



Recommendations

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



More Information

- Huawei learning APP



Huawei Certification

HCIP-Routing & Switching-IENP

**Improving Enterprise Network Performance
Lab Guide**



Huawei Technologies Co., Ltd

Copyright © Huawei Technologies Co., Ltd. 2017. All rights reserved.

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

Trademarks and Permissions



and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd. All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

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



Huawei Certification

Improving Enterprise Network Performance

Lab Guide

Edition 2.5

Huawei Certification System

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

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

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

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

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

About This Document

Overview

This document is HCIP-Improving Enterprise Network Performance (HCIP-IENP) certification training material. It is intended for those who are preparing for the HCIP-IENP exam and those who want to improve enterprise network performance and Huawei Versatile Routing Platform (VRP) implementation.

Chapter 1 introduces principles and configurations of MPLS and MPLS VPN, and helps readers master methods to improve enterprise network service bearer capabilities.

Chapters 2, 3, 4, 5, and 6 illustrate principles and configurations of DHCP, QoS, network security basics, VRRP, and BFD, and help readers to comprehensively master corresponding theories and practices to improve enterprise network service quality and enhance enterprise network security and reliability.

This course helps readers gradually understand routing technologies and how these technologies are implemented on Huawei products.

Background Knowledge Required

To fully understand this document, readers should:

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

Icons



Router



Switch



Firewall



Cloud



Ethernet link



Serial link

CONTENTS

CHAPTER 1 MPLS VPN CONFIGURATION.....	1
LAB 1-1 MPLS LDP CONFIGURATION.....	1
LAB 1-2 MPLS VPN CONFIGURATION.....	22
CHAPTER 2 DHCP CONFIGURATION	46
LAB 2-1 DHCP CONFIGURATION.....	46
CHAPTER 3 SERVICE QUALITY AND TRAFFIC CONTROL.....	74
LAB 3-1 QOS BASICS.....	74
CHAPTER 4 FIREWALL CONFIGURATION.....	103
LAB 4-1 FIREWALL ZONE AND SECURITY POLICY CONFIGURATION.....	103
LAB 4-2 FIREWALL NAT CONFIGURATION.....	124
CHAPTER 5 VRRP CONFIGURATION.....	147
LAB 5-1 VRRP CONFIGURATION.....	147
CHAPTER 6 BFD CONFIGURATION	179
LAB 6-1 ASSOCIATION BETWEEN BFD AND STATIC ROUTES.....	179
LAB 6-2 ASSOCIATION BETWEEN BFD AND OSPF.....	196
LAB 6-3 ASSOCIATION BETWEEN BFD AND VRRP.....	208

Chapter 1 MPLS VPN Configuration

Lab 1-1 MPLS LDP Configuration

Learning Objectives

The objectives of this lab are to learn and understand:

- How to enable and disable MPLS
- How to enable and disable MPLS LDP
- How to configure LSPs using MPLS LDP
- How to configure the LDP LSP trigger policy on an MPLS router

Topology




Figure 1-1 MPLS LDP topology

Scenario

Assume that you are a network administrator of an enterprise. Your enterprise uses an IP network with poor forwarding performance. You need to use MPLS to improve the forwarding rate of routers. Static LSPs are configured manually, while LDP is a protocol developed for label distribution. To perform flexible configuration, use LDP to set up MPLS LSPs.

Tasks

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

Configure IP addresses and masks for all routers.

```
<Huawei>system-view
```

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname S1
```

```
[S1]interface Vlanif 1
```

```
[S1-Vlanif1]ip address 10.0.1.2 24
```

```
<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 10.0.1.1 24
```

```
[R1-GigabitEthernet0/0/1]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 2.2.2.2 24
```

```
<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 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 3.3.3.3 24

<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 10.0.2.1 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 loopback 0

[R3-LoopBack0]ip address 4.4.4.4 24

<Huawei>system-view

Enter system view, return user view with Ctrl+Z.

[Huawei]sysname S2

[S2]interface Vlanif 1

[S2-Vlanif1]ip address 10.0.2.2 24

After the configurations are complete, test the connectivity of direct links.

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

Add 10.0.12.0/24, 10.0.23.0/24, 10.0.1.0/24, and 10.0.2.0/24 to OSPF area 0.

[S1]ospf 1 router-id 1.1.1.1

[S1-ospf-1]area 0

[S1-ospf-1-area-0.0.0.0]network 10.0.1.0 0.0.0.255

```
[R1]ospf 1 router-id 2.2.2.2
```

```
[R1-ospf-1]area 0
```

```
[R1-ospf-1-area-0.0.0.0]network 10.0.1.0 0.0.0.255
```

```
[R1-ospf-1-area-0.0.0.0]network 10.0.12.0 0.0.0.255
```

```
[R1-ospf-1-area-0.0.0.0]network 2.2.2.0 0.0.0.255
```

```
[R2]ospf 1 router-id 3.3.3.3
```

```
[R2-ospf-1]area 0
```

```
[R2-ospf-1-area-0.0.0.0]network 10.0.12.0 0.0.0.255
```

```
[R2-ospf-1-area-0.0.0.0]network 10.0.23.0 0.0.0.255
```

```
[R2-ospf-1-area-0.0.0.0]network 3.3.3.0 0.0.0.255
```

```
[R3]ospf 1 router-id 4.4.4.4
```

```
[R3-ospf-1]area 0
```

```
[R3-ospf-1-area-0.0.0.0]network 10.0.23.0 0.0.0.255
```

```
[R3-ospf-1-area-0.0.0.0]network 10.0.2.0 0.0.0.255
```

```
[R3-ospf-1-area-0.0.0.0]network 4.4.4.0 0.0.0.255
```

```
[S2]ospf 1 router-id 5.5.5.5
```

```
[S2-ospf-1]area 0
```

```
[S2-ospf-1-area-0.0.0.0]network 10.0.2.0 0.0.0.255
```

Check the routing table and test connectivity on the entire network.

```
[R2]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=253 time=36 ms
```

```
    Reply from 10.0.1.2: bytes=56 Sequence=2 ttl=253 time=31 ms
```

Reply from 10.0.1.2: bytes=56 Sequence=3 ttl=253 time=31 ms

Reply from 10.0.1.2: bytes=56 Sequence=4 ttl=253 time=31 ms

Reply from 10.0.1.2: bytes=56 Sequence=5 ttl=253 time=31 ms

--- 10.0.1.2 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

round-trip min/avg/max = 31/32/36 ms

[R2]ping 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=253 time=38 ms

Reply from 10.0.2.2: bytes=56 Sequence=2 ttl=253 time=33 ms

Reply from 10.0.2.2: bytes=56 Sequence=3 ttl=253 time=33 ms

Reply from 10.0.2.2: bytes=56 Sequence=4 ttl=253 time=33 ms

Reply from 10.0.2.2: bytes=56 Sequence=5 ttl=253 time=33 ms

--- 10.0.2.2 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

round-trip min/avg/max = 33/34/38 ms

Run the **display ip routing-table** command to check the OSPF routing table.

[R2]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
2.2.2.2/32	OSPF	10	1562	D	10.0.12.1	Serial1/0/0
3.3.3.0/24	Direct	0	0	D	3.3.3.3	LoopBack0
3.3.3.3/32	Direct	0	0	D	127.0.0.1	InLoopBack0
3.3.3.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
4.4.4.4/32	OSPF	10	1562	D	10.0.23.3	Serial2/0/0
10.0.1.0/24	OSPF	10	1563	D	10.0.12.1	Serial1/0/0
10.0.2.0/24	OSPF	10	1563	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	InLoopBack0
10.0.12.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.23.0/24	Direct	0	0	D	10.0.23.2	Serial2/0/0
10.0.23.2/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.23.3/32	Direct	0	0	D	10.0.23.3	Serial2/0/0
10.0.23.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.0/8	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
127.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
255.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0

Step 3 Configure MPLS LDP.

Configure MPLS and LDP globally on MPLS routers.

[R1]mpls lsr-id 2.2.2.2

[R1]mpls

Info: Mpls starting, please wait... OK!

[R1-mpls]mpls ldp

[R2]mpls lsr-id 3.3.3.3

[R2]mpls

Info: Mpls starting, please wait... OK!

[R2-mpls]mpls ldp

[R3]mpls lsr-id 4.4.4.4

[R3]mpls

Info: Mpls starting, please wait... OK!

[R3-mpls]mpls ldp

Configure MPLS and LDP on interfaces of MPLS routers.

[R1]interface Serial 1/0/0

[R1-Serial1/0/0]mpls

[R1-Serial1/0/0]mpls ldp

[R2]interface Serial 1/0/0

[R2-Serial1/0/0]mpls

[R2-Serial1/0/0]mpls ldp

[R2-Serial1/0/0]quit

[R2]interface Serial 2/0/0

[R2-Serial2/0/0]mpls

[R2-Serial2/0/0]mpls ldp

[R3]interface Serial 2/0/0

[R3-Serial2/0/0]mpls

[R3-Serial2/0/0]mpls ldp

After the configurations are complete, run the **display mpls ldp session** command on Routers. You can see that the status of local LDP sessions between R1 and R2 and between R1 and R3 are **Operational**.

[R1]display mpls ldp session

LDP Session(s) in Public Network

Codes: LAM(Label Advertisement Mode), SsnAge Unit(DDDD:HH:MM)

A '*' before a session means the session is being deleted.

```

-----
PeerID           Status      LAM  SsnRole  SsnAge      KASent/Rcv
-----
3.3.3.3:0        Operational DU   Passive 0000:00:10 41/41
-----
TOTAL: 1 session(s) Found.
  
```

[R2]display mpls ldp session

LDP Session(s) in Public Network

Codes: LAM(Label Advertisement Mode), SsnAge Unit(DDDD:HH:MM)

A '*' before a session means the session is being deleted.

```

-----
PeerID           Status      LAM  SsnRole  SsnAge      KASent/Rcv
-----
2.2.2.2:0        Operational DU   Active  0000:00:11 46/46
4.4.4.4:0        Operational DU   Passive 0000:00:10 43/43
-----
  
```

TOTAL: 2 session(s) Found.

[R3]display mpls ldp session

LDP Session(s) in Public Network

Codes: LAM(Label Advertisement Mode), SsnAge Unit(DDDD:HH:MM)

A '*' before a session means the session is being deleted.

```

-----
PeerID           Status      LAM  SsnRole  SsnAge      KASent/Rcv
-----
3.3.3.3:0        Operational DU   Active  0000:00:11  46/46
-----

```

TOTAL: 1 session(s) Found.

Step 4 Establish LDP LSPs.

All LSRs are triggered to establish LDP LSPs based on the host route, which is the default trigger policy.

Run the **display mpls ldp lsp** command on LSRs. All host routes are triggered to establish LDP LSPs.

[R1]display mpls ldp lsp

LDP LSP Information

```

-----
DestAddress/Mask  In/OutLabel  UpstreamPeer  NextHop      OutInterface
-----
2.2.2.2/32        3/NULL       3.3.3.3       127.0.0.1    InLoop0
*2.2.2.2/32        Liberal/1024  DS/3.3.3.3
3.3.3.3/32        NULL/3        -             10.0.12.2    S1/0/0
3.3.3.3/32        1024/3        3.3.3.3       10.0.12.2    S1/0/0
4.4.4.4/32        NULL/1025    -             10.0.12.2    S1/0/0
4.4.4.4/32        1025/1025    3.3.3.3       10.0.12.2    S1/0/0
-----

```

TOTAL: 5 Normal LSP(s) Found.

TOTAL: 1 Liberal LSP(s) Found.

TOTAL: 0 Frr LSP(s) Found.

A '*' before an LSP means the LSP is not established

A '*' before a Label means the USCB or DSCB is stale

A '*' before a UpstreamPeer means the session is in GR state

A '*' before a DS means the session is in GR state

A '*' before a NextHop means the LSP is FRR LSP

[R2]display mpls ldp lsp

LDP LSP Information

DestAddress/Mask	In/OutLabel	UpstreamPeer	NextHop	OutInterface
2.2.2.2/32	NULL/3	-	10.0.12.1	S1/0/0
2.2.2.2/32	1024/3	2.2.2.2	10.0.12.1	S1/0/0
2.2.2.2/32	1024/3	4.4.4.4	10.0.12.1	S1/0/0
*2.2.2.2/32	Liberal/1024		DS/4.4.4.4	
3.3.3.3/32	3/NULL	2.2.2.2	127.0.0.1	InLoop0
3.3.3.3/32	3/NULL	4.4.4.4	127.0.0.1	InLoop0
*3.3.3.3/32	Liberal/1024		DS/2.2.2.2	
*3.3.3.3/32	Liberal/1025		DS/4.4.4.4	
4.4.4.4/32	NULL/3	-	10.0.23.3	S2/0/0
4.4.4.4/32	1025/3	2.2.2.2	10.0.23.3	S2/0/0
4.4.4.4/32	1025/3	4.4.4.4	10.0.23.3	S2/0/0
*4.4.4.4/32	Liberal/1025		DS/2.2.2.2	

TOTAL: 8 Normal LSP(s) Found.

TOTAL: 4 Liberal LSP(s) Found.

TOTAL: 0 Frr LSP(s) Found.

A '*' before an LSP means the LSP is not established

A '*' before a Label means the USCB or DSCB is stale

A '*' before a UpstreamPeer means the session is in GR state

A '*' before a DS means the session is in GR state

A '*' before a NextHop means the LSP is FRR LSP

[R3]display mpls ldp lsp

LDP LSP Information

DestAddress/Mask	In/OutLabel	UpstreamPeer	NextHop	OutInterface
2.2.2.2/32	NULL/1024	-	10.0.23.2	S2/0/0
2.2.2.2/32	1024/1024	3.3.3.3	10.0.23.2	S2/0/0
3.3.3.3/32	NULL/3	-	10.0.23.2	S2/0/0
3.3.3.3/32	1025/3	3.3.3.3	10.0.23.2	S2/0/0
4.4.4.4/32	3/NULL	3.3.3.3	127.0.0.1	InLoop0
*4.4.4.4/32	Liberal/1025		DS/3.3.3.3	

TOTAL: 5 Normal LSP(s) Found.

TOTAL: 1 Liberal LSP(s) Found.

TOTAL: 0 Frr LSP(s) Found.

A '*' before an LSP means the LSP is not established

A '*' before a Label means the USCB or DSCB is stale

A '*' before a UpstreamPeer means the session is in GR state

A '*' before a DS means the session is in GR state

A '*' before a NextHop means the LSP is FRR LSP

In most cases, the default trigger policy is used. The establishment of an LDP LSP is triggered in Host mode.

Change the trigger policy to All on LSRs so that all static routes and IGP entries can trigger the establishment of the LDP LSPs.

```
[R1]mpls
```

```
[R1-mpls]lsp-trigger all
```

```
[R2]mpls
```

```
[R2-mpls]lsp-trigger all
```

```
[R3]mpls
```

```
[R3-mpls]lsp-trigger all
```

After the configuration is complete, run the **display mpls ldp lsp** command on each node to view the established LDP LSPs.

```
[R1]display mpls ldp lsp
```

```
LDP LSP Information
```

DestAddress/Mask	In/OutLabel	UpstreamPeer	NextHop	OutInterface
2.2.2.0/24	3/NULL	3.3.3.3	2.2.2.2	Loop0
2.2.2.2/32	3/NULL	3.3.3.3	127.0.0.1	InLoop0
*2.2.2.2/32	Liberal/1024		DS/3.3.3.3	
*3.3.3.0/24	Liberal/3		DS/3.3.3.3	
3.3.3.3/32	NULL/3	-	10.0.12.2	S1/0/0
3.3.3.3/32	1024/3	3.3.3.3	10.0.12.2	S1/0/0
4.4.4.4/32	NULL/1025	-	10.0.12.2	S1/0/0
4.4.4.4/32	1025/1025	3.3.3.3	10.0.12.2	S1/0/0

10.0.1.0/24	3/NULL	3.3.3.3	10.0.1.1	GE0/0/1
*10.0.1.0/24	Liberal/1026		DS/3.3.3.3	
10.0.2.0/24	NULL/1027	-	10.0.12.2	S1/0/0
10.0.2.0/24	1027/1027	3.3.3.3	10.0.12.2	S1/0/0
10.0.12.0/24	3/NULL	3.3.3.3	10.0.12.1	S1/0/0
*10.0.12.0/24	Liberal/3		DS/3.3.3.3	
10.0.23.0/24	NULL/3	-	10.0.12.2	S1/0/0
10.0.23.0/24	1026/3	3.3.3.3	10.0.12.2	S1/0/0

TOTAL: 12 Normal LSP(s) Found.

TOTAL: 4 Liberal LSP(s) Found.

TOTAL: 0 Frr LSP(s) Found.

A '*' before an LSP means the LSP is not established

A '*' before a Label means the USCB or DSCB is stale

A '*' before a UpstreamPeer means the session is in GR state

A '*' before a DS means the session is in GR state

A '*' before a NextHop means the LSP is FRR LSP

[R2]display mpls ldp lsp

LDP LSP Information

DestAddress/Mask	In/OutLabel	UpstreamPeer	NextHop	OutInterface

*2.2.2.0/24	Liberal/3		DS/2.2.2.2	
2.2.2.2/32	NULL/3	-	10.0.12.1	S1/0/0
2.2.2.2/32	1024/3	2.2.2.2	10.0.12.1	S1/0/0
2.2.2.2/32	1024/3	4.4.4.4	10.0.12.1	S1/0/0
*2.2.2.2/32	Liberal/1024		DS/4.4.4.4	
3.3.3.0/24	3/NULL	2.2.2.2	3.3.3.3	Loop0

HCIP-Routing & Switching-IENP Lab Guide

3.3.3.0/24	3/NULL	4.4.4.4	3.3.3.3	Loop0
3.3.3.3/32	3/NULL	2.2.2.2	127.0.0.1	InLoop0
3.3.3.3/32	3/NULL	4.4.4.4	127.0.0.1	InLoop0
*3.3.3.3/32	Liberal/1024		DS/2.2.2.2	
*3.3.3.3/32	Liberal/1025		DS/4.4.4.4	
*4.4.4.0/24	Liberal/3		DS/4.4.4.4	
4.4.4.4/32	NULL/3	-	10.0.23.3	S2/0/0
4.4.4.4/32	1025/3	2.2.2.2	10.0.23.3	S2/0/0
4.4.4.4/32	1025/3	4.4.4.4	10.0.23.3	S2/0/0
*4.4.4.4/32	Liberal/1025		DS/2.2.2.2	
10.0.1.0/24	NULL/3	-	10.0.12.1	S1/0/0
10.0.1.0/24	1026/3	2.2.2.2	10.0.12.1	S1/0/0
10.0.1.0/24	1026/3	4.4.4.4	10.0.12.1	S1/0/0
*10.0.1.0/24	Liberal/1026		DS/4.4.4.4	
10.0.2.0/24	NULL/3	-	10.0.23.3	S2/0/0
10.0.2.0/24	1027/3	2.2.2.2	10.0.23.3	S2/0/0
10.0.2.0/24	1027/3	4.4.4.4	10.0.23.3	S2/0/0
*10.0.2.0/24	Liberal/1027		DS/2.2.2.2	
10.0.12.0/24	3/NULL	2.2.2.2	10.0.12.2	S1/0/0
10.0.12.0/24	3/NULL	4.4.4.4	10.0.12.2	S1/0/0
*10.0.12.0/24	Liberal/3		DS/2.2.2.2	
*10.0.12.0/24	Liberal/1027		DS/4.4.4.4	
10.0.23.0/24	3/NULL	2.2.2.2	10.0.23.2	S2/0/0
10.0.23.0/24	3/NULL	4.4.4.4	10.0.23.2	S2/0/0
*10.0.23.0/24	Liberal/1026		DS/2.2.2.2	
*10.0.23.0/24	Liberal/3		DS/4.4.4.4	

TOTAL: 20 Normal LSP(s) Found.

TOTAL: 12 Liberal LSP(s) Found.

TOTAL: 0 Frr LSP(s) Found.

A '*' before an LSP means the LSP is not established

A '*' before a Label means the USCB or DSCB is stale

A '*' before a UpstreamPeer means the session is in GR state

A '*' before a DS means the session is in GR state

A '*' before a NextHop means the LSP is FRR LSP

[R3]display mpls ldp lsp

LDP LSP Information

DestAddress/Mask	In/OutLabel	UpstreamPeer	NextHop	OutInterface
2.2.2.2/32	NULL/1024	-	10.0.23.2	S2/0/0
2.2.2.2/32	1024/1024	3.3.3.3	10.0.23.2	S2/0/0
*3.3.3.0/24	Liberal/3		DS/3.3.3.3	
3.3.3.3/32	NULL/3	-	10.0.23.2	S2/0/0
3.3.3.3/32	1025/3	3.3.3.3	10.0.23.2	S2/0/0
4.4.4.0/24	3/NULL	3.3.3.3	4.4.4.4	Loop0
4.4.4.4/32	3/NULL	3.3.3.3	127.0.0.1	InLoop0
*4.4.4.4/32	Liberal/1025		DS/3.3.3.3	
10.0.1.0/24	NULL/1026	-	10.0.23.2	S2/0/0
10.0.1.0/24	1026/1026	3.3.3.3	10.0.23.2	S2/0/0
10.0.2.0/24	3/NULL	3.3.3.3	10.0.2.1	GE0/0/2
*10.0.2.0/24	Liberal/1027		DS/3.3.3.3	
10.0.12.0/24	NULL/3	-	10.0.23.2	S2/0/0
10.0.12.0/24	1027/3	3.3.3.3	10.0.23.2	S2/0/0
10.0.23.0/24	3/NULL	3.3.3.3	10.0.23.3	S2/0/0
*10.0.23.0/24	Liberal/3		DS/3.3.3.3	

TOTAL: 12 Normal LSP(s) Found.

TOTAL: 4 Liberal LSP(s) Found.

TOTAL: 0 Frr LSP(s) Found.

A '*' before an LSP means the LSP is not established

A '*' before a Label means the USCB or DSCB is stale

A '*' before a UpstreamPeer means the session is in GR state

A '*' before a DS means the session is in GR state

A '*' before a NextHop means the LSP is FRR LSP

Step 5 Configure the LDP inbound policy.

If labels received on R1 are not controlled, R1 will establish a large number of LSPs, consuming large memory.

After an inbound LDP policy is configured, R1 receives label mapping messages only from R2 and establishes LSPs to R2, saving resources.

Run the **display mpls lsp** command on R1. Information about established LSPs is displayed.

[R1]display mpls lsp

LSP Information: LDP LSP

FEC	In/Out Label	In/Out IF	Vrf Name
3.3.3.3/32	NULL/3	-/S1/0/0	
3.3.3.3/32	1024/3	-/S1/0/0	
2.2.2.2/32	3/NULL	-/	
4.4.4.4/32	NULL/1025	-/S1/0/0	
4.4.4.4/32	1025/1025	-/S1/0/0	
10.0.12.0/24	3/NULL	-/	
10.0.1.0/24	3/NULL	-/	
2.2.2.0/24	3/NULL	-/	
10.0.23.0/24	NULL/3	-/S1/0/0	

10.0.23.0/24	1026/3	-/S1/0/0
10.0.2.0/24	NULL/1027	-/S1/0/0
10.0.2.0/24	1027/1027	-/S1/0/0

You can see that LSPs to R2 and R3 are established on R1. Configure the inbound policy on R1 to allow only the routes to R2.

```
[R1]ip ip-prefix prefix1 permit 10.0.12.0 24
```

```
[R1]mpls ldp
```

```
[R1-mpls-ldp]inbound peer 3.3.3.3 fec ip-prefix prefix1
```

```
[R1-mpls-ldp]quit
```

```
[R1]display mpls lsp
```

LSP Information: LDP LSP

FEC	In/Out Label	In/Out IF	Vrf Name
2.2.2.2/32	3/NULL	-/-	
10.0.12.0/24	3/NULL	-/-	
10.0.1.0/24	3/NULL	-/-	
2.2.2.0/24	3/NULL	-/-	

----End

Additional Exercise: Analysis and Verification

How can you configure R1 to receive Label Mapping messages from R1 to R3?

Device Configuration

```
<S1>display current-configuration
```

```
!Software Version V200R008C00SPC500
```

```
#
```

```
sysname S1
```

```
#
```

```
interface Vlanif1
```

```
ip address 10.0.1.2 255.255.255.0
```

```
#
```

```
ospf 1 router-id 1.1.1.1
```

```
area 0.0.0.0
```

```
network 10.0.1.0 0.0.0.255
```

```
#
```

```
return
```

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

```
[V200R007C00SPC600]
```

```
#
```

```
sysname R1
```

```
#
```

```
mpls lsr-id 2.2.2.2
```

```
mpls
```

```
lsp-trigger all
```

```
#
```

```
mpls ldp
```

```
inbound peer 3.3.3.3 fec ip-prefix prefix1
```

```
#
```

```
interface Serial1/0/0
```

```
link-protocol ppp
```

```
ip address 10.0.12.1 255.255.255.0
```

```
mpls
```

```
mpls ldp
```

```
#
```

```
interface GigabitEthernet0/0/1
  ip address 10.0.1.1 255.255.255.0
#
interface LoopBack0
  ip address 2.2.2.2 255.255.255.0
#
ospf 1 router-id 2.2.2.2
  area 0.0.0.0
    network 10.0.1.0 0.0.0.255
    network 10.0.12.0 0.0.0.255
    network 2.2.2.0 0.0.0.255
#
  ip ip-prefix prefix1 index 10 permit 10.0.12.0 24
#
return

<R2> display current-configuration
[V200R007C00SPC600]
#
sysname R2
#
mpls lsr-id 3.3.3.3
mpls
  lsp-trigger all
#
mpls ldp
#
interface Serial1/0/0
  link-protocol ppp
```

```
ip address 10.0.12.2 255.255.255.0
```

```
mpls
```

```
mpls ldp
```

```
#
```

```
interface Serial2/0/0
```

```
link-protocol ppp
```

```
ip address 10.0.23.2 255.255.255.0
```

```
mpls
```

```
mpls ldp
```

```
#
```

```
interface LoopBack0
```

```
ip address 3.3.3.3 255.255.255.0
```

```
#
```

```
ospf 1 router-id 3.3.3.3
```

```
area 0.0.0.0
```

```
network 10.0.12.0 0.0.0.255
```

```
network 10.0.23.0 0.0.0.255
```

```
network 3.3.3.0 0.0.0.255
```

```
#
```

```
return
```

```
<R3> display current-configuration
```

```
[V200R007C00SPC600]
```

```
#
```

```
sysname R3
```

```
#
```

```
mpls lsr-id 4.4.4.4
```

```
mpls
```

```
lsp-trigger all
```

```
#
mpls ldp
#
interface Serial2/0/0
  link-protocol ppp
  ip address 10.0.23.3 255.255.255.0
  mpls
  mpls ldp
#
interface GigabitEthernet0/0/2
  ip address 10.0.2.1 255.255.255.0
#
interface LoopBack0
  ip address 4.4.4.4 255.255.255.0
#
ospf 1 router-id 4.4.4.4
  area 0.0.0.0
    network 10.0.2.0 0.0.0.255
    network 10.0.23.0 0.0.0.255
    network 4.4.4.0 0.0.0.255
#
return

<S2>display current-configuration
!Software Version V200R008C00SPC500
#
sysname S2
#
interface Vlanif1
```

```
ip address 10.0.2.2 255.255.255.0
```

```
#
```

```
ospf 1 router-id 5.5.5.5
```

```
area 0.0.0.0
```

```
network 10.0.2.0 0.0.0.255
```

```
#
```

```
return
```

Lab 1-2 MPLS VPN Configuration

Learning Objectives

The objectives of this lab are to learn and understand:

- How to configure MPLS VPN instances
- How to configure MP-BGP
- How to configure MPLS LDP
- MPLS VPN route transmission and data forwarding processes

Topology




Figure 1-2 MPLS VPN topology

Scenario

An enterprise has networks A and B. Employees on the two networks are required to communicate through VPN routes. The edge device needs to use the Border

Gateway Protocol (BGP) to advertise VPN routes to the carrier network. The carrier uses MP-BGP to transmit VPN routes on the public network, and ensures security and privacy of customer network information through MPLS VPN.

