D     10.0.23.2          Serial2/0/0
       10.0.2.2/32  OSPF      10     1562          D     10.0.23.2          Serial2/0/0
       10.0.3.3/32  Direct     0      0             D     127.0.0.1         LoopBack0
      10.0.12.0/24  OSPF      10     3124          D     10.0.23.2          Serial2/0/0
      10.0.23.0/24  Direct      0      0             D     10.0.23.3          Serial2/0/0
      10.0.23.2/32  Direct      0      0             D     10.0.23.2          Serial2/0/0
```

```
      10.0.23.3/32  Direct   0    0       D   127.0.0.1       Serial2/0/0
    10.0.23.255/32  Direct   0    0       D   127.0.0.1       Serial2/0/0
      10.0.35.0/24  Direct   0    0       D   10.0.35.3       Serial3/0/0
      10.0.35.3/32  Direct   0    0       D   127.0.0.1       Serial3/0/0
      10.0.35.5/32  Direct   0    0       D   10.0.35.5       Serial3/0/0
    10.0.35.255/32  Direct   0    0       D   127.0.0.1       Serial3/0/0
      127.0.0.0/8   Direct   0    0       D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct   0    0       D   127.0.0.1       InLoopBack0
127.255.255.255/32Direct   0    0       D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct   0    0       D   127.0.0.1       InLoopBack0
```

The IP routing tables of R1, R2, and R3 show that each router can learn routes to the network segment connected to Loopback0 of the other two routers.

Step 3 **Configure IBGP peers.**

Configure IBGP full mesh on R1, R2, and R3. Use Loopback0 address as a source address.

```
[R1]bgp 64512
[R1-bgp]peer 10.0.2.2 as-number 64512
[R1-bgp]peer 10.0.2.2 connect-interface LoopBack 0
[R1-bgp]peer 10.0.3.3 as-number 64512
[R1-bgp]peer 10.0.3.3 connect-interface LoopBack 0
[R1-bgp]quit

[R2]bgp 64512
[R2-bgp]peer 10.0.1.1 as-number 64512
[R2-bgp]peer 10.0.1.1 connect-interface loopback 0
[R2-bgp]peer 10.0.3.3 as-number 64512
```

[R2-bgp]peer 10.0.3.3 connect-interface LoopBack 0

[R2-bgp]quit

[R3]bgp 64512

[R3-bgp]peer 10.0.1.1 as-number 64512

[R3-bgp]peer 10.0.1.1 connect-interface loopback 0

[R3-bgp]peer 10.0.2.2 as-number 64512

[R3-bgp]peer 10.0.2.2 connect-interface LoopBack 0

[R3-bgp]quit

Run the **display tcp status** command to check the TCP port connection status.

[R2]display tcp status

TCPCB	Tid/Soid	Local Add:port	Foreign Add:port	VPNID	State
37a32f14	76 /1	0.0.0.0:80	0.0.0.0:0	23553	Listening
37a33b34	239/2	0.0.0.0:179	10.0.1.1:0	0	Listening
39052914	239/6	0.0.0.0:179	10.0.3.3:0	0	Listening
37a3321c	76 /3	0.0.0.0:443	0.0.0.0:0	23553	Listening
39052c1c	239/11	10.0.2.2:179	10.0.3.3:54086	0	Established
3905260c	239/5	10.0.2.2:61635	10.0.1.1:179	0	Established

The preceding command output shows that the **Local Add** field displays 10.0.2.2 (Loopback0 address of R2) and port number is 179 (TCP port number of BGP). The neighbor state with 10.0.3.3 and 10.0.1.1 is Established, indicating that R2 has established a TCP connection with R1 and R3.

Run the **display bgp peer** command to check BGP peer relationships of routers.

[R1]display bgp peer

BGP local router ID : 10.0.1.1

Local AS number : 64512

```
Total number of peers : 2                Peers in established state : 2

 Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State  PrefRcv

 10.0.2.2        4       64512      273      277     0 02:15:53 Established        0
 10.0.3.3        4       64512      276      276     0 02:15:53 Established        0

[R2]display bgp peer

 BGP local router ID : 10.0.2.2
 Local AS number : 64512
 Total number of peers : 2               Peers in established state : 2

 Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State  PrefRcv

 10.0.1.1        4       64512       38       38     0 00:18:02 Established        0
 10.0.3.3        4       64512     1000     1000     0 16:38:38 Established        0

[R3]display bgp peer

 BGP local router ID : 10.0.3.3
 Local AS number : 64512
 Total number of peers : 2               Peers in established state : 2

 Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State  PrefRcv

 10.0.1.1        4       64512       39       39     0 00:18:35 Established        0
 10.0.2.2        4       64512     1001     1001     0 16:39:11 Established        0
```

The preceding command output shows that BGP peer relationships between R1, R2, and R3 are in Established state, indicating that BGP peer relationships have been established.

On R1, run the timer command in the BGP process to change the Keepalive time to 30s and hold time to 90s. Check whether the BGP peer relationship between R1 and R2 is established and run the **display bgp peer verbose** command to check the negotiation interval after the BGP peer relationship is established.

```
[R1-bgp] bgp 64512
[R1-bgp] timer keepalive 30 hold 90
Warning: Changing the parameter in this command resets the peer session. Continue?[Y/N]:y
[R1-bgp]quit
```

Note that changing the Keepalive time and hold time will restart the BGP session.

```
[R2]display bgp peer verbose

    BGP Peer is 10.0.1.1,   remote AS 64512
    Type: IBGP link
    BGP version 4, Remote router ID 10.0.1.1
    Update-group ID: 1
    BGP current state: Established, Up for 00h07m19s
    BGP current event: KATimerExpired
    BGP last state: OpenConfirm
    BGP Peer Up count: 2
    Received total routes: 0
    Received active routes total: 0
    Advertised total routes: 0
    Port:   Local - 50117      Remote - 179
    Configured: Connect-retry Time: 32 sec
```

Configured: Active Hold Time: 180 sec Keepalive Time:60 sec

Received : Active Hold Time: 90 sec

Negotiated: Active Hold Time: 90 sec Keepalive Time:30 sec

Peer optional capabilities:

Peer supports bgp multi-protocol extension

Peer supports bgp route refresh capability

Peer supports bgp 4-byte-as capability

Address family IPv4 Unicast: advertised and received

Received: Total 16 messages

Update messages 0

Open messages 1

KeepAlive messages 15

Notification messages 0

Refresh messages 0

Sent: Total 16 messages

Update messages 0

Open messages 1

KeepAlive messages 15

Notification messages 0

Refresh messages 0

Authentication type configured: None

Last keepalive received: 2011/12/07 08:33:52

Minimum route advertisement interval is 15 seconds

Optional capabilities:

Route refresh capability has been enabled

4-byte-as capability has been enabled

Connect-interface has been configured

Peer Preferred Value: 0

Routing policy configured:

No routing policy is configured

BGP Peer is 10.0.3.3, remote AS 64512

Type: IBGP link

BGP version 4, Remote router ID 10.0.3.3

Update-group ID: 1

BGP current state: Established, Up for 16h28m14s

BGP current event: RecvKeepalive

BGP last state: OpenConfirm

BGP Peer Up count: 1

Received total routes: 0

Received active routes total: 0

Advertised total routes: 0

Port: Local - 179 Remote - 49663

Configured: Connect-retry Time: 32 sec

Configured: Active Hold Time: 180 sec Keepalive Time:60 sec

Received : Active Hold Time: 180 sec

Negotiated: Active Hold Time: 180 sec Keepalive Time:60 sec

Peer optional capabilities:

Peer supports bgp multi-protocol extension

Peer supports bgp route refresh capability

Peer supports bgp 4-byte-as capability

Address family IPv4 Unicast: advertised and received

Received: Total 990 messages

Update messages 0

Open messages 1

KeepAlive messages 989

Notification messages 0

Refresh messages 0

```
Sent: Total 990 messages
        Update messages          0
        Open messages            1
        KeepAlive messages       989
        Notification messages    0
        Refresh messages         0
Authentication type configured: None
Last keepalive received: 2011/12/07 08:34:17
Minimum route advertisement interval is 15 seconds
Optional capabilities:
Route refresh capability has been enabled
4-byte-as capability has been enabled
Connect-interface has been configured
Peer Preferred Value: 0
Routing policy configured:
No routing policy is configured
```

The preceding command output of R2 shows that the default parameter Active Hold Time is 180s and Keepalive Time is 60s.

After parameters of R1 are changed, the Active Hold Time of packets received by R2 becomes 90s. The negotiated parameters use the smaller value. Therefore, the Active Hold Time and Keepalive Timer that are negotiated between R2 and R1 are 90s and 30s respectively, but the parameters of R3 still use the default values.

Therefore, the negotiated parameters are the same as the configured parameters. That is, the Active Hold Time and Keepalive Timer on R2 and R3 are 180s and 60s respectively.

Step 4 **Configure EBGP peers.**

Configure BGP on R4, set the local AS number to 64513, and establish an EBGP peer relationship between R4 and R1. During EBGP peer relationship establishment, specify the address of Loopback0 as the source address and set **ebgp-max-hop** to 2. Add a 32-bit static route to the Loopback0 address of the peer device to ensure that an EBGP peer relationship can be established normally.

```
[R1]ip route-static 10.0.4.4 32 10.0.14.4


[R4]ip route-static 10.0.1.1 32 10.0.14.1


[R1]bgp 64512
[R1-bgp]peer 10.0.4.4 as-number 64513
[R1-bgp]peer 10.0.4.4 ebgp-max-hop 2
[R1-bgp]peer 10.0.4.4 connect-interface LoopBack0
[R1-bgp]quit


[R4]router id 10.0.4.4
[R4]bgp 64513
[R4-bgp]peer 10.0.1.1 as-number 64512
[R4-bgp]peer 10.0.1.1 ebgp-max-hop 2
[R4-bgp]peer 10.0.1.1 connect-interface LoopBack0
[R4-bgp]quit
```

After an EBGP peer relationship is established, run the **display bgp peer** command to check the peer relationship status.

```
[R4]display bgp peer


 BGP local router ID : 10.0.4.4
 Local AS number : 64513
 Total number of peers : 1                Peers in established state : 1
```

Peer	V	AS	MsgRcvd	MsgSent	OutQ	Up/Down	State	PrefRcv
10.0.1.1	4	64512	4	5	0	00:01:18	Established	0

Run the **debugging ip packet verbose** command on R4 to check the TTL value of Keepalive packets.

<R4>terminal monitor

<R4>terminal debugging

<R4>debugging ip packet

<R4>

Oct 31 2016 17:22:44.900.2+00:00 R4 IP/7/debug_case:

Receiving, interface = Serial1/0/0, version = 4, headlen = 20, tos = 192,

pktlen = 40, pktid = 429, offset = 0, ttl = 2, protocol = 6,

checksum = 40287, s = 10.0.1.1, d = 10.0.4.4

prompt: IP Process By Board Begin!

45 c0 00 28 01 ad 00 00 02 06 9d 5f 0a 00 01 01

0a 00 04 04

Oct 31 2016 17:22:44.900.3+00:00 R4 IP/7/debug_case:

Receiving, interface = Serial1/0/0, version = 4, headlen = 20, tos = 192,

pktlen = 40, pktid = 429, offset = 0, ttl = 2, protocol = 6,

checksum = 40287, s = 10.0.1.1, d = 10.0.4.4

prompt: Before search fib per flow in IP Forward.

The preceding command output shows that the TTL value of received packets is 2.

Establish an EBGP peer relationship between R3 and R5. Use physical interface addresses to establish a connection.

```
[R3]bgp 64512
[R3-bgp]peer 10.0.35.5 as-number 64514
[R3-bgp]quit

[R5]router id 10.0.5.5
[R5]bgp 64514
[R5-bgp]peer 10.0.35.3 as-number 64512
[R5-bgp]quit

[R5]display bgp peer

BGP local router ID : 10.0.5.5
 Local AS number : 64514
 Total number of peers : 1                 Peers in established state : 1

  Peer            V          AS     MsgRcvd  MsgSent  OutQ  Up/Down   State PrefRcv

  10.0.35.3       4       64512        2         3       0 00:00:46 Established       0
```

Step 5 **Run the network command to advertise routing information.**

Configure Loopback1 on R4 and configure an IP address 10.1.4.4/24 for Loopback1. Run the **network** command to advertise the network segment of Loopback1 into BGP.

```
[R4]interface LoopBack 1
[R4-LoopBack1]ip address 10.1.4.4 24
[R4-LoopBack1]quit
```

```
[R4]bgp 64513
[R4-bgp]network 10.1.4.4 24
[R4-bgp]quit
```

Check the IP routing tables of R1 and R3 to check whether the route to 10.1.4.4/24 exists.

Check the BGP routing table of R3 to analyze next-hop information of this route.

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 18        Routes : 18

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.1/32  Direct  0    0           D   127.0.0.1       LoopBack0
       10.0.2.2/32  OSPF    10   1562        D   10.0.12.2       Serial1/0/0
       10.0.3.3/32  OSPF    10   3124        D   10.0.12.2       Serial1/0/0
       10.0.4.4/32  Static  60   0          RD   10.0.14.4       Serial3/0/0
      10.0.12.0/24  Direct  0    0           D   10.0.12.1       Serial1/0/0
      10.0.12.1/32  Direct  0    0           D   127.0.0.1       Serial1/0/0
      10.0.12.2/32  Direct  0    0           D   10.0.12.2       Serial1/0/0
    10.0.12.255/32  Direct  0    0           D   127.0.0.1       Serial1/0/0
      10.0.14.0/24  Direct  0    0           D   10.0.14.1       Serial3/0/0
      10.0.14.1/32  Direct  0    0           D   127.0.0.1       Serial3/0/0
      10.0.14.4/32  Direct  0    0           D   10.0.14.4       Serial3/0/0
    10.0.14.255/32  Direct  0    0           D   127.0.0.1       Serial3/0/0
      10.0.23.0/24  OSPF    10   3124        D   10.0.12.2       Serial1/0/0
```

```
        10.1.4.0/24  EBGP   255  0        RD   10.0.4.4        Serial3/0/0
       127.0.0.0/8   Direct 0    0        D    127.0.0.1       InLoopBack0
       127.0.0.1/32  Direct 0    0        D    127.0.0.1       InLoopBack0
127.255.255.255/32Direct    0    0        D    127.0.0.1       InLoopBack0
255.255.255.255/3  Direct   0    0        D    127.0.0.1       InLoopBack0
```

The command output shows that R1 has learned the EBGP route to 10.1.4.0/24.

Check whether R3 has the route to 10.1.4.0/24.

```
[R3]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 16        Routes : 16

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        10.0.1.1/32  OSPF   10   3124      D    10.0.23.2       Serial2/0/0
        10.0.2.2/32  OSPF   10   1562      D    10.0.23.2       Serial2/0/0
        10.0.3.3/32  Direct 0    0         D    127.0.0.1       LoopBack0
       10.0.12.0/24  OSPF   10   3124      D    10.0.23.2       Serial2/0/0
       10.0.23.0/24  Direct 0    0         D    10.0.23.3       Serial2/0/0
       10.0.23.2/32  Direct 0    0         D    10.0.23.2       Serial2/0/0
       10.0.23.3/32  Direct 0    0         D    127.0.0.1       Serial2/0/0
     10.0.23.255/32 Direct  0    0         D    127.0.0.1       Serial2/0/0
       10.0.35.0/24  Direct 0    0         D    10.0.35.3       Serial3/0/0
       10.0.35.3/32  Direct 0    0         D    127.0.0.1       Serial3/0/0
       10.0.35.5/32  Direct 0    0         D    10.0.35.5       Serial3/0/0
     10.0.35.255/32 Direct  0    0         D    127.0.0.1       Serial3/0/0
```

```
        127.0.0.0/8    Direct   0     0          D    127.0.0.1        InLoopBack0
        127.0.0.1/32  Direct   0     0          D    127.0.0.1        InLoopBack0
127.255.255.255/32Direct   0     0          D    127.0.0.1        InLoopBack0
255.255.255.255/32Direct   0     0          D    127.0.0.1        InLoopBack0
```

The preceding command output shows that R3 does not have any BGP route to 10.1.4.4.

Check the BGP routing table of R3.

```
[R3]display bgp routing-table

 BGP Local router ID is 10.0.3.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 1
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

   i  10.1.4.0/24        10.0.4.4       0          100        0       64513i
```

The preceding command output shows that there is a BGP route to 10.1.4.0/24, but this route is not marked with *, indicating that this route is not preferred. The **NextHop** field of this route displays 10.0.4.4, but R3 does not have the route to 10.0.4.4. According to BGP route selection rules, when the next hop of a BGP route is unreachable, this route is ignored.

Run the **next-hop-local** command on R1 and check the BGP routing table of R3 again.

```
[R1]bgp 64512
```

```
[R1-bgp]peer 10.0.3.3  next-hop-local
[R1-bgp]peer 10.0.2.2  next-hop-local
[R1-bgp]quit

[R3]display bgp routing-table

 BGP Local router ID is 10.0.3.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 1
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.1.4.0/24        10.0.1.1       0          100       0       64513i
```

The preceding command output shows that the next hop of the BGP route 10.1.4.0/24 is 10.0.1.1 and this route is marked with * and >, indicating that this route is correct and the optimal route.

Check the IP routing table of R3.

```
[R3]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 17        Routes : 17

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

10.0.1.1/32	OSPF	10	3124	D	10.0.23.2	Serial2/0/0
10.0.2.2/32	OSPF	10	1562	D	10.0.23.2	Serial2/0/0
10.0.3.3/32	Direct	0	0	D	127.0.0.1	LoopBack0
10.0.12.0/24	OSPF	10	3124	D	10.0.23.2	Serial2/0/0
10.0.23.0/24	Direct	0	0	D	10.0.23.3	Serial2/0/0
10.0.23.2/32	Direct	0	0	D	10.0.23.2	Serial2/0/0
10.0.23.3/32	Direct	0	0	D	127.0.0.1	Serial2/0/0
10.0.23.255/32	Direct	0	0	D	127.0.0.1	Serial2/0/0
10.0.35.0/24	Direct	0	0	D	10.0.35.3	Serial3/0/0
10.0.35.3/32	Direct	0	0	D	127.0.0.1	Serial3/0/0
10.0.35.5/32	Direct	0	0	D	10.0.35.5	Serial3/0/0
10.0.35.255/32	Direct	0	0	D	127.0.0.1	Serial3/0/0
10.1.4.0/24	IBGP	255	0	RD	10.0.1.1	Serial2/0/0
127.0.0.0/8	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
127.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
255.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0

The IP routing table of R3 contains the route 10.1.4.0/24.

Create Loopback1 on R5, configure an IP address 10.1.5.5/24 for Loopback1, advertise it into BGP, and configure **next-hop-local**.

[R5]interface LoopBack 1

[R5-LoopBack1]ip address 10.1.5.5 24

[R5-LoopBack1]quit

[R5]bgp 64514

[R5-bgp]network 10.1.5.0 24

```
[R3]bgp 64512
[R3-bgp]peer 10.0.1.1 next-hop-local
[R3-bgp]peer 10.0.2.2 next-hop-local
```

Check the BGP routing table of R4 to determine whether R4 learns a route to the network segment connected to Loopback1 of R5. Analyze the **display bgp routing-table** command output.

```
[R4]display bgp routing-table

 BGP Local router ID is 10.0.4.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.1.4.0/24        0.0.0.0        0                     0       i
 *>   10.1.5.0/24        10.0.1.1                             0       64512 64514i
```

Perform a ping with the source address on R5 to test connectivity to Loopback1 address of R4.

```
[R5]ping -c 1 -a  10.1.5.5 10.1.4.4
  PING 10.1.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.1.4.4: bytes=56 Sequence=1 ttl=252 time=125 ms

  --- 10.1.4.4 ping statistics ---
    1 packet(s) transmitted
```

```
1 packet(s) received
0.00% packet loss
round-trip min/avg/max = 125/125/125 ms
```

Additional Exercises: Analysis and Verification

In which situation is it suitable to establish an EBGP peer relationship using physical addresses?

Why does the TTL value of packets sent to EBGP peers default to 1? What is the default configuration of the **peer** *group_name* **ebgp-max-hop** *[hop-count]* command?

Device Configurations

```
[R1]display current-configuration
[V200R007C00SPC600]
#
 sysname R1
#
 router id 10.0.1.1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
```

```
#
bgp 64512
 timer keepalive 30 hold 90
 peer 10.0.2.2 as-number 64512
 peer 10.0.2.2 connect-interface LoopBack0
 peer 10.0.3.3 as-number 64512
 peer 10.0.3.3 connect-interface LoopBack0
 peer 10.0.4.4 as-number 64513
 peer 10.0.4.4 ebgp-max-hop 2
 peer 10.0.4.4 connect-interface LoopBack0
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.2.2 enable
  peer 10.0.2.2 next-hop-local
  peer 10.0.3.3 enable
  peer 10.0.3.3 next-hop-local
  peer 10.0.4.4 enable
#
ospf 1
 area 0.0.0.0
  network 10.0.12.0 0.0.0.255
  network 10.0.1.1 0.0.0.0
#
 ip route-static 10.0.4.4 255.255.255.255 10.0.14.4
return

[R2]display current-configuration
[V200R007C00SPC600]
```

```
#
 sysname R2
#
 router id 10.0.2.2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.255
#
bgp 64512
 peer 10.0.1.1 as-number 64512
 peer 10.0.1.1 connect-interface LoopBack0
 peer 10.0.3.3 as-number 64512
 peer 10.0.3.3 connect-interface LoopBack0
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.1.1 enable
  peer 10.0.3.3 enable
#
ospf 1
 area 0.0.0.0
```

```
  network 10.0.12.0 0.0.0.255
  network 10.0.23.0 0.0.0.255
  network 10.0.2.2 0.0.0.0
return

[R3]display current-configuration
[V200R007C00SPC600]
#
 sysname R3
#
 router id 10.0.3.3
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
#
bgp 64512
 peer 10.0.1.1 as-number 64512
 peer 10.0.1.1 connect-interface LoopBack0
 peer 10.0.2.2 as-number 64512
 peer 10.0.2.2 connect-interface LoopBack0
 peer 10.0.35.5 as-number 64514
```

```
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.1.1 enable
  peer 10.0.1.1 next-hop-local
  peer 10.0.2.2 enable
  peer 10.0.2.2 next-hop-local
  peer 10.0.35.5 enable
#
ospf 1
 area 0.0.0.0
  network 10.0.23.0 0.0.0.255
  network 10.0.3.3 0.0.0.0
return

[R4]display current-configuration
[V200R007C00SPC600]
#
 sysname R4
#
 router id 10.0.4.4
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
#
```

```
interface LoopBack1
 ip address 10.1.4.4 255.255.255.0
#
bgp 64513
 peer 10.0.1.1 as-number 64512
 peer 10.0.1.1 ebgp-max-hop 2
 peer 10.0.1.1 connect-interface LoopBack0
 #
 ipv4-family unicast
  undo synchronization
  network 10.0.4.0 255.255.255.0
  network 10.1.4.0 255.255.255.0
  peer 10.0.1.1 enable
#
 ip route-static 10.0.1.1 255.255.255.255 10.0.14.1
return

[R5]display current-configuration
[V200R007C00SPC600]
#
 sysname R5
#
 router id 10.0.5.5
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
#
interface LoopBack0
```

```
 ip address 10.0.5.5 255.255.255.255
#
interface LoopBack1
 ip address 10.1.5.5 255.255.255.0
#
bgp 64514
 peer 10.0.35.3 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.5.0 255.255.255.0
  peer 10.0.35.3 enable
return
```

Lab 3-2 BGP Route Summarization

Learning Objectives

The objectives of this lab are to learn and understand:

- How to run the **network** command to implement BGP route summarization
- How to configure BGP route summarization to suppress specific routes
- How to change attributes of summarized routes
- How to use the AS_Set during route summarization

Topology

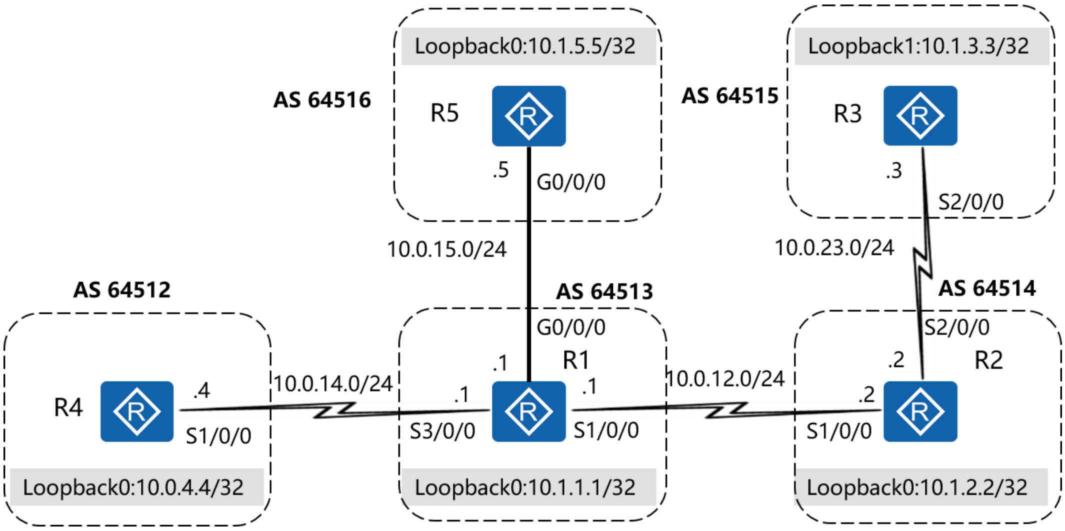

Figure 3-2 BGP route summarization

Scenario

You are a network administrator of a company. The company' s network uses BGP as the routing protocol. This network consists of multiple ASs, and different branches use different AS numbers. As the company expands, routers have more and more routing tables. It is urgent to summarize BGP routes to reduce the routing table size. You test several route summarization methods and select a suitable method to summarize routes.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

Configure IP addresses and masks for physical interfaces and loopback interfaces of all the routers. Each loopback interface address uses the 32-bit mask.

```
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 255.255.255.0
[R1-Serial1/0/0]quit
[R1]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 255.255.255.0
```

```
[R1-Serial3/0/0]quit
[R1]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]ip add 10.0.15.1 255.255.255.0
[R1-GigabitEthernet0/0/0]quit
[R1]interface   LoopBack 0
[R1-LoopBack0]ip address 10.1.1.1 255.255.255.255
[R1-LoopBack0]quit

[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 255.255.255.0
[R2-Serial1/0/0]quit
[R2]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 255.255.255.0
[R2-Serial2/0/0]quit
[R2]interface loopback 0
[R2-LoopBack0]ip address 10.1.2.2 255.255.255.255
[R2-LoopBack0]quit

[R3]interface Serial 2/0/0
[R3-Serial2/0/0]quit
[R3]ip address 10.0.23.3 255.255.255.0
[R3-Serial2/0/0]quit
[R3]interface loopback 1
[R3-LoopBack1]ip address 10.1.3.3 255.255.255.255
[R3-LoopBack1]quit

[R4]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 255.255.255.0
[R4-Serial1/0/0]quit
```

```
[R4]interface loopback 0
[R4-LoopBack0]ip address 10.0.4.4 255.255.255.255
[R4-LoopBack0]quit

[R5]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.15.5 255.255.255.0
[R5-GigabitEthernet0/0/0]quit
[R5]interface loopback 0
[R5-LoopBack0]ip address 10.1.5.5 255.255.255.255
[R5-LoopBack0]quit
```

After the configurations are complete, test direct link connectivity.

```
[R1]ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=41 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 41/41/41 ms

[R1]ping -c 1 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=41 ms

  --- 10.0.14.4 ping statistics ---
    1 packet(s) transmitted
```

```
     1 packet(s) received
     0.00% packet loss
     round-trip min/avg/max = 41/41/41 ms

[R1]ping -c 1 10.0.15.5
  PING 10.0.15.5: 56  data bytes, press CTRL_C to break
     Reply from 10.0.15.5: bytes=56 Sequence=1 ttl=255 time=34 ms

  --- 10.0.15.5 ping statistics ---
     1 packet(s) transmitted
     1 packet(s) received
     0.00% packet loss
     round-trip min/avg/max = 34/34/34 ms

[R2]ping -c 1 10.0.23.3
  PING 10.0.23.3: 56  data bytes, press CTRL_C to break
     Reply from 10.0.23.3: bytes=56 Sequence=1 ttl=255 time=34 ms

  --- 10.0.23.3 ping statistics ---
     1 packet(s) transmitted
     1 packet(s) received
     0.00% packet loss
     round-trip min/avg/max = 34/34/34 ms
```

Step 2 **Configure EBGP and advertise routes.**

Configure directly connected routers to establish BGP peer relationships using physical interface addresses.

```
[R1]router id 10.1.1.1
```

```
[R1]bgp 64513
[R1-bgp]peer 10.0.12.2 as-number 64514
[R1-bgp]peer 10.0.14.4 as-number 64512
[R1-bgp]peer 10.0.15.5 as-number 64516
[R1-bgp]quit

[R2]router id 10.1.2.2
[R2]bgp 64514
[R2-bgp]peer 10.0.12.1 as-number 64513
[R2-bgp]peer 10.0.23.3 as-number 64515
[R2-bgp]quit

[R3]router id 10.1.3.3
[R3]bgp 64515
[R3-bgp]peer 10.0.23.2 as-number 64514
[R3-bgp]quit

[R4]router id 10.0.4.4
[R4]bgp 64512
[R4-bgp]peer 10.0.14.1 as-number 64513
[R4-bgp]quit

[R5]router id 10.1.5.5
[R5]bgp 64516
[R5-bgp]peer 10.0.15.1 as-number 64513
[R5-bgp]quit
```

After the configurations are complete, check BGP peer relationships.

```
[R1]display bgp peer

BGP local router ID : 10.1.1.1
 Local AS number : 64513
 Total number of peers : 3                 Peers in established state : 3

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State  PrefRcv

  10.0.12.2       4       64514        4        6     0 00:02:19 Established        0
  10.0.14.4       4       64512        2        4     0 00:00:40 Established        0
  10.0.15.5       4       64516        2        4     0 00:00:17 Established        0

[R2]display bgp peer

BGP local router ID : 10.1.2.2
 Local AS number : 64514
 Total number of peers : 2                 Peers in established state : 2

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.0.12.1       4       64513        5        6     0 00:03:42 Established        0
  10.0.23.3       4       64515        4        6     0 00:02:25 Established        0

[R3]display bgp peer

BGP local router ID : 10.1.3.3
 Local AS number : 64515
 Total number of peers : 1                 Peers in established state : 1
```

```
 Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

 10.0.23.2       4       64514        6        7     0 00:04:55 Established       0

[R4]display bgp peer

BGP local router ID : 10.0.4.4
 Local AS number : 64512
 Total number of peers : 1                Peers in established state : 1

 Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

 10.0.14.1       4       64513        7        8     0 00:05:11 Established       0

[R5]display bgp peer

BGP local router ID : 10.1.5.5
 Local AS number : 64516
 Total number of peers : 1                Peers in established state : 1

 Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

 10.0.15.1       4       64513        7        8     0 00:05:16 Established       0
```

The preceding command output shows that all BGP peer relationships are in Established state.

Run the **network** command to advertise the network segment of the loopback interface on each router into BGP.

[R1]bgp 64513

[R1-bgp]network 10.1.1.1 255.255.255.255

[R1-bgp]quit

[R2]bgp 64514

[R2-bgp]network 10.1.2.2 255.255.255.255

[R2-bgp]quit

[R3]bgp 64515

[R3-bgp]network 10.1.3.3 255.255.255.255

[R3-bgp]quit

[R4]bgp 64512

[R4-bgp]network 10.0.4.4 255.255.255.255

[R4-bgp]quit

[R5]bgp 64516

[R5-bgp]network 10.1.5.5 255.255.255.255

[R5-bgp]quit

Check the BGP routing table of R4 and observe the AS_Path attribute.

[R4]display bgp routing-table

BGP Local router ID is 10.0.4.4

Status codes: * - valid, > - best, d - damped,

h - history, i - internal, s - suppressed, S - Stale

Origin : i - IGP, e - EGP, ? - incomplete

```
 Total Number of Routes: 5
       Network            NextHop        MED        LocPrf PrefVal Path/Ogn

 *>    10.0.4.4/32        0.0.0.0        0                 0       i
 *>    10.1.1.1/32        10.0.14.1      0                 0       64513i
 *>    10.1.2.2/32        10.0.14.1                        0       64513 64514i
 *>    10.1.3.3/32        10.0.14.1                        0       64513 64514 64515i
 *>    10.1.5.5/32        10.0.14.1                        0       64513 64516i
```

Step 3 **Run the network command to summarize BGP routes.**

You need to summarize routes on R1.

On R1, add the static route 10.1.0.0/16 pointing to interface Null0 and run the **network** command to advertise this route.

```
[R1]ip route-static 10.1.0.0 16 NULL 0
[R1]bgp 64513
[R1-bgp]network 10.1.0.0 255.255.0.0
[R1-bgp]quit
```

Check the BGP routing table of R4 to determine whether the summarized route exists.

```
<R4>display bgp routing-table

 BGP Local router ID is 10.0.4.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete
```

```
Total Number of Routes: 6
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.0.4.4/32        0.0.0.0        0                    0       i
 *>   10.1.0.0/16        10.0.14.1      0                    0       64513i
 *>   10.1.1.1/32        10.0.14.1      0                    0       64513i
 *>   10.1.2.2/32        10.0.14.1                           0       64513 64514i
 *>   10.1.3.3/32        10.0.14.1                           0       64513 64514 64515i
 *>   10.1.5.5/32        10.0.14.1                           0       64513 64516i
```

Set an IP prefix list named **pref_detail_control** to filter the routes to be sent to the BGP peer R4 and prevent the specific routes from being sent out.

```
[R1]ip ip-prefix pref_detail_control index 10 permit 10.1.0.0 8 less-equal 24
[R1]bgp 64513
[R1-bgp]peer 10.0.14.4 ip-prefix pref_detail_control export
[R1-bgp]quit
```

Check the BGP routing table of R4 again. Observe the AS_Path attribute of the summarized route.

```
<R4>display bgp routing-table

 BGP Local router ID is 10.0.4.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn
```

```
*>   10.0.4.4/32    0.0.0.0      0             0    i
*>   10.1.0.0/16    10.0.14.1    0             0    64513i
```

Step 4 **Run the aggregate command to summarize BGP routes.**

Delete the IP prefix list configured in step 3 and the summarized route advertised using the **network** command.

Run the **aggregate** command to summarize routes to 10.1.0.0/16 using the default mode.

```
[R1]bgp 64513
[R1-bgp]undo network 10.1.0.0 255.255.0.0
[R1-bgp]undo peer 10.0.14.4 ip-prefix pref_detail_control export
[R1-bgp]quit
[R1]undo ip ip-prefix pref_detail_control
[R1]undo ip route-static 10.1.0.0 16 NULL 0

[R1]bgp 64513
[R1-bgp]aggregate 10.1.0.0 255.255.0.0
[R1-bgp]quit
```

Check the BGP routing tables of R1 and R4 and observe the Origin attribute of the summarized route.

```
[R1]display bgp routing-table

 BGP Local router ID is 10.1.1.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
```

```
              Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 6
       Network            NextHop        MED        LocPrf    PrefVal Path/Ogn


 *>    10.0.4.4/32        10.0.14.4      0                     0       64512i
 *>    10.1.0.0/16        127.0.0.1                            0       i
 *>    10.1.1.1/32        0.0.0.0        0                     0       i
 *>    10.1.2.2/32        10.0.12.2      0                     0       64514i
 *>    10.1.3.3/32        10.0.12.2                            0       64514 64515i
 *>    10.1.5.5/32        10.0.15.5      0                     0       64516i

<R4>display bgp routing-table


 BGP Local router ID is 10.0.4.4
 Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 6
       Network            NextHop        MED    LocPrf    PrefVal Path/Ogn


 *>    10.0.4.4/32        0.0.0.0        0                 0       i
 *>    10.1.0.0/16        10.0.14.1                        0       64513i
 *>    10.1.1.1/32        10.0.14.1      0                 0       64513i
 *>    10.1.2.2/32        10.0.14.1                        0       64513 64514i
 *>    10.1.3.3/32        10.0.14.1                        0       64513 64514 64515i
 *>    10.1.5.5/32        10.0.14.1                        0       64513 64516i
```

The preceding command output shows that the Origin attribute of the summarized route retains unchanged and is still IGP.

When configuring route summarization on R1, suppress specific routes and advertise only the summarized route.

```
[R1]bgp 64513
[R1-bgp]aggregate 10.1.0.0 255.255.0.0 detail-suppressed
[R1-bgp]quit
```

Check the BGP routing table of R4.

```
[R4]display bgp routing-table

 BGP Local router ID is 10.0.4.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.0.4.4/32        0.0.0.0        0                     0       i
 *>   10.1.0.0/16        10.0.14.1                            0       64513i
```

The preceding command output shows that the BGP routing table of R4 does not contain specific routes.

Check the IP routing table of R1 to view the next hop of the route to 10.1.0.0/16.

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
```

```
------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 21        Routes : 21

Destination/Mask     Proto    Pre   Cost      Flags NextHop        Interface

       10.0.4.4/32  EBGP    255   0            D   10.0.14.4       Serial3/0/0
      10.0.12.0/24  Direct  0     0            D   10.0.12.1       Serial1/0/0
      10.0.12.1/32  Direct  0     0            D   127.0.0.1       Serial1/0/0
      10.0.12.2/32  Direct  0     0            D   10.0.12.2       Serial1/0/0
    10.0.12.255/32  Direct  0     0            D   127.0.0.1       Serial1/0/0
      10.0.14.0/24  Direct  0     0            D   10.0.14.1       Serial3/0/0
      10.0.14.1/32  Direct  0     0            D   127.0.0.1       Serial3/0/0
      10.0.14.4/32  Direct  0     0            D   10.0.14.4       Serial3/0/0
    10.0.14.255/32  Direct  0     0            D   127.0.0.1       Serial3/0/0
      10.0.15.0/24  Direct  0     0            D   10.0.15.1       GigabitEthernet0/0/0
      10.0.15.1/32  Direct  0     0            D   127.0.0.1       GigabitEthernet0/0/0
    10.0.15.255/32  Direct  0     0            D   127.0.0.1       GigabitEthernet0/0/0
       10.1.0.0/16  IBGP    255   0            D   0.0.0.0         NULL0
       10.1.1.1/32  Direct  0     0            D   127.0.0.1       LoopBack0
       10.1.2.2/32  EBGP    255   0            D   10.0.12.2       Serial1/0/0
       10.1.3.3/32  EBGP    255   0            D   10.0.12.2       Serial1/0/0
       10.1.5.5/32  EBGP    255   0            D   10.0.15.5       GigabitEthernet0/0/0
      127.0.0.0/8   Direct  0     0            D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0     0            D   127.0.0.1       InLoopBack0
127.255.255.255/32Direct  0     0            D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct  0     0            D   127.0.0.1       InLoopBack0
```

The route to 10.1.0.0/16 is a summarized route, which is configured on R1. Therefore, the outbound interface is interface Null0. This configuration can prevent routing loops.

Check the BGP routing table of R1 to view specific routes.

```
[R1]display bgp routing-table

 BGP Local router ID is 10.1.1.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 6
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.0.4.4/32        10.0.14.4      0                     0       64512i
 *>   10.1.0.0/16        127.0.0.1                           0       i
 s>   10.1.1.1/32        0.0.0.0        0                    0       i
 s>   10.1.2.2/32        10.0.12.2      0                    0       64514i
 s>   10.1.3.3/32        10.0.12.2                           0       64514 64515i
 s>   10.1.5.5/32        10.0.15.5      0                    0       64516i
```

If the **detail-suppressed** parameter is specified, only summarized routes are sent. Specific routes are marked with s, indicating that they are suppressed during route summarization.

Step 5 **Change the attributes of summarized routes.**

By default, BGP does not advertise the Community attribute to any peer.

Configure R5 and R1 to advertise the Community attribute to R1 and R4 respectively.

[R5]bgp 64516

[R5-bgp]peer 10.0.15.1 advertise-community

[R5-bgp]quit

[R1]bgp 64513

[R1-bgp]peer 10.0.14.4 advertise-community

[R1-bgp]quit

Verify that the Community attribute disappears after routes are summarized.

On R5, add the Community attribute 100 to the route 10.1.5.5/32 advertised by R5 and advertise this route to R1.

[R5]acl number 2000

[R5-acl-basic-2000]rule 0 permit source 10.1.5.5 0

[R5-acl-basic-2000]quit

[R5]route-policy set_comm permit node 10

[R5-route-policy]if-match acl 2000

[R5-route-policy]apply community 100

[R5-route-policy]quit

[R5]bgp 64516

[R5-bgp]peer 10.0.15.1 route-policy set_comm export

[R5-bgp]quit

On R1, check whether the route 10.1.5.5/32 carries the Community attribute 100.

<R1>display bgp routing-table community

BGP Local router ID is 10.1.1.1

Status codes: * - valid, > - best, d - damped,

h - history, i - internal, s - suppressed, S - Stale

```
        Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Community


 *>   10.1.5.0/24        10.0.15.5                            0       <0:100>
```

On R4, check whether the summarized route carries the Community attribute 100.

```
<R4>display bgp routing-table community


 Total Number of Routes: 0
```

The preceding command output shows that R4 does not have any route that carries the Community attribute.

Configure a route policy **add_comm** on R1 to add the Community attribute 100:2 to the summarized route.

```
[R1]acl number 2000
[R1-acl-basic-2000]rule 0 permit source 10.1.0.0 0.0.255.255
[R1-acl-basic-2000]quit
[R1]route-policy add_comm permit node 10
[R1-route-policy]if-match acl 2000
[R1-route-policy]apply community 100:2
[R1-route-policy]quit
[R1]bgp 64513
[R1-bgp]aggregate 10.1.0.0 255.255.0.0 attribute-policy add_comm
```

On R4, check whether the summarized route carries the Community attribute 100:2.

```
<R4>display bgp routing-table community

BGP Local router ID is 10.0.4.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Community

 *>   10.1.0.0/16        10.0.14.1                            0       <100:2>
 *>   10.1.5.5/32        10.0.14.1                            0       <0:100>
```

The preceding command output shows that the summarized route learned by R4 carries the Community attribute 100:2.

Step 6 **Use the AS_Set attribute to configure the AS_Path attribute.**

After routes are summarized, the AS_Path attribute of the summarized route is discarded by default, which may cause a routing loop. To eliminate this risk, add the AS_Set attribute to the summarized route.

Configure R1 to add the AS_Set attribute to the summarized route during route summarization.

```
[R1]bgp 64513
[R1-bgp]aggregate 10.1.0.0 255.255.0.0 detail-suppressed as-set
[R1-bgp]quit
```

Check the AS_Path attribute of the summarized route in the BGP routing tables of R1 and R4.

```
[R1]display bgp routing-table

 BGP Local router ID is 10.1.1.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 6
      Network            NextHop        MED     LocPrf    PrefVal Path/Ogn

 *>   10.0.4.4/32        10.0.14.4       0                  0      64512i
 *>   10.1.0.0/16        127.0.0.1                          0      {64514 64515 64516}i
 s>   10.1.1.1/32        0.0.0.0         0                  0      i
 s>   10.1.2.2/32        10.0.12.2       0                  0      64514i
 s>   10.1.3.3/32        10.0.12.2                          0      64514 64515i
 s>   10.1.5.5/32        10.0.15.5       0                  0      64516i

<R4>display bgp routing-table

 BGP Local router ID is 10.0. 4.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 2
      Network            NextHop        MED LocPrf PrefVal Path/Ogn
```

```
*>   10.0.4.4/32      0.0.0.0     0         0     i
*>   10.1.0.0/16      10.0.14.1             0     64513 {64514 64515 64516}i
```

The AS_Path attribute of the summarized route to which the AS_Set attribute is added contains AS path information of specific routes.

On R3, stop advertising the route 10.1.3.3/32 and reset the peer relationship.

```
[R3]bgp 64515
[R3-bgp]undo network 10.1.3.3 255.255.255.255
[R3-bgp]return
<R3>reset bgp all
```

After the peer relationship is established again, check the AS_Path attribute of the summarized route learned by R4.

```
<R4>display bgp routing-table

 BGP Local router ID is 10.0.4.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 2
      Network            NextHop        MED   LocPrf    PrefVal Path/Ogn

*>   10.0.4.4/32      0.0.0.0        0                0     i
*>   10.1.0.0/16      10.0.14.1                       0     64513 {64514 64516}i
```

The preceding command output shows that the AS_Path attribute does not contain the AS number 64515.

----End

Additional Exercises: Analysis and Verification

After step 6 is complete, can R5 access the loopback interface address of R3?

What are the differences between the **aggregate** and **summary automatic** commands?

Device Configurations

```
<R1>display current-configuration
[V200R007C00SPC600]
#
 sysname R1
#
router id 10.1.1.1
#
acl number 2000
 rule 0 permit source 10.1.0.0 0.0.255.255
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
#
interface GigabitEthernet0/0/0
```

```
 ip address 10.0.15.1 255.255.255.0
#
interface NULL0
#
interface LoopBack0
 ip address 10.1.1.1 255.255.255.255
#
bgp 64513
 peer 10.0.12.2 as-number 64514
 peer 10.0.14.4 as-number 64512
 peer 10.0.15.5 as-number 64516
 #
 ipv4-family unicast
  undo synchronization
  aggregate 10.1.0.0 255.255.0.0 as-set detail-suppressed
  network 10.1.1.1 255.255.255.255
  peer 10.0.12.2 enable
  peer 10.0.14.4 enable
  peer 10.0.14.4 advertise-community
  peer 10.0.15.5 enable
#
route-policy add_comm permit node 10
 if-match acl 2000
 apply community 100:2
#
return

<R2>display current-configuration
[V200R007C00SPC600]
```

```
#
 sysname R2
#
router id 10.1.2.2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface LoopBack0
 ip address 10.1.2.2 255.255.255.255
#
bgp 64514
 peer 10.0.12.1 as-number 64513
 peer 10.0.23.3 as-number 64515
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.2.2 255.255.255.255
  peer 10.0.12.1 enable
  peer 10.0.23.3 enable
#
return

<R3>display current-configuration
```

```
[V200R007C00SPC600]
#
 sysname R3
#
router id 10.1.3.3
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface LoopBack1
 ip address 10.1.3.3 255.255.255.255
#
bgp 64515
 peer 10.0.23.2 as-number 64514
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.23.2 enable
#
return

<R4>display current-configuration
[V200R007C00SPC600]
#
 sysname R4
#
router id 10.0.4.4
#
```

```
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
#
bgp 64512
 peer 10.0.14.1 as-number 64513
 #
 ipv4-family unicast
  undo synchronization
  network 10.0.4.4 255.255.255.255
  peer 10.0.14.1 enable
#
return

<R5>display current-configuration
[V200R007C00SPC600]
#
 sysname R5
#
router id 10.1.5.5
#
acl number 2000
 rule 0 permit source 10.1.5.5 0
#
interface GigabitEthernet0/0/0
 ip address 10.0.15.5 255.255.255.0
```

```
#
interface LoopBack0
 ip address 10.1.5.5 255.255.255.255
#
bgp 64516
 peer 10.0.15.1 as-number 64513
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.5.5 255.255.255.255
  peer 10.0.15.1 enable
  peer 10.0.15.1 advertise-community
peer 10.0.15.1 route-policy set_comm export
#
route-policy set_comm permit node 10
 if-match acl 2000
 apply community 100
#
return
```

Lab 3-3 BGP Attributes and Route Selection 1

Learning Objectives

The objectives of this lab are to learn and understand:

- How to configure the AS_Path attribute to affect route selection
- How to modify the Origin attribute to affect route selection
- How to modify the Local_Pref attribute to affect route selection
- How to modify the MED attribute to affect route selection

Topology

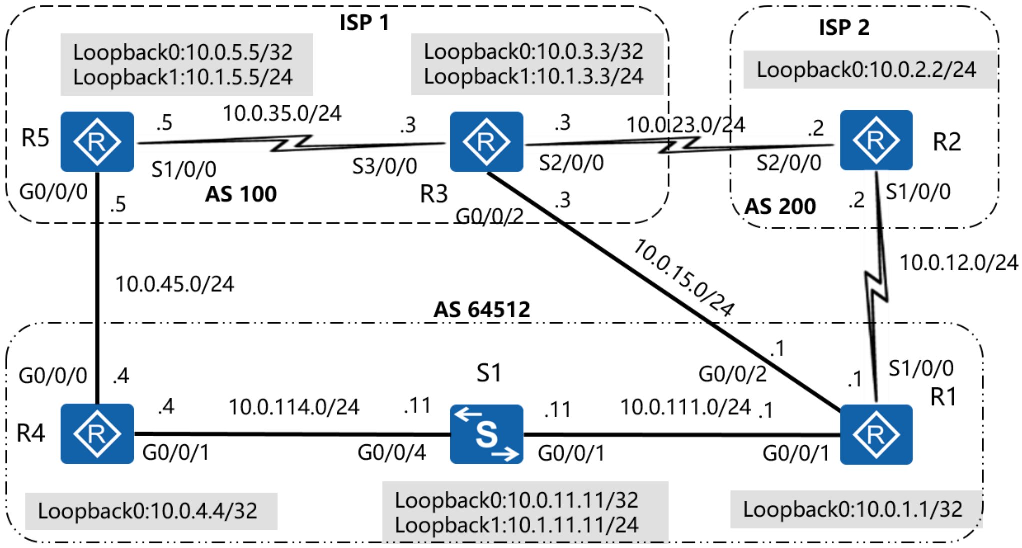

Figure 3-3 BGP attributes and route selection

Scenario

You are a network administrator of a company. The company' s network uses BGP to connect to two Internet Service Providers (ISPs). The company uses a private AS number 64512 and connects to ISP1 through two links, and ISP1 uses the AS number 100. ISP2 uses the AS number 200, and the company leases a line to connect to ISP2. Some Internet users reflect that access to the company website is slow. You change various BGP attributes to optimize route selection.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

Configure IP addresses and masks for physical interfaces and loopback interfaces of all the routers. Each Loopback0 uses the 32-bit mask.

```
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]quit
[R1]interface GigabitEthernet 0/0/2
```

```
[R1-GigabitEthernet0/0/2]ip address 10.0.15.1 24
[R1-GigabitEthernet0/0/2]quit
[R1]interface GigabitEthernet 0/0/1
[R1-GigabitEthernet0/0/1]ip address 10.0.111.1 24
[R1-GigabitEthernet0/0/1]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 32
[R1-LoopBack0]quit

[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]quit
[R2]interface  Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 24
[R2-Serial2/0/0]quit
[R2]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24
[R2-LoopBack0]quit

[R3]interface GigabitEthernet 0/0/2
[R3-GigabitEthernet0/0/2]ip address 10.0.15.3 24
[R3-GigabitEthernet0/0/2]quit
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 24
[R3-Serial2/0/0]quit
[R3]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 24
[R3-Serial3/0/0]quit
[R3]interface loopback 0
```

```
[R3-LoopBack0]ip address 10.0.3.3 32
[R3-LoopBack0]quit

[R4]interface GigabitEthernet 0/0/1
[R4-GigabitEthernet0/0/1]ip address 10.0.114.4 24
[R4-GigabitEthernet0/0/1]quit
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ip address 10.0.45.4 24
[R4-GigabitEthernet0/0/0]quit
[R4]interface loopback 0
[R4-LoopBack0]ip address 10.0.4.4 32
[R4-LoopBack0]quit

[R5]interface Serial 1/0/0
[R5-Serial1/0/0]ip address 10.0.35.5 24
[R5-Serial1/0/0]quit
[R5]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.45.5 24
[R5-GigabitEthernet0/0/0]quit
[R5]interface loopback 0
[R5-LoopBack0]ip address 10.0.5.5 32
[R5-LoopBack0]quit
```

After the configurations are complete, test direct link connectivity.

```
<R1>ping -c 1 10.0.12.2
  PING 10.0.12.2: 56   data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=29 ms
```

```
  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 29/29/29 ms

[R1]ping -c 1 10.0.15.3
  PING 10.0.15.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.15.3: bytes=56 Sequence=1 ttl=255 time=59 ms

  --- 10.0.15.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 59/59/59 ms

<R2>ping -c 1 10.0.23.3
  PING 10.0.23.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.3: bytes=56 Sequence=1 ttl=255 time=32 ms

  --- 10.0.23.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 32/32/32 ms

[R3]ping -c 1 10.0.35.5
  PING 10.0.35.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.35.5: bytes=56 Sequence=1 ttl=255 time=36 ms
```

```
--- 10.0.35.5 ping statistics ---
  1 packet(s) transmitted
  1 packet(s) received
  0.00% packet loss
  round-trip min/avg/max = 36/36/36 ms

<R4>ping -c 1 10.0.45.5
PING 10.0.45.5: 56  data bytes, press CTRL_C to break
  Reply from 10.0.45.5: bytes=56 Sequence=1 ttl=255 time=11 ms

--- 10.0.45.5 ping statistics ---
  1 packet(s) transmitted
  1 packet(s) received
  0.00% packet loss
round-trip min/avg/max = 11/11/11 ms
```

Step 2 **Configure IGP and BGP.**

Configure OSPF in AS 64512 and configure all devices to belong to Area 0.

