华为认证系列教程

HCIP-IERS

部署企业级路由交换网络

实验指导书

HUAWEI

华为技术有限公司

版权声明

华为认证系列教程

HCIP-IERS部署企业级路由交换网络

实验指导书

第2.5版本

华为认证体系介绍

华为认证是华为凭借多年信息通信技术人才培养经验及对行业发展的深刻理解，基于ICT产业链人才职业发展生命周期，以学院化的职业技术认证为指引，搭载华为“云 - 管 - 端”融合技术，推出覆盖IP、IT、CT技术领域的认证体系，是业界唯一的ICT全技术领域认证体系。

基于IP、IT、CT技术，华为公司提供了工程师、资深工程师和专家三类技术认证等级，为ICT从业者提供了层次化的培训认证。华为认证包括10个领域，12个技术方向的认证，是业界唯一覆盖ICT全技术领域的认证体系。

HCIA 是对企业网络初级知识和技能的认证。证明您具备配置和维护小型企业网络的能力。HCIA 认证考查工程师协助设计、部署小型企业网络和基本网络运维的能力。目的是考察企业网络工程师使用华为网络设备搭建小型企业路由交换网络的能力，使之能承载基本的语音、无线、云、安全和存储等网络应用，满足企业对网络的使用需求。 HCIA 定位于企业网络技术领域具备初级知识和技能水平的专业人士。侧重于对初级企业网络技术的考察和认证。具备 HCIA 证书的工程师是公认的具备小型企业网络通用技术和基本设计能力的专业人士。

HCIP-R&S 是对企业网络高级知识和技能的认证。目的是帮助企业网络工程师使用华为网络设备搭建完整的中小型企业网络，并支撑企业所需的语音、无线、云、安全和存储等应用全面地集成到网络之中，满足企业各种应用对网络的使用需求，并提供较高的安全性、可用性和可靠性。HCIP-R&S 定位于企业网络技术领域具备高级知识和技能水平的专业人士。侧重于对中小型企业网络技术的考察和认证。具备 HCIP-R&S 证书的工程师是公认的具备中小型企业网络构建和管理能力的专业人士。

HCIE-R&S 是对企业网络专家级知识和技能的认证。目的是帮助企业网络高级工程师搭建完整的大型复杂企业网络，支撑企业所需的语音、无线、云、安全和存储等应用全面集成到网络之中，满足企业各种应用对网络的使用需求。同时能够提供完整的故障排除能力，可根据企业和网络技术发展来规划企业网络，并提高安全性、可用性和可靠性。HCIE-R&S 定位于企业网络技术领域中具备专家知识和技能水平的专业人士。侧重于对大型复杂企业网络技术的考察和认证。具备 HCIE-R&S 证书的工程师是公认的具备大型复杂企业网络构建、优化和管理能力的专业人士。

华为认证协助您打开行业之窗，开启改变之门，屹立在ICT世界的潮头浪尖！

前言

简介

本书为HCIP-IERS认证培训教程，适用于准备参加HCIP-IERS考试的学员或者希望系统掌握通用路由协议原理以及在华为通用路由平台VRP上的实现的读者。

内容描述

Module 1、2、3分别系统而详尽地介绍内部网关协议OSPF、IS-IS以及外部网关协议BGP的工作原理以及在VRP上的配置和实现。帮助读者全面深入地掌握IPv4路由协议知识；

Module 4简要介绍了组播地址、IGMP、PIM-SM等，帮助读者了解组播基础知识、常用组播协议基本原理以及组播应用。

Module 5通过丰富的实例分析阐述了如何灵活使用各种工具实现路由的控制和选择。帮助读者提高综合规划和灵活使用路由协议的能力。

Module 6详细介绍了VLAN原理及其实现，包括VLAN二层互通和三层路由等，帮助读者全面深入地了解VLAN工作原理及其在VRP上的配置；

Module 7对STP协议进行了详细的描述，包括STP、RSTP、MSTP三部分内容，使读者掌握STP协议工作原理以及在VRP上的实现。

本书引导读者循序渐进地掌握路由技术在华为产品中的实现，读者也可以根据自身情况选择感兴趣的章节阅读。

读者知识背景

为了更好地掌握本书内容，阅读本书的读者应首先具备以下基本条件之一：

1) 参加过HCIA培训

2) 通过HCIA考试

3) 熟悉TCP/IP协议栈工作原理，熟悉IP地址

4) 熟悉以太网交换机基本工作原理

本书常用图标


通用路由器


通用交换机


防火墙


网络云

以太网线缆


串口线缆

实验环境说明

eNSP 安装

1. 下载 eNSP：

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

2. 进入链接后点击下载 eNSP 最新版本

CE.zip	564.58MB	2019/03/08	7104	
CX.zip	405.65MB	2019/03/08	4993	
NE40E.zip	405.69MB	2019/03/08	5541	
NE5000E.zip	405.19MB	2019/03/08	4899	
NE9000.zip	405.48MB	2019/03/08	4792	
USG6000V.zip	344.93MB	2019/03/08	6885	
eNSP V100R003C00SPC100 Setup.zip	542.52MB	2019/03/08	20904	

下载

3. 具体软件安装见下方的软件安装指南

文档名称	下载
eNSP V100R003C00SPC100 版本说明书	
eNSP V100R003C00SPC100 软件安装指南	
eNSP V100R003C00SPC100 Virus Scan Report	

下载

4. 关于在 eNSP 中使用几个特殊设备时的操作：

   1） 只要用到以下设备，都需去官网上下载相应的镜像文件

软件名称	文件大小	发布时间	下载次数	下载
CE.zip	564.58MB	2019/03/08	7104	
CX.zip	405.65MB	2019/03/08	4993	
NE40E.zip	405.69MB	2019/03/08	5541	
NE5000E.zip	405.19MB	2019/03/08	4899	
NE9000.zip	405.48MB	2019/03/08	4792	
USG6000V.zip	344.93MB	2019/03/08	6885	

   2） 例如在 eNSP 中选用 USG6000V 后，右键点击启动：

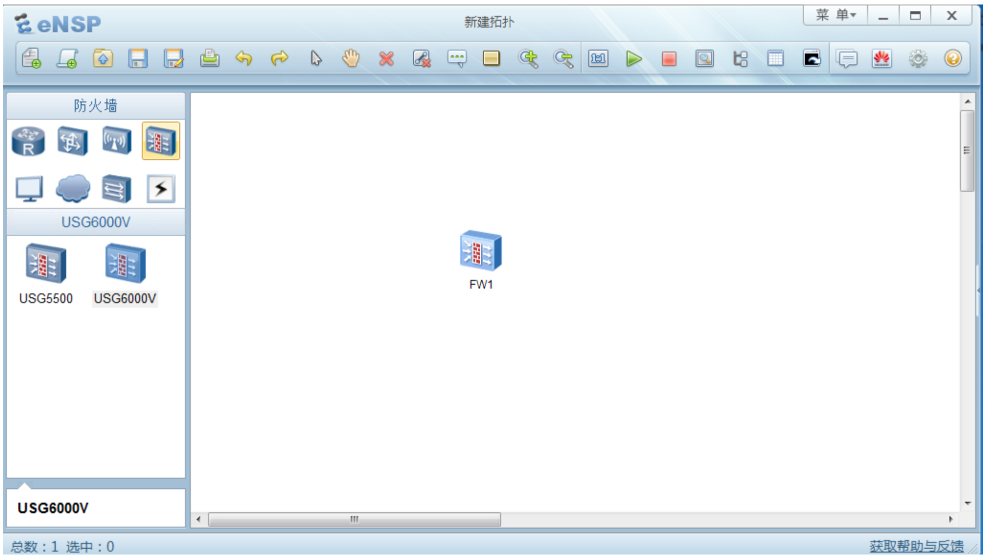

3） 启动设备后会弹出导入设备包的对话框：

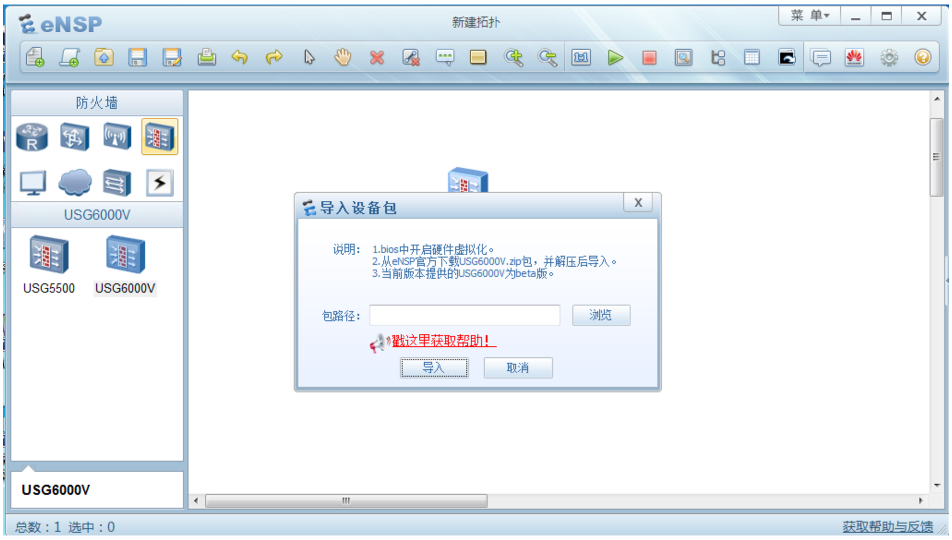

4） 点击“浏览”——找到从官网上下载的镜像文件导入即可。

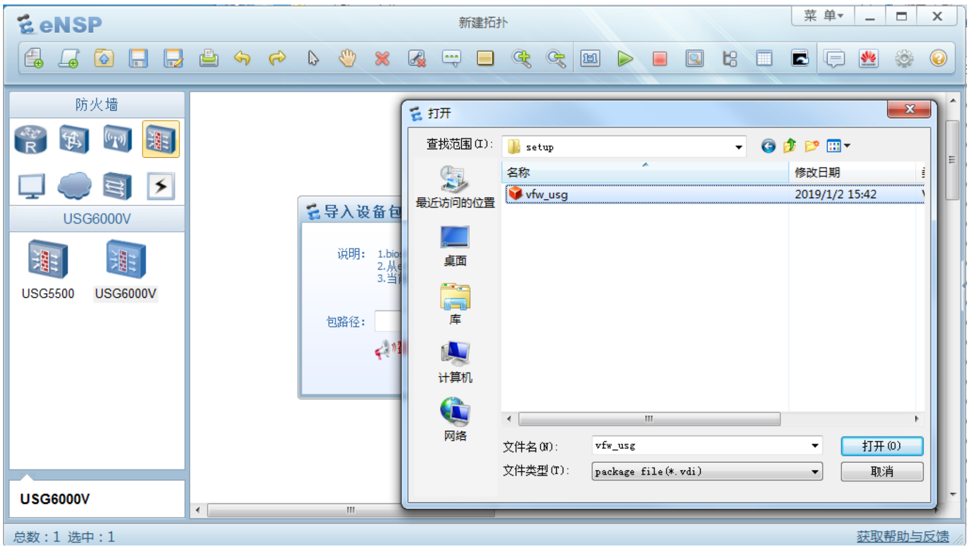

组网介绍

本实验环境面向准备HCIP-IERS考试的网络工程师，实验设备包括路由器5台，交换机4台，防火墙1台。每套实验环境适用于2名学员同时上机操作。

设备介绍

为了满足HCIP-IERS实验需要，建议每套实验环境采用以下配置：

设备名称、型号与版本的对应关系如下：

设备名称	设备型号	软件版本
R1	AR 2220E	V2R7
R2	AR 2220E	V2R7
R3	AR 2220E	V2R7
R4	AR 2220E	V2R7
R5	AR 2220E	V2R7
S1	S5720-36C-EI-AC	V2R8
S2	S5720-36C-EI-AC	V2R8
S3	S5720-36C-EI-AC	V2R8
S4	S5720-36C-EI-AC	V2R8
FW1	USG6306E-AC	V100R001C30

目录

第一章 OSPF协议特性与配置

实验 1-1 OSPF 单区域

学习目的

- 掌握单区域OSPF的配置方法
- 掌握OSPF区域认证的配置方法
- 了解OSPF在多路访问网络邻居关系建立的过程
- 理解OSPF对Loopback接口所连接网络的掩码发布的形式
- 掌握对OSPF接口代价值进行修改的方法
- 掌握OSPF中Silent-interface的配置方法
- 掌握使用Display查看OSPF各种状态的方法
- 掌握使用Debug命令查看OSPF邻接关系和进行故障排除的方法

拓扑图

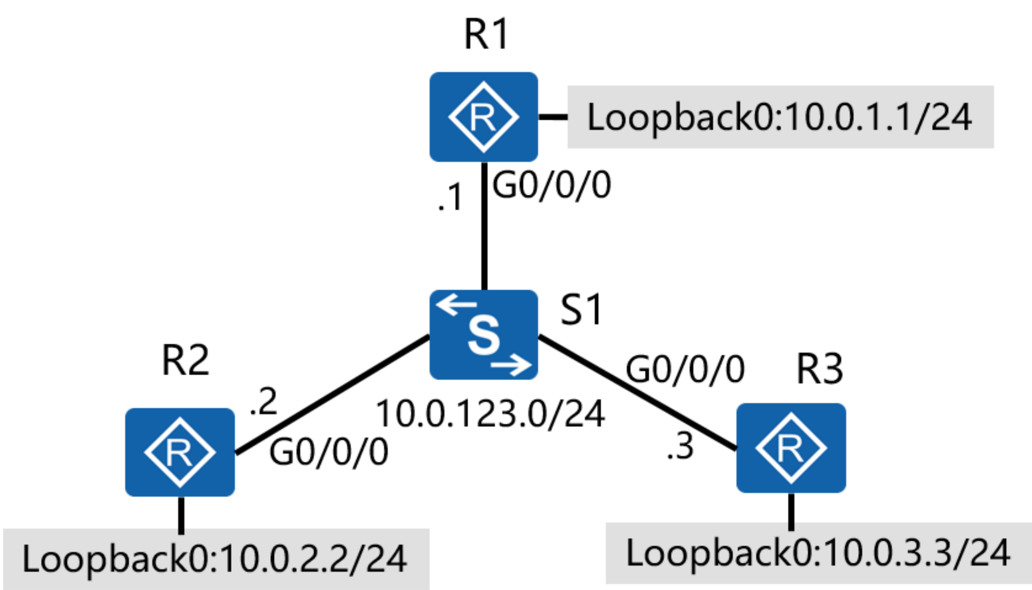

图1-1 OSPF单区域

场景

你是公司的网络管理员。现在公司的网络中有三台ARG3路由器，通过以太网实现相互的连通。在以太网这样的广播式多路访问网络上，可能存在安全隐患，所有你选择采用OSPF区域认证的方法来避免恶意的路由攻击。在部署网络的过程中，出现了网络连通性的问题，你通过使用**display**和**debug**命令进行了故障排除。

学习任务

步骤一. 基础配置与 IP 编址

给R1、R2和R3配置IP地址和掩码。配置时Loopback接口配置掩码为24位，模拟成一个单独的网段。配置完成后，测试直连链路的连通性。

```
<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
```

配置完各接口地址之后验证路由器之间的连通性。

```
[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
```

步骤二. 配置单区域的 OSPF

配置单区域OSPF。所有路由器属于区域0，配置使用OSPF进程1。同时配置区域认证，使用密码“huawei ”。在区域中，华为的设备支持使用明文或MD5值进行认证，在这里，我们仅使用明文进行认证。

注意在使用**network**命令时，通配符掩码使用0.0.0.0。为了保证路由器的Router ID稳定，我们在启动OSPF进程时使用**router-id**参数静态指定路由器的Router ID。

```
[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
```

配置完成后，查看设备的路由表，并测试全网的连通性。我们首先在R1上查看路由表。

```
 [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
```

从输出中我们可以看到R1从OSPF学习到了2条路由，10.0.2.2/32和10.0.3.3/32，下一跳分别是10.0.123.2和10.0.123.3。然后分别检查从R1到达R2及R3的Loopback地址的连通性。

```
[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
```

使用display ospf brief命令查看路由器运行的基本OSPF信息。

```
[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
```

从上面的输出中我们可以看到区域0开启了明文认证(Authtype: Simple) ,共有两个接口参加了OSPF的运行 :GigabitEthernet0/0/0和LoopBack0。其中，GigabitEthernet0/0/0为广播型网络（Broadcast），开销（Cost）为1，优先级(Priority)为1 ,R1自己的角色为DR ,后面列出了该网络上的BDR(10.0.123.2)。另外一个运行OSPF的接口LoopBack0的网络类型为P2P。

使用display ospf peer brief命令查看路由器的OSPF邻居关系建立情况。

```
[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
```

从上面的输出中我们可以看到在区域0.0.0.0中，R1有两个邻居，邻居的Router ID分别为10.0.2.2和10.0.3.3，他们之间的状态为Full。

使用display ospf lsdb命令查看路由器的OSPF数据库信息。

```
[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
 Network      10.0.123.110.0.1.1            1567   36     80000004       0
```

在这里我们一共可以看到4条LSA，前3条为第一类LSA，分别由R1、R2和R3产生，我们可以通过AdvRouter判断该LSA是由哪台路由器生成的。第四条为第二类LSA，是由一个网段的DR产生的。在这里，R1是10.0.123.0/24这个网段的DR，所以我们可以看到这条LSA的AdvRouter为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
```

上面的输出是R1产生的Router LSA的详细信息，我们可以看到这条LSA一共描述了2个网络，第一个网络为Loopback接口所在网段，链路类型为StubNet，Link ID和Data分别是该Stub网段的IP地址和掩码。第二个网络为三台路由器的互联网段，链路类型为TransNet，可以看到Link ID为DR的端口地址10.0.123.1，Data为该网段上本地端口的IP地址10.0.123.1；

```
[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
```

上面的输出是R1产生的Network LSA的详细信息，我们可以看到第二类LSA描述了DR所在网段的邻居信息。

步骤三.　观察路由器在以太网上邻接关系的建立过程

首先查看在10.0.123.0/24网段，OSPF邻居关系中DR和BDR选举的情况，并分析为什么会这样？以及是否所有人在做这个实验时，结果都是一样的？

我们首先查看在10.0.123.0/24网段，OSPF邻居关系中DR和BDR选举的情

况。从下面的输出中，我们可以得知现在该网段的DR的接口IP为10.0.123.1，BDR的接口IP为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 ]
```

有可能每个人的实验结果输出不一样。因为在OSPF中，DR的选举不是抢占的，即网络中存在DR或BDR时，新进入网络的路由器不能抢占DR或BDR的角色。在这个网络中，先启动OSPF进程或先接入该网络的路由器成为了该网段上的DR，其他路由器成为的BDR或DROther。

当DR发生故障后，BDR就会接替DR的位置，我们在实验中可以通过重置OSPF进程的方法来观察DR角色的改变，在这里，我们重置R1的OSPF进程。

```
<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 ]
```

当重置R1的OSPF进程以后 ,原来该网络上的BDR 10.0.123.2成为了新的DR，原来的DROther 10.0.123.3成为了新的BDR。

下面我们关闭R1、R2与R3的G0/0/0接口，使用命令**debugging ospf 1 event**准备查看OSPF邻接关系建立的具体过程。然后尽量同时打开R1、R2与R3的G0/0/0接口。观察在广播式多路访问网络上邻居状态的变化过程和DR和BDR选举的过程。

```
<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
```

在R2和R3上进行相同的操作，查看R3的debug信息。由于所有路由器默认的接口优先级都是1，所以在DR选举的时候会参考路由器的Router ID，在这三台路由器中，R3的Router ID是最大的，所以R3成为了该网段上的DR。

```
[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
```

当刚打开接口时，接口状态由Down变为Waiting，此时路由器开始交互Hello数据包，等待约40秒以后，R3的接口由Waiting变为DR。

步骤四. 配置 OSPF 中 Loopback 接口的网络类型

观察R1的路由表，关注这两条路由：10.0.2.2/32和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
```

在配置R2和R3的Loopback接口地址时，使用的掩码是24位，分析为什么这里路由表中显示的是32位掩码的路由？

使用命令**display ospf interface LoopBack 0 verbose**查看OSPF在Loopback 0接口运行的状态信息。

```
[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
```

可以看到对于Loopback接口，OSPF知道该网段只可能有一个IP地址，所以发布的路由的子网掩码是32位的。

修改R2的Loopback0接口的网络类型为Broadcast，OSPF在发布这个接口的网络信息时，就会使用24位掩码进行发布。

```
[R2]interface LoopBack 0
[R2-LoopBack0]ospf network-type broadcast
```

这时我们看到R2发布的Loopback 0地址的路由子网掩码为24位。

```
[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
```

使用命令**display ospf interface LoopBack 0 verbose**查看Loopback接口的运行状态可以看到，该接口网络类型为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
```

步骤五. 修改接口的 OSPF 代价值

首先在R1上查看R1到达R3的Loopback0接口路由的代价值，我们可以看到到达10.0.3.3/32的代价值为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
127.255.255.255/32Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct  0    0           D   127.0.0.1       InLoopBack0
```

修改R1的G0/0/0接口代价值为20，修改R3的G0/0/0接口代价值为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
```

重新查看R1到达R3的Loopback0接口路由的代价值，可以看到，到达10.0.3.3/32的代价值为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   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
127.255.255.255/32Direct  0    0          D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct  0    0          D   127.0.0.1       InLoopBack0
```

在R3上查看10.0.1.1/32的代价值，可以看到值为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
127.255.255.255/32Direct   0    0         D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct   0    0         D   127.0.0.1       InLoopBack0
```

步骤六. 配置 OSPF 的 Silent-interface

配置R1的G0/0/0接口为Silent-interface。

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

查看R1的邻居关系建立和路由表学习情况可发现，路由表中从OSPF学习到的路由条目消失了。

```
[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
127.255.255.255/32Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32Direct  0    0           D   127.0.0.1       InLoopBack0
```

查看R1的邻居列表可以看到R1和R2、R3之间的邻居关系也消失了。在RIP中将一个接口置为Silent-interface以后，该接口不再发送RIP更新；但在OSPF中，路由器之间需要建立邻居关系之后才会交互路由信息，当一个接口被设置为Silent-interface以后，该接口不再接收或发送Hello包，造成该接口不能和其他路由器形成邻居关系。

```
 [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
```

恢复R1的G0/0/0接口为默认状态，将三个路由器的Loopback0接口配置为Silent-interface。

```
[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
```

检查R1的路由表可见，将Loopback设为Silent-interface以后不影响该接口路由的发布。

```
[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
```

附加实验：思考并验证

为什么在配置OSPF时，使用的通配符掩码是0.0.0.0，实际的配置中，也可以使用通配符掩码0.0.0.255，思考一下，这两种表达形式有什么差异？

分析在实际的网络中，哪些类型的接口应该配置为Silent-interface接口？

最终设备配置

```
<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
```

实验 1-2 OSPF 多区域

学习目的

- 掌握OSPF配置指定Router ID的方法
- 掌握多区域OSPF的配置方法
- 掌握OSPF区域之间路由汇总的配置方法
- 掌握OSPF参考带宽的配置方法
- 掌握OSPF引入外部路由的配置方法
- 掌握OSPF引入的外部路由时进行路由汇总的方法
- 掌握向OSPF导入缺省路由的方法
- 掌握对OSPF中各类路由的管理距离的修改方法

拓扑图

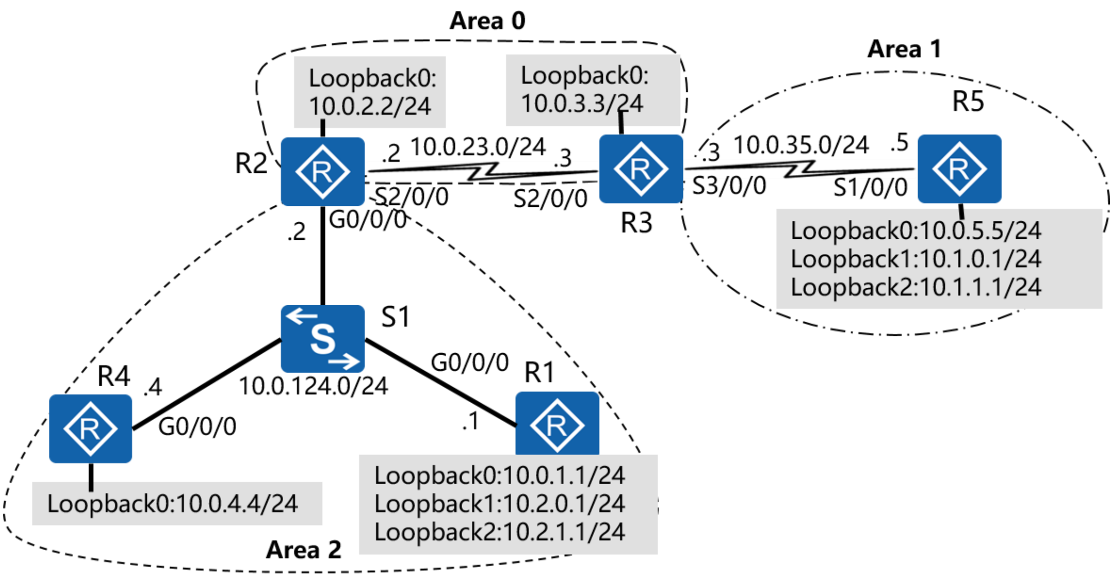

图1-2 OSPF多区域

场景

你是公司的网络管理员。现在公司的网络中有五台ARG3路由器，其中R1、R2和R4在公司总部，通过以太网互联。R3、与R5在公司分部，R3通过专线与公司总部的R2相连，R5与R3之间也通过专线相连。由于网络规模较大，为了控制LSA的洪泛，你设计了多区域的OSPF互联方式。

其中R2与R3的Loopback0接口、互联接口属于区域0；R3与R5互联的网段、R5的Loopback0/1/2接口属于区域1；R1、R2与R4互联的网段以及R1、R4的Loopback0接口属于区域2。

同时为了明确设备的Router-ID，你配置设备使用固定的地址作为Router ID。

为了使路由器进行路由转发时效率更高，你在区域的边界配置了自动汇总。

R1路由器连接到公司以外的网络，你配置将这些OSPF区域之外的路由信息引入到OSPF区域。

R4路由器连接到Internet，你需要配置一条缺省路由，引入到OSPF区域，以便于OSPF区域的所有路由器都知道如何访问Internet。

同时OSPF路由信息中区分了内部路由和外部路由，你修改了OSPF路由信息的优先级信息，以避免潜在的风险。

OSPF中特定路由信息的度量值是将到达目的网络经过的所有链路的代价值进行累加得到的。而链路的代价值是路由器将接口带宽与参考带宽进行对比得到。参考带宽值为100Mbps，实际接口带宽可能为1000Mbps，而度量值都是整数，所以快速以太网接口和千兆以太网接口的OSPF代价值均为1。为了能够相互区分这些链路，你定义参考带宽值为10Gbps。

在配置设备的同时，出现了一些网络故障，你通过使用**display**和**debug**命令进行了故障排除。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置IP地址和掩码。配置时注意所有的Loopback接口配置掩码均为24位，模拟成一个单独的网段。

```
<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
```