Tasks

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

Configure IP addresses and masks for all routers.

```
<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 10.1.12.1 24
```

```
[R1-Serial1/0/0]quit
```

```
[R1]interface Serial 3/0/0
```

```
[R1-Serial3/0/0]ip address 10.1.14.1 24
```

```
[R1-Serial3/0/0]quit
```

```
[R1]interface LoopBack 0
```

```
[R1-LoopBack0]ip address 1.1.1.1 32
```

```
<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 10.1.12.2 24
```

```
[R2-Serial1/0/0]quit
```

```
[R2]interface Serial 2/0/0
```

```
[R2-Serial2/0/0]ip address 10.1.23.2 24
```

```
[R1-Serial2/0/0]quit
```

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

[R2-LoopBack0]ip address 2.2.2.2 32

<Huawei>system-view

Enter system view, return user view with Ctrl+Z.

[Huawei]sysname R3

[R3]interface Serial 2/0/0

[R3-Serial2/0/0]ip address 10.1.23.3 24

[R3-Serial2/0/0]quit

[R3]interface Serial 3/0/0

[R3-Serial3/0/0]ip address 10.1.35.3 24

[R3-Serial3/0/0]quit

[R3]interface LoopBack 0

[R3-LoopBack0]ip address 3.3.3.3 32

<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 10.1.14.4 24

[R4-Serial1/0/0]quit

[R4]interface LoopBack 0

[R4-LoopBack0]ip address 192.168.1.1 24

<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 10.1.35.5 24

[R5-Serial1/0/0]quit

```
[R5]interface LoopBack 0
```

```
[R5-LoopBack0]ip address 192.168.2.1 24
```

Test link connectivity after the configurations are complete.

Step 2 **Configure a single OSPF area on the carrier network.**

Add 10.1.12.0/24, 10.1.23.0/24, and addresses of Loopback0 interfaces on the carrier network to OSPF area 0.

```
[R1]router id 1.1.1.1
```

```
[R1]ospf 1
```

```
[R1-ospf-1]area 0
```

```
[R1-ospf-1-area-0.0.0.0]network 10.1.12.0 0.0.0.255
```

```
[R1-ospf-1-area-0.0.0.0]network 1.1.1.1 0.0.0.0
```

```
[R2]router id 2.2.2.2
```

```
[R2]ospf 1
```

```
[R2-ospf-1]area 0
```

```
[R2-ospf-1-area-0.0.0.0]network 10.1.12.0 0.0.0.255
```

```
[R2-ospf-1-area-0.0.0.0]network 10.1.23.0 0.0.0.255
```

```
[R2-ospf-1-area-0.0.0.0]network 2.2.2.2 0.0.0.0
```

```
[R3]router id 3.3.3.3
```

```
[R3]ospf 1
```

```
[R3-ospf-1]area 0
```

```
[R3-ospf-1-area-0.0.0.0]network 10.1.23.0 0.0.0.255
```

```
[R3-ospf-1-area-0.0.0.0]network 3.3.3.3 0.0.0.0
```

Check the OSPF neighbor relationship on R1, R2, and R3 after the configurations are complete.

[R1]display ospf peer brief

OSPF Process 1 with Router ID 1.1.1.1

Peer Statistic Information

Area Id	Interface	Neighbor id	State
0.0.0.0	Serial1/0/0	2.2.2.2	Full

Total Peer(s): 1

[R2]display ospf peer brief

OSPF Process 1 with Router ID 2.2.2.2

Peer Statistic Information

Area Id	Interface	Neighbor id	State
0.0.0.0	Serial1/0/0	1.1.1.1	Full
0.0.0.0	Serial2/0/0	3.3.3.3	Full

Total Peer(s): 2

[R3]display ospf peer brief

OSPF Process 1 with Router ID 3.3.3.3

Peer Statistic Information

Area Id	Interface	Neighbor id	State
0.0.0.0	Serial2/0/0	2.2.2.2	Full

Total Peer(s): 1

Step 3 **Configure VPN instances on edge devices of the carrier network.**

Configure VPN instances for network A and network B on R1 and R3 respectively. Set the VPN instance to **VPN1**, router distinguisher (RD) to 1:1, and route target to 1:2 for network A. Set the VPN instance to **VPN2**, RD to 2:2, and route target to 1:2 for network B.

```
[R1]ip vpn-instance VPN1
```

```
[R1-vpn-instance-VPN1]route-distinguisher 1:1
```

```
[R1-vpn-instance-VPN1-af-ipv4]vpn-target 1:2 both
```

```
[R1-vpn-instance-VPN1-af-ipv4]quit
```

```
[R1-vpn-instance-VPN1]quit
```

```
[R1]interface Serial 3/0/0
```

```
[R1-Serial3/0/0]ip binding vpn-instance VPN1
```

Info: All IPv4 related configurations on this interface are removed!

Info: All IPv6 related configurations on this interface are removed!

```
[R1-Serial3/0/0] ip address 10.1.14.1 24
```

```
[R3]ip vpn-instance VPN2
```

```
[R3-vpn-instance-VPN2]route-distinguisher 2:2
```

```
[R3-vpn-instance-VPN2-af-ipv4]vpn-target 1:2 both
```

```
[R3-vpn-instance-VPN2-af-ipv4]quit
```

```
[R3-vpn-instance-VPN2]quit
```

```
[R3]interface Serial 3/0/0
```

```
[R3-Serial3/0/0]ip binding vpn-instance VPN2
```

Info: All IPv4 related configurations on this interface are removed!

Info: All IPv6 related configurations on this interface are removed!

[R3-Serial3/0/0]ip address 10.1.35.3 24

Check VPN instances on R1 and R3 after the configurations are complete.

[R1]display ip vpn-instance verbose

Total VPN-Instances configured : 1

Total IPv4 VPN-Instances configured : 1

Total IPv6 VPN-Instances configured : 0

VPN-Instance Name and ID : VPN1, 1

Interfaces : Serial3/0/0

Address family ipv4

Create date : 2016/09/20 14:51:08

Up time : 0 days, 00 hours, 09 minutes and 34 seconds

Route Distinguisher : 1:1

Export VPN Targets : 1:2

Import VPN Targets : 1:2

Label Policy : label per route

Log Interval : 5

[R3]display ip vpn-instance verbose

Total VPN-Instances configured : 1

Total IPv4 VPN-Instances configured : 1

Total IPv6 VPN-Instances configured : 0

VPN-Instance Name and ID : VPN2, 1

Interfaces : Serial3/0/0

Address family ipv4

Create date : 2016/09/20 15:02:52

Up time : 0 days, 00 hours, 05 minutes and 32 seconds

Route Distinguisher : 2:2

Export VPN Targets : 1:2

Import VPN Targets : 1:2

Label Policy : label per route

Log Interval : 5

Step 4 Configure BGP to transmit routes on edge devices of the customer networks(CE) and carrier network(PE).

Set AS numbers of network A, carrier network, and network B to 14, 123, and 35 respectively. Establish BGP neighbor relationships between CE and PE to advertise customer VPN routes to PE using BGP.

```
[R1]bgp 123
```

```
[R1-bgp]ipv4-family vpn-instance VPN1
```

```
[R1-bgp-VPN1]peer 10.1.14.4 as-number 14
```

```
[R3]bgp 123
```

```
[R3-bgp]ipv4-family vpn-instance VPN2
```

```
[R3-bgp-VPN2]peer 10.1.35.5 as-number 35
```

```
[R4]bgp 14
```

```
[R4-bgp]peer 10.1.14.1 as-number 123
```

```
[R4-bgp]network 192.168.1.0 24
```

```
[R5]bgp 35
```

```
[R5-bgp]peer 10.1.35.3 as-number 123
```

```
[R5-bgp]network 192.168.2.0 24
```

Check the OSPF neighbor relationship between R1 and R4 and between R3 and R5 after the configurations are complete.

[R1]display bgp vpnv4 vpn-instance VPN1 peer

BGP local router ID : 1.1.1.1

Local AS number : 123

VPN-Instance VPN1, Router ID 1.1.1.1:

Total number of peers : 1 Peers in established state : 1

Peer	V	AS	MsgRcvd	MsgSent	OutQ	Up/Down	State	PrefRcv
10.1.14.4	4	14	7	8	0	00:05:21	Established	0

[R4]display bgp peer

BGP local router ID : 10.1.14.4

Local AS number : 14

Total number of peers : 1 Peers in established state : 1

Peer	V	AS	MsgRcvd	MsgSent	OutQ	Up/Down	State	PrefRcv
10.1.14.1	4	123	4	6	0	00:02:56	Established	0

[R3]display bgp vpnv4 vpn-instance VPN2 peer

BGP local router ID : 3.3.3.3

Local AS number : 123

VPN-Instance VPN2, Router ID 3.3.3.3:

Total number of peers : 1 Peers in established state : 1

Peer	V	AS	MsgRcvd	MsgSent	OutQ	Up/Down	State	PrefRcv
------	---	----	---------	---------	------	---------	-------	---------

10.1.35.5	4	35	7	8	0 00:05:16	Established	0
-----------	---	----	---	---	------------	-------------	---

[R5]display bgp peer

BGP local router ID : 192.168.1.1

Local AS number : 35

Total number of peers : 1

Peers in established state : 1

Peer	V	AS	MsgRcvd	MsgSent	OutQ	Up/Down	State	PrefRcv
10.1.35.3	4	123	8	10	0 00:06:04	Established	0	

Check VPN routes learned from customer networks in VPN routing table on R1 and R3.

[R1]display ip routing-table vpn-instance VPN1

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

Routing Tables: VPN1

Destinations : 6 Routes : 6

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
10.1.14.0/24	Direct	0	0	D	10.1.14.1	Serial3/0/0
10.1.14.1/32	Direct	0	0	D	127.0.0.1	Serial3/0/0
10.1.14.4/32	Direct	0	0	D	10.1.14.4	Serial3/0/0
10.1.14.255/32	Direct	0	0	D	127.0.0.1	Serial3/0/0
192.168.1.0/24	EBGP	255	0	D	10.1.14.4	Serial3/0/0
255.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0

[R3]display ip routing-table vpn-instance VPN2

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

Routing Tables: VPN2

Destinations : 6 Routes : 6

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
10.1.35.0/24	Direct	0	0	D	10.1.35.3	Serial3/0/0
10.1.35.3/32	Direct	0	0	D	127.0.0.1	Serial3/0/0
10.1.35.5/32	Direct	0	0	D	10.1.35.5	Serial3/0/0
10.1.35.255/32	Direct	0	0	D	127.0.0.1	Serial3/0/0
192.168.2.0/24	EBGP	255	0	D	10.1.35.5	Serial3/0/0
255.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0

Step 5 Configure devices on the carrier network to transmit customer VPN routes using MP-BGP.

Establish the IBGP neighbor relationship between R1 and R3, and transmit customer VPN routes to the remote PE using MP-BGP.

[R1]bgp 123

[R1-bgp]peer 3.3.3.3 as-number 123

[R1-bgp]peer 3.3.3.3 connect-interface LoopBack 0

[R1-bgp]ipv4-family vpnv4 unicast

[R1-bgp-af-vpnv4]peer 3.3.3.3 enable

[R3]bgp 123

[R3-bgp]peer 1.1.1.1 as-number 123

[R3-bgp]peer 1.1.1.1 connect-interface LoopBack 0

```
[R3-bgp]ipv4-family vpnv4 unicast
```

```
[R3-bgp-af-vpnv4]peer 1.1.1.1 enable
```

Check the MP-BGP neighbor relationship on R1 and R3 after the configurations are complete.

```
[R1]display bgp vpnv4 all peer
```

BGP local router ID : 1.1.1.1

Local AS number : 123

Total number of peers : 2

Peers in established state : 2

Peer	V	AS	MsgRcvd	MsgSent	OutQ	Up/Down	State	PrefRcv
3.3.3.3	4	123	4	7	0	00:02:10	Established	0

```
[R3]display bgp vpnv4 all peer
```

BGP local router ID : 3.3.3.3

Local AS number : 123

Total number of peers : 2

Peers in established state : 2

Peer	V	AS	MsgRcvd	MsgSent	OutQ	Up/Down	State	PrefRcv
1.1.1.1	4	123	5	6	0	00:03:22	Established	0

Step 6 **Configure devices on the carrier network to forward customer VPN data using MPLS LDP.**

Enable MPLS LDP on each device of the carrier network, and use labels to forward customer VPN data to isolate customer data from other network data.

```
[R1]mpls lsr-id 1.1.1.1
[R1]mpls
[R1-mpls]mpls ldp
[R1-mpls-ldp]quit
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]mpls
[R1-Serial1/0/0]mpls ldp
```

```
[R2]mpls lsr-id 2.2.2.2
[R2]mpls
[R2-mpls]mpls ldp
[R2-mpls-ldp]quit
[R2]interface s1/0/0
[R2-Serial1/0/0]mpls
[R2-Serial1/0/0]mpls ldp
[R2-Serial1/0/0]quit
[R2]interface s2/0/0
[R2-Serial2/0/0]mpls
[R2-Serial2/0/0]mpls ldp
```

```
[R3]mpls lsr-id 3.3.3.3
[R3]mpls
[R3-mpls]mpls ldp
[R3-mpls-ldp]quit
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]mpls
[R3-Serial2/0/0]mpls ldp
```

Check the MPLS LDP neighbor relationship on R1, R2, and R3 after the configurations are complete.

[R1]display mpls ldp peer

LDP Peer Information in Public network

A '*' before a peer means the peer is being deleted.

```
-----  
PeerID           TransportAddress  DiscoverySource  
-----  
2.2.2.2:0        2.2.2.2          Serial1/0/0  
-----
```

TOTAL: 1 Peer(s) Found.

[R2]display mpls ldp peer

LDP Peer Information in Public network

A '*' before a peer means the peer is being deleted.

```
-----  
PeerID           TransportAddress  DiscoverySource  
-----  
1.1.1.1:0        1.1.1.1          Serial1/0/0  
3.3.3.3:0        3.3.3.3          Serial2/0/0  
-----
```

TOTAL: 2 Peer(s) Found.

[R3]display mpls ldp peer

LDP Peer Information in Public network

A '*' before a peer means the peer is being deleted.

```
-----  
PeerID           TransportAddress  DiscoverySource  
-----
```

2.2.2.2:0 2.2.2.2 Serial2/0/0

TOTAL: 1 Peer(s) Found.

Step 7 Test the connectivity between network A and network B on CEs.

Use Loopback0 to simulate the user network on R4 and R5 respectively, and run the **ping** command to test connectivity between network A and network B.

<R4>ping -a 192.168.1.1 192.168.2.1

PING 192.168.2.1: 56 data bytes, press CTRL_C to break

Reply from 192.168.2.1: bytes=56 Sequence=1 ttl=252 time=106 ms

Reply from 192.168.2.1: bytes=56 Sequence=2 ttl=252 time=107 ms

Reply from 192.168.2.1: bytes=56 Sequence=3 ttl=252 time=106 ms

Reply from 192.168.2.1: bytes=56 Sequence=4 ttl=252 time=105 ms

Reply from 192.168.2.1: bytes=56 Sequence=5 ttl=252 time=106 ms

--- 192.168.2.1 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

round-trip min/avg/max = 105/106/107 ms

<R5>ping -a 192.168.2.1 192.168.1.1

PING 192.168.1.1: 56 data bytes, press CTRL_C to break

Reply from 192.168.1.1: bytes=56 Sequence=1 ttl=252 time=107 ms

Reply from 192.168.1.1: bytes=56 Sequence=2 ttl=252 time=105 ms

Reply from 192.168.1.1: bytes=56 Sequence=3 ttl=252 time=106 ms

Reply from 192.168.1.1: bytes=56 Sequence=4 ttl=252 time=106 ms

Reply from 192.168.1.1: bytes=56 Sequence=5 ttl=252 time=106 ms

--- 192.168.1.1 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

round-trip min/avg/max = 105/106/107 ms

Check routes learned from remote networks on R4 and R5.

<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	
10.1.14.0/24	Direct	0	0	D 10.1.14.4	Serial1/0/0	
10.1.14.1/32	Direct	0	0	D 10.1.14.1	Serial1/0/0	
10.1.14.4/32	Direct	0	0	D 127.0.0.1	Serial1/0/0	
10.1.14.255/32	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	
192.168.1.0/24	Direct	0	0	D 192.168.1.1	LoopBack0	
192.168.1.1/32	Direct	0	0	D 127.0.0.1	LoopBack0	
192.168.1.255/32	Direct	0	0	D 127.0.0.1	LoopBack0	
192.168.2.0/24	EBGP	255	0	D 10.1.14.1	Serial1/0/0	
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 : 12 Routes : 12

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
10.1.35.0/24	Direct	0	0	D	10.1.35.5	Serial1/0/0
10.1.35.3/32	Direct	0	0	D	10.1.35.3	Serial1/0/0
10.1.35.5/32	Direct	0	0	D	127.0.0.1	Serial1/0/0
10.1.35.255/32	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
192.168.1.0/24	EBGP	255	0	D	10.1.35.3	Serial1/0/0
192.168.2.0/24	Direct	0	0	D	192.168.2.1	LoopBack0
192.168.2.1/32	Direct	0	0	D	127.0.0.1	LoopBack0
192.168.2.255/32	Direct	0	0	D	127.0.0.1	LoopBack0
255.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0

----End

Additional Exercise: Analysis and Verification

When another MPLS VPN is added on R1, how R1 is configured to enable communication between the two VPNs?

Device Configuration

<R1>display current-configuration

[V200R007C00SPC600]

#

sysname R1

```
#
router id 1.1.1.1

#
ip vpn-instance VPN1
    ipv4-family
        route-distinguisher 1:1
        vpn-target 1:2 export-extcommunity
        vpn-target 1:2 import-extcommunity
#
mpls lsr-id 1.1.1.1
mpls
#
mpls ldp
#
interface Serial1/0/0
    link-protocol ppp
    ip address 10.1.12.1 255.255.255.0
    mpls
    mpls ldp
#
interface Serial3/0/0
    link-protocol ppp
    ip binding vpn-instance VPN1
    ip address 10.1.14.1 255.255.255.0
#
interface LoopBack0
    ip address 1.1.1.1 255.255.255.255
#
bgp 123
```

```
peer 3.3.3.3 as-number 123

peer 3.3.3.3 connect-interface LoopBack0

#

ipv4-family unicast

undo synchronization

peer 3.3.3.3 enable

#

ipv4-family vpnv4

policy vpn-target

peer 3.3.3.3 enable

#

ipv4-family vpn-instance VPN1

peer 10.1.14.4 as-number 14

#

ospf 1

area 0.0.0.0

network 1.1.1.1 0.0.0.0

network 10.1.12.0 0.0.0.255

#

return
```

<R2>display current-configuration

```
[V200R007C00SPC600]
```

```
#

sysname R2

#

router id 2.2.2.2

#
```

```
mpls lsr-id 2.2.2.2

mpls
#

mpls ldp
#

interface Serial1/0/0

link-protocol ppp

ip address 10.1.12.2 255.255.255.0

mpls

mpls ldp
#

interface Serial2/0/0

link-protocol ppp

ip address 10.1.23.2 255.255.255.0

mpls

mpls ldp
#

interface LoopBack0

ip address 2.2.2.2 255.255.255.255
#

ospf 1

area 0.0.0.0

network 2.2.2.2 0.0.0.0

network 10.1.12.0 0.0.0.255

network 10.1.23.0 0.0.0.255
#

return
```

<R3>display current-configuration

[V200R007C00SPC600]

#

sysname R3

#

router id 3.3.3.3

#

ip vpn-instance VPN2

ipv4-family

route-distinguisher 2:2

vpn-target 1:2 export-extcommunity

vpn-target 1:2 import-extcommunity

#

mpls lsr-id 3.3.3.3

mpls

#

mpls ldp

#

interface Serial2/0/0

link-protocol ppp

ip address 10.1.23.3 255.255.255.0

mpls

mpls ldp

#

interface Serial3/0/0

link-protocol ppp

ip binding vpn-instance VPN2

ip address 10.1.35.3 255.255.255.0

#

```
interface LoopBack0
  ip address 3.3.3.3 255.255.255.255
#
bgp 123
  peer 1.1.1.1 as-number 123
  peer 1.1.1.1 connect-interface LoopBack0
#
  ipv4-family unicast
    undo synchronization
    peer 1.1.1.1 enable
#
  ipv4-family vpnv4
    policy vpn-target
    peer 1.1.1.1 enable
#
  ipv4-family vpn-instance VPN2
    peer 10.1.35.5 as-number 35
#
ospf 1
  area 0.0.0.0
    network 3.3.3.3 0.0.0.0
    network 10.1.23.0 0.0.0.255
#
return
```

<R4> display current-configuration

[V200R007C00SPC600]

#

```
sysname R4

#

interface Serial1/0/0

  link-protocol ppp

  ip address 10.1.14.4 255.255.255.0

#

interface LoopBack0

  ip address 192.168.1.1 255.255.255.0

#

bgp 14

  peer 10.1.14.1 as-number 123

#

  ipv4-family unicast

    undo synchronization

    network 192.168.1.0

    peer 10.1.14.1 enable

#

return

<R5>display current-configuration

[V200R007C00SPC600]

#

sysname R5

#

interface Serial1/0/0

  link-protocol ppp

  ip address 10.1.35.5 255.255.255.0

#

interface LoopBack0
```

```
ip address 192.168.2.1 255.255.255.0
```

```
#
```

```
bgp 35
```

```
peer 10.1.35.3 as-number 123
```

```
#
```

```
ipv4-family unicast
```

```
undo synchronization
```

```
network 192.168.2.0
```

```
peer 10.1.35.3 enable
```

```
#
```

```
return
```

Chapter 2 DHCP Configuration

Lab 2-1 DHCP Configuration

Learning Objectives

The objectives of this lab are to learn and understand:

- How to configure an IP address pool
- How to configure the DHCP server
- How to configure the DHCP client
- How to configure the DHCP relay
- How to configure basic functions of DHCP snooping

Topology




Figure 2-1 DHCP configuration

Scenario

Assume that you are a network administrator of an enterprise. A DHCP server needs to be configured on the network because it is difficult to manage many hosts with static addresses.

R1 functions as the DHCP server, R4 as the DHCP client, and R2 as the gateway for devices on S1. DHCP Discover packets are broadcast packets and cannot pass

through routers, so a DHCP relay agent is deployed to send DHCP Request packets from R2 to R1. S2 does not require any configuration, and only transparently transmit packets.

To improve network security and prevent DHCP clients from obtaining incorrect IP addresses from other DHCP servers, deploy DHCP snooping on S1 so that R4 obtains the IP address of R1 (DHCP server 1) but not the IP address of R3 (DHCP server 2). To further enhance security, enable some features of DHCP snooping to prevent DHCP exhaustion and DHCP man-in-the-middle attacks.

Tasks

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

Configure IP addresses and masks for all routers.

```
<Huawei>system-view
```

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname R1
```

```
[R1]interface GigabitEthernet 0/0/2
```

```
[R1-GigabitEthernet0/0/2]ip address 10.0.12.1 24
```

```
[R1-GigabitEthernet0/0/2]quit
```

```
[R1]interface loopback 0
```

```
[R1-LoopBack0]ip address 1.1.1.1 32
```

```
<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 10.0.12.2 24
```

```
[R2]interface GigabitEthernet 0/0/1
```

```
[R2-GigabitEthernet0/0/1]ip address 10.10.10.1 24
```

```
<Huawei>system-view
```

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname R3
```

```
[R3]interface GigabitEthernet 0/0/1
```

```
[R3-GigabitEthernet0/0/1]ip address 192.168.1.1 24
```

To enable an interface of R4 to obtain IP addresses using DHCP, enable the DHCP client function on the interface.

```
<Huawei>system-view
```

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname R4
```

```
[R4]dhcp enable
```

```
[R4]interface GigabitEthernet 0/0/1
```

```
[R4-GigabitEthernet0/0/1] ip address dhcp-alloc
```

Configure the switch name and disable unnecessary interfaces.

```
<Huawei>system-view
```

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname S1
```

```
[S1]interface GigabitEthernet 0/0/9
```

```
[S1-GigabitEthernet0/0/9]shutdown
```

```
[S1-GigabitEthernet0/0/9]quit
```

```
[S1]interface GigabitEthernet 0/0/10
```

```
[S1-GigabitEthernet0/0/10]shutdown
```

```
[S1-GigabitEthernet0/0/10]quit
```

```
[S1]interface GigabitEthernet 0/0/13
```

```
[S1-GigabitEthernet0/0/13]shutdown
```

```
[S1-GigabitEthernet0/0/13]quit
```

```
[S1]interface GigabitEthernet 0/0/14
```

```
[S1-GigabitEthernet0/0/14]shutdown
```

```
<Huawei>system-view
```

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname S2
```

```
[S2]interface GigabitEthernet 0/0/6
```

```
[S2-GigabitEthernet0/0/6]shutdown
```

```
[S2-GigabitEthernet0/0/6]quit
```

```
[S2]interface GigabitEthernet 0/0/7
```

```
[S2-GigabitEthernet0/0/7]shutdown
```

Test the connectivity between R2 and R1.

```
[R1]ping 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=1 ms
```

```
Reply from 10.0.12.2: bytes=56 Sequence=2 ttl=255 time=1 ms
```

```
Reply from 10.0.12.2: bytes=56 Sequence=3 ttl=255 time=1 ms
```

```
Reply from 10.0.12.2: bytes=56 Sequence=4 ttl=255 time=1 ms
```

```
Reply from 10.0.12.2: bytes=56 Sequence=5 ttl=255 time=1 ms
```

```
--- 10.0.12.2 ping statistics ---
```

```
5 packet(s) transmitted
```

```
5 packet(s) received
```

```
0.00% packet loss
```

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

Step 2 **Configure a route between R1 and R2.**

R1 advertises the route of its loopback interface to R2, and R2 advertises the route of

the interface connected to S1 to R1 so that the LAN gateway and external network can communicate.

```
[R1]ospf 1
```

```
[R1-ospf-1]area 0
```

```
[R1-ospf-1-area-0.0.0.0]network 1.1.1.1 0.0.0.0
```

```
[R1-ospf-1-area-0.0.0.0]network 10.0.12.0 0.0.0.255
```

```
[R2]ospf 1
```

```
[R2-ospf-1]silent-interface GigabitEthernet 0/0/1
```

```
[R2-ospf-1]area 0
```

```
[R2-ospf-1-area-0.0.0.0]network 10.10.10.0 0.0.0.255
```

```
[R2-ospf-1-area-0.0.0.0]network 10.0.12.0 0.0.0.255
```

Configure the interface connected to R2 as a silent interface to ensure route advertisement of the network segment. However, no neighbor relationship is established on this interface. Test the connectivity between the two networks.

```
[R2]ping -a 10.10.10.1 1.1.1.1
```

```
  PING 1.1.1.1: 56  data bytes, press CTRL_C to break
```

```
    Reply from 1.1.1.1: bytes=56 Sequence=1 ttl=255 time=1 ms
```

```
    Reply from 1.1.1.1: bytes=56 Sequence=2 ttl=255 time=1 ms
```

```
    Reply from 1.1.1.1: bytes=56 Sequence=3 ttl=255 time=1 ms
```

```
    Reply from 1.1.1.1: bytes=56 Sequence=4 ttl=255 time=1 ms
```

```
    Reply from 1.1.1.1: bytes=56 Sequence=5 ttl=255 time=1 ms
```

```
--- 1.1.1.1 ping statistics ---
```

```
  5 packet(s) transmitted
```

```
  5 packet(s) received
```

```
  0.00% packet loss
```

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

Step 3 **Configure IP address pools.**

Create IP address pools on R1 and R3 respectively. The IP address pool on R1 ranges from 10.10.10.0 to 10.10.10.24, the IP address of G0/0/0 on R2 is 10.10.10.1, and the DNS server address is 1.1.1.1. To prevent some static addresses on this network from being allocated, reserve 10.10.10.2 to 10.10.10.10 from being dynamically allocated through DHCP. The IP address pool on R3 ranges from 192.168.1.0 to 192.168.1.24, the IP address of G0/0/0 on R3 is 192.168.1.1, and the DNS server address is 192.168.1.1. To reserve IP addresses 192.168.1.2 to 192.168.1.10 from being dynamically allocated through DHCP, set the leases of two servers to 3 days.

```
[R1]ip pool DHCP
```

```
[R1-ip-pool-DHCP]gateway-list 10.10.10.1
```

```
[R1-ip-pool-DHCP]network 10.10.10.0 mask 255.255.255.0
```

```
[R1-ip-pool-DHCP]excluded-ip-address 10.10.10.2 10.10.10.10
```

```
[R1-ip-pool-DHCP]dns-list 1.1.1.1
```

```
[R1-ip-pool-DHCP]lease day 3
```

```
[R3]ip pool DHCP
```

```
[R3-ip-pool-DHCP]gateway-list 192.168.1.1
```

```
[R3-ip-pool-DHCP]network 192.168.1.0 mask 255.255.255.0
```

```
[R3-ip-pool-DHCP]excluded-ip-address 192.168.1.2 192.168.1.10
```

```
[R3-ip-pool-DHCP]dns-list 192.168.1.1
```

```
[R1-ip-pool-DHCP]lease day 3
```

Verify the IP address pool configuration.