Run OSPF on the network segments connected to G0/0/1 and Loopback0 of R1.

```
[R1]router id 10.0.1.1
[R1]ospf 1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.0.111.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]quit
 [R1-ospf-1]quit
```

Create VLAN 111 on S1 and configure a VLANIF address for interconnection with R1.

Create VLAN 114 on S1 and configure a VLANIF address for interconnection with R4.

Set the link type of interconnected interfaces to access, and run OSPF on the network segments connected to S1's VLANIF 111, VLANIF 114, and Loopback0.

```
[S1]router id 10.0.11.11
[S1]vlan batch 111 114
[S1]interface vlan 111
[S1-Vlanif111]ip address 10.0.111.11 24
[S1-Vlanif111]quit
[S1]interface  vlan 114
[S1-Vlanif114]ip address 10.0.114.11 24
[S1-Vlanif114]quit
[S1]interface loopback 0
[S1-LoopBack0]ip address 10.0.11.11 32
[S1-LoopBack0]quit
[S1]interface GigabitEthernet 0/0/1
[S1-GigabitEthernet0/0/1]port link-type access
[S1-GigabitEthernet0/0/1]port default vlan 111
[S1-GigabitEthernet0/0/1]quit
[S1]interface GigabitEthernet 0/0/4
[S1-GigabitEthernet0/0/4]port link-type access
[S1-GigabitEthernet0/0/4]port default vlan 114
[S1-GigabitEthernet0/0/4]quit
[S1]ospf 1
[S1-ospf-1]area 0
[S1-ospf-1-area-0.0.0.0]network 10.0.111.11 0.0.0.0
[S1-ospf-1-area-0.0.0.0]network 10.0.114.11 0.0.0.0
[S1-ospf-1-area-0.0.0.0]network 10.0.11.11 0.0.0.0
```

```
[S1-ospf-1-area-0.0.0.0]quit
[S1-ospf-1]quit
```

Run OSPF on the network segments connected to G0/0/1 and Loopback0 of R4.

```
[R4]router id 10.0.4.4
[R4]ospf 1
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]network 10.0.114.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]network 10.0.4.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]quit
[R4-ospf-1]quit
```

Check whether the devices learn the network segment where Loopback0 of other devices resides.

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 18        Routes : 18

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.1/32  Direct  0    0           D   127.0.0.1       LoopBack0
       10.0.4.4/32  OSPF    10   2           D   10.0.111.11     GigabitEthernet0/0/1
     10.0.11.11/32  OSPF    10   1           D   10.0.111.11     GigabitEthernet0/0/1
      10.0.12.0/24  Direct  0    0           D   10.0.12.1       Serial1/0/0
      10.0.12.1/32  Direct  0    0           D   127.0.0.1       Serial1/0/0
      10.0.12.2/32  Direct  0    0           D   10.0.12.2       Serial1/0/0
```

```
     10.0.12.255/32 Direct   0    0        D   127.0.0.1       Serial1/0/0
      10.0.15.0/24  Direct   0    0        D   10.0.15.1    GigabitEthernet0/0/2
      10.0.15.1/32  Direct   0    0        D   127.0.0.1    GigabitEthernet0/0/2
     10.0.15.255/32 Direct   0    0        D   127.0.0.1    GigabitEthernet0/0/2
     10.0.111.0/24  Direct   0    0        D   10.0.111.1   GigabitEthernet0/0/1
     10.0.111.1/32  Direct   0    0        D   127.0.0.1    GigabitEthernet0/0/1
   10.0.111.255/32  Direct   0    0        D   127.0.0.1    GigabitEthernet0/0/1
     10.0.114.0/24  OSPF    10    2        D   10.0.111.11  GigabitEthernet0/0/1
      127.0.0.0/8   Direct   0    0        D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct   0    0        D   127.0.0.1       InLoopBack0
127.255.255.255/32Direct     0    0        D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct     0    0        D   127.0.0.1       InLoopBack0

[S1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 9        Routes : 9

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.1/32  OSPF    10   1           D   10.0.111.1      Vlanif111
       10.0.4.4/32  OSPF    10   1           D   10.0.114.4      Vlanif114
     10.0.11.11/32  Direct  0    0           D   127.0.0.1       LoopBack0
     10.0.111.0/24  Direct  0    0           D   10.0.111.11     Vlanif111
    10.0.111.11/32  Direct  0    0           D   127.0.0.1       Vlanif111
     10.0.114.0/24  Direct  0    0           D   10.0.114.11     Vlanif114
    10.0.114.11/32  Direct  0    0           D   127.0.0.1       Vlanif114
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
```

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


<R4>display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 14        Routes : 14

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.1/32  OSPF   10    2           D   10.0.114.11 GigabitEthernet0/0/1
       10.0.4.4/32  Direct  0    0           D   127.0.0.1       LoopBack0
      10.0.11.11/32 OSPF   10    1           D   10.0.114.11 GigabitEthernet0/0/1
      10.0.45.0/24  Direct  0    0           D   10.0.45.4    GigabitEthernet0/0/0
      10.0.45.4/32  Direct  0    0           D   127.0.0.1    GigabitEthernet0/0/0
    10.0.45.255/32  Direct  0    0           D   127.0.0.1    GigabitEthernet0/0/0
     10.0.111.0/24  OSPF   10    2           D   10.0.114.11 GigabitEthernet0/0/1
     10.0.114.0/24  Direct  0    0           D   10.0.114.4   GigabitEthernet0/0/1
     10.0.114.4/32  Direct  0    0           D   127.0.0.1    GigabitEthernet0/0/1
   10.0.114.255/32  Direct  0    0           D   127.0.0.1    GigabitEthernet0/0/1
       127.0.0.0/8  Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

Configure BGP on R1, R4, and S1, enable them to establish BGP peer relationships using their Loopback0 interfaces, and configure a peer group named **as64512**.

By default, BGP load balancing is disabled. Enable BGP load balancing on all the routers and set the maximum number of equal-cost routes to 4 for load balancing.

```
[R1]bgp 64512
[R1-bgp]group as64512 internal
[R1-bgp]peer 10.0.11.11 group as64512
[R1-bgp]peer 10.0.11.11 connect-interface LoopBack 0
[R1-bgp]maximum load-balancing 4
[R1-bgp]quit

[S1]bgp 64512
[S1-bgp]group as64512 internal
[S1-bgp]peer 10.0.4.4 group as64512
[S1-bgp]peer 10.0.4.4 connect-interface LoopBack 0
[S1-bgp]maximum load-balancing 4
[S1-bgp]peer 10.0.1.1 group as64512
[S1-bgp]peer 10.0.1.1 connect-interface LoopBack 0
[S1-bgp]quit

[R4]bgp 64512
[R4-bgp]group as64512 internal
[R4-bgp]peer 10.0.11.11 group as64512
[R4-bgp]peer 10.0.11.11 connect-interface LoopBack 0
[R4-bgp]maximum load-balancing 4
[R4-bgp]quit
```

Configure EBGP on R1, R2, R3, R4, and R5, and enable these routers to use physical interfaces to establish EBGP peer relationships according to the topology.

```
[R1]bgp 64512
```

```
[R1-bgp]peer 10.0.12.2 as-number 200
[R1-bgp]peer 10.0.15.3 as-number 100
[R1-bgp]quit

[R2]router id 10.0.2.2
[R2]bgp 200
[R2-bgp]peer 10.0.12.1 as-number 64512
[R2-bgp]peer 10.0.23.3 as-number 100
[R2-bgp]maximum load-balancing 4
[R2-bgp]quit

[R3]router id 10.0.3.3
[R3]bgp 100
[R3-bgp]peer 10.0.23.2 as-number 200
[R3-bgp]peer 10.0.35.5 as-number 100
[R3-bgp]peer 10.0.15.1 as-number 64512
[R3-bgp]maximum load-balancing 4
[R3-bgp]quit

[R4]bgp 64512
[R4-bgp]peer 10.0.45.5 as-number 100
[R4-bgp]quit

[R5]router id 10.0.5.5
[R5]bgp 100
[R5-bgp]peer 10.0.35.3 as-number 100
[R5-bgp]peer 10.0.45.4 as-number 64512
[R5-bgp]maximum load-balancing 4
[R5-bgp]quit
```

Step 3 **Configure the AS_Path attribute.**

Create Loopback1 on S1, assign an address 10.1.11.11/24 to Loopback1, and use the **network** command to advertise this address into BGP.

```
[S1]interface loopback 1
[S1-LoopBack1]ip address 10.1.11.11 24
[S1-LoopBack1]quit
[S1]bgp 64512
[S1-bgp]network 10.1.11.11 255.255.255.0
[S1]quit
```

Check the BGP routing table of R2. The following command output shows that the next hop for the route 10.1.11.0/24 is selected based on the AS_Path attribute.

```
[R2]display bgp routing-table

 BGP Local router ID is 10.0.2.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 2
      Network          NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.1.11.0/24     10.0.12.1                            0      64512i
 *                     10.0.23.3                            0      100 64512i
```

Bandwidth between R1 and R4 is limited. It is expected that R2 accesses the network segment 10.1.11.0/24 through AS 100.

You can use the AS_Path attribute to affect route selection.

Create a route policy **as_path** on R1 to add two duplicate AS numbers to the route 10.1.11.0/24.

```
[R1]acl number 2001
[R1-acl-basic-2001]rule 5 permit source 10.1.11.0 0.0.0.255
[R1-acl-basic-2001]quit
[R1]route-policy as_path permit node 10
[R1-route-policy]if-match acl 2001
[R1-route-policy]apply as-path 64512 64512 additive
[R1-route-policy]quit
```

Apply this route policy to R1 so that the AS_Path attribute of the route learned by R2 from R1 has three values.

```
[R1]bgp 64512
[R1-bgp]peer 10.0.12.2 route-policy as_path  export
[R1-bgp]quit
```

Check the BGP routing table of R2.

```
<R2>display bgp routing-table

 BGP Local router ID is 10.0.2.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete
```

```
Total Number of Routes: 2
      Network            NextHop         MED     LocPrf     PrefVal Path/Ogn

*>    10.1.11.0/24       10.0.23.3                          0        100 64512i
*                        10.0.12.1                          0        64512 64512 64512i
```

The preceding command output shows that R2 accesses the network segment 10.1.11.0/24 through AS 100.

Step 4 **Configure the Origin attribute.**

Check the BGP routing table of R3.

```
<R3>display bgp routing-table

 BGP Local router ID is 10.0.3.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 2
      Network            NextHop         MED        LocPrf     PrefVal Path/Ogn

*>    10.1.11.0/24       10.0.15.1                             0        64512i
* i                      10.0.35.5                  100        0        64512i
```

The preceding command output shows that the next hop for the route 10.1.11.0/24 is R1. Analyze the cause.

R3 needs to access AS 64512 through R5. The Origin attribute of the route 10.1.11.0/24 is **IGP**.

Configure a route-policy **22** and change the Origin attribute of the route advertised from R1 to R3 to **incomplete**.

```
[R1]route-policy 22 permit node 10
[R1-route-policy]if-match acl 2001
[R1-route-policy]apply origin incomplete
[R1-route-policy]quit
[R1]bgp 64512
[R1-bgp]peer 10.0.15.3 route-policy 22 export
[R1-bgp]quit
```

After the configured route-policy takes effect, check the BGP routing table of R3.

```
<R3>display bgp routing-table

 BGP Local router ID is 10.0.3.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.1.11.0/24       10.0.35.5                 100        0      64512i
 *                       10.0.15.1                            0      64512?
```

The preceding command output shows that the next hop for the route 10.1.11.0/24 is R5.

Step 5 **Configure the Local_Pref attribute.**

The Local_Pref attribute has a high priority in route selection.

You can change the Local_Pref attribute to affect route selection.

Create Loopback1 on R3, assign an address 10.1.3.3/24 to Loopback1, and advertise this address into BGP.

```
[R3]interface loopback 1
[R3-LoopBack1]ip address 10.1.3.3 255.255.255.0
[R3-LoopBack1]quit
[R3]bgp 100
[R3-bgp]network 10.1.3.3 255.255.255.0
[R3-bgp]quit
```

Create Loopback1 on R5, assign an address 10.1.5.5/24 to Loopback1, and advertise this address into BGP.

```
[R5]interface loopback 1
[R5-LoopBack1]ip address 10.1.5.5 255.255.255.0
[R5-LoopBack1]quit
[R5]bgp 100
[R5-bgp]network 10.1.5.5 24
[R5-bgp]quit
```

Check the BGP routing table of S1.

```
[S1]display bgp routing-table

 BGP Local router ID is 10.0.11.11
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete
```

```
Total Number of Routes: 5
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.1.3.0/24   10.0.1.1        0          100        0      100i
 * i                10.0.4.4                   100        0      100i
 *>i  10.1.5.0/24   10.0.1.1                   100        0      100i
 * i                10.0.4.4        0          100        0      100i
 *>   10.1.11.0/24   0.0.0.0        0                     0      i
```

Traffic to the network segment 10.1.5.0/24 needs to be sent from R4, and traffic to the network segment 10.1.3.0/24 needs to be sent from R1.

Create a route-policy **Pref4** on R4 to match the route 10.1.5.0/24 and change its Local_Pref attribute to 110.

On R1, create a route-policy **Pref1** to match the route 10.1.3.0/24, change its Local_Pref attribute to 110, and apply the route-policy to the IBGP peer group.

```
[R4]acl number 2001
[R4-acl-basic-2001]rule 5 permit source 10.1.5.0 0.0.0.255
[R4-acl-basic-2001]quit
[R4]route-policy Pref4 permit node 10
[R4-route-policy]if-match acl 2001
[R4-route-policy]apply local-preference 110
[R4-route-policy]quit
[R4]route-policy Pref4 permit node 20
[R4-route-policy]quit
[R4]bgp 64512
[R4-bgp]peer as64512 route-policy Pref4 export
[R4-bgp]quit
```

```
[R1]acl number 2002
[R1-acl-basic-2002]rule 5 permit source 10.1.3.0 0.0.0.255
[R1-acl-basic-2002]quit
[R1]route-policy Pref1 permit node 10
[R1-route-policy]if-match acl 2002
[R1-route-policy]apply local-preference 110
[R1-route-policy]quit
[R1]route-policy Pref1 permit node 20
[R1-route-policy]quit
[R1]bgp 64512
[R1-bgp]peer as64512 route-policy Pref1 export
[R1-bgp]quit
```

Check the BGP routing table of S1.

```
[S1]display bgp routing-table

 BGP Local router ID is 10.0.11.11
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? – incomplete

Total Number of Routes: 5
      Network            NextHop        MED        LocPrf     PrefVal Path/Ogn

 *>i  10.1.3.0/24        10.0.1.1       0          110        0       100i
* i                      10.0.4.4                  100        0       100i
 *>i  10.1.5.0/24        10.0.4.4       0          110        0       100i
* i                      10.0.1.1       0          100        0       100i
```

```
*>   10.1.11.0/24       0.0.0.0         0                     0      i
```

The preceding command output shows that routes are selected based on the Local_Pref attribute, and the route with the highest Local_Pref is preferred.

Step 6 **Configure the MED attribute.**

Delete the route-policy configured in step 4 and change the MED attribute to affect route selection.

```
[R1]bgp 64512
[R1-bgp]undo peer 10.0.15.3 route-policy 22 export
[R1-bgp]quit
[R1]undo route-policy 22
```

Create a route-policy **med** on R1 to match the route 10.1.11.0/24, change the MED attribute of this route to 100, and apply this route-policy to R3.

```
[R1]route-policy med permit node 10
[R1-route-policy]if-match acl 2001
[R1-route-policy]apply cost 100
[R1-route-policy]quit
[R1]bgp 64512
[R1-bgp]peer 10.0.15.3 route-policy med export
[R1-bgp]quit
```

Check the BGP routing table of R3.

```
<R3>display bgp routing-table

 BGP Local router ID is 10.0.15.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
```

```
              Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 4
      Network            NextHop          MED          LocPrf     PrefVal Path/Ogn

 *>    10.1.3.0/24  0.0.0.0           0                          0        i
 *>i  10.1.5.0/24   10.0.35.5         0            100          0        i
 *>i  10.1.11.0/24  10.0.35.5                      100          0        64512i
 *                  10.0.15.1         100                        0        64512i

[R3]display bgp routing-table 10.1.11.0

BGP local router ID : 10.0.3.3
Local AS number : 100
Paths: 2 available, 1 best, 1 select
BGP routing table entry information of 10.1.11.0/24:
From: 10.0.35.5 (10.0.5.5)
Route Duration: 00h00m33s
Relay IP Nexthop: 0.0.0.0
Relay IP Out-Interface: Serial3/0/0
Original nexthop: 10.0.35.5
Qos information : 0x0
AS-path 64512, origin igp, localpref 100, pref-val 0, valid, internal, best, select, active, pre 255
Advertised to such 2 peers:
10.0.23.2
10.0.15.1
BGP routing table entry information of 10.1.11.0/24:
From: 10.0.15.1 (10.0.1.1)
Route Duration: 18h52m36s
```

Direct Out-interface: GigabitEthernet0/0/2

Original nexthop: 10.0.15.1

Qos information : 0x0

AS-path 64512, origin igp, MED 100, pref-val 0, valid, external, pre 255, not preferred for MED

Not advertised to any peer yet

The route with the smallest MED value is preferred.

The route selection result in step 6 is the same as that in step 4.

----End

Additional Exercises: Analysis and Verification

After step 6 is complete and S1/0/0 of R1 is shut down, what is the MED value of the route 10.1.11.0/24 learned on R2?

Whether a route-policy can be used to delete an AS from the AS_Path attribute?

Device Configurations

```
<R1>display current-configuration
[V200R007C00SPC600]
#
 sysname R1
#
router id 10.0.2.2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
```

```
 ip address 10.0.14.1 255.255.255.0
#
interface GigabitEthernet0/0/1
 ip address 10.0.111.1 255.255.255.0
#
interface GigabitEthernet0/0/2
 ip address 10.0.15.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
#
bgp 64512
 peer 10.0.12.2 as-number 200
 peer 10.0.15.3 as-number 100
 group as64512 internal
 peer 10.0.11.11 as-number 64512
 peer 10.0.11.11 group as64512
 peer 10.0.11.11 connect-interface LoopBack0
 #
 ipv4-family unicast
  undo synchronization
  maximum load-balancing 4
  peer 10.0.12.2 enable
  peer 10.0.12.2 route-policy as_path export
  peer 10.0.15.3 enable
  peer 10.0.15.3 route-policy med export
  peer as64512 enable
  peer as64512 route-policy Pref1 export
  peer 10.0.11.11 enable
```

```
  peer 10.0.11.11 group as64512
#
ospf 1
 area 0.0.0.0
  network 10.0.1.1 0.0.0.0
  network 10.0.111.1 0.0.0.0
#
route-policy as_path permit node 10
 if-match acl 2001
 apply as-path 64512 64512 additive
#
route-policy Pref1 permit node 10
 if-match acl 2002
 apply local-preference 110
#
route-policy Pref1 permit node 20
#
route-policy med permit node 10
 if-match acl 2001
 apply cost 100
#
return

<R2>display current-configuration
[V200R007C00SPC600]
#
 sysname R2
#
router id 10.0.2.2
```

```
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.0
#
bgp 200
 peer 10.0.12.1 as-number 64512
 peer 10.0.23.3 as-number 100
 #
 ipv4-family unicast
  undo synchronization
  maximum load-balancing 4
  peer 10.0.12.1 enable
  peer 10.0.23.3 enable
#
return

<R3>display current-configuration
[V200R007C00SPC600]
#
 sysname R3
#
```

```
router id 10.0.3.3
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
#
interface GigabitEthernet0/0/2
 ip address 10.0.15.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
#
interface LoopBack1
 ip address 10.1.3.3 255.255.255.0
#
bgp 100
 peer 10.0.15.1 as-number 64512
 peer 10.0.23.2 as-number 200
 peer 10.0.35.5 as-number 100
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.3.0 255.255.255.0
  maximum load-balancing 4
  peer 10.0.15.1 enable
```

```
  peer 10.0.23.2 enable
  peer 10.0.35.5 enable
#
return

<R4>display current-configuration
[V200R007C00SPC600]
#
 sysname R4
#
router id 10.0.4.4
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.45.4 255.255.255.0
#
interface GigabitEthernet0/0/1
 ip address 10.0.114.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
#
bgp 64512
 peer 10.0.45.5 as-number 100
 group as64512 internal
 peer 10.0.11.11 as-number 64512
```

```
 peer 10.0.11.11 group as64512
 peer 10.0.11.11 connect-interface LoopBack0
 #
 ipv4-family unicast
  undo synchronization
  maximum load-balancing 4
  peer 10.0.45.5 enable
  peer as64512 enable
  peer as64512 route-policy Pref4 export
  peer 10.0.11.11 enable
  peer 10.0.11.11 group as64512
#
ospf 1
 area 0.0.0.0
  network 10.0.114.4 0.0.0.0
  network 10.0.4.4 0.0.0.0
#
route-policy Pref4 permit node 10
 if-match acl 2001
 apply local-preference 110
#
route-policy Pref4 permit node 20
#
return

<R5>display current-configuration
[V200R007C00SPC600]
#
 sysname R5
```

```
#
router id 10.0.5.5
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.45.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.255
#
interface LoopBack1
 ip address 10.1.5.5 255.255.255.0
#
bgp 100
 peer 10.0.35.3 as-number 100
 peer 10.0.45.4 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.5.0 255.255.255.0
  maximum load-balancing 4
  peer 10.0.35.3 enable
  peer 10.0.45.4 enable
#
return
```

Lab 3-4 BGP Attributes and Route Selection 2 (Optional)

Learning Objectives

The objectives of this lab are to learn and understand:

- How to modify the Community attribute to affect route selection
- How to use a route-policy to filter BGP routes

Topology

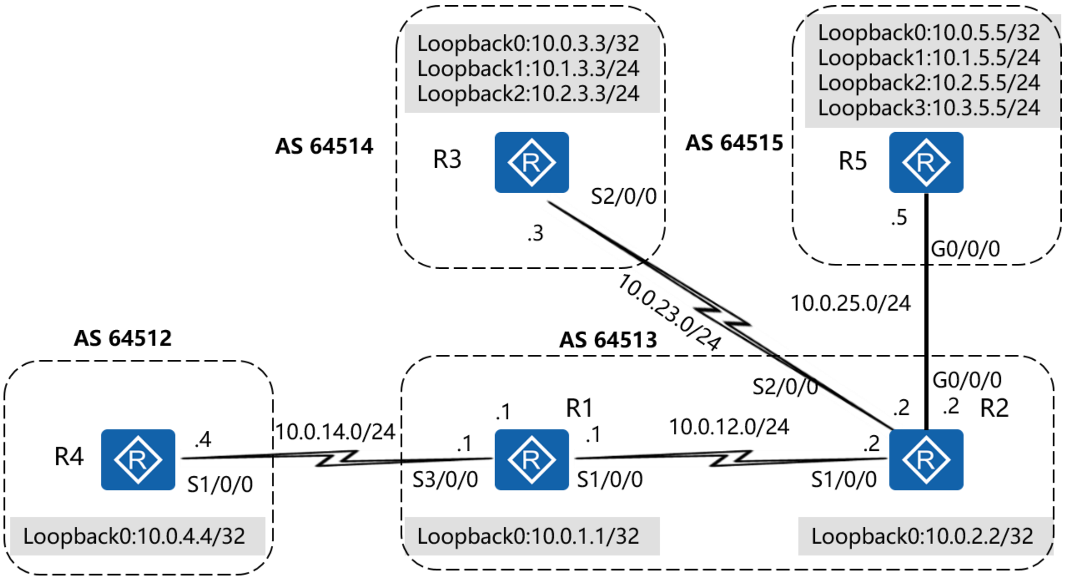

Figure 3-4 BGP attributes and route selection 2

Scenario

You are a network administrator of a company. The company's network uses BGP for interconnection. Design AS numbers according to the topology. To ensure network security, some departments of branches cannot communicate with each other. To control routing information transmission, you need to use the Community attribute to filter BGP routes.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

Configure IP addresses and masks for physical interfaces and Loopback0 of all the routers. Each Loopback0 uses the 32-bit mask.

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 255.255.255.0
[R1-Serial1/0/0]quit
[R1]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 255.255.255.0
[R1-Serial3/0/0]quit
[R1]interface loopback 0
[R1-LoopBack0]ip address 10.0.1.1 255.255.255.255
[R1-LoopBack0]quit

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 255.255.255.0
[R2-Serial1/0/0]quit
[R2]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 255.255.255.0
[R2-Serial2/0/0]quit
[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.25.2 255.255.255.0
[R2-GigabitEthernet0/0/0]quit
[R2]interface loopback 0
[R2-LoopBack0]ip address 10.0.2.2 255.255.255.255
[R2-LoopBack0]quit
```

```
<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 255.255.255.0
[R3-Serial2/0/0]
[R3]interface loopback 0
[R3-LoopBack0]ip address 10.0.3.3 255.255.255.255
[R3-LoopBack0]quit

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 255.255.255.0
[R4-Serial1/0/0]quit
[R4]interface loopback 0
[R4-LoopBack0]ip address 10.0.4.4 255.255.255.255
[R4-LoopBack0]quit

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.25.5 255.255.255.0
[R5-GigabitEthernet0/0/0]quit
[R5]interface loopback 0
[R5-LoopBack0]ip address 10.0.5.5 255.255.255.255
[R5-LoopBack0]quit
```

After the configurations are complete, test direct link connectivity.

```
<R1>ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=40 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 40/40/40 ms

<R1>ping -c 1 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=61 ms

  --- 10.0.14.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 61/61/61 ms

<R2>ping -c 1 10.0.25.5
  PING 10.0.25.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.25.5: bytes=56 Sequence=1 ttl=255 time=14 ms

  --- 10.0.25.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 14/14/14 ms
```

```
<R2>ping -c 1 10.0.23.3
  PING 10.0.23.2: 56   data bytes, press CTRL_C to break
    Reply from 10.0.23.3: bytes=56 Sequence=1 ttl=255 time=2 ms
  --- 10.0.23.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 2/2/2 ms
```

Step 2 **Configure BGP.**

Establish an IBGP peer relationship between R1 and R2 and establish EBGP peer relationships between other routers.

```
[R1]router id 10.0.1.1
[R1]bgp 64513
[R1-bgp]peer 10.0.12.2 as-number 64513
[R1-bgp]peer 10.0.14.4 as-number 64512
[R1-bgp]quit

[R2]router id 10.0.2.2
[R2]bgp 64513
[R2-bgp]peer 10.0.12.1 as-number 64513
[R2-bgp]peer 10.0.23.3 as-number 64514
[R2-bgp]peer 10.0.25.5 as-number 64515
[R2-bgp]quit

[R3]router id 10.0.3.3
[R3]bgp 64514
```

[R3-bgp]peer 10.0.23.2 as-number 64513

[R3-bgp]quit

[R4]router id 10.0.4.4

[R4]bgp 64512

[R4-bgp]peer 10.0.14.1 as-number 64513

[R4-bgp]quit

[R5]router id 10.0.5.5

[R5]bgp 64515

[R5-bgp]peer 10.0.25.2 as-number 64513

[R5-bgp]quit

After the configurations are complete, check whether BGP peer relationships are established between routers.

[R1]display bgp peer

BGP local router ID : 10.0.1.1

Local AS number : 64513

Total number of peers : 2 Peers in established state : 2

Peer	V	AS	MsgRcvd	MsgSent	OutQ	Up/Down	State	PrefRcv
10.0.12.2	4	64513	5	6	0	00:03:28	Established	0
10.0.14.4	4	64512	2	3	0	00:00:39	Established	0

[R2]display bgp peer

```
 BGP local router ID : 10.0.2.2
 Local AS number : 64513
 Total number of peers : 3                Peers in established state : 3

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State  PrefRcv

  10.0.12.1       4       64513        6        5     0 00:04:00 Established        0
  10.0.23.3       4       64514        4        6     0 00:02:44 Established        0
  10.0.25.5       4       64515        2        3     0 00:00:41 Established        0

[R3]display bgp peer

 BGP local router ID : 10.0.3.3
 Local AS number : 64514
 Total number of peers : 1                Peers in established state : 1

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State  PrefRcv

  10.0.23.2       4       64513        4        4     0 00:02:59 Established        0

 [R4]display bgp peer

 BGP local router ID : 10.0.4.4
 Local AS number : 64512
 Total number of peers : 1                Peers in established state : 1

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State  PrefRcv

  10.0.14.1       4       64513        3        3     0 00:01:40 Established        0
```

```
[R5]display bgp peer
 BGP local router ID : 10.0.5.5
 Local AS number : 64515
 Total number of peers : 1                Peers in established state : 1

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State  PrefRcv

  10.0.25.2       4       64513        3        3     0  00:01:23 Established        0
```

The preceding command output shows that all BGP peer relationships are in Established state.

Step 3 **Configure a common Community attribute.**

Create Loopback1, Loopack2, and Loopback3 on R5, assign addresses 10.1.5.5/24, 10.2.5.5/24, and 10.3.5.5/24 to the three loopback interfaces respectively, and run the **network** command to advertise these addresses into BGP.

```
[R5]interface loopback 1
[R5-LoopBack1]ip address 10.1.5.5 255.255.255.0
[R5-LoopBack1]quit
[R5]interface loopback 2
[R5-LoopBack2]ip address 10.2.5.5 255.255.255.0
[R5-LoopBack2]quit
[R5]interface loopback 3
[R5-LoopBack3]ip address 10.3.5.5 255.255.255.0
[R5-LoopBack3]quit
[R5]bgp 64515
[R5-bgp]network 10.1.5.5 255.255.255.0
```

[R5-bgp]network 10.2.5.5 255.255.255.0

[R5-bgp]network 10.3.5.5 255.255.255.0

[R5-bgp]quit

[R2]bgp 64513

[R2-bgp]peer 10.0.12.1 next-hop-local

[R2-bgp]quit

Check whether the routing information is correctly transmitted on R2 and R4.

[R2]**d**isplay bgp routing-table

BGP Local router ID is 10.0.2.2

Status codes: * - valid, > - best, d - damped,

h - history, i - internal, s - suppressed, S - Stale

Origin : i - IGP, e - EGP, ? - incomplete

Total Number of Routes: 3

	Network	NextHop	MED	LocPrf	PrefVal	Path/Ogn
*>	10.1.5.0/24	10.0.25.5	0		0	64515i
*>	10.2.5.0/24	10.0.25.5	0		0	64515i
*>	10.3.5.0/24	10.0.25.5	0		0	64515i

[R4]display bgp routing-table

BGP Local router ID is 10.0.4.4

Status codes: * - valid, > - best, d - damped,

h - history, i - internal, s - suppressed, S - Stale

```
        Origin : i - IGP, e - EGP, ? - incomplete


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


 *>   10.1.5.0/24        10.0.14.1                             0       64513 64515i
 *>   10.2.5.0/24        10.0.14.1                             0       64513 64515i
 *>   10.3.5.0/24        10.0.14.1                             0       64513 64515i
```

On R5, create a route-policy **comm_r5** to add the Community attribute 100 to the route 10.1.5.0/24.

```
[R5]acl number 2000
[R5-acl-basic-2000]rule 0 permit source 10.1.5.0 0.0.0.255
[R5-acl-basic-2000]quit
[R5]route-policy comm_r5 permit node 10
[R5-route-policy]if-match acl 2000
[R5-route-policy]apply community 100
[R5-route-policy]quit

[R5]bgp 64515
[R5-bgp]peer 10.0.25.2 route-policy comm_r5 export
[R5-bgp]quit
```

Configure all BGP peers to advertise the Community attribute between each other.

```
[R1]bgp 64513
[R1-bgp]peer 10.0.14.4 advertise-community
[R1-bgp]peer 10.0.12.2 advertise-community
[R1-bgp]quit
```

```
[R2]bgp 64513
[R2-bgp]peer 10.0.12.1 advertise-community
[R2-bgp]peer 10.0.23.3 advertise-community
[R2-bgp]peer 10.0.25.5 advertise-community
[R2-bgp]quit

[R3]bgp 64514
[R3-bgp]peer 10.0.23.2 advertise-community
[R3-bgp]quit

[R4]bgp 64512
[R4-bgp]peer 10.0.14.1 advertise-community
[R4-bgp]quit

[R5]bgp 64515
[R5-bgp]peer 10.0.25.2 advertise-community
[R5-bgp]quit
```

On R2 and R4, check whether the Community attribute is transmitted normally.

```
<R2>display bgp routing-table community

 BGP Local router ID is 10.0.2.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 5
```

```
      Network            NextHop         MED        LocPrf     PrefVal Community

*>    10.1.5.0/24        10.0.25.5       0                      0       <0:100>

<R4>display bgp routing-table community

 BGP Local router ID is 10.0.4.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 5
      Network            NextHop         MED        LocPrf     PrefVal Community

*>    10.1.5.0/24        10.0.25.5       0                      0       <0:100>
```

Step 4 **Configure a special Community attribute.**

Use a route-policy on R5 to add a special Community attribute **no-export** for the route 10.2.5.0/24 and add a special Community attribute **no-advertise** for the route 10.3.5.0/24.

You only need to add two new nodes and if-match clauses to the route-policy **comm_r5** created on R5.

```
[R5]acl 2001
[R5-acl-basic-2001]rule 0 permit source 10.2.5.0 0.0.0.255
[R5-acl-basic-2001]quit
[R5]route-policy comm_r5 permit node 20
[R5-route-policy]if-match acl 2001
[R5-route-policy]apply community no-export
```

```
[R5-route-policy]quit
[R5]acl number 2002
[R5-acl-basic-2002]rule 0 permit source 10.3.5.0 0.0.0.255
[R5-acl-basic-2002]quit
[R5]route-policy comm_r5 permit node 30
[R5-route-policy]if-match acl 2002
[R5-route-policy]apply community no-advertise
[R5-route-policy]quit
```

Check the Community attribute of the routes learned by R2.

```
<R2>dis bgp routing-table community

 BGP Local router ID is 10.0.2.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 4
      Network            NextHop        MED        LocPrf    PrefVal Community

 *>   10.1.5.0/24        10.0.25.5      0                    0       <0:100>
 *>   10.2.5.0/24        10.0.25.5      0                    0       no-export
 *>   10.3.5.0/24        10.0.25.5      0                    0       no-advertise
```

Check the BGP routing tables of R2, R1, and R4 to observe transmission of the routes 10.1.5.0/24, 10.2.5.0/24, and 10.3.5.0/24.

```
<R2>display bgp routing-table
```

```
 BGP Local router ID is 10.0.2.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

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

 *>   10.1.5.0/24        10.0.25.5      0                     0      64515i
 *>   10.2.5.0/24        10.0.25.5      0                     0      64515i
 *>   10.3.5.0/24        10.0.25.5      0                     0      64515i

<R1>display bgp routing-table

 BGP Local router ID is 10.0.1.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.1.5.0/24        10.0.12.2      0          100        0      64515i
 *>i  10.2.5.0/24        10.0.12.2      0          100        0      64515i

[R4]display bgp routing-table

 BGP Local router ID is 10.0.4.4
 Status codes: * - valid, > - best, d - damped,
```

h - history, i - internal, s - suppressed, S - Stale

Origin : i - IGP, e - EGP, ? - incomplete

Total Number of Routes: 1

	Network	NextHop	MED	LocPrf	PrefVal	Path/Ogn
*>	10.1.5.0/24	10.0.14.1			0	64513 64515i

The preceding command output shows that R2 does not advertise the route 10.2.5.0/24 carrying the special Community attribute **no-export** outside its AS but advertises it to R1 in the same AS. R2 does not advertise the route 10.3.5.0/24 carrying the special Community attribute **no-advertise** to any BGP peer.

Step 5 **Configure the Community attribute for route summarization.**

Create Loopback1 and Loopack2 on R3, assign addresses 10.1.3.3/24 and 10.2.3.3/24 to the two loopback interfaces respectively, and run the **network** command to advertise these addresses into BGP.

```
[R3]interface LoopBack 1
[R3-LoopBack1]ip address 10.1.3.3 255.255.255.0
[R3-LoopBack1]quit
[R3]interface loopback 2
[R3-LoopBack2]ip address 10.2.3.3 255.255.255.0
[R3-LoopBack2]quit
[R3]bgp 64514
[R3-bgp]network 10.1.3.3 255.255.255.0
[R3-bgp]network 10.2.3.3 255.255.255.0
[R3-bgp]quit
```

Currently, the following requirements need to be met: The route 10.1.5.0/24 advertised by R5 and the route 10.2.3.0/24 advertised by R3 need to be summarized into a Class A network segment 10.0.0.0/8. The specific routes need to be suppressed when the summarized route is advertised, and the Community attribute of the summarized route advertised to R4 must be 200. Specific routes 10.1.3.0/24 must be advertised to R4.

To meet these requirements, create a route-policy **comm_r3** on R3 to add the Community attribute 100 to the route 10.2.3.0/24 advertised by R3.

```
[R3]acl number 2001
[R3-acl-basic-2001]rule 0 permit source 10.2.3.0 0.0.0.255
[R3-acl-basic-2001]quit
[R3]route-policy comm_r3 permit node 10
[R3-route-policy]if-match acl 2001
[R3-route-policy]apply community 100
[R3-route-policy]quit
[R3]route-policy comm_r3 permit node 20
[R3-route-policy]quit
[R3]bgp 64514
[R3-bgp]peer 10.0.23.2 route-policy comm_r3 export
[R3-bgp]quit
```

On R1, check whether the learned routes 10.1.5.0/24 and 10.2.3.0/24 carry the Community attribute 100.

```
<R1>display bgp routing-table community

 BGP Local router ID is 10.0.1.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete
```

Total Number of Routes: 3

	Network	NextHop	MED	LocPrf	PrefVal	Community
*>i	10.1.5.0/24	10.0.12.2	0	100	0	<0:100>
*>i	10.2.3.0/24	10.0.12.2	0	100	0	<0:100>
*>i	10.2.5.0/24	10.0.12.2	0	100	0	no-export

Create a community filter to filter the route with the Community attribute 100.

[R1]ip community-filter 1 permit 100

Create a route-policy **match_comm** to match the route with the Community attribute 100.

[R1]route-policy match_comm permit node 10

[R1-route-policy]if-match community-filter 1

[R1-route-policy]quit

Create a route-policy **add_comm** to add the Community attribute 200:1 to the summarized route.

[R1]route-policy add_comm permit node 10

[R1-route-policy]apply community 200:1 additive

[R1-route-policy]quit

On R1, summarize the route matching the route-policy **match_comm** and use the route-policy **add_comm** to add the Community attribute.

[R1]bgp 64513

[R1-bgp]aggregate 10.0.0.0 255.0.0.0 detail-suppressed origin-policy match_comm attribute-policy add_comm

```
[R1-bgp]quit
```

Check the BGP routing table of R4.

```
<R4>display bgp routing-table

 BGP Local router ID is 10.0.4.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.0.0.0           10.0.14.1                            0      64513i
 *>   10.1.3.0/24        10.0.14.1                            0      64513 64514i
```

On R4, check the Community attribute of the summarized route.

```
<R4>display bgp routing-table community

 BGP Local router ID is 10.0.4.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 1
      Network            NextHop        MED        LocPrf    PrefVal Community

 *>   10.0.0.0           10.0.14.1                            0      <200:1>
```

----End

Additional Exercises: Analysis and Verification

If in step 4, the Community attribute of the route 10.2.5.0/24 is also changed to no advertise, check the BGP routing tables of R2, R1, and R4 to analyze transmission of the routes 10.1.5.0/24, 10.2.5.0/24, and 10.3.5.0/24.

Consider how to retain specific routes of the two routes 10.1.3.0/24 and 10.2.3.0/24 and suppress only specific routes of the route 10.1.5.0/24 on R4.

Device Configurations

```
<R1>display current-configuration
[V200R007C00SPC600]
#
 sysname R1
#
router id 10.0.1.1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
#
bgp 64513
```

```
 peer 10.0.12.2 as-number 64513
 peer 10.0.14.4 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  aggregate 10.0.0.0 255.0.0.0 detail-suppressed origin-policy match_comm attribute-policy add_comm
  peer 10.0.12.2 enable
  peer 10.0.12.2 advertise-community
  peer 10.0.14.4 enable
  peer 10.0.14.4 advertise-community
#
route-policy match_comm permit node 10
 if-match community-filter 1
#
route-policy add_comm permit node 10
 apply community 200:1 additive
#
 ip community-filter 1 permit 100
#
return

<R2>display current-configuration
[V200R007C00SPC600]
#
 sysname R2
#
router id 10.0.2.2
#
interface Serial1/0/0
```

```
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.255
#
bgp 64513
 peer 10.0.12.1 as-number 64513
 peer 10.0.23.3 as-number 64514
 peer 10.0.25.5 as-number 64515
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.12.1 enable
  peer 10.0.12.1 next-hop-local
  peer 10.0.12.1 advertise-community
  peer 10.0.23.3 enable
  peer 10.0.23.3 advertise-community
  peer 10.0.25.5 enable
  peer 10.0.25.5 advertise-community
#
return
```

```
<R3>display current-configuration
[V200R007C00SPC600]
#
 sysname R3
#
router id 10.0.3.3
#
acl number 2001
 rule 0 permit source 10.2.3.0 0.0.0.255
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
#
interface LoopBack1
 ip address 10.1.3.3 255.255.255.0
#
interface LoopBack2
 ip address 10.2.3.3 255.255.255.0
#
bgp 64514
 peer 10.0.23.2 as-number 64513
 #
 ipv4-family unicast
  undo synchronization
```

```
  network 10.1.3.0 255.255.255.0
  network 10.2.3.0 255.255.255.0
  peer 10.0.23.2 enable
  peer 10.0.23.2 route-policy comm_r3 export
  peer 10.0.23.2 advertise-community
#
route-policy comm_r3 permit node 10
 if-match acl 2001
 apply community 100
#
route-policy comm_r3 permit node 20
#
return

<R4>display current-configuration
[V200R007C00SPC600]
#
 sysname R4
#
router id 10.0.4.4
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
#
bgp 64512
```

```
 peer 10.0.14.1 as-number 64513
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.14.1 enable
  peer 10.0.14.1 advertise-community
#
Return

<R5>display current-configuration
[V200R007C00SPC600]
#
 sysname R5
#
router id 10.0.5.5
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.255
#
interface LoopBack1
 ip address 10.1.5.5 255.255.255.0
#
interface LoopBack2
 ip address 10.2.5.5 255.255.255.0
#
interface LoopBack3
```

```
 ip address 10.3.5.5 255.255.255.0
#
bgp 64515
 peer 10.0.25.2 as-number 64513
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.5.0 255.255.255.0
  network 10.2.5.0 255.255.255.0
  network 10.3.5.0 255.255.255.0
  peer 10.0.25.2 enable
  peer 10.0.25.2 route-policy comm_r5 export
  peer 10.0.25.2 advertise-community
#
route-policy comm_r5 permit node 10
 if-match acl 2000
 apply community 100
#
route-policy comm_r5 permit node 20
 if-match acl 2001
 apply community no-export
#
route-policy comm_r5 permit node 30
 if-match acl 2002
 apply community no-advertise
#
return
```

Lab 3-5 BGP Multi-homing

Learning Objectives

The objectives of this lab are to learn and understand:

- How to use only default routes in BGP multi-homing scenario
- How to use default routes to filter some routes in BGP multi-homing scenario
- How to use only BGP routes in BGP multi-homing scenario

Topology

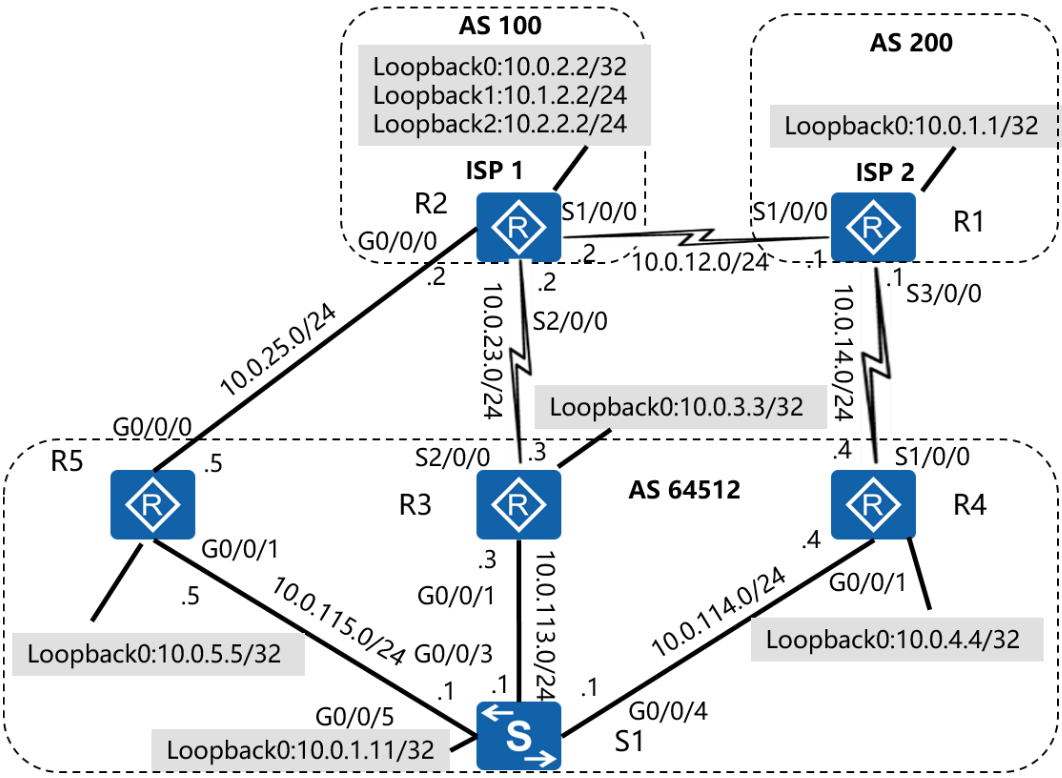

Figure 3-5 BGP multi-homing

Scenario

You are a network administrator of a company. The company's network uses BGP to connect to ISP1. The company uses a private AS number 64512 and connects to ISP1 through two routers, and ISP1 uses the AS number 100. The company initially used default routes to connect to the Internet through ISP1. With the development of the company, this Internet access mode cannot meet requirements. You need to import some Internet routes into the AS of the company. After a certain period, the company leases a line to connect to ISP2, whose AS number is 200. Finally, the company builds a BGP multi-homing network.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

Configure IP addresses and masks for physical interfaces and loopback interfaces of all the routers. Each Loopback0 uses the 32-bit mask.