配置完成后，测试直连链路的连通性。

```
[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
```

步骤二. 配置多区域 OSPF

为保证OSPF的Router ID稳定，我们通常手工指定路由器的Router ID。有2种方法可以手工指定一台路由器运行OSPF的Router ID，第一种方式是在系统视图下使用**router id**的命令。

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

第二种方式是在启动OSPF进程时加上参数**router-id**。

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

当路由器上同时配置了这两条命令以后，路由器最终会选取第二种方式配置的值作为Router ID。如果在一台路由器上需要起多个OSPF进程，且每个OSPF进程的Router ID需要不一样时，我们只能使用第二种方式来指定Router ID。

在R1上配置Loopback 0接口及GigabitEthernet 0/0/0属于区域2。这里我们将所有OSPF区域的Loopback接口，修改其OSPF网络类型为Broadcast类型，以便于OSPF发布Loopback口的真实掩码信息。

```
[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
```

在R2上配置Loopback 0和Serial 2/0/0接口属于区域0，GigabitEthernet

0/0/0属于区域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
```

在R3上配置Loopback 0和Serial 2/0/0接口属于区域0，Serial 3/0/0属于区域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
```

在R4上配置Loopback 0及GigabitEthernet 0/0/0属于区域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
```

在R5上配置所有的Loopback接口及Serial 1/0/0属于区域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
```

配置完成后，在R1上查看路由表。

```
[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
```

该路由器上已拥有全网所有的路由条目。

在R1上测试到其他路由器Loopback接口的连通性。

```
[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
```

我们使用**display ospf brief**命令在R2上查看路由器运行的基本OSPF信息。

```
[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
```

第一行Border Router: AREA 表示该路由器是一台ABR；如果路由器是一台区域内路由器，该值为空；如果路由器是一台ASBR，该值为 AS。

该路由器共有三个接口参加OSPF运算，我们已手工将Loopback 0接口的网络类型修改为Broadcast。Serial2/0/0的封装类型为PPP，所以默认的网络类型为点对点。另外GigabitEthernet 0/0/0连接到区域2，是广播型网络。

我们在R2上使用**display ospf peer brief**命令查看路由器的OSPF邻居关系建立情况。可以看到，在区域0，R2有一个邻居10.0.3.3，在区域2，R2有2个邻居：10.0.1.1和10.0.4.4，R2与他们都形成了邻接关系（Full）。

```
[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
 ----------------------------------------------------------------------------
```

我们在R2上使用**display ospf lsdb**命令查看路由器的OSPF数据库信息。我们可以发现由于R2是一台ABR，所以在该路由器上维护了2个LSDB，分别用来描述区域0和区域2的路由。

```
[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
```

步骤三. 配置 OSPF 区域之间的路由汇总

首先查看R2和R3的OSPF路由表。

```
[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
```

10.1.0.0/24和10.1.1.0/24两条路由信息均以详细条目出现。

对于这样的路由信息，可以进行汇总，再向其他区域发送。一方面减少其他区域的路由表条目，另外一方面还可以减少路由振荡情况的发生。我们可在R3上使用**abr-summary**的命令将R5的Loopback1和Loopback2接口的网段进行汇总发送。

```
[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
```

配置完成后在R3和R2上分别查看汇总路由信息。

```
[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
```

通过输出可以看到，在R3的路由表中，这2条路由仍以明细路由的形式出现，在R2上，仅存在汇总路由10.1.0.0/23。

配置完成后，测试其他路由器与网络10.1.0.0/24与10.1.1.0/24的连通性。

```
[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
```

步骤四. 修改 OSPF 的参考带宽值

在实际网络我们可能使用了千兆甚至万兆以太网。但是由于OSPF的默认参考带宽值为100Mbps，并且接口代价值仅为整数，所以OSPF无法在带宽上区分百兆以太网和千兆以太网。

在R2上修改OSPF的参考带宽值为10Gbps。这里，使用命令**bandwidth-reference**进行修改，相应带宽参数值的单位为Mbps。

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

在R2上查看OSPF邻居关系，以及路由信息学习情况，我们可以看到，在路由表中，Cost值已经发生了变化。

```
[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
```

在运行OSPF的多个区域内，OSPF的参考带宽值必须一致，否则OSPF无法正常工作。修改所有路由器的OSPF参考带宽值为10Gbps。

```
[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
```

在R2上查看邻居列表、路由表，观察OSPF邻居关系以及路由信息是否正常。

```
[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
```

如上所示，路由信息正常。另外可测试网络的连通性。

步骤五. 配置将直连路由汇总并引入到 OSPF 区域

R1的Loopback1和Loopback2接口不属于OSPF区域。将这两条直连路由引入到OSPF区域，并在R1上执行路由汇总。

```
[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
```

在R1上查看外部路由信息。

```
[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通过一条第五类LSA向其他路由器通告了网段10.2.0.0，子网掩码是255.255.254.0。

在其他路由器上查看汇总路由，并测试网络连通性。

```
[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
```

在R2上可以看到一条掩码为23位的汇总路由。

将R1的Loopback 2接口删除，查看R2上路由条目变化情况。我们可以看到，当Loopback 2接口不存在了，汇总路由仍然存在。

```
[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
```

在R5设备上向10.2.1.1地址发送Tracert数据包。

```
<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  *  *  *
...
```

我们可以看到虽然Loopback 2接口被删除了，到达该目的地址的数据包仍然被R2和R3转发，直到R1上该数据包被丢弃。

步骤六.　OSPF 引入缺省路由

R4的Loopback0接口连接到Internet。在R4上配置缺省路由，下一跳指向Loopback0。

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

将这条缺省路由引入到OSPF区域，定义类型为1，Cost值为10，并且定义为永久引入。

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

在R2上查看缺省路由的学习情况。我们可以看到R2通过第五类LSA学习到了一条默认路由，下一跳是R4的接口地址。

```
[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
```

最后测试R5路由器与10.0.4.4之间的连通性。

```
[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
```

步骤七. 修改 OSPF 中两类路由的优先级

查看R1的路由表，关注OSPF不同类型路由的优先级信息。

```
[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
```

默认情况下，OSPF区域内和区域之间的路由，优先级为10。OSPF外部路由，优先级为150。

修改R1和R4路由器上的OSPF区域内和区域之间的路由优先级为20，修改OSPF外部路由的优先级为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
```

查看路由表中OSPF内部路由及外部路由的优先级，确认已修改成功。

```
[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
```

路由优先级仅在本地有效，用于衡量在本地通过多种方式学到的路由的优略程度。本地区域的不同路由器，如果优先级信息不同，也可以正常工作。

附加实验：思考并验证

思考在步骤六中，定义缺省路由的永久发布的作用是什么？有哪些优点和缺点？

路由汇总就像一把双刃剑，有好处也有坏处。思考并总结使用路由汇总的好处和坏处，并分析如何避免这些坏处。

最终设备配置

```
<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

实验 1-3 OSPF 的邻接关系和 LSA

学习目的

- 了解四个OSPF邻居路由器在以太网上邻居关系建立的过程
- 掌握对OSPF的DR的选举进行干预的方法
- 观察5种类型的LSA的内容，以及它们的作用
- 了解OSPF的LSR、LSU、LSAck数据包的相互发送情况

拓扑图

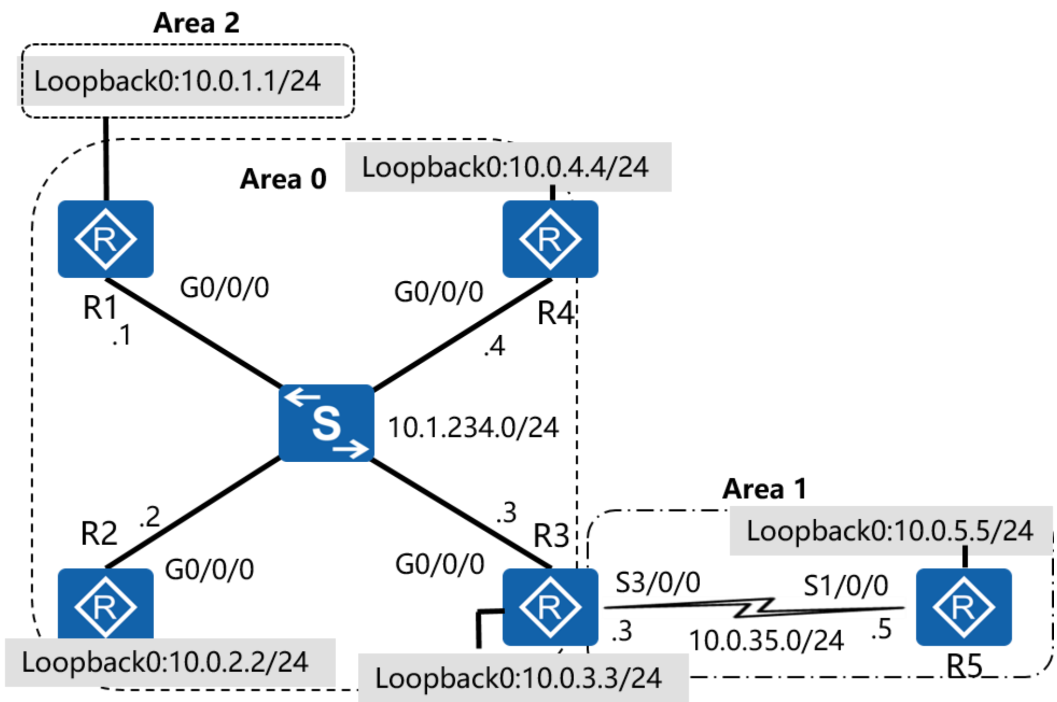

图1-3 OSPF的邻接关系和LSA

场景

你是公司的网络管理员。现在公司的网络中有五台AR G3路由器，其中R1、R2、R3和R4在公司总部，通过以太网互联。R5在公司分部，R3通过专线与公

司总部的R3相连。由于网络规模较大，为了控制LSA的洪泛，你设计了多区域的OSPF互联方式。

其中R1的Loopback0接口属于区域2。R2、R3、R4的Loopback0接口与10.1.234.0/24网段属于区域0。R3与R5之间互联的网段属于区域1。R5的Loopback0接口属于OSPF外部网络。

同时为了明确设备的Router ID，你配置设备使用固定的地址作为Router ID。

在R1、R2、R3与R4之间互联的网络上，需要干预DR与BDR的选举。实际使用中将R3定义为DR、R2定义为BDR。R4设备定义为DROther。

学习任务

步骤一.　基础配置与 IP 编址

给所有路由器配置IP地址和掩码。配置时注意所有的Loopback接口配置掩码均为24位，模拟成一个单独的网段。

```
<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
```

配置完成后，测试直连链路的连通性。

```
[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
```

步骤二. 配置多区域 OSPF

在R1上配置GigabitEthernet 0/0/0属于区域0，Loopback 0属于区域2。对所有OSPF区域的Loopback接口，修改其OSPF网络类型为Broadcast类型，以便于OSPF发布Loopback口的真实掩码信息。

```
[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
```

R2、R4的所有接口均位于区域0中。

```
[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
```

在R3上配置Loopback 0和GigabitEthernet 0/0/0属于区域0，Serial 3/0/0属于区域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
```

在R5上配置Serial 1/0/0属于区域1，Loopback 0不属于任何区域。

```
[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
```

配置完成后，在R1查看设备的路由表。

```
[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
```

除了没有发布进OSPF的网络10.0.5.5/24，在R1上已拥有全网的路由表。

测试网络的连通性。

```
[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
```

在R1使用**display ospf brief**命令查看路由器运行的基本OSPF信息。我们可以看到，由于R1的Loopback 0接口位于区域2中，所以R1成为了一台ABR。R1的GigabitEthernet 0/0/0接口所连接的网络为广播型网络，且R1为这个网段的DR。

```
[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
```

使用**display ospf peer brief**命令查看路由器的OSPF邻居关系建立情况。由于R1是DR，所以它与该网段的所有路由器形成邻接关系。在R3上查看邻居就可以发现R3和R4之间仅存在邻居关系，而没有邻接关系。

```
[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
 ----------------------------------------------------------------------------
```

在R5上使用**display ospf lsdb**命令查看路由器的OSPF数据库信息。

```
[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
```

可以看到由于在区域1中仅存在2台路由器，所以在R5的lsdb中，仅存在2条第一类LSA，剩余的5条第三类LSA是由R3向R5通告的区域间路由。

在R2上使用**display ospf lsdb**命令查看路由器的OSPF数据库信息。

```
[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
```

在R2上除了4条第一类LSA以外，还有一条第二类LSA。R2的GigabitEthernet 0/0/0所连接的是一个广播型网络，该网络上的DR会产生一条第二类LSA来描述所有的邻居。在这里可以从AdvRouter字段得知生成这条LSA的路由器是R1，符合R1是该网段DR的结果。

在R1上使用**display ospf lsdb**命令查看路由器的OSPF数据库信息。

```
[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
```

由于R1的Loopback 0接口位于区域2中，所以R1上有2个区域的LSDB，分别是区域0和区域2的。

在R4上使用**display ospf lsdb**命令查看路由器的OSPF数据库信息。

```
[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
```

注意由于OSPF路由器的角色差异，OSPF链路状态数据库内容也会有所差异。对比分析R5、R2、R1与R4链路状态数据库的差异。

步骤三. 修改路由器 OSPF 接口优先级，影响 DR 选举

配置R3的G0/0/0接口优先级为255，确保R3成为10.1.234.0/24网段的DR。修改R2的G0/0/0接口优先级为254，确保R2成为10.1.234.0/24网段的BDR。修改R4的G0/0/0接口优先级为0，确保R4不参加DR/BDR选举，而成为10.1.234.0/24网段的DROther。

```
[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
```

配置完成后，由于DR/BDR已经选举，并且DR/BDR角色不能抢占。所以必须关闭R1、R2、R3、R4的G0/0/0接口，并依次打开R3、R2、R1和R4的G0/0/0接口。

```
[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
```

查看网段10.1.234.0/24网段的DR/BDR选举情况。

```
[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 ]
```

在重启接口后R3成为了该网段的DR，R2成为了BDR。

查看R4与R1的邻居关系。

```
[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 ]
```

当邻居关系稳定以后，由于R1和R4均为DROther路由器，所以他们之间仅形成邻居关系，保持在2-way状态。

步骤四. 配置将直连路由汇总并引入到 OSPF 区域

R5的Loopback0接口不属于OSPF区域。将这条直连路由引入到OSPF区域。

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

在R1和R3上查看引入的外部路由。

```
[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
```

在R1和R3上均看到2条外部路由，分别是10.0.5.0/24和10.0.35.3/32。10.0.5.0/24为R5的Loopback接口，但为什么还有一条10.0.35.3/32呢？

查看R5的路由表，由于R3和R5之间是以PPP的形式封装的，R3的Serial 3/0/0的接口地址会以直连路由的形式出现在R5的路由表里，所以在R5上运行**import-route direct**以后该路由条目也被发布出去了（下面的输出略去了其他路由条目）。

```
[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
```

最后测试网络连通性。

```
[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
```

在R1上查看OSPF外部路由在链路状态数据库中的情况。我们可以看到，R1的LSDB中一共有3条外部路由：10.0.5.0/24、10.0.35.0/24、10.0.35.3/32。

在R1的路由表中看见的外部路由只有2条，另一条不见了。

```
 [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
```

经过比较后，可以发现10.0.35.0/24这条路由是以内部路由的形式出现在路由表中的。

检查R1的LSDB中得第三类LSA就可以看到这个条目：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
```

可以看出，当第三类和第五类LSA通告路由的网络位和掩码相同的情况下，OSPF优选第三类LSA通告的路由加到路由表里。

步骤五. 查看各种类型的 LSA

在R1上查看一类LSA 10.0.1.0在Area0和Area2的详细内容。

```
[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
```

对于一类LSA来说，Ls id字段表示生成这条LSA的路由器的Router ID。

R1共生成了两条第一类LSA，一条在区域0中泛洪。R1在区域0中与一个Transit网段相连，所以Link Type字段为TransNet。对于TransNet，Link ID字段为该网段上DR的接口IP地址，Data字段为本地接口的IP地址。

第二条一类LSA是R1向区域2中泛洪的，R1与区域2通过Loopback接口相连。对于Loopback接口，Link Type为StubNet，此时Link ID表示该Stub网段的IP网络地址，Data表示该Stub网段的网络掩码。

在R2、R3和R4上分别查看二类LSA 10.1.234.0在Area0的详细内容。

```
[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
```

发现在R2、R3和R4上看到的这条LSA是一样的。

同样，可以通过Adv rtr字段得知这条LSA是由R3生成的。第二类LSA 的Ls id描述的是该网段上DR的接口IP地址，Attached Router为该网段上所有路由器的Router ID。

在R1和R3上查看三类LSA 10.0.35.0/24在Area0的详细内容。

```
[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
```

从输出中可以看到该路由是由R3向区域0中通告的。Ls id就是通告的目的网段的网络地址，Net mask描述了目的网段的掩码信息。

```
[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
```

在R1上共有2条描述10.0.35.0/24的第三类LSA。其中，从Adv rtr字段我们可以得知区域0中这条LSA是R3产生的。由于R1本身也是一台ABR，所以R1收到这条LSA以后又产生了一条LSA，向区域2中通告。

在R1上查看四类LSA 10.0.5.0在Area2的详细内容。第四类LSA用于描述如何到达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
```

从输出中可以看到，R1从R3收到了一条第四类LSA。Ls id用于描述ASBR的Router ID。由于这类LSA不能跨区域泛洪，所以R1又生成了一条第四类LSA向区域2中泛洪。

在R2、R4以及R3的区域0的LSDB中均存在这条LSA，因为这些路由器和ASBR（R5）不在同一个区域，他们需要通过第四类LSA来得知ASBR的位置。

```
[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
```

在区域1中就没有这条第四类LSA，同一个区域的路由器，不需要依赖这条LSA来得知ASBR的位置。

步骤六.　观察 LSR、LSU 和 LSAck

我们首先观察OSPF的Update数据包及ACK数据包发送的过程。在R1上打开**debugging ospf packet update**、**debugging ospf packet ack**。

```
<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
```

默认情况下，网络稳定运行，OSPF路由器每30分钟更新一次。为触发查询和更新信息，我们将R3的Loopback 0接口删除。

```
[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
```

我们可以观察到首先在R1上接收到10.1.234.3发来的Update消息，消息的目的地址为224.0.0.5（即所有OSPF路由器），描述了一个网段（# Links: 1），后面有该网段的LinkID和LinkData。

```
<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
```

最后是R1自己发送的ACK报文。报文源地址为R1 GigabitEthernet 0/0/0的接口地址，目的地址为224.0.0.6。该报文是发送给DR和BDR的。该报文的序列号也是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
```

接下来恢复R3上删除的Loopback0接口。

```
[R3]interface loopback 0
[R3-LoopBack0]ip address 10.0.3.3 24
[R3-LoopBack0]quit
```

和刚才一样，R1首先收到来自R3的Update报文，但这次在报文中通告了一个新的网段，所以# Links这里值为2，后面有新通告的网段的网络号和掩码。

```
<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首先收到BDR的ACK报文。

```
<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

最后是R1自己发送的ACK报文。

<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
```

在下面一个步骤中我们看Request报文。正常情况下，路由器不会主动发送该报文，为观察该报文的发送，我们将R1的OSPF进程重启。在路由器上观察到的是R1向R2发起了LS Request。

```
<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

随后R1收到了来自R3的LS Request。

<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
```

附加实验：思考并验证

假设区域2存在一台路由器R6。它计算到达10.0.5.0/24网段的路由信息与R2、R3计算该信息的步骤有什么差异？

类型4的LSA什么时候会出现？

实验中如果将R1和R4都配置成DROther，会有什么隐患？

最终设备配置

```
<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

实验 1-4 OSPF Stub 区域与 NSSA 区域

学习目的

- 掌握OSPF的Stub区域的配置
- 掌握OSPF的NSSA区域的配置
- 观察LSA Type7的内容
- 理解LSA Type7与Type5之间的转化关系

拓扑图

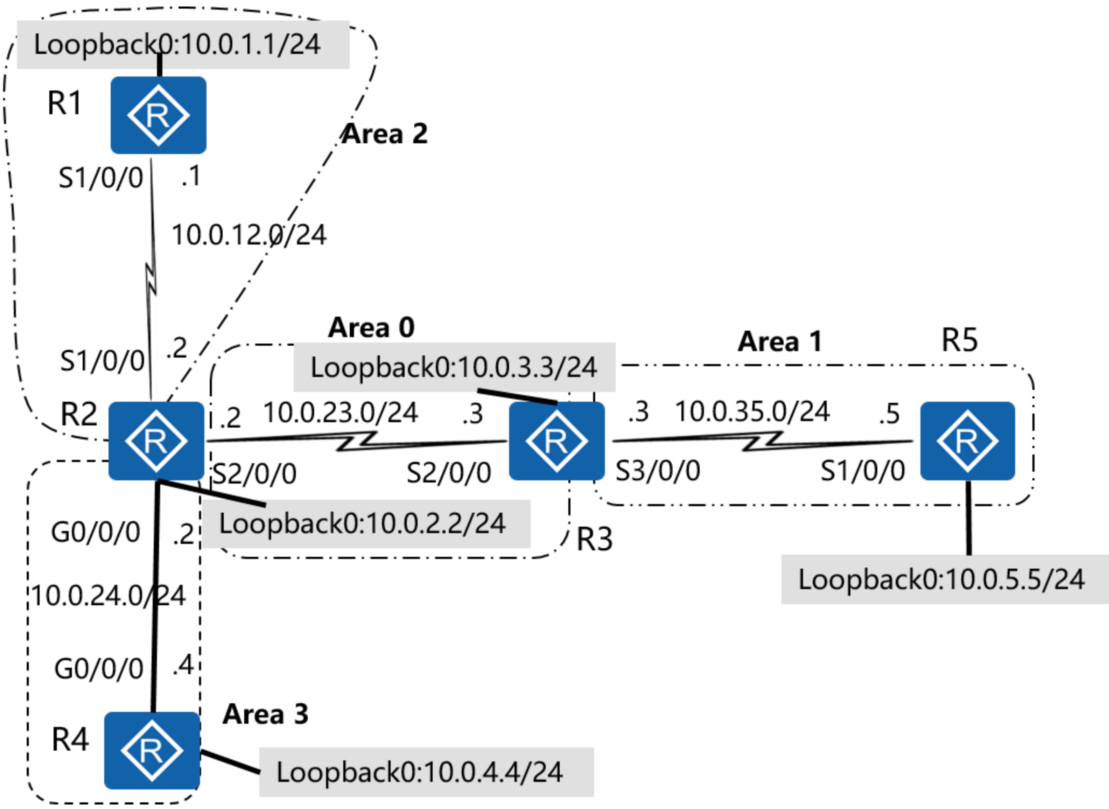

图1-4 OSPF Stub区域与NSSA区域

场景

你是公司的网络管理员。现在公司的网络中有五台AR G3路由器，其中R2、R3和R4在公司总部。R5在公司分部。R5通过专线与公司总部的R3相连。R1在公司的另外一个分部，通过专线与公司总部的R2相连。

网段10.0.23.0/24、10.0.2.0/24、10.0.3.0/24属于区域0。

网段10.0.35.0/24属于区域1，区域1为NSSA区域。R5的Loopback0接口不属于OSPF区域。

网段10.0.24.0/24属于区域3。R4的Loopback0接口连接到Internet，需要配置一条缺省路由。

网段10.0.12.0/24、10.0.1.0/24属于区域2，区域2为Stub区域。

同时为了明确设备的Router-ID，你配置设备使用固定的地址作为Router-ID。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置IP地址和掩码。配置时注意所有的Loopback接口配置掩码均为24位，模拟成一个单独的网段。

```
<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
```

测试直连链路的连通性。

```
[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
```

步骤二. 配置多区域 OSPF

在R1上配置Serial 1/0/0及Loopback 0属于区域2，并对所有OSPF区域的Loopback接口，修改其OSPF网络类型为Broadcast类型，以便于OSPF发布Loopback口的真实掩码信息，所有的路由器使用Loopback 0的IP地址作为Router ID。

```
[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
```

在R2上配置接口Serial 2/0/0及Loopback 0属于区域0，接口Serial 1/0/0属于区域2，接口GigabitEthernet 0/0/0属于区域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
```

在R3上配置接口Serial 2/0/0及Loopback 0属于区域0，接口Serial 3/0/0属于区域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
```

在R4上配置接口GigabitEthernet 0/0/0属于区域3，接口Loopback 0不属于任何区域。

```
[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
```

在R5上配置接口Serial 1/0/0属于区域1，接口Loopback 0不属于OSPF区域。

```
[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
```

配置完成后，在R1上查看路由器的路由表。

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

测试全网的连通性。

```
[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
```

步骤三. 配置将外部路由引入到 OSPF

将R5的Loopback0接口网段10.0.5.0/24引入到OSPF区域。使用默认配置进行路由引入。

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

配置完成后，在R1上查看该路由，并测试网络连通性。

```
[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
```

在R4上配置缺省路由，下一跳为Loopback0接口。并将该缺省路由引入到OSPF区域，定义为类型1，代价值为20，不使用永久发布。

```
[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
```

配置完成后，在R1上查看该缺省路由学习的情况。并测试网络的连通性。

```
[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
```

步骤四. 配置区域 2 为 Stub 区域

在R1上查看路由信息。注意刚才看到的默认路由是外部路由（O_ASE），是通过R4发布的第五类LSA学习到的。

```
[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
```

在R1和R2上配置区域2为Stub区域。

```
[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
```

配置完成后，在R1上对比之前的路由表，查看路由信息学习情况。这时可以看到，刚才的外部路由消失了，默认路由也变成了内部路由。

```
[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
```

查看R1的LSDB，可以看到描述外部路由的LSA也消失了，默认路由是由一条第三类LSA学习到的。

```
[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
```

查看这条LSA的详细信息，可以发现这条默认路由是由R2发布的，这就验证了将一个区域配置为Stub区域以后，ABR会阻断第四、五类LSA向该区域发送，并通过三类LSA向该区域内泛洪一条默认路由指向ABR自己。

```
[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
```

在R2上将区域2配置为**no-summary**的完全Stub区域。

```
[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
```

查看R1的路由表，这时发现通过OSPF学习到的路由条目只剩一条默认路由了。

```
[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
```

查看R1的LSDB信息，发现除了R1和R2产生的第一类LSA以外，只剩一条由R2发布的三类LSA。

验证了在完全Stub区域中ABR会阻断了第三、四、五类LSA，并生成一条三类LSA，通告一条指向自己的默认路由。

```
[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
```

步骤五. 配置区域 1 为 NSSA 区域

查看R3的路由表，R5发布的10.0.5.0/24是以外部路由的形式出现的。

```
[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
```

在查看R5的路由表和LSDB信息。R5从R4学习到一条外部路由，其余的路由均是内部路由。R5通过第五类LSA向外发布了网络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
```

配置R3和R5的区域1为NSSA区域。

```
[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
```

待邻居关系重新建立后，在R3上查看路由表。

```
[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
```

