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

HCIE-Data Center Lab Guide
V2.0

Huawei Technologies Co., Ltd.

Trademarks and Permissions


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

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

Notice

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

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

Huawei Technologies Co., Ltd.

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

Website: http://e.huawei.com

Huawei Certificate System

Relying on the strong technical strength and professional training system, Huawei provides a practical and professional four-level certificate system to meet various customer requirements on different WLAN technologies.

Huawei Certified ICT Associate-Wireless Local Area Network (HCIA-WLAN) is designed for Huawei local offices, online engineers in representative offices, and readers who want to understand Huawei WLAN products and technology. HCIA-WLAN covers WLAN basics, Control and Provisioning of Wireless Access Points (CAPWAP) protocol, WLAN networking, Huawei WLAN product features, security configuration, WLAN advanced technology, antennas, WLAN network planning and optimization, and WLAN fault troubleshooting.

The HCIA-WLAN certificate system introduces you to the industry and market, helps you in innovation, and enables you to stand atop the WLAN frontiers.

Contents

1 HCIE-Data Center Lab Environment Preparation 1

1.1 About the Environment 1

1.2 Lab Networking 2

1.3 Device Version and Configuration Information 2

1.4 Device Login Information 3

2 VXLAN Feature Experiment 4

2.1 Objectives 4

2.2 Configuring VXLAN Layer 2 Interconnection (No-Tunnel Mode) 5

2.2.1 Networking and Service Description 5

2.2.2 Configuration Guideline 5

2.2.3 Configuration Procedure 5

2.2.4 Verifying the Configuration Result 6

2.2.5 Complete Configuration 7

2.3 Configuring VXLAN Layer 2 Interconnection (Tunnel Mode) 17

2.3.1 Networking and Service Description 17

2.3.2 Configuration Guideline 18

2.3.3 Configuration Procedure 18

2.3.4 Verifying the Configuration Result 20

2.3.5 Complete Configuration 20

2.4 Configuring VXLAN Layer 3 Interconnection (No-Tunnel Mode) 29

2.4.1 Networking and Service Description 29

2.4.2 Configuration Guideline 30

2.4.3 Configuration Procedure 30

2.4.4 Verifying the Configuration Result 31

2.4.5 Complete Configuration 32

2.5 VXLAN BGP EVPN Distributed Gateway (Communication in the Same Subnet) 41

2.5.1 Networking and Service Description 41

2.5.2 Configuration Guideline 43

2.5.3 Configuration Procedure 43

2.5.4 Verifying the Configuration Result 47

2.5.5 Complete Configuration 48

2.6 VXLAN BGP EVPN Distributed Gateway (Communication Between Different Subnets) 59

2.6.1 Networking and Service Description 59

2.6.2 Configuration Guideline 61

2.6.3 Configuration Procedure 61

2.6.4 Verifying the Configuration Result 65
2.6.5 Complete Configuration 66
2.7 VXLAN Layer 3 Interconnection (Intra-AS Cross-DC Three-Segment VXLAN) 77
2.7.1 Networking and Service Description 77
2.7.2 Configuration Guideline 79
2.7.3 Configuration Procedure 79
2.7.4 Verifying the Configuration Result 85
2.7.5 Complete Configuration 86
3 Underlay Network Configuration 102
3.1 Objectives 102
3.2 Networking and Service Description 103
3.3 Configure the iStack Stacking on the Access Switch 104
3.3.1 Networking and Configuration Guideline 104
3.3.2 Configure the Stack Attributes for Leaf-1A and Leaf-1B 104
3.3.3 Configuring the Stack Ports 105
3.3.4 Checking the Stack Configuration 106
3.3.5 Saving the Configurations and Restarting the Device 106
3.3.6 Connecting Stack Cables to Set Up the Stack 107
3.3.7 Verifying the Configuration Result 107
3.3.8 Configuring DAD 108
3.3.9 Saving the Stack Configuration 108
3.3.10 Complete Stack Configuration 109
3.4 Layer-3 Data Center Network Configuration 110
3.4.1 Objectives 110
3.4.2 Networking and Service Description 110
3.4.3 Configuration Guideline 110
3.4.4 Experiment IP Address Planning 110
3.4.5 Configuring OSPF 112
3.4.6 OSPF Complete Configuration 115
3.5 Configuring M-LAG Work Groups for Gateway Nodes 117
3.5.1 Objectives 117
3.5.2 Networking and Service Description 117
3.5.3 Configuration Guideline 118
3.5.4 Configuring M-LAG for Spines 118
3.5.5 Configuring M-LAG for Server Leaf 123
3.6 Configuring Firewalls 125
3.6.1 Objectives 125
3.6.2 Networking and Service Description 126
3.6.3 Configuration Guideline 127
3.6.4 Device Physical Connection 127

3.6.5 Complete Configuration 130
3.7 Pre-configuring the Agile Data Center VXLAN 132
3.7.1 Objectives 132
3.7.2 Networking and Service Description 133
3.7.3 Configuring an Active-Active Loopback Interface 133
3.7.4 Configuring the NVO3 Extension Function of the Gateway 133
3.7.5 Performing Basic Configurations of VXLAN 134
3.7.6 Configuring BGP EVPN 135
3.7.7 Verifying the Active-Active Result (See section 3.5 for the configurations of DFS) 138
3.8 Configuring SNMP 138
3.8.1 Objectives 138
3.8.2 Configuring SNMP Parameters for the Switch 139
3.8.3 Configuring the SNMP Parameters for the Firewall 141
3.9 Configuring NETCONF 143
3.9.1 Objectives 143
3.9.2 Configuring NETCONF Parameters for a Switch 143
3.9.3 Configuring the NETCONF Parameters for the Firewall 144
4 Agile Controller-DCN Pre-Configuration 145
4.1 Preparing for Management 145
4.2 Networking and Agile Controller-DCN Parameter Planning 145
4.3 Discovering Devices and Links 147
4.4 Adding Device Groups 150
4.5 Creating a Fabric 152
4.6 Creating an L4-L7 Resource Pool 153
4.7 Configuring Interconnection Resources and Global Resources 155
4.8 Collecting Device Alarms 156
4.9 Creating an SFC Template 158
4.10 Deploying the Egress Network 159
4.10.1 Configuring the Access to the Public Network (Default VRF) 159
4.10.2 Configuring the Intranet of the Tenant Router 163
5 FusionStorage Block Configuration 164
5.1 Objectives 164
5.2 Networking and Service Description 164
5.3 Performing the Converged Deployment Configuration 164
5.3.1 Checking the Computing Cluster Configuration 164
5.3.2 Installing FSA 165
5.3.3 Configuring FusionStorage Block 167
5.4 Interconnecting FusionCompute to FusionStorage 170
5.4.1 Configuring the CNA Storage Port 170
5.4.2 Adding Storage Resources 171

5.4.3 Adding a Data Storage Resource 172
6 FusionSphere OpenStack Resource Configuration 176
6.1 Objectives 176
6.2 Configuring FusionCompute Resource Pools 176
6.2.1 Checking the Distributed Switch Configuration 176
6.2.2 Interconnecting VRM to FusionSphere OpenStack 176
6.3 Configuring FusionSphere OpenStack Resource Pools 179
6.3.1 Configuring Computing Clusters 179
6.3.2 Creating a Storage Cluster 181
6.4 Configuring the FusinSphere OpenStack OM 183
6.4.1 Creating an External Network 183
6.4.2 Creating a Host Group 187
6.4.3 Registering a VM Image (Configured) 190
6.4.4 Creating the VM Flavors (Configured) 191
7 SDN Cloud-Network Synergy Service Operation 194
7.1 Configuring ManageOne ServiceCenter 194
7.1.1 Configuring a Cloud Resource Pool 194
7.1.2 Create an Organization and Its Users 197
7.1.3 Creating a VDC and Its Users 200
7.1.4 Associating a VDC with an External Network 204
7.1.5 Creating a VPC 204
7.2 Commissioning the Routed Network 207
7.2.1 Prerequisites 207
7.2.2 Applying for a router 208
7.2.3 Applying for a Virtual Firewall 212
7.2.4 Creating a Routed Network 214
7.3 Commissioning the Cloud Host 220
7.3.1 Creating the Cloud Host Service 220
7.3.2 Applying for a Cloud Host 223
7.3.3 Mutual Access between Two Cloud Hosts Using the Same Subnet 230
7.3.4 Mutual Access between Two Cloud Hosts Using Different Subnets 233
7.4 Commissioning SNAT 238
7.4.1 Applying for SNAT 238
7.4.2 Configuring the Firewall Policy 242
7.5 Security Group 244
7.5.1 Creating a Security Group 244
7.5.2 Adding a Cloud Host to a Security Group 245
7.5.3 Adding a Security Group Rule 246
8 Experiment on Cloud Host Backup in Cloud Data Centers 248
8.1 Objectives 248

8.2 Networking and Service Description 248
8.3 Planning Interconnection Data 249
8.4 Configuring eBackup Manager 249
8.5 Configuring Production Storage 252
8.5.1 Configuring eBackup Server 252
8.5.2 Installing FusionStorage Agent on eBackup Server 253
8.6 Configuring Backup Storage 256
8.6.1 Configuring OceanStor V3 NAS 256
8.6.2 Configuring Backup Storage 260
8.7 Configuring CSBS Karbor 261
8.8 Applying for Cloud Host Backup 262
8.8.1 Applying for Cloud Host Backup 262
8.8.2 Using Cloud Host Backup 264
9 O&M Experiment in Cloud Data Centers 266
9.1 Performing eSight O&M 266
9.1.1 About This Experiment 266
9.1.2 Objectives 267
9.1.3 Configuring Experiment Tasks 267
9.1.4 Verifying the Result 276
9.2 Performing an O&M Experiment on the ManageOne OperationCenter System 278
9.2.1 Objectives 278
9.2.2 Planning Networking 278
9.2.3 Preparing for the Installation 279
9.2.4 Configuring the Interconnection 280
9.2.5 Installing Analysis Tools 309
9.2.6 Installing UVP VMTools 317
9.2.7 Backing Up and Recovering OperationCenter 321
9.2.8 Adding Analysis Tools 326
9.3 Agile Controller-DCN Network Service O&M and Monitoring 331
9.3.1 Three-level Topology Visibility 331
9.3.2 Network Path Detection 337
9.3.3 Loop Detection 340
9.4 Agile Controller-DCN Data Inconsistency Discovery 342

1 HCIE-Data Center Lab Environment Preparation

1.1 About the Environment

This document is a training material for trainees planning to earn the HCIE-DC V2.0 certification and engineers willing to learn technologies of Huawei cloud data center. This guide describes how to perform experiments around technologies including Large L2, SDN cloud-network synergy, distributed storage, cloud computing, and cloud host backup in the converged cloud data center scenario.

This guide includes the following experiments:

- VXLAN Feature Experiment
- Underlay Network Pre-configuration
- Agile Controller-DCN 3.0 Pre-configuration
- FusionStorage Block Configuration
- FusionSphere OpenStack Resource Configuration
- SDN Cloud-Network Synergy Service Operation
- Cloud Data Center Cloud Host Backup Experiment
- Cloud Data Center O&M Experiment

1.2 Lab Networking

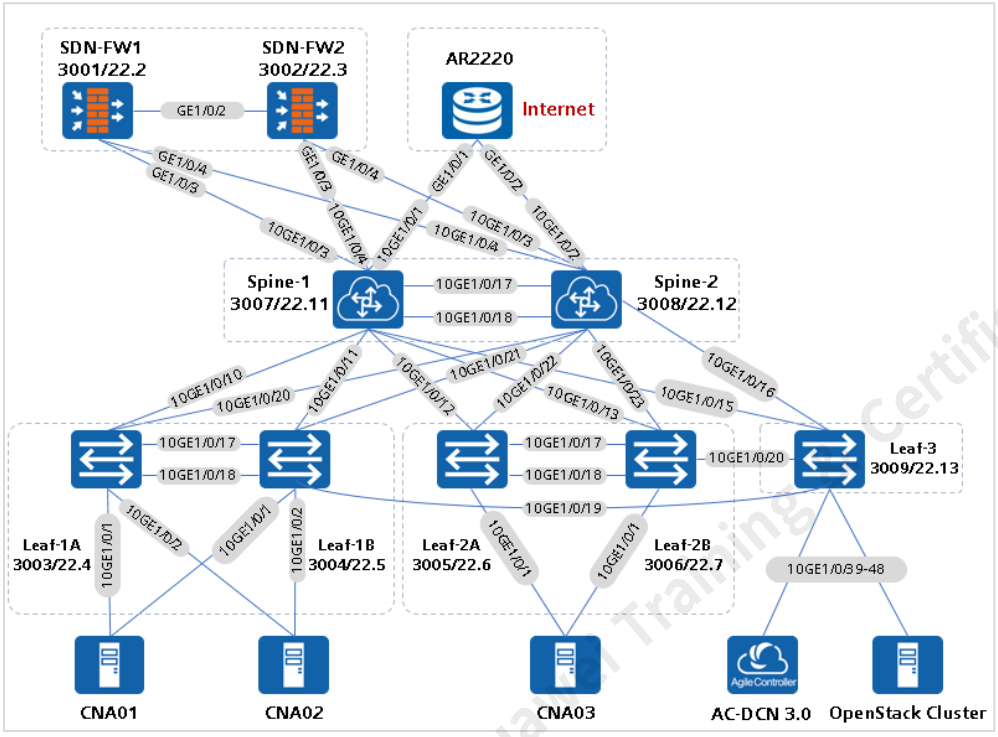

Figure 1-1 Overall lab network topology

1.3 Device Version and Configuration Information

Table 1-1 Lab device list

Device and Software Type		Device Model and Software Name	Quantity	Version
Hardware	Network	CE6850-HI	4	V200R002C50
		CE6851-HI	1	
		CE12800	2	
	Security	USG 6620	2	V500R001C60
	Server	RH 2288V3	8	-
Software	Controller	Agile Controller-DCN 3.0	1	V300R001C20SPC800
	Cloud computing	FusionSphere OpenStack	1	V100R006C10SPC600
		eSight	1	V300R007C00SPC500

Device and Software Type		Device Model and Software Name	Quantity	Version
		CNA	3	V100R006C10SPC101
		FusionStorage	1	V100R006C10SPC200
		ManageOne OC	1	3.0.9

1.4 Device Login Information

Table 1-2 Device login information

Device Information	Console Port and Login Password	Management Network Port	Management IP Address
Console Server	Huawei@123	Vlanif4060	172.21.22.1
Spine-1	3007/Huawei@123	MEth0/0/0	172.21.22.11
Spine-2	3008/Huawei@123	MEth0/0/0	172.21.22.12
SDN-FW1	3001/Huawei@123	GE1/0/1	172.21.22.2
SDN-FW2	3002/Huawei@123	GE1/0/1	172.21.22.3
Leaf-1A	3003/Huawei@123	MEth0/0/0	172.21.22.4
Leaf-1B	3004/Huawei@123	MEth0/0/0	172.21.22.5
Leaf-2A	3005/Huawei@123	MEth0/0/0	172.21.22.6
Leaf-2B	3006/Huawei@123	MEth0/0/0	172.21.22.7
Leaf-3	300/Huawei@Admin	Vlanif4060	172.21.22.13
FusionSphere CPS	admin/Huawei12#$	-	https://192.168.0.2:8890
FusionSphere OM	cloud_admin/FusionSphere123	-	https://192.168.0.150:643
AC-DCN	admin/Huawei@123	-	https://192.168.4.4:18002
ManageOne SC	cloud_admin/FusionSphere123	-	https://192.168.0.170
ManageOne OC	admin/Huawei@123	-	https://192.168.0.180
FusionStorage	admin/Huawei12#$	-	https://10.1.0.190:28443

Device Information	Console Port and Login Password	Management Network Port	Management IP Address
FusionCompute	admin/Huawei@123	-	https://10.1.0.160:8443

The local PC resides on the 172.21.0.0/16 network segment of the lab network. You can log in to the network device using the console client. The Leaf-3 is the redirection device of all IT software. You need to add a route to the specified network segment locally as follows:

To add a route to FusionSphere CPS, enter **route add 192.168.0.0 mask 255.255.255.0 172.21.51.7** in Windows CMD.

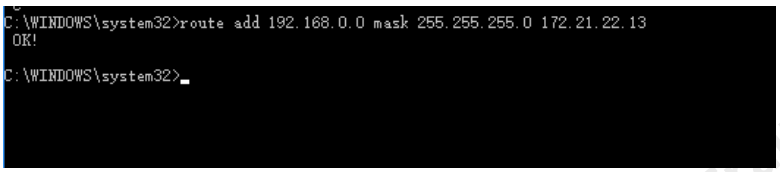

Add the routes to the network segments 192.168.4.0/24 and 10.9.1.0/24 in the same way.

NOTE

Run the CMD as an administrator.

2 VXLAN Feature Experiment

2.1 Objectives

- Understand how to configure VXLAN Layer 2 interconnection.
- Understand how to configure VXLAN Layer 3 interconnection.
- Understand how to configure the centralized gateway of the VXLAN BGP EVPN.
- Understand how to configure the distributed gateway of the VXLAN BGP EVPN.

2.2 Configuring VXLAN Layer 2 Interconnection (No-Tunnel Mode)

2.2.1 Networking and Service Description

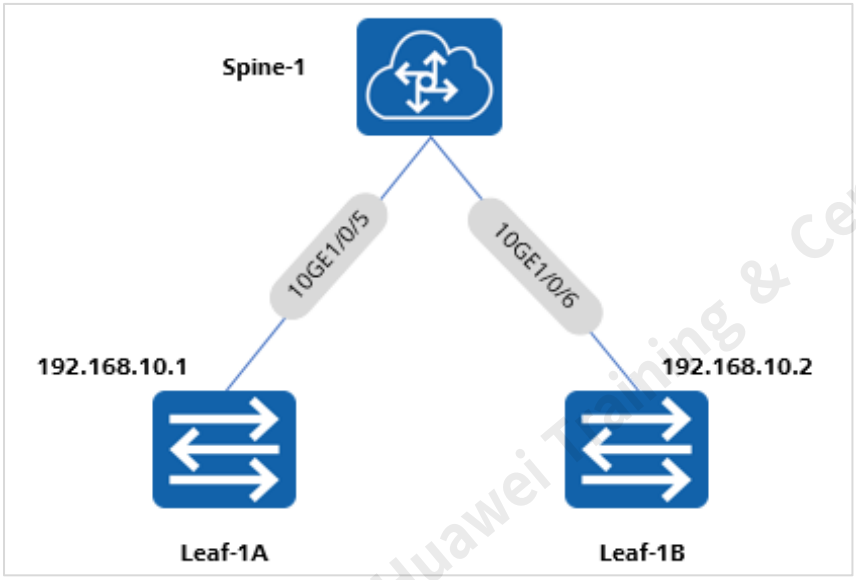

Figure 2-1 Network diagram of the VXLAN Layer 2 interconnection experiment

Based on user requirements, VXLAN is deployed on Spine-1 only, and two switches Leaf-1A and Leaf-1B are used to simulate PCs so that VXLAN can work as VLAN, implementing Layer 2 intercommunication.

2.2.2 Configuration Guideline

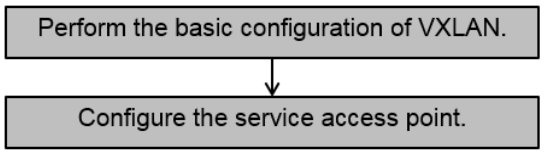

2.2.3 Configuration Procedure

Step 1 **Configuring the IP addresses of Leaf-1 and Leaf-2**

```
<HUAWEI> system-view
[~HUAWEI] sysname Leaf-1
[*Leaf-1A] commit
[~Leaf-1A] interface 10ge1/0/5
[*Leaf-1A-10GE1/0/5] undo portswitch
```

```
[*Leaf-1A-10GE1/0/5] ip address 192.168.10.1 24
[*Leaf-1A-10GE1/0/5] quit
[*Leaf-1A] commit

<HUAWEI> system-view
[~HUAWEI] sysname Leaf-1B
[*Leaf-1B] commit
[~Leaf-1B] interface 10ge1/0/6
[*Leaf-1B-10GE1/0/6] undo portswitch
[*Leaf-1B-10GE1/0/6] ip address 192.168.10.2 24
[*Leaf-1B-10GE1/0/6] quit
[*Leaf-1B] commit
```

Step 2 **Performing the basic configurations of VXLAN for Spine-1**

```
<HUAWEI> system-view
[~HUAWEI] sysname Spine-1
[*Spine-1] commit
[~Spine-1] bridge-domain 10
[*Spine-1-bd10] vxlan vni 5000
[*Spine-1-bd10] quit
[*Spine-1] commit
```

Step 3 **Configuring the service access point of Spine-1**

```
[~Spine-1] bridge-domain 10
[*Spine-1-bd10] quit
[*Spine-1] interface 10ge 1/0/5.1 mode l2
[*Spine-1-10GE1/0/5.1] encapsulation  untag
[*Spine-1-10GE1/0/5.1] bridge-domain 10
[*Spine-1-10GE1/0/5.1] quit
[*Spine-1] interface 10ge 1/0/6.1 mode l2
[*Spine-1-10GE1/0/6.1] encapsulation  untag
[*Spine-1-10GE1/0/6.1] bridge-domain 10
[*Spine-1-10GE1/0/6.1] quit
[*Spine-1] commit
----End
```

2.2.4 Verifying the Configuration Result

Run the Ping command to test the network connectivity. The following uses Leaf-1A Ping Leaf-1B as an example:

```
<Leaf-1A>ping 192.168.10.2
  PING 192.168.10.2: 56  data bytes, press CTRL_C to break
    Reply from 192.168.10.2: bytes=56 Sequence=1 ttl=255 time=60 ms
    Reply from 192.168.10.2: bytes=56 Sequence=2 ttl=255 time=70 ms
    Reply from 192.168.10.2: bytes=56 Sequence=3 ttl=255 time=50 ms
    Reply from 192.168.10.2: bytes=56 Sequence=4 ttl=255 time=50 ms
    Reply from 192.168.10.2: bytes=56 Sequence=5 ttl=255 time=50 ms

  --- 192.168.10.2 ping statistics ---
    5 packet(s) transmitted
```

```
    5 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 50/56/70 ms
```

2.2.5 Complete Configuration

Complete configuration of Spine-1:

```
#
sysname Spine-1
#
system tcam ED-extend slot 1
#
system tcam ED-extend slot 2
#
device board 5 board-type CE-MPUA-S
device board 6 board-type CE-MPUA-S
device board 1 board-type CE-L24XS-ED
device board 2 board-type CE-L24XS-ED
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
#
bridge-domain 10
 vxlan vni 5000
#
aaa
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface MEth0/0/0
#
interface 10GE1/0/0
#
interface 10GE1/0/1
 device transceiver 10GBASE-COPPER
```

```
#
interface 10GE1/0/2
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/3
 device transceiver 10GBASE-COPPER
#
          interface 10GE1/0/4
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/5
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/5.1 mode l2
 encapsulation dot1q untag
 bridge-domain 10
#
interface 10GE1/0/6
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/6.1 mode l2
 encapsulation dot1q untag
 bridge-domain 10
#
interface 10GE1/0/7
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/8
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/9
#
interface 10GE1/0/10
#
interface 10GE1/0/11
#
interface 10GE1/0/12
#
interface 10GE1/0/13
#
interface 10GE1/0/14
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
#
interface 10GE1/0/16
#
interface 10GE1/0/17
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
```

```
interface 10GE1/0/20
#
interface 10GE1/0/21
#
interface 10GE1/0/22
#
interface 10GE1/0/23
#
interface 10GE2/0/0
#
interface 10GE2/0/1
#
interface 10GE2/0/2
#
interface 10GE2/0/3
#
interface 10GE2/0/4
#
interface 10GE2/0/5
#
interface 10GE2/0/6
#
interface 10GE2/0/7
#
interface 10GE2/0/8
#
interface 10GE2/0/9
#
interface 10GE2/0/10
#
interface 10GE2/0/11
#
interface 10GE2/0/12
#
interface 10GE2/0/13
#
interface 10GE2/0/14
 device transceiver 10GBASE-COPPER
#
interface 10GE2/0/15
#
interface 10GE2/0/16
#
interface 10GE2/0/17
#
interface 10GE2/0/18
#
interface 10GE2/0/19
#
interface 10GE2/0/20
#
interface 10GE2/0/21
#
interface 10GE2/0/22
#
```

```
interface 10GE2/0/23
#
interface Sip5/0/0
#
interface Sip5/0/1
#
interface Sip6/0/0
#
interface Sip6/0/1
#
interface NULL0
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher
$1c$gu5KXvs`=A$ulD11S\`=TB4B_<5r~1%X}jnDz^1{6qM!jDc~WuK$
#
vm-manager
#
Return
```

Complete configuration of Leaf-1A:

```
#
sysname Leaf-1A
#
system resource standard
#
device board 1 board-type CE6851-48S6Q-HI
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
#
aaa
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
```

```
 #
 domain default
 #
 domain default_admin
#
stack
#
interface MEth0/0/0
#
interface 10GE1/0/1
 device transceiver 1000BASE-T
#
interface 10GE1/0/2
#
interface 10GE1/0/3
 device transceiver 1000BASE-T
#
interface 10GE1/0/4
 device transceiver 1000BASE-T
#
interface 10GE1/0/5
 undo portswitch
 ip address 192.168.10.1 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/6
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/7
#
interface 10GE1/0/8
#
interface 10GE1/0/9
#
interface 10GE1/0/10
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/11
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/12
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/13
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/14
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/16
 device transceiver 10GBASE-COPPER
#
```

```
interface 10GE1/0/17
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
interface 10GE1/0/20
#
interface 10GE1/0/21
#
interface 10GE1/0/22
#
interface 10GE1/0/23
#
interface 10GE1/0/24
#
interface 10GE1/0/25
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/26
#
interface 10GE1/0/27
#
interface 10GE1/0/28
#
interface 10GE1/0/29
#
interface 10GE1/0/30
#
interface 10GE1/0/31
#
interface 10GE1/0/32
#
interface 10GE1/0/33
#
interface 10GE1/0/34
#
interface 10GE1/0/35
#
interface 10GE1/0/36
#
interface 10GE1/0/37
#
interface 10GE1/0/38
#
interface 10GE1/0/39
#
interface 10GE1/0/40
#
interface 10GE1/0/41
#
interface 10GE1/0/42
#
```

```
interface 10GE1/0/43
#
interface 10GE1/0/44
#
interface 10GE1/0/45
#
interface 10GE1/0/46
#
interface 10GE1/0/47
#
interface 10GE1/0/48
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface 40GE1/0/3
#
interface 40GE1/0/4
#
interface 40GE1/0/5
#
interface 40GE1/0/6
#
interface NULL0
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher
$1c$9c'I!EAXZ6$#_5B%G|9uTT[z9B\_AQ9anfe;hd(OW~i`^/XW21P$
#
vm-manager
#
Return
```

Complete configuration of Leaf-1B:

```
#
sysname Leaf-1B
#
system resource standard
#
device board 1 board-type CE6851-48S6Q-HI
#
drop-profile default
#
dcb pfc
#
```

```
dcb ets-profile default
#
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
#
aaa
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface MEth0/0/0
#
interface 10GE1/0/1
#
interface 10GE1/0/2
 device transceiver 1000BASE-T
#
interface 10GE1/0/3
 device transceiver 1000BASE-T
#
interface 10GE1/0/4
 device transceiver 1000BASE-T
#
interface 10GE1/0/5
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/6
 undo portswitch
 ip address 192.168.10.2 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/7
#
interface 10GE1/0/8
#
interface 10GE1/0/9
#
interface 10GE1/0/10
#
interface 10GE1/0/11
#
interface 10GE1/0/12
#
interface 10GE1/0/13
```

```
#
interface 10GE1/0/14
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/16
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
interface 10GE1/0/20
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/21
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/22
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/23
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/24
#
interface 10GE1/0/25
#
interface 10GE1/0/26
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/27
#
interface 10GE1/0/28
#
interface 10GE1/0/29
#
interface 10GE1/0/30
#
interface 10GE1/0/31
#
interface 10GE1/0/32
#
interface 10GE1/0/33
#
interface 10GE1/0/34
#
interface 10GE1/0/35
#
```

```
interface 10GE1/0/36
#
interface 10GE1/0/37
#
interface 10GE1/0/38
#
interface 10GE1/0/39
#
interface 10GE1/0/40
#
interface 10GE1/0/41
#
interface 10GE1/0/42
#
interface 10GE1/0/43
#
interface 10GE1/0/44
#
interface 10GE1/0/45
#
interface 10GE1/0/46
#
interface 10GE1/0/47
#
interface 10GE1/0/48
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface 40GE1/0/3
#
interface 40GE1/0/4
#
interface 40GE1/0/5
#
interface 40GE1/0/6
#
interface NULL0
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher
$1c$4;5h;(QD"&$u.nRPe#@=TyyJ<#`4>|CC}$*'9@:[~X"ms=RfCpB$
#
vm-manager
#
return
```

2.3 Configuring VXLAN Layer 2 Interconnection (Tunnel Mode)

2.3.1 Networking and Service Description

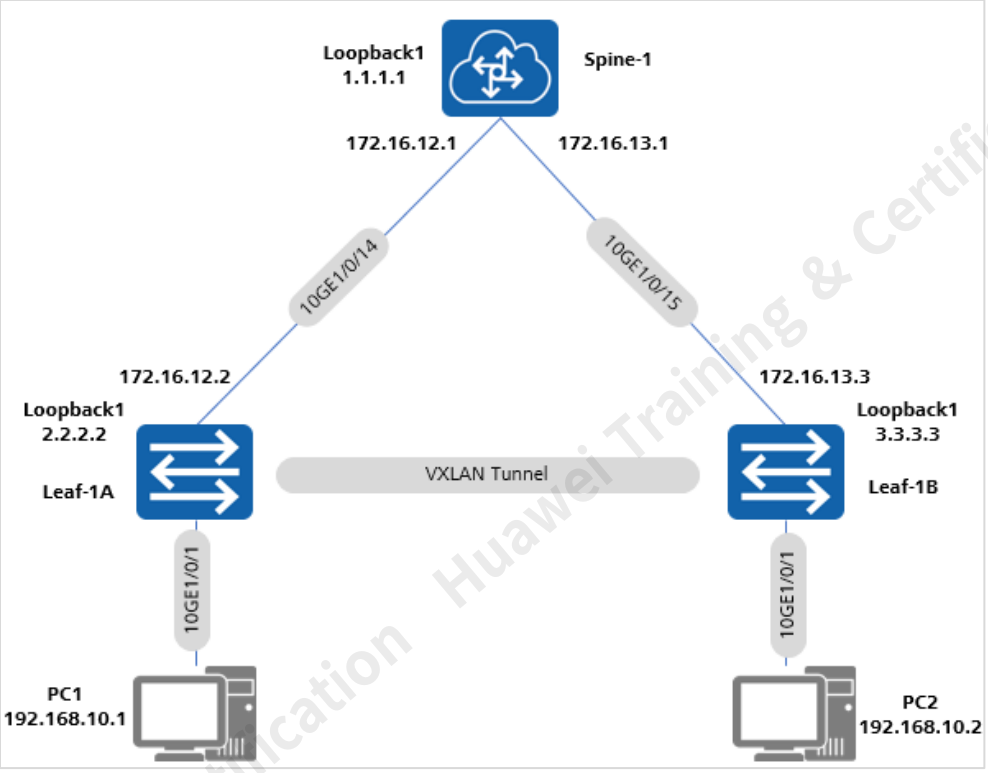

Figure 2-2 Network diagram of the VXLAN Layer 2 interconnection experiment

Based on user requirements, VXLAN is deployed on Leaf-1A and Leaf-1B, and static VXLAN tunnels are established between Leaf-1A and Leaf-1B to enable the interconnection between PC1 and PC2 on the same network segment.

2.3.2 Configuration Guideline

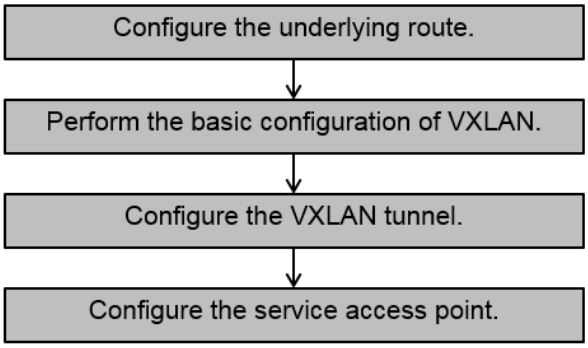

2.3.3 Configuration Procedure

Step 1 **Configuring the IP addresses of Spine-1, Leaf-1 and Leaf-2**

```
<HUAWEI>system-view immediately
[HUAWEI] sysname Spine-1
[Spine-1]interface  LoopBack  1
[Spine-1-LoopBack1]ip address  1.1.1.1 32
[Spine-1-LoopBack1]quit
[Spine-1]interface 10GE 1/0/14
[Spine-1-10GE1/0/14]undo portswitch
[Spine-1-10GE1/0/14]ip address  172.16.12.1 24
[Spine-1-10GE1/0/14]quit
[Spine-1]interface 10GE 1/0/15
[Spine-1-10GE1/0/15]undo portswitch
[Spine-1-10GE1/0/15]ip address  172.16.13.1 24
[Spine-1-10GE1/0/15]quit

<HUAWEI>system-view immediately
[HUAWEI]sysname Leaf-1A
[Leaf-1A]interface LoopBack 1
[Leaf-1A-LoopBack1]ip address 2.2.2.2 32
[Leaf-1A-LoopBack1]quit
[Leaf-1A]interface 10
[Leaf-1A]interface 10GE 1/0/14
[Leaf-1A-10GE1/0/14]undo portswitch
[Leaf-1A-10GE1/0/14]ip address  172.12.12.2 24
[Leaf-1A-10GE1/0/14]quit

<HUAWEI>system-view immediately
[HUAWEI]sysname Leaf-1B
[Leaf-1B] interface LoopBack  1
[Leaf-1B-LoopBack1]ip address  3.3.3.3 32
[Leaf-1B-LoopBack1]quit
[Leaf-1B]interface 10GE 1/0/15
[Leaf-1B-10GE1/0/15]undo portswitch
[Leaf-1B-10GE1/0/15]ip address 172.16.13.2 24
[Leaf-1B-10GE1/0/15]quit
```

Step 2 Configuring OSPF

```
Configure the loopback interface LoopBack1 on Leaf-1A and Leaf-1B, and
advertise its IP address (considered as the VTEP IP) using the underlay
routing protocol.
<Spine-1>system-view  immediately
[Spine-1]ospf 1
[Spine-1-ospf-1]area 0
[Spine-1-ospf-1-area-0.0.0.0]network  1.1.1.1 0.0.0.0
[Spine-1-ospf-1-area-0.0.0.0]network  172.16.12.0 0.0.0.255
[Spine-1-ospf-1-area-0.0.0.0]network  172.16.13.0 0.0.0.255
[Spine-1-ospf-1-area-0.0.0.0]return

<Leaf-1A>system-view  immediately
[Leaf-1A]ospf 1
[Leaf-1A-ospf-1]area 0
[Leaf-1A-ospf-1-area-0.0.0.0]net
[Leaf-1A-ospf-1-area-0.0.0.0]network  2.2.2.2 0.0.0.0
[Leaf-1A-ospf-1-area-0.0.0.0]network  172.16.12.0 0.0.0.255
[Leaf-1A-ospf-1-area-0.0.0.0]return

<Leaf-1B>system-view  immediately
[Leaf-1B]ospf 1
[Leaf-1B-ospf-1]area 0
[Leaf-1B-ospf-1-area-0.0.0.0]network 3.3.3.3 0.0.0.0
[Leaf-1B-ospf-1-area-0.0.0.0]network  172.16.13.0 0.0.0.255
[Leaf-1B-ospf-1-area-0.0.0.0]return
```

Step 3 Performing the basic configurations of VXLAN for Leaf-1A and Leaf-1B

```
[Leaf-1A]bridge-domain 10
[Leaf-1A-bd10]vxlan vni 5000
[Leaf-1A-bd10]quit
[Leaf-1A]interface Nve 1
[Leaf-1A-Nve1]source  2.2.2.2
[Leaf-1A-Nve1]vni 5000 head-end peer-list 3.3.3.3
[Leaf-1A-Nve1]quit

[Leaf-1B]bridge-domain 10
[Leaf-1B-bd10]vxlan vni 5000
[Leaf-1B-bd10]quit
[Leaf-1B]interface Nve 1
[Leaf-1B-Nve1]source 3.3.3.3
[Leaf-1B-Nve1]vni 5000 head-end peer-list 2.2.2.2
[Leaf-1B-Nve1]quit
```

Step 4 Configuring the service access points for Leaf-1A and Leaf-1B

```
[Leaf-1A] interface 10ge 1/0/1.1 mode l2
[Leaf-1A-10GE1/0/1.1] encapsulation dot1q vid 10
[Leaf-1A-10GE1/0/1.1] bridge-domain 10
[Leaf-1A-10GE1/0/1.1] quit
```

```
[Leaf-1B] interface 10ge 1/0/1.1 mode l2
[Leaf-1B-10GE1/0/1.1] encapsulation dot1q vid 10
[Leaf-1B-10GE1/0/1.1] bridge-domain 10
[Leaf-1B-10GE1/0/1.1] quit
----End
```

2.3.4 Verifying the Configuration Result

Run the **display vxlan tunnel** command to view VXLAN tunnel information. The following uses the command output of Spine-1 as an example:

```
<Leaf-1A>display vxlan vni
Number of vxlan vni : 1
VNI          Bridge-domain-ID          State
--------------------------------------
5000         10            up
<Leaf-1A>display vxlan tunnel
Number of vxlan tunnel : 1
Tunnel ID   Source              Destination         State  Type     Uptime
--------------------------------------------------------------------------------
--
4026531841  2.2.2.2             3.3.3.3             up     static   00:19:21
```

Run the Ping command to test the network connectivity. The following uses PC1 Ping PC2 as an example:

```
C:\Users\admin\Desktop>ping 192.168.10.2

Pinging 192.168.10.2 with 32 bytes of data:
Reply from 192.168.10.2: bytes=32 time<1ms TTL=128
Reply from 192.168.10.2: bytes=32 time<1ms TTL=128
Reply from 192.168.10.2: bytes=32 time<1ms TTL=128
Reply from 192.168.10.2: bytes=32 time<1ms TTL=128

Ping statistics for 192.168.10.2:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

2.3.5 Complete Configuration

Complete configuration of Spine-1:

```
#
sysname Spine-1
#
system resource standard
#
device board 1 board-type CE6851-48S6Q-HI
#
drop-profile default
#
dcb pfc
```

```
#
dcb ets-profile default
#
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
#
aaa
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface MEth0/0/0
#
interface 10GE1/0/1
#
interface 10GE1/0/2
 device transceiver 1000BASE-T
#
interface 10GE1/0/3
 device transceiver 1000BASE-T
#
interface 10GE1/0/4
 device transceiver 1000BASE-T
#
interface 10GE1/0/5
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/6
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/7
#
interface 10GE1/0/8
#
interface 10GE1/0/9
#
interface 10GE1/0/10
#
interface 10GE1/0/11
#
interface 10GE1/0/12
#
interface 10GE1/0/13
#
```

```
interface 10GE1/0/14
 undo portswitch
 ip address 172.16.12.1 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
 undo portswitch
 ip address 172.16.13.1 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/16
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
interface 10GE1/0/20
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/21
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/22
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/23
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/24
#
interface 10GE1/0/25
#
interface 10GE1/0/26
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/27
#
interface 10GE1/0/28
#
interface 10GE1/0/29
#
interface 10GE1/0/30
#
interface 10GE1/0/31
#
interface 10GE1/0/32
#
interface 10GE1/0/33
#
interface 10GE1/0/34
```

```
#
interface 10GE1/0/35
#
interface 10GE1/0/36
#
interface 10GE1/0/37
#
interface 10GE1/0/38
#
interface 10GE1/0/39
#
interface 10GE1/0/40
#
interface 10GE1/0/41
#
interface 10GE1/0/42
#
interface 10GE1/0/43
#
interface 10GE1/0/44
#
interface 10GE1/0/45
#
interface 10GE1/0/46
#
interface 10GE1/0/47
#
interface 10GE1/0/48
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface 40GE1/0/3
#
interface 40GE1/0/4
#
interface 40GE1/0/5
#
interface 40GE1/0/6
#
interface LoopBack1
 ip address 1.1.1.1 255.255.255.255
#
interface NULL0
#
ospf 1
 area 0.0.0.0
  network 1.1.1.1 0.0.0.0
  network 172.16.12.0 0.0.0.255
  network 172.16.13.0 0.0.0.255
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
```

```
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher
$1c$4;5h;(QD"&$u.nRPe#@=TyyJ<#`4>|CC}$*'9@:[~X"ms=RfCpB$
#
vm-manager
#
Return
```

Complete configuration of Leaf-1A:

```
#
sysname Leaf-1A
#
system resource standard
#
device board 1 board-type CE6850U-24S2Q-HI
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
#
bridge-domain 10
 vxlan vni 5000
#
aaa
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface MEth0/0/0
#
interface 10GE1/0/1
 device transceiver 10GBASE-COPPER
```

```
#
interface 10GE1/0/1.1 mode l2
 encapsulation dot1q vid 10
 bridge-domain 10
#
interface 10GE1/0/2
#
interface 10GE1/0/3
#
interface 10GE1/0/4
#
interface 10GE1/0/5
#
interface 10GE1/0/6
#
interface 10GE1/0/7
#
interface 10GE1/0/8
#
interface 10GE1/0/9
#
interface 10GE1/0/10
#
interface 10GE1/0/11
#
interface 10GE1/0/12
#
interface 10GE1/0/13
#
interface 10GE1/0/14
 undo portswitch
 ip address 172.16.12.2 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/16
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
interface 10GE1/0/20
#
interface 10GE1/0/21
#
interface 10GE1/0/22
#
interface 10GE1/0/23
```

```
#
interface 10GE1/0/24
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface LoopBack0
 ip address 2.2.2.2 255.255.255.255
#
interface Nve1
 source 2.2.2.2
 vni 5000 head-end peer-list 3.3.3.3
#
interface NULL0
#
ospf 1
 area 0.0.0.0
  network 2.2.2.2 0.0.0.0
  network 172.16.12.0 0.0.0.255
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher
$1c$Y@4\VN@!#P$#W=zV;c6j([Uu6FCkrLIlz>='%|rAHA4-EXzRrg($
#
vm-manager
#
Return
```

Complete configuration of Leaf-1B:

```
#
sysname Leaf-1B
#
system resource standard
#
device board 1 board-type CE6850U-24S2Q-HI
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
telnet server disable
telnet ipv6 server disable
```

```
#
diffserv domain default
#
bridge-domain 10
 vxlan vni 5000
#
aaa
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface MEth0/0/0
#
interface 10GE1/0/1
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/1.1 mode l2
 encapsulation dot1q vid 10
 bridge-domain 10
#
interface 10GE1/0/2
#
interface 10GE1/0/3
#
interface 10GE1/0/4
#
interface 10GE1/0/5
#
interface 10GE1/0/6
#
interface 10GE1/0/7
#
interface 10GE1/0/8
#
interface 10GE1/0/9
#
interface 10GE1/0/10
#
interface 10GE1/0/11
#
interface 10GE1/0/12
#
interface 10GE1/0/13
#
interface 10GE1/0/14
 device transceiver 10GBASE-COPPER
```

```
#
interface 10GE1/0/15
 undo portswitch
 ip address 172.16.13.2 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/16
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
interface 10GE1/0/20
#
interface 10GE1/0/21
#
interface 10GE1/0/22
#
interface 10GE1/0/23
#
interface 10GE1/0/24
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface LoopBack1
 ip address 3.3.3.3 255.255.255.255
#
interface Nve1
 source 3.3.3.3
 vni 5000 head-end peer-list 2.2.2.2
#
interface NULL0
#
ospf 1
 area 0.0.0.0
  network 3.3.3.3 0.0.0.0
  network 172.16.13.0 0.0.0.255
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
```

```
 set authentication password cipher
$1c$^y4|C=)iR6$tB)qKME3!!R8ZNLHgoMPYMW!=5c3j@(q);D-S"q0$
#
vm-manager
#
Return
```

2.4 Configuring VXLAN Layer 3 Interconnection (No-Tunnel Mode)

2.4.1 Networking and Service Description

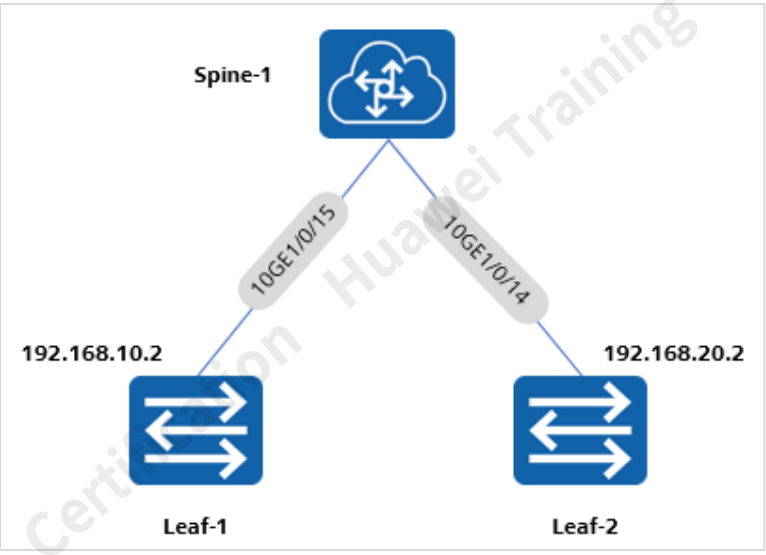

Figure 2-3 Network diagram of the VXLAN Layer 3 interconnection experiment

Based on user requirements, VXLAN is deployed on Spine-1 only, and two switches Leaf-1 and Leaf-2 are used to simulate PCs so that VXLAN can work as VLAN, implementing Layer 3 intercommunication.

2.4.2 Configuration Guideline

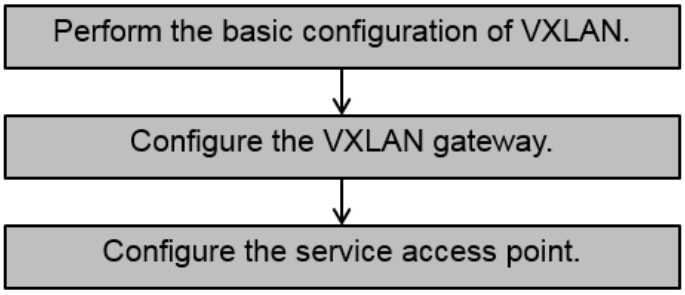

2.4.3 Configuration Procedure

Step 1 **Configuring the IP addresses of Leaf-1 and Leaf-2**

```
<HUAWEI> system-view
[~HUAWEI] sysname Leaf-1
[*Leaf-1] commit
[~Leaf-1] interface 10ge1/0/15
[*Leaf-1-10GE1/0/15] undo portswitch
[*Leaf-1-10GE1/0/15] ip address 192.168.10.2 24
[*Leaf-1-10GE1/0/15] quit
[*Leaf-1] commit

<HUAWEI> system-view
[~HUAWEI] sysname Leaf-2
[*Leaf-2] commit
[~Leaf-2] interface 10ge1/0/14
[*Leaf-2-10GE1/0/14] undo portswitch
[*Leaf-2-10GE1/0/14] ip address 192.168.20.2 24
[*Leaf-2-10GE1/0/14] quit
[*Leaf-2] commit
```

Step 2 **Performing the basic configurations of VXLAN for Spine-1**

```
<HUAWEI> system-view
[~HUAWEI] sysname Spine-1
[*Spine-1] commit
[~Spine-1] bridge-domain 10
[*Spine-1-bd10] vxlan vni 10
[*Spine-1-bd10] quit
[*Spine-1] commit
[~Spine-1] bridge-domain 20
[*Spine-1-bd20] vxlan vni 20
[*Spine-1-bd20] quit
[*Spine-1] commit
[Spine-1]interface Vbdif10
[Spine-1-Vbdif10]ip address  192.168.10.1 24
[Spine-1-Vbdif10]quit
```

```
[Spine-1]interface Vbdif20
[Spine-1-Vbdif20]ip address  192.168.20.1 24
[Spine-1-Vbdif20]quit
```

Step 3 **Configuring the default routes for the Leaf nodes**

```
[Leaf-1]ip route-static 0.0.0.0 0 192.168.10.1
[Leaf-2]ip route-static 0.0.0.0 0 192.168.20.2
```

Step 4 **Configuring the service access point of Spine-1**

```
[~Spine-1] bridge-domain 10
[*Spine-1-bd10] quit
[*Spine-1] interface 10ge 1/0/15.1 mode l2
[*Spine-1-10GE1/0/15.1] encapsulation  untag
[*Spine-1-10GE1/0/15.1] bridge-domain 10
[*Spine-1-10GE1/0/15.1] quit
[*Spine-1] interface 10ge 1/0/14.1 mode l2
[*Spine-1-10GE1/0/14.1] encapsulation  untag
[*Spine-1-10GE1/0/14.1] bridge-domain 20
[*Spine-1-10GE1/0/14.1] quit
[*Spine-1] commit
```

Enable the interface loopback function on the distributed gateway (not required for CE12800). If the gateway is deployed on the ToR CE6850, enable the interface loopback function on the device and add an unused interface to the Eth-Trunk. Otherwise, the gateway cannot be accessed.

```
[Spine-1] interface Eth-Trunk1
[Spine-1] trunkport 10ge 1/0/6
[Spine-1-Eth-Trunk1] service type tunnel
----End
```

2.4.4 Verifying the Configuration Result

Run the Ping command to test the network connectivity. The following uses Leaf-1 Ping Leaf-2 as an example:

```
<Leaf-1>ping 192.168.20.2
  PING 192.168.20.2: 56  data bytes, press CTRL_C to break
    Reply from 192.168.20.2: bytes=56 Sequence=1 ttl=255 time=60 ms
    Reply from 192.168.20.2: bytes=56 Sequence=2 ttl=255 time=70 ms
    Reply from 192.168.20.2: bytes=56 Sequence=3 ttl=255 time=50 ms
    Reply from 192.168.20.2: bytes=56 Sequence=4 ttl=255 time=50 ms
    Reply from 192.168.20.2: bytes=56 Sequence=5 ttl=255 time=50 ms

  --- 192.168.20.2 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 50/56/70 ms
```

2.4.5 Complete Configuration

Complete configuration of Spine-1:

```
#
sysname Spine-1
#
system resource standard
#
device board 1 board-type CE6850U-24S2Q-HI
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
#
bridge-domain 10
 vxlan vni 10
#
bridge-domain 20
 vxlan vni 20
#
aaa
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface Vbdif10
 ip address 192.168.10.1 255.255.255.0
#
interface Vbdif20
 ip address 192.168.20.1 255.255.255.0
#
interface MEth0/0/0
#
interface Eth-Trunk1
 service type tunnel
#
interface 10GE1/0/1
```

```
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/2
#
interface 10GE1/0/3
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/4
#
interface 10GE1/0/5
#
interface 10GE1/0/6
 eth-trunk 1
#
interface 10GE1/0/7
#
interface 10GE1/0/8
#
interface 10GE1/0/9
#
interface 10GE1/0/10
#
interface 10GE1/0/11
#
interface 10GE1/0/12
#
interface 10GE1/0/13
#
interface 10GE1/0/14
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/14.1 mode l2
 encapsulation untag
 bridge-domain 20
#
interface 10GE1/0/15
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15.1 mode l2
 encapsulation untag
 bridge-domain 10
#
interface 10GE1/0/16
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 shutdown
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 shutdown
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
```

```
interface 10GE1/0/20
#
interface 10GE1/0/21
#
interface 10GE1/0/22
#
interfce 10GE1/0/23
#
interface 10GE1/0/24
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface NULL0
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher $1c$6X>#PH`<!X$\L`VWwG8D$[4q^:oAf,6-
+op0=Fh5A@W2P6NO[dT$
#
vm-manager
#
Return
```

Complete configuration of Leaf-1:

```
#
sysname Leaf-1
#
system resource standard
#
device board 1 board-type CE6851-48S6Q-HI
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
#
aaa
 #
 authentication-scheme default
```

```
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface MEth0/0/0
#
interface 10GE1/0/1
 device transceiver 1000BASE-T
#
interface 10GE1/0/2
#
interface 10GE1/0/3
 device transceiver 1000BASE-T
#
interface 10GE1/0/4
 device transceiver 1000BASE-T
#
interface 10GE1/0/5
device transceiver 10GBASE-COPPER
#
interface 10GE1/0/6
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/7
#
interface 10GE1/0/8
#
interface 10GE1/0/9
#
interface 10GE1/0/10
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/11
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/12
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/13
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/14
device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
undo portswitch
 ip address 192.168.10.1 255.255.255.0
 device transceiver 10GBASE-COPPER
```

```
#
interface 10GE1/0/16
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
interface 10GE1/0/20
#
interface 10GE1/0/21
#
interface 10GE1/0/22
#
interface 10GE1/0/23
#
interface 10GE1/0/24
#
interface 10GE1/0/25
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/26
#
interface 10GE1/0/27
#
interface 10GE1/0/28
#
interface 10GE1/0/29
#
interface 10GE1/0/30
#
interface 10GE1/0/31
#
interface 10GE1/0/32
#
interface 10GE1/0/33
#
interface 10GE1/0/34
#
interface 10GE1/0/35
#
interface 10GE1/0/36
#
interface 10GE1/0/37
#
interface 10GE1/0/38
#
interface 10GE1/0/39
#
interface 10GE1/0/40
#
```

```
interface 10GE1/0/41
#
interface 10GE1/0/42
#
interface 10GE1/0/43
#
interface 10GE1/0/44
#
interface 10GE1/0/45
#
interface 10GE1/0/46
#
interface 10GE1/0/47
#
interface 10GE1/0/48
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface 40GE1/0/3
#
interface 40GE1/0/4
#
interface 40GE1/0/5
#
interface 40GE1/0/6
#
interface NULL0
#
ip route-static 0.0.0.0 0.0.0.0 192.168.10.1
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher
$1c$9c'I!EAXZ6$#_5B%G|9uTT[z9B\_AQ9anfe;hd(OW~i`^/XW21P$
#
vm-manager
#
Return
```

Complete configuration of Leaf-2:

```
#
sysname Leaf-2
#
system resource standard
```

```
#
device board 1 board-type CE6851-48S6Q-HI
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
#
aaa
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface MEth0/0/0
#
interface 10GE1/0/1
#
interface 10GE1/0/2
 device transceiver 1000BASE-T
#
interface 10GE1/0/3
 device transceiver 1000BASE-T
#
interface 10GE1/0/4
 device transceiver 1000BASE-T
#
interface 10GE1/0/5
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/6
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/7
#
interface 10GE1/0/8
#
interface 10GE1/0/9
#
interface 10GE1/0/10
#
```

```
interface 10GE1/0/11
#
interface 10GE1/0/12
#
interface 10GE1/0/13
#
interface 10GE1/0/14
undo portswitch
 ip address 192.168.10.2 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
device transceiver 10GBASE-COPPER
#
interface 10GE1/0/16
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
interface 10GE1/0/20
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/21
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/22
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/23
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/24
#
interface 10GE1/0/25
#
interface 10GE1/0/26
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/27
#
interface 10GE1/0/28
#
interface 10GE1/0/29
#
interface 10GE1/0/30
#
interface 10GE1/0/31
#
interface 10GE1/0/32
```

```
#
interface 10GE1/0/33
#
interface 10GE1/0/34
#
interface 10GE1/0/35
#
interface 10GE1/0/36
#
interface 10GE1/0/37
#
interface 10GE1/0/38
#
interface 10GE1/0/39
#
interface 10GE1/0/40
#
interface 10GE1/0/41
#
interface 10GE1/0/42
#
interface 10GE1/0/43
#
interface 10GE1/0/44
#
interface 10GE1/0/45
#
interface 10GE1/0/46
#
interface 10GE1/0/47
#
interface 10GE1/0/48
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface 40GE1/0/3
#
interface 40GE1/0/4
#
interface 40GE1/0/5
#
interface 40GE1/0/6
#
interface NULL0
#
ip route-static 0.0.0.0 0.0.0.0 192.168.20.1
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
```

```
#
user-interface con 0
 authentication-mode password
 set authentication password cipher
$1c$4;5h;(QD"&$u.nRPe#@=TyyJ<#`4>|CC}$*'9@:[~X"ms=RfCpB$
#
vm-manager
#
return
```

2.5 VXLAN BGP EVPN Distributed Gateway (Communication in the Same Subnet)

2.5.1 Networking and Service Description

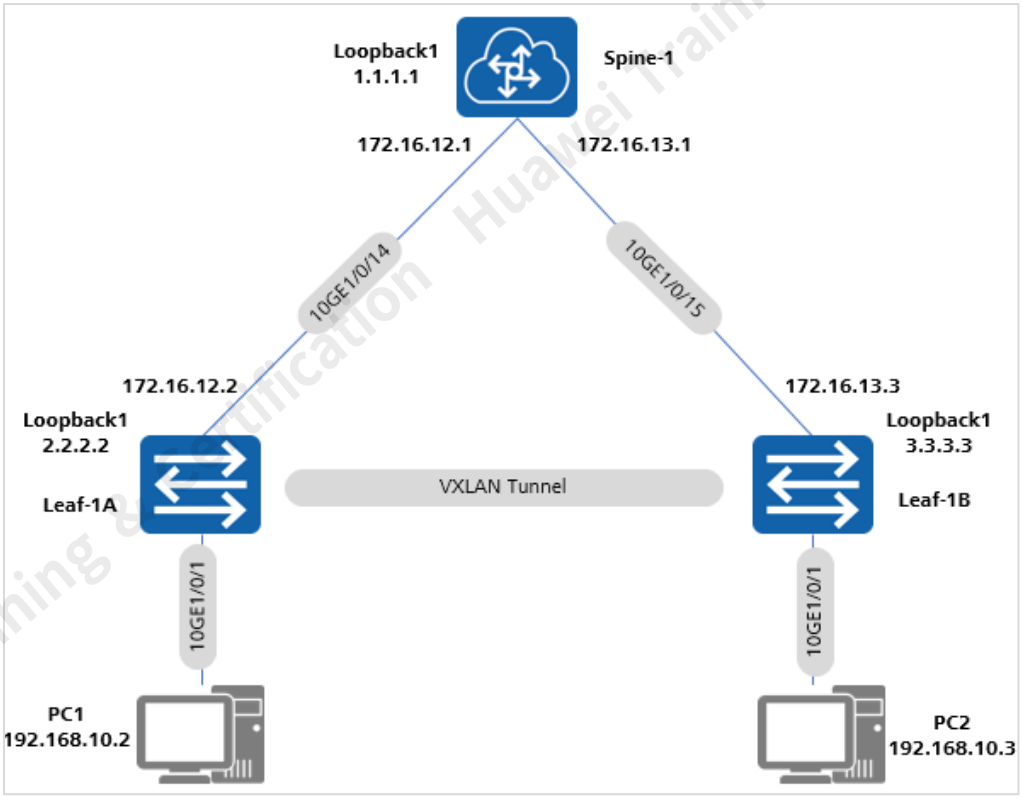

Figure 2-4 VXLAN BGP EVPN distributed gateway

Based on the user requirements, use Spine-1 as the spine node, and use Leaf-1A and Leaf-1B to deploy the VXLAN Layer 2 and Layer 3 gateways. Leaf-1A and Leaf-1B use BGP EVPN to dynamically create VXLAN tunnels, enabling PC1 and PC2 on the same network segment to communicate with each other.

Set the Bridge-domain and VNI parameters for the VXLAN to which the two VMs connect as follows:

Table 2-1 Bridge-domain and VLAN parameter planning for server access

VM Name	VLAN	Bridge-domain	VNI
PC1	10	10	10
PC2	10	10	10

Configure VPN instances and VNI mappings as follows:

Table 2-2 VPN and EVPN instance parameter planning

VM	VPN Instance	VNI	RD	vpn-target import	vpn-target export
PC1	vpn1	5000	11:11	1:1	1:1
PC2	vpn1	5000	11:11	1:1	1:1

On the Leaf node, set the parameter mappings between Bridge-domain and RT, and between Bridge-domain and RD for VXLAN EVPN as follows:

Table 2-3 EVPN instance parameter planning

VM	Bridge-domain	VNI	RD	vpn-target import	vpn-target export
PC1	10	10	10:1	10:1	11:1
PC2	10	10	10:1	10:1	11:1

Deploy the distributed gateway on the two Leaf nodes and configure the gateway as follows:

Table 2-4 VXLAN IP parameter planning

Device	VPN Instance	Vbdif	IP Address
Leaf-1A	vpn1	10	192.168.10.1/24
Leaf-1B	vpn1	10	192.168.10.1/24

Configure the connectivity between the tenant server and the distributed EVPN gateway.

2.5.2 Configuration Guideline

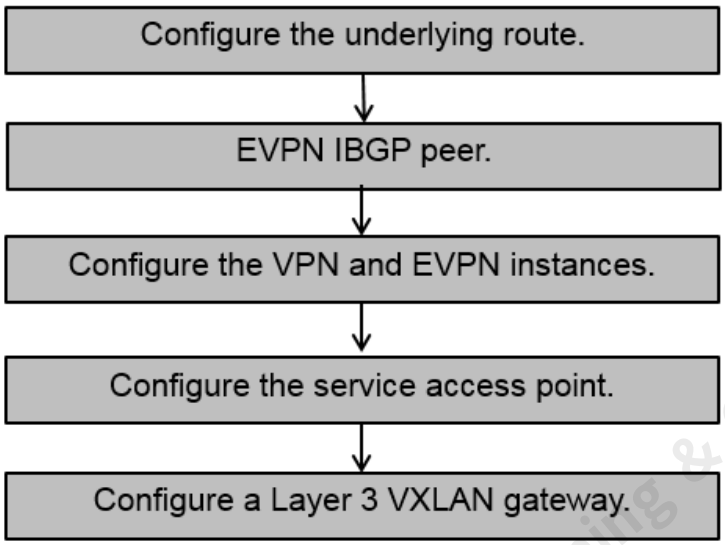

2.5.3 Configuration Procedure

Step 1 **Configuring the IP addresses of Spine-1, Leaf-1A and Leaf-1B**

```
<HUAWEI>system-view immediately
[HUAWEI] sysname Spine-1
[Spine-1]interface  LoopBack 1
[Spine-1-LoopBack1]ip address  1.1.1.1 32
[Spine-1-LoopBack1]quit
[Spine-1]interface 10GE 1/0/14
[Spine-1-10GE1/0/14]undo portswitch
[Spine-1-10GE1/0/14]ip address  172.16.12.1 24
[Spine-1-10GE1/0/14]quit
[Spine-1]interface 10GE 1/0/15
[Spine-1-10GE1/0/15]undo portswitch
[Spine-1-10GE1/0/15]ip address  172.16.13.1 24
[Spine-1-10GE1/0/15]quit

<HUAWEI>system-view immediately
[HUAWEI]sysname Leaf-1A
[Leaf-1A]interface LoopBack 1
[Leaf-1A-LoopBack1]ip address 2.2.2.2 32
[Leaf-1A-LoopBack1]quit
[Leaf-1A]interface 10
[Leaf-1A]interface 10GE 1/0/14
[Leaf-1A-10GE1/0/14]undo portswitch
[Leaf-1A-10GE1/0/14]ip address  172.12.12.2 24
[Leaf-1A-10GE1/0/14]quit

<HUAWEI>system-view immediately
[HUAWEI]sysname Leaf-1B
[Leaf-1B] interface LoopBack  1
```

```
[Leaf-1B-LoopBack1]ip address  3.3.3.3 32
[Leaf-1B-LoopBack1]quit
[Leaf-1B]interface 10GE 1/0/15
[Leaf-1B-10GE1/0/15]undo portswitch
[Leaf-1B-10GE1/0/15]ip address 172.16.13.2 24
[Leaf-1B-10GE1/0/15]quit
```

Step 2 **Configuring OSPF**

```
Configure the loopback interface LoopBack1 on Leaf-1A and Leaf-1B, and
advertise its IP address (considered as the VTEP IP) using the underlay
routing protocol.
<Spine-1>system-view  immediately
[Spine-1]ospf 1
[Spine-1-ospf-1]area 0
[Spine-1-ospf-1-area-0.0.0.0]network  1.1.1.1 0.0.0.0
[Spine-1-ospf-1-area-0.0.0.0]network  172.16.12.0 0.0.0.255
[Spine-1-ospf-1-area-0.0.0.0]network  172.16.13.0 0.0.0.255
[Spine-1-ospf-1-area-0.0.0.0]return

<Leaf-1A>system-view  immediately
[Leaf-1A]ospf 1
[Leaf-1A-ospf-1]area 0
[Leaf-1A-ospf-1-area-0.0.0.0]net
[Leaf-1A-ospf-1-area-0.0.0.0]network  2.2.2.2 0.0.0.0
[Leaf-1A-ospf-1-area-0.0.0.0]network  172.16.12.0 0.0.0.255
[Leaf-1A-ospf-1-area-0.0.0.0]return

<Leaf-1B>system-view  immediately
[Leaf-1B]ospf 1
[Leaf-1B-ospf-1]area 0
[Leaf-1B-ospf-1-area-0.0.0.0]network 3.3.3.3 0.0.0.0
[Leaf-1B-ospf-1-area-0.0.0.0]network  172.16.13.0 0.0.0.255
[Leaf-1B-ospf-1-area-0.0.0.0]return
```

Step 3 **Configuring BGP EVPN**

Enable EVPN as the VXLAN control plane.