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]quit
[R1]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 24
[R1-Serial3/0/0]quit
[R1]interface loopback 0
[R1-LoopBack0]ip address 10.0.1.1 32
[R1-LoopBack0]quit

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]quit
[R2]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 24
```

```
[R2-Serial2/0/0]quit
[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.25.2 24
[R2-GigabitEthernet0/0/0]quit
[R2]interface loopback 0
[R2-LoopBack0]ip address 10.0.2.2 32
[R2-LoopBack0]quit

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 24
[R3-Serial2/0/0]quit
[R3]interface GigabitEthernet 0/0/1
[R3-GigabitEthernet0/0/1]ip address 10.0.113.3 24
[R3-GigabitEthernet0/0/1]quit
[R3]interface loopback 0
[R3-LoopBack0]ip address 10.0.3.3 32
[R3-LoopBack0]quit

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 24
[R4-Serial1/0/0]quit
[R4]interface GigabitEthernet 0/0/1
[R4-GigabitEthernet0/0/1]ip address 10.0.114.4 24
[R4-GigabitEthernet0/0/1]quit
[R4]interface LoopBack 0
```

```
[R4-LoopBack0]ip address 10.0.4.4 32
[R4-LoopBack0]quit

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.25.5 24
[R5-GigabitEthernet0/0/0]quit
[R5]interface GigabitEthernet 0/0/1
[R5-GigabitEthernet0/0/1]ip address 10.0.115.5 24
[R5-GigabitEthernet0/0/1]quit
[R5]interface loopback 0
[R5-LoopBack0]ip address 10.0.5.5 32
[R5-LoopBack0]quit
```

After the configurations are complete, test direct link connectivity.

```
<R1>ping -c 1 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=33 ms

  --- 10.0.14.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 33/33/33 ms

<R1>ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=34 ms
```

```
  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 34/34/34 ms

<R2>ping -c 1 10.0.25.5
  PING 10.0.25.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.25.5: bytes=56 Sequence=1 ttl=255 time=13 ms

  --- 10.0.25.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 13/13/13 ms

<R2>ping -c 1 10.0.23.3
  PING 10.0.23.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.3: bytes=56 Sequence=1 ttl=255 time=39 ms

  --- 10.0.23.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 39/39/39 ms
```

Step 2 **Configure IGP and BGP.**

Configure OSPF in AS 64512 and configure all devices to belong to Area 0.

Run OSPF on the network segments connected to G0/0/1 and Loopback0 of R3.

```
[R3]router id 10.0.3.3
[R3]ospf 1
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]network 10.0.113.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]quit
[R3-ospf-1]quit
```

Run OSPF on the network segments connected to G0/0/1 and Loopback0 of R4.

```
[R4]router id 10.0.4.4
[R4]ospf 1
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]network 10.0.114.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]network 10.0.4.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]quit
[R4-ospf-1]quit
```

Run OSPF on the network segments connected to G0/0/1 and Loopback0 of R5.

```
[R5]router id 10.0.5.5
[R5]ospf 1
[R5-ospf-1]area 0
[R5-ospf-1-area-0.0.0.0]network 10.0.115.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]network 10.0.5.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]quit
[R5-ospf-1]quit
```

Create VLAN 13 on S1 and configure a VLANIF address for interconnection with R3.

Create VLAN 14 on S1 and configure a VLANIF address for interconnection with R4.

Create VLAN 15 on S1 and configure a VLANIF address for interconnection with R5.

Set the link type of interconnected interfaces to access, and run OSPF on the network segments connected to S1's VLANIF 13, VLANIF 14, VLANIF 15, and Loopback0.

```
[S1]vlan batch 13 to 15
[S1]interface vlan 13
[S1-Vlanif13]ip address 10.0.113.1 255.255.255.0
[S1-Vlanif13]quit
[S1]interface vlan 14
[S1-Vlanif14]ip address 10.0.114.1 255.255.255.0
[S1-Vlanif14]quit
[S1]interface vlan 15
[S1-Vlanif15]ip address 10.0.115.1 255.255.255.0
[S1-Vlanif15]quit
[S1]interface GigabitEthernet 0/0/3
[S1-GigabitEthernet0/0/3]port link-type access
[S1-GigabitEthernet0/0/3]port default vlan 13
[S1-GigabitEthernet0/0/3]quit
[S1]interface GigabitEthernet 0/0/4
[S1-GigabitEthernet0/0/4]port link-type access
[S1-GigabitEthernet0/0/4]port default vlan 14
[S1-GigabitEthernet0/0/4]quit
[S1]interface GigabitEthernet 0/0/5
[S1-GigabitEthernet0/0/5]port link-type access
[S1-GigabitEthernet0/0/5]port default vlan 15
[S1-GigabitEthernet0/0/5]quit
```

```
[S1]interface loopback 0
[S1-LoopBack0]ip address 10.0.1.11 32
[S1-LoopBack0]quit
[S1]router id 10.0.1.11
[S1]ospf 1
[S1-ospf-1]area 0
[S1-ospf-1-area-0.0.0.0]network 10.0.113.1 0.0.0.0
[S1-ospf-1-area-0.0.0.0]network 10.0.114.1 0.0.0.0
[S1-ospf-1-area-0.0.0.0]network 10.0.115.1 0.0.0.0
[S1-ospf-1-area-0.0.0.0]network 10.0.1.11 0.0.0.0
[S1-ospf-1-area-0.0.0.0]quit
[S1-ospf-1]quit
```

Check whether the devices learn the network segment connected to Loopback0 of other devices.

```
<R3>display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 17          Routes : 17

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.11/32  OSPF    10    1           D    10.0.113.1   GigabitEthernet0/0/1
       10.0.3.3/32  Direct  0     0           D    127.0.0.1        LoopBack0
       10.0.4.4/32  OSPF    10    2           D    10.0.113.1   GigabitEthernet0/0/1
       10.0.5.5/32  OSPF    10    2           D    10.0.113.1   GigabitEthernet0/0/1
      10.0.23.0/24  Direct  0     0           D    10.0.23.3        Serial2/0/0
```

```
      10.0.23.2/32  Direct   0    0           D   10.0.23.2       Serial2/0/0
      10.0.23.3/32  Direct   0    0           D   127.0.0.1       Serial2/0/0
    10.0.23.255/32  Direct   0    0           D   127.0.0.1       Serial2/0/0
     10.0.113.0/24  Direct   0    0           D   10.0.113.3   GigabitEthernet0/0/1
     10.0.113.3/32  Direct   0    0           D   127.0.0.1    GigabitEthernet0/0/1
   10.0.113.255/32  Direct   0    0           D   127.0.0.1    GigabitEthernet0/0/1
     10.0.114.0/24  OSPF     10   2           D   10.0.113.1   GigabitEthernet0/0/1
     10.0.115.0/24  OSPF     10   2           D   10.0.113.1   GigabitEthernet0/0/1
      127.0.0.0/8   Direct   0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct   0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32Direct     0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct     0    0           D   127.0.0.1       InLoopBack0

<R4>display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 17        Routes : 17

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.11/32  OSPF     10   1           D   10.0.114.1   GigabitEthernet0/0/1
       10.0.3.3/32  OSPF     10   2           D   10.0.114.1   GigabitEthernet0/0/1
       10.0.4.4/32  Direct   0    0           D   127.0.0.1       LoopBack0
       10.0.5.5/32  OSPF     10   2           D   10.0.114.1   GigabitEthernet0/0/1
      10.0.14.0/24  Direct   0    0           D   10.0.14.4       Serial1/0/0
      10.0.14.1/32  Direct   0    0           D   10.0.14.1       Serial1/0/0
      10.0.14.4/32  Direct   0    0           D   127.0.0.1       Serial1/0/0
    10.0.14.255/32  Direct   0    0           D   127.0.0.1       Serial1/0/0
```

```
    10.0.113.0/24 OSPF    10    2           D   10.0.114.1   GigabitEthernet0/0/1
    10.0.114.0/24 Direct  0     0           D   10.0.114.4   GigabitEthernet0/0/1
    10.0.114.4/32 Direct  0     0           D   127.0.0.1    GigabitEthernet0/0/1
  10.0.114.255/32 Direct  0     0           D   127.0.0.1    GigabitEthernet0/0/1
    10.0.115.0/24 OSPF    10    2           D   10.0.114.1   GigabitEthernet0/0/1
      127.0.0.0/8 Direct  0     0           D   127.0.0.1        InLoopBack0
     127.0.0.1/32 Direct  0     0           D   127.0.0.1        InLoopBack0
127.255.255.255/32Direct  0     0           D   127.0.0.1        InLoopBack0
255.255.255.255/32Direct  0     0           D   127.0.0.1        InLoopBack0

<R5>display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 16       Routes : 16

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

     10.0.1.11/32 OSPF    10    1           D   10.0.115.1   GigabitEthernet0/0/1
      10.0.3.3/32 OSPF    10    2           D   10.0.115.1   GigabitEthernet0/0/1
      10.0.4.4/32 OSPF    10    2           D   10.0.115.1   GigabitEthernet0/0/1
      10.0.5.5/32 Direct  0     0           D   127.0.0.1        LoopBack0
     10.0.25.0/24 Direct  0     0           D   10.0.25.5    GigabitEthernet0/0/0
     10.0.25.5/32 Direct  0     0           D   127.0.0.1    GigabitEthernet0/0/0
   10.0.25.255/32 Direct  0     0           D   127.0.0.1    GigabitEthernet0/0/0
    10.0.113.0/24 OSPF    10    2           D   10.0.115.1   GigabitEthernet0/0/1
    10.0.114.0/24 OSPF    10    2           D   10.0.115.1   GigabitEthernet0/0/1
    10.0.115.0/24 Direct  0     0           D   10.0.115.5   GigabitEthernet0/0/1
    10.0.115.5/32 Direct  0     0           D   127.0.0.1    GigabitEthernet0/0/1
```

```
10.0.115.255/32 Direct  0    0        D   127.0.0.1     GigabitEthernet0/0/1
      127.0.0.0/8   Direct  0    0        D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0        D   127.0.0.1       InLoopBack0
127.255.255.255/32Direct  0    0        D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct  0    0        D   127.0.0.1       InLoopBack0

[S1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 12        Routes : 12

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.11/32  Direct  0    0           D   127.0.0.1       LoopBack0
       10.0.3.3/32  OSPF    10   1           D   10.0.113.3      Vlanif13
       10.0.4.4/32  OSPF    10   1           D   10.0.114.4      Vlanif14
       10.0.5.5/32  OSPF    10   1           D   10.0.115.5      Vlanif15
     10.0.113.0/24  Direct  0    0           D   10.0.113.1      Vlanif13
     10.0.113.1/32  Direct  0    0           D   127.0.0.1       Vlanif13
     10.0.114.0/24  Direct  0    0           D   10.0.114.1      Vlanif14
     10.0.114.1/32  Direct  0    0           D   127.0.0.1       Vlanif14
     10.0.115.0/24  Direct  0    0           D   10.0.115.1      Vlanif15
     10.0.115.1/32  Direct  0    0           D   127.0.0.1       Vlanif15
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

Configure EBGP on R2, R3, and R5.

Enable R2, R3, and R5 to establish EBGP peer relationships using physical interfaces according to the topology. Do not run BGP on SI.

```
[R2]router id 10.0.2.2
[R2]bgp 100
[R2-bgp]peer 10.0.25.5 as-number 64512
[R2-bgp]peer 10.0.23.3 as-number 64512
[R2-bgp]quit

[R3]bgp 64512
[R3-bgp]peer 10.0.23.2 as-number 100
[R3-bgp]quit

[R5]bgp 64512
[R5-bgp]peer 10.0.25.2 as-number 100
[R5-bgp]quit
```

After the configurations are complete, check whether EBGP peer relationships are established.

```
[R2]display bgp peer

 BGP local router ID : 10.0.2.2
 Local AS number : 100
 Total number of peers : 2                 Peers in established state : 2

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down    State       PrefRcv

  10.0.23.3       4       64512        7        9     0 00:05:55 Established        0
  10.0.25.5       4       64512        6        7     0 00:04:17 Established        0
```

[R3]display bgp peer

BGP local router ID : 10.0.3.3

Local AS number : 64512

Total number of peers : 1 Peers in established state : 1

Peer	V	AS	MsgRcvd	MsgSent	OutQ	Up/Down	State	PrefRcv
10.0.23.2	4	100	8	8	0	00:06:09	Established	0

<R5>display bgp peer

BGP local router ID : 10.0.5.5

Local AS number : 64512

Total number of peers : 1 Peers in established state : 1

Peer	V	AS	MsgRcvd	MsgSent	OutQ	Up/Down	State	PrefRcv
10.0.25.2	4	100	7	7	0	00:05:31	Established	0

Step 3 **Use only default routes to connect to a single ISP.**

By default, BGP load balancing is disabled. Enable BGP load balancing on all the routers and set the maximum number of equal-cost routes to 4 for load balancing.

[R1]router id 10.0.1.1

[R1]bgp 200

[R1-bgp]maximum load-balancing 4

[R1-bgp]quit

```
[R2]bgp 100
[R2-bgp]maximum load-balancing 4
[R2-bgp]quit

[R3]bgp 64512
[R3-bgp]maximum load-balancing 4
[R3-bgp]quit

[R4]bgp 64512
[R4-bgp]maximum load-balancing 4
[R4-bgp]quit

[R5]bgp 64512
[R5-bgp]maximum load-balancing 4
[R5-bgp]quit
```

Create Loopback1 and Loopback2 on R2 and assign addresses 10.1.2.2/24 and 10.2.2.2/24 to them. Run the **network** command to advertise the two addresses into BGP.

```
[R2]interface LoopBack 1
[R2-LoopBack1]ip address 10.1.2.2 24
[R2-LoopBack1]quit
[R2]interface LoopBack 2
[R2-LoopBack2]ip address 10.2.2.2 24
[R2-LoopBack2]quit

[R2]bgp 100
```

```
[R2-bgp]network 10.1.2.0 255.255.255.0
[R2-bgp]network 10.2.2.0 255.255.255.0
[R2-bgp]quit
```

Check whether R3 and R5 learn the two routes to 10.1.2.2/24 and 10.2.2.2/24.

```
[R3]display bgp routing-table

 BGP Local router ID is 10.0.3.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.1.2.0/24        10.0.23.2      0                     0       100i
 *>   10.2.2.0/24        10.0.23.2      0                     0       100i

<R5>display bgp  routing-table

 BGP Local router ID is 10.0.5.5
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn
```

*>	10.1.2.0/24	10.0.25.2	0	0	100i
*>	10.2.2.0/24	10.0.25.2	0	0	100i

The link from R1 to ISP1 is the primary link, and the link from R5 to ISP1 is the backup link. On R3 and R5, run the **import-route** command to import OSPF routes into BGP.

```
[R3]bgp 64512
[R3-bgp]import-route ospf 1
[R3-bgp]quit

[R5]bgp 64512
[R5-bgp]import-route ospf 1
[R5-bgp]quit
```

On R3 and R5, configure forcible advertisement of default routes into Area 0 and set the route type as Type 1 external route. Set the costs of default routes advertised by R3 and R5 to 20 and 40 respectively.

```
[R3]ospf 1
[R3-ospf-1]default-route-advertise always cost 20 type 1
[R3-ospf-1]quit

[R5]ospf 1
[R5-ospf-1]default-route-advertise always cost 40 type 1
[R5-ospf-1]quit
```

Check the IP routing table of S1.

```
[S1]display ip routing-table
Route Flags: R - relay, D - download to fib
```

```
------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 14        Routes : 14

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        0.0.0.0/0   O_ASE   150  21          D   10.0.113.3      Vlanif13
      10.0.1.0/24   Direct  0    0           D   10.0.1.11       LoopBack0
     10.0.1.11/32   Direct  0    0           D   127.0.0.1       LoopBack0
      10.0.3.3/32   OSPF    10   1           D   10.0.113.3      Vlanif13
      10.0.4.4/32   OSPF    10   1           D   10.0.114.4      Vlanif14
      10.0.5.5/32   OSPF    10   1           D   10.0.115.5      Vlanif15
    10.0.113.0/24   Direct  0    0           D   10.0.113.1      Vlanif13
    10.0.113.1/32   Direct  0    0           D   127.0.0.1       Vlanif13
    10.0.114.0/24   Direct  0    0           D   10.0.114.1      Vlanif14
    10.0.114.1/32   Direct  0    0           D   127.0.0.1       Vlanif14
    10.0.115.0/24   Direct  0    0           D   10.0.115.1      Vlanif15
    10.0.115.1/32   Direct  0    0           D   127.0.0.1       Vlanif15
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
```

On S1, check the route to 10.1.2.2.

```
[S1]tracert 10.1.2.2
 traceroute to  10.1.2.2(10.1.2.2), max hops: 30 ,packet length: 40
 1 10.0.113.3 10 ms  1 ms  1 ms
 2 10.0.23.2 40 ms  20 ms  20 ms
```

The preceding command output shows that S1 uses the default route learned from R3. That is, S1 accesses 10.1.2.2 through the primary link.

Shut down S2/0/0 of R3 to simulate a failure of the link from the company to ISP1.

```
[R3]interface s2/0/0

[R3-Serial2/0/0]shutdown

[R3-Serial2/0/0]quit
```

After route convergence is complete, check the IP routing table of S1. Check connectivity to 10.1.2.2.

```
[S1]display ip routing-table

Route Flags: R - relay, D - download to fib

------------------------------------------------------------------------------

Routing Tables: Public

         Destinations : 13          Routes : 13


Destination/Mask    Proto   Pre  Cost        Flags NextHop         Interface


         0.0.0.0/0   O_ASE   150   21              D    10.0.113.3      Vlanif13

       10.0.1.11/32  Direct   0    0               D    127.0.0.1       InLoopBack0

        10.0.3.3/32  OSPF     10   1               D    10.0.113.3      Vlanif13

        10.0.4.4/32  OSPF     10   1               D    10.0.114.4      Vlanif14

        10.0.5.5/32  OSPF     10   1               D    10.0.115.5      Vlanif15

      10.0.113.0/24  Direct   0    0               D    10.0.113.1      Vlanif13

      10.0.113.1/32  Direct   0    0               D    127.0.0.1       InLoopBack0

      10.0.114.0/24  Direct   0    0               D    10.0.114.1      Vlanif14

      10.0.114.1/32  Direct   0    0               D    127.0.0.1       InLoopBack0

      10.0.115.0/24  Direct   0    0               D    10.0.115.1      Vlanif15

      10.0.115.1/32  Direct   0    0               D    127.0.0.1       InLoopBack0
```

```
    127.0.0.0/8   Direct   0    0         D   127.0.0.1       InLoopBack0
    127.0.0.1/32  Direct   0    0         D   127.0.0.1       InLoopBack0

[S1]ping 10.1.2.2
  PING 10.1.2.2: 56   data bytes, press CTRL_C to break
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out

  --- 10.1.2.2 ping statistics ---
    5 packet(s) transmitted
    0 packet(s) received
    100.00% packet loss
```

The preceding command output shows that the IP routing table of S1 does not change and S1 still accesses the target network through R3.

When the uplink fails, S1 selects the default route advertised by R3 after comparing the costs of the default routes advertised by R3 and R5. Therefore, the network cannot operate normally.

Restore S2/0/0 of R3 and shut down G0/0/1 of R3 to simulate a downlink failure of R3.

Check route convergence and connectivity.

```
[R3]interface s2/0/0
[R3-Serial2/0/0]undo shutdown
[R3-Serial2/0/0]quit
```

```
[R3]interface g0/0/1
[R3-GigabitEthernet0/0/1]shutdown
[R3-GigabitEthernet0/0/1]quit

[S1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 10        Routes : 10

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        0.0.0.0/01  O_ASE   150  41          D   10.0.115.5      Vlanif15
      10.0.1.11/32  Direct  0    0           D   127.0.0.1       InLoopBack0
       10.0.4.4/32  OSPF    10   1           D   10.0.114.4      Vlanif14
       10.0.5.5/32  OSPF    10   1           D   10.0.115.5      Vlanif15
     10.0.114.0/24  Direct  0    0           D   10.0.114.1      Vlanif14
     10.0.114.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.115.0/24  Direct  0    0           D   10.0.115.1      Vlanif15
     10.0.115.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
       127.0.0.0/8  Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0

[S1]ping 10.1.2.2
  PING 10.1.2.2: 56  data bytes, press CTRL_C to break
    Reply from 10.1.2.2: bytes=56 Sequence=1 ttl=254 time=1 ms
    Reply from 10.1.2.2: bytes=56 Sequence=2 ttl=254 time=1 ms
    Reply from 10.1.2.2: bytes=56 Sequence=3 ttl=254 time=1 ms
    Reply from 10.1.2.2: bytes=56 Sequence=4 ttl=254 time=1 ms
```

```
    Reply from 10.1.2.2: bytes=56 Sequence=5 ttl=254 time=1 ms

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

S1 selects the default route learned from R5. That is, S1 accesses the target network through the backup link.

Restore G0/0/1 of R3.

```
[R3]interface g0/0/1
[R3-GigabitEthernet0/0/1]undo shutdown
[R3-GigabitEthernet0/0/1]quit
```

Step 4 **Connect to a single ISP using default routes and some filtered routes.**

Establish IBGP peer relationships between R3 and S1, between R4 and S1, and between R5 and S1, and specify the **next-hop-local** parameter to ensure that S1 can learn the route Update message sent from the ISP.

```
[R3]bgp 64512
[R3-bgp]peer 10.0.113.1 as-number 64512
[R3-bgp]peer 10.0.113.1 next-hop-local
[R3-bgp]quit

[R4]bgp 64512
[R4-bgp]peer 10.0.114.1 as-number 64512
```

```
[R4-bgp]peer 10.0.114.1 next-hop-local
[R4-bgp]quit

[R5]bgp 64512
[R5-bgp]peer 10.0.115.1 as-number 64512
[R5-bgp]peer 10.0.115.1 next-hop-local
[R5-bgp]quit

[S1]bgp 64512
[S1-bgp]peer 10.0.113.3 as-number 64512
[S1-bgp]peer 10.0.114.4 as-number 64512
[S1-bgp]peer 10.0.115.5 as-number 64512
[S1-bgp]quit
```

Check whether S1 learns the routes 10.1.2.0/24 and 10.2.2.0/24.

```
[S1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 15        Routes : 15

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        0.0.0.0/0   O_ASE   150  21          D   10.0.113.3      Vlanif13
     10.0.1.11/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.3.3/32   OSPF    10   1           D   10.0.113.3      Vlanif13
      10.0.4.4/32   OSPF    10   1           D   10.0.114.4      Vlanif14
      10.0.5.5/32   OSPF    10   1           D   10.0.115.5      Vlanif15
```

10.0.113.0/24	Direct	0	0	D	10.0.113.1	Vlanif13
10.0.113.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.114.0/24	Direct	0	0	D	10.0.114.1	Vlanif14
10.0.114.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.115.0/24	Direct	0	0	D	10.0.115.1	Vlanif15
10.0.115.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.1.2.0/24	BGP	255	0	RD	10.0.113.3	Vlanif13
10.2.2.0/24	BGP	255	0	RD	10.0.113.3	Vlanif13
127.0.0.0/8	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0

To affect route selection through BGP, configure a route-policy **policy_r3** on R3 to filter the route 10.1.2.0/24.

[R3]acl number 2001

[R3-acl-basic-2001]rule 0 permit source 10.1.2.0 0.0.0.255

[R3-acl-basic-2001]quit

[R3]route-policy policy_r3 deny node 10

[R3-route-policy]if-match acl 2001

[R3-route-policy]quit

[R3]route-policy policy_r3 permit node 20

[R3-route-policy]bgp 64512

[R3-bgp]peer 10.0.113.1 route-policy policy_r3 export

[R3-bgp]quit

Configure a route-policy **policy_r5** on R5 to filter the route 10.2.2.0/24.

[R5]acl number 2001

[R5-acl-basic-2001]rule 0 permit source 10.2.2.0 0.0.0.255

[R5-acl-basic-2001]quit

[R5]route-policy policy_r5 deny node 10

[R5-route-policy]if-match acl 2001

[R5-route-policy]quit

[R5]route-policy policy_r5 permit node 20

[R5-route-policy]quit

[R5]bgp 64512

[R5-bgp]peer 10.0.115.1 route-policy policy_r5 export

[R5-bgp]quit

Check the IP routing table of S1.

```
[S1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 15        Routes : 15

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        0.0.0.0/0   O_ASE   150  21          D   10.0.113.3      Vlanif13
      10.0.1.11/32  Direct  0    0           D   127.0.0.1       InLoopBack0
       10.0.3.3/32  OSPF    10   1           D   10.0.113.3      Vlanif13
       10.0.4.4/32  OSPF    10   1           D   10.0.114.4      Vlanif14
       10.0.5.5/32  OSPF    10   1           D   10.0.115.5      Vlanif15
     10.0.113.0/24  Direct  0    0           D   10.0.113.1      Vlanif13
     10.0.113.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.114.0/24  Direct  0    0           D   10.0.114.1      Vlanif14
     10.0.114.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.115.0/24  Direct  0    0           D   10.0.115.1      Vlanif15
```

```
10.0.115.1/32 Direct   0    0     D   127.0.0.1      InLoopBack0
10.1.2.0/24  BGP      255  0     RD  10.0.115.5     Vlanif15
10.2.2.0/24  BGP      255  0     RD  10.0.113.3     Vlanif13
127.0.0.0/8   Direct   0    0     D   127.0.0.1      InLoopBack0
127.0.0.1/32  Direct   0    0     D   127.0.0.1      InLoopBack0
```

The preceding command output shows that the next hop for the route 10.1.2.0/24 is R5 and the next hop for the route 10.2.2.0/24 is R3.

Shut down S2/0/0 of R3.

```
[R3]interface s2/0/0
[R3-Serial2/0/0]shutdown
[R3-Serial2/0/0]quit
```

Check the IP routing table of S1 and test connectivity to 10.1.2.2.

```
[S1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 14        Routes : 14

Destination/Mask    Proto   Pre  Cost       Flags NextHop         Interface

        0.0.0.0/0   O_ASE   150  21          D   10.0.113.3      Vlanif13
     10.0.1.11/32  Direct   0    0           D   127.0.0.1       InLoopBack0
      10.0.3.3/32  OSPF     10   1           D   10.0.113.3      Vlanif13
      10.0.4.4/32  OSPF     10   1           D   10.0.114.4      Vlanif14
      10.0.5.5/32  OSPF     10   1           D   10.0.115.5      Vlanif15
    10.0.113.0/24  Direct   0    0           D   10.0.113.1      Vlanif13
```

10.0.113.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.114.0/24	Direct	0	0	D	10.0.114.1	Vlanif14
10.0.114.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.115.0/24	Direct	0	0	D	10.0.115.1	Vlanif15
10.0.115.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.1.2.0/24	BGP	255	0	RD	10.0.115.5	Vlanif15
127.0.0.0/8	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0

The IP routing table of S1 has only one BGP route 10.1.2.0/24 because a route-policy is configured on R5 to filter the BGP route 10.2.2.0/24.

```
[S1]ping 10.1.2.2
  PING 10.1.2.2: 56   data bytes, press CTRL_C to break
    Reply from 10.1.2.2: bytes=56 Sequence=1 ttl=254 time=1 ms
    Reply from 10.1.2.2: bytes=56 Sequence=2 ttl=254 time=1 ms
    Reply from 10.1.2.2: bytes=56 Sequence=3 ttl=254 time=1 ms
    Reply from 10.1.2.2: bytes=56 Sequence=4 ttl=254 time=1 ms
    Reply from 10.1.2.2: bytes=56 Sequence=5 ttl=254 time=1 ms

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

Restore S2/0/0 of R3.

```
[R3]interface s2/0/0
[R3-Serial2/0/0]undo shutdown
```

[R3-Serial2/0/0]quit

Step 5 **Connect to multiple ISPs using BGP routes.**

The company applies for an Internet line to connect to ISP2. To select routes among BGP routes, delete the default routes advertised by OSPF.

```
[R3]ospf
[R3-ospf-1]undo default-route-advertise
[R3-ospf-1]quit

[R5]ospf
[R5-ospf-1]undo default-route-advertise
[R5-ospf-1]quit
```

Delete the route-policies configured on R3 and R5.

```
[R3]undo route-policy policy1

[R5]undo route-policy policy2
```

On R3 and R5, delete the command used to import OSPF routes into BGP.

```
[R3]bgp 64512
[R3-bgp]undo import-route ospf 1
[R3-bgp]quit

[R5]bgp 64512
[R5-bgp]undo import-route ospf 1
[R5-bgp]quit
```

Establish EBGP peer relationships between R1 and R2 and between R1 and R4 so that ISP2 can also transmit the routes 10.1.2.0/24 and 10.2.2.0/24.

```
[R1]bgp 200
[R1-bgp]peer 10.0.12.2 as-number 100
[R1-bgp]peer 10.0.14.4 as-number 64512
[R1-bgp]quit

[R2]bgp 100
[R2-bgp]peer 10.0.12.1 as-number 200
[R2-bgp]quit

[R4]bgp 64512
[R4-bgp]peer 10.0.14.1 as-number 200
[R4-bgp]quit
```

On S1, check the routes 10.1.2.0/24 and 10.2.2.0/24 and analyze current route selection rules.

```
[S1]display bgp routing-table

 Total Number of Routes: 6

 BGP Local router ID is 10.0.1.11
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn
```

```
*>i  10.1.2.0/24      10.0.113.3     0          100        0      100i
* i                   10.0.115.5     0          100        0      100i
* i                   10.0.114.4                100        0      200 100i
*>i  10.2.2.0/24      10.0.113.3     0          100        0      100i
* i                   10.0.115.5     0          100        0      100i
* i                   10.0.114.4                100        0      200 100i
```

The company wants to access 10.2.2.0/24 through the new line connected to ISP2. Configure a route-policy **policy_r4** on R4 and change the Local_Pref attribute of the route 10.2.2.0/24 to 150.

```
[R4]acl number 2001
[R4-acl-basic-2001]rule 0 permit source 10.2.2.0 0.0.0.255
[R4-acl-basic-2001]quit
[R4]route-policy policy_r4 permit node 10
[R4-route-policy]if-match acl 2001
[R4-route-policy]apply local-preference 150
[R4-route-policy]quit
[R4]route-policy policy_r4 permit node 20
[R4-route-policy]quit
```

Configure R4 to advertise this route-policy to S1.

```
[R4]bgp 64512
[R4-bgp]peer 10.0.114.1 route-policy policy_r4 export
[R4-bgp]quit
```

Check the BGP routing table of S1.

```
[S1]display bgp routing-table
```

```
Total Number of Routes: 6

BGP Local router ID is 10.0.1.11
Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete

     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

*>i  10.1.2.0/24        10.0.113.3     0          100       0       100i
* i                     10.0.115.5     0          100       0       100i
* i                     10.0.114.4                100       0       200 100i
*>i  10.2.2.0/24        10.0.114.4                150       0       200 100i
* i                     10.0.113.3     0          100       0       100i
* i                     10.0.115.5     0          100       0       100i
```

The preceding command output shows that S1 accesses 10.2.2.0/24 through the route obtained from ISP2 connected to R4.

Shut down S1/0/0 of R4 to simulate a failure.

```
[R4]interface s1/0/0
[R4-Serial1/0/0]shutdown
[R4-Serial1/0/0]quit
```

Check the BGP routing table of S1.

```
[S1]display bgp routing-table

Total Number of Routes: 4
```

BGP Local router ID is 10.0.1.11

Status codes: * - valid, > - best, d - damped,

h - history, i - internal, s - suppressed, S - Stale

Origin : i - IGP, e - EGP, ? - incomplete

	Network	NextHop	MED	LocPrf	PrefVal	Path/Ogn
*>i	10.1.2.0/24	10.0.113.3	0	100	0	100i
* i		10.0.115.5	0	100	0	100i
*>i	10.2.2.0/24	10.0.113.3	0	100	0	100i
* i		10.0.115.5	0	100	0	100i

The preceding command output shows that S1 obtains the routes 10.1.2.0/24 and 10.2.2.0/24 through ISP1 connected to R3.

Enable S1/0/0 of R4.

```
[R4]interface s1/0/0
[R4-Serial1/0/0]undo shutdown
[R4-Serial1/0/0]quit
```

Check the BGP routing table of S1 to determine whether the failure is rectified.

```
[S1]display bgp routing-table
```

Total Number of Routes: 6

BGP Local router ID is 10.0.1.11

Status codes: * - valid, > - best, d - damped,

h - history, i - internal, s - suppressed, S - Stale

```
         Origin : i - IGP, e - EGP, ? - incomplete

     Network          NextHop       MED        LocPrf     PrefVal Path/Ogn

 *>i  10.1.2.0/24     10.0.113.3    0          100        0       100i
 * i                  10.0.115.5    0          100        0       100i
 * i                  10.0.114.4               100        0       200 100i
 *>i  10.2.2.0/24     10.0.114.4               150        0       200 100i
 * i                  10.0.113.3    0          100        0       100i
 * i                  10.0.115.5    0          100        0       100i
```

----End

Additional Exercises: Analysis and Verification

After S2/0/0 of R3 is shut down in step 3, and the primary link from the company to ISP1 fails, the backup link between R5 and ISP1 still works normally. How to address the connectivity problem?

This company is dual-homed to two ISPs. Analyze how to load balance incoming traffic of the same network segment.

Device Configurations

```
<R1>display current-configuration
[V200R007C00SPC600]
#
 sysname R1
#
router id 10.0.1.1
#
interface Serial1/0/0
```

```
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
#
bgp 200
 peer 10.0.12.2 as-number 100
 peer 10.0.14.4 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  maximum load-balancing 4
  peer 10.0.12.2 enable
  peer 10.0.14.4 enable
#
return

<R2>display current-configuration
[V200R007C00SPC600]
#
 sysname R2
#
router id 10.0.2.2
```

```
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.255
#
interface LoopBack1
 ip address 10.1.2.2 255.255.255.0
#
interface LoopBack2
 ip address 10.2.2.2 255.255.255.0
#
bgp 100
 peer 10.0.12.1 as-number 200
 peer 10.0.23.3 as-number 64512
 peer 10.0.25.5 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.2.0 255.255.255.0
```

```
  network 10.2.2.0 255.255.255.0
  maximum load-balancing 4
  peer 10.0.12.1 enable
  peer 10.0.23.3 enable
  peer 10.0.25.5 enable
#
return

<R3>display current-configuration
[V200R007C00SPC600]
#
 sysname R3
#
router id 10.0.3.3
#
acl number 2001
 rule 0 permit source 10.1.2.0 0.0.0.255
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface GigabitEthernet0/0/1
 ip address 10.0.113.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
#
```

```
bgp 64512
 peer 10.0.23.2 as-number 100
 peer 10.0.113.1 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  maximum load-balancing 4
  peer 10.0.23.2 enable
  peer 10.0.113.1 enable
  peer 10.0.113.1 next-hop-local
#
ospf 1
 area 0.0.0.0
  network 10.0.113.3 0.0.0.0
  network 10.0.3.3 0.0.0.0
#
return

<R4>display current-configuration
[V200R007C00SPC600]
#
 sysname R4
#
router id 10.0.4.4
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
```

```
#
interface GigabitEthernet0/0/1
 ip address 10.0.114.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
#
bgp 64512
 peer 10.0.14.1 as-number 200
 peer 10.0.114.1 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  maximum load-balancing 4
  peer 10.0.14.1 enable
  peer 10.0.114.1 enable
  peer 10.0.114.1 route-policy policy_r4 export
  peer 10.0.114.1 next-hop-local
#
ospf 1
 area 0.0.0.0
  network 10.0.114.4 0.0.0.0
  network 10.0.4.4 0.0.0.0
#
route-policy policy_r4 permit node 10
 if-match acl 2001
 apply local-preference 150
route-policy policy_r4 permit node 20
#
```

Return

```
<R5>display current-configuration
[V200R007C00SPC600]
#
 sysname R5
#
router id 10.0.5.5
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.5 255.255.255.0
#
interface GigabitEthernet0/0/1
 ip address 10.0.115.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.255
#
bgp 64512
 peer 10.0.25.2 as-number 100
 peer 10.0.115.1 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  maximum load-balancing 4
  peer 10.0.25.2 enable
  peer 10.0.115.1 enable
  peer 10.0.115.1 next-hop-local
```

```
#
ospf 1
 area 0.0.0.0
  network 10.0.115.5 0.0.0.0
  network 10.0.5.5 0.0.0.0
#
return

<S1>display current-configuration
#
!Software Version V100R005C01SPC100
 sysname S1
#
router id 10.0.1.11
#
interface Vlanif13
 ip address 10.0.113.1 255.255.255.0
#
interface Vlanif14
 ip address 10.0.114.1 255.255.255.0
#
interface Vlanif15
 ip address 10.0.115.1 255.255.255.0
#
interface GigabitEthernet0/0/3
 port link-type access
 port default vlan 13
#
```

```
interface GigabitEthernet0/0/4
 port link-type access
 port default vlan 14
#
interface GigabitEthernet0/0/5
 port link-type access
 port default vlan 15
#
interface LoopBack0
 ip address 10.0.1.11 255.255.255.255
#
bgp 64512
 peer 10.0.113.3 as-number 64512
 peer 10.0.114.4 as-number 64512
 peer 10.0.115.5 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.113.3 enable
  peer 10.0.114.4 enable
  peer 10.0.115.5 enable
#
ospf 1
 area 0.0.0.0
  network 10.0.113.1 0.0.0.0
  network 10.0.114.1 0.0.0.0
  network 10.0.115.1 0.0.0.0
  network 10.0.1.11 0.0.0.0
#
```

```
return
```

Lab 3-6 BGP Troubleshooting

Learning Objectives

The objectives of this lab are to learn and understand:

- How to troubleshoot a failure to establish a BGP peer relationship
- How to use BGP debugging commands

Topology

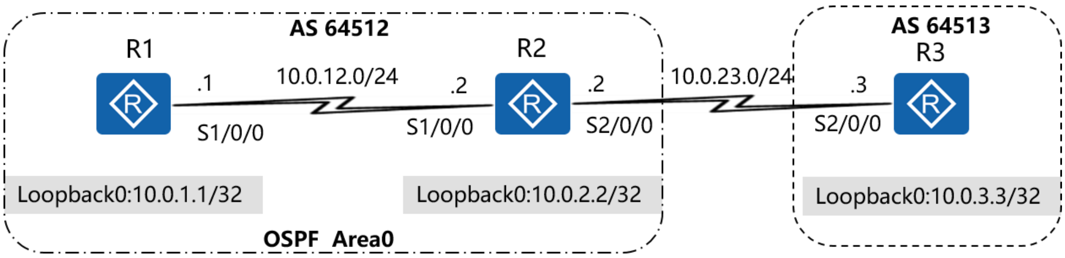

Figure 3-6 BGP troubleshooting

Scenario

You are a network administrator of a company. The company's network uses BGP as the routing protocol. This network consists of multiple ASs, and different branches use different AS numbers. You have finished building the company's network. During BGP configurations, you encountered many problems and have rectified all network failures.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

Configure IP addresses and masks for physical interfaces and loopback interfaces of all the routers. Each Loopback0 uses the 32-bit mask.

```
<R1>system-view
```

```
Enter system view, return user view with Ctrl+Z.
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ip add 10.0.1.1 32
[R1-LoopBack0]quit

[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]quit
[R2]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 24
[R2-Serial2/0/0]quit
[R2]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 32
[R2-LoopBack0]quit

[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 24
[R3-Serial2/0/0]quit
[R3]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 32
[R3-LoopBack0]quit
```

After the configurations are complete, test direct link connectivity.

```
[R2]ping -c 1 10.0.12.1
  PING 10.0.12.1: 56  data bytes, press CTRL_C to break
```

```
    Reply from 10.0.12.1: bytes=56 Sequence=1 ttl=255 time=40 ms

  --- 10.0.12.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 40/40/40 ms

[R2]ping -c 1 10.0.23.3
  PING 10.0.23.3: 56   data bytes, press CTRL_C to break
    Reply from 10.0.23.3: bytes=56 Sequence=1 ttl=255 time=38 ms

  --- 10.0.23.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 38/38/38 ms
```

The preceding command output shows that direct link connectivity is normal.

Step 2 **Configure IGP and BGP.**

Configure OSPF in AS 64512 and configure all devices to belong to Area 0. Configure each router to use Loopback0 address as its router ID. Run OSPF on the network segments connected to S1/0/0 and Loopback0 of R1.

```
[R1]router-id 10.0.1.1
[R1]ospf 1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.1.1 0.0.0.0
```

```
[R1-ospf-1-area-0.0.0.0]quit
[R1-ospf-1]quit
```

Run OSPF on the network segments connected to S1/0/0 and Loopback0 of R2.

```
[R2]router id 10.0.2.2
[R2]ospf 1
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.12.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]quit
[R2-ospf-1]quit
```

After the configurations are complete, check connectivity between Loopback0 addresses of R1 and R2.

```
[R1]ping -c 1 -a 10.0.1.1 10.0.2.2
  PING 10.0.2.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.2.2: bytes=56 Sequence=1 ttl=255 time=40 ms

  --- 10.0.2.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 40/40/40 ms
```

Configure IBGP between R1 and R2, configure EBGP between R2 and R3, and configure these routers to establish BGP peer relationships using loopback interface addresses. To ensure normal transmission of routing information, configure

next-hop-local on R2 and specify R1’s address as the peer address and incorrectly set the AS number of the peer 10.0.2.2 to 64514 on R3.

```
[R1]bgp 64512
[R1-bgp]peer 10.0.2.2 as-number 64512
[R1-bgp]quit

[R2]bgp 64512
[R2-bgp]peer 10.0.1.1 as-number 64512
[R2-bgp]peer 10.0.1.1 next-hop-local
[R2-bgp]peer 10.0.3.3 as-number 64513
[R2-bgp]quit

[R3]router id 10.0.3.3
[R3]bgp 64513
[R3-bgp]peer 10.0.2.2 as-number 64514
[R3-bgp]quit
```

Step 3 **Troubleshoot the failure to establish BGP peer relationships.**

After the configurations are complete, you can see that BGP peer relationships between routers are not established. Check the BGP peer relationships of R2 first.

```
[R2]display bgp peer

 BGP local router ID : 10.0.2.2
 Local AS number : 64512
 Total number of peers : 2                 Peers in established state : 0

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv
```

```
10.0.1.1        4       64512       0       0    0  00:05:36     Active      0
10.0.3.3        4       64513       0       0    0  00:05:21       Idle      0
```

The preceding command output shows that the **State** field of 10.0.1.1 displays **Active** and the **State** field of 10.0.2.2 displays **idle**. If a BGP peer relationship is established normally, the **State** field displays **Established**. If the **State** field remains another state for a long period, a failure occurs and needs to be rectified.

Generally, when a peer IP address is unreachable for a local router, the peer status displays Idle. That is, this router does not initiate a TCP connection with the peer. When the peer IP address is reachable but an error occurs during the establishment of a TCP connection, you can see that the peer status remains Active.

First check the BGP peer relationship between R2 and R3 and check connectivity between loopback interface addresses of R2 and R3.

```
[R2]ping -c 1 -a 10.0.2.2 10.0.3.3
  PING 10.0.3.3: 56   data bytes, press CTRL_C to break
    Request time out

  --- 10.0.3.3 ping statistics ---
    1 packet(s) transmitted
    0 packet(s) received
100.00% packet loss
```

The preceding command output shows that connectivity between loopback interface addresses of R2 and R3 is abnormal.

Check the IP routing table of R2.

```
[R2]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
```

```
Routing Tables: Public
         Destinations : 14        Routes : 14

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.1/32  OSPF    10   1562        D   10.0.12.1       Serial1/0/0
       10.0.2.2/32  Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.12.0/24  Direct  0    0           D   10.0.12.2       Serial1/0/0
      10.0.12.1/32  Direct  0    0           D   10.0.12.1       Serial1/0/0
      10.0.12.2/32  Direct  0    0           D   127.0.0.1       InLoopBack0
    10.0.12.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.23.0/24  Direct  0    0           D   10.0.23.2       Serial2/0/0
      10.0.23.2/32  Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.23.3/32  Direct  0    0           D   10.0.23.3       Serial2/0/0
    10.0.23.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32Direct    0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct    0    0           D   127.0.0.1       InLoopBack0
```

The preceding command output shows that there is no route to 10.0.3.3 of R3.

Check the IP routing table of R3.

```
[R3]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 9         Routes : 9
```

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

       10.0.3.3/32  Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.23.0/24  Direct  0    0           D   10.0.23.3       Serial2/0/0
      10.0.23.2/32  Direct  0    0           D   10.0.23.2       Serial2/0/0
      10.0.23.3/32  Direct  0    0           D   127.0.0.1       InLoopBack0
    10.0.23.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct  0    0           D   127.0.0.1       InLoopBack0
```

The preceding command output shows that there is no route to 10.0.2.2 of R2.

For different ASs, using static routes can ensure connectivity between loopback interface addresses of neighboring routers.

On R2 and R3, you need to configure static routes to the network segments connected to the loopback interfaces of R3 and R2 respectively.

```
[R2]ip route-static 10.0.3.3 32 10.0.23.3


[R3]ip route-static 10.0.2.2 32 10.0.23.2
```

Check connectivity between R2 and R3.

```
[R2]ping -c 1 -a 10.0.2.2 10.0.3.3
  PING 10.0.3.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.3.3: bytes=56 Sequence=1 ttl=255 time=30 ms


  --- 10.0.3.3 ping statistics ---
```

```
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 30/30/30 ms
```

Check BGP peer relationships of R2.

```
[R2]display bgp peer

 BGP local router ID : 10.0.12.2
 Local AS number : 64512
 Total number of peers : 2                Peers in established state : 0

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.0.1.1        4       64512        0        0     0 05:23:27     Active        0
  10.0.3.3        4       64513        0        0     0 05:23:02     Active        0
```

The BGP peer relationship between R2 and R3 changes from Idle to Active state.

Check the BGP peer relationship between R1 and R2. After OSPF is configured, connectivity between loopback interface addresses of R1 and R2 has been tested.

BGP uses TCP port 179 for communication. Check whether port 179 is enabled on the routers.

On R1 and R2, check the TCP status.

```
[R1]display tcp status
TCPCB     Tid/Soid  Local Add:port       Foreign Add:port      VPNID   State
194b9500  8   /2     0.0.0.0:22            0.0.0.0:0            23553    Listening
194b939c  8   /1     0.0.0.0:23            0.0.0.0:0            23553    Listening
```

```
194b90d4   106/1     0.0.0.0:80      0.0.0.0:0       0       Listening
194b9a90   234/2     0.0.0.0:179     10.0.2.2:0      0       Listening
194b9664   8   /3    0.0.0.0:830     0.0.0.0:0       23553   Listening
194b9238   6   /1    0.0.0.0:7547    0.0.0.0:0       0       Listening

[R2]display tcp status
TCPCB      Tid/Soid  Local Add:port  Foreign Add:port  VPNID  State
1949a048   234/5     0.0.0.0:0       0.0.0.0:0       0       Closed
19499d80   8   /2    0.0.0.0:22      0.0.0.0:0       23553   Listening
19499c1c   8   /1    0.0.0.0:23      0.0.0.0:0       23553   Listening
19499954   106/1     0.0.0.0:80      0.0.0.0:0       0       Listening
1949a474   234/2     0.0.0.0:179     10.0.1.1:0      0       Listening
1949a310   234/4     0.0.0.0:179     10.0.3.3:0      0       Listening
19499ee4   8   /3    0.0.0.0:830     0.0.0.0:0       23553   Listening
19499ab8   6   /1    0.0.0.0:7547    0.0.0.0:0       0       Listening
```

The preceding command output shows that port 179 for corresponding peer address is in Listening state. BGP works normally on a single router.

Run the debugging command on R1 to check whether R1 receives BGP packets sent from R2.

```
<R1>terminal monitor
<R1>terminal debugging
<R1>debugging tcp packet
Dec   7 2011 10:08:16.620.1+00:00 R1 SOCKET/7/TCP PACKET:
TCP debug packet information:
1323252496: Input: no port,
(src = 10.0.12.2:52688,dst = 10.0.1.1:179,VrfIndex = 0,seq = 2254758724,
ack = 0,datalen = 0,optlen = 4,flag =   SYN ,window = 16384,ttl = 0,tos = 0,MSS = 0)
```

```
Dec  7 2011 10:08:16.620.2+00:00 R1 SOCKET/7/TCP PACKET:
TCP debug packet information:
1323252496: Output: task = (0), socketid = 0,
(src = 10.0.1.1:179,dst = 10.0.12.2:52688,VrfIndex = 0,seq = 0,
ack = 2254758725,datalen = 0,optlen = 0,flag =  ACK  RST ,window = 0,ttl = 255,tos = 0,MSS = 0)

<R1>undo debugging all
Info: All possible debugging has been turned off
```

The preceding command output shows that the source address of the packet with the destination port number 179 is 10.0.12.2. After checking the topology, you can see that 10.0.12.2 is the address of R2's Serial1/0/0.

When establishing BGP peer relationships, you use the loopback interface address of R2. As a result, the BGP peer relationship between R1 and R2 cannot be established. Therefore, you need to use **connect-interface** to specify the source address during establishment of BGP peer relationships.

Similarly, this problem also exists between R2 and R3. Therefore, you need to use **connect-interface** to specify the source address during establishment of BGP peer relationships.