```
<R1>display ip pool
```

```
-----
```

```

Pool-name      : DHCP
Pool-No       : 0
Lease         : 3 Days 0 Hours 0 Minutes
Position      : Local          Status          : Unlocked
Gateway-0     : 10.10.10.1
Network       : 10.10.10.0
Mask          : 255.255.255.0
VPN instance   : --
Address Statistic: Total      :253      Used      :0
                  Idle       :244      Expired    :0
                  Conflict   :0        Disable    :9
  
```

IP address Statistic

```

Total      :253
Used       :0      Idle      :244
Expired    :0      Conflict   :0      Disable   :9
  
```

<R3>display ip pool

```

-----
Pool-name      : DHCP
Pool-No       : 0
Lease         : 3 Days 0 Hours 0 Minutes
Position      : Local          Status          : Unlocked
Gateway-0     : 192.168.1.1
Network       : 192.168.1.0
Mask          : 255.255.255.0
VPN instance   : --
  
```

Address Statistic: Total	:253	Used	:0
Idle	:244	Expired	:0
Conflict	:0	Disable	:9

IP address Statistic

Total	:253		
Used	:0	Idle	:244
Expired	:0	Conflict	:0
		Disable	:9

Step 4 Configure a DHCP server based on the global address pool.

DHCP address pool parameters have been configured, but cannot be used by clients. DHCP needs to be enabled globally and on the interface.

```
[R3]dhcp enable
```

```
[R3]interface GigabitEthernet 0/0/1
```

```
[R3-GigabitEthernet0/0/1]dhcp select global
```

After DHCP is configured on R3, R4 can obtain the IP address normally.

```
<R4>display ip interface GigabitEthernet 0/0/1
```

```
GigabitEthernet0/0/1 current state : UP
```

```
Line protocol current state : UP
```

```
The Maximum Transmit Unit : 1500 bytes
```

```
input packets : 0, bytes : 0, multicasts : 0
```

```
output packets : 17, bytes : 5605, multicasts : 0
```

```
Directed-broadcast packets:
```

```
received packets:      0, sent packets:      17
```

```
forwarded packets:    0, dropped packets:    0
```

ARP packet input number: 0

Request packet: 0

Reply packet: 0

Unknown packet: 0

Internet Address is allocated by DHCP, 192.168.1.254/24

Broadcast address : 192.168.1.255

TTL being 1 packet number: 0

TTL invalid packet number: 0

ICMP packet input number: 0

Echo reply: 0

Unreachable: 0

Source quench: 0

Routing redirect: 0

Echo request: 0

Router advert: 0

Router solicit: 0

Time exceed: 0

IP header bad: 0

Timestamp request: 0

Timestamp reply: 0

Information request: 0

Information reply: 0

Netmask request: 0

Netmask reply: 0

Unknown type: 0

The IP address of this interface is 192.168.1.254, which is obtained through DHCP.

Step 5 **Configure the DHCP relay.**

The configuration of R3 as a temporary DHCP server is complete. The actual DHCP

server is R1. Because DHCP Discover packets cannot be directly sent from DHCP clients to R1. In this case, configure R2 as the DHCP relay agent and specify R2 as the gateway of the LAN connected to S1. Then R2 can transmit DHCP Request packets of DHCP clients.

Enable DHCP on R1.

```
[R1]dhcp enable
[R1]interface GigabitEthernet 0/0/2
[R1-GigabitEthernet0/0/2]dhcp select global
```

On R2, specify the IP address of the DHCP server as 10.0.12.1 and configure the DHCP relay on the interface.

```
[R2]dhcp enable
[R2]dhcp server group DHCP
[R2-dhcp-server-group-DHCP]dhcp-server 10.0.12.1
[R2-dhcp-server-group-DHCP]quit
[R2]interface GigabitEthernet 0/0/1
[R2-GigabitEthernet0/0/1]dhcp select relay
[R2-GigabitEthernet0/0/1]dhcp relay server-select DHCP
```

Verify the DHCP relay configuration on R2.

```
[R2]display dhcp server group

Group-name          : DHCP
(0)  Server-IP      : 10.0.12.1
Gateway             : --
VPN instance        : --
1 DHCP server group(s) in total
```

```
[R2]display dhcp relay all
```

DHCP relay agent running information of interface GigabitEthernet0/0/1 :

Server group name : DHCP

Gateway address in use : 10.10.10.1

The DHCP server group is configured on R2 and the IP address of the DHCP server in the DHCP server group is 10.0.12.1. The DHCP relay function is enabled on G0/0/1 of R2, and the DHCP relay agent sends DHCP Request packets to the DHCP server.

To further verify whether the DHCP relay is deployed successfully, disable the interface on R3 (to prevent R2 from obtaining an IP address from R3), disable the interface on R4, and then enable the interfaces again. Normally, R4 can obtain the subnet address of 10.10.10.0/24.

```
[R3]interface GigabitEthernet 0/0/1
```

```
[R3-GigabitEthernet0/0/1]shutdown
```

```
[R4]interface GigabitEthernet 0/0/1
```

```
[R4-GigabitEthernet0/0/1]shutdown
```

```
[R4-GigabitEthernet0/0/1]undo shutdown
```

```
[R4]display ip interface GigabitEthernet 0/0/1
```

GigabitEthernet0/0/1 current state : UP

Line protocol current state : UP

The Maximum Transmit Unit : 1500 bytes

input packets : 0, bytes : 0, multicasts : 0

output packets : 36, bytes : 11866, multicasts : 0

Directed-broadcast packets:

received packets: 0, sent packets: 36

forwarded packets: 0, dropped packets: 0

ARP packet input number: 0

Request packet: 0

Reply packet: 0

Unknown packet: 0

Internet Address is allocated by DHCP, 10.10.10.254/24

Broadcast address : 10.10.10.255

TTL being 1 packet number: 0

TTL invalid packet number: 0

ICMP packet input number: 0

Echo reply: 0

Unreachable: 0

Source quench: 0

Routing redirect: 0

Echo request: 0

Router advert: 0

Router solicit: 0

Time exceed: 0

IP header bad: 0

Timestamp request: 0

Timestamp reply: 0

Information request: 0

Information reply: 0

Netmask request: 0

Netmask reply: 0

Unknown type: 0

R4 successfully obtains the IP address 10.10.10.254. Check statistics on R2 and the IP address pool status on R1.

<R2> display dhcp relay statistics

The statistics of DHCP RELAY:

```

DHCP packets received from clients    : 2
    DHCP DISCOVER packets received    : 1
    DHCP REQUEST packets received     : 1
    DHCP RELEASE packets received     : 0
    DHCP INFORM packets received      : 0
    DHCP DECLINE packets received     : 0
DHCP packets sent to clients          : 2
    Unicast packets sent to clients    : 2
    Broadcast packets sent to clients  : 0
DHCP packets received from servers    : 2
    DHCP OFFER packets received       : 1
    DHCP ACK packets received         : 1
    DHCP NAK packets received         : 0
DHCP packets sent to servers          : 2
DHCP Bad packets received             : 0
    
```

<R1>display ip pool

```

-----
Pool-name      : DHCP
Pool-No       : 0
Lease         : 3 Days 0 Hours 0 Minutes
Position      : Local          Status      : Unlocked
Gateway-0     : 10.10.10.1
Network       : 10.10.10.0
Mask          : 255.255.255.0
VPN instance   : --
Address Statistic: Total      :253      Used      :1
                  Idle        :243      Expired   :0
                  Conflict    :0        Disable   :9
    
```

IP address Statistic

```
Total      :253
Used       :1      Idle      :243
Expired    :0      Conflict  :0      Disable   :9
```

Check the route of R4, and test the connectivity from R4 to the loopback interface of R1.

<R4>display ip routing-table

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

Routing Tables: Public

Destinations : 8 Routes : 8

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
0.0.0.0/0	Unr	60	0	D	10.10.10.1	GigabitEthernet0/0/0
10.10.10.0/24	Direct	0	0	D	10.10.10.254	GigabitEthernet0/0/0
10.10.10.254/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/0
10.10.10.255/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/0
127.0.0.0/8	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
127.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
255.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0

<R4>ping 1.1.1.1

PING 1.1.1.1: 56 data bytes, press CTRL_C to break

Reply from 1.1.1.1: bytes=56 Sequence=1 ttl=254 time=1 ms

Reply from 1.1.1.1: bytes=56 Sequence=2 ttl=254 time=1 ms

Reply from 1.1.1.1: bytes=56 Sequence=3 ttl=254 time=1 ms

Reply from 1.1.1.1: bytes=56 Sequence=4 ttl=254 time=1 ms

Reply from 1.1.1.1: bytes=56 Sequence=5 ttl=254 time=1 ms

--- 1.1.1.1 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

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

Step 6 **Configure DHCP snooping and attack defense features.**

In the preceding step, the interface on R3 is disabled temporarily. In this case, R4 obtains the IP address only from R1 through R2. Which configuration is required when the interface on R3 needs to be enabled and R4 does not obtain an IP address from R3? In particular, DHCP is enabled on SOHO-level routers by default. If they are connected to the network, there are serious security risks. You are advised to enable DHCP snooping on S1 to prevent unauthorized DHCP servers from interfering with hosts on the LAN.

On S1, defense against DHCP exhaustion attacks and DHCP man-in-the-middle attacks can be enabled to further protect the network where devices obtain IP addresses through DHCP.

Configure DHCP snooping to prevent unauthorized DHCP servers from providing IP addresses.

```
[S1]dhcp enable
```

```
[S1]dhcp snooping enable
```

```
[S1]interface GigabitEthernet 0/0/3
```

```
[S1-GigabitEthernet0/0/3]dhcp snooping enable
```

```
[S1-GigabitEthernet0/0/3]quit
```

```
[S1]interface GigabitEthernet 0/0/2
```

```
[S1-GigabitEthernet0/0/2]dhcp snooping enable
```

By default, the interface where DHCP snooping is enabled is an untrusted interface.

```
[S1]display dhcp snooping
```

DHCP snooping global running information :

DHCP snooping	: Enable	
Static user max number	: 1024	
Current static user number	: 0	
Dhcp user max number	: 1024	(default)
Current dhcp user number	: 0	
Arp dhcp-snooping detect	: Disable	(default)
Alarm threshold	: 100	(default)
Check dhcp-rate	: Disable	(default)
Dhcp-rate limit(pps)	: 100	(default)
Alarm dhcp-rate	: Disable	(default)
Alarm dhcp-rate threshold	: 100	(default)
Discarded dhcp packets for rate limit	: 0	
Bind-table autosave	: Disable	(default)
Offline remove mac-address	: Disable	(default)
Client position transfer allowed	: Enable	(default)

DHCP snooping running information for interface GigabitEthernet0/0/2 :

DHCP snooping	: Enable	
Trusted interface	: No	
Dhcp user max number	: 1024	(default)
Current dhcp user number	: 0	

Check dhcp-giaddr	: Disable (default)
Check dhcp-chaddr	: Disable (default)
Alarm dhcp-chaddr	: Disable (default)
Check dhcp-request	: Disable (default)
Alarm dhcp-request	: Disable (default)
Check dhcp-rate	: Disable (default)
Alarm dhcp-rate	: Disable (default)
Alarm dhcp-rate threshold	: 100
Discarded dhcp packets for rate limit	: 0
Alarm dhcp-reply	: Disable (default)

DHCP snooping running information for interface GigabitEthernet0/0/3 :

DHCP snooping	: Enable
Trusted interface	: No
Dhcp user max number	: 1024 (default)
Current dhcp user number	: 0
Check dhcp-giaddr	: Disable (default)
Check dhcp-chaddr	: Disable (default)
Alarm dhcp-chaddr	: Disable (default)
Check dhcp-request	: Disable (default)
Alarm dhcp-request	: Disable (default)
Check dhcp-rate	: Disable (default)
Alarm dhcp-rate	: Disable (default)
Alarm dhcp-rate threshold	: 100
Discarded dhcp packets for rate limit	: 0
Alarm dhcp-reply	: Disable (default)

Enable the interface on R4 again. In this case, R4 cannot obtain an IP address from

any DHCP server because the interfaces connected to the two servers are untrusted interfaces.

```
[R4]interface GigabitEthernet 0/0/1
```

```
[R4-GigabitEthernet0/0/1]shutdown
```

```
[R4-GigabitEthernet0/0/1]undo shutdown
```

```
[r4]display ip interface GigabitEthernet 0/0/1
```

GigabitEthernet0/0/1 current state : UP

Line protocol current state : DOWN

The Maximum Transmit Unit : 1500 bytes

input packets : 0, bytes : 0, multicasts : 0

output packets : 8, bytes : 2624, multicasts : 0

Directed-broadcast packets:

received packets:	0, sent packets:	8
-------------------	------------------	---

forwarded packets:	0, dropped packets:	0
--------------------	---------------------	---

ARP packet input number:	0
--------------------------	---

Request packet:	0
-----------------	---

Reply packet:	0
---------------	---

Unknown packet:	0
-----------------	---

Internet protocol processing : disabled

Broadcast address : 0.0.0.0

TTL being 1 packet number:	0
----------------------------	---

TTL invalid packet number:	0
----------------------------	---

ICMP packet input number:	0
---------------------------	---

Echo reply:	0
-------------	---

Unreachable:	0
--------------	---

Source quench:	0
----------------	---

Routing redirect:	0
-------------------	---

Echo request:	0
---------------	---

Router advert:	0
Router solicit:	0
Time exceed:	0
IP header bad:	0
Timestamp request:	0
Timestamp reply:	0
Information request:	0
Information reply:	0
Netmask request:	0
Netmask reply:	0
Unknown type:	0

Because R1 is the authorized DHCP server, the interface on the switch connected to R2 can be configured as the trusted interface.

```
[S1]interface GigabitEthernet 0/0/2
```

```
[S1-GigabitEthernet0/0/2]dhcp snooping trusted
```

Check DHCP snooping on this interface.

```
[S1]display dhcp snooping interface GigabitEthernet 0/0/2
```

DHCP snooping running information for interface GigabitEthernet0/0/2 :

DHCP snooping	: Enable
Trusted interface	: Yes
Dhcp user max number	: 1024 (default)
Current dhcp user number	: 0
Check dhcp-giaddr	: Disable (default)
Check dhcp-chaddr	: Disable (default)
Alarm dhcp-chaddr	: Disable (default)
Check dhcp-request	: Disable (default)

Alarm dhcp-request	: Disable (default)
Check dhcp-rate	: Disable (default)
Alarm dhcp-rate	: Disable (default)
Alarm dhcp-rate threshold	: 100
Discarded dhcp packets for rate limit	: 0
Alarm dhcp-reply	: Disable (default)

R4 can obtain an IP address again.

[R4]display ip interface GigabitEthernet 0/0/0

GigabitEthernet0/0/0 current state : UP

Line protocol current state : UP

The Maximum Transmit Unit : 1500 bytes

input packets : 0, bytes : 0, multicasts : 0

output packets : 94, bytes : 30832, multicasts : 0

Directed-broadcast packets:

received packets:	0, sent packets:	94
forwarded packets:	0, dropped packets:	0
ARP packet input number:	0	
Request packet:	0	
Reply packet:	0	
Unknown packet:	0	

Internet Address is allocated by DHCP, 10.10.10.254/24

Broadcast address : 10.10.10.255

TTL being 1 packet number: 0

TTL invalid packet number: 0

ICMP packet input number: 0

Echo reply: 0

Unreachable: 0

Source quench:	0
Routing redirect:	0
Echo request:	0
Router advert:	0
Router solicit:	0
Time exceed:	0
IP header bad:	0
Timestamp request:	0
Timestamp reply:	0
Information request:	0
Information reply:	0
Netmask request:	0
Netmask reply:	0
Unknown type:	0

The preceding configurations are complete. Assume that R4 is an untrusted host that may send many DHCP Request packets to exhaust available IP addresses. Defense against DHCP exhaustion attacks is enabled on the interface of S1 connected to R4.

```
[S1]interface GigabitEthernet 0/0/4
```

```
[S1-GigabitEthernet0/0/4]dhcp snooping check dhcp-chaddr enable
```

Check whether the value of **Check dhcp-chaddr** is **Enable**. If the value of **Check dhcp-chaddr** is **Enable**, the switch checks the CHADDR field in the received DHCP Request packet and determines whether the value of the CHADDR field is consistent with the hardware address of the host. If the value of the CHADDR field is inconsistent with the hardware address of the host, the interface does not forward the DHCP Request packet.

```
[S1]display dhcp snooping interface GigabitEthernet 0/0/4
```

DHCP snooping running information for interface GigabitEthernet0/0/4 :

DHCP snooping	: Disable (default)
Trusted interface	: No
Dhcp user max number	: 1024 (default)
Current dhcp user number	: 0
Check dhcp-giaddr	: Disable (default)
Check dhcp-chaddr	: Enable
Alarm dhcp-chaddr	: Disable (default)
Check dhcp-request	: Disable (default)
Alarm dhcp-request	: Disable (default)
Check dhcp-rate	: Disable (default)
Alarm dhcp-rate	: Disable (default)
Alarm dhcp-rate threshold	: 100
Discarded dhcp packets for rate limit	: 0
Alarm dhcp-reply	: Disable (default)

Enable defense against DHCP man-in-the-middle attacks.

[S1]arp dhcp-snooping-detect enable

Check global DHCP snooping.

[S1]display dhcp snooping

DHCP snooping global running information :

DHCP snooping	: Enable
Static user max number	: 1024
Current static user number	: 0
Dhcp user max number	: 1024 (default)
Current dhcp user number	: 0
Arp dhcp-snooping detect	: Enable

Alarm threshold	: 100	(default)
Check dhcp-rate	: Disable	(default)
Dhcp-rate limit(pps)	: 100	(default)
Alarm dhcp-rate	: Disable	(default)
Alarm dhcp-rate threshold	: 100	(default)
Discarded dhcp packets for rate limit	: 0	
Bind-table autosave	: Disable	(default)
Offline remove mac-address	: Disable	(default)
Client position transfer allowed	: Enable	(default)

Association between ARP and DHCP snooping is enabled. By default, association between ARP and DHCP snooping is disabled. The DHCP security defense configurations are complete.

----End

Configuration Reference

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

```
[V200R007C00SPC600]
```

```
#
```

```
sysname R1
```

```
#
```

```
dhcp enable
```

```
#
```

```
ip pool DHCP
```

```
gateway-list 10.10.10.1
```

```
network 10.10.10.0 mask 255.255.255.0
```

```
excluded-ip-address 10.10.10.2 10.10.10.10
```

```
lease day 3 hour 0 minute 0
```

```
dns-list 1.1.1.1
```

```
interface GigabitEthernet0/0/2
ip address 10.0.12.1 255.255.255.0
dhcp select global
#
```

```
interface LoopBack0
ip address 1.1.1.1 255.255.255.255
#
```

```
ospf 1
area 0.0.0.0
network 1.1.1.1 0.0.0.0
network 10.0.12.0 0.0.0.255
#
```

```
return
```

```
<R2> display current-configuration
[V200R007C00SPC600]
```

```
#
```

```
sysname R2
```

```
#
```

```
dhcp enable
```

```
#
```

```
dhcp server group DHCP
```

```
dhcp-server 10.0.12.1 0
```

```
#
```

```
interface GigabitEthernet0/0/1
ip address 10.10.10.1 255.255.255.0
dhcp select relay
dhcp relay server-select DHCP
```

```
#
interface GigabitEthernet0/0/2
 ip address 10.0.12.2 255.255.255.0
#
ospf 1
 silent-interface GigabitEthernet0/0/1
 area 0.0.0.0
  network 10.0.12.0 0.0.0.255
  network 10.10.10.0 0.0.0.255
#
return

<R3>display current-configuration
[V200R007C00SPC600]
#
 sysname R3
#
 dhcp enable
#
 ip pool DHCP
 gateway-list 192.168.1.1
 network 192.168.1.0 mask 255.255.255.0
 excluded-ip-address 192.168.1.2 192.168.1.10
 lease day 3 hour 0 minute 0
 dns-list 192.168.1.1
#
 interface GigabitEthernet0/0/1
 ip address 192.168.1.1 255.255.255.0
 dhcp select global
```

#

return

<R4>display current-configuration

[V200R007C00SPC600]

#

sysname R4

#

dhcp enable

#

interface GigabitEthernet0/0/1

ip address dhcp-alloc

#

return

<SW1>display current-configuration

!Software Version V200R008C00SPC500

#

sysname S1

#

dhcp enable

#

dhcp snooping enable

arp dhcp-snooping-detect enable

#

#

interface GigabitEthernet0/0/2

dhcp snooping enable

dhcp snooping trusted

```
#  
interface GigabitEthernet0/0/3  
    dhcp snooping enable  
#  
interface GigabitEthernet0/0/4  
    dhcp snooping check dhcp-chaddr enable  
#  
interface GigabitEthernet0/0/9  
    shutdown  
#  
interface GigabitEthernet0/0/10  
    shutdown  
#  
interface GigabitEthernet0/0/13  
    shutdown  
#  
interface GigabitEthernet0/0/14  
    shutdown  
#  
return
```

```
<SW2>display current-configuration  
!Software Version V200R008C00SPC500
```

```
#  
sysname SW2  
#  
interface GigabitEthernet0/0/6  
    shutdown  
#
```

```
interface GigabitEthernet0/0/7  
    shutdown  
  
#  
  
return
```

Chapter 3 Service Quality and Traffic Control

Lab 3-1 QoS Basics

Learning Objectives

The objectives of this lab are to learn and understand:

- How to analyze the SLA using NQA
- How to implement priority mapping and traffic policing
- How to configure traffic shaping
- How to implement congestion management based on queues and traffic classifiers
- How to configure congestion avoidance based on WRED

Topology




Figure 3-1 QoS

Scenario

Assume that you are a network administrator of an enterprise. R1 and S1 are located in the enterprise headquarters, and R2 and S2 are located in the enterprise branch. The headquarters and branch are connected through a leased line.

The intranet bandwidth increases gradually, but the bandwidth of the leased line does not increase. In this case, important services are delayed or some services are unavailable.

You can use differentiated services of QoS and adjust QoS parameters to ensure that important service data is first sent to the destination.

In the lab, S3 and S4 use NQA to exchange a large number of data flows. R3, R4, and R5 simulate the clients and server to check whether important applications are available.

Tasks

Step 1 Perform basic configurations and configure IP addresses.

Configure IP addresses and masks for all the routers and switches S3 and S4.

Set the baud rate of S1/0/0 on R1 to 72000, and simulate congestion on the WAN link because of insufficient bandwidth.

```
<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 10.0.12.1 255.255.255.0
```

```
[R1-Serial1/0/0]baudrate 72000
```

```
[R1-Serial1/0/0]interface GigabitEthernet 0/0/1
```

```
[R1-GigabitEthernet0/0/1]ip address 10.0.145.1 255.255.255.0
```

```
<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 10.0.12.2 255.255.255.0
```

[R2-Serial1/0/0]interface GigabitEthernet 0/0/2

[R2-GigabitEthernet0/0/2]ip address 10.0.34.2 255.255.255.0

<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 10.0.34.3 255.255.255.0

<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 10.0.145.4 255.255.255.0

<Huawei>system-view

Enter system view, return user view with Ctrl+Z.

[Huawei]sysname R5

[R5]interface GigabitEthernet 0/0/1

[R5-GigabitEthernet0/0/1]ip address 10.0.145.5 255.255.255.0

<Huawei>system-view

Enter system view, return user view with Ctrl+Z.

[Huawei]sysname S3

[S3]interface Vlanif 1

[S3-Vlanif1]ip address 10.0.145.3 255.255.255.0

<Huawei>system-view

Enter system view, return user view with Ctrl+Z.

[Huawei]sysname S4

[S4]interface Vlanif 1

[S4-Vlanif1]ip address 10.0.34.4 255.255.255.0

After the configurations are complete, test the connectivity of direct links.

[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=36 ms

--- 10.0.12.2 ping statistics ---

1 packet(s) transmitted

1 packet(s) received

0.00% packet loss

round-trip min/avg/max = 36/36/36 ms

[R1]ping -c 1 10.0.145.3

PING 10.0.145.3: 56 data bytes, press CTRL_C to break

Reply from 10.0.145.3: bytes=56 Sequence=1 ttl=255 time=35 ms

--- 10.0.145.3 ping statistics ---

1 packet(s) transmitted

1 packet(s) received

0.00% packet loss

round-trip min/avg/max = 35/35/35 ms

[R1]ping -c 1 10.0.145.4

PING 10.0.145.4: 56 data bytes, press CTRL_C to break

Reply from 10.0.145.4: bytes=56 Sequence=1 ttl=255 time=6 ms

--- 10.0.145.4 ping statistics ---

1 packet(s) transmitted

1 packet(s) received

0.00% packet loss

round-trip min/avg/max = 6/6/6 ms

[R1]ping -c 1 10.0.145.5

PING 10.0.145.5: 56 data bytes, press CTRL_C to break

Reply from 10.0.145.5: bytes=56 Sequence=1 ttl=255 time=6 ms

--- 10.0.145.5 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 10.0.34.3

PING 10.0.34.3: 56 data bytes, press CTRL_C to break

Reply from 10.0.34.3: bytes=56 Sequence=1 ttl=255 time=5 ms

--- 10.0.34.3 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 10.0.34.4

PING 10.0.34.4: 56 data bytes, press CTRL_C to break

Reply from 10.0.34.4: bytes=56 Sequence=1 ttl=255 time=36 ms

--- 10.0.34.4 ping statistics ---

1 packet(s) transmitted

1 packet(s) received

0.00% packet loss

round-trip min/avg/max = 36/36/36 ms

Step 2 **Configure static routes and NQA.**

Configure static routes for all the routers and switches S3 and S4.

```
[R1]ip route-static 10.0.34.0 255.255.255.0 10.0.12.2
```

```
[R2]ip route-static 10.0.145.0 255.255.255.0 10.0.12.1
```

```
[R3]ip route-static 0.0.0.0 0.0.0.0 10.0.34.2
```

```
[R4]ip route-static 0.0.0.0 0.0.0.0 10.0.145.1
```

```
[R5]ip route-static 0.0.0.0 0.0.0.0 10.0.145.1
```

```
[S3]ip route-static 0.0.0.0 0.0.0.0 10.0.145.1
```

```
[S4]ip route-static 0.0.0.0 0.0.0.0 10.0.34.2
```

After the configurations are complete, test the network connectivity.

```
[S3]ping -c 1 10.0.34.4
```

PING 10.0.34.4: 56 data bytes, press CTRL_C to break

Reply from 10.0.34.4: bytes=56 Sequence=1 ttl=252 time=40 ms

--- 10.0.34.4 ping statistics ---

1 packet(s) transmitted

1 packet(s) received

0.00% packet loss

round-trip min/avg/max = 40/40/40 ms

[R4]ping -c 1 10.0.34.3

PING 10.0.145.4: 56 data bytes, press CTRL_C to break

Reply from 10.0.145.4: bytes=56 Sequence=1 ttl=255 time=3 ms

--- 10.0.145.4 ping statistics ---

1 packet(s) transmitted

1 packet(s) received

0.00% packet loss

round-trip min/avg/max = 3/3/3 ms

[R5]ping -c 1 10.0.34.3

PING 10.0.34.3: 56 data bytes, press CTRL_C to break

Reply from 10.0.34.3: bytes=56 Sequence=1 ttl=253 time=44 ms

--- 10.0.34.3 ping statistics ---

1 packet(s) transmitted

1 packet(s) received

0.00% packet loss

round-trip min/avg/max = 44/44/44 ms

The links between S3 and S4, between R4 and R3, and between R5 and R3 are reachable, indicating that network communication is normal.

Congestion easily occurs on the 72 kbit/s serial link between the headquarters and branch.

Use NQA to generate traffic. S4 functions as the NQA server and S3 functions as the NQA client.

Create NQA UDP and jitter test instances to simulate data and voice traffic respectively.

Set parameters in NQA test instances to simulate the environment where congestion does not occur if there is only data or voice traffic and congestion occurs if there is data and voice traffic.

Configure S4 as the NQA server, and set the IP address of the interface used for monitoring UDP services to 10.0.34.4 and port number to 6000.