此时，R5通告的外部路由是以O_NSSA的形式出现在路由表里的。

再查看R5的路由表。

```
[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
```

刚才默认路由是以外部路由（O_ASE）的形式出现的，现在该默认路由变成了NSSA区域的外部路由（O_NSSA）。

查看R5的LSDB。

```
[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
```

发现，刚才的第五类LSA都消失了，外部路由以第七类LSA的形式向外通告。

查看默认路由的明细信息。

```
[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
```

刚才R5上的默认路由是R4通告给它的 ,而现在这条默认路由的通告者是R3。

从上面的结果我们可知，NSSA区域阻断了外部的第四、五类LSA进入，并且ABR会以第七类LSA的形式，向区域内通告一条默认路由。本区域的外部路由会以第七类LSA的形式，由ASBR向NSSA区域内通告。

NSSA和Stub区域的根本区别是，NSSA区域允许引入外部路由，而Stub区域不可以。

步骤六.　观察 NSSA 给 OSPF 带来的变化

使用**display ospf brief**命令查看R3目前所处的OSPF角色，可以看到，在Border Router这个字段有三个值：AREA　AS　NSSA。AREA表示该路由器是一台ABR，AS表示该路由器是一台ASBR，NSSA表示该路由器至少有一个接口位于NSSA区域。

```
[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
```

在NSSA区域中，由于不允许第五类LSA存在，所以ASBR是以第七类LSA的形式，向区域内通告外部路由的。但第七类LSA仅允许在NSSA区域内泛洪，NSSA区域的ABR收到这个第七类的LSA后，会将该第七类LSA转换成第五类LSA，然后向其他普通区域发布。

接下来我们在R3上观察7类LSA与5类LSA的转换过程。以10.0.5.0/24为例观察路由信息的传递。对于第七类LSA，Ls id描述了目的网段，Net mask描述了目的网段对应的掩码。Options字段为NP表示该LSA可以被ABR转化成一条第五类LSA，如果Options字段显示此LSA不可以被转换成第五类LSA，则Forwarding Address可以被设置成0.0.0.0；如果Options字段显示此LSA可以被转换成第五类LSA，则Forwarding Address不能被设置0.0.0.0。

在这里，所引入外部路由的下一跳不在OSPF路由域内，Forwarding Address设置为该ASBR上某个OSPF路由域内的Stub网段的接口IP地址。这里使用的地址为R5的Serial 1/0/0的接口地址。

```
[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
```

查看R3生成的用于描述10.0.5.0/24的第五类LSA。

```
[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
```

Ls id、Network Mask和Forwarding Address这几个字段的值直接从原来第七类LSA中拷贝。这样，10.0.5.0/24这个网段就被通告到其他区域了。

附加实验：思考并验证

NSSA区域类型适合用在哪些场景？

分析为什么R3路由器被定义为ASBR？

最终设备配置

```
<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
```

实验 1-5 OSPF 虚电路和区域路由过滤

学习目的

- 掌握使用OSPF虚电路来连接不连续的区域0的配置方法
- 掌握使用OSPF虚电路将非骨干区域连接到区域0的配置方法
- 掌握区域之间进行路由过滤和路由控制的方法

拓扑图

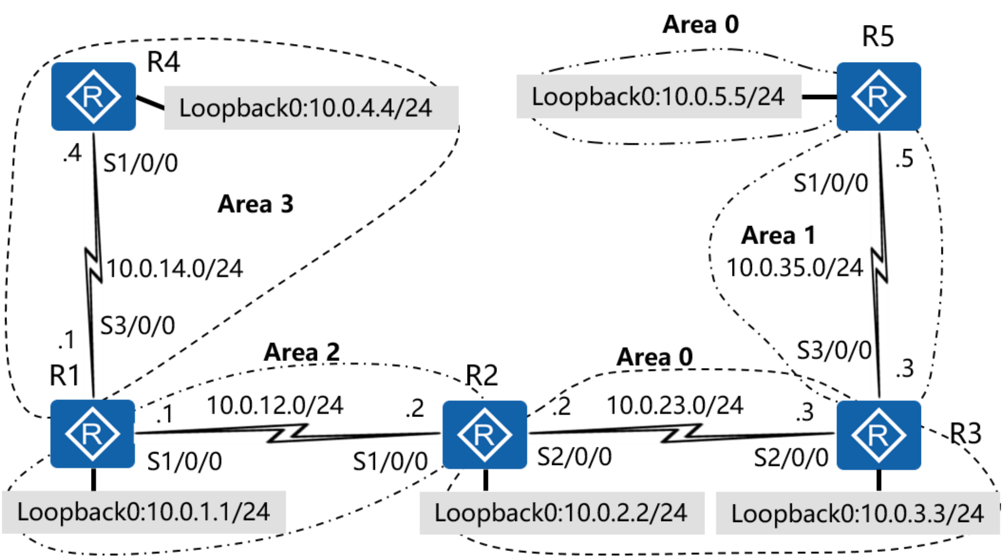

图1-5 OSPF虚电路和区域路由过滤

场景

你是公司的网络管理员。公司最近收购了两家小公司，他们的路由器是R4和R5。为了尽快合并网络，你决定先不去重新规划网络，而是使用虚电路实现网络互联。

网络直接相连后，你发现存在不连续的区域0，另外区域3与区域0没有直接连接。所以你决定在R1和R2之间建立虚电路，实现区域3与区域0的直接连接。另外在R3和R5之间建立虚电路，将不连续的区域0连接到一块。

同时为了明确设备的Router-ID，你配置设备使用固定的地址作为Router-ID。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置IP地址和掩码。配置时注意所有的Loopback接口配置掩码均为24位，模拟成一个单独的网段。

```
<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
```

配置完成后，测试直连链路的连通性。

```
[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
```

步骤二. 配置多区域 OSPF

在R1上配置Serial 1/0/0及Loopback 0属于区域2，Serial 3/0/0属于区域3。并对所有OSPF区域的Loopback接口，修改其OSPF网络类型为Broadcast类型，以便于OSPF发布Loopback口的真实掩码信息。所有的路由器使用Loopback 0的IP地址作为Router ID。

```
[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
```

在R2上配置Serial 2/0/0及Loopback 0属于区域0，Serial 1/0/0属于区域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
```

在R3上配置Serial 2/0/0及Loopback 0属于区域0，Serial 3/0/0属于区域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
```

在R4上配置Serial 1/0/0及Loopback 0属于区域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
```

在R5上配置Serial 1/0/0属于区域1，Loopback 0属于区域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
```

步骤三.　查看每个路由器的路由表

查看R4的路由表，R4虽然与R1建立了邻接关系，没有学习到任何OSPF路由。

```
[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 ]
```

再查看R4的LSDB发现仅存在第一类LSA，也就是说R1没有将其他区域的路由通告进区域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
```

查看R1的路由表，缺失了10.0.5.0/24。至于缺少这条路由的原因，我们分析完R3的LSDB就明白了。

```
[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
```

下面我们先来看一下R1的LSDB。为了避免区域间的环路，OSPF规定不允许直接在两个非骨干区域之间发布路由信息。从LSDB的角度来看，可以发现，ABR不会转发从非骨干区域收到的第三类LSA。

在R1上我们可以看到，在区域2的LSDB中有4条区域间路由，该路由是从R2（10.0.2.2）上学习到的，R1并没有将这些LSA转发到区域3里，所以R4学习不到非本区域的路由。

ABR也不会将从非骨干区域中学习到的路由转发给另一个非骨干区域，这里R1从R4这里学习到的路由不会以第三类LSA的形式通告进区域2，所以R2、R3、R5均学习不到区域3内的路由。

```
[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
```

查看R2的路由表，在R2的路由表中缺失了到达网络10.0.4.0/24、10.0.5.0/24、10.0.14.0/24的三条路由。

```
[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
```

查看R2的LSDB，可以发现R1没有将区域3的路由通告给R2。

R2上会缺失到达网络10.0.4.0/24、10.0.14.0/24的路由。

在区域0中，R3也没有将10.0.5.0网络的路由通告给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
```

查看R3的路由表，缺失了到达网络10.0.4.0/24、10.0.5.0/24、10.0.14.0/24的路由。

```
[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
```

查看R3的LSDB，我们可以发现，R3从区域1中收到了R5发来的第三类

10.0.5.0。根据前面的原则，从非骨干区域收到的第三类LSA不会被转发。

R3没有将这条LSA再次发送到区域0中，这也正是R1和R2中没有10.0.5.0/24这条路由的原因。

```
[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
```

在这里我们注意到，R3的LSDB中已经有了R5发来的用于描述10.0.5.0/24的第三类LSA，但在R3的路由表上并没有出现这条路由。

查看R5的路由表。

```
[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缺失了到达网络10.0.4.0/24、10.0.14.0/24的路由。

同时可以看到，R5上存在到达R3 Loopback 0的路由。

分析原因可知，R3存在连接到区域0的物理接口，即可以与区域0中其他的路由器交互路由信息，这时在R3上不会将非骨干区域发来的第三类LSA学习到的路由加到路由表里。R5虽然有接口在区域0中，但该接口为Loopback接口，Loopback接口在OSPF进行路由计算时为StubNet的链路类型。

查看R3的产生的类型一的LSA（此处已略去其他输出信息）：

```
[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
```

从上面的输出中可以看到，R3与R2相连的链路的类型为P-2-P。类型为P-2-P、TransNet和Virtual类型的链路，路由器均认为该接口与其他路由器会交互路由信息。对于存在这三类链路连接到骨干区域的路由器不会将非骨干区域发来的第三类LSA加到路由表中。

```
[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
```

在R5上，骨干区域中仅有一个Loopback 0，在描述这个接口的LSA中，链路的类型是StubNet，即末节网络，表示该接口不会再连接其他路由器，这时，路由器就会采用从非骨干区域发来的第三类LSA。

步骤四. 将不连续的区域 0 连在一块

在R3和R5上配置虚电路，注意在配置虚电路的时候在**vlink-peer**中配置的是对端ABR的Router ID。

```
[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
```

然后检查虚电路邻居的状态是否为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
```

观察路由信息发生的变化。

```
[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
```

在R3上查看路由表发现已经学习到了10.0.5.0/24这条路由。

测试网络的连通性，R3可以与R5的Loopback 0连接的网段通讯。

```
[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
```

查看R3的LSDB。

```
<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

看到在R3上共收到了2条来自R5的第一类LSA。第一条是在区域0中收到的，虚电路属于区域0，所以该LSA实际上是通过虚电路学习到的。另一条第一类LSA是在区域1中学习到的，这条LSA刚才没建虚电路的时候就有。路由表中的10.0.5.0/24路由是通过区域0学习到的LSA计算出的。

查看R3的LSDB中关于10.0.5.5这条第一类LSA的详细信息。

```
[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
```

可以看到这条LSA中描述了网络10.0.5.0/24，所以在R3上就有了这条路由。而从区域1中学习到的这条第一类LSA仅描述了R3和R5的互联网段。

查看R5的LSDB。

```
[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

可以发现和R3的LSDB是一样的。建立的虚电路以后，R3和R5均有接口属于区域0了，所以LSDB是同步的。

步骤五. 区域 3 通过虚电路连接到区域 0

在R1和R2上配置虚电路。

```
[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
```

查看R4的OSPF路由表。

```
[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
```

发现该路由器已拥有了全网路由。

测试网络的连通性。

```
[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
```

查看R1的LSDB。

```
[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

由于创建了虚电路，R1中有了区域0的LSA，这样区域0和区域3就可以直接交互路由信息了。R1把区域0中的路由信息以第三类LSA的形式通告进了区域3。

查看R4的LSDB。

```
[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
```

可以看到该路由器学习到了R1发布的第三类LSA。

R4有其他区域的路由。

步骤六. 配置区域之间的路由过滤

控制10.0.4.0/24网段的路由信息的发布。使R1可以学到该路由，但R2、R3、R5学不到这条路由。

设置一个访问控制列表。

```
[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
```

R1配置针对类型3的LSA的过滤，配置在区域3向其他区域发送更新时进行过滤。

```
[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
```

在R2上查看路由信息过滤的情况。

```
[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已经学习不到路由10.0.4.0/24了。

此时R1的路由表中仍然有该条目。因为R1和R4是同一个区域的，R4通过第一类LSA把该路由通告给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
```

附加实验：思考并验证

OSPF协议中为什么区域0必须连续？从当前的OSPF设计的角度来看，是否

可以对类型1和类型2的LSA进行过滤？

最终设备配置

```
<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
```

实验 1-6 OSPF 故障排除

学习目的

- 掌握对单区域OSPF中区域号码不匹配进行故障排除的方法
- 掌握对单区域OSPF中掩码不匹配进行故障排除的方法
- 掌握对单区域OSPF中Hello时间不匹配进行故障排除的方法
- 掌握对单区域OSPF中Router-id冲突进行故障排除的方法
- 掌握OSPF认证相关的故障排除方法
- 掌握OSPF汇总相关的故障排除方法
- 掌握虚电路相关的故障排除方法

拓扑图

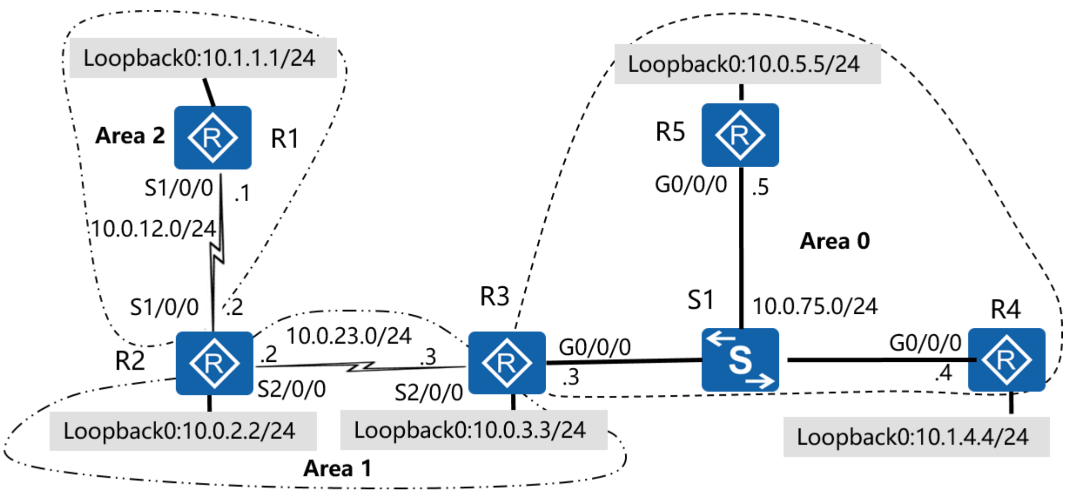

图1-6 OSPF 故障排除

场景

你是公司的网络管理员。公司的网络采用了OSPF协议作为路由协议。OSPF协议功能强大，但是相应的配置也较为复杂。并且在网络规划中，你使用了OSPF的各种特性，同时也使用了虚链路。在实施过程中，你碰到很多的网络通讯问题。不过庆幸的是，通过使用故障排除的思想和方法，你成功的找到了各种错误，并实现了网络的恢复。

学习任务

步骤一.　基础配置与 IP 编址

给所有路由器配置IP地址和掩码。配置时注意所有的Loopback接口配置掩码均为24位，模拟成一个单独的网段。

```
<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
```

为模拟相应的错误，R3的G0/0/0接口配置IP地址为10.0.75.3/25，其余接口地址按照拓扑图中的标识进行配置。

```
<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
```

配置完成后，测试直连链路的连通性。

```
[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
```

步骤二. 配置多区域 OSPF

配置R1的接口Serial 1/0/0和Loopback 0属于区域2 ,使用接口Loopback 0的地址作为Router ID。注意对所有OSPF区域的Loopback接口，修改其OSPF网络类型为Broadcast类型，以便于OSPF发布Loopback口的真实掩码信息。

```
[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
```

在R2上配置接口Serial 2/0/0及Loopback 0属于区域1，接口Serial 1/0/0

属于区域2，在启用OSPF时没有静态指定Router ID。

```
[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
```

在R3上配置接口Serial 2/0/0及Loopback 0属于区域1，接口GigabitEthernet 0/0/0属于区域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
```

在R4上配置接口GigabitEthernet 0/0/0属于区域1，接口Loopback 0不属于任何区域。在配置OSPF进程时使用**ospf 1 router-id**指定R4的Router ID为10.0.5.5。

```
[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
```

在R5上配置接口GigabitEthernet 0/0/0和Loopback 0属于区域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
```

步骤三. 在区域内排除 OSPF 故障

查看R4邻居列表，发现R4没有与其他路由器建立邻居关系。

```
[R4]display ospf peer

          OSPF Process 1 with Router ID 10.0.5.5
```

在R3、R4、R5上分别运行**display ospf error**查看OSPF发生的错误。

```
[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
```

从上面的输出中，我们可以看到，在R3、R4、R5这三台路由器之间一共发生了五种错误：Router ID冲突（Router id confusion）、子网掩码不匹配（Netmask mismatch）、错误的区域号（Bad area id）、错误的数据包（Bad packet）、错误的虚电路（Bad virtual link）。

在这里我们并没有配置虚电路，在这种情况下错误的虚电路其实就是区域号错误。对于R4来说，它在一个区域号为1的接口上收到了一个区域号为0的OSPF数据包，它会认为这是一个通过虚电路发送过来的数据包。路由器本身没有配置虚电路，就发生了这种错误。

子网掩码错误也是错误的数据包的一部分，我们可以先修正子网掩码的问题再观察是否还有错误的数据包。

首先排除Router ID冲突的问题。我们可以依次查看每台路由器的Router ID来手工找出发生冲突的路由器，也可以通过系统日志来查找。通过**display logbuffer**翻阅路由器当前的系统日志。

```
[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)
```

从上面可以看到，与R5发生冲突的路由器接口的IP地址是10.0.75.4。查看拓扑，发现10.0.75.4是R4的接口地址。然后查看R4的Router ID，可以看到该路由器的Router ID和R5是一样的。同时还发现R4的区域号配置也有误。

```
[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
```

修改R4的Router ID和区域号。

```
[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
```

修改完成以后通过命令**reset ospf counter**清空OSPF计数器。

注意reset命令需在用户视图下运行。

```
<R4>reset ospf counters
```

重置后，稍等片刻，再运行display ospf error检查该问题是否消失。

```
<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
```

可以看到在修改完配置以后，Route ID冲突和区域号错误的问题消失了，还剩下子网掩码不匹配的问题。为了找出是哪台路由器配置了错误的子网掩码，我们在R4上查看Debug信息。

```
<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
```

从上面的信息我们可以看出，从10.0.75.3发来的Hello包中子网掩码是255.255.255.128。查看拓扑，发现R3的对应接口配置错误。

```
[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
```

再次清空OSPF计数器，查看是否还存在错误。

```
<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
```

在R3上检查邻居列表，发现各邻居的状态已正常。

```
[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
 ----------------------------------------------------------------------------
```

接下来我们修改R4的GigabitEthernet 0/0/0接口的Hello间隔为5秒，观察邻居关系是否可以形成。

```
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ospf timer hello 5
[R4-GigabitEthernet0/0/0]quit
```

经过约半分钟以后，可以观察到R4的邻居都消失了。

```
[R4]display ospf peer brief

          OSPF Process 1 with Router ID 10.1.4.4
                   Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                     Neighbor id      State
 ----------------------------------------------------------------------------
```

清空R4 OSPF计数器，查看OSPF的错误。

```
<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
```

可以看到有Hello时间不匹配的错误出现，说明OSPF要求邻居间Hello间隔一样。

取消Hello间隔的修改。再次检查邻居列表。

```
[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
 ----------------------------------------------------------------------------
```

发现邻居关系已恢复正常。

步骤四. OSPF 认证故障排除

在R1和R2上配置基于接口的认证。

其中R1采用simple方式，密钥为123。

R2采用MD5方式，密钥为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
```

配置完成以后在R1上清空OSPF计数器，可以查看到OSPF的错误。

```
<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
```

将R1的认证方式配置为MD5后，查看是否还存在错误。

```
[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
```

可以看到该问题还存在。

将R1的密钥也改成huawei，观察邻居关系。

```
[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
 ----------------------------------------------------------------------------
```

可见，R1与R2已建立邻接关系。

步骤五. 虚电路故障排除

为保证区域2与区域0之间的连通性，在R2和R3之间创建虚电路。

```
[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
```

观察虚电路建立是否正常，以及R1是否学习到了全网路由。

```
[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
```

在R1上测试连通性，证实可以到达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
```

由于测试的需要，删除R2的loopback0接口。

```
[R2]undo interface LoopBack 0
```

后来由于一次偶然的事故，路由器重启了。在这里我们通过重启OSPF进程的方法模拟路由器重启。

```
<R2>reset ospf process
Warning: The OSPF process will be reset. Continue? [Y/N]:y
```

这时连接到R1的用户发现自己无法访问区域外的地址。管理员登录到R1上

发现无法与R5的Loopback地址通讯。

```
[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
```

检查R2和R3之间的虚电路之后发现状态不正常，同时发现R2的Router ID发生了变化。

```
[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
```

由于虚电路的建立是基于对端设备的Router ID的。R2的Router ID发生了变化，所以虚电路发生了故障。

通常我们在启动OSPF进程时指定该进程的Router ID，就是为了防止路由器在运行过程中Router ID发生变化。

下面我们将R2的Router ID固定为10.0.2.2，并将Loopback地址添加回去，然后重启OSPF进程。

```
[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
```

再次查看虚电路状态。

```
[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
```

此时虚电路已恢复正常。

管理员出于安全的考虑，在区域0使用了基于区域的认证，启用了MD5对报文进行加密，密钥为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
```

这时，管理员再次发现区域2中的用户无法访问区域外的网络，检查虚电路后发现虚电路又出于故障的状态。

```
[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
```

检查OSPF的错误发现有认证错误发生。

```
<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
```

OSPF的虚电路属于区域0。区域0打开了基于区域的认证，虚电路上也需要打开认证。

```
[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
```

这时虚电路的状态恢复了正常，R1也能正常访问其他区域了。

```
[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
```

步骤六.　OSPF 路由汇总故障排除

首先在R4上以外部路由的形式引入Loopback 0接口地址，并进行地址汇总，汇总后的子网掩码为16位。

```
[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
```

一段时间之后，管理员在R2上配置了区域间汇总，将R1的Loopback 0接口连接的网段汇总成16位掩码的路由。

```
[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
```

这时，除了连接到R4的用户以外，全网所有用户均反馈不能访问R4的Loopback地址10.1.4.4。

检查与R4同一区域的路由器R5的路由表发现，若要到达10.1.4.4，匹配到路由条目10.1.0.0/16，而该路由的下一跳是10.0.75.3。

为何会产生这样一个错误的条目呢？

```
[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
```

我们先来观察一下R5的LSDB。

```
[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
```

在LSDB中我们看到有2条描述10.1.0.0的路由，接下来查看LSA的详细信息。下面这条第三类LSA是由R2始发的，而这条第五类LSA是由R5始发的。这两条LSA描述了一个完全相同的网段信息。

```
[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
```

在OSPF中，第三类LSA始终优于第五类LSA，所以在R5路由表里出现的10.1.0.0/16这条路由的下一跳会是R3。

为了避免这类问题的发生，我们在R4上取消原来对外部路由的汇总，这样这条路由就会再次在其他路由器的路由表中出现。

```
[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
```

这时我们可以看到，在R5上已学习到了一条关于10.1.4.4/24正确的路由。这时我们在R1上测试连通性。

```
[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
```

可见，网络恢复正常。

附加实验：思考并验证

可否在一个区域中即打开基于区域的认证，又打开基于接口的认证？

非骨干区域的区域号可否一样？

最终设备配置

```
<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
```

第二章 IS-IS协议特性与配置

实验 2-1 IS-IS 配置实验

实验目标

- 掌握IS-IS协议基本配置
- 掌握IS-IS协议DIS优先级修改方式
- 掌握IS-IS协议网络类型修改方式
- 掌握IS-IS协议外部路由引入
- 掌握IS-IS接口cost修改方式
- 掌握IS-IS路由渗透配置方式

拓扑图

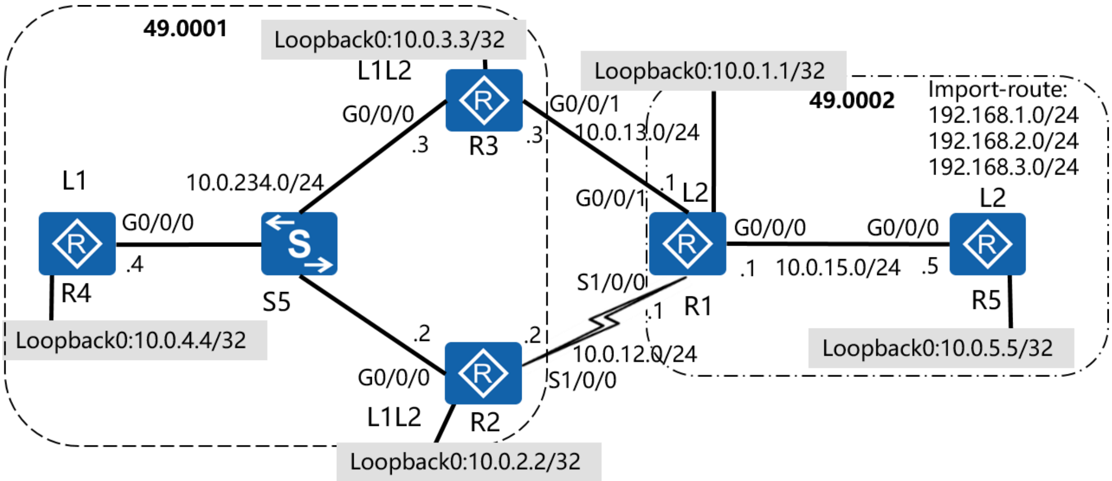

图2-1 IS-IS实验拓扑

场景

使用IS-IS协议作为某网络的IGP，R1和R5运行在49.0002区域Level-2，R2、 R3和R4运行在49.0001区域，R4是Level-1路由器，R2和R3作为Level-1-2路由器。要求：R4可以访问R5引入的外部路由；R4连接S5的接口作为DIS；R1和R5之间使用点到点链路；从R4访问R5的上下行流量应该从以太网接口进行转发，可以通过cost和路由渗透进行选路控制；交换机不做额外的配置，仅透明转发。

学习任务

步骤一. IP 编址与基本配置

给所有路由器配置IP地址信息。

```
[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
```

配置完成后，在R1上测试到R2、R3和R5直连接口的连通性。

```
[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
```

测试R4到R2和R3之间直连接口的连通性。

```
[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
```

步骤二. IS-IS 协议基本配置

按照拓扑设计逐台配置路由器的isis进程，进程号使用1，以R1为例：

由于R1在区域49.0002，使用network-entity 49.0002.0000.0000.0001。

```
[R1]isis
[R1-isis-1]network-entity 49.0002.0000.0000.0001.00
```

默认情况下启用IS-IS进程后，路由器工作在Level-1-2模式，按照规划R1应该是L2模式，因此进行修改：

```
[R1-isis-1]is-level level-2
[R1-isis-1]quit
```