```
[R1]bgp 64512
[R1-bgp]peer 10.0.2.2 connect-interface LoopBack 0
[R1-bgp]quit

[R2]bgp 64512
[R2-bgp]peer 10.0.1.1 connect-interface LoopBack 0
[R2-bgp]peer 10.0.3.3 connect-interface LoopBack 0
[R2-bgp]quit
```

[R3]bgp 64513

[R3-bgp]peer 10.0.2.2 connect-interface LoopBack 0

[R3-bgp]quit

After the modifications are complete, check BGP peer relationships of R2 again.

[R2]display bgp peer

BGP local router ID : 10.0.2.2

Local AS number : 64512

Total number of peers : 2 Peers in established state : 1

Peer	V	AS	MsgRcvd	MsgSent	OutQ	Up/Down	State	PrefRcv
10.0.1.1	4	64512	16	17	0	00:14:18	Established	0
10.0.3.3	4	64513	0	0	0	00:14:35	Active	0

The preceding command output shows that the BGP peer relationship between R1 and R2 is in Established state.

Run the debugging command on R3 to check whether R3 receives any BGP packet and check the content of the packet.

<R3>terminal monitor

<R3>terminal debugging

<R3>debugging ip packet

Dec 7 2011 10:51:44.30.5+00:00 R3 IP/7/debug_case:

Delivering, interface = S2/0/0, version = 4, headlen = 20, tos = 192,

pktlen = 40, pktid = 4752, offset = 0, ttl = 1, protocol = 6,

checksum = 36220, s = 10.0.2.2, d = 10.0.3.3

prompt: Packet is before IP_Reass before really deliver to up.

```
Dec  7 2011 10:51:44.30.6+00:00 R3 IP/7/debug_case:
Sending, interface = S2/0/0, version = 4, headlen = 20, tos = 0,
pktlen = 40, pktid = 9953, offset = 0, ttl = 255, protocol = 6,
checksum = 31722, s = 10.0.3.3, d = 10.0.2.2
prompt: Sending the packet from local at S2/0/0

<R3>undo debugging all
Info: All possible debugging has been turned off
```

The preceding command output shows that the TTL of the BGP packet received by R3 is 1. For EBGP, the default TTL of the packet sent from a router to its peer is 1.

In this scenario, R2 and R3 establish a BGP peer relationship using loopback interface addresses. There are two hops from the loopback interface address of R2 to that of R3. Therefore, before this BGP packet reaches the loopback interface address of R2, it is discarded because of TTL expiry.

To address this problem, change the TTL of the packet sent between two EBGP peers.

```
[R2]bgp 64512
[R2-bgp]peer 10.0.3.3 ebgp-max-hop 2
[R2-bgp]quit

[R3]bgp 64513
[R3-bgp]peer 10.0.2.2 ebgp-max-hop 2
[R3-bgp]quit
```

After the configurations are complete, check the BGP peer relationship of R2 again.

```
[R2]display bgp peer
```

```
BGP local router ID : 10.0.2.2
Local AS number : 64512
Total number of peers : 2          Peers in established state : 1

 Peer        V      AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

 10.0.1.1    4   64512        3        4     0 00:01:34  Established       0
 10.0.3.3    4   64513        0        1     0 00:00:44       Active       0
```

The preceding command output shows that the BGP peer relationship between R2 and R3 remains Active.

Check BGP errors on R3.

```
[R3]display bgp error
 Error Type      : Peer Error
 Date/Time       : 2011/12/07 11:24:37
 Peer Address    : 10.0.2.2
 VRF Name        : Public
 Error Info      : Incorrect remote AS

 Error Type      : Peer Error
 Date/Time       : 2011/12/07 11:25:09
 Peer Address    : 10.0.2.2
 VRF Name        : Public
 Error Info      : Incorrect remote AS

 Error Type      : Peer Error
 Date/Time       : 2011/12/07 11:25:41
 Peer Address    : 10.0.2.2
```

```
 VRF Name         : Public
 Error Info       : Incorrect remote AS

<R3>terminal debugging
<R3>debugging bgp packet verbose
Dec   7 2011 11:31:01.540.1+00:00 R3 RM/6/RMDEBUG:
 BGP.Public: Err/SubErr: 2/2 Errdata: 41040000fc00
 Identified in OPEN MSG from 10.0.2.2.

Dec   7 2011 11:31:01.540.2+00:00 R3 RM/6/RMDEBUG:

Dec   7 2011 11:31:01.540.3+00:00 R3 RM/6/RMDEBUG:
 BGP.Public: Err/SubErr: 2/2 Errdata: 41040000fc00
 Identified in OPEN MSG from 10.0.2.2.
```

The preceding command output shows an incorrect AS number message.

Run the debugging command to troubleshoot this error.

The preceding command output shows that the error code/suberror code is 2. This error indicates incorrect AS number. You need to change the peer AS number on R3.

```
[R3]bgp 64513
[R3-bgp]undo peer 10.0.2.2
[R3-bgp]peer 10.0.2.2 as-number 64512
[R3-bgp]peer 10.0.2.2 ebgp-max-hop 2
[R3-bgp]peer 10.0.2.2 connect-interface LoopBack0
[R3-bgp]quit
```

Check the BGP peer relationship between R2 and R3.

```
[R2]display bgp peer

 BGP local router ID : 10.0.2.2
 Local AS number : 64512
 Total number of peers : 2                 Peers in established state : 2

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.0.1.1        4       64512       81       82     0 01:19:18 Established       0
  10.0.3.3        4       64513        3        4     0 00:01:12 Established       0
```

Step 4 **Configure BGP security.**

BGP is often used in a backbone network, so BGP security is important. If a BGP router is attacked, large-scale network breakdown occurs.

To prevent malicious users from forging valid routers to establish BGP peer relationships with BGP routers, configure MD5 authentication between BGP peers.

Enable MD5 authentication between R1 and R2. First, configure an incorrect password. Configure the password to huawei on R1 and that to 123 on R2 and then check the BGP peer relationship changes.

```
[R1]bgp 64512
[R1-bgp]peer 10.0.2.2 password simple huawei
[R1-bgp]quit

[R2]bgp 64512
[R2-bgp]peer 10.0.1.1 password simple 123
[R2-bgp]quit
```

Reset the BGP peer relationship of R1. The following command output shows that the peer relationship between R1 and R2 remains Connect and Active and cannot enter the Established state.

```
<R1>reset bgp 10.0.2.2

[R1]display bgp peer

 BGP local router ID : 10.0.1.1
 Local AS number : 64512
 Total number of peers : 1                 Peers in established state : 0

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.0.2.2        4       64512        0        0     0 00:03:39     Connect       0
```

Change the password of R2 to huawei.

```
[R2]bgp 64512
[R2-bgp]undo peer 10.0.1.1 password
[R2-bgp]peer 10.0.1.1 password simple huawei
[R2-bgp]quit
```

Wait for about 30 seconds, and then check the peer relationship again.

```
 [R2]display bgp peer

 BGP local router ID : 10.0.2.2
 Local AS number : 64512
 Total number of peers : 2                 Peers in established state : 2
```

Peer	V	AS	MsgRcvd	MsgSent	OutQ	Up/Down	State	PrefRcv
10.0.1.1	4	64512	2	2	0	00:00:34	Established	0
10.0.3.3	4	64513	166	167	0	02:44:05	Established	0

The preceding command output shows that the peer relationship between R1 and R2 has reached the Established state.

In this scenario, the administrator of AS 64512 does not want the routers in AS 64513 to view its actual AS number.

The **fake-as** parameter can achieve this purpose. It can specify a fake AS number for the peer.

Configure this command on R2, specify R3's address as the peer address, and set the fake AS number to 100.

On R3, change the AS number of the BGP peer R2.

```
[R2]bgp 64512
[R2-bgp]peer 10.0.3.3 fake-as 100
[R2-bgp]quit

[R3]bgp 64513
[R3-bgp]undo peer 10.0.2.2
[R3-bgp]peer 10.0.2.2 as-number 100
[R3-bgp]peer 10.0.2.2 ebgp-max-hop 2
[R3-bgp]peer 10.0.2.2 connect-interface LoopBack0
[R3-bgp]quit
```

Check the BGP peer of R3. The following command output shows that the AS number of R2 changes to 100.

```
[R3]display bgp peer

 BGP local router ID : 10.0.3.3
 Local AS number : 64513
 Total number of peers : 1                 Peers in established state : 1

  Peer          V        AS  MsgRcvd  MsgSent  OutQ  Up/Down       State  PrefRcv

  10.0.2.2      4       100        2        2     0 00:00:28   Established      0
```

On R2, advertise the network segment where its Loopback0 resides and observe the AS_Path attribute of the BGP route learned from R3.

```
[R2]bgp 64512
[R2-bgp]network 10.0.2.2 32
[R2-bgp]quit

[R3]display bgp routing-table

 BGP Local router ID is 10.0.3.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 1
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

     10.0.2.2/32        10.0.2.2        0                     0      100i
```

The preceding command output shows that the AS_Path attribute of the BGP route 10.0.2.2/32 learned from R3 is 100. That is, R3 considers that this route is originated from AS 100. In this situation, AS 64512 is not displayed.

BGP also provides a security mechanism: Generalized TTL Security Mechanism (GTSM). GTSM protects routers by checking whether the TTL value in the IP header is within a specified range. That is, if the TTL value of the received BGP packet exceeds the specified range, this packet is discarded. Both the GTSM and ebgp-max-hop functions will affect the TTL value of sent BGP packets and the two functions are mutually exclusive. You can only enable one of the GTSM and ebgp-max-hop functions on one peer or peer group.

In this scenario, enable GTSM on the link between R2 and R3 to observe BGP packet exchange. First, in the system view of R2, configure the default action to be taken on BGP packets whose TTL value is not within the specified range. Here, you configure the default action to **drop**. That is, the BGP packets whose TTL value is not within the specified range will be discarded.

```
[R2]gtsm default-action drop
```

In the BGP view of R2, specify R3's address as the peer address and enable GTSM. Before performing this configuration, you need to delete the **ebgp-max-hop** configuration. Because R2 and R3 are directly connected, the **valid-ttl-hops** parameter is 1.

```
[R2]bgp 64512
[R2-bgp]undo peer 10.0.3.3 ebgp-max-hop
[R2-bgp]peer 10.0.3.3 valid-ttl-hops 1
[R2-bgp]peer 10.0.1.1 valid-ttl-hops 1
[R2-bgp]quit
```

Perform the same operation on R1 and R3.

```
[R1]gtsm default-action drop
```

[R1]bgp 64512

[R1-bgp]peer 10.0.2.2 valid-ttl-hops 1

[R1-bgp]quit

[R3]gtsm default-action drop

[R3]bgp 64513

[R3-bgp]undo peer 10.0.2.2 ebgp-max-hop

[R3-bgp]peer 10.0.2.2 valid-ttl-hops 1

[R3-bgp]quit

Check the BGP peer relationship between R2 and R3.

[R3]dis bgp peer

BGP local router ID : 10.0.3.3

Local AS number : 64513

Total number of peers : 1 Peers in established state : 1

Peer	V	AS	MsgRcvd	MsgSent	OutQ	Up/Down	State	PrefRcv
10.0.2.2	4	100	3	2	0	00:00:06	Established	1

On R3, observe the TTL value of BGP packets.

<R3>terminal monitor

<R3>terminal debugging

<R3>debugging ip packet

Dec 7 2011 16:34:51.10.1+00:00 R3 IP/7/debug_case:

Receiving, interface = S2/0/0, version = 4, headlen = 20, tos = 192,

pktlen = 59, pktid = 8820, offset = 0, ttl = 255, protocol = 6,

checksum = 32644, s = 10.0.2.2, d = 10.0.3.3

prompt: Receiving IP packet from S2/0/0

Dec 7 2011 16:34:51.10.2+00:00 R3 IP/7/debug_case:

Receiving, interface = Serial2/0/0, version = 4, headlen = 20, tos = 192,

pktlen = 59, pktid = 8820, offset = 0, ttl = 255, protocol = 6,

checksum = 32644, s = 10.0.2.2, d = 10.0.3.3

prompt: IP_ProcessByBoard Begin!

<R3>undo debugging all

Info: All possible debugging has been turned off

The preceding command output shows that the TTL value of packets received by R3 from R2 is 255 instead of the default value 1. To confirm that GTSM discards the BGP packets whose TTL value is not within the specified range, enable the GTSM log function on R3. When BGP packets are discarded by GTSM, a log is recorded.

[R3]gtsm log drop-packet all

Run the **ebgp-max-hop** command on R2 to ensure that the TTL value of BGP packets sent from R2 to R3 is less than 254.

[R2]bgp 64512

[R2-bgp]undo peer 10.0.3.3 valid-ttl-hops

[R2-bgp]peer 10.0.3.3 ebgp-max-hop 253

[R2-bgp]quit

After waiting for a certain period, you can see that the BGP peer relationship between R2 and R3 is in Idle state. Check GTSM statistics on R3. The following command output shows that some BGP packets are discarded by GTSM.

Dec 7 2011 16:48:34+00:00 R3 %%01BGP/3/STATE_CHG_UPDOWN(l)[4]:The status of the peer 10.0.2.2 changed from ESTABLISHED to IDLE. (InstanceName=Public, StateChangeReason=Hold Timer Expired)

```
[R3]display gtsm statistics all
GTSM Statistics Table
----------------------------------------------------------------
SlotId  Protocol  Total Counters  Drop Counters  Pass Counters
----------------------------------------------------------------
 0      BGP        83              27              56
 0      OSPF       0               0               0
 0      LDP        0               0               0
----------------------------------------------------------------
```

Run the following commands on R2 to change the current configuration to the previous configuration. Wait for a certain period and then check whether BGP packets are discarded.

```
[R2]bgp 64512
[R2-bgp]undo peer 10.0.3.3 ebgp-max-hop
[R2-bgp]peer 10.0.3.3 valid-ttl-hops 1
[R2-bgp]quit
```

```
[R3]display gtsm statistics all
GTSM Statistics Table
----------------------------------------------------------------
SlotId  Protocol  Total Counters  Drop Counters  Pass Counters
----------------------------------------------------------------
 0      BGP        89              27              62
 0      OSPF       0               0               0
```

```
 0        LDP              0                     0                   0
--------------------------------------------------------------
```

The preceding command output shows that no more BGP packets are discarded.

----End

Additional Exercises: Analysis and Verification

What type of attacks is GTSM mainly used to defend against?

Can fake-as be used in a confederation?

Device Configurations

```
<R1>display current-configuration
[V200R007C00SPC600]
#
 sysname R1
#
router id 10.0.1.1
#
gtsm default-action drop
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
#
bgp 64512
```

```
 peer 10.0.2.2 as-number 64512
 peer 10.0.2.2 connect-interface LoopBack0
 peer 10.0.2.2 password simple huawei
 peer 10.0.2.2 valid-ttl-hops 1
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.2.2 enable
#
ospf 1 router-id 10.0.1.1
 area 0.0.0.0
  network 10.0.12.1 0.0.0.0
  network 10.0.1.1 0.0.0.0
#
return

<R2>display current-configuration
[V200R007C00SPC600]
#
 sysname R2
#
router id 10.0.2.2
#
gtsm default-action drop
#
acl number 2001
 rule 5 permit source 10.0.2.2 0
#
```

```
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.255
#
bgp 64512
 peer 10.0.1.1 as-number 64512
 peer 10.0.1.1 connect-interface LoopBack0
 peer 10.0.1.1 password simple huawei
 peer 10.0.1.1 valid-ttl-hops 1
 peer 10.0.3.3 as-number 64513
 peer 10.0.3.3 connect-interface LoopBack0
 peer 10.0.3.3 fake-as 100
 peer 10.0.3.3 valid-ttl-hops 1
 #
 ipv4-family unicast
  undo synchronization
  network 10.0.2.2 255.255.255.255
  peer 10.0.1.1 enable
  peer 10.0.1.1 next-hop-local
  peer 10.0.3.3 enable
#
ospf 1 router-id 10.0.2.2
```

```
 area 0.0.0.0
  network 10.0.12.2 0.0.0.0
  network 10.0.2.2 0.0.0.0
#
route-policy change_origin deny node 10
 if-match acl 2001
apply origin egp 100
#
 ip route-static 10.0.3.3 255.255.255.255 10.0.23.3
#
return

<R3>display current-configuration
[V200R007C00SPC600]
#
 sysname R3
#
router id 10.0.3.3
#
gtsm default-action drop
 gtsm log drop-packet all
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
```

```
#
bgp 64513
 peer 10.0.2.2 as-number 100
 peer 10.0.2.2 connect-interface LoopBack0
 peer 10.0.2.2 valid-ttl-hops 1
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.2.2 enable
#
 ip route-static 10.0.2.2 255.255.255.255 10.0.23.2
#
return
```

Chapter 4 Multicast Protocols

Lab 4-1 Multicast, IGMP, and PIM DM Protocols

Learning Objectives

The objectives of this lab are to learn and understand:

- How to enable multicast routing on routers
- How to configure IGMP on interfaces
- How to configure PIM-DM
- How to check and test multicast
- How to configure advanced PIM functions

Topology

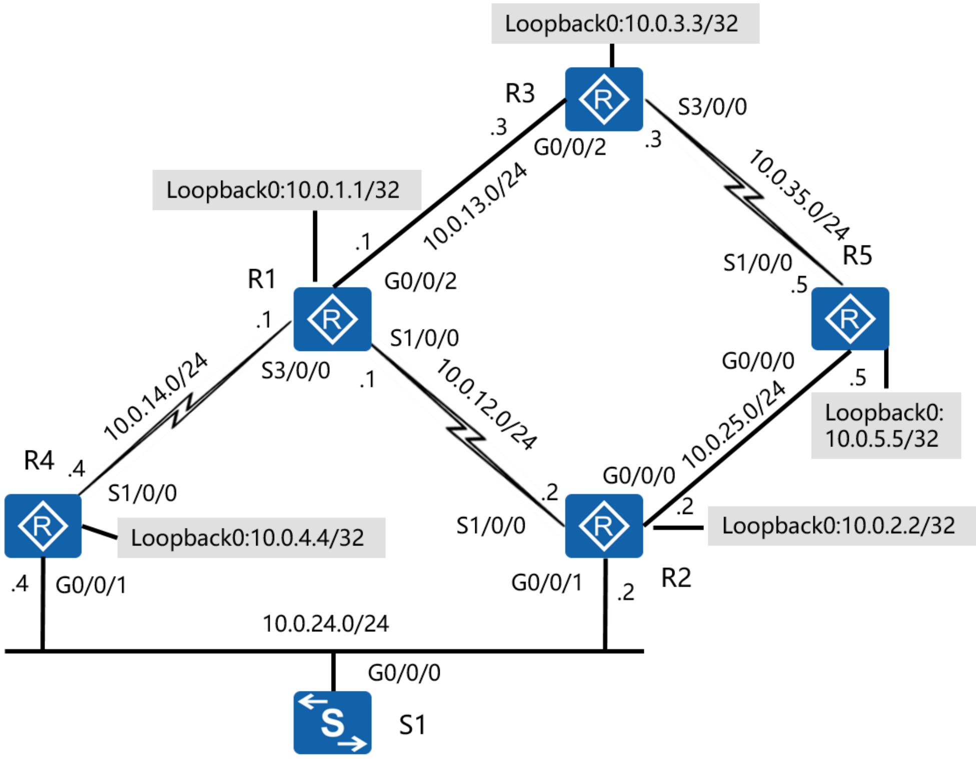

Figure 4-1 Multicast, IGMP, and PIM DM

Scenario

You are a network administrator of a company. Multicast needs to be configured to forward some services. The network size is small, so you can configure PIM DM to implement multicast route learning. To implement multicast forwarding, you must consider whether the host applications are compatible with different IGMP versions and choose appropriate methods to test the multicast service. To improve network efficiency and security, you can use the PIM DM methods, such as PIM neighbor control and graft. You may encounter network failures and need to rectify the faults.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

S2 is located between R1 and R3, but does not need to be configured. Before the experiment, clear the configuration on S2 and restart S2.

Configure IP addresses and masks for all the routers. All loopback interfaces must have 24-bit masks.

```
<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R1
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]quit
[R1]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 24
[R1-Serial3/0/0]quit
[R1]interface GigabitEthernet 0/0/2
[R1-GigabitEthernet0/0/2]ip address 10.0.13.1 24
[R1-GigabitEthernet0/0/2]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 24
```

```
[R1-LoopBack0]quit

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R2
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]quit
[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.25.2 24
[R2-GigabitEthernet0/0/0]quit
[R2]interface GigabitEthernet 0/0/1
[R2-GigabitEthernet0/0/1]ip address 10.0.24.2 24
[R2-GigabitEthernet0/0/1]quit
[R2]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24
[R2-LoopBack0]quit

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R3
[R3]interface GigabitEthernet 0/0/2
[R3-GigabitEthernet0/0/2]ip address 10.0.13.3 24
[R3-GigabitEthernet0/0/2]quit
[R3]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 24
[R3-Serial3/0/0]quit
[R3]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 24
```

```
[R3-LoopBack0]quit

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R4
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 24
[R4-Serial1/0/0]quit
[R4]interface GigabitEthernet 0/0/1
[R4-GigabitEthernet0/0/1]ip address 10.0.24.4 24
[R4-GigabitEthernet0/0/1]
[R4]interface LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 24
[R4-LoopBack0]quit

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R5
[R5]interface Serial 1/0/0
[R5-Serial1/0/0]ip address 10.0.35.5 24
[R5-Serial1/0/0]quit
[R5]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.25.5 24
[R5-GigabitEthernet0/0/0]quit
[R5]interface LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 24
[R5-LoopBack0]quit
```

After the configurations are complete, test the connectivity between routers.

```
[R1]ping -c 1 10.0.13.3
  PING 10.0.13.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.13.3: bytes=56 Sequence=1 ttl=255 time=5 ms

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

[R1]ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=37 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 37/37/37 ms

[R1]ping -c 1 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=38 ms

  --- 10.0.14.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
```

```
round-trip min/avg/max = 38/38/38 ms

[R5]ping -c 1 10.0.35.3
  PING 10.0.35.3: 56 data bytes, press CTRL_C to break
    Reply from 10.0.35.3: bytes=56 Sequence=1 ttl=255 time=33 ms

  --- 10.0.35.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 33/33/33 ms

[R5]ping -c 1 10.0.25.2
  PING 10.0.25.2: 56 data bytes, press CTRL_C to break
    Reply from 10.0.25.2: bytes=56 Sequence=1 ttl=255 time=10 ms

  --- 10.0.25.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 10/10/10 ms
```

Step 2 **Enable multicast routing on all routers.**

Enable multicast routing on R1, R2, R3, R4, and R5. To enable multicast, run the **multicast routing-enable** command in the system view.

By default, the multicast function is disabled on VRP. Before using PIM or IGMP, enable multicast globally.

```
[R1]multicast routing-enable
```

To run PIM DM on an interface, run the **pim dm** command in the interface view.

[R1]interface GigabitEthernet 0/0/2

[R1-GigabitEthernet0/0/2]pim dm

[R1-GigabitEthernet0/0/2]quit

[R1]interface Serial 1/0/0

[R1-Serial1/0/0]pim dm

[R1-Serial1/0/0]quit

[R1]interface Serial 3/0/0

[R1-Serial3/0/0]pim dm

[R1-Serial3/0/0]quit

Perform the same configurations on R2, R3, R4, and R5. Enable PIM DM on the interfaces between routers.

[R2]multicast routing-enable

[R2]interface Serial 1/0/0

[R2-Serial1/0/0]pim dm

[R2-Serial1/0/0]quit

[R2]interface GigabitEthernet 0/0/0

[R2-GigabitEthernet0/0/0]pim dm

[R2-GigabitEthernet0/0/0]quit

[R3]multicast routing-enable

[R3]interface GigabitEthernet 0/0/2

[R3-GigabitEthernet0/0/2]pim dm

[R3-GigabitEthernet0/0/2]quit

[R3]interface Serial 3/0/0

[R3-Serial3/0/0]pim dm

[R3-Serial3/0/0]quit

[R4]multicast routing-enable

[R4]interface Serial 1/0/0

[R4-Serial1/0/0]pim dm

[R4-Serial1/0/0]quit

[R4]interface GigabitEthernet 0/0/1

[R4-GigabitEthernet0/0/1]pim dm

[R4-GigabitEthernet0/0/1]quit

[R5]multicast routing-enable

[R5]interface Serial 1/0/0

[R5-Serial1/0/0]pim dm

[R5-Serial1/0/0]quit

[R5]interface GigabitEthernet 0/0/0

[R5-GigabitEthernet0/0/0]pim dm

[R5-GigabitEthernet0/0/0]quit

After the configuration, check the PIM status on interfaces.

[R1]display pim interface

VPN-Instance: public net

Interface	State	NbrCnt	HelloInt	DR-Pri	DR-Address
GE0/0/2	up	1	30	1	10.0.13.3
S1/0/0	up	1	30	1	10.0.12.2
S3/0/0	up	1	30	1	10.0.14.4

R1 has three interfaces running PIM and each interface has a neighbor (NbrCnt). On a network segment, the router with a larger interface IP address becomes the DR.

Check detailed PIM information on R1's G0/0/2.

```
[R1]display pim interface GigabitEthernet 0/0/2 verbose
 VPN-Instance: public net
 Interface: GigabitEthernet0/0/2, 10.0.13.1
     PIM version: 2
     PIM mode: Dense
     PIM state: up
     PIM DR: 10.0.13.3
     PIM DR Priority (configured): 1
     PIM neighbor count: 1
     PIM hello interval: 30 s
     PIM LAN delay (negotiated): 500 ms
     PIM LAN delay (configured): 500 ms
     PIM hello override interval (negotiated): 2500 ms
     PIM hello override interval (configured): 2500 ms
     PIM Silent: disabled
     PIM neighbor tracking (negotiated): disabled
     PIM neighbor tracking (configured): disabled
     PIM generation ID: 0X5325911
     PIM require-GenID: disabled
     PIM hello hold interval: 105 s
     PIM assert hold interval: 180 s
     PIM triggered hello delay: 5 s
     PIM J/P interval: 60 s
     PIM J/P hold interval: 210 s
     PIM state-refresh processing: enabled
     PIM state-refresh interval: 60 s
     PIM graft retry interval: 3 s
     PIM state-refresh capability on link: capable
```

```
PIM dr-switch-delay timer : not configured
Number of routers on link not using DR priority: 0
Number of routers on link not using LAN delay: 0
Number of routers on link not using neighbor tracking: 2
ACL of PIM neighbor policy: -
ACL of PIM ASM join policy: -
ACL of PIM SSM join policy: -
ACL of PIM join policy: -
```

By default, the hello interval of PIM DM is 30s, the hello hold time is 3.5 times of the hello interval (105s).

Check the neighbor list of R1. Three routers established PIM neighbor relationships with R1. The default DR priority of neighbors is 1.

```
[R1]display pim neighbor
 VPN-Instance: public net
 Total Number of Neighbors = 3

 Neighbor        Interface          Uptime   Expires  Dr-Priority
 10.0.13.3       GE0/0/2            01:40:27 00:01:18       1
 10.0.12.2       S1/0/0             01:42:21 00:01:24       1
 10.0.14.4       S3/0/0             01:38:02 00:01:16       1
```

Check details about neighbor R3. **Uptime** indicates the neighbor relationship setup time, **Expiry time** indicates the remaining time of the PIM neighbor, **LAN delay** indicates the delay in transmitting the prune messages, and **Override interval** indicates the interval for overriding the prune messages.

```
[R1]display pim neighbor 10.0.13.3 verbose
 VPN-Instance: public net
```

Neighbor: 10.0.13.3

Interface: GigabitEthernet0/0/2

Uptime: 01:41:00

Expiry time: 00:01:45

DR Priority: 1

Generation ID: 0XD1A5CA9

Holdtime: 105 s

LAN delay: 500 ms

Override interval: 2500 ms

State refresh interval: 60 s

Neighbor tracking: Disabled

Step 3 **Configure IGMP.**

In this experiment, multicast users are connected to S1. Enable IGMP on G0/0/1 of R2 and R4. To enable IGMP, run the **igmp enable** command in the interface view.

```
[R2]interface GigabitEthernet 0/0/1
[R2-GigabitEthernet0/0/1]igmp enable
[R2-GigabitEthernet0/0/1]quit

[R4]interface GigabitEthernet 0/0/1
[R4-GigabitEthernet0/0/1]igmp enable
[R4-GigabitEthernet0/0/1]quit
```

Add static multicast groups to G0/0/1 of R2 and R4. Then the interfaces always forward multicast traffic with destination address 225.1.1.1.

```
[R2]interface GigabitEthernet 0/0/1
[R2-GigabitEthernet0/0/1]igmp static-group 225.1.1.1
[R2-GigabitEthernet0/0/0]quit
```

[R4]interface GigabitEthernet 0/0/1

[R4-GigabitEthernet0/0/1]igmp static-group 225.1.1.1

[R4-GigabitEthernet0/0/1]quit

By default, VRP uses IGMPv2. The command output shows that R2 (10.0.24.2) is the querier of the network segment where G0/0/1 is located. IGMPv2 selects the router with a smaller IP address as the querier.

```
[R2]display igmp interface GigabitEthernet 0/0/1
Interface information
 GigabitEthernet0/0/1(10.0.24.2):
   IGMP is enabled
   Current IGMP version is 2
   IGMP state: up
   IGMP group policy: none
   IGMP limit: -
   Value of query interval for IGMP (negotiated): -
   Value of query interval for IGMP (configured): 60 s
   Value of other querier timeout for IGMP: 0 s
   Value of maximum query response time for IGMP: 10 s
   Querier for IGMP: 10.0.24.2 (this router)

[R4]display igmp interface GigabitEthernet 0/0/1
Interface information
 GigabitEthernet0/0/1(10.0.24.4):
   IGMP is enabled
   Current IGMP version is 2
   IGMP state: up
```

IGMP group policy: none

IGMP limit: -

Value of query interval for IGMP (negotiated): -

Value of query interval for IGMP (configured): 60 s

Value of other querier timeout for IGMP: 123 s

Value of maximum query response time for IGMP: 10 s

Querier for IGMP: 10.0.24.2

Check static IGMP groups on the interfaces. Group 225.1.1.1 is the manually added multicast group.

```
[R2]display igmp group static
Static join group information
 Total 1 entry, Total 1 active entry
  Group Address    Source Address   Interface          State    Expires
  225.1.1.1          0.0.0.0          GE0/0/1            UP       never
```

Check IGMP routing table on the interface.

```
[R2]display igmp routing-table
Routing table
 Total 1 entry

 00001. (*, 225.1.1.1)
       List of 1 downstream interface
         GigabitEthernet0/0/1 (10.0.24.2),
                    Protocol: STATIC
```

IGMP routing entries are generated on an interface only when the interface has IGMP but not PIM enabled, and the interface is an IGMP querier. The routing entries are not displayed on R4 because R2 is the querier of network segment 10.0.24.0/24.

By default, the query interval of the querier is 60s. To increase the speed of user addition to multicast groups, run the **igmp timer query** command to shorten the interval for sending query packets.

```
[R2]interface GigabitEthernet 0/0/1
[R2-GigabitEthernet0/0/1]igmp timer query 20
[R2-GigabitEthernet0/0/1]quit
```

Verify the configuration.

```
[R2]display igmp interface GigabitEthernet 0/0/1
Interface information of VPN-Instance: public net
 GigabitEthernet0/0/1(10.0.24.2):
   IGMP is enabled
   Current IGMP version is 2
   IGMP state: up
   IGMP group policy: none
   IGMP limit: -
   Value of query interval for IGMP (negotiated): -
   Value of query interval for IGMP (configured): 20 s
   Value of other querier timeout for IGMP: 0 s
   Value of maximum query response time for IGMP: 10 s
   Querier for IGMP: 10.0.24.2 (this router)
```

The debugging information shows that the interface sends a General-Query message every 20s.

```
<R1>terminal monitor
```

<R2>terminal debugging

<R2>

Sep 14 2016 14:19:53.740.1+00:00 R2 IGMP/7/QUERY:(public net): Send version 2 general query on

GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073969)

<R2>

Sep 14 2016 14:20:13.830.1+00:00 R2 IGMP/7/QUERY:(public net): Send version 2 general query on

GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073969)

<R2>

Sep 14 2016 14:20:33.770.1+00:00 R2 IGMP/7/QUERY:(public net): Send version 2 general query on

GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073969)

<R2>

Sep 14 2016 14:20:53.760.1+00:00 R2 IGMP/7/QUERY:(public net): Send version 2 general query on

GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073969)

<R2>undo debugging all

Info: All possible debugging has been turned off

The default robustness variable of a router is 2. Shut down the interface to test robustness.

Observe the IGMP query message sending in default settings.

<R2>terminal monitor

<R2>terminal debugging

<R2>debugging igmp query send

Sep 14 2016 14:26:13.880.1+00:00 R2 IGMP/7/QUERY:(public net): Send version 2 general query on

GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073969)

<R2>

Sep 14 2016 14:26:33.890.1+00:00 R2 IGMP/7/QUERY:(public net): Send version 2 general query on

GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073969)

```
<R2>system-view
[R2]interface GigabitEthernet 0/0/1
[R2-GigabitEthernet0/0/1]shutdown
[R2-GigabitEthernet0/0/1]undo shutdown
Sep 14 2016 14:26:51.810.1+00:00 R2 IGMP/7/QUERY:(public net): Send version 2 general query on
GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073969)
[R2-GigabitEthernet0/0/1]
Sep 14 2016 14:26:56.790.1+00:00 R2 IGMP/7/QUERY:(public net): Send version 2 general query on
GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073969)
[R2-GigabitEthernet0/0/1]
Sep 14 2016 14:27:16.790.1+00:00 R2 IGMP/7/QUERY:(public net): Send version 2 general query on
GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073969)
[R2-GigabitEthernet0/0/1]
Sep 14 2016 14:27:36.770.1+00:00 R2 IGMP/7/QUERY:(public net): Send version 2 general query on
GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073969)
```

When the interface is not shut down, the router's interface sends a General Query message every 20s. when the interface is shut down and enabled, the interval for sending the first two query messages is 5s. When a router starts, it sends *N* General Query messages (*N* is the robustness variable). The message sending interval is 1/4 of the configured interval for sending General Query messages.

Run the **robust-count** command to set the IGMP robustness variable. This parameter is valid only for IGMPv2 and IGMPv3. Change the robustness variable on R2's G0/0/1 to 3.

```
[R2-GigabitEthernet0/0/1]igmp robust-count 3
```

Enable debugging and observe General Query message sending.

```
[R2-GigabitEthernet0/0/1]shutdown
[R2-GigabitEthernet0/0/1]undo shutdown
```

```
Sep 14 2016 14:33:07.420.1+00:00 R2 IGMP/7/QUERY:(public net): Send version 2 general query on
GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073969)
[R2-GigabitEthernet0/0/1]
Sep 14 2016 14:33:12.340.1+00:00 R2 IGMP/7/QUERY:(public net): Send version 2 general query on
GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073969)
[R2-GigabitEthernet0/0/1]
Sep 14 2016 14:33:17.340.1+00:00 R2 IGMP/7/QUERY:(public net): Send version 2 general query on
GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073969)
[R2-GigabitEthernet0/0/1]
Sep 14 2016 14:33:37.420.1+00:00 R2 IGMP/7/QUERY:(public net): Send version 2 general query on
GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073969)

[R2-GigabitEthernet0/0/1]return
<R2>undo debugging all
Info: All possible debugging has been turned off
```

After the robustness variable is changed to 3, the interval for sending the first three General Query messages is 5s, and the interval for sending later messages is 20s.

Step 4 **Observe the multicast routing table.**

To observe multicast routing, configure OSPF as the unicast routing protocol.

```
[R1]ospf 1 router-id 10.0.1.1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.0.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.14.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.13.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]quit
[R1-ospf-1]quit
```

```
[R2]ospf 1 router-id 10.0.2.2
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.25.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.12.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]quit
[R2-ospf-1]quit

[R3]ospf 1 router-id 10.0.3.3
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.13.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.35.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]quit
[R3-ospf-1]quit

[R4]ospf 1 router-id 10.0.4.4
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]network 10.0.4.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]network 10.0.14.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]quit
[R4-ospf-1]quit

[R5]ospf 1 router-id 10.0.5.5
[R5-ospf-1]area 0
[R5-ospf-1-area-0.0.0.0]network 10.0.5.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]network 10.0.25.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]network 10.0.35.5 0.0.0.0
```

```
[R5-ospf-1-area-0.0.0.0]quit
[R5-ospf-1]quit
```

After the configuration, check whether the routers can learn the loopback addresses of other routers.

```
[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 7        Routes : 8

OSPF routing table status : <Active>
         Destinations : 7        Routes : 8

Destination/Mask   Proto Pre  Cost      Flags NextHop         Interface

      10.0.1.1/32  OSPF   10   1562        D   10.0.12.1       Serial1/0/0
      10.0.3.3/32  OSPF   10   1563        D   10.0.12.1       Serial1/0/0
                   OSPF   10   1563        D   10.0.25.5       GigabitEthernet0/0/0
      10.0.4.4/32  OSPF   10   3124        D   10.0.12.1       Serial1/0/0
      10.0.5.5/32  OSPF   10   1           D   10.0.25.5       GigabitEthernet0/0/0
     10.0.13.0/24  OSPF   10   1563        D   10.0.12.1       Serial1/0/0
     10.0.14.0/24  OSPF   10   3124        D   10.0.12.1       Serial1/0/0
     10.0.35.0/24  OSPF   10   1563        D   10.0.25.5       GigabitEthernet0/0/0

OSPF routing table status : <Inactive>
         Destinations : 0        Routes : 0
```

To simulate multicast information transmission, configure R3's loopback interface address as the source address to send ping packets to destination address 225.1.1.1.

```
[R3]ping -a 10.0.3.3 -c 300 225.1.1.1
```

After several minutes, the multicast routing entries can be seen on all the other routers. View the multicast routing table on R2.

```
[R2]display pim routing-table
 VPN-Instance: public net
 Total 1(*, G) entry; 1 (S, G) entry

(*, 225.1.1.1)
     Protocol: pim-dm, Flag: WC EXT
     UpTime: 00:09:04
     Upstream interface: NULL
         Upstream neighbor: NULL
         RPF prime neighbor: NULL
     Downstream interface(s) information: None

 (10.0.3.3, 225.1.1.1)
     Protocol: pim-dm, Flag:
     UpTime: 00:00:52
     Upstream interface: GigabitEthernet0/0/0
         Upstream neighbor: 10.0.25.5
         RPF prime neighbor: 10.0.25.5
     Downstream interface(s) information: None
```

There are two entries:

The first entry (*, 225.1.1.1) is generated after static IGMP group is configured on an interface.

The second entry (10.0.3.3, 225.1.1.1) is generated after multicast traffic is spread on the router.

For R2, the upstream router of multicast traffic is 10.0.25.5.

After PIM is enabled, routers use the unicast routing table for RPF check. The command output shows that the RPF neighbor of multicast source 10.0.3.3 is 10.0.25.5.

```
[R2]display multicast rpf-info 10.0.3.3
 VPN-Instance: public net
 RPF information about source: 10.0.3.3
     RPF interface: GigabitEthernet0/0/0, RPF neighbor: 10.0.25.5
     Referenced route/mask: 10.0.3.3/32
     Referenced route type: unicast
     Route selection rule: preference-preferred
     Load splitting rule: disable
```

Step 5 **Adjust PIM DM parameters.**

If you do not want the traffic to be transmitted to the destination through a unicast route, run the **rpf-route-static** command to change the RPF path. In this experiment, change the RPF path from 10.0.25.5 to 10.0.12.1.

```
[R2]ip rpf-route-static 10.0.3.0 255.255.255.0 10.0.12.1
```

After the configuration, the RPF neighbor is changed to 10.0.12.1.

```
[R2]display multicast rpf-info 10.0.3.3
 VPN-Instance: public net
 RPF information about source: 10.0.3.3
```

RPF interface: Serial1/0/0, RPF neighbor: 10.0.12.1

Referenced route/mask: 10.0.3.0/24

Referenced route type: mstatic

Route selection rule: preference-preferred

Load splitting rule: disable

To observe the PIM prune and graft messages, delete and add static IGMP groups to simulate the user deletion and addition operations. Enable debugging on R2.

<R1>terminal monitor

<R2>terminal debugging

<R2>debugging pim join-prune

Delete static IGMP group 225.1.1.1 from R2.

<R2>system-view

[R2]interface GigabitEthernet 0/0/1

[R2-GigabitEthernet0/0/1]undo igmp static-group 225.1.1.1

Dec 31 2011 15:00:05.300.1+00:00 R2 PIM/7/JP:(public net): PIM ver 2 JP sending 10.0.12.2 -> 224.0.0.13 on Serial1/0/0 (P012689)

Dec 31 2011 15:00:05.300.2+00:00 R2 PIM/7/JP:(public net): Upstream 10.0.12.1, Groups 1, Holdtime 210 (P012693)

Dec 31 2011 15:00:05.300.3+00:00 R2 PIM/7/JP:(public net): Group: 225.1.1.1/32 --- 0 joins 1 prunes (P012701)

Dec 31 2011 15:00:05.310.1+00:00 R2 PIM/7/JP:(public net): Prune: 10.0.3.3/32 (P012707)

Dec 31 2011 15:00:05.350.1+00:00 R2 PIM/7/JP:(public net): PIM ver 2 JP receiving 10.0.12.1 -> 224.0.0.13 on Serial1/0/0 (P012689)

Dec 31 2011 15:00:05.350.2+00:00 R2 PIM/7/JP:(public net): Upstream 10.0.12.1, Groups 1, Holdtime 207 (P012693)

Dec 31 2011 15:00:05.350.3+00:00 R2 PIM/7/JP:(public net): Group: 225.1.1.1/32 --- 0 joins 1 prunes (P012701)

Dec 31 2011 15:00:05.350.4+00:00 R2 PIM/7/JP:(public net): Prune: 10.0.3.3/32 (P012707)

R2 uses multicast address 224.0.0.13 to send prune messages to upstream interfaces. The IP address of the upstream router is 10.0.12.1. The multicast group 225.1.1.1 has been pruned. R1 sends a prune acknowledgement message to R2.

Add the deleted static IGMP multicast group again.

```
[R2-GigabitEthernet0/0/1]igmp static-group 225.1.1.1
Dec 31 2011 15:00:19.440.1+00:00 R2 PIM/7/JP:(public net): PIM ver 2 GFT sending 10.0.12.2 -> 10.0.12.1 on
Serial1/0/0   (P012633)
Dec 31 2011 15:00:19.440.2+00:00 R2 PIM/7/JP:(public net): Upstream 10.0.12.1, Groups 1, Holdtime 0 (P012639)
Dec 31 2011 15:00:19.440.3+00:00 R2 PIM/7/JP:(public net): Group: 225.1.1.1/32 --- 1 joins 0 prunes (P012648)
Dec 31 2011 15:00:19.440.4+00:00 R2 PIM/7/JP:(public net): Join: 10.0.3.3/32   (P012654)
Dec 31 2011 15:00:19.480.1+00:00 R2 PIM/7/JP:(public net): PIM ver 2 GAK receiving 10.0.12.1 -> 10.0.12.2 on
Serial1/0/0   (P012633)
Dec 31 2011 15:00:19.480.2+00:00 R2 PIM/7/JP:(public net): Upstream 10.0.12.2, Groups 1, Holdtime 0 (P012639)
Dec 31 2011 15:00:19.480.3+00:00 R2 PIM/7/JP:(public net): Group: 225.1.1.1/32 --- 1 joins 0 prunes (P012648)
Dec 31 2011 15:00:19.480.4+00:00 R2 PIM/7/JP:(public net): Join: 10.0.3.3/32   (P012654)

[R2-GigabitEthernet0/0/1]return
<R2>undo debugging all
Info: All possible debugging has been turned off
```

R2 immediately sends a prune message to the upstream device through unicast to join the multicast group 225.1.1.1. R1 returns a prune acknowledgement to R2 through unicast.

The prune message is sent from multicast address 224.0.0.13, while prune messages are sent to the upstream device through unicast.

If you need to transmit multicast traffic within a specified range, you can run the **multicast boundary** command on an interface to specify the range of a multicast group or multicast address segment.

Prevent the traffic from multicast group 225.1.1.2 from being transmitted to R4. Perform the following configuration on R1's interface connected to R4:

```
[R1]interface Serial 3/0/0
[R1-Serial3/0/0]multicast boundary 225.1.1.2 255.255.255.255
[R1-Serial3/0/0]quit
```

Simulate the multicast traffic with destination address 225.1.1.2 on R3.

```
[R3]ping -a 10.0.3.3 -c 300 225.1.1.2
```

Check the multicast routing table on R2 and R4. R2 has the entry (10.0.3.3, 225.1.1.2), while R4 does not have a routing entry of this multicast group. This indicates that multicast traffic is not spread to R4.

```
[R2]display pim routing-table
 VPN-Instance: public net
 Total 1 (*, G) entry; 2 (S, G) entries

 (*, 225.1.1.1)
     Protocol: pim-dm, Flag: WC EXT
     UpTime: 00:09:04
     Upstream interface: NULL
         Upstream neighbor: NULL
         RPF prime neighbor: NULL
     Downstream interface(s) information: None

 (10.0.3.3, 225.1.1.1)
```

```
     Protocol: pim-dm, Flag: EXT
     UpTime: 00:02:11
     Upstream interface: Serial1/0/0
         Upstream neighbor: 10.0.12.1
         RPF prime neighbor: 10.0.12.1
     Downstream interface(s) information: None

 (10.0.3.3, 225.1.1.2)
     Protocol: pim-dm, Flag:
     UpTime: 00:00:08
     Upstream interface: Serial1/0/0
         Upstream neighbor: 10.0.12.1
         RPF prime neighbor: 10.0.12.1
     Downstream interface(s) information: None

[R4]display pim routing-table
 VPN-Instance: public net
 Total 1 (*, G) entry; 1 (S, G) entry

 (*, 225.1.1.1)
     Protocol: pim-dm, Flag: WC
     UpTime: 00:08:03
     Upstream interface: NULL
         Upstream neighbor: NULL
         RPF prime neighbor: NULL
     Downstream interface(s) information:
     Total number of downstreams: 1
         1: GigabitEthernet0/0/1
             Protocol: static, UpTime: 00:08:03, Expires: never
```

```
(10.0.3.3, 225.1.1.1)
    Protocol: pim-dm, Flag:
    UpTime: 00:02:43
    Upstream interface: Serial1/0/0
        Upstream neighbor: 10.0.14.1
        RPF prime neighbor: 10.0.14.1
    Downstream interface(s) information:
    Total number of downstreams: 1
        1: GigabitEthernet0/0/1
            Protocol: pim-dm, UpTime: 00:02:43, Expires:   -
```

By default, PIM DM selects the router connected to the interface with a greater IP address as the DR.

```
[R2]display pim interface
 VPN-Instance: public net
 Interface           State NbrCnt HelloInt   DR-Pri     DR-Address
 GE0/0/0              up      1        30          1            10.0.25.5
 S1/0/0               up      1        30          1            10.0.12.2          (local)
```

Check the interface status on R2. R5 is the DR. You can change the interface priority to affect the DR election result. The priority is a 32-bit value. The default value is 1. In the following example, change the priority of the R2's interface connected to R5 to 100.

```
[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]pim hello-option dr-priority 100
[R2-GigabitEthernet0/0/0]quit
```

```
[R2]display pim interface
 VPN-Instance: public net
 Interface           State NbrCnt HelloInt DR-Pri     DR-Address
 GE0/0/0               up    1       30       100        10.0.25.2      (local)
 S1/0/0                up    1       30       1          10.0.12.2      (local)
```

After the interface priority is changed to 100, R2 becomes the DR.

For security purposes, you can disable the user-side interface from sending and receiving PIM hello packets by running the **pim silent** command.

```
[R4]interface GigabitEthernet 0/0/1
[R4-GigabitEthernet0/0/1]pim silent
[R4-GigabitEthernet0/0/1]quit
```

After the configuration, check whether PIM silent takes effect.