```
[Spine-1] evpn-overlay enable
[Leaf-1A] evpn-overlay enable
[Leaf-1B] evpn-overlay enable
```

Use Spine-1 to configure the RR and establish the BGP EVPN peer relationship between Leaf-1A and Leaf-1B. As the RR, Spine-1 needs to transparently transmit community attributes, which requires you to run the **undo policy vpn-target** command.

```
[Spine-1] bgp 100 instance evpn1
[Spine-1-bgp-instance-evpn1] peer 2.2.2.2 as-number 100
[Spine-1-bgp-instance-evpn1] peer 2.2.2.2 connect-interface LoopBack1
[Spine-1-bgp-instance-evpn1] peer 3.3.3.3 as-number 100
[Spine-1-bgp-instance-evpn1] peer 3.3.3.3 connect-interface LoopBack1
[Spine-1-bgp-instance-evpn1] l2vpn-family evpn
```

```
[Spine-1-bgp-instance-evpn1-af-evpn] peer 2.2.2.2 enable
[Spine-1-bgp-instance-evpn1-af-evpn] peer 2.2.2.2 reflect-client
[Spine-1-bgp-instance-evpn1-af-evpn] peer 2.2.2.2 advertise irb
[Spine-1-bgp-instance-evpn1-af-evpn] peer 3.3.3.3 enable
[Spine-1-bgp-instance-evpn1-af-evpn] peer 3.3.3.3 reflect-client
[Spine-1-bgp-instance-evpn1-af-evpn] peer 3.3.3.3 advertise irb
[Spine-1-bgp-instance-evpn1-af-evpn] undo policy vpn-target
[Spine-1-bgp-instance-evpn1-af-evpn] quit
[Spine-1-bgp-instance-evpn1] quit
```

Configure and advertise IRB routes on Leaf-1A:

```
[Leaf-1A] bgp 100 instance evpn1
[Leaf-1A-bgp-instance-evpn1] peer 1.1.1.1 as-number 100
[Leaf-1A-bgp-instance-evpn1] peer 1.1.1.1 connect-interface LoopBack1
[Leaf-1A-bgp-instance-evpn1] l2vpn-family evpn
[Leaf-1A-bgp-instance-evpn1-af-evpn] peer 1.1.1.1 enable
[Leaf-1A-bgp-instance-evpn1-af-evpn] peer 1.1.1.1 advertise irb
[Leaf-1A-bgp-instance-evpn1-af-evpn] quit
[Leaf-1A-bgp-instance-evpn1] quit
```

Configure and advertise IRB routes on Leaf-1B:

```
[Leaf-1B] bgp 100 instance evpn1
[Leaf-1B-bgp-instance-evpn1] peer 1.1.1.1 as-number 100
[Leaf-1B-bgp-instance-evpn1] peer 1.1.1.1 connect-interface LoopBack1
[Leef-1B-bgp-instance-evpn1] l2vpn-family evpn
[Leaf-1B-bgp-instance-evpn1-af-evpn] peer 1.1.1.1 enable
[Leaf-1B-bgp-instance-evpn1-af-evpn] peer 1.1.1.1 advertise irb
[Leaf-1B-bgp-instance-evpn1-af-evpn] quit
[Leaf-1B-bgp-instance-evpn1] quit
```

Configure the VPN instance and EVPN instances. Deploy the distributed gateway on the Leaf node. Configure the VPN instance of the EVPN instance cross route and use the VPN instance to differentiate tenants and store tenant routes.

Leaf-1A

```
[Leaf-1A] ip vpn-instance vpn1
[Leaf-1A-vpn-instance-vpn1] vxlan vni 5000
[Leaf-1A-vpn-instance-vpn1] ipv4-family
[Leaf-1A-vpn-instance-vpn1-af-ipv4] route-distinguisher 11:11
[Leaf-1A-vpn-instance-vpn1-af-ipv4] vpn-target 1:1
[Leaf-1A-vpn-instance-vpn1-af-ipv4] vpn-target 11:1 evpn
[Leaf-1A-vpn-instance-vpn1-af-ipv4] quit
[Leaf-1A-vpn-instance-vpn1] quit
[Leaf-1A] bridge-domain 10
[Leaf-1A-bd10] vxlan vni 10
[Leaf-1A-bd10] evpn
[Leaf-1A-bd10-evpn] route-distinguisher 10:1
[Leaf-1A-bd10-evpn] vpn-target 11:1
[Leaf-1A-bd10-evpn] quit
[Leaf-1A-bd10] quit
```

Leaf-1B

```
[Leaf-1B] ip vpn-instance vpn1
[Leaf-1B-vpn-instance-vpn1] vxlan vni 5000
[Leaf-1B-vpn-instance-vpn1] ipv4-family
[Leaf-1B-vpn-instance-vpn1-af-ipv4] route-distinguisher 11:11
[Leaf-1B-vpn-instance-vpn1-af-ipv4] vpn-target 1:1
[Leaf-1B-vpn-instance-vpn1-af-ipv4] vpn-target 11:1 evpn
[Leaf-1B-vpn-instance-vpn1-af-ipv4] quit
[Leaf-1B-vpn-instance-vpn1] quit
[Leaf-1B] bridge-domain 10
[Leaf-1B-bd10] vxlan vni 10
[Leaf-1B-bd10] evpn
[Leaf-1B-bd10-evpn] route-distinguisher 10:1
[Leaf-1B-bd10-evpn] vpn-target 11:1
[Leaf-1B-bd10-evpn] quit
[Leaf-1A-bd10] quit
```

Configure the ingress replication function for Leaf-1A and Leaf-1B:

```
[Leaf-1A] interface nve 1
[Leaf-1A-Nve1] source 2.2.2.2
[Leaf-1A-Nve1] vni 10 head-end peer-list protocol bgp
[Leaf-1A-Nve1] quit
[Leaf-1A] commit

[Leaf-1B] interface nve 1
[Leaf-1B-Nve1] source 3.3.3.3
[Leaf-1B-Nve1] vni 10 head-end peer-list protocol bgp
[Leaf-1B-Nve1] quit
[Leaf-1B] commit
```

Step 4 Configuring the service access points for Leaf-1A and Leaf-1B

```
[Leaf-1A] interface 10ge 1/0/1.1 mode l2
[Leaf-1A-10GE1/0/1.1] encapsulation dot1q vid 10
[Leaf-1A-10GE1/0/1.1] bridge-domain 10
[Leaf-1A-10GE1/0/1.1] quit

[Leaf-1B] interface 10ge 1/0/1.1 mode l2
[Leaf-1B-10GE1/0/1.1] encapsulation dot1q vid 10
[Leaf-1B-10GE1/0/1.1] bridge-domain 10
[Leaf-1B-10GE1/0/1.1] quit
```

Step 5 Configuring the active-active gateway

Configure VXLAN Layer 3 active-active gateways, ensure that the gateways can send different routing information. Enable the distributed gateway function, and configure the switch to advertise host routes. Deploy distribute gateways on Leaf-1A and Leaf-1B, configure gateways, and bind VPN instances.

```
[Leaf-1A] interface Vbdif10
```

```
[Leaf-1A-Vbdif10] ip binding vpn-instance vpn1
[Leaf-1A-Vbdif10] ip address 192.168.10.1 255.255.255.0
[Leaf-1A-Vbdif10] mac-address 0000-5e00-0101
[Leaf-1A-Vbdif10] vxlan anycast-gateway enable
[Leaf-1A-Vbdif10] arp collect host enable
[Leaf-1A-Vbdif10] quit

[Leaf-1B] interface Vbdif10
[Leaf-1B-Vbdif10] ip binding vpn-instance vpn1
[Leaf-1B-Vbdif10] ip address 192.168.10.1 255.255.255.0
[Leaf-1B-Vbdif10] mac-address 0000-5e00-0101
[Leaf-1B-Vbdif10] vxlan anycast-gateway enable
[Leaf-1B-Vbdif10] arp collect host enable
[Leaf-1B-Vbdif10] quit
```

Enable the interface loopback function on the distributed gateway (not required for CE12800). If the gateway is deployed on the ToR CE6850, enable the interface loopback function on the device and add an unused interface to the Eth-Trunk. Otherwise, the gateway cannot be accessed.

```
[Leaf-1A] interface Eth-Trunk1
[Leaf-1A-Eth-Trunk1] trunkport 10ge 1/0/16
[Leaf-1A-Eth-Trunk1] service type tunnel
[Leaf-1B] interface Eth-Trunk1
[Leaf-1B-Eth-Trunk1] trunkport 10ge 1/0/16
[Leaf-1B-Eth-Trunk1] service type tunnel
----End
```

2.5.4 Verifying the Configuration Result

After the configuration is complete, check if the VXLAN tunnel is successfully established.

```
[Leaf-1A]display vxlan tunnel
Number of vxlan tunnel : 1
Tunnel ID   Source             Destination         State  Type    Uptime
--------------------------------------------------------------------------------
------
4026531841  2.2.2.2            3.3.3.3             up     dynamic  001h21m
```

After the configuration is complete, check the BGP EVPN peer relationship.

```
Leaf-1A]display bgp instance evpn1 evpn peer
 BGP local router ID       : 2.2.2.2
 Local AS number          : 100
 Total number of peers     : 1
 Peers in established state : 1

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

```
  3.3.3.3         4          100     1342     1274     0 18:20:12 Established
3
```

After the configuration is complete, check the host routes in the VPN instance.

```
<Leaf-1A>display  ip routing-table vpn-instance vpn1
Proto: Protocol       Pre: Preference
Route Flags: R - relay, D - download to fib, T - to vpn-instance, B - black
hole route
------------------------------------------------------------------------------
-
Routing Table : vpn1
        Destinations : 5       Routes : 5

Destination/Mask    Proto   Pre  Cost       Flags NextHop         Interface

   192.168.10.0/24  Direct  0    0           D   192.168.10.1    Vbdif10
   192.168.10.1/32  Direct  0    0           D   127.0.0.1       Vbdif10
 192.168.10.255/32  Direct  0    0           D   127.0.0.1       Vbdif10
   192.168.10.3/32  IBGP    255  0           RD  2.2.2.2         VXLAN
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

Run the Ping command to test the network connectivity. The following uses PC1 Ping PC2 as an example:

```
C:\Users\admin\Desktop>ping 192.168.10.3

Pinging 192.168.10.3 with 32 bytes of data:
Reply from 192.168.10.3: bytes=32 time=1ms TTL=254
Reply from 192.168.10.3: bytes=32 time=1ms TTL=254
Reply from 192.168.10.3: bytes=32 time=1ms TTL=254
Reply from 192.168.10.3: bytes=32 time=1ms TTL=254

Ping statistics for 192.168.10.3:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 1ms, Maximum = 1ms, Average = 1ms
```

2.5.5 Complete Configuration

Complete configuration of Spine-1:

```
#
sysname Spine-1
#
system resource standard
#
device board 1 board-type CE6851-48S6Q-HI
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
evpn-overlay enable
```

```
#
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
#
aaa
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface MEth0/0/0
#
interface 10GE1/0/1
#
interface 10GE1/0/2
 device transceiver 1000BASE-T
#
interface 10GE1/0/3
 device transceiver 1000BASE-T
#
interface 10GE1/0/4
 device transceiver 1000BASE-T
#
interface 10GE1/0/5
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/6
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/7
#
interface 10GE1/0/8
#
interface 10GE1/0/9
#
interface 10GE1/0/10
#
interface 10GE1/0/11
#
interface 10GE1/0/12
#
interface 10GE1/0/13
#
interface 10GE1/0/14
 undo portswitch
```

```
 ip address 172.16.12.1 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
 undo portswitch
 ip address 172.16.13.1 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/16
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
interface 10GE1/0/20
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/21
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/22
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/23
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/24
#
interface 10GE1/0/25
#
interface 10GE1/0/26
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/27
#
interface 10GE1/0/28
#
interface 10GE1/0/29
#
interface 10GE1/0/30
#
interface 10GE1/0/31
#
interface 10GE1/0/32
#
interface 10GE1/0/33
#
interface 10GE1/0/34
#
interface 10GE1/0/35
```

```
#
interface 10GE1/0/36
#
interface 10GE1/0/37
#
interface 10GE1/0/38
#
interface 10GE1/0/39
#
interface 10GE1/0/40
#
interface 10GE1/0/41
#
interface 10GE1/0/42
#
interface 10GE1/0/43
#
interface 10GE1/0/44
#
interface 10GE1/0/45
#
interface 10GE1/0/46
#
interface 10GE1/0/47
#
interface 10GE1/0/48
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface 40GE1/0/3
#
interface 40GE1/0/4
#
interface 40GE1/0/5
#
interface 40GE1/0/6
#
interface LoopBack1
 ip address 1.1.1.1 255.255.255.255
#
interface NULL0
#
bgp 100 instance evpn1
 peer 2.2.2.2 as-number 100
 peer 2.2.2.2 connect-interface LoopBack1
 peer 3.3.3.3 as-number 100
 peer 3.3.3.3 connect-interface LoopBack1
 #
 l2vpn-family evpn
  undo policy vpn-target
  peer 2.2.2.2 enable
  peer 2.2.2.2 advertise irb
  peer 2.2.2.2 reflect-client
```

```
  peer 3.3.3.3 enable
  peer 3.3.3.3 advertise irb
  peer 3.3.3.3 reflect-client
#
ospf 1
 area 0.0.0.0
  network 1.1.1.1 0.0.0.0
  network 172.16.12.0 0.0.0.255
  network 172.16.13.0 0.0.0.255
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher
$1c$4;5h;(QD"&$u.nRPe#@=TyyJ<#`4>|CC}$*'9@:[~X"ms=RfCpB$
#
vm-manager
#
Return
```

Complete configuration of Leaf-1A:

```
sysname Leaf-1A
#
system resource standard
#
device board 1 board-type CE6850U-24S2Q-HI
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
evpn-overlay enable
#
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
#
ip vpn-instance vpn1
 ipv4-family
  route-distinguisher 11:11
  vpn-target 1:1 export-extcommunity
  vpn-target 11:1 export-extcommunity evpn
  vpn-target 1:1 import-extcommunity
  vpn-target 11:1 import-extcommunity evpn
 vxlan vni 5000
```

```
#
bridge-domain 10
 vxlan vni 10
evpn
  route-distinguisher 10:1
  vpn-target 11:1 export-extcommunity
  vpn-target 11:1 import-extcommunity#
aaa
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface Vbdif10
 ip binding vpn-instance vpn1
 ip address 192.168.10.1 255.255.255.0
 mac-address 0000-5e00-0101
 vxlan anycast-gateway enable
 arp collect host enable
#
interface MEth0/0/0
#
interface Eth-Trunk1
 service type tunnel
#
interface 10GE1/0/1
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/1.1 mode l2
 encapsulation dot1q vid 10
 bridge-domain 10
#
interface 10GE1/0/2
#
interface 10GE1/0/3
#
interface 10GE1/0/4
#
interface 10GE1/0/5
#
interface 10GE1/0/6
#
interface 10GE1/0/7
#
interface 10GE1/0/8
#
interface 10GE1/0/9
```

```
#
interface 10GE1/0/10
#
interface 10GE1/0/11
#
interface 10GE1/0/12
#
interface 10GE1/0/13
#
interface 10GE1/0/14
 undo portswitch
 ip address 172.16.12.2 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/16
 eth-trunk 1
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
interface 10GE1/0/20
#
interface 10GE1/0/21
#
interface 10GE1/0/22
#
interface 10GE1/0/23
#
interface 10GE1/0/24
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface LoopBack1
 ip address 2.2.2.2 255.255.255.255
#
interface Nve1
 source 2.2.2.2
 vni 10 head-end peer-list protocol bgp
#
interface NULL0
#
bgp 100 instance evpn1
 peer 1.1.1.1 as-number 100
 peer 1.1.1.1 connect-interface LoopBack1
```

```
 #
 l2vpn-family evpn
  policy vpn-target
  peer 1.1.1.1 enable
  peer 1.1.1.1 advertise irb
#
ospf 1
 area 0.0.0.0
  network 2.2.2.2 0.0.0.0
  network 172.16.12.0 0.0.0.255
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher
$1c$Y@4\VN@!#P$#W=zV;c6j([Uu6FCkrLIlz>='%|rAHA4-EXzRrg($
#
vm-manager
#
Return
```

Complete configuration of Leaf-1B:

```
#
sysname Leaf-1B
#
system resource standard
#
device board 1 board-type CE6850U-24S2Q-HI
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
evpn-overlay enable
#
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
#
ip vpn-instance vpn1
 ipv4-family
  route-distinguisher 11:11
  vpn-target 1:1 export-extcommunity
  vpn-target 11:1 export-extcommunity evpn
  vpn-target 1:1 import-extcommunity
```

```
  vpn-target 11:1 import-extcommunity evpn
 vxlan vni 5000
#
bridge-domain 10
 vxlan vni 10
  evpn
  route-distinguisher 10:1
  vpn-target 11:1 export-extcommunity
  vpn-target 11:1 import-extcommunity#
aaa
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface Vbdif10
 ip binding vpn-instance vpn1
 ip address 192.168.10.1 255.255.255.0
 mac-address 0000-5e00-0101
 vxlan anycast-gateway enable
 arp collect host enable
#
interface MEth0/0/0
#
interface Eth-Trunk1
 service type tunnel
#
interface 10GE1/0/1
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/1.1 mode l2
 encapsulation dot1q vid 10
 bridge-domain 10
#
interface 10GE1/0/2
#
interface 10GE1/0/3
#
interface 10GE1/0/4
#
interface 10GE1/0/5
#
interface 10GE1/0/6
#
interface 10GE1/0/7
#
interface 10GE1/0/8
```

```
#
interface 10GE1/0/9
#
interface 10GE1/0/10
#
interface 10GE1/0/11
#
interface 10GE1/0/12
#
interface 10GE1/0/13
#
interface 10GE1/0/14
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
 undo portswitch
 ip address 172.16.13.2 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/16
 eth-trunk 1
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
interface 10GE1/0/20
#
interface 10GE1/0/21
#
interface 10GE1/0/22
#
interface 10GE1/0/23
#
interface 10GE1/0/24
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface LoopBack1
 ip address 3.3.3.3 255.255.255.255
#
interface Nve1
 source 3.3.3.3
 vni 10 head-end peer-list protocol bgp
#
interface NULL0
#
bgp 100 instance evpn1
```

```
 peer 1.1.1.1 as-number 100
 peer 1.1.1.1 connect-interface LoopBack1
 #
 l2vpn-family evpn
  policy vpn-target
  peer 1.1.1.1 enable
  peer 1.1.1.1 advertise irb
#
ospf 1
 area 0.0.0.0
  network 3.3.3.3 0.0.0.0
  network 172.16.13.0 0.0.0.255
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher
$1c$^y4|C=)iR6$tB)qKME3!!R8ZNLHgoMPYMW!=5c3j@(q);D-S"q0$
#
vm-manager
#
return
```

2.6 VXLAN BGP EVPN Distributed Gateway (Communication Between Different Subnets)

2.6.1 Networking and Service Description

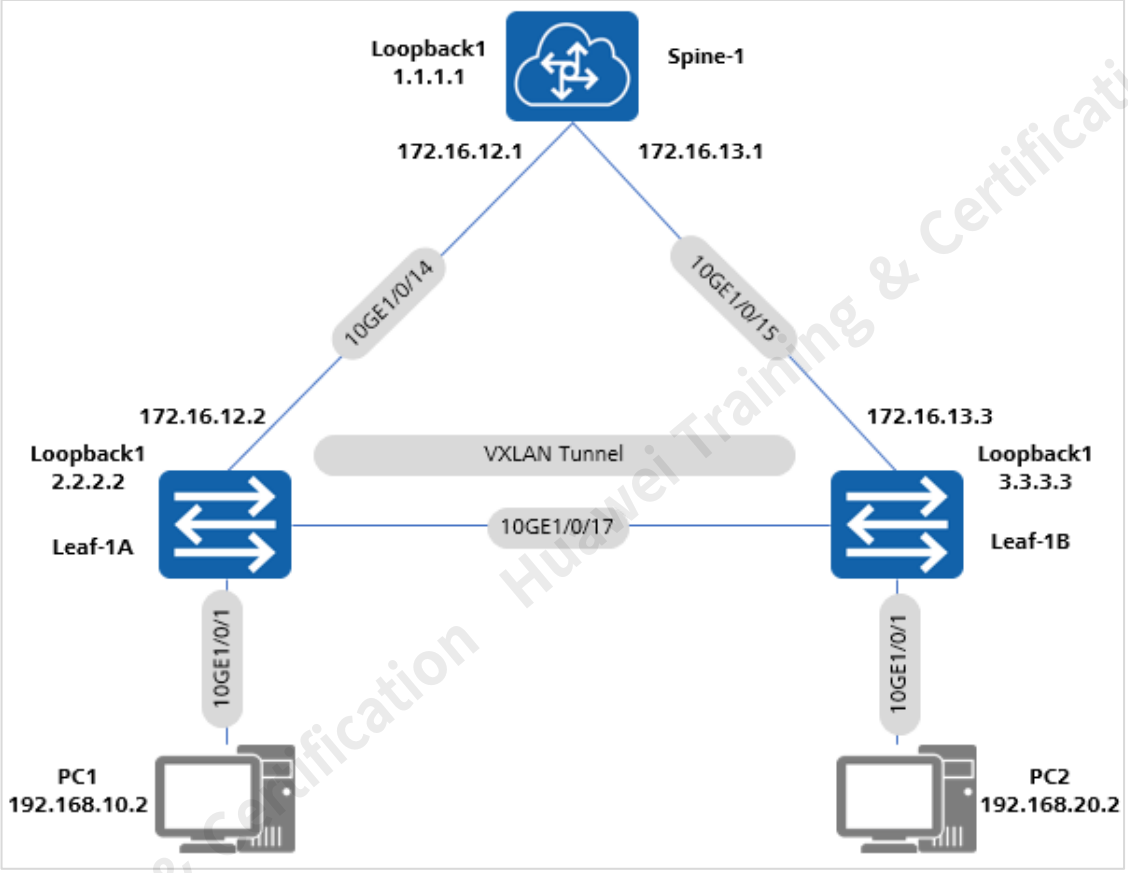

Figure 2-5 Network diagram of the VXLAN Layer 2 interconnection experiment

Based on the user requirements, use Spine-1 as the spine node, and use Leaf-1A and Leaf-1B to deploy the VXLAN Layer 2 and Layer 3 gateways. Leaf-1A and Leaf-1B use BGP EVPN to dynamically create VXLAN tunnels, enabling PC1 and PC2 on different network segments to communicate with each other.

Set the Bridge-domain and VNI parameters for the VXLAN to which the two VMs connect as follows:

Table 2-5 Bridge-domain and VLAN parameter planning for server access

VM	VLAN	Bridge-domain	VNI
PC1	10	10	10

VM	VLAN	Bridge-domain	VNI
PC2	20	20	20

Configure VPN instances and VNI mappings as follows:

Table 2-6 VPN and EVPN instance parameter planning

VM	VPN Instance	VNI	RD	vpn-target import	vpn-target export
PC1	vpn1	5000	11:11	1:1	1:1
PC2	vpn2	5000	22:22	2:2	2:2

On the Leaf node, set the parameter mappings between Bridge-domain and RT, and between Bridge-domain and RD for VXLAN EVPN as follows:

Table 2-7 EVPN instance parameter planning

VM	Bridge-domain	VNI	RD	vpn-target import	vpn-target export
PC1	10	10	10:1	11:1	11:1
PC2	20	20	20:1	11:1	11:1

Deploy the distributed gateway on the two Leaf nodes and configure the gateway as follows:

Table 2-8 VXLAN IP parameter planning

Device	VPN Instance	Vbdif	IP Address
Leaf-1A	vpn1	10	192.168.10.1/24
Leaf-1B	vpn1	20	192.168.20.1/24

Configure the connectivity between the tenant server and the distributed EVPN gateway.

2.6.2 Configuration Guideline

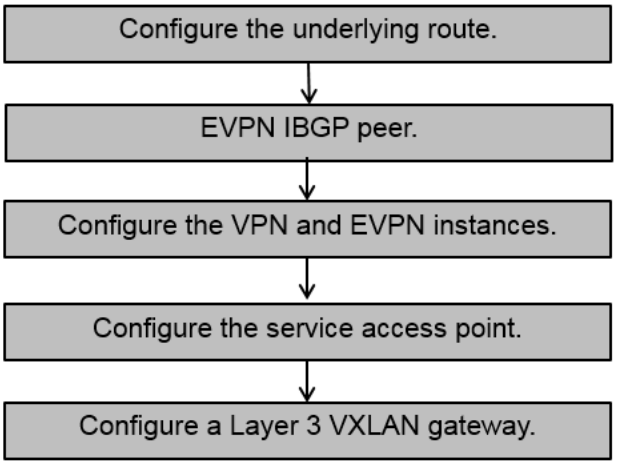

2.6.3 Configuration Procedure

Step 1 **Configuring the IP addresses of Spine-1, Leaf-1A and Leaf-1B**

```
<HUAWEI>system-view immediately
[HUAWEI] sysname Spine-1
[Spine-1]interface  LoopBack  1
[Spine-1-LoopBack1]ip address  1.1.1.1 32
[Spine-1-LoopBack1]quit
[Spine-1]interface 10GE 1/0/14
[Spine-1-10GE1/0/14]undo portswitch
[Spine-1-10GE1/0/14]ip address  172.16.12.1 24
[Spine-1-10GE1/0/14]quit
[Spine-1]interface 10GE 1/0/15
[Spine-1-10GE1/0/15]undo portswitch
[Spine-1-10GE1/0/15]ip address  172.16.13.1 24
[Spine-1-10GE1/0/15]quit

<HUAWEI>system-view immediately
[HUAWEI]sysname Leaf-1A
[Leaf-1A]interface LoopBack 1
[Leaf-1A-LoopBack1]ip address 2.2.2.2 32
[Leaf-1A-LoopBack1]quit
[Leaf-1A]interface 10
[Leaf-1A]interface 10GE 1/0/14
[Leaf-1A-10GE1/0/14]undo portswitch
[Leaf-1A-10GE1/0/14]ip address  172.12.12.2 24
[Leaf-1A-10GE1/0/14]quit

<HUAWEI>system-view immediately
[HUAWEI]sysname Leaf-1B
[Leaf-1B] interface LoopBack  1
[Leaf-1B-LoopBack1]ip address  3.3.3.3 32
[Leaf-1B-LoopBack1]quit
```

```
[Leaf-1B]interface 10GE 1/0/15
[Leaf-1B-10GE1/0/15]undo portswitch
[Leaf-1B-10GE1/0/15]ip address 172.16.13.2 24
[Leaf-1B-10GE1/0/15]quit
```

Step 2 **Configuring OSPF**

```
Configure the loopback interface LoopBack1 on Leaf-1A and Leaf-1B, and
advertise its IP address (considered as the VTEP IP) using the underlay
routing protocol.
<Spine-1>system-view  immediately
[Spine-1]ospf 1
[Spine-1-ospf-1]area 0
[Spine-1-ospf-1-area-0.0.0.0]network  1.1.1.1 0.0.0.0
[Spine-1-ospf-1-area-0.0.0.0]network  172.16.12.0 0.0.0.255
[Spine-1-ospf-1-area-0.0.0.0]network  172.16.13.0 0.0.0.255
[Spine-1-ospf-1-area-0.0.0.0]return

<Leaf-1A>system-view  immediately
[Leaf-1A]ospf 1
[Leaf-1A-ospf-1]area 0
[Leaf-1A-ospf-1-area-0.0.0.0]net
[Leaf-1A-ospf-1-area-0.0.0.0]network  2.2.2.2 0.0.0.0
[Leaf-1A-ospf-1-area-0.0.0.0]network  172.16.12.0 0.0.0.255
[Leaf-1A-ospf-1-area-0.0.0.0]return

<Leaf-1B>system-view  immediately
[Leaf-1B]ospf 1
[Leaf-1B-ospf-1]area 0
[Leaf-1B-ospf-1-area-0.0.0.0]network 3.3.3.3 0.0.0.0
[Leaf-1B-ospf-1-area-0.0.0.0]network  172.16.13.0 0.0.0.255
[Leaf-1B-ospf-1-area-0.0.0.0]return
```

Step 3 **Configuring BGP EVPN**

Enable EVPN as the VXLAN control plane.

```
[Spine-1] evpn-overlay enable
[Leaf-1A] evpn-overlay enable
[Leaf-1B] evpn-overlay enable
```

Use Spine-1 to configure the RR and establish the BGP EVPN peer relationship between Leaf-1A and Leaf-1B. As the RR, Spine-1 needs to transparently transmit community attributes, which requires you to run the **undo policy vpn-target** command.

```
[Spine-1] bgp 100 instance evpn1
[Spine-1-bgp-instance-evpn1] peer 2.2.2.2 as-number 100
[Spine-1-bgp-instance-evpn1] peer 2.2.2.2 connect-interface LoopBack1
[Spine-1-bgp-instance-evpn1] peer 3.3.3.3 as-number 100
[Spine-1-bgp-instance-evpn1] peer 3.3.3.3 connect-interface LoopBack1
[Spine-1-bgp-instance-evpn1] l2vpn-family evpn
[Spine-1-bgp-instance-evpn1-af-evpn] peer 2.2.2.2 enable
[Spine-1-bgp-instance-evpn1-af-evpn] peer 2.2.2.2 reflect-client
```

```
[Spine-1-bgp-instance-evpn1-af-evpn] peer 2.2.2.2 advertise irb
[Spine-1-bgp-instance-evpn1-af-evpn] peer 3.3.3.3 enable
[Spine-1-bgp-instance-evpn1-af-evpn] peer 3.3.3.3 reflect-client
[Spine-1-bgp-instance-evpn1-af-evpn] peer 3.3.3.3 advertise irb
[Spine-1-bgp-instance-evpn1-af-evpn] undo policy vpn-target
[Spine-1-bgp-instance-evpn1-af-evpn] quit
[Spine-1-bgp-instance-evpn1] quit
```

Configure and advertise IRB routes on Leaf-1A:

```
[Leaf-1A] bgp 100 instance evpn1
[Leaf-1A-bgp-instance-evpn1] peer 1.1.1.1 as-number 100
[Leaf-1A-bgp-instance-evpn1] peer 1.1.1.1 connect-interface LoopBack1
[Leaf-1A-bgp-instance-evpn1] l2vpn-family evpn
[Leaf-1A-bgp-instance-evpn1-af-evpn] peer 1.1.1.1 enable
[Leaf-1A-bgp-instance-evpn1-af-evpn] peer 1.1.1.1 advertise irb
[Leaf-1A-bgp-instance-evpn1-af-evpn] quit
[Leaf-1A-bgp-instance-evpn1] quit
```

Configure and advertise IRB routes on Leaf-1B:

```
[Leaf-1B] bgp 100 instance evpn1
[Leaf-1B-bgp-instance-evpn1] peer 1.1.1.1 as-number 100
[Leaf-1B-bgp-instance-evpn1] peer 1.1.1.1 connect-interface LoopBack1
[Leef-1B-bgp-instance-evpn1] l2vpn-family evpn
[Leaf-1B-bgp-instance-evpn1-af-evpn] peer 1.1.1.1 enable
[Leaf-1B-bgp-instance-evpn1-af-evpn] peer 1.1.1.1 advertise irb
[Leaf-1B-bgp-instance-evpn1-af-evpn] quit
[Leaf-1B-bgp-instance-evpn1] quit
```

Configure the VPN instance and EVPN instances. Deploy the distributed gateway on the Leaf node. Configure the VPN instance of the EVPN instance cross route and use the VPN instance to differentiate tenants and store tenant routes.

Leaf-1A

```
[Leaf-1A] ip vpn-instance vpn1
[Leaf-1A-vpn-instance-vpn1] vxlan vni 5000
[Leaf-1A-vpn-instance-vpn1] ipv4-family
[Leaf-1A-vpn-instance-vpn1-af-ipv4] route-distinguisher 11:11
[Leaf-1A-vpn-instance-vpn1-af-ipv4] vpn-target 1:1
[Leaf-1A-vpn-instance-vpn1-af-ipv4] vpn-target 11:1 evpn
[Leaf-1A-vpn-instance-vpn1-af-ipv4] quit
[Leaf-1A-vpn-instance-vpn1] quit
[Leaf-1A] bridge-domain 10
[Leaf-1A-bd10] vxlan vni 10
[Leaf-1A-bd10] evpn
[Leaf-1A-bd10-evpn] route-distinguisher 10:1
[Leaf-1A-bd10-evpn] vpn-target 11:1
[Leaf-1A-bd10-evpn] quit
[Leaf-1A-bd10] quit
```

Leaf-1B

```
[Leaf-1B] ip vpn-instance vpn1
[Leaf-1B-vpn-instance-vpn1] vxlan vni 5000
[Leaf-1B-vpn-instance-vpn1] ipv4-family
[Leaf-1B-vpn-instance-vpn1-af-ipv4] route-distinguisher 22:22
[Leaf-1B-vpn-instance-vpn1-af-ipv4] vpn-target 2:2
[Leaf-1B-vpn-instance-vpn1-af-ipv4] vpn-target 11:1 evpn
[Leaf-1B-vpn-instance-vpn1-af-ipv4] quit
[Leaf-1B-vpn-instance-vpn1] quit
[Leaf-1B] bridge-domain 20
[Leaf-1B-bd20] vxlan vni 20
[Leaf-1B-bd20] evpn
[Leaf-1B-bd20-evpn] route-distinguisher 20:1
[Leaf-1B-bd20-evpn] vpn-target 11:1
[Leaf-1B-bd20-evpn] quit
[Leaf-1B-bd20] quit
```

Configure the ingress replication function for Leaf-1A and Leaf-1B:

```
[Leaf-1A] interface nve 1
[Leaf-1A-Nve1] source 2.2.2.2
[Leaf-1A-Nve1] vni 10 head-end peer-list protocol bgp
[Leaf-1A-Nve1] quit
[Leaf-1A] commit

[Leaf-1B] interface nve 1
[Leaf-1B-Nve1] source 3.3.3.3
[Leaf-1B-Nve1] vni 20 head-end peer-list protocol bgp
[Leaf-1B-Nve1] quit
[Leaf-1B] commit
```

Step 4 Configuring the service access points for Leaf-1A and Leaf-1B

```
[Leaf-1A] interface 10ge 1/0/1.1 mode l2
[Leaf-1A-10GE1/0/1.1] encapsulation dot1q vid 10
[Leaf-1A-10GE1/0/1.1] bridge-domain 10
[Leaf-1A-10GE1/0/1.1] quit

[Leaf-1B] interface 10ge 1/0/1.1 mode l2
[Leaf-1B-10GE1/0/1.1] encapsulation dot1q vid 20
[Leaf-1B-10GE1/0/1.1] bridge-domain 20
[Leaf-1B-10GE1/0/1.1] quit
```

Step 5 Configuring the Layer 3 VXLAN gateway

Configure VXLAN Layer 3 active-active gateways, ensure that the gateways can send different routing information. Enable the distributed gateway function, and configure the switch to advertise host routes. Deploy distribute gateways on Leaf-1A and Leaf-1B, configure gateways, and bind VPN instances.

```
[Leaf-1A] interface Vbdif10
[Leaf-1A-Vbdif10] ip binding vpn-instance vpn1
[Leaf-1A-Vbdif10] ip address 192.168.10.1 255.255.255.0
```

```
[Leaf-1A-Vbdif10] arp distribute-gateway enable
[Leaf-1A-Vbdif10] arp collect host enable
[Leaf-1A-Vbdif10] quit

[Leaf-1B] interface Vbdif20
[Leaf-1B-Vbdif20] ip binding vpn-instance vpn1
[Leaf-1B-Vbdif20] ip address 192.168.20.1 255.255.255.0
[Leaf-1B-Vbdif20] arp distribute-gateway enable
[Leaf-1B-Vbdif20] arp collect host enable
[Leaf-1B-Vbdif20] quit
```

Enable the interface loopback function on the distributed gateway (not required for CE12800). If the gateway is deployed on the ToR CE6850, enable the interface loopback function on the device and add an unused interface to the Eth-Trunk. Otherwise, the gateway cannot be accessed.

```
[Leaf-1A] interface Eth-Trunk1
[Leaf1-Eth-Trunk1] trunkport 10ge 1/0/16
[Leaf1-Eth-Trunk1] service type tunnel

[Leaf-1B] interface Eth-Trunk1
[Leaf-1B-Eth-Trunk1] trunkport 10ge 1/0/16
[Leaf-1B-Eth-Trunk1] service type tunnel
----End
```

2.6.4 Verifying the Configuration Result

After the configuration is complete, check if the VXLAN tunnel is successfully established.

```
<Leaf-1A>display vxlan tunnel
Number of vxlan tunnel : 1
Tunnel ID   Source              Destination         State  Type     Uptime
--------------------------------------------------------------------------------
------
4026531845  2.2.2.2             3.3.3.3             up     dynamic  00:33:42
```

After the configuration is complete, check the BGP EVPN peer relationship.

```
<Leaf-1A>display bgp instance evpn1 evpn peer
 BGP local router ID       : 2.2.2.2
 Local AS number          : 100
 Total number of peers     : 1
 Peers in established state : 1

  Peer          V         AS  MsgRcvd  MsgSent  OutQ  Up/Down      State
PrefRcv
  1.1.1.1        4        100     1107     1053     0 15:08:43 Established
4
<Leaf-1A>
```

After the configuration is complete, check the host routes in the VPN instance.

```
<Leaf-1A>display  ip routing-table vpn-instance vpn1
Proto: Protocol       Pre: Preference
Route Flags: R - relay, D - download to fib, T - to vpn-instance, B - black
hole route
------------------------------------------------------------------------------
-
Routing Table : vpn1
        Destinations : 5       Routes : 5

Destination/Mask    Proto   Pre  Cost       Flags NextHop         Interface

   192.168.10.0/24  Direct  0    0           D   192.168.10.1    Vbdif10
   192.168.10.1/32  Direct  0    0           D   127.0.0.1       Vbdif10
 192.168.10.255/32  Direct  0    0           D   127.0.0.1       Vbdif10
   192.168.20.2/32  IBGP    255  0           RD  3.3.3.3         VXLAN
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

Run the Ping command to test the network connectivity. The following uses PC1 Ping PC2 as an example:

```
C:\Users\admin\Desktop>ping 192.168.20.2

Pinging 192.168.20.2 with 32 bytes of data:
Reply from 192.168.20.2: bytes=32 time=1ms TTL=126
Reply from 192.168.20.2: bytes=32 time<1ms TTL=126
Reply from 192.168.20.2: bytes=32 time<1ms TTL=126
Reply from 192.168.20.2: bytes=32 time<1ms TTL=126

Ping statistics for 192.168.20.2:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

2.6.5 Complete Configuration

Complete configuration of Spine-1:

```
#
sysname Spine-1
#
system resource standard
#
device board 1 board-type CE6851-48S6Q-HI
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
evpn-overlay enable
#
```

```
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
#
aaa
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface MEth0/0/0
#
interface 10GE1/0/1
#
interface 10GE1/0/2
 device transceiver 1000BASE-T
#
interface 10GE1/0/3
 device transceiver 1000BASE-T
#
interface 10GE1/0/4
 device transceiver 1000BASE-T
#
interface 10GE1/0/5
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/6
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/7
#
interface 10GE1/0/8
#
interface 10GE1/0/9
#
interface 10GE1/0/10
#
interface 10GE1/0/11
#
interface 10GE1/0/12
#
interface 10GE1/0/13
#
interface 10GE1/0/14
 undo portswitch
 ip address 172.16.12.1 255.255.255.0
```

```
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
 undo portswitch
 ip address 172.16.13.1 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/16
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
interface 10GE1/0/20
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/21
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/22
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/23
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/24
#
interface 10GE1/0/25
#
interface 10GE1/0/26
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/27
#
interface 10GE1/0/28
#
interface 10GE1/0/29
#
interface 10GE1/0/30
#
interface 10GE1/0/31
#
interface 10GE1/0/32
#
interface 10GE1/0/33
#
interface 10GE1/0/34
#
interface 10GE1/0/35
#
```

```
interface 10GE1/0/36
#
interface 10GE1/0/37
#
interface 10GE1/0/38
#
interface 10GE1/0/39
#
interface 10GE1/0/40
#
interface 10GE1/0/41
#
interface 10GE1/0/42
#
interface 10GE1/0/43
#
interface 10GE1/0/44
#
interface 10GE1/0/45
#
interface 10GE1/0/46
#
interface 10GE1/0/47
#
interface 10GE1/0/48
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface 40GE1/0/3
#
interface 40GE1/0/4
#
interface 40GE1/0/5
#
interface 40GE1/0/6
#
interface LoopBack1
 ip address 1.1.1.1 255.255.255.255
#
interface NULL0
#
bgp 100 instance evpn1
 peer 2.2.2.2 as-number 100
 peer 2.2.2.2 connect-interface LoopBack1
 peer 3.3.3.3 as-number 100
 peer 3.3.3.3 connect-interface LoopBack1
 #
 l2vpn-family evpn
  undo policy vpn-target
  peer 2.2.2.2 enable
  peer 2.2.2.2 advertise irb
  peer 2.2.2.2 reflect-client
  peer 3.3.3.3 enable
```

```
  peer 3.3.3.3 advertise irb
  peer 3.3.3.3 reflect-client
#
ospf 1
 area 0.0.0.0
  network 1.1.1.1 0.0.0.0
  network 172.16.12.0 0.0.0.255
  network 172.16.13.0 0.0.0.255
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher
$1c$4;5h;(QD"&$u.nRPe#@=TyyJ<#`4>|CC}$*'9@:[~X"ms=RfCpB$
#
vm-manager
#
Return
```

Complete configuration of Leaf-1A:

```
<Leaf-1A>display current-configuration
!Software Version V200R002C50
!Last configuration was updated at 2018-06-01 10:02:04+00:00
#
sysname Leaf-1A
#
system resource standard
#
device board 1 board-type CE6850U-24S2Q-HI
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
evpn-overlay enable
#
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
#
ip vpn-instance vpn1
 ipv4-family
  route-distinguisher 11:11
  vpn-target 1:1 export-extcommunity
  vpn-target 11:1 export-extcommunity evpn
```

```
  vpn-target 1:1 import-extcommunity
  vpn-target 11:1 import-extcommunity evpn
 vxlan vni 5000
#
bridge-domain 10
 vxlan vni 10
 evpn
  route-distinguisher 10:1
  vpn-target 11:1 export-extcommunity
  vpn-target 11:1 import-extcommunity
#
aaa
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface Vbdif10
 ip binding vpn-instance vpn1
 ip address 192.168.10.1 255.255.255.0
 vxlan anycast-gateway enable
 arp collect host enable
#
interface MEth0/0/0
#
interface Eth-Trunk1
 service type tunnel
#
interface 10GE1/0/1
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/1.1 mode l2
 encapsulation dot1q vid 10
 bridge-domain 10
#
interface 10GE1/0/2
#
interface 10GE1/0/3
#
interface 10GE1/0/4
#
interface 10GE1/0/5
#
interface 10GE1/0/6
#
interface 10GE1/0/7
#
```

```
interface 10GE1/0/8
#
interface 10GE1/0/9
#
interface 10GE1/0/10
#
interface 10GE1/0/11
#
interface 10GE1/0/12
#
interface 10GE1/0/13
#
interface 10GE1/0/14
 undo portswitch
 ip address 172.16.12.2 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/16
 eth-trunk 1
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
interface 10GE1/0/20
#
interface 10GE1/0/21
#
interface 10GE1/0/22
#
interface 10GE1/0/23
#
interface 10GE1/0/24
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface LoopBack1
 ip address 2.2.2.2 255.255.255.255
#
interface Nve1
 source 2.2.2.2
 vni 10 head-end peer-list protocol bgp
#
interface NULL0
#
```

```
bgp 100 instance evpn1
 peer 1.1.1.1 as-number 100
 peer 1.1.1.1 connect-interface LoopBack1
 #
 l2vpn-family evpn
  policy vpn-target
  peer 1.1.1.1 enable
  peer 1.1.1.1 advertise irb
#
ospf 1
 area 0.0.0.0
  network 2.2.2.2 0.0.0.0
  network 172.16.12.0 0.0.0.255
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher
$1c$Y@4\VN@!#P$#W=zV;c6j([Uu6FCkrLIlz>='%|rAHA4-EXzRrg($
#
vm-manager
#
Return
```

Complete configuration of Leaf-1B:

```
#
sysname Leaf-1B
#
system resource standard
#
device board 1 board-type CE6850U-24S2Q-HI
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
evpn-overlay enable
#
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
#
ip vpn-instance vpn1
 ipv4-family
  route-distinguisher 22:22
```

```
  vpn-target 2:2 export-extcommunity
  vpn-target 11:1 export-extcommunity evpn
  vpn-target 2:2 import-extcommunity
  vpn-target 11:1 import-extcommunity evpn
 vxlan vni 5000
#
bridge-domain 20
 vxlan vni 20
 evpn
  route-distinguisher 20:1
  vpn-target 20:1 export-extcommunity
  vpn-target 11:1 export-extcommunity
  vpn-target 20:1 import-extcommunity
#
aaa
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface Vbdif20
 ip binding vpn-instance vpn1
 ip address 192.168.20.1 255.255.255.0
 vxlan anycast-gateway enable
 arp collect host enable
#
interface MEth0/0/0
#
interface Eth-Trunk1
 service type tunnel
#
interface 10GE1/0/1
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/1.1 mode l2
 encapsulation dot1q vid 20
 bridge-domain 20
#
interface 10GE1/0/2
#
interface 10GE1/0/3
#
interface 10GE1/0/4
#
interface 10GE1/0/5
#
interface 10GE1/0/6
```

```
#
interface 10GE1/0/7
#
interface 10GE1/0/8
#
interface 10GE1/0/9
#
interface 10GE1/0/10
#
interface 10GE1/0/11
#
interface 10GE1/0/12
#
interface 10GE1/0/13
#
interface 10GE1/0/14
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
 undo portswitch
 ip address 172.16.13.2 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/16
 eth-trunk 1
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
interface 10GE1/0/20
#
interface 10GE1/0/21
#
interface 10GE1/0/22
#
interface 10GE1/0/23
#
interface 10GE1/0/24
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface LoopBack1
 ip address 3.3.3.3 255.255.255.255
#
interface Nve1
 source 3.3.3.3
 vni 20 head-end peer-list protocol bgp
```

```
#
interface NULL0
#
bgp 100 instance evpn1
 peer 1.1.1.1 as-number 100
 peer 1.1.1.1 connect-interface LoopBack1
 #
 l2vpn-family evpn
  policy vpn-target
  peer 1.1.1.1 enable
  peer 1.1.1.1 advertise irb
#
ospf 1
 area 0.0.0.0
  network 3.3.3.3 0.0.0.0
  network 172.16.13.0 0.0.0.255
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher
$1c$^y4|C=)iR6$tB)qKME3!!R8ZNLHgoMPYMW!=5c3j@(q);D-S"q0$
#
vm-manager
#
Return
```

2.7 VXLAN Layer 3 Interconnection (Intra-AS Cross-DC Three-Segment VXLAN)

2.7.1 Networking and Service Description

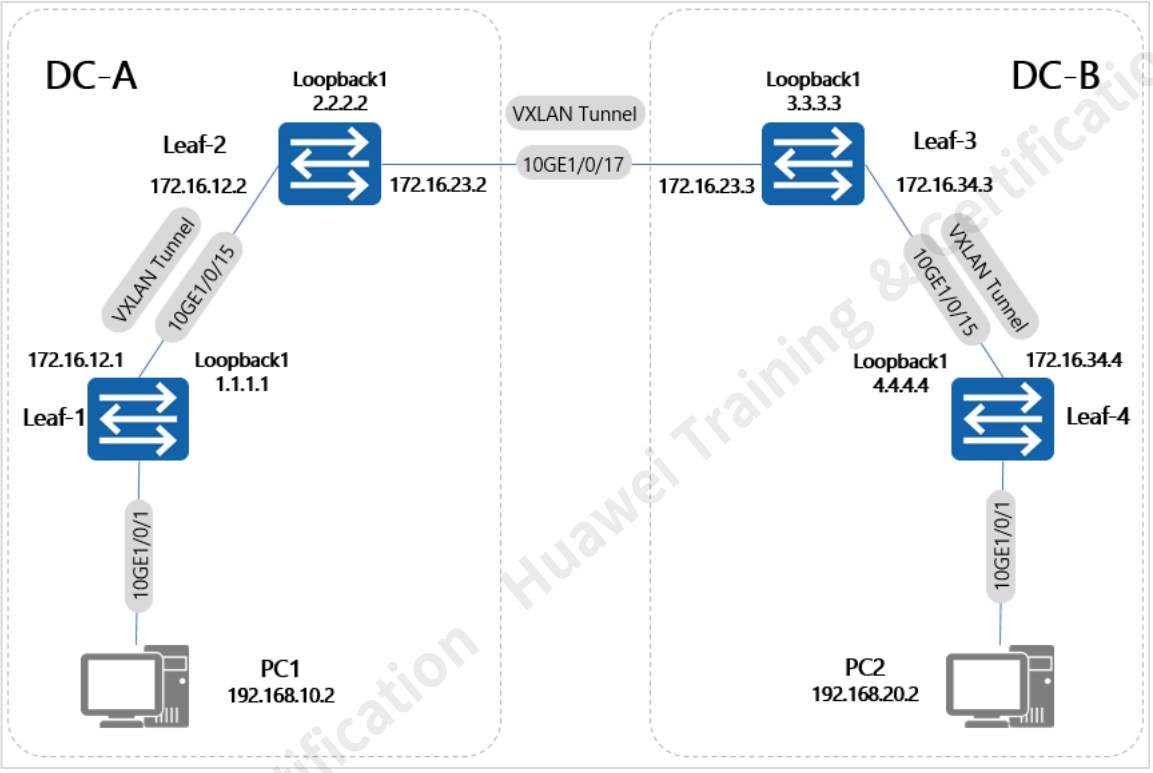

Figure 2-6 VXLAN BGP EVPN distributed gateway

Based on user requirements, data center A and data center B are planned in the same BGP AS. The BGP EVPN protocol is configured in each data center to create VXLAN tunnels for distributed gateways, and BGP EVPN protocol is configured between Leaf-2 and Leaf-3 to create VXLAN tunnels. Thus PC 1 and PC 2 on the same network segment between data center A and data center B can communicate with each other.

Set the Bridge-domain and VNI parameters for the VXLAN to which the two VMs in DCA and DCB connect as follows:

Table 2-9 Bridge-domain and VLAN parameter planning for server access

VM	VLAN	Bridge-domain	VNI
PC1	10	10	10
PC2	20	20	20

Configure VPN instances and VNI mappings as follows:

Table 2-10 VPN and EVPN instance parameter planning

Device	VPN Instance	VNI	RD	vpn-target import	vpn-target export
Leaf1	vpn1	5010	11:11	1:1	1:1
Leaf2	vpn1	5020	11:12	1:2	1:2
Leaf3	vpn1	5010	11:13	1:3	1:3
Leaf4	vpn1	5020	11:14	1:4	1:4

On the Leaf node, set the parameter mappings between Bridge-domain and RT, and between Bridge-domain and RD for VXLAN EVPN as follows:

Table 2-11 EVPN instance parameter planning

Device	Bridge-domain	VNI	RD	vpn-target import	vpn-target export
Leaf1	10	10	10:1	10:1	11:1
Leaf2	20	20	10:2	20:1	11:1 33:3
Leaf3	10	10	10:3	30:1	22:2 33:3
Leaf4	20	20	10:4	40:1	22:2

Deploy the distributed gateway on the two Leaf nodes and configure the gateway as follows:

Table 2-12 VXLAN IP parameter planning

Device	VPN Instance	Vbdif	IP Address
Leaf1	vpn1	10	192.168.10.1/24
Leaf4	vpn1	20	192.168.20.1/24

Configure the connectivity between the tenant server and the distributed EVPN gateway.

2.7.2 Configuration Guideline

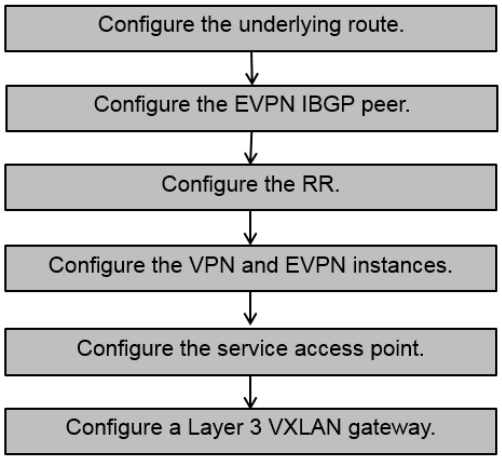

2.7.3 Configuration Procedure

Step 1 **Configuring IP addresses for all Leafs**

```
<HUAWEI>system-view immediately
[HUAWEI] sysname Leaf-1
[Leaf-1]interface  LoopBack  1
[Leaf-1-LoopBack1]ip address  1.1.1.1 32
[Leaf-1-LoopBack1]quit
[Leaf-1]interface 10GE 1/0/15
[Leaf-1-10GE1/0/15]undo portswitch
[Leaf-1-10GE1/0/15]ip address  172.16.12.1 24
[Leaf-1-10GE1/0/15]quit

<HUAWEI>system-view immediately
[HUAWEI] sysname Leaf-2
[Leaf-2]interface  LoopBack  1
[Leaf-2-LoopBack1]ip address  2.2.2.2 32
[Leaf-2-LoopBack1]quit
[Leaf-2]interface 10GE 1/0/15
[Leaf-2-10GE1/0/15]undo portswitch
[Leaf-2-10GE1/0/15]ip address  172.16.12.2 24
[Leaf-2-10GE1/0/15]quit
[Leaf-2]interface 10GE 1/0/17
[Leaf-2-10GE1/0/17]undo portswitch
[Leaf-2-10GE1/0/17]ip address  172.16.23.2 24
[Leaf-2-10GE1/0/17]quit

<HUAWEI>system-view immediately
[HUAWEI] sysname Leaf-3
[Leaf-3]interface  LoopBack  1
[Leaf-3-LoopBack1]ip address  3.3.3.3 32
[Leaf-3-LoopBack1]quit
[Leaf-3]interface 10GE 1/0/15
```

```
[Leaf-3-10GE1/0/15]undo portswitch
[Leaf-3-10GE1/0/15]ip address  172.16.34.3 24
[Leaf-3-10GE1/0/15]quit
[Leaf-3]interface 10GE 1/0/17
[Leaf-3-10GE1/0/17]undo portswitch
[Leaf-3-10GE1/0/17]ip address  172.16.23.3 24
[Leaf-3-10GE1/0/17]quit

<HUAWEI>system-view immediately
[HUAWEI] sysname Leaf-4
[Leaf-4]interface  LoopBack  1
[Leaf-4-LoopBack1]ip address  4.4.4.4 32
[Leaf-4-LoopBack1]quit
[Leaf-4]interface 10GE 1/0/15
[Leaf-4-10GE1/0/15]undo portswitch
[Leaf-4-10GE1/0/15]ip address  172.16.34.4 24
[Leaf-4-10GE1/0/15]quit
```

Step 2 **Configuring OSPF**

```
Configure the loopback interface LoopBack1 on all leafs, and advertise its IP
address (considered as the VTEP IP) using the underlay routing protocol.
<Leaf-1>system-view  immediately
[Leaf-1]ospf 1
[Leaf-1-ospf-1]area 0
[Leaf-1-ospf-1-area-0.0.0.0]network  1.1.1.1 0.0.0.0
[Leaf-1-ospf-1-area-0.0.0.0]network  172.16.12.0 0.0.0.255
[Leaf-1-ospf-1-area-0.0.0.0]return

<Leaf-2>system-view  immediately
[Leaf-2]ospf 1
[Leaf-2-ospf-1]area 0
[Leaf-2-ospf-1-area-0.0.0.0]network  2.2.2.2 0.0.0.0
[Leaf-2-ospf-1-area-0.0.0.0]network  172.16.12.0 0.0.0.255
[Leaf-2-ospf-1-area-0.0.0.0]return

<Leaf-3>system-view  immediately
[Leaf-3]ospf 1
[Leaf-3-ospf-1]area 0
[Leaf-3-ospf-1-area-0.0.0.0]network  3.3.3.3 0.0.0.0
[Leaf-3-ospf-1-area-0.0.0.0]network  172.16.34.0 0.0.0.255
[Leaf-3-ospf-1-area-0.0.0.0]return

<Leaf-4>system-view  immediately
[Leaf-4]ospf 1
[Leaf-4-ospf-1]area 0
[Leaf-4-ospf-1-area-0.0.0.0]network  3.3.3.3 0.0.0.0
[Leaf-4-ospf-1-area-0.0.0.0]network  172.16.34.0 0.0.0.255
[Leaf-4-ospf-1-area-0.0.0.0]return
```

Step 3 **Configuring BGP EVPN**

Enable EVPN as the VXLAN control plane.