将相应的接口启用IS-IS，包括LoopBack接口，不指定IS-IS进程号的情况下默认在进程1下启用：

```
[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
```

检查IS-IS在接口的启用状况：

```
[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 --
```

可以看到**ISIS(1)**共启用了4个接口，在IPV4地址下状态为Up。

用相同的方式配置其他路由器，R2和R3工作在Level-1-2模式，因此不需要修改is-level：

```
[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
```

在R2上检查IS-IS在接口的启用状况：

```
[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 --
```

在R3上配置IS-IS：

```
[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
```

在R3上检查IS-IS在接口的启用状况：

```
[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

在R4上配置IS-IS：

```
[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
```

在R4上检查IS-IS在接口的启用状况：

```
[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
```

在R5上配置IS-IS：

```
[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
```

在R5上检查IS-IS在接口的启用状况：

```
[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
```

配置完成后，观察各设备的邻居状态，以R1为例，应该有3个邻居，分别是R2、R3和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
```

这里的System Id，类似于其他协议的Router Id，可以看到R2、R3和R5的状态都是Up，邻居状态正常。

继续检查其他设备的邻居状态：

```
[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
```

步骤三. 修改 DIS 优先级

R2、R3和R4在一个广播网络下建立邻居，因此需要选举DIS，默认情况下，DIS优先级都为64，如果优先级相同MAC地址大的接口将成为DIS，我们希望选举更为清晰，因此通过修改R4的DIS优先级来保证其成为DIS。

```
[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 –
```

在R2和R3上查看邻居的DIS优先级：

```
[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
```

步骤四. 配置 IS-IS 网络类型

在广播网络中，IS-IS默认会将接口的circuit-type设置为广播模式，并参与DIS的选举。但在拓扑中R1和R5之间的以太网只有两台路由器，我们可以将这两台路由器之间的互联接口设置为点到点模式，进行优化。

```
[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

修改circuit-type过程中，邻居会重新建立，检查配置效果，Circuit Id会改变格式，接口的详细信息也会显示p2p，以R1为例：

```
[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
```

步骤五. 配置 IS-IS 外部路由引入

在引入路由之前，先检查当前所有路由的学习情况，R1去往R4的路由呈现负载均衡的状态，数据包将会在GE0/0/1和S1/0/0之间均匀分布：

```
[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

因为R2是LEVEL-1-2路由器，因此在两个level生成不同的路由，注意在level1会有默认路由的条目，指向空接口，R3和R2的情况相同：

[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
```

由于R4是L1路由器，只能和同区域的L1或者LEVEL-1-2路由器建立邻居，并且默认情况下L1路由器无法学到L2的路由信息，只能够通过默认路由访问外部，R4可以看到两条默认路由指向R2和R3，呈现负载均衡。

[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

R5在引入外部路由前，路由学习情况：

```
[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
```

在R5上创建新的LoopBack口，并引入到isis进程：

```
[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
```

在R5上直连路由引入level-2，查看效果：

```
[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
```

此时在R4上再次查看路由表，发现并没有变化，原因是在默认情况下，L2路由不会渗透进入L1路由器，但通过默认路由，R4可访问 192.168.1.0/24、192.168.2.0/24和192.168.3.0/24。

```
[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
```

步骤六.　修改 IS-IS 接口 Cost 值

默认情况下，IS-IS接口cost值为10，不会基于带宽自动计算，对于R1来说，去往R4的流量会在R2和R3上负载均衡，但由于R1和R2之间使用Series接口，带宽较小，容易出现瓶颈，因此可以通过修改相应的cost值来控制R1的选路。

增大出口的cost值：

```
[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
```

可以看到在R1上负载均衡已经消失，去往R4的流量将从以太网接口转发。

步骤七.　配置 IS-IS 路由渗透

观察R4的数据转发，由于R4对Level-2网络状况一无所知，因此使用负载均衡转发数据到R2和R3，如果我们希望R4不使用R2到R1之间的链路，可以通过路由渗透将Level-2的路由引入到Level-1中，通过路由最长匹配选路，从R3转发去往R5的数据包。

测试tracert之前，需要在R5上开启端口不可达的响应，默认是关闭的，如果不开启，则最后一跳会超时。

在开启前：

```
[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   *   *   *
```

在R5上开启后：

```
[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
```

可以看到每次发出的tracert包会在两个下一跳地址负载均衡到达R5，此时我们在R3开启路由渗透，使得R3为更优的下一跳：

```
[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        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
```

最终，通过tracert测试验证了路由渗透对IS-IS选路效果的影响。

配置文件参考

```
<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
```

第三章 BGP协议特性与配置

实验 3-1 IBGP 与 EBGP

学习目的

- 掌握区域内部BGP的配置方法
- 掌握多区域BGP的配置方法
- 观察BGP的邻居表和数据库
- 掌握BGP更新源的配置方法
- 掌握EBGP多跳的配置方法
- 观察IBGP和EBGP中路由的下一跳的变化
- 掌握IBGP中下一跳的配置
- 掌握BGP的Network命令的配置方法

拓扑图

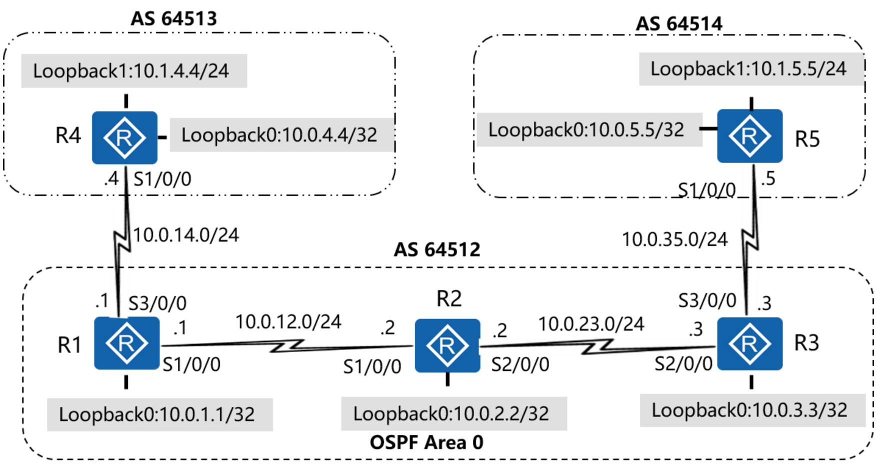

图3-1 IBGP与EBGP

场景

你是公司的网络管理员。公司的网络采用了BGP协议作为路由协议。公司的网络由多个自治系统组成，不同的分支机构使用了不同的AS号，现在你需要完成公司网络的搭建工作。在公司总部使用了OSPF作为IGP，公司内部不同分支机构使用的是私有的BGP AS号。在完成网络搭建以后，你还需要观察BGP路由信息的传递。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置IP地址和掩码，其中R4和R5的loopback 1接口掩码为24位，模拟用户网络。

```
<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
```

测试各直连链路的连通性。

```
<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
```

显示直连联络连通性正常。

步骤二. 配置区域内 IGP

在AS 64512中使用OSPF作为IGP ,将Loopback 0连接的网段发布进OSPF。R1的S1/0/0连接的网段运行OSPF。

```
[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
```

R2的S1/0/0和S2/0/0连接的网段运行OSPF。

```
[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
```

R3的S2/0/0连接的网段运行OSPF。

```
[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
```

注意在使用network命令时，通配符掩码使用0.0.0.0。

配置完成后察看OSPF的邻居关系是否建立。

```
[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 ]
```

查看所有路由器的路由表。检查是否学习到对端设备Loopback接口网段的

路由。

```
[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
255.255.255.255/3  Direct  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
10.0.12.255/32          Direct  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
10.0.23.255/32      Direct  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
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

从R1，R2，R3的路由表中，可以看到学习每一台路由器都能够学到其他2台路由器的loopback 0接口连接的网段的路由。

步骤三. 建立 IBGP 对等体

在R1、R2、R3上配置IBGP全互联。使用Loopback0地址作为更新源。

```
[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
```

使用**display tcp status**查看TCP端口连接状态。

```
[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
```

从表项中我们可以观察到Local Add为10.0.2.2（即为R2的Loopback0接口地址），端口号为179（BGP协议的TCP端口号）。与10.0.3.3和10.0.1.1的状态已经为Established，说明R2和R1，R3的TCP连接已建立。

使用**display bgp peer**察看各路由器BGP邻居关系状态。

```
[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
```

可以看到3台路由器之间的BGP邻居关系处于Established状态，说明邻居关系已建立。

在R1上在BGP进程下使用**timer**修改BGP的keep alive时间为30秒，hold时间为90秒。观察R1与R2的对等体关系建立是否有问题，使用**display bgp peer verbose**命令观察建立以后协商的间隔时间是多少。

```
[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
```

注意：修改此参数将引起bgp邻居重启。

[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
```

可以看到在R2上默认的配置参数Active Hold Time为180s，Keepalive Time为60s。

当R1的参数修改之后，R2收到数据包的Active Hold Time为90s。协商的参数取值数值小的参数，所以R2与R1的协商的结果Active Hold Time为90s，keepalive Time为30s，而R3的参数仍然为默认参数。

R2和R3一样，所以协商的结果配置参数和协商参数一致，Active Hold Time为180s，Keepalive Time为60s。

步骤四. 配置 EBGP 对等体

在R4上配置BGP，本地AS号为64513，与R1建立EBGP对等体关系。在建立对等体关系时，指定更新源为Loopback 0接口的地址，并指定**ebgp-max-hop**为2。添加到对端Loopback 0接口地址的32位的静态路由，使之能正常建立对等体关系。

```
[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
```

对等体关系建立完成后，使用**display bgp peer**检查对等体关系状态。

```
[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
```

在R4上使用**debugging ip packet verbose**查看keepalive报文的TTL值。

```
<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.
```

可以看到收到的报文中TTL都为2。

在R3和R5之间也建立EBGP对等体关系。直接使用物理接口地址建立连接。

```
[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
```

步骤五. 使用 Network 命令发布路由信息

在R4上配置Loopback1，地址为10.1.4.4/24。使用**network**命令将该网段发布进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
```

在R1和R3上全局路由表分别查看该路由是否存在。

查看R3上BGP路由表，分析该路由的下一跳信息。

```
[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
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

```
255.255.255.255/3        Direct   0     0             D    127.0.0.1        InLoopBack0
```

可以看到在R1上已经学到10.1.4.0/24的EBGP路由。

在R3上查看是否有到达网络10.1.4.0/24的路由。

```
[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    0    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
255.255.255.255/3  Direct  0     0            D    127.0.0.1       InLoopBack0
```

在R3上并没有10.1.4.4的bgp路由。

查看R3的BGP表。

```
[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
```

可以在R3的BGP路由表中看到，但是这条BGP路由没有*号标识，说明这条路由并没有被优选。因为这条路由的NextHop为10.0.4.4，而R3上并没有到达地址10.0.4.4的路由。根据BGP选路原则，当BGP路由的下一跳不可达时，忽略此路由。

在R1上配置**next-hop-local**，再次在R3上查看该路由表。

```
[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
```

可以看到BGP路由10.1.4.0/24的下一跳为10.0.1.1，同时具有*号和>号，说明这条路由是正确并且最优的。

查看R3的路由表。

```
[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
127.255.255.255/32Direct  0    0        D   127.0.0.1       InLoopBack0
255.255.255.255/3  Direct  0    0        D   127.0.0.1       InLoopBack0
```

路由表出现路由10.1.4.0/24。

在R5上创建Loopback1，地址为10.1.5.5/24，发布进BGP，在R3上配置**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
```

在R4上查看是否学习到R5的Loopback 1连接网络的路由。分析**display bgp routing-table**的输出。

```
[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
```

在R5上使用带源地址ping测试到R4的Loopback1地址的连通性。

```
[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
```

附加实验：思考并验证

什么情况下用物理地址直接建立EBGP邻居关系比较恰当？

为何默认发给EBGP邻居报文的TTL值为1？ 运行**peer *group_name* ebgp-max-hop *[hop-count]***的默认值是多少？

最终设备配置

[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
```

实验 3-2 BGP 路由汇总

学习目的

- 掌握使用**network**命令实现的BGP路由汇总的配置方法
- 掌握实现抑制具体路由的BGP路由汇总的配置方法
- 掌握改变汇总路由属性的配置方法
- 掌握在路由汇总时使用AS-SET的配置方法

拓扑图

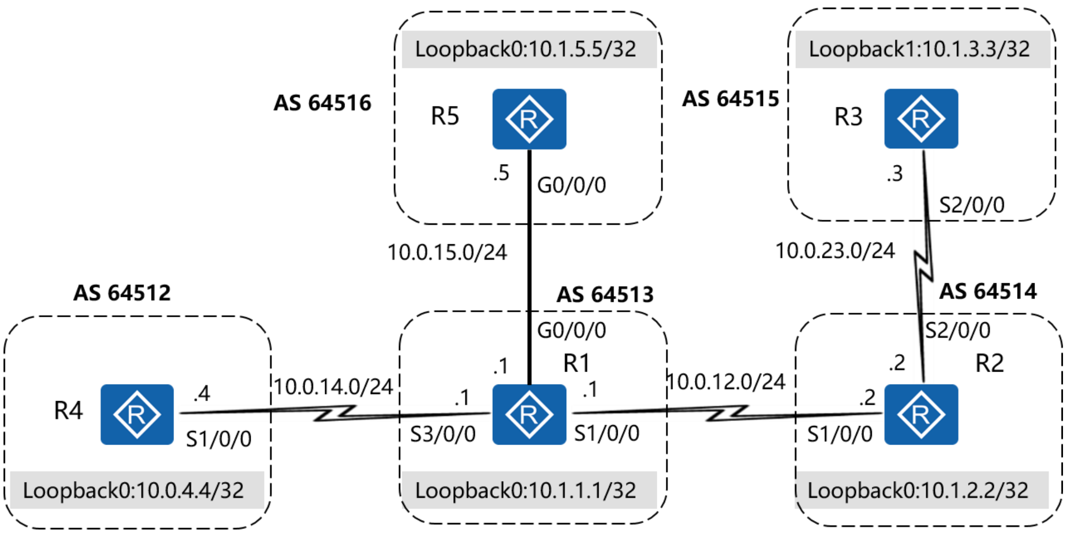

图3-2 BGP路由汇总

场景

你是公司的网络管理员。公司的网络采用了BGP协议作为路由协议。公司的网络由多个自制系统组成，不同的分支机构使用了不同的AS号。随着公司规模的扩大，路由器中已经有越来越多的路由表，进行BGP的路由汇总迫在眉睫。你测试了几种进行路由汇总的方法，最终选择了合适的方式实现了路由汇总。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置物理接口及Loopback接口的IP地址和掩码。注意各Loopback接口地址均为32位掩码。

```
[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
```

配置完成后，测试直连链路的连通性。

```
[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
```

步骤二. 配置 EBGP 及发布路由

各直连路由器之间直接使用物理接口地址建立BGP对等体关系。

```
[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
```

配置完成以后检查对等体关系，

```
[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
```

BGP邻居状态此时全部都是Established状态。

使用**network**命令将各个路由器的Loopback接口网段发布进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

在R4上查看BGP路由表，观察AS-PATH属性的信息。

[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
```

步骤三.　使用 Network 实现对 BGP 路由的汇总

现在需要在R1上进行路由汇总。

在R1上添加指向Null0接口的静态路由10.1.0.0/16 ,并使用**network**命令发布该路由。

```
[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
```

在R4上查看路由表，观察汇总路由是否存在。

```
<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
```

设置名为**pref_detail_control**的前缀列表，对向对等体R4发送的路由进行过滤，不允许汇总路由中包括的详细路由被发送过去。

```
[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
```

在R4上再次查看BGP路由表。注意观察汇总后路由的As-path属性。

```
<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
```

步骤四. 使用 Aggregate 实现对 BGP 路由的汇总

删去步骤三中使用的前缀列表及**network**命令发布的汇总路由。

使用**Aggregate**命令对路由10.1.0.0/16进行汇总，使用默认方式进行。

```
[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
```

在R1、R4查看BGP路由表。观察汇总路由的Origin属性。

```
[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
```

汇总路由的**origin**属性没有改变，仍旧是IGP。

在R1配置路由汇总时，抑制明细路由，仅通告聚合路由。

```
[R1]bgp 64513
[R1-bgp]aggregate 10.1.0.0 255.255.0.0 detail-suppressed
[R1-bgp]quit
```

在R4查看BGP路由表。

```
[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
```

这时R4上已经看不到明细路由。

观察R1的全局路由表，查看路由10.1.0.0/16的下一跳。

```
[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
```

因为这条聚合路由是R1上配置的，所以出接口为Null0，这样定义的汇总路由有利于避免路由环路的产生。

观察R1的BGP路由表，观察明细路由。

```
[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
```

使用了**detail-suppressed**参数，对外仅发送汇总路由，这时在明细路由条目前面多了个标识s，代表在路由汇总时，明细路由被抑制。

步骤五. 改变汇总路由的属性

缺省情况下，BGP不将团体属性发布给任何对等体。

配置R5向R1、R1向R4通告团体属性。

```
[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
```

验证进行路由汇总后团体属性会丢失。

在R5上对R5通告的10.1.5.5/32加上100的团体属性，并向R1通告。

```
[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
```

在R1上查看该路由是否携带该团体属性。

```
<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>
```

在R4上查看汇总后路由是否携带该属性。

```
<R4>display bgp routing-table community

 Total Number of Routes: 0
```

R4没有任何携带团体属性的路由。

在R1上配置路由策略**add_comm**，将100:2的团体属性加到汇总后的路由上。

```
[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
```

在R4上观察该汇总路由是否携带100:2团体属性。

```
<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>
```

R4学习到的汇总路由携带该属性。

步骤六. 使用 AS-SET 配置 AS-PATH 属性

路由在汇总后，默认会丢弃AS-PATH属性信息，AS-PATH的丢失可能会造成环路。为避免信息丢失带来的风险，汇总路由可以在汇总时添加AS-Set属性。

配置R1在执行路由汇总时添加AS-Set属性。

```
[R1]bgp 64513
[R1-bgp]aggregate 10.1.0.0 255.255.0.0 detail-suppressed as-set
[R1-bgp]quit
```

观察R1、R4中BGP路由表中汇总路由的AS-PATH属性信息。

```
[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
```

加了AS-Set属性之后的汇总路由的AS-PATH包含了具体路由的AS路径信息。

在R3上停止通告10.1.3.3/32，重置对等体关系。

```
[R3]bgp 64515
[R3-bgp]undo network 10.1.3.3 255.255.255.255
[R3-bgp]return
<R3>reset bgp all
```

待邻居关系重新建立后，观察R4上学习到的汇总路由的AS-PATH属性信息。

```
<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
```

这时发现AS-PATH属性中已经没有AS号码64515。

附加实验：思考并验证

这个例子中，完成步骤六以后，R5是否能访问到R3的Loopback地址？

Aggregate和**Summary automatic**有什么区别？

最终设备配置

```
<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
```

实验 3-3 BGP 属性与路径选择 1

学习目的

- 掌握通过配置AS-Path属性来影响路径选择的方法
- 掌握通过修改Origin属性来影响路径选择的配置方法
- 掌握通过修改Local-Pref属性来影响路径选择的配置方法
- 掌握通过修改MED属性来影响路径选择的配置方法

拓扑图

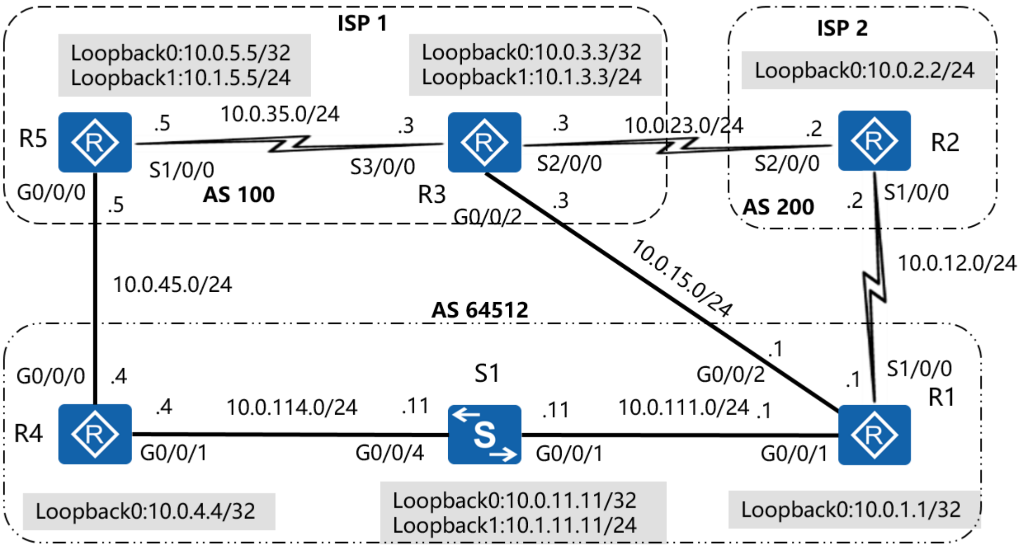

图3-3 BGP属性与路径选择

场景

你是公司的网络管理员。公司的网络采用了BGP协议接入了两个服务运营商。公司自己采用了私有的AS号64512，ISP1的AS号为100，公司共有2条链路接入ISP1。ISP2的AS号为200，公司租用了一条线路接入ISP2。现在Internet上的部分用户反应访问公司网站的速度较慢，你通过改变BGP的各种属性达到了调整路由走向的目的。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置物理接口及Loopback接口的IP地址和掩码。注意各Loopback 0接口均使用32位掩码。

```
[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
```

配置完成后，测试直连链路的连通性。

```
<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
```

步骤二. 配置 IGP 及 BGP

在AS 64512内部使用OSPF作为IGP，所有设备属于区域0。

R1的G0/0/1和Loopback 0连接的网段运行OSPF。

```
[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
```

在S1上创建Vlan 111，配置Vlanif地址与R1进行互联。

创建Vlan114，配置Vlanif地址与R4进行互联。

互联接口使用Access模式，S1的Vlanif 111、Vlanif 114 和Loopback 0连接的网段运行OSPF。

```
[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
```

R4的G0/0/1和Loopback 0连接的网段运行OSPF，

```
[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
```

检查是否学习到其他设备的Loopback 0地址所在的网段。

```
[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
127.255.255.255/32Direct  0      0              D     127.0.0.1         InLoopBack0
255.255.255.255/32Direct  0      0              D     127.0.0.1         InLoopBack0
```

在R1、R4、S1上配置BGP，要求使用各自的Loopback 0接口建立对等体连接，配置的时候使用对等体组，组名为**as64512。**

默认情况下，BGP的负载分担是关闭的。在所有的路由器上打开负载分担，设置最大同时使用4条等价路径。

```
[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
```

在R1、R2、R3、R4、R5上配置EBGP，AS的规划如图所示，EBGP全部使用物理接口地址建立对等体关系。

```
[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
```

步骤三. 配置 AS-Path 属性

在S1上创建Loopback 1，地址为10.1.11.11/24，使用**network**命令发布到BGP中。

```
[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
```

在R2上观察BGP路由表，可看出这时10.1.11.0/24路由是依据AS-Path属性来选择下一跳的。

```
[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
```

由于R1到R4之间的带宽有限，现在希望R2能经由AS100来访问10.1.11.0/24。

这里通过AS-Path来影响选路。

在R1上创建路由策略**as_path**，针对10.1.11.0/24这条路由增加2个重复的AS号。

```
[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
```

然后把该策略应用在R1上，让R2从R1上学习到的这条路由的AS-Path有3个值。

```
[R1]bgp 64512
[R1-bgp]peer 10.0.12.2 route-policy as_path  export
[R1-bgp]quit
```

在R2上观察BGP路由表。

```
<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
```

此时R2经由AS100访问10.1.11.0/24网段。

步骤四. 配置 Origin 属性

观察R3的路由表。

```
<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
```

到达10.1.11.0/24的下一跳是R1，分析原因。

我们希望R3通过R5访问AS 64512，查看10.1.11.0/24原来的Origin属性是IGP。

这时，我们将R1对R3通告的该路由修改为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
```

策略生效后观察R3的BGP路由表。

```
<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?
```

这时R3到达网络10.1.11.0/24的下一跳是R5。

步骤五. 配置 Local-Pref 属性

本地优先属性在选路中有很高的优先级。

通过改变本地优先属性可以影响选路。

在R3上创建Loopback 1，地址为10.1.3.3/24，发布进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
```

在R5上创建Loopback 1，地址为10.1.5.5/24，发布进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
```

在S1上观察路由表。

```
[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
```

现在希望到达网络10.1.5.0/24的流量从R4发送到目标，到达网络10.1.3.0/24的流量从R1发送到目标。

在R4上创建路由策略Pref4，匹配路由10.1.5.0/24，将其本地优先属性修改为110。

R1上创建路由策略Pref1，匹配路由10.1.3.0/24，将其本地优先属性修改为110，然后将策略应用到IBGP的对等体组上。

```
[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
```

在S1上查看BGP路由表。

```
[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
```

可以观察到，此时根据Local-Pref属性进行选路，越高越优先。

步骤六. 配置 MED 属性

删除步骤四中通过修改Origin来影响AS100中对10.1.11.0/24选路的路由策略，本实验中通过修改MED值来影响选路。

```
[R1]bgp 64512
[R1-bgp]undo peer 10.0.15.3 route-policy 22 export
[R1-bgp]quit
[R1]undo route-policy 22
```

在R1上创建路由策略**med**，针对10.1.11.0/24，将MED值修改为100，将该策略应用到对等体R3上。