```
[R4]display pim interface GigabitEthernet 0/0/1 verbose
 VPN-Instance: public net
 Interface: GigabitEthernet0/0/1, 10.0.24.4
     PIM version: 2
     PIM mode: Dense
     PIM state: up
     PIM DR: 10.0.24.4 (local)
     PIM DR Priority (configured): 1
     PIM neighbor count: 0
     PIM hello interval: 30 s
     PIM LAN delay (negotiated): 500 ms
     PIM LAN delay (configured): 500 ms
     PIM hello override interval (negotiated): 2500 ms
     PIM hello override interval (configured): 2500 ms
```

PIM Silent: enabled

PIM neighbor tracking (negotiated): disabled

PIM neighbor tracking (configured): disabled

PIM generation ID: 0XAD457D14

PIM require-GenID: disabled

PIM hello hold interval: 105 s

PIM assert hold interval: 180 s

PIM triggered hello delay: 5 s

PIM J/P interval: 60 s

PIM J/P hold interval: 210 s

PIM state-refresh processing: enabled

PIM state-refresh interval: 60 s

PIM graft retry interval: 3 s

PIM state-refresh capability on link: capable

PIM dr-switch-delay timer : not configured

Number of routers on link not using DR priority: 0

Number of routers on link not using LAN delay: 0

Number of routers on link not using neighbor tracking: 1

ACL of PIM neighbor policy: -

ACL of PIM ASM join policy: -

ACL of PIM SSM join policy: -

ACL of PIM join policy: -

----End

Additional Exercises: Analysis and Verification

PIM DM is applicable to the high user density scenarios.

For which networks you will configure PIM DM? What are the characteristics of these networks?

What are the advantages and disadvantages of configuring PIM DM on a large-sized network?

Device Configurations

```
<R1>display current-configuration
[V200R007C00SPC600]
#
 sysname R1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
 pim dm
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
 pim dm
 multicast boundary 225.1.1.2 32
#
ip address 10.0.13.1 255.255.255.0
 pim dm
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
#
ospf 1 router-id 10.0.1.1
```

```
 area 0.0.0.0
  network 10.0.1.1 0.0.0.0
  network 10.0.14.1 0.0.0.0
  network 10.0.13.1 0.0.0.0
  network 10.0.12.1 0.0.0.0
#
return

<R2>display current-configuration
[V200R007C00SPC600]
#
 sysname R2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
 pim dm
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.2 255.255.255.0
 pim hello-option dr-priority 100
 pim dm
#
interface GigabitEthernet0/0/1
 ip address 10.0.24.2 255.255.255.0
 igmp enable
 igmp robust-count 3
 igmp timer query 20
 igmp static-group 225.1.1.1
```

```
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.255
#
ospf 1 router-id 10.0.2.2
 area 0.0.0.0
  network 10.0.2.2 0.0.0.0
  network 10.0.25.2 0.0.0.0
  network 10.0.12.2 0.0.0.0
#
 ip rpf-route-static 10.0.3.0 24 10.0.12.1
#
return

<R3>display current-configuration
[V200R007C00SPC600]
#
 sysname R3
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
 pim dm
#
interface GigabitEthernet0/0/2
 ip address 10.0.13.3 255.255.255.0
 pim dm
#
interface LoopBack0
```

```
 ip address 10.0.3.3 255.255.255.255
#
ospf 1 router-id 10.0.3.3
 area 0.0.0.0
  network 10.0.3.3 0.0.0.0
  network 10.0.13.3 0.0.0.0
  network 10.0.35.3 0.0.0.0
#
return

<R4>display current-configuration
[V200R007C00SPC600]
#
 sysname R4
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
 pim dm
#
interface GigabitEthernet0/0/1
 ip address 10.0.24.4 255.255.255.0
pim silent
 igmp enable
 igmp static-group 225.1.1.1
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
#
```

```
ospf 1 router-id 10.0.4.4
 area 0.0.0.0
  network 10.0.4.4 0.0.0.0
  network 10.0.14.4 0.0.0.0
#
return

<R5>display current-configuration
[V200R007C00SPC600]
#
 sysname R5
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
 pim dm
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.5 255.255.255.0
 pim dm
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.255
#
ospf 1 router-id 10.0.5.5
 area 0.0.0.0
  network 10.0.5.5 0.0.0.0
  network 10.0.25.5 0.0.0.0
  network 10.0.35.5 0.0.0.0
```

#

return

Lab 4-2 PIM SM and Static RP

Learning Objectives

The objectives of this lab are to learn and understand:

- How to configure PIM SM
- How to configure static RP and RP load balancing

Topology

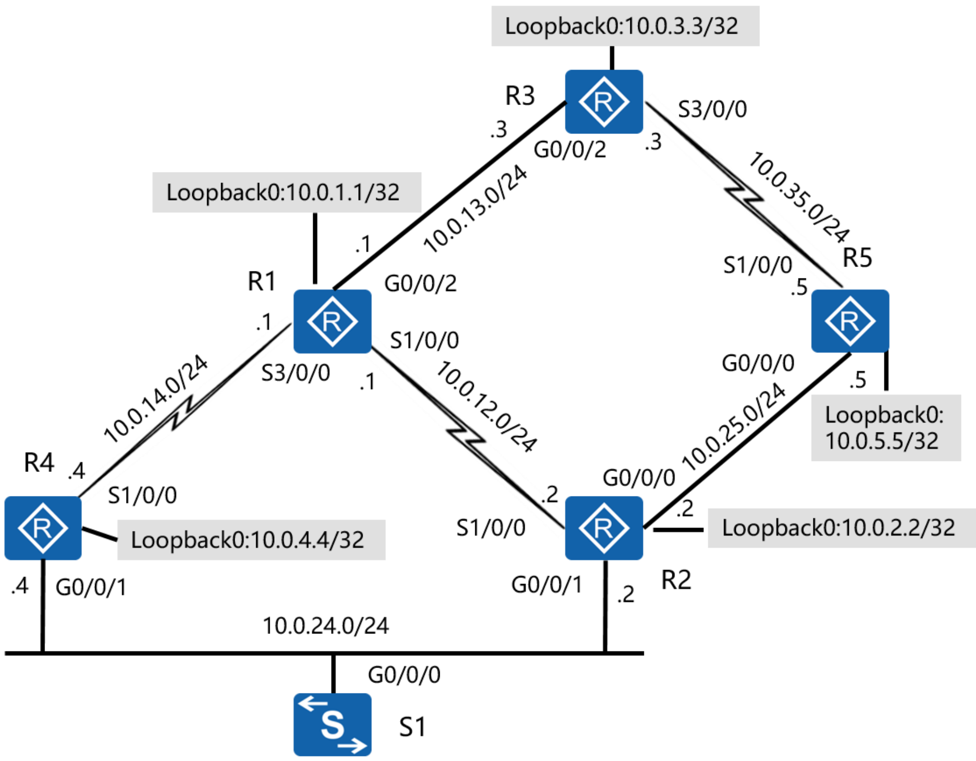

Figure 4-2 PIM SM and dynamic RP

Scenario

You are a network administrator of a company. The PIM DM has been configured on the company's network. However, when more and more multicast users are dispersed on the network, multicast service quality degrades. To improve multicast reliability, security, and efficiency, you can configure PIM SM.

In the PIM SM mode, you need to define the RP, which is used as the root of the shared tree in SM mode. In addition, RPs need to perform load balancing.

You may encounter network failures and need to rectify the faults.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

S2 participates in the experiment, but you do not need to configure S2. Before the experiment, clear the configuration on S2 and restart S2.

Configure IP addresses and masks for all the routers. All loopback interfaces must have 24-bit masks.

```
[R1]interface GigabitEthernet0/0/2
[R1-GigabitEthernet0/0/2]ip address 10.0.13.1 24
[R1-GigabitEthernet0/0/2]quit
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]quit
[R1]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 24
[R1-Serial3/0/0]quit
[R1]interface loopback 0
[R1-LoopBack0]ip address 10.0.1.1 24
[R1-LoopBack0]quit
```

```
[R2]interface GigabitEthernet0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.25.2 24
[R2-GigabitEthernet0/0/0]quit
[R2]interface GigabitEthernet0/0/1
[R2-GigabitEthernet0/0/1]ip address 10.0.24.2 24
[R2-GigabitEthernet0/0/1]quit
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]quit
[R2]interface loopback 0
[R2-LoopBack0]ip address 10.0.2.2 24
[R2-LoopBack0]quit

[R3]interface GigabitEthernet0/0/2
[R3-GigabitEthernet0/0/2]ip address 10.0.13.3 24
[R3-GigabitEthernet0/0/2]quit
[R3]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 24
[R3-Serial3/0/0]quit
[R3]interface loopback 0
[R3-LoopBack0]ip address 10.0.3.3 24
[R3-LoopBack0]quit

[R4]interface GigabitEthernet0/0/1
[R4-GigabitEthernet0/0/1]ip address 10.0.24.4 24
[R4-GigabitEthernet0/0/1]quit
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 24
[R4-Serial1/0/0]quit
```

[R4]interface loopback 0

[R4-LoopBack0]ip address 10.0.4.4 24

[R4-LoopBack0]quit

[R5]interface GigabitEthernet0/0/0

[R5-GigabitEthernet0/0/0]ip address 10.0.25.5 24

[R5-GigabitEthernet0/0/0]quit

[R5]interface Serial 1/0/0

[R5-Serial1/0/0]ip address 10.0.35.5 24

[R5-Serial1/0/0]quit

[R5]interface loopback 0

[R5-LoopBack0]ip address 10.0.5.5 24

[R5-LoopBack0]quit

[S1]interface Vlanif 1

[S1-Vlanif1]ip address 10.0.24.1 24

[S1-Vlanif1]interface loopback 0

[S1-LoopBack0]ip address 10.0.11.11 24

[S1-LoopBack0]quit

After the configurations are complete, test link connectivity.

[R1]ping -c 1 10.0.12.2

  PING 10.0.12.2: 56 data bytes, press CTRL_C to break

    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=41 ms

  --- 10.0.12.2 ping statistics ---

    1 packet(s) transmitted

    1 packet(s) received

```
      0.00% packet loss
      round-trip min/avg/max = 41/41/41 ms

[R1]ping -c 1 10.0.13.3
   PING 10.0.13.3: 56   data bytes, press CTRL_C to break
      Reply from 10.0.13.3: bytes=56 Sequence=1 ttl=255 time=5 ms

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

[R1]ping -c 1 10.0.14.4
   PING 10.0.14.4: 56   data bytes, press CTRL_C to break
      Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=62 ms

   --- 10.0.14.4 ping statistics ---
      1 packet(s) transmitted
      1 packet(s) received
      0.00% packet loss
      round-trip min/avg/max = 62/62/62 ms

[R5]ping -c 1 10.0.25.2
   PING 10.0.25.2: 56 data bytes, press CTRL_C to break
      Reply from 10.0.25.2: bytes=56 Sequence=1 ttl=255 time=7 ms

   --- 10.0.25.2 ping statistics ---
      1 packet(s) transmitted
```

```
      1 packet(s) received
      0.00% packet loss
      round-trip min/avg/max = 7/7/7 ms

[R5]ping -c 1 10.0.35.3
   PING 10.0.35.3: 56 data bytes, press CTRL_C to break
      Reply from 10.0.35.3: bytes=56 Sequence=1 ttl=255 time=37 ms

   --- 10.0.35.3 ping statistics ---
      1 packet(s) transmitted
      1 packet(s) received
      0.00% packet loss
round-trip min/avg/max = 37/37/37 ms

[S1]ping -c 1 10.0.24.2
   PING 10.0.24.2: 56   data bytes, press CTRL_C to break
      Reply from 10.0.24.2: bytes=56 Sequence=1 ttl=255 time=1 ms

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

Configure OSPF on R1, R2, R3, R4, R5, and S1. Implement network connectivity.

```
[R1]ospf 1 router-id 10.0.1.1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.0.14.1 0.0.0.0
```

```
[R1-ospf-1-area-0.0.0.0]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.13.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]quit
[R1-ospf-1]quit

[R2]ospf 1 router-id 10.0.2.2
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.12.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.24.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.25.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]quit
[R2-ospf-1]quit

[R3]ospf 1 router-id 10.0.3.3
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]network 10.0.13.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.35.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]quit
[R3-ospf-1]quit

[R4]ospf 1 router-id 10.0.4.4
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]network 10.0.14.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]network 10.0.24.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]network 10.0.4.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]quit
```

```
[R4-ospf-1]quit

[R5]ospf 1 router-id 10.0.5.5
[R5-ospf-1]area 0
[R5-ospf-1-area-0.0.0.0]network 10.0.25.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]network 10.0.35.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]network 10.0.5.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]quit
[R5-ospf-1]quit

[S1]ospf 1 router-id 10.0.11.11
[S1-ospf-1]area 0
[S1-ospf-1-area-0.0.0.0]network 10.0.24.1 0.0.0.0
[S1-ospf-1-area-0.0.0.0]network 10.0.11.11 0.0.0.0
[S1-ospf-1-area-0.0.0.0]quit
[S1-ospf-1]quit
```

After the configuration, wait until the OSPF neighbor relationship is set up. When route information exchange is complete, test the connectivity between S1 and Loopback0 of R3.

```
[S1]ping -c 1 10.0.3.3
  PING 10.0.3.3: 56    data bytes, press CTRL_C to break
    Reply from 10.0.3.3: bytes=56 Sequence=1 ttl=253 time=37 ms

  --- 10.0.3.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
```

round-trip min/avg/max = 37/37/37 ms

The test result shows that the network works normally.

Step 2 **Enable PIM SM on all routers.**

Enable multicast routing on R1, R2, R3, R4, R5, and S1.

[R1]multicast routing-enable

[R2]multicast routing-enable

[R3]multicast routing-enable

[R4]multicast routing-enable

[R5]multicast routing-enable

[S1]multicast routing-enable

Configure PIM SM on all interfaces of all devices.

[R1]interface GigabitEthernet0/0/2

[R1-GigabitEthernet0/0/2]pim sm

[R1-GigabitEthernet0/0/2]quit

[R1]interface Serial 1/0/0

[R1-Serial1/0/0]pim sm

[R1-Serial1/0/0]quit

[R1]interface Serial 3/0/0

[R1-Serial3/0/0]pim sm

[R1-Serial3/0/0]quit

```
[R1]interface loopback 0
[R1-LoopBack0]pim sm
[R1-LoopBack0]quit

[R2]interface GigabitEthernet0/0/0
[R2-GigabitEthernet0/0/0]pim sm
[R2-GigabitEthernet0/0/0]quit
[R2]interface GigabitEthernet0/0/1
[R2-GigabitEthernet0/0/1]pim sm
[R2-GigabitEthernet0/0/1]quit
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]pim sm
[R2-Serial1/0/0]quit
[R2]interface loopback 0
[R2-LoopBack0]pim sm
[R2-LoopBack0]quit

[R3]interface GigabitEthernet0/0/2
[R3-GigabitEthernet0/0/2]pim sm
[R3-GigabitEthernet0/0/2]quit
[R3]interface Serial 3/0/0
[R3-Serial3/0/0]pim sm
[R3-Serial3/0/0]quit
[R3]interface loopback 0
[R3-LoopBack0]pim sm
[R3-LoopBack0]quit

[R4]interface GigabitEthernet0/0/1
[R4-GigabitEthernet0/0/1]pim sm
```

```
[R4-GigabitEthernet0/0/1]quit
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]pim sm
[R4-Serial1/0/0]quit
[R4]interface loopback 0
[R4-LoopBack0]pim sm
[R4-LoopBack0]quit

[R5]interface GigabitEthernet0/0/0
[R5-GigabitEthernet0/0/0]pim sm
[R5-GigabitEthernet0/0/0]quit
[R5]interface Serial 1/0/0
[R5-Serial1/0/0]pim sm
[R5-Serial1/0/0]quit
[R5]interface loopback 0
[R5-LoopBack0]pim sm
[R5-LoopBack0]quit

[S1]interface Vlanif 1
[S1-Vlanif1]pim sm
[S1-Vlanif1]quit
[S1]interface loopback 0
[S1-LoopBack0]pim sm
[S1-LoopBack0]quit
```

After the configuration, check PIM neighbor learning information on R1, R5, and S1.

```
<R1>display pim neighbor
 VPN-Instance: public net
```

```
 Total Number of Neighbors = 3

 Neighbor        Interface           Uptime   Expires  Dr-Priority
 10.0.13.3       GE0/0/2             00:08:52 00:01:23 1
 10.0.12.2       S1/0/0              00:40:44 00:01:30 1
 10.0.14.4       S3/0/0              00:07:53 00:01:23 1

[R5]display pim neighbor
 VPN-Instance: public net
 Total Number of Neighbors = 2

 Neighbor        Interface           Uptime   Expires  Dr-Priority
 10.0.25.2       GE0/0/0             00:08:38 00:01:30 1
 10.0.35.3       S1/0/0              00:08:38 00:01:28 1

[S1]display pim neighbor
 VPN-Instance: public net
 Total Number of Neighbors = 2

 Neighbor        Interface           Uptime   Expires  Dr-Priority  BFD-Session
 10.0.24.4       Vlanif1             00:01:24 00:01:23 1            N
 10.0.24.2       Vlanif1             00:01:22 00:01:17 1            N
```

The command output shows that the PIM protocol has been running.

Step 3 **Implement load balance between static RPs.**

Specify static RPs to control multicast data flows on the network.

Configure R1's S3/0/0 as the static RP on all devices.

[R1]pim

[R1-pim]static-rp 10.0.14.1

[R1-pim]quit

[R2]pim

[R2-pim]static-rp 10.0.14.1

[R2-pim]quit

[R3]pim

[R3-pim]static-rp 10.0.14.1

[R3-pim]quit

[R4]pim

[R4-pim]static-rp 10.0.14.1

[R4-pim]quit

[R5]pim

[R5-pim]static-rp 10.0.14.1

[R5-pim]quit

[S1]pim

[S1-pim]static-rp 10.0.14.1

[S1-pim]quit

S1 is the Layer 3 switch connected to multicast users. Enable IGMP on Loopback 0 of S1.

[S1]interface LoopBack 0

[S1-LoopBack0]igmp enable

[S1-LoopBack0]quit

Add S1's Loopback 0 to the multicast group 225.0.0.1 to simulate multicast users of group 225.0.0.1.

[S1]interface LoopBack 0

[S1-LoopBack0]igmp static-group 225.0.0.1

[S1-LoopBack0]quit

Run the **display pim routing-table** command on R1, R4, and S1 to check the PIM routing table.

```
[R1]display pim routing-table
 VPN-Instance: public net
 Total 1 (*, G) entry; 0 (S, G) entry

 (*, 225.0.0.1)
     RP: 10.0.14.1 (local)
     Protocol: pim-sm, Flag: WC
     UpTime: 00:02:40
     Upstream interface: Register
         Upstream neighbor: NULL
         RPF prime neighbor: NULL
     Downstream interface(s) information:
     Total number of downstreams: 1
         1: Serial3/0/0
             Protocol: pim-sm, UpTime: 00:02:40, Expires: 00:02:50

[R4]display pim routing-table
 VPN-Instance: public net
```

```
 Total 1 (*, G) entry; 0 (S, G) entry

 (*, 225.0.0.1)
     RP: 10.0.14.1
     Protocol: pim-sm, Flag: WC
     UpTime: 00:01:46
     Upstream interface: Serial1/0/0
         Upstream neighbor: 10.0.14.1
         RPF prime neighbor: 10.0.14.1
     Downstream interface(s) information:
     Total number of downstreams: 1
         1: GigabitEthernet0/0/1
             Protocol: pim-sm, UpTime: 00:01:46, Expires: 00:02:43

[S1-LoopBack0]display pim routing-table
 VPN-Instance: public net
 Total 1 (*, G) entry; 0 (S, G) entry

 (*, 225.0.0.1)
     RP: 10.0.14.1
     Protocol: pim-sm, Flag: WC
     UpTime: 00:01:19
     Upstream interface: Vlanif1
         Upstream neighbor: 10.0.24.4
         RPF prime neighbor: 10.0.24.4
     Downstream interface(s) information:
     Total number of downstreams: 1
         1: LoopBack0
             Protocol: static, UpTime: 00:01:19, Expires: -
```

The command output shows that R1 is the RP on the network. S1 generates a multicast path to R1 through R4.

Create an ACL and apply the ACL to static RP. Specify R1 as the RP serving the network segment 225.0.0.0/24. Specify R5 as the RP serving the network segment 225.0.1.0/24.

```
[R1]acl 2000
[R1-acl-basic-2000]rule permit source 225.0.0.0 0.0.0.255
[R1-acl-basic-2000]quit
[R1]acl 2001
[R1-acl-basic-2001]rule permit source 225.0.1.0 0.0.0.255
[R1-acl-basic-2001]quit
[R1]pim
[R1-pim]static-rp 10.0.14.1 2000
[R1-pim]static-rp 10.0.25.5 2001
[R1-pim]quit

[R2]acl 2000
[R2-acl-basic-2000]rule permit source 225.0.0.0 0.0.0.255
[R2-acl-basic-2000]quit
[R2]acl 2001
[R2-acl-basic-2001]rule permit source 225.0.1.0 0.0.0.255
[R2-acl-basic-2001]quit
[R2]pim
[R2-pim]static-rp 10.0.14.1 2000
[R2-pim]static-rp 10.0.25.5 2001
[R2-pim]quit
```

```
[R3]acl 2000
[R3-acl-basic-2000]rule permit source 225.0.0.0 0.0.0.255
[R3-acl-basic-2000]quit
[R3]acl 2001
[R3-acl-basic-2001]rule permit source 225.0.1.0 0.0.0.255
[R3-acl-basic-2001]quit
[R3]pim
[R3-pim]static-rp 10.0.14.1 2000
[R3-pim]static-rp 10.0.25.5 2001
[R3-pim]quit

[R4]acl 2000
[R4-acl-basic-2000]rule permit source 225.0.0.0 0.0.0.255
[R4-acl-basic-2000]quit
[R4]acl 2001
[R4-acl-basic-2001]rule permit source 225.0.1.0 0.0.0.255
[R4-acl-basic-2001]quit
[R4]pim
[R4-pim]static-rp 10.0.14.1 2000
[R4-pim]static-rp 10.0.25.5 2001
[R4-pim]quit

[R5]acl 2000
[R5-acl-basic-2000]rule permit source 225.0.0.0 0.0.0.255
[R5-acl-basic-2000]quit
[R5]acl 2001
[R5-acl-basic-2001]rule permit source 225.0.1.0 0.0.0.255
[R5-acl-basic-2001]quit
[R5]pim
```

[R5-pim]static-rp 10.0.14.1 2000

[R5-pim]static-rp 10.0.25.5 2001

[R5-pim]quit

[S1]acl 2000

[S1-acl-basic-2000]rule permit source 225.0.0.0 0.0.0.255

[S1-acl-basic-2000]quit

[S1]acl 2001

[S1-acl-basic-2001]rule permit source 225.0.1.0 0.0.0.255

[S1-acl-basic-2001]quit

[S1]pim

[S1-pim]static-rp 10.0.14.1 2000

[S1-pim]static-rp 10.0.25.5 2001

Add S1's Loopback 0 to the multicast group 225.0.1.1 to simulate multicast users of group 225.0.1.1.

[S1]interface LoopBack 0

[S1-LoopBack0]igmp static-group 225.0.1.1

[S1-LoopBack0]quit

Run the **display pim routing-table** command on S1, R2, and R5 to check the PIM routing table.

[R5]display pim routing-table

VPN-Instance: public net

Total 1 (*, G) entry; 0 (S, G) entry

(*, 225.0.1.1)

RP: 10.0.25.5 (local)

```
     Protocol: pim-sm, Flag: WC
     UpTime: 00:03:13
     Upstream interface: Register
         Upstream neighbor: NULL
         RPF prime neighbor: NULL
     Downstream interface(s) information:
     Total number of downstreams: 1
         1: GigabitEthernet0/0/0
             Protocol: pim-sm, UpTime: 00:03:13, Expires: 00:03:17

[R2]display pim routing-table
 VPN-Instance: public net
 Total 1 (*, G) entry; 0 (S, G) entry

 (*, 225.0.1.1)
     RP: 10.0.25.5
     Protocol: pim-sm, Flag: WC
     UpTime: 00:03:41
     Upstream interface: GigabitEthernet0/0/0
         Upstream neighbor: 10.0.25.5
         RPF prime neighbor: 10.0.25.5
     Downstream interface(s) information:
     Total number of downstreams: 1
         1: GigabitEthernet0/0/1
             Protocol: pim-sm, UpTime: 00:03:41, Expires: 00:02:48

[S1]display pim routing-table
 VPN-Instance: public net
 Total 2 (*, G) entries; 0 (S, G) entry
```

```
(*, 225.0.0.1)
    RP: 10.0.14.1
    Protocol: pim-sm, Flag: WC
    UpTime: 00:17:09
    Upstream interface: Vlanif1
        Upstream neighbor: 10.0.24.4
        RPF prime neighbor: 10.0.24.4
    Downstream interface(s) information:
    Total number of downstreams: 1
        1: LoopBack0
            Protocol: static, UpTime: 00:17:09, Expires: -

(*, 225.0.1.1)
    RP: 10.0.25.5
    Protocol: pim-sm, Flag: WC
    UpTime: 00:03:58
    Upstream interface: Vlanif1
        Upstream neighbor: 10.0.24.2
        RPF prime neighbor: 10.0.24.2
    Downstream interface(s) information:
    Total number of downstreams: 1
        1: LoopBack0
            Protocol: static, UpTime: 00:03:58, Expires: -
```

The command output shows that S1 generates two multicast paths for 225.0.0.1 and 225.0.1.1. The multicast path of 225.0.1.1 reaches R5 through R2.

----End

Additional Exercises: Analysis and Verification

PIM SM is applicable to the scenarios where users are dispersed.

For which networks you will configure PIM SM? What are the characteristics of these networks?

Device Configurations

```
[R1]display current-configuration
[V200R007C00SPC600]
#
 sysname R1
#
 board add 0/1 1SA
 board add 0/2 1SA
 board add 0/3 1SA
 board add 0/4 1SA
#
 drop illegal-mac alarm
#
multicast routing-enable
#
acl number 2000
 rule 5 permit source 225.0.0.0 0.0.0.255
acl number 2001
 rule 5 permit source 225.0.1.0 0.0.0.255
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
 pim sm
```

```
#
interface Serial2/0/0
 link-protocol ppp
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
 pim sm
#
interface Serial4/0/0
 link-protocol ppp
#
interface GigabitEthernet0/0/0
#
interface GigabitEthernet0/0/1
#
interface GigabitEthernet0/0/2
 ip address 10.0.13.1 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/3
 description VirtualPort
#
interface Cellular0/0/0
#
interface Cellular0/0/1
#
interface NULL0
#
```

```
interface LoopBack0
 ip address 10.0.1.1 255.255.255.0
 pim sm
#
ospf 1 router-id 10.0.1.1
 area 0.0.0.0
  network 10.0.1.1 0.0.0.0
  network 10.0.12.1 0.0.0.0
  network 10.0.13.1 0.0.0.0
  network 10.0.14.1 0.0.0.0
#
pim
 static-rp 10.0.14.1 2000
 static-rp 10.0.25.5 2001
#
return

[R2]display current-configuration
[V200R007C00SPC600]
#
 sysname R2
#
 board add 0/1 1SA
 board add 0/2 1SA
 board add 0/3 1SA
 board add 0/4 2FE
#
 drop illegal-mac alarm
#
```

```
multicast routing-enable
#
acl number 2000
 rule 5 permit source 225.0.0.0 0.0.0.255
acl number 2001
 rule 5 permit source 225.0.1.0 0.0.0.255
#
interface Ethernet4/0/0
#
interface Ethernet4/0/1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
 pim sm
#
interface Serial2/0/0
 link-protocol ppp
#
interface Serial3/0/0
 link-protocol ppp
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.2 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/1
 ip address 10.0.24.2 255.255.255.0
 pim sm
```

```
#
interface GigabitEthernet0/0/2
#
interface GigabitEthernet0/0/3
 description VirtualPort
#
interface Cellular0/0/0
#
interface Cellular0/0/1
#
interface NULL0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.0
 pim sm
#
ospf 1 router-id 10.0.2.2
 area 0.0.0.0
  network 10.0.2.2 0.0.0.0
  network 10.0.12.2 0.0.0.0
  network 10.0.24.2 0.0.0.0
  network 10.0.25.2 0.0.0.0
#
pim
 static-rp 10.0.14.1 2000
 static-rp 10.0.25.5 2001
#
return
```

```
[R3]display current-configuration
[V200R007C00SPC600]
#
 sysname R3
#
 board add 0/1 1SA
 board add 0/2 1SA
 board add 0/3 1SA
 board add 0/4 2FE
#
 drop illegal-mac alarm
#
multicast routing-enable
#
acl number 2000
 rule 5 permit source 225.0.0.0 0.0.0.255
acl number 2001
 rule 5 permit source 225.0.1.0 0.0.0.255
#
interface Ethernet4/0/0
#
interface Ethernet4/0/1
#
interface Serial1/0/0
 link-protocol ppp
#
interface Serial2/0/0
 link-protocol ppp
#
```

```
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/0
 ip address 192.168.1.1 255.255.255.0
#
interface GigabitEthernet0/0/1
#
interface GigabitEthernet0/0/2
 ip address 10.0.13.3 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/3
 description VirtualPort
#
interface Cellular0/0/0
#
interface Cellular0/0/1
#
interface NULL0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.0
 pim sm
#
ospf 1 router-id 10.0.3.3
 area 0.0.0.0
```

```
  network 10.0.3.3 0.0.0.0
  network 10.0.13.3 0.0.0.0
  network 10.0.35.3 0.0.0.0
#
pim
 static-rp 10.0.14.1 2000
 static-rp 10.0.25.5 2001
#
return

[R4]display current-configuration
[V200R007C00SPC600]
#
 sysname R4
#
 board add 0/1 2SA
 board add 0/2 2FE
#
 drop illegal-mac alarm
#
multicast routing-enable
#
acl number 2000
 rule 5 permit source 225.0.0.0 0.0.0.255
acl number 2001
 rule 5 permit source 225.0.1.0 0.0.0.255
#
interface Ethernet2/0/0
#
```

```
interface Ethernet2/0/1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
 pim sm
#
interface Serial1/0/1
 link-protocol ppp
#
interface GigabitEthernet0/0/0
#
interface GigabitEthernet0/0/1
 ip address 10.0.24.4 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/2
#
interface GigabitEthernet0/0/3
 description VirtualPort
#
interface Cellular0/0/0
#
interface Cellular0/0/1
#
interface NULL0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.0
```

```
 pim sm
#
ospf 1 router-id 10.0.4.4
 area 0.0.0.0
  network 10.0.4.4 0.0.0.0
  network 10.0.14.4 0.0.0.0
  network 10.0.24.4 0.0.0.0
#
pim
 static-rp 10.0.14.1 2000
 static-rp 10.0.25.5 2001
#
return

[R5]display current-configuration
[V200R007C00SPC600]
#
 sysname R5
#
 board add 0/1 2SA
 board add 0/2 2FE
#
 drop illegal-mac alarm
#
multicast routing-enable
#
acl number 2000
 rule 5 permit source 225.0.0.0 0.0.0.255
acl number 2001
```

```
 rule 5 permit source 225.0.1.0 0.0.0.255
#
interface Ethernet2/0/0
#
interface Ethernet2/0/1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
 pim sm
#
interface Serial1/0/1
 link-protocol ppp
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.5 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/1
#
interface GigabitEthernet0/0/2
#
interface GigabitEthernet0/0/3
 description VirtualPort
#
interface Cellular0/0/0
#
interface Cellular0/0/1
#
```

```
interface NULL0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.0
 pim sm
#
ospf 1 router-id 10.0.5.5
 area 0.0.0.0
  network 10.0.5.5 0.0.0.0
  network 10.0.25.5 0.0.0.0
  network 10.0.35.5 0.0.0.0
#
pim
 static-rp 10.0.14.1 2000
 static-rp 10.0.25.5 2001
#
return

[S1]display current-configuration
!Software Version V200R008C00SPC500
#
sysname S1
#
multicast routing-enable
#
diffserv domain default
#
acl number 2000
 rule 5 permit source 225.0.0.0 0.0.0.255
```

```
acl number 2001
 rule 5 permit source 225.0.1.0 0.0.0.255
#
interface Vlanif1
 ip address 10.0.24.1 255.255.255.0
 pim sm
#
interface LoopBack0
 ip address 10.0.11.11 255.255.255.0
 pim sm
 igmp enable
 igmp static-group 225.0.0.1
 igmp static-group 225.0.1.1
#
ospf 1
 area 0.0.0.0
  network 10.0.11.11 0.0.0.0
  network 10.0.24.1 0.0.0.0
#
pim
 static-rp 10.0.14.1 2000
 static-rp 10.0.25.5 2001
#
return
```

Chapter 5 Route Control

Lab 5-1 Route Import and Control

Learning Objectives

The objectives of this lab are to learn and understand:

- How to import OSPF and ISIS routes to each other
- How to configure route filtering based on IP prefix list
- How to filter routes using a route-policy

Topology

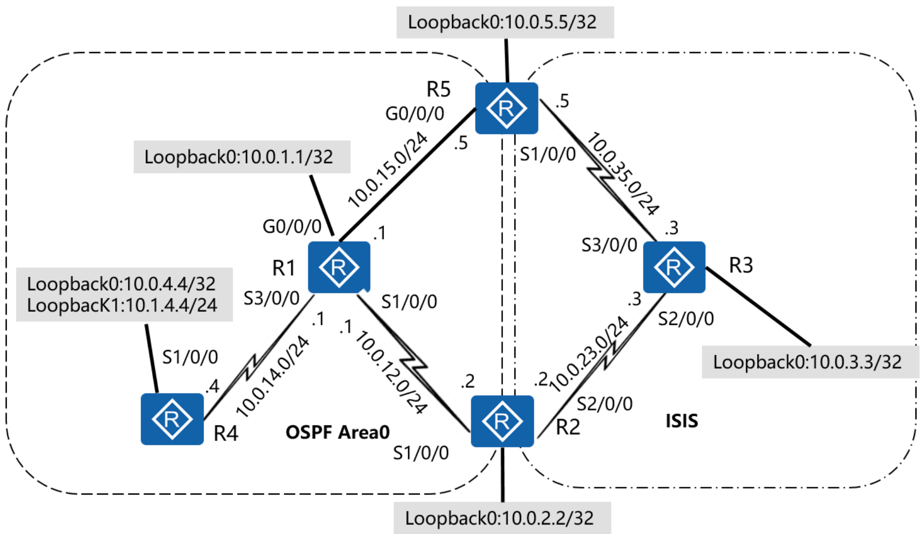

Figure 5-1 Route import and control

Scenario

You are a network administrator of a company. The company's network has OSPF areas and ISIS areas. To implement network connectivity, configure route import. When two routing protocols on two devices import routes from each other, some problems may occur, such as routing loops and sub-optimal routes. To prevent these problems, you can configure the IP prefix list and route-policy to control routes.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

Configure IP addresses and masks for physical interfaces and loopback interfaces of all routers. Each Loopback0 uses the 32-bit mask.

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface  Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 255.255.255.0
[R1-Serial1/0/0]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]ip address 10.0.15.1 255.255.255.0
[R1-GigabitEthernet0/0/0]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 255.255.255.0
[R1-Serial3/0/0]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 255.255.255.255

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 255.255.255.0
[R2-Serial1/0/0]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 255.255.255.0
[R2-Serial2/0/0]interface LoopBack 0
[R2-LoopBack0]ip add 10.0.2.2 255.255.255.255

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 255.255.255.0
```

```
[R3-Serial2/0/0]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 255.255.255.0
[R3-Serial3/0/0]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 255.255.255.255

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 255.255.255.0
[R4-Serial1/0/0]interface LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 255.255.255.255

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface Serial 1/0/0
[R5-Serial1/0/0]ip address 10.0.35.5 255.255.255.0
[R5-Serial1/0/0]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.15.5 255.255.255.0
[R5-GigabitEthernet0/0/0]interface LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 255.255.255.255
```

After the configurations are complete, test link connectivity.

```
[R1]ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=38 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
```

```
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 38/38/38 ms

[R1]ping -c 1 10.0.15.5
  PING 10.0.15.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.15.5: bytes=56 Sequence=1 ttl=255 time=12 ms

  --- 10.0.15.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 12/12/12 ms

[R1]ping -c 1 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=33 ms

  --- 10.0.14.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 33/33/33 ms

[R2]ping -c 1 10.0.23.3
  PING 10.0.23.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.3: bytes=56 Sequence=1 ttl=255 time=34 ms

  --- 10.0.23.3 ping statistics ---
```

```
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 34/34/34 ms

[R3]ping -c 1 10.0.35.5
  PING 10.0.35.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.35.5: bytes=56 Sequence=1 ttl=255 time=39 ms

  --- 10.0.35.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 39/39/39 ms
```

Step 2 **Configure an IGP protocol.**

Run OSPF on R1, R2, R4, and R5. All devices belong to area 0.

Run OSPF on the network segments connected to R1's S1/0/0, S3/0/0, G0/0/0, and Loopback0.

```
[R1]ospf 1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.15.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.14.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.1.1 0.0.0.0
```

Run OSPF on the network segment connected to R2's S1/0/0.

```
[R2]ospf 1
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.12.2 0.0.0.0
```

Run OSPF on the network segment connected to R4's S1/0/0 and Loopback0.

```
[R4]ospf 1
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]network 10.0.14.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]network 10.0.4.4 0.0.0.0
```

Run OSPF on the network segment connected to R5's G0/0/0.

```
[R5]ospf 1
[R5-ospf-1]area 0
[R5-ospf-1-area-0.0.0.0]network 10.0.15.5 0.0.0.0
```

Check whether the routers can learn the routes from the network segments connected to Loopback0 of other routers.

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 17          Routes : 17

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
       10.0.4.4/32  OSPF    10   1562        D   10.0.14.4       Serial3/0/0
```

```
      10.0.12.0/24  Direct  0    0        D   10.0.12.1        Serial1/0/0
      10.0.12.1/32  Direct  0    0        D   127.0.0.1        InLoopBack0
      10.0.12.2/32  Direct  0    0        D   10.0.12.2        Serial1/0/0
    10.0.12.255/32  Direct  0    0        D   127.0.0.1        InLoopBack0
      10.0.14.0/24  Direct  0    0        D   10.0.14.1        Serial3/0/0
      10.0.14.1/32  Direct  0    0        D   127.0.0.1        InLoopBack0
      10.0.14.4/32  Direct  0    0        D   10.0.14.4        Serial3/0/0
    10.0.14.255/32  Direct  0    0        D   127.0.0.1        InLoopBack0
      10.0.15.0/24  Direct  0    0        D   10.0.15.1    GigabitEthernet0/0/0
      10.0.15.1/32  Direct  0    0        D   127.0.0.1        InLoopBack0
    10.0.15.255/32  Direct  0    0        D   127.0.0.1        InLoopBack0
      127.0.0.0/8   Direct  0    0        D   127.0.0.1        InLoopBack0
      127.0.0.1/32  Direct  0    0        D   127.0.0.1        InLoopBack0
127.255.255.255/32Direct  0    0        D   127.0.0.1        InLoopBack0
255.255.255.255/32Direct  0    0        D   127.0.0.1        InLoopBack0

[R2]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 17        Routes : 17

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        10.0.1.1/32  OSPF    10   1562        D   10.0.12.1       Serial1/0/0
        10.0.2.2/32  Direct  0    0           D   127.0.0.1       InLoopBack0
        10.0.4.4/32  OSPF    10   3124        D   10.0.12.1       Serial1/0/0
      10.0.12.0/24  Direct   0    0           D   10.0.12.2       Serial1/0/0
      10.0.12.1/32  Direct   0    0           D   10.0.12.1       Serial1/0/0
```

```
       10.0.12.2/32  Direct   0    0           D   127.0.0.1       InLoopBack0
     10.0.12.255/32  Direct   0    0           D   127.0.0.1       InLoopBack0
       10.0.14.0/24  OSPF     10   3124        D   10.0.12.1       Serial1/0/0
       10.0.15.0/24  OSPF     10   1563        D   10.0.12.1       Serial1/0/0
       10.0.23.0/24  Direct   0    0           D   10.0.23.2       Serial2/0/0
       10.0.23.2/32  Direct   0    0           D   127.0.0.1       InLoopBack0
       10.0.23.3/32  Direct   0    0           D   10.0.23.3       Serial2/0/0
     10.0.23.255/32  Direct   0    0           D   127.0.0.1       InLoopBack0
       127.0.0.0/8   Direct   0    0           D   127.0.0.1       InLoopBack0
       127.0.0.1/32  Direct   0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32Direct   0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct   0    0           D   127.0.0.1       InLoopBack0

[R4]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 12        Routes : 12

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        10.0.1.1/32  OSPF     10   1562        D   10.0.14.1       Serial1/0/0
        10.0.4.4/32  Direct   0    0           D   127.0.0.1       InLoopBack0
       10.0.12.0/24  OSPF     10   3124        D   10.0.14.1       Serial1/0/0
       10.0.14.0/24  Direct   0    0           D   10.0.14.4       Serial1/0/0
       10.0.14.1/32  Direct   0    0           D   10.0.14.1       Serial1/0/0
       10.0.14.4/32  Direct   0    0           D   127.0.0.1       InLoopBack0
     10.0.14.255/32  Direct   0    0           D   127.0.0.1       InLoopBack0
       10.0.15.0/24  OSPF     10   1563        D   10.0.14.1       Serial1/0/0
```

```
        127.0.0.0/8   Direct  0    0        D   127.0.0.1       InLoopBack0
        127.0.0.1/32  Direct  0    0        D   127.0.0.1       InLoopBack0
127.255.255.255/32Direct  0    0        D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct  0    0        D   127.0.0.1       InLoopBack0

[R5]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 16        Routes : 16

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

         10.0.1.1/32  OSPF   10   1           D   10.0.15.1   GigabitEthernet0/0/0
         10.0.4.4/32  OSPF   10   1563        D   10.0.15.1   GigabitEthernet0/0/0
         10.0.5.5/32  Direct  0    0          D   127.0.0.1       InLoopBack0
        10.0.12.0/24  OSPF   10   1563        D   10.0.15.1   GigabitEthernet0/0/0
        10.0.14.0/24  OSPF   10   1563        D   10.0.15.1   GigabitEthernet0/0/0
        10.0.15.0/24  Direct  0    0          D   10.0.15.5   GigabitEthernet0/0/0
        10.0.15.5/32  Direct  0    0          D   127.0.0.1       InLoopBack0
      10.0.15.255/32 Direct  0    0          D   127.0.0.1       InLoopBack0
        10.0.35.0/24  Direct  0    0          D   10.0.35.5       Serial1/0/0
        10.0.35.3/32  Direct  0    0          D   10.0.35.3       Serial1/0/0
        10.0.35.5/32  Direct  0    0          D   127.0.0.1       InLoopBack0
      10.0.35.255/32 Direct  0    0          D   127.0.0.1       InLoopBack0
        127.0.0.0/8   Direct  0    0          D   127.0.0.1       InLoopBack0
        127.0.0.1/32  Direct  0    0          D   127.0.0.1       InLoopBack0
127.255.255.255/32Direct  0    0          D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct  0    0          D   127.0.0.1       InLoopBack0
```

Configure ISIS on R2, R3, and R5.

Run ISIS on the network segments connected to R2's S2/0/0 and Loopback0.

```
[R2]isis 1
[R2-isis-1]network-entity 49.0001.0000.0000.0002.00
[R2-isis-1]is-level level-2
[R2-isis-1]interface serial2/0/0
[R2-Serial2/0/0]isis enable 1
[R2]interface loopback0
[R2-LoopBack0]isis enable 1
```

Run ISIS on the network segments connected to R3's S2/0/0, S3/0/0, and Loopback0.

```
[R3]isis 1
[R3-isis-1]network-entity 49.0001.0000.0000.0003.00
[R3-isis-1]is-level level-2
[R3-isis-1]interface serial2/0/0
[R3-Serial2/0/0]isis enable 1
[R3-Serial2/0/0]interface serial3/0/0
[R3-Serial3/0/0]isis enable 1
[R3-Serial3/0/0]interface loopback0
[R3-LoopBack0]isis enable 1
```

Run ISIS on the network segments connected to R5's S1/0/0 and Loopback0.

```
[R5]isis 1
[R5-isis-1]network-entity 49.0001.0000.0000.0005.00
[R5-isis-1]is-level level-2
[R5-isis-1]interface serial1/0/0
[R5-Serial1/0/0]isis enable 1
```

```
[R5-Serial1/0/0]interface loopback 0
[R5-LoopBack0]isis enable 1
```

Check whether the routers learn Loopback0 addresses of other devices.

```
[R2]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 20        Routes : 20
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
       10.0.1.1/32  OSPF    10   48          D   10.0.12.1       Serial1/0/0
       10.0.2.2/32  Direct  0    0           D   127.0.0.1       LoopBack0
       10.0.3.3/32  ISIS-L2 15   10          D   10.0.23.3       Serial2/0/0
       10.0.4.4/32  OSPF    10   96          D   10.0.12.1       Serial1/0/0
       10.0.5.5/32  ISIS-L2 15   20          D   10.0.23.3       Serial2/0/0
      10.0.12.0/24  Direct  0    0           D   10.0.12.2       Serial1/0/0
      10.0.12.1/32  Direct  0    0           D   10.0.12.1       Serial1/0/0
      10.0.12.2/32  Direct  0    0           D   127.0.0.1       Serial1/0/0
    10.0.12.255/32  Direct  0    0           D   127.0.0.1       Serial1/0/0
      10.0.14.0/24  OSPF    10   96          D   10.0.12.1       Serial1/0/0
      10.0.15.0/24  OSPF    10   49          D   10.0.12.1       Serial1/0/0
      10.0.23.0/24  Direct  0    0           D   10.0.23.2       Serial2/0/0
      10.0.23.2/32  Direct  0    0           D   127.0.0.1       Serial2/0/0
      10.0.23.3/32  Direct  0    0           D   10.0.23.3       Serial2/0/0
    10.0.23.255/32  Direct  0    0           D   127.0.0.1       Serial2/0/0
      10.0.35.0/24  ISIS-L2 15   20          D   10.0.23.3       Serial2/0/0
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

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


[R3]dis ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 15        Routes : 15

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        10.0.2.2/32  ISIS-L2 15   10          D   10.0.23.2       Serial2/0/0
        10.0.3.3/32  Direct  0    0           D   127.0.0.1       LoopBack0
        10.0.5.5/32  ISIS-L2 15   10          D   10.0.35.5       Serial3/0/0
       10.0.23.0/24  Direct  0    0           D   10.0.23.3       Serial2/0/0
       10.0.23.2/32  Direct  0    0           D   10.0.23.2       Serial2/0/0
       10.0.23.3/32  Direct  0    0           D   127.0.0.1       Serial2/0/0
     10.0.23.255/32  Direct  0    0           D   127.0.0.1       Serial2/0/0
       10.0.35.0/24  Direct  0    0           D   10.0.35.3       Serial3/0/0
       10.0.35.3/32  Direct  0    0           D   127.0.0.1       Serial3/0/0
       10.0.35.5/32  Direct  0    0           D   10.0.35.5       Serial3/0/0
     10.0.35.255/32  Direct  0    0           D   127.0.0.1       Serial3/0/0
       127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
       127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0


[R5]display   ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
```

```
Routing Tables: Public
         Destinations : 19        Routes : 19

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.1/32  OSPF    10   1           D   10.0.15.1       GigabitEthernet0/0/0
       10.0.2.2/32  ISIS-L2 15   20          D   10.0.35.3       Serial1/0/0
       10.0.3.3/32  ISIS-L2 15   10          D   10.0.35.3       Serial1/0/0
       10.0.4.4/32  OSPF    10   49          D   10.0.15.1       GigabitEthernet0/0/0
       10.0.5.5/32  Direct  0    0           D   127.0.0.1       LoopBack0
      10.0.12.0/24  OSPF    0    49          D   10.0.15.1       GigabitEthernet0/0/0
      10.0.14.0/24  OSPF    10   49          D   10.0.15.1       GigabitEthernet0/0/0
      10.0.15.0/24  Direct  0    0           D   10.0.15.5       GigabitEthernet0/0/0
      10.0.15.5/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
    10.0.15.255/32  Direct  0    0           D   127.0.0.1       GigabitEthernet0/0/0
      10.0.23.0/24  ISIS-L2 15   20          D   10.0.35.3       Serial1/0/0
      10.0.35.0/24  Direct  0    0           D   10.0.35.5       Serial1/0/0
      10.0.35.3/32  Direct  0    0           D   10.0.35.3       Serial1/0/0
      10.0.35.5/32  Direct  0    0           D   127.0.0.1       Serial1/0/0
    10.0.35.255/32  Direct  0    0           D   127.0.0.1       Serial1/0/0
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

Step 3 **Configure a prefix list to filter routes.**

Create static routes 1.1.1.1/32, 1.1.1.0/24, 1.1.1.0/25, 1.1.0.0/16, and 1.0.0.0/8 pointing to NULL 0 on R1. Run the **import-route static** command to import routes to OSPF.

[R1]ip route-static 1.1.1.1 255.255.255.255 NULL 0

[R1]ip route-static 1.1.1.0 255.255.255.0 NULL 0

[R1]ip route-static 1.1.1.0 255.255.255.128 NULL 0

[R1]ip route-static 1.1.0.0 255.255.0.0 NULL 0

[R1]ip route-static 1.0.0.0 255.0.0.0 NULL 0

[R1]ospf 1

[R1-ospf-1]import-route static

Check whether R4 receives the static routes added to R1.

[R4]display ip routing-table

Route Flags: R - relay, D - download to fib

Routing Tables: Public

Destinations : 17 Routes : 17

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
1.0.0.0/8	O_ASE	150	1	D	10.0.14.1	Serial1/0/0
1.1.0.0/16	O_ASE	150	1	D	10.0.14.1	Serial1/0/0
1.1.1.0/24	O_ASE	150	1	D	10.0.14.1	Serial1/0/0
1.1.1.0/25	O_ASE	150	1	D	10.0.14.1	Serial1/0/0
1.1.1.1/32	O_ASE	150	1	D	10.0.14.1	Serial1/0/0
10.0.1.1/32	OSPF	10	1562	D	10.0.14.1	Serial1/0/0
10.0.4.4/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.12.0/24	OSPF	10	3124	D	10.0.14.1	Serial1/0/0
10.0.14.0/24	Direct	0	0	D	10.0.14.4	Serial1/0/0
10.0.14.1/32	Direct	0	0	D	10.0.14.1	Serial1/0/0
10.0.14.4/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.14.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0

```
     10.0.15.0/24     OSPF   10    1563       D    10.0.14.1       Serial1/0/0
     127.0.0.0/8      Direct  0     0          D    127.0.0.1       InLoopBack0
     127.0.0.1/32     Direct  0     0          D    127.0.0.1       InLoopBack0
127.255.255.255/32    Direct  0     0          D    127.0.0.1       InLoopBack0
255.255.255.255/32    Direct  0     0          D    127.0.0.1       InLoopBack0
```

Configure the prefix list **pref_r1** on R1 to match the route 1.1.1.0/24.

```
[R1]ip ip-prefix pref_r1 index 10 permit 1.1.1.0 24 greater-equal 24 less-equal 24
```

Create the routing policy **policy_r1** and invoke the prefix list **pref_r1** to control static route import on R1.

```
[R1]route-policy policy_r1 permit node 10
[R1-route-policy]if-match ip-prefix pref_r1
[R1-route-policy]ospf
[R1-ospf-1]import-route static route-policy policy_r1
```

Check the routing table on R4.

```
[R4]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 13        Routes : 13

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        1.1.1.0/24     O_ASE   150   1           D    10.0.14.1       Serial1/0/0
       10.0.1.1/32     OSPF    10    1562        D    10.0.14.1       Serial1/0/0
```

10.0.4.4/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.12.0/24	OSPF	10	3124	D	10.0.14.1	Serial1/0/0
10.0.14.0/24	Direct	0	0	D	10.0.14.4	Serial1/0/0
10.0.14.1/32	Direct	0	0	D	10.0.14.1	Serial1/0/0
10.0.14.4/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.14.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.15.0/24	OSPF	10	1563	D	10.0.14.1	Serial1/0/0
127.0.0.0/8	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
127.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
255.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0

Step 4 **Use the route policy to filter routes, avoiding loops.**

Create Loopback 1 with address 10.1.4.4/24 on R4 and run the **import-route direct** command to import routes to OSPF.

[R4]interface LoopBack 1

[R4-LoopBack1]ip address 10.1.4.4 255.255.255.0

[R4-LoopBack1]ospf 1

[R4-ospf-1]import-route direct

Import OSPF routes on R2 to ISIS, and import ISIS routes on R5 to OSPF.

[R2]isis 1

[R2-isis-1]import-route ospf

[R5]ospf

[R5-ospf-1]import-route isis

Test the connectivity between R1 and 10.1.4.4.