```
[S4]nqa-server udpecho 10.0.34.4 6000
```

On S3, configure an NQA UDP test instance to simulate data traffic, and set the ToS to 28, packet size to 5800 bytes, interval at which packets are sent to 1s, interval for the NQA test to 3s, and timeout interval for the NQA test to 1s, and start the NQA UDP test instance.

```
[S3]nqa test-instance admin udp
```

```
[S3-nqa-admin-udp]test-type udp
```

```
[S3-nqa-admin-udp]destination-address ipv4 10.0.34.4
```

```
[S3-nqa-admin-udp]destination-port 6000
```

```
[S3-nqa-admin-udp]tos 28
```

```
[S3-nqa-admin-udp]datasize 5800
```

```
[S3-nqa-admin-udp]interval seconds 1
```

```
[S3-nqa-admin-udp]frequency 3
```

```
[S3-nqa-admin-udp]timeout 1
```

[S3-nqa-admin-udp]start now

Check the NQA UDP test instance.

[S3]display nqa results test-instance admin udp

1 . Test 2 result The test is finished

Send operation times: 3	Receive response times: 3
Completion:success	RTD OverThresholds number: 0
Attempts number:1	Drop operation number:0
Disconnect operation number:0	Operation timeout number:0
System busy operation number:0	Connection fail number:0
Operation sequence errors number:0	RTT Stats errors number:0
Destination ip address:10.0.34.4	
Min/Max/Average Completion Time: 930/950/943	
Sum/Square-Sum Completion Time: 2830/2669900	
Last Good Probe Time: 2010-10-10 18:10:02.4	
Lost packet ratio: 0 %	

No packet is discarded and congestion does not occur. Shut down the NQA UDP test instance.

[S3]nqa test-instance admin udp

[S3-nqa-admin-udp]stop

On S3, configure an NQA jitter test instance to simulate voice traffic, and set the ToS to 46, packet size to 90 bytes, interval at which packets are sent to 20 ms, interval for the NQA test to 3s, and timeout interval for the NQA test to 1s, and start the NQA jitter test instance.

[S3]nqa test-instance admin jitter

[S3-nqa-admin-jitter]test-type jitter

[S3-nqa-admin-jitter]destination-address ipv4 10.0.34.4

[S3-nqa-admin-jitter]destination-port 6000

[S3-nqa-admin-jitter]tos 46

[S3-nqa-admin-jitter]datasize 90

[S3-nqa-admin-jitter]interval milliseconds 20

[S3-nqa-admin-jitter]frequency 3

[S3-nqa-admin-jitter]timeout 1

[S3-nqa-admin-jitter]start now

Check the NQA jitter test instance.

[S3]display nqa results test-instance admin jitter

NQA entry(admin, jitter) :testflag is active ,testtype is jitter

1 . Test 1 result The test is finished

SendProbe:60	ResponseProbe:60
Completion:success	RTD OverThresholds number:0
Min/Max/Avg/Sum RTT:40/70/54/3260	RTT Square Sum:179800
NumOfRTT:60	Drop operation number:0
Operation sequence errors number:0	RTT Stats errors number:0
System busy operation number:0	Operation timeout number:0
Min Positive SD:10	Min Positive DS:10
Max Positive SD:10	Max Positive DS:10
Positive SD Number:5	Positive DS Number:11
Positive SD Sum:50	Positive DS Sum:110
Positive SD Square Sum:500	Positive DS Square Sum:1100
Min Negative SD:10	Min Negative DS:10
Max Negative SD:10	Max Negative DS:20
Negative SD Number:4	Negative DS Number:10

Negative SD Sum:40	Negative DS Sum:110
Negative SD Square Sum:400	Negative DS Square Sum:1300
Min Delay SD:20	Min Delay DS:19
Avg Delay SD:27	Avg Delay DS:26
Max Delay SD:35	Max Delay DS:34
Packet Loss SD:0	Packet Loss DS:0
Packet Loss Unknown:0	jitter out value:0.0937500
jitter in value:0.2291667	NumberOfOWD:60
OWD SD Sum:1630	OWD DS Sum:1570
TimeStamp unit: ms	

No packet is discarded and congestion does not occur. Shut down the NQA jitter test instance.

```
[S3]nqa test-instance admin jitter
[S3-nqa-admin-jitter]stop
```

Step 3 **Configure priority mapping.**

Run the **ping** command to simulate traffic of less important services, and map DSCP priorities of traffic to BE without QoS guarantee.

Configure G0/0/1 and S1/0/0 on R1 to trust DSCP priorities of packets.

```
[R1]interface GigabitEthernet 0/0/1
[R1-GigabitEthernet0/0/1]trust dscp override
[R1-GigabitEthernet0/0/1]interface Serial 1/0/0
[R1-Serial1/0/0]trust dscp
```

Specify **override** in the **trust** command on G0/0/1 so that DSCP priorities are changed to mapped values after priority mapping is configured on R1.

Run the **ping** command on R4 to simulate the traffic destined for R3 and set the ToS to 26.

```
[R4]ping -tos 26 10.0.34.3
```

Configure priority mapping on R1 and map DSCP priority 26 to 0.

```
[R1]qos map-table dscp-dscp
```

```
[R1-maptbl-dscp-dscp]input 26 output 0
```

Check the priority mapping configuration on R1.

```
[R1]display qos map-table dscp-dscp
```

Input DSCP	DSCP
------------	------

0	0
1	1
2	2
3	3
4	4
5	5
6	6
7	7
8	8
9	9
10	10
11	11
12	12
13	13
14	14
15	15

16	16
17	17
18	18
19	19
20	20
21	21
22	22
23	23
24	24
25	25
26	0
27	27
28	28
29	29
30	30

The preceding information shows that DSCP priority 26 is mapped to 0 and other DSCP priorities use default values.

Step 4 **Configure traffic shaping and traffic policing.**

Start NQA UDP and jitter test instances on S3 to simulate congestion on the 72 kbit/s link between the headquarters and branch.

```
[S3]nqa test-instance admin udp
```

```
[S3-nqa-admin-udp]start now
```

```
[S3-nqa-admin-udp]quit
```

```
[S3]nqa test-instance admin jitter
```

```
[S3-nqa-admin-jitter]start now
```

On R4, run the **ping** command with the packet size of 700 bytes and packet count of 10 to simulate the traffic destined for R3.

```
[R4]ping -s 700 -c 10 10.0.34.3
```

```
PING 10.0.34.3: 700 data bytes, press CTRL_C to break
```

```
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 10.0.34.3: bytes=700 Sequence=9 ttl=253 time=1944 ms
```

```
Request time out
```

```
--- 10.0.34.3 ping statistics ---
```

```
10 packet(s) transmitted
```

```
1 packet(s) received
```

```
90.00% packet loss
```

```
round-trip min/avg/max = 1944/1944/1944 ms
```

Severe congestion occurs on the link between the headquarters and branch, causing serious packet loss. There is a long delay for forwarded data packets. In this case, R4 cannot communicate with R3.

The following describes how to configure traffic policing and traffic shaping to eliminate congestion on the link so that R4 in headquarters can communicate with R3 on the branch.

Configure traffic policing to eliminate congestion. On S1, configure traffic policing

on G0/0/13 and set the CIR to 64 kbit/s.

```
[S1]interface GigabitEthernet 0/0/13
```

```
[S1-GigabitEthernet0/0/13]qos lr inbound cir 64
```

Check the traffic policing configuration on S1.

```
[S1]display qos lr inbound interface GigabitEthernet 0/0/13
```

```
GigabitEthernet0/0/13 lr inbound:
```

```
  cir: 64 Kbps, cbs: 8000 Byte
```

On R4, run the **ping** command with the packet size of 700 bytes and packet count of 10 to simulate the traffic destined for R3.

```
[R4]ping -s 700 -c 10 10.0.34.3
```

```
PING 10.0.34.3: 700  data bytes, press CTRL_C to break
```

```
  Reply from 10.0.34.3: bytes=700 Sequence=1 ttl=253 time=1412 ms
```

```
  Reply from 10.0.34.3: bytes=700 Sequence=2 ttl=253 time=255 ms
```

```
  Reply from 10.0.34.3: bytes=700 Sequence=3 ttl=253 time=736 ms
```

```
  Reply from 10.0.34.3: bytes=700 Sequence=4 ttl=253 time=1746 ms
```

```
  Reply from 10.0.34.3: bytes=700 Sequence=5 ttl=253 time=246 ms
```

```
  Reply from 10.0.34.3: bytes=700 Sequence=6 ttl=253 time=746 ms
```

```
  Reply from 10.0.34.3: bytes=700 Sequence=7 ttl=253 time=1736 ms
```

```
  Reply from 10.0.34.3: bytes=700 Sequence=8 ttl=253 time=258 ms
```

```
  Reply from 10.0.34.3: bytes=700 Sequence=9 ttl=253 time=766 ms
```

```
  Reply from 10.0.34.3: bytes=700 Sequence=10 ttl=253 time=1736 ms
```

```
--- 10.0.34.3 ping statistics ---
```

```
  10 packet(s) transmitted
```

```
  10 packet(s) received
```

```
  0.00% packet loss
```

round-trip min/avg/max = 246/963/1746 ms

No packets are discarded and R4 can communicate with R3 normally, indicating that traffic policing takes effect.

Delete the traffic policing configuration from S1.

```
[S1]interface GigabitEthernet 0/0/13
```

```
[S1-GigabitEthernet0/0/13]undo qos lr inbound
```

The following uses traffic shaping to eliminate congestion. On S3, configure traffic shaping on GE0/0/13 and set the CIR to 64 kbit/s.

```
[S3]interface GigabitEthernet0/0/13
```

```
[S3-GigabitEthernet0/0/13]qos lr outbound cir 64
```

On R4, run the **ping** command with the packet size of 700 bytes and packet count of 10 to simulate the traffic destined for R3.

```
[R4]ping -s 700 -c 10 10.0.34.3
```

```
PING 10.0.34.3: 700 data bytes, press CTRL_C to break
```

```
Reply from 10.0.34.3: bytes=700 Sequence=1 ttl=253 time=240 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=2 ttl=253 time=284 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=3 ttl=253 time=334 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=4 ttl=253 time=224 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=5 ttl=253 time=344 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=6 ttl=253 time=275 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=7 ttl=253 time=534 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=8 ttl=253 time=184 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=9 ttl=253 time=204 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=10 ttl=253 time=314 ms
```

--- 10.0.34.3 ping statistics ---

10 packet(s) transmitted

10 packet(s) received

0.00% packet loss

round-trip min/avg/max = 184/293/534 ms

No packets are discarded and R4 can communicate with R3 normally, indicating that traffic shaping takes effect.

Delete the traffic shaping configuration from S3.

```
[S3]interface GigabitEthernet0/0/13
```

```
[S3-GigabitEthernet0/0/13]undo qos lr outbound
```

On R4, run the **ping** command with the packet size of 700 bytes and packet count of 10 to simulate the traffic destined for R3.

```
[R4]ping -s 700 -c 10 10.0.34.3
```

PING 10.0.34.3: 700 data bytes, press CTRL_C to break

Reply from 10.0.34.3: bytes=700 Sequence=1 ttl=253 time=1918 ms

Request time out

Reply from 10.0.34.3: bytes=700 Sequence=3 ttl=253 time=1762 ms

Request time out

Request time out

Request time out

Request time out

Request time out

Request time out

Request time out

--- 10.0.34.3 ping statistics ---

10 packet(s) transmitted

2 packet(s) received

80.00% packet loss

round-trip min/avg/max = 1762/1840/1918 ms

After the configuration is deleted, many packets are discarded and forwarded data packets are delayed. R4 cannot communicate with R3.

Step 5 **Configure queue-based congestion management and congestion avoidance.**

To prevent network congestion on the link between the headquarters and branch, configure queue-based congestion management and congestion avoidance.

On R1, create a WRED drop profile named **data** based on DSCP priorities and set the upper drop threshold to 90, lower drop threshold to 50, and maximum drop probability to 30.

```
[R1]drop-profile data
```

```
[R1-drop-profile-data]wred dscp
```

```
[R1-drop-profile-data]dscp af32 low-limit 50 high-limit 90 discard-percentage 30
```

Create a queue profile named **queue-profile1** on R1, put data traffic into WFQ queues, bind the queue profile to the WRED drop profile **data**, and put high-priority and delay-sensitive voice traffic to PQ queues.

```
[R1]qos queue-profile queue-profile1
```

```
[R1-qos-queue-profile-queue-profile1]schedule wfq 3 pq 5
```

```
[R1-qos-queue-profile-queue-profile1]queue 3 drop-profile data
```

Apply the queue profile to S1/0/0 of R1.

```
[R1]interface Serial 1/0/0
```

```
[R1-Serial1/0/0]qos queue-profile queue-profile1
```

Check the queue profile configuration.

```
[R1]display qos queue-profile queue-profile1
```

```
Queue-profile: queue-profile1
```

```
Queue  Schedule  Weight  Length(Bytes/Packets)  Gts(CIR/CBS)
```

```
-----
```

3	WFQ	10	0/0	-/-
5	PQ	-	0/0	-/-

Data traffic and voice traffic enter WFQ and PQ queues respectively.

Check the drop profile configuration.

```
[R1]display drop-profile data
```

```
Drop-profile[1]: data
```

```
DSCP          Low-limit  High-limit  Discard-percentage
```

```
-----
```

default	30	100	10
1	30	100	10
2	30	100	10
3	30	100	10
4	30	100	10
5	30	100	10
6	30	100	10
7	30	100	10
cs1	30	100	10
9	30	100	10

af11	30	100	10
11	30	100	10
af12	30	100	10
13	30	100	10
af13	30	100	10
15	30	100	10
cs2	30	100	10
17	30	100	10
af21	30	100	10
19	30	100	10
af22	30	100	10
21	30	100	10
af23	30	100	10
23	30	100	10
cs3	30	100	10
25	30	100	10
af31	30	100	10
27	30	100	10
af32	50	90	30
29	30	100	10
af33	30	100	10
31	30	100	10
cs4	30	100	10
33	30	100	10
af41	30	100	10

Parameters in the WRED drop profile **data** take effect, and other parameters use default values.

Step 6 **Configure flow-based congestion management and congestion avoidance.**

To prevent network congestion on the link between the headquarters and branch, configure flow-based congestion management and congestion avoidance.

Define the traffic exchanged between R4 in the headquarters and R3 on the branch as important traffic and perform QoS guarantee for the traffic so that R4 can communicate with R3.

Delete the queue profile from S1/0/0 on R1.

```
[R1]interface GigabitEthernet 0/0/1
```

```
[R1-GigabitEthernet0/0/1]undo qos queue-profile
```

On R4, run the **ping** command with the source address of 10.0.145.4, packet size of 700 bytes, and packet count of 10 to test the connectivity between R4 and R3.

```
[R4]ping -a 10.0.145.4 -s 700 -c 10 10.0.34.3
```

```
PING 10.0.34.3: 700 data bytes, press CTRL_C to break
```

```
Reply from 10.0.34.3: bytes=700 Sequence=1 ttl=253 time=1279 ms
```

```
Request time out
```

```
Reply from 10.0.34.3: bytes=700 Sequence=3 ttl=253 time=1587 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=4 ttl=253 time=1827 ms
```

```
Request time out
```

```
Reply from 10.0.34.3: bytes=700 Sequence=6 ttl=253 time=1717 ms
```

```
Request time out
```

```
Request time out
```

```
Request time out
```

```
Request time out
```

```
--- 10.0.34.3 ping statistics ---
```

```
10 packet(s) transmitted
```

4 packet(s) received

60.00% packet loss

round-trip min/avg/max = 1279/1602/1827 ms

Congestion has occurred on the link between the headquarters and branch, a large number of packets are discarded, and R4 cannot communicate with R3.

Create ACL 3001 on R1 to match the traffic sent from 10.0.145.4 to 10.0.34.3.

```
[R1]acl number 3001
```

```
[R1-acl-adv-3001]rule 0 per ip source 10.0.145.4 0.0.0.0 destination 10.0.34.3 0.0.0.0
```

Create a traffic classifier **class-ef**, reference ACL 3001 in the traffic classifier, create a traffic behavior **behavior-ef**, set the queue scheduling mode to EF, and set the bandwidth to 10 kbit/s.

```
[R1]traffic classifier class-ef
```

```
[R1-classifier-class-ef]if-match acl 3001
```

```
[R1-classifier-class-ef]quit
```

```
[R1]traffic behavior behavior-ef
```

```
[R1-behavior-behavior-ef]queue ef bandwidth 10
```

Create a traffic classifier **class-af32** to match data traffic with the DSCP priority of AF32, create a traffic behavior **behavior-af32**, set the queue scheduling mode to AF, set the bandwidth to 30 kbit/s, and bind the traffic behavior to the drop profile **data**.

```
[R1]traffic classifier class-af32
```

```
[R1-classifier-class-af32]if-match dscp af32
```

```
[R1-classifier-class-af32]quit
```

```
[R1]traffic behavior behavior-af32
```

```
[R1-behavior-behavior-af32]queue af bandwidth 30
```

```
[R1-behavior-behavior-af32]drop-profile data
```

Create a traffic policy **policy-1**, associate the traffic policy with the traffic classifier **class-ef** and traffic behavior **behavior-ef**, and the traffic classifier **class-af32** and traffic behavior **behavior-af32**, and apply the traffic policy to S1/0/0 on R1.

```
[R1]traffic policy policy-1
```

```
[R1-trafficpolicy-policy-1]classifier class-ef behavior behavior-ef
```

```
[R1-trafficpolicy-policy-1]classifier class-af32 behavior behavior-af32
```

```
[R1-trafficpolicy-policy-1]quit
```

```
[R1]interface Serial 1/0/0
```

```
[R1-Serial1/0/0]traffic-policy policy-1 outbound
```

On R4, run the **ping** command with the source address of 10.0.145.4, packet size of 700 bytes, and packet count of 10 to test the connectivity between R4 and R3.

```
[R4]ping -a 10.0.145.4 -s 700 -c 10 10.0.34.3
```

```
PING 10.0.34.3: 700 data bytes, press CTRL_C to break
```

```
Reply from 10.0.34.3: bytes=700 Sequence=1 ttl=253 time=694 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=2 ttl=253 time=391 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=3 ttl=253 time=361 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=4 ttl=253 time=671 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=5 ttl=253 time=211 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=6 ttl=253 time=611 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=7 ttl=253 time=688 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=8 ttl=253 time=391 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=9 ttl=253 time=301 ms
```

```
Reply from 10.0.34.3: bytes=700 Sequence=10 ttl=253 time=651 ms
```

```
--- 10.0.34.3 ping statistics ---
```

10 packet(s) transmitted

10 packet(s) received

0.00% packet loss

round-trip min/avg/max = 211/497/694 ms

Configure traffic from R4 to R3 to enter EF queues. Then R4 can communicate with R3.

----End

Additional Exercise: Analysis and Verification

QoS uses differentiated services to ensure bandwidth and shorten the delay for various services. Does increased bandwidth solve service quality problems so that QoS is not required?

After the lab, recollect and summarize the QoS process.

Device Configuration

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

```
[V200R007C00SPC600]
```

```
#
```

```
sysname R1
```

```
#
```

```
acl number 3001
```

```
rule 0 permit ip source 10.0.145.4 0 destination 10.0.34.3 0
```

```
#
```

```
drop-profile data
```

```
wred dscp
```

```
dscp af32 low-limit 50 high-limit 90 discard-percentage 30
```

```
#
```

```

qos queue-profile queue-profile1
    queue 3 drop-profile data
    schedule wfq 3 pq 5
#
qos map-table dscp-dscp
    input 26 output 0
#
traffic classifier class-ef operator or
    if-match acl 3001
traffic classifier class-af32 operator or
    if-match dscp af32
#
traffic behavior behavior-ef
    queue ef bandwidth 10 cbs 250
traffic behavior behavior-af32
    queue af bandwidth 30
    drop-profile data
traffic behavior behavior-af32
    queue af bandwidth 30
#
traffic policy policy-1
    classifier class-ef behavior behavior-ef
    classifier class-af32 behavior behavior-af32
#
interface Serial1/0/0
    link-protocol ppp
    ip address 10.0.12.1 255.255.255.0
    trust dscp
    traffic-policy policy-1 outbound

```

```
baudrate 72000

#

interface GigabitEthernet0/0/1

ip address 10.0.145.1 255.255.255.0

trust dscp override

#

ip route-static 10.0.34.0 255.255.255.0 10.0.12.2

#

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 GigabitEthernet0/0/2

ip address 10.0.34.2 255.255.255.0

#

ip route-static 10.0.145.0 255.255.255.0 10.0.12.1

#

return
```

<R3>display current-configuration

[V200R007C00SPC600]

```
#
```

```

sysname R3

#

interface GigabitEthernet0/0/2

ip address 10.0.34.3 255.255.255.0

#

ip route-static 0.0.0.0 0.0.0.0 10.0.34.2

#

return

```

```

<R4> display current-configuration
[V200R007C00SPC600]

#

sysname R4

#

interface GigabitEthernet0/0/1

ip address 10.0.145.4 255.255.255.0

#

ip route-static 0.0.0.0 0.0.0.0 10.0.145.1

#

return

```

```

<R5> display current-configuration
[V200R007C00SPC600]

#

sysname R5

#

interface GigabitEthernet0/0/1

ip address 10.0.145.5 255.255.255.0

#

```

```
ip route-static 0.0.0.0 0.0.0.0 10.0.145.1
```

```
#
```

```
return
```

```
<S3>display current-configuration
```

```
#
```

```
!Software Version V200R008C00SPC500
```

```
sysname S3
```

```
#
```

```
interface Vlanif1
```

```
ip address 10.0.145.3 255.255.255.0
```

```
#
```

```
ip route-static 0.0.0.0 0.0.0.0 10.0.145.1
```

```
#
```

```
nqa test-instance admin udp
```

```
test-type udp
```

```
destination-address ipv4 10.0.34.4
```

```
destination-port 6000
```

```
tos 28
```

```
frequency 3
```

```
interval seconds 1
```

```
timeout 1
```

```
datasize 5800
```

```
start now
```

```
nqa test-instance admin jitter
```

```
test-type jitter
```

```
destination-address ipv4 10.0.34.4
```

```
destination-port 6000
```

```
tos 46
```

frequency 3

interval milliseconds 20

timeout 1

datasize 90

start now

#

return

<S4> display current-configuration

#

!Software Version V200R008C00SPC500

sysname S4

#

interface Vlanif1

ip address 10.0.34.4 255.255.255.0

#

nqa-server udpecho 10.0.34.4 6000

#

ip route-static 0.0.0.0 0.0.0.0 10.0.34.2

#

return

Chapter 4 Firewall Configuration

Lab 4-1 Firewall Zone and Security Policy Configuration

Learning Objectives

The objectives of this lab are to learn and understand:

- How to configure firewall zones
- How to configure security policies

Topology




Figure 4-1 Firewall zone configuration

Scenario

Assume that you are a network administrator of an enterprise. The headquarters network is divided into three zones: trust, untrust, and DMZ. The firewall is used to control data, ensure internal network security, and provide services for external networks through the DMZ.

Tasks

Step 1 **Log in to the device using the console port.**

1. Connect cables of configuration ports.
 - Disable the firewall and configure a power supply for the terminal.
 - Connect the RS-232 serial port of the terminal to the console port of the firewall through the configuration cable.
 - Power on the device after checking the installation.
2. Configure HyperTerminal software. (You can obtain free HyperTerminal software such as PuTTY from the Internet.)
 - Download the PuTTY software to the local PC and double-click it to run the software.
 - Select Session and set Connection type to Serial.
 - Set parameters for connecting the serial port to the device. Figure 1-2 shows parameter settings.




Figure 4-2 Setting PuTTY parameters for connecting the serial port to the firewall

- Click **Open**.
- 3. Press **Enter**, and enter the default administrator account **admin** and password **Admin@123**.
- 4. Modify the password of the default administrator account, and enter the CLI.

To ensure security, the password must meet the minimum complexity requirement. That is, the password must contain at least three combinations of uppercase letters (A to Z), lowercase letters (a to z), digits (0 to 9), and special characters (such as !, @, #, \$, and %).

Remember the new password for future logins.

Step 2 Perform basic configurations and configure IP addresses.

Configure IP addresses and static routes for routers and the firewall, and configure

VLANs on the switch.

```
<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
```

```
[R1]interface loopback 0
```

```
[R1-LoopBack0]ip address 10.0.1.1 24
```

```
<Huawei>system-view
```

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname R2
```

```
[R2]interface GigabitEthernet0/0/1
```

```
[R2-GigabitEthernet0/0/1]ip address 10.0.20.1 24
```

```
[R2-GigabitEthernet0/0/1]quit
```

```
[R2]interface loopback 0
```

```
[R2-LoopBack0]ip address 10.0.2.2 24
```

```
<Huawei>system-view
```

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname R3
```

```
[R3]interface GigabitEthernet 0/0/1
```

```
[R3-GigabitEthernet0/0/1]ip address 10.0.30.1 24
```

```
[R3-GigabitEthernet0/0/1]quit
```

```
[R3]interface loopback 0
```

```
[R3-LoopBack0]ip address 10.0.3.3 24
```

By default, the firewall has configured the IP address of GigabitEthernet0/0/0. You can delete it to prevent address conflicts.

```
<USG6300>system-view
```

Enter system view, return user view with Ctrl+Z.

```
[USG6300]sysname FW
```

```
[FW]int GigabitEthernet 0/0/0
```

```
[FW-GigabitEthernet0/0/0]undo ip address
```

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

```
[FW]interface GigabitEthernet 1/0/0
```

```
[FW-GigabitEthernet1/0/0]ip address 10.0.10.254 24
```

```
[FW-GigabitEthernet1/0/0]quit
```

```
[FW]interface GigabitEthernet 1/0/1
```

```
[FW-GigabitEthernet1/0/1]ip address 10.0.20.254 24
```

```
[FW-GigabitEthernet1/0/1]quit
```

```
[FW]interface GigabitEthernet 1/0/2
```

```
[FW-GigabitEthernet1/0/2]ip address 10.0.30.254 24
```

```
[FW-GigabitEthernet1/0/2]quit
```

Configure VLANs on the switch as required.

```
[Quidway]sysname S1
```

```
[S1]vlan batch 11 to 13
```

```
[S1]interface GigabitEthernet 0/0/1
```

```
[S1-GigabitEthernet0/0/1]port link-type access
```

```
[S1-GigabitEthernet0/0/1]port default vlan 11
```

```
[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 12
```

```
[S1-GigabitEthernet0/0/2]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 0/0/21
[S1-GigabitEthernet0/0/21]port link-type access
[S1-GigabitEthernet0/0/21]port default vlan 11
[S1-GigabitEthernet0/0/21]quit
[S1]interface GigabitEthernet 0/0/22
[S1-GigabitEthernet0/0/22]port link-type access
[S1-GigabitEthernet0/0/22]port default vlan 12
[S1-GigabitEthernet0/0/22]quit
[S1]interface GigabitEthernet 0/0/23
[S1-GigabitEthernet0/0/23]port link-type access
[S1-GigabitEthernet0/0/23]port default vlan 13
```

Configure default routes on R1, R2, and R3 and specific static routes on the firewall to implement connectivity of three network segments that are connected by three Loopback0 interfaces.

```
[R1]ip route-static 0.0.0.0 0 10.0.10.254

[R2]ip route-static 0.0.0.0 0 10.0.20.254

[R3]ip route-static 0.0.0.0 0 10.0.30.254

[FW]ip route-static 10.0.1.0 24 10.0.10.1
[FW]ip route-static 10.0.2.0 24 10.0.20.1
```

```
[FW]ip route-static 10.0.3.0 24 10.0.30.1
```

After the configuration is complete, check routing information on the firewall.

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

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

Routing Tables: Public

Destinations : 11 Routes : 11

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
10.0.1.0/24	Static	60	0	RD	10.0.10.1	GigabitEthernet1/0/0
10.0.2.0/24	Static	60	0	RD	10.0.20.1	GigabitEthernet1/0/1
10.0.3.0/24	Static	60	0	RD	10.0.30.1	GigabitEthernet1/0/2
10.0.10.0/24	Direct	0	0	D	10.0.10.254	GigabitEthernet1/0/0
10.0.10.254/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.20.0/24	Direct	0	0	D	10.0.20.254	GigabitEthernet1/0/1
10.0.20.254/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.30.0/24	Direct	0	0	D	10.0.30.254	GigabitEthernet1/0/2
10.0.30.254/32	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.0/8	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0

Step 3 **Configure firewall zones.**

The firewall has four zones by default: local zone, trust zone, untrust zone, and DMZ. Here, the trust zone, untrust zone, and DMZ are used. Add interfaces to zones. To

prevent address conflicts, delete GE0/0/0 because GE0/0/0 is added to the trust zone by default.