```
[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
```

在R3上查看BGP路由表。

```
<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

MED值越小越优先。

最后观察现象，可以和步骤四中达到同样的选路效果。

附加实验：思考并验证

思考在完成了步骤六以后，关闭R1的S1/0/0接口，那么在R2上学习到的关于10.1.11.0/24的MED值是多少？

思考可否使用路由策略将AS-Path属性里的某个AS删除？

最终设备配置

<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
```

实验 3-4 BGP 属性与路径选择 2(选做)

学习目的

- 掌握通过修改团体属性来影响路径选择的配置方法
- 掌握使用Route-policy来过滤BGP路由信息的配置方法

拓扑图

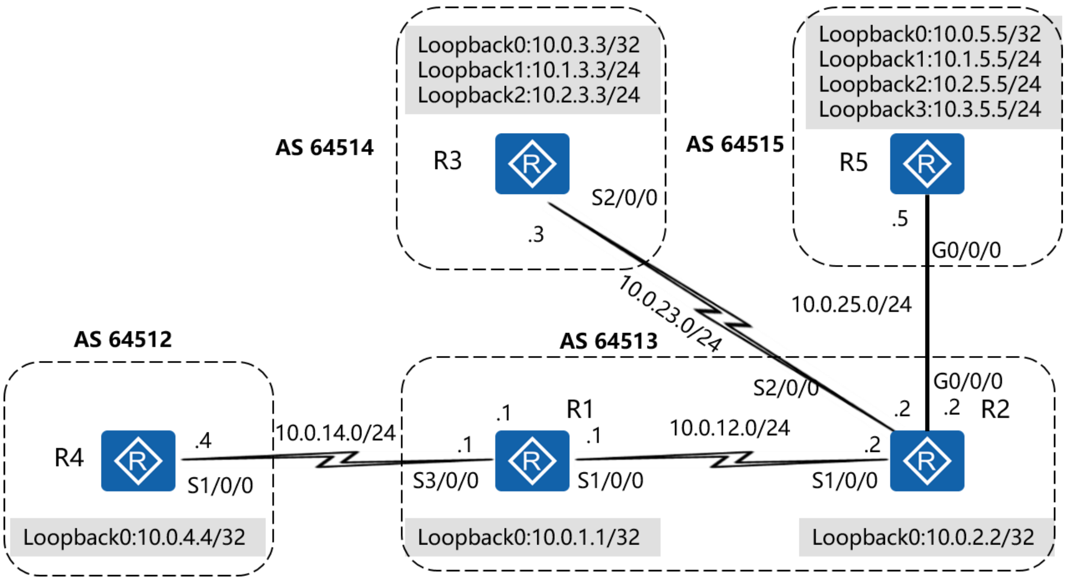

图3-4 BGP属性与路径选择2

场景

你是公司的网络管理员。公司的网络采用了BGP进行互联，BGP的AS号规划如拓扑图中所示。为了公司网络的安全，并非所有分支机构之间的所有部门都能互访。为了控制路由信息的传递，现在你需要使用BGP的团体属性对BGP的路由进行过滤。

学习任务

步骤一． 基础配置与 IP 编址

给所有路由器配置物理接口及Loopback 0接口的IP地址和掩码。注意各Loopback 0接口地址均使用32位掩码。

```
<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
```

配置完成后，测试直连链路的连通性。

```
<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
```

步骤二.　配置 BGP

R1与R2之间为IBGP邻居关系，其他路由器之间均为EBGP邻居关系，

```
[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
```

在配置完BGP之后检查路由器之间的邻居关系建立情况，

```
[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
```

此时BGP邻居状态全部都是Established状态。

步骤三. 配置普通的团体属性

在R5上创建Loopback1、Loopback2和Loopback3，地址分别为10.1.5.5/24、10.2.5.5/24、10.3.5.5/24，并通过**network**命令发布到BGP中。

```
[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
```

在R2和R4上检查该路由信息是否被正确传递。

```
[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

[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
```

在R5上创建路由策略**comm_r5**，对10.1.5.0/24这条路由添加团体属性为100。

```
[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
```

为了后面实验的需要，需要允许路由器之间通告团体属性，配置所有BGP邻居之间通告Community属性。

```
[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
```

在R2和R4上查看该属性是否被正常传递。

```
<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>
```

步骤四. 配置特殊的团体属性值

在R5上使用路由策略为路由10.2.5.0/24添加特殊的团体属性no-export ,使用路由策略为10.3.5.0/24添加特殊的团体属性no-advertise。

这时只需要在R5已创建的路由策略**comm_r5**的基础上添加入两个新节点语句即可。

```
[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
```

在R2上使用查看学习到的路由的团体属性信息。

```
<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
```

在R2、R1、R4上分别观察BGP路由表观察这几条路由的传递情况。

```
<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
```

对于使用了特殊团体属性no-export后的BGP路由10.2.5.0/24，R2不向其AS外发送，但会发布给其AS内的R1。而对于使用no-advertise属性的BGP路由10.3.5.0/24，R2则不向任何对等体发送。

步骤五.　配置团体属性在地址汇总中的应用

在R3上创建Loopback 1和Loopback2，地址分别配置为10.1.3.3/24和10.2.3.3/24，并通过**network**命令发布到BGP中。

```
[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
```

现在有一个需求，我们需要将R5发布的10.1.5.0/24和R3发布的10.2.3.0/24汇总成一个A类网段10.0.0.0/8。通告时抑制明细路由，并且该汇总路由最后通告给R4时携带的团体属性为200。对路由10.1.3.0/24保留明细通告给R4。

为了实现这个需求，我们在R3上创建一个名为**comm_r3**的路由策略，对R3发布的路由10.2.3.0/24加上100的团体属性。

```
[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

在R1上观察学习到的10.1.5.0/24和10.2.3.0/24是否都带上了100的团体属性。

```
<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
```

创建一个团体属性过滤列表，过滤出团体属性为100的路由。

```
[R1]ip community-filter 1 permit 100
```

创建一个名为**match_comm**的路由策略，匹配团体属性为100的路由。

```
[R1]route-policy match_comm permit node 10
[R1-route-policy]if-match community-filter 1
[R1-route-policy]quit
```

创建一个名为**add_comm**的路由策略，为汇总后路由添加团体属性为200:1。

```
[R1]route-policy add_comm permit node 10
[R1-route-policy]apply community 200:1 additive
[R1-route-policy]quit
```

在R1上进行地址汇总，定义对匹配策略**match_comm**的路由进行汇总，并使用策略**add_comm**添加团体属性。

```
[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
```

在R4上查看BGP路由表。

```
<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
```

在R4上观察汇总路由的团体属性。

```
<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>
```

附加实验：思考并验证

在步骤四中，如果将10.2.5.0/24这条路由的属性也改成no advertise，思考此时在R2、R1、R4上再分别查看BGP路由表，这几条路由的传递情况。

思考如何在R4上实现同时保留10.1.3.0/24和10.2.3.0/24这两条路由的明细，仅抑制路由10.1.5.0/24的明细。

最终设备配置

```
<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
```

实验 3-5 BGP 多宿主

学习目的

- 掌握BGP多宿主环境中仅使用缺省路由时的配置方法
- 掌握BGP多宿主环境中使用缺省路由过滤部分路由的配置方法
- 掌握BGP多宿主环境中，仅使用BGP路由时的配置方法

拓扑图

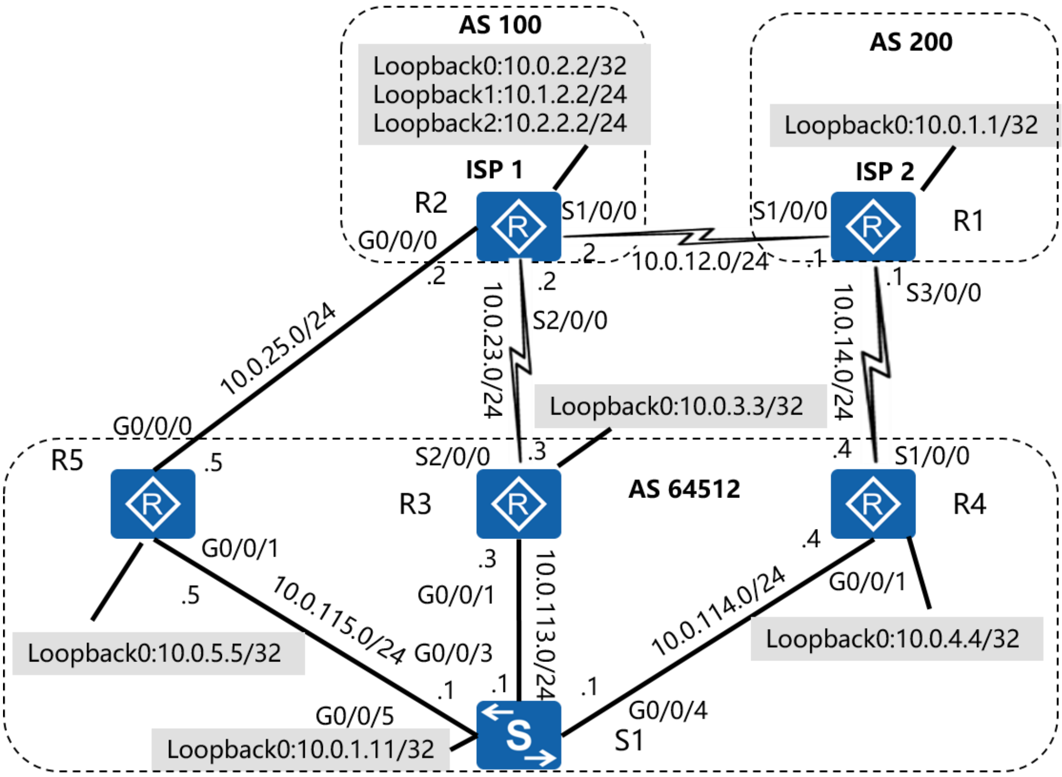

图3-5 BGP多宿主

场景

你是公司的网络管理员。公司的网络采用了BGP协议接入了运营商ISP1。公司自己采用了私有的AS号64512，ISP1的AS号为100，公司从两台路由器分别接入ISP1。起初公司采用默认路由的方式通过运营商接入Internet，随着公司的

发展，这种默认路由的接入方式已不能满足选路的需求，你需要把一部分Internet路由引入到公司的AS中。一段时间以后，公司又租用了一条线路接入到ISP2，ISP2的AS号为200。最终公司实现了由BGP选路的多宿主网络。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置物理接口及Loopback接口的IP地址和掩码。注意各Loopback 0接口均使用32位掩码。

```
<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
```

配置完成后，测试直连链路的连通性。

```
<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
```

步骤二.　配置 IGP 及 BGP

在AS 64512内部使用OSPF作为IGP，所有设备属于区域0。

R3的G0/0/1和Loopback 0连接的运行OSPF。

```
[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
```

R4的G0/0/1和Loopback 0连接的网段运行OSPF。

```
[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
```

R5的G0/0/1和Loopback 0连接的网段运行OSPF，

```
[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
```

在S1上创建Vlan13，配置Vlanif地址与R3进行互联。

创建Vlan14，配置Vlanif地址与R4进行互联。

创建Vlan15，配置Vlanif地址与R5进行互联。

互联接口使用Access模式，Vlanif 13、Vlanif 14 、Vlanif 15和Loopback 0连接的网段运行OSPF。

```
[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
```

检查是否学习到其他设备的Loopback 0接口连接网段的路由。

```
<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
```

在R2、R3、R5上配置EBGP。

AS的规划如图所示，EBGP全部使用物理接口地址建立对等体关系。S1不运行BGP。

```
[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
```

完成后检查BGP邻居是否完全建立。

```
[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
```

步骤三. 配置仅使用缺省路由归属到单个 ISP

默认情况下，BGP的负载分担是关闭的。在所有的路由器上打开负载分担，设置为最大4条路径。

```
[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
```

在R2上创建Loopback 1和Loopback 2，地址分别为10.1.2.2/24和10.2.2.2/24。使用**network**命令将这两个网段发布到BGP中，

```
[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
```

在R3和R5上分别检查是否学习到该路由。

```
[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
```

在这里，R3接入到ISP1的线路为主用线路，R5接入到ISP1的线路为备用线路。在R3和R5上分别将OSPF的路由通过**import-route**命令引入进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
```

在R3和R5上向区域0内强制发布默认路由，并采用第一类外部路由类型发布。设置R3发布的默认路由cost为20，R5发布的默认路由cost为40。

```
[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
```

在S1上查看路由表。

```
[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       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
```

在S1上查看到达地址10.1.2.2经过的路径。

缺省情况下，设备的ICMP端口不可达报文的发送功能处于未使能状态，所以我们首先在R2上执行命令**icmp port-unreachable send**，使能设备的ICMP端口不可达报文的发送功能。

```
[R2]icmp port-unreachable send
```

```
[S1]tracert  10.1.2.2
 traceroute to  10.1.2.2(10.1.2.2), max hops: 30 ,packet length: 40,press CTRL_C to break
 1 10.0.113.3 3 ms  2 ms  48 ms
 2 10.1.2.2 19 ms  19 ms  18 ms
```

此时S1选择使用从R3上学习到的默认路由，即通过主用线路访问10.1.2.2。

关闭R3的S2/0/0，模拟公司到运营商的线路故障。

```
[R3]interface s2/0/0
[R3-Serial2/0/0]shutdown
[R3-Serial2/0/0]quit
```

待路由收敛后在S1上查看路由表。并检查到10.1.2.2的连通性。

```
[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       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

 [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
```

可以观察到S1的路由表没有变化，仍旧是通过R3访问目标网络。

由于是上联链路故障，而下联链路中S1是通过比较R3与R5下发默认路由的cost值的大小，最终选择R3下发的默认路由，二者互不影响，所以网络无法正常工作。

恢复R3的S2/0/0，关闭R3的G0/0/1，模拟R3的下联线路故障。

查看路由收敛情况，并检查连通性。

```
[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 : 12        Routes : 12

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       LoopBack0
       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   10.0.114.1      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

[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
```

此时S1通过R5学习到默认路由，即通过备用链路访问目标网络。

恢复R3的G0/0/1端口。

```
[R3]interface g0/0/1
[R3-GigabitEthernet0/0/1]undo shutdown
[R3-GigabitEthernet0/0/1]quit
```

步骤四. 配置使用缺省路由和部分过滤路由归属到单个 ISP

配置R3、R4、R5到S1的IBGP邻居关系，并添加**next-hop-local**参数，保证S1能学习到从ISP发来的路由更新消息。

```
[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
```

观察S1是否学习到了10.1.2.0/24和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       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
       10.1.2.0/24  IBGP    255  0            RD  10.0.113.3      Vlanif13
       10.2.2.0/24  IBGP    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
```

这时我们希望通过BGP来影响选路，在R3上添加路由策略**policy_r3**，过滤掉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]quit
[R3]bgp 64512
[R3-bgp]peer 10.0.113.1 route-policy policy_r3 export
[R3-bgp]quit
```

在R5上添加路由策略**policy_r5**，过滤掉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
```

在S1上观察路由表的变化。

```
[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       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
      10.1.2.0/24  IBGP     255  0            RD   10.0.115.5      Vlanif15
      10.2.2.0/24  IBGP     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
```

此时到达网络10.1.2.0/24的下一跳是R5，到达网络10.2.2.0/24的下一跳是R3。

关闭R3的S2/0/0。

```
[R3]interface s2/0/0
[R3-Serial2/0/0]shutdown
[R3-Serial2/0/0]quit
```

观察S1的路由变化，测试到10.1.2.2的连通性。

```
[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       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
      10.1.2.0/24   IBGP    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
```

此时只有路由10.1.2.0/24，因为R5上做了策略过滤掉了路由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
```

恢复R3的S2/0/0接口。

```
[R3]interface s2/0/0
[R3-Serial2/0/0]undo shutdown
[R3-Serial2/0/0]quit
```

步骤五. 配置使用 BGP 路由归属到多个 ISP

这时公司又申请了一条Internet线路接入ISP2，为了采用BGP来进行路径选择，这里先删除前面试验中OSPF发布的默认路由。

```
[R3]ospf 1
[R3-ospf-1]undo default-route-advertise
[R3-ospf-1]quit

[R5]ospf 1
[R5-ospf-1]undo default-route-advertise
[R5-ospf-1]quit
```

删除R3和R5上对路由进行过滤的策略。

```
[R3]bgp 64512
[R3-bgp]undo peer 10.0.113.1 route-policy policy_r3 export
[R3-bgp]quit
[R3]undo route-policy policy_r3

[R5]bgp 64512
[R5-bgp]undo peer 10.0.115.1 route-policy policy_r5 export
[R5-bgp]quit
[R5]undo route-policy policy_r5
```

删除R3和R5上将OSPF路由引入到BGP中的命令。

```
[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
```

在R1和R2、R1和R4之间建立EBGP邻居关系，使ISP2也能传送10.1.2.0/24和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
```

在S1上观察路由10.1.2.0/24和10.2.2.0/24，注意当前的选路原则。

```
[S1]display bgp routing-table

 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

Total Number of Routes: 6
      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
```

现在我们希望公司使用连接到ISP2的新线路来访问网络10.2.2.0/24。在R4上设置路由策略**policy_r4**，将该路由的本地优先属性改为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
```

将该路由策略对S1发布。

```
[R4]bgp 64512
```

```
[R4-bgp]peer 10.0.114.1 route-policy policy_r4 export
[R4-bgp]quit
```

在S1上检查BGP路由表。

```
[S1]display bgp routing-table

 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

Total Number of Routes: 6
      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
```

此时S1使用通过R4连接的ISP2获得路由到达网段10.2.2.0/24。

关闭R4的S1/0/0端口模拟故障。

```
[R4]interface s1/0/0
[R4-Serial1/0/0]shutdown
[R4-Serial1/0/0]quit
```

查看S1上BGP路由表的变化。

```
[S1]display bgp routing-table

 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

 Total Number of Routes: 4
      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
```

此时S1通过R3所连接的ISP1获得路由10.1.2.0/24和10.2.2.0/24。

打开R4的S1/0/0接口。

```
[R4]interface s1/0/0
[R4-Serial1/0/0]undo shutdown
[R4-Serial1/0/0]quit
```

检查S1上的BGP路由表，查看是否恢复。

```
[S1]display bgp routing-table

 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

Total Number of Routes: 6
      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
```

附加实验：思考并验证

思考在步骤三中，关闭R3的S2/0/0接口后，虽然公司到运营商的主用链路发生故障，但是此时R5与ISP1之间的备用链路仍然正常，该如何解决此时的连通问题？

在这个例子中，多归属到两个运营商，怎样实现对同一网段入流量的负载分担？

最终设备配置

```
<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
#
 icmp port-unreachable send
#
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
```

实验 3-6 BGP 故障排除

学习目的

- 掌握BGP邻居无法建立的故障排除
- 掌握BGP相关Debug命令的使用

拓扑图

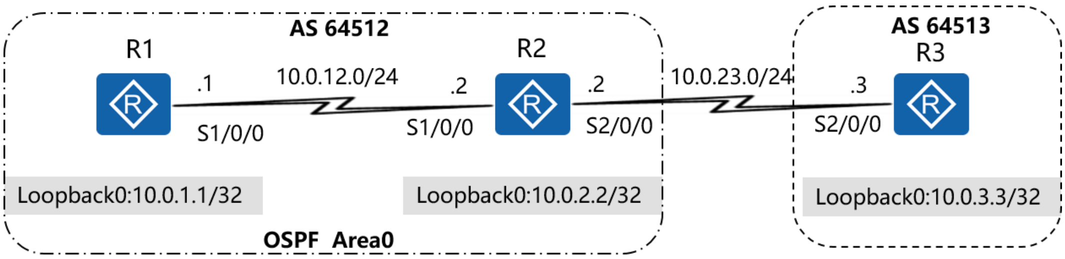

图3-6 BGP故障排除

场景

你是公司的网络管理员。公司的网络采用了BGP协议作为路由协议。公司的网络由多个自治系统组成，不同的分支机构使用了不同的AS号。现在你已经完成公司网络的搭建工作，可是在配置BGP的过程中你遇到了不少问题，最后你通过故障排除的思想和方法，你成功的找到了各种错误，并进行了网络的故障排除。

学习任务

步骤一. 基础配置与 IP 编址

首先给所有路由器配置物理接口及Loopback接口的IP地址和掩码，各Loopback接口均为32位掩码，地址的规划如拓扑图中所示。

```
<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
```

配置完成后，测试直连链路的连通性。

```
[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
```

结果显示直连链路连通性正常。

步骤二. 配置 IGP 及 BGP

在这个场景中，在AS 64512内部我们使用OSPF作为IGP，所有设备属于区域0。各路由器使用Loopback0的地址作为Router ID。R1的S1/0/0和Loopback 0所连接的网段运行OSPF。

```
[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
```

R2的S1/0/0和Loopback 0所连接的网段运行OSPF。

```
[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
```

配置完成以后检查R1到R2的Loopback0地址的连通性。

```
[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
```

在R1、R2之间配置IBGP，在R2和R3之间配置EBGP，均使用Loopback地址建立对等体关系，为保证路由信息的正常传递，在R2上针对R1配置**Next-hop-local**，在R3上你错误的将对等体10.0.2.2的AS号配置为64514了。

```
[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
```

步骤三. 排除对等建立的故障

在完成上面的配置以后，你发现路由器之间对等体关系并没有建立起来，首先在R2上观察对等体关系。

```
[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
```

从上面的输出中我们可以看到，10.0.1.1的状态为Active，10.0.3.3的状态为Idle。BGP对等体关系建立正常的情况下，State的状态应为**Established**。长时间停留在其他状态均为故障的现象，下面将逐步排除各种故障。

一般情况下，当对等体IP地址对于本地路由器是不可达的时，对等体状态会显示为Idle。也就是路由器根本没有向对等体发起TCP连接。当对等体IP地址可达，但TCP连接建立存在问题时，就可以看到对等体之间停留在Active状态。

首先考虑R2与R3之间的对等体关系的问题，先检查R2和R3之间Loopback地址之间的连通性。

```
[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
```

发现R2与R3的Loopback0接口地址之间的通讯存在问题。

检查R2的路由表。

```
[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       LoopBack0
      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
      10.0.23.0/24  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
      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
```

发现没有到达地址10.0.3.3的路由。

检查R3的路由表。

```
[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      LoopBack0
    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
    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
```

也没有到达R2的Loopback0地址10.0.2.2的路由。

对于不同的AS，使用静态路由可以实现相邻路由器的Loopback接口地址之间的可达性。

在这里，我们需要在R2和R3上分别添加到对端Loopback接口连接网段的静态路由。

```
[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
```

检查R2到R3之间的连通性。

```
[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
```

检查R2的BGP对等体关系。

```
[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 05:23:27      Active       0
```

```
10.0.3.3        4        64513        0        0    0 05:23:02    Active        0
```

R2和R3之间的状态由刚才的Idle状态转变为Active状态。

接下来考虑R1和R2之间的对等体关系，刚才在配置OSPF之后已经验证了R1的Loopback接口地址到R2的Loopback接口地址之间的连通性。

BGP是通过TCP 179端口进行通信的。首先检查路由器的179端口是否处于打开状态。

在R1和R2上分别检查TCP状态。

```
[R1]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           10.0.2.2:0            0       Listening
37a3321c   76  /3     0.0.0.0:443           0.0.0.0:0             23553   Listening

[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
390888bc 239/4      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
```

可以看到，对等体之间对应地址的179端口均处于监听状态。显示BGP协议在单个路由器上工作正常。

然后通过**debug**命令在R1上观察是否收到R2发来的BGP数据包。

```
<R1>terminal monitor
<R1>terminal debugging
<R1>debugging tcp packet
Nov 16 2016 16:34:56.300.1+00:00 R1 SOCKET/7/TCP PACKET:
TCP debug packet information:
1479314096: Input: no port,
(src = 10.0.12.2:64316,dst = 10.0.1.1:179,VrfIndex = 0,seq = 4034991477,
ack = 0,datalen = 0,optlen = 4,flag =   SYN ,window = 16384,ttl = 0,tos = 0,MSS = 0)

Nov 16 2016 16:34:56.300.2+00:00 R1 SOCKET/7/TCP PACKET:
TCP debug packet information:
1479314096: Output: task = (0), socketid = 0,
(src = 10.0.1.1:179,dst = 10.0.12.2:64316,VrfIndex = 0,seq = 0,
ack = 4034991478,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
```

从上面的输出中可以看到R2发来的目的端口号为179的数据包的源地址为10.0.12.2。查看拓扑后发现该地址为R2的Serial1/0/0的接口地址。

在建立对等体关系时，使用的是R2的Loopback接口地址，这就造成R1和R2之间的对等体关系建立不起来。所以我们需要在建立对等体关系时需使用**connect-interface**指定更新源地址。

同样的，R2和R3之间也存在这个问题。在这里，一并进行修改。

```
[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
```

修改完成以后再次在R2上检查对等体关系。

```
[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         Idle        0
```

看到R1与R2之间的状态已经是Established状态了。

在R3上使用**debug**命令测试是否收到BGP数据包，并检查数据包的内容。

```
<R3>terminal monitor
<R3>terminal debugging
<R3>debugging ip packet
Nov 16 2016 16:52:11.880.4+00:00 R3 IP/7/debug_case:
```

```
Receiving, interface = Serial2/0/0, version = 4, headlen = 20, tos = 192,
pktlen = 40, pktid = 3446, offset = 0, ttl = 1, protocol = 6,
checksum = 37526, s = 10.0.2.2, d = 10.0.3.3
prompt: Out interface if is loopback0 and packet is going to IP Distribute.

<R3>undo debugging all
Info: All possible debugging has been turned off
```

注意到R3收到的BGP数据包TTL值为1。对于EBGP来说，路由器向对等体发送的数据包默认的TTL就是1。

在这个场景中，R2和R3之间使用Loopback地址建立对等体关系。从R2的Loopback地址到R3的Loopback地址有2跳。所以该数据包在还没有到达R2的Loopback地址之前就因为TTL超时被丢弃了。

为了解决这个问题，我们需要修改EBGP对等体之间发送的数据包的TTL值。

```
[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
```

配置完成以后，在R2上重新检查对等体之间的关系。

```
[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        Idle       0
```

R2和R3之间仍停留在Idle状态。

在R3上检查BGP的错误。

```
[R3]display bgp error
Error Type      : Peer Error
 Date/Time       : 2016/11/16 17:01:48
 Peer Address    : 10.0.2.2
```

```
 VRF Name         : Public
 Error Info      : Incorrect remote AS

 Error Type      : Peer Error
 Date/Time        : 2016/11/16 17:03:06
 Peer Address   : 10.0.2.2
 VRF Name         : Public
 Error Info      : Incorrect remote AS

 Error Type      : Peer Error
 Date/Time        : 2016/11/16 17:03:38
 Peer Address   : 10.0.2.2
 VRF Name         : Public
 Error Info      : Incorrect remote AS

<R3>terminal debugging
<R3>debugging bgp packet
Nov 16 2016 17:10:02.860.2+00:00 R3 RM/6/RMDEBUG:
 BGP.Public: Err/SubErr: 2/2 Errdata: 41040000fc00
 Identified in OPEN MSG from 10.0.2.2.