```
[Leaf-1] evpn-overlay enable
[Leaf-1] evpn-overlay enable
[Leaf-1] evpn-overlay enable
[Leaf-1] evpn-overlay enable
```

After receiving EVPN routes from IBGP EVPN peers, Leaf-2 or Leaf-3 will not send EVPN routes to other IBGP EVPN peers. Therefore, Leaf-2 and Leaf-3 need to be configured as RRs. As RRs, Leaf-2 and Leaf-3 need to transparently transmit the community attributes, which requires you to run the **undo policy vpn-target** command.

Configure and advertise IRB routes on Leaf-1:

```
[Leaf-1] bgp 100 instance evpn1
[Leaf-1-bgp-instance-evpn1] peer 2.2.2.2 as-number 100
[Leaf-1-bgp-instance-evpn1] peer 2.2.2.2 connect-interface LoopBack1
[Leaf-1-bgp-instance-evpn1] l2vpn-family evpn
[Leaf-1-bgp-instance-evpn1-af-evpn] peer 2.2.2.2 enable
[Leaf-1-bgp-instance-evpn1-af-evpn] peer 2.2.2.2 advertise irb
[Leaf-1-bgp-instance-evpn1-af-evpn] quit
[Leaf-1-bgp-instance-evpn1] quit
```

Configure Leaf-2 as an RR:

```
[Leaf-2] bgp 100 instance evpn1
[Leaf-2-bgp-instance-evpn1] peer 1.1.1.1 as-number 100
[Leaf-2-bgp-instance-evpn1] peer 1.1.1.1 connect-interface LoopBack1
[Leaf-2-bgp-instance-evpn1] peer 172.16.23.3 as-number 100
[Leaf-2-bgp-instance-evpn1] l2vpn-family evpn
[Leaf-2-bgp-instance-evpn1-af-evpn] peer 1.1.1.1 enable
[Leaf-2-bgp-instance-evpn1-af-evpn] peer 1.1.1.1 reflect-client
[Leaf-2-bgp-instance-evpn1-af-evpn] peer 1.1.1.1 advertise irb
[Leaf-2-bgp-instance-evpn1-af-evpn] peer 1.1.1.1 import reoriginate
[Leaf-2-bgp-instance-evpn1-af-evpn] peer 1.1.1.1 advertise route-reoriginated
evpn mac-ip
[Leaf-2-bgp-instance-evpn1-af-evpn] peer 1.1.1.1 advertise route-reoriginated
evpn ip
[Leaf-2-bgp-instance-evpn1-af-evpn] peer 172.16.23.3 enable
[Leaf-2-bgp-instance-evpn1-af-evpn] peer 172.16.23.3 reflect-client
[Leaf-2-bgp-instance-evpn1-af-evpn] peer 172.16.23.3 advertise irb
[Leaf-2-bgp-instance-evpn1-af-evpn] peer 172.16.23.3 import reoriginate
[Leaf-2-bgp-instance-evpn1-af-evpn] peer 172.16.23.3 advertise route-
reoriginated evpn mac-ip
[Leaf-2-bgp-instance-evpn1-af-evpn] peer 172.16.23.3 advertise route-
reoriginated evpn ip
[Leaf-2-bgp-instance-evpn1-af-evpn] undo policy vpn-target
[Leaf-2-bgp-instance-evpn1-af-evpn] quit
[Leaf-2-bgp-instance-evpn1] quit
```

Configure Leaf-3 as an RR:

```
[Leaf-3] bgp 100 instance evpn1
```

```
[Leaf-3-bgp-instance-evpn1] peer 4.4.4.4 as-number 100
[Leaf-3-bgp-instance-evpn1] peer 4.4.4.4 connect-interface LoopBack1
[Leaf-3-bgp-instance-evpn1] peer 172.16.23.2 as-number 100
[Leaf-3-bgp-instance-evpn1] l2vpn-family evpn
[Leaf-3-bgp-instance-evpn1-af-evpn] peer 4.4.4.4 enable
[Leaf-3-bgp-instance-evpn1-af-evpn] peer 4.4.4.4 reflect-client
[Leaf-3-bgp-instance-evpn1-af-evpn] peer 4.4.4.4 advertise irb
[Leaf-3-bgp-instance-evpn1-af-evpn] peer 4.4.4.4 import reoriginate
[Leaf-3-bgp-instance-evpn1-af-evpn] peer 4.4.4.4 advertise route-reoriginated
evpn mac-ip
[Leaf-3-bgp-instance-evpn1-af-evpn] peer 4.4.4.4 advertise route-reoriginated
evpn ip
[Leaf-3-bgp-instance-evpn1-af-evpn] peer 172.16.23.2 enable
[Leaf-3-bgp-instance-evpn1-af-evpn] peer 172.16.23.2 reflect-client
[Leaf-3-bgp-instance-evpn1-af-evpn] peer 172.16.23.2 advertise irb
[Leaf-3-bgp-instance-evpn1-af-evpn] peer 172.16.23.2 import reoriginate
[Leaf-3-bgp-instance-evpn1-af-evpn] peer 172.16.23.2 advertise route-
reoriginated evpn mac-ip
[Leaf-3-bgp-instance-evpn1-af-evpn] peer 172.16.23.2 advertise route-
reoriginated evpn ip
[Leaf-3-bgp-instance-evpn1-af-evpn] undo policy vpn-target
[Leaf-3-bgp-instance-evpn1-af-evpn] quit
[Leaf-3-bgp-instance-evpn1] quit
```

Configure and advertise IRB routes on Use Leaf-4:

```
[Leaf-4] bgp 100 instance evpn1
[Leaf-4-bgp-instance-evpn1] peer 3.3.3.3 as-number 100
[Leaf-4-bgp-instance-evpn1] peer 3.3.3.3 connect-interface LoopBack1
[Leaf-4-bgp-instance-evpn1] l2vpn-family evpn
[Leaf-4-bgp-instance-evpn1-af-evpn] peer 3.3.3.3 enable
[Leaf-4-bgp-instance-evpn1-af-evpn] peer 3.3.3.3 advertise irb
[Leaf-4-bgp-instance-evpn1-af-evpn] quit
[Leaf-4-bgp-instance-evpn1] quit
```

Configure the VPN instance and EVPN instances. Deploy the distributed gateway on the Leaf node. Configure the VPN instance of the EVPN instance cross route and use the VPN instance to differentiate tenants and store tenant routes.

Leaf-1

```
[Leaf-1] ip vpn-instance vpn1
[Leaf-1-vpn-instance-vpn1] vxlan vni 5010
[Leaf-1-vpn-instance-vpn1] ipv4-family
[Leaf-1-vpn-instance-vpn1-af-ipv4] route-distinguisher 11:11
[Leaf-1-vpn-instance-vpn1-af-ipv4] vpn-target 1:1
[Leaf-1-vpn-instance-vpn1-af-ipv4] vpn-target 11:1 evpn
[Leaf-1-vpn-instance-vpn1-af-ipv4] quit
[Leaf-1-vpn-instance-vpn1] quit
[Leaf-1] bridge-domain 10
[Leaf-1-bd10] vxlan vni 10
[Leaf-1-bd10] evpn
[Leaf-1-bd10-evpn] route-distinguisher 10:1
[Leaf-1-bd10-evpn] vpn-target 10:1
```

```
[Leaf-1-bd10-evpn] vpn-target 11:1 export-extcommunity
[Leaf-1-bd10-evpn] quit
[Leaf-1-bd10] quit
```

Leaf-2

```
[Leaf-2] ip vpn-instance vpn1
[Leaf-2-vpn-instance-vpn1] vxlan vni 5020
[Leaf-2-vpn-instance-vpn1] ipv4-family
[Leaf-2-vpn-instance-vpn1-af-ipv4] route-distinguisher 11:12
[Leaf-2-vpn-instance-vpn1-af-ipv4] vpn-target 1:2
[Leaf-2-vpn-instance-vpn1-af-ipv4] vpn-target 11:1 evpn
[Leaf-2-vpn-instance-vpn1-af-ipv4] vpn-target 33:3 evpn
[Leaf-2-vpn-instance-vpn1-af-ipv4] quit
[Leaf-2-vpn-instance-vpn1] quit
[Leaf-2] bridge-domain 20
[Leaf-2-bd20] vxlan vni 20
[Leaf-2-bd20] evpn
[Leaf-2-bd20-evpn] route-distinguisher 10:2
[Leaf-2-bd20-evpn] vpn-target 20:1
[Leaf-2-bd20-evpn] vpn-target 11:1 export-extcommunity
[Leaf-2-bd20-evpn] vpn-target 33:3 export-extcommunity
[Leaf-2-bd20-evpn] quit
[Leaf-2-bd20] quit
```

Leaf-3

```
[Leaf-3] ip vpn-instance vpn1
[Leaf-3-vpn-instance-vpn1] vxlan vni 5010
[Leaf-3-vpn-instance-vpn1] ipv4-family
[Leaf-3-vpn-instance-vpn1-af-ipv4] route-distinguisher 11:13
[Leaf-3-vpn-instance-vpn1-af-ipv4] vpn-target 1:3
[Leaf-3-vpn-instance-vpn1-af-ipv4] vpn-target 22:2 evpn
[Leaf-3-vpn-instance-vpn1-af-ipv4] vpn-target 33:3 evpn
[Leaf-3-vpn-instance-vpn1-af-ipv4] quit
[Leaf-3-vpn-instance-vpn1] quit
[Leaf-3] bridge-domain 10
[Leaf-3-bd10] vxlan vni 10
[Leaf-3-bd10] evpn
[Leaf-3-bd10-evpn] route-distinguisher 10:3
[Leaf-3-bd10-evpn] vpn-target 30:1
[Leaf-3-bd10-evpn] vpn-target 22:2 export-extcommunity
[Leaf-3-bd10-evpn] vpn-target 33:3 export-extcommunity
[Leaf-3-bd10-evpn] quit
[Leaf-3-bd10] quit
```

Leaf-4

```
[Leaf-4] ip vpn-instance vpn1
[Leaf-4-vpn-instance-vpn1] vxlan vni 5020
[Leaf-4-vpn-instance-vpn1] ipv4-family
[Leaf-4-vpn-instance-vpn1-af-ipv4] route-distinguisher 11:14
[Leaf-4-vpn-instance-vpn1-af-ipv4] vpn-target 1:4
[Leaf-4-vpn-instance-vpn1-af-ipv4] vpn-target 22:2 evpn
```

```
[Leaf-4-vpn-instance-vpn1-af-ipv4] quit
[Leaf-4-vpn-instance-vpn1] quit
[Leaf-4] bridge-domain 20
[Leaf-4-bd20] vxlan vni 20
[Leaf-4-bd20] evpn
[Leaf-4-bd20-evpn] route-distinguisher 10:4
[Leaf-4-bd20-evpn] vpn-target 40:1
[Leaf-4-bd20-evpn] vpn-target 22:2 export-extcommunity
[Leaf-4-bd20-evpn] quit
[Leaf-4-bd20] quit
```

Configure the ingress replication function for Leaf-1A and Leaf-1B.

```
[Leaf-1] interface nve 1
[Leaf-1-Nve1] source 1.1.1.1
[Leaf-1-Nve1] vni 10 head-end peer-list protocol bgp
[Leaf-1-Nve1] quit

[Leaf-2] interface nve 1
[Leaf-2-Nve1] source 2.2.2.2
[Leaf-2-Nve1] vni 20 head-end peer-list protocol bgp
[Leaf-2-Nve1] quit
[Leaf-3] interface nve 1
[Leaf-3-Nve1] source 3.3.3.3
[Leaf-3-Nve1] vni 10 head-end peer-list protocol bgp
[Leaf-3-Nve1] quit
[Leaf-4] interface nve 1
[Leaf-4-Nve1] source 4.4.4.4
[Leaf-4-Nve1] vni 20 head-end peer-list protocol bgp
[Leaf-4-Nve1] quit
```

Step 4 Configuring the service access points for Leaf-1 and Leaf-4

```
[Leaf-1] interface 10ge 1/0/1.1 mode l2
[Leaf-1-10GE1/0/1.1] encapsulation dot1q vid 10
[Leaf-1-10GE1/0/1.1] bridge-domain 10
[Leaf-1-10GE1/0/1.1] quit
[Leaf-4] interface 10ge 1/0/1.1 mode l2
[Leaf-4-10GE1/0/1.1] encapsulation dot1q vid 20
[Leaf-4-10GE1/0/1.1] bridge-domain 20
[Leaf-4-10GE1/0/1.1] quit
```

Step 5 Configuring the VXLAN Gateway

Configure the Layer 3 VXLAN gateways, ensure that the gateways can send different routing information. Enable the distributed gateway function, and configure the switch to advertise host routes. Deploy distribute gateways on Leaf-1 and Leaf-4, configure gateways, and bind VPN instances.

```
[Leaf-1] interface Vbdif10
[Leaf-1-Vbdif10] ip binding vpn-instance vpn1
[Leaf-1-Vbdif10] ip address 192.168.10.1 255.255.255.0
```

```
[Leaf-1-Vbdif10] vxlan anycast-gateway enable
[Leaf-1-Vbdif10] arp collect host enable
[Leaf-1-Vbdif10] quit
[Leaf-4] interface Vbdif20
[Leaf-4-Vbdif20] ip binding vpn-instance vpn1
[Leaf-4-Vbdif20] ip address 192.168.20.1 255.255.255.0
[Leaf-4-Vbdif20] vxlan anycast-gateway enable
[Leaf-4-Vbdif20] arp collect host enable
[Leaf-4-Vbdif20] quit
```

Enable the interface loopback function on the distributed gateway (not required for CE12800). If the gateway is deployed on the ToR CE6850, enable the interface loopback function on the device and add an unused interface to the Eth-Trunk. Otherwise, the gateway cannot be accessed.

```
[Leaf-1] interface Eth-Trunk1
[Leaf-1-Eth-Trunk1] trunkport 10ge 1/0/16
[Leaf-1-Eth-Trunk1] service type tunnel

[Leaf-4] interface Eth-Trunk1
[Leaf-4-Eth-Trunk1] trunkport 10ge 1/0/16
[Leaf-4-Eth-Trunk1] service type tunnel
----End
```

2.7.4 Verifying the Configuration Result

After the configuration is complete, check if the VXLAN tunnel is successfully established.

```
[Leaf-1]display vxlan tunnel
Number of vxlan tunnel : 1
Tunnel ID   Source              Destination         State  Type    Uptime
-----------------------------------------------------------------------------
--
4026531846  1.1.1.1             2.2.2.2             up     dynamic  00:05:49
```

After the configuration is complete, check the BGP EVPN peer relationship.

```
[Leaf-2]display bgp instance evpn1 evpn peer
 BGP local router ID       : 2.2.2.2
 Local AS number          : 100
 Total number of peers     : 2
 Peers in established state : 2

  Peer          V        AS  MsgRcvd  MsgSent  OutQ  Up/Down      State
PrefRcv
  1.1.1.1       4       100     1342     1274     0 18:20:12 Established
3
  3.3.3.3       4       100     1107     1053     0 15:08:43 Established
4
```

After the configuration is complete, check the host routes in the VPN instance.

```
[Leaf-1]display  ip routing-table vpn-instance vpn1
Proto: Protocol       Pre: Preference
Route Flags: R - relay, D - download to fib, T - to vpn-instance, B - black
hole route
------------------------------------------------------------------------------
-
Routing Table : vpn1
        Destinations : 5       Routes : 5

Destination/Mask    Proto   Pre  Cost       Flags NextHop         Interface

   192.168.10.0/24  Direct  0    0           D   192.168.10.1    Vbdif10
   192.168.10.1/32  Direct  0    0           D   127.0.0.1       Vbdif10
   192.168.20.2/32  IBGP    255  0           RD  2.2.2.2         VXLAN
 192.168.10.255/32  Direct  0    0           D   127.0.0.1       Vbdif10
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

Run the Ping command to test the network connectivity. The following uses PC1 Ping PC2 as an example:

```
C:\Users\admin\Desktop>ping 192.168.20.2

Pinging 192.168.20.2 with 32 bytes of data:
Reply from 192.168.20.2: bytes=32 time=1ms TTL=126
Reply from 192.168.20.2: bytes=32 time<1ms TTL=126
Reply from 192.168.20.2: bytes=32 time<1ms TTL=126
Reply from 192.168.20.2: bytes=32 time<1ms TTL=126

Ping statistics for 192.168.20.2:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

2.7.5 Complete Configuration

Complete configuration of Leaf-1:

```
#
sysname Leaf-1
#
system resource standard
#
device board 1 board-type CE6850U-24S2Q-HI
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
evpn-overlay enable
#
```

```
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
ip vpn-instance vpn1
 ipv4-family
  route-distinguisher 11:11
  vpn-target 1:1 export-extcommunity
  vpn-target 11:1 export-extcommunity evpn
  vpn-target 1:1 import-extcommunity
  vpn-target 11:1 import-extcommunity evpn
 vxlan vni 5010
#
bridge-domain 10
 vxlan vni 10
 evpn
  route-distinguisher 10:1
  vpn-target 10:1 export-extcommunity
  vpn-target 11:1 export-extcommunity
  vpn-target 10:1 import-extcommunity
#
aaa
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface Vbdif10
 ip binding vpn-instance vpn1
 ip address 192.168.10.1 255.255.255.0
 vxlan anycast-gateway enable
 arp collect host enable
#
interface MEth0/0/0
#
interface Eth-Trunk1
 service type tunnel
#
interface 10GE1/0/1
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/1.1 mode l2
 encapsulation dot1q vid 10
 bridge-domain 10
#
interface 10GE1/0/2
 shutdown
```

```
#
interface 10GE1/0/3
 shutdown
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/4
 shutdown
#
interface 10GE1/0/5
 shutdown
#
interface 10GE1/0/6
 shutdown
 eth-trunk 1
#
interface 10GE1/0/7
 shutdown
#
interface 10GE1/0/8
 shutdown
#
interface 10GE1/0/9
 shutdown
#
interface 10GE1/0/10
 shutdown
#
interface 10GE1/0/11
 shutdown
#
interface 10GE1/0/12
 shutdown
#
interface 10GE1/0/13
 shutdown
#
interface 10GE1/0/14
 shutdown
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
 undo portswitch
 ip address 172.16.12.1 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/16
 shutdown
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 shutdown
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 shutdown
```

```
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
 shutdown
#
interface 10GE1/0/20
 shutdown
#
interface 10GE1/0/21
 shutdown
#
interface 10GE1/0/22
 shutdown
#
interface 10GE1/0/23
 shutdown
#
interface 10GE1/0/24
 shutdown
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface LoopBack1
 ip address 1.1.1.1 255.255.255.255
#
interface Nve1
 source 1.1.1.1
 vni 10 head-end peer-list protocol bgp
#
interface NULL0
#
bgp 100
 #
 ipv4-family unicast
#
bgp 100 instance evpn1
 peer 2.2.2.2 as-number 100
 peer 2.2.2.2 connect-interface LoopBack1
 #
 l2vpn-family evpn
  policy vpn-target
  peer 2.2.2.2 enable
  peer 2.2.2.2 advertise irb
#
ospf 1
 area 0.0.0.0
  network 1.1.1.1 0.0.0.0
  network 172.16.12.0 0.0.0.255
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
```

```
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher
$1c$j\VBJ9#V>*$u",pMYU$@0j2[B<Ul;9(Hm\,,;qrh:&\W0EZEM`*$
#
vm-manager
#
Return
```

Complete configuration of Leaf-2:

```
#
sysname Leaf-2
#
system resource standard
#
device board 1 board-type CE6851-48S6Q-HI
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
evpn-overlay enable
#
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
#
ip vpn-instance vpn1
 ipv4-family
  route-distinguisher 11:12
  vpn-target 1:2 export-extcommunity
  vpn-target 11:1 export-extcommunity evpn
  vpn-target 33:3 export-extcommunity evpn
  vpn-target 1:2 import-extcommunity
  vpn-target 11:1 import-extcommunity evpn
  vpn-target 33:3 import-extcommunity evpn
 vxlan vni 5020
#
bridge-domain 20
 vxlan vni 20
 evpn
  route-distinguisher 10:2
  vpn-target 20:1 export-extcommunity
  vpn-target 11:1 export-extcommunity
  vpn-target 33:3 export-extcommunity
  vpn-target 20:1 import-extcommunity
#
aaa
```

```
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface MEth0/0/0
#
interface Eth-Trunk1
 service type tunnel
#
interface 10GE1/0/1
 device transceiver 1000BASE-T
#
interface 10GE1/0/2
 shutdown
#
interface 10GE1/0/3
 shutdown
 device transceiver 1000BASE-T
#
interface 10GE1/0/4
 shutdown
 device transceiver 1000BASE-T
#
interface 10GE1/0/5
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/6
 shutdown
 eth-trunk 1
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/7
 shutdown
#
interface 10GE1/0/8
 shutdown
#
interface 10GE1/0/9
 shutdown
#
interface 10GE1/0/10
 shutdown
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/11
 shutdown
```

```
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/12
 shutdown
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/13
 shutdown
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/14
 shutdown
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
 undo portswitch
 ip address 172.16.12.2 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/16
 shutdown
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 undo portswitch
 ip address 172.16.23.2 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 shutdown
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
 shutdown
#
interface 10GE1/0/20
 shutdown
#
interface 10GE1/0/21
 shutdown
#
interface 10GE1/0/22
 shutdown
#
interface 10GE1/0/23
 shutdown
#
interface 10GE1/0/24
 shutdown
#
interface 10GE1/0/25
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/26
#
```

```
interface 10GE1/0/27
#
interface 10GE1/0/28
#
interface 10GE1/0/29
#
interface 10GE1/0/30
#
interface 10GE1/0/31
#
interface 10GE1/0/32
#
interface 10GE1/0/33
#
interface 10GE1/0/34
#
interface 10GE1/0/35
#
interface 10GE1/0/36
#
interface 10GE1/0/37
#
interface 10GE1/0/38
#
interface 10GE1/0/39
#
interface 10GE1/0/40
#
interface 10GE1/0/41
#
interface 10GE1/0/42
#
interface 10GE1/0/43
#
interface 10GE1/0/44
#
interface 10GE1/0/45
#
interface 10GE1/0/46
#
interface 10GE1/0/47
#
interface 10GE1/0/48
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface 40GE1/0/3
#
interface 40GE1/0/4
#
interface 40GE1/0/5
#
interface 40GE1/0/6
```

```
#
interface LoopBack1
 ip address 2.2.2.2 255.255.255.255
#
interface Nve1
 source 2.2.2.2
 vni 20 head-end peer-list protocol bgp
#
interface NULL0
#
bgp 100
 #
 ipv4-family unicast
#
bgp 100 instance evpn1
 peer 1.1.1.1 as-number 100
 peer 1.1.1.1 connect-interface LoopBack1
 peer 172.16.23.3 as-number 100
 #
 l2vpn-family evpn
  undo policy vpn-target
  peer 1.1.1.1 enable
  peer 1.1.1.1 advertise irb
  peer 1.1.1.1 reflect-client
  peer 1.1.1.1 import reoriginate
  peer 1.1.1.1 advertise route-reoriginated evpn mac-ip
  peer 1.1.1.1 advertise route-reoriginated evpn ip
  peer 172.16.23.3 enable
  peer 172.16.23.3 advertise irb
  peer 172.16.23.3 import reoriginate
  peer 172.16.23.3 advertise route-reoriginated evpn mac-ip
  peer 172.16.23.3 advertise route-reoriginated evpn ip
#
ospf 1
 area 0.0.0.0
  network 2.2.2.2 0.0.0.0
  network 172.16.12.0 0.0.0.255
#
lldp enable
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher
$1c$S\Z~G>LDI=$]"r"YOYnFVsoJ#YtM<.BP)uYVX`"R//W1VLk{)8@$
#
vm-manager
#
Return
```

Complete configuration of Leaf-3:

```
#
sysname Leaf-3
#
system resource standard
#
device board 1 board-type CE6851-48S6Q-HI
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
evpn-overlay enable
#
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
ip vpn-instance vpn1
 ipv4-family
  route-distinguisher 11:13
  vpn-target 1:3 export-extcommunity
  vpn-target 22:2 export-extcommunity evpn
  vpn-target 33:3 export-extcommunity evpn
  vpn-target 1:3 import-extcommunity
  vpn-target 22:2 import-extcommunity evpn
  vpn-target 33:3 import-extcommunity evpn
 vxlan vni 5010
#
bridge-domain 10
 vxlan vni 10
 evpn
  route-distinguisher 30:1
  vpn-target 30:1 export-extcommunity
  vpn-target 22:2 export-extcommunity
  vpn-target 33:3 export-extcommunity
  vpn-target 30:1 import-extcommunity
#
aaa
 #
 authentication-scheme default
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
```

```
interface MEth0/0/0
#
interface 10GE1/0/1
#
interface 10GE1/0/2
 device transceiver 1000BASE-T
#
interface 10GE1/0/3
 device transceiver 1000BASE-T
#
interface 10GE1/0/4
 device transceiver 1000BASE-T
#
interface 10GE1/0/5
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/6
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/7
#
interface 10GE1/0/8
#
interface 10GE1/0/9
#
interface 10GE1/0/10
#
interface 10GE1/0/11
#
interface 10GE1/0/12
#
interface 10GE1/0/13
#
interface 10GE1/0/14
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
 undo portswitch
 ip address 172.16.34.3 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/16
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 undo portswitch
 ip address 172.16.23.3 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
interface 10GE1/0/20
```

```
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/21
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/22
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/23
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/24
#
interface 10GE1/0/25
#
interface 10GE1/0/26
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/27
#
interface 10GE1/0/28
#
interface 10GE1/0/29
#
interface 10GE1/0/30
#
interface 10GE1/0/31
#
interface 10GE1/0/32
#
interface 10GE1/0/33
#
interface 10GE1/0/34
#
interface 10GE1/0/35
#
interface 10GE1/0/36
#
interface 10GE1/0/37
#
interface 10GE1/0/38
#
interface 10GE1/0/39
#
interface 10GE1/0/40
#
interface 10GE1/0/41
#
interface 10GE1/0/42
#
interface 10GE1/0/43
#
interface 10GE1/0/44
#
interface 10GE1/0/45
```

```
#
interface 10GE1/0/46
#
interface 10GE1/0/47
#
interface 10GE1/0/48
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface 40GE1/0/3
#
interface 40GE1/0/4
#
interface 40GE1/0/5
#
interface 40GE1/0/6
#
interface LoopBack1
 ip address 3.3.3.3 255.255.255.255
#
interface Nve1
 source 3.3.3.3
 vni 10 head-end peer-list protocol bgp
#
interface NULL0
#
bgp 100 instance evpn1
 peer 4.4.4.4 as-number 100
 peer 4.4.4.4 connect-interface LoopBack1
 peer 172.16.23.2 as-number 100
 #
 l2vpn-family evpn
  policy vpn-target
  peer 4.4.4.4 enable
  peer 4.4.4.4 advertise irb
  peer 4.4.4.4 reflect-client
  peer 4.4.4.4 import reoriginate
  peer 4.4.4.4 advertise route-reoriginated evpn mac-ip
  peer 4.4.4.4 advertise route-reoriginated evpn ip
  peer 172.16.23.2 enable
  peer 172.16.23.2 advertise irb
  peer 172.16.23.2 reflect-client
  peer 172.16.23.2 import reoriginate
  peer 172.16.23.2 advertise route-reoriginated evpn mac-ip
  peer 172.16.23.2 advertise route-reoriginated evpn ip
#
ospf 1
 area 0.0.0.0
  network 3.3.3.3 0.0.0.0
  network 172.16.34.0 0.0.0.255
#
ssh authorization-type default aaa
#
```

```
ssh server cipher aes256_ctr aes128_ctr
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher $1c$O.#mLm])f7$h]!y3MQ-
r$RFvvOe{=*;,LYH3.>Q:Es6It@-+=]A$
#
vm-manager
#
Return
```

Complete configuration of Leaf-4:

```
sysname Leaf-4
#
system resource standard
#
device board 1 board-type CE6850U-24S2Q-HI
#
drop-profile default
#
dcb pfc
#
dcb ets-profile default
#
evpn-overlay enable
#
telnet server disable
telnet ipv6 server disable
#
diffserv domain default
#
ip vpn-instance vpn1
 ipv4-family
  route-distinguisher 11:14
  vpn-target 1:4 export-extcommunity
  vpn-target 22:2 export-extcommunity evpn
  vpn-target 1:4 import-extcommunity
  vpn-target 22:2 import-extcommunity evpn
 vxlan vni 5020
#
bridge-domain 20
 vxlan vni 20
 evpn
  route-distinguisher 10:4
  vpn-target 40:1 export-extcommunity
  vpn-target 22:2 export-extcommunity
  vpn-target 40:1 import-extcommunity
#
aaa
 #
 authentication-scheme default
```

```
 #
 authorization-scheme default
 #
 accounting-scheme default
 #
 domain default
 #
 domain default_admin
#
stack
#
interface Vbdif20
 ip binding vpn-instance vpn1
 ip address 192.168.20.1 255.255.255.0
 vxlan anycast-gateway enable
 arp collect host enable
#
interface MEth0/0/0
#
interface 10GE1/0/1
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/1.1 mode l2
 encapsulation dot1q vid 20
 bridge-domain 20
#
interface 10GE1/0/2
#
interface 10GE1/0/3
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/4
#
interface 10GE1/0/5
#
interface 10GE1/0/6
#
interface 10GE1/0/7
#
interface 10GE1/0/8
#
interface 10GE1/0/9
#
interface 10GE1/0/10
#
interface 10GE1/0/11
#
interface 10GE1/0/12
#
interface 10GE1/0/13
#
interface 10GE1/0/14
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/15
```

```
 undo portswitch
 ip address 172.16.34.4 255.255.255.0
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/16
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/17
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/18
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/19
#
interface 10GE1/0/20
#
interface 10GE1/0/21
#
interface 10GE1/0/22
#
interface 10GE1/0/23
#
interface 10GE1/0/24
#
interface 40GE1/0/1
#
interface 40GE1/0/2
#
interface LoopBack1
 ip address 4.4.4.4 255.255.255.255
#
interface Nve1
 source 4.4.4.4
 vni 20 head-end peer-list protocol bgp
#
interface NULL0
#
bgp 100 instance evpn1
 peer 3.3.3.3 as-number 100
 peer 3.3.3.3 connect-interface LoopBack1
 #
 l2vpn-family evpn
  policy vpn-target
  peer 3.3.3.3 enable
  peer 3.3.3.3 advertise irb
#
ospf 1
 area 0.0.0.0
  network 4.4.4.4 0.0.0.0
  network 172.16.34.0 0.0.0.255
#
ssh authorization-type default aaa
#
ssh server cipher aes256_ctr aes128_ctr
```

```
ssh server hmac sha2_256_96 sha2_256 sha1_96
ssh server key-exchange dh_group_exchange_sha256 dh_group_exchange_sha1
ecdh_sha2_nistp256 ecdh_sha2_nistp384 ecdh_sha2_nistp521 sm2_kep
#
user-interface con 0
 authentication-mode password
 set authentication password cipher $1c$=!pEPOGO:K$}@RJP$l'6W-
O!+!}FY$0Gs#%Vu(*h<Kc!7(v2.6S$
#
vm-manager
#
return
```

3 Underlay Network Configuration

3.1 Objectives

- Understand how to deploy and design a Layer 2 agile data center network.
- Understand how to configure advanced features such as chain stacking and M-LAG on the CE6800.
- Understand how to configure SDN Underlay network devices.
- Learn the basic configuration of the Hot Standby Backup (HSB) of the firewall.

3.2 Networking and Service Description

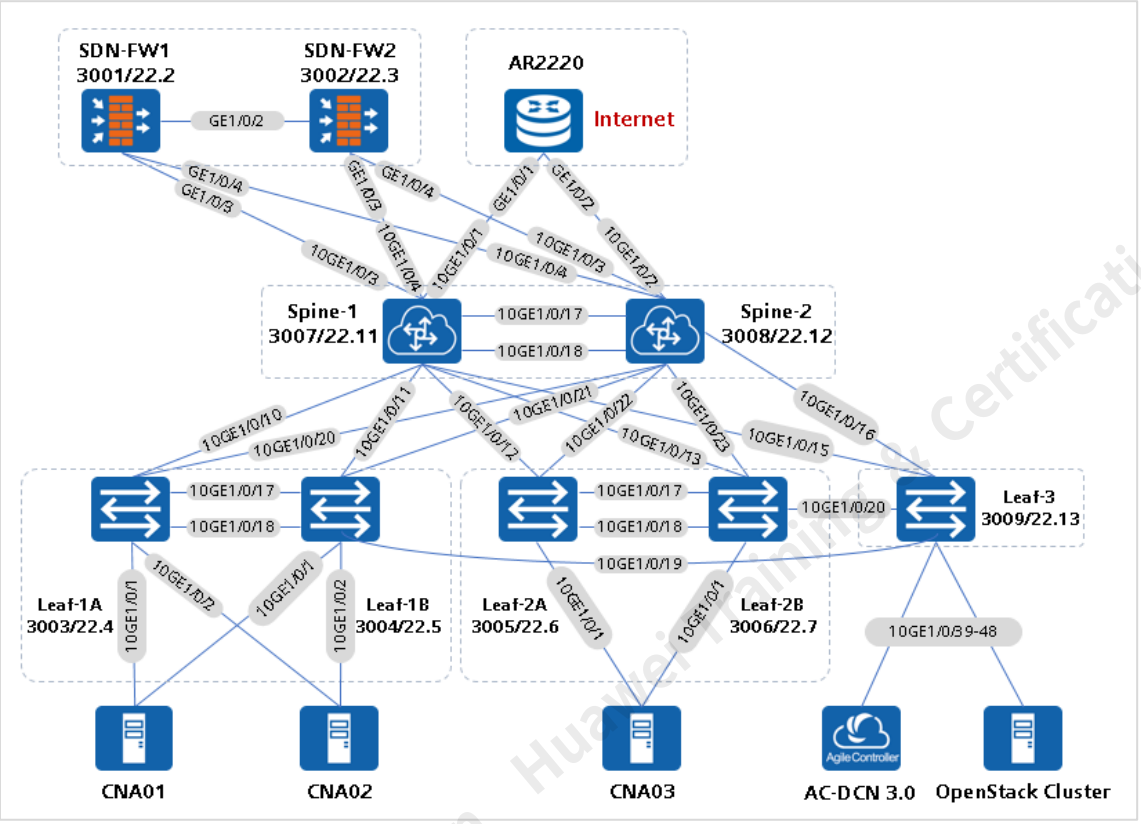

Figure 3-1 Underlay experimental networking

In the data center network, it is required that the core layer be highly reliable. The core switch functions as both the backbone and gateway nodes to deploy VXLAN all-active gateways. The Underlay network uses technologies such as stacking and M-LAG to implement high availability, simplify the network structure, and facilitate management and O&M.

3.3 Configure the iStack Stacking on the Access Switch

3.3.1 Networking and Configuration Guideline

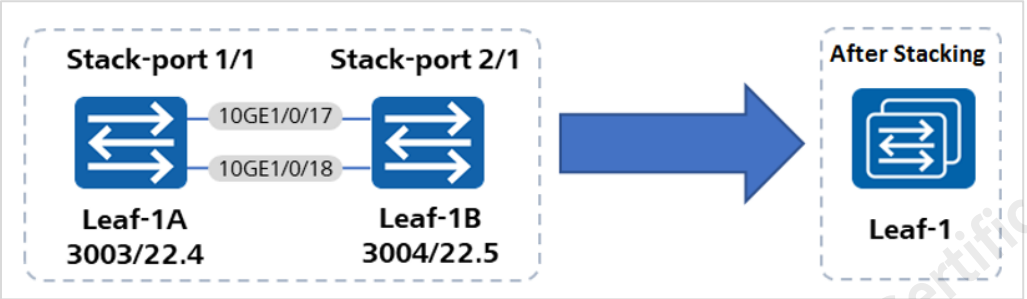

Figure 3-2 Network diagram of ToR iStack Experiment

Based on the user requirements, Leaf-1A and Leaf-1B adopt stacked networking, and 10GE1/0/17 and 10GE1/0/18 are added to stack ports. The two switches form a chain stack system.

The ToR switch model at the access layer is CE6851.

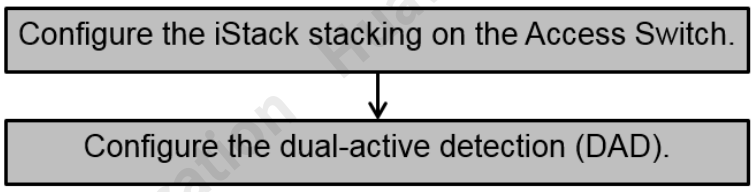

To stack devices and make the new stack different from other stacks on the network, configure the stack attributes for the switches, including the stack member ID, priority, and stack domain ID.

To forward data packets between the stack member devices, configure a stack port. Add multiple ports of the physical stack member ports to the stack port to increase the bandwidth and reliability of the stack link.

To make the configuration take effect and set up a stack, you can save the configuration and connect stack cables after restarting the device.

3.3.2 Configure the Stack Attributes for Leaf-1A and Leaf-1B

Set the stack member ID of Leaf-1A to **1**, priority to **150**, and domain ID to **10**.

```
<HUAWEI> system-view
[~HUAWEI] sysname Leaf-1A
[Leaf-1A] commit
[~Leaf-1A] stack
[*Leaf-1A-stack] stack member 1 priority 150
```

```
[*Leaf-1A-stack] stack member 1 domain 10
[*Leaf-1A-stack] quit
[*Leaf-1A] commit
```

Remarks: By default, the stack member ID is 1. By default, the stack member ID of Leaf-1A is **1**, which does not need to be configured.

Set the stack member ID of Leaf-1B to **2**, domain ID to **10**, and priority to **100** (which is the default value).

```
<HUAWEI> system-view
[~HUAWEI] sysname Leaf-1B
[Leaf-1B] commit
[~Leaf-1B] stack
[*Leaf-1B-stack] stack member 1 renumber 2 inherit-config
Warning: The stack configuration of member ID 1 will be inherited to member
ID 2 after the device resets. Continue? [Y/N]: y
[*Leaf-1B-stack] stack member 1 domain 10
[*Leaf-1B-stack] quit
[~Leaf-1B] commit
```

When running the **stack renumber new-member-id [inherit-config]** command to change the stack member ID:

If the parameter is set to **inherit-config**, the stack configuration (including the stack priority, domain ID, and physical stack member port) is inherited after restart of the switches. However, the configuration of common service ports will be lost, for example, the port splitting configuration (while the configurations of physical stack member ports will be inherited).

If the parameter is not set to **inherit-config**, the configurations of the physical stack member ports and common service ports will be lost after the restart, and the configurations of the new stack member ID in the configuration file will take effect.

If offline configurations related to the new stack member ID exist in the switches, the **inherit-config** parameter cannot be specified.

3.3.3 Configuring the Stack Ports

Add service ports 10GE1/0/17 to 10GE1/0/18 of Leaf-1A to the stack port 1/1.

```
[~Leaf-1A] interface stack-port 1/1
[*Leaf-1A-Stack-Port1/1] port member-group interface 10ge 1/0/17 to 1/0/18
Warning: The interface(s) (10GE1/0/17-1/0/18) will be converted to stack
mode. [Y/N]: y
[*Leaf-1A-Stack-Port1/1] quit
[*Leaf-1A] commit
[~Leaf-1A] quit
```

Add service ports 10GE1/0/17 to 10GE1/0/18 of Leaf-1B to the stack port 1/1.

```
[*Leaf-1B] interface stack-port 1/1
[*Leaf-1B-Stack-Port1/1] port member-group interface 10ge 1/0/17 to 1/0/18
Warning: The interface(s) (10GE1/0/17-1/0/18) will be converted to stack
mode. [Y/N]: y
[*Leaf-1B-Stack-Port1/1] quit
[*Leaf-1B] commit
[~Leaf-1B] quit
```

3.3.4 Checking the Stack Configuration

After the preceding configuration is complete, you are advised to run the **display stack configuration** command to check whether the configuration is the same as expected. The following uses Leaf-1A as an example.

```
<Leaf-1A> display stack configuration
Oper : Operation
Conf : Configuration
: Offline configuration
Isolated Port: The port is in stack mode, but does not belong to any Stack-
Port

Attribute Configuration:
----------------------------------------
 MemberID  Domain Priority
Oper(Conf)   Oper(Conf)  Oper(Conf)
----------------------------------------
1(1) --(10)  100(150)
----------------------------------------

Stack-Port Configuration:
------------------------------------------------------------------------
Stack-Port   Member Ports
------------------------------------------------------------------------
Stack-Port1/110GE1/0/17 10GE1/0/18
------------------------------------------------------------------------
```

3.3.5 Saving the Configurations and Restarting the Device

Save the configurations of Leaf-1A and Leaf-1B, and then restart the device.

```
<Leaf-1A> save
Warning: The current configuration will be written to the device. Continue?
[Y/N]: y
<Leaf-1A> reboot
Warning: The system will reboot. Continue? [Y/N]: y
```

3.3.6 Connecting Stack Cables to Set Up the Stack

3.3.7 Verifying the Configuration Result

Check information about the stack members.

```
<Leaf-1>dis stack
--------------------------------------------------------------------------------
--
MemberID Role    MAC            Priority   DeviceType        Description
--------------------------------------------------------------------------------
---
1       Master   c4ff-1f3d-e290   150       CE6851-48S6Q-HI
2       Standby  c4ff-1f3d-e320   120       CE6851-48S6Q-HI
--------------------------------------------------------------------------------
-- indicates the device through which the user logs in.
```

Check the stack topology information.

```
<Leaf-1>dis stack topology
Stack Topology:
-----------------------------------------------
          Stack-Port 1     Stack-Port 2
MemberID   Status Neighbor   Status Neighbor
-----------------------------------------------
1         up     2         --     --
2         up     1         --     --
-----------------------------------------------
Stack Link:
--------------------------------------------------------------------------------
Stack-Port      Port             Status     PeerPort          PeerStatus
--------------------------------------------------------------------------------
Stack-Port1/1   10GE1/0/17         up         10GE2/0/17         up
Stack-Port1/1   10GE1/0/18         up         10GE2/0/18         up
Stack-Port2/1   10GE2/0/17         up         10GE1/0/17         up
Stack-Port2/1   10GE2/0/18         up         10GE1/0/18         up
--------------------------------------------------------------------------------
```

Check the complete configuration of the stack.

```
<Leaf-1>dis stack configuration
Oper         : Operation
Conf         : Configuration
*            : Offline configuration
Isolated Port : The port is in stack mode, but does not belong to any Stack-
Port
System Forwarding Model:
-------------------------------
Oper             Conf
-------------------------------
hybrid           hybrid
-------------------------------
```

```
Attribute Configuration:
-------------------------------------------------------------------------------
 MemberID      Domain        Priority      Switch Mode     Uplink Port
Oper(Conf)   Oper(Conf)      Oper(Conf)      Oper(Conf)      Oper(Conf)
-------------------------------------------------------------------------------
1(1)         10(10)         150(150)       Auto(Auto)      6*40GE(6*40GE)
2(2)         10(10)         120(120)       Auto(Auto)      6*40GE(6*40GE)
-------------------------------------------------------------------------------
Stack-Port Configuration:
-------------------------------------------------------------------------------
--
Stack-Port     Member Ports
-------------------------------------------------------------------------------
-
Stack-Port1/1   10GE1/0/17         10GE1/0/18
Stack-Port2/1   10GE2/0/17         10GE2/0/18
-------------------------------------------------------------------------------
```

3.3.8 Configuring DAD

Configure DAD (Dual-active Detection) through the management network port.

When the management network ports of the Leaf-1A and Leaf-1B stack member switches are connected to the management network, DAD can be implemented using the management network ports to quickly detect and handle the dual-active conflict after the stack splits. This mode does not require additional interfaces or relay devices.

To implement DAD through management network ports, IP addresses must be configured for the stack management network ports.

When DAD is implemented through management network ports, a dual-active conflict is detected if different stacks have management network ports connected to the same management network and have the same stack domain ID and management IP address. As a result, interfaces on the low-priority device will become Error-Down.

```
<Leaf-1A>system-view
Enter system view, return user view with return command.
[~Leaf-1A]interface MEth 0/0/0
[~Leaf-1A-MEth0/0/0]dual-active detect enable
[~Leaf-1A-MEth0/0/0]quit
[~Leaf-1A]commit
[*Leaf-1A]quit
```

3.3.9 Saving the Stack Configuration

After the stack is successfully set up, you are advised to run the **save** command immediately to save the stack configuration.

```
<Leaf-1A> save
```

```
Warning: The current configuration will be written to the device. Continue?
[Y/N]: y
```

3.3.10 Complete Stack Configuration

Complete configuration of Leaf-1A:

```
#
sysname Leaf-1A
#
stack
 #
 stack member 1 domain 10
 stack member 1 priority 150
 #
 stack member 2 domain 10
 #
interface MEth0/0/0
 ip address 172.21.22.4 255.255.0.0
 dual-active detect enable
#
interface Stack-Port1/1
#
interface Stack-Port2/1
#
interface 10GE1/0/17
 port mode stack
 stack-port 1/1
 port crc-statistics trigger error-down
 device transceiver 40GBASE-COPPER
#
interface 10GE1/0/18
 port mode stack
 stack-port 1/1
 port crc-statistics trigger error-down
 device transceiver 40GBASE-COPPER
#
interface 10GE2/0/17
 port mode stack
 stack-port 2/1
 port crc-statistics trigger error-down
 device transceiver 40GBASE-COPPER
#
interface 10GE2/0/18
 port mode stack
 stack-port 2/1
 port crc-statistics trigger error-down
 device transceiver 40GBASE-COPPER
#
Return
#
```

3.4 Layer-3 Data Center Network Configuration

3.4.1 Objectives

- Understand how to configure the Layer 3 data center network.
- Understand how to configure BGP on CE series switches and AR routers and how to rectify common faults that may occur during BGP configuration.

3.4.2 Networking and Service Description

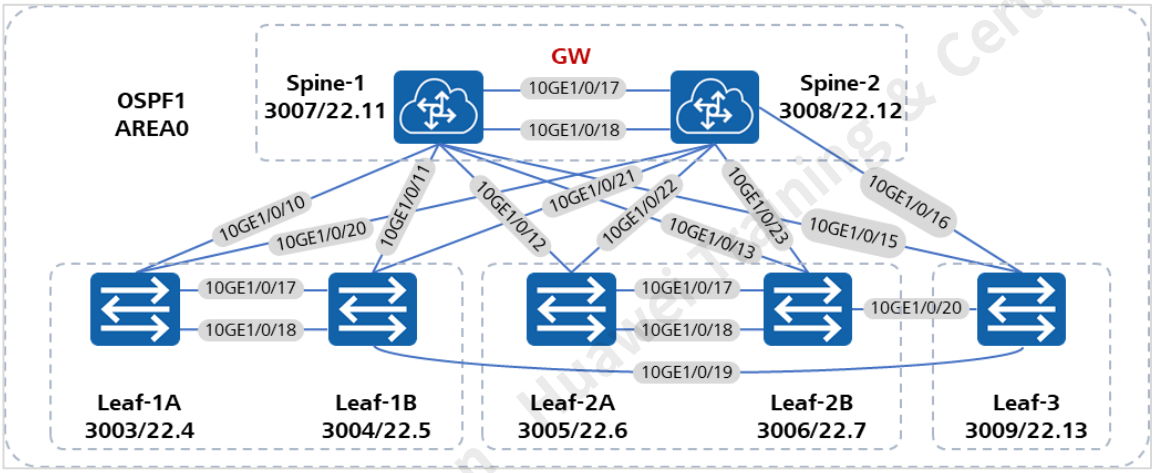

Figure 3-3 Layer 3 data center network topology

3.4.3 Configuration Guideline

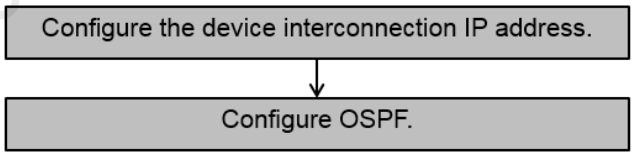

3.4.4 Experiment IP Address Planning

Table 3-1 Layer 3 IP address planning

Device Name	Interface Name	IP Address and Subnet Mask	Device Name	Interface Name	IP Address and Subnet Mask
Spine-1	10GE1/0/17	172.16.1.45/30	Spine-2	10GE1/0/17	172.16.1.46/30
	10GE1/0/10	172.16.1.13/30		10GE1/0/16	172.16.1.37/30
	10GE1/0/11	172.16.1.49/30		10GE1/0/23	172.16.1.33/30

Device Name	Interface Name	IP Address and Subnet Mask	Device Name	Interface Name	IP Address and Subnet Mask
	10GE1/0/12	172.16.1.17/30		10GE1/0/22	172.16.1.29/30
	10GE1/0/13	172.16.1.21/30		10GE1/0/21	172.16.1.25/30
	10GE1/0/15	172.16.1.57/30		10GE1/0/20	172.16.1.53/30
	Loopback0	10.45.45.45/32		Loopback0	10.45.45.45/32
	Loopback1	11.1.1.1/32		Loopback1	11.2.2.2./32
Leaf-2A	10GE1/0/12	172.16.1.18/30	Leaf-1	10GE1/0/10	172.16.1.14/30
	10GE1/0/22	172.16.1.30/30		10GE1/0/20	172.16.1.54/30
	10GE1/0/17	172.16.1.41/30		10GE2/0/11	172.16.1.50/30
	Loopback0	10.2.2.2/32		10GE2/0/21	172.16.1.26/30
	Loopback1	11.4.4.4/32		Loopback1	11.3.3.3/32
Leaf-2B	10GE1/0/13	172.16.1.22/30		Loopback0	10.1.1.1/32
	10GE1/0/23	172.16.1.34/30	Leaf-3	10GE1/0/15	172.16.1.58/30
	10GE1/0/17	172.16.1.42/30		10GE1/0/16	172.16.1.38/30
	Loopback0	10.2.2.2/32		Loopback0	10.3.3.3/32
	Loopback1	11.5.5.5/32		Loopback1	11.6.6.6/32

The following takes the configuration of the Layer 3 IP address for Spine-1 as an example. The configuration of IP addresses for other devices is similar and thus no more details are provided here.

```
<Spine-1>system-view immediately
Enter system view, return user view with return command.
[Spine-1]interface 10ge1/0/17
[Spine-1-10GE1/0/17]undo portswitch
[Spine-1-10GE1/0/17]ip address 172.16.1.45 30
[Spine-1-10GE1/0/17]quit
[Spine-1]interface 10ge1/0/10
[Spine-1-10GE1/0/10]undo portswitch
[Spine-1-10GE1/0/10]ip address 172.16.1.13 30
[Spine-1-10GE1/0/10]quit
[Spine-1]interface 10ge1/0/11
[Spine-1-10GE1/0/11]undo portswitch
[Spine-1-10GE1/0/11]ip address 172.16.1.49 30
[Spine-1-10GE1/0/11]quit
[Spine-1]interface 10ge1/0/12
[Spine-1-10GE1/0/12]undo portswitch
```

```
[Spine-1-10GE1/0/12]ip address 172.16.1.17 30
[Spine-1-10GE1/0/12]quit
[Spine-1]interface 10ge1/0/13
[Spine-1-10GE1/0/13]undo portswitch
[Spine-1-10GE1/0/13]ip address 172.16.1.21 30
[Spine-1-10GE1/0/13]quit
[Spine-1]interface 10ge1/0/15
[Spine-1-10GE1/0/25]undo portswitch
[Spine-1-10GE1/0/25]ip address 172.16.1.57 30
[Spine-1-10GE1/0/25]quit
[Spine-1]interface loopback0
[Spine-1-loopback0]ip address 10.45.45.45 32
[Spine-1-loopback0]quit
```

3.4.5 Configuring OSPF

Table 3-2 OSPF planning

Device Name	AREA	Router-id	OSPF Network Segment
Leaf-1	0	11.3.3.3	172.16.1.1/32, 172.16.1.12/30, 11.3.3.3/32, 10.1.1.1/32
			172.16.1.24/30, 172.16.1.48/30, 172.16.1.52/30
Spine-2	0	11.2.2.2	10.45.45.45/32, 172.16.1.52/32, 172.16.1.44/30, 172.16.1.36/30, 11.2.2.2/32
			10.1.1.24/30, 172.16.1.28/30, 172.16.1.32/30
Spine-1	0	11.1.1.1	10.45.45.45/32, 172.16.1.56/30, 172.16.1.44/30, 172.16.1.12/30, 11.1.1.1/32
			172.16.1.16/30, 172.16.1.20/30, 172.16.1.48/30
Leaf-2A	0	11.4.4.4	172.16.1.16/30, 172.16.1.28/30
			172.16.1.40/30, 10.2.2.2/32, 11.4.4.4/32
Leaf-2B	0	11.5.5.5	172.16.1.20/30, 172.16.1.32/30
			172.16.1.40/30, 10.2.2.2/32,11.5.5.5/32
Leaf-3	0	11.6.6.6	172.16.1.36/30, 10.3.3.3/32, 172.16.1.56/30, 11.6.6.6/32

Configure OSPF for Leaf-1:

```
<Leaf-1>system-view
[Leaf-1]ospf 1 router-id 11.3.3.3
[Leaf-1-ospf-1] stub-router on-startup 600 include-stub // Configure a Stub
router to improve network convergence performance in fault scenarios. When
the Spine switch is a CE12800 series switch, set the on-startup interval to
3000s. When the spine switch is a ToR switch, set the interval to 600s.
```

```
[Leaf-1-ospf-1]area 0
[Leaf-1-ospf-1-area-0.0.0.0]network 10.1.1.1 0.0.0.0
[Leaf-1-ospf-1-area-0.0.0.0]network 11.3.3.3 0.0.0.0
[Leaf-1-ospf-1-area-0.0.0.0]network 172.16.1.12 0.0.0.3
[Leaf-1-ospf-1-area-0.0.0.0]network 172.16.1.24 0.0.0.3
[Leaf-1-ospf-1-area-0.0.0.0]network 172.16.1.48 0.0.0.3
[Leaf-1-ospf-1-area-0.0.0.0]network 172.16.1.52 0.0.0.3
[Leaf-1-ospf-1-area-0.0.0.0]quit
<Leaf-1>
```

Configure OSPF for Spine-1:

```
[Spine-1]ospf 1 router-id 11.1.1.1
[Spine-1-ospf-1]stub-router on-startup 600 include-stub
[Spine-1-ospf-1]area 0
[Spine-1-ospf-1-area-0.0.0.0]network 172.16.1.44 0.0.0.3
[Spine-1-ospf-1-area-0.0.0.0]network 172.16.1.12 0.0.0.3
[Spine-1-ospf-1-area-0.0.0.0]network 172.16.1.16 0.0.0.3
[Spine-1-ospf-1-area-0.0.0.0]network 172.16.1.20 0.0.0.3
[Spine-1-ospf-1-area-0.0.0.0]network 172.16.1.48 0.0.0.3
[Spine-1-ospf-1-area-0.0.0.0]network 172.16.1.56 0.0.0.3
[Spine-1-ospf-1-area-0.0.0.0]network 10.45.45.45 0.0.0.0
[Spine-1-ospf-1-area-0.0.0.0]network 11.1.1.1 0.0.0.0
[Spine-1-OSPF-1-area-0.0.0.0]quit
```

Configure OSPF for Spine-2:

```
[Spine-2]ospf 1 router-id 11.2.2.2
[Spine-2-ospf-1]stub-router on-startup 600 include-stub
[Spine-2-ospf-1]area 0
[Spine-2-ospf-1-area-0.0.0.0]network 172.16.1.44 0.0.0.3
[Spine-2-ospf-1-area-0.0.0.0]network 172.16.1.24 0.0.0.3
[Spine-2-ospf-1-area-0.0.0.0]network 172.16.1.28 0.0.0.3
[Spine-2-ospf-1-area-0.0.0.0]network 172.16.1.32 0.0.0.3
[Spine-2-ospf-1-area-0.0.0.0]network 172.16.1.36 0.0.0.3
[Spine-2-ospf-1-area-0.0.0.0]network 172.16.1.52 0.0.0.3
[Spine-2-ospf-1-area-0.0.0.0]network 10.45.45.45 0.0.0.0
[Spine-2-ospf-1-area-0.0.0.0]network 11.2.2.2 0.0.0.0
[Spine-2-OSPF-1-area-0.0.0.0]quit
```

Configure OSPF for Leaf-1:

```
[Leaf-2A]ospf 1 router-id 11.4.4.4
[Leaf-2A-ospf-1]stub-router on-startup 600 include-stub
[Leaf-2A-ospf-1]area 0
[Leaf-2A-ospf-1-area-0.0.0.0]network 172.16.1.16 0.0.0.3
[Leaf-2A-ospf-1-area-0.0.0.0]network 172.16.1.28 0.0.0.3
[Leaf-2A-ospf-1-area-0.0.0.0]network 172.16.1.40 0.0.0.3
[Leaf-2A-ospf-1-area-0.0.0.0]network 10.2.2.2 0.0.0.0
[Leaf-2A-ospf-1-area-0.0.0.0]network 11.4.4.4 0.0.0.0
```

Configure OSPF for Leaf-2B:

```
[Leaf-2B]ospf 1 router-id 11.5.5.5
[Leaf-2B-ospf-1]stub-router on-startup 600 include-stub
[Leaf-2B-ospf-1]area 0
[Leaf-2B-ospf-1-area-0.0.0.0]network 172.16.1.20 0.0.0.3
[Leaf-2B-ospf-1-area-0.0.0.0]network 172.16.1.32 0.0.0.3
[Leaf-2B-ospf-1-area-0.0.0.0]network 172.16.1.40 0.0.0.3
[Leaf-2B-ospf-1-area-0.0.0.0]network 10.2.2.2 0.0.0.0
[Leaf-2B-ospf-1-area-0.0.0.0]network 11.5.5.5 0.0.0.0
```

Configure OSPF for Leaf-3:

```
[Leaf-3]ospf 1 router-id 11.6.6.6
[Leaf-3-ospf-1]stub-router on-startup 600 include-stub
[Leaf-3-ospf-1]area 0
[Leaf-3-ospf-1-area-0.0.0.0]network 172.16.1.36 0.0.0.3
[Leaf-3-ospf-1-area-0.0.0.0]network 172.16.1.56 0.0.0.3
[Leaf-3-ospf-1-area-0.0.0.0]network 10.3.3.3 0.0.0.0
[Leaf-3-ospf-1-area-0.0.0.0]network 11.6.6.6 0.0.0.0
```

Check the OSPF neighbor relationship. The following takes Spine-2 as an example:

```
<Spine-2>display ospf peer brief
OSPF Process 1 with Router ID 10.0.0.3
                Peer Statistic Information
Total number of peer(s): 6
 Peer(s) in full state: 6
--------------------------------------------------------------------------------
 Area Id         Interface                Neighbor id        State
 0.0.0.0         10GE1/0/16               11.6.6.6           Full
 0.0.0.0         10GE1/0/17               11.1.1.1           Full
 0.0.0.0         10GE1/0/20               11.3.3.3           Full
 0.0.0.0         10GE1/0/21               11.3.3.3           Full
 0.0.0.0         10GE1/0/22               11.4.4.4           Full
 0.0.0.0         10GE1/0/23               11.5.5.5           Full
--------------------------------------------------------------------------------
```

View the OSPF route table. The following takes Spine-2 as an example:

```
<Spine-2>display  ip routing-table protocol  ospf
Proto: Protocol       Pre: Preference
Route Flags: R - relay, D - download to fib, T - to vpn-instance, B - black
hole route
------------------------------------------------------------------------------
-
_public_ Routing Table : OSPF
        Destinations : 22      Routes : 33

OSPF routing table status : <Active>
        Destinations : 14      Routes : 25

Destination/Mask    Proto   Pre  Cost       Flags NextHop         Interface

      10.1.1.1/32  OSPF    10   1           D   172.16.1.26     10GE1/0/21
                   OSPF    10   1           D   172.16.1.54     10GE1/0/20
```

```
      10.2.2.2/32  OSPF    10   1           D   172.16.1.34     10GE1/0/23
                   OSPF    10   1         D   172.16.1.30     10GE1/0/22
      10.3.3.3/32  OSPF    10   1           D   172.16.1.38     10GE1/0/16
      11.1.1.1/32  OSPF    10   1           D   172.16.1.45     10GE1/0/17
      11.3.3.3/32  OSPF    10   1           D   172.16.1.26     10GE1/0/21
                   OSPF    10   1         D   172.16.1.54     10GE1/0/20
      11.4.4.4/32  OSPF    10   1           D   172.16.1.30     10GE1/0/22
      11.5.5.5/32  OSPF    10   1           D   172.16.1.34     10GE1/0/23
      11.6.6.6/32  OSPF    10   1           D   172.16.1.38     10GE1/0/16
   172.16.1.12/30  OSPF    10   2           D   172.16.1.26     10GE1/0/21
                   OSPF    10   2         D   172.16.1.54     10GE1/0/20
                   OSPF    10   2         D   172.16.1.45     10GE1/0/17
   172.16.1.16/30  OSPF    10   2           D   172.16.1.30     10GE1/0/22
                   OSPF    10   2         D   172.16.1.45     10GE1/0/17
   172.16.1.20/30  OSPF    10   2           D   172.16.1.34     10GE1/0/23
                   OSPF    10   2         D   172.16.1.45     10GE1/0/17
   172.16.1.40/30  OSPF    10   2           D   172.16.1.34     10GE1/0/23
                   OSPF    10   2         D   172.16.1.30     10GE1/0/22
   172.16.1.48/30  OSPF    10   2           D   172.16.1.26     10GE1/0/21
                   OSPF    10   2         D   172.16.1.54     10GE1/0/20
                   OSPF    10   2         D   172.16.1.45     10GE1/0/17
   172.16.1.56/30  OSPF    10   2           D   172.16.1.45     10GE1/0/17
                   OSPF    10   2         D   172.16.1.38     10GE1/0/16

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

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

   10.45.45.45/32  OSPF    10   0             10.45.45.45     LoopBack0
     11.2.2.2/32  OSPF    10   0            11.2.2.2        LoopBack1
   172.16.1.24/30  OSPF    10   1             172.16.1.25     10GE1/0/21
   172.16.1.28/30  OSPF    10   1             172.16.1.29     10GE1/0/22
   172.16.1.32/30  OSPF    10   1             172.16.1.33     10GE1/0/23
   172.16.1.36/30  OSPF    10   1             172.16.1.37     10GE1/0/16
   172.16.1.44/30  OSPF    10   1             172.16.1.46     10GE1/0/17
   172.16.1.52/30  OSPF    10   1             172.16.1.53     10GE1/0/20
```

3.4.6 OSPF Complete Configuration

Complete configuration of Leaf-1:

```
#
ospf 1 router-id 11.3.3.3
 stub-router on-startup 600 include-stub
 area 0.0.0.0
  network 10.1.1.1 0.0.0.0
  network 11.3.3.3 0.0.0.0
  network 172.16.1.12 0.0.0.3
  network 172.16.1.24 0.0.0.3
  network 172.16.1.48 0.0.0.3
  network 172.16.1.52 0.0.0.3
#
```

Complete configuration of Leaf-3:

```
ospf 1 router-id 11.6.6.6
 stub-router on-startup 600 include-stub
 area 0.0.0.0
  network 172.16.1.36 0.0.0.3
network 172.16.1.56 0.0.0.3
  network 10.3.3.3 0.0.0.0
  network 11.6.6.6 0.0.0.0
#
```

Complete configuration of Spine-1:

```
#
ospf 1 router-id 11.1.1.1
 stub-router on-startup 600 include-stub
 area 0.0.0.0
  network 172.16.1.12 0.0.0.3
  network 172.16.1.16 0.0.0.3
  network 172.16.1.20 0.0.0.3
  network 172.16.1.44 0.0.0.3
  network 172.16.1.48 0.0.0.3
network 172.16.1.56 0.0.0.3
  network 10.45.45.45 0.0.0.0
  network 11.1.1.1 0.0.0.0
#
```

Complete configuration of Spine-2:

```
#
ospf 1 router-id 11.2.2.2
 stub-router on-startup 600 include-stub
 area 0.0.0.0
  network 172.16.1.24 0.0.0.3
  network 172.16.1.28 0.0.0.3
  network 172.16.1.32 0.0.0.3
  network 172.16.1.36 0.0.0.3
  network 172.16.1.44 0.0.0.3
  network 172.16.1.52 0.0.0.3
  network 10.45.45.45 0.0.0.0
  network 11.2.2.2 0.0.0.0
#
```

Complete configuration of Leaf-2A:

```
#
ospf 1 router-id 11.4.4.4
 stub-router on-startup 600 include-stub
 area 0.0.0.0
  network 172.16.1.16 0.0.0.3
  network 172.16.1.28 0.0.0.3
  network 172.16.1.40 0.0.0.3
  network 10.2.2.2 0.0.0.0
```

```
  network 11.4.4.4 0.0.0.0
#
```

Complete configuration of Leaf-2B:

```
ospf 1 router-id 11.5.5.5
 stub-router on-startup 600 include-stub
 area 0.0.0.0
  network 172.16.1.20 0.0.0.3
  network 172.16.1.32 0.0.0.3
  network 172.16.1.40 0.0.0.3
  network 10.2.2.2 0.0.0.0
  network 11.5.5.5 0.0.0.0
#
```

3.5 Configuring M-LAG Work Groups for Gateway Nodes

3.5.1 Objectives

- Understand the basic functions of M-LAG.
- Understand how to configure the M-LAG active-active.

3.5.2 Networking and Service Description

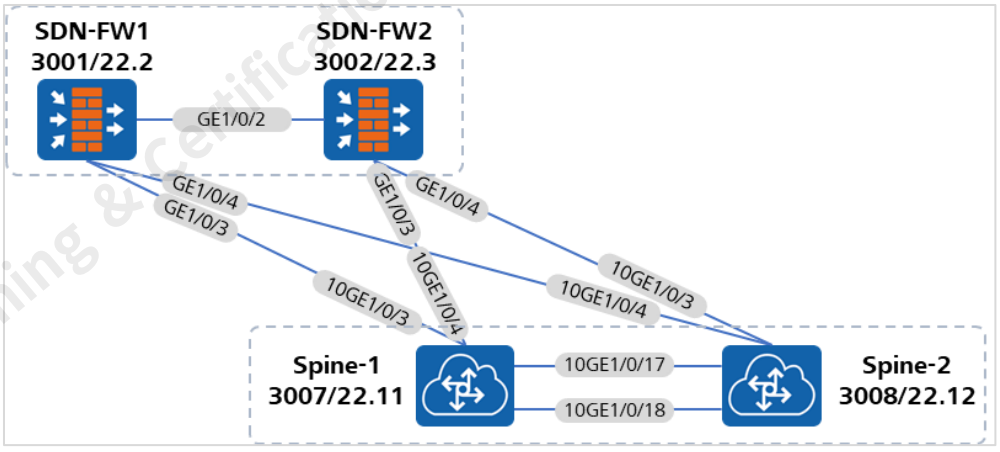

Figure 3-4 Network diagram of M-LAG experiment

Configure M-LAG on Spine-1 and Spine-2 to implement dual-homing access of firewalls.

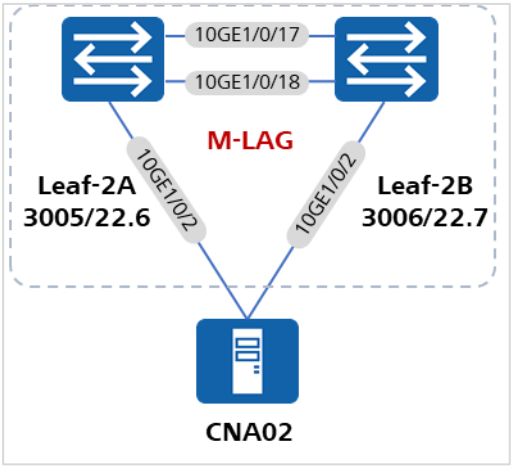

Figure 3-5 Configure M-LAG on hosts to implement dual-homing access of hosts.

3.5.3 Configuration Guideline

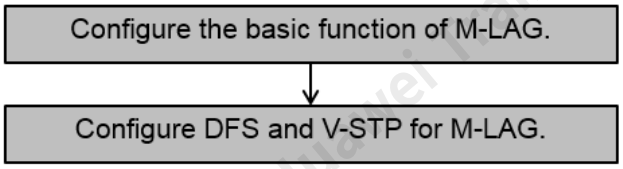

3.5.4 Configuring M-LAG for Spines

Configure the DFS Group, peer-link, and M-LAG interfaces for Spine-1 and Spine-2.

Configurations for Spine-1:

```
<Spine-1>system-view immediately
[Spine-1] dfs-group 1
[Spine-1-dfs-group-1] source ip 172.16.1.45
[Spine-1-dfs-group-1] priority 150
[Spine-1-dfs-group-1]active-active-gateway
[Spine-1-dfs-group-1-active-active-gateway]peer 172.16.1.46 // Configure an
ARP synchronization address for the active-active gateway. Set Spine-1 to
172.16.1.45 and Spine-2 to 172.16.1.46.
[Spine-1-dfs-group-1] quit
[Spine-1] stp mode rstp
[Spine-1] stp v-stp enable
[Spine-1] lacp m-lag priority 10
[Spine-1] lacp m-lag system-id 00e0-fc00-0101
[Spine-1] interface eth-trunk 23
[Spine-1-Eth-Trunk23] trunkport 10ge 1/0/18 // 10GE1/0/18 is a Peer-link
synchronization interface of DFS.
[Spine-1-Eth-Trunk23] mode lacp-static
[Spine-1-Eth-Trunk23] peer-link 1
[Spine-1-Eth-Trunk23] quit
[Spine-1] interface Eth-Trunk24
```

```
[Spine-1-Eth-Trunk24] description to-USG6620-1-GE1/0/3
[Spine-1-Eth-Trunk24] trunkport 10ge 1/0/3 // 10GE1/0/3 is the interface
connected to the firewall. Add it to the M-LAG synchronization group and
configure it as a trunk interface.
[Spine-1-Eth-Trunk24] port link-type trunk
[Spine-1-Eth-Trunk24] undo port trunk allow-pass vlan 1
[Spine-1-Eth-Trunk24] dfs-group 1 m-lag 1
[Spine-1-Eth-Trunk24] interface Eth-Trunk34
[Spine-1-Eth-Trunk34] description to-USG6620-2-GE1/0/4
[Spine-1-Eth-Trunk34] trunkport 10ge 1/0/4 // 10GE1/0/4 is the interface
connected to the firewall. Add it to the M-LAG synchronization group and
configure it as a trunk interface.
[Spine-1-Eth-Trunk34] port link-type trunk
[Spine-1-Eth-Trunk34] undo port trunk allow-pass vlan 1
[Spine-1-Eth-Trunk34] dfs-group 1 m-lag 2
```

Configurations for Spine-2:

```
<Spine-2>system-view immediately
[Spine-2] dfs-group 1
[Spine-2-dfs-group-1] source ip 172.16.1.46
[Spine-2-dfs-group-1] priority 120
[Spine-2-dfs-group-1] active-active-gateway
[Spine-2-dfs-group-1-active-active-gateway] peer 172.16.1.45
[Spine-2-dfs-group-1] quit
[Spine-2] stp mode rstp
[Spine-2] stp v-stp enable
[Spine-2] lacp m-lag priority 10
[Spine-2] lacp m-lag system-id 00e0-fc00-0101
[Spine-2] interface eth-trunk 23
[Spine-2-Eth-Trunk23] trunkport 10ge 1/0/18
[Spine-2-Eth-Trunk23] mode lacp-static
[Spine-2-Eth-Trunk23] peer-link 1
[Spine-2-Eth-Trunk23] quit
[Spine-2] interface Eth-Trunk24
[Spine-2-Eth-Trunk24] description to-USG6620-1-GE1/0/3
[Spine-2-Eth-Trunk24] trunkport 10ge 1/0/3
[Spine-2-Eth-Trunk24] port link-type trunk
[Spine-2-Eth-Trunk24] undo port trunk allow-pass vlan 1
[Spine-2-Eth-Trunk24] dfs-group 1 m-lag 1
[Spine-2-Eth-Trunk24] interface Eth-Trunk34
[Spine-2-Eth-Trunk34] description to-USG6620-2-GE1/0/4
[Spine-2-Eth-Trunk34] trunkport 10ge 1/0/4
[Spine-2-Eth-Trunk34] port link-type trunk
[Spine-2-Eth-Trunk34] undo port trunk allow-pass vlan 1
[Spine-2-Eth-Trunk34] dfs-group 1 m-lag 2
```

3.5.4.1 Verifying the M-LAG Result

Run the **display dfs-group** command to check the M-LAG information. Spine-1 is used as an example.