```
[FW]firewall zone dmz
```

```
[FW-zone-dmz]add interface GigabitEthernet 1/0/2
```

```
[FW-zone-dmz]quit
```

```
[FW]firewall zone trust
```

```
[FW-zone-trust]add interface GigabitEthernet 1/0/1
```

```
[FW-zone-trust]undo add interface GigabitEthernet 0/0/0
```

```
[FW-zone-trust]quit
```

```
[FW]firewall zone untrust
```

```
[FW-zone-untrust]add interface GigabitEthernet 1/0/0
```

```
[FW-zone-untrust]quit
```

Check zones where interfaces belong.

```
[FW]display zone interface
```

```
local
```

```
#
```

```
trust
```

```
interface of the zone is (1):
```

```
GigabitEthernet1/0/1
```

```
#
```

```
untrust
```

```
interface of the zone is (1):
```

```
GigabitEthernet1/0/0
```

```
#
```

```
dmz
```

```
interface of the zone is (1):
```

```
GigabitEthernet1/0/2
```

```
#
```

Check the priority of each zone.

```
[FW]display zone
```

```
local
```

```
priority is 100
```

```
#
```

```
trust
```

```
priority is 85
```

```
interface of the zone is (1):
```

```
GigabitEthernet1/0/1
```

```
#
```

```
untrust
```

```
priority is 5
```

```
interface of the zone is (1):
```

```
GigabitEthernet1/0/0
```

```
#
```

```
dmz
```

```
priority is 50
```

```
interface of the zone is (1):
```

```
GigabitEthernet1/0/2
```

```
#
```

You can see that three interfaces have been added to corresponding zones. By default, interfaces in different zones cannot communicate with each other. Traffic between routers cannot pass through zones, so inter-zone security policies are

required to allow traffic to pass.

Step 4 **Configure a security policy.**

If no inter-zone security policy is configured on the firewall or no security policy is matched, the default packet filtering policy is used by default. That is, all traffic is denied.

Configure a security policy to enable devices in the trust zone to access devices in other zones and prevent access between other zones.

```
[FW]security-policy
[FW-policy-security]rule name policy_sec_1
[FW-policy-security-rule-policy_sec_1]source-zone trust
[FW-policy-security-rule-policy_sec_1]destination-zone untrust
[FW-policy-security-rule-policy_sec_1]action permit
[FW-policy-security-rule-policy_sec_1]quit
[FW-policy-security]rule name policy_sec_2
[FW-policy-security-rule-policy_sec_2]source-zone trust
[FW-policy-security-rule-policy_sec_2]destination-zone dmz
[FW-policy-security-rule-policy_sec_2]action permit
[FW-policy-security-rule-policy_sec_2]quit
[FW-policy-security]quit
```

Verify the configuration.

```
[FW]display security-policy all
```

Total:3

RULE ID	RULE NAME	STATE	ACTION	HITTED
---------	-----------	-------	--------	--------

0	default	enable	deny	0
1	policy_sec_1	enable	permit	0
2	policy_sec_2	enable	permit	0

[FW]display security-policy rule policy_sec_1

(0 times matched)

rule name policy_sec_1

source-zone trust

destination-zone untrust

action permit

[FW]display security-policy rule policy_sec_2

(0 times matched)

rule name policy_sec_2

source-zone trust

destination-zone dmz

action permit

Check the connectivity from the trust zone to the untrust zone and DMZ.

[R2]ping -a 10.0.2.2 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=254 time=1 ms

Reply from 10.0.1.1: bytes=56 Sequence=2 ttl=254 time=1 ms

Reply from 10.0.1.1: bytes=56 Sequence=3 ttl=254 time=1 ms

Reply from 10.0.1.1: bytes=56 Sequence=4 ttl=254 time=1 ms

Reply from 10.0.1.1: bytes=56 Sequence=5 ttl=254 time=1 ms

--- 10.0.1.1 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

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

[R2]ping -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=254 time=1 ms

Reply from 10.0.3.3: bytes=56 Sequence=2 ttl=254 time=1 ms

Reply from 10.0.3.3: bytes=56 Sequence=3 ttl=254 time=1 ms

Reply from 10.0.3.3: bytes=56 Sequence=4 ttl=254 time=1 ms

Reply from 10.0.3.3: bytes=56 Sequence=5 ttl=254 time=1 ms

--- 10.0.3.3 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

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

Check the connectivity from the untrust zone to the trust zone and DMZ.

[R1]ping -a 10.0.1.1 10.0.2.2

PING 10.0.2.2: 56 data bytes, press CTRL_C to break

Request time out

Request time out

Request time out

Request time out

Request time out

--- 10.0.2.2 ping statistics ---

5 packet(s) transmitted

0 packet(s) received

100.00% packet loss

[R1]ping -a 10.0.1.1 10.0.3.3

PING 10.0.3.3: 56 data bytes, press CTRL_C to break

Request time out

Request time out

Request time out

Request time out

Request time out

--- 10.0.3.3 ping statistics ---

5 packet(s) transmitted

0 packet(s) received

100.00% packet loss

Check the connectivity from the DMZ to the untrust zone and trust zone.

[R3]ping -a 10.0.3.3 10.0.1.1

PING 10.0.1.1: 56 data bytes, press CTRL_C to break

Request time out

Request time out

Request time out

Request time out

Request time out

--- 10.0.1.1 ping statistics ---

5 packet(s) transmitted

0 packet(s) received

100.00% packet loss

[R3]ping -a 10.0.3.3 10.0.2.2

PING 10.0.2.2: 56 data bytes, press CTRL_C to break

Request time out

Request time out

Request time out

Request time out

Request time out

--- 10.0.2.2 ping statistics ---

5 packet(s) transmitted

0 packet(s) received

100.00% packet loss

Through verification, devices in the trust zone can access the untrust zone and the DMZ, but devices in other zones cannot access each other.

Configure an inter-zone packet filtering policy to allow devices in the untrust zone to access the specified server in the DMZ.

The Telnet service is enabled for the untrust zone on the server at 10.0.3.3 in the DMZ. Enable ICMP ping to test network connectivity.

[FW]security-policy

[FW-policy-security]rule name policy_sec_3

```
[FW-policy-security-rule-policy_sec_3]source-zone untrust
[FW-policy-security-rule-policy_sec_3]destination-zone dmz
[FW-policy-security-rule-policy_sec_3]destination-address 10.0.3.3 mask 255.255.255.255
[FW-policy-security-rule-policy_sec_3]service icmp
[FW-policy-security-rule-policy_sec_3]service telnet
[FW-policy-security-rule-policy_sec_3]action permit
```

Enable the Telnet function on R3 to perform the Telnet test.

```
[R3]telnet server enable
[R3]aaa
[R3-aaa]local-user test password irreversible-cipher Admin@123
[R3-aaa]local-user test service-type telnet
[R3-aaa]quit
[R3]user-interface vty 0 4
[R3-ui-vty0-4]authentication-mode aaa
[R3-ui-vty0-4]protocol inbound telnet
```

Perform ping and Telnet operations from R1 (untrust zone) to R3 (DMZ).

```
<R1>ping 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=1 ms
  Reply from 10.0.3.3: bytes=56 Sequence=2 ttl=254 time=1 ms
  Reply from 10.0.3.3: bytes=56 Sequence=3 ttl=254 time=1 ms
  Reply from 10.0.3.3: bytes=56 Sequence=4 ttl=254 time=1 ms
  Reply from 10.0.3.3: bytes=56 Sequence=5 ttl=254 time=1 ms

--- 10.0.3.3 ping statistics ---
  5 packet(s) transmitted
```

5 packet(s) received

0.00% packet loss

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

<R1>ping 10.0.30.1

PING 10.0.30.1: 56 data bytes, press CTRL_C to break

Request time out

Request time out

Request time out

Request time out

Request time out

--- 10.0.30.1 ping statistics ---

5 packet(s) transmitted

0 packet(s) received

100.00% packet loss

<R1>telnet 10.0.3.3

Press CTRL_] to quit telnet mode

Trying 10.0.3.3 ...

Connected to 10.0.3.3 ...

Login authentication

Username:test

Password:

User last login information:

Access Type: Telnet

IP-Address : 10.0.10.1

Time : 2016-09-25 03:29:23+00:00

<R3>quit

Info:Configuration console exit, please retry to log on

The connection was closed by the remote host

<R1>telnet 10.0.30.1

Press CTRL_] to quit telnet mode

Trying 10.0.30.1 ...

Error: Can't connect to the remote host

<R1>

Through verification, only ICMP and Telnet packets with the specified IP address can pass, and other traffic are denied.

Device Configuration

<S1>display current-configuration

!Software Version V200R008C00SPC500

#

sysname S1

#

vlan batch 11 to 13

#

interface GigabitEthernet0/0/1

port link-type access

```
port default vlan 11
#
interface GigabitEthernet0/0/2
port link-type access
port default vlan 12
#
interface GigabitEthernet0/0/3
port link-type access
port default vlan 13
#
interface GigabitEthernet0/0/21
port link-type access
port default vlan 11
#
interface GigabitEthernet0/0/22
port link-type access
port default vlan 12
#
interface GigabitEthernet0/0/23
port link-type access
port default vlan 13
#
return

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

```
interface GigabitEthernet0/0/1
  ip address 10.0.10.1 255.255.255.0
#
interface LoopBack0
  ip address 10.0.1.1 255.255.255.0
#
ip route-static 0.0.0.0 0.0.0.0 10.0.10.254
#
return
```

```
<R2>display current-configuration
[V200R007C00SPC600]
#
sysname R2
#
interface GigabitEthernet0/0/1
  ip address 10.0.20.1 255.255.255.0
#
interface LoopBack0
  ip address 10.0.2.2 255.255.255.0
#
ip route-static 0.0.0.0 0.0.0.0 10.0.20.254
#
return
```

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

```
#  
  
aaa  
  
local-user test password irreversible-cipher Admin@123  
  
local-user test privilege level 0  
  
local-user test service-type telnet
```

```
#  
  
interface GigabitEthernet0/0/1  
  
ip address 10.0.30.1 255.255.255.0
```

```
#  
  
interface LoopBack0  
  
ip address 10.0.3.3 255.255.255.0
```

```
#  
  
telnet server enable
```

```
#  
  
ip route-static 0.0.0.0 0.0.0.0 10.0.30.254
```

```
#  
  
user-interface vty 0 4  
  
authentication-mode aaa  
  
protocol inbound telnet
```

```
#  
  
return
```

```
<FW> display current-configuration
```

```
#  
  
sysname FW
```

```
#  
  
interface GigabitEthernet1/0/0  
  
ip address 10.0.10.254 255.255.255.0  
  
#
```

```
interface GigabitEthernet1/0/1
  ip address 10.0.20.254 255.255.255.0
#
interface GigabitEthernet1/0/2
  ip address 10.0.30.254 255.255.255.0
#
firewall zone local
  set priority 100
#
firewall zone trust
  set priority 85
  add interface GigabitEthernet1/0/1
#
firewall zone untrust
  set priority 5
  add interface GigabitEthernet1/0/0
#
firewall zone dmz
  set priority 50
  add interface GigabitEthernet1/0/2
#
ip route-static 10.0.1.0 255.255.255.0 10.0.10.1
ip route-static 10.0.2.0 255.255.255.0 10.0.20.1
ip route-static 10.0.3.0 255.255.255.0 10.0.30.1
#
security-policy
  rule name policy_sec_1
    source-zone trust
    destination-zone untrust
```

```
action permit
rule name policy_sec_2
source-zone trust
destination-zone dmz
action permit
rule name policy_sec_3
source-zone untrust
destination-zone dmz
destination-address 10.0.3.3 mask 255.255.255.255
service icmp
service telnet
action permit
#
return
```

Lab 4-2 Firewall NAT Configuration

Learning Objectives

The objectives of this lab are to learn and understand:

- How to configure Network Address and Port Translation (NAPT) based on the address pool on the firewall
- How to configure the NAT server on the firewall

Topology




Figure 4-2 NAT configuration on the firewall

Scenario

Assume that you are a network administrator of an enterprise. The enterprise network is isolated into three zones by the firewall. Users in the trust zone are required to access the untrust zone. You need to deliver Telnet and FTP services provided by the server at 10.0.4.4 in the DMZ, and set the public IP address to 1.1.1.254/24.

Tasks

Step 1 Perform basic configurations and configure IP addresses.

Configure IP addresses and static routes for routers and the firewall, and configure

VLANs on the switch.

```
<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 1.1.1.1 24
```

```
[R1-GigabitEthernet0/0/1]quit
```

```
[R1]interface loopback 0
```

```
[R1-LoopBack0]ip address 11.11.11.11 24
```

```
<Huawei>system-view
```

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname R2
```

```
[R2]interface GigabitEthernet0/0/1
```

```
[R2-GigabitEthernet0/0/1]ip address 10.0.20.2 24
```

```
[R2-GigabitEthernet0/0/1]quit
```

```
[R2]interface loopback 0
```

```
[R2-LoopBack0]ip address 10.0.2.2 24
```

```
<Huawei>system-view
```

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname R3
```

```
[R3]interface GigabitEthernet0/0/1
```

```
[R3-GigabitEthernet0/0/1]ip address 10.0.20.3 24
```

```
[R3-GigabitEthernet0/0/1]quit
```

```
[R3]interface loopback 0
```

```
[R3-LoopBack0]ip address 10.0.3.3 24
```

```
<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 10.0.40.4 24
```

```
[R4-GigabitEthernet0/0/1]quit
```

```
[R4]interface loopback 0
```

```
[R4-LoopBack0]ip address 10.0.4.4 24
```

By default, the firewall has configured the IP address of GigabitEthernet0/0/0. You can delete it to prevent address conflicts.

```
<USG6300>system-view
```

Enter system view, return user view with Ctrl+Z.

```
[USG6300]sysname FW
```

```
[FW]interface GigabitEthernet 0/0/0
```

```
[FW-GigabitEthernet0/0/0]undo ip address
```

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

```
[FW]interface GigabitEthernet 1/0/0
```

```
[FW-GigabitEthernet1/0/0]ip address 10.0.10.254 24
```

```
[FW-GigabitEthernet1/0/0]ip address 1.1.1.254 24 sub
```

```
[FW-GigabitEthernet1/0/0]quit
```

```
[FW]interface GigabitEthernet 1/0/1
```

```
[FW-GigabitEthernet1/0/1]ip address 10.0.20.254 24
```

```
[FW-GigabitEthernet1/0/1]quit
```

```
[FW]interface GigabitEthernet 1/0/2
```

```
[FW-GigabitEthernet1/0/2]ip address 10.0.40.254 24
```

```
[FW-GigabitEthernet1/0/2]quit
```

Configure VLANs on the switch as required.

```
[Quidway]sysname S1
[S1]vlan batch 11 to 13
[S1]interface GigabitEthernet 0/0/1
[S1-GigabitEthernet0/0/1]port link-type access
[S1-GigabitEthernet0/0/1]port default vlan 11
[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 12
[S1-GigabitEthernet0/0/2]quit
[S1]interface GigabitEthernet 0/0/3
[S1-GigabitEthernet0/0/3]port link-type access
[S1-GigabitEthernet0/0/3]port default vlan 12
[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 13
[S1-GigabitEthernet0/0/4]quit
[S1]interface GigabitEthernet 0/0/21
[S1-GigabitEthernet0/0/21]port link-type access
[S1-GigabitEthernet0/0/21]port default vlan 11
[S1-GigabitEthernet0/0/21]quit
[S1]interface GigabitEthernet 0/0/22
[S1-GigabitEthernet0/0/22]port link-type access
[S1-GigabitEthernet0/0/22]port default vlan 12
[S1-GigabitEthernet0/0/22]quit
[S1]interface GigabitEthernet 0/0/23
[S1-GigabitEthernet0/0/23]port link-type access
[S1-GigabitEthernet0/0/23]port default vlan 13
```

Configure default routes on R2, R3, and R4 and specific static routes on the firewall to implement connectivity of network segments that are connected four Loopback0 interfaces. The default route does not need to be defined on R1 used as an Internet device because R1 does not need to know any private network information about the trust zone and DMZ.

```
[R2]ip route-static 0.0.0.0 0 10.0.20.254
```

```
[R3]ip route-static 0.0.0.0 0 10.0.20.254
```

```
[R4]ip route-static 0.0.0.0 0 10.0.40.254
```

```
[FW]ip route-static 10.0.2.0 24 10.0.20.2
```

```
[FW]ip route-static 10.0.3.0 24 10.0.20.3
```

```
[FW]ip route-static 10.0.4.0 24 10.0.40.4
```

```
[FW]ip route-static 0.0.0.0 0 1.1.1.1
```

After the configuration is complete, check routing information on the firewall.

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

06:44:57 2016/09/25

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

Routing Tables: Public

Destinations : 12 Routes : 12

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
0.0.0.0/0	Static	60	0	RD	1.1.1.1	GigabitEthernet1/0/0

1.1.1.0/24	Direct	0	0	D	1.1.1.254	GigabitEthernet1/0/0
1.1.1.254/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.2.0/24	Static	60	0	RD	10.0.20.2	GigabitEthernet1/0/1
10.0.3.0/24	Static	60	0	RD	10.0.20.3	GigabitEthernet1/0/1
10.0.4.0/24	Static	60	0	RD	10.0.40.4	GigabitEthernet1/0/2
10.0.20.0/24	Direct	0	0	D	10.0.20.254	GigabitEthernet1/0/1
10.0.20.254/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.40.0/24	Direct	0	0	D	10.0.40.254	GigabitEthernet1/0/2
10.0.40.254/32	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.0/8	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0

Step 2 Add interfaces to zones.

The firewall has four zones by default: local zone, trust zone, untrust zone, and DMZ. Here, the trust zone, untrust zone, and DMZ are used. Add interfaces to zones. To prevent address conflicts, delete GE0/0/0 because GE0/0/0 is added to the trust zone by default.

```
[FW]firewall zone dmz
```

```
[FW-zone-dmz]add interface GigabitEthernet 1/0/2
```

```
[FW-zone-dmz]quit
```

```
[FW]firewall zone trust
```

```
[FW-zone-trust]add interface GigabitEthernet 1/0/1
```

```
[FW-zone-trust]undo add interface GigabitEthernet 0/0/0
```

```
[FW-zone-trust]quit
```

```
[FW]firewall zone untrust
```

```
[FW-zone-untrust]add interface GigabitEthernet 1/0/0
```

```
[FW-zone-untrust]quit
```

Check zones where interfaces belong.

[FW]display zone interface

local

#

trust

interface of the zone is (1):

GigabitEthernet1/0/1

#

untrust

interface of the zone is (1):

GigabitEthernet1/0/0

#

dmz

interface of the zone is (1):

GigabitEthernet1/0/2

#

Check the priority of each zone.

[FW]display zone

local

priority is 100

#

trust

priority is 85

interface of the zone is (1):

GigabitEthernet1/0/1

#

untrust

priority is 5

interface of the zone is (1):

GigabitEthernet1/0/0

#

dmz

priority is 50

interface of the zone is (1):

GigabitEthernet1/0/2

#

You can see that three interfaces have been added to corresponding zones. By default, interfaces in different zones cannot communicate with each other. Traffic between routers cannot pass through zones, so inter-zone security policies are required to allow traffic to pass.

Step 3 **Configure a security policy.**

If no inter-zone security policy is configured on the firewall or no security policy is matched, the default packet filtering policy is used by default. That is, all traffic is denied.

Configure the firewall to permit data packets sent from network segments 10.0.2.0 and 10.0.3.0 in the trust zone to the untrust zone, and allow Telnet and FTP requests sent from the untrust zone to the destination server at 10.0.4.4 in the DMZ.

```
[FW]security-policy
```

```
[FW-policy-security]rule name policy_sec_1
```

```
[FW-policy-security-rule-policy_sec_1]source-zone trust
```

```
[FW-policy-security-rule-policy_sec_1]destination-zone untrust
```

```
[FW-policy-security-rule-policy_sec_1]source-address 10.0.2.0 mask 255.255.255.0
```

```
[FW-policy-security-rule-policy_sec_1]source-address 10.0.3.0 mask 255.255.255.0
[FW-policy-security-rule-policy_sec_1]action permit
[FW-policy-security-rule-policy_sec_1]quit
[FW-policy-security]rule name policy_sec_2
[FW-policy-security-rule-policy_sec_2]source-zone untrust
[FW-policy-security-rule-policy_sec_2]destination-zone dmz
[FW-policy-security-rule-policy_sec_2]destination-address 10.0.4.4 mask 255.255.255.255
[FW-policy-security-rule-policy_sec_2]service ftp
[FW-policy-security-rule-policy_sec_2]service telnet
[FW-policy-security-rule-policy_sec_2]action permit
```

Step 4 Configure NAT based on source IP addresses.

Use the public IP address 1.1.1.254 to translate the source IP address.

```
[FW]nat address-group group1
[FW-nat-address-group-group1]section 1.1.1.254 1.1.1.254
```

After the configuration is complete, check the address pool status.

```
[FW]display nat address-group
```

NAT address-group information:

ID	: 0	name	: group1
sectionID	: 0	sectionName	: ---
startaddr	: 1.1.1.254	endaddr	: 1.1.1.254
excludeIP	: 0	excludePort	: 0
reference	: 0	vrrp	: ---
vpninstance	: root	natMode	: pat
description	: ---		

Total 1 address-groups

Configure a NAT policy.

```
[FW]nat-policy
[FW-policy-nat]rule name policy_nat_1
[FW-policy-nat-rule-policy_nat_1]source-zone trust
[FW-policy-nat-rule-policy_nat_1]destination-zone untrust
[FW-policy-nat-rule-policy_nat_1]source-address 10.0.2.2 24
[FW-policy-nat-rule-policy_nat_1]source-address 10.0.3.3 24
[FW-policy-nat-rule-policy_nat_1]action nat address-group group1
```

Test the connectivity.

```
[R2]ping 11.11.11.11
PING 11.11.11.11: 56 data bytes, press CTRL_C to break

Request time out
Request time out
Request time out
Request time out
Request time out

--- 11.11.11.11 ping statistics ---
5 packet(s) transmitted
0 packet(s) received
100.00% packet loss

[R2]ping -a 10.0.2.2 1.1.1.1
PING 1.1.1.1: 56 data bytes, press CTRL_C to break
```

Reply from 1.1.1.1: bytes=56 Sequence=1 ttl=254 time=1 ms

Reply from 1.1.1.1: bytes=56 Sequence=2 ttl=254 time=1 ms

Reply from 1.1.1.1: bytes=56 Sequence=3 ttl=254 time=1 ms

Reply from 1.1.1.1: bytes=56 Sequence=4 ttl=254 time=1 ms

Reply from 1.1.1.1: bytes=56 Sequence=5 ttl=254 time=1 ms

--- 1.1.1.1 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

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

[R3]ping 11.11.11.11

PING 11.11.11.11: 56 data bytes, press CTRL_C to break

Request time out

Request time out

Request time out

Request time out

Request time out

--- 11.11.11.11 ping statistics ---

5 packet(s) transmitted

0 packet(s) received

100.00% packet loss

[R3]ping -a 10.0.3.3 11.11.11.11

PING 11.11.11.11: 56 data bytes, press CTRL_C to break

Reply from 11.11.11.11: bytes=56 Sequence=1 ttl=254 time=1 ms

Reply from 11.11.11.11: bytes=56 Sequence=2 ttl=254 time=1 ms

Reply from 11.11.11.11: bytes=56 Sequence=3 ttl=254 time=1 ms

Reply from 11.11.11.11: bytes=56 Sequence=4 ttl=254 time=1 ms

Reply from 11.11.11.11: bytes=56 Sequence=5 ttl=254 time=1 ms

--- 11.11.11.11 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

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

When you directly test the connectivity between R2 and 11.11.11.11 and between R3 and 11.11.11.11, the connectivity cannot be implemented. Perform the ping operation with the source IP address specified. The connectivity is implemented.

The cause is that R2 directly sends data packets with the source IP address of 10.0.20.2 to the firewall and this IP address is not within the NAT address range, so does R3.

[FW]display nat-policy all

Total:2

RULE ID	RULE NAME	STATE	ACTION	HITTED

0	default	enable	no-nat	0
1	policy_nat_1	enable	nat	2

[FW]display nat-policy rule policy_nat_1

(2 times matched)

```
rule name policy_nat_1
  source-zone trust
  destination-zone untrust
  source-address 10.0.2.0 mask 255.255.255.0
  source-address 10.0.3.0 mask 255.255.255.0
  action nat address-group group1
```

Step 5 Configure the NAT server and NAT based on source IP addresses to deliver services provided by the server.

Set the public IP address of the NAT server to 1.1.1.254, Telnet port number to 2323, and FTP port number to 2121.

```
[FW]nat server policy_natserver_1 protocol tcp global 1.1.1.254 2323 inside 10.0.4.4 telnet no-reverse
```

```
[FW]nat server policy_natserver_2 protocol tcp global 1.1.1.254 2121 inside 10.0.4.4 ftp no-reverse
```

```
[FW]display nat server
```

Server in private network information:

name	: policy_natserver_1		
zone	: ---		
interface	: ---		
global-start-addr	: 1.1.1.254	global-end-addr	: ---
inside-start-addr	: 10.0.4.4	inside-end-addr	: ---
global-start-port	: 2323	global-end-port	: ---
insideport	: 23(telnet)		
globalvpn	: public	insidevpn	: public
protocol	: tcp	vrrp	: ---
no-reverse	: yes		

```
name          : policy_natserver_2
zone          : ---
interface     : ---
global-start-addr : 1.1.1.254      global-end-addr  : ---
inside-start-addr : 10.0.4.4      inside-end-addr  : ---
global-start-port : 2121          global-end-port   : ---
insideport     : 21(ftp)
globalvpn      : public           insidevpn        : public
protocol       : tcp              vrrp             : ---
no-reverse     : yes
```

Total 2 NAT servers

Enable Telnet and FTP services on R4.

```
[R4]telnet server enable
[R4]ftp server enable
[R4]user-interface vty 0 4
[R4-ui-vty0-4]authentication-mode aaa
[R4-ui-vty0-4]protocol inbound telnet
[R4-ui-vty0-4]quit
[R4]aaa
[R4-aaa]local-user test password irreversible-cipher Admin@123
[R4-aaa]local-user test service telnet ftp
[R4-aaa]local-user test ftp-directory flash:/
[R4-aaa]local-user test privilege level 3
[R4-aaa]quit
```

You need to configure NAT Application Level Gateway (NAT ALG) during address

translation because FTP is a multi-channel protocol.

Configure NAT ALG in the DMZ and the untrust zone to so that the NAT server can provide FTP services for external users.

```
[FW]firewall interzone dmz untrust
```

```
[FW-interzone-dmz-untrust]detect ftp
```

Test the configuration results on R1.

```
<R1>telnet 1.1.1.254 2323
```

Press CTRL_] to quit telnet mode

Trying 1.1.1.254 ...

Connected to 1.1.1.254 ...

Login authentication

Username:test

Password:

User last login information:

Access Type: Telnet

IP-Address : 1.1.1.1

Time : 2016-09-25 07:45:45+00:00

```
<R4>quit
```

```
<R1>ftp 1.1.1.254 2121
```

Trying 1.1.1.254 ...

Press CTRL+K to abort

Connected to 1.1.1.254.

220 FTP service ready.

User(1.1.1.254:(none)):test

331 Password required for test.

Enter password:

230 User logged in.

[R1-ftp]

The untrust zone can access Telnet and FTP services provided by the DMZ.