```
[R1]ping 10.1.4.4
  PING 10.1.4.4: 56  data bytes, press CTRL_C to break
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out

  --- 10.1.4.4 ping statistics ---
    5 packet(s) transmitted
    0 packet(s) received
    100.00% packet loss
```

The connection is abnormal. Check the routing table on R1.

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 28        Routes : 28
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
        1.0.0.0/8   Static  60   0           D   0.0.0.0         NULL0
        1.1.0.0/16  Static  60   0           D   0.0.0.0         NULL0
        1.1.1.0/24  Static  60   0           D   0.0.0.0         NULL0
        1.1.1.0/25  Static  60   0           D   0.0.0.0         NULL0
        1.1.1.1/32  Static  60   0           D   0.0.0.0         NULL0
       10.0.1.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
       10.0.2.2/32  O_ASE   150  1           D   10.0.15.5       GigabitEthernet0/0/0
```

10.0.3.3/32	O_ASE	150	1	D	10.0.15.5	GigabitEthernet0/0/0
10.0.4.4/32	OSPF	10	1562	D	10.0.14.4	Serial3/0/0
10.0.5.5/32	O_ASE	150	1	D	10.0.15.5	GigabitEthernet0/0/0
10.0.12.0/24	Direct	0	0	D	10.0.12.1	Serial1/0/0
10.0.12.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.12.2/32	Direct	0	0	D	10.0.12.2	Serial1/0/0
10.0.12.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.14.0/24	Direct	0	0	D	10.0.14.1	Serial3/0/0
10.0.14.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.14.4/32	Direct	0	0	D	10.0.14.4	Serial3/0/0
10.0.14.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.15.0/24	Direct	0	0	D	10.0.15.1	GigabitEthernet0/0/0
10.0.15.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.15.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.23.0/24	O_ASE	150	1	D	10.0.15.5	GigabitEthernet0/0/0
10.0.35.0/24	O_ASE	150	1	D	10.0.15.5	GigabitEthernet0/0/0
10.1.4.0/24	O_ASE	150	1	D	10.0.15.5	GigabitEthernet0/0/0
127.0.0.0/8	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
127.255.255.255/32	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 10.1.4.0/24 on R1 is R5.

Check the 10.1.4.0/24 route in the routing tables on R2, R3, and R5.

[R2]display ip routing-table 10.1.4.0

Route Flags: R - relay, D - download to fib

Routing Table : Public

```
Summary Count : 1
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
      10.1.4.0/24  O_ASE   150  1           D   10.0.12.1       Serial1/0/0
[R3]display ip routing-table 10.1.4.0
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Table : Public
Summary Count : 1
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
      10.1.4.0/24   ISIS-L2 15   74          D   10.0.23.2       Serial2/0/0

 [R5]display ip routing-table 10.1.4.0
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Table : Public
Summary Count : 1
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
      10.1.4.0/24   ISIS-L2 15   84          D   10.0.35.3       Serial1/0/0
```

Run the **tracert** command on R1 to check the path to 10.1.4.4.

```
[R1]tracert 10.1.4.4
 traceroute to  10.1.4.4(10.1.4.4), max hops: 30 ,packet length: 40,press CTRL_C to break
 1 10.0.15.5 61 ms  2 ms  2 ms
 2 10.0.35.3 29 ms  28 ms  29 ms
 3 10.0.23.2 31 ms  36 ms  36 ms
 4 10.0.12.1 34 ms  36 ms  36 ms
 5 10.0.15.5 34 ms  37 ms  37 ms
 6 10.0.35.3 55 ms  59 ms  59 ms
```

```
7 10.0.23.2 60 ms   66 ms   66 ms

8 10.0.12.1 63 ms   66 ms   66 ms

9 10.0.15.5 65 ms   67 ms   67 ms
```

A loop is detected.

The reason is: After route import is configured, R5 can learn the 10.1.4.0/24 route in both the ISIS and OSPF domains.

The ISIS routes have higher priority than OSPF external routes, so R5 uses the routes learned from the ISIS domain.

R1 can learn this route from both R5 and R4. The two routes are OSPF external routes. Therefore, the cost values of the two routes are compared. R1 is connected to R5 through GE links, which is better than the serial link between R1 and R4. Therefore, R1 uses the routes learned from R5, causing the loop.

Apply the route policy **policy_r5** to R5 and add tag 100 to the route 10.1.4.0/24.

```
[R5]acl number 2001

[R5-acl-basic-2001]rule 0 permit source 10.1.4.0 0.0.0.255

[R5-acl-basic-2001]route-policy add_tag permit node 10

[R5-route-policy]if-match acl 2001

[R5-route-policy]apply tag 100

[R5-route-policy]route-policy add_tag permit node 20

[R5-route-policy]ospf

[R5-ospf-1]import-route rip route-policy add_tag
```

Check OSPF routing information on R1.

```
[R1]display ospf routing


        OSPF Process 1 with Router ID 10.0.12.1
```

Routing Tables

Routing for Network

Destination	Cost	Type	NextHop	AdvRouter	Area
10.0.1.1/32	0	Stub	10.0.1.1	10.0.12.1	0.0.0.0
10.0.12.0/24	1562	Stub	10.0.12.1	10.0.12.1	0.0.0.0
10.0.14.0/24	1562	Stub	10.0.14.1	10.0.12.1	0.0.0.0
10.0.15.0/24	1	Transit	10.0.15.1	10.0.12.1	0.0.0.0
10.0.4.4/32	1562	Stub	10.0.14.4	10.0.14.4	0.0.0.0

Routing for ASEs

Destination	Cost	Type	Tag	NextHop	AdvRouter
1.1.1.0/24	1	Type2	1	10.0.15.5	10.0.35.5
10.0.2.2/32	1	Type2	1	10.0.15.5	10.0.35.5
10.0.3.3/32	1	Type2	1	10.0.15.5	10.0.35.5
10.0.5.5/32	1	Type2	1	10.0.15.5	10.0.35.5
10.0.14.1/32	1	Type2	1	10.0.15.5	10.0.35.5
10.0.23.0/24	1	Type2	1	10.0.15.5	10.0.35.5
10.0.35.0/24	1	Type2	1	10.0.15.5	10.0.35.5
10.1.4.0/24	1	Type2	100	10.0.15.5	10.0.35.5

Total Nets: 13

Intra Area: 5 Inter Area: 0 ASE: 8 NSSA: 0

R1 has the 10.1.4.0/24 route with tag 100, indicating that the route is obtained from R5.

To address the loop problem, filter out the 10.1.4.0/24 route when R5 imports ISIS routes to OSPF.

Configure the route policy **route_delete** on R5 to control ISIS route import to OSPF.

[R5]route-policy route_delete deny node 10

[R5-route-policy]if-match acl 2001

[R5-route-policy]route-policy route_delete permit node 20

[R5-route-policy]ospf 1

[R5-ospf-1]import-route rip route-policy route_delete

Check the routing table on R1.

[R1]display ip routing-table

Route Flags: R - relay, D - download to fib

Routing Tables: Public

Destinations : 28 Routes : 28

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
1.0.0.0/8	Static	60	0	D	0.0.0.0	NULL0
1.1.0.0/16	Static	60	0	D	0.0.0.0	NULL0
1.1.1.0/24	Static	60	0	D	0.0.0.0	NULL0
1.1.1.0/25	Static	60	0	D	0.0.0.0	NULL0
1.1.1.1/32	Static	60	0	D	0.0.0.0	NULL0
10.0.1.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.2.2/32	O_ASE	150	1	D	10.0.15.5	GigabitEthernet0/0/0
10.0.3.3/32	O_ASE	150	1	D	10.0.15.5	GigabitEthernet0/0/0
10.0.4.4/32	OSPF	10	1562	D	10.0.14.4	Serial3/0/0
10.0.5.5/32	O_ASE	150	1	D	10.0.15.5	GigabitEthernet0/0/0
10.0.12.0/24	Direct	0	0	D	10.0.12.1	Serial1/0/0
10.0.12.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0

```
      10.0.12.2/32   Direct  0    0           D   10.0.12.2       Serial1/0/0
    10.0.12.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.14.0/24   Direct  0    0           D   10.0.14.1       Serial3/0/0
      10.0.14.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.14.4/32   Direct  0    0           D   10.0.14.4       Serial3/0/0
    10.0.14.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.15.0/24   Direct  0    0           D   10.0.15.1   GigabitEthernet0/0/0
      10.0.15.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
    10.0.15.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.23.0/24   O_ASE 150    1           D  10.0.15.5   GigabitEthernet0/0/0
      10.0.35.0/24   O_ASE 150    1           D  10.0.15.5   GigabitEthernet0/0/0
       10.1.4.0/24   O_ASE 150    1           D  10.0.14.4        Serial3/0/0
       127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
```

The routing information is correct.

Test the connectivity between R1 and 10.1.4.4.

```
[R1]ping 10.1.4.4
  PING 10.1.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.1.4.4: bytes=56 Sequence=1 ttl=255 time=33 ms
    Reply from 10.1.4.4: bytes=56 Sequence=2 ttl=255 time=29 ms
    Reply from 10.1.4.4: bytes=56 Sequence=3 ttl=255 time=29 ms
    Reply from 10.1.4.4: bytes=56 Sequence=4 ttl=255 time=29 ms
    Reply from 10.1.4.4: bytes=56 Sequence=5 ttl=255 time=29 ms

  --- 10.1.4.4 ping statistics ---
```

```
5 packet(s) transmitted
5 packet(s) received
0.00% packet loss
round-trip min/avg/max = 29/29/33 ms
```

Test the path from R1 to 10.1.4.4.

```
[R1]tracert 10.1.4.4
 traceroute to  10.1.4.4(10.1.4.4), max hops: 30 ,packet length: 40,press CTRL_C to break
 1 10.0.14.4 61 ms  29 ms  29 ms
```

The loop is removed. R4 is the next hop of the 10.1.4.0/24 route on R1.

Step 5 **Use the route-policy to change route priorities, avoiding loops.**

Check the IP routing table of R5. Observe the next hop of the route 10.1.4.0/24.

```
[R5]display  ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 22        Routes : 22
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       1.1.1.0/24  ISIS-L2  15   84          D   10.0.35.3       Serial1/0/0
      10.0.1.1/32  OSPF     10   1           D   10.0.15.1       GigabitEthernet0/0/0
      10.0.2.2/32  ISIS-L2  15   20          D   10.0.35.3       Serial1/0/0
      10.0.3.3/32  ISIS-L2  15   10          D   10.0.35.3       Serial1/0/0
      10.0.4.4/32  OSPF     10   49          D   10.0.15.1       GigabitEthernet0/0/0
      10.0.5.5/32  Direct   0    0           D   127.0.0.1       LoopBack0
```

10.0.12.0/24	OSPF	10	49	D	10.0.15.1	GigabitEthernet0/0/0
10.0.14.0/24	OSPF	10	49	D	10.0.15.1	GigabitEthernet0/0/0
10.0.14.1/32	ISIS-L2	15	84	D	10.0.35.3	Serial1/0/0
10.0.15.0/24	Direct	0	0	D	10.0.15.5	GigabitEthernet0/0/0
10.0.15.5/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/0
10.0.15.255/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/0
10.0.23.0/24	ISIS-L2	15	20	D	10.0.35.3	Serial1/0/0
10.0.35.0/24	Direct	0	0	D	10.0.35.5	Serial1/0/0
10.0.35.3/32	Direct	0	0	D	10.0.35.3	Serial1/0/0
10.0.35.5/32	Direct	0	0	D	127.0.0.1	Serial1/0/0
10.0.35.255/32	Direct	0	0	D	127.0.0.1	Serial1/0/0
10.1.4.0/24	ISIS-L2	15	84	D	10.0.35.3	Serial1/0/0
127.0.0.0/8	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
127.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
255.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0

Test the path from R5 to 10.1.4.4.

```
[R5]tracert 10.1.4.4
 traceroute to  10.1.4.4(10.1.4.4), max hops: 30 ,packet length: 40,press CTRL_C to break
 1 10.0.35.3 62 ms  24 ms  24 ms
 2 10.0.23.2 43 ms  44 ms  44 ms
 3 10.0.12.1 33 ms  33 ms  33 ms
 4 10.0.14.4 74 ms  55 ms  55 ms
```

Although route filtering can fix the loop problem, R5 still learns the 10.1.4.0/24 route from the ISIS domain. The sub-optimal route problem still exists.

To fix both the loop problem and sub-optimal route problem, R5 must learn the 10.1.4.0/24 route from the OSPF domain.

Delete the policy **route_delete** on R5.

[R5]undo route-policy route_delete

Configure the route-policy **route_pref** on R5. Change the priority of the 10.1.4.0/24 route to 180, which is smaller than the priority of OSPF external routes.

[R5]route-policy route_pref permit node 10

[R5-route-policy]if-match acl 2001

[R5-route-policy]apply preference 180

Use the route-policy **route_pref** to control the ISIS routes imported into OSPF.

[R5]isis

[R5-isis-1]preference route-policy route_pref

Check the IP routing tables of R5 and R1. Observe the next hops of the 10.1.4.0/24 routes.

[R5]display ip routing-table 10.1.4.0

Route Flags: R - relay, D - download to fib

Routing Table : Public

Summary Count : 1

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
10.1.4.0/24	O_ASE	150	1	D	10.0.15.1	GigabitEthernet0/0/0

[R1]display ip routing-table 10.1.4.0

Route Flags: R - relay, D - download to fib

```
------------------------------------------------------------------------------
Routing Table : Public
Summary Count : 1
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
       10.1.4.0/24  O_ASE   150  1           D   10.0.14.4       Serial3/0/0
```

Test the connectivity between R1 and 10.1.4.4.

```
[R1]ping 10.1.4.4
  PING 10.1.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.1.4.4: bytes=56 Sequence=1 ttl=255 time=39 ms
    Reply from 10.1.4.4: bytes=56 Sequence=2 ttl=255 time=35 ms
    Reply from 10.1.4.4: bytes=56 Sequence=3 ttl=255 time=35 ms
    Reply from 10.1.4.4: bytes=56 Sequence=4 ttl=255 time=35 ms
    Reply from 10.1.4.4: bytes=56 Sequence=5 ttl=255 time=35 ms

  --- 10.1.4.4 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 35/35/39 ms
```

Test the path from R1 to 10.1.4.4.

```
[R1]tracert 10.1.4.4
 traceroute to  10.1.4.4(10.1.4.4), max hops: 30 ,packet length: 40,press CTRL_C to break
 1 10.0.14.4 61 ms  25 ms  25 ms
```

Test the path from R5 to 10.1.4.4.

```
[R5]tracert 10.1.4.4
```

```
traceroute to 10.1.4.4(10.1.4.4), max hops: 30 ,packet length: 40,press CTRL_C to break
1 10.0.15.1 61 ms 2 ms 2 ms
2 10.0.14.4 41 ms 28 ms 27 ms
```

The loop problem is fixed.

R4 is the next hop of the route 10.1.4.0/24 on R1. R1 is the next hop of the route 10.1.4.0/24 on R5. The sub-optimal route problem is also fixed.

----End

Additional Exercises: Analysis and Verification

Can you use an ACL to achieve the same effect as that in step 3? What is the difference between using an ACL and a prefix list?

In the R3's routing table in step 5, why the 10.0.15.0/24 route has two next hops but the 10.0.12.0/24 route has only one next hop?

Device Configurations

```
<R1>display current-configuration
[V200R007C00SPC600]
#
 sysname R1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
```

```
#
interface GigabitEthernet0/0/0
 ip address 10.0.15.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
#
ospf 1
 import-route static route-policy policy_r1
 area 0.0.0.0
  network 10.0.12.1 0.0.0.0
  network 10.0.15.1 0.0.0.0
  network 10.0.14.1 0.0.0.0
  network 10.0.1.1 0.0.0.0
#
route-policy policy_r1 permit node 10
 if-match ip-prefix pref_r1
#
 ip ip-prefix pref_r1 index 10 permit 1.1.1.0 24 greater-equal 24 less-equal 24
#
 ip route-static 1.0.0.0 255.0.0.0 NULL0
 ip route-static 1.1.0.0 255.255.0.0 NULL0
 ip route-static 1.1.1.0 255.255.255.0 NULL0
 ip route-static 1.1.1.0 255.255.255.128 NULL0
 ip route-static 1.1.1.1 255.255.255.255 NULL0
#
return
```

```
<R2>display current-configuration
```

```
[V200R007C00SPC600]
#
 sysname R2
#
isis 1
 is-level level-2
 network-entity 49.0001.0000.0000.0002.00
 import-route ospf 1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
isis enable 1
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.255
isis enable 1
#
ospf 1
 area 0.0.0.0
  network 10.0.12.2 0.0.0.0
#
return
```

<R3>**display current-configuration**

```
[V200R007C00SPC600]
#
 sysname R3
#
isis 1
 is-level level-2
 network-entity 49.0001.0000.0000.0003.00
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
isis enable 1
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
isis enable 1
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
isis enable 1
```

```
<R4>display current-configuration
[V200R007C00SPC600]
#
 sysname R4
#
 icmp port-unreachable send
```

```
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
#
interface LoopBack1
 ip address 10.1.4.4 255.255.255.0
#
ospf 1
 import-route direct
 area 0.0.0.0
  network 10.0.14.4 0.0.0.0
  network 10.0.4.4 0.0.0.0
#
return
```

```
<R5>display current-configuration
[V200R007C00SPC600]
#
 sysname R5
#
isis 1
 is-level level-2
 network-entity 49.0001.0000.0000.0005.00
 preference route-policy route_pref
#
```

```
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
isis enable 1

#
interface GigabitEthernet0/0/0
 ip address 10.0.15.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.255
  isis enable 1
#
ospf 1
 import-route rip 1 route-policy route_delete
 area 0.0.0.0
  network 10.0.15.5 0.0.0.0
#
route-policy add_tag permit node 10
 if-match acl 2001
 apply tag 100
#
route-policy add_tag permit node 20
#
route-policy route_pref permit node 10
 if-match acl 2001
 apply preference 180
#
return
```

Chapter 6 VLAN Features and Configurations

Lab 6-1 VLAN Configurations

Learning Objectives

The objectives of this lab are to learn and understand:

- How to configure VLANs
- How to configure Eth-Trunk

Topology

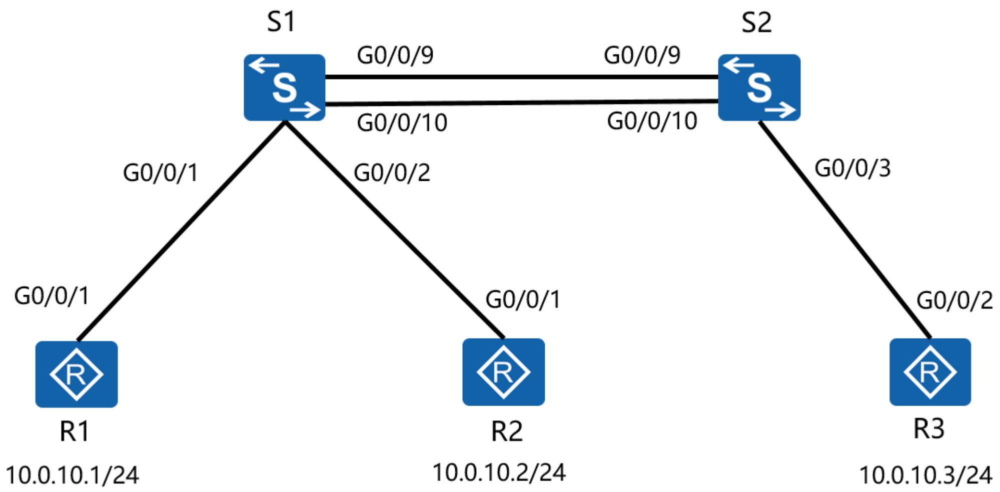

Figure 6-1 VLAN configuration

Scenario

You are a network administrator of a company. The company's network is an Ethernet that has two switches. In the preceding figure, the routers simulate the computers, and R3 is a server. To optimize the network, you need to improve the link speed and reliability between S1 and S2. Two VLANs are required to isolate broadcast storms. R2 and R3 are on the same VLAN. Ensure that R1 can access R3.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

Configure IP addresses and masks for all devices.

```
<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R1
[R1]interface GigabitEthernet 0/0/1
[R1-GigabitEthernet0/0/1]ip address 10.0.10.1 24
[R1-GigabitEthernet0/0/1]quit

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R2
[R2]interface GigabitEthernet 0/0/1
[R2-GigabitEthernet0/0/1]ip address 10.0.10.2 24
[R2-GigabitEthernet0/0/1]quit

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R3
[R3]interface GigabitEthernet 0/0/2
[R3-GigabitEthernet0/0/2]ip address 10.0.10.3 24
[R3-GigabitEthernet0/0/2]quit
```

Set names for switches.

```
<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
```

[Huawei]sysname S1

<Huawei>system-view

Enter system view, return user view with Ctrl+Z.

[Huawei]sysname S2

Step 2 **Configure an Eth-Trunk.**

You can bind two or more links into an Eth-Trunk to improve link bandwidth and reliability. Add G0/0/9 and G0/0/10 and S1 and S2 to an Eth-Trunk group.

Create an Eth-Trunk.

[S1]interface Eth-Trunk 1

[S1-Eth-Trunk1]

[S2]interface Eth-Trunk 1

[S2-Eth-Trunk1]

Set the working mode of the Eth-Trunk to LACP.

[S1-Eth-Trunk1]mode lacp

[S1-Eth-Trunk1]quit

[S2-Eth-Trunk1]mode lacp

[S2-Eth-Trunk1]quit

Add G0/0/9 and G0/0/10 of S1 and S2 to an Eth-Trunk.

[S1]interface GigabitEthernet 0/0/9

[S1-GigabitEthernet0/0/9]eth-trunk 1

[S1-GigabitEthernet0/0/9]quit

```
[S1]interface GigabitEthernet 0/0/10
[S1-GigabitEthernet0/0/10]eth-trunk 1
[S1-GigabitEthernet0/0/10]quit

[S2]interface GigabitEthernet 0/0/9
[S2-GigabitEthernet0/0/9]eth-trunk 1
[S2-GigabitEthernet0/0/9]quit
[S2]interface GigabitEthernet 0/0/10
[S2-GigabitEthernet0/0/10]eth-trunk 1
[S2-GigabitEthernet0/0/10]quit
```

Run the **display eth-trunk** command to check configurations.

```
[S1]display eth-trunk
Eth-Trunk1's state information is:
Local:
LAG ID: 1                    WorkingMode: LACP
Preempt Delay: Disabled      Hash arithmetic: According to SIP-XOR-DIP
System Priority: 32768       System ID: d0d0-4ba6-aab0
Least Active-linknumber: 1   Max Active-linknumber: 8
Operate status: up           Number Of Up Port In Trunk: 2
--------------------------------------------------------------------------------
ActorPortName          Status   PortType PortPri PortNo PortKey PortState Weight
GigabitEthernet0/0/9   Selected 1GE      32768   1      305     10111100  1
GigabitEthernet0/0/10  Selected 1GE      32768   2      305     10111100  1

Partner:
--------------------------------------------------------------------------------
ActorPortName          SysPri   SystemID        PortPri PortNo PortKey PortState
```

```
GigabitEthernet0/0/9    32768    d0d0-4ba6-ac20   32768    1     305     10111100
GigabitEthernet0/0/10   32768    d0d0-4ba6-ac20   32768    2     305     10111100
```

The command output shows that the Eth-Trunk working mode is LACP, and the threshold of active interfaces is 8. G0/0/9 and G0/0/10 are active.

Change the threshold of active interfaces.

```
[S1]interface Eth-Trunk 1
[S1-Eth-Trunk1]max active-linknumber 1
[S1-Eth-Trunk1]quit

[S2]interface Eth-Trunk 1
[S2-Eth-Trunk1]max active-linknumber 1
[S2-Eth-Trunk1]quit
```

Check the Eth-Trunk configurations.

```
[S1]display eth-trunk 1
Eth-Trunk1's state information is:
Local:
LAG ID: 1                        WorkingMode: LACP
Preempt Delay: Disabled          Hash arithmetic: According to SIP-XOR-DIP
System Priority: 32768           System ID: d0d0-4ba6-aab0
Least Active-linknumber: 1       Max Active-linknumber: 1
Operate status: up               Number Of Up Port In Trunk: 1
--------------------------------------------------------------------------------
ActorPortName            Status     PortType PortPri PortNo PortKey PortState Weight
GigabitEthernet0/0/9     Selected 1GE      32768    1      305     10111100   1
GigabitEthernet0/0/10    Unselect 1GE      32768    2      305     10100000   1
```

```
Partner:
--------------------------------------------------------------------------
ActorPortName          SysPri   SystemID         PortPri PortNo PortKey PortState
GigabitEthernet0/0/9   32768    d0d0-4ba6-ac20   32768   1      305     10111100
GigabitEthernet0/0/10  32768    d0d0-4ba6-ac20   32768   2      305     10100000
```

The status of G0/0/10 changes to Unselect. One link in the Eth-Trunk transmits data, and the other link is the backup. Network reliability is improved.

Shut down G0/0/9 of S1 to verify link backup.

```
[S1]interface GigabitEthernet 0/0/9
[S1-GigabitEthernet0/0/9]shutdown
[S1-GigabitEthernet0/0/9]quit
```

Check the Eth-Trunk information.

```
[S1]display eth-trunk 1
Eth-Trunk1's state information is:
Local:
LAG ID: 1                      WorkingMode: LACP
Preempt Delay: Disabled        Hash arithmetic: According to SIP-XOR-DIP
System Priority: 32768         System ID: d0d0-4ba6-aab0
Least Active-linknumber: 1     Max Active-linknumber: 1
Operate status: up             Number Of Up Port In Trunk: 1
------------------------------------------------------------------------------
ActorPortName          Status   PortType PortPri PortNo PortKey PortState Weight
GigabitEthernet0/0/9   Unselect 1GE      32768   1      305     10100010  1
GigabitEthernet0/0/10  Selected 1GE      32768   2      305     10111100  1

Partner:
```

```
----------------------------------------------------------------------------
ActorPortName          SysPri   SystemID        PortPri PortNo PortKey PortState
GigabitEthernet0/0/9   0        0000-0000-0000  0       0      0       10100011
GigabitEthernet0/0/10  32768    d0d0-4ba6-ac20  32768   2      305     10111100
```

The command output shows that the status of G0/0/9 in the Eth-Trunk changes to Unselect, and the status of G0/0/10 changes from Unselect to Selected and G0/0/10 resumes data forwarding. Link backup is successful.

Device Configurations

```
[S1]display current-configuration
!Software Version V200R008C00SPC500
#
sysname SW1
#
diffserv domain default
#
drop-profile default
#
aaa
 authentication-scheme default
 authorization-scheme default
 accounting-scheme default
 domain default
 domain default_admin
 local-user admin password
irreversible-cipher %^%#tK;J&jw0HG8<9-"zX!kHwzXRNjuXn96[vN47F$*L~pXcROEP3!>c)NV+:`i;%^%#
 local-user admin service-type http
#
```

```
interface Vlanif1
#
interface MEth0/0/1
#
interface Eth-Trunk1
 mode lacp
 max active-linknumber 1
#
interface GigabitEthernet0/0/1
#
interface GigabitEthernet0/0/2
#
interface GigabitEthernet0/0/3
#
interface GigabitEthernet0/0/4
#
interface GigabitEthernet0/0/5
#
interface GigabitEthernet0/0/6
#
interface GigabitEthernet0/0/7
#
interface GigabitEthernet0/0/8
#
interface GigabitEthernet0/0/9
 shutdown
 eth-trunk 1
#
interface GigabitEthernet0/0/10
```

```
 eth-trunk 1
#
interface GigabitEthernet0/0/11
#
interface GigabitEthernet0/0/12
#
interface GigabitEthernet0/0/13
#
interface GigabitEthernet0/0/14
#
interface GigabitEthernet0/0/15
#
interface GigabitEthernet0/0/16
#
interface GigabitEthernet0/0/17
#
interface GigabitEthernet0/0/18
#
interface GigabitEthernet0/0/19
#
interface GigabitEthernet0/0/20
#
interface GigabitEthernet0/0/21
#
interface GigabitEthernet0/0/22
#
interface GigabitEthernet0/0/23
#
interface GigabitEthernet0/0/24
```

```
#
interface GigabitEthernet0/0/25
#
interface GigabitEthernet0/0/26
#
interface GigabitEthernet0/0/27
#
interface GigabitEthernet0/0/28
#
interface XGigabitEthernet0/0/1
#
interface XGigabitEthernet0/0/2
#
interface XGigabitEthernet0/0/3
#
interface XGigabitEthernet0/0/4
#
interface NULL0
#
user-interface con 0
 authentication-mode password
 set authentication password cipher $1a$fcjGHMtb0U$^GKZ+`,g@DfG$:T/P,R~iJ&')|!O":$b4)0*~&c-$
 idle-timeout 0 0
user-interface vty 0 4
user-interface vty 16 20
#
return

[S2]display current-configuration
```

```
!Software Version V200R008C00SPC500
#
sysname SW2
#
diffserv domain default
#
drop-profile default
#
aaa
 authentication-scheme default
 authorization-scheme default
 accounting-scheme default
 domain default
 domain default_admin
 local-user admin password
irreversible-cipher %^%#gI/bO8qF$HkpAPUgNd'GiYR4TC!>EK#oG("Wl4_#$G*OKo-'7*R[h3+49<Z2%^%#
 local-user admin service-type http
#
interface Vlanif1
#
interface MEth0/0/1
#
interface Eth-Trunk1
 mode lacp
 max active-linknumber 1
#
interface GigabitEthernet0/0/1
#
interface GigabitEthernet0/02
```

```
#
interface GigabitEthernet0/0/3
#
interface GigabitEthernet0/0/4
#
interface GigabitEthernet0/0/5
#
interface GigabitEthernet0/0/6
#
interface GigabitEthernet0/0/7
#
interface GigabitEthernet0/0/8
#
interface GigabitEthernet0/0/9
 eth-trunk 1
#
interface GigabitEthernet0/0/10
 eth-trunk 1
#
interface GigabitEthernet0/0/11
#
interface GigabitEthernet0/0/12
#
interface GigabitEthernet0/0/13
#
interface GigabitEthernet0/0/14
#
interface GigabitEthernet0/0/15
#
```

```
interface GigabitEthernet0/0/16
#
interface GigabitEthernet0/0/17
#
interface GigabitEthernet0/0/18
#
interface GigabitEthernet0/0/19
#
interface GigabitEthernet0/0/20
#
interface GigabitEthernet0/0/21
#
interface GigabitEthernet0/0/22
#
interface GigabitEthernet0/0/23
#
interface GigabitEthernet0/0/24
#
interface GigabitEthernet0/0/25
#
interface GigabitEthernet0/0/26
#
interface GigabitEthernet0/0/27
#
interface GigabitEthernet0/0/28
#
interface XGigabitEthernet0/0/1
#
interface XGigabitEthernet0/0/2
```

```
#
interface XGigabitEthernet0/0/3
#
interface XGigabitEthernet0/0/4
#
interface NULL0
#
user-interface con 0
 authentication-mode password
 set authentication password cipher $1a$5"l`L7$/5T$,KFQ9dEy~'IggWOa7V(C+9fQOd*M;U6q,.Sl1y'H$
 idle-timeout 0 0
user-interface vty 0 4
user-interface vty 16 20
#
Return
```

Lab 6-2 MUX VLAN

Learning Objectives

The objectives of this lab are to learn and understand:

- How to configure MUX VLAN

Topology

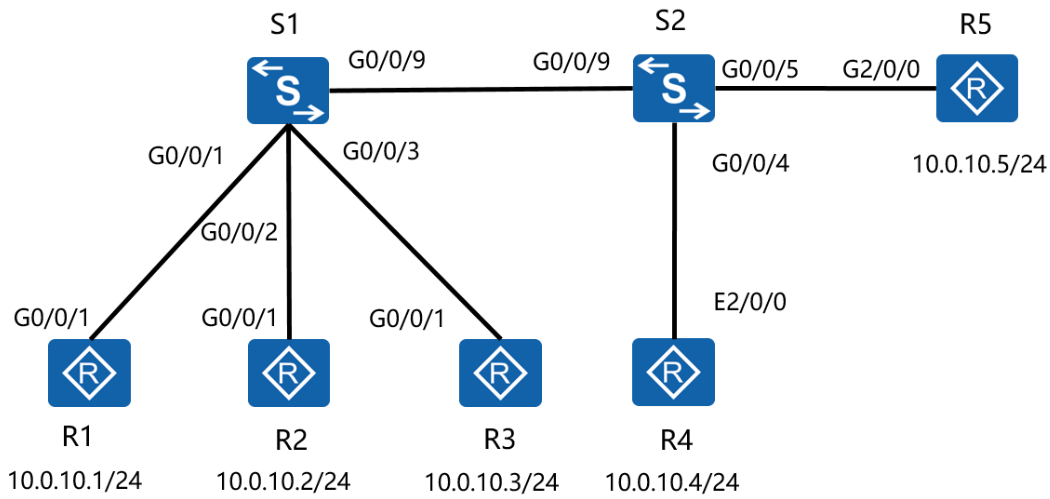

Figure 6-2 MUX VLAN configuration

Scenario

You are a network administrator of a company. The company's network is an Ethernet that has two switches. In the preceding figure, routers simulate the computers. To optimize the network, you need to isolate the broadcast domains. R1 and R2 are on the same VLAN, and R3 and R4 are on another VLAN. The company requires that all PCs can access R5, and R3 and R4 cannot communicate with R1 and R2 or access each other. In the future network plan, S2's G0/0/24 will be connected to voice devices. Therefore, plan the voice VLAN and related configurations.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

Configure IP addresses and masks for all devices.

```
<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R1
[R1]interface g0/0/1
[R1-GigabitEthernet0/0/1]ip address 10.0.10.1 24
[R1-GigabitEthernet0/0/1]quit

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R2
[R2]interface g0/0/1
[R2-GigabitEthernet0/0/1]ip address 10.0.10.2 24
[R2-GigabitEthernet0/0/1]quit

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R3
[R3]interface g0/0/1
[R3-GigabitEthernet0/0/1]ip address 10.0.10.3 24
[R3-GigabitEthernet0/0/1]quit

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R4
[R4]interface Ethernet2/0/0
[R4-Ethernet2/0/0]ip address 10.0.10.4 24
[R4-GigabitEthernet2/0/0]quit
```

```
<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R5
[R5]interface Ethernet2/0/0
[R5-Ethernet2/0/0]ip address 10.0.10.5 24
[R1-GigabitEthernet0/0/1]quit

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname S1
[S1]

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname S2
[S2]
```

Test the connectivity between R1 and R2, R3, R4, as well as R5.

```
[R1]ping -c 1 10.0.10.2
  PING 10.0.10.2: 56   data bytes, press CTRL_C to break
    Reply from 10.0.10.2: bytes=56 Sequence=1 ttl=255 time=14 ms

  --- 10.0.10.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 14/14/14 ms
```

```
[R1]ping -c 1 10.0.10.3
  PING 10.0.10.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.10.3: bytes=56 Sequence=1 ttl=255 time=5 ms

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

[R1]ping -c 1 10.0.10.4
  PING 10.0.10.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.10.4: bytes=56 Sequence=1 ttl=255 time=15 ms

  --- 10.0.10.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 15/15/15 ms

[R1]ping -c 1 10.0.10.5
  PING 10.0.10.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.10.5: bytes=56 Sequence=1 ttl=255 time=6 ms

  --- 10.0.10.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 6/6/6 ms
```

Step 2 **Configure MUX VLAN.**

After the devices on the same network segment are added to different VLANs, MUX VLAN forbids Layer 2 communication between them and allows them to communicate with the specified VLAN. In addition, MUX VLAN forbids communication between the devices in the same VLAN.

Configure VLAN 100 as the principle VLAN of MUX VLAN, and VLAN 10 as well as VLAN 20 as subordinate VLANs.

Set the interface types between PCs and switches. Allow all PCs to communicate with R4. Prevent R3 and R4 from communicating with other VLANs or accessing each other.

Configure VLAN 100 as the principle VLAN and configure subordinate VLANs.

```
[S1]vlan 10
[S1-vlan10]quit
[S1]vlan 20
[S1-vlan20]quit
[S1]vlan 100
[S1-vlan100]mux-vlan
[S1-vlan100]subordinate group 10
[S1-vlan100]subordinate separate 20
[S1-vlan100]quit

[S2]vlan 10
[S2-vlan10]quit
[S2]vlan 20
[S2-vlan20]quit
[S2]vlan 100
[S2-vlan100]mux-vlan
```

[S2-vlan100]subordinate group 10

[S2-vlan100]subordinate separate 20

[S2-vlan100]quit

Add G0/0/5 between R5 and S2 to VLAN 100 and enable MUX VLAN.

[S2]interface GigabitEthernet 0/0/5

[S2-GigabitEthernet0/0/5]port link-type access

[S2-GigabitEthernet0/0/5]port default vlan 100

[S2-GigabitEthernet0/0/5]port mux-vlan enable vlan 100

[S2-GigabitEthernet0/0/5]quit

Add G0/0/1 between R1 and S1 and G0/0/2 between R2 and S1 to VLAN 10, and enable MUX VLAN.

[S1]interface GigabitEthernet 0/0/1

[S1-GigabitEthernet0/0/1]port link-type access

[S1-GigabitEthernet0/0/1]port default vlan 10

[S1-GigabitEthernet0/0/1]port mux-vlan enable vlan 10

[S1-GigabitEthernet0/0/1]quit

[S1]interface GigabitEthernet 0/0/2

[S1-GigabitEthernet0/0/2]port link-type access

[S1-GigabitEthernet0/0/2]port default vlan 10

[S1-GigabitEthernet0/0/2]port mux-vlan enable vlan 10

[S1-GigabitEthernet0/0/2]quit

Add G0/0/3 between R3 and S1 and G0/0/4 between R4 and S2 to VLAN 20, and enable MUX VLAN.

[S1]interface GigabitEthernet 0/0/3

```
[S1-GigabitEthernet0/0/3]port link-type access
[S1-GigabitEthernet0/0/3]port default vlan 20
[S1-GigabitEthernet0/0/3]port mux-vlan enable vlan 20
[S1-GigabitEthernet0/0/3]quit

[S2]interface GigabitEthernet 0/0/4
[S2-GigabitEthernet0/0/4]port link-type access
[S2-GigabitEthernet0/0/4]port default vlan 20
[S2-GigabitEthernet0/0/4]port mux-vlan enable vlan 20
[S2-GigabitEthernet0/0/4]quit
```

Run the **display mux-vlan** command to display MUX VLAN information.

```
[S1]display mux-vlan
Principal Subordinate Type       Interface
--------------------------------------------------------------------------
100       -          principal
100       20         separate    GigabitEthernet0/0/3
100       10         group       GigabitEthernet0/0/1 GigabitEthernet0/0/2

[S2]display mux-vlan
Principal Subordinate Type       Interface
--------------------------------------------------------------------------
100       -          principal   GigabitEthernet0/0/5
100       20         separate    GigabitEthernet0/0/4
100       10         group
--------------------------------------------------------------------------
```

Ping R2, R3, R4, and R5 from R1.

```
[R1]ping -c 1 10.0.10.2
  PING 10.0.10.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.10.2: bytes=56 Sequence=1 ttl=255 time=3 ms

  --- 10.0.10.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 3/3/3 ms

[R1]ping -c 1 10.0.10.3
  PING 10.0.10.3: 56  data bytes, press CTRL_C to break
    Request time out

  --- 10.0.10.3 ping statistics ---
    1 packet(s) transmitted
    0 packet(s) received
    100.00% packet loss

[R1]ping -c 1 10.0.10.4
  PING 10.0.10.4: 56  data bytes, press CTRL_C to break
    Request time out

  --- 10.0.10.4 ping statistics ---
    1 packet(s) transmitted
    0 packet(s) received
    100.00% packet loss

[R1]ping -c 1 10.0.10.5
```

```
  PING 10.0.10.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.10.5: bytes=56 Sequence=1 ttl=255 time=3 ms

  --- 10.0.10.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 3/3/3 ms
```

Ping R2, R4, and R5 from R3.

```
[R3]ping -c 1 10.0.10.2
  PING 10.0.10.2: 56  data bytes, press CTRL_C to break
    Request time out

  --- 10.0.10.2 ping statistics ---
    1 packet(s) transmitted
    0 packet(s) received
    100.00% packet loss

[R3]ping -c 1 10.0.10.4
  PING 10.0.10.4: 56  data bytes, press CTRL_C to break
    Request time out

  --- 10.0.10.4 ping statistics ---
    1 packet(s) transmitted
    0 packet(s) received
    100.00% packet loss
```

```
[R3]ping -c 1 10.0.10.5
  PING 10.0.10.5: 56   data bytes, press CTRL_C to break
    Reply from 10.0.10.5: bytes=56 Sequence=1 ttl=255 time=3 ms

  --- 10.0.10.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 3/3/3 ms
```

The ping command output shows that R1 and R2 in VLAN 10 of MUX VLAN can communicate with R5, and also access each other. R3 and R4 in VLAN 20 can only communicate with R5.

----End

Additional Exercises: Analysis and Verification

Can the users belonging to different MUX VLANs communicate with each other?

Device Configurations

```
[S1]display current-configuration
!Software Version V200R008C00SPC500
#
sysname S1
#
vlan batch 10 20 100
#
diffserv domain default
#
drop-profile default
```

```
#
vlan 100
 mux-vlan
 subordinate separate 20
 subordinate group 10
#
aaa
 authentication-scheme default
 authorization-scheme default
 accounting-scheme default
 domain default
 domain default_admin
 local-user admin password
irreversible-cipher %^%#tK;J&jw0HG8<9-"zX!kHwzXRNjuXn96[vN47F$*L~pXcROEP3!>c)NV+:`i;%^%#
 local-user admin service-type http
#
interface Vlanif1
#
interface MEth0/0/1
#
interface GigabitEthernet0/0/1
 port link-type access
 port default vlan 10
 port mux-vlan enable vlan 10
#
interface GigabitEthernet0/0/2
 port link-type access
 port default vlan 10
 port mux-vlan enable vlan 10
```

```
#
interface GigabitEthernet0/0/3
 port link-type access
 port default vlan 20
 port mux-vlan enable vlan 20
#
interface GigabitEthernet0/0/4
#
interface GigabitEthernet0/0/5
#
interface GigabitEthernet0/0/6
#
interface GigabitEthernet0/0/7
#
interface GigabitEthernet0/0/8
#
interface GigabitEthernet0/0/9
#
interface GigabitEthernet0/0/10
#
interface GigabitEthernet0/0/11
#
interface GigabitEthernet0/0/12
#
interface GigabitEthernet0/0/13
#
interface GigabitEthernet0/0/14
#
interface GigabitEthernet0/0/15
```

```
#
interface GigabitEthernet0/0/16
#
interface GigabitEthernet0/0/17
#
interface GigabitEthernet0/0/18
#
interface GigabitEthernet0/0/19
#
interface GigabitEthernet0/0/20
#
interface GigabitEthernet0/0/21
#
interface GigabitEthernet0/0/22
#
interface GigabitEthernet0/0/23
#
interface GigabitEthernet0/0/24
#
interface GigabitEthernet0/0/25
#
interface GigabitEthernet0/0/26
#
interface GigabitEthernet0/0/27
#
interface GigabitEthernet0/0/28
#
interface XGigabitEthernet0/0/1
#
```

```
interface XGigabitEthernet0/0/2
#
interface XGigabitEthernet0/0/3
#
interface XGigabitEthernet0/0/4
#
interface NULL0
#
user-interface con 0
 authentication-mode password
 set authentication password cipher $1a$fcjGHMtb0U$^GKZ+`,g@DfG$:T/P,R~iJ&')|!O":$b4)0*~&c-$
 idle-timeout 0 0
user-interface vty 0 4
user-interface vty 16 20
#
return

[S2]display current-configuration
!Software Version V200R008C00SPC500
#
sysname S2
#
vlan batch 10 20 100
#
diffserv domain default
#
drop-profile default
#
vlan 100
```

```
 mux-vlan
 subordinate separate 20
 subordinate group 10
#
aaa
 authentication-scheme default
 authorization-scheme default
 accounting-scheme default
 domain default
 domain default_admin
 local-user admin password
irreversible-cipher %^%#gI/bO8qF$HkpAPUgNd'GiYR4TC!>EK#oG("Wl4_#$G*OKo-'7*R[h3+49<Z2%^%#
 local-user admin service-type http
#
interface Vlanif1
#
interface MEth0/0/1
#
interface GigabitEthernet0/0/1
#
interface GigabitEthernet0/0/2
#
interface GigabitEthernet0/0/3
#
interface GigabitEthernet0/0/4
 port link-type access
 port default vlan 20
 port mux-vlan enable vlan 20
#
```

```
interface GigabitEthernet0/0/5
 port link-type access
 port default vlan 100
 port mux-vlan enable vlan 100
#
interface GigabitEthernet0/0/6
#
interface GigabitEthernet0/0/7
#
interface GigabitEthernet0/0/8
#
interface GigabitEthernet0/0/9
#
interface GigabitEthernet0/0/10
#
interface GigabitEthernet0/0/11
#
interface GigabitEthernet0/0/12
#
interface GigabitEthernet0/0/13
#
interface GigabitEthernet0/0/14
#
interface GigabitEthernet0/0/15
#
interface GigabitEthernet0/0/16
#
interface GigabitEthernet0/0/17
#
```

```
interface GigabitEthernet0/0/18
#
interface GigabitEthernet0/0/19
#
interface GigabitEthernet0/0/20
#
interface GigabitEthernet0/0/21
#
interface GigabitEthernet0/0/22
#
interface GigabitEthernet0/0/23
#
interface GigabitEthernet0/0/24
#
interface GigabitEthernet0/0/25
#
interface GigabitEthernet0/0/26
#
interface GigabitEthernet0/0/27
#
interface GigabitEthernet0/0/28
#
interface XGigabitEthernet0/0/1
#
interface XGigabitEthernet0/0/2
#
interface XGigabitEthernet0/0/3
#
interface XGigabitEthernet0/0/4
```

```
#
interface NULL0
#
user-interface con 0
 authentication-mode password
 set authentication password cipher $1a$5"l`L7$/5T$,KFQ9dEy~'IggWOa7V(C+9fQOd*M;U6q,.Sl1y'H$
 idle-timeout 0 0
user-interface vty 0 4
user-interface vty 16 20
#
```

Lab 6-3 Inter-VLAN Communication

Learning Objectives

The objectives of this lab are to learn and understand:

- How to configure multi-armed route (connect switches to independent physical interfaces of a router)
- How to configure router-on-a-stick (connect switches to subinterfaces belonging to the same physical interface on a router)
- Method of implementing inter-VLAN communication
- How to configure VLAN aggregation

Topology

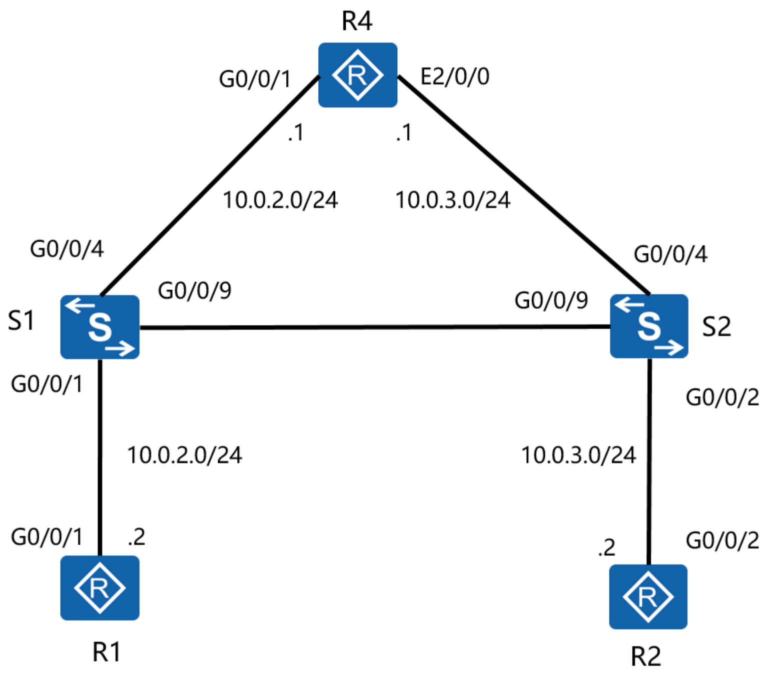

Figure 6-3 Inter-VLAN communication

Scenario

You are a network administrator of a company. The company's network is an Ethernet with two switches and one router. In the preceding figure, R1 and R2 represent PCs of different departments, and are added to two VLANs respectively. R1 and R2 need to communicate with each other. The company used the multi-armed method. That is, the switches are connected to different physical interfaces of the router. To conserve interfaces, the company decided to use the router-on-a-stick method.

Then, due to the network structure change, more traffic is transmitted between VLANs. Therefore, the company required multi-level switching. To facilitate network management, VLAN aggregation is needed.