```
[Spine-1]dis dfs-group 1 m-lag
*              : Local node
```

```
Heart beat state : OK
Node 1 *
  Dfs-Group ID   : 1
  Priority       : 150
  Address        : ip address 172.21.1.45
  State          : Master
  Causation      : -
  System ID      : c81f-be6f-ba21
  SysName        : Spine-1
  Version        : V200R002C50
  Device Type    : CE6851HI
Node 2
  Dfs-Group ID   : 1
  Priority       : 120
  Address        : ip address 172.21.1.46
  State          : Backup
  Causation      : -
  System ID      : c81f-be6f-b9c1
  SysName        : Spine-2
  Version        : V200R002C50
  Device Type    : CE12800
```

Check M-LAG information on Spine-1.

```
[Spine-1]dis dfs-group 1 node 1 m-lag brief
* - Local node

M-Lag ID    Interface     Port State   Status
      1    Eth-Trunk 24   Up          active(*)-active
      2    Eth-Trunk 34   Up          active(*)-active
```

Check M-LAG information on Spine-2.

```
<Spine-2>display dfs-group 1 node 1 m-lag brief
* - Local node
M-Lag ID Interface  Port StateStatus
   1 Eth-Trunk 24   active -active(*)
   2 Eth-Trunk 34   active -active(*)
```

Check the link aggregation information on Spine-1.

```
<Spine-1>display eth-trunk 24
Eth-Trunk24's state information is:
Working Mode: NormalHash Arithmetic: profile default
Least Active-linknumber: 1  Max Bandwidth-affected-linknumber: 16
Operating Status: upNumber of Up Ports in Trunk: 1
--------------------------------------------------------------------------------
PortName   Status  Weight
10GE1/0/3  Up  1

<Spine-1>display eth-trunk 34
Eth-Trunk34's state information is:
Working Mode: NormalHash Arithmetic: profile default
Least Active-linknumber: 1  Max Bandwidth-affected-linknumber: 16
```

```
Operating Status: upNumber of Up Ports in Trunk: 1
--------------------------------------------------------------------------------
PortName  Status  Weight
10GE1/0/4 Up  1
```

Check the link aggregation information on Spine-2.

```
<Spine-2>display eth-trunk 24
Eth-Trunk24's state information is:
Working Mode: NormalHash Arithmetic: profile default
Least Active-linknumber: 1  Max Bandwidth-affected-linknumber: 16
Operating Status: upNumber of Up Ports in Trunk: 1
--------------------------------------------------------------------------------
PortName  Status  Weight
10GE1/0/3 Up  1

<Spine-2> display eth-trunk 34
Eth-Trunk34's state information is:
Working Mode: NormalHash Arithmetic: profile default
Least Active-linknumber: 1  Max Bandwidth-affected-linknumber: 16
Operating Status: upNumber of Up Ports in Trunk: 1
--------------------------------------------------------------------------------
PortName  Status  Weight
10GE1/0/4 Up  1
```

3.5.4.2 M-LAG Complete Configuration

Configurations for Spine-1:

```
#
dfs-group 1
 priority 150
 source ip 172.16.1.45
 #
 active-active-gateway
  peer 172.16.1.46
#
lacp m-lag system-id 00e0-fc00-0101
lacp m-lag priority 10
#
interface Eth-Trunk23
 mode lacp-static
 peer-link 1
#
interface Eth-Trunk24
 description to-USG6620-1-GE1/0/3
 port link-type trunk
 undo port trunk allow-pass vlan 1
 dfs-group 1 m-lag 1
#
interface Eth-Trunk34
 description to-USG6620-2-GE1/0/4
 port link-type trunk
 undo port trunk allow-pass vlan 1
```

```
 dfs-group 1 m-lag 2
#
interface 10GE1/0/18
 eth-trunk 23
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/3
 eth-trunk 24
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/4
 eth-trunk 34
 device transceiver 10GBASE-COPPER
```

Complete configuration of Spine-2:

```
#
dfs-group 1
 priority 120
 source ip 172.16.1.46
 #
 active-active-gateway
  peer 172.16.1.45
#
lacp m-lag system-id 00e0-fc00-0101
lacp m-lag priority 10
#
interface Eth-Trunk23
 mode lacp-static
 peer-link 1
#
interface Eth-Trunk24
 description to-USG6620-1-GE1/0/3
 port link-type trunk
 undo port trunk allow-pass vlan 1
 dfs-group 1 m-lag 1
#
interface Eth-Trunk34
 description to-USG6620-2-GE1/0/4
 port link-type trunk
 undo port trunk allow-pass vlan 1
 dfs-group 1 m-lag 2
#
interface 10GE1/0/18
 eth-trunk 23
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/3
 eth-trunk 24
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/4
 eth-trunk 34
 device transceiver 10GBASE-COPPER
```

3.5.5 Configuring M-LAG for Server Leaf

Configure DFS-Group, Peer-Link, and M-LAG interfaces for Leaf-2A and Leaf-2B.

Configurations of Leaf-2A:

```
<Leaf-2A>system-view immediately
[Leaf-2A] dfs-group 1
[Leaf-2A-dfs-group-1] source ip 172.16.1.41
[Leaf-2A-dfs-group-1] priority 150
[Leaf-2A-dfs-group-1] quit
[Leaf-2A] stp mode rstp
[Leaf-2A] stp v-stp enable
[Leaf-2A] lacp m-lag priority 10
[Leaf-2A] lacp m-lag system-id a008-6fe0-fc30 // It is recommended that the
MAC system (master) in the M-LAG be used as the system-id. The system-id
configurations on the peer device must be the same. You can run the display
system mac-address command to check the MAC system.
[Leaf-2A] interface eth-trunk 23
[Leaf-2A-Eth-Trunk23] trunkport 10ge 1/0/18 // 10GE1/0/18 is a Peer-link
synchronization interface of DFS.
[Leaf-2A-Eth-Trunk23] mode lacp-static
[Leaf-2A-Eth-Trunk23] peer-link 1
[Leaf-2A-Eth-Trunk23] quit
[Leaf-2A] interface Eth-Trunk24
[Leaf-2A-Eth-Trunk24] trunkport 10ge 1/0/2 // 10GE1/0/2 is the interface
connected to the host. Add it to the M-LAG synchronization group and
configure it as a trunk interface.
[Leaf-2A-Eth-Trunk24] port link-type trunk
[Leaf-2A-Eth-Trunk24] port default vlan 4005 4011
[Leaf-2A-Eth-Trunk24] stp edged-port enable
[Leaf-2A-Eth-Trunk24] dfs-group 1 m-lag 1
```

Configurations of Leaf-2B:

```
<Leaf-2B>system-view immediately
[Leaf-2B] dfs-group 1
[Leaf-2B-dfs-group-1] source 172.16.1.42
[Leaf-2B-dfs-group-1] quit
[Leaf-2B] stp mode rstp
[Leaf-2B] stp v-stp enable
[Leaf-2B] lacp m-lag system-id a008-6fe0-fc30
[Leaf-2B] interface eth-trunk 23
[Leaf-2B-Eth-Trunk23] trunkport 10ge 1/0/18
[Leaf-2B-Eth-Trunk23] mode lacp-static
[Leaf-2B-Eth-Trunk23] peer-link 1
[Leaf-2B-Eth-Trunk23] quit
[Leaf-2B] interface Eth-Trunk24
[Leaf-2B-Eth-Trunk24] trunkport 10ge 1/0/2
[Leaf-2B-Eth-Trunk24] port default vlan 4005 4011
[Leaf-2B-Eth-Trunk24] stp edged-port enable
[Leaf-2B-Eth-Trunk24] dfs-group 1 m-lag 1
```

3.5.5.1 Verifying the M-LAG Result

Run the **display dfs-group** command to check the M-LAG information. Leaf-2A is used as an example.

```
[Leaf-2A]display dfs-group 1 m-lag
*              : Local node
Heart beat state : OK
Node 1 *
  Dfs-Group ID   : 1
  Priority       : 150
  Address        : ip address 172.16.1.41
  State          : Master
  Causation      : -
  System ID      : 8866-394a-80e1
  SysName        : Leaf-2A
  Version        : V200R002C50
  Device Type    : CE6851HI
Node 2
  Dfs-Group ID   : 1
  Priority       : 100
  Address        : ip address 172.16.1.42
  State          : Backup
  Causation      : -
  System ID      : 8866-394a-80f1
  SysName        : Leaf-2B
  Version        : V200R002C50
  Device Type    : CE6851HI
```

Check M-LAG information on Leaf-2A.

```
[Leaf-2A]dis dfs-group 1 node 1 m-lag brief
* - Local node

M-Lag ID    Interface      Port State    Status
      1     Eth-Trunk 24   Up           active(*)-active
```

3.5.5.2 M-LAG Complete Configuration

Configurations of Leaf-2A:

```
#
dfs-group 1
 priority 150
 source ip 172.16.1.41
 #
lacp m-lag system-id a008-6fe0-fc30
lacp m-lag priority 10
#
interface Eth-Trunk23
 mode lacp-static
 peer-link 1
```

```
#
interface Eth-Trunk24
port default vlan 4005 4011
stp edged-port enable
  dfs-group 1 m-lag 1
#
interface 10GE1/0/18
 eth-trunk 23
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/1
 eth-trunk 24
 device transceiver 10GBASE-COPPER
#
```

Complete configuration of Leaf-2B:

```
#
dfs-group 1
 priority 120
 source ip 172.16.1.42
#
lacp m-lag system-id a008-6fe0-fc30
#
interface Eth-Trunk23
 mode lacp-static
 peer-link 1
#
interface Eth-Trunk24
port default vlan 4005 4011
stp edged-port enable
  dfs-group 1 m-lag 1
#
interface 10GE1/0/18
 eth-trunk 23
 device transceiver 10GBASE-COPPER
#
interface 10GE1/0/1
 eth-trunk 24
 device transceiver 10GBASE-COPPER
#
```


3.6 Configuring Firewalls

3.6.1 Objectives

- Learn the zone division and security protection functions of NGFW.
- Learn the HA HSB function of NGFW.
- Learn the security policies and filtering functions of NFGW.

3.6.2 Networking and Service Description

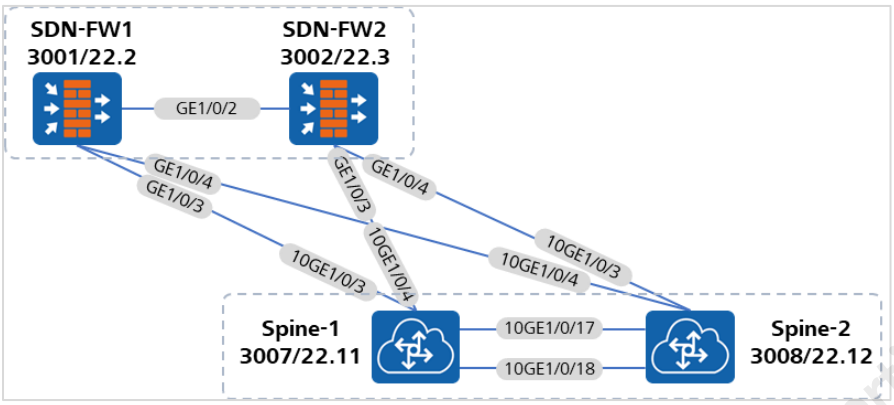

Figure 3-6 Security protection topology

Two physical firewalls have been connected to the Agile Controller-DCN through the management network. The Layer 3 link between the firewall and the gateway switch is automatically configured when the Agile Controller-DCN delivers services.

In the Underlay pre-configuration phase, you only need to configure the HSB mirroring mode of the firewall and the Eth-Trunk connected to the gateway switch. The switch has been configured in the M-LAG.

Two firewalls in active/standby mirroring mode must use the same interface to connect to the same GW group device. The SDN-FW1 uses the 10GE 1/0/3 interface to connect to the 10GE1/0/3 of the Spine-1, and the SDN-FW2 uses the 10GE 1/0/3 interface to connect to the 10GE1/0/4 of the Spine-1.

3.6.3 Configuration Guideline

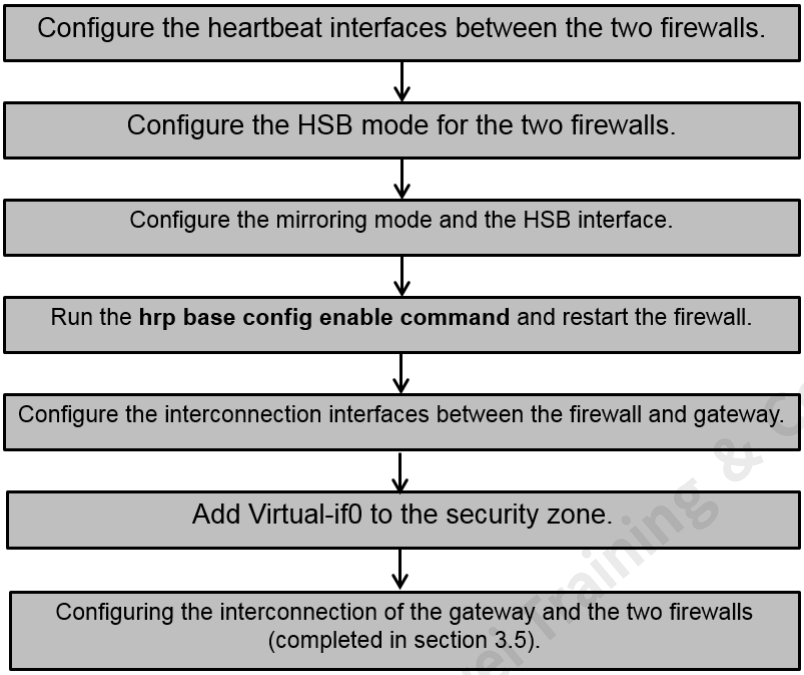

3.6.4 Device Physical Connection

Table 3-3 Physical connection

Device Name	Interface Name	Peer Device	Peer Interface
Spine-1	10GE1/0/3	SDN-FW1	GE1/0/3
	10GE1/0/4	SDN-FW2	GE1/0/3
Spine-2	10GE1/0/3	SDN-FW1	GE1/0/4
	10GE1/0/4	SDN-FW2	GE1/0/4
SDN-FW1	G1/0/2	SDN-FW2	G1/0/2

Configure the heartbeat interfaces G1/0/2 of the two firewalls and add them to the DMZ zone.

```
<SDN-FW1>system-view
Enter system view, return user view with Ctrl+Z.
[SDN-FW1]interface GigabitEthernet 1/0/2
[SDN-FW1-GigabitEthernet1/0/2] ip address 172.2.2.3 255.255.255.0
[SDN-FW1-GigabitEthernet1/0/2]quit
[SDN-FW1]firewall zone dmz
[SDN-FW1-zone-dmz]add interface GigabitEthernet 1/0/2
[SDN-FW1-zone-dmz]quit
```

```
[SDN-FW1]interface GigabitEthernet 1/0/1
[SDN-FW1-GigabitEthernet1/0/1] ip address 172.21.22.3 255.255.0.0
[SDN-FW1-GigabitEthernet1/0/1] service-manage http permit
[SDN-FW1-GigabitEthernet1/0/1] service-manage https permit
[SDN-FW1-GigabitEthernet1/0/1] service-manage ping permit
[SDN-FW1-GigabitEthernet1/0/1] service-manage ssh permit
[SDN-FW1-GigabitEthernet1/0/1] service-manage snmp permit
[SDN-FW1-GigabitEthernet1/0/1] service-manage telnet permit
[SDN-FW1-GigabitEthernet1/0/1] service-manage netconf permit
[SDN-FW1-GigabitEthernet1/0/1]quit
[SDN-FW1]firewall zone trust
[SDN-FW1-zone-trust]add interface GigabitEthernet 1/0/1
[SDN-FW1-zone-trust]quit
[SDN-FW1]

<SDN-FW2>system-view
Enter system view, return user view with Ctrl+Z.
[SDN-FW2]interface GigabitEthernet 1/0/2
[SDN-FW2-GigabitEthernet1/0/2] ip address 172.2.2.2 255.255.255.0
[SDN-FW2-GigabitEthernet1/0/2]quit
[SDN-FW2]firewall zone dmz
[SDN-FW2-zone-dmz]add interface GigabitEthernet 1/0/2
[SDN-FW2-zone-dmz]quit
[SDN-FW2]interface GigabitEthernet 1/0/1
[SDN-FW2-GigabitEthernet1/0/1] ip address 172.21.22.2 255.255.0.0
[SDN-FW2-GigabitEthernet1/0/1] service-manage http permit
[SDN-FW2-GigabitEthernet1/0/1] service-manage https permit
[SDN-FW2-GigabitEthernet1/0/1] service-manage ping permit
[SDN-FW2-GigabitEthernet1/0/1] service-manage ssh permit
[SDN-FW2-GigabitEthernet1/0/1] service-manage snmp permit
[SDN-FW2-GigabitEthernet1/0/1] service-manage telnet permit
[SDN-FW2-GigabitEthernet1/0/1] service-manage netconf permit
[SDN-FW2-GigabitEthernet1/0/1]quit
[SDN-FW2]firewall zone trust
[SDN-FW2-zone-trust]add interface GigabitEthernet 1/0/1
[SDN-FW2-zone-trust]quit
[SDN-FW2]
```

Configure the port for connecting to the gateway switches Spine-1 and Spine-2.

```
<SDN-FW1> system-view
[SDN-FW1] interface Eth-Trunk1
[SDN-FW1-Eth-Trunk0] quit
[SDN-FW1] interface GigabitEthernet1/0/3
[SDN-FW1-GigabitEthernet1/0/3] description To-Spine-1
[SDN-FW1-GigabitEthernet1/0/3] undo shutdown
[SDN-FW1-GigabitEthernet1/0/3] eth-trunk 1
[SDN-FW1-GigabitEthernet1/0/3] quit
[SDN-FW1] interface GigabitEthernet1/0/4
[SDN-FW1-GigabitEthernet1/0/4] description To-Spine-2
[SDN-FW1-GigabitEthernet1/0/4] undo shutdown
[SDN-FW1-GigabitEthernet1/0/4] eth-trunk 1
[SDN-FW1-GigabitEthernet1/0/4] quit
[SDN-FW1]interface Eth-Trunk 1
[SDN-FW1]portswitch
```

```
[SDN-FW1]port link-type trunk
[SDN-FW1]undo port trunk allow-pass vlan 1
```

Perform the same configurations for SDN-FW2.

Configure the HSB mode of the firewall. Set SDN-FW1 as the active node and SDN-FW2 as the standby node.

Configure the HSB mode on SDN-FW1.

```
<SDN-FW1> system-view
[SDN-FW1] hrp track interface Eth-Trunk1
[SDN-FW1] hrp mgt interface GigabitEthernet1/0/1
[SDN-FW1] hrp interface GigabitEthernet1/0/2 remote 172.2.2.2
[SDN-FW1] hrp enable
```

Configure the HSB mode on SDN-FW2.

```
<SDN-FW2> system-view
[SDN-FW2] hrp track interface Eth-Trunk1
[SDN-FW2] hrp mgt-interface GigabitEthernet1/0/1
[SDN-FW2] hrp interface GigabitEthernet1/0/2 remote 172.2.2.3
[SDN-FW2] hrp enable
```

Configure the mirroring mode and the management interface in HSB mode.

Configure the mirroring mode and the management interface in HSB mode on SDN-FW1.

```
HRP_M[SDN-FW1] hrp mirror config enable
HRP_M[SDN-FW1] hrp mirror session enable
HRP_M[SDN-FW1] hrp standby config enable
HRP_M[SDN-FW1] undo hrp preempt
```

Configure the mirroring mode and the management interface in HSB mode on SDN-FW2.

```
HRP_S[SDN-FW2] hrp mirror config enable
HRP_S[SDN-FW2] hrp mirror session enable
HRP_S[SDN-FW2] hrp standby config enable
HRP_S[SDN-FW2] undo hrp preempt
```

Restart the firewall in the HSB mode and synchronize common service configurations from another firewall.

Run the hrp base config enable command.

```
HRP_M[SDN-FW1] hrp base config enable
HRP_S[SDN-FW2] hrp base config enable
```

Restart the firewall to enable the configurations.

```
HRP_M[SDN-FW1]  reboot
  System will reboot! Do you want to save the running configuration? [Y/N]:n
  System will reboot! Continue? [Y/N]:y

HRP_M[SDN-FW2]  reboot
  System will reboot! Do you want to save the running configuration? [Y/N]:n
  System will reboot! Continue? [Y/N]:y
```

Add the port to the security zone and configure the default security policy.

Add **Virtual-if0** to the security zone for traffic diversion between the root firewall and the virtual firewall.

```
HRP_M[SDN-FW1] firewall zone untrust
HRP_M[SDN-FW1-zone-untrust] add interface Virtual-if0
HRP_M[SDN-FW1-zone-untrust] quit
```

Configure the default security policy to **permit**.

```
HRP_M[SDN-FW1] security-policy
HRP_M[SDN-FW1-security-policy] default action permit
HRP_M[SDN-FW1]vsys enable // Enable the vsys function on the firewall.
HRP_M[SDN-FW1]interface Virtual-if api transform // Enable the Virtual-if
name conversion function of the northbound interface.
```

Configure static routes to enable Agile Controller-DCN in-band management.

```
HRP_M[SDN-FW1]ip route-static 192.168.4.0 24 172.21.22.13
```

3.6.5 Complete Configuration

SDN-FW1 complete configuration

```
#
hrp enable
 hrp mirror config enable
 hrp interface GigabitEthernet1/0/2 remote  172.2.2.2
 hrp mgt-interface GigabitEthernet1/0/1
 hrp base config enable
 hrp mirror session enable
 hrp standby config enable
 hrp configuration auto-check 1440
 undo hrp preempt
 hrp track interface Eth-Trunk
#
interface Eth-Trunk1
 portswitch
 port link-type trunk
 undo port trunk allow-pass vlan 1
#
interface GigabitEthernet1/0/2
 undo shutdown
```

```
 ip address 10.1.2.3 255.255.255.0
#
interface GigabitEthernet1/0/1
 undo shutdown
 ip address 172.21.22.2 255.255.0.0
 service-manage http permit
 service-manage https permit
 service-manage ping permit
 service-manage ssh permit
 service-manage snmp permit
 service-manage telnet permit
 service-manage netconf permit
#
interface GigabitEthernet1/0/3
 undo shutdown
 eth-trunk 1
#
interface GigabitEthernet1/0/4
 undo shutdown
 eth-trunk 1
#
interface Virtual-if0
#
firewall zone trust
 set priority 85
 add interface GigabitEthernet0/0/0
 add interface GigabitEthernet1/0/1
#
firewall zone untrust
 set priority 5
 add interface Virtual-if0
#
firewall zone dmz
 set priority 50
 add interface GigabitEthernet1/0/2
#
ip route-static 192.168.4.0 255.255.255.0 172.21.22.13
```

SDN-FW2 complete configuration

```
#
 hrp enable
 hrp mirror config enable
 hrp interface GigabitEthernet1/0/2 remote 172.2.2.3
 hrp mgt-interface GigabitEthernet1/0/1
 hrp base config enable
 hrp mirror session enable
 hrp standby config enable
 hrp configuration auto-check 1440
 undo hrp preempt
 hrp track interface Eth-Trunk
#
interface Eth-Trunk1
 portswitch
 port link-type trunk
```

```
 undo port trunk allow-pass vlan 1
#
interface GigabitEthernet1/0/2
 undo shutdown
 ip address 172.2.2.2 255.255.255.0
#
interface GigabitEthernet1/0/1
 undo shutdown
 ip address 172.21.51.2 255.255.0.0
 service-manage http permit
 service-manage https permit
 service-manage ping permit
 service-manage ssh permit
 service-manage snmp permit
 service-manage telnet permit
 service-manage netconf permit
#
interface GigabitEthernet1/0/3
 undo shutdown
 eth-trunk 1
#
interface GigabitEthernet1/0/4
 undo shutdown
 eth-trunk 1
#
interface Virtual-if0
#
firewall zone trust
 set priority 85
 add interface GigabitEthernet0/0/0
 add interface GigabitEthernet1/0/1
#
firewall zone untrust
 set priority 5
 add interface Virtual-if0
#
firewall zone dmz
 set priority 50
 add interface GigabitEthernet1/0/2
#
ip route-static 192.168.4.0 255.255.255.0 172.21.51.7
```


3.7 Pre-configuring the Agile Data Center VXLAN

3.7.1 Objectives

- Understand how to pre-configure VXLAN on the SDN network.
- Understand how to configure the VXLAN all-active gateway.

3.7.2 Networking and Service Description

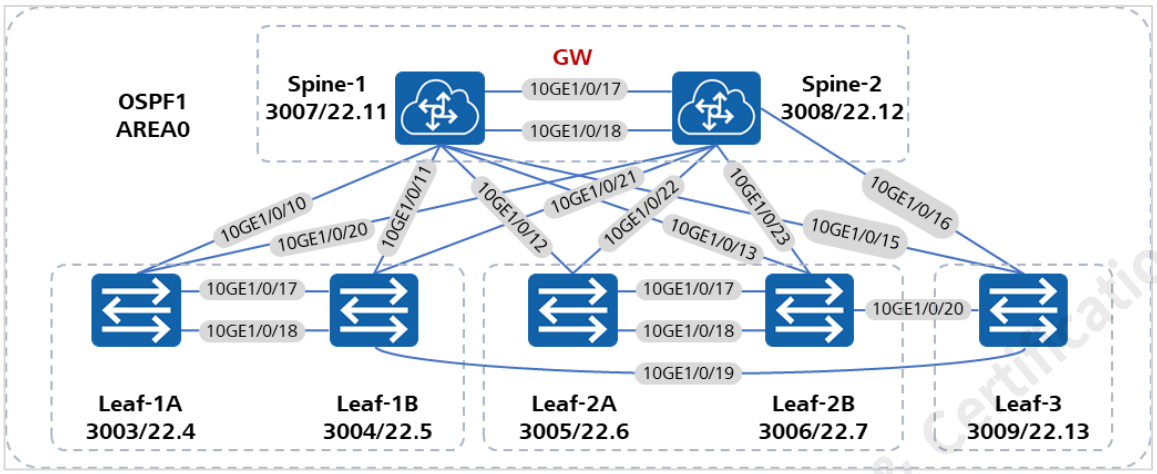

Figure 3-7 Network diagram of the VXLAN BGP EVPN

3.7.3 Configuring an Active-Active Loopback Interface

Spine-1 and Spine-2 have the same Loopback0 interface, use the same IP address as the logical address, and advertise the logical address to BGP to ensure that the Leaf is reachable.

The configurations of Spine-1 and Spine-2 are as follows. The configurations of Spine-1 and Spine-2 have been completed in OSPF.

```
[Spine-1] interface loopback 0
[Spine-1-Loopback1]ip address 10.45.45.45 32
[Spine-1-Loopback1]quit
[Spine-1]OSPF 1
[Spine-1-ospf-1]area 0
[Spine-1-ospf-1-area-0.0.0.0]network 10.45.45.45 0.0.0.0
```

Similarly, Leaf-2A and Leaf-2B have the same loopback interface whose IP address 10.2.2.2/32 is used as the VTEP IP and be advertised to the IGP.

3.7.4 Configuring the NVO3 Extension Function of the Gateway

Only when the CE12800 functions as a gateway, the NVO3 extension function is disabled by default. After the NVO3 service is deployed on the device, there is a high possibility that other services (such as MQC, simplified ACL, traffic monitoring, Bridge-domain traffic statistics, and DHCP) fail to be added as required.

You can use the following methods on the NVO3 device to reduce the failures for adding other services:

```
<Spine-1>system-view immediately
Enter system view, return user view with return command.
[Spine-1] assign forward nvo3 service extend enable
[Spine-1] assign forward nvo3 acl extend enable
```

Configuring the Enhanced Mode of the NVO3 Gateway

Only when the CE12800 functions as a VXLAN Layer 3 gateway, you need to configure the enhanced mode of the Layer 3 gateway. In this mode, the gateway forwards VXLAN packets that need to be encapsulated or decapsulated without consuming the forwarding capability of the board.

```
[Spine-1] assign forward nvo3-gateway enhanced l3
```

3.7.5 Performing Basic Configurations of VXLAN

Configuring Leaf-1:

```
[Leaf-1] interface nve 1
[Leaf-1-Nve1] source 10.1.1.1
[Leaf-1-Nve1] quit
```

Configuring Leaf-2A and Leaf-2B:

```
[Leaf-2A] interface nve 1
[Leaf-2A-Nve1] source 10.2.2.2
[leaf-2A-Nve1] mac-address 0000-5e00-0101
[Leaf-2A-Nve1] quit
```

Configuring Leaf-3:

```
[Leaf-3] interface nve 1
[Leaf-3-Nve1] source 10.3.3.3
[Leaf-3-Nve1] quit
```

Configuring Spine-1:

```
[Spine-1] interface nve 1
[Spine-1-Nve1] source 10.45.45.45
[Spine-1-Nve1] mac-address 0000-5e00-0102
[Spine-1-Nve1] quit
```

Configuring Spine-2:

```
[Spine-2] interface nve 1
[Spine-2-Nve1] source 10.45.45.45
[Spine-2-Nve1] mac-address 0000-5e00-0102
[Spine-2-Nve1] quit
```

3.7.6 Configuring BGP EVPN

When the EVPN on all devices functions as the VXLAN control plane, the configurations of all switches are as follows:

```
[Spine-1]evpn-overlay enable
```

Configure the Spine to establish the peer relationship with the Leaf node and BGP EVPN respectively.

Configuring Spine-1:

```
[Spine-1]bgp 100
[Spine-1-bgp]router-id 11.1.1.1
[Spine-1-bgp]peer 172.16.1.50 as-number 100
[Spine-1-bgp]peer 172.16.1.14 as-number 100
[Spine-1-bgp]peer 172.16.1.18 as-number 100
[Spine-1-bgp]peer 172.16.1.22 as-number 100
[Spine-1-bgp]peer 172.16.1.58 as-number 100
[Spine-1-bgp]peer 172.16.1.58 connect-interface LoopBack 0
[Spine-1-bgp]l2vpn-family evpn // Enable and enter the BGP-EVPN address
family view.
[Spine-1-bgp-af-evpn]undo policy vpn-target // Disable the VPN target
filtering function for VPN routes, that is, accepting all the VPN routes.
[Spine-1-bgp-af-evpn]peer 172.16.1.50 enable
[Spine-1-bgp-af-evpn]peer 172.16.1.50 advertise arp// Advertise the ARP
routes to the BGP EVPN peer.
[Spine-1-bgp-af-evpn]peer 172.16.1.26 reflect-client // Use the spine nodes
as RRs and configure Leaf devices as its clients.
[Spine-1-bgp-af-evpn]peer 172.16.1.14 enable
[Spine-1-bgp-af-evpn]peer 172.16.1.14 advertise arp
[Spine-1-bgp-af-evpn]peer 172.16.1.14 reflect-client
[Spine-1-bgp-af-evpn]peer 172.16.1.18 enable
[Spine-1-bgp-af-evpn]peer 172.16.1.18 advertise arp
[Spine-1-bgp-af-evpn]peer 172.16.1.18 reflect-client
[Spine-1-bgp-af-evpn]peer 172.16.1.22 enable
[Spine-1-bgp-af-evpn]peer 172.16.1.22 advertise arp
[Spine-1-bgp-af-evpn]peer 172.16.1.22 reflect-client
[Spine-1-bgp-af-evpn]peer 172.16.1.58 enable
[Spine-1-bgp-af-evpn]peer 172.16.1.58 advertise arp
[Spine-1-bgp-af-evpn]peer 172.16.1.58 reflect-client
```

Configuring Spine-2

```
[Spine-2]evpn-overlay enable
[Spine-2]bgp 100
[Spine-2-bgp]router-id 11.2.2.2
[Spine-2-bgp]peer 172.16.1.45 as-number 100
[Spine-2-bgp]peer 172.16.1.54 as-number 100
[Spine-2-bgp]peer 172.16.1.26 as-number 100
[Spine-2-bgp]peer 172.16.1.30 as-number 100
[Spine-2-bgp]peer 172.16.1.34 as-number 100
```

```
[Spine-2-bgp]peer 172.16.1.38 as-number 100
[Spine-2-bgp]l2vpn-family evpn // Enable and enter the BGP-EVPN address
family view.
[Spine-2-bgp-af-evpn]undo policy vpn-target // Disable the VPN target
filtering function for VPN routes, that is, accepting all the VPN routes.
[Spine-2-bgp-af-evpn]peer 172.16.1.45 enable
[Spine-2-bgp-af-evpn]peer 172.16.1.45 advertise arp // Advertise the ARP
routes to the BGP EVPN peer.
[Spine-2-bgp-af-evpn]peer 172.16.1.45 reflect-client // Use the spine nodes
as RRs and configure Leaf devices as its clients.
[Spine-2-bgp-af-evpn]peer 172.16.1.54 enable
[Spine-2-bgp-af-evpn]peer 172.16.1.54 advertise arp
[Spine-2-bgp-af-evpn]peer 172.16.1.54 reflect-client
[Spine-2-bgp-af-evpn]peer 172.16.1.26 enable
[Spine-2-bgp-af-evpn]peer 172.16.1.26 advertise arp
[Spine-2-bgp-af-evpn]peer 172.16.1.26 reflect-client
[Spine-2-bgp-af-evpn]peer 172.16.1.30 enable
[Spine-2-bgp-af-evpn]peer 172.16.1.30 advertise arp
[Spine-2-bgp-af-evpn]peer 172.16.1.30 reflect-client
[Spine-2-bgp-af-evpn]peer 172.16.1.34 enable
[Spine-2-bgp-af-evpn]peer 172.16.1.34 advertise arp
[Spine-2-bgp-af-evpn]peer 172.16.1.34 reflect-client
[Spine-2-bgp-af-evpn]peer 172.16.1.38 enable
[Spine-2-bgp-af-evpn]peer 172.16.1.38 advertise arp
[Spine-2-bgp-af-evpn]peer 172.16.1.38 reflect-client
```

Configuring Leaf-1

```
[Leaf-1]evpn-overlay enable
[Leaf-1]bgp 100
[Leaf-1-bgp]router-id 11.3.3.3
[Leaf-1-bgp]peer 172.16.1.13 as-number 100// Establish the BGP EVPN peer
relationship with two Spine switches.
[Leaf-1-bgp]peer 172.16.1.25 as-number 100
[Leaf-1-bgp]peer 172.16.1.49 as-number 100
[Leaf-1-bgp]peer 172.16.1.53 as-number 100
[Leaf-1-bgp]l2vpn-family evpn // Enable and enter the BGP-EVPN address family
view.
[Leaf-1-bgp-af-evpn]peer 172.16.1.13 enable
[Leaf-1-bgp-af-evpn]peer 172.16.1.13 advertise arp // Advertise the ARP
routes to the BGP EVPN peer.
[Leaf-1-bgp-af-evpn]peer 172.16.1.25 enable
[Leaf-1-bgp-af-evpn]peer 172.16.1.25 advertise arp
[Leaf-1-bgp-af-evpn]peer 172.16.1.49 enable
[Leaf-1-bgp-af-evpn]peer 172.16.1.49 advertise arp
[Leaf-1-bgp-af-evpn]peer 172.16.1.53 enable
[Leaf-1-bgp-af-evpn]peer 172.16.1.53 advertise arp
```

Configuring Leaf-2A

```
[Leaf-2A]evpn-overlay enable
[Leaf-2A]bgp 100
[Leaf-2A-bgp]router-id 11.4.4.4
```

```
[Leaf-2A-bgp]peer 172.16.1.17 as-number 100 // Establish the BGP EVPN peer
relationship with two Spine switches.
[Leaf-2A-bgp]peer 172.16.1.29 as-number 100
[Leaf-2A-bgp]l2vpn-family evpn // Enable and enter the BGP-EVPN address
family view.
[Leaf-2A-bgp-af-evpn]peer 172.16.1.17 enable
[Leaf-2A-bgp-af-evpn]peer 172.16.1.17 advertise arp // Advertise the ARP
routes to the BGP EVPN peer.
[Leaf-2A-bgp-af-evpn]peer 172.16.1.29 enable
[Leaf-2A-bgp-af-evpn]peer 172.16.1.29 advertise arp
```

Configuring Leaf-2B

```
[Leaf-2B]evpn-overlay enable
[Leaf-2B]bgp 100
[Leaf-2B-bgp]router-id 11.5.5.5
[Leaf-2B-bgp]peer 10.1.1.21 as-number 100 // Establish the BGP EVPN peer
relationship with two Spine switches.
[Leaf-2B-bgp]peer 10.1.1.33 as-number 100
[Leaf-2B-bgp]l2vpn-family evpn // Enable and enter the BGP-EVPN address
family view.
[Leaf-2B-bgp-af-evpn]peer 10.1.1.21 enable
[Leaf-2B-bgp-af-evpn]peer 10.1.1.21 advertise arp // Advertise the ARP routes
to the BGP EVPN peer.
[Leaf-2B-bgp-af-evpn]peer 10.1.1.33 enable
[Leaf-2B-bgp-af-evpn]peer 10.1.1.33 advertise arp
```

Leaf-3

```
[Leaf-3]evpn-overlay enable
[Leaf-3]bgp 100
[Leaf-3-bgp]router-id 11.6.6.6
[Leaf-3-bgp]peer 172.16.1.37 as-number 100
[Leaf-3-bgp]peer 172.16.1.57 as-number 100
[Leaf-3-bgp]l2vpn-family evpn // Enable and enter the BGP-EVPN address family
view.
[Leaf-3-bgp-af-evpn]peer 172.16.1.37 enable
[Leaf-3-bgp-af-evpn]peer 172.16.1.37 advertise arp
[Leaf-3-bgp-af-evpn]peer 172.16.1.57 enable
[Leaf-3-bgp-af-evpn]peer 172.16.1.57 advertise arp
```

Configuring the ARP resource allocation mode (only for the CE12800)

```
[Spine-2]arp resource-mode extend
Info: Please save the configuration and reboot the system to enable the
operation.
```

Set the ARP resource allocation mode to extend. In specific scenarios, you need to increase the number of ARP entries on the device.

3.7.7 Verifying the Active-Active Result (See section 3.5 for the configurations of DFS)

Run the **display dfs-group 1 active-active-gateway** command to check information about the DFS Group all-active gateway. The Spine-1 command output is used as an example.

```
[Spine-1] display dfs-group 1 active-active-gateway
A:Active    I:Inactive
------------------------------------------------------------------------
Peer           System name        State        Duration
10.1.1.46       Spine-2             A           0:0:53
```

Check the BGP EVPN peer relationship:

```
[Spine-1]dis bgp evpn peer
 BGP local router ID       : 172.21.51.3
 Local AS number          : 100
 Total number of peers     : 5
 Peers in established state : 5
  Peer          V         AS  MsgRcvd  MsgSent  OutQ  Up/Down      State
PrefRcv
  172.16.1.14      4        100      21       21     0 00:14:35 Established
0
  172.16.1.18      4        100      11       12     0 00:06:37 Established
0
  172.16.1.22      4        100      10       10     0 00:05:22 Established
0
  172.16.1.50      4        100      20       20     0 00:14:17 Established
0
  172.16.1.58      4        100       5        4     0 00:00:53 Established
0
```

3.8 Configuring SNMP

3.8.1 Objectives

- Learn the basic configuration of the SNMP protocol.

  To enable the Agile Controller-DCN to discover and add devices through the SNMP protocol, you need to configure SNMP parameters on the device. The configured SNMP parameters must be the same as those set in the interconnection operation on the Agile Controller-DCN.

  The configuration of SNMP parameters on the CE switch is different from that on the firewall.

NOTE

The Agile Controller-DCN can only establish the SNMP connections with the device through the more secure SNMPv3 protocol.

The Agile Controller-DCN supports the following MIB trees: nt iso, rd iso, wt iso, and iso-view iso.

Before configuring SNMP parameters, ensure that the info-center is enabled. Otherwise, the device cannot report traps to the Agile Controller-DCN.

Run the **display info-center** command to check whether the info-center is enabled.

If **Information Center: enable** is displayed in the command output, the info-center has been enabled and SNMP parameters can be configured.

If **Information Center: disable** is displayed in the command output, the info-center is disabled. Run the **info-center enable** command to enable the info-center.

3.8.2 Configuring SNMP Parameters for the Switch

Configure the read/write community name of SNMPv3 for the network device, enable the **Trap** notification function and define the source interface as the management interface. The address of the management interface is as follows:

The SNMPv3 parameters are as follows:

Table 3-4 Planning of SNMP parameters

Parameter	Value	Description
snmp-agent udp-port	161	The default UDP port number used for the interconnection between the SNMP Agent (CE switch) and Agile Controller is **161**.
snmp-agent group	dc-admin	Name of an SNMPv3 user group
snmp-agent usm-user	admin	SNMPv3 user
snmp-agent usm-user authentication-mode	SHA	User authentication mode
authentication password	Huawei@123	User authentication password
privacy-mode	AES128	Encryption mode for authentication
privacy password	Huawei@123	Encryption password

The address of the device management interface is as follows:

Table 3-5 IP address planning for the device management interface

Device Name	Interface	Management IP address
SDN-FW1	GE1/0/1	172.21.22.2
SDN-FW2	GE1/0/1	172.21.22.3
Spine-1	MEth0/0/0	172.21.22.11
Spine-2	MEth0/0/0	172.21.22.12
Leaf-1	MEth0/0/0	172.21.22.4
Leaf-2A	MEth0/0/0	172.21.22.6
Leaf-2B	MEth0/0/0	172.21.22.7
Leaf-3	Vlanif4060	172.21.22.13

The configuration of Spine-1, Spine-2 and Leaf-4 is the same. The Spine-1 is used as an example.

Run the **system-view** command to go to the system view.

Change the number of the port used for connecting the SNMP Agent to Agile Controller-DCN. By default, the SNMP Agent uses port 161 to interconnect with Agile Controller-DCN.

```
[Spine-1] snmp-agent udp-port 161
```

Configure an SNMPv3 user group and user, and set the authentication and encryption modes. In this example, the user group is **dc-admin**, the username is **admin**, the authentication mode is SHA, and the encryption mode is AES128.

```
[Spine-1] snmp-agent usm-user v3 admin group dc-admin
[Spine-1] snmp-agent usm-user v3 admin authentication-mode sha
Please configure the authentication password (8-255)
Enter Password: // Enter the authentication password. In this example, the
authentication password is Huawei@123.
Confirm Password: // Confirm the authentication password. In this example,
the authentication password is Huawei@123.
[Spine-1] snmp-agent usm-user v3 admin privacy-mode aes128
Please configure the privacy password (8-255)
Enter Password: // Enter the encryption password. In this example, the
encryption password is Huawei@123.
Confirm Password: // Confirm the encryption password. In this example, the
encryption password is Huawei@123.
```

Configure the device to send traps to Agile Controller-DCN using SNMPv3.

```
[Spine-1] snmp-agent trap enable feature-name trunk
[Spine-1] snmp-agent trap enable // Enable the switch to send trap packets.
[Spine-1] snmp-agent trap source meth0/0/0 // The port indicates the name of
the interface where the IP address of the device connected to the Agile
Controller-DCN is located. The interface must have been configured with an IP
address.
```

Set a MIB view and add it to the attribute list of the user group, so that the user group has the read, write, and trap reporting functions.

NOTE

The Agile Controller-DCN uses the specified MIB view defined in SNMP to obtain LLDP link information from the device. The MIB view defined in SNMP is iso-view, and the OID MIB sub-tree of the MIB objects is iso.

```
[Spine-1] snmp-agent mib-view included iso-view iso
[Spine-1] snmp-agent mib-view included nt iso
[Spine-1] snmp-agent mib-view included rd iso
[Spine-1] snmp-agent mib-view included wt iso
[Spine-1] snmp-agent group v3 dc-admin privacy read-view rd write-view wt
notify-view nt
```

3.8.3 Configuring the SNMP Parameters for the Firewall

The firewall configuration is different from that of the switch. The SDN-FW1 configuration is as follows:

Change the number of the port used for connecting the SNMP Agent to Agile Controller-DCN. By default, the SNMP Agent uses port 161 to interconnect with Agile Controller-DCN.

```
HRP_M[SDN-FW1] snmp-agent udp-port 161
```

Configure an SNMP user group.

```
HRP_M[SDN-FW1] snmp-agent group v3 dc-admin privacy read-view rd write-view
wt notify-view nt // The rd, wt, and nt are MIB views, the names of which
must be the same as those of the pre-configured views on the device.
```

Configure an SNMPv3 user group and user, and set the authentication and encryption modes. In this example, the user group is **dc-admin**, the username is **admin**, the authentication mode is SHA, and the encryption mode is AES128.

```
HRP_M[SDN-FW1] snmp-agent usm-user v3 admin group dc-admin
HRP_M[SDN-FW1]] snmp-agent usm-user v3 admin authentication-mode sha
Please configure the authentication password (8-255)
Enter Password: // Enter the authentication password. In this example, the
authentication password is Huawei@123.
Confirm Password: // Confirm the authentication password. In this example,
the authentication password is Huawei@123.
```

```
HRP_M[FW1] snmp-agent usm-user v3 admin privacy-mode aes128
Please configure the privacy password (8-255)
Enter Password: // Enter the encryption password. In this example, the
encryption password is Huawei@123.
Confirm Password: // Confirm the encryption password. In this example, the
encryption password is Huawei@123.
```

Configure a set of SNMP parameters whose user name is **ACTrap_huawei**, authentication mode and password are **SHA/Huawei@123**, and encryption mode and encryption password are **AES128/Huawei@123**. The AC can obtain the system startup time of the firewall through SNMP.

```
HRP_M[SDN-FW1] snmp-agent usm-user v3 ACTrap_huawei
HRP_M[SDN-FW1] snmp-agent usm-user v3 ACTrap_huawei group ACTRAP
HRP_M[SDN-FW1] snmp-agent usm-user v3 ACTrap_huawei authentication-mode sha
Please configure the authentication password (8-64)
Enter Password:
Confirm Password:
// Enter the authentication password twice. In this example, the
authentication password is Huawei@123.
HRP_M[SDN-FW1] snmp-agent usm-user v3 ACTrap_huawei privacy-mode aes128
Please configure the authentication password (8-64)
Enter Password:
Confirm Password:
// Enter the authentication password twice. In this example, the
authentication password is Huawei@123.
```

Configure the device to send traps to Agile Controller-DCN using SNMPv3.

```
HRP_M[SDN-FW1] snmp-agent trap enable feature-name trunk
HRP_M[SDN-FW1] snmp-agent trap enable // Enable the switch to send trap
packets.
HRP_M[SDN-FW1] snmp-agent trap source GigabitEthernet1/0/1 // Specify the
name of the interface where the IP address of the device connected to the
Agile Controller-DCN is located.
HRP_M[SDN-FW1] snmp-agent target-host trap address udp-domain 192.168.4.4
udp-port 1025 params securityname ACTrap_huawei v3 privacy private-netmanager
// The IP address is the Agile Controller-DCN node address.
```

Set a MIB view and add it to the attribute list of the user group, so that the user group has the read, write, and trap reporting functions.

```
HRP_M[SDN-FW1] snmp-agent mib-view included iso-view iso
HRP_M[SDN-FW1] snmp-agent group v3 ACTRAP privacy read-view rd write-view wt
notify-view nt
HRP_M[SDN-FW1] snmp-agent mib-view included nt iso
HRP_M[SDN-FW1] snmp-agent mib-view included rd iso
HRP_M[SDN-FW1] snmp-agent mib-view included wt iso
```

3.9 Configuring NETCONF

3.9.1 Objectives

- NETCONF parameters must be configured on a switch before the Agile Controller-DCN can deliver service configuration to the network device or obtain configuration information from the switch. The configured NETCONF parameters must be the same as those set in the interconnection operation on the Agile Controller-DCN console.
- NETCONF parameters are configured differently on CE switches and firewalls.

3.9.2 Configuring NETCONF Parameters for a Switch

Configure the VTY user interface of devices to support the SSH protocol.

```
<Spine-1> system-view
[Spine-1] user-interface vty 0 4
[Spine-1-ui-vty0-4] authentication-mode aaa
[Spine-1-ui-vty0-4] protocol inbound ssh
```

NOTE

After SSH is configured as the login protocol, the switch automatically disables Telnet. The Telnet protocol has security risks. You are not advised to run the **protocol inbound all** command which enables SSH and Telnet simultaneously.

Deploy SSH on the device.

Create an SSH user.

Create a local SSH user, and set the user name to **client**, domain name to **huawei.com**, and the password to **Huawei@123**.

```
[Spine-1] aaa
[Spine-1-aaa] local-user client@huawei.com password irreversible-cipher
Huawei@123
[Spine-1-aaa] local-user client@huawei.com service-type ssh
[Spine-1-aaa] local-user client@huawei.com level 3
```

Generate a local RSA key pair.

```
[Spine-1] rsa local-key-pair create
The key name will be: Spine-1_Host
The range of public key size is (512 ~ 2048).
NOTE: If the key modulus is greater than 512,
 It will take a few minutes.
Input the bits in the modulus [default = 512] :
```

Run the **display rsa local-key-pair public** command to check the public key in the local RSA key pair.

NOTE

After you run this command, the generated key pair is saved in the device and will not be lost after the device restarts.

This command is not saved in the configuration file.

Set the authentication mode for the SSH user to **Password**.

```
[Spine-1] ssh user client@huawei.com authentication-type password
```

Configure the service mode for the SSH user.

```
[Spine-1] ssh user client@huawei.com service-type snetconf
```

Enable the NETCONF function. After the SNETCONF service is enabled, the device enables the NETCONF service for the SSH server on the port.

```
[Spine-1] snetconf server enable
```

3.9.3 Configuring the NETCONF Parameters for the Firewall

The firewall configuration is different from that of the switch. The SDN-FW1 configuration is as follows:

Configure the management port to permit NETCONF.

```
HRP_M<SDN-FW1> system-view
HRP_M[SDN-FW1] interface GigabitEthernet1/0/1 // The interface that connects
to the management network can be an Eth-Trunk or a physical interface.
HRP_M[SDN-FW1-GigabitEthernet1/0/1] service-manage enable
HRP_M[SDN-FW1-GigabitEthernet1/0/1] service-manage netconf permit
HRP_M[SDN-FW1-GigabitEthernet1/0/1] quit
```

Configure the administrator, service type, level, and authentication type. Set the account to **netconf-admin**, and password to **Huawei@123**.

```
HRP_M[SDN-FW1] aaa
HRP_M[SDN-FW1-aaa] manager-user netconf-admin
HRP_M[SDN-FW1-aaa-manager-user-client] password
Enter Password:// Enter the password Huawei@123.
Confirm Password: // Confirm the password as Huawei@123.
HRP_M[SDN-FW1-aaa-manager-user-client] service-type api
HRP_M[SDN-FW1-aaa-manager-user-client] level 15
HRP_M[SDN-FW1-aaa-manager-user-client] authentication-scheme admin_local
HRP_M[SDN-FW1-aaa-manager-user-client] quit
```

Configure the NETCONF port number and enable the NETCONF service.