Device Configuration

```
<S1>display current-configuration
```

```
!Software Version V200R008C00SPC500
```

```
#
```

```
sysname S1
```

```
#
```

```
vlan batch 11 to 13
```

```
#
```

```
interface GigabitEthernet0/0/1
```

```
port link-type access
```

```
port default vlan 11
```

```
#
```

```
interface GigabitEthernet0/0/2
```

```
port link-type access
```

```
port default vlan 12
```

```
#
```

```
interface GigabitEthernet0/0/3
```

```
port link-type access
port default vlan 12
#
interface GigabitEthernet0/0/4
port link-type access
port default vlan 13
#
interface GigabitEthernet0/0/21
port link-type access
port default vlan 11
#
interface GigabitEthernet0/0/22
port link-type access
port default vlan 12
#
interface GigabitEthernet0/0/23
port link-type access
port default vlan 13
#
return

<R1> display current-configuration
[V200R007C00SPC600]
#
sysname R1
#
interface GigabitEthernet0/0/1
ip address 1.1.1.1 255.255.255.0
#
```

```
interface LoopBack0  
  ip address 11.11.11.11 255.255.255.0
```

```
#
```

```
return
```

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

```
[V200R007C00SPC600]
```

```
#
```

```
sysname R2
```

```
#
```

```
interface GigabitEthernet0/0/1
```

```
  ip address 10.0.20.2 255.255.255.0
```

```
#
```

```
interface LoopBack0
```

```
  ip address 10.0.2.2 255.255.255.0
```

```
#
```

```
ip route-static 0.0.0.0 0.0.0.0 10.0.20.254
```

```
#
```

```
return
```

```
<R3> display current-configuration
```

```
[V200R007C00SPC600]
```

```
#
```

```
sysname R3
```

```
#
```

```
interface GigabitEthernet0/0/1
```

```
  ip address 10.0.20.3 255.255.255.0
```

```
#
```

```
interface LoopBack0
```

```
ip address 10.0.3.3 255.255.255.0
```

```
#
```

```
ip route-static 0.0.0.0 0.0.0.0 10.0.20.254
```

```
#
```

```
return
```

```
<R4> display current-configuration
```

```
[V200R007C00SPC600]
```

```
#
```

```
sysname R4
```

```
#
```

```
aaa
```

```
local-user test password irreversible-cipher Admin@123
```

```
local-user test privilege level 3
```

```
local-user test ftp-directory flash:/
```

```
local-user test service-type telnet ftp
```

```
#
```

```
interface GigabitEthernet0/0/1
```

```
ip address 10.0.40.4 255.255.255.0
```

```
#
```

```
interface LoopBack0
```

```
ip address 10.0.4.4 255.255.255.0
```

```
#
```

```
ftp server enable
```

```
#
```

```
telnet server enable
```

```
#
```

```
ip route-static 0.0.0.0 0.0.0.0 10.0.40.254
```

```
#
```

```
user-interface vty 0 4
  authentication-mode aaa
  protocol inbound telnet
#
return

<FW> display current-configuration
#
nat server policy_natserver_1 protocol tcp global 1.1.1.254 2323 inside 10.0.4.4 telnet no-reverse
nat server policy_natserver_2 protocol tcp global 1.1.1.254 2121 inside 10.0.4.4 ftp no-reverse
#
sysname FW
#
interface GigabitEthernet1/0/0
  ip address 10.0.10.254 255.255.255.0
  ip address 1.1.1.254 255.255.255.0 sub
#
interface GigabitEthernet1/0/1
  ip address 10.0.20.254 255.255.255.0
#
interface GigabitEthernet1/0/2
  ip address 10.0.40.254 255.255.255.0
#
firewall zone local
  set priority 100
#
firewall zone trust
  set priority 85
```

```
add interface GigabitEthernet1/0/1
#
firewall zone untrust
set priority 5
add interface GigabitEthernet1/0/0
#
firewall zone dmz
set priority 50
add interface GigabitEthernet1/0/2
#
firewall interzone dmz untrust
detect ftp
#
ip route-static 0.0.0.0 0.0.0.0 1.1.1.1
ip route-static 10.0.2.0 255.255.255.0 10.0.20.2
ip route-static 10.0.3.0 255.255.255.0 10.0.20.3
ip route-static 10.0.4.0 255.255.255.0 10.0.40.4
#
nat address-group group1
section 0 1.1.1.254 1.1.1.254
#
security-policy
rule name policy_sec_1
source-zone trust
destination-zone untrust
source-address 10.0.2.0 mask 255.255.255.0
source-address 10.0.3.0 mask 255.255.255.0
action permit
rule name policy_sec_2
```

```
source-zone untrust
destination-zone dmz
destination-address 10.0.4.4 mask 255.255.255.255
service ftp
service telnet
action permit
#
nat-policy
rule name policy_nat_1
source-zone trust
destination-zone untrust
source-address 10.0.2.0 mask 255.255.255.0
source-address 10.0.3.0 mask 255.255.255.0
action nat address-group group1
#
return
```

Chapter 5 VRRP Configuration

Lab 5-1 VRRP Configuration

Learning Objectives

The objectives of this lab are to learn and understand:

- How to configure a Virtual Router Redundancy Protocol (VRRP) group and the virtual IP address
- How to configure the VRRP priority
- How to check the VRRP configuration
- How to configure VRRP to monitor the uplink
- How to configure load balancing for multiple VRRP groups

Topology




Figure 5-1 VRRP topology

Scenario

R1 functions as the gateway for a LAN and an external network. R1 connects to R2 and R3 through S5, and R2, R3, R4, and R5 connect to a LAN through S1. VRRPv2 needs to be enabled on interfaces of R2 and R3 connecting to S1 to implement first-hop redundancy. R2 is the master router and R3 is the backup router. No extra configurations are required for switches. The switches only transparently forward packets.

Tasks

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

Configure IP addresses for all routers.

<Huawei>system-view

Enter system view, return user view with Ctrl+Z.

[Huawei]sysname R1

[R1]interface loopback 0

[R1-LoopBack0]ip address 1.1.1.1 32

[R1-LoopBack0]quit

[R1]interface GigabitEthernet 0/0/0

[R1-GigabitEthernet0/0/0]ip address 10.0.123.1 24

<Huawei>system-view

Enter system view, return user view with Ctrl+Z.

[Huawei]sysname R2

[R2]interface GigabitEthernet 0/0/0

[R2-GigabitEthernet0/0/0]ip address 10.0.123.2 24

[R2-GigabitEthernet0/0/0]quit

[R2]interface GigabitEthernet 0/0/1

[R2-GigabitEthernet0/0/1]ip address 192.168.1.2 24

<Huawei>system-view

Enter system view, return user view with Ctrl+Z.

[Huawei]sysname R3

[R3]interface GigabitEthernet 0/0/0

[R3-GigabitEthernet0/0/0]ip address 10.0.123.3 24

[R3-GigabitEthernet0/0/0]quit

[R3]interface GigabitEthernet 0/0/1

[R3-GigabitEthernet0/0/1]ip address 192.168.1.3 24

<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 192.168.1.4 24

<Huawei>system-view

Enter system view, return user view with Ctrl+Z.

[Huawei]sysname R5

[R5]interface GigabitEthernet 0/0/1

[R5-GigabitEthernet0/0/1]ip address 192.168.1.5 24

After the configuration is complete, test the connectivity between R1 and R2 and between R1 and R3.

[R1]ping 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=1 ms

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

Reply from 10.0.123.2: bytes=56 Sequence=3 ttl=255 time=1 ms

Reply from 10.0.123.2: bytes=56 Sequence=4 ttl=255 time=1 ms

Reply from 10.0.123.2: bytes=56 Sequence=5 ttl=255 time=1 ms

--- 10.0.123.2 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

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

[R1]ping 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=1 ms

Reply from 10.0.123.3: bytes=56 Sequence=2 ttl=255 time=1 ms

Reply from 10.0.123.3: bytes=56 Sequence=3 ttl=255 time=1 ms

Reply from 10.0.123.3: bytes=56 Sequence=4 ttl=255 time=1 ms

Reply from 10.0.123.3: bytes=56 Sequence=5 ttl=255 time=1 ms

--- 10.0.123.3 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

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

Test the connectivity between R2, R3, R4, and R5. R2 is used as an example.

[R2]ping 192.168.1.3

PING 192.168.1.3: 56 data bytes, press CTRL_C to break

Reply from 192.168.1.3: bytes=56 Sequence=1 ttl=255 time=27 ms

Reply from 192.168.1.3: bytes=56 Sequence=2 ttl=255 time=1 ms

Reply from 192.168.1.3: bytes=56 Sequence=3 ttl=255 time=1 ms

Reply from 192.168.1.3: bytes=56 Sequence=4 ttl=255 time=1 ms

Reply from 192.168.1.3: bytes=56 Sequence=5 ttl=255 time=1 ms

--- 192.168.1.3 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

round-trip min/avg/max = 1/6/27 ms

[R2]ping 192.168.1.4

PING 192.168.1.4: 56 data bytes, press CTRL_C to break

Reply from 192.168.1.4: bytes=56 Sequence=1 ttl=255 time=1 ms

Reply from 192.168.1.4: bytes=56 Sequence=2 ttl=255 time=1 ms

Reply from 192.168.1.4: bytes=56 Sequence=3 ttl=255 time=1 ms

Reply from 192.168.1.4: bytes=56 Sequence=4 ttl=255 time=1 ms

Reply from 192.168.1.4: bytes=56 Sequence=5 ttl=255 time=1 ms

--- 192.168.1.4 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

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

[R2]ping 192.168.1.5

PING 192.168.1.5: 56 data bytes, press CTRL_C to break

Reply from 192.168.1.5: bytes=56 Sequence=1 ttl=255 time=1 ms

Reply from 192.168.1.5: bytes=56 Sequence=2 ttl=255 time=1 ms

Reply from 192.168.1.5: bytes=56 Sequence=3 ttl=255 time=1 ms

Reply from 192.168.1.5: bytes=56 Sequence=4 ttl=255 time=1 ms

Reply from 192.168.1.5: bytes=56 Sequence=5 ttl=255 time=1 ms

--- 192.168.1.5 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

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

Step 2 **Configure OSPF and static routes.**

The loopback interface on R1 and connected interfaces on R1, R2, and R3 run in OSPF area 0. Routes of interfaces on R2 and R3 connecting to S1 are advertised to OSPF, but no OSPF neighbor relationships. The silent mode is therefore used.

To simulate PCs, R4 and R5 use default static routes pointing to 192.168.1.1 (VRRP virtual IP address).

Enable R1 to learn routes to 192.168.1.0, and enable R2 and R3 to learn routes to 1.1.1.1.

```
[R1]ospf 1
```

```
[R1-ospf-1]area 0
```

```
[R1-ospf-1-area-0.0.0.0]network 10.0.123.0 0.0.0.255
```

```
[R1-ospf-1-area-0.0.0.0]network 1.1.1.1 0.0.0.0
```

```
[R2]ospf 1
```

```
[R2-ospf-1]silent-interface GigabitEthernet 0/0/1
```

```
[R2-ospf-1]area 0
```

```
[R2-ospf-1-area-0.0.0.0]network 10.0.123.0 0.0.0.255
```

```
[R2-ospf-1-area-0.0.0.0]network 192.168.1.0 0.0.0.255
```

```
[R3]ospf 1
```

```
[R3-ospf-1]silent-interface GigabitEthernet 0/0/1
```

```
[R3-ospf-1]area 0
```

```
[R3-ospf-1-area-0.0.0.0]network 10.0.123.0 0.0.0.255
```

```
[R3-ospf-1-area-0.0.0.0]network 192.168.1.0 0.0.0.255
```

```
[R4]ip route-static 0.0.0.0 0.0.0.0 192.168.1.1
```

```
[R5]ip route-static 0.0.0.0 0.0.0.0 192.168.1.1
```

After the configuration is complete, check the routing table of R1, R2, and R4.

[R1]display ip routing-table

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

Routing Tables: Public

Destinations : 9 Routes : 10

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
1.1.1.1/32	Direct	0	0	D	127.0.0.1	LoopBack0
10.0.123.0/24	Direct	0	0	D	10.0.123.1	GigabitEthernet0/0/0
10.0.123.1/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/0
10.0.123.255/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/0
127.0.0.0/8	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
127.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
192.168.1.0/24	OSPF	10	2	D	10.0.123.3	GigabitEthernet0/0/0
	OSPF	10	2	D	10.0.123.2	GigabitEthernet0/0/0
255.255.255.255/32	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 : 12 Routes : 12

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

1.1.1.1/32	OSPF	10	1	D	10.0.123.1	GigabitEthernet0/0/0
10.0.0.2/32	Direct	0	0	D	127.0.0.1	LoopBack0
10.0.123.0/24	Direct	0	0	D	10.0.123.2	GigabitEthernet0/0/0
10.0.123.2/32	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
192.168.1.0/24	Direct	0	0	D	192.168.1.2	GigabitEthernet0/0/1
192.168.1.2/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/1
192.168.1.255/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/1
255.255.255.255/32	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 : 9 Routes : 9

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
0.0.0.0/0	Static	60	0	RD	192.168.1.1	GigabitEthernet0/0/1
10.0.0.4/32	Direct	0	0	D	127.0.0.1	LoopBack0
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
192.168.1.0/24	Direct	0	0	D	192.168.1.4	GigabitEthernet0/0/1
192.168.1.4/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/1
192.168.1.255/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/1

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

The preceding output shows that R1 can learn routes to 192.168.1.0/24, R2 can learn routes to 1.1.1.1/32, and R4 has a default static route to 192.168.1.1.

Step 3 **Configure a VRRP group and the virtual IP address.**

Enable VRRP on interfaces of R2 and R3, and configure the VRID and virtual IP address.

```
[R2]interface GigabitEthernet 0/0/1
```

```
[R2-GigabitEthernet0/0/1]vrrp vrid 1 virtual-ip 192.168.1.1
```

R2 is configured first and becomes the master router if there is no other member in the VRRP group after a period of time.

```
[R3]interface GigabitEthernet 0/0/1
```

```
[R3-GigabitEthernet0/0/1]vrrp vrid 1 virtual-ip 192.168.1.1
```

After the configuration is complete, check the VRRP status on R2 and R3.

```
[R2]display vrrp
```

```
GigabitEthernet0/0/1 | Virtual Router 1
```

```
State : Master
```

```
Virtual IP : 192.168.1.1
```

```
Master IP : 192.168.1.2
```

```
PriorityRun : 100
```

```
PriorityConfig : 100
```

```
MasterPriority : 100
```

```
Preempt : YES Delay Time : 0 s
```

```
TimerRun : 1 s
```

TimerConfig : 1 s

Auth type : NONE

Virtual MAC : 0000-5e00-0101

Check TTL : YES

Config type : normal-vrrp

Backup-forward : disabled

Create time : 2016-07-22 18:00:03

Last change time : 2016-07-22 18:00:07

[R3]display vrrp

GigabitEthernet0/0/1 | Virtual Router 1

State : Backup

Virtual IP : 192.168.1.1

Master IP : 192.168.1.2

PriorityRun : 100

PriorityConfig : 100

MasterPriority : 100

Preempt : YES Delay Time : 0 s

TimerRun : 1 s

TimerConfig : 1 s

Auth type : NONE

Virtual MAC : 0000-5e00-0101

Check TTL : YES

Config type : normal-vrrp

Backup-forward : disabled

Create time : 2016-07-22 18:03:16

Last change time : 2016-07-22 18:03:16

R2 is selected as the master router and R3 as the backup router. The priorities of master and slave routers are both 100. When R3 is started first, it becomes the master router, which is not expected.

Configure the VRRP priority and verify the active/standby switchover.

Configure VRRP priorities on R2 and R3. A greater priority value indicates a higher priority. Set VRRP priorities of R2 and R3 to 120 and 110, respectively.

```
[R2]interface GigabitEthernet 0/0/1
```

```
[R2-GigabitEthernet0/0/1]vrrp vrid 1 priority 120
```

```
[R3]interface GigabitEthernet 0/0/1
```

```
[R3-GigabitEthernet0/0/1]vrrp vrid 1 priority 110
```

Check the configuration after priorities are changed.

```
[R2]display vrrp
```

```
GigabitEthernet0/0/1 | Virtual Router 1
```

```
State : Master
```

```
Virtual IP : 192.168.1.1
```

```
Master IP : 192.168.1.2
```

```
PriorityRun : 120
```

```
PriorityConfig : 120
```

```
MasterPriority : 120
```

```
Preempt : YES    Delay Time : 0 s
```

```
TimerRun : 1 s
```

```
TimerConfig : 1 s
```

```
Auth type : NONE
```

```
Virtual MAC : 0000-5e00-0101
```

```
Check TTL : YES
```

Config type : normal-vrrp

Backup-forward : disabled

Create time : 2016-07-22 18:00:03

Last change time : 2016-07-22 18:00:07

[R3]display vrrp

GigabitEthernet0/0/1 | Virtual Router 1

State : Backup

Virtual IP : 192.168.1.1

Master IP : 192.168.1.2

PriorityRun : 110

PriorityConfig : 110

MasterPriority : 120

Preempt : YES Delay Time : 0 s

TimerRun : 1 s

TimerConfig : 1 s

Auth type : NONE

Virtual MAC : 0000-5e00-0101

Check TTL : YES

Config type : normal-vrrp

Backup-forward : disabled

Create time : 2016-07-22 18:03:16

Last change time : 2016-07-22 18:03:16

The preceding output shows that priorities of R2 and R3 have been changed successfully. By default, VRRP preemption is enabled. When the priority of R3 is changed to be higher, an active/standby switchover will be triggered.

Test the connectivity between R4 and R1.

[R4]ping 1.1.1.1

PING 1.1.1.1: 56 data bytes, press CTRL_C to break

Reply from 1.1.1.1: bytes=56 Sequence=1 ttl=254 time=57 ms

Reply from 1.1.1.1: bytes=56 Sequence=2 ttl=254 time=1 ms

Reply from 1.1.1.1: bytes=56 Sequence=3 ttl=254 time=1 ms

Reply from 1.1.1.1: bytes=56 Sequence=4 ttl=254 time=1 ms

Reply from 1.1.1.1: bytes=56 Sequence=5 ttl=254 time=1 ms

--- 1.1.1.1 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

round-trip min/avg/max = 1/12/57 ms

The preceding output shows that the virtual gateway works properly and can forward data of the LAN where R4 is located to R1. Normally, the master router forwards data, so traffic passes through R2. To verify the switching status, perform the ping operation from R4 to R1 for a long time and shut down the interface of R2 connected to S1.

[R2]interface GigabitEthernet 0/0/1

[R2-GigabitEthernet0/0/1]shutdown

R4 discards two data packets during the switchover, and subsequent data is forwarded normally.

[R4]ping -c 1000 1.1.1.1

PING 1.1.1.1: 56 data bytes, press CTRL_C to break

Reply from 1.1.1.1: bytes=56 Sequence=1 ttl=254 time=1 ms

Reply from 1.1.1.1: bytes=56 Sequence=2 ttl=254 time=1 ms

Reply from 1.1.1.1: bytes=56 Sequence=3 ttl=254 time=1 ms

```

Reply from 1.1.1.1: bytes=56 Sequence=4 ttl=254 time=1 ms
Reply from 1.1.1.1: bytes=56 Sequence=5 ttl=254 time=1 ms
Reply from 1.1.1.1: bytes=56 Sequence=6 ttl=254 time=1 ms
Request time out
Request time out
Reply from 1.1.1.1: bytes=56 Sequence=9 ttl=254 time=1 ms
Reply from 1.1.1.1: bytes=56 Sequence=10 ttl=254 time=1 ms
Reply from 1.1.1.1: bytes=56 Sequence=11 ttl=254 time=1 ms
Reply from 1.1.1.1: bytes=56 Sequence=12 ttl=254 time=1 ms
Reply from 1.1.1.1: bytes=56 Sequence=13 ttl=254 time=1 ms
Reply from 1.1.1.1: bytes=56 Sequence=14 ttl=254 time=1 ms
Reply from 1.1.1.1: bytes=56 Sequence=15 ttl=254 time=1 ms
Reply from 1.1.1.1: bytes=56 Sequence=16 ttl=254 time=1 ms
Reply from 1.1.1.1: bytes=56 Sequence=17 ttl=254 time=1 ms
Reply from 1.1.1.1: bytes=56 Sequence=18 ttl=254 time=1 ms
Reply from 1.1.1.1: bytes=56 Sequence=19 ttl=254 time=1 ms
Reply from 1.1.1.1: bytes=56 Sequence=20 ttl=254 time=1 ms

```

```

--- 1.1.1.1 ping statistics ---

```

```

20 packet(s) transmitted
18 packet(s) received
10.00% packet loss
round-trip min/avg/max = 1/1/1 ms

```

R3 becomes the master router after the switchover.

```

[R3]display vrrp

GigabitEthernet0/0/1 | Virtual Router 1

State : Master

```

Virtual IP : 192.168.1.1

Master IP : 192.168.1.3

PriorityRun : 110

PriorityConfig : 110

MasterPriority : 110

Preempt : YES Delay Time : 0 s

TimerRun : 1 s

TimerConfig : 1 s

Auth type : NONE

Virtual MAC : 0000-5e00-0101

Check TTL : YES

Config type : normal-vrrp

Backup-forward : disabled

Create time : 2016-07-22 18:03:16

Last change time : 2016-07-22 18:29:41

Configure VRRP to monitor the uplink.

The VRRP active/standby switchover is implemented by listening to Advertisement packets. If the backup router cannot listen to messages of the master router or has a higher priority, the backup router preempts to be the master router (no preemption delay by default).

If the fault occurs on the uplink link, the active/standby switchover is not performed. In this case, all Internet access traffic cannot be forwarded after reaching R2. VRRP is enabled to monitor the uplink. When the uplink fails, R2 automatically reduces its priority. R3 preempts to be the master router, and traffic is switched to the backup router and backup uplink.

Before configuring VRRP to monitor the uplink, restore the link that is shut down.

Configure VRRP to monitor the uplink interface and set the value by which the

priority decreases to 30. That is, when the link fails, the priority of R2 becomes 90, which is lower than the priority of R3 (110).

```
[R2]interface GigabitEthernet 0/0/1
```

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

```
[R2-GigabitEthernet0/0/1]vrrp vrid 1 track interface GigabitEthernet 0/0/0 reduced 30
```

Check the configuration.

```
[R2]display vrrp
```

GigabitEthernet0/0/1 | Virtual Router 1

State : Master

Virtual IP : 192.168.1.1

Master IP : 192.168.1.2

PriorityRun : 120

PriorityConfig : 120

MasterPriority : 120

Preempt : YES Delay Time : 0 s

TimerRun : 1 s

TimerConfig : 1 s

Auth type : NONE

Virtual MAC : 0000-5e00-0101

Check TTL : YES

Config type : normal-vrrp

Backup-forward : disabled

Track IF : GigabitEthernet0/0/0 Priority reduced : 30

IF state : UP

Create time : 2016-07-25 17:14:56 UTC-08:00

Last change time : 2016-07-25 17:32:27 UTC-08:00

Perform the ping operation on R4 for a long time and shut down the uplink interface on R2.

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

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

```
[R2]display vrrp
```

GigabitEthernet0/0/1 | Virtual Router 1

State : Backup

Virtual IP : 192.168.1.1

Master IP : 192.168.1.3

PriorityRun : 90

PriorityConfig : 120

MasterPriority : 110

Preempt : YES Delay Time : 0 s

TimerRun : 1 s

TimerConfig : 1 s

Auth type : NONE

Virtual MAC : 0000-5e00-0101

Check TTL : YES

Config type : normal-vrrp

Backup-forward : disabled

Track IF : GigabitEthernet0/0/0 Priority reduced : 30

IF state : DOWN

Create time : 2016-07-25 17:14:56 UTC-08:00

Last change time : 2016-07-25 19:57:46 UTC-08:00

```
[R3]display vrrp
```

GigabitEthernet0/0/1 | Virtual Router 1

State : Master

Virtual IP : 192.168.1.1
Master IP : 192.168.1.3
PriorityRun : 110
PriorityConfig : 110
MasterPriority : 110
Preempt : YES Delay Time : 0 s
TimerRun : 1 s
TimerConfig : 1 s
Auth type : NONE
Virtual MAC : 0000-5e00-0101
Check TTL : YES
Config type : normal-vrrp
Backup-forward : disabled
Create time : 2016-07-25 17:20:00 UTC-08:00
Last change time : 2016-07-25 19:56:24 UTC-08:00

R3 becomes the master router and traffic is successfully switched to R3.

Restore the uplink and priority of R2. R2 preempts to be the master router again. During preemption, few packets are discarded on R4. This is because OSPF routes are not converged rapidly. For details about route convergence acceleration, see the OSPF experiment.

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

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

```
[R2]display vrrp
```

```
GigabitEthernet0/0/1 | Virtual Router 1
```

```
State : Master
```

```
Virtual IP : 192.168.1.1
```

Master IP : 192.168.1.2

PriorityRun : 120

PriorityConfig : 120

MasterPriority : 120

Preempt : YES Delay Time : 0 s

TimerRun : 1 s

TimerConfig : 1 s

Auth type : NONE

Virtual MAC : 0000-5e00-0101

Check TTL : YES

Config type : normal-vrrp

Backup-forward : disabled

Track IF : GigabitEthernet0/0/0 Priority reduced : 30

IF state : UP

Create time : 2016-07-25 17:14:56 UTC-08:00

Last change time : 2016-07-25 20:04:40 UTC-08:00

When the interface goes Up, the OSPF neighbor relationship needs to be reestablished on the uplink interface of R2. If OSPF fast convergence is not configured, data cannot be forwarded for several seconds. You are advised to set the preemption delay to be longer than the OSPF convergence time during the switchback.

```
[R2]interface GigabitEthernet 0/0/1
```

```
[R2-GigabitEthernet0/0/1]vrrp vrid 1 preempt-mode timer delay 10
```

Check the VRRP configurations again. You can see that the preemption delay has been configured successfully.

```
[R2]display vrrp
```

GigabitEthernet0/0/1 | Virtual Router 1

State : Master

Virtual IP : 192.168.1.1

Master IP : 192.168.1.2

PriorityRun : 120

PriorityConfig : 120

MasterPriority : 120

Preempt : YES Delay Time : 10 s

TimerRun : 1 s

TimerConfig : 1 s

Auth type : NONE

Virtual MAC : 0000-5e00-0101

Check TTL : YES

Config type : normal-vrrp

Backup-forward : disabled

Track IF : GigabitEthernet0/0/0 Priority reduced : 30

IF state : UP

Create time : 2016-07-25 17:14:56 UTC-08:00

Last change time : 2016-07-25 20:04:40 UTC-08:00

Step 4 **Configure load balancing of multiple VRRP groups.**

Normally, the master device forwards all traffic, and the backup device is idle.

To implement dual-gateway load balancing, configure multiple VRRP groups.

Configure VRRP group 1 on R2 and R3, set the virtual IP address to 192.168.1.1, and configure R2 as the master device. Configure VRRP group 2, set the virtual address to 192.168.1.254, and configure R3 as the master device. Configure the default gateway address pointing to 192.168.1.1 for R4 and the default gateway address pointing to 192.168.1.254 for R5. Internet access traffic can be load balanced to two

gateways.

The configuration is as follows.

```
[R2]interface GigabitEthernet 0/0/1
```

```
[R2-GigabitEthernet0/0/1]vrrp vrid 2 virtual-ip 192.168.1.254
```

```
[R2-GigabitEthernet0/0/1]vrrp vrid 2 priority 110
```

```
[R3]interface GigabitEthernet 0/0/1
```

```
[R3-GigabitEthernet0/0/1]vrrp vrid 2 virtual-ip 192.168.1.254
```

```
[R3-GigabitEthernet0/0/1]vrrp vrid 2 priority 120
```

```
[R3-GigabitEthernet0/0/1]vrrp vrid 2 track interface GigabitEthernet0/0/0 reduced 30
```

```
[R5]undo ip route-static 0.0.0.0 0.0.0.0 192.168.1.1
```

```
[R5]ip route-static 0.0.0.0 0.0.0.0 192.168.1.254
```

Check load balancing of two VRRP groups on R2 and R3.

```
<R2>display vrrp
```

```
GigabitEthernet0/0/1 | Virtual Router 1
```

```
State : Master
```

```
Virtual IP : 192.168.1.1
```

```
Master IP : 192.168.1.2
```

```
PriorityRun : 120
```

```
PriorityConfig : 120
```

```
MasterPriority : 120
```

```
Preempt : YES    Delay Time : 10 s
```

```
TimerRun : 1 s
```

```
TimerConfig : 1 s
```

```
Auth type : NONE
```

```
Virtual MAC : 0000-5e00-0101
```

Check TTL : YES

Config type : normal-vrrp

Backup-forward : disabled

Track IF : GigabitEthernet0/0/0 Priority reduced : 30

IF state : UP

Create time : 2016-07-25 17:14:56 UTC-08:00

Last change time : 2016-07-25 20:04:40 UTC-08:00

GigabitEthernet0/0/1 | Virtual Router 2

State : Backup

Virtual IP : 192.168.1.254

Master IP : 192.168.1.3

PriorityRun : 110

PriorityConfig : 110

MasterPriority : 120

Preempt : YES Delay Time : 0 s

TimerRun : 1 s

TimerConfig : 1 s

Auth type : NONE

Virtual MAC : 0000-5e00-0102

Check TTL : YES

Config type : normal-vrrp

Backup-forward : disabled

Create time : 2016-07-25 17:15:54 UTC-08:00

Last change time : 2016-07-25 17:20:30 UTC-08:00

<R3>display vrrp

GigabitEthernet0/0/1 | Virtual Router 1

State : Backup

Virtual IP : 192.168.1.1
 Master IP : 192.168.1.2
 PriorityRun : 110
 PriorityConfig : 110
 MasterPriority : 120
 Preempt : YES Delay Time : 0 s
 TimerRun : 1 s
 TimerConfig : 1 s
 Auth type : NONE
 Virtual MAC : 0000-5e00-0101
 Check TTL : YES
 Config type : normal-vrrp
 Backup-forward : disabled
 Create time : 2016-07-25 17:20:00 UTC-08:00
 Last change time : 2016-07-25 20:03:15 UTC-08:00

GigabitEthernet0/0/1 | Virtual Router 2

State : Master
 Virtual IP : 192.168.1.254
 Master IP : 192.168.1.3
 PriorityRun : 120
 PriorityConfig : 120
 MasterPriority : 120
 Preempt : YES Delay Time : 0 s
 TimerRun : 1 s
 TimerConfig : 1 s
 Auth type : NONE
 Virtual MAC : 0000-5e00-0102
 Check TTL : YES

Config type : normal-vrrp

Backup-forward : disabled

Track IF : GigabitEthernet0/0/0 Priority reduced : 30

IF state : UP

Create time : 2016-07-25 17:20:14 UTC-08:00

Last change time : 2016-07-25 17:20:23 UTC-08:00

Perform the tracert operation to check the gateways that process data destined for the two default routes. You can see that data sent by R4 is forwarded by the master device in VRRP group 1 and data sent by R5 is forwarded by the master device in VRRP group 2.