Nov 16 2016 17:10:02.860.3+00:00 R3 RM/6/RMDEBUG:

Nov 16 2016 17:10:02.860.4+00:00 R3 RM/6/RMDEBUG:
 BGP.Public: Err/SubErr: 2/2 Errdata: 41040000fc00
 Identified in OPEN MSG from 10.0.2.2.
```

出现的是AS号错误的消息。

使用**debug**命令来诊断该错误。

从上面的输出中我们可以看到，错误号/子错误号是2。翻阅文档可知，该错误就是AS号不匹配。接下来我们在R3上修改对等体的AS号。

```
[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
```

这时再检查R2和R3之间的对等体关系。

```
[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
```

步骤四. BGP 安全

BGP通常用在骨干网上，其安全性显得尤为重要。一旦BGP路由器被攻击，将会造成大面积网络瘫痪。

通常为了防止非法恶意用户冒充合法路由器与BGP路由器建立对等体关系，在BGP对等体会话之间会设置MD5认证。

R1和R2之间开启MD5的认证。首先设置一个错误的密码。在R1上设置密码为huawei，在R2上设置密码为123，观察对等体关系的变化情况。

```
[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
```

在R1上重置BGP对等体关系，这时可以观察到R1与R2的对等体关系已停留在Connect和Active状态。无法进入Established状态。

```
<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
```

我们把R2的密码改成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
```

等待约半分钟，再次查看对等体关系。

```
[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
```

这时，R1和R2之间的对等体关系已经到达了Established状态。

在这个场景中AS 64512的管理员不希望AS64513中的路由器看自己的真实AS号。

fake-as的功能就能达到这个目的，该命令可以用来为本端对等体指定一个伪AS号。

我们在R2上针对R3配置这条命令，伪装以后的AS号为100。

同时，也要修改R3配置的BGP对等体R2的AS号。

```
[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
```

在R3上检查对等体，发现R2的AS号已经变成了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
```

在R2上发布自己的Loopback 0接口所在网段，观察在R3上学习到的BGP路由的AS-Path属性。

```
[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
```

可以看到在R3上学习到的路由10.0.2.2/32的AS-Path是100，即R3认为该路由是在AS100始发的，这样就达到了隐藏AS 64512的目的。

BGP还提供了一种安全特性：GTSM。该特性通过检测IP报文头中的TTL值是否在一个预先定义好的范围内来对路由器进行保护。也就是说，若收到的BGP报文的TTL值超出了预先设定的值的范围以后，就丢弃该报文。注意GTSM和ebgp-max-hop功能均会影响到发送出去的BGP报文的TTL值，存在冲突。只能对同一对等体或对等体组使能两种功能中的一种。

在这个场景中，我们对R2和R3之间的链路开启该特性，观察BGP报文的交互。首先在R2的系统视图下设置违背GTSM规则的缺省动作，在这里我们选择**drop**，即不符合要求的BGP数据包将会被丢弃。

```
[R2]gtsm default-action drop
```

然后在BGP视图下针对R3开启GTSM特性。在配置之前需先要删除**ebgp-max-hop**的配置，注意这里从R2和R3是直连的，所以这里的**valid-ttl-hops**参数填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
```

在R1和R3上进行相同的操作。

```
[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
```

检查R2和R3之间的对等体关系。

```
[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
```

在R3上观察使用了该命令以后BGP报文TTL值的变化。

```
<R3>terminal monitor
<R3>terminal debugging
<R3>debugging ip packet
Nov 16 2016 17:24:45.340.1+00:00 R3 IP/7/debug_case:
Receiving, interface = S2/0/0, version = 4, headlen = 20, tos = 192,
pktlen = 59, pktid = 4015, offset = 0, ttl = 255, protocol = 6,
checksum = 37449, s = 10.0.2.2, d = 10.0.3.3
prompt: Receiving IP packet from S2/0/0

45 c0 00 3b 0f af 00 00 ff 06 92 49 0a 00 02 02
0a 00 03 03
```

```
Nov 16 2016 17:24:45.340.2+00:00 R3 IP/7/debug_case:
Receiving, interface = Serial2/0/0, version = 4, headlen = 20, tos = 192,
pktlen = 59, pktid = 4015, offset = 0, ttl = 255, protocol = 6,
checksum = 37449, s = 10.0.2.2, d = 10.0.3.3
prompt: IP Process By Board Begin!

45 c0 00 3b 0f af 00 00 ff 06 92 49 0a 00 02 02
0a 00 03 03

<R3>undo debugging all
Info: All possible debugging has been turned off
```

这时看到R3接收到的从R2发来的数据包的TTL值为255，而不是EBGP的默认值1。为了确认GTSM会丢弃TTL不合规定的BGP数据包，我们首先在R3上打开GTSM日志记录功能，当有数据包被GTSM丢弃时就记录相应信息。

```
[R3]gtsm log drop-packet all
```

然后在R2上配置**ebgp-max-hop**命令，使R2发给R3的BGP数据包的TTL值小于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
```

等待一段时间以后，可以看到R2和R3之间的对等体关系进入了IDLE状态。现在R3上查看GTSM统计信息可以发现，已经有数据包被GTSM丢弃了。

```
[R3]

Nov 16 2016 17:30:44+00:00 R3 %%01BGP/3/STATE_CHG_UPDOWN(l)[6]: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               29              15              14
 0      BGPv6             0               0               0
 0      OSPF              0               0               0
 0      LDP               0               0               0
```

```
 0       OSPFv3              0              0              0
 0       RIP                 0              0              0
--------------------------------------------------------------
```

把R2的配置修改回来，间隔一段时间以后再观察是否有数据包被丢弃掉。

```
[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                 69             37             32
 0       BGPv6               0              0              0
 0       OSPF                0              0              0
 0       LDP                 0              0              0
 0       OSPFv3              0              0              0
 0       RIP                 0              0              0
--------------------------------------------------------------
```

这时我们看到已经没有新的数据包被丢弃掉了。

附加实验：思考并验证

GTSM主要用于防止哪种类型的攻击？

在联盟环境中是否可以使用fake-as？

最终设备配置

```
<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
```

第四章 路由控制

实验 4-1 路由引入与路由控制

学习目的

- 掌握OSPF与ISIS相互路由引入的配置方法
- 掌握通过地址前缀列表过滤路由信息的配置方法
- 掌握通过Route-policy过滤路由信息的配置方法

拓扑图

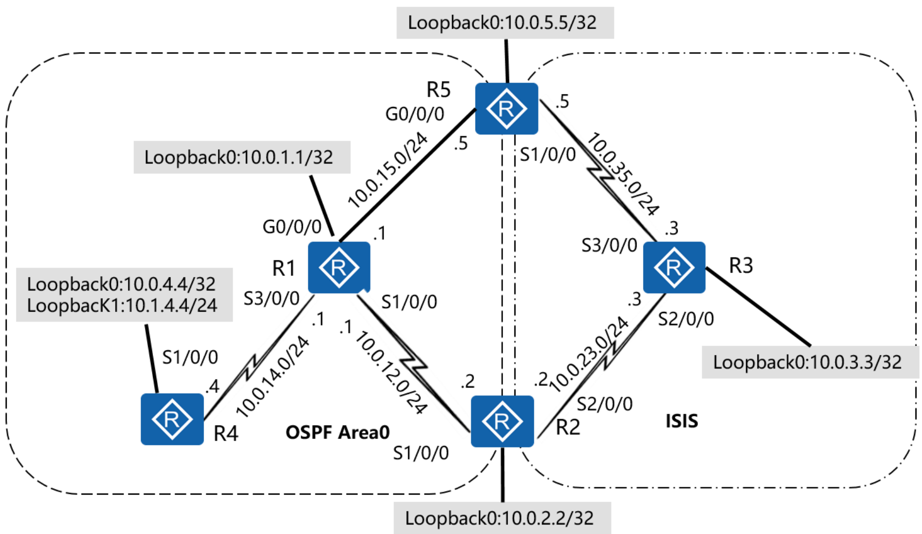

图4-1 路由引入与路由控制

场景

你是你们公司的网络管理员。公司网络中有两部分路由区域，一部分运行OSPF,另外一部分运行RIP。为了实现网络互通，你需要配置路由相互引入。在典型的双点双向路由引入中，会出现一些问题。为了解决这些问题，避免可能出现的路由环路和次优路径的产生，现在你需要使用到前缀列表和路由策略对路由进行控制。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置物理接口以及Loopback接口的IP地址和掩码。注意各Loopback 0接口均使用32位掩码。

```
<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
```

配置完成后，测试直连链路的连通性。

```
[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
```

步骤二. 配置 IGP

在R1，R2，R4，R5上运行OSPF，所有设备属于区域0。

配置R1的S1/0/0，S3/0/0，G0/0/0和Loopback0连接的网段运行OSPF，

```
[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
```

配置R2的S1/0/0连接的网段运行OSPF。

```
[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
```

配置R4的S1/0/0和Loopback 0连接的网段运行OSPF。

```
[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
```

配置R5的G0/0/0连接的网段运行OSPF。

```
[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
```

检查是否学习到其他设备上的Loopback0接口连接网段的路由。

```
[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       LoopBack0
       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       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
      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       LoopBack0
        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       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   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       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
        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       LoopBack0
       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       Serial1/0/0
     10.0.14.255/32  Direct  0    0         D   127.0.0.1       Serial1/0/0
       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	LoopBack0
10.0.5.5/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/0
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	GigabitEthernet0/0/0
10.0.15.255/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/0
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

在R2，R3，R5上配置RIP。

R2的S2/0/0和Loopback0连接的网段运行ISIS，

```
[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
```

R3的S2/0/0，S3/0/0和Loopback0连接的网段运行ISIS，

```
[R3]rip
[R3-rip-1]version 2
[R3-rip-1]network 10.0.0.0
[R3-rip-1]quit
[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
```

R5的S1/0/0和Loopback0连接的网段运行ISIS，

```
[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
```

检查是否学习到其他设备上的Loopback0地址。

```
 [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
```

步骤三. 配置前缀列表过滤路由信息

在R1上创建静态路由1.1.1.1/32，1.1.1.0/24，1.1.1.0/25，1.1.0.0/16，1.0.0.0/8分别指向NULL 0接口。并使用**import-route static**命令将这些路由引入OSPF中。

```
[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
```

在R4上检查是否接收到R1上添加的这几条静态路由。

```
[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       LoopBack0
     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       Serial1/0/0
```

10.0.14.255/32	Direct	0	0	D	127.0.0.1	Serial1/0/0
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

在R1上配置前缀列表pref_r1，匹配路由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

创建路由策略policy_r1，调用前缀列表pref_r1，控制R1上引入的静态路由信息。

[R1]route-policy policy_r1 permit node 10
[R1-route-policy]if-match ip-prefix pref_r1
[R1-route-policy]quit
[R1]ospf
[R1-ospf-1]import-route static route-policy policy_r1
[R1-ospf-1]quit

在R4上查看路由表。

[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	LoopBack0
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	Serial1/0/0
10.0.14.255/32	Direct	0	0	D	127.0.0.1	Serial1/0/0
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

步骤四. 使用 Route-policy 过滤路由信息避免路由环路

在R4上创建Loopback1，地址为10.1.4.4/24，使用**import-route direct**命令引入OSPF中。

```
[R4]interface LoopBack 1
[R4-LoopBack1]ip address 10.1.4.4 255.255.255.0
[R4-LoopBack1]quit
[R4]ospf 1
[R4-ospf-1]import-route direct
[R4-ospf-1]quit
```

在R2上将OSPF引入到ISIS 中，在R5上将ISIS引入到OSPF中。

```
[R2]rip
[R2-rip-1]import-route ospf
[R2]isis 1
[R2-isis-1]import-route  ospf
[R2-rip-1]quit

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

在R1上测试到10.1.4.4的连通性。

```
[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
```

发现不通，查看R1路由表。

```
[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	LoopBack0
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	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
10.0.23.0/24	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

此时R1上路由10.1.4.0/24的下一跳是R5。

再分别查看R2，R3，R5上路由表中的10.1.4.0/24这条路由。

[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
```

在R1上使用**tracert**命令查看到地址10.1.4.4的路径。

```
[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
```

此时发现环路。

原因是配置了路由引入后，R5既可以在ISIS域中学习到10.1.4.0/24，也可以在OSPF域中学习到这条路由。

由于ISIS的路由优先级比OSPF外部路由的优先级高，从而导致R5使用了从

ISIS域中学习到的路由。

R1能同时从R5和R4学习到这条路由。对于这两条外部路由，同是OSPF外部路由，所以比较二者的cost值。由于R1连接R5的是千兆链路，优于R1连接R4的串行链路，使R1选择使用从R5学习的路由，最终导致环路。

在R5上使用路由策略policy_r5，给路由10.1.4.0/24打上标记100。

```
[R5]acl number 2001
[R5-acl-basic-2001]rule 0 permit source 10.1.4.0 0.0.0.255
[R5-acl-basic-2001]quit
[R5]route-policy add_tag permit node 10
[R5-route-policy]if-match acl 2001
[R5-route-policy]apply tag 100
[R5-route-policy]quit
[R5]route-policy add_tag permit node 20
[R5-route-policy]quit
[R5]ospf 1
[R5-ospf-1]import-route rip route-policy add_tag
[R5-ospf-1]quit
```

在R1上查看OSPF路由信息。

```
[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
```

在R1上发现携带标记100的10.1.4.0/24路由，再次R1验证是从R5获得这条路由。

为了解决环路问题，在R5上向OSPF引入ISIS路由时，过滤掉路由10.1.4.0/24即可。

在R5上配置路由策略route_delete，控制向OSPF引入的RIP路由。

```
[R5]route-policy route_delete deny node 10
[R5-route-policy]if-match acl 2001
[R5-route-policy]quit
[R5]route-policy route_delete permit node 20
[R5-route-policy]quit
[R5]ospf 1
[R5-ospf-1]import-route rip route-policy route_delete
[R5-ospf-1]quit
```

在R1上查看路由表，

```
[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       LoopBack0
      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       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
  10.0.23.0/24      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
```

显示路由信息正常。

测试R1与地址10.1.4.4之间的连通性。

```
[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
```

在R1上检测到达地址10.1.4.4的路径。

缺省情况下，设备的ICMP端口不可达报文的发送功能处于未使能状态，所以我们首先在R4上执行命令icmp port-unreachable send，使能设备的ICMP端口不可达报文的发送功能。

```
[R4]icmp port-unreachable send


[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
```

此时环路问题解决，R1上路由10.1.4.0/24的下一跳为R4。

步骤五. 使用 Route-policy 修改路由优先级避免路由环路

查看R5的路由表。观察路由10.1.4.0/24的下一跳。

```
[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        GigabitEthernet
0/0/0
        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        GigabitEthernet
0/0/0
        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        GigabitEthernet
0/0/0
       10.0.14.0/24   OSPF      10     49            D     10.0.15.1        GigabitEthernet
0/0/0
       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         GigabitEthernet
0/0/0
       10.0.15.5/32   Direct   0      0            D     127.0.0.1         GigabitEthernet
0/0/0
     10.0.15.255/32   Direct   0      0            D     127.0.0.1         GigabitEthernet
0/0/0
       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
```

在R5上检测到达地址10.1.4.4的路径。

```
[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
```

发现通过路由过滤的方法虽然能解决环路问题，但是R5仍旧是从ISIS域学习到路由10.1.4.0/24。显然次优路径问题仍然没有得到解决。

为了解决环路问题以及次优路径问题，阻止R5从ISIS域中获得路由10.1.4.0/24，使R5从OSPF域学习到路由10.1.4.0/24。

删除R5上的策略route_delete。

```
[R5]ospf 1
[R5-ospf-1]undo import-route isis
[R5-ospf-1]quit
[R5]undo route-policy route_delete
[R5]ospf 1
[R5-ospf-1]import-route isis
[R5-ospf-1]quit
```

在R5上配置策略route_pref，将10.1.4.0/24的路由优先级修改为180，使其小于OSPF外部路由的路由优先级。

```
[R5]route-policy route_pref permit node 10
[R5-route-policy]if-match acl 2001
[R5-route-policy]apply preference 180
[R5-route-policy]quit
```

使用路由策略route_pref来控制RIP发布给OSPF的路由。

```
[R5]isis 1
[R5-isis-1]preference route-policy route_pref
```

检查R5，R1的路由表，观察路由10.1.4.0/24的下一跳。

```
[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
```

检查R1与地址10.1.4.4的连通性。

```
[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
```

在R1上检测到达地址10.1.4.4的路径。

```
[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
```

在R5上检测到达地址10.1.4.4的路径。

```
[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
```

此时环路问题解决。

在R1上，路由10.1.4.0/24的下一跳是R4。在R5上，路由10.1.4.0/24的下一跳是R1。次优路由也得到解决。

附加实验：思考并验证

在步骤三中，思考通过访问控制列表是否可以实现同样的效果，它和前缀列表的区别又是什么。

在步骤五中，思考为什么R3路由表中的10.0.15.0/24出现有两个下一跳，而10.0.12.0/24只有一个下一跳。

最终设备配置

```
<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
```

第五章 组播协议

实验 5-1 组播、IGMP 及 PIM DM 协议

学习目的

- 掌握路由器启动组播路由功能的配置方法
- 掌握配置接口IGMP功能的方法
- 掌握PIM DM的配置方法
- 掌握查看和测试组播的方法
- 掌握PIM一些高级特性的配置方法

拓扑图

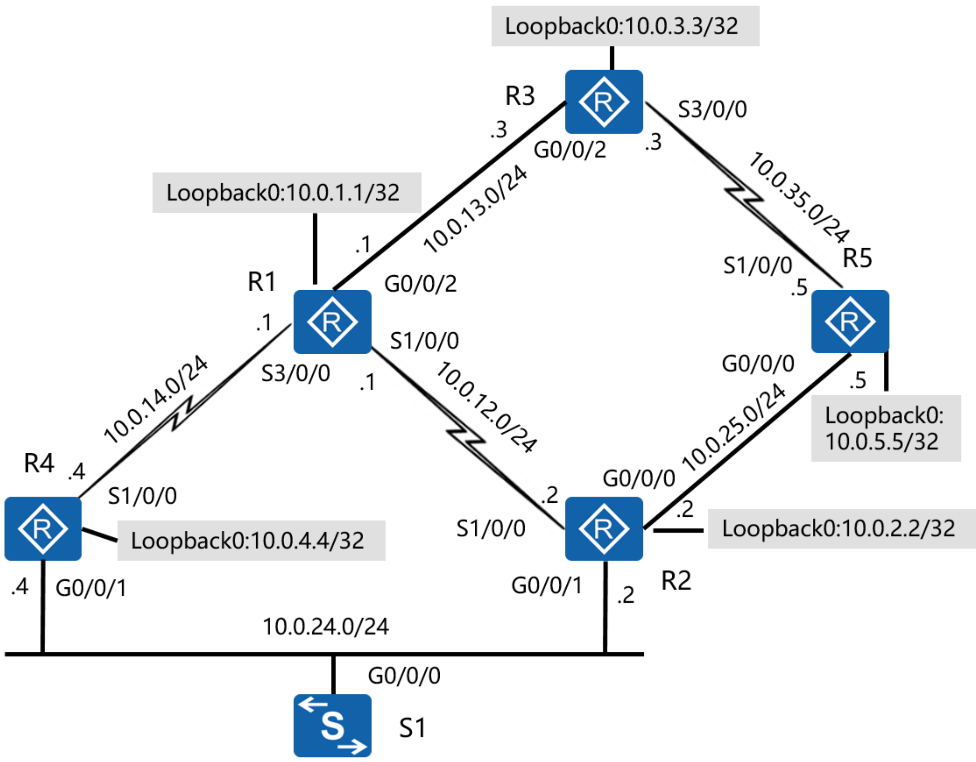

图5-1 组播、IGMP及PIM DM协议

场景

你是公司的网络管理员。公司准备使用组播来进行一些业务的转发。在当前网络上，网络规模较小，你决定使用PIM的DM模式来实现组播路由信息的学习。组播转发的实现过程中，你需要考虑到主机应用对IGMP不同版本的兼容，同时需要考虑使用合适的方式测试网络中组播是否正常工作。为了提升网络的效率和安全性，你采用了PIM DM的一些手段，包括PIM邻居的控制、嫁接和其他安全措施。同时实现网络的组播转发之前，你也遇到了一些网络故障，经过一些故障排除步骤，最终网络正常工作了。

学习任务

步骤一． 基础配置与 IP 编址

首先，如果采用真机的实验环境，为防止路由器收到其他网段的组播报文，建议对S2交换机（实现R1和R3的互联）做如下配置（如果是eNSP模拟器环境，则不需要进行下列配置）：

```
<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname S2
[S2]vlan 13
[S2-vlan13]quit
[S2]int GigabitEthernet 0/0/1
[S2-GigabitEthernet0/0/1]port link-t access
[S2-GigabitEthernet0/0/1]port default vlan 13
[S2-GigabitEthernet0/0/1]quit
[S2]int GigabitEthernet 0/0/3
[S2-GigabitEthernet0/0/3]port link-t access
[S2-GigabitEthernet0/0/3]port default vlan 13
[S2-GigabitEthernet0/0/3]quit
```

给所有路由器配置IP地址和掩码。配置时注意所有的Loopback接口配置掩码均为24位。

```
<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 32
[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 32
[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 32
[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]quit
[R4]interface LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 32
[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 32
[R5-LoopBack0]quit
```

配置完成以后，验证路由器之间的连通性。

```
[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
```

步骤二. 配置所有路由器启用组播路由功能

启用R1、R2、R3、R4及R5的组播路由功能。要开启组播功能，首先在系统视图下运行命令**multicast routing-enable**。

默认情况下，VRP平台的组播功能是关闭的。无论要使用PIM还是IGMP都需要先在全局下开启组播功能。

```
[R1]multicast routing-enable
```

对于要运行PIM DM的接口，在接口视图下运行**pim dm**开启组播路由协议。

```
[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
```

在R2、R3、R4、R5上进行相同的配置，注意在路由器与路由器的互联接口上开启PIM DM的功能。

```
[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
```

配置完成后，查看路由器PIM在接口上的运行状态。

```
[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
```

可以看到R1有3个接口运行了PIM，并且在每个接口上各有一个邻居（NbrCnt），同时我们还可以发现，在一个网段上接口IP地址较大的路由器将成为这个网段的DR。

查看R1上接口G0/0/2的PIM详细信息。

```
[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 join attribute (negotiated): disabled
     PIM generation ID: 0X53201C1B
     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 BFD: disabled
```

```
    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
    Number of routers on link not using join attribute: 2
    ACL of PIM neighbor policy: -
    ACL of PIM ASM join policy: -
    ACL of PIM SSM join policy: -
    ACL of PIM join policy: -
```

可以看到PIM DM默认的Hello间隔是30秒，Hello的保持时间是Hello间隔的3.5倍，也就是105秒。

查看R1的邻居列表，共有3台路由器与R1形成PIM邻居关系，邻居默认的DR优先级均为1。

```
[R1]display pim neighbor
VPN-Instance: public net
 Total Number of Neighbors = 3

 Neighbor        Interface           Uptime     Expires   Dr-Priority   BFD-Session
 10.0.13.3       GE0/0/2             00:09:02 00:01:16 1                  N
 10.0.12.2       S1/0/0              00:17:01 00:01:43 1                  N
 10.0.14.4       S3/0/0              00:16:08 00:01:37 1                  N
```

查看邻居R3的详细信息，Uptime表示邻居关系已经建立的时间，Expiry time表示PIM邻居还有多少时间就要超时，LAN delay表示传递Prune剪枝消息的延迟时间，Override interval表示否决Prune剪枝的时间间隔。

```
[R1]display pim neighbor 10.0.13.3 verbose
 VPN-Instance: public net
 Neighbor: 10.0.13.3
    Interface: GigabitEthernet0/0/2
    Uptime: 00:40:11
    Expiry time: 00:01:37
    DR Priority: 1
    Generation ID: 0X9AF72165
    Holdtime: 105 s
    LAN delay: 500 ms
    Override interval: 2500 ms
    State refresh interval: 60 s
    Neighbor tracking: Disabled
    PIM join attribute: Disabled
    PIM BFD-Session: N
```

步骤三. 配置 IGMP

在这个实验中，我们模拟组播用户连接在交换机S1上。在R2和R4的G0/0/1接口开启IGMP功能。要开启IGMP的功能，在接口模式下运行**igmp enable**。

```
[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
```

为了实验需要，在R2和R4的G0/0/1接口添加静态组播组。这样，该接口始终会转发目的地址为225.1.1.1的组播流量。

```
[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
```

默认情况下，VRP平台使用的IGMP版本为v2。从下面的输出中可以看到现在G0/0/1接口所在网段的查询器是10.0.24.2，即为R2。对于IGMP v2来说，选取网段上IP地址较小的那台路由器作为查询器。

```
[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): 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 of VPN-Instance: public net
```

```
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: 104 s
   Value of maximum query response time for IGMP: 10 s
   Querier for IGMP: 10.0.24.2
```

在R2上查看接口的静态IGMP组，可以看到225.1.1.1这个组是刚才我们手工添加的组播组。

```
[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
```

在接口下查看IGMP路由表。

```
[R2]display igmp routing-table
Routing table of VPN-Instance: public net
 Routing table
 Total 1 entry

 00001. (*, 225.1.1.1)
        List of 1 downstream interface
          GigabitEthernet0/0/1 (10.0.24.2),
                       Protocol: STATIC
```

若在接口上只配置了IGMP，没有配置PIM，且接口为查询器的情况下，才会生成IGMP路由表项。该路由表条目在R4上是看不到的，因为现在R2是网段10.0.24.0/24的查询器。

默认情况下，查询器的查询周期为60秒，为了加快用户加入组播组的速度，我们可以通过**igmp timer query**修改发送查询报文的时间间隔。

```
[R2]interface GigabitEthernet 0/0/1
[R2-GigabitEthernet0/0/1]igmp timer query 20
[R2-GigabitEthernet0/0/1]quit
```

配置完成以后，验证配置已生效。

```
[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)
```

开启Debugging后可以看到接口每隔20秒发送一次成员关系查询(general query)。

```
<R2>terminal monitor
<R2>terminal debugging
<R2>debugging igmp query send
Nov 17 2016 12:33:03.390.1+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)

Nov 17 2016 12:33:23.380.1+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)

Nov 17 2016 12:33:43.360.1+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
```

路由器的健壮系数描述了IGMP路由器的健壮程度。路由器默认的健壮系数为2，这里通过关闭接口的方式测试健壮系数。首先观察默认情况下IGMP查询消息的间隔。

```
<R2>terminal monitor
<R2>terminal debugging
<R2>debugging igmp query send
Nov 17 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>
Nov 17 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
Nov 17 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]
Nov 17 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]
Nov 17 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]
Nov 17 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)
```

在没有关闭接口之前，路由器的接口仍按照每20秒一次的间隔发送普遍查询消息，当接口被关闭又重新打开之后，前面2个查询报文的时间间隔为5秒。当路由器启动时会发送“健壮系数”次的“普遍组查询消息”，发送间隔是“IGMP普遍组查询消息的发送间隔”的1/4。

执行命令**robust-count**可配置IGMP健壮系数，注意该参数只有在IGMP v2和IGMP v3中才有效，在R2的G0/0/1上将健壮系数修改为3。

```
[R2-GigabitEthernet0/0/1]igmp robust-count 3
```

再使用Debugging观察普遍查询消息的发送。

```
[R2-GigabitEthernet0/0/1]shutdown
[R2-GigabitEthernet0/0/1]undo shutdown
Nov 17 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]
Nov 17 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]
Nov 17 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]
Nov 17 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
```

可以看到当健壮系数修改为3以后，当接口启用后前3个普遍组查询消息的时间间隔为5秒，从第四个普遍组查询消息开始时间间隔为20秒。

步骤四. 观察组播路由表

为了观察组播路由的传递，在该拓扑上启用OSPF作为单播路由协议。

```
[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
```

配置完成以后，检查各路由器已能学习到其他路由器的Loopback地址。

```
[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
```

为了模拟组播信息的传递，我们在R3上以自己的Loopback接口作为源地址，向目的地址225.1.1.1发送Ping数据包，模拟组播源。

```
[R3]ping -a 10.0.3.3 -c 300 225.1.1.1
```

间隔几分钟后，我们可以在其他所有路由器上看到组播路由表。在R2上查看组播路由表。