```
HRP_M[SDN-FW1] api
The default port number of the HRP_M[SDN-FW1-api] api netconf port 1025 //
The default port number of the switch is 22, and the minimum port number of
the firewall is 830. Change the port number of the firewall to 1025.
HRP_M[SDN-FW1-api] api netconf enable
HRP_M[SDN-FW1-api] quit
```

4 Agile Controller-DCN Pre-Configuration

4.1 Preparing for Management

- Underlay network pre-configuration
- Agile Controller-DCN installation and configuration

4.2 Networking and Agile Controller-DCN Parameter Planning

During the Agile Controller-DCN configuration, plan parameters such as parameters related to interconnection resources and global resources. When the Agile Controller-DCN connects to FusionSphere OpenStack, ensure that parameter settings are consistent between the Agile Controller-DCN and FusionSphere OpenStack.

Parameter	Data	Description
Interconnection VLAN	3000-3499	Used to connect a firewall to a gateway in bypass mode. A maximum of 128 VLAN IDs can be added. The VLAN ID ranges from 2 to 4094.
Interconnection IP address	10.125.97.240-10.125.97.255/30	Used to create a VLANIF interface when the gateway connects to the firewall. A maximum of 1000 segments can be added. Class A/B/C addresses are supported, excluding 0.0.0.0/8 and 127.0.0.0/8.
Bridge Domain	5000	Bridge domain (Bridge-domain) on the VXLAN network.

Parameter	Data	Description
		The Agile Controller-DCN automatically allocates the Bridge-domains to switches.
VLAN	1000-2999	Used to connect VMs to logical switches. A maximum of 128 segments can be added. In a FusionSphere scenario, the global VLAN segments cannot be overlapped with VLAN segments of physical networks. If the cloud platform delivers VXLANs to the switch, global VLAN resources configured on the Agile Controller-DCN are used. If the cloud platform delivers VLANs to the switch, VLAN ranges of physical networks on the cloud platform are used. The method of calculating the global VLAN range on the Agile Controller-DCN is as follows: Number of required VLANs = Number of compute nodes under a port on a switch x Number of VMs that can be created on a single compute node x Number of VM NICs
VNI	4000-10000 10001-20000	Range of VNIs that tenant services can use, including Layer 2 and Layer 3 VNIs. A maximum of 128 segments can be added. The VNI value ranges from 1 to 16000000.
Public Network IP Address	100.11.11.0/30 100.11.11.4/30	Used for the Floating IP service. When the virtual route in the gateway route provides services externally, a public IP address is required. Cloud-network integration is delivered by the cloud platform.
Interconnection IP address	10.125.91.0-10.125.91.255/30	-
Eth-Trunk ID	40-127	-
M-Lag ID	20-2048	-
Sub-interface	200-4094	Used for VXLAN service access. The sub-interface range is the range of Layer 2 sub-interfaces that can be created on a single physical interface. Do not conflict with the sub-interfaces that have been occupied by the switch.
Loopback Interface	100-400	Used for the DHCP service. The loopback interface range is the range of loopback interfaces that can be created on a single device. Do not conflict with the loopback interfaces that have been occupied by the switch.

4.3 Discovering Devices and Links

To implement management of devices and services on the network on the Agile Controller-DCN, you need to add the devices and servers on the network to the Agile Controller-DCN, that is, to discover the devices and servers to the Agile Controller-DCN.

Huawei switches and firewalls are discovered through automatic discovery, batch import, and device registration (for CE1800V only). Third-party firewalls or load balancers are discovered by creating or importing dumb devices.

Step 1 **Log in to the Agile Controller-DCN.**

Open a browser and enter https://192.168.4.4:18002/ in the address box, and press **Enter**. Log in to the AC-DCN using the account **admin** and the password **Huawei@123**.

Step 2 **Add network devices.**

Choose Network > Device Management > Discover Device.

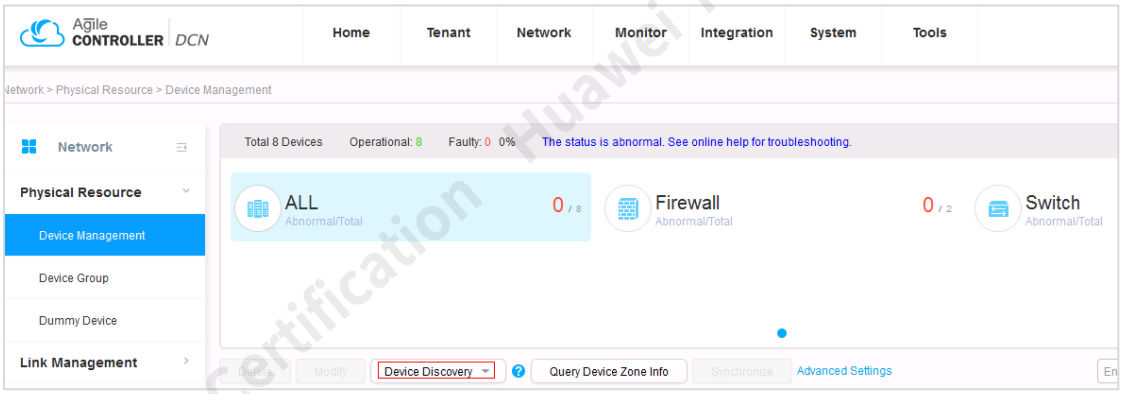

Enter the network segment information of underlay configuration, SNMP, and NETCONF configuration information, click **Start**, and check if the device is successfully discovered.

For details about the configuration, see section 3.8 "**Configuring SNMP**" and section 3.9 "**Configuring NETCONF**".

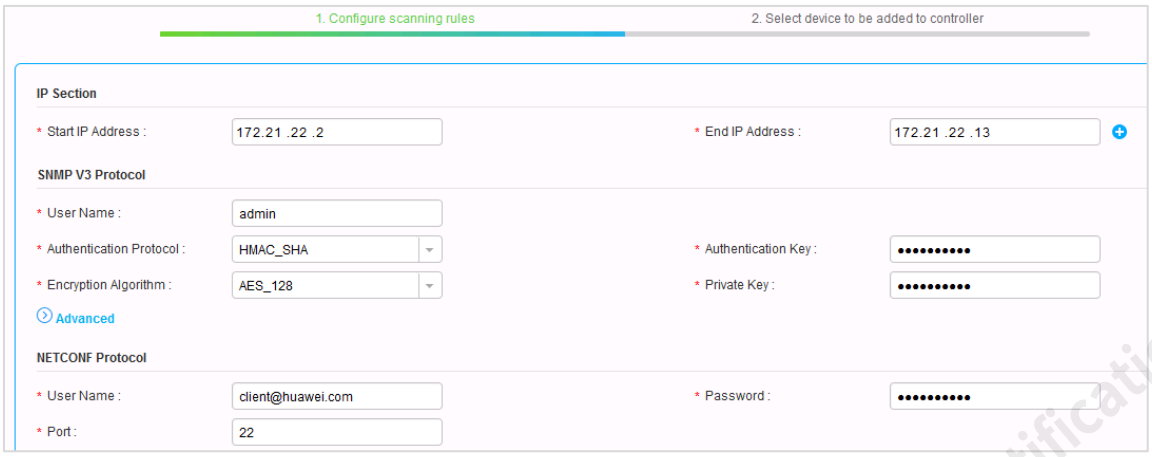

After the device is successfully discovered, click **Add** and check if the firewall is successfully added.

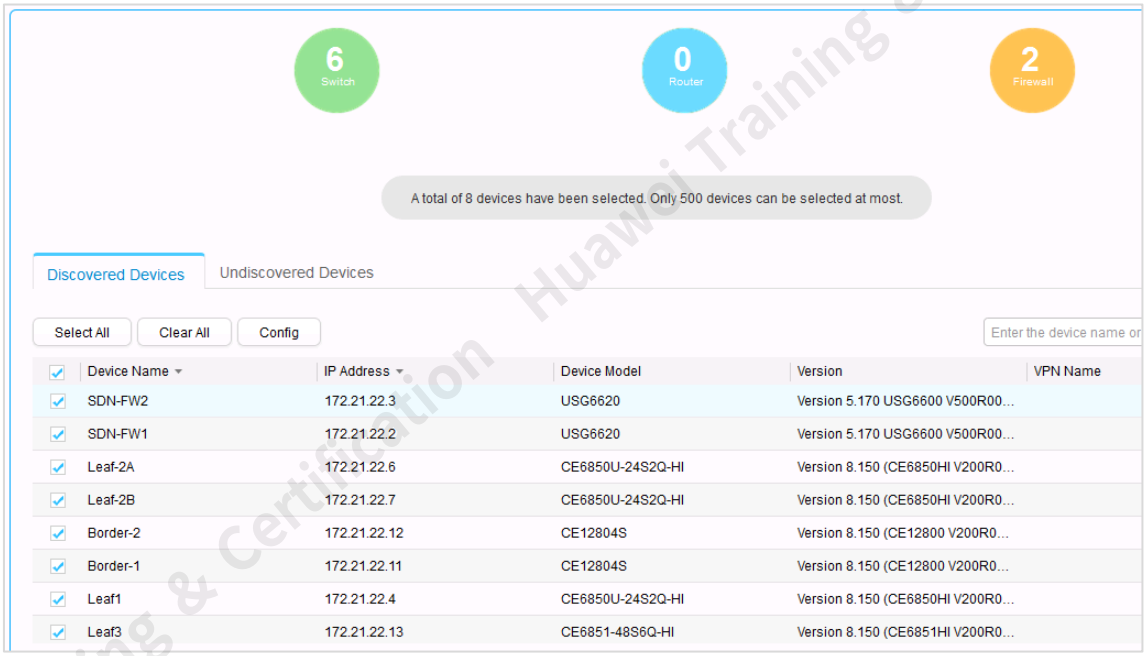

Failed to add the firewall. Change the NETCONF port number to **1025** and SNMP password to **Huawei@123**.

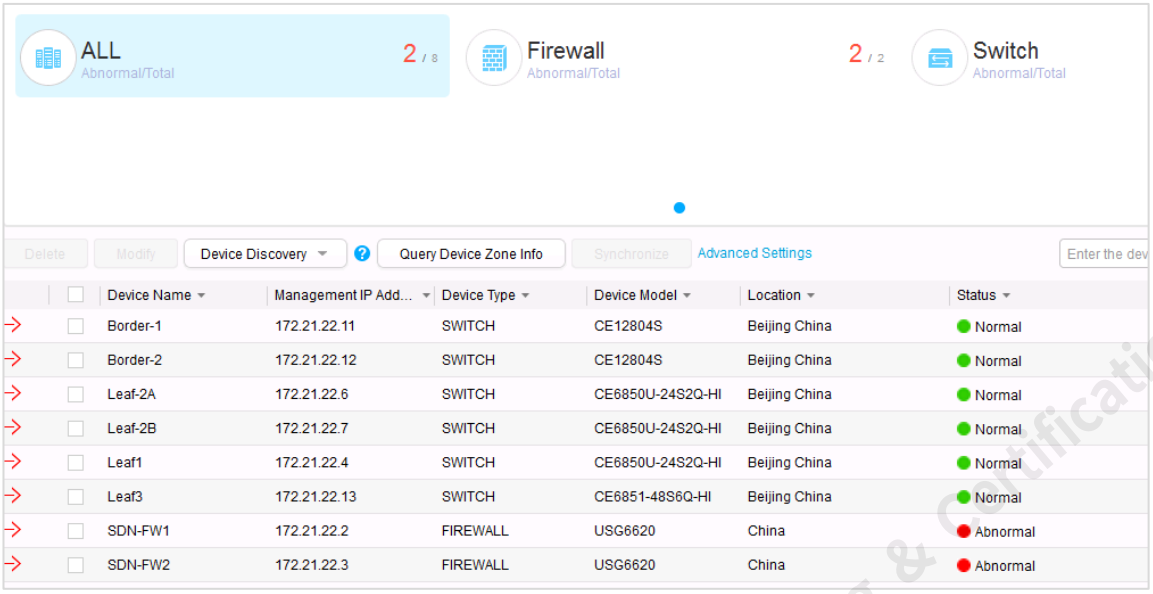

Change the NETCONF account to **netconf-admin**.

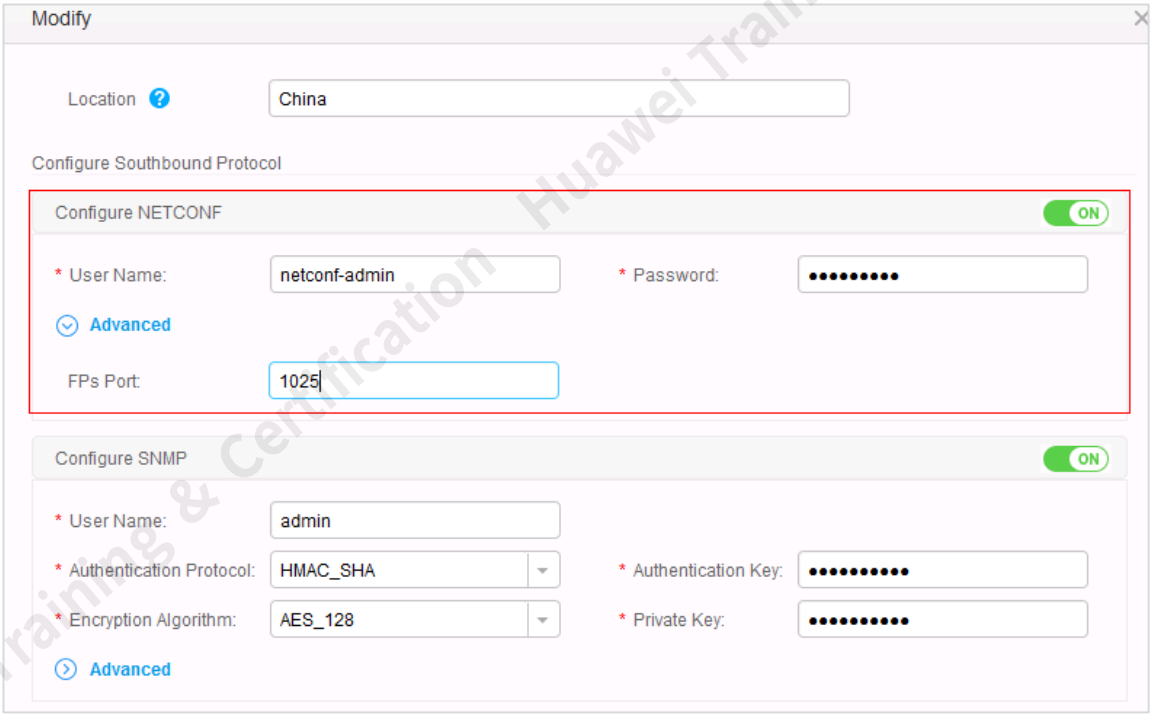

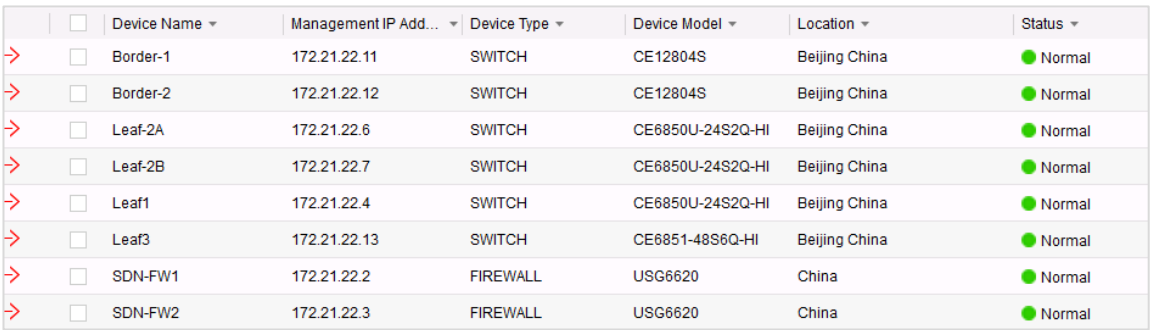

Step 3 **Discover and view inter-devices links.**

Automatically discovers inter-device links.

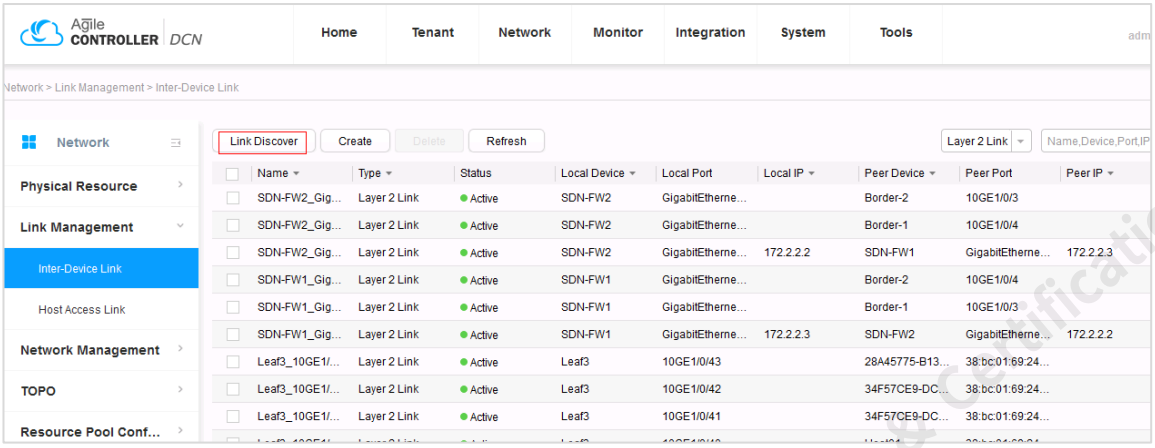

After a link is successfully discovered, the link status is normal.

----End

4.4 Adding Device Groups

Configure the active-active gateway and firewall as device groups and click **Add DeviceGroup**.

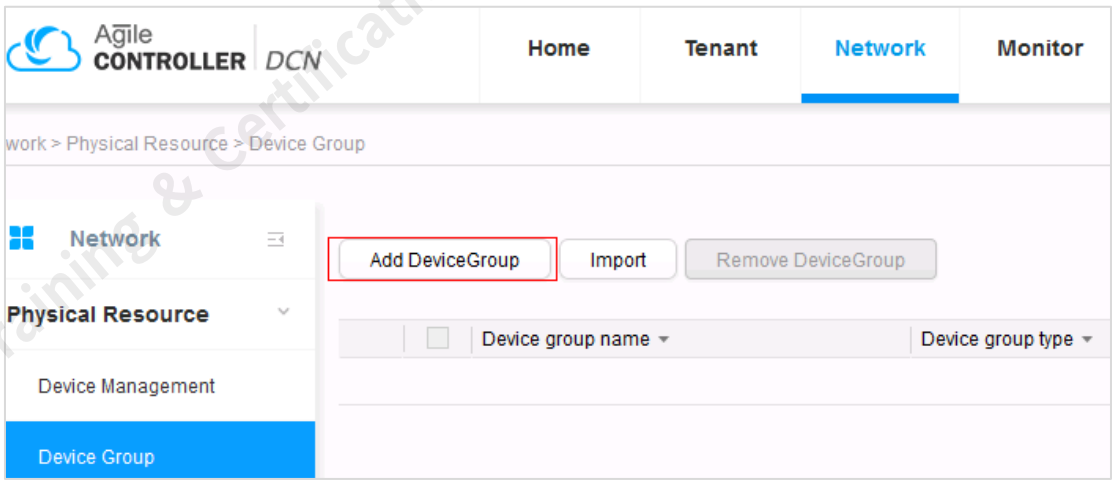

Set **Type** to **Multi-active** and click **Add** to add the GW device.

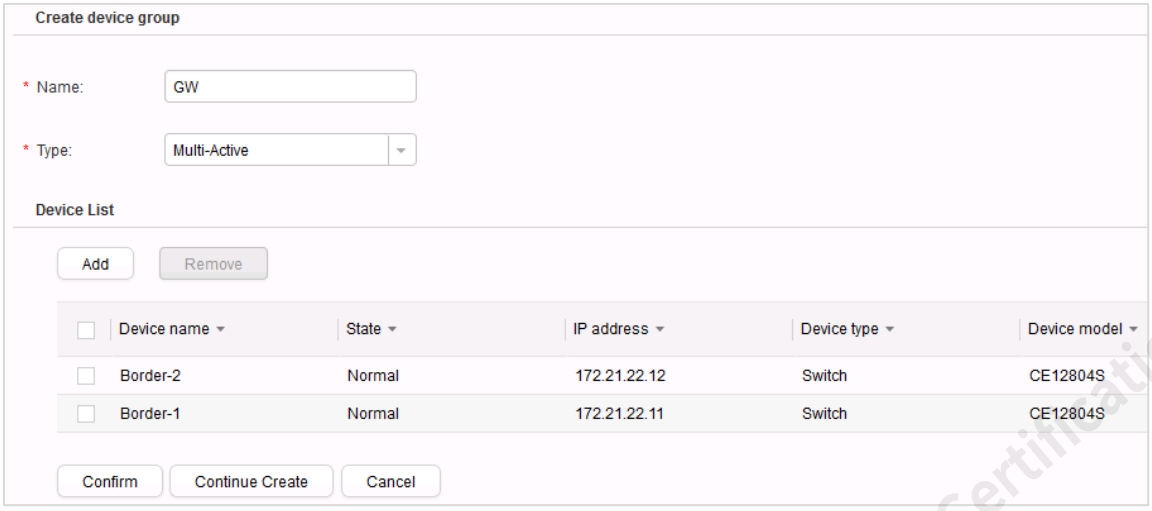

Similarly, create a firewall device group and a Leaf-2 group.

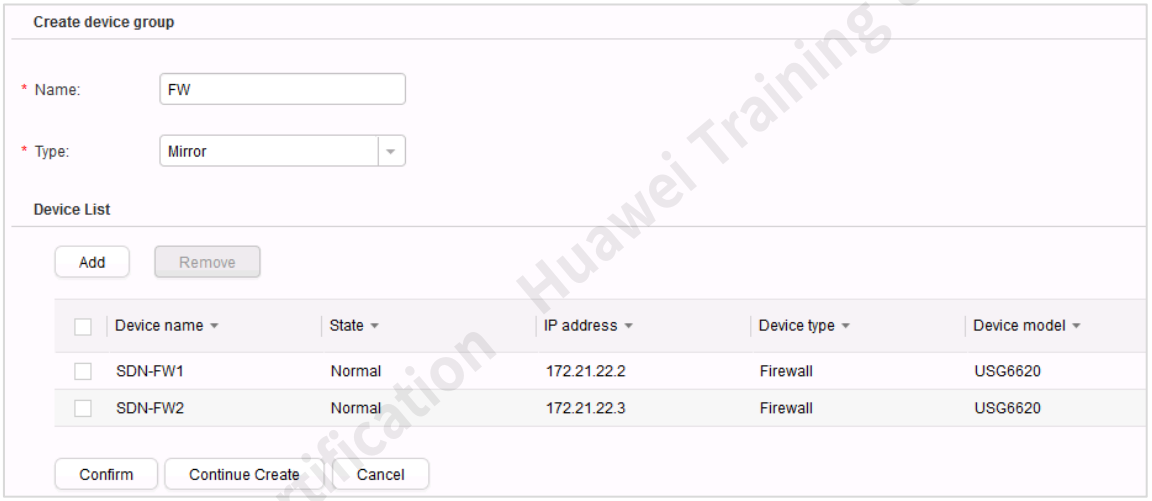

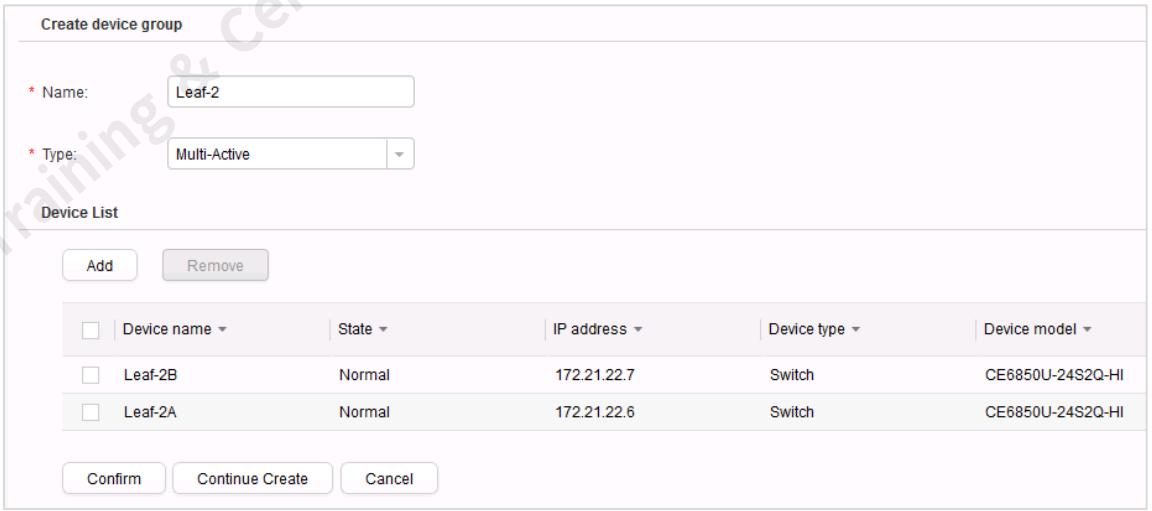

4.5 Creating a Fabric

When devices and links in a network are added to the Agile Controller-DCN, you need to add the devices and links to a fabric network so that tenants can use them flexibly.

Choose Network > Network Management > Fabric to enter the Fabric page.

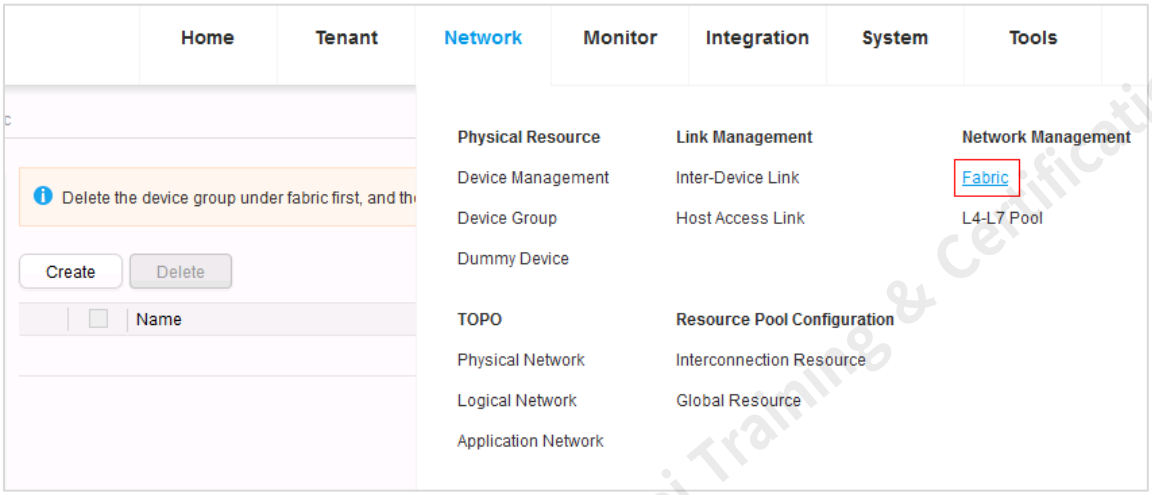

Set the basic information and network features and click **Confirm**.

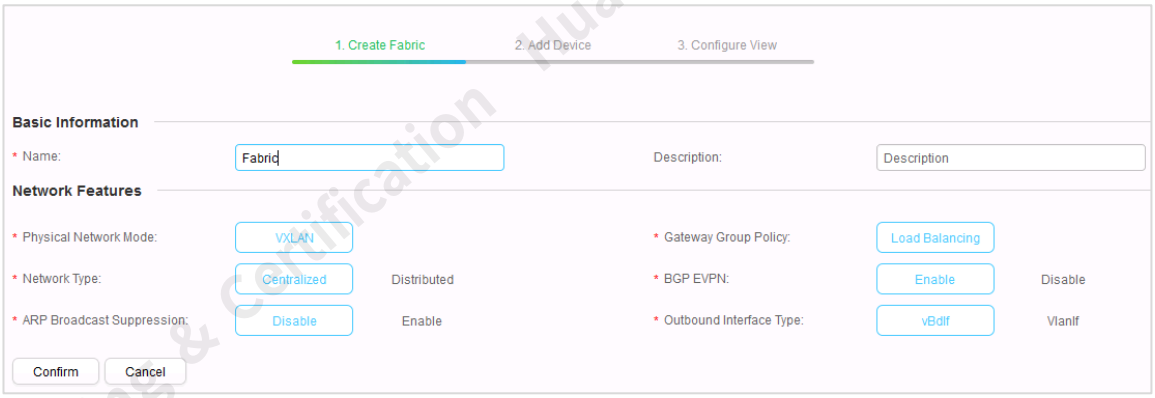

Click **Add Device**, select the switch devices to be added to the fabric, and click **Confirm**.

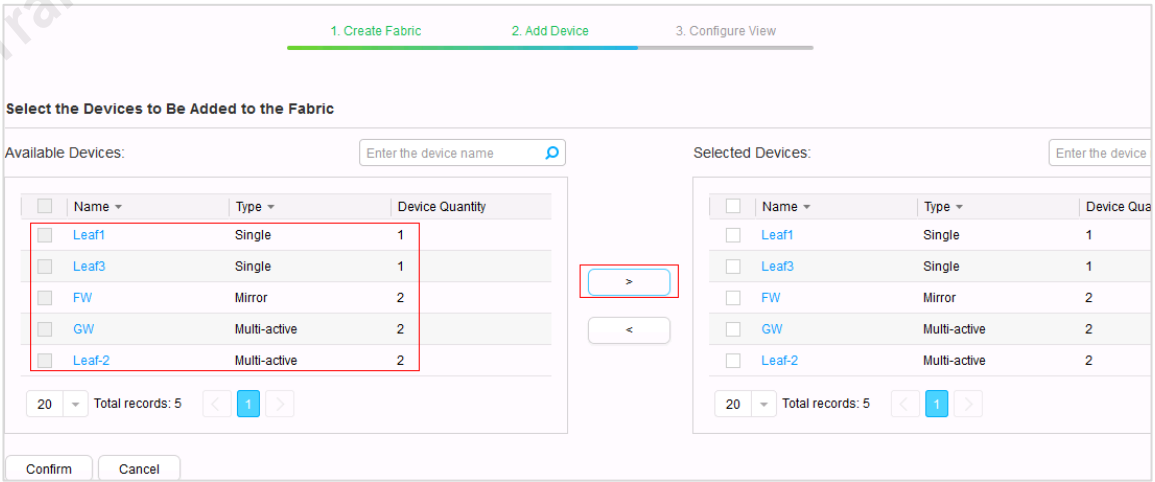

Select a Border Leaf device, click **Configure Egress Gateway**, and then click **OK**.

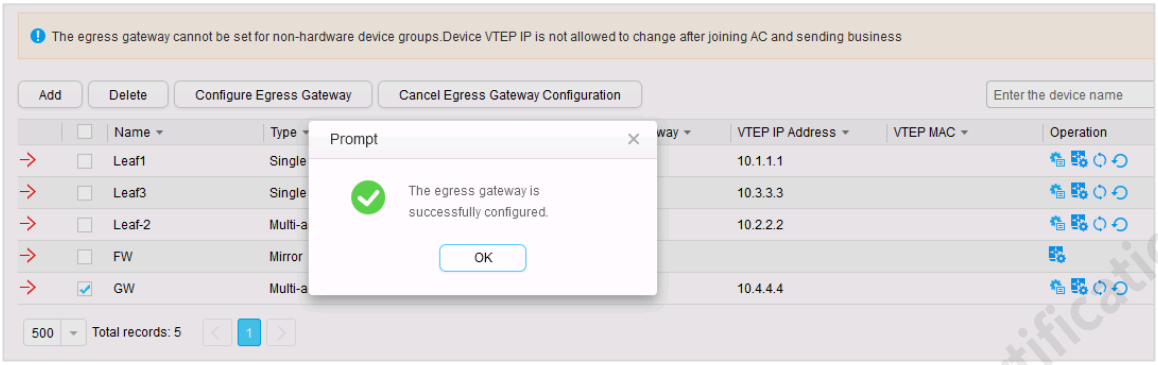

4.6 Creating an L4-L7 Resource Pool

Firewalls with the same network access mode can be added to a resource pool for physical resource virtualization.

Set the role of the link between the firewall and the gateway. Choose **Network** > **TOPO** > **Physical Network** from the main menu. Double-click the link between the firewall and the gateway. The link is zoomed in. In this case, you can view the link role.

On the **Set the link role** page, set the link role according to the following rule table. Set roles for all links between the firewall and the gateway. **Internal link** indicates the traffic from the gateway to the firewall. **External link** indicates the traffic from the firewall to the gateway. **Internal and external link** indicates the bidirectional traffic between the gateway and the firewall.

Set all links of the firewall as the internal and external links.

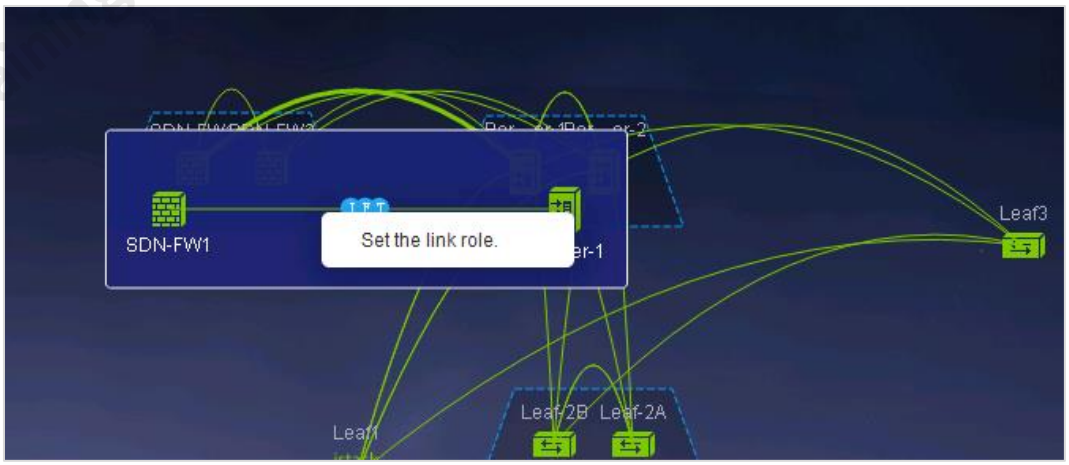

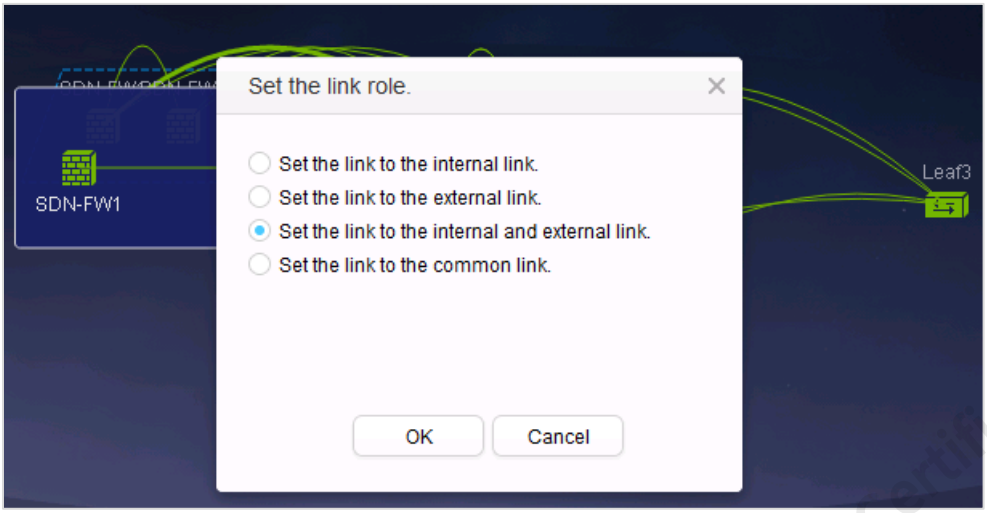

Choose **Network** > **Network Management** > **L4-L7 Pool**, and click **Create**.

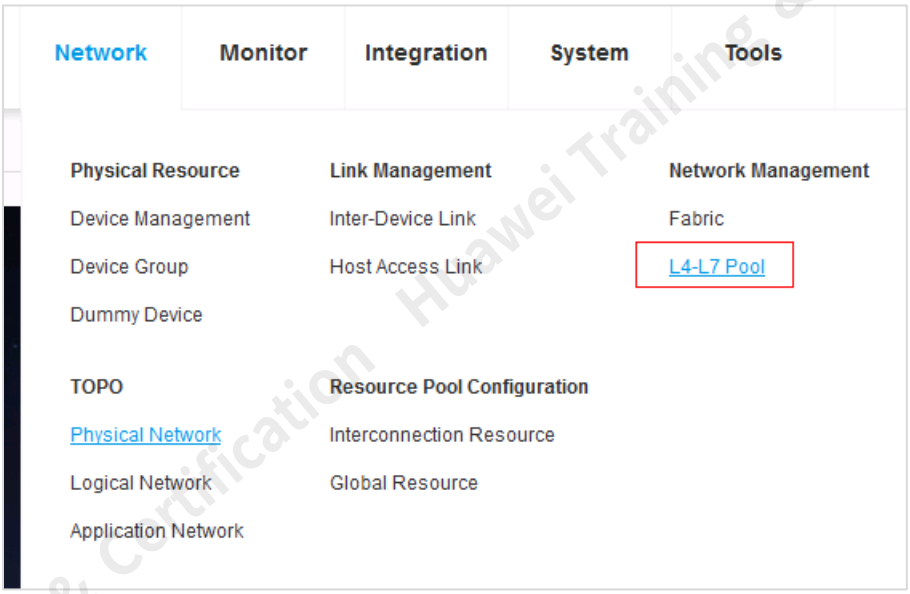

Select the L4-L7 resource pool attribute.

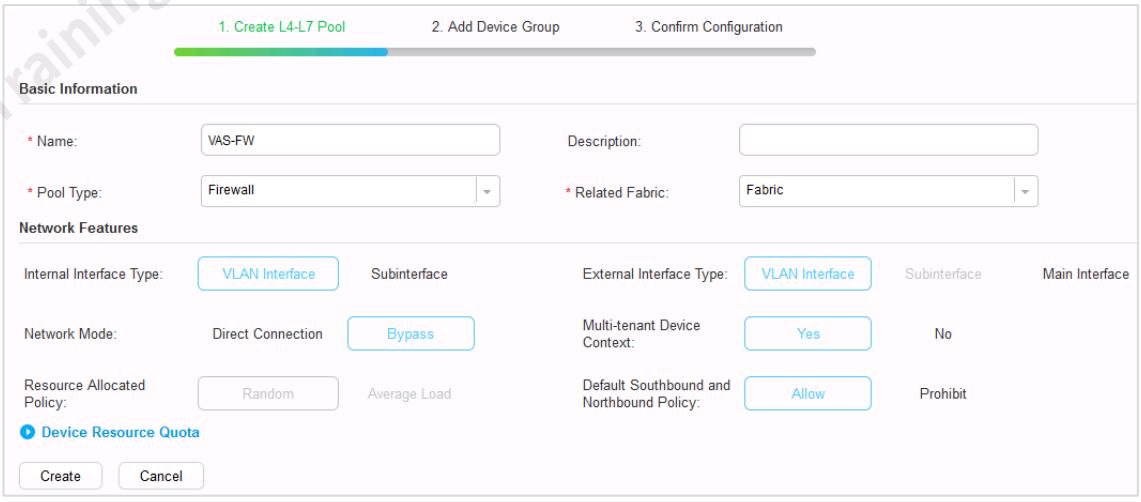

Add a firewall group to the resource pool.

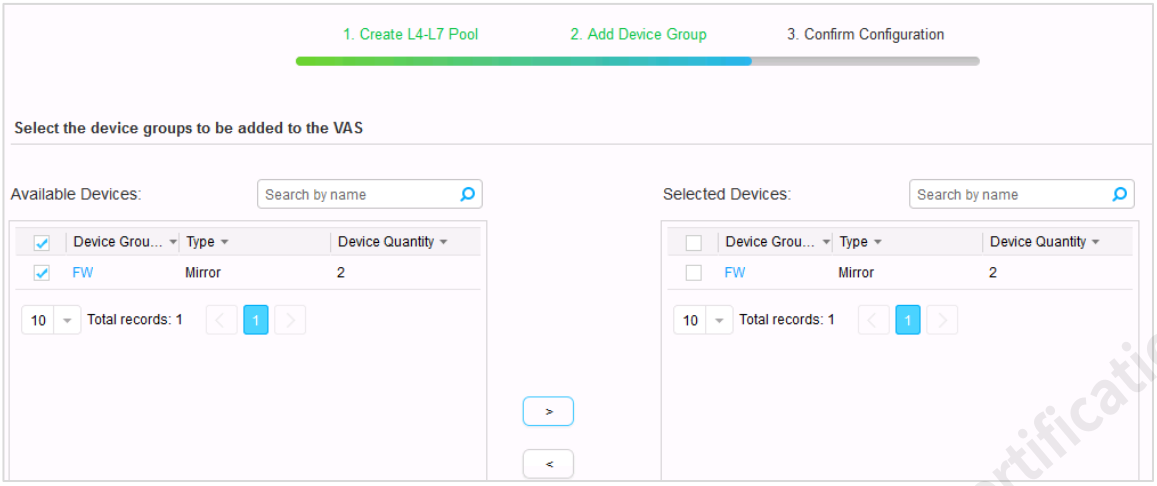

4.7 Configuring Interconnection Resources and Global Resources

VLANs and IP addresses need to be configured for interconnection between gateways and firewalls, implementing communication between VPCs through firewalls.

Choose **Network** > **Resource Pool Configuration** > **Interconnection Resource**.

For details about the parameter settings in this section, see the AC-DCN parameter planning in section 4.2 "**Networking and Agile Controller-DCN Parameter Planning**".

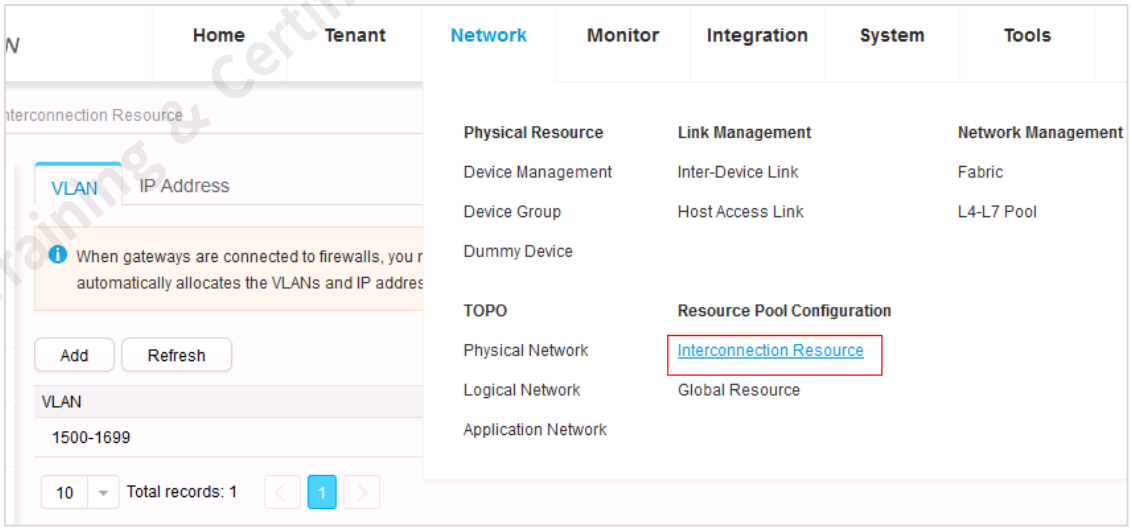

Add the user-defined interconnection VLAN range and IP address range.

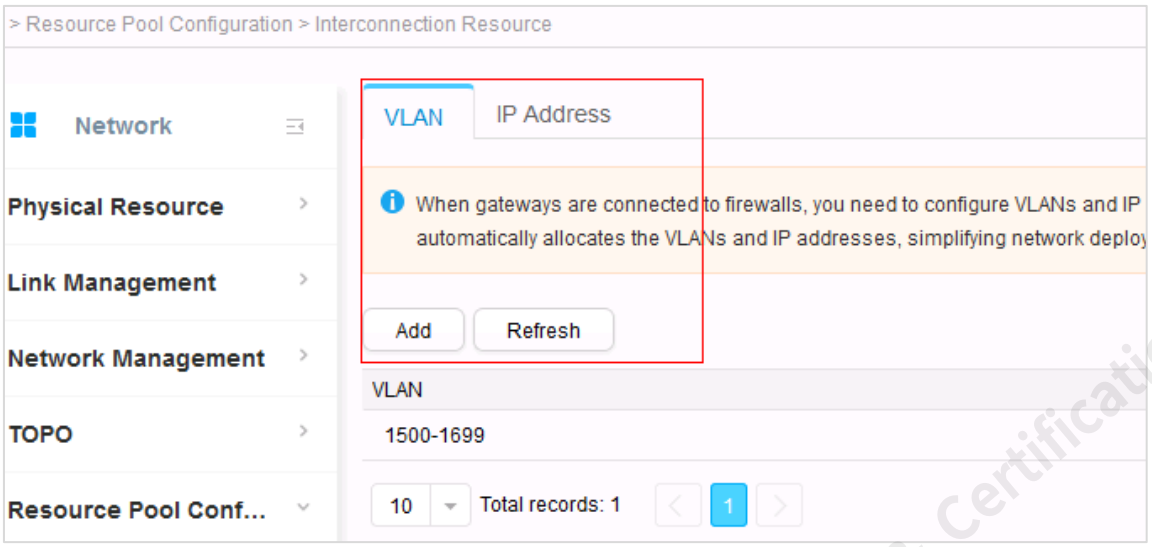

Similarly, when planning a network, you need to pre-configure global resources, such as the Bridge-domains, global VNI, global VLAN, public IP address, and interconnection IP address, so that the Agile Controller-DCN can automatically allocate them to tenants.

Configure tenant resource information under global resources.

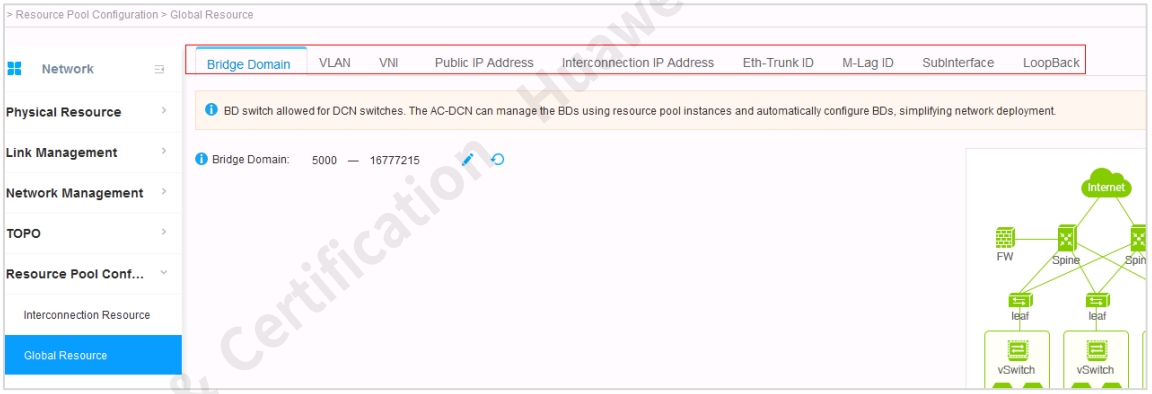

4.8 Collecting Device Alarms

To monitor the status of switches and firewalls in real time, the Agile Controller-DCN can collect device alarms for fault location. Enable SNMP Trap on Agile Controller-DCN.

Choose **System** > **System Settings** > **Southbound Protocol** to enter the page for configuring the southbound protocol.

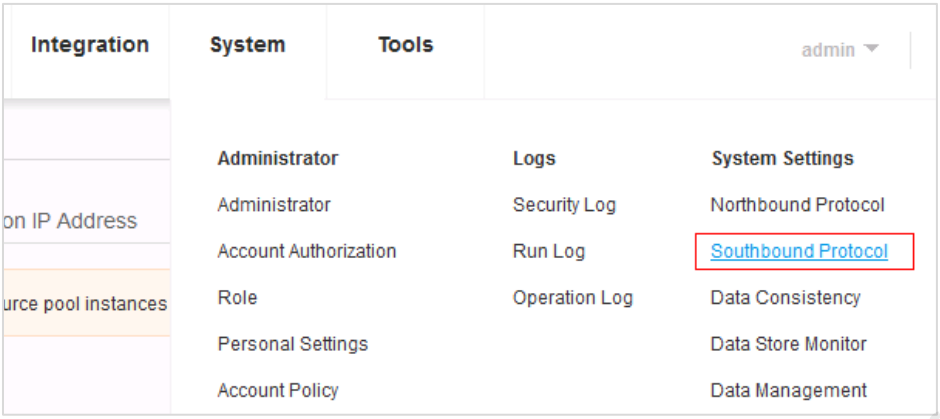

Choose **SNMP** > **SNMP Configuration**. Configure the trap service parameters. Click **Yes** to enable the trap service.

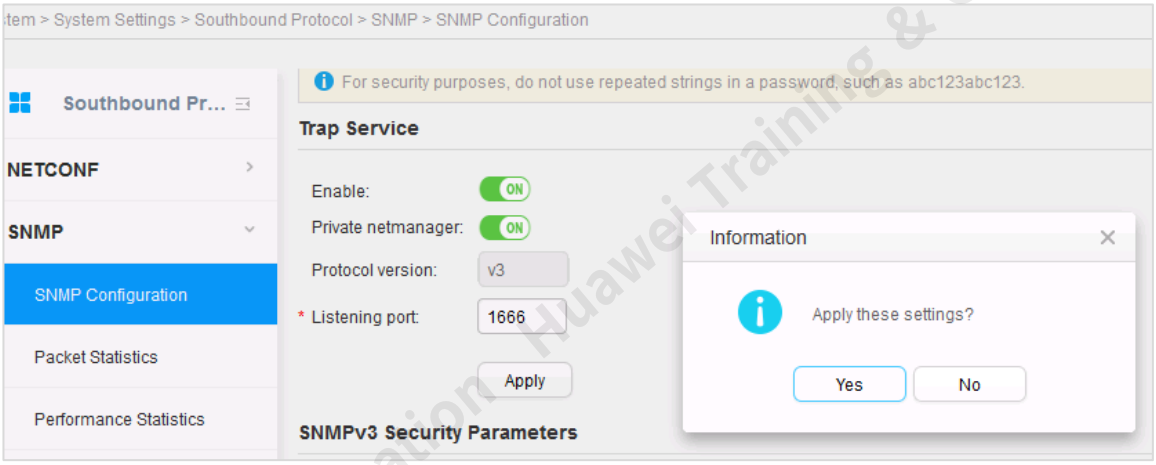

Click **Create**. Enter the user name, authentication protocol, authentication key, encryption algorithm, encryption key, and the password **Huawei@123**. The security parameter sent by Agile Controller-DCN to switches (used for device alarm encryption) can be the same as the parameter for device discovery. The firewall needs to be manually configured. For details, see section 3.8.3 "**Configuring the SNMP Parameters for the Firewall**".

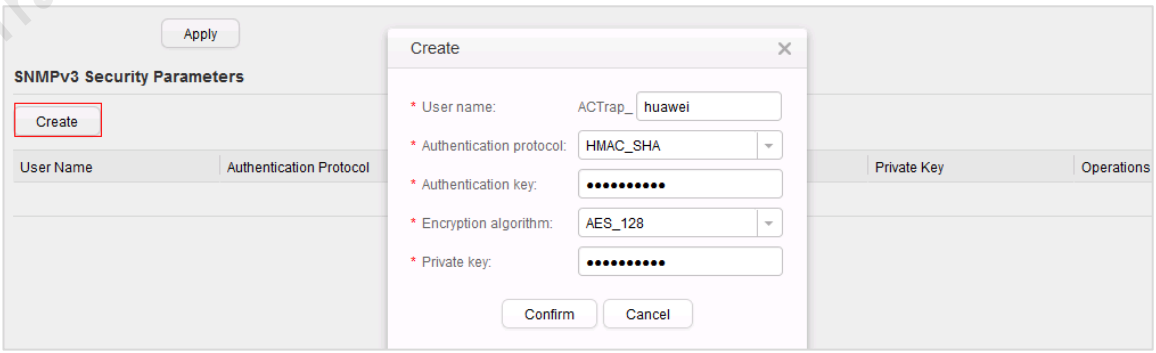

4.9 Creating an SFC Template

Create two SFC templates to enable the services to pass through the firewall and access the external network (one uses SNAT and the other does not use SNAT), facilitating the invocation during network service provisioning.

Choose **Tenant** > **SFC**, and click **Create**.

On the page that is displayed, set the parameters based on data planning:

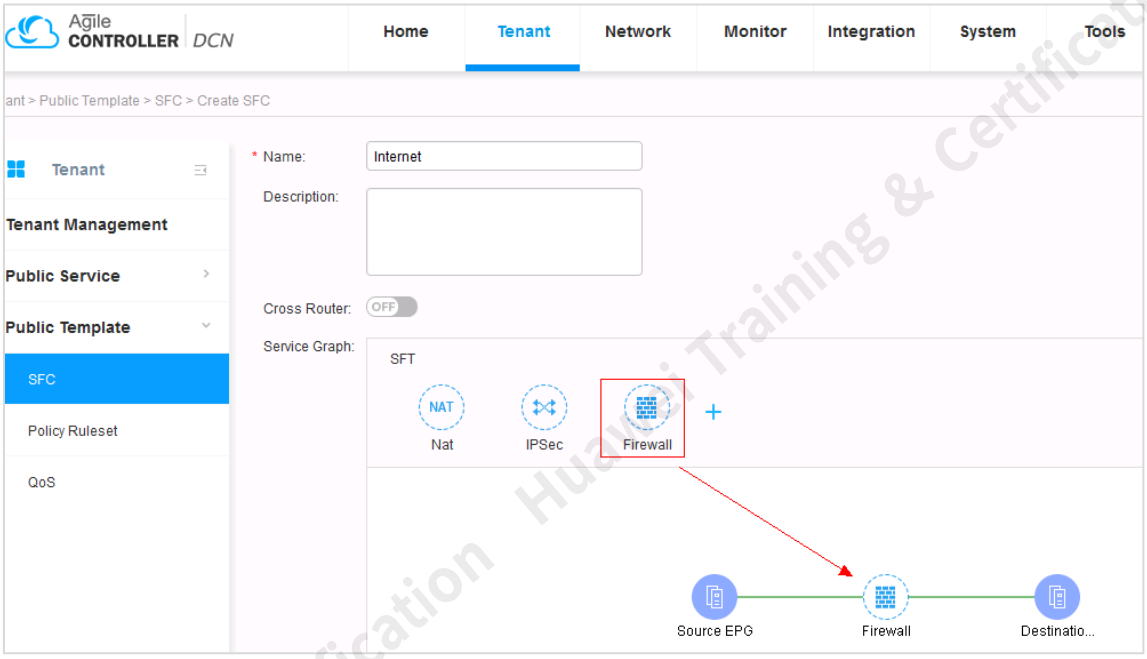

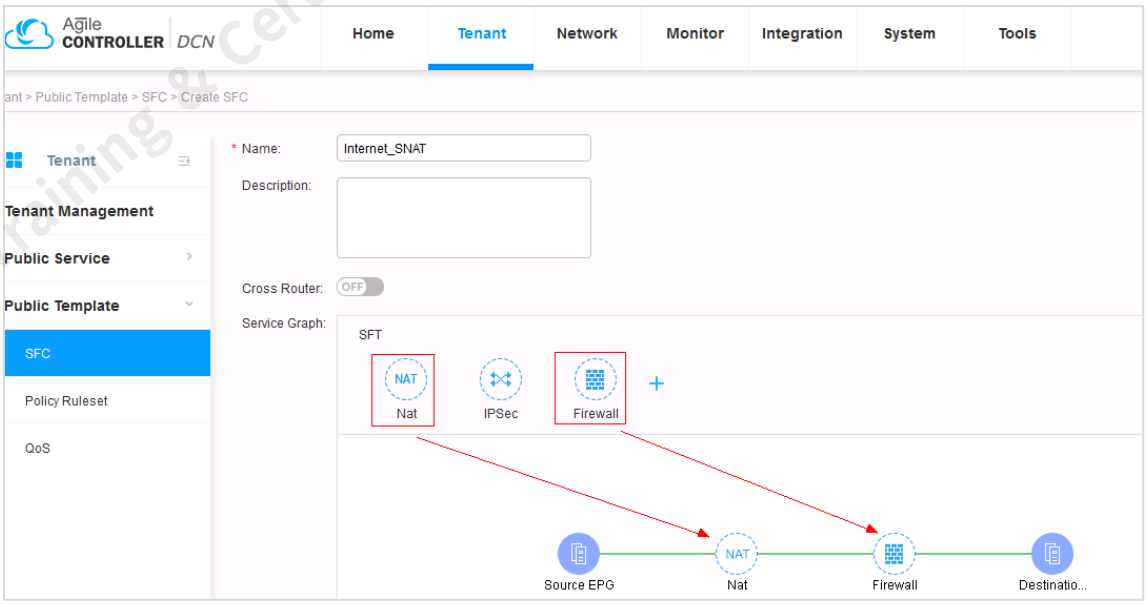

The two SFCs Internet and Internet_SNAT have been created.

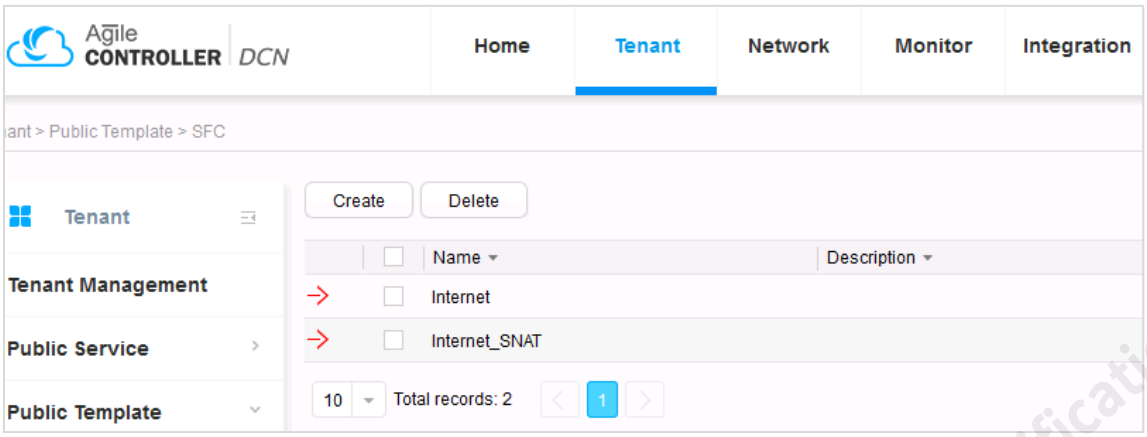

4.10 Deploying the Egress Network

Before VMs can access the Internet, the administrator needs to configure a public external gateway on the Agile Controller-DCN in advance. The external gateway must be consistent with the external network of the cloud platform for tenant router application and public network access.

In this example, two external networks are created. One is used for the tenant router to apply for the Intranet, and the other is used for the public network to access the Internet.

4.10.1 Configuring the Access to the Public Network (Default VRF)

Parameter	Data	Description
External network name	Public_Internet	
External network type	L3	Select **L3** for an L3 gateway in the SDN scenario. The value **L2** indicates an L2 gateway used to transmit traffic between LPUs and PE nodes in the NFV scenario. This value is used in scenarios where no gateway is provided on the network.
External gateway type	General	Used to access the public network.

Parameter	Data	Description
Fabric	Fabric	
Public IP address	NA	Not required in the cloud-network integration scenario.
Default VRF	ON	Preferentially use public VRF as the egress gateway. If more egresses need to be expanded or multiple isolated egresses need to be used, egresses providing an independent VRF can be created on the Agile Controller-DCN.
Device group name	GW	

Create a gateway that uses public VRF to access the external network, applying to the single-egress scenario when accessing the public network.

Step 1 **Creating an external network**

Choose Tenant > Public Service > External Gateway.

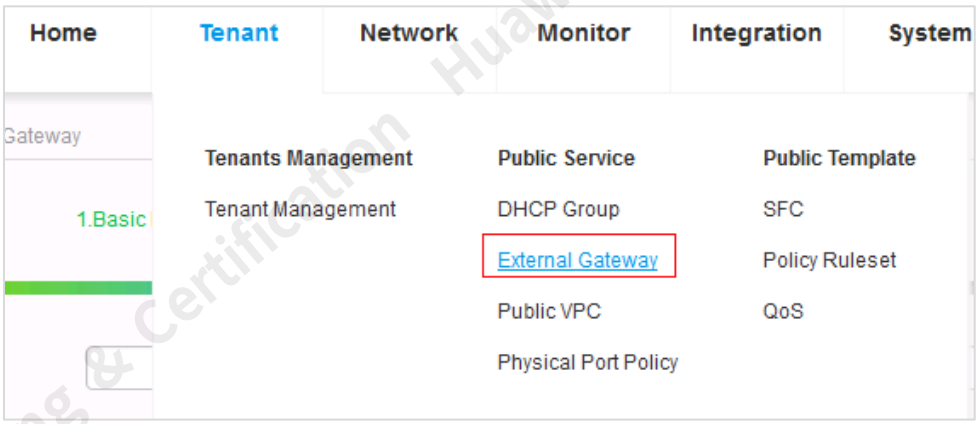

Click **Create**, enter basic information based on the prepared data, and click **Next**. Set **Default VRF** to **ON**, indicating that the public VRF on the device is used as an egress gateway.

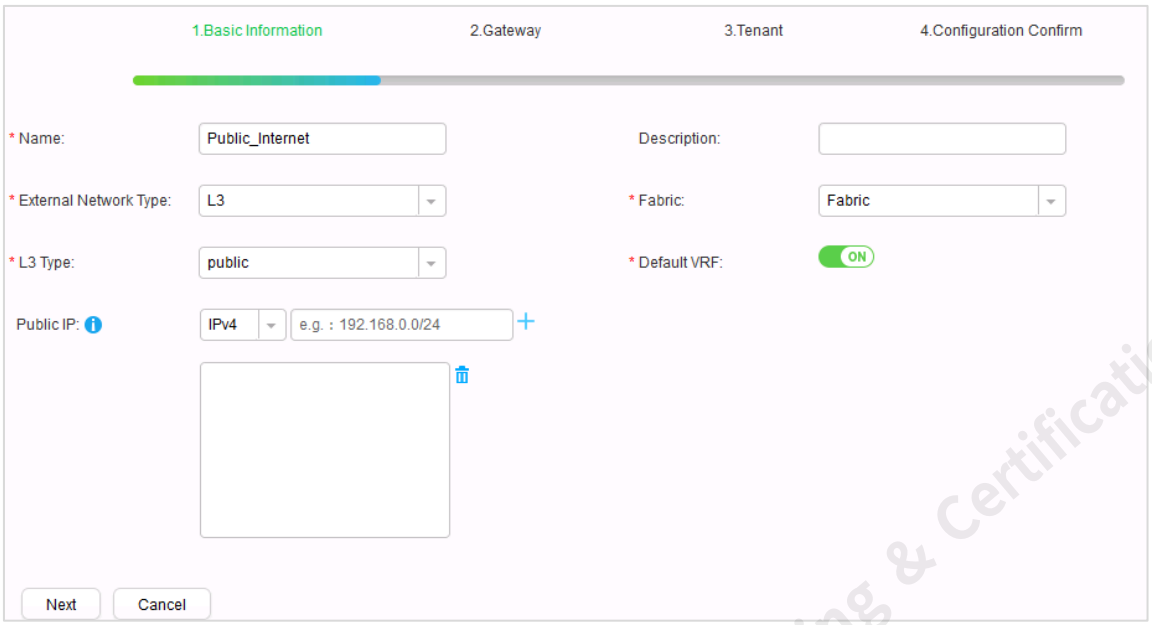

In the spine active-active scenario, the outbound interface is manually deployed. Therefore, set **Enable Outbound Interface** to **OFF**.

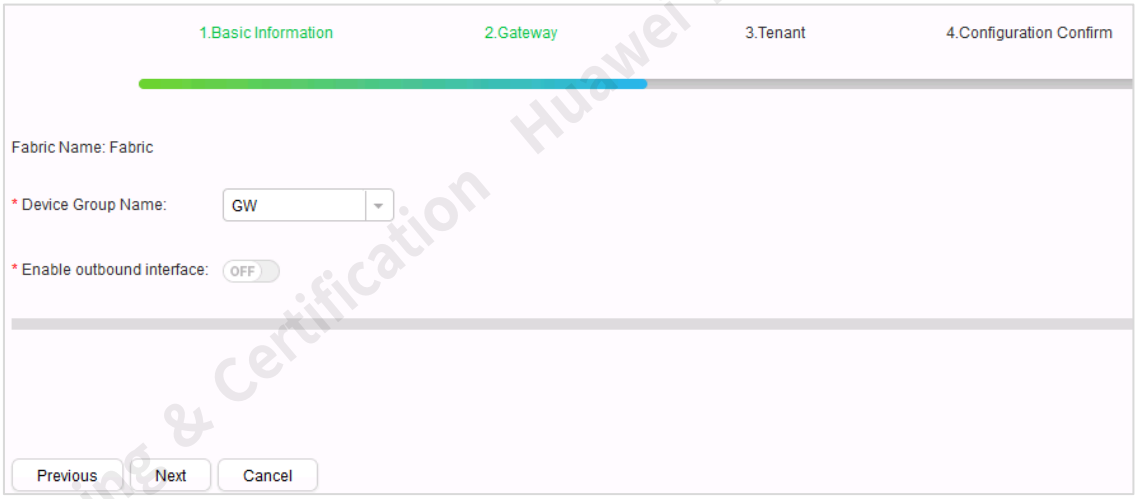

Set **Tenant Share** to **ON** and click **Next**.

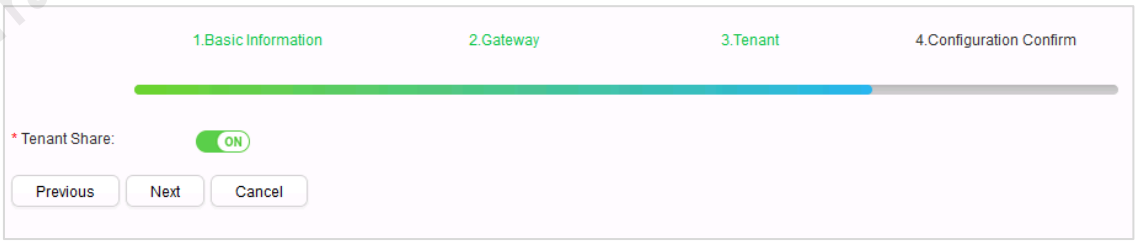

Click **Complete** to start creating the external network.

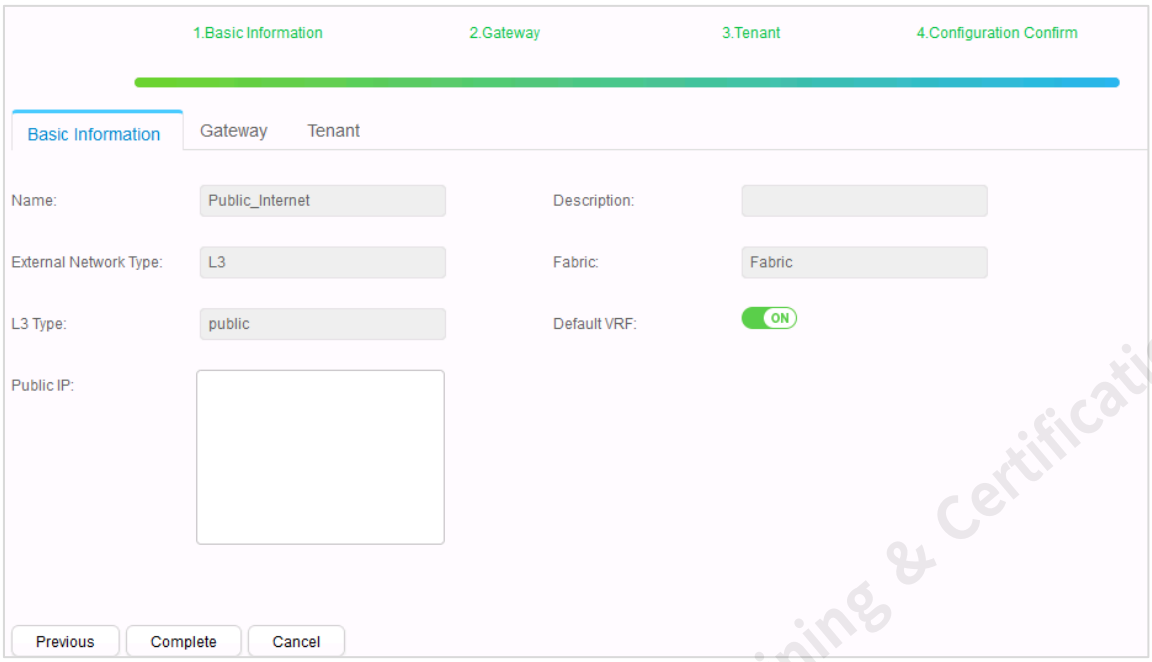

Step 2 **In this case, the Agile Controller-DCN does not deliver any configuration to the device. Manually configure the interconnection interfaces and routes between spine nodes and PE nodes.**

Configure the PE interface.