Tasks

Step 1 **Set basic parameters and configure IP addresses.**

Configure IP addresses and masks for all devices.

```
<huawei>system-view
Enter system view, return user view with Ctrl+Z.
[huawei]sysname R1
[R1]interface GigabitEthernet 0/0/1
[R1-GigabitEthernet0/0/1]ip address 10.0.2.2 24
[R1-GigabitEthernet0/0/1]quit

<huawei>system-view
Enter system view, return user view with Ctrl+Z.
[huawei]sysname R2
[R2]interface GigabitEthernet 0/0/2
[R2-GigabitEthernet0/0/2]ip address 10.0.3.2 24
[R2-GigabitEthernet0/0/2]quit

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname S1

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname S2

<huawei>system-view
Enter system view, return user view with Ctrl+Z.
[huawei]sysname R4
[R4]interface GigabitEthernet 0/0/1
[R4-GigabitEthernet0/0/1]ip address 10.0.2.1 24
[R4-GigabitEthernet0/0/1]quit
```

```
[R4]interface Ethernet2/0/0
[R4-Ethernet2/0/0]ip address 10.0.3.1 24
[R4-Ethernet2/0/0]quit
```

Ping R4's G0/0/1 from R1's G0/0/1.

```
[R1]ping -c 1 10.0.2.1
  PING 10.0.2.1: 56  data bytes, press CTRL_C to break
    Reply from 10.0.2.1: bytes=56 Sequence=1 ttl=255 time=4 ms

  --- 10.0.2.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 4/4/4 ms
```

Ping R4 from R2.

```
[R2]ping -c 1 10.0.3.1
  PING 10.0.3.1: 56  data bytes, press CTRL_C to break
    Reply from 10.0.3.1: bytes=56 Sequence=1 ttl=255 time=3 ms

  --- 10.0.3.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 3/3/3 ms
```

Step 2 **Configure multi-armed route.**

R1 and R2 belong to different VLANs.

The R1's gateway address is the address of R4's G0/0/1, and the R2's gateway address is the address of R4's G0/0/2.

R4 provides multiple physical interfaces to implement inter-VLAN communication. This is multi-armed route.

Create VLAN 2 and VLAN 3 on S1 and S2.

```
[S1]vlan batch 2 3

Info: This operation may take a few seconds. Please wait for a moment...done.


[S2]vlan batch 2 3

Info: This operation may take a few seconds. Please wait for a moment...done.
```

Add R1 to VLAN 2, R2 to VLAN 3, R4's G0/0/1 to VLAN 2, and R4's G0/0/2 to VLAN 3.

```
[S1]interface GigabitEthernet 0/0/1

[S1-GigabitEthernet0/0/1]port link-type access

[S1-GigabitEthernet0/0/1]port default vlan 2

[S1-GigabitEthernet0/0/1]quit

[S1]interface GigabitEthernet 0/0/4

[S1-GigabitEthernet0/0/4]port link-type access

[S1-GigabitEthernet0/0/4]port default vlan 2

[S1-GigabitEthernet0/0/4]quit


[S2]interface GigabitEthernet 0/0/2

[S2-GigabitEthernet0/0/2]port link-type access

[S2-GigabitEthernet0/0/2]port default vlan 3

[S2-GigabitEthernet0/0/2]

[S2]interface GigabitEthernet 0/0/2

[S2-GigabitEthernet0/0/4]port link-type access
```

[S2-GigabitEthernet0/0/4]port default vlan 3

[S2-GigabitEthernet0/0/4]quit

Configure the gateway addresses on R1 and R2. The gateway addresses are the R4 interface addresses belonging to their respective VLANs.

[R1]ip route-static 0.0.0.0 0 10.0.2.1

[R2]ip route-static 0.0.0.0 0 10.0.3.1

Run the **display vlan** command to check configurations.

```
[S1]display vlan 2
--------------------------------------------------------------------------------
U: Up;          D: Down;          TG: Tagged;          UT: Untagged;
MP: Vlan-mapping;                 ST: Vlan-stacking;
#: ProtocolTransparent-vlan;    *: Management-vlan;
--------------------------------------------------------------------------------
VID  Type     Ports
--------------------------------------------------------------------------------
2    common  UT:GE0/0/1(U)       GE0/0/4(U)
VID  Status  Property      MAC-LRN Statistics Description
--------------------------------------------------------------------------------
2    enable  default        enable  disable     VLAN 0002

[S2]display vlan 3
--------------------------------------------------------------------------------
U: Up;          D: Down;          TG: Tagged;          UT: Untagged;
MP: Vlan-mapping;                 ST: Vlan-stacking;
#: ProtocolTransparent-vlan;    *: Management-vlan;
```

```
------------------------------------------------------------------------
VID   Type      Ports
------------------------------------------------------------------------
3     common   UT:GE0/0/2(U)         GE0/0/4(U)
VID   Status   Property        MAC-LRN Statistics Description
------------------------------------------------------------------------
3     enable   default         enable   disable     VLAN 0003
```

Test the connectivity between R1 and R2.

```
[R1]ping -c 1 10.0.3.2
  PING 10.0.3.2: 56   data bytes, press CTRL_C to break
    Reply from 10.0.3.2: bytes=56 Sequence=1 ttl=254 time=3 ms

  --- 10.0.3.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 3/3/3 ms

[R2]ping -c 1 10.0.2.2
  PING 10.0.2.2: 56   data bytes, press CTRL_C to break
    Reply from 10.0.2.2: bytes=56 Sequence=1 ttl=254 time=3 ms

  --- 10.0.2.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 3/3/3 ms
```

Step 3 **Configure router-on-a-stick.**

R4 provides two subinterfaces belonging to the same physical interface to implement inter-VLAN communication.

This is router-on-a-stick.

Shut down S2's G0/0/4.

```
[S2]interface GigabitEthernet 0/0/4
[S2-GigabitEthernet0/0/4]shutdown
[S2-GigabitEthernet0/0/4]quit
```

Add G0/0/9 of S1 and S2 to VLAN 3.

```
[S2]interface GigabitEthernet 0/0/9
[S2-GigabitEthernet0/0/9]port link-type access
[S2-GigabitEthernet0/0/9]port default vlan 3
[S2-GigabitEthernet0/0/9]quit

[S1]interface GigabitEthernet 0/0/9
[S1-GigabitEthernet0/0/9]port link-type access
[S1-GigabitEthernet0/0/9]port default vlan 3
[S1-GigabitEthernet0/0/9]quit
```

Change the access type of S1's G0/0/4 to Trunk, allowing VLAN 2 and VLAN 3.

```
[S1]interface GigabitEthernet 0/0/4
[S1-GigabitEthernet0/0/4]port default vlan 1
[S1-GigabitEthernet0/0/4]port link-type trunk
[S1-GigabitEthernet0/0/4]port trunk allow-pass vlan 2 3
```

[S1-GigabitEthernet0/0/4]quit

Create two subinterfaces on R4's G0/0/1. Assign IP addresses to subinterfaces and encapsulate VID.

[R4]inter GigabitEthernet 0/0/1.2

[R4-GigabitEthernet0/0/1.2]dot1q termination vid 2

[R4-GigabitEthernet0/0/1.2]arp broadcast enable

[R4-GigabitEthernet0/0/1.2]ip address 10.0.20.1 24

[R4-GigabitEthernet0/0/1.2]quit

[R4]interface GigabitEthernet 0/0/1.3

[R4-GigabitEthernet0/0/1.3]dot1q termination vid 3

[R4-GigabitEthernet0/0/1.3]arp broadcast enable

[R4-GigabitEthernet0/0/1.3]ip address 10.0.30.1 24

[R4-GigabitEthernet0/0/1.3]quit

Run the **display ip interface brief** command to check subinterface information on R4.

[R4]display ip interface brief

*down: administratively down

^down: standby

(l): loopback

(s): spoofing

(E): E-Trunk down

The number of interface that is UP in Physical is 7

The number of interface that is DOWN in Physical is 6

The number of interface that is UP in Protocol is 5

The number of interface that is DOWN in Protocol is 8

Interface	IP Address/Mask	Physical	Protocol
Cellular0/0/0	unassigned	down	down
Cellular0/0/1	unassigned	down	down
Ethernet2/0/0	10.0.3.1/24	down	down
Ethernet2/0/1	unassigned	down	down
GigabitEthernet0/0/0	unassigned	up	down
GigabitEthernet0/0/1	10.0.2.1/24	up	up
GigabitEthernet0/0/1.2	10.0.20.1/24	up	up
GigabitEthernet0/0/1.3	10.0.30.1/24	up	up
GigabitEthernet0/0/2	unassigned	down	down
GigabitEthernet0/0/3	unassigned	up	down
NULL0	unassigned	up	up(s)
Serial1/0/0	unassigned	up	up
Serial1/0/1	unassigned	down	down

Change the IP addresses and gateway addresses of R1 and R2.

```
[R1]interface GigabitEthernet 0/0/1
[R1-GigabitEthernet0/0/1]ip address 10.0.20.2 24
[R1-GigabitEthernet0/0/1]quit
[R1]undo ip route-static 0.0.0.0 0 10.0.2.1
[R1]ip route-static 0.0.0.0 0 10.0.20.1

[R2]interface GigabitEthernet 0/0/2
[R2-GigabitEthernet0/0/2]ip address 10.0.30.2 24
[R2-GigabitEthernet0/0/2]quit
[R2]undo ip route-static 0.0.0.0 0 10.0.3.1
[R2]ip route-static 0.0.0.0 0 10.0.30.1
```

Test the connectivity between R1 and R2.

```
[R1]ping -c 1 10.0.30.2
  PING 10.0.30.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.30.2: bytes=56 Sequence=1 ttl=254 time=3 ms


  --- 10.0.30.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 3/3/3 ms
```

The **ping** command output shows that the computers in VLAN 2 and VLAN 3 successfully communicate with each other.

Compared with multi-armed route, this method reduces investment on routers.

However, in the router-on-a-stick method, all data is transmitted through the same interface. When the number of VLANs increases, the load on a single link increases. This link potentially causes a single-point failure.

Step 4 **Configure Layer 3 switching.**

In Layer 3 switching, each VLAN has a VLANIF interface, which functions as a router to allow inter-VLAN communication.

Shut down S1's G0/0/4.

```
[S1]interface GigabitEthernet 0/0/4
[S1-GigabitEthernet0/0/4]shutdown
[S1-GigabitEthernet0/0/4]quit
```

Change the access types of S1's G0/0/9 and S2's G0/0/9 to Trunk, allowing VLAN 2 and VLAN 3.

```
[S1]interface GigabitEthernet 0/0/9
[S1-GigabitEthernet0/0/9]port default vlan 1
[S1-GigabitEthernet0/0/9]port link-type trunk
[S1-GigabitEthernet0/0/9]port trunk allow-pass vlan 2 3
[S1-GigabitEthernet0/0/9]quit

[S2]interface GigabitEthernet 0/0/9
[S2-GigabitEthernet0/0/9]port default vlan 1
[S2-GigabitEthernet0/0/9]port link-type trunk
[S2-GigabitEthernet0/0/9]port trunk allow-pass vlan 2 3
[S2-GigabitEthernet0/0/9]quit
```

Create VLANIF2 and VLANIF3 on S1 and assign IP addresses to them.

```
[S1]interface Vlanif 2
[S1-Vlanif2]ip address 10.0.20.1 24
[S1-Vlanif2]quit
[S1]inter Vlanif 3
[S1-Vlanif3]ip address 10.0.30.1 24
[S1-Vlanif3]quit
```

Test the connectivity between R1 and R2.

```
[R1]ping -c 1 10.0.30.2
  PING 10.0.30.2: 56   data bytes, press CTRL_C to break
    Reply from 10.0.30.2: bytes=56 Sequence=1 ttl=254 time=2 ms

  --- 10.0.30.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
```

```
0.00% packet loss
round-trip min/avg/max = 2/2/2 ms
```

The **ping** command output shows that the computers on VLAN 2 and VLAN 3 implement Layer 3 communication through two VLANIF interfaces of S1.

Compared with the router-on-a-stick method, Layer 3 switching is more extensible. The increasing number of VLANs has little impact on services.

On a network where most traffic is inter-VLAN traffic, the network can fully support the services.

Step 5 **Configure VLAN aggregation.**

Similar to Layer 3 switching, VLAN aggregation can implement inter-VLAN communication. Different from Layer 3 switching, VLAN aggregation places all VLANs in the same network segment, reducing the number of required IP network segments and implementing unified gateway configuration.

Create VLAN 10, 20, and 100 on S1 and S2.

```
[S1]vlan batch 10 20 100
Info: This operation may take a few seconds. Please wait for a moment...done.

[S2]vlan batch 10 20 100
Info: This operation may take a few seconds. Please wait for a moment...done.
```

Allow VLAN 10 and VLAN 20 on G0/0/9 of S1 and S2.

```
[S1]interface GigabitEthernet 0/0/9
[S1-GigabitEthernet0/0/9]port trunk allow-pass vlan 10 20
[S1-GigabitEthernet0/0/9]quit

[S2]interface GigabitEthernet 0/0/9
```

[S2-GigabitEthernet0/0/9]port trunk allow-pass vlan 10 20

[S2-GigabitEthernet0/0/9]quit

Add S1's G0/0/1 and S2's G0/0/2 to VLAN 10 and VLAN 20 respectively.

[S1]interface GigabitEthernet 0/0/1

[S1-GigabitEthernet0/0/1]port default vlan 10

[S1-GigabitEthernet0/0/1]quit

[S2]interface GigabitEthernet 0/0/2

[S2-GigabitEthernet0/0/1]port default vlan 20

[S2-GigabitEthernet0/0/1]quit

Configure VLAN 100 as a super VLAN, and add VLAN 10 and VLAN 20 to VLAN 100 as sub VLANs.

[S1]vlan 100

[S1-vlan100]aggregate-vlan

[S1-vlan100]access-vlan 10 20

[S1-Vlan100]quit

Configure a VLANIF interface for VLAN 100 and enable ARP proxy.

[S1]interface Vlanif 100

[S1-Vlanif100]ip address 10.0.100.1 24

[S1-Vlanif100]arp-proxy inter-sub-vlan-proxy enable

[S1-Vlanif100]quit

Change the IP addresses of R1 and R2 to make them on the same network segment as VLANIF 100. Configure the VLANIF 100 address as the gateway address.

```
[R1]interface GigabitEthernet 0/0/1
[R1-GigabitEthernet0/0/1]ip address 10.0.100.2 24
[R1-GigabitEthernet0/0/1]quit
[R1]undo ip route-static 0.0.0.0 0 10.0.20.1
[R1]ip route-static 0.0.0.0 0 10.0.100.1

[R2]interface GigabitEthernet 0/0/2
[R2-GigabitEthernet0/0/2]ip address 10.0.100.3 24
[R2-GigabitEthernet0/0/2]quit
[R2]undo ip route-static 0.0.0.0 0 10.0.30.1
[R2]ip route-static 0.0.0.0 0 10.0.100.1
```

Test the connectivity between R1 and R2 and S1's VLANIF 100.

```
[R1]ping -c 1 10.0.100.1
  PING 10.0.100.1: 56  data bytes, press CTRL_C to break
    Reply from 10.0.100.1: bytes=56 Sequence=1 ttl=254 time=3 ms

  --- 10.0.100.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 3/3/3 ms

[R1]ping -c 1 10.0.100.3
  PING 10.0.100.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.100.3: bytes=56 Sequence=1 ttl=254 time=2 ms

  --- 10.0.100.3 ping statistics ---
```

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

[R2]pin -c 1 10.0.100.1
  PING 10.0.100.1: 56   data bytes, press CTRL_C to break
    Reply from 10.0.100.1: bytes=56 Sequence=1 ttl=254 time=3 ms

  --- 10.0.100.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 3/3/3 ms
```

The **ping** command output shows that R1 and R2 can communicate with S1's VLANIF 100. With VLAN aggregation, different VLANs can use the same gateway to communicate with each other. This conserves IP addresses and improves management efficiency. However, the computers on the same network segment communicate with each other through the same VLANIF interface. This interface bears large loads.

----End

Additional Exercises: Analysis and Verification

What are the characteristics, advantages, disadvantages, and usage scenarios of multi-armed route, router-on-a-stick, inter-VLAN communication, and VLAN aggregation?

Device Configurations

```
[S1]display current-configuration
!Software Version V200R008C00SPC500
#
sysname S1
#
vlan batch 2 to 3 10 20 100
#
diffserv domain default
#
drop-profile default
#
vlan 100
 aggregate-vlan
 access-vlan 10 20
#
aaa
 authentication-scheme default
 authorization-scheme default
 accounting-scheme default
 domain default
 domain default_admin
 local-user admin password
irreversible-cipher %^%#tK;J&jw0HG8<9-"zX!kHwzXRNjuXn96[vN47F$*L~pXcROEP3!>c)NV+:`i;%^%#
 local-user admin service-type http
#
interface Vlanif1
#
interface Vlanif2
```

```
 ip address 10.0.20.1 255.255.255.0
#
interface Vlanif3
 ip address 10.0.30.1 255.255.255.0
#
interface Vlanif100
 ip address 10.0.100.1 255.255.255.0
 arp-proxy inter-sub-vlan-proxy enable
#
interface MEth0/0/1
#
interface GigabitEthernet0/0/1
 port link-type access
 port default vlan 10
#
interface GigabitEthernet0/0/2
#
interface GigabitEthernet0/0/3
#
interface GigabitEthernet0/0/4
 shutdown
 port link-type trunk
 port trunk allow-pass vlan 2 to 3
#
interface GigabitEthernet0/0/5
#
interface GigabitEthernet0/0/6
#
interface GigabitEthernet0/0/7
```

```
#
interface GigabitEthernet0/0/8
#
interface GigabitEthernet0/0/9
 port link-type trunk
 port trunk allow-pass vlan 2 to 3 10 20
#
interface GigabitEthernet0/0/10
#
interface NULL0
#
user-interface con 0
 authentication-mode password
 set authentication password cipher $1a$fcjGHMtb0U$^GKZ+`,g@DfG$:T/P,R~iJ&')|!O":$b4)0*~&c-$
 idle-timeout 0 0
user-interface vty 0 4
user-interface vty 16 20
#
return

[S2]display current-configuration
!Software Version V200R008C00SPC500
#
sysname S2
#
vlan batch 2 to 3 10 20 100
#
diffserv domain default
#
```

```
drop-profile default
#
aaa
 authentication-scheme default
 authorization-scheme default
 accounting-scheme default
 domain default
 domain default_admin
 local-user admin password
irreversible-cipher %^%#gI/bO8qF$HkpAPUgNd'GiYR4TC!>EK#oG("Wl4_#$G*OKo-'7*R[h3+49<Z2%^%#
 local-user admin service-type http
#
interface Vlanif1
#
interface MEth0/0/1
#
interface GigabitEthernet0/0/1
#
interface GigabitEthernet0/0/2
 port link-type access
 port default vlan 20
#
interface GigabitEthernet0/0/3
#
interface GigabitEthernet0/0/4
 shutdown
 port link-type access
 port default vlan 3
#
```

```
interface GigabitEthernet0/0/5
#
interface GigabitEthernet0/0/6
#
interface GigabitEthernet0/0/7
#
interface GigabitEthernet0/0/8
#
interface GigabitEthernet0/0/9
 port link-type trunk
 port trunk allow-pass vlan 2 to 3 10 20
#
interface GigabitEthernet0/0/10
#
user-interface con 0
 authentication-mode password
 set authentication password cipher $1a$5"l`L7$/5T$,KFQ9dEy~'IggWOa7V(C+9fQOd*M;U6q,.Sl1y'H$
 idle-timeout 0 0
user-interface vty 0 4
user-interface vty 16 20
#
return

[R4]display current-configuration
[V200R007C00SPC600]
#
 sysname R4
#
 board add 0/1 2SA
```

```
 board add 0/2 2FE
#
 drop illegal-mac alarm
#
pki realm default
 enrollment self-signed
#
ssl policy default_policy type server
 pki-realm default
#
aaa
 authentication-scheme default
 authorization-scheme default
 accounting-scheme default
 domain default
 domain default_admin
 local-user admin password
irreversible-cipher %^%#`S|f)zA5xQeP^7UA/d/LH:}m3<KxR6fH,g5a%d)'zc,T/&qu:XPCg7))ihy5%^%#
 local-user admin privilege level 15
 local-user admin service-type terminal http
#
firewall zone Local
 priority 64
#
interface Ethernet2/0/0
 ip address 10.0.3.1 255.255.255.0
#
interface Ethernet2/0/1
#
```

```
interface Serial1/0/0
 link-protocol ppp
#
interface Serial1/0/1
 link-protocol ppp
#
interface GigabitEthernet0/0/0
#
interface GigabitEthernet0/0/1
 ip address 10.0.2.1 255.255.255.0
#
interface GigabitEthernet0/0/1.2
 dot1q termination vid 2
 ip address 10.0.20.1 255.255.255.0
#
interface GigabitEthernet0/0/1.3
 dot1q termination vid 3
 ip address 10.0.30.1 255.255.255.0
#
interface GigabitEthernet0/0/2
#
interface GigabitEthernet0/0/3
 description VirtualPort
#
interface Cellular0/0/0
#
interface Cellular0/0/1
#
interface NULL0
```

```
#
 snmp-agent local-engineid 800007DB03D0D04B03D43B
#
 http secure-server ssl-policy default_policy
 http server enable
 http secure-server enable
#
user-interface con 0
 authentication-mode aaa
 idle-timeout 0 0
user-interface vty 0
 authentication-mode aaa
 user privilege level 15
user-interface vty 1 4
#
wlan ac
#
voice
 #
 diagnose
#
ops
#
autostart
#
return
```

Chapter 7 STP Configurations

Lab 7-1 STP, RSTP, and MSTP

Learning Objectives

The objectives of this lab are to learn and understand:

- Differences between STP, RSTP, and MSTP
- How to change the bridge priority to control root bridge election
- How to change the port priority to control election of the root port and designated port
- How to configure RSTP and compatibility between STP and RSTP
- How to configure MSTP to implement VLAN load balancing

Topology

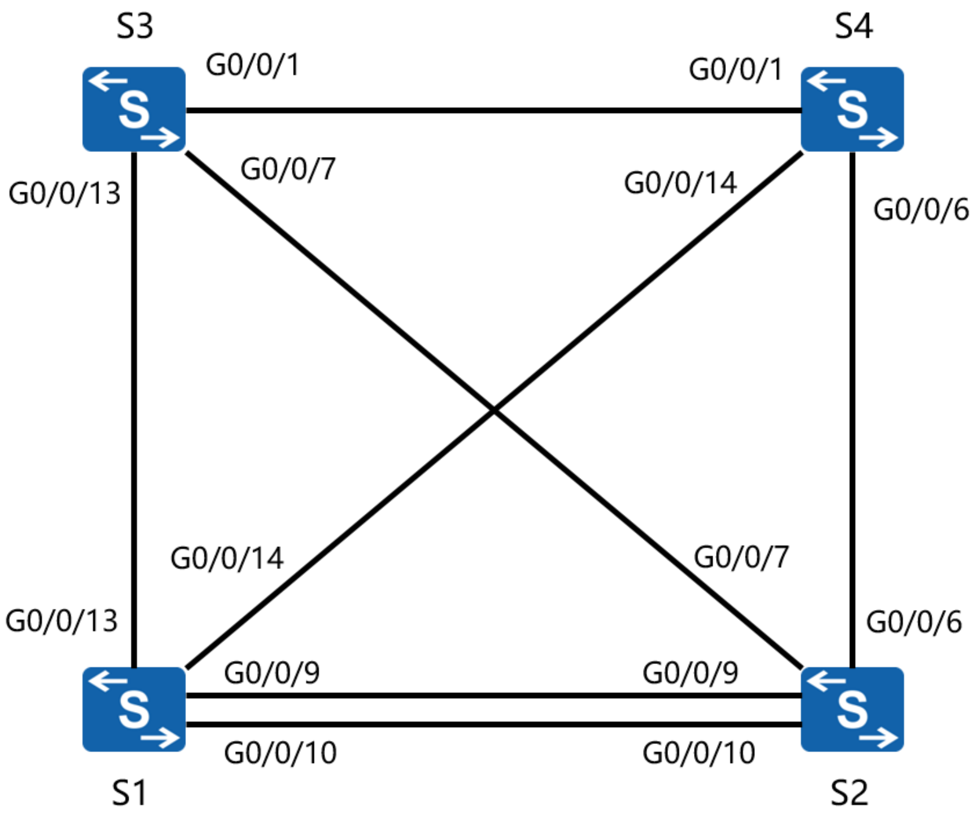

Figure 7-1 STP, RSTP, and MSTP topology

Scenario

You are a network administrator of a company. The company uses a backup network, and configure STP to prevent loops. STP convergence on interfaces requires a long time. To speed up convergence, the company needs to use RSTP. All VLANs share an STP tree. To load balance traffic between VLANs, the company needs to use MSTP.

Tasks

Step 1 **Configure STP and verify the STP configuration.**

If STP is not enabled, enable it.

```
[S1]stp enable

[S2]stp enable

[S3]stp enable

[S4]stp enable
```

Configure STP.

```
[S1]stp mode stp

[S2]stp mode stp

[S3]stp mode stp

[S4]stp mode stp
```

Check STP status.

```
[S1]display stp
-------[CIST Global Info][Mode STP]-------
CIST Bridge         :32768.4c1f-cc45-aadc
Bridge Times        :Hello 2s MaxAge 20s FwDly 15s MaxHop 20
CIST Root/ERPC      :32768.4c1f-cc45-aac1 / 20000
CIST RegRoot/IRPC   :32768.4c1f-cc45-aadc / 0
CIST RootPortId     :128.9
BPDU-Protection     :Disabled
TC or TCN received  :36
TC count per hello  :2
STP Converge Mode   :Normal
Share region-configuration :Enabled
Time since last TC  :0 days 0h:0m:1s
...output omit...

[S2]display stp
-------[CIST Global Info][Mode STP]-------
CIST Bridge         :32768.4c1f-cc45-aac1
Bridge Times        :Hello 2s MaxAge 20s FwDly 15s MaxHop 20
CIST Root/ERPC      :32768.4c1f-cc45-aac1 / 0
CIST RegRoot/IRPC   :32768.4c1f-cc45-aac1 / 0
CIST RootPortId     :0.0
BPDU-Protection     :Disabled
TC or TCN received  :20
TC count per hello  :0
STP Converge Mode   :Normal
Share region-configuration :Enabled
Time since last TC  :0 days 0h:1m:4s
```

```
...output omit...

[S1]display stp brief
 MSTID  Port                    Role  STP State     Protection
   0    GigabitEthernet0/0/9    ROOT  FORWARDING    NONE
   0    GigabitEthernet0/0/10   ALTE  DISCARDING    NONE
   0    GigabitEthernet0/0/13   DESI  FORWARDING    NONE
   0    GigabitEthernet0/0/14   DESI  FORWARDING    NONE

[S2]display stp brief
 MSTID  Port                    Role  STP State     Protection
   0    GigabitEthernet0/0/9    DESI  FORWARDING    NONE
   0    GigabitEthernet0/0/10   DESI  FORWARDING    NONE
   0    GigabitEthernet0/0/13   DESI  FORWARDING    NONE
   0    GigabitEthernet0/0/14   DESI  FORWARDING    NONE
```

S2 is the root bridge, and all ports are designated ports.

The switch MAC addresses are not fixed, so the actual lab test result may be different.

Step 2 **Control root bridge election.**

Configure S1 as the primary root bridge and S2 as the backup root bridge.

```
[S1]stp root primary

[S2]stp root secondary
```

Check the STP configuration.

```
[S1]display stp
```

```
-------[CIST Global Info][Mode STP]-------
CIST Bridge         :0     .4c1f-cc45-aadc
Bridge Times        :Hello 2s MaxAge 20s FwDly 15s MaxHop 20
CIST Root/ERPC      :0     .4c1f-cc45-aadc / 0
CIST RegRoot/IRPC   :0     .4c1f-cc45-aadc / 0
CIST RootPortId     :0.0
BPDU-Protection     :Disabled
CIST Root Type      :Primary root
TC or TCN received  :67
TC count per hello  :0
STP Converge Mode   :Normal
Share region-configuration :Enabled
Time since last TC  :0 days 0h:0m:15s
...output omit...

[S2]display stp
-------[CIST Global Info][Mode STP]-------
CIST Bridge         :4096 .4c1f-cc45-aac1
Bridge Times        :Hello 2s MaxAge 20s FwDly 15s MaxHop 20
CIST Root/ERPC      :0     .4c1f-cc45-aadc / 20000
CIST RegRoot/IRPC   :4096 .4c1f-cc45-aac1 / 0
CIST RootPortId     :128.9
BPDU-Protection     :Disabled
CIST Root Type      :Secondary root
TC or TCN received  :26
TC count per hello  :0
STP Converge Mode   :Normal
Share region-configuration :Enabled
Time since last TC  :0 days 0h:0m:1s
```

...output omit...

S1 is the primary root bridge and S2 is the backup root bridge.

A small bridge priority value indicates a high priority. Change the bridge priority of S1 to 8192 and bridge priority of S2 to 4096.

[S1]undo stp root

[S1]stp priority 8192

[S2]undo stp root

[S2]stp priority 4096

Check STP information.

[S1]display stp

-------[CIST Global Info][Mode STP]-------

CIST Bridge :8192 .4c1f-cc45-aadc

Bridge Times :Hello 2s MaxAge 20s FwDly 15s MaxHop 20

CIST Root/ERPC :4096 .4c1f-cc45-aac1 / 20000

CIST RegRoot/IRPC :8192 .4c1f-cc45-aadc / 0

CIST RootPortId :128.9

BPDU-Protection :Disabled

TC or TCN received :79

TC count per hello :1

STP Converge Mode :Normal

Share region-configuration :Enabled

Time since last TC :0 days 0h:0m:0s

...output omit...

[S2]display stp

```
-------[CIST Global Info][Mode STP]-------
CIST Bridge              :4096 .4c1f-cc45-aac1
Bridge Times             :Hello 2s MaxAge 20s FwDly 15s MaxHop 20
CIST Root/ERPC           :4096 .4c1f-cc45-aac1 / 0
CIST RegRoot/IRPC        :4096 .4c1f-cc45-aac1 / 0
CIST RootPortId          :0.0
BPDU-Protection          :Disabled
TC or TCN received       :88
TC count per hello       :0
STP Converge Mode        :Normal
Share region-configuration :Enabled
Time since last TC       :0 days 0h:0m:9s
...output omit...
```

The priority of S1 is 8192, the priority of S2 is 4096, and S2 is the root bridge.

Step 3 **Control root port election.**

Check port roles on S1.

```
[S1]display stp brief
 MSTID  Port                    Role  STP State     Protection
    0   GigabitEthernet0/0/9    ROOT  FORWARDING      NONE
    0   GigabitEthernet0/0/10   ALTE  DISCARDING      NONE
    0   GigabitEthernet0/0/13   DESI  FORWARDING      NONE
    0   GigabitEthernet0/0/14   DESI  FORWARDING      NONE
```

GigabitEthernet0/0/9 of S1 is the root port.

The default port priority is 128. A large value indicates a low priority.

S1 and S2 are connected through G0/0/9 and G0/0/10.

On S2, set the priority of G0/0/9 to 32 and the priority of G0/0/10 to 16.

```
[S2]interface GigabitEthernet 0/0/9
[S2-GigabitEthernet0/0/9]stp port priority 32
[S2-GigabitEthernet0/0/9]quit
[S2]interface GigabitEthernet 0/0/10
[S2-GigabitEthernet0/0/10]stp port priority 16
[S2-GigabitEthernet0/0/10]quit
```

Note: The port priorities are changed on S2, not S1.

Check port roles on S1.

```
[S1]display stp brief
 MSTID  Port                     Role  STP State     Protection
   0    GigabitEthernet0/0/9     ALTE  DISCARDING      NONE
   0    GigabitEthernet0/0/10    ROOT  DISCARDING      NONE
   0    GigabitEthernet0/0/13    DESI  FORWARDING      NONE
   0    GigabitEthernet0/0/14    DESI  FORWARDING      NONE
```

GigabitEthernet0/0/10 of S1 is the root port.

Step 4 **Control designated port election.**

Check the status of interfaces directly connected between S3 and S4.

```
 [S3]display stp interface GigabitEthernet 0/0/1
----[CIST][Port1(GigabitEthernet0/0/1)][DISCARDING]----
 Port Protocol          :Enabled
 Port Role              :Alternate Port
 Port Priority          :128
 Port Cost(Dot1T )      :Config=auto / Active=199999
 Designated Bridge/Port     :32768.5489-98ec-f00a / 128.1
```

```
 Port Edged            :Config=default / Active=disabled
 Point-to-point        :Config=auto / Active=true
 Transit Limit         :147 packets/hello-time
 Protection Type       :None
 Port STP Mode         :STP
 Port Protocol Type    :Config=auto / Active=dot1s
 PortTimes             :Hello 2s MaxAge 20s FwDly 15s RemHop 0
 TC or TCN send        :17
 TC or TCN received    :52
 BPDU Sent             :172
          TCN: 0, Config: 172, RST: 0, MST: 0
 BPDU Received         :206
          TCN: 0, Config: 206, RST: 0, MST: 0

[S4]display stp interface GigabitEthernet 0/0/6
----[CIST][Port24(GigabitEthernet0/0/6)][DISCARDING]----
 Port Protocol         :Enabled
 Port Role             :Designated Port
 Port Priority         :128
 Port Cost(Dot1T )     :Config=auto / Active=199999
 Designated Bridge/Port    :32768.5489-98ec-f00a / 128.1
 Port Edged            :Config=default / Active=disabled
 Point-to-point        :Config=auto / Active=true
 Transit Limit         :147 packets/hello-time
 Protection Type       :None
 Port STP Mode         :STP
 Port Protocol Type    :Config=auto / Active=dot1s
 PortTimes             :Hello 2s MaxAge 20s FwDly 15s RemHop 20
 TC or TCN send        :37
```

TC or TCN received :17

BPDU Sent :181

TCN: 0, Config: 181, RST: 0, MST: 0

BPDU Received :172

TCN: 0, Config: 172, RST: 0, MST: 0

S3's GigabitEthernet 0/0/1 is an alternate port. S4's GigabitEthernet 0/0/1 is a designated port. Change the path cost of S4's GE0/0/6 to 2000000.

[S4]interface GigabitEthernet0/0/6

[S4-GigabitEthernet0/0/6]stp cost 2000000

[S4-GigabitEthernet0/0/6]quit

Check current port roles.

[S3]display stp interface GigabitEthernet 0/0/1

----[CIST][Port1(GigabitEthernet0/0/1)][FORWARDING]----

Port Protocol :Enabled

Port Role :Designated Port

Port Priority :128

Port Cost(Dot1T) :Config=auto / Active=199999

Designated Bridge/Port :32768.5489-98ec-f022 / 128.1

Port Edged :Config=default / Active=disabled

Point-to-point :Config=auto / Active=true

Transit Limit :147 packets/hello-time

Protection Type :None

Port STP Mode :STP

Port Protocol Type :Config=auto / Active=dot1s

PortTimes :Hello 2s MaxAge 20s FwDly 15s RemHop 20

TC or TCN send :52

```
 TC or TCN received   :52
 BPDU Sent             :284
          TCN: 0, Config: 284, RST: 0, MST: 0
 BPDU Received         :380
          TCN: 0, Config: 380, RST: 0, MST: 0

[S4]display stp interface GigabitEthernet 0/0/6
----[CIST][Port6(GigabitEthernet0/0/6)][DISCARDING]----
 Port Protocol         :Enabled
 Port Role             :Alternate Port
 Port Priority         :128
 Port Cost(Dot1T )     :Config=2000000 / Active=2000000
 Designated Bridge/Port    :4096.4c1f-cc45-aac1 / 128.24
 Port Edged            :Config=default / Active=disabled
 Point-to-point        :Config=auto / Active=true
 Transit Limit         :147 packets/hello-time
 Protection Type       :None
 Port STP Mode         :STP
 Port Protocol Type    :Config=auto / Active=dot1s
 PortTimes             :Hello 2s MaxAge 20s FwDly 15s RemHop 0
 TC or TCN send        :7
 TC or TCN received    :162
 BPDU Sent             :8
          TCN: 7, Config: 1, RST: 0, MST: 0
 BPDU Received         :1891
          TCN: 0, Config: 1891, RST: 0, MST: 0
```

S3's GigabitEthernet 0/0/1 is a designated port. S4's GigabitEthernet 0/0/1 is an alternate port.

Step 5 **Configure RSTP and verify the RSTP configuration.**

Configure VLANIF 1 addresses on S1 and S2. Test the connectivity between S1 and S2.

```
[S1]interface Vlanif 1
[S1-Vlanif1]ip address 10.0.1.1 24
[S1-Vlanif1]quit

[S2]interface Vlanif 1
[S2-Vlanif1]ip address 10.0.1.2 24
[S2-Vlanif1]quit

[S1]ping 10.0.1.2
  PING 10.0.1.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.1.2: bytes=56 Sequence=1 ttl=255 time=9 ms
    Reply from 10.0.1.2: bytes=56 Sequence=2 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=3 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=4 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=5 ttl=254 time=1 ms

  --- 10.0.1.2 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 1/2/9 ms
```

Check port roles on S1.

```
[S1]display stp brief
 MSTID  Port                    Role  STP State     Protection
   0    GigabitEthernet0/0/9    ALTE  DISCARDING    NONE
   0    GigabitEthernet0/0/10   ROOT  FORWARDING    NONE
   0    GigabitEthernet0/0/13   DESI  FORWARDING    NONE
   0    GigabitEthernet0/0/14   DESI  FORWARDING    NONE
```

S1's GigabitEthernet0/0/10 is the root port. Ping S2 from S1 20 times to test connectivity.

Note: After S1 performs ping, shut down S2's GigabitEthernet 0/0/10 immediately.

```
 [S1]ping -c 20 10.0.1.2
  PING 10.0.1.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.1.2: bytes=56 Sequence=1 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=2 ttl=254 time=1 ms
Dec 21 2011 16:20:44-05:13 S1 %%01IFNET/4/IF_STATE(l)[5]:Interface GigabitEthernet0/0/10 has turned into
DOWN state.
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
```

```
    Request time out
    Request time out
    Request time out
    Reply from 10.0.1.2: bytes=56 Sequence=18 ttl=255 time=15 ms
    Reply from 10.0.1.2: bytes=56 Sequence=19 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=20 ttl=254 time=1 ms

  --- 10.0.1.2 ping statistics ---
    20 packet(s) transmitted
    5 packet(s) received
    75.00% packet loss
    round-trip min/avg/max = 1/3/15 ms

[S2]interface GigabitEthernet 0/0/10
[S2-GigabitEthernet0/0/10]shutdown
[S2-GigabitEthernet0/0/10]quit
```

Check port roles on S1.

```
[S1]display stp brief
 MSTID  Port                        Role  STP State     Protection
   0    GigabitEthernet0/0/9        ROOT  FORWARDING      NONE
   0    GigabitEthernet0/0/13       DESI  FORWARDING      NONE
   0    GigabitEthernet0/0/14       DESI  FORWARDING      NONE
```

S1's GigabitEthernet0/0/9 becomes the root port, and the port enters the Forwarding state. There are 15 timeout packets, and network convergence time is 30s.

Enable S2's GigabitEthernet 0/0/10.

[S2]interface GigabitEthernet 0/0/10

[S2-GigabitEthernet0/0/10]undo shutdown

[S2-GigabitEthernet0/0/10]quit

Set the STP mode to RSTP.

[S1]stp mode rstp

[S2]stp mode rstp

[S3]stp mode rstp

[S4]stp mode rstp

Check port roles on S1.

[S1]display stp brief

MSTID	Port	Role	STP State	Protection
0	GigabitEthernet0/0/9	ALTE	DISCARDING	NONE
0	GigabitEthernet0/0/10	ROOT	FORWARDING	NONE
0	GigabitEthernet0/0/13	DESI	FORWARDING	NONE
0	GigabitEthernet0/0/14	DESI	FORWARDING	NONE

S1's GigabitEthernet0/0/10 is the root port. Ping S2 from S1 20 times to test connectivity.

Note: After S1 performs ping, shut down S2's GigabitEthernet 0/0/10 immediately.

[S1]ping -c 20 10.0.1.2

PING 10.0.1.2: 56 data bytes, press CTRL_C to break

Reply from 10.0.1.2: bytes=56 Sequence=1 ttl=254 time=1 ms

```
    Reply from 10.0.1.2: bytes=56 Sequence=2 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=3 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=4 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=5 ttl=254 time=1 ms
Dec 21 2011 16:37:10-05:13 S1 %%01IFNET/4/IF_STATE(l)[7]:Interface GigabitEthernet0/0/10 has turned into
DOWN state.
    Request time out
    Reply from 10.0.1.2: bytes=56 Sequence=7 ttl=255 time=10 ms
    Reply from 10.0.1.2: bytes=56 Sequence=8 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=9 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=10 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=11 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=12 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=13 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=14 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=15 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=16 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=17 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=18 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=19 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=20 ttl=254 time=1 ms

  --- 10.0.1.2 ping statistics ---
    20 packet(s) transmitted
    19 packet(s) received
    5.00% packet loss
round-trip min/avg/max = 1/1/10 ms

[S2]interface GigabitEthernet 0/0/10
```

```
[S2-GigabitEthernet0/0/10]shutdown
[S2-GigabitEthernet0/0/10]quit
```

Check port roles on S1.

```
[S1]display stp brief
 MSTID  Port                        Role  STP State     Protection
   0    GigabitEthernet0/0/9        ROOT  FORWARDING      NONE
   0    GigabitEthernet0/0/13       DESI  FORWARDING      NONE
   0    GigabitEthernet0/0/14       DESI  FORWARDING      NONE
```

S1's GigabitEthernet0/0/9 becomes the root port and enters the Forwarding state. There is one timeout packet, and network convergence time is 2s.

Enable S2's GigabitEthernet 0/0/10.

```
[S2]interface GigabitEthernet 0/0/10
[S2-GigabitEthernet0/0/10]undo shutdown
[S2-GigabitEthernet0/0/10]quit
```

Step 6 **Check compatibility between RSTP and STP.**

Configure STP mode on S1 to STP and retain other configurations unchanged.

```
[S1]stp mode stp
```

Check port roles on S1.

```
[S1]display stp brief
 MSTID  Port                        Role  STP State     Protection
   0    GigabitEthernet0/0/9        ALTE  DISCARDING      NONE
   0    GigabitEthernet0/0/10       ROOT  FORWARDING      NONE
```

```
0     GigabitEthernet0/0/13       DESI  FORWARDING      NONE
0     GigabitEthernet0/0/14       DESI  FORWARDING      NONE
```

S1's GigabitEthernet0/0/10 is the root port. Ping S2 from S1 20 times to test connectivity.

Note: After S1 performs ping, shut down S2's GigabitEthernet 0/0/10 immediately.

```
[S1]ping -c 20 10.0.1.2
  PING 10.0.1.2: 56   data bytes, press CTRL_C to break
    Reply from 10.0.1.2: bytes=56 Sequence=1 ttl=254 time=1 ms
    Reply from 10.0.1.2: bytes=56 Sequence=2 ttl=254 time=1 ms
Dec 21 2011 16:20:44-05:13 S1 %%01IFNET/4/IF_STATE(l)[5]:Interface GigabitEthernet0/0/10 has turned into
DOWN state.
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out
    Reply from 10.0.1.2: bytes=56 Sequence=18 ttl=255 time=15 ms
```

```
  Reply from 10.0.1.2: bytes=56 Sequence=19 ttl=254 time=1 ms
  Reply from 10.0.1.2: bytes=56 Sequence=20 ttl=254 time=1 ms

 --- 10.0.1.2 ping statistics ---
  20 packet(s) transmitted
  5 packet(s) received
  75.00% packet loss
  round-trip min/avg/max = 1/3/15 ms

[S2]interface GigabitEthernet 0/0/10
[S2-GigabitEthernet0/0/10]shutdown
```

Check port roles on S1.

```
[S1]display stp brief
 MSTID  Port                        Role  STP State     Protection
   0    GigabitEthernet0/0/9        ROOT  FORWARDING      NONE
   0    GigabitEthernet0/0/13       DESI  FORWARDING      NONE
   0    GigabitEthernet0/0/14       DESI  FORWARDING      NONE
```

S1's GigabitEthernet0/0/9 becomes the root port and enters the Forwarding state. There are 15 timeout packets, and network convergence time is 30s.

RSTP is compatible with STP, but the convergence mode is STP.

Enable S2's GigabitEthernet 0/0/10.

```
[S2]interface GigabitEthernet 0/0/10
[S2-GigabitEthernet0/0/10]undo shutdown
[S2-GigabitEthernet0/0/10]quit
```

Step 7 **Configure MSTP and verify the MSTP configuration.**

Create VLANs 2 to 20 and add related interfaces to the VLANs.

[S1]vlan batch 2 to 20

Info: This operation may take a few seconds. Please wait for a moment...done.

[S1]interface GigabitEthernet 0/0/9

[S1-GigabitEthernet0/0/9]port link-type trunk

[S1-GigabitEthernet0/0/9]port trunk allow-pass vlan 1 TO 20

[S1-GigabitEthernet0/0/9]quit

[S1]interface GigabitEthernet 0/0/10

[S1-GigabitEthernet0/0/10]port link-type trunk

[S1-GigabitEthernet0/0/10]port trunk allow-pass vlan 1 TO 20

[S1-GigabitEthernet0/0/10]quit

[S1]interface GigabitEthernet 0/0/13

[S1-GigabitEthernet0/0/13]port link-type trunk

[S1-GigabitEthernet0/0/13]port trunk allow-pass vlan 1 TO 20

[S1-GigabitEthernet0/0/13]quit

[S1]interface GigabitEthernet 0/0/14

[S1-GigabitEthernet0/0/14]port link-type trunk

[S1-GigabitEthernet0/0/14]port trunk allow-pass vlan 1 TO 20

[S1-GigabitEthernet0/0/14]quit

[S2]vlan batch 1 to 20

Info: This operation may take a few seconds. Please wait for a moment...done.

[S2]interface GigabitEthernet 0/0/9

[S2-GigabitEthernet0/0/9]port link-type trunk

[S2-GigabitEthernet0/0/9]port trunk allow-pass vlan 1 TO 20

[S2-GigabitEthernet0/0/9]quit

[S2]interface GigabitEthernet 0/0/10

```
[S2-GigabitEthernet0/0/10]port link-type trunk
[S2-GigabitEthernet0/0/10]port trunk allow-pass vlan 1 TO 20
[S2-GigabitEthernet0/0/10]quit
[S2]interface GigabitEthernet 0/0/6
[S2-GigabitEthernet0/0/6]port link-type trunk
[S2-GigabitEthernet0/0/6]port trunk allow-pass vlan 1 TO 20
[S2-GigabitEthernet0/0/6]quit
[S2]interface GigabitEthernet 0/0/7
[S2-GigabitEthernet0/0/7]port link-type trunk
[S2-GigabitEthernet0/0/7]port trunk allow-pass vlan 1 TO 20
[S2-GigabitEthernet0/0/7]quit

[S3]vlan batch 1 to 20
Info: This operation may take a few seconds. Please wait for a moment...done.
[S3]interface GigabitEthernet0/0/1
[S3-GigabitEthernet0/0/1]port link-type trunk
[S3-GigabitEthernet0/0/1]port trunk allow-pass vlan 1 TO 20
[S3-GigabitEthernet0/0/1]quit
[S3]interface GigabitEthernet0/0/13
[S3-GigabitEthernet0/0/13]port link-type trunk
[S3-GigabitEthernet0/0/13]port trunk allow-pass vlan 1 TO 20
[S3-GigabitEthernet0/0/13]quit
[S3]interface GigabitEthernet0/0/7
[S3-GigabitEthernet0/0/7]port link-type trunk
[S3-GigabitEthernet0/0/7]port trunk allow-pass vlan 1 TO 20
[S3-GigabitEthernet0/0/7]quit

[S4]vlan batch 1 to 20
Info: This operation may take a few seconds. Please wait for a moment...done.
```

```
[S4]interface GigabitEthernet0/0/1
[S4-GigabitEthernet0/0/1]port link-type trunk
[S4-GigabitEthernet0/0/1]port trunk allow-pass vlan 1 TO 20
[S4-GigabitEthernet0/0/1]quit
[S4]interface GigabitEthernet0/0/14
[S4-GigabitEthernet0/0/14]port link-type trunk
[S4-GigabitEthernet0/0/14]port trunk allow-pass vlan 1 TO 20
[S4-GigabitEthernet0/0/14]quit
[S4]interface GigabitEthernet0/0/6
[S4-GigabitEthernet0/0/6]port link-type trunk
[S4-GigabitEthernet0/0/6]port trunk allow-pass vlan 1 TO 20
[S4-GigabitEthernet0/0/6]quit
```

Configure MSTP.

Add VLANs 1-10 to instance 1 and VLANs 11-20 to instance 2.

```
[S1]stp mode mstp
[S1]stp region-configuration
[S1-mst-region]region-name RG1
[S1-mst-region]instance 1 vlan 1 TO 10
[S1-mst-region]instance 2 vlan 11 to 20
[S1-mst-region]active region-configuration
Info: This operation may take a few seconds. Please wait for a moment....done.
[S1-mst-region]quit

[S2]stp mode mstp
[S2]stp region-configuration
[S2-mst-region]region-name RG1
[S2-mst-region]instance 1 vlan 1 TO 10
```

[S2-mst-region]instance 2 vlan 11 to 20

[S2-mst-region]active region-configuration

Info: This operation may take a few seconds. Please wait for a moment....done.

[S2-mst-region]quit

[S3]STP mode mstp

Info: This operation may take a few seconds. Please wait for a moment.....done.

[S3]stp region-configuration

[S3-mst-region]region-name RG1

[S3-mst-region]instance 1 vlan 1 to 10

[S3-mst-region]instance 2 vlan 11 to 20

[S3-mst-region]quit

[S4]STP mode mstp

Info: This operation may take a few seconds. Please wait for a moment.....done.

[S4]stp region-configuration

[S4-mst-region]region-name RG1

[S4-mst-region]instance 1 vlan 1 to 10

[S4-mst-region]instance 2 vlan 11 to 20

[S4-mst-region]quit

Check the mappings between MSTP instances and VLANs.

[S1]display stp region-configuration

Oper configuration

Format selector :0

Region name :RG1

Revision level :0

Instance VLANs Mapped

0	21 to 4094
1	1 to 10
2	11 to 20

Set the S1 priority in instance 1 to 4096 and the S1 priority in instance 2 to 8192.

Set the S2 priority in instance 2 to 4096 and the S2 priority in instance 1 to 8192.