Enable R1 to send ICMP Port Unreachable packets.

[R1]icmp port-unreachable send

<R4>tracert 1.1.1.1

tracert to 1.1.1.1(1.1.1.1), max hops: 30 ,packet length: 40,press CTRL_C to break

1 192.168.1.2 80 ms 40 ms 40 ms

2 10.0.123.1 100 ms 70 ms 70 ms

<R5>tracert 1.1.1.1

tracert to 1.1.1.1(1.1.1.1), max hops: 30 ,packet length: 40,press CTRL_C to break

1 192.168.1.3 50 ms 30 ms 50 ms

2 10.0.123.1 60 ms 90 ms 60 ms

Check the switchover when the uplink fails.

[R2]interface GigabitEthernet 0/0/0

[R2-GigabitEthernet0/0/0]shutdown

<R4>tracert 1.1.1.1

tracert to 1.1.1.1(1.1.1.1), max hops: 30 ,packet length: 40,press CTRL_C t
o break

1 192.168.1.3 50 ms 40 ms 50 ms

2 10.0.123.1 70 ms 80 ms 50 ms

<R5>tracert 1.1.1.1

tracert to 1.1.1.1(1.1.1.1), max hops: 30 ,packet length: 40,press CTRL_C t
o break

1 192.168.1.3 40 ms 50 ms 40 ms

2 10.0.123.1 70 ms 100 ms 90 ms

Check the status of the two VRRP groups.

<R2>display vrrp

GigabitEthernet0/0/1 | Virtual Router 1

State : Backup

Virtual IP : 192.168.1.1

Master IP : 192.168.1.3

PriorityRun : 90

PriorityConfig : 120

MasterPriority : 110

Preempt : YES Delay Time : 10 s

TimerRun : 1 s

TimerConfig : 1 s

Auth type : NONE

Virtual MAC : 0000-5e00-0101

Check TTL : YES

Config type : normal-vrrp

Backup-forward : disabled

Track IF : GigabitEthernet0/0/0 Priority reduced : 30

IF state : DOWN

Create time : 2016-07-25 17:14:56 UTC-08:00

Last change time : 2016-07-25 20:48:28 UTC-08:00

GigabitEthernet0/0/1 | Virtual Router 2

State : Backup

Virtual IP : 192.168.1.254

Master IP : 192.168.1.3

PriorityRun : 110

PriorityConfig : 110

MasterPriority : 120

Preempt : YES Delay Time : 0 s

TimerRun : 1 s

TimerConfig : 1 s

Auth type : NONE

Virtual MAC : 0000-5e00-0102

Check TTL : YES

Config type : normal-vrrp

Backup-forward : disabled

Create time : 2016-07-25 17:15:54 UTC-08:00

Last change time : 2016-07-25 17:20:30 UTC-08:00

<R3>display vrrp

GigabitEthernet0/0/1 | Virtual Router 1

State : Master

Virtual IP : 192.168.1.1

Master IP : 192.168.1.3

PriorityRun : 110

PriorityConfig : 110

MasterPriority : 110

Preempt : YES Delay Time : 0 s

TimerRun : 1 s

TimerConfig : 1 s

Auth type : NONE

Virtual MAC : 0000-5e00-0101

Check TTL : YES

Config type : normal-vrrp

Backup-forward : disabled

Create time : 2016-07-25 17:20:00 UTC-08:00

Last change time : 2016-07-25 20:46:42 UTC-08:00

GigabitEthernet0/0/1 | Virtual Router 2

State : Master

Virtual IP : 192.168.1.254

Master IP : 192.168.1.3

PriorityRun : 120

PriorityConfig : 120

MasterPriority : 120

Preempt : YES Delay Time : 0 s

TimerRun : 1 s

TimerConfig : 1 s

Auth type : NONE

Virtual MAC : 0000-5e00-0102

Check TTL : YES

Config type : normal-vrrp

Backup-forward : disabled

Track IF : GigabitEthernet0/0/0 Priority reduced : 30

IF state : UP

Create time : 2016-07-25 17:20:14 UTC-08:00

Last change time : 2016-07-25 17:20:23 UTC-08:00

Normally, R2 and R3 load balance traffic. When R2 is faulty, R3 takes over all traffic on R2. In this case, load balancing of two VRRP groups is configured.

Device Configuration

<R1>display current-configuration

```
#
```

```
sysname R1
```

```
#
```

```
interface GigabitEthernet0/0/0
  ip address 10.0.123.1 255.255.255.0
#
interface LoopBack0
  ip address 1.1.1.1 255.255.255.255
#
ospf 1
  area 0.0.0.0
    network 1.1.1.1 0.0.0.0
    network 10.0.123.0 0.0.0.255
#
return
```

<R2> display current-configuration

```
#
sysname R2
#
interface GigabitEthernet0/0/0
  shutdown
  ip address 10.0.123.2 255.255.255.0
#
interface GigabitEthernet0/0/1
  ip address 192.168.1.2 255.255.255.0
  vrrp vrid 1 virtual-ip 192.168.1.1
  vrrp vrid 1 priority 120
  vrrp vrid 1 preempt-mode timer delay 10
  vrrp vrid 1 track interface GigabitEthernet0/0/0 reduced 30
  vrrp vrid 2 virtual-ip 192.168.1.254
```

```

vrp vrid 2 priority 110

#

ospf 1

silent-interface GigabitEthernet0/0/1

area 0.0.0.0

network 10.0.123.0 0.0.0.255

network 192.168.1.0 0.0.0.255

#

return

```

<R3>display current-configuration

```

#

sysname R3

#

interface GigabitEthernet0/0/0

ip address 10.0.123.3 255.255.255.0

#

interface GigabitEthernet0/0/1

ip address 192.168.1.3 255.255.255.0

vrp vrid 1 virtual-ip 192.168.1.1

vrp vrid 1 priority 110

vrp vrid 2 virtual-ip 192.168.1.254

vrp vrid 2 priority 120

vrp vrid 2 track interface GigabitEthernet0/0/0 reduced 30

#

ospf 1

silent-interface GigabitEthernet0/0/1

area 0.0.0.0

```

```
network 10.0.123.0 0.0.0.255
```

```
network 192.168.1.0 0.0.0.255
```

```
#
```

```
return
```

<R4>display current-configuration

```
#
```

```
sysname R4
```

```
#
```

```
interface GigabitEthernet0/0/1
```

```
ip address 192.168.1.4 255.255.255.0
```

```
#
```

```
ip route-static 0.0.0.0 0.0.0.0 192.168.1.1
```

```
#
```

```
return
```

<R5>display current-configuration

```
#
```

```
sysname R5
```

```
#
```

```
interface GigabitEthernet0/0/1
```

```
ip address 192.168.1.5 255.255.255.0
```

```
#
```

```
ip route-static 0.0.0.0 0.0.0.0 192.168.1.254
```

```
#
```

```
return
```

Chapter 6 BFD Configuration

Lab 6-1 Association Between BFD and Static Routes

Learning Objectives

The objectives of this lab are to learn and understand:

- How to associate Bidirectional Forwarding Detection (BFD) with static routes to implement floating routes

Topology




Figure 6-1 Networking of association between BFD and static routes

Scenario

R1 is connected to R2 and R3 through S1 and S2. Devices are connected through static routes, and packets can reach the target network 23.23.23.23/32 through R2 or

R3. R2 is the active next hop, and R3 is the standby next hop. The link is not the direct one, so the interface status does not affect reachability of the static routes. BFD is used to detect reachability of the static routes. When detection fails, the backup static route is used to forward data.

Tasks

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

Configure IP addresses for all routers and check the addresses.

```
<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 10.0.12.1 24
```

```
[R1-GigabitEthernet0/0/1]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]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 9

The number of interface that is DOWN in Physical is 3

The number of interface that is UP in Protocol is 6

The number of interface that is DOWN in Protocol is 6

Interface	IP Address/Mask	Physical	Protocol
Cellular0/0/0	unassigned	down	down
Cellular0/0/1	unassigned	down	down
GigabitEthernet0/0/0	unassigned	up	down
GigabitEthernet0/0/1	10.0.12.1/24	up	up
GigabitEthernet0/0/2	10.0.13.1/24	up	up
GigabitEthernet0/0/3	unassigned	up	down
LoopBack0	10.0.1.1/32	up	up(s)
NULL0	unassigned	up	up(s)
Serial1/0/0	unassigned	up	up
Serial2/0/0	unassigned	up	down
Serial3/0/0	unassigned	up	up
Serial4/0/0	unassigned	down	down

<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 10.0.12.2 24

[R2-GigabitEthernet0/0/1]quit

[R2]interface LoopBack 0

[R2-LoopBack0]ip address 23.23.23.23 32

[R2-LoopBack0]quit

[R2]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 9

The number of interface that is DOWN in Physical is 4

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
Ethernet4/0/0	unassigned	down	down
Ethernet4/0/1	unassigned	down	down
GigabitEthernet0/0/0	unassigned	up	down
GigabitEthernet0/0/1	10.0.12.2/24	up	up
GigabitEthernet0/0/2	unassigned	up	down
GigabitEthernet0/0/3	unassigned	up	down
LoopBack0	23.23.23.23/32	up	up(s)
NULL0	unassigned	up	up(s)
Serial1/0/0	unassigned	up	up
Serial2/0/0	unassigned	up	up
Serial3/0/0	unassigned	up	down

<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 10.0.13.2 24
[R3-GigabitEthernet0/0/2]quit
[R3]interface LoopBack 0
[R3-LoopBack0]ip address 23.23.23.23 32
[R3-LoopBack0]quit
```

```
[R3]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 9

The number of interface that is DOWN in Physical is 4

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
Ethernet4/0/0	unassigned	down	down
Ethernet4/0/1	unassigned	down	down
GigabitEthernet0/0/0	unassigned	up	down
GigabitEthernet0/0/1	unassigned	up	down
GigabitEthernet0/0/2	10.0.13.2/24	up	up
GigabitEthernet0/0/3	unassigned	up	down
LoopBack0	23.23.23.23/32	up	up(s)

NULL0	unassigned	up	up(s)
Serial1/0/0	unassigned	up	down
Serial2/0/0	unassigned	up	up
Serial3/0/0	unassigned	up	up

Check the connectivity between R1 and R2 and between R1 and R3.

[R1]ping 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=1 ms

Reply from 10.0.12.2: bytes=56 Sequence=2 ttl=255 time=1 ms

Reply from 10.0.12.2: bytes=56 Sequence=3 ttl=255 time=1 ms

Reply from 10.0.12.2: bytes=56 Sequence=4 ttl=255 time=1 ms

Reply from 10.0.12.2: bytes=56 Sequence=5 ttl=255 time=1 ms

--- 10.0.12.2 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

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

[R1]ping 10.0.13.2

PING 10.0.13.2: 56 data bytes, press CTRL_C to break

Reply from 10.0.13.2: bytes=56 Sequence=1 ttl=255 time=1 ms

Reply from 10.0.13.2: bytes=56 Sequence=2 ttl=255 time=1 ms

Reply from 10.0.13.2: bytes=56 Sequence=3 ttl=255 time=1 ms

Reply from 10.0.13.2: bytes=56 Sequence=4 ttl=255 time=1 ms

Reply from 10.0.13.2: bytes=56 Sequence=5 ttl=255 time=1 ms

--- 10.0.13.2 ping statistics ---

5 packet(s) transmitted

5 packet(s) received

0.00% packet loss

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

Step 2 Configure BFD.

Enable BFD on the active path and check the interface on R1 connected to R2.

[R1]bfd

[R1-bfd]quit

[R1]bfd 1 bind peer-ip 10.0.12.2 source-ip 10.0.12.1 auto

[R1-bfd-session-1]commit

[R1-bfd-session-1]quit

[R2]bfd

[R2-bfd]quit

[R2]bfd 1 bind peer-ip 10.0.12.1 source-ip 10.0.12.2 auto

[R2-bfd-session-1]commit

[R2-bfd-session-1]quit

Check BFD session information.

[R1]display bfd session all

Local	Remote	PeerIpAddr	State	Type	InterfaceName

8192	8192	10.0.12.2	Up	S_AUTO_PEER	-

Total UP/DOWN Session Number : 1/0

[R2]display bfd session all

```

-----
Local Remote      PeerIpAddr      State      Type      InterfaceName
-----
8192  8192      10.0.12.1      Up        S_AUTO_PEER      -
-----

```

Total UP/DOWN Session Number : 1/0

Step 3 Configure association between BFD and static routes.

On R2 and R3, configure static routes to the loopback interface on R1.

[R2]ip route-static 10.0.0.0 8 10.0.12.1

[R2]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
-----
10.0.0.0/8          Static   60   0        RD    10.0.12.1         GigabitEthernet0/0/1
10.0.12.0/24        Direct   0     0        D     10.0.12.2         GigabitEthernet0/0/1
10.0.12.2/32        Direct   0     0        D     127.0.0.1         GigabitEthernet0/0/1
10.0.12.255/32      Direct   0     0        D     127.0.0.1         GigabitEthernet0/0/1
23.23.23.23/32      Direct   0     0        D     127.0.0.1         LoopBack0
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]ip route-static 10.0.0.0 8 10.0.13.1
```

```
[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
10.0.0.0/8	Static	60	0	RD	10.0.13.1	GigabitEthernet0/0/2
10.0.13.0/24	Direct	0	0	D	10.0.13.2	GigabitEthernet0/0/2
10.0.13.2/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/2
10.0.13.255/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/2
23.23.23.23/32	Direct	0	0	D	127.0.0.1	LoopBack0
127.0.0.0/8	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
127.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
255.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0

Configure two static routes on R1 and associate them with BFD.

```
[R1]ip route-static 0.0.0.0 0.0.0.0 10.0.12.2 track bfd-session 1
```

```
[R1]ip route-static 0.0.0.0 0.0.0.0 10.0.13.2 preference 100
```

The route to R3 has the priority of 100 and is lower than the route to R2 (60). The routing table is as follows.

```
[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
0.0.0.0/0	Static	60	0	RD	10.0.12.2	GigabitEthernet0/0/1
10.0.1.1/32	Direct	0	0	D	127.0.0.1	LoopBack0
10.0.12.0/24	Direct	0	0	D	10.0.12.1	GigabitEthernet0/0/1
10.0.12.1/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/1
10.0.12.255/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/1
10.0.13.0/24	Direct	0	0	D	10.0.13.1	GigabitEthernet0/0/2
10.0.13.1/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/2
10.0.13.255/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/2
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

```
[R1]display ip routing-table 0.0.0.0 0.0.0.0 verbose
```

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

Routing Table : Public

Summary Count : 2

Destination: 0.0.0.0/0

Protocol: Static

Process ID: 0

Preference: 60 Cost: 0
 NextHop: 10.0.12.2 Neighbour: 0.0.0.0
 State: Active Adv Relied Age: 00h01m19s
 Tag: 0 Priority: medium
 Label: NULL QoSInfo: 0x0

IndirectID: 0x80000001

RelayNextHop: 0.0.0.0 Interface: GigabitEthernet0/0/1
 TunnelID: 0x0 Flags: RD

Destination: 0.0.0.0/0

Protocol: Static Process ID: 0
 Preference: 100 Cost: 0
 NextHop: 10.0.13.2 Neighbour: 0.0.0.0
 State: Inactive Adv Relied Age: 00h01m03s
 Tag: 0 Priority: medium
 Label: NULL QoSInfo: 0x0

IndirectID: 0x80000002

RelayNextHop: 0.0.0.0 Interface: GigabitEthernet0/0/2
 TunnelID: 0x0 Flags: R

Check the connectivity in the normal situation.

[R1]ping -a 10.0.1.1 23.23.23.23

PING 23.23.23.23: 56 data bytes, press CTRL_C to break

Reply from 23.23.23.23: bytes=56 Sequence=1 ttl=255 time=1 ms

Reply from 23.23.23.23: bytes=56 Sequence=2 ttl=255 time=1 ms

Reply from 23.23.23.23: bytes=56 Sequence=3 ttl=255 time=1 ms

Reply from 23.23.23.23: bytes=56 Sequence=4 ttl=255 time=1 ms

Reply from 23.23.23.23: bytes=56 Sequence=5 ttl=255 time=1 ms

```
--- 23.23.23.23 ping statistics ---
```

```
5 packet(s) transmitted
```

```
5 packet(s) received
```

```
0.00% packet loss
```

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

Perform the ping operation on R1 for a long time and shut down the interface on R2.

```
[R1]ping -c 100 23.23.23.23
```

```
[R2]interface GigabitEthernet 0/0/1
```

```
[R2-GigabitEthernet0/0/1]shutdown
```

Check the configuration on R1.

```
[R1]ping -c 100 23.23.23.23
```

```
PING 23.23.23.23: 56 data bytes, press CTRL_C to break
```

```
Reply from 23.23.23.23: bytes=56 Sequence=1 ttl=255 time=1 ms
```

```
Reply from 23.23.23.23: bytes=56 Sequence=2 ttl=255 time=1 ms
```

```
Reply from 23.23.23.23: bytes=56 Sequence=3 ttl=255 time=1 ms
```

```
Reply from 23.23.23.23: bytes=56 Sequence=4 ttl=255 time=1 ms
```

```
Reply from 23.23.23.23: bytes=56 Sequence=5 ttl=255 time=1 ms
```

```
Reply from 23.23.23.23: bytes=56 Sequence=6 ttl=255 time=1 ms
```

```
Reply from 23.23.23.23: bytes=56 Sequence=7 ttl=255 time=1 ms
```

```
Reply from 23.23.23.23: bytes=56 Sequence=8 ttl=255 time=1 ms
```

```
Reply from 23.23.23.23: bytes=56 Sequence=9 ttl=255 time=1 ms
```

```
Request time out
```

```
Request time out
```

```
Reply from 23.23.23.23: bytes=56 Sequence=12 ttl=255 time=1 ms
```

Reply from 23.23.23.23: bytes=56 Sequence=13 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=14 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=15 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=16 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=17 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=18 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=19 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=20 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=21 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=22 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=23 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=24 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=25 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=26 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=27 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=28 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=29 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=30 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=31 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=32 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=33 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=34 ttl=255 time=1 ms
 Reply from 23.23.23.23: bytes=56 Sequence=35 ttl=255 time=1 ms

--- 23.23.23.23 ping statistics ---

35 packet(s) transmitted

33 packet(s) received

5.71% packet loss

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

Check the BFD session.

[R1]display bfd session all

```
-----  
Local Remote      PeerIpAddr      State      Type      InterfaceName  
-----  
8192  0           10.0.12.2       Down       S_AUTO_PEER  -  
-----
```

Total UP/DOWN Session Number : 0/1

Check routing information on R1.

[R1]display ip routing-table

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

Routing Tables: Public

Destinations : 12 Routes : 12

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
0.0.0.0/0	Static	100	0	RD	10.0.13.2	GigabitEthernet0/0/2
10.0.1.1/32	Direct	0	0	D	127.0.0.1	LoopBack0
10.0.12.0/24	Direct	0	0	D	10.0.12.1	GigabitEthernet0/0/1
10.0.12.1/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/1
10.0.12.255/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/1
10.0.13.0/24	Direct	0	0	D	10.0.13.1	GigabitEthernet0/0/2
10.0.13.1/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/2
10.0.13.255/32	Direct	0	0	D	127.0.0.1	GigabitEthernet0/0/2
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

[R1]display ip routing-table 0.0.0.0 verbose

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

Routing Table : Public

Summary Count : 2

Destination: 0.0.0.0/0

Protocol: Static

Process ID: 0

Preference: 60

Cost: 0

NextHop: 10.0.12.2

Neighbour: 0.0.0.0

State: Invalid Adv Relied

Age: 00h05m27s

Tag: 0

Priority: medium

Label: NULL

QoSInfo: 0x0

IndirectID: 0x80000001

RelayNextHop: 0.0.0.0

Interface: GigabitEthernet0/0/1

TunnelID: 0x0

Flags: R

Destination: 0.0.0.0/0

Protocol: Static

Process ID: 0

Preference: 100

Cost: 0

NextHop: 10.0.13.2

Neighbour: 0.0.0.0

State: Active Adv Relied

Age: 00h05m11s

Tag: 0

Priority: medium

Label: NULL

QoSInfo: 0x0

IndirectID: 0x80000002

RelayNextHop: 0.0.0.0

Interface: GigabitEthernet0/0/2

TunnelID: 0x0

Flags: RD

If BFD is not configured, no mechanism on R1 can detect reachability of static routes. That is, BFD is important in such scenarios.

Device Configuration

```
<R1>display current-configuration
[V200R007C00SPC600]
#
 sysname R1
#
bfd
#
interface GigabitEthernet0/0/1
 ip address 10.0.12.1 255.255.255.0
#
interface GigabitEthernet0/0/2
 ip address 10.0.13.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
#
bfd 1 bind peer-ip 10.0.12.2 source-ip 10.0.12.1 auto
commit
#
ip route-static 0.0.0.0 0.0.0.0 10.0.12.2 track bfd-session 1
ip route-static 0.0.0.0 0.0.0.0 10.0.13.2 preference 100
```

```
#  
  
return  
  
<R2>display current-configuration  
[V200R007C00SPC600]  
  
#  
  
sysname R2  
  
#  
  
bfd  
  
#  
  
interface GigabitEthernet0/0/1  
  
ip address 10.0.12.2 255.255.255.0  
  
#  
  
interface LoopBack0  
  
ip address 23.23.23.23 255.255.255.255  
  
#  
  
bfd 1 bind peer-ip 10.0.12.1 source-ip 10.0.12.2 auto  
  
commit  
  
#  
  
ip route-static 10.0.0.0 255.0.0.0 10.0.12.1  
  
#  
  
return  
  
<R3>display current-configuration  
[V200R007C00SPC600]  
  
#  
  
sysname R3  
  
#  
  
interface GigabitEthernet0/0/2
```

```
ip address 10.0.13.2 255.255.255.0
#
interface LoopBack0
ip address 23.23.23.23 255.255.255.255
#
ip route-static 10.0.0.0 255.0.0.0 10.0.13.1
#
return
```

Lab 6-2 Association Between BFD and OSPF

Learning Objectives

The objectives of this lab are to learn and understand:

- How to rapidly configure BFD in the Open Shortest Path First (OSPF) environment

Topology




Figure 6-2 Networking of association between BFD and OSPF

Scenario

R1 connects to R2 through S1. Interfaces on R1 and R2 run in OSPF area 0, and are not directly connected. When one interface goes Down, the other interface cannot

detect the fault immediately and has to wait for four times the Hello time to delete the neighbor. During this period, data is forwarded abnormally. In this case, configure BFD to accelerate the OSPF convergence speed.

Tasks

Step 1 Perform basic configurations and configure IP addresses.

Configure IP addresses for all routers.

```
<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 10.0.12.1 24
```

```
[R1-GigabitEthernet0/0/1]quit
```

```
[R1]interface loopback 0
```

```
[R1-LoopBack0]ip address 10.0.1.1 32
```

```
[R1-LoopBack0]quit
```

After the configuration is complete, check the IP address of each router.

```
[R1]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 3

The number of interface that is DOWN in Physical is 9

The number of interface that is UP in Protocol is 3

The number of interface that is DOWN in Protocol is 9

Interface	IP Address/Mask	Physical	Protocol
Cellular0/0/0	unassigned	down	down
Cellular0/0/1	unassigned	down	down
GigabitEthernet0/0/0	unassigned	*down	down
GigabitEthernet0/0/1	10.0.12.1/24	up	up
GigabitEthernet0/0/2	unassigned	*down	down
GigabitEthernet0/0/3	unassigned	*down	down
LoopBack0	10.0.1.1/32	up	up(s)
NULL0	unassigned	up	up(s)
Serial1/0/0	unassigned	*down	down
Serial2/0/0	unassigned	*down	down
Serial3/0/0	unassigned	*down	down
Serial4/0/0	unassigned	*down	down

<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 10.0.12.2 24

[R2-GigabitEthernet0/0/1]quit

[R2]interface loopback 0

[R2-LoopBack0]ip address 10.0.2.2 32

[R2-LoopBack0]quit

[R2]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 3

The number of interface that is DOWN in Physical is 10

The number of interface that is UP in Protocol is 3

The number of interface that is DOWN in Protocol is 10

Interface	IP Address/Mask	Physical	Protocol
Cellular0/0/0	unassigned	down	down
Cellular0/0/1	unassigned	down	down
Ethernet4/0/0	unassigned	*down	down
Ethernet4/0/1	unassigned	*down	down
GigabitEthernet0/0/0	unassigned	*down	down
GigabitEthernet0/0/1	10.0.12.2/24	up	up
GigabitEthernet0/0/2	unassigned	*down	down
GigabitEthernet0/0/3	unassigned	*down	down
LoopBack0	10.0.2.2/32	up	up(s)
NULL0	unassigned	up	up(s)
Serial1/0/0	unassigned	*down	down
Serial2/0/0	unassigned	*down	down
Serial3/0/0	unassigned	*down	down

Check the connectivity between R1 and R2.

[R1]ping 10.0.12.2

PING 10.0.12.2: 56 data bytes, press CTRL_C to break

Request time out

Reply from 10.0.12.2: bytes=56 Sequence=2 ttl=255 time=1 ms

Reply from 10.0.12.2: bytes=56 Sequence=3 ttl=255 time=1 ms

Reply from 10.0.12.2: bytes=56 Sequence=4 ttl=255 time=1 ms

Reply from 10.0.12.2: bytes=56 Sequence=5 ttl=255 time=1 ms

--- 10.0.12.2 ping statistics ---

5 packet(s) transmitted

4 packet(s) received

20.00% packet loss

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

Step 2 **Configure OSPF.**

Assign interfaces of R1 and R2 including Loopback0 interfaces to the OSPF area 0 based on the topology.