```
interface GigabitEthernet0/0/1
 ip address 10.125.97.2 255.255.255.252
#
interface GigabitEthernet0/0/2
 ip address 10.125.97.6 255.255.255.252
#
```

Configure the Spine-1 egress.

```
[Spine-1] interface 10GE 1/0/1
[Spine-1] undo portswitch
[Spine-1] ip address 10.125.97.1 255.255.255.252
[Spine-1] quit
[Spine-1] ip route-static 0.0.0.0 0.0.0.0 10.125.97.2 preference 120
[Spine-1] ip route-static 0.0.0.0 0.0.0.0 10.1.1.46 preference 150 //
Configure the best-effort link.
```

Configure the Spine-2 egress.

```
[Spine-2]interface 10GE 1/0/2
[Spine-2] undo portswitch
[Spine-2] ip address 10.125.97.5 255.255.255.252
[Spine-2] quit
[Spine-2] ip route-static 0.0.0.0 0.0.0.0 10.125.97.6 preference 120
[Spine-2] ip route-static 0.0.0.0 0.0.0.0 10.1.1.45 preference 150 //
Configure the best-effort link.
```

Test the connectivity.

```
<Spine-1>ping 10.125.97.2
  PING 10.125.97.2: 56  data bytes, press CTRL_C to break
    Reply from 10.125.97.2: bytes=56 Sequence=1 ttl=255 time=1 ms
    Reply from 10.125.97.2: bytes=56 Sequence=2 ttl=255 time=1 ms
```

----End

4.10.2 Configuring the Intranet of the Tenant Router

The procedure is the same as that for configuring the public network to access the Internet. The only difference operation is that in step 2, set **L3 Type** to **private** and set **Fabric** to a specified value to enable the router correspond to the network when the tenant applies for the router.

Create an external network.

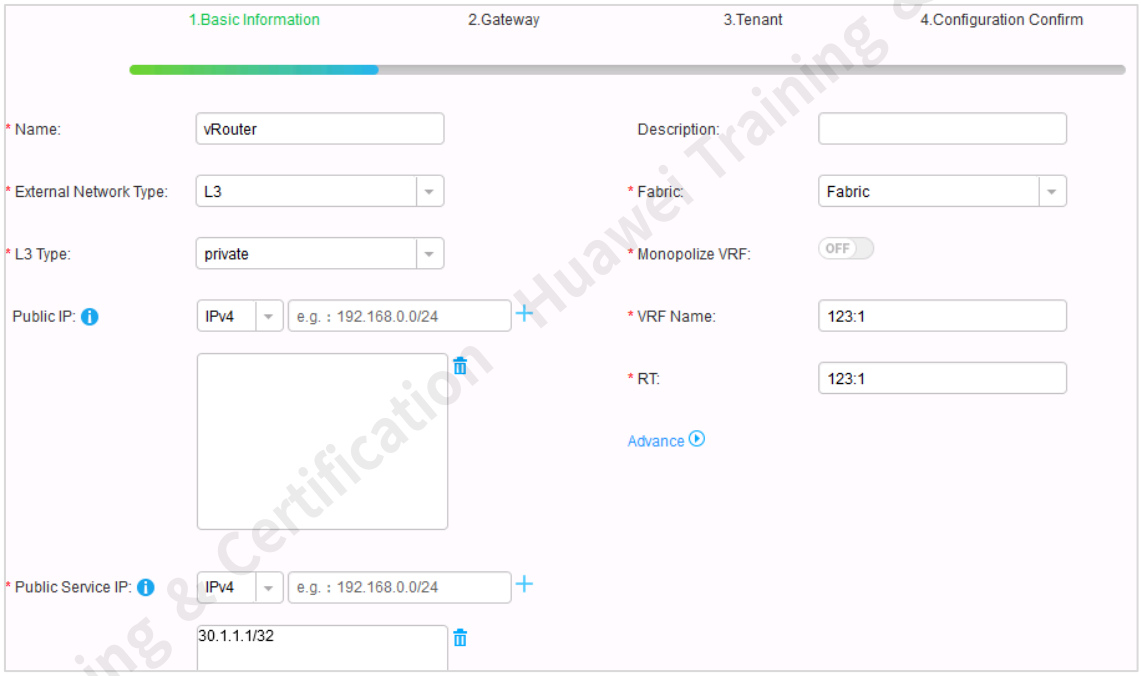

The external network is successfully created.

5 FusionStorage Block Configuration

5.1 Objectives

- Understand how to configure the FusionStorage Block converged deployment solution.
- Understand how to configure the storage resource to access the virtualization environment.

5.2 Networking and Service Description

FSM is deployed on the VMs of the FusionSphere OpenStack. In the converged deployment scenario where FSA nodes and CNA components are deployed, deploy resource pools and block storage clients. Connect the storage resources to the virtualization environment.

Network Plane	VLAN	IP
Production Storage	4011	10.11.0.0/24
Management Name	4005	10.1.0.0/24

5.3 Performing the Converged Deployment Configuration

5.3.1 Checking the Computing Cluster Configuration

Cluster1 and Cluster2 have been created on the VRM. Add CNA01 and CNA02 to Cluster1, and add CNA03 to Cluster2.

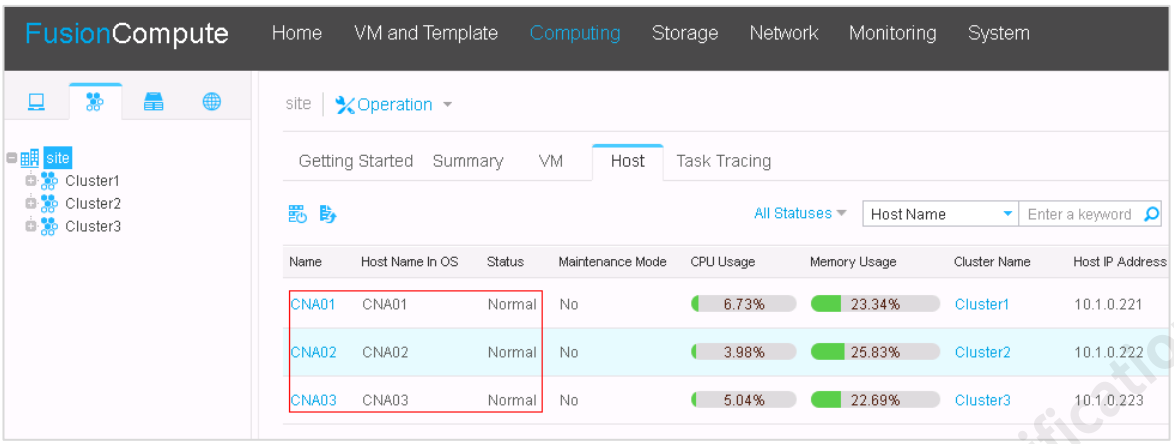

5.3.2 Installing FSA

Step 1 **Connect to servers.**

Click **Single Add** and enter the management address of the converged deployment server.

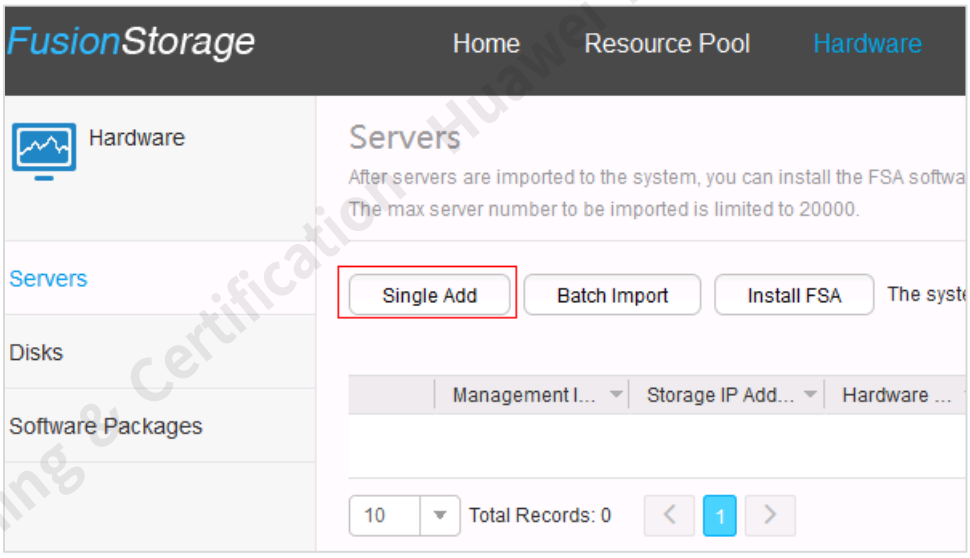

Connect the three servers in sequence.

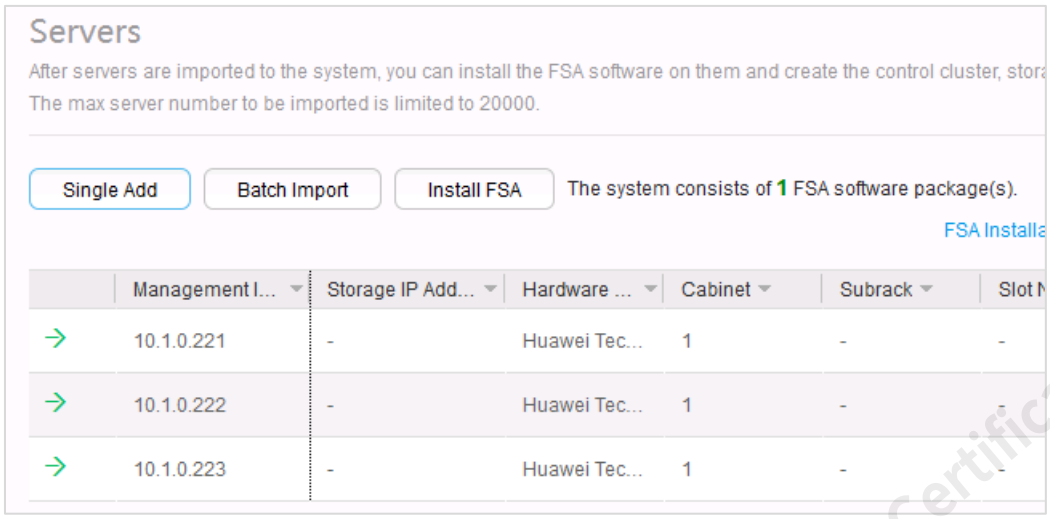

Step 2 **Install FSA.**

Select a software package on the **Install FSA** page.

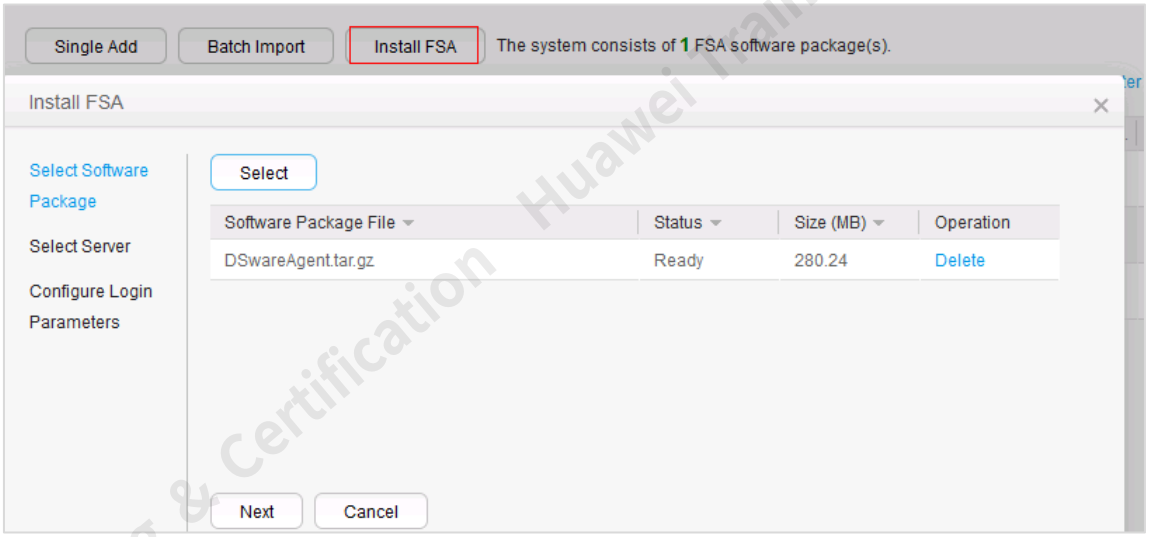

Select servers.

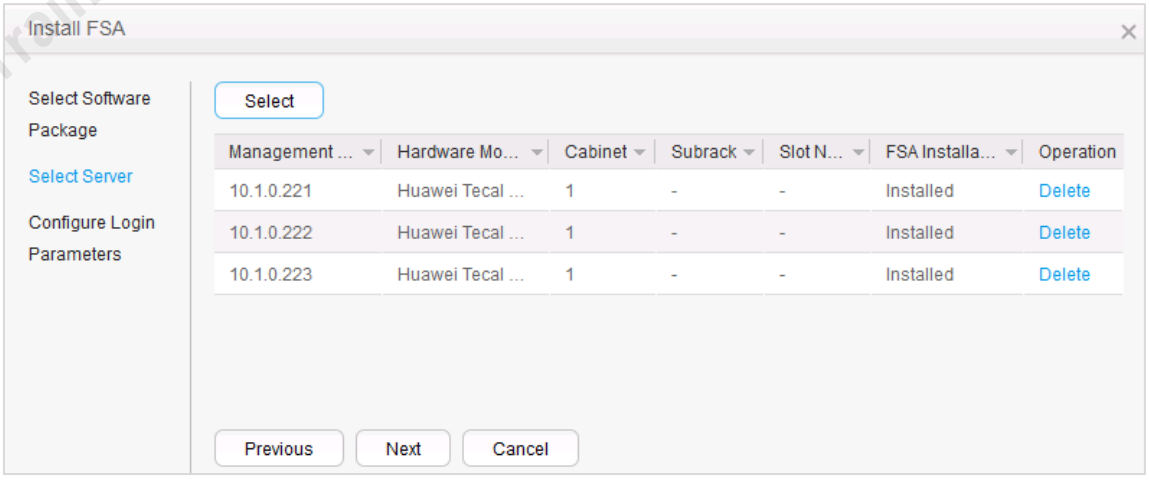

Enter the password for logging in to the server. Enter the username **gandalf,** the password **Huawei@CLOUD8**, and the root password.

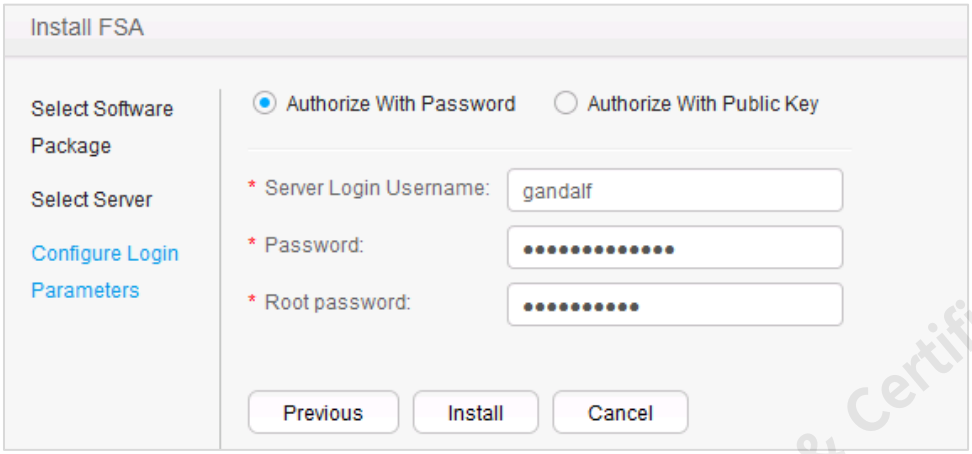

The FSA is successfully installed.

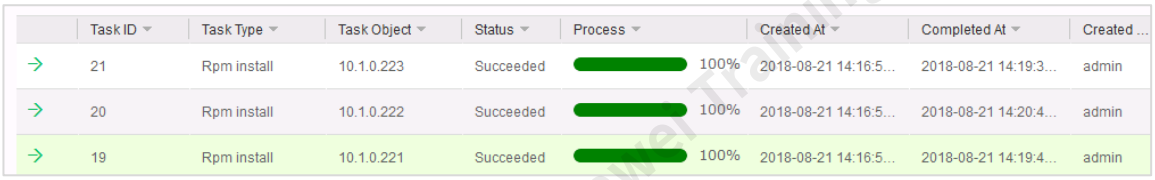

----End

5.3.3 Configuring FusionStorage Block

Step 1 **Configure a resource pool.**

Click **Resource Pool** and then click **Configure** on the pop-up window.

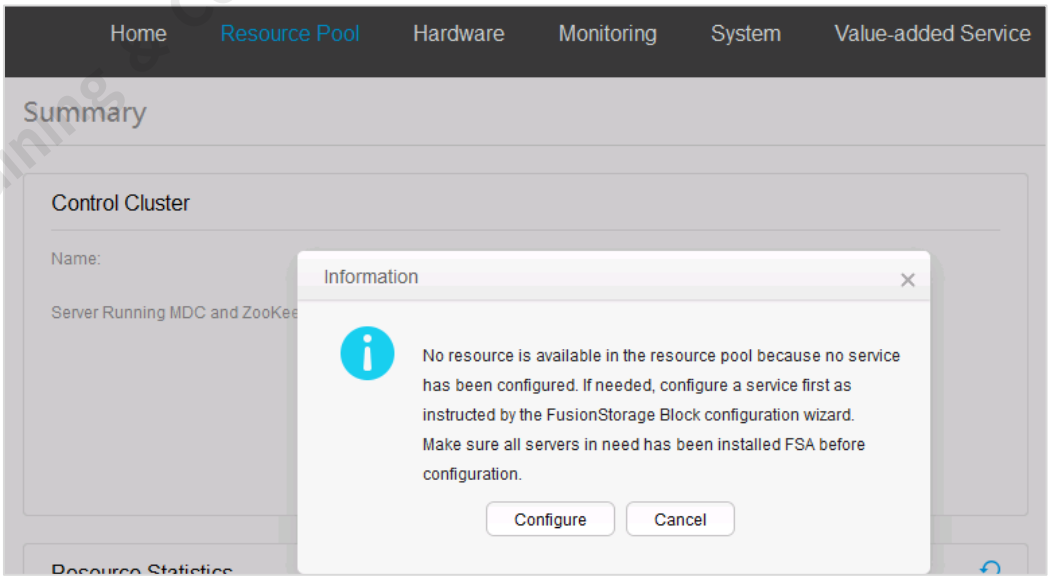

Configure a storage port, set the name of network port on the CNA to **iscsi-0**, and click **Next**.

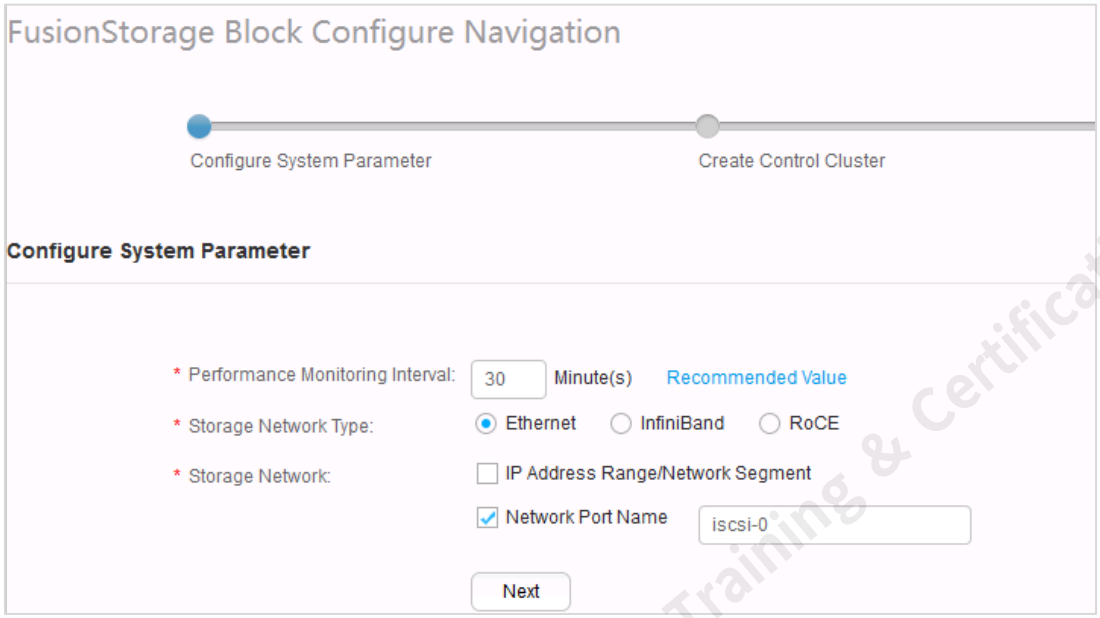

Select a cluster name and a metadata disk, and click **Create**.

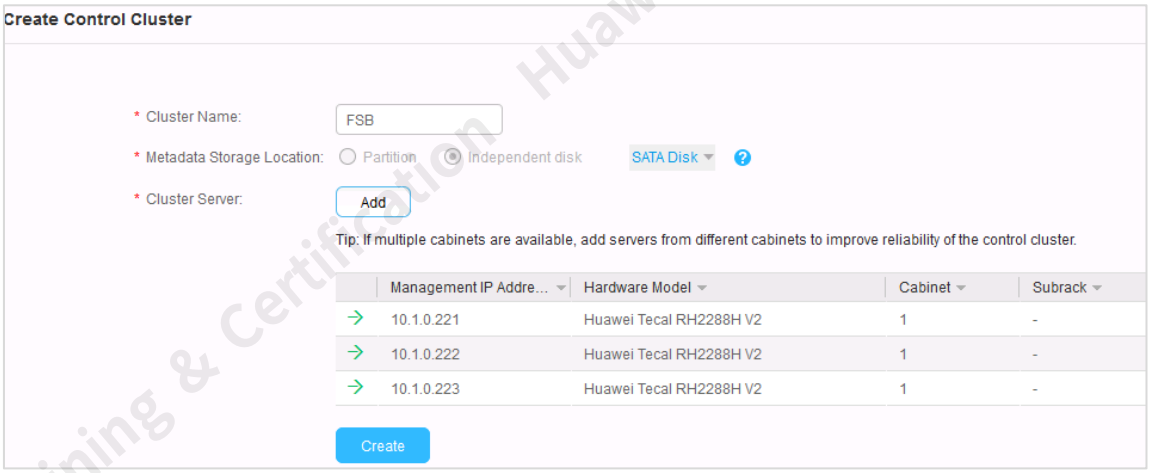

Select main storage and cache disks on the **Create Storage Pool** page.

Calculate the cache/main storage ratio. The default algorithm is as follows: Total cache capacity on each server/ (Number of OSD nodes on each server x Capacity of a single primary storage device).

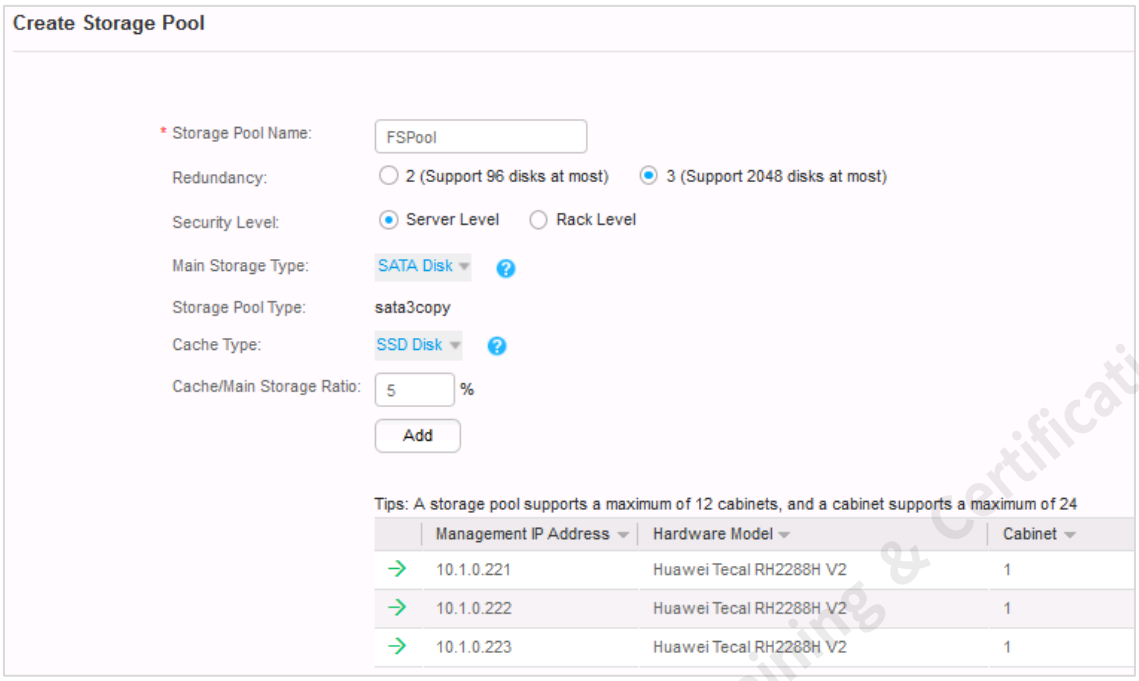

The resource pool is successfully created.

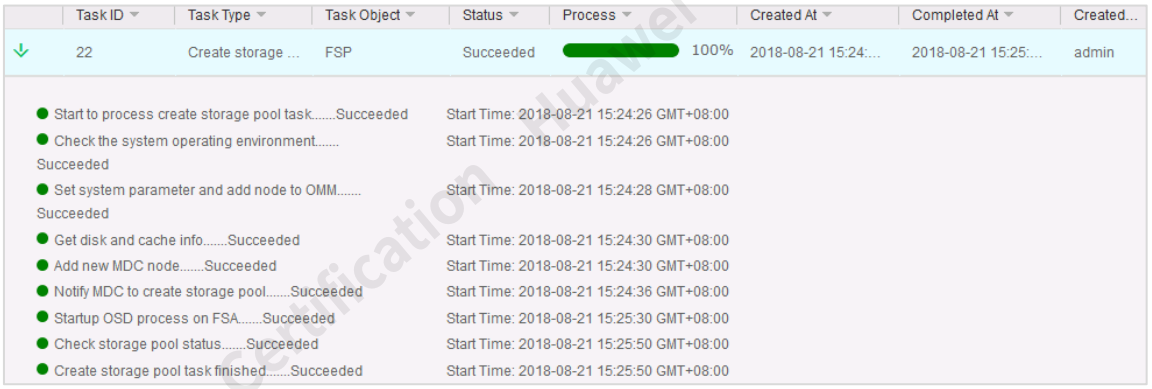

Step 2 **Create a block storage client.**

After the resource pool is created, create a block storage client.

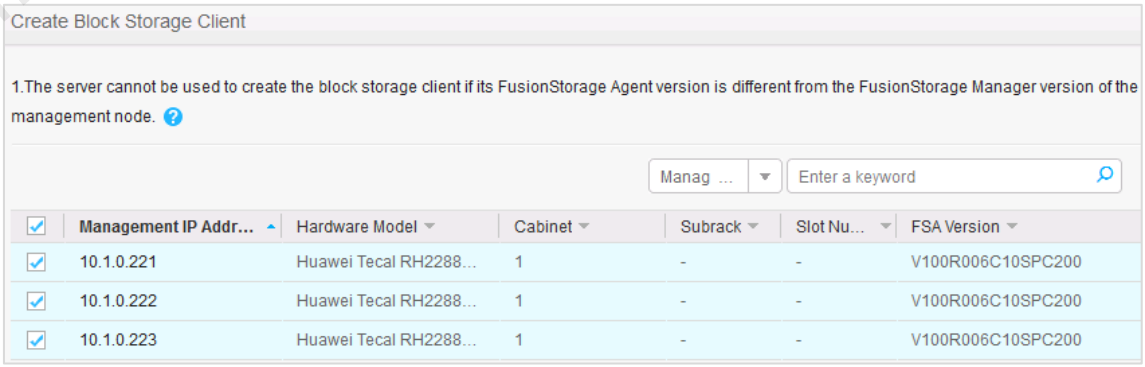

Select a server and the block storage client is created successfully.

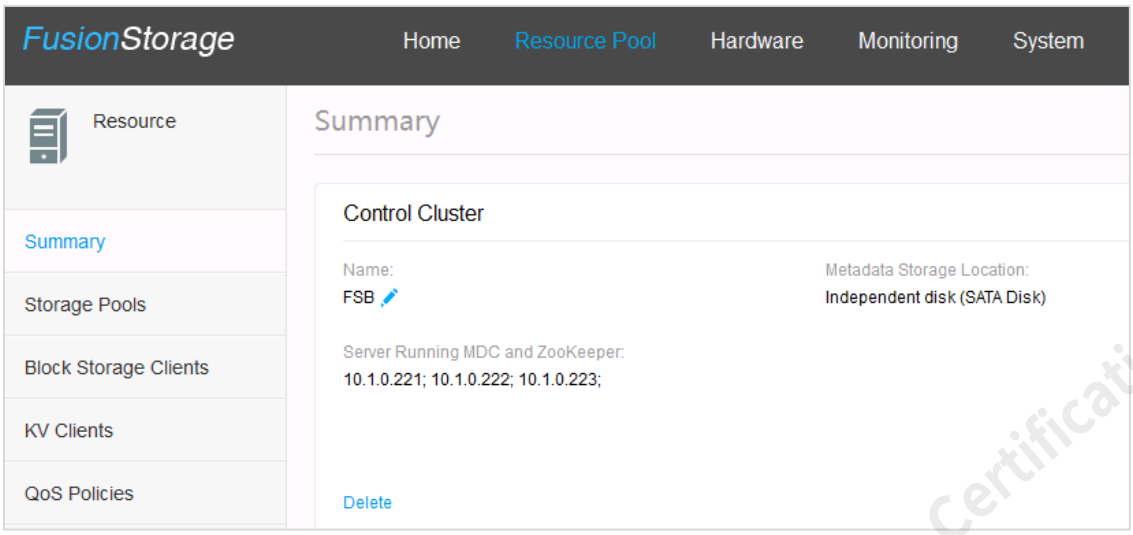

----End

5.4 Interconnecting FusionCompute to FusionStorage

5.4.1 Configuring the CNA Storage Port

Choose **Configuration** > **Add Storage Port** to add storage ports for the compute nodes.

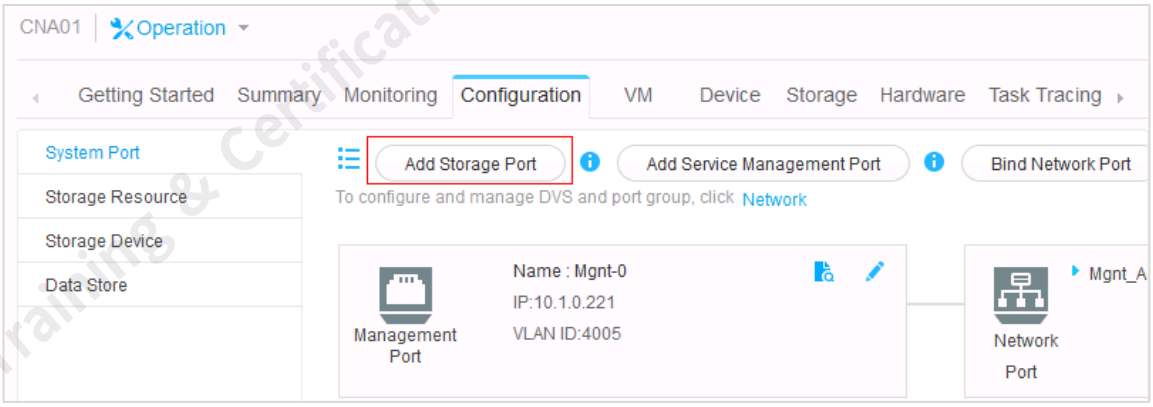

If there are only two NICs, reuse the management network ports in active/standby mode.

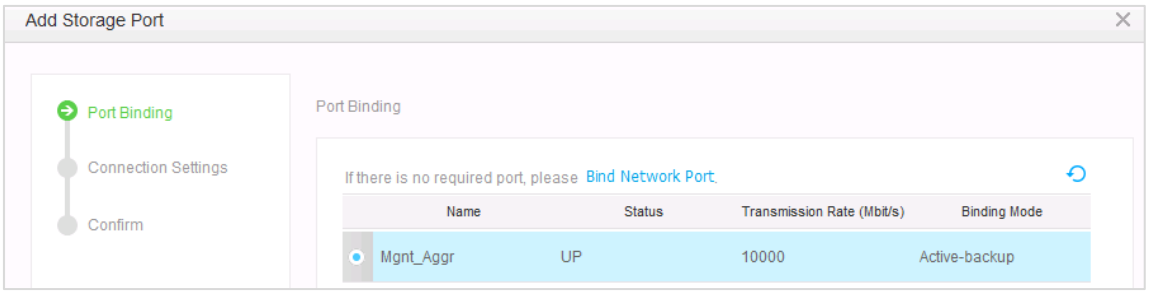

Enter the storage port IP address and VLAN information based on the plan.

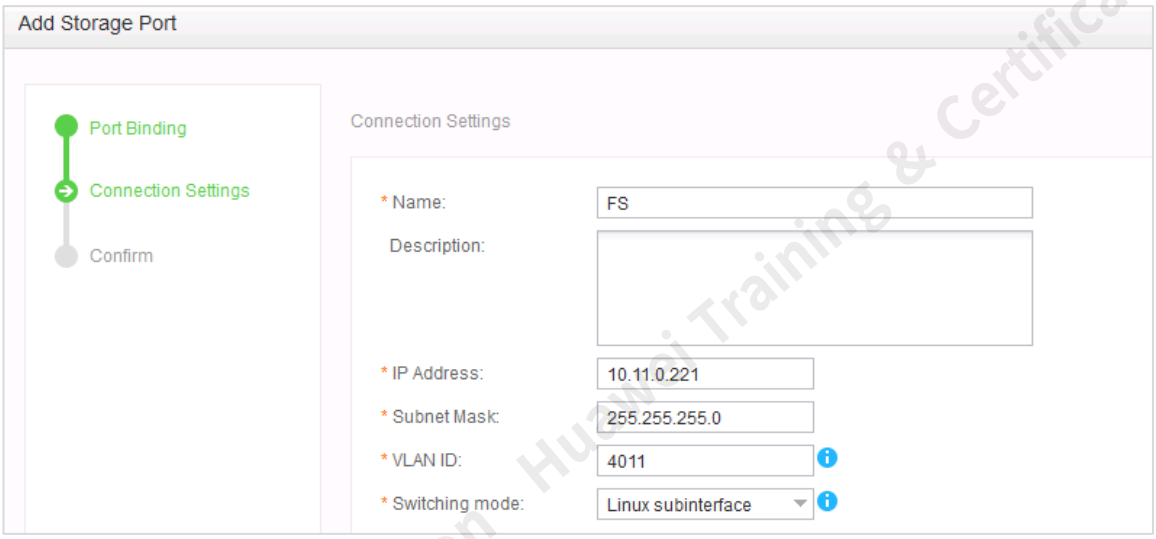

Note that the **Switching Mode** must be set to **Linux subinterface** to configure the storage ports of CNA01, 02, and 03. The storage ports of the CNA03 are bonded in active/standby mode. Note that the bond NICs are distinguished from those using M-LAG to connect to the ToR.

The underlying VLAN4011 storage network plane needs to be connected.

5.4.2 Adding Storage Resources

Choose **Storage** > **Storage Resource**, and click **Add** to add storage resources.

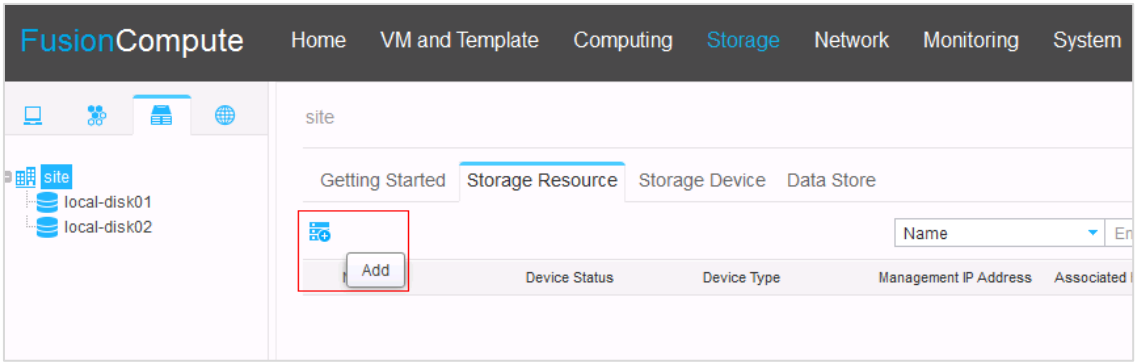

Select FusionStorage.

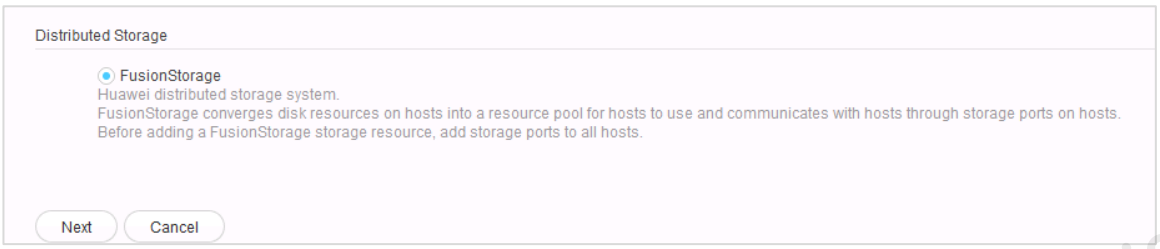

Enter the Management IP Address of the FusionStorage Manager and click **Finish**.

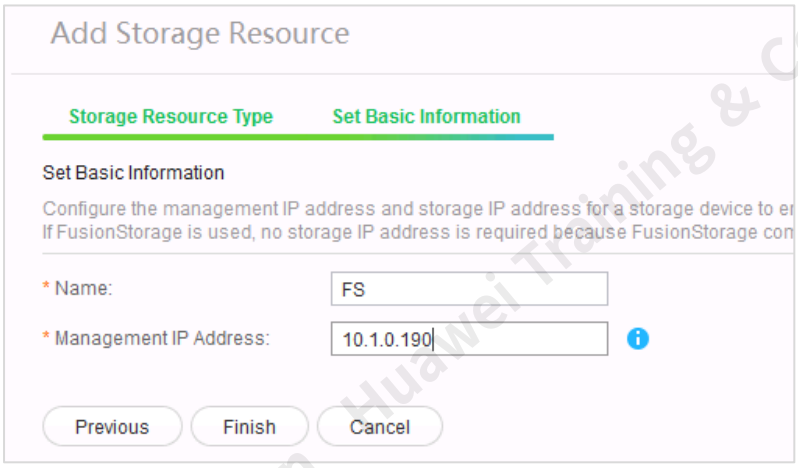

5.4.3 Adding a Data Storage Resource

Step 1 **Associate host**

View the value of the associated host. Choose **More** > **Associate Host**.

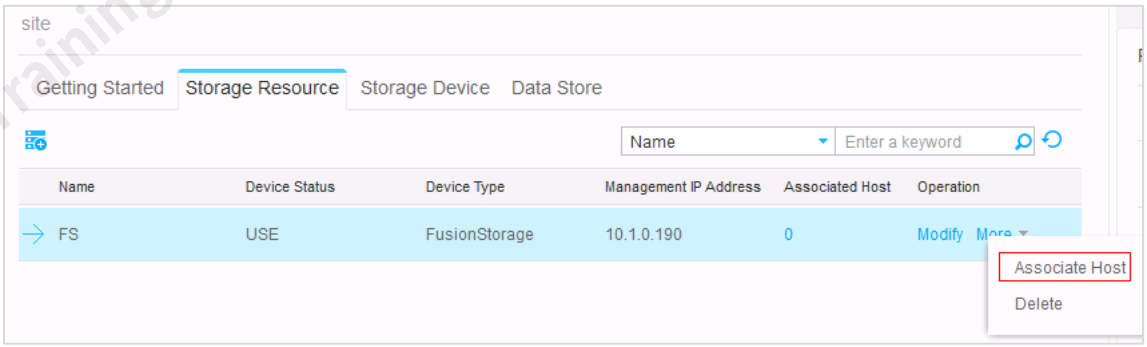

Associate all compute nodes.

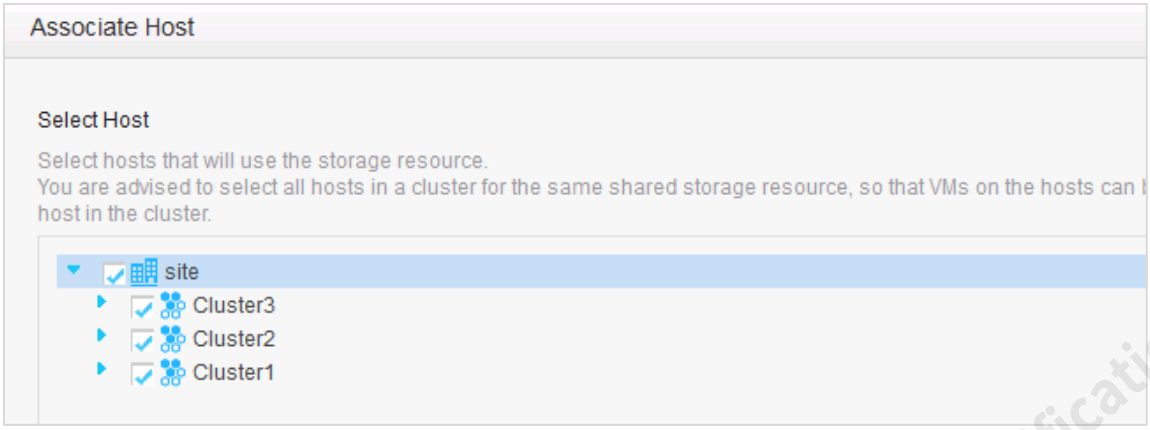

Step 2 **Scan the storage devices.**

Choose **Storage Resource** > **Scan** to scan the storage devices. The host reads the FusionStorage block resources.

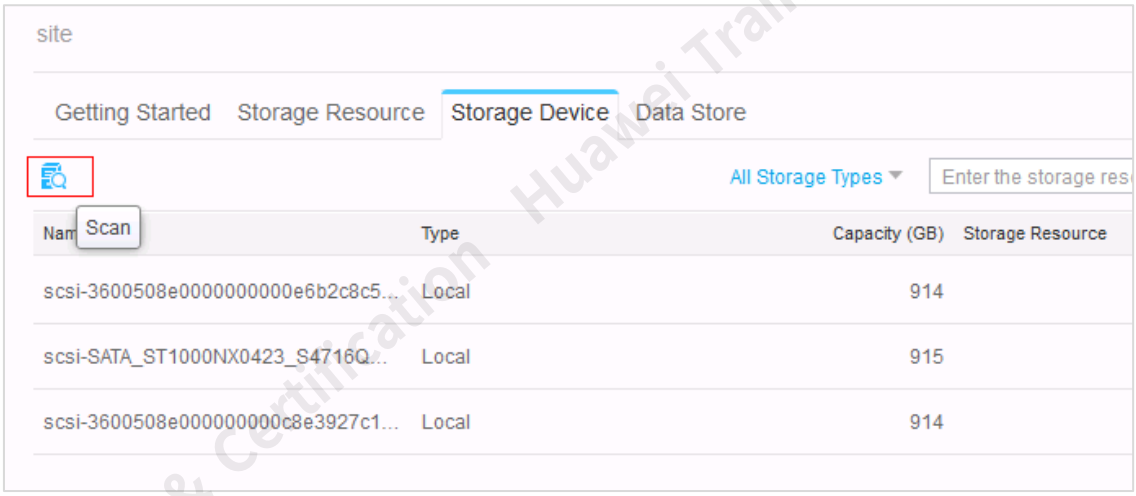

Select the hosts to scan.

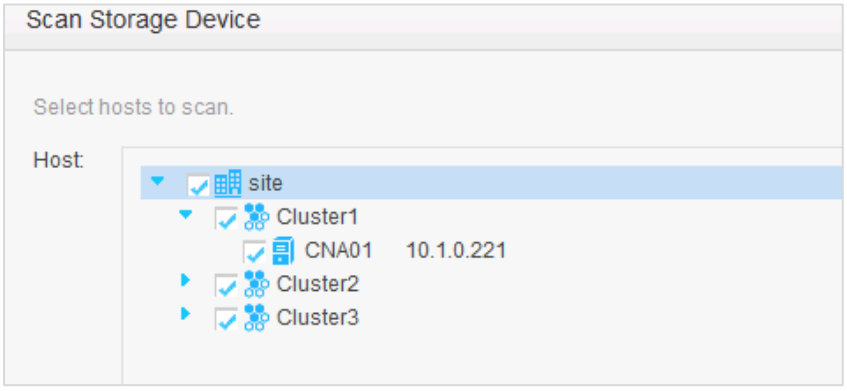

FusionStorage is successfully scanned.

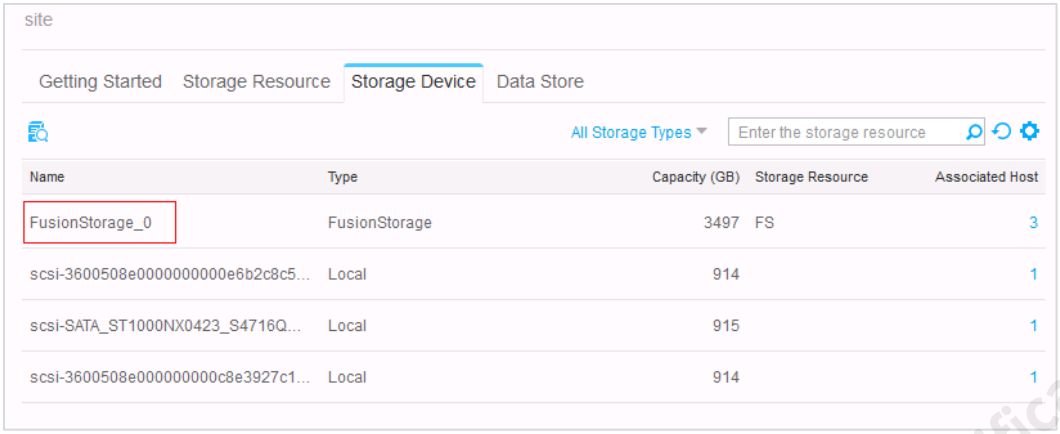

Step 3 **Add a data storage resource**

Add the scanned block storage to **Data Store**.

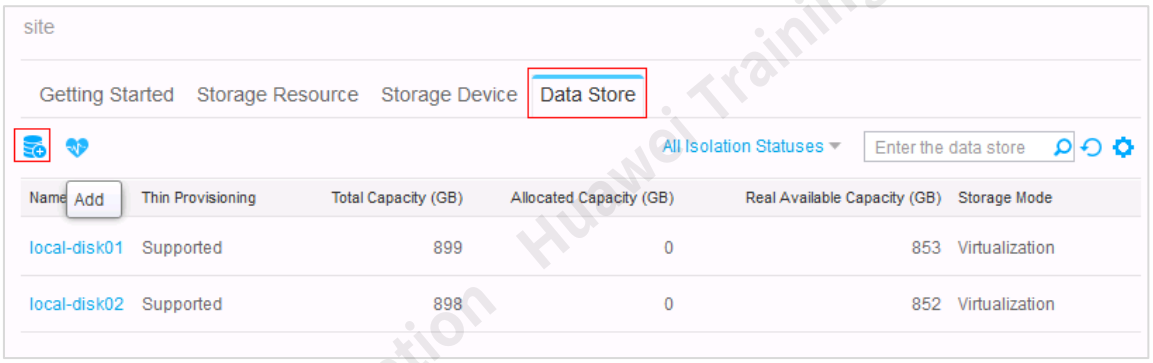

Select FusionStorage.

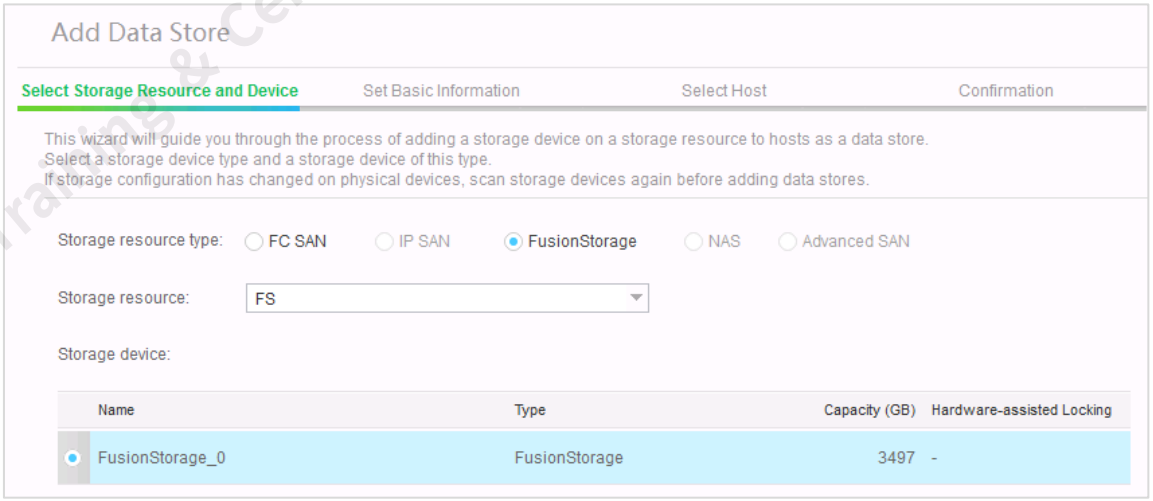

Enter the data store name.

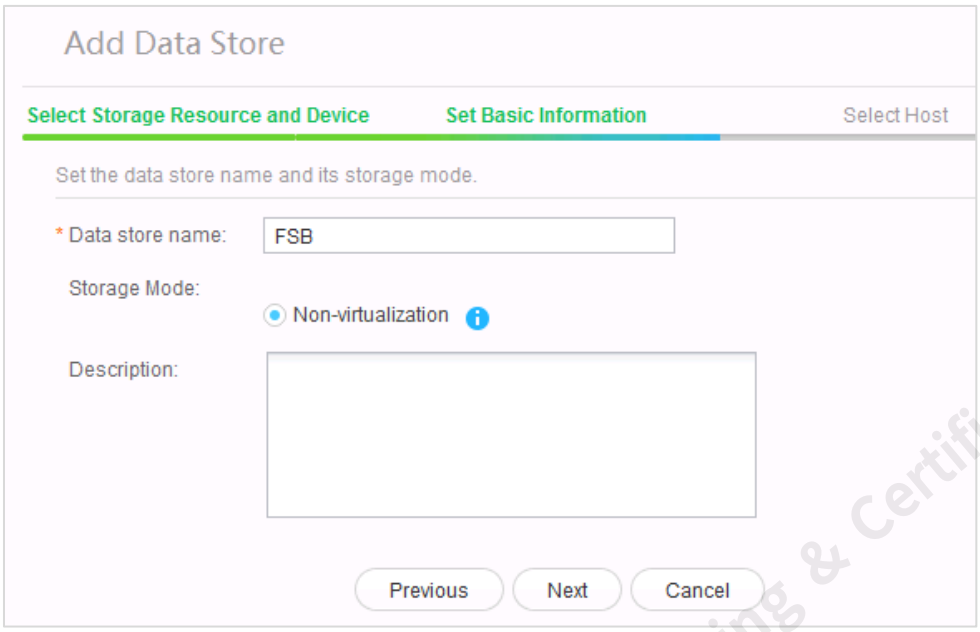

Add a host that will use the data store.

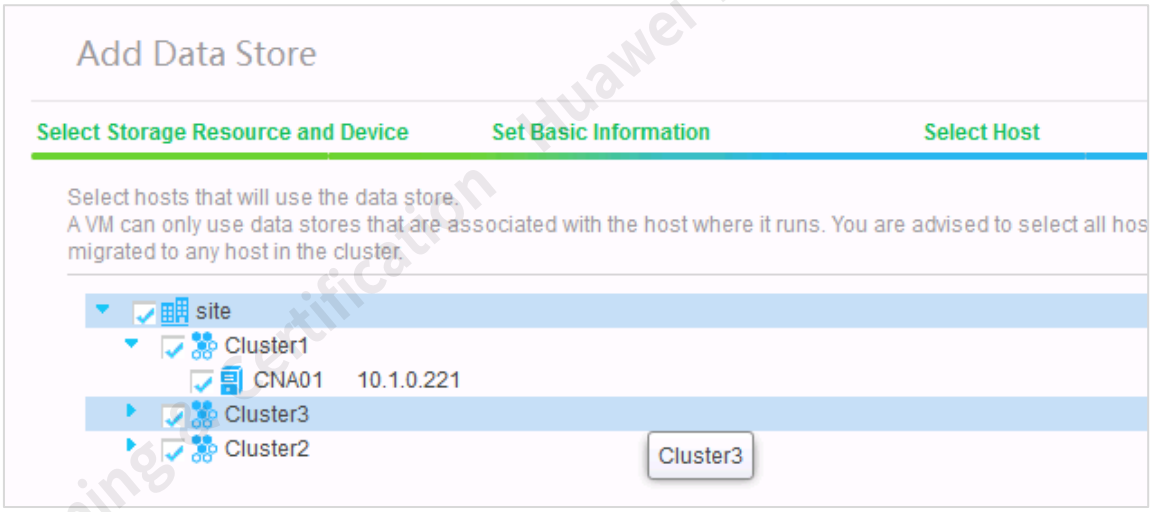

The data store is created successfully.