```
[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
```

可以看到2个条目。

第一个条目(*, 225.1.1.1)为该接口配置了静态IGMP组产生的。

第二个条目(10.0.3.3, 225.1.1.1)为组播流量进行扩散后在该路由器上产生的条目。

从输出中我们还可以看到对于R2来说，该组播流的上游路由器为10.0.25.5。

启用了PIM以后，路由器会采用单播路由表进行RPF检查，从下面的输出中可以看到，对于组播源10.0.3.3，RPF的邻居是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
```

步骤五.　调整 PIM DM 参数

有时我们希望流量不按照单播路由的路径流向目的地，就可以通过**rpf-route-static**静态修改RPF路径。这个试验中，我们把RPF路径由原来的10.0.25.5修改为10.0.12.1。

```
[R2]ip rpf-route-static 10.0.3.0 255.255.255.0 10.0.12.1
```

配置完成以后，可验证RPF邻居已变成了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
```

为了观察PIM的剪枝及嫁接消息，我们通过删除及添加IGMP静态组的方式来模拟用户的离开及加入。首先在R2上打开Debugging。

```
<R1>terminal monitor
<R2>terminal debugging
<R2>debugging pim join-prune
```

然后把R2的静态IGMP组225.1.1.1删除。

```
<R2>system-view
[R2]interface GigabitEthernet 0/0/1
[R2-GigabitEthernet0/0/1]undo   igmp static-group 225.1.1.1
Nov 17 2016 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)
Nov 17 2016 15:00:05.300.2+00:00 R2 PIM/7/JP:(public net): Upstream 10.0.12.1, Groups 1, Holdtime 210 (P012693)
Nov 17 2016 15:00:05.300.3+00:00 R2 PIM/7/JP:(public net): Group: 225.1.1.1/32 --- 0 joins 1 prunes (P012701)
Nov 17 2016 15:00:05.310.1+00:00 R2 PIM/7/JP:(public net): Prune: 10.0.3.3/32   (P012707)
Nov 17 2016 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)
Nov 17 2016 15:00:05.350.2+00:00 R2 PIM/7/JP:(public net): Upstream 10.0.12.1, Groups 1, Holdtime 207 (P012693)
Nov 17 2016 15:00:05.350.3+00:00 R2 PIM/7/JP:(public net): Group: 225.1.1.1/32 --- 0 joins 1 prunes (P012701)
Nov 17 2016 15:00:05.350.4+00:00 R2 PIM/7/JP:(public net): Prune: 10.0.3.3/32   (P012707)
```

可以看到R2立刻以组播地址224.0.0.13向上游接口发送剪枝消息，上游路由器的地址为10.0.12.1，此时225.1.1.1这个组播组已被剪枝。随后R1向R2发送消息确认剪枝。

然后再把刚才删除的静态IGMP组播组添加回去。

```
[R2-GigabitEthernet0/0/1]igmp static-group 225.1.1.1
Nov 17 2016 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)
Nov 17 2016 15:00:19.440.2+00:00 R2 PIM/7/JP:(public net): Upstream 10.0.12.1, Groups 1, Holdtime 0 (P012639)
```

Nov 17 2016 15:00:19.440.3+00:00 R2 PIM/7/JP:(public net): Group: 225.1.1.1/32 --- 1 joins 0 prunes (P012648)
Nov 17 2016 15:00:19.440.4+00:00 R2 PIM/7/JP:(public net): Join: 10.0.3.3/32 (P012654)
Nov 17 2016 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)
Nov 17 2016 15:00:19.480.2+00:00 R2 PIM/7/JP:(public net): Upstream 10.0.12.2, Groups 1, Holdtime 0 (P012639)
Nov 17 2016 15:00:19.480.3+00:00 R2 PIM/7/JP:(public net): Group: 225.1.1.1/32 --- 1 joins 0 prunes (P012648)
Nov 17 2016 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
```

这时R2立刻向上游以单播的形式发送了嫁接消息，加入225.1.1.1，同时R1也以单播的形式向R2回应了嫁接确认。

从这里可以总结出：**剪枝消息是以组播地址224.0.0.13发送的，而嫁接消息是以单播向上游发送的。**

有时我们希望组播流量只在规定的范围内传递，这时候可以在接口下通过**multicast boundary**为某个特定的组播组或组播地址段定界。

控制组播组225.1.1.2的流量不要传递到R4上，在R1连接到R4的接口上增加如下配置。

```
[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
```

在R3上模拟目的地址为225.1.1.2的组播流量。

```
[R3]ping -a 10.0.3.3 -c 300 225.1.1.2
```

等待在R2和R4上分别查看组播路由表，可以看到在R2上存在表项(10.0.3.3, 225.1.1.2)，而在R4上没有该组播组路由条目，说明组播流量并没有扩散到R4上。

```
[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:   -
```

默认情况下PIM DM选取接口IP地址较大的路由器作为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)
```

在R2上查看接口状态可以看到在与R5连接的接口上，R5是DR。我们可以通过修改接口的优先级来影响DR的选举，该优先级值是一个32bit长度的数值，默认值为1。在下面的例子中，将R2连接到R5的接口的优先级改成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)
```

可以看到当把路由器接口优先级更改到100以后，R2立刻抢占了DR的位置。

有时为了安全需要，我们希望面向用户侧的接口上不再收发PIM的Hello包，使用**pim silent**可实现该功能。

```
[R4]interface GigabitEthernet 0/0/1
[R4-GigabitEthernet0/0/1]pim silent
[R4-GigabitEthernet0/0/1]quit
```

配置完成以后检查PIM Silent已生效。

```
[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 join attribute (negotiated): disabled
PIM generation ID: 0X1420A003
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 BFD: disabled
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
Number of routers on link not using join attribute: 1
ACL of PIM neighbor policy: -
ACL of PIM ASM join policy: -
ACL of PIM SSM join policy: -
ACL of PIM join policy: -

附加实验：思考并验证

PIM的DM模式适合于用户比较多，比较密集的场景。

思考一下生活中哪些网络应用适合使用PIM DM模式的组播来实现数据转发？它们的特点是什么？

PIM的DM模式在应用在大规模网络上，有哪些劣势？

最终设备配置

```
<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
```

实验 5-2 PIM SM 及静态 RP

学习目的

- 掌握PIM SM协议的配置方法
- 掌握静态RP和RP负载均衡的配置方法

拓扑图

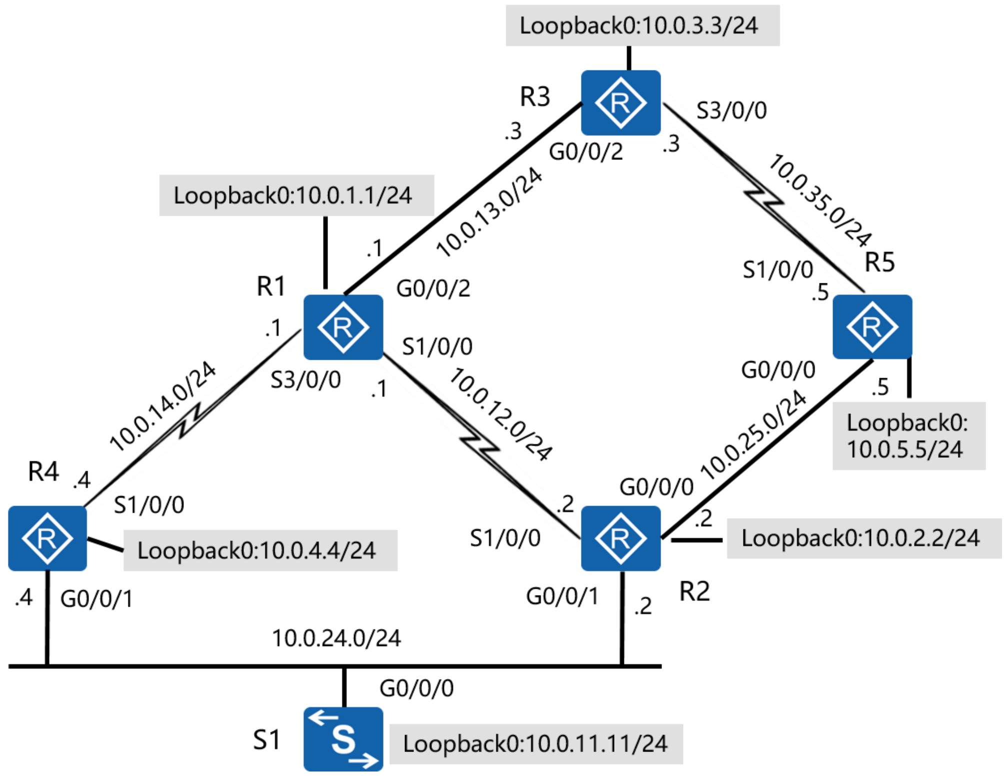

图5-2 PIM SM及动态RP

场景

你是公司的网络管理员。公司以前使用过PIM的DM模式来实现组播路由学习，但是后来发现随着组播应用的发展，组播用户分部越来越广，这时组播质量显示出一定程度的下降。为了提高组播的可靠性、安全性和效率，你决定使用PIM的SM模式来实现组播路由学习。

在PIM的SM模式中，你需要定义RP，作为SM模式的共享树的树根。但是在

实际的使用中，只是这样简单的配置是不够的，组播还需要实现RP之间的负载分担。

在实现网络的组播转发之前，你也遇到了一些网络故障，经过一些故障排除步骤，最终网络正常工作了。

学习任务

步骤一. 基础配置与 IP 编址

S2参与到本次实验（实现R1和R3的互联）。

```
<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname S2
[S2]vlan 13
[S2-vlan13]quit
[S2]int GigabitEthernet 0/0/1
[S2-GigabitEthernet0/0/1]port link-t access
[S2-GigabitEthernet0/0/1]port default vlan 13
[S2-GigabitEthernet0/0/1]quit
[S2]int GigabitEthernet 0/0/3
[S2-GigabitEthernet0/0/3]port link-t access
[S2-GigabitEthernet0/0/3]port default vlan 13
[S2-GigabitEthernet0/0/3]quit
```

给所有路由器配置IP地址和掩码。配置时注意所有的Loopback接口配置掩码均为24位。

```
[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]quit
[S1]interface loopback 0
[S1-LoopBack0]ip address 10.0.11.11 24
[S1-LoopBack0]quit
```

配置完成后，测试直连链路的连通性。

```
[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
```

在R1、R2、R3、R4、 R5和S1上启用OSPF路由协议。实现所有网络互通。

```
[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
```

配置完成后，稍等片刻，待OSPF邻居关系连接，路由信息交互完成后，测试S1与路由器R3的Loopback 0地址之间的连通性。

```
[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
```

如测试显示，网络工作正常。

步骤二. 配置所有路由器启用 PIM SM

启用R1、R2、R3、R4、R5和S1的组播路由功能。

```
[R1]multicast routing-enable

[R2]multicast routing-enable

[R3]multicast routing-enable

[R4]multicast routing-enable

[R5]multicast routing-enable

[S1]multicast routing-enable
```

在所有设备的所有接口上配置运行SM模式的PIM。

```
[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
```

配置完成后，查看R1、R5和S1的PIM邻居学习的情况。

```
<R1>display pim neighbor
VPN-Instance: public net
```

```
 Total Number of Neighbors = 3

 Neighbor        Interface           Uptime    Expires   Dr-Priority   BFD-Session
 10.0.13.3       GE0/0/2             00:01:25 00:01:20 1               N
 10.0.12.2       S1/0/0              00:01:47 00:01:28 1               N
 10.0.14.4       S3/0/0              00:00:56 00:01:19 1               N

[R5]display pim neighbor
 VPN-Instance: public net
 Total Number of Neighbors = 2

 Neighbor        Interface           Uptime    Expires   Dr-Priority   BFD-Session
 10.0.25.2       GE0/0/0             00:01:24 00:01:29 1               N
 10.0.35.3       S1/0/0              00:01:19 00:01:39 1               N

[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 00:01:28 1               N
 10.0.24.2       Vlanif1             00:01:24 00:01:41 1               N
```

从命令输出结果可以看到，PIM协议已经在网络中正确运行。

步骤三. 静态 RP 和静态 RP 负载均衡

通过给网络手动指定静态RP来控制网络中组播数据流。

在所有设备配置R1的S3/0/0接口作为网络中的静态RP。

```
[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模拟网络中组播用户接入的三层交换机，在S1的Loopback 0接口上开启IGMP功能。

```
[S1]interface LoopBack 0
[S1-LoopBack0]igmp enable
[S1-LoopBack0]quit
```

将S1的Loopback 0接口静态加入225.0.0.1组播组，模拟连接有225.0.0.1组播组的用户。

```
[S1]interface LoopBack 0
[S1-LoopBack0]igmp static-group 225.0.0.1
[S1-LoopBack0]quit
```

在R1、R4和S1上使用命令**display pim routing-table**查看PIM路由表。

```
[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]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: -
```

从命令输出结果可以看到，在指定静态RP的网络中，R1是网络中的RP。S1生成了一条经过R4到达RP路由器R1的组播路径。

创建ACL并应用到静态RP上，定义R1作为RP，服务的组播范围是225.0.0.0/24网段，定义R5作为RP，服务的组播范围是225.0.1.0/24网段。

```
[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
```

将S1的Loopback 0接口静态加入225.0.1.1组播组，模拟连接有225.0.1.1组播组的用户。

```
[S1]interface LoopBack 0
[S1-LoopBack0]igmp static-group 225.0.1.1
[S1-LoopBack0]quit
```

在S1、R2和 R5上使用命令**display pim routing-table**查看PIM路由表。

```
[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: -
```

从命令输出结果可以看到，S1针对225.0.0.1和225.0.1.1生成了二条组播路径。225.0.1.1组播路径经过R2到达RP路由器R5。

附加实验：思考并验证

PIM的SM模式适合于用户比较分散的场景。

思考一下生活中哪些网络应用适合使用PIM SM模式的组播来实现数据转发？它们的特点是什么？

最终设备配置

```
[R1]display current-configuration
[V200R007C00SPC600]
#
 sysname R1
#
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 Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/2
 ip address 10.0.13.1 255.255.255.0
 pim sm
#
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 10.0.25.25 2001
#
return
```

[R2]**display current-configuration**

```
[V200R007C00SPC600]
#
 sysname R2
#
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 10.0.12.2 255.255.255.0
 pim sm
#
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 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
#
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 Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/2
 ip address 10.0.13.3 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/3
 description VirtualPort
#
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
#
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 10.0.14.4 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/1
 ip address 10.0.24.4 255.255.255.0
 pim sm
#
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
#
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 10.0.35.5 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.5 255.255.255.0
 pim sm
#
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
```

```
#
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 router-id 10.0.11.11
 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
```

第六章 VLAN特性与配置

实验 6-1 VLAN 配置

学习目的

- 掌握Eth-trunk的配置方法

拓扑图

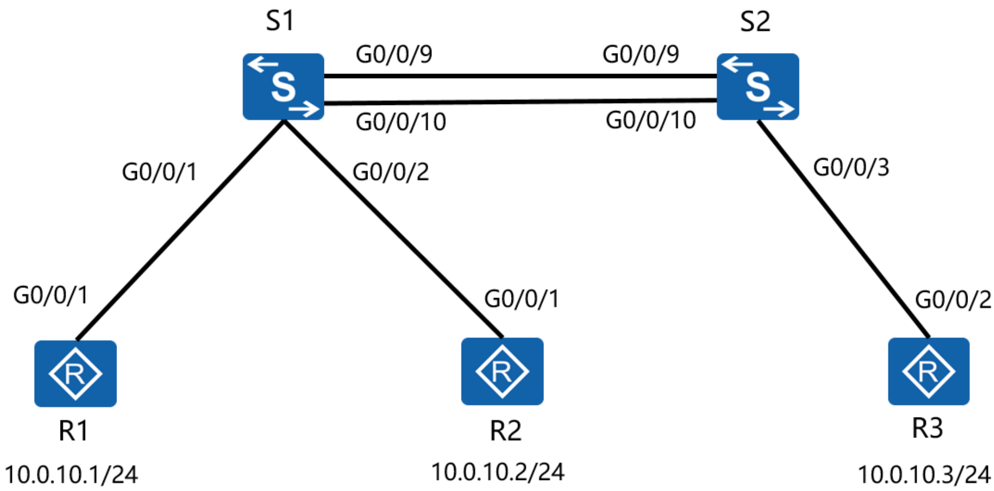

图6-1 VLAN配置

场景

你是公司的网络管理员。现在公司网络是由二台交换机组成的以太网环境。图中路由器模拟网络中的计算机，R3是一台服务器。为了优化这个网络，需要你提升S1和S2相连链路的速度和可靠性。

学习任务

步骤一. 基础配置与 IP 编址

给所有设备配置IP地址和掩码。

```
<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
```

给交换机配置名称。

```
<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
```

步骤二. Eth-trunk 链路聚合

Eth-trunk可以将二条或多条链路捆绑成一条链路以提升链路带宽和可靠性。将S1和S2的G0/0/9和G0/0/10接口加入到同一个Eth-trunk组可以实现企业需

求。

创建Eth-trunk接口。

```
[S1]interface Eth-Trunk 1
[S1-Eth-Trunk1]

[S2]interface Eth-Trunk 1
[S2-Eth-Trunk1]
```

配置Eth-trunk的工作模式为LACP模式。

```
[S1-Eth-Trunk1]mode lacp
[S1-Eth-Trunk1]quit

[S2-Eth-Trunk1]mode lacp
[S2-Eth-Trunk1]quit
```

将S1和S2的G0/0/9和G0/0/10接口加入到Eth-trunk接口。

```
[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
```

使用命令**dispay eth-trunk**查看配置情况。

```
[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
```

输出信息显示，此时链路运行模式为LACP模式，并且最大活动接口数阈值是8条链路。同时G0/0/9和G0/0/10接口都处于活动状态。

使用命令更改活动接口数阈值。

```
[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
```

查看Eth-trunk链路配置情况。

```
[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
```

从输出信息可以看到 ,G0/0/10接口状态改变为Unselect。实现了Eth-trunk聚合链路下一条链路传输数据，另一条链路备份的功能，提升了网络可靠性。

通过关闭S1的G0/0/9接口验证备份链路功能。

```
[S1]interface GigabitEthernet 0/0/9
[S1-GigabitEthernet0/0/9]shutdown
[S1-GigabitEthernet0/0/9]quit
```

查看Eth-trunk链路信息。

```
[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
```

从输出信息看到Eth-trunk链路中G0/0/9已经变为Unselect状态，G0/0/10由Unselect状态自动转换为Selected状态继续传输数据。由此可知，实现了链路备份功能。

最终设备配置

```
[S1]display current-configuration
!Software Version V200R008C00SPC500
#
sysname S1
#
interface Eth-Trunk1
 mode lacp
 max active-linknumber 1
#
```

```
interface GigabitEthernet0/0/9
 shutdown
 eth-trunk 1
#
interface GigabitEthernet0/0/10
 eth-trunk 1
#
return
```

[S2]**display current-configuration**

```
!Software Version V200R008C00SPC500
#
sysname S2
#
interface Eth-Trunk1
 mode lacp
 max active-linknumber 1
#
interface GigabitEthernet0/0/9
 shutdown
 eth-trunk 1
#
interface GigabitEthernet0/0/10
 eth-trunk 1
#
Return
```

实验 6-2 MUX VLAN (选做)

学习目的

- 掌握MUX VLAN的配置方法

拓扑图

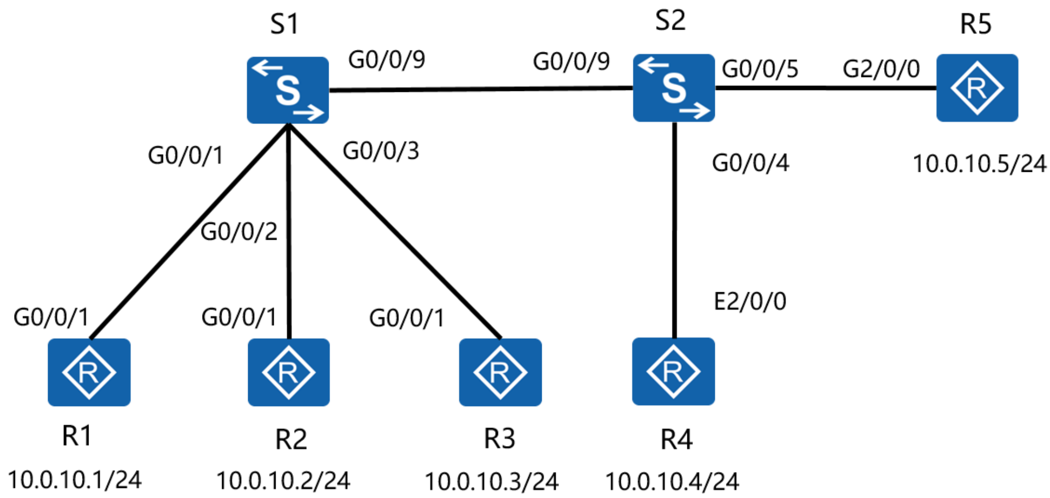

图6-2 MUX VLAN配置

场景

你是公司的网络管理员。现在公司网络是由二台交换机组成的以太网环境。图中路由器代表网络中的计算机。为了优化这个网络，需要你实现广播域的互相隔离。R1和R2处于相同的VLAN中，R3和R4分别处另一个VLAN中。公司策略需要所有PC均可以访问R5，R3和R4除了不能与R1、R2通信外也不能互相访问。

学习任务

步骤一. 基础配置与 IP 编址

给所有设备配置IP地址和掩码。

```
<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]
```

在R1上测试与R2、R3、R4和R5的连通性。

```
[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
```

步骤二. MUX VLAN

MUX VLAN可以实现处于相同网段的设备划入不同VLAN后，虽然二层通信是隔离的，但是还可以和同一个指定VLAN通信。并且还能实现禁止相同VLAN内不同设备之间的通信。

将VLAN 100配置为MUX VLAN的主VLAN ,VLAN 10和20配置为从VLAN。

通过配置各PC与交换机相连接口的类型实现所有PC均可和R4通信 ,R3和R4不能和其他VLAN通信的同时也不能互相通信。

配置VLAN 100为主VLAN并添加从VLAN配置。

```
[S1]vlan batch 10 20 100
[S1]vlan 100
[S1-vlan100]mux-vlan
[S1-vlan100]subordinate group 10
[S1-vlan100]subordinate separate 20
[S1-vlan100]quit

[S2]vlan batch 10 20 100
[S2]vlan 100
[S2-vlan100]mux-vlan
[S2-vlan100]subordinate group 10
[S2-vlan100]subordinate separate 20
[S2-vlan100]quit
```

将R5与S2连接的G0/0/5接口加入VLAN 100并开启MUX VLAN功能。

```
[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
```

将R1与S1连接的G0/0/1和R2与S1连接的G0/0/2接口加入VLAN 10并开启MUX VLAN功能。

```
[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
```

将R3与S1的G0/0/3和R4与S2的G0/0/4接口加入VLAN 20并开启MUX VLAN功能。

```
[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
```

使用命令**display mux-vlan**查看所有MUX VLAN信息。

```
[S1]display mux-vlan
Principal Subordinate Type          Interface
---------------------------------------------------------------------------
100         -              principal
100         20             separate          GE0/0/3
100         10             group             GE0/0/1            GE0/0/2
---------------------------------------------------------------------------

[S2]display mux-vlan
Principal Subordinate Type          Interface
---------------------------------------------------------------------------100        -
principal     GE0/0/5
100         20             separate      GE0/0/4
100         10             group
---------------------------------------------------------------------------
```

使用**ping**命令测试R1与R2、R3、R4、R5的连通性。

```
[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
```

使用**ping**命令测试R3与R2、R4、R5的连通性。

```
[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
```

从**ping**命令输出结果可知。 MUX VLAN中VLAN 10内部的R1和R2除了能够和R5通信外还能互相通信。VLAN 20内部的R3和R4仅仅能够和R5通信。

附加实验：思考并验证

属于二个不同的MUX VLAN之间的用户，互相通信是否能实现？

最终设备配置

```
[S1]display current-configuration
!Software Version V200R008C00SPC500
#
sysname S1
#
vlan batch 10 20 100
#
vlan 100
 mux-vlan
 subordinate separate 20
 subordinate group 10
#
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
#
return
```

[S2]**display current-configuration**

```
!Software Version V200R008C00SPC500
#
sysname SW2
#
vlan batch 10 20 100
#
vlan 100
 mux-vlan
 subordinate separate 20
 subordinate group 10
#
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
#
return
```

实验 6-3 VLAN 间通信

学习目的

- 掌握多臂路由的配置方法
- 掌握单臂路由的配置方法
- 掌握VLAN间通信的配置方法
- 掌握VLAN聚合的配置方法

拓扑图

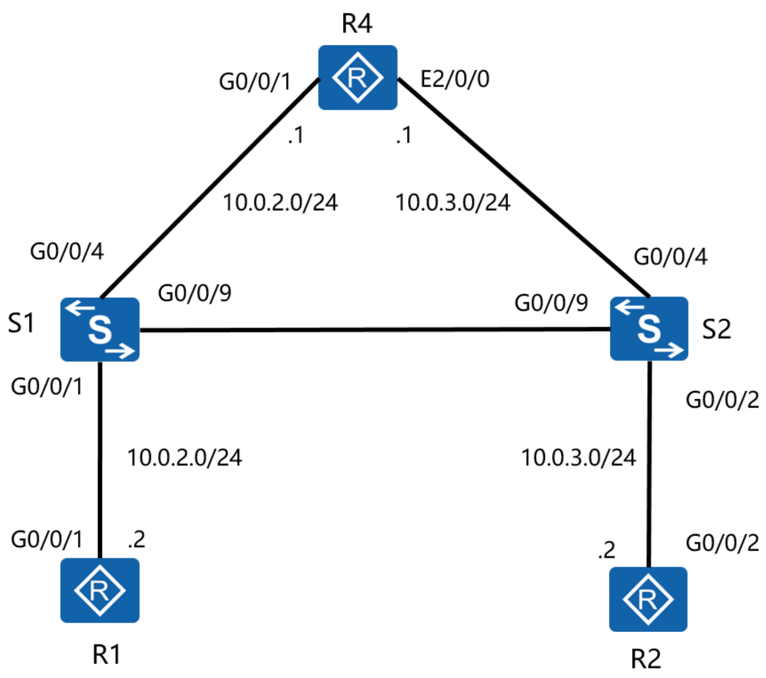

图6-3 VLAN间通信

场景

你是公司的网络管理员。现在公司网络是由二台交换机和一台路由器组成的以太网环境。图中R1和R2代表公司不同部门的PC ,分别加入了二个不同的VLAN。现在需要你实现R1和R2之间的通信。公司最初使用的是多臂路由，后来为了节

省成本使用单臂路由。

再后来，因为网络架构的变化，流量更多是在VLAN间传输，所以采用了多层交换。最后，因为为了方便网络管理采用VLAN聚合的技术。

学习任务

步骤一. 基础配置与 IP 编址

给所有设备配置IP地址和掩码。

```
<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
```

使用**ping**命令测试R1与R4接口G0/0/1的地址的连通性。

```
[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
```

使用**ping**命令测试R2和R4直连接口的连通性。

```
[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
```