```
[S1]stp instance 1 priority 4096
[S1]stp instance 2 priority 8192

[S2]stp instance 2 priority 4096
[S2]stp instance 1 priority 8192
```

Check the status of instance 1 and instance 2.

```
[S1]display stp instance 1
-------[MSTI 1 Global Info]-------
MSTI Bridge ID          :4096.4c1f-cc45-aadc
MSTI RegRoot/IRPC       :4096.4c1f-cc45-aadc / 0
MSTI RootPortId         :0.0
Master Bridge           :4096.4c1f-cc45-aac1
Cost to Master          :20000
TC received             :20
TC count per hello      :0

[S2]display stp instance 2
-------[MSTI 2 Global Info]-------
MSTI Bridge ID          :4096.4c1f-cc45-aac1
MSTI RegRoot/IRPC       :4096.4c1f-cc45-aac1 / 0
```

MSTI RootPortId :0.0

Master Bridge :4096.4c1f-cc45-aac1

Cost to Master :0

TC received :16

TC count per hello :0

S1 is the root bridge of instance 1 and S2 is the root bridge of instance 2.

Check port roles in MSTP instance 1.

[S1]display stp instance 1 brief

MSTID	Port	Role	STP State	Protection
1	GigabitEthernet0/0/9	DESI	FORWARDING	NONE
1	GigabitEthernet0/0/10	DESI	FORWARDING	NONE
1	GigabitEthernet0/0/13	DESI	FORWARDING	NONE
1	GigabitEthernet0/0/14	DESI	FORWARDING	NONE

[S2]display stp instance 1 brief

MSTID	Port	Role	STP State	Protection
1	GigabitEthernet0/0/9	ROOT	FORWARDING	NONE
1	GigabitEthernet0/0/10	ALTE	DISCARDING	NONE
1	GigabitEthernet0/0/6	DESI	FORWARDING	NONE
1	GigabitEthernet0/0/7	DESI	FORWARDING	NONE

[S3]display stp instance 1 brief

MSTID	Port	Role	STP State	Protection
1	GigabitEthernet0/0/1	ALTE	DISCARDING	NONE
1	GigabitEthernet0/0/7	ALTE	DISCARDING	NONE
1	GigabitEthernet0/0/13	ROOT	FORWARDING	NONE

```
[S4]display stp instance 1 brief
 MSTID  Port                        Role  STP State     Protection
   1    GigabitEthernet0/0/1        DESI  FORWARDING      NONE
   1    GigabitEthernet0/0/6        ALTE  DISCARDING      NONE
   1    GigabitEthernet0/0/14       ROOT  FORWARDING      NONE
```

S1 in instance 1 is the root bridge. The users in VLAN 1-10 on S3 communicate with the users in VLAN 1-10 on S1, S2, and S4 through Ethernet0/0/13.

Check port roles in MSTP instance 2.

```
[S1]display stp instance 2 brief
 MSTID  Port                        Role  STP State     Protection
   2    GigabitEthernet0/0/9        ROOT  FORWARDING      NONE
   2    GigabitEthernet0/0/10       ALTE  DISCARDING      NONE
   2    GigabitEthernet0/0/13       DESI  FORWARDING      NONE
```

```
[S2]display stp instance 2 brief
 MSTID  Port                        Role  STP State     Protection
   2    GigabitEthernet0/0/9        DESI  FORWARDING      NONE
   2    GigabitEthernet0/0/10       DESI  FORWARDING      NONE
   2    GigabitEthernet0/0/6        DESI  FORWARDING      NONE
   2    GigabitEthernet0/0/7        DESI  FORWARDING      NONE
```

```
[S3]display stp instance 2 brief
 MSTID  Port                        Role  STP State     Protection
   2    GigabitEthernet0/0/1        ALTE  DISCARDING      NONE
   2    GigabitEthernet0/0/7        ROOT  FORWARDING      NONE
   2    GigabitEthernet0/0/13       ALTE  DISCARDING      NONE
```

```
[S4]display stp instance 2 brief
 MSTID  Port                        Role  STP State     Protection
    2   GigabitEthernet0/0/1        DESI  FORWARDING      NONE
    2   GigabitEthernet0/0/14       DESI  FORWARDING      NONE
    2   GigabitEthernet0/0/6        ROOT  FORWARDING      NONE
```

S2 in instance 2 is the root bridge. The users in VLAN 11-20 on S3 communicate with the users in VLAN 11-20 on S1, S2, and S4 through Ethernet0/0/23.

----End

Additional Exercises: Analysis and Verification

How can MSTP load balance data from different VLANs in multiple regions?

What's the reason that RSTP is able to implement fast forwarding?

Device Configurations

```
[S1]display current-configuration
#
!Software Version V200R008C00SPC500
 sysname S1
#
vlan batch 2 to 20
#
 stp instance 0 priority 8192
 stp instance 1 priority 4096
stp instance 2 priority 8192
#
 stp region-configuration
  region-name RG1
```

```
  instance 1 vlan 1 to 10
  instance 2 vlan 11 to 20
  active region-configuration
#
interface Vlanif1
 ip address 10.0.1.1 255.255.255.0
#
interface GigabitEthernet0/0/9
 port link-type trunk
 port trunk allow-pass vlan 2 to 20
#
interface GigabitEthernet0/0/10
 port link-type trunk
 port trunk allow-pass vlan 2 to 20
#
interface GigabitEthernet0/0/13
 port link-type trunk
 port trunk allow-pass vlan 2 to 20
#
Return

[S2]display current-configuration
#
!Software Version V200R008C00SPC500
 sysname S2
#
 vlan batch 2 to 20
#
stp instance 0 priority 4096
```

```
stp instance 1 priority 8192
stp instance 2 priority 4096
#
 stp region-configuration
  region-name RG1
  instance 1 vlan 1 to 10
  instance 2 vlan 11 to 20
  active region-configuration
#
interface Vlanif1
 ip address 10.0.1.2 255.255.255.0
#
interface GigabitEthernet0/0/6
 port link-type trunk
 port trunk allow-pass vlan 2 to 20
#
interface GigabitEthernet0/0/7
 port link-type trunk
 port trunk allow-pass vlan 2 to 20
#
interface GigabitEthernet0/0/9
 port link-type trunk
 port trunk allow-pass vlan 2 to 20
 stp instance 0 port priority 32
#
interface GigabitEthernet0/0/10
 port link-type trunk
 port trunk allow-pass vlan 2 to 20
 stp instance 0 port priority 16
```

```
#
Return

[S3]display current-configuration
#
!Software Version V200R008C00SPC500
 sysname S3
#
vlan batch 2 to 20
#
stp region-configuration
  region-name RG1
  instance 1 vlan 1 to 10
  instance 2 vlan 11 to 20
  active region-configuration
#
interface GigabitEthernet0/0/1
 port link-type trunk
 port trunk allow-pass vlan 2 to 20
#
interface GigabitEthernet0/0/13
 port link-type trunk
 port trunk allow-pass vlan 2 to 20
#
interface GigabitEthernet0/0/7
 port link-type trunk
 port trunk allow-pass vlan 2 to 20
#
```

```
Return

[S4]display current-configuration
#
!Software Version V200R008C00SPC500
 sysname S4
#
 vlan batch 2 to 20
#
 stp region-configuration
  region-name RG1
  instance 1 vlan 1 to 10
  instance 2 vlan 11 to 20
  active region-configuration
#
interface GigabitEthernet0/0/1
 port link-type trunk
 port trunk allow-pass vlan 2 to 20
#
interface GigabitEthernet0/0/14
 port link-type trunk
 port trunk allow-pass vlan 2 to 20
#
interface Ethernet0/0/23
#
interface GigabitEthernet0/0/6
 port link-type trunk
 port trunk allow-pass vlan 2 to 20
 stp instance 0 cost 2000000
```

#

Return

Lab 7-2 Compatibility Between MST Multi-Region and STP (Optional)

Learning Objectives

The objectives of this lab are to learn and understand:

- How to configure MST multi-instance and multi-region
- How to configure compatibility between MSTP and STP
- How to configure protection for MSTP edge ports, designated ports, loop, and TC-BPDU

Topology

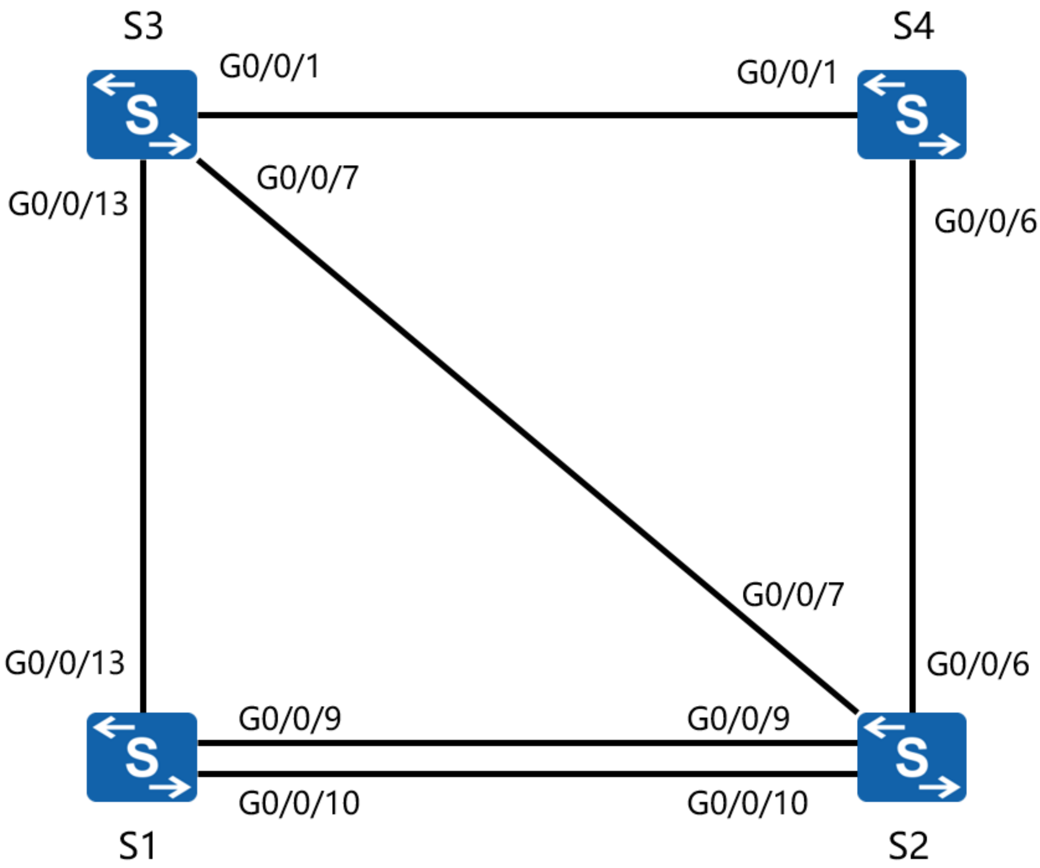

Figure 7-2 Compatibility between MSTP multi-region and STP

Scenario

You are a network administrator of a company. In the Layer 2 network structure, a single spanning tree may congest some VLAN paths and cause the second optimal path to be used. MSTP can address these problems and implement load balancing. In addition, MSTP is compatible with traditional spanning tree modes.

Tasks

Step 1 **Set basic parameters.**

Before this lab test, shut down undesired interfaces.

```
<S1>system-view
```

Enter system view, return user view with Ctrl+Z.

[S1]interface GigabitEthernet 0/0/9

[S1-GigabitEthernet0/0/9]shutdown

[S1-GigabitEthernet0/0/9]quit

<S3>system-view

Enter system view, return user view with Ctrl+Z.

[S3]interface GigabitEthernet 0/0/6

[S3-GigabitEthernet0/0/6]shutdown

[S3-GigabitEthernet0/0/6]quit

<S4>system-view

Enter system view, return user view with Ctrl+Z.

[S4]interface GigabitEthernet 0/0/14

[S4-GigabitEthernet0/0/14]shutdown

[S4-GigabitEthernet0/0/14]quit

Create VLANs 3, 4, 5, 6, 7, and 8 on all switches.

[S1]vlan batch 3 to 8

[S2]vlan batch 3 to 8

[S3]vlan batch 3 to 8

[S4]vlan batch 3 to 8

Check VLAN configurations.

[S1]display vlan

```
* : management-vlan
---------------------
The total number of vlans is : 7
VLAN ID Type        Status   MAC Learning Broadcast/Multicast/Unicast Property
--------------------------------------------------------------------------------
1       common      enable   enable       forward   forward   forward default
3       common      enable   enable       forward   forward   forward default
4       common      enable   enable       forward   forward   forward default
5       common      enable   enable       forward   forward   forward default
6       common      enable   enable       forward   forward   forward default
7       common      enable   enable       forward   forward   forward default
8       common      enable   enable       forward   forward   forward default

[S2]display vlan
* : management-vlan
---------------------
The total number of vlans is : 7
VLAN ID Type        Status   MAC Learning Broadcast/Multicast/Unicast Property
--------------------------------------------------------------------------------
1       common      enable   enable       forward   forward   forward default
3       common      enable   enable       forward   forward   forward default
4       common      enable   enable       forward   forward   forward default
5       common      enable   enable       forward   forward   forward default
6       common      enable   enable       forward   forward   forward default
7       common      enable   enable       forward   forward   forward default
8       common      enable   enable       forward   forward   forward default

[S3]display vlan
* : management-vlan
```

```
--------------------
The total number of vlans is : 7
VLAN ID Type        Status   MAC Learning Broadcast/Multicast/Unicast Property
--------------------------------------------------------------------------------
1       common      enable   enable       forward   forward   forward default
3       common      enable   enable       forward   forward   forward default
4       common      enable   enable       forward   forward   forward default
5       common      enable   enable       forward   forward   forward default
6       common      enable   enable       forward   forward   forward default
7       common      enable   enable       forward   forward   forward default
8       common      enable   enable       forward   forward   forward default

[S4]display vlan
* : management-vlan
--------------------
The total number of vlans is : 7
VLAN ID Type        Status   MAC Learning Broadcast/Multicast/Unicast Property
--------------------------------------------------------------------------------
1       common      enable   enable       forward   forward   forward default
3       common      enable   enable       forward   forward   forward default
4       common      enable   enable       forward   forward   forward default
5       common      enable   enable       forward   forward   forward default
6       common      enable   enable       forward   forward   forward default
7       common      enable   enable       forward   forward   forward default
8       common      enable   enable       forward   forward   forward default
```

Set the access types of the links among all switches to Trunk to receive BPDUs. Allow all VLANs. You do not need to configure the direct link between S2 and S3.

```
[S1]interface GigabitEthernet 0/0/13
[S1-GigabitEthernet0/0/13]port link-type trunk
[S1-GigabitEthernet0/0/13]port trunk allow-pass vlan all
[S1-GigabitEthernet0/0/13]bpdu enable
[S1-GigabitEthernet0/0/13]quit
[S1]interface GigabitEthernet 0/0/10
[S1-GigabitEthernet0/0/10]port link-type trunk
[S1-GigabitEthernet0/0/10]port trunk allow-pass vlan all
[S1-GigabitEthernet0/0/10]bpdu enable
[S1-GigabitEthernet0/0/10]quit

[S2]interface GigabitEthernet 0/0/6
[S2-GigabitEthernet0/0/6]port link-type trunk
[S2-GigabitEthernet0/0/6]port trunk allow-pass vlan all
[S2-GigabitEthernet0/0/6]bpdu enable
[S2-GigabitEthernet0/0/6]quit
[S2]interface GigabitEthernet 0/0/10
[S2-GigabitEthernet0/0/10]port link-type trunk
[S2-GigabitEthernet0/0/10]port trunk allow-pass vlan all
[S2-GigabitEthernet0/0/10]bpdu enable
[S2-GigabitEthernet0/0/10]quit

[S3]interface  GigabitEthernet0/0/1
[S3-GigabitEthernet0/0/1]port link-type trunk
[S3-GigabitEthernet0/0/1]port trunk allow-pass vlan all
[S3-GigabitEthernet0/0/1]bpdu enable
[S3-GigabitEthernet0/0/1]quit
[S3]interface GigabitEthernet0/0/13
[S3-GigabitEthernet0/0/13]port link-type trunk
```

```
[S3-GigabitEthernet0/0/13]port trunk allow-pass vlan all
[S3-GigabitEthernet0/0/13]bpdu enable
[S3-GigabitEthernet0/0/13]quit

[S4]interface  GigabitEthernet0/0/1
[S4-GigabitEthernet0/0/1]port link-type trunk
[S4-GigabitEthernet0/0/1]port trunk allow-pass vlan all
[S4-GigabitEthernet0/0/1]bpdu enable
[S4-GigabitEthernet0/0/1]quit
[S4]interface GigabitEthernet0/0/6
[S4-GigabitEthernet0/0/6]port link-type trunk
[S4-GigabitEthernet0/0/6]port trunk allow-pass vlan all
[S4-GigabitEthernet0/0/6]bpdu enable
[S4-GigabitEthernet0/0/6]quit
```

Step 2 **Configure MST multi-instance.**

Enable MSTP in the system view.

```
[S1]stp enable
[S1]stp mode mstp

[S2]stp enable
[S2]stp mode mstp

[S3]stp enable
[S3]stp mode mstp

[S4]stp enable
[S4]stp mode mstp
```

Allocate all switches to the same region RG1 and set the revision level to 1. Map instance 1 to VLANs 3, 4, and 5. Create instance 2 and map it to VLANs 6, 7, and 8. Activate region configuration.

```
[S1]stp region-configuration
[S1-mst-region]region-name RG1
[S1-mst-region]revision-level 1
[S1-mst-region]instance 1 vlan 3 4 5
[S1-mst-region]instance 2 vlan 6 7 8
[S1-mst-region]active region-configuration
[S1-mst-region]quit

[S2]stp  region-configuration
[S2-mst-region]region-name RG1
[S2-mst-region]revision-level 1
[S2-mst-region]instance 1 vlan 3 4 5
[S2-mst-region]instance 2 vlan 6 7 8
[S2-mst-region]active region-configuration
[S2-mst-region]quit

[S3]stp  region-configuration
[S3-mst-region]region-name RG1
[S3-mst-region]revision-level 1
[S3-mst-region]instance 1 vlan 3 4 5
[S3-mst-region]instance 2 vlan 6 7 8
[S3-mst-region]active region-configuration
[S3-mst-region]quit
```

[S4]stp region-configuration

[S4-mst-region]region-name RG1

[S4-mst-region]revision-level 1

[S4-mst-region]instance 1 vlan 3 4 5

[S4-mst-region]instance 2 vlan 6 7 8

[S4-mst-region]active region-configuration

[S4-mst-region]quit

Check MSTP information.

[S1]display stp brief

MSTID	Port	Role	STP State	Protection
0	GigabitEthernet0/0/10	DESI	FORWARDING	NONE
0	GigabitEthernet0/0/13	DESI	FORWARDING	NONE
1	GigabitEthernet0/0/10	DESI	FORWARDING	NONE
1	GigabitEthernet0/0/13	DESI	FORWARDING	NONE
2	GigabitEthernet0/0/10	DESI	FORWARDING	NONE
2	GigabitEthernet0/0/13	DESI	FORWARDING	NONE

[S2]display stp brief

MSTID	Port	Role	STP State	Protection
0	GigabitEthernet0/0/6	DESI	FORWARDING	NONE
0	GigabitEthernet0/0/10	ROOT	FORWARDING	NONE
1	GigabitEthernet0/0/6	DESI	FORWARDING	NONE
1	GigabitEthernet0/0/10	ROOT	FORWARDING	NONE
2	GigabitEthernet0/0/6	DESI	FORWARDING	NONE
2	GigabitEthernet0/0/10	ROOT	FORWARDING	NONE

[S3]display stp brief

```
 MSTID  Port                        Role  STP State     Protection
   0    GigabitEthernet0/0/1        DESI  FORWARDING      NONE
   0    GigabitEthernet0/0/13       ROOT  FORWARDING      NONE
   1    GigabitEthernet0/0/1        DESI  FORWARDING      NONE
   1    GigabitEthernet0/0/13       ROOT  FORWARDING      NONE
   2    GigabitEthernet0/0/1        DESI  FORWARDING      NONE
   2    GigabitEthernet0/0/13       ROOT  FORWARDING      NONE

[S4]display stp brief
 MSTID  Port                        Role  STP State     Protection
   0    GigabitEthernet0/0/1        ROOT  FORWARDING      NONE
   0    GigabitEthernet0/0/6        ALTE  DISCARDING      NONE
   1    GigabitEthernet0/0/1        ROOT  FORWARDING      NONE
   1    GigabitEthernet0/0/6        ALTE  DISCARDING      NONE
   2    GigabitEthernet0/0/1        ROOT  FORWARDING      NONE
   2    GigabitEthernet0/0/6        ALTE  DISCARDING      NONE
```

S1 is the root switch. S4's E0/0/24 is the backup port for all MST processes.

In instance 2, the S2's priority is 0, the S1's priority is 4096, and the S4's priority is 8192; therefore, S2 becomes the root switch in instance 2.

```
[S2]stp instance 2 priority 0

[S1]stp instance 2 priority 4096

[S4]stp instance 2 priority 8192
```

After the configurations are complete, check the MSTP basic information.

```
[S1]display stp brief
```

```
MSTID  Port                    Role  STP State     Protection
  0    GigabitEthernet0/0/10   DESI  FORWARDING      NONE
  0    GigabitEthernet0/0/13   DESI  FORWARDING      NONE
  1    GigabitEthernet0/0/10   DESI  FORWARDING      NONE
  1    GigabitEthernet0/0/13   DESI  FORWARDING      NONE
  2    GigabitEthernet0/0/10   ROOT  FORWARDING      NONE
  2    GigabitEthernet0/0/13   DESI  FORWARDING      NONE

[S2]display stp brief
MSTID  Port                    Role  STP State     Protection
  0    GigabitEthernet0/0/6    DESI  FORWARDING      NONE
  0    GigabitEthernet0/0/10   ROOT  FORWARDING      NONE
  1    GigabitEthernet0/0/6    DESI  FORWARDING      NONE
  1    GigabitEthernet0/0/10   ROOT  FORWARDING      NONE
  2    GigabitEthernet0/0/6    DESI  FORWARDING      NONE
  2    GigabitEthernet0/0/10   DESI  FORWARDING      NONE

[S3]display stp brief
MSTID  Port                    Role  STP State     Protection
  0    GigabitEthernet0/0/1    DESI  FORWARDING      NONE
  0    GigabitEthernet0/0/13   ROOT  FORWARDING      NONE
  1    GigabitEthernet0/0/1    DESI  FORWARDING      NONE
  1    GigabitEthernet0/0/13   ROOT  FORWARDING      NONE
  2    GigabitEthernet0/0/1    ALTE  DISCARDING      NONE
  2    GigabitEthernet0/0/13   ROOT  FORWARDING      NONE

[S4]display stp brief
MSTID  Port                    Role  STP State     Protection
  0    GigabitEthernet0/0/1    ALTE  DISCARDING      NONE
```

```
0    GigabitEthernet0/0/6          ROOT   FORWARDING     NONE
1    GigabitEthernet0/0/1          ALTE   DISCARDING     NONE
1    GigabitEthernet0/0/6          ROOT   FORWARDING     NONE
2    GigabitEthernet0/0/1          DESI   FORWARDING     NONE
2    GigabitEthernet0/0/6          ROOT   FORWARDING     NONE
```

S2 becomes the root switch in instance 2, and S3's E0/0/1 becomes the alternate port in instance 2. However, the status of switches in instance 1 is not changed. Each MST instance independently calculates the spanning tree.

Step 3 **Configure MST multi-region.**

Delete the MST regions and priority settings of all switches in step 2.

```
[S1]undo stp region-configuration
[S1]undo stp instance 2 priority

[S2]undo stp region-configuration
[S2]undo stp instance 2 priority

[S3]undo stp region-configuration

[S4]undo stp region-configuration
[S4]undo stp instance 2 priority
```

Add S1 and S3 to the same MST region. The region name is RG1 and revision version is 1.

Create instance 1 and map it to VLANs 3, 4, and 5.

Create instance 2 and map it to VLANs 6, 7, and 8.

```
[S1]stp region-configuration
```

```
[S1-mst-region]region-name RG1
[S1-mst-region]revision-level 1
[S1-mst-region]instance 1 vlan 3 4 5
[S1-mst-region]instance 2 vlan 6 7 8
[S1-mst-region]active region-configuration
[S1-mst-region]quit

[S3]stp region-configuration
[S3-mst-region]region-name RG1
[S3-mst-region]revision-level 1
[S3-mst-region]instance 1 vlan 3 4 5
[S3-mst-region]instance 2 vlan 6 7 8
[S3-mst-region]active region-configuration
[S3-mst-region]quit
```

Add S2 and S4 to another MST region. The region name is RG2 and revision version is 2.

Create instance 1 and map it to VLANs 3, 4, and 5.

Create instance 2 and map it to VLANs 6, 7, and 8. Activate all region configurations.

```
[S2]stp region-configuration
[S2-mst-region]region-name RG2
[S2-mst-region]revision-level 2
[S2-mst-region]instance 1 vlan 3 4 5
[S2-mst-region]instance 2 vlan 6 7 8
[S2-mst-region]active region-configuration
[S2-mst-region]quit

[S4]stp region-configuration
```

```
[S4-mst-region]region-name RG2
[S4-mst-region]revision-level 2
[S4-mst-region]instance 1 vlan 3 4 5
[S4-mst-region]instance 2 vlan 6 7 8
[S4-mst-region]active region-configuration
[S4-mst-region]quit
```

After the configurations are complete, check the MSTP basic information.

```
[S1]display stp brief
 MSTID  Port                        Role  STP State     Protection
   0    GigabitEthernet0/0/10       DESI  FORWARDING      NONE
   0    GigabitEthernet0/0/13       DESI  FORWARDING      NONE
   1    GigabitEthernet0/0/10       DESI  FORWARDING      NONE
   1    GigabitEthernet0/0/13       DESI  FORWARDING      NONE
   2    GigabitEthernet0/0/10       DESI  FORWARDING      NONE
   2    GigabitEthernet0/0/13       DESI  FORWARDING      NONE

[S2]display stp brief
 MSTID  Port                        Role  STP State     Protection
   0    GigabitEthernet0/0/6        DESI  FORWARDING      NONE
   0    GigabitEthernet0/0/10       ROOT  FORWARDING      NONE
   1    GigabitEthernet0/0/6        DESI  FORWARDING      NONE
   1    GigabitEthernet0/0/10       MAST  FORWARDING      NONE
   2    GigabitEthernet0/0/6        DESI  FORWARDING      NONE
   2    GigabitEthernet0/0/10       MAST  FORWARDING      NONE

[S3]display stp brief
 MSTID  Port                        Role  STP State     Protection
```

```
  0     Gigabit Ethernet0/0/1           DESI   FORWARDING      NONE
  0     Gigabit Ethernet0/0/13          ROOT   FORWARDING      NONE
  1     Gigabit Ethernet0/0/1           DESI   FORWARDING      NONE
  1     Gigabit Ethernet0/0/13          ROOT   FORWARDING      NONE
  2     Gigabit Ethernet0/0/1           DESI   FORWARDING      NONE
  2     Gigabit Ethernet0/0/13          ROOT   FORWARDING      NONE
```

```
[S4]display stp brief
 MSTID  Port                        Role  STP State     Protection
  0     GigabitEthernet0/0/1        ALTE  DISCARDING    NONE
  0     GigabitEthernet0/0/6        ROOT  FORWARDING    NONE
  1     GigabitEthernet0/0/1        ALTE  DISCARDING    NONE
  1     GigabitEthernet0/0/6        ROOT  FORWARDING    NONE
  2     GigabitEthernet0/0/1        ALTE  DISCARDING    NONE
  2     GigabitEthernet0/0/6        ROOT  FORWARDING    NONE
```

S1 is the root switch and S4's E0/0/1 is the alternate port.

Set the S3 priority in instance 0 to 0 so that S3 becomes the CIST root. Set the S3 priority in instance 1 to 0 so that S3 becomes the root of instance 1. Set the S4 priority in instance 1 to 0 so that S4 becomes the root of instance 1.

```
[S3]stp instance 0 priority 0
[S3]stp instance 1 priority 0


[S4]stp instance 1 priority 0
```

After the configurations are complete, check the MSTP basic information.

```
[S1]display stp brief
 MSTID  Port                        Role  STP State     Protection
```

```
0     GigabitEthernet0/0/10      DESI   FORWARDING      NONE
0     GigabitEthernet0/0/13      ROOT   FORWARDING      NONE
1     GigabitEthernet0/0/10      DESI   FORWARDING      NONE
1     GigabitEthernet0/0/13      ROOT   FORWARDING      NONE
2     GigabitEthernet0/0/10      DESI   FORWARDING      NONE
2     GigabitEthernet0/0/13      DESI   FORWARDING      NONE

[S2]display stp brief
 MSTID  Port                        Role  STP State     Protection
   0    GigabitEthernet0/0/6        DESI  FORWARDING      NONE
   0    GigabitEthernet0/0/10       ROOT  FORWARDING      NONE
   1    GigabitEthernet0/0/6        ROOT  FORWARDING      NONE
   1    GigabitEthernet0/0/10       MAST  FORWARDING      NONE
   2    GigabitEthernet0/0/6        DESI  FORWARDING      NONE
   2    GigabitEthernet0/0/10       MAST  FORWARDING      NONE

[S3]display stp brief
 MSTID  Port                        Role  STP State     Protection
   0    GigabitEthernet0/0/1        DESI  FORWARDING      NONE
   0    GigabitEthernet0/0/13       DESI  FORWARDING      NONE
   1    GigabitEthernet0/0/1        DESI  FORWARDING      NONE
   1    GigabitEthernet0/0/13       DESI  FORWARDING      NONE
   2    GigabitEthernet0/0/1        DESI  FORWARDING      NONE
   2    GigabitEthernet0/0/13       ROOT  FORWARDING      NONE

[S4]display stp brief
 MSTID  Port                        Role  STP State     Protection
   0    GigabitEthernet0/0/1        ALTE  DISCARDING      NONE
   0    GigabitEthernet0/0/6        ROOT  FORWARDING      NONE
```

```
1    GigabitEthernet0/0/1         ALTE   DISCARDING     NONE
1    GigabitEthernet0/0/6         DESI   FORWARDING     NONE
2    GigabitEthernet0/0/1         ALTE   DISCARDING     NONE
2    GigabitEthernet0/0/6         ROOT   FORWARDING     NONE
```

Delete MSTP configuration on S2 and S4, and S2 and S4 to another MST region. The region name is RG2 and revision version is 2. Create instance 1 and map instance 1 to VLANs 6, 7, and 8. Create instance 2 and map it to VLANs 3, 4, and 5. Activate region configuration.

```
[S2]undo stp region-configuration

[S3]undo stp instance 0 priority
[S3]undo stp instance 1 priority

[S4]undo stp region-configuration
[S4]undo stp instance 1 priority

[S2]stp region-configuration
[S2-mst-region]region-name RG2
[S2-mst-region]revision-level 2
[S2-mst-region]instance 1 vlan 6 7 8
[S2-mst-region]instance 2 vlan 3 4 5
[S2-mst-region]active region-configuration
[S2-mst-region]quit

[S4]stp region-configuration
[S4-mst-region]region-name RG2
[S4-mst-region]revision-level 2
[S4-mst-region]instance 1 vlan 6 7 8
```

```
[S4-mst-region]instance 2 vlan 3 4 5
[S4-mst-region]active region-configuration
[S4-mst-region]quit
```

After the configurations are complete, check the MSTP basic information.

```
[S1]display stp brief
 MSTID  Port                    Role  STP State     Protection
   0    GigabitEthernet0/0/10   DESI  FORWARDING      NONE
   0    GigabitEthernet0/0/13   DESI  FORWARDING      NONE
   1    GigabitEthernet0/0/10   DESI  FORWARDING      NONE
   1    GigabitEthernet0/0/13   DESI  FORWARDING      NONE
   2    GigabitEthernet0/0/10   DESI  FORWARDING      NONE
   2    GigabitEthernet0/0/13   DESI  FORWARDING      NONE

[S2]display stp brief
 MSTID  Port                    Role  STP State     Protection
   0    GigabitEthernet0/0/6    DESI  FORWARDING      NONE
   0    GigabitEthernet0/0/10   ROOT  FORWARDING      NONE
   1    GigabitEthernet0/0/6    DESI  FORWARDING      NONE
   1    GigabitEthernet0/0/10   MAST  FORWARDING      NONE
   2    GigabitEthernet0/0/10   MAST  FORWARDING      NONE

[S3]display stp brief
 MSTID  Port                    Role  STP State     Protection
   0    GigabitEthernet0/0/1    DESI  FORWARDING      NONE
   0    GigabitEthernet0/0/13   ROOT  FORWARDING      NONE
   1    GigabitEthernet0/0/1    DESI  FORWARDING      NONE
   1    GigabitEthernet0/0/13   ROOT  FORWARDING      NONE
```

```
  2    GigabitEthernet0/0/1             DESI   FORWARDING       NONE
  2    GigabitEthernet0/0/13            ROOT   FORWARDING       NONE

[S4]display stp brief
 MSTID  Port                        Role  STP State      Protection
  0    GigabitEthernet0/0/1             ALTE   DISCARDING       NONE
  0    GigabitEthernet0/0/6             ROOT   FORWARDING       NONE
  1    GigabitEthernet0/0/1             ALTE   DISCARDING       NONE
  1    GigabitEthernet0/0/6             ROOT   FORWARDING       NONE
  2    GigabitEthernet0/0/1             ALTE   DISCARDING       NONE
```

All MST instances in each region, except instance 0, independently calculate spanning trees, no matter whether repeated VLANs are included and whether VLANs map the instances. The spanning tree calculations in a region do not affect each other.

Step 4 **Configure compatibility between MSTP and STP.**

Add S1, S2, and S3 to the same MST region. Configure STP on S4.

Delete and reconfigure MSTP on S2. The region name is RG1. Create instance 1 and map it to VLANs 3, 4, and 5. Create instance 2 and map it to VLANs 6, 7, and 8. Activate all region configurations.

```
[S2]undo stp region-configuration
[S2]stp region-configuration
[S2-mst-region]region-name RG1
[S2-mst-region]revision-level 1
[S2-mst-region]instance 1 vlan 3 4 5
[S2-mst-region]instance 2 vlan   6 7 8
[S2-mst-region]active region-configuration
[S2-mst-region]quit
```

Enable S2's S0/0/23 and S3's E0/0/23.

Set the access type of the direct link between S2 and S3 to Trunk to receive BPDUs. All VLANs are allowed.

```
[S2]interface  GigabitEthernet 0/0/7
[S2-GigabitEthernet0/0/7]undo shutdown
[S2-GigabitEthernet0/0/7]port link-type trunk
[S2-GigabitEthernet0/0/7]port trunk all vlan all
[S2-GigabitEthernet0/0/7]bpdu enable
[S2-GigabitEthernet0/0/7]quit

[S3]interface  GigabitEthernet 0/0/7
[S3- GigabitEthernet0/0/7]undo shutdown
[S3- GigabitEthernet0/0/7]port link-type trunk
[S3- GigabitEthernet0/0/7]port trunk allow-pass vlan all
[S3- GigabitEthernet0/0/7]bpdu enable
[S3- GigabitEthernet0/0/7]quit
```

Delete MSTP and enable STP on S4.

```
[S4]undo stp region-configuration
[S4]stp mode stp
```

After the configurations are complete, check the STP basic information.

```
[S1]display stp brief
 MSTID  Port                    Role  STP State     Protection
   0    GigabitEthernet0/0/10   DESI  FORWARDING      NONE
   0    GigabitEthernet0/0/13   DESI  FORWARDING      NONE
```

```
 1    GigabitEthernet0/0/10     DESI   FORWARDING     NONE
 1    GigabitEthernet0/0/13     DESI   FORWARDING     NONE
 2    GigabitEthernet0/0/10     DESI   FORWARDING     NONE
 2    GigabitEthernet0/0/13     DESI   FORWARDING     NONE

[S2]display stp brief
 MSTID  Port                     Role  STP State     Protection
   0    GigabitEthernet0/0/6      DESI  FORWARDING    NONE
   0    GigabitEthernet0/0/7      DESI  FORWARDING    NONE
   0    GigabitEthernet0/0/10    ROOT  FORWARDING    NONE
   1    GigabitEthernet0/0/6     DESI  FORWARDING    NONE
   1    GigabitEthernet0/0/7     DESI  FORWARDING    NONE
   1    GigabitEthernet0/0/10    ROOT  FORWARDING    NONE
   2    GigabitEthernet0/0/6     DESI  FORWARDING    NONE
   2    GigabitEthernet0/0/7     DESI  FORWARDING    NONE
   2    GigabitEthernet0/0/10    ROOT  FORWARDING    NONE

[S3]display stp brief
 MSTID  Port                     Role  STP State     Protection
 0    GigabitEthernet0/0/1          DESI  FORWARDING    NONE
 0    GigabitEthernet0/0/7         ALTE  DISCARDING    NONE
 0    GigabitEthernet0/0/13         ROOT  FORWARDING    NONE
 1    GigabitEthernet0/0/1          DESI  FORWARDING    NONE
 1    GigabitEthernet0/0/7         ALTE  DISCARDING    NONE
 1    GigabitEthernet0/0/13         ROOT  FORWARDING    NONE
 2    GigabitEthernet0/0/1          DESI  FORWARDING    NONE
 2    GigabitEthernet0/0/7        ALTE  DISCARDING    NONE
 2    GigabitEthernet0/0/13         ROOT  FORWARDING    NONE
```

```
[S4]display stp brief
 MSTID  Port                    Role  STP State     Protection
   0    GigabitEthernet0/0/1    ROOT  FORWARDING    NONE
   0    GigabitEthernet0/0/6    ALTE  DISCARDING    NONE
```

Instance 0 on S4 running STP and instance 0 on S1, S2, and S3 running MSTP calculate CIST together. In this situation, S1 is the root of CIST.

Set the S4's priority to 4096 so that S4 becomes the root of CIST.

```
[S4]stp priority 4096
```

Check STP basic information.

```
[S1]display stp brief
 MSTID  Port                    Role  STP State     Protection
   0    GigabitEthernet0/0/10   ROOT  FORWARDING    NONE
   0    GigabitEthernet0/0/13   DESI  FORWARDING    NONE
   1    GigabitEthernet0/0/10   DESI  FORWARDING    NONE
   1    GigabitEthernet0/0/13   DESI  FORWARDING    NONE
   2    GigabitEthernet0/0/10   DESI  FORWARDING    NONE
   2    GigabitEthernet0/0/13   DESI  FORWARDING    NONE

[S2]display stp brief
 MSTID  Port                    Role  STP State     Protection

   0    GigabitEthernet0/0/6    ROOT  FORWARDING    NONE
   0    GigabitEthernet0/0/7    DESI  FORWARDING    NONE
   0    GigabitEthernet0/0/10   DESI  FORWARDING    NONE
   1    GigabitEthernet0/0/6    MAST  FORWARDING    NONE
   1    GigabitEthernet0/0/7    DESI  FORWARDING    NONE
```

```
1    GigabitEthernet0/0/10      ROOT   FORWARDING      NONE
2    GigabitEthernet0/0/6       MAST   FORWARDING      NONE
2    GigabitEthernet0/0/7       DESI   FORWARDING      NONE
2    GigabitEthernet0/0/10      ROOT   FORWARDING      NONE
```

```
[S3]display stp brief
MSTID  Port                       Role  STP State      Protection
  0    GigabitEthernet0/0/1       ALTE  DISCARDING     NONE
  0    GigabitEthernet0/0/7       ROOT  FORWARDING     NONE
  0    GigabitEthernet0/0/13      ALTE  DISCARDING     NONE
  1    GigabitEthernet0/0/1       ALTE  DISCARDING     NONE
  1    GigabitEthernet0/0/7       ALTE  DISCARDING     NONE
  1    GigabitEthernet0/0/13      ROOT  FORWARDING     NONE
  2    GigabitEthernet0/0/1       ALTE  DISCARDING     NONE
  2    GigabitEthernet0/0/7       ALTE  DISCARDING     NONE
  2    GigabitEthernet0/0/13      ROOT  FORWARDING     NONE
```

```
[S4]display stp brief
MSTID  Port                       Role  STP State      Protection
 0     GigabitEthernet0/0/1       DESI  FORWARDING     NONE
 0     GigabitEthernet0/0/6       DESI  FORWARDING     NONE
```

S4 becomes the root of CIST, and all ports on S4 are designated ports.

Step 5 **Configure designated port protection.**

Configure designated port protection for E0/0/1 and E0/0/24 of S4.

```
[S4]interface GigabitEthernet0/0/1
[S4-GigabitEthernet0/0/1]stp root-protection
[S4-GigabitEthernet0/0/1]quit
```

[S4]interface GigabitEthernet0/0/6

[S4-GigabitEthernet0/0/6]stp root-protection

[S4-GigabitEthernet0/0/6]quit

Check STP basic information on S4.

[S4]display stp brief

MSTID	Port	Role	STP State	Protection
0	GigabitEthernet0/0/1	DESI	FORWARDING	ROOT
0	GigabitEthernet0/0/6	DESI	DISCARDING	ROOT

Set the priority of instance 0 on S2 to 0 to simulate CIST root preemption.

[S2]stp instance 0 priority 0

Check STP information on S2 and S4.

[S2]display stp brief

MSTID	Port	Role	STP State	Protection
0	GigabitEthernet0/0/6	DESI	LEARNING	NONE
0	GigabitEthernet0/0/7	DESI	FORWARDING	NONE
0	GigabitEthernet0/0/10	DESI	FORWARDING	NONE
1	GigabitEthernet0/0/6	DESI	LEARNING	NONE
1	GigabitEthernet0/0/7	DESI	LEARNING	NONE
1	GigabitEthernet0/0/10	ROOT	FORWARDING	NONE
2	GigabitEthernet0/0/6	DESI	LEARNING	NONE
2	GigabitEthernet0/0/7	DESI	LEARNING	NONE
2	GigabitEthernet0/0/10	ROOT	FORWARDING	NONE

[S4]display stp brief

MSTID	Port	Role	STP State	Protection
0	GigabitEthernet0/0/1	DESI	DISCARDING	ROOT
0	GigabitEthernet0/0/6	DESI	DISCARDING	ROOT

The S4's ports enter the Discarding state, and do not forward packets. This indicates that the S4's port status does not change and S4 is still the root switch.

Delete the priority setting of instance 0 on S2.

[S2]undo stp instance 0 priority

Check STP information on S2 and S4.

[S2]display stp brief

MSTID	Port	Role	STP State	Protection
0	GigabitEthernet0/0/6	ROOT	FORWARDING	NONE
0	GigabitEthernet0/0/7	DESI	FORWARDING	NONE
0	GigabitEthernet0/0/10	DESI	FORWARDING	NONE
1	GigabitEthernet0/0/6	MAST	FORWARDING	NONE
1	GigabitEthernet0/0/7	DESI	FORWARDING	NONE
1	GigabitEthernet0/0/10	ROOT	FORWARDING	NONE
2	GigabitEthernet0/0/6	MAST	FORWARDING	NONE
2	GigabitEthernet0/0/7	DESI	FORWARDING	NONE
2	GigabitEthernet0/0/10	ROOT	FORWARDING	NONE

[S4]display stp brief

MSTID	Port	Role	STP State	Protection
0	GigabitEthernet0/0/1	DESI	FORWARDING	ROOT
0	GigabitEthernet0/0/6	DESI	FORWARDING	ROOT

The port recovers to the normal state if the port does not receive packets of higher priorities for a fixed period (Max Age, default value 20s).

Step 6 **Configure edge port protection.**

Enable S2's G0/0/9.

```
[S2]interface  GigabitEthernet 0/0/9
[S2-GigabitEthernet0/0/9]undo shutdown
[S2-GigabitEthernet0/0/9]quit
```

Configure S1's G0/0/9 as an edge port. Enable edge port protection globally.

```
[S1]interface  GigabitEthernet 0/0/9
[S1-GigabitEthernet0/0/9]undo shutdown
[S1-GigabitEthernet0/0/9]stp edged-port enable
[S1-GigabitEthernet0/0/9]quit
[S1]stp bpdu-protection
```

Check STP information on S1.

```
[S1]display stp interface GigabitEthernet 0/0/9 brief
 MSTID  Port                        Role  STP State     Protection
   0    GigabitEthernet0/0/9        DESI  FORWARDING      BPDU
```

Enable S1's G0/0/9 so that the edge port can receive BPDUs. Simulate an attack on the switch.

```
[S1]interface  GigabitEthernet 0/0/9
[S1-GigabitEthernet0/0/9]undo shutdown
[S1-GigabitEthernet0/0/9]quit
```

Observe S1.

Dec 21 2011 08:39:51-05:13 S1 %%01IFNET/4/IF_STATE(l)[3]:Interface GigabitEthernet0/0/9 has turned into UP state.

Dec 21 2011 08:39:51-05:13 S1 %%01MSTP/4/BPDU_PROTECTION(l)[4]:This edged-port GigabitEthernet0/0/9 that enabled BPDU-Protection will be shutdown, because it received BPDU packet!

Dec 21 2011 08:39:52-05:13 S1 %%01IFNET/4/IF_STATE(l)[5]:Interface GigabitEthernet0/0/9 has turned into DOWN state.

After edge port protection is configured, the edge port is shut down once it receives a BPDU.

Step 7 **Configure loop prevention.**

Configure loop prevention on S3's E0/0/23.

```
[S3]interface GigabitEthernet0/0/7
[S3-GigabitEthernet0/0/7]stp loop-protection
[S3-GigabitEthernet0/0/7]quit
```

Check STP information on S3's E0/0/23.

```
[S3]display stp interface GigabitEthernet 0/0/7 brief
 MSTID  Port                        Role  STP State     Protection
   0    GigabitEthernet0/0/7        ROOT  FORWARDING    LOOP
   1    GigabitEthernet0/0/7        ALTE  DISCARDING    LOOP
   2    GigabitEthernet0/0/7        ALTE  DISCARDING    LOOP
```

Step 8 **Configure TC BPDU protection.**

Enable TC BPDU protection on S1.

```
[S1]stp tc-protection
```

----End

Additional Exercises: Analysis and Verification

If the switches have the same MSTP region name, can the revision versions be different?

In step 4, if the priority of instance 1 on S3 is changed to 0, how will the port status on four switches be changed?

Device Configurations

```
<S1>display current-configuration
#
!Software Version V200R008C00SPC500
 sysname S1
#
 vlan batch 3 to 8
#
 stp bpdu-protection
 stp tc-protection
#
 stp region-configuration
  region-name RG1
  revision-level 1
  instance 1 vlan 3 to 5
  instance 2 vlan 6 to 8
  active region-configuration
#
interface GigabitEthernet0/0/9
 shutdown
stp edged-port enable
```

```
#
interface GigabitEthernet0/0/10
 port link-type trunk
 port trunk allow-pass vlan 2 to 4094
#
interface GigabitEthernet0/0/13
 port link-type trunk
 port trunk allow-pass vlan 2 to 4094
#
return
<S2>display current-configuration
#
!Software Version V200R008C00SPC500
 sysname S2
#
 vlan batch 3 to 8
#
 stp region-configuration
  region-name RG1
  revision-level 1
  instance 1 vlan 3 to 5
  instance 2 vlan 6 to 8
  active region-configuration
#
interface GigabitEthernet0/0/9
#
interface GigabitEthernet0/0/6
 port link-type trunk
 port trunk allow-pass vlan 2 to 4094
```

```
#
interface GigabitEthernet0/0/7
 port link-type trunk
 port trunk allow-pass vlan 2 to 4094
 #
interface GigabitEthernet0/0/10
 port link-type trunk
 port trunk allow-pass vlan 2 to 4094
  #
return
<S3>display current-configuration
#
!Software Version V200R008C00SPC500
 sysname S3
#
 vlan batch 3 to 8
#
 stp region-configuration
  region-name RG1
  revision-level 1
  instance 1 vlan 3 to 5
  instance 2 vlan 6 to 8
  active region-configuration
#
interface GigabitEthernet0/0/1
 port link-type trunk
 port trunk allow-pass vlan 2 to 4094
#
interface GigabitEthernet0/0/7
```

```
 port link-type trunk
 port trunk allow-pass vlan 2 to 4094
 stp loop-protection
#
interface GigabitEthernet0/0/13
 port link-type trunk
 port trunk allow-pass vlan 2 to 4094
#
return
<S4>display current-configuration
#
!Software Version V200R008C00SPC500
 sysname S4
#
 vlan batch 3 to 8 30
#
 stp mode stp
 stp instance 0 priority 4096
#
interface Vlanif30
 ip address 100.100.100.8 255.255.255.0
#
interface GigabitEthernet0/0/1
 port link-type trunk
 port trunk allow-pass vlan 2 to 4094
 stp root-protection
 undo ntdp enable
 undo ndp enable
#
```

```
interface GigabitEthernet0/0/6
 port link-type trunk
 port trunk allow-pass vlan 2 to 4094
 stp root-protection
 undo ntdp enable
 undo ndp enable
#
interface GigabitEthernet0/0/7
 port link-type access
 port default vlan 30
 undo ntdp enable
 undo ndp enable
 bpdu disable
#
Return
```

Recommendations

- Huawei Learning Website
  - http://learning.huawei.com/en
- Huawei e-Learning
  - https://ilearningx.huawei.com/portal/#/portal/ebg/51
- Huawei Certification
  - http://support.huawei.com/learning/NavigationAction!createNavi?navId=_31&lang=en
- Find Training
  - http://support.huawei.com/learning/NavigationAction!createNavi?navId=_trainingsearch&lang=en

More Information

- Huawei learning APP