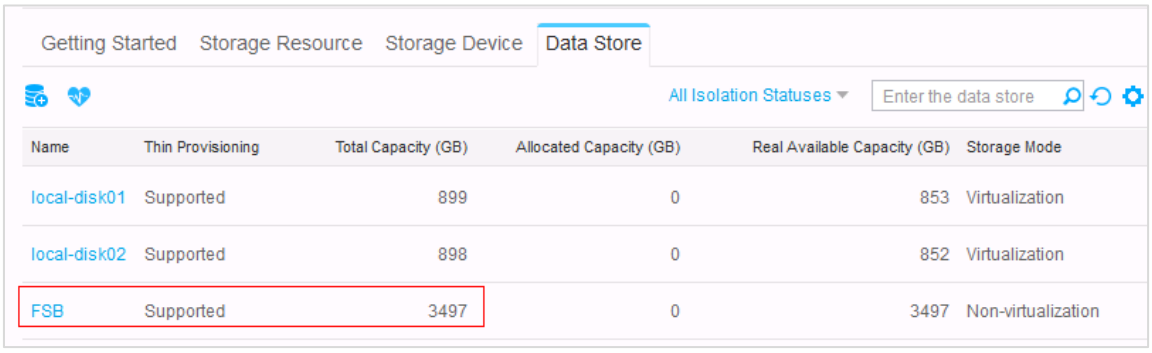

----End

6 FusionSphere OpenStack Resource Configuration

6.1 Objectives

- Understand how to configure the FusionCompute resource pool.
- Learn the configuration and commissioning of FusinSphere OpenStack resources.

6.2 Configuring FusionCompute Resource Pools

6.2.1 Checking the Distributed Switch Configuration

Pre-create the distributed switch vSW. The physical network name must be the same as that of the CPS physical network.

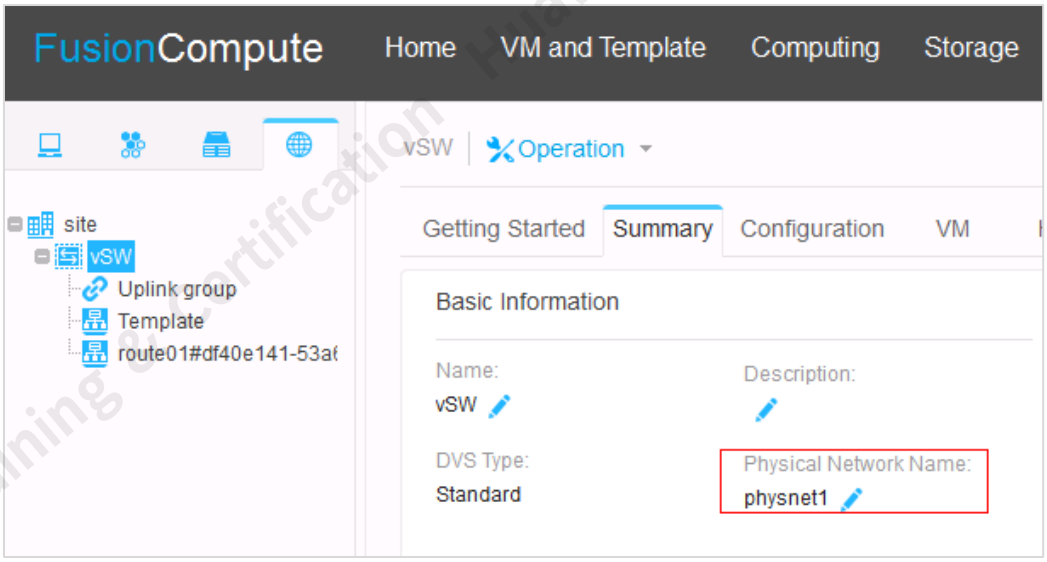

6.2.2 Interconnecting VRM to FusionSphere OpenStack

Step 1 **Interconnecting VRM to FusionSphere OpenStack.**

Choose **System** > **Connect To** > **OpenStack**, and check the parameters for interconnecting VRM to OpenStack.

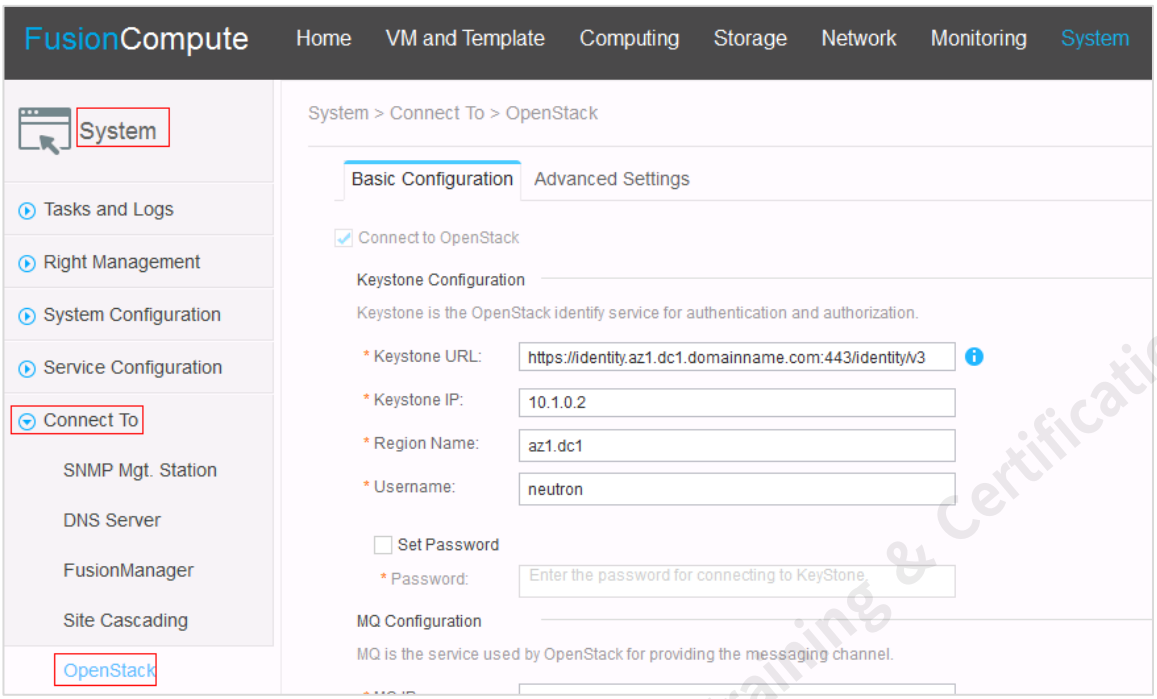

Step 2 **Creating an interconnection account on the VRM.**

Choose System > Right Management > User Management, and click Add User.

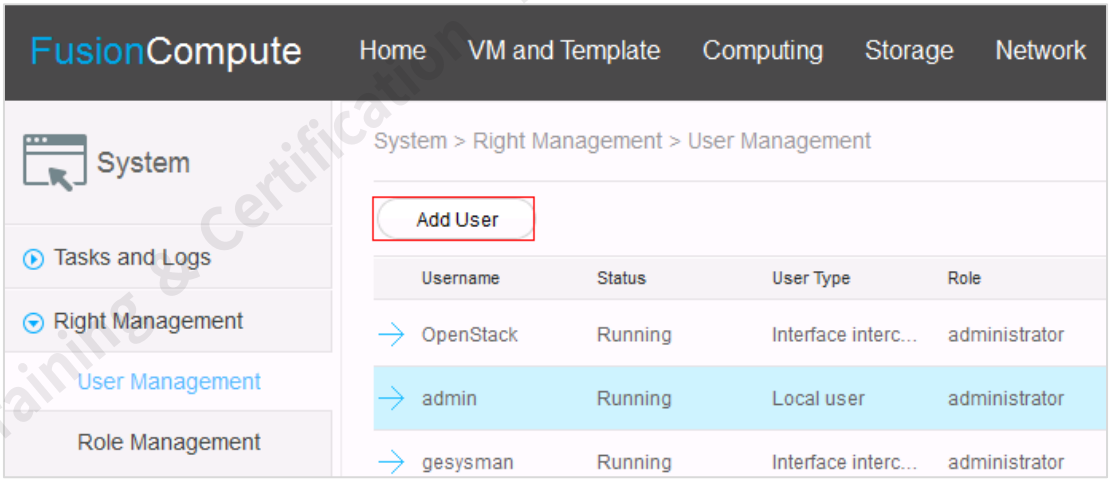

Enter the username **OpenStack** and the password **Huawei@123**.

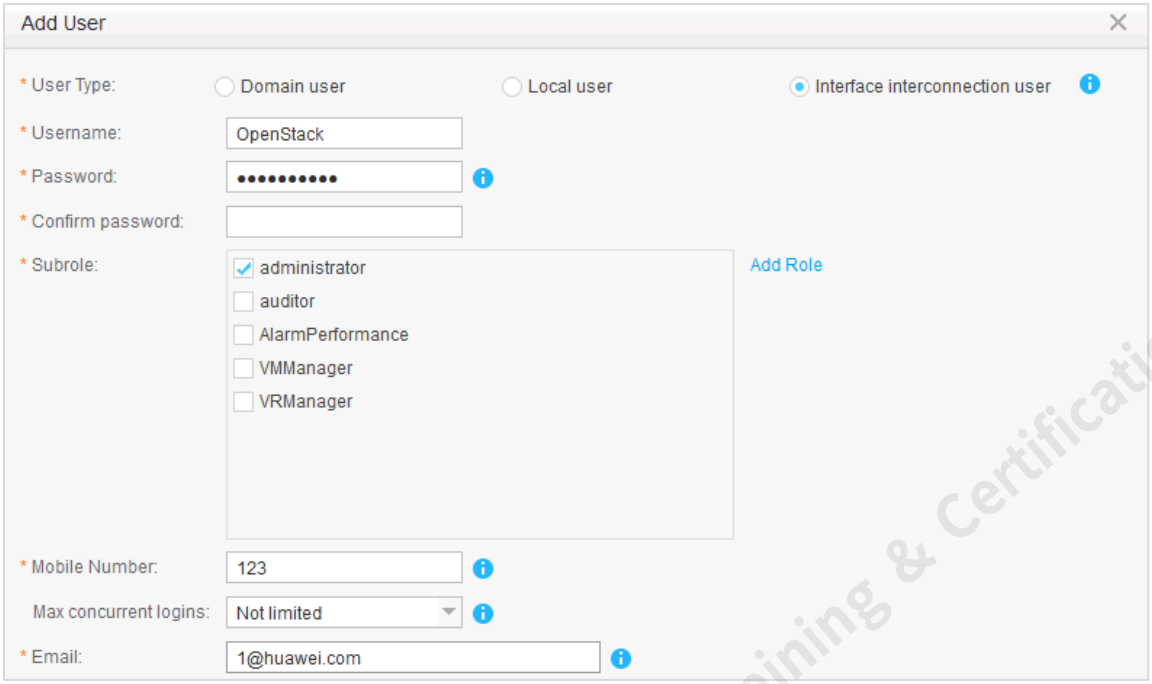

Step 3 **Interconnect CPS to VRM.**

Log in to the CPS, click **Resource Pool**, and connect to FusionCompute.

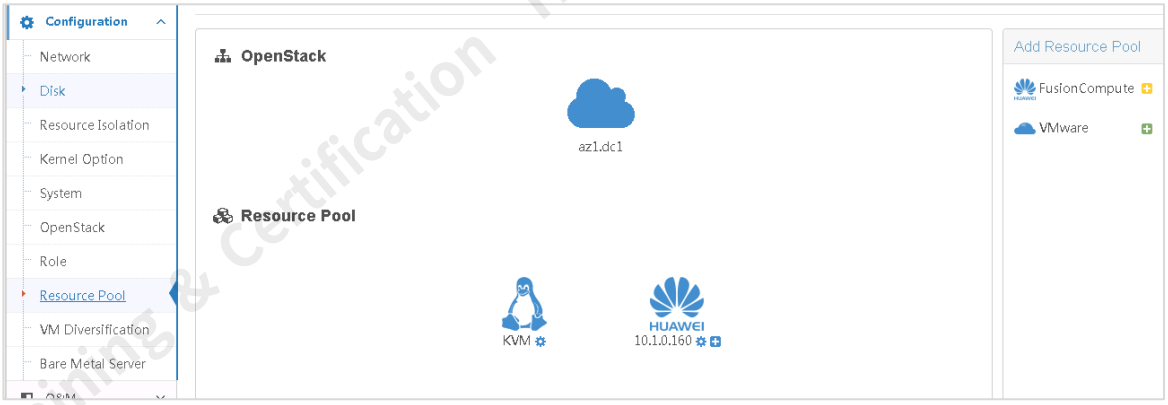

Configure the interconnection parameters.

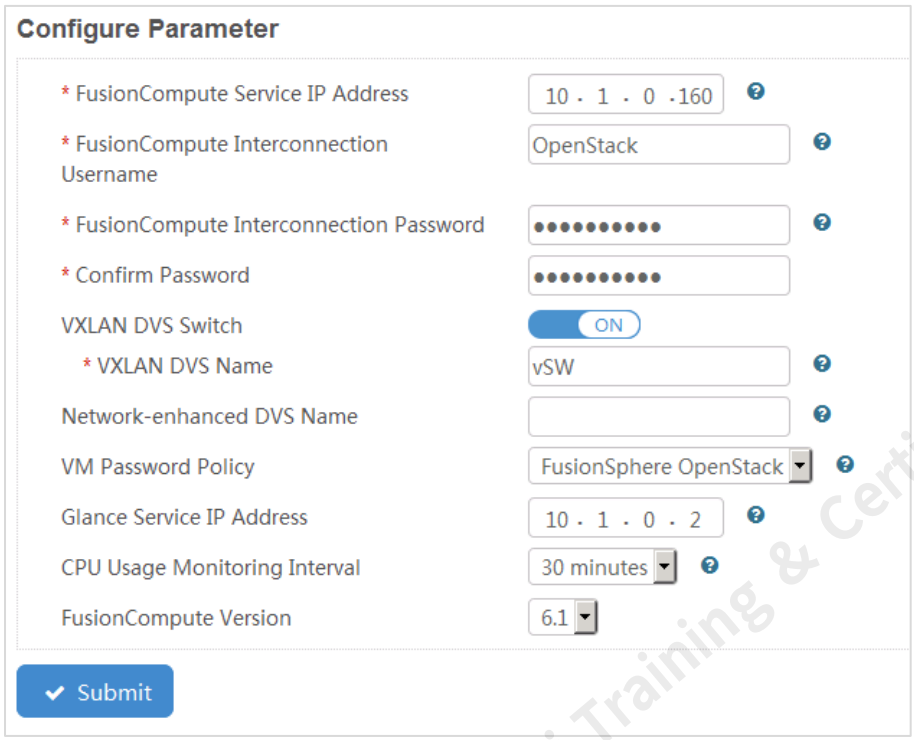

The interconnection is successful.

----End

6.3 Configuring FusionSphere OpenStack Resource Pools

6.3.1 Configuring Computing Clusters

Click + on the **Resource Pool** page to configure resources.

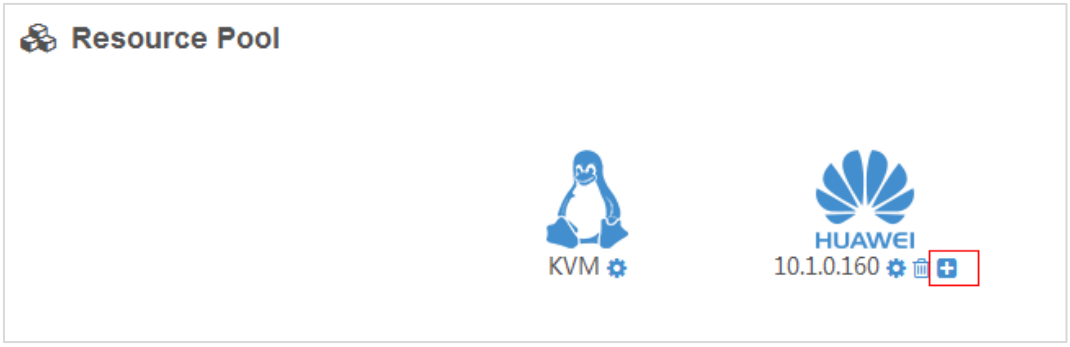

Add a computing cluster.

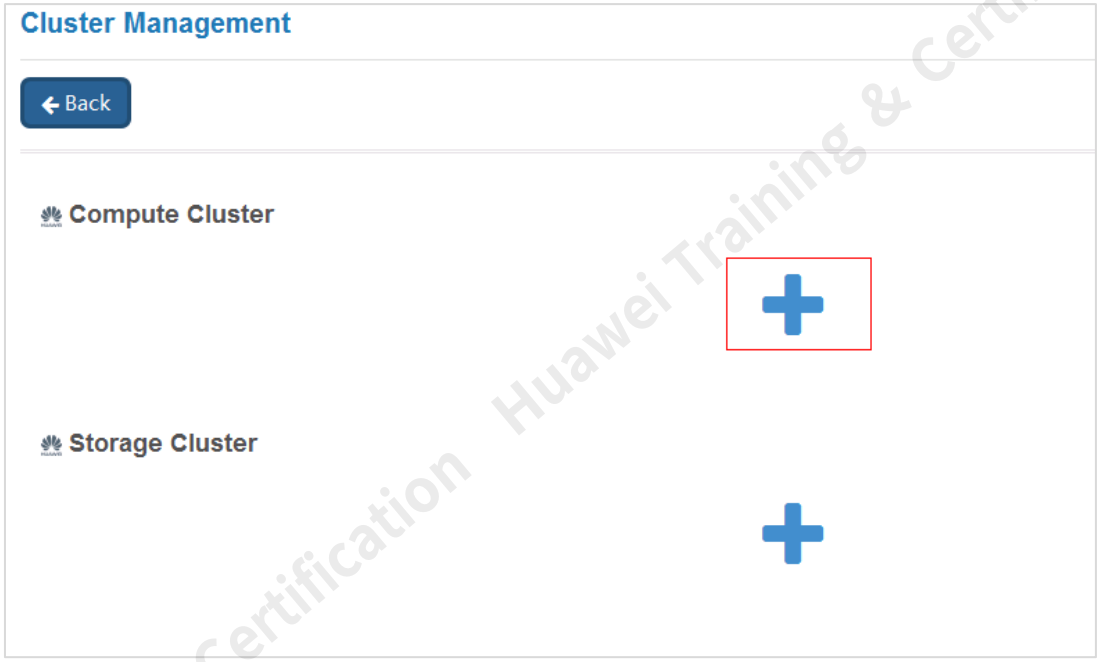

Configure the mapping between FusionSphere OpenStack hosts and FusionCompute clusters. Map Cluster1, Cluster2 and Cluster3 to the same host. Set **vCPU Overcommitment Ratio** to **3**.

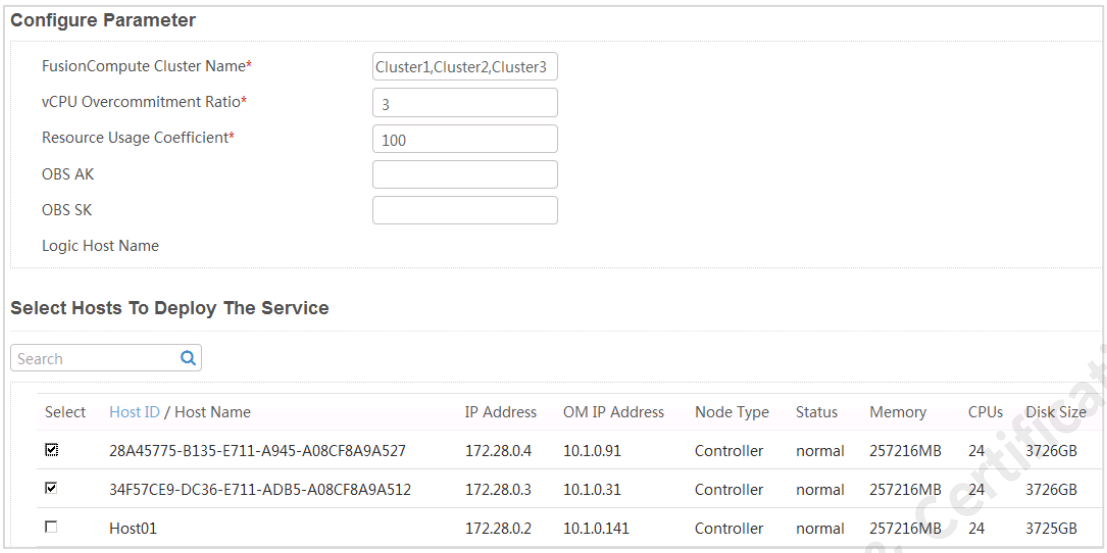

The fc-nova-compute001 host is successfully created.

6.3.2 Creating a Storage Cluster

Click +.

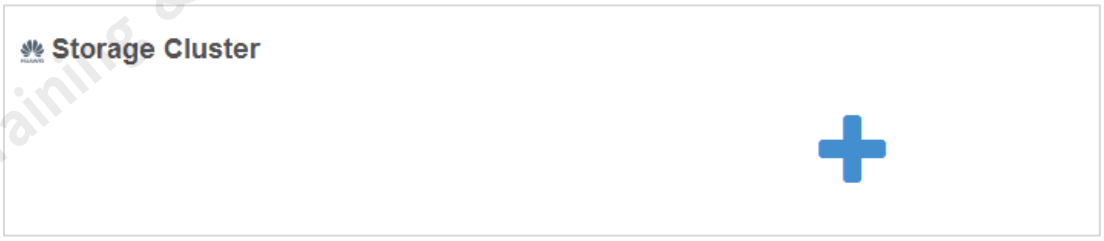

Enter the AZ name and configure the backend storage.

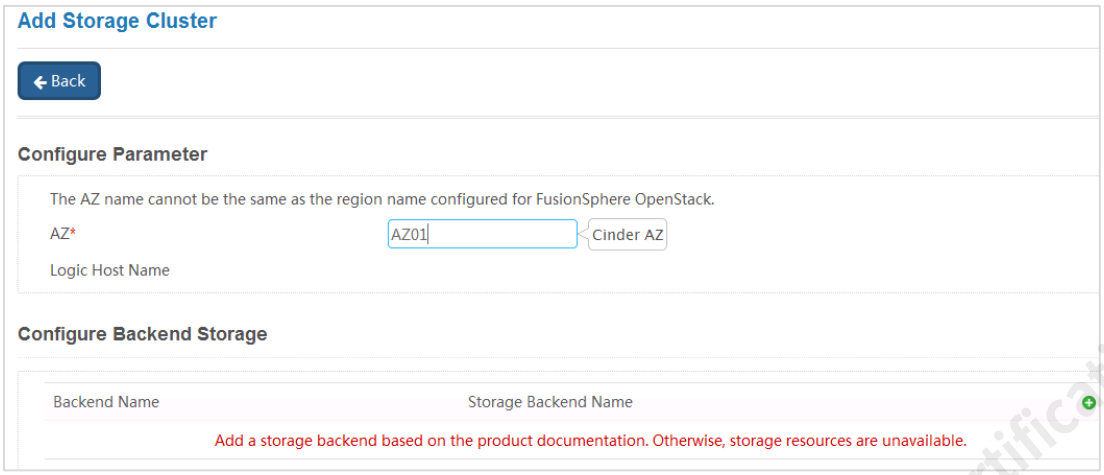

Enter the name of the Cinder backend storage and that of the data store on FusionCompute.

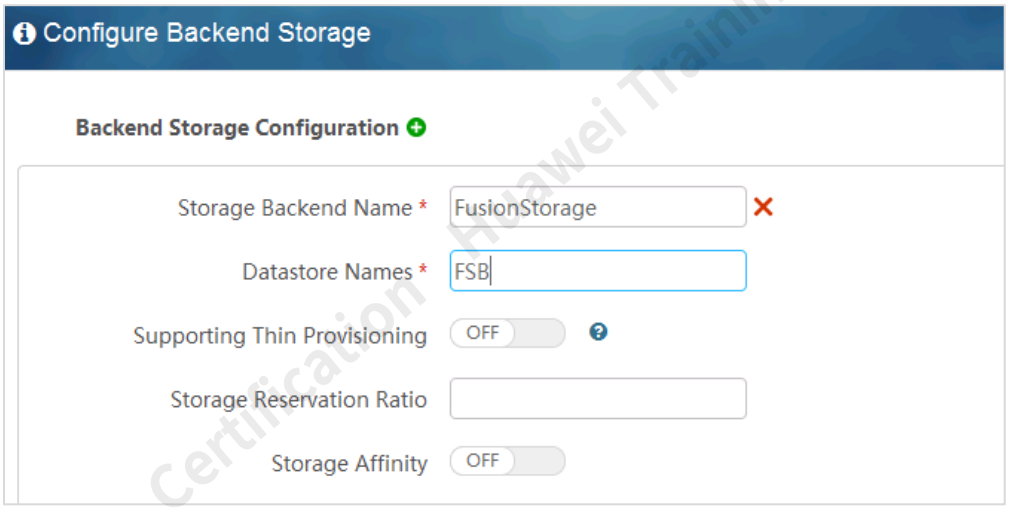

Select all hosts and click **Submit**.

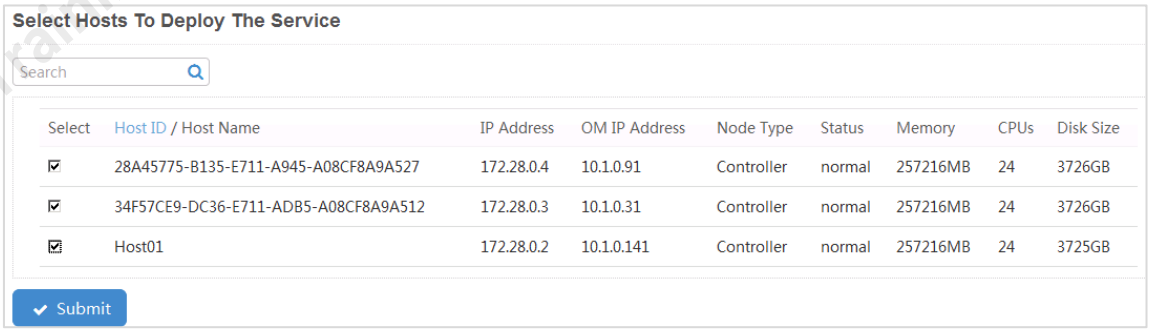

The Cinder-vrm001 is successfully created.

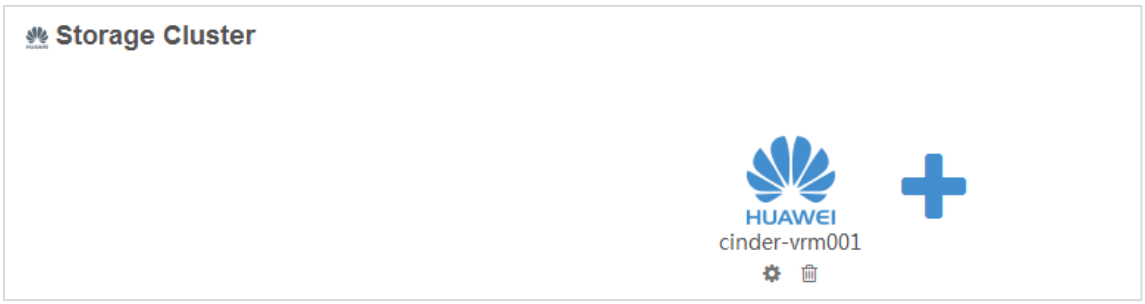

6.4 Configuring the FusinSphere OpenStack OM

6.4.1 Creating an External Network

The external network of the cloud platform needs to be consistent with the external gateway of the Agile Controller with the same name. In this example, two external networks are created. One is used for public network access (elastic IP/SNAT address acquisition), and the other is used for tenant router application (router address acquisition), which correspond to the Internet and Intranet respectively.

Step 1 **Create an external network.**

Choose Resources > Networks > External Networks, and click Create.

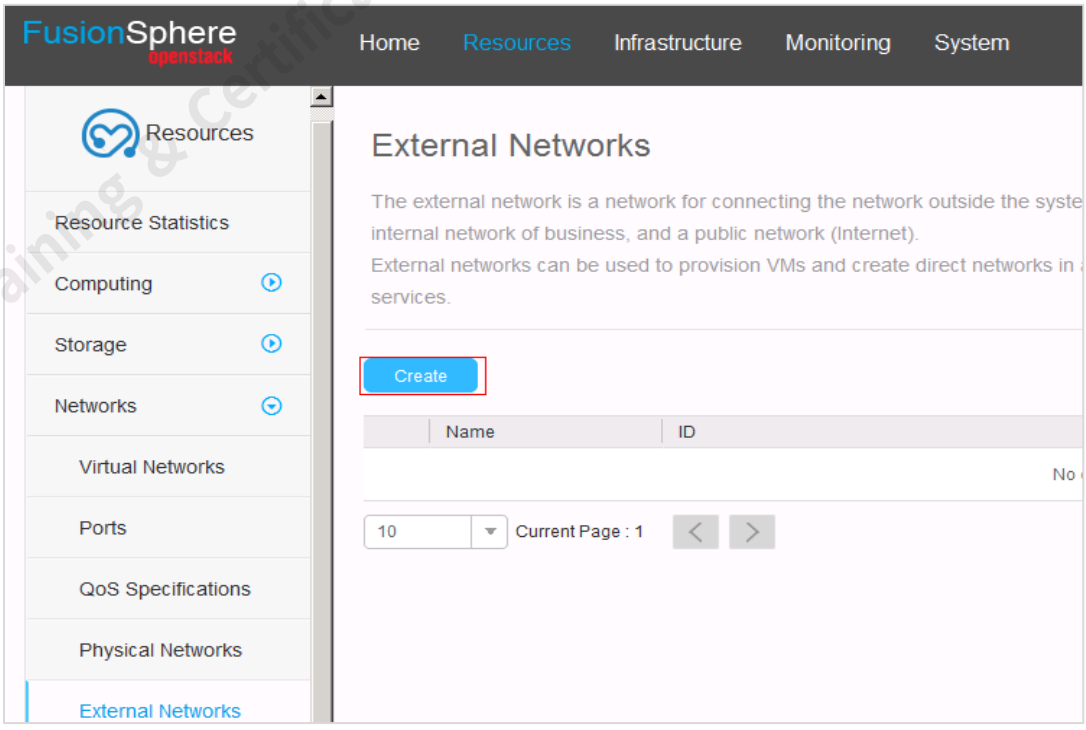

Enter the network name, select the network type and physical network, enter the VLAN ID, and click **Create**.

In the SDN scenario, the network name must be the same as that of the external network created on the Agile Controller. Ensure that the VLAN is not within the available range of the VLAN of the Agile Controller, that is, the available range for configuring the VNI/VLAN/Bridge-domain.

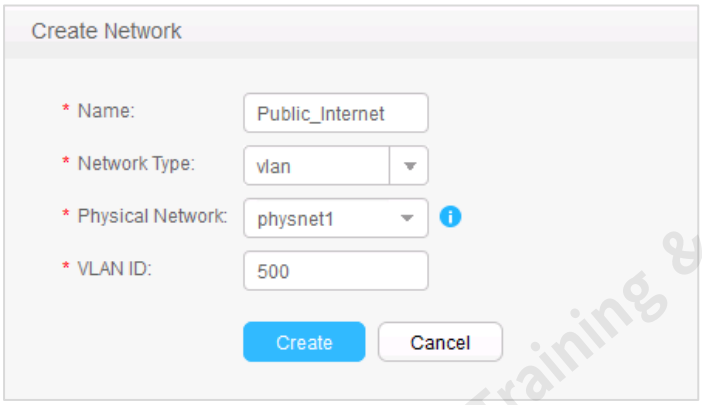

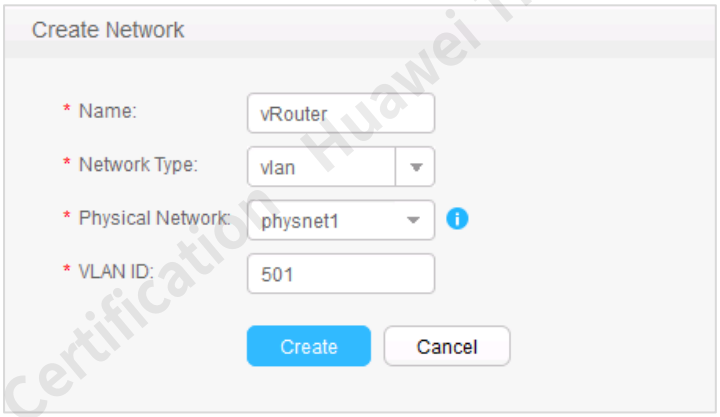

Step 2 **Create a subnet.**

Locate the row that contains the external network, click **Modify More**, and then click **Create Subnet**.

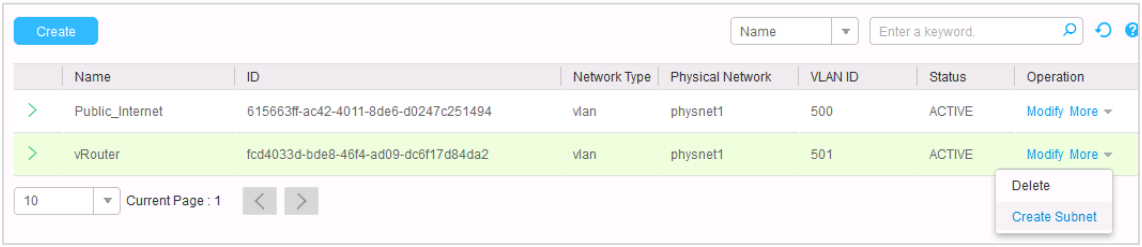

Create the Public_Internet network and vRouter network for the SNAT/elastic IP address and router to obtain IP addresses.

Create a public network subnet.

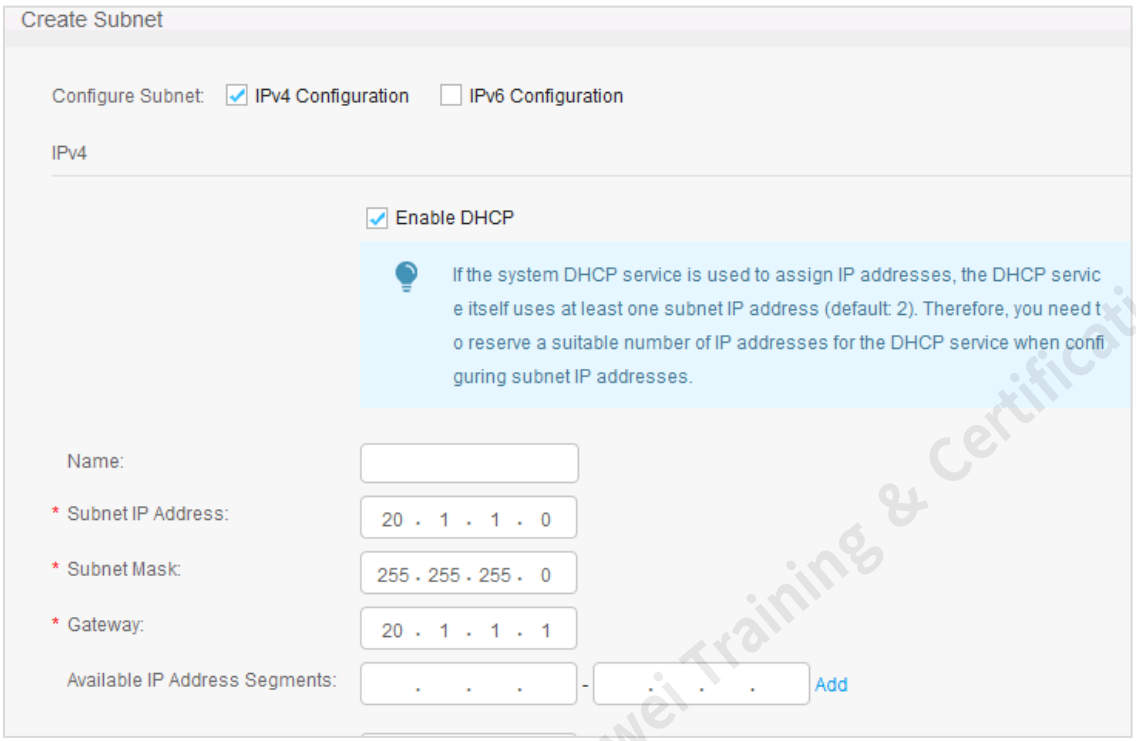

This subnet is mainly used for elastic IP addresses and SNAT addresses of service VMs.

Create a vRouter subnet.

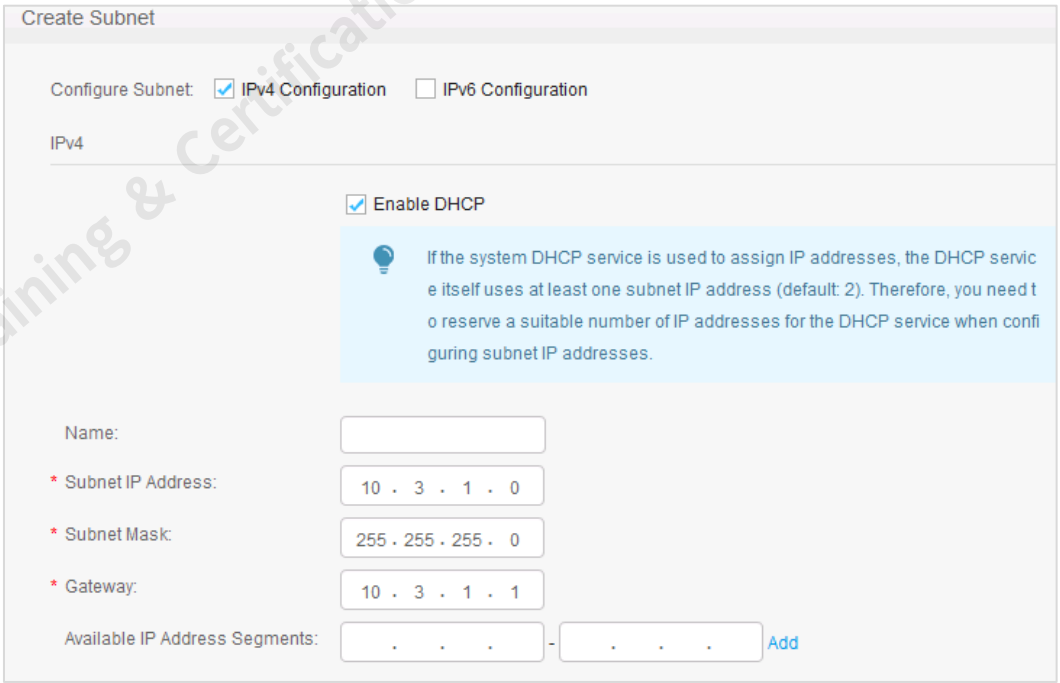

This subnet is used to apply for router addresses.

Step 3 **Check whether the configuration is delivered properly.**

Log in to the Agile Controller to check the delivered configurations.

Choose **System** > **Logs** > **Operation Log**. The configuration is successfully delivered.

Time	Severity	VM/Host	Service/Module	Ten...	Oper...	Operation	Operation O...	Terminal	Result	AdditionInfo	Link
2018-08-21 1...	Infor...	controller-192-...	neutron	NAAS	fsp@...	Modify port	Port(Status R...	192.16...	Succ...	Update a port. (Port id ...	
2018-08-21 1...	Infor...	controller-192-...	neutron	NAAS	fsp@...	Modify port	Port	192.16...	Succ...	Update a port. (Port id ...	
2018-08-21 1...	Infor...	controller-192-...	neutron	NAAS	fsp@...	Modify port	Port(Status R...	192.16...	Succ...	Update a port. (Port id ...	
2018-08-21 1...	Infor...	controller-192-...	neutron	NAAS	fsp@...	Modify port	Port	192.16...	Succ...	Update a port. (Port id ...	
2018-08-21 1...	Infor...	controller-192-...	neutron	NAAS	fsp@...	Modify port	Port(Status R...	192.16...	Succ...	Update a port. (Port id ...	
2018-08-21 1...	Infor...	controller-192-...	neutron	NAAS	fsp@...	Modify port	Port	192.16...	Succ...	Update a port. (Port id ...	
2018-08-21 1...	Infor...	controller-192-...	neutron	NAAS	fsp@...	Create port	Port(Status R...	192.16...	Succ...	Create a port. (Port id ...	

Config Record

Enter the device name or device IP | Please select result

Configuration Time	Device Name	Device IP	Operation Description	Message Body	Result
2018-08-21 15:55:46	Leaf3	172.21.22.13	create BD 5001.	View	Succeed
2018-08-21 15:55:46	Leaf3	172.21.22.13	bind vni 10003 to BD 5001.	View	Succeed
2018-08-21 15:55:46	Leaf3	172.21.22.13	create evpn instance 7:100...	View	Succeed
2018-08-21 15:55:46	Leaf3	172.21.22.13	create interface 10GE1/0/3...	View	Succeed
2018-08-21 15:55:46	Leaf3	172.21.22.13	set if 10GE1/0/39.201 acce...	View	Succeed
2018-08-21 15:55:46	Leaf3	172.21.22.13	bind bd 5001 to if 10GE1/0/...	View	Succeed
2018-08-21 15:55:46	Leaf3	172.21.22.13	create interface 10GE1/0/4...	View	Succeed
2018-08-21 15:55:46	Leaf3	172.21.22.13	set if 10GE1/0/40.201 acce...	View	Succeed
2018-08-21 15:55:46	Leaf3	172.21.22.13	bind bd 5001 to if 10GE1/0/...	View	Succeed
2018-08-21 15:55:46	Leaf3	172.21.22.13	set vni 10003.	View	Succeed

20 Total records: 10 1

Log in to the Leaf-3 node and check the delivered configuration.