```
[R1]ospf 1
```

```
[R1-ospf-1]area 0
```

```
[R1-ospf-1-area-0.0.0.0]network 10.0.12.0 0.0.0.255
```

```
[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
```

```
[R2-ospf-1]area 0
```

```
[R2-ospf-1-area-0.0.0.0]network 10.0.12.0 0.0.0.255
```

```
[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
```

Check the OSPF interface status and neighbor status.

[R1]display ospf interface

OSPF Process 1 with Router ID 10.0.12.1

Interfaces

Area: 0.0.0.0 (MPLS TE not enabled)

IP Address	Type	State	Cost	Pri	DR	BDR
10.0.12.1	Broadcast	BDR	1	1	10.0.12.2	10.0.12.1
10.0.1.1	P2P	P-2-P	0	1	0.0.0.0	0.0.0.0

[R2]display ospf interface

OSPF Process 1 with Router ID 10.0.12.2

Interfaces

Area: 0.0.0.0 (MPLS TE not enabled)

IP Address	Type	State	Cost	Pri	DR	BDR
10.0.12.2	Broadcast	DR	1	1	10.0.12.2	10.0.12.1
10.0.2.2	P2P	P-2-P	0	1	0.0.0.0	0.0.0.0

Check the OSPF neighbor relationship status.

[R1]display ospf peer brief

OSPF Process 1 with Router ID 10.0.12.1

Peer Statistic Information

Area Id	Interface	Neighbor id	State
---------	-----------	-------------	-------

0.0.0.0	GigabitEthernet0/0/1	10.0.12.2	Full
---------	----------------------	-----------	------

Total Peer(s): 1

[R2]display ospf peer brief

OSPF Process 1 with Router ID 10.0.12.2

Peer Statistic Information

Area Id	Interface	Neighbor id	State
0.0.0.0	GigabitEthernet0/0/1	10.0.12.1	Full

Total Peer(s): 1

When the OSPF neighbor relationship status is full, the OSPF configuration is complete.

Step 3 **Configure BFD sessions.**

Enable BFD globally and in an OSPF area.

[R1]bfd

[R1-bfd]quit

[R1]ospf 1

[R1-ospf-1]bfd all-interfaces enable

[R1-ospf-1]quit

[R2]bfd

[R2-bfd]quit

[R2]ospf 1

[R2-ospf-1]bfd all-interfaces enable

```
[R2-ospf-1]quit
```

After the configuration is complete, check the BFD session status.

```
[R1]display bfd session all
```

```
-----
Local Remote      PeerIpAddr      State      Type      InterfaceName
-----
8192  8192      10.0.12.2      Up          D_IP_IF      GigabitEthernet0/0/1
-----
```

Total UP/DOWN Session Number : 1/0

```
[R2]display bfd session all
```

```
-----
Local Remote      PeerIpAddr      State      Type      InterfaceName
-----
8192  8192      10.0.12.1      Up          D_IP_IF      GigabitEthernet0/0/1
-----
```

Total UP/DOWN Session Number : 1/0

```
[R1]display ospf bfd session all
```

OSPF Process 1 with Router ID 10.0.12.1

Area 0.0.0.0 interface 10.0.12.1(GigabitEthernet0/0/1)'s BFD Sessions

NeighborId:10.0.12.2	AreaId:0.0.0.0	Interface:GigabitEthernet0/0/1
BFDState:up	rx :1000	tx :1000
Multiplier:3	BFD Local Dis:8192	LocalIpAdd:10.0.12.1
RemoteIpAdd:10.0.12.2	Diagnostic Info:No diagnostic information	

To test the BFD effect, shut down the interface on R2.

```
[R2]interface GigabitEthernet 0/0/1
[R2-GigabitEthernet0/0/1]shutdown
```

Enable the debugging function on R1. The following debugging information is displayed on R1.

```
<R1>debug ospf bfd
```

```
Sep 23 2016 03:39:25+00:00 R1 %%01BFD/4/STACHG_TODWN(l)[23]:BFD session changed to Down.
(SlotNumber=0, Discriminator=8192, Diagnostic=DetectDown, Applications=OSPF, ProcessPST=False,
BindInterfaceName=GigabitEthernet0/0/1, InterfacePhysicalState=Up, InterfaceProtocolState=Up)
```

```
<R1>
```

```
Sep 23 2016 03:39:25+00:00 R1 %%01OSPF/3/NBR_CHG_DOWN(l)[24]:Neighbor event:neighbor state changed
to Down. (ProcessId=1, NeighborAddress=10.0.12.2, NeighborEvent=KillNbr, NeighborPreviousState=Full,
NeighborCurrentState=Down)
```

```
<R1>
```

```
Sep 23 2016 03:39:25+00:00 R1 %%01OSPF/3/NBR_DOWN_REASON(l)[25]:Neighbor state leaves full or changed
to Down. (ProcessId=1, NeighborRouterId=10.0.12.2, NeighborAreaId=0,
NeighborInterface=GigabitEthernet0/0/1,NeighborDownImmediate reason=Neighbor Down Due to Kill
Neighbor, NeighborDownPrimeReason=BFD Session Down, NeighborChangeTime=2016-09-23 03:39:25)
```

Other association logs are not displayed here. Focus on the preceding important logs.

Restart the interface.

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

Check the BFD session status and OSPF neighbor relationship status again.

[R1]display bfd session all

Local Remote		PeerIpAddr	State	Type	InterfaceName
8193	8193	10.0.12.2	Up	D_IP_IF	GigabitEthernet0/0/1

Total UP/DOWN Session Number : 1/0

[R1]display ospf bfd session all

OSPF Process 1 with Router ID 10.0.12.1

Area 0.0.0.0 interface 10.0.12.1(GigabitEthernet0/0/1)'s BFD Sessions

```
NeighborId:10.0.12.2      Areaid:0.0.0.0      Interface:GigabitEthernet0/0/1
BFDState:up              rx      :1000      tx      :1000
Multiplier:3             BFD Local Dis:8193      LocalIpAdd:10.0.12.1
RemoteIpAdd:10.0.12.2     Diagnostic Info:No diagnostic information
```

[R2]display bfd session all

Local Remote		PeerIpAddr	State	Type	InterfaceName
8193	8193	10.0.12.1	Up	D_IP_IF	GigabitEthernet0/0/1

Total UP/DOWN Session Number : 1/0

[R2]display ospf bfd session all

OSPF Process 1 with Router ID 10.0.12.2

Area 0.0.0.0 interface 10.0.12.2(GigabitEthernet0/0/1)'s BFD Sessions

NeighborId:10.0.12.1	AreaId:0.0.0.0	Interface:GigabitEthernet0/0/1
BFDState:up	rx :1000	tx :1000
Multiplier:3	BFD Local Dis:8193	LocalIpAdd:10.0.12.2
RemoteIpAdd:10.0.12.1	Diagnostic Info:No diagnostic information	

BFD sessions are established again.

Device Configuration

<R1>display current-configuration

[V200R007C00SPC600]

#

sysname R1

#

bfd

#

interface GigabitEthernet0/0/1

ip address 10.0.12.1 255.255.255.0

#

interface LoopBack0

ip address 10.0.1.1 255.255.255.255

#

ospf 1

bfd all-interfaces enable

area 0.0.0.0

```
network 10.0.1.1 0.0.0.0

network 10.0.12.0 0.0.0.255

#

return

<R2>display current-configuration
[V200R007C00SPC600]

#

sysname R2

#

bfd

#

interface GigabitEthernet0/0/1

ip address 10.0.12.2 255.255.255.0

#

interface LoopBack0

ip address 10.0.2.2 255.255.255.255

#

ospf 1

bfd all-interfaces enable

area 0.0.0.0

network 10.0.2.2 0.0.0.0

network 10.0.12.0 0.0.0.255

#

return
```

Lab 6-3 Association Between BFD and VRRP

Learning Objectives

The learning objectives of this lab are to learn and understand:

- How to check indirectly connected interfaces through association between BFD and VRRP

Topology




Figure 6-3 Networking of association between BFD and VRRP

Scenario

R1 connects to R2 and R3 through S1 and S2. R2 and R3 are configured with VRRP and function as gateways of R4 and R5 on a LAN. R2 is used as the master device, and R3 is used as the backup device. When the indirectly connected uplink of R2 goes Down, uplink traffic is still forwarded through R2, causing blackhole routes. To prevent such a problem, configure association between BFD and VRRP. When

connected interfaces on R1 and R2 go Down, the priority of the VRRP group of R2 is reduced immediately and R3 functions as the master router to forward uplink traffic.

Tasks

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

Configure IP addresses for all routers.

```
<Huawei>system-view
```

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname R1
```

```
[R1]interface LoopBack 0
```

```
[R1-LoopBack0]ip address 10.0.1.1 24
```

```
[R1-LoopBack0]quit
```

```
[R1]interface GigabitEthernet 0/0/1
```

```
[R1-GigabitEthernet0/0/1]ip address 10.0.12.1 24
```

```
[R1-GigabitEthernet0/0/1]quit
```

```
[R1]interface GigabitEthernet 0/0/2
```

```
[R1-GigabitEthernet0/0/2]ip address 10.0.13.1 24
```

```
[R1-GigabitEthernet0/0/2]quit
```

```
<Huawei>system-view
```

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname R2
```

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

```
[R2-LoopBack0]ip address 10.0.2.2 24
```

```
[R2-LoopBack0]quit
```

```
[R2]interface GigabitEthernet 0/0/1
```

```
[R2-GigabitEthernet0/0/1]ip address 10.0.12.2 24
```

```
[R2-GigabitEthernet0/0/1]quit
```

```
[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 192.168.45.2 24
[R2-GigabitEthernet0/0/0]quit
```

<Huawei>system-view

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname R3
[R3]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 24
[R3-LoopBack0]quit
[R3]interface GigabitEthernet 0/0/2
[R3-GigabitEthernet0/0/2]ip address 10.0.13.2 24
[R3-GigabitEthernet0/0/2]quit
[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]ip address 192.168.45.3 24
[R3-GigabitEthernet0/0/0]quit
```

<Huawei>system-view

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname R4
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ip address 192.168.45.4 24
[R4-GigabitEthernet0/0/0]quit
```

<Huawei>system-view

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname R5
[R5]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 192.168.45.5 24
```

```
[R5-GigabitEthernet0/0/0]quit
```

Configure VLANs on SW1 and SW2 respectively to avoid conflicts.

```
<Huawei>system-view
```

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname SW1
```

```
[SW1]vlan 12
```

```
[SW1-vlan12]quit
```

```
[SW1]interface GigabitEthernet 0/0/1
```

```
[SW1-GigabitEthernet0/0/1]port link-type access
```

```
[SW1-GigabitEthernet0/0/1]port default vlan 12
```

```
[SW1-GigabitEthernet0/0/1]quit
```

```
[SW1]interface GigabitEthernet 0/0/2
```

```
[SW1-GigabitEthernet0/0/2]port link-type access
```

```
[SW1-GigabitEthernet0/0/2]port default vlan 12
```

```
[SW1-GigabitEthernet0/0/2]quit
```

```
<Huawei>system-view
```

Enter system view, return user view with Ctrl+Z.

```
[Huawei]sysname SW2
```

```
[SW2]vlan 13
```

```
[SW2-vlan13]quit
```

```
[SW2]interface GigabitEthernet 0/0/1
```

```
[SW2-GigabitEthernet0/0/1]port link-type access
```

```
[SW2-GigabitEthernet0/0/1]port default vlan 13
```

```
[SW2-GigabitEthernet0/0/1]quit
```

```
[SW2]interface GigabitEthernet 0/0/3
```

```
[SW2-GigabitEthernet0/0/3]port link-type access
```

```
[SW2-GigabitEthernet0/0/3]port default vlan 13
```

```
[SW2-GigabitEthernet0/0/3]quit
```

Check IP addresses after the configuration is complete.

```
[R1]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 9

The number of interface that is DOWN in Physical is 3

The number of interface that is UP in Protocol is 6

The number of interface that is DOWN in Protocol is 6

Interface	IP Address/Mask	Physical	Protocol
Cellular0/0/0	unassigned	down	down
Cellular0/0/1	unassigned	down	down
GigabitEthernet0/0/0	unassigned	up	down
GigabitEthernet0/0/1	10.0.12.1/24	up	up
GigabitEthernet0/0/2	10.0.13.1/24	up	up
GigabitEthernet0/0/3	unassigned	up	down
LoopBack0	10.0.1.1/24	up	up(s)
NULL0	unassigned	up	up(s)
Serial1/0/0	unassigned	up	up
Serial2/0/0	unassigned	up	down
Serial3/0/0	unassigned	up	up
Serial4/0/0	unassigned	down	down

[R2]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 9

The number of interface that is DOWN in Physical is 4

The number of interface that is UP in Protocol is 6

The number of interface that is DOWN in Protocol is 7

Interface	IP Address/Mask	Physical	Protocol
Cellular0/0/0	unassigned	down	down
Cellular0/0/1	unassigned	down	down
Ethernet4/0/0	unassigned	down	down
Ethernet4/0/1	unassigned	down	down
GigabitEthernet0/0/0	192.168.45.2/24	up	up
GigabitEthernet0/0/1	10.0.12.2/24	up	up
GigabitEthernet0/0/2	unassigned	up	down
GigabitEthernet0/0/3	unassigned	up	down
LoopBack0	10.0.2.2/24	up	up(s)
NULL0	unassigned	up	up(s)
Serial1/0/0	unassigned	up	up
Serial2/0/0	unassigned	up	up
Serial3/0/0	unassigned	up	down

[R3]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 9

The number of interface that is DOWN in Physical is 4

The number of interface that is UP in Protocol is 6

The number of interface that is DOWN in Protocol is 7

Interface	IP Address/Mask	Physical	Protocol
Cellular0/0/0	unassigned	down	down
Cellular0/0/1	unassigned	down	down
Ethernet4/0/0	unassigned	down	down
Ethernet4/0/1	unassigned	down	down
GigabitEthernet0/0/0	192.168.45.3/24	up	up
GigabitEthernet0/0/1	unassigned	up	down
GigabitEthernet0/0/2	10.0.13.2/24	up	up
GigabitEthernet0/0/3	unassigned	up	down
LoopBack0	10.0.3.3/24	up	up(s)
NULL0	unassigned	up	up(s)
Serial1/0/0	unassigned	up	down
Serial2/0/0	unassigned	up	up
Serial3/0/0	unassigned	up	up

[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 6

The number of interface that is DOWN in Physical is 5

The number of interface that is UP in Protocol is 3

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	unassigned	up	down
Ethernet2/0/1	unassigned	down	down
GigabitEthernet0/0/0	192.168.45.4/24	up	up
GigabitEthernet0/0/1	unassigned	up	down
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

[R5]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 6

The number of interface that is DOWN in Physical is 5

The number of interface that is UP in Protocol is 3

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	unassigned	up	down
Ethernet2/0/1	unassigned	down	down
GigabitEthernet0/0/0	192.168.45.5/24	up	up
GigabitEthernet0/0/1	unassigned	up	down
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

Check the connectivity between R1 and R2 and between R1 and R3.

[R1]ping 10.0.12.2

PING 10.0.12.2: 56 data bytes, press CTRL_C to break

Request time out

Reply from 10.0.12.2: bytes=56 Sequence=2 ttl=255 time=1 ms

Reply from 10.0.12.2: bytes=56 Sequence=3 ttl=255 time=1 ms

Reply from 10.0.12.2: bytes=56 Sequence=4 ttl=255 time=1 ms

Reply from 10.0.12.2: bytes=56 Sequence=5 ttl=255 time=1 ms

--- 10.0.12.2 ping statistics ---

5 packet(s) transmitted

4 packet(s) received

20.00% packet loss

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

[R1]ping 10.0.13.2

PING 10.0.13.2: 56 data bytes, press CTRL_C to break

Request time out

Reply from 10.0.13.2: bytes=56 Sequence=2 ttl=255 time=1 ms

Reply from 10.0.13.2: bytes=56 Sequence=3 ttl=255 time=1 ms

Reply from 10.0.13.2: bytes=56 Sequence=4 ttl=255 time=1 ms

Reply from 10.0.13.2: bytes=56 Sequence=5 ttl=255 time=1 ms

--- 10.0.13.2 ping statistics ---

5 packet(s) transmitted

4 packet(s) received

20.00% packet loss

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

Step 2 **Configure OSPF and static routes.**

Configure OSPF on R1, R2, and R3 according to the topology. Use network commands to import OSPF routes and enable the silent interface.

[R1]ospf 1

[R1-ospf-1]area 0

[R1-ospf-1-area-0.0.0.0]network 10.0.1.0 0.0.0.255

[R1-ospf-1-area-0.0.0.0]network 10.0.12.0 0.0.0.255

[R1-ospf-1-area-0.0.0.0]network 10.0.13.0 0.0.0.255

[R1-ospf-1-area-0.0.0.0]quit

[R1-ospf-1]quit

Modify the OSPF cost on R1 so that downlink traffic is forwarded through R2.

```
[R1]interface GigabitEthernet 0/0/1
```

```
[R1-GigabitEthernet0/0/1]ospf cost 90
```

```
[R1-GigabitEthernet0/0/1]quit
```

```
[R1]interface GigabitEthernet 0/0/2
```

```
[R1-GigabitEthernet0/0/2]ospf cost 100
```

```
[R1-GigabitEthernet0/0/2]quit
```

```
[R2]ospf 1
```

```
[R2-ospf-1]silent-interface GigabitEthernet 0/0/0
```

```
[R2-ospf-1]area 0
```

```
[R2-ospf-1-area-0.0.0.0]network 10.0.2.0 0.0.0.255
```

```
[R2-ospf-1-area-0.0.0.0]network 10.0.12.0 0.0.0.255
```

```
[R2-ospf-1-area-0.0.0.0]network 192.168.45.0 0.0.0.255
```

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

```
[R2-ospf-1]quit
```

```
[R2]interface GigabitEthernet 0/0/1
```

```
[R2-GigabitEthernet0/0/1]ospf cost 90
```

```
[R2-GigabitEthernet0/0/1]quit
```

```
[R3]ospf 1
```

```
[R3-ospf-1]silent-interface GigabitEthernet 0/0/0
```

```
[R3-ospf-1]area 0
```

```
[R3-ospf-1-area-0.0.0.0]network 10.0.3.0 0.0.0.255
```

```
[R3-ospf-1-area-0.0.0.0]network 192.168.45.0 0.0.0.255
```

```
[R3-ospf-1-area-0.0.0.0]network 10.0.13.0 0.0.0.255
```

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

```
[R3-ospf-1]quit
[R3]interface GigabitEthernet 0/0/2
[R3-GigabitEthernet0/0/2]ospf cost 100
[R3-GigabitEthernet0/0/2]quit
```

Check the routing information after OSPF convergence is complete.

```
[R1]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
10.0.2.2/32	OSPF	10	90	D	10.0.12.2	GigabitEthernet0/0/1
10.0.3.3/32	OSPF	10	100	D	10.0.13.2	GigabitEthernet0/0/2
192.168.45.0/24	OSPF	10	91	D	10.0.12.2	GigabitEthernet0/0/1

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 : 3 Routes : 3

OSPF routing table status : <Active>

Destinations : 3 Routes : 3

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
10.0.1.1/32	OSPF	10	90	D	10.0.12.1	GigabitEthernet0/0/1
10.0.3.3/32	OSPF	10	190	D	10.0.12.1	GigabitEthernet0/0/1
10.0.13.0/24	OSPF	10	190	D	10.0.12.1	GigabitEthernet0/0/1

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 : 3 Routes : 3

OSPF routing table status : <Active>

Destinations : 3 Routes : 3

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
10.0.1.1/32	OSPF	10	100	D	10.0.13.1	GigabitEthernet0/0/2
10.0.2.2/32	OSPF	10	190	D	10.0.13.1	GigabitEthernet0/0/2
10.0.12.0/24	OSPF	10	190	D	10.0.13.1	GigabitEthernet0/0/2

OSPF routing table status : <Inactive>

Destinations : 0 Routes : 0

Configure default routes to the VRRP virtual IP address on R4 and R5.

```
[R4]ip route-static 0.0.0.0 0 192.168.45.1
```

```
[R5]ip route-static 0.0.0.0 0 192.168.45.1
```

Step 3 **Configure VRRP.**

Configure VRRP on downlink interfaces of R2 and R3.

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

```
[R2-GigabitEthernet0/0/0]vrrp vrid 45 virtual-ip 192.168.45.1
```

```
[R2-GigabitEthernet0/0/0]vrrp vrid 45 priority 150
```

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

```
[R3]interface GigabitEthernet 0/0/0
```

```
[R3-GigabitEthernet0/0/0]vrrp vrid 45 virtual-ip 192.168.45.1
```

```
[R3-GigabitEthernet0/0/0]quit
```

Check the VRRP status of devices.

```
[R2]display vrrp
```

```
GigabitEthernet0/0/0 | Virtual Router 45
```

```
State : Master
```

```
Virtual IP : 192.168.45.1
```

```
Master IP : 192.168.45.2
```

```
PriorityRun : 150
```

```
PriorityConfig : 150
```

MasterPriority : 150
Preempt : YES Delay Time : 0 s
TimerRun : 1 s
TimerConfig : 1 s
Auth type : NONE
Virtual MAC : 0000-5e00-012d
Check TTL : YES
Config type : normal-vrrp
Backup-forward : disabled
Create time : 2016-09-25 15:18:54
Last change time : 2016-09-25 15:18:57

[R3]display vrrp

GigabitEthernet0/0/0 | Virtual Router 45

State : Backup

Virtual IP : 192.168.45.1
Master IP : 192.168.45.2
PriorityRun : 100
PriorityConfig : 100
MasterPriority : 150
Preempt : YES Delay Time : 0 s
TimerRun : 1 s
TimerConfig : 1 s
Auth type : NONE
Virtual MAC : 0000-5e00-012d
Check TTL : YES
Config type : normal-vrrp
Backup-forward : disabled
Create time : 2016-09-25 15:21:49

Last change time : 2016-09-25 15:21:49

Step 4 **Configure BFD association.**

Enable BFD on R1 and R2, and associate BFD with VRRP. When BFD detects a fault, the priority of the VRRP group is reduced immediately.

```
[R1]bfd
```

```
[R1-bfd]quit
```

```
[R1]bfd 1 bind peer-ip 192.168.45.2 source-ip 10.0.12.1 auto
```

```
[R1-bfd-session-1]commit
```

```
[R1-bfd-session-1]quit
```

```
[R2]bfd
```

```
[R2-bfd]quit
```

```
[R2]bfd 1 bind peer-ip 10.0.12.1 source-ip 192.168.45.2 auto
```

```
[R2-bfd-session-1]commit
```

```
[R2-bfd-session-1]quit
```

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

```
[R2-GigabitEthernet0/0/0]vrrp vrid 45 track bfd-session session-name 1 reduce 60
```

Check association configuration.

```
[R2]display vrrp
```

```
GigabitEthernet0/0/0 | Virtual Router 45
```

```
State : Master
```

```
Virtual IP : 192.168.45.1
```

```
Master IP : 192.168.45.2
```

```
PriorityRun : 150
```

```
PriorityConfig : 150
```

MasterPriority : 150
Preempt : YES Delay Time : 0 s
TimerRun : 1 s
TimerConfig : 1 s
Auth type : NONE
Virtual MAC : 0000-5e00-012d
Check TTL : YES
Config type : normal-vrrp
Backup-forward : disabled
Track BFD : 1 Priority reduced : 60
BFD-session state : UP
Create time : 2016-09-25 15:18:54
Last change time : 2016-09-25 15:18:57

Check the BFD session.

[R2]display bfd session all

```
-----  
Local Remote    PeerIpAddr    State    Type    InterfaceName  
-----  
8192   8192    10.0.12.1    Up    S_AUTO_PEER    -  
-----
```

Total UP/DOWN Session Number : 1/0

Test BFD effects. Perform the ping operation for a long time on R4 and shut down the interface on R1.

[R4]ping -c 100 10.0.1.1

[R1]interface GigabitEthernet 0/0/1

```
[R1-GigabitEthernet0/0/1]shutdown
```

Check the ping operation result on R4.

```
[R4]ping -c 100 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=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=2 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=3 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=4 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=5 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=6 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=7 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=8 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=9 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=10 ttl=254 time=1 ms
```

```
Request time out
```

```
Request time out
```

```
Reply from 10.0.1.1: bytes=56 Sequence=13 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=14 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=15 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=16 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=17 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=18 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=19 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=20 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=21 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=22 ttl=254 time=1 ms
```

```
Reply from 10.0.1.1: bytes=56 Sequence=23 ttl=254 time=1 ms
```

Reply from 10.0.1.1: bytes=56 Sequence=24 ttl=254 time=1 ms

Reply from 10.0.1.1: bytes=56 Sequence=25 ttl=254 time=1 ms

Reply from 10.0.1.1: bytes=56 Sequence=26 ttl=254 time=1 ms

Reply from 10.0.1.1: bytes=56 Sequence=27 ttl=254 time=1 ms

Reply from 10.0.1.1: bytes=56 Sequence=28 ttl=254 time=1 ms

Reply from 10.0.1.1: bytes=56 Sequence=29 ttl=254 time=1 ms

Reply from 10.0.1.1: bytes=56 Sequence=30 ttl=254 time=1 ms

Reply from 10.0.1.1: bytes=56 Sequence=31 ttl=254 time=1 ms

Reply from 10.0.1.1: bytes=56 Sequence=32 ttl=254 time=1 ms

--- 10.0.1.1 ping statistics ---

32 packet(s) transmitted

30 packet(s) received

6.25% packet loss

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

Check the VRRP status.

[R2]display vrrp

GigabitEthernet0/0/0 | Virtual Router 45

State : Backup

Virtual IP : 192.168.45.1

Master IP : 192.168.45.3

PriorityRun : 90

PriorityConfig : 150

MasterPriority : 100

Preempt : YES Delay Time : 0 s

TimerRun : 1 s

TimerConfig : 1 s

Auth type : NONE

Virtual MAC : 0000-5e00-012d

Check TTL : YES

Config type : normal-vrrp

Backup-forward : disabled

Track BFD : 1 Priority reduced : 60

BFD-session state : DOWN

Create time : 2016-09-25 15:18:54

Last change time : 2016-09-25 15:27:26

Association between BFD and VRRP used to detect indirectly connected uplink detections is successful. The result on R5 is similar to that on R4, and the verification is not provided.

Device Configuration

<R1>display current-configuration

[V200R007C00SPC600]

#

sysname R1

#

bfd

#

interface GigabitEthernet0/0/1

ip address 10.0.12.1 255.255.255.0

ospf cost 90

#

interface GigabitEthernet0/0/2

ip address 10.0.13.1 255.255.255.0

ospf cost 100

```
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.0
#
bfd 1 bind peer-ip 192.168.45.2 source-ip 10.0.12.1 auto
 commit
#
ospf 1
 area 0.0.0.0
  network 10.0.1.0 0.0.0.255
  network 10.0.12.0 0.0.0.255
  network 10.0.13.0 0.0.0.255
#
return

<R2> display current-configuration
[V200R007C00SPC600]
#
 sysname R2
#
bfd
#
interface GigabitEthernet0/0/0
 ip address 192.168.45.2 255.255.255.0
 vrrp vrid 45 virtual-ip 192.168.45.1
 vrrp vrid 45 priority 150
 vrrp vrid 45 track bfd-session session-name 1 reduced 60
#
interface GigabitEthernet0/0/1
```

```
ip address 10.0.12.2 255.255.255.0

ospf cost 90

#

interface LoopBack0

ip address 10.0.2.2 255.255.255.0

#

bfd 1 bind peer-ip 10.0.12.1 source-ip 192.168.45.2 auto

commit

#

ospf 1

silent-interface GigabitEthernet0/0/0

area 0.0.0.0

network 10.0.2.0 0.0.0.255

network 10.0.12.0 0.0.0.255

network 192.168.45.0 0.0.0.255

#

return

<R3>display current-configuration

[V200R007C00SPC600]

#

sysname R3

#

interface GigabitEthernet0/0/0

ip address 192.168.45.3 255.255.255.0

vrrp vrid 45 virtual-ip 192.168.45.1

#

interface GigabitEthernet0/0/2

ip address 10.0.13.2 255.255.255.0
```

```
ospf cost 100

#

interface LoopBack0

ip address 10.0.3.3 255.255.255.0

#

ospf 1

silent-interface GigabitEthernet0/0/0

area 0.0.0.0

network 10.0.3.0 0.0.0.255

network 10.0.13.0 0.0.0.255

network 192.168.45.0 0.0.0.255

#

return
```

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

#

sysname R4

#

interface GigabitEthernet0/0/0

ip address 192.168.45.4 255.255.255.0

#

ip route-static 0.0.0.0 0.0.0.0 192.168.45.1

#

return
```

```
<R5> display current-configuration
[V200R007C00SPC600]

#
```

```
sysname R5

#

interface GigabitEthernet0/0/0

ip address 192.168.45.5 255.255.255.0

#

ip route-static 0.0.0.0 0.0.0.0 192.168.45.1

#

return
```

```
<SW1>display current-configuration
!Software Version V200R008C00SPC500

#

sysname SW1

#

vlan batch 12

#

interface GigabitEthernet0/0/1

port link-type access

port default vlan 12

#

interface GigabitEthernet0/0/2

port link-type access

port default vlan 12

#

return
```

```
<SW2>display current-configuration
!Software Version V200R008C00SPC500

#
```

```
sysname SW2

#

vlan batch 13

#

interface GigabitEthernet0/0/1

    port link-type access

    port default vlan 13

#

interface GigabitEthernet0/0/3

    port link-type access

    port default vlan 13

#

return
```



Recommendations

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



More Information

- Huawei learning APP