步骤二. 多臂路由

R1和R2分别处于不同的VLAN中。

R1的网关使用R4的G0/0/1接口地址，R2的网关使用R4的E2/0/0接口地址。

由R4的多个接口提供VLAN间通信的服务就叫做多臂路由。

在交换机S1和S2上创建VLAN2和VLAN3。

```
[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.
```

将R1加入VLAN2，R2加入VLAN3，R4的G0/0/1加入VLAN2，E2/0/0加入VLAN3。

```
[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]quit
[S2]interface GigabitEthernet 0/0/4
[S2-GigabitEthernet0/0/4]port link-type access
[S2-GigabitEthernet0/0/4]port default vlan 3
[S2-GigabitEthernet0/0/4]quit
```

在R1和R2上配置网关，分别使用所在VLAN的R4接口地址。

```
[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
```

使用命令**display vlan**查看并确认配置。

```
[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)
               TG:GE0/0/9(U)        GE0/0/10(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)
                TG:GE0/0/9(U)        GE0/0/10(U)

VID   Status   Property         MAC-LRN Statistics Description
--------------------------------------------------------------------------
3     enable   default          enable   disable      VLAN 0003
```

测试R1和R2之间的连通性。

```
[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
```

步骤三.　单臂路由

在R4的一个物理接口上创建二个子接口，VLAN间的通讯通过对应的子接口完成。

这种方法叫做单臂路由。

关闭S2的G0/0/4接口。

```
[S2]interface GigabitEthernet 0/0/4
[S2-GigabitEthernet0/0/4]shutdown
```

```
[S2-GigabitEthernet0/0/4]quit
```

将S1和S2的G0/0/9接口加入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
```

将S1的G0/0/4接口模式改为Trunk模式，并允许VLAN2和VLAN3通过。

```
[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
```

在R4上为G0/0/1接口创建两个子接口。同时子接口上配置IP地址，并封装相应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
```

使用**display ip interface brief**命令查看R4上子接口配置信息。

```
[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
```

更改R1和R2的IP地址和网关。

```
[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
```

测试R1和R2之间的连通性。

```
[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
```

从**ping**命令的输出结果可知，VLAN 2的计算机和VLAN 3的计算机成功通信。

这种方案相对于多臂路由方案可以节省企业购买路由器接口的资金。

但相对的，单臂路由由于所有数据都在同一个接口上传输，随着VLAN数量的增加将会增大这条链路的带宽压力。同时这条链路也成为了企业网络的单点故障，一旦出现问题则整个网络都无法通信。

步骤四. 三层交换

三层交换指的是不需要路由器帮助，每个VLAN都有一个Vlanif接口充当路由器接口的角色来实现不同VLAN间通信的方法。

关闭S1的G0/0/4接口。

```
[S1]interface GigabitEthernet 0/0/4
[S1-GigabitEthernet0/0/4]shutdown
[S1-GigabitEthernet0/0/4]quit
```

更改S1的G0/0/9接口和S2的G0/0/9接口的模式为Trunk模式，允许VLAN2和VLAN3通过。

```
[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
```

在S1上创建Vlanif 2接口和Vlanif 3接口，并配置IP地址。

```
[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
```

测试R1和R2之间的连通性。

```
[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
```

从**ping**命令的输出结果可知，VLAN2 和VLAN 3的计算机通过交换机S1的二个Vlanif接口实现了三层数据通信。

相比单臂路由方案，三层交换更具有扩展性，即使VLAN增加也不会对其性能造成很大影响。

在VLAN间通信占企业大部分流量的网络中能够很好的承担服务压力。

步骤五.　VLAN 聚合

VLAN聚合和三层交换类似，都可以实现交换机上不同VLAN之间的通信。相比三层交换的方案它能将所有VLAN都放置在同一个网段中，达到减少IP网段使用和统一网关配置的效果。

在S1和S2上创建VLAN 10、20、100。

```
[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.
```

配置S1和S2的G0/0/9接口允许VLAN10、20通过。

```
[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
```

将S1的G0/0/1接口和S2的G0/0/2接口分别加入VLAN 10和VLAN 20。

```
[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
```

将VLAN100配置为Super-VLAN，并将VLAN 10和VLAN 20作为Sub-VLAN加入VLAN 100。

```
[S1]vlan 100
[S1-vlan100]aggregate-vlan
[S1-vlan100]access-vlan 10 20
[S1-Vlan100]quit
```

配置VLAN 100的Vlanif接口，启用ARP Proxy功能。

```
[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
```

更改R1和R2的IP地址，使其与Vlanif 100接口在同一个网段。并且将网关配置为Vlanif 100的接口地址。

```
[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
```

测试R1、R2和S1的Vlanif100接口之间的连通性。

```
[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
```

从**ping**命令的输出结果可知 ,R1、R2和S1的Vlanif 100接口可以互相通信。相比三层交换方案，VLAN聚合方案能够实现不同VLAN都使用相同的网关通信的功能，有效的减少IP地址的浪费和提高管理效率。但相对的，相同网段的计算机之间互访都依靠同一个Vlanif接口，也使这个接口的压力增大。

附加实验：思考并验证

多臂路由、单臂路由、VLAN间通信和VLAN聚合这四种方案各自的特点、优缺点和适用场景是怎么样的？

最终设备配置

[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
```

第七章 STP配置

实验 7-1 STP、RSTP 与 MSTP 的配置

学习目的

- 了解STP、RSTP与MSTP的差异
- 掌握修改网桥优先级影响根网桥选举的方法
- 掌握修改端口优先级影响根端口与指定端口选举的方法
- 掌握配置RSTP的方法以及STP与RSTP的相互兼容问题
- 掌握配置MSTP实现不同VLAN负载均衡的方法

拓扑图

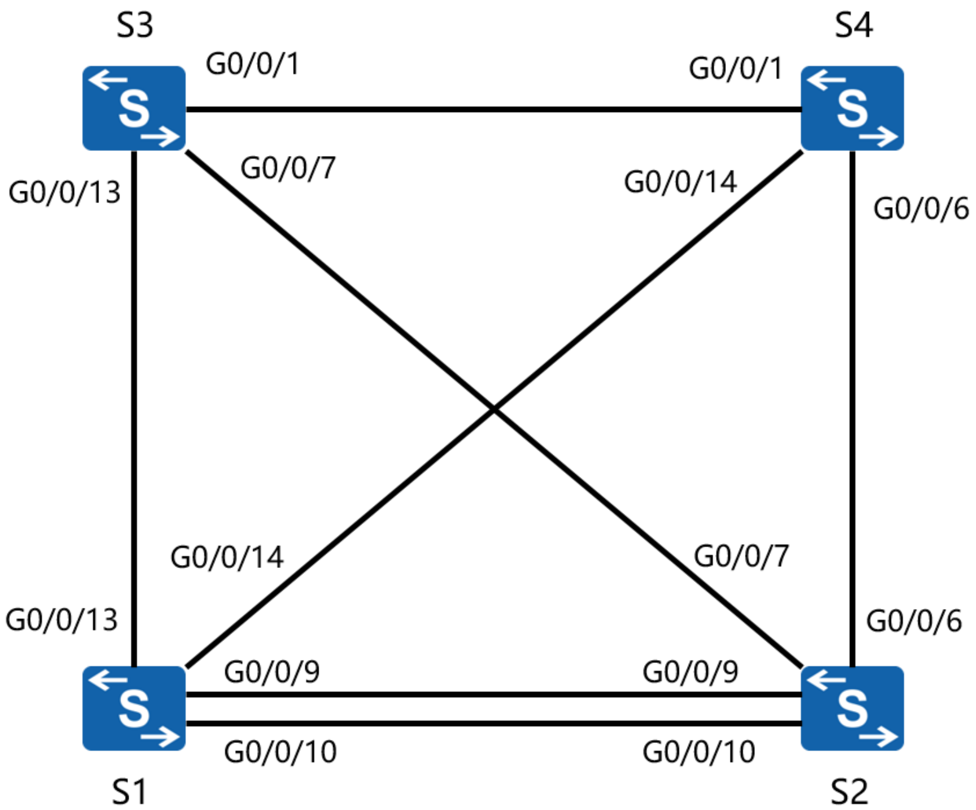

图7-1 STP、RSTP与MSTP实验拓扑图

场景

你是公司的网络管理员。公司的网络采用了备份网络，为避免环路问题，使用STP来进行环路控制。接口在STP的收敛时，所需时间较长，为了加快网络的收敛速度，可以配置RSTP来实现。所有的VLAN共享一棵STP生成树，为了实现VLAN间数据流量的负载均衡，可以配置MSTP来实现。

学习任务

步骤一. STP 配置及验证

如果设备默认生成树没有开启，使用以下命令开启。

```
[S1]stp enable

[S2]stp enable

[S3]stp enable

[S4]stp enable
```

配置使用传统生成树。

```
[S1]stp mode stp

[S2]stp mode stp

[S3]stp mode stp

[S4]stp mode stp
```

查看STP状态信息。

```
[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    GigabitEthernet0/0/6        DESI  FORWARDING      NONE
   0    GigabitEthernet0/0/7        DESI  FORWARDING      NONE
   0    GigabitEthernet0/0/9        DESI  FORWARDING      NONE
   0    GigabitEthernet0/0/10       DESI  FORWARDING      NONE
```

S2为根网桥，所有端口都为指定端口。

实际使用中，由于交换机MAC地址的不可确定性，实际的实验结果可能与如上结果有差异。

步骤二. 根桥选举控制

配置S1为主根网桥，S2为备份根网桥。

```
[S1]stp root primary

[S2]stp root secondary
```

查看STP配置信息。

```
[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为主根网桥，S2为备份根网桥。

桥优先级数值越小的优先级越高，将S1的桥优先级修改为8192，将S2的桥优先级修改为4096。

```
[S1]undo stp root
[S1]stp priority 8192

[S2]undo stp root
[S2]stp priority 4096
```

查看STP信息。

```
[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……
```

S1优先级为8192，S2优先级为4096，S2成为根网桥。

步骤三. 根端口选举控制

在S1上查看当前端口角色信息。

```
[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
```

S1的GigabitEthernet0/0/9为根端口。

端口优先级默认值为128，数值越大优先级越小。

S1与S2通过G0/0/9与G0/0/10接口互联。

将S2的端口G0/0/9端口优先级设置为32，G0/0/10端口优先级设置为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
```

提示：此处是修改S2的端口优先级，而不是修改S1的端口优先级。

在S1上查看当前端口角色信息。

```
[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
```

S1的GigabitEthernet0/0/10成为根端口。

步骤四. 指定端口选举控制

查看S3和S4直连接口状态。

```
 [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的GigabitEthernet 0/0/1为替代端口。S4的GigabitEthernet 0/0/1是指定端口。修改S4端口GE0/0/6路径开销为2000000。

```
[S4]interface GigabitEthernet 0/0/6
[S4-GigabitEthernet0/0/6]stp cost 2000000
[S4-GigabitEthernet0/0/6]quit
```

查看当前端口角色信息。

```
[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][Port24(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的GigabitEthernet 0/0/1是指定端口。S4的GigabitEthernet 0/0/1为替代端口。

步骤五. RSTP 配置及验证

配置S1和S2的VLANIF 1接口地址。测试S1到S2的连通性。

```
[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
```

查看S1端口角色信息。

```
[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的GigabitEthernet0/0/10为根端口，用ping测试S1到S2的连通性20次。

提示：S1执行ping操作之后立刻关闭S2的GigabitEthernet 0/0/10接口。

```
[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
```

查看S1端口角色信息。

```
[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的GigabitEthernet0/0/9接口成为根端口，端口进入FORWARDING状态，15个包超时，网络收敛时间为30秒。

恢复S2的GigabitEthernet 0/0/10接口。

```
[S2]interface GigabitEthernet 0/0/10
[S2-GigabitEthernet0/0/10]undo shutdown
[S2-GigabitEthernet0/0/10]quit
```

配置快速生成树。

```
[S1]stp mode rstp

[S2]stp mode rstp

[S3]stp mode rstp

[S4]stp mode rstp
```

查看S1端口角色信息。

```
[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的GigabitEthernet0/0/10为根端口，用**ping**测试S1到S2的连通性20次。

提示：S1上执行ping之后立刻关闭S2的GigabitEthernet 0/0/10。

```
[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
```

查看S1端口角色信息。

```
[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的GigabitEthernet0/0/9接口成为根端口，进入Forwarding状态。1个包超时，网络收敛时间为2秒。

开启S2的GigabitEthernet 0/0/10接口。

```
[S2]interface GigabitEthernet 0/0/10
[S2-GigabitEthernet0/0/10]undo shutdown
[S2-GigabitEthernet0/0/10]quit
```

步骤六. RSTP 与 STP 的兼容

配置S1为STP生成树，其他配置保持不变。

```
[S1]stp mode stp
```

查看S1端口角色信息。

```
[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的GigabitEthernet0/0/10为根端口，用ping测试S1到S2的连通性20次。

提示：S1上执行ping之后立刻关闭S2的GigabitEthernet 0/0/10。

```
 [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
```

查看S1端口角色信息。

```
[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的GigabitEthernet0/0/9接口成为根端口，进入Forwarding状态。15个包超时，网络收敛时间为30秒。

RSTP兼容STP，但收敛方式以STP模式运行。

恢复S2的GigabitEthernet 0/0/10接口。

```
[S2]interface GigabitEthernet 0/0/10
[S2-GigabitEthernet0/0/10]undo shutdown
[S2-GigabitEthernet0/0/10]quit
```

步骤七. MSTP 的配置与验证

创建VLAN 2到20，并将相应的接口加入到VLAN中。

```
[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 0/0/7
[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
[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

[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 GigabitEthernet0/0/6
[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
[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
```

配置MSTP。

定义VLAN1-10属于INSTANCE 1，VLAN11-20属于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
```

查看MSTP实例和VLAN的映射关系。

```
[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
```

配置S1在实例1中的优先级为4096，在实例2中的优先级为8192。

配置S2在实例2中的优先级为4096，在实例1中的优先级为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
```

查看实例1和实例2的状态信息。

```
[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为实例1的根桥，S2为实例2的根桥。

查看MSTP实例1端口角色。

```
[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     GigabitEthernet0/0/6          DESI   FORWARDING        NONE
   1     GigabitEthernet0/0/7          DESI   FORWARDING        NONE
   1     GigabitEthernet0/0/9          ROOT   FORWARDING        NONE
   1     GigabitEthernet0/0/10         ALTE   DISCARDING        NONE

[S3]display stp instance 1 brief
 MSTID   Port                          Role   STP State       Protection
   1     GigabitEthernet0/0/1            ALTE   DISCARDING        NONE
   1     GigabitEthernet0/0/13           ROOT FORWARDING        NONE
   1     GigabitEthernet0/0/7              ALTE   DISCARDING        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
```

实例1中S1为根网桥，S3的VLAN 1到VLAN10的用户经过Ethernet0/0/13接口和S1、S2、S4的VLAN 1到VLAN10的用户通讯。

查看MSTP实例2端口角色。

```
[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     GigabitEthernet0/0/6       DESI   FORWARDING        NONE
   2     GigabitEthernet0/0/7        DESI   FORWARDING        NONE
   2     GigabitEthernet0/0/9        DESI   FORWARDING        NONE
   2     GigabitEthernet0/0/10       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     GigabitEthernet0/0/6          ROOT FORWARDING         NONE
   2     GigabitEthernet0/0/14         DESI   FORWARDING       NONE
```

实例2中S2为根网桥，S3的VLAN 11到VLAN 20经过Ethernet0/0/23和S1、S2、S4的VLAN 11到VLAN 20通讯。

附加实验：思考并验证

MSTP如何在多个区域实现不同VLAN数据传输的均衡？

RSTP快速转发的原因是什么？

最终设备配置

```
[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 V100R006C05
 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 GigabitEthernet0/0/7
 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

[S4]display current-configuration
#
!Software Version V100R006C05
 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 GigabitEthernet0/0/6
 port link-type trunk
 port trunk allow-pass vlan 2 to 20
 stp instance 0 cost 2000000
#
interface GigabitEthernet0/0/7
#
interface GigabitEthernet0/0/14
 port link-type trunk
 port trunk allow-pass vlan 2 to 20
#

Return
```

实验 7-2 MSTP 多区域与 STP 的兼容（选做）

学习目的

- 掌握MSTP多实例以及多区域的配置方法
- 掌握实现MSTP与相STP兼容的配置方法
- 掌握MSTP边缘端口保护、指定端口保护、环路保护、TC-BPDU保护的配置方法

拓扑图

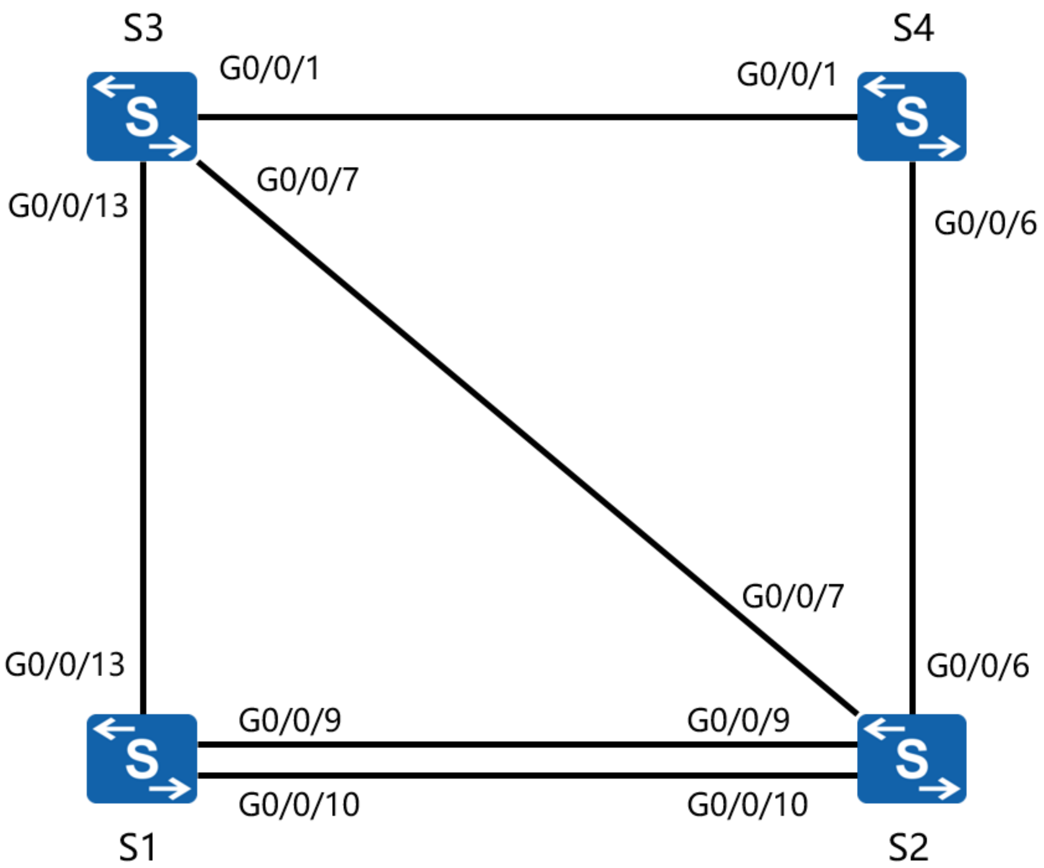

图7-2 MSTP多区域与STP的兼容

场景

你是公司的网络管理员。在公司的二层网络架构中，为了克服单个生成树的会造成部分VLAN路径不通及次优路径等种种弊端，并且实现流量的分担。在网络中部署MSTP协议，并能实现与传统单生成树之间的相互兼容。

学习任务

步骤一. 基本配置

在本实验之前，需要关闭一些不使用的接口。

```
<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 0/0/7
[S3-GigabitEthernet0/0/7]shutdown
[S3-GigabitEthernet0/0/7]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
```

所有交换机上都创建VLAN 3,4,5,6,7,8，

```
[S1]vlan batch 3 to 8

[S2]vlan batch 3 to 8

[S3]vlan batch 3 to 8

[S4]vlan batch 3 to 8
```

查看创建的VLAN，

```
[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
```

设置所有交换机之间链路类型全部为Trunk链路，接收BPDU，并且允许所有VLAN通过。注意S2与S3之间的直连链路此时不做任何配置。

```
[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
```

步骤二. 配置 MSTP 多实例

在系统视图下，启用MSTP，

```
[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
```

配置所有交换机都属于同一个区域RG1，修订级别为1。建实例1和VLAN 3,4,5建立映射关系，创建实例2和VLAN 6,7,8建立映射关系，并激活域配置。

```
[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
```

完成后查看MSTP信息，

[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

此时S1为根交换机，S4上的E0/0/24端口为所有MSTP进程的备用端口。

在实例2中，设置S2的优先级为0，S1的优先级为4096，S4的优先级为8192，使S2成为实例2中的根交换机。

[S2]stp instance 2 priority 0

[S1]stp instance 2 priority 4096

[S4]stp instance 2 priority 8192

配置完成后，查看MSTP的基本信息。

[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
```

此时S2成为了实例2中的根交换机，S3的E0/0/1成为实例2中的预备端口。但各交换机在实例1中的状态没有改变，验证了每个MST实例都独立计算生成树，互不影响。

步骤三.　配置 MSTP 多区域

删除步骤二中所有交换机上的MSTP区域以及优先级的配置。

```
[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
```

配置S1和S3到一个MSTP域内，域名为RG1，修订版本号为1。

创建实例1，与VLAN 3,4,5，建立映射关系。

创建实例2，与VLAN 6,7,8,建立映射关系。

```
[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
```

将S2和S4配置到另一个MSTP域内，域名为RG2，修订版本号为2。

创建实例1，和VLAN 3,4,5建立映射关系。

创建实例2，与VLAN 6,7,8建立映射关系，并全部激活域配置，

```
[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

配置完成后，查看MSTP基本信息。

```
[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
   2    GigabitEthernet0/0/6     ROOT  FORWARDING      NONE
```

此时S1为根交换机，S4上的E0/0/1为预备端口。

配置S3在实例0中优先级为0，使其成为CIST的总根，配置S3在实例1中的优先级为0，使其成为实例1的域根，配置S4在实例1中的优先级为0，使其成为实

例1的域根。

```
[S3]stp instance 0 priority 0
[S3]stp instance 1 priority 0

[S4]stp instance 1 priority 0
```

配置完成后，查看MSTP基本信息。

```
[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
```

删除S2和S4上的MSTP配置，将S2和S4配置到另一个MSTP域内，域名为RG2，修订版本号为2。创建实例1，和VLAN 6,7,8建立映射关系，创建实例2，与VLAN 3,4,5建立映射关系，并全部激活域配置，

```
[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
```

配置完成后，查看MSTP基本信息。

[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    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
```

除了Instance 0之外，每个区域的MST实例都独立计算生成树，不管是否包含相同的VLAN，不管VLAN是否和实例映射一致，区域内的生成树计算相互之间互不影响。

步骤四. 配置 MSTP 与 STP 兼容

将S1，S2，S3配置在一个MSTP域中，S4单独配置STP。

删除S2上的MSTP配置，在S2上重新创建MSTP。域名为RG1，创建实例1与VLAN 3,4,5，建立映射关系。创建实例2与VLAN6,7,8建立映射关系，并激活域配置。

```
[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
```

开启S2的S0/0/23和S3的E0/0/23端口。

设置S2 ,S3之间的直连链路类为trunk链路 ,接收BPDU ,并且允许所有VLAN通过。

```
[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   GigabitEthernet0/0/7
[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
```

删除S4上的MSTP配置，并在S4上使能STP，

```
[S4]undo stp region-configuration
[S4]stp mode stp
```

配置完成，查看STP基本信息。

[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
```

运行STP的S4和运行MSTP的S1、S2、S3中的实例0共同计算CIST。此时S1为CIST的总根。

配置S4优先级为4096，使其成为CIST的总根，

```
[S4]stp priority 4096
```

查看STP基本信息。

```
[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
```

此时S4成为了CIST的总根，S4上的端口都是指定端口。

步骤五. 配置指定端口保护

在S4的接口E0/0/1和E0/0/6上配置指定端口保护，

```
[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
```

查看此时S4的STP基本信息，

```
[S4]display stp brief
MSTID  Port                        Role  STP State     Protection
```

```
 0     GigabitEthernet0/0/1            DESI   FORWARDING       ROOT
 0     GigabitEthernet0/0/6            DESI   DISCARDING       ROOT
```

配置S2在实例0中的优先级为0，模拟抢占CIST根交换机。

```
[S2]stp instance 0 priority 0
```

查看S2，S4上的STP信息。

```
[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
```

此时S4上的端口状态进入Discarding状态，不转发报文。说明此时S4上的端口角色并没有变化，仍然是根交换机。

删除S2上的实例0优先级配置。

```
[S2]undo stp instance 0 priority
```

查看S2，S4上的STP信息。

```
[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
```

当在足够长的时间内（Max Age，默认20秒）没有收到更优的配置消息时，端口会恢复原来的正常状态，进入转发状态。

步骤六. 配置边缘端口保护

开启S2的G0/0/9接口，

```
[S2]interface  GigabitEthernet 0/0/9
[S2-GigabitEthernet0/0/9]undo shutdown
[S2-GigabitEthernet0/0/9]quit
```

将S1的接口G0/0/9配置成边缘端口，并在全局模式下开启边缘端口的保护功能。

```
[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
```

查看S1上的STP信息。

```
[S1]display stp interface GigabitEthernet 0/0/9 brief
 MSTID  Port                        Role  STP State       Protection
  0     GigabitEthernet0/0/9        DESI  FORWARDING         BPDU
```

开启S1的G0/0/9接口，使边缘端口端口能收到BPDU，模拟交换机受到攻击。

```
[S1]interface  GigabitEthernet 0/0/9
[S1-GigabitEthernet0/0/9]undo shutdown
[S1-GigabitEthernet0/0/9]quit
```

观察S1上的现象。

```
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.
```

配置边缘端口保护之后，一旦启用了边缘端口，使其收到BPDU，会自动关闭该端口。

步骤七. 配置环路保护

在S3的接口E0/0/23上配置环路保护。

```
[S3]interface GigabitEthernet0/0/7
[S3-GigabitEthernet0/0/7]stp loop-protection
[S3-GigabitEthernet0/0/7]quit
```

查看S3上该接口的STP信息。

```
[S3]display stp interface Ethernet 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
```

步骤八. 配置 TC-BPDU 保护

在S1上开启对TC类型BPDU报文的保护功能。

```
[S1]stp tc-protection
```

附加实验：思考并验证

思考不同交换机上的MSTP域名一致，修订版本号不同会不会产生影响。

思考在步骤四中，如果改变S3中实例1上的优先级为0，则此时四台交换机端口的状态会产生哪些变化。

最终设备配置

```
<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 GigabitEthernet0/0/6
 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 GigabitEthernet0/0/9
#
interface GigabitEthernet0/0/10
 port link-type trunk
 port trunk allow-pass vlan 2 to 4094
#


return
```

<S3>**display current-configuration**

```
#
!Software Version V100R006C05
 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 V100R006C05
 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
```

```
#
interface GigabitEthernet0/0/14
 shutdown
 undo ntdp enable
 undo ndp enable
 bpdu disable
#

Return
```