```
<Leaf-3>display configuration commit changes last 1
Building configuration
  #
+ bridge-domain 5001
+  vxlan vni 10003
+  evpn
+   route-distinguisher 2:10003
+   vpn-target 0:10003 export-extcommunity
+   vpn-target 0:10003 import-extcommunity
  #
+ interface 10GE1/0/45.2 mode l2
+  encapsulation dot1q vid 501
+  bridge-domain 5001
  #
+ interface 10GE1/0/46.2 mode l2
+  encapsulation dot1q vid 501
+  bridge-domain 5001
  #
  interface Nve1
+  vni 10003 head-end peer-list protocol bgp
  #
<Leaf-3>
```

10GE1/0/45 and 10GE1/0/46 are interfaces for connecting to the FSP server.

----End

6.4.2 Creating a Host Group

Step 1 **Login**

FusionSphere OpenStack OM is deployed on VMs of the cloud services.

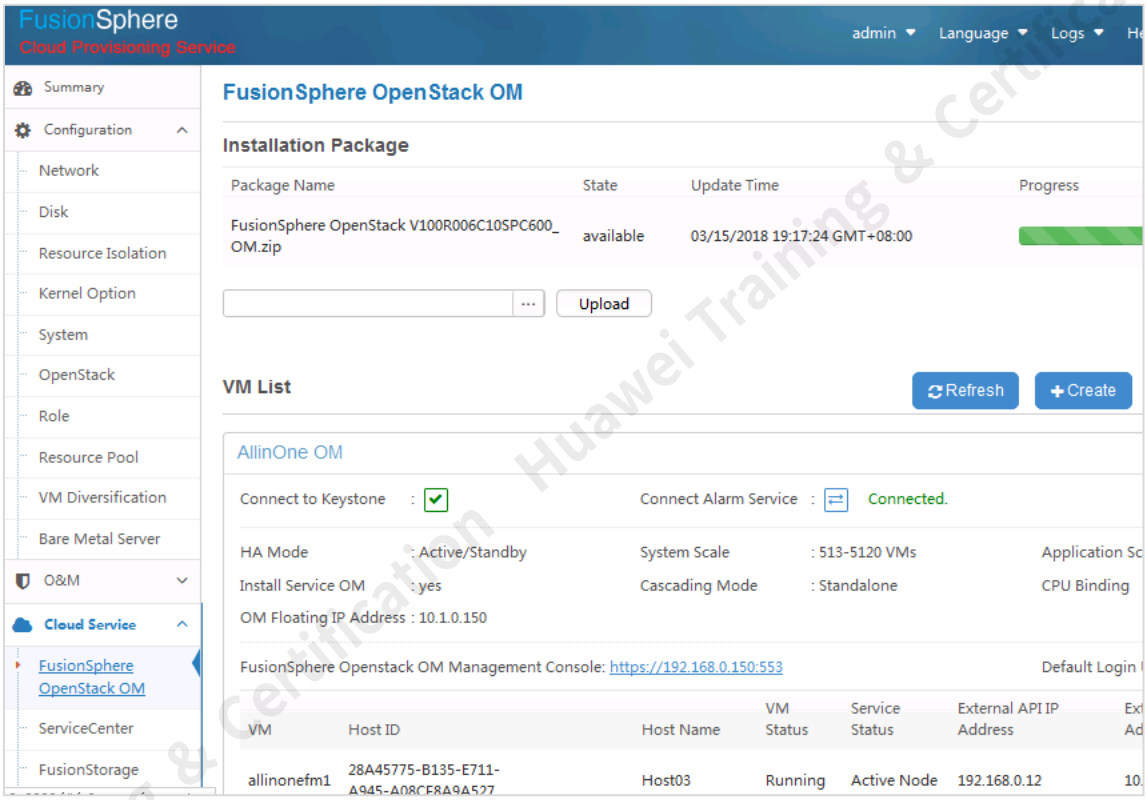

Enter the username **account cloud_admin** and password **FusionSphere123**.

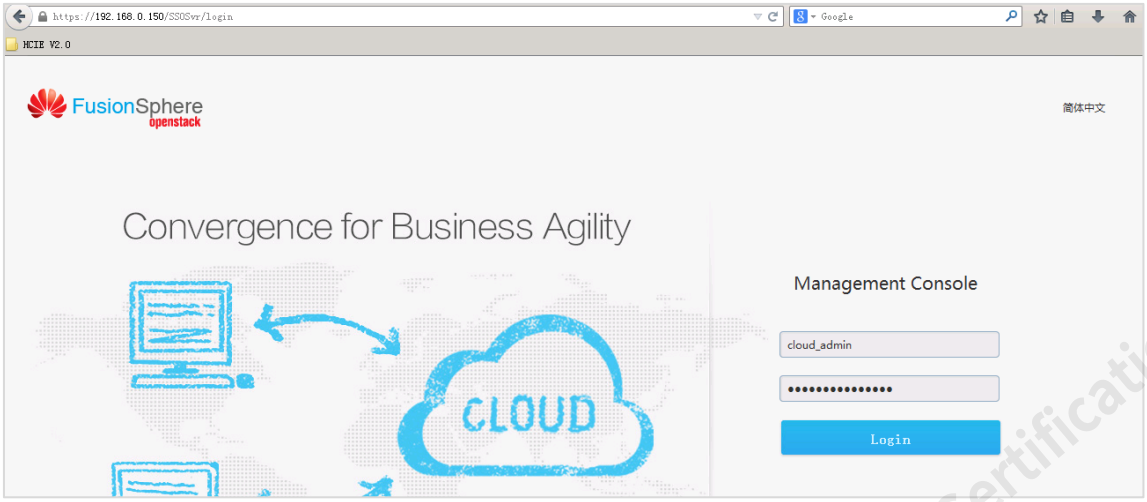

Step 2 **Create a host group.**

Choose Resources > Host Groups.

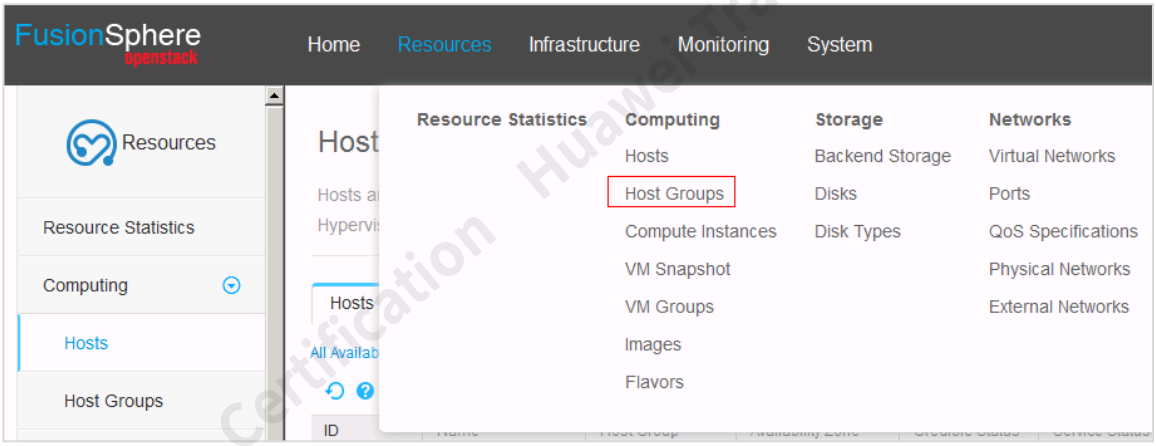

Click **Create**, enter the host group name **FC**, and set **Availability Zone** to **AZ01**.

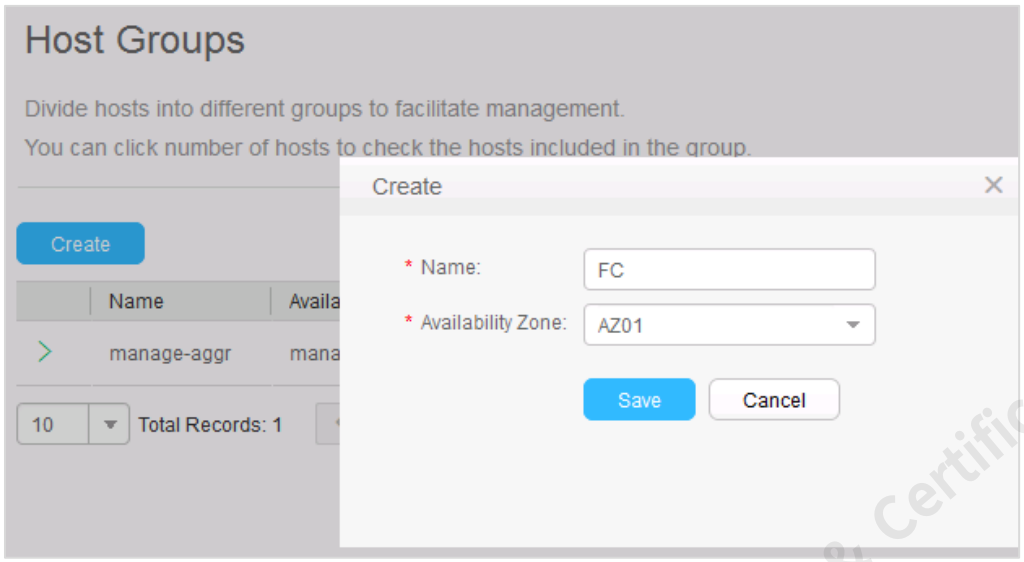

AZ01 is the name of the storage resource used by the cloud platform to connect to the FusionCompute resource pool.

Step 3 **Adding a host.**

Click Host Management.

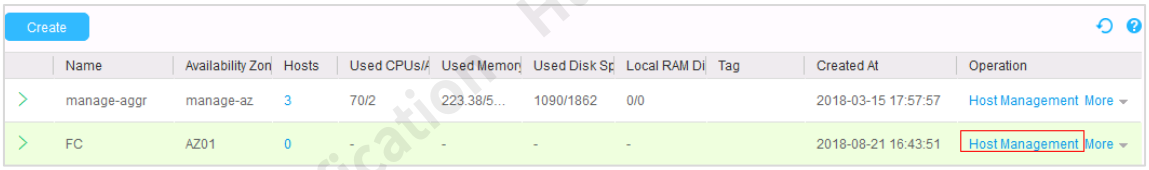

Add **fc-nova-compute001** to the host group and click **Save**.

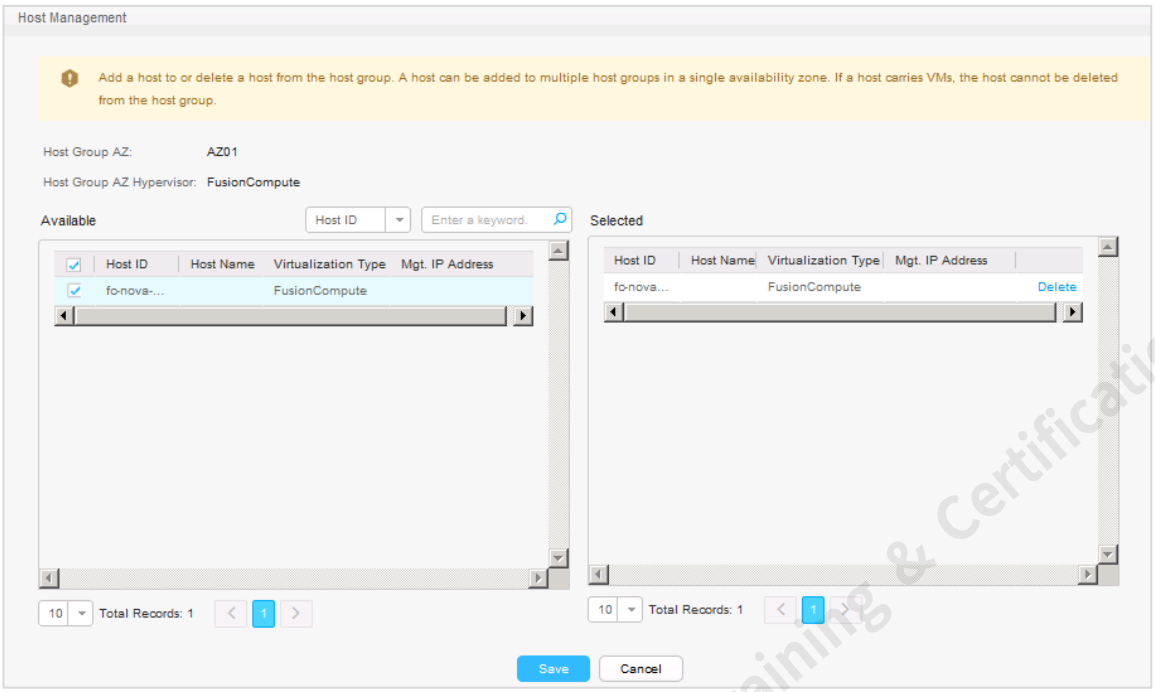

Fc-nova-compute001 is the computing resource used by the cloud platform to connect to FusionCompute.

----End

6.4.3 Registering a VM Image (Configured)

Choose Resources > Computing > Images, and click Register.

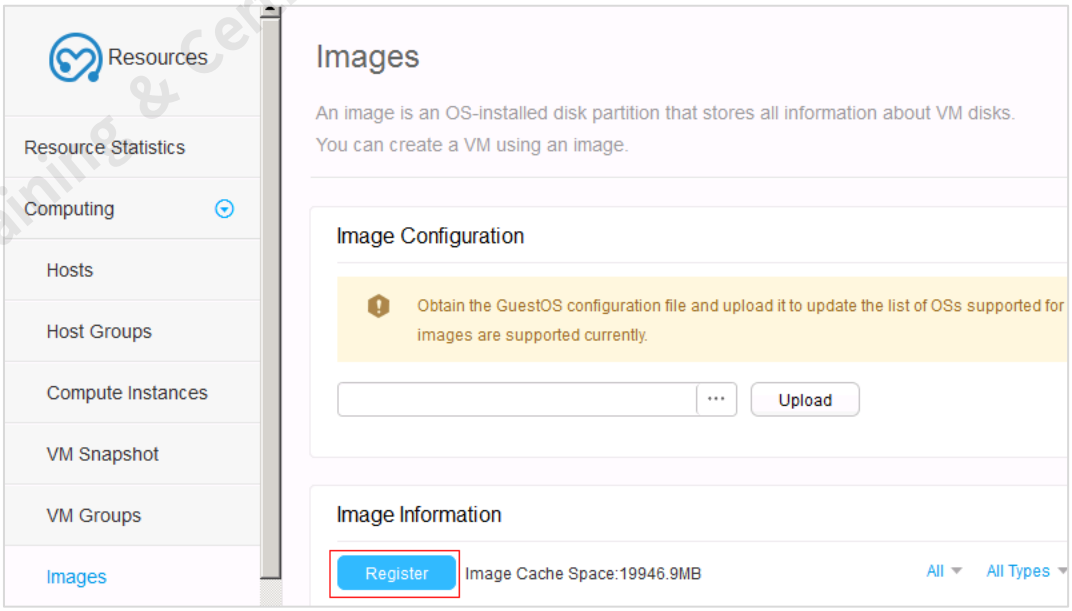

Enter the image information and click **Register**.

a 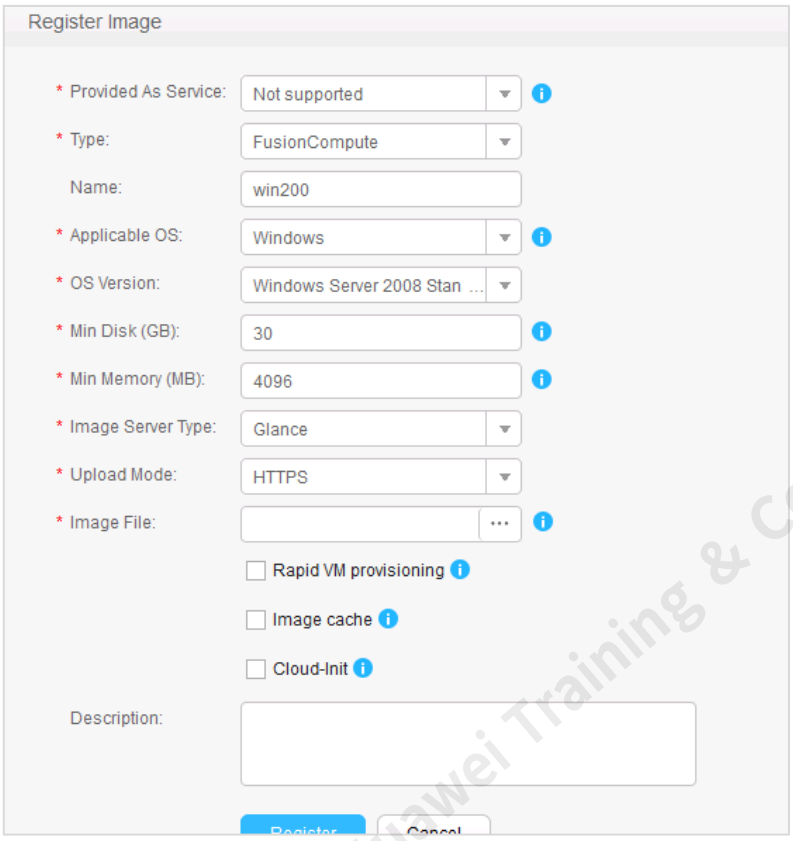

A VM image is successfully registered.

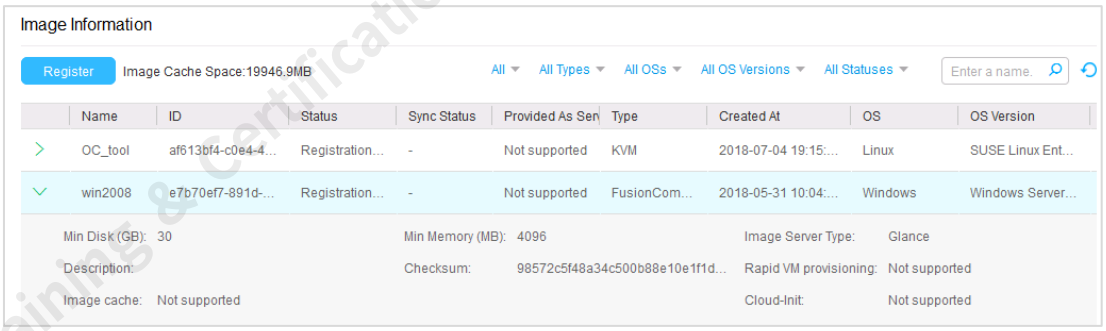

6.4.4 Creating the VM Flavors (Configured)

Choose Resources > Computing > Flavors.

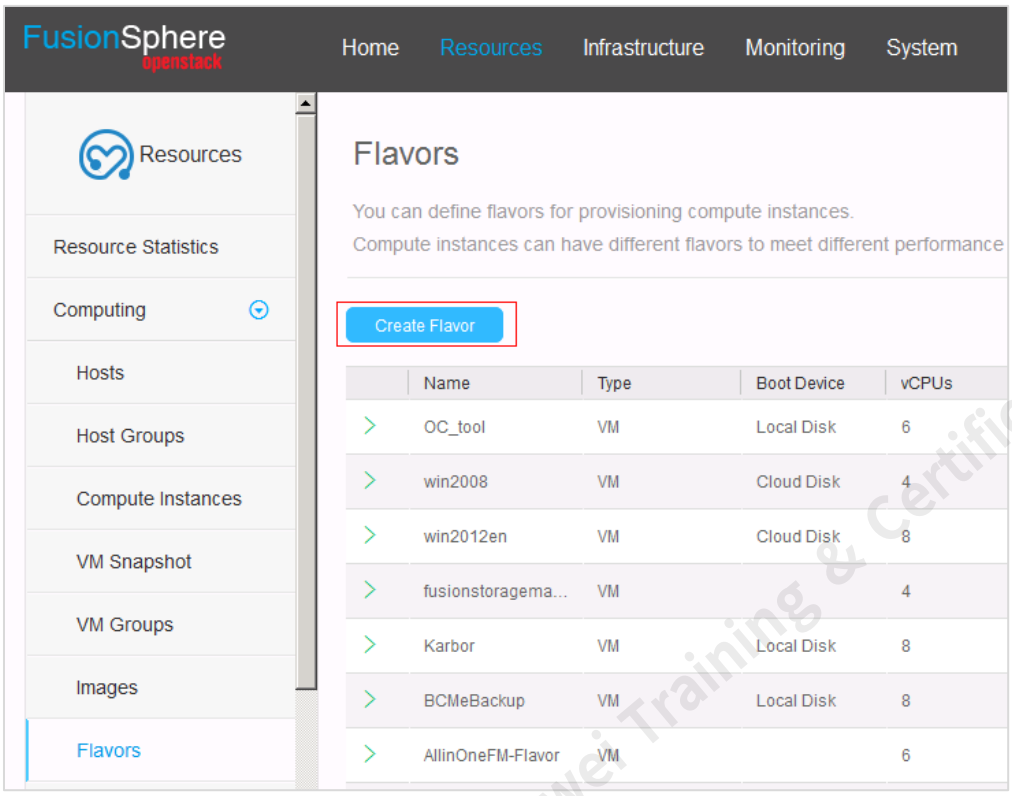

Click **Create Flavor**, and enter the VM flavors in the dialog box that is displayed, click **OK**.

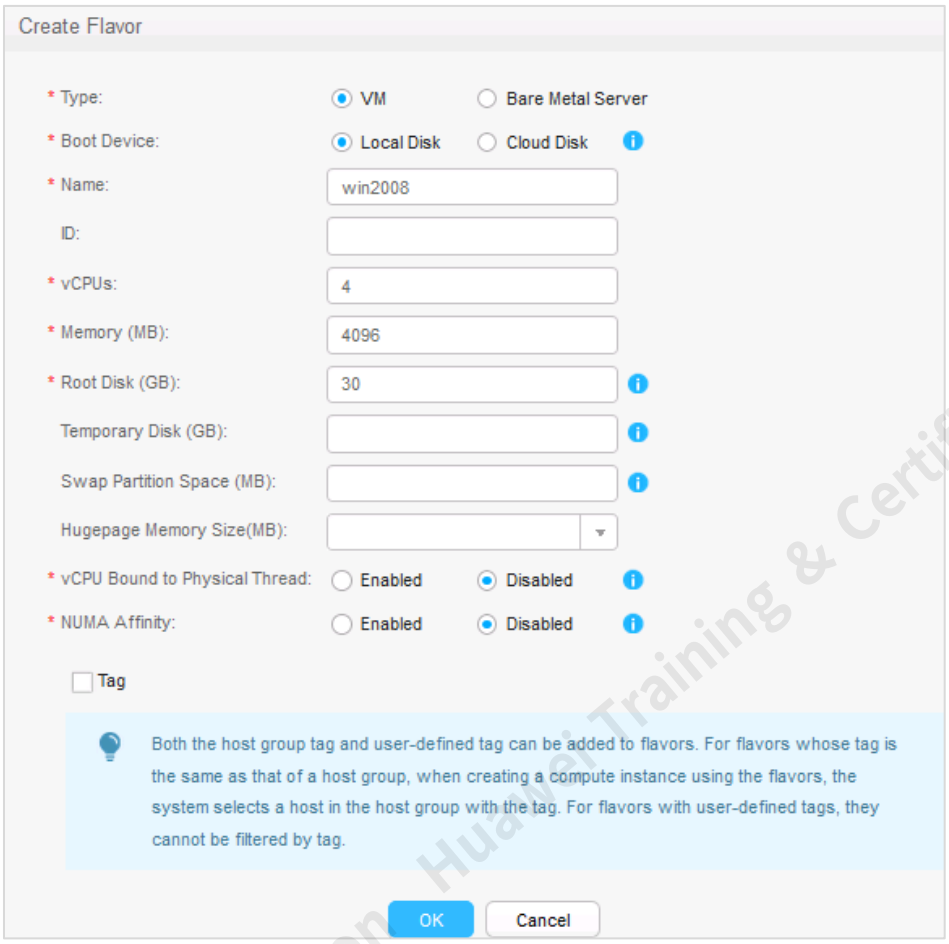

The VM flavor is successfully created.

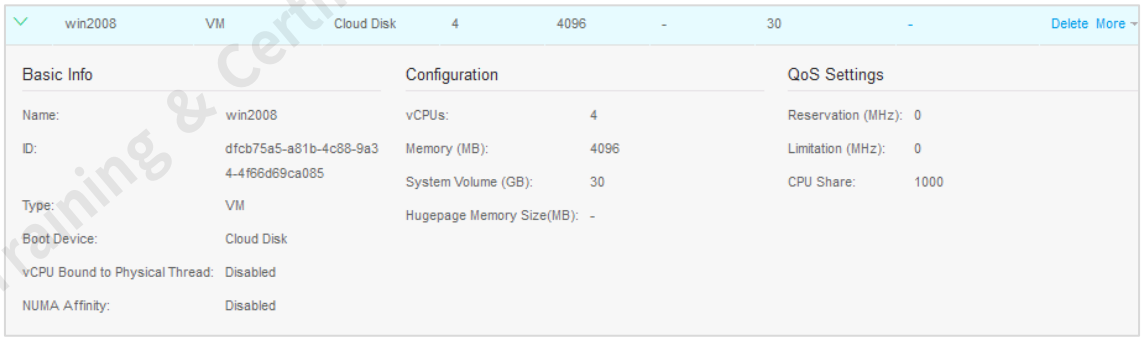

7 SDN Cloud-Network Synergy Service Operation

7.1 Configuring ManageOne ServiceCenter

This section describes how to configure cloud computing pools, apply for cloud hosts, and commission cloud host networks.

7.1.1 Configuring a Cloud Resource Pool

Step 1 **Log in to the ManageOne ServiceCenter.**

Choose Cloud Service > ServiceCenter, and click ServiceCenter Page Link.

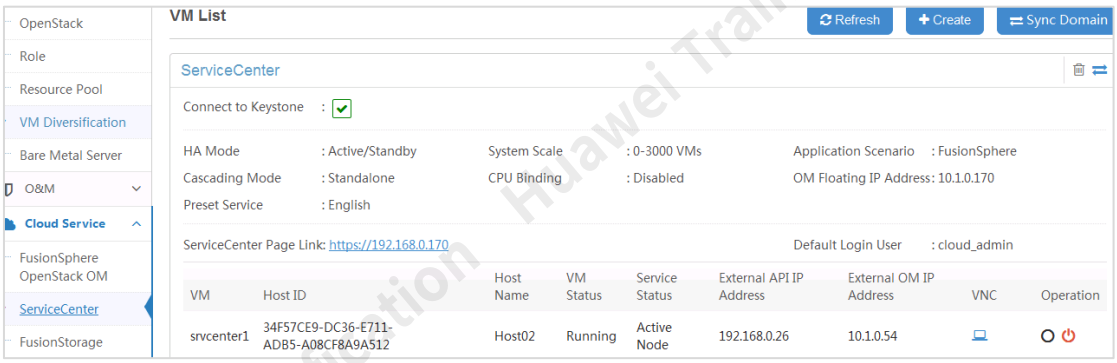

Enter the username **account cloud_admin** and password **FusionSphere123**.

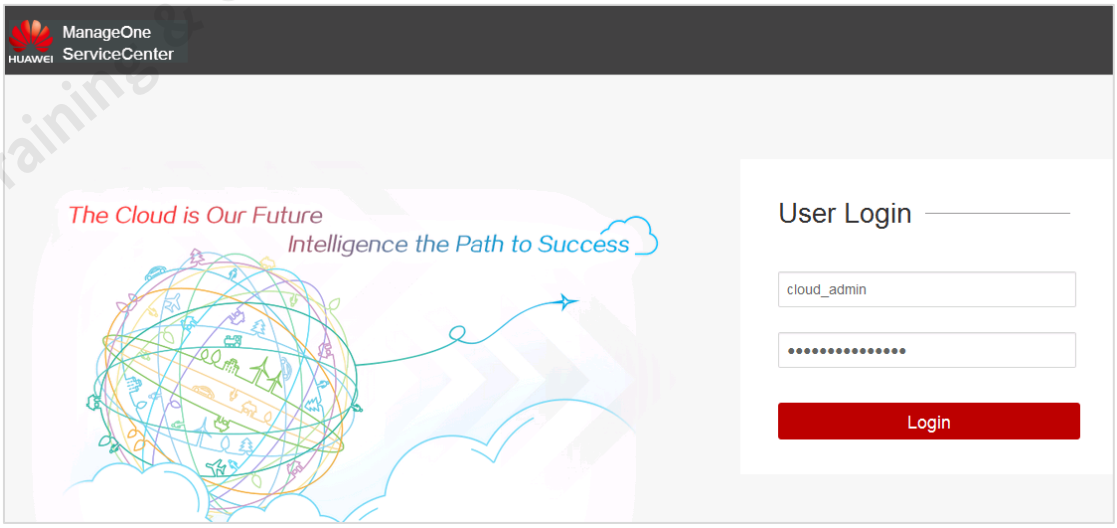

Step 2 **Configure a cloud resource pool.**

Locate the row that contains the added resource pool, click **More**, and then click **Advanced Settings**.

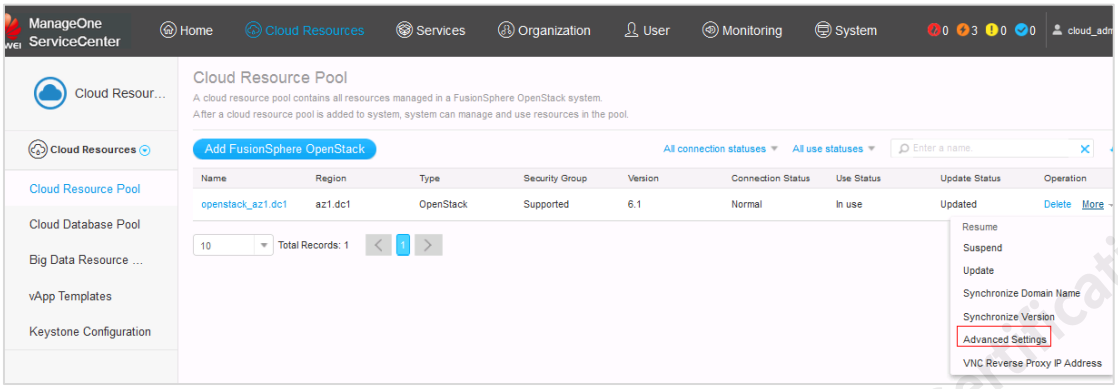

Enable network automation capabilities in the SDN scenario.

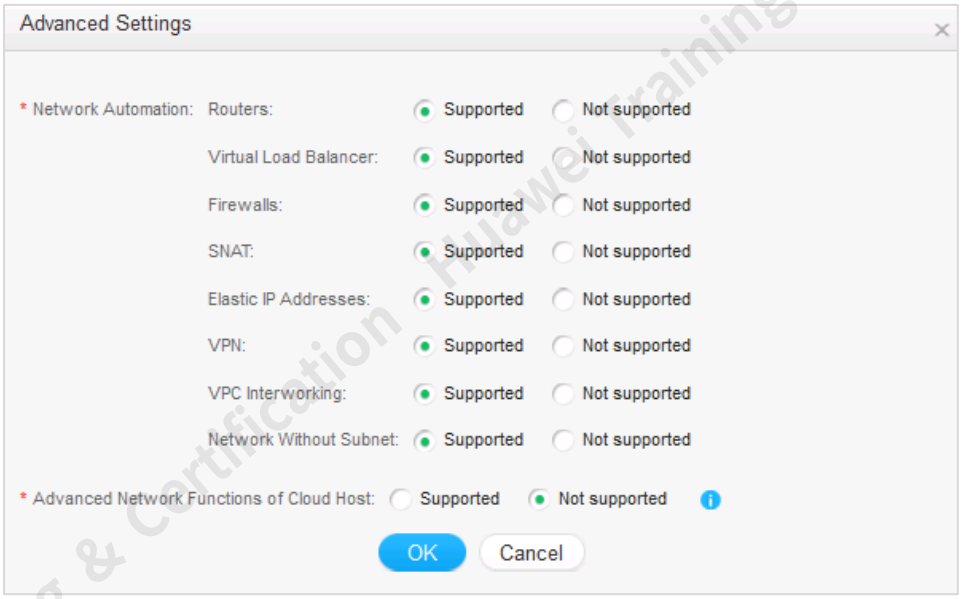

Step 3 **Configure the AZ.**

Select a cloud resource pool in the cloud resource pool list.

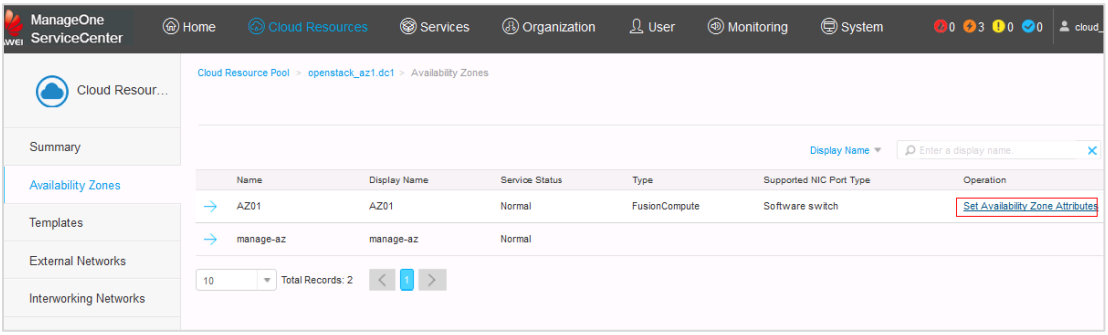

Select **Availability Zones**. Locate the row that contains the target AZ, and click **Set Availability Zone Attributes**.

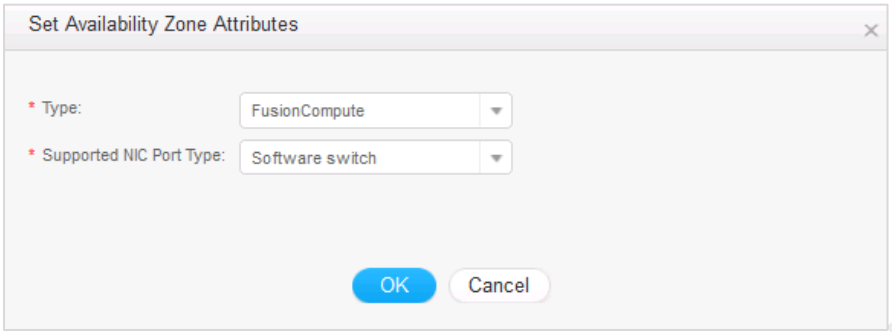

AZ01 is the FusionCompute resource pool. Click **OK**.

Step 4 **Configure an external network.**

Configure the application of the external network on OM.

Click **External Networks** to configure the external network.

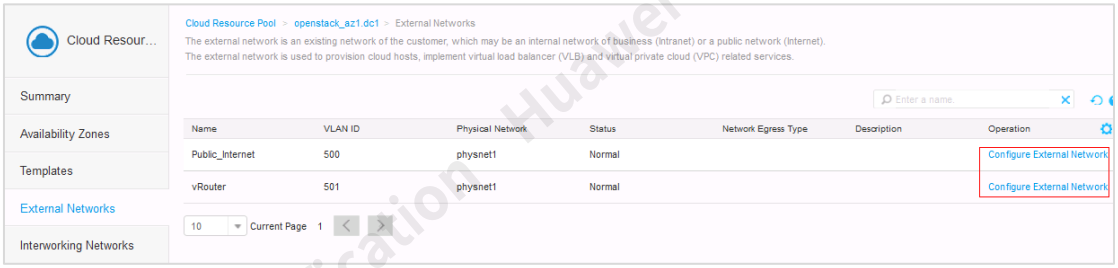

Set the Public_Internet network egress type to **Internet**.

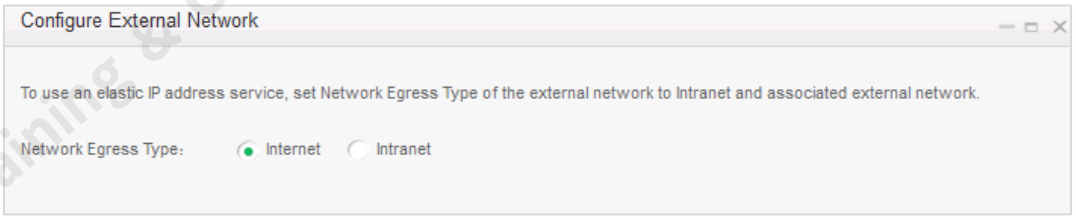

Set the vRouter network egress type to **Intranet**, click **Associate** to associate with the Public_internet network, and then click **OK**.

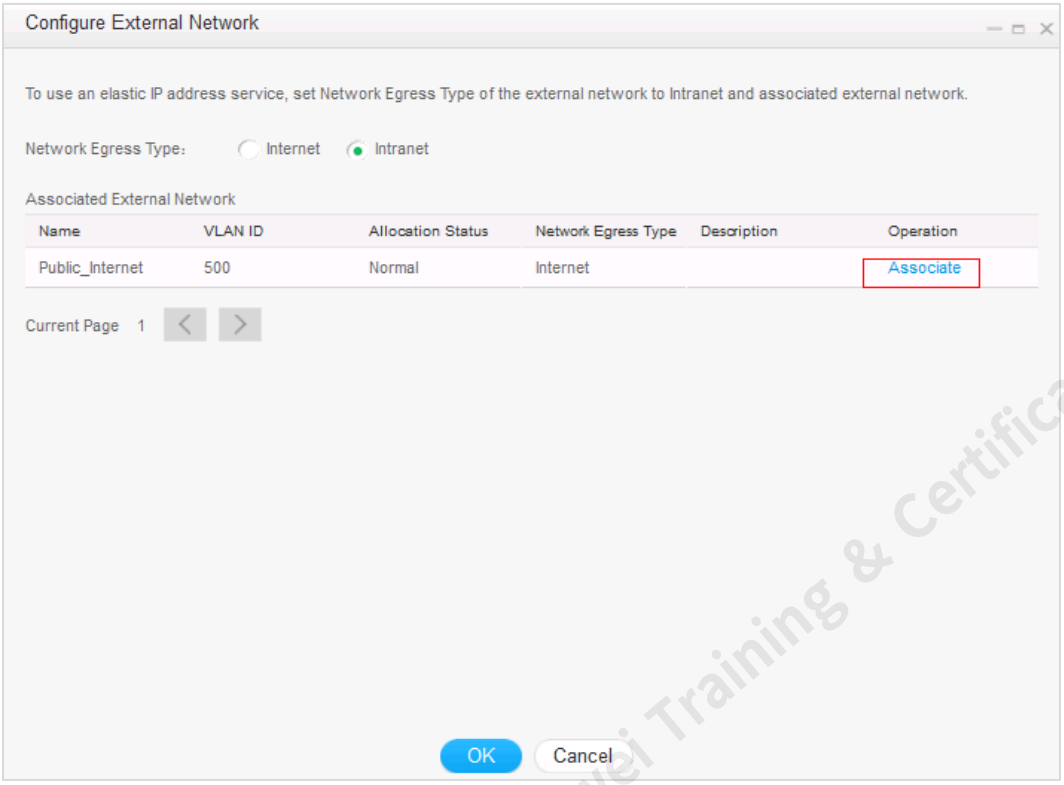

The external network is successfully configured.

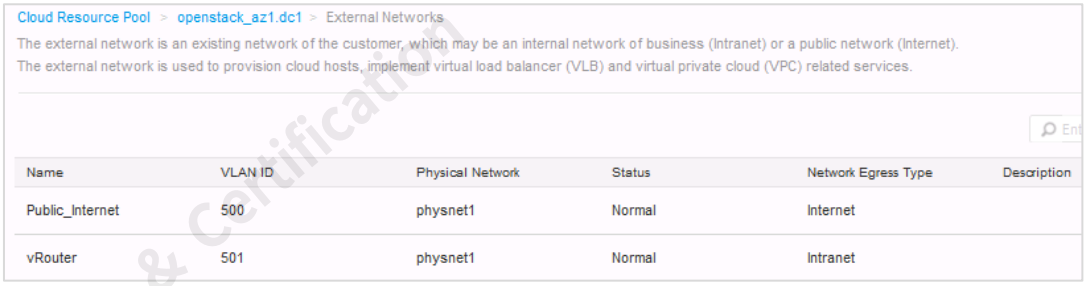

----End

7.1.2 Create an Organization and Its Users

Step 1 **Create an organization.**

Choose Organization > Organization Management, and click Create.

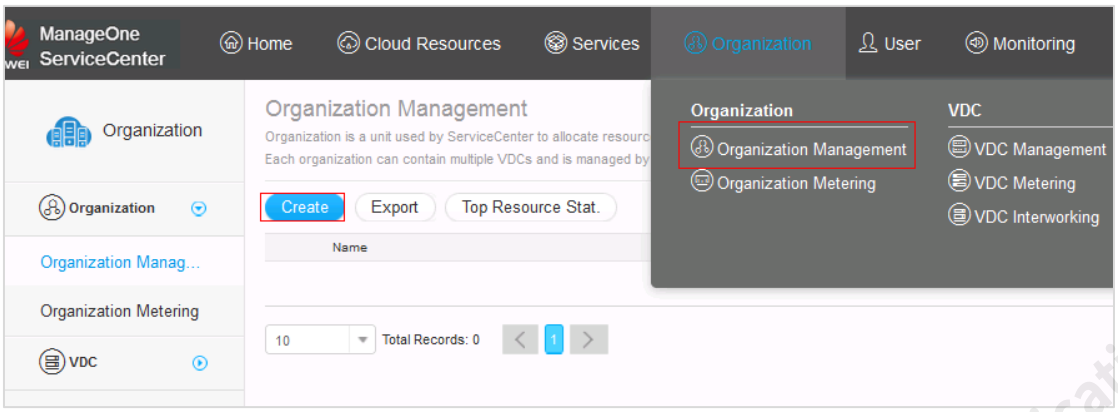

Enter the organization name.

Configure an AZ.

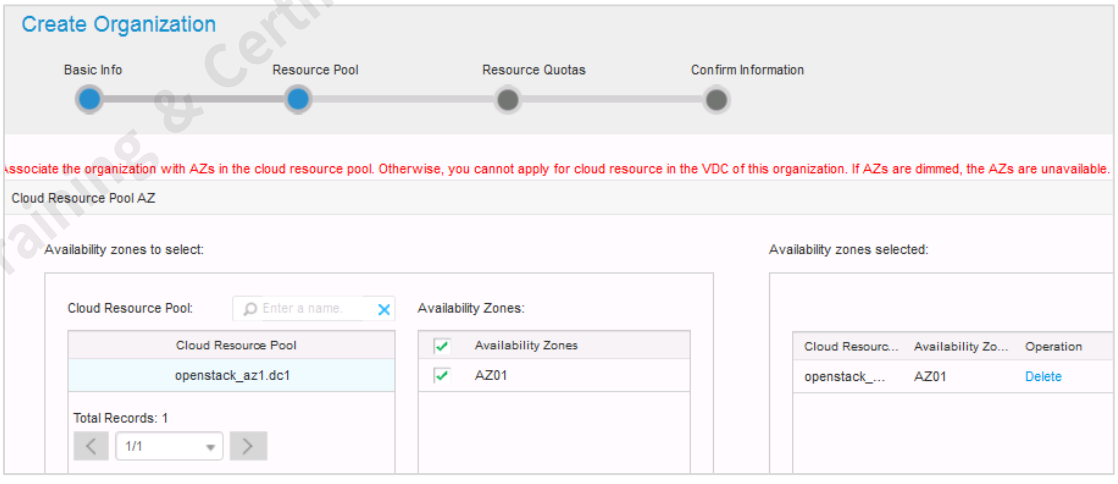

Configure the resource quotas.

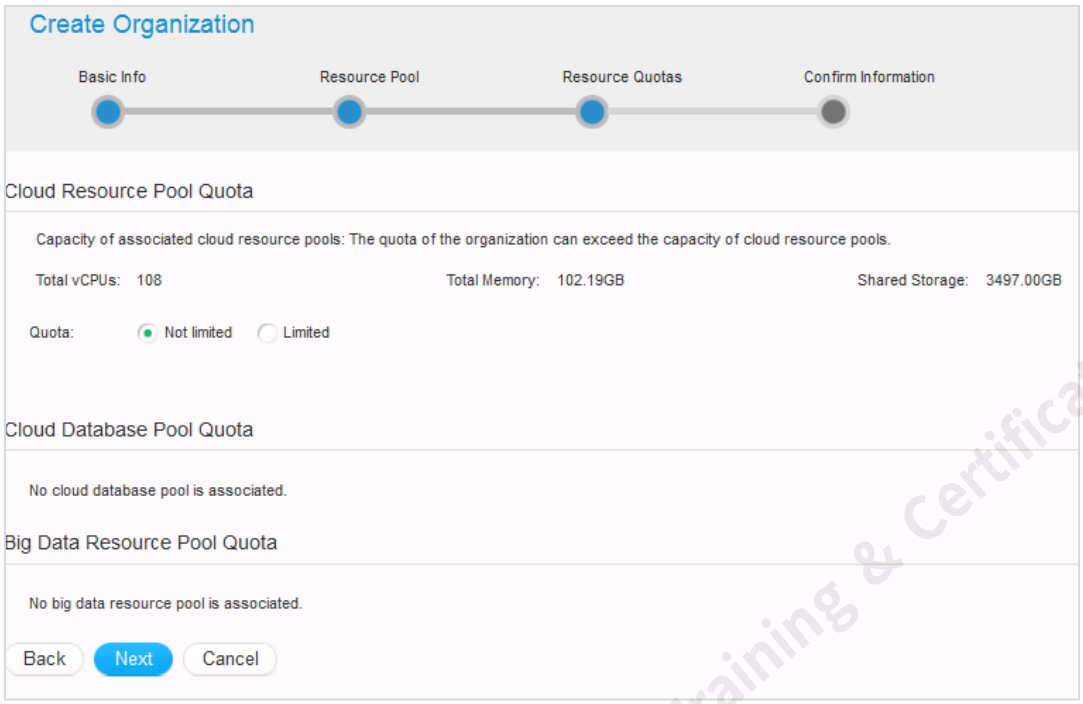

The organization is successfully created.

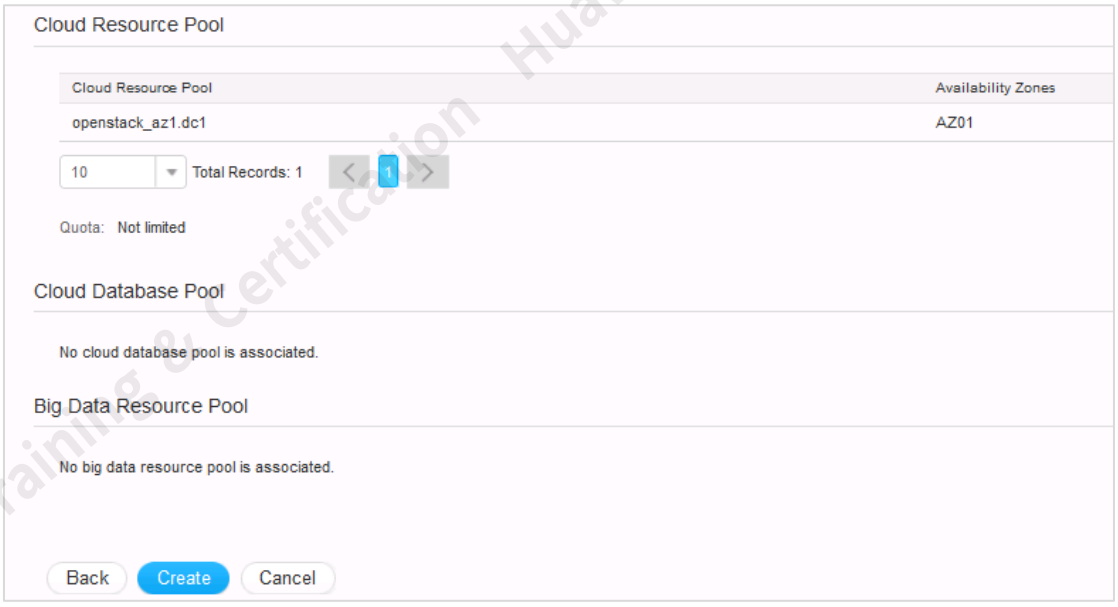

Step 2 **Create an organization administrator.**

Choose User > User management and click Create User.

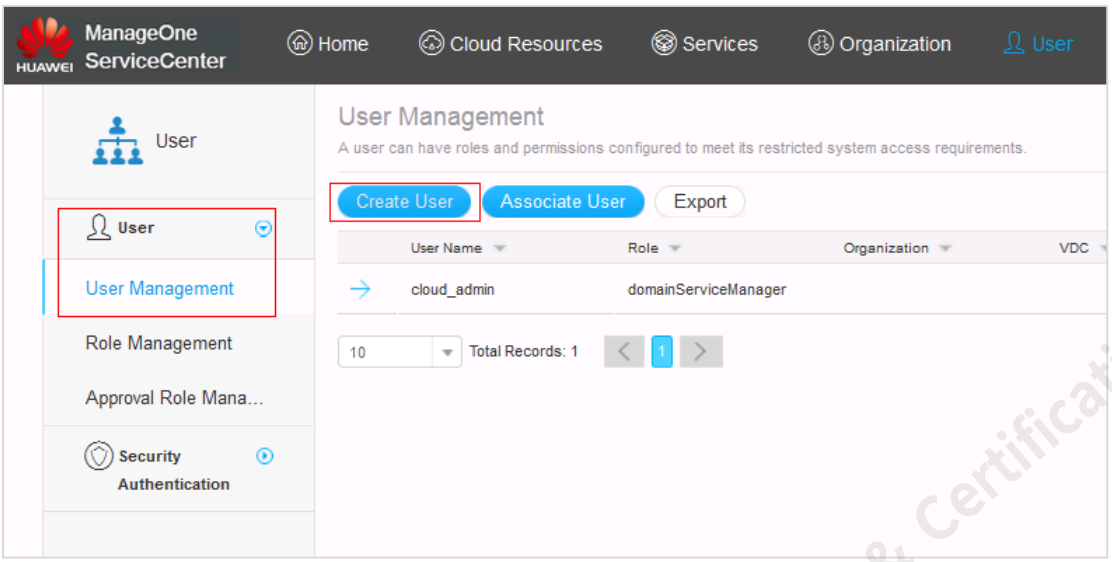

Enter the user name and password, set the **Role** to **orgServiceManager**, and select the created organization.

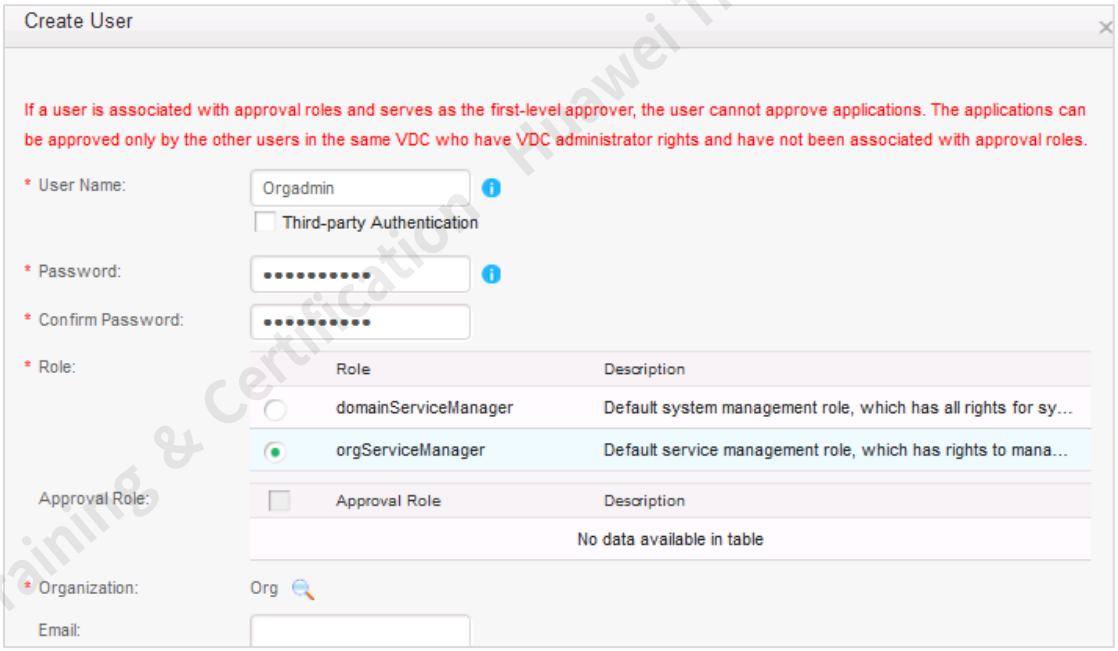

----End

7.1.3 Creating a VDC and Its Users

Step 1 **Create a VDC.**

Choose Organization > VDC > VDC Management and click Create VDC.

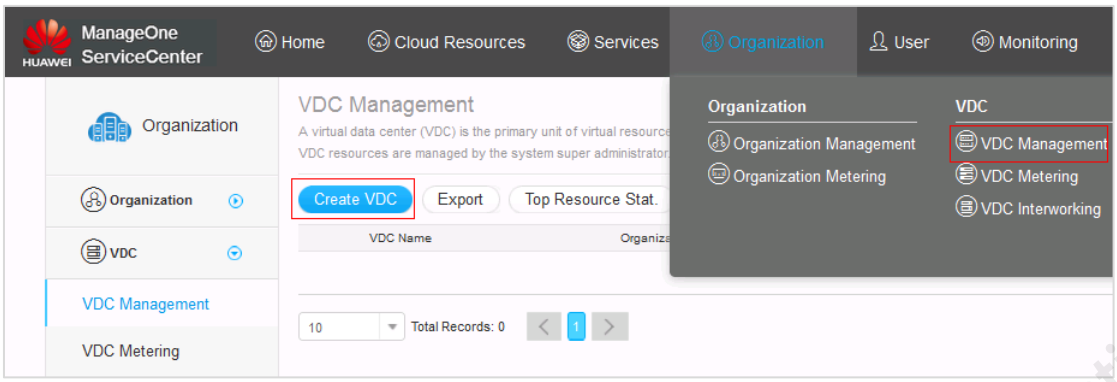

Enter the VDC name, select the organization to which the current organization administrator belongs, and click **Next**.

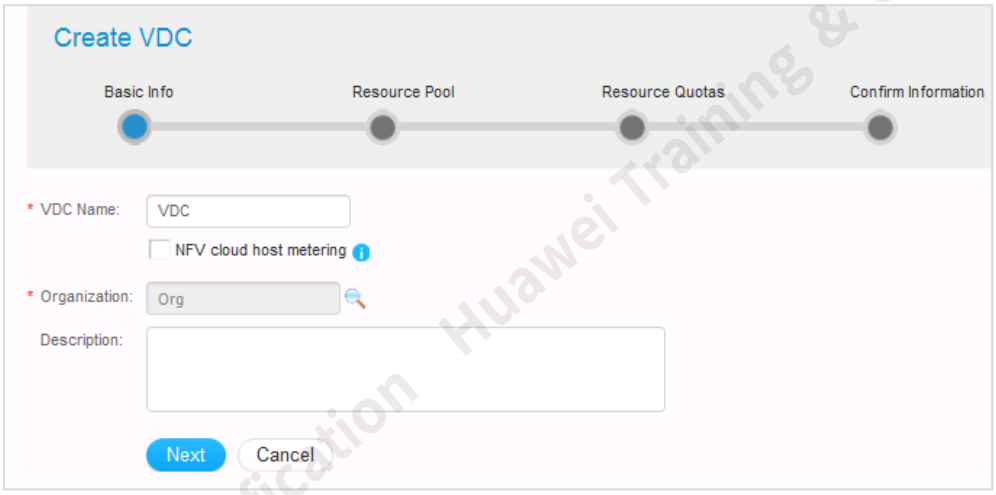

Select a cloud resource pool, select an AZ, and click **Next**.

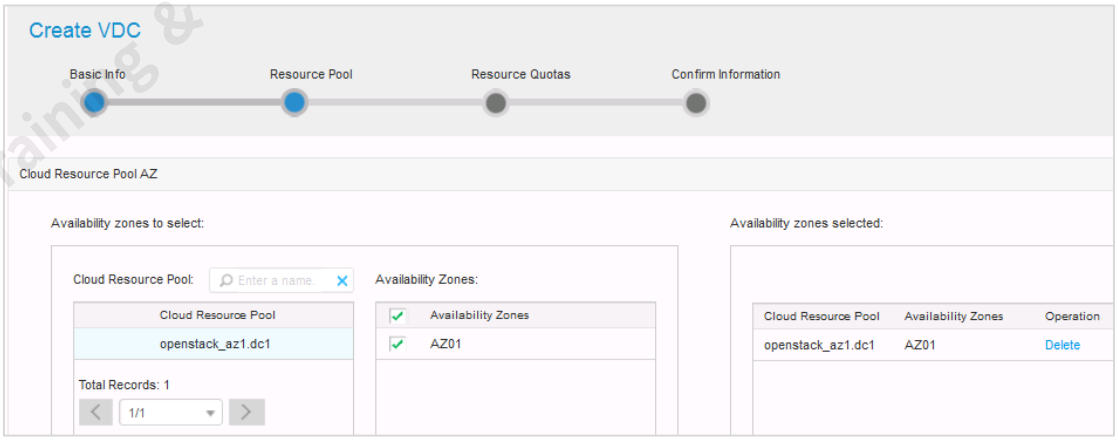

Select the VDC quotas.

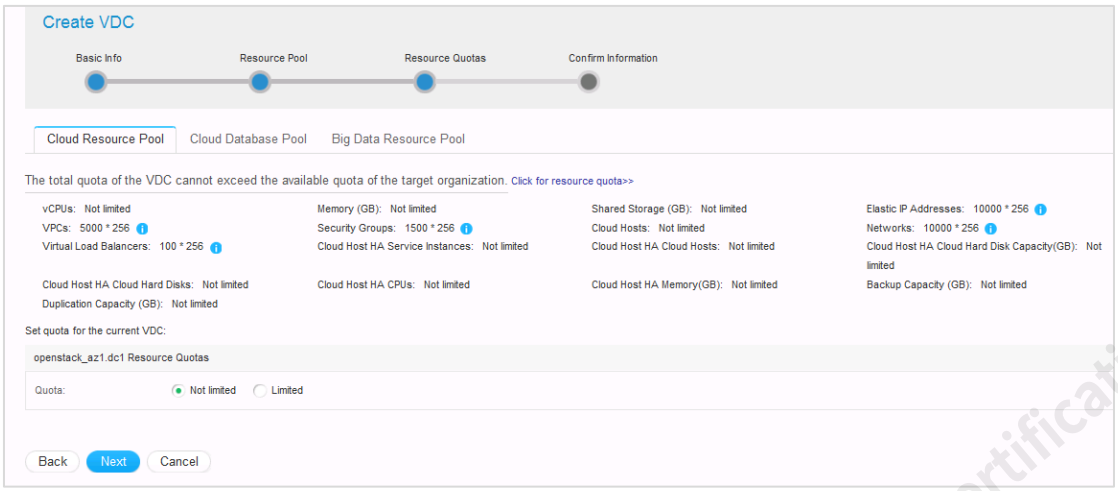

The VDC is successfully created.

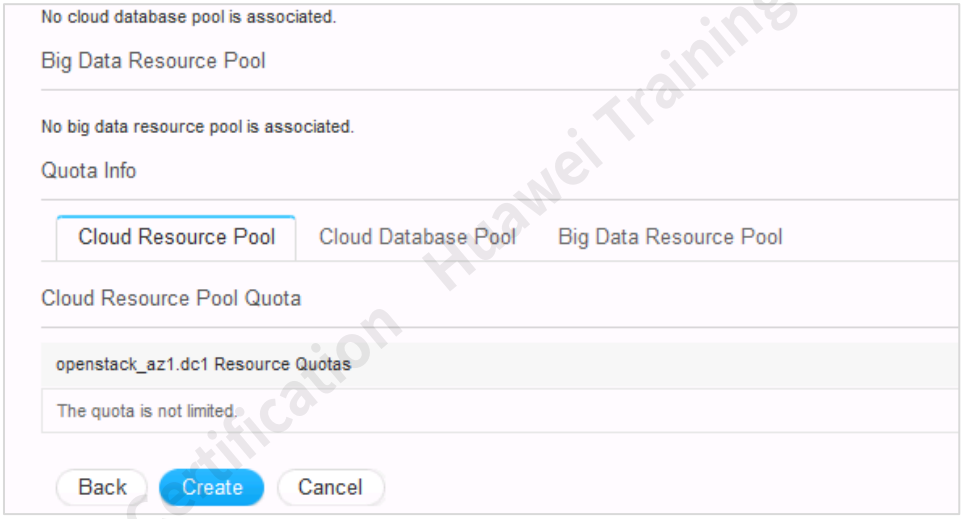

Step 2 **Create a VDC administrator.**

Choose Create VDC > Quota Management, and click Member Management.

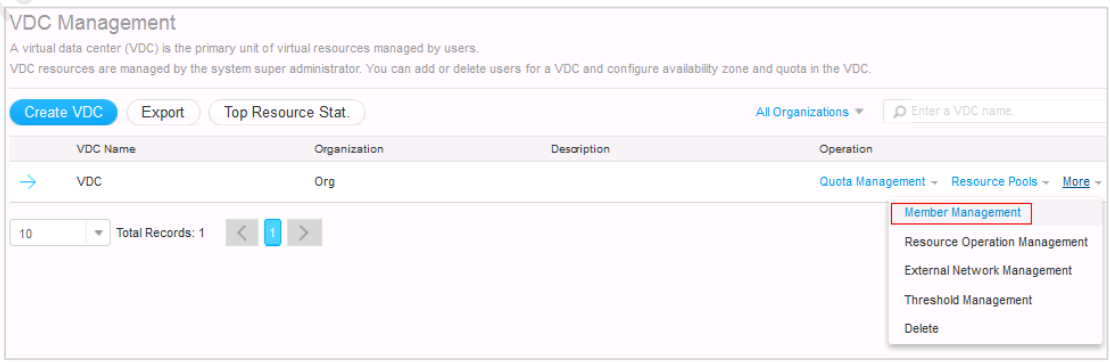

Click Create User.

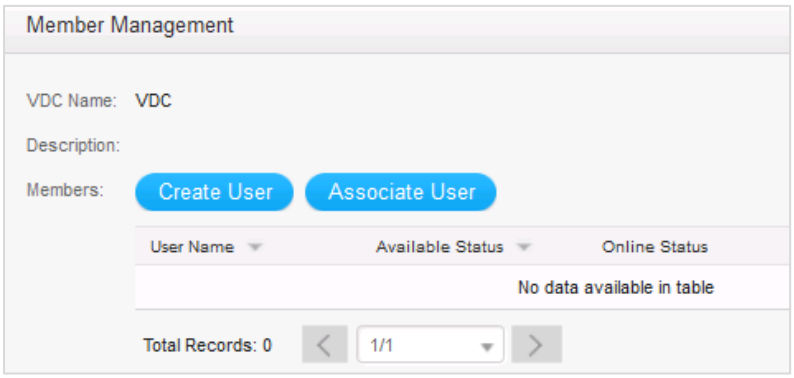

Create a VDC administrator, and enter the username **vdcadmin** and the initial password **Huawei@123**. Change the password to Huawei@1234 upon the first login.

Set the Role to **vdcServiceManager** and click **OK**.

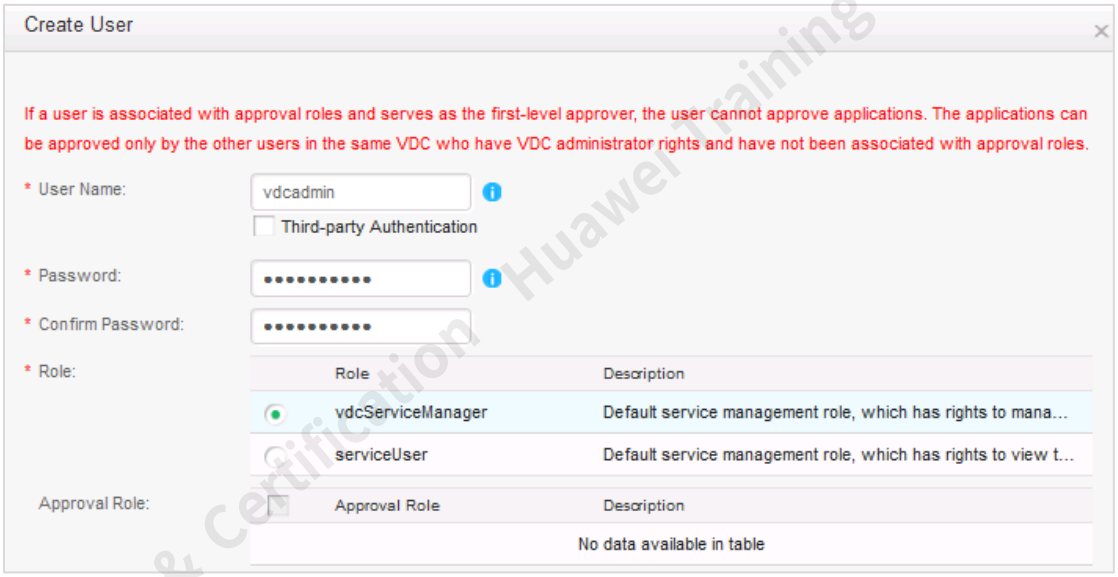

Step 3 **Create a VDC user.**

Create a VDC user and enter the username **vdcuser01** and the initial password **Huawei@123**. Change the password to **Huawei@1234** upon the first login.

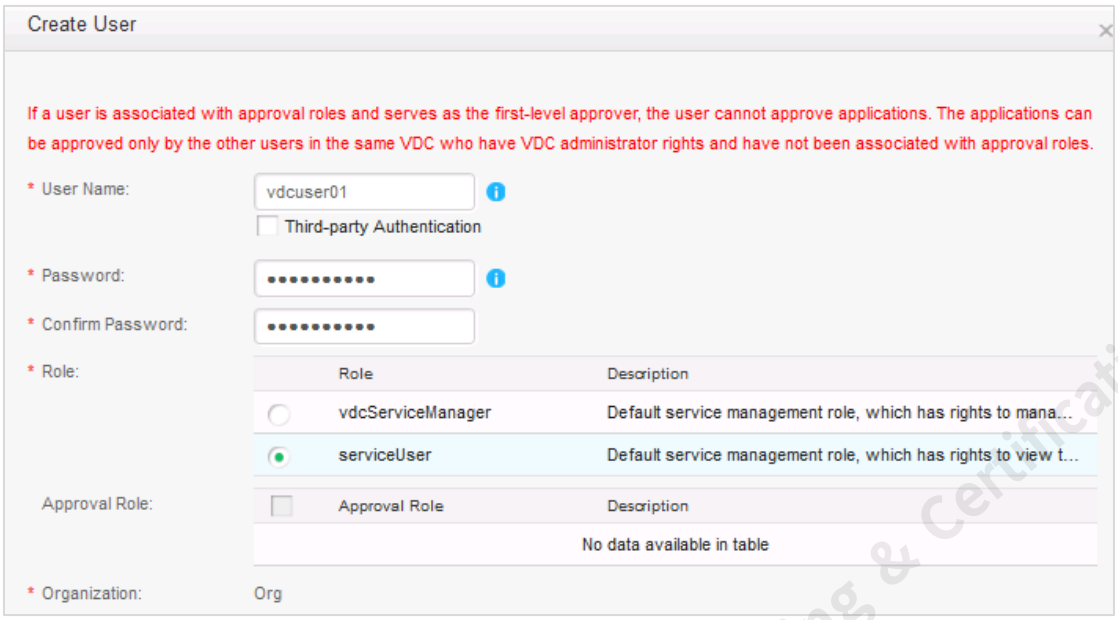

----End

7.1.4 Associating a VDC with an External Network

Click External Network Management.

Click Allocate.

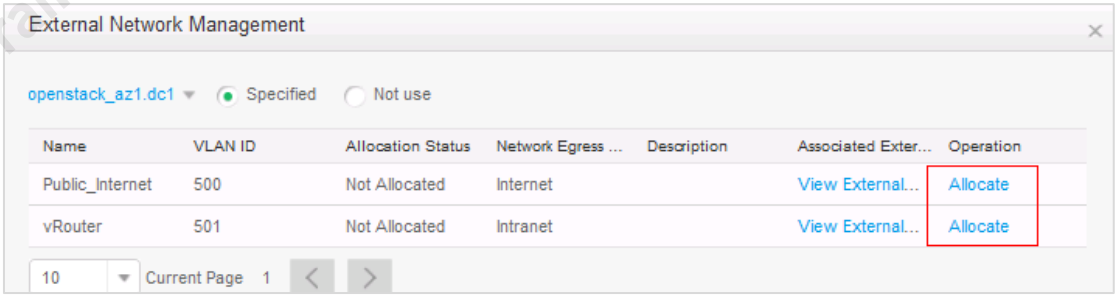

7.1.5 Creating a VPC

Use the **vdcadmin** account to log in to ServiceCenter as a VDC administrator.

Change the password to Huawei@1234 upon the first login.

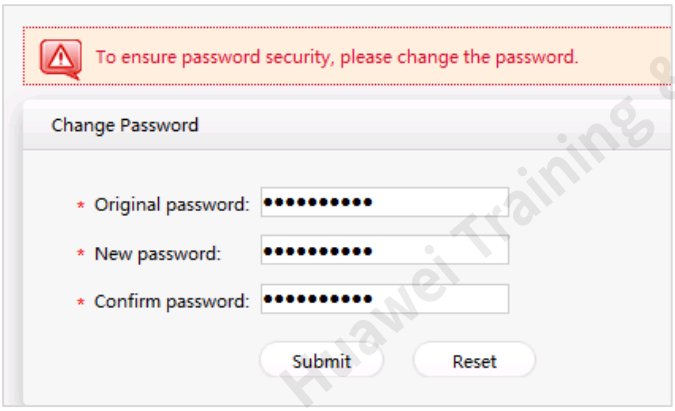

Choose Console > Networks > VPCs, and click Create VPC.

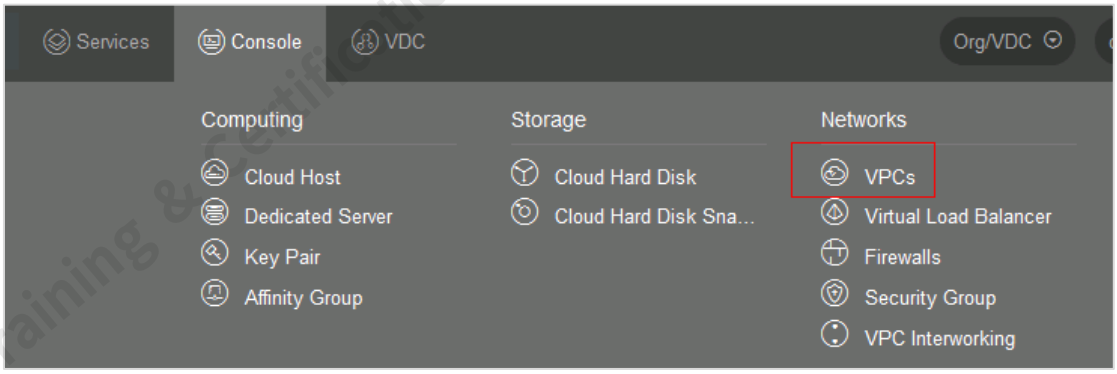

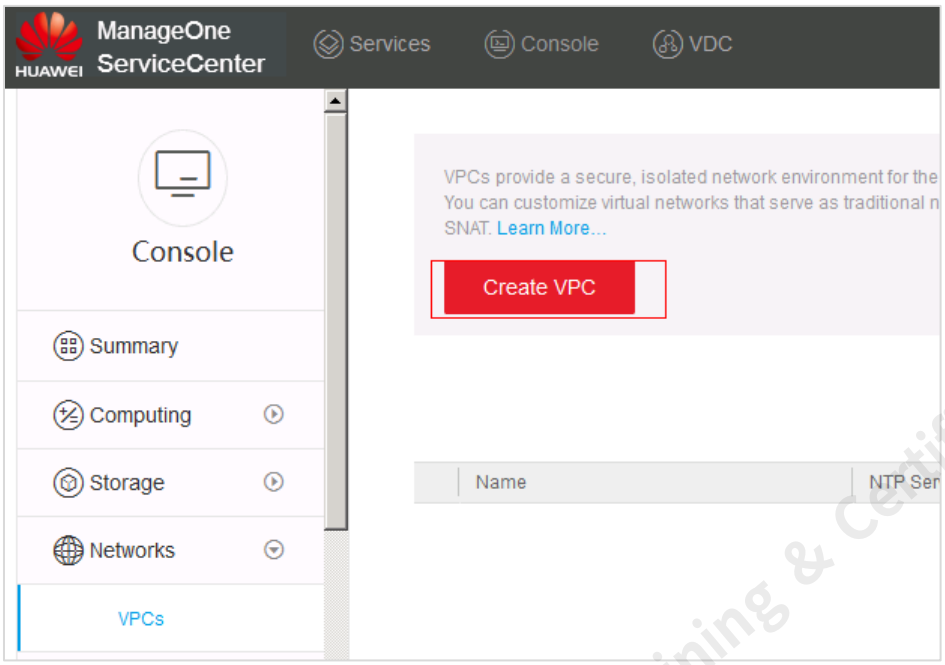

Enter the VPC Name to VPC and click **Next**.

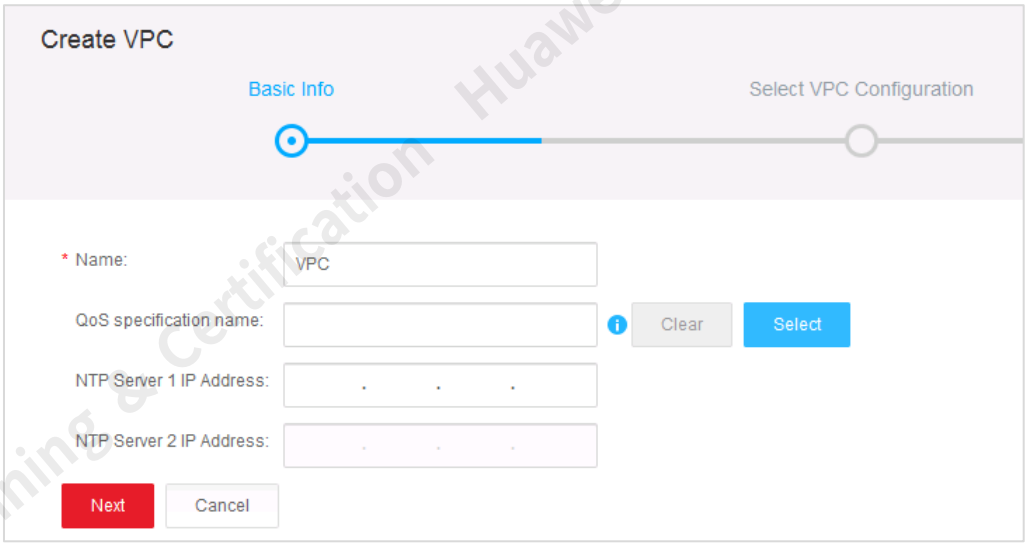

Customize networks after the VPC is created.

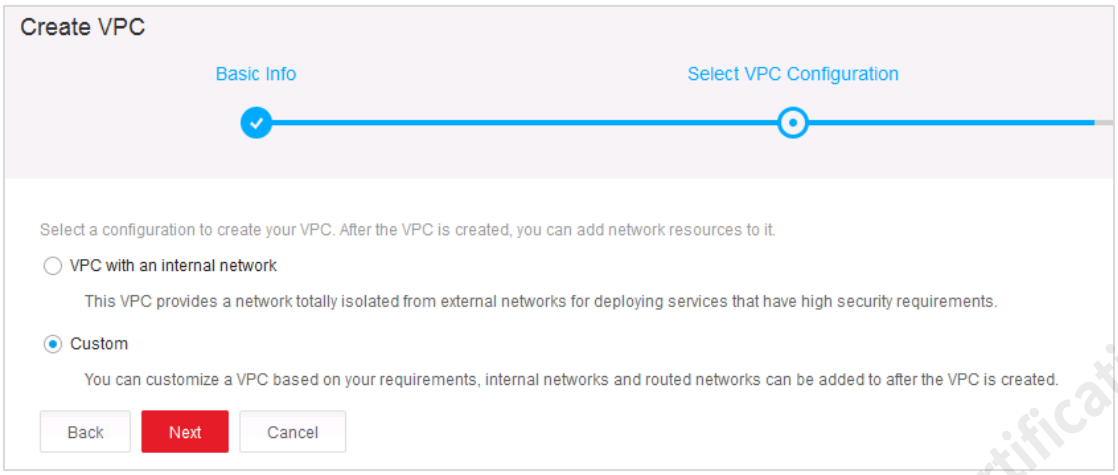

The VPC is successfully created.

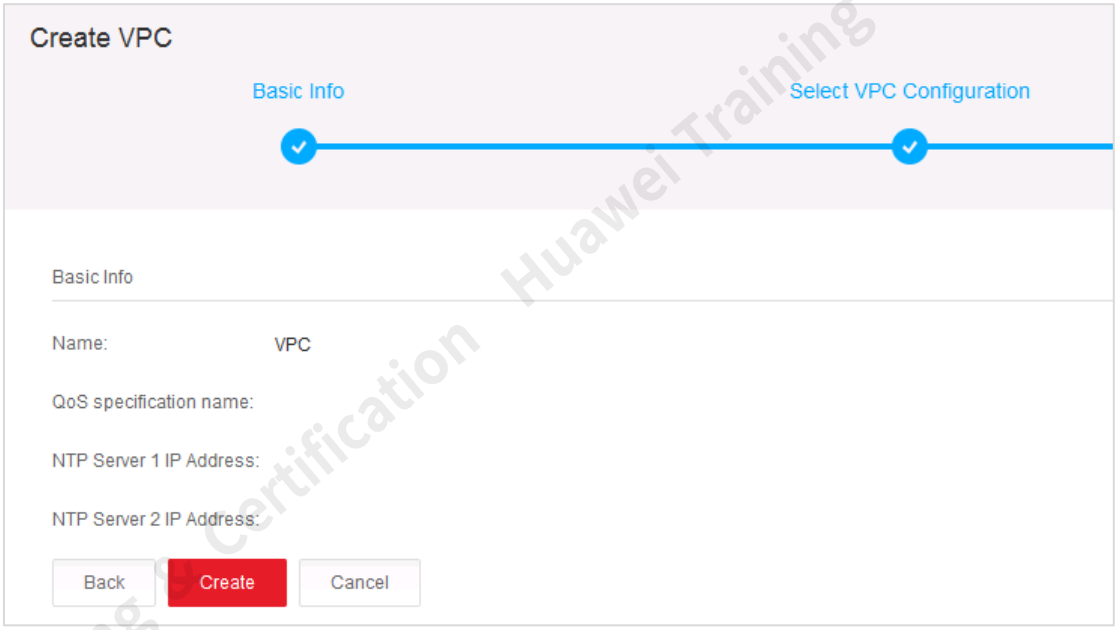

7.2 Commissioning the Routed Network

A routed network supports multiple service functions and flexible interworking capabilities. Virtual routers in a VPC can communicate with public networks based on elastic IP address (EIP), source network address translation (SNAT), or virtual private network (VPN), or communicate with other routed networks in the VPC.

7.2.1 Prerequisites

- The organization is created.

- The VDC/VPC is created.
- The cloud resource pool and VDC have been associated with external networks.

7.2.2 Applying for a router

Before applying for computing resources, you need to apply for network resources. The routed network is the most commonly used network in the cloud-network integration. The router application is the basis for creating the routed network.

Step 1 **Apply for a router.**

Log in to ServiceCenter as the VDC administrator and choose **Console** > **Network** > **VPCs**.

Click **VPC**, click **Routers**, and click **Apply for Router**.

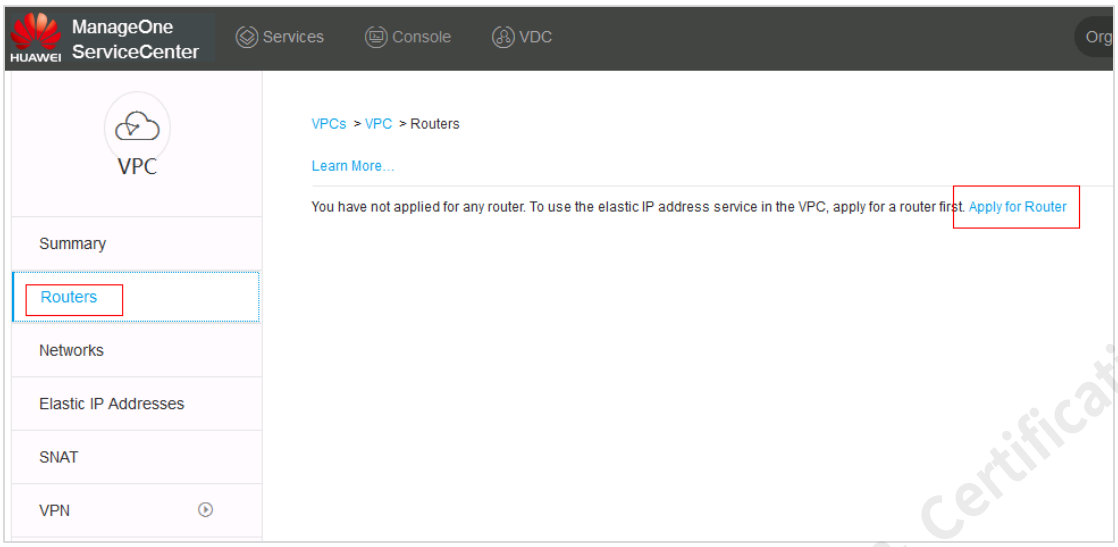

Select an external network and click **Apply**.

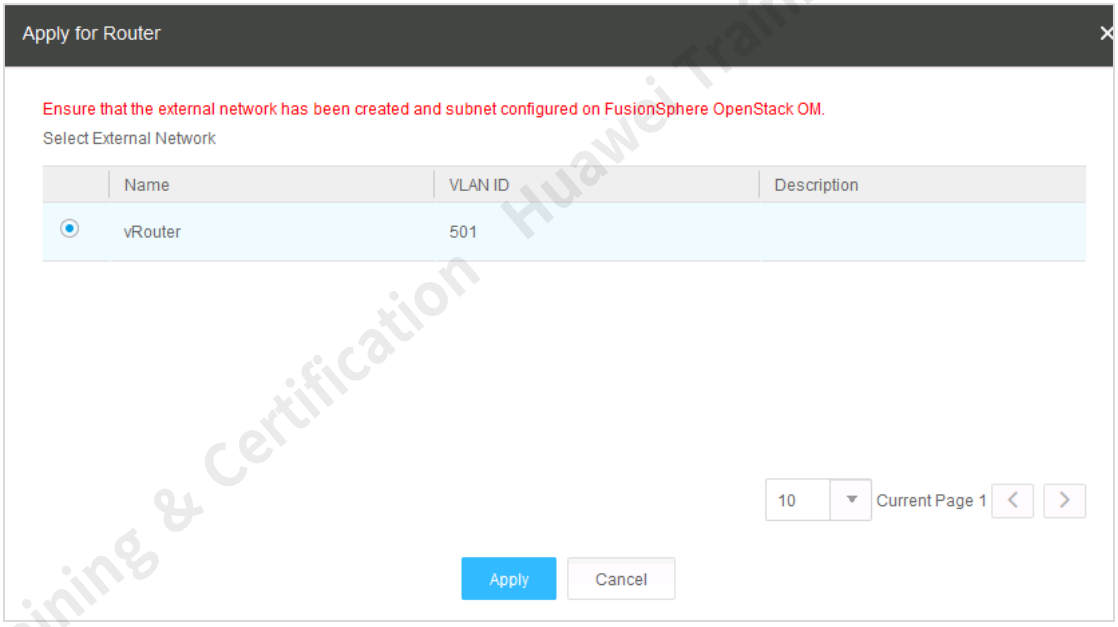

After application is successful, the router obtains an IP address from the external network segment.

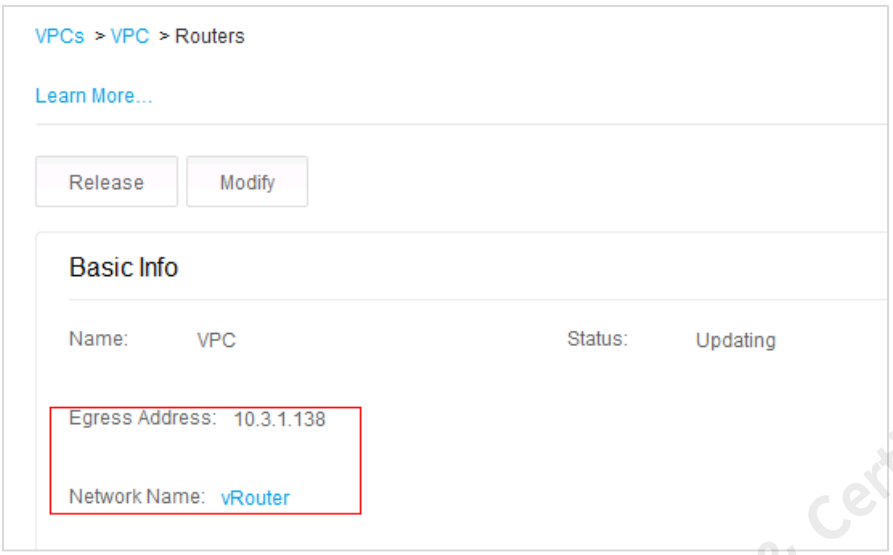

Step 2 **Check the configuration of the cloud- network configuration.**

Log in to the Agile Controller to view the operation logs. The router and the port are successfully created.

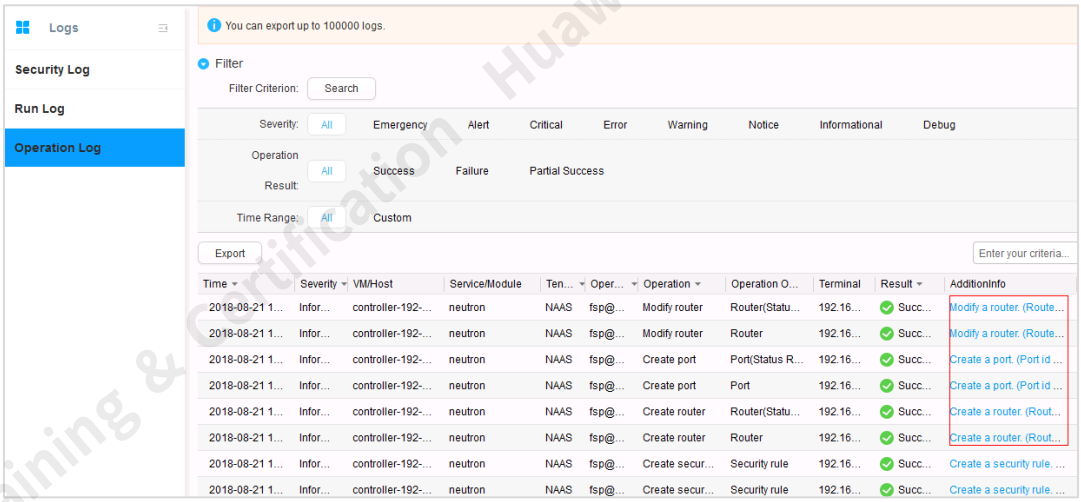

Choose **Tenant** > **Tenant Management**, and the VDC information shows that the VDC created on the cloud platform corresponds to the tenant created on the Agile Controller.

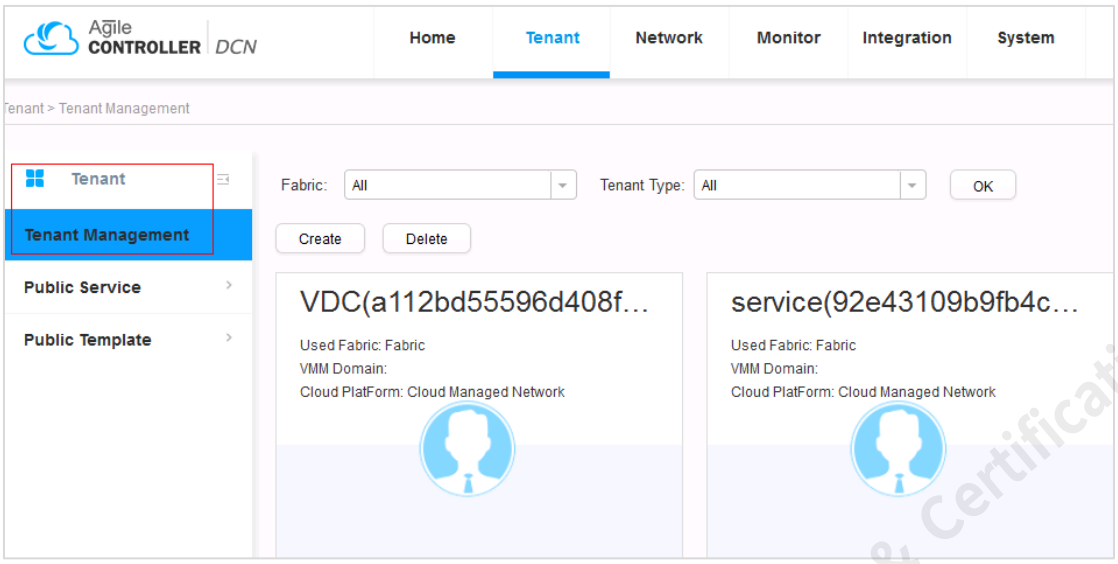

On the VDC page, choose **Overview** > **VDC**, and check the basic information of the VPC network deployed by the Agile Controller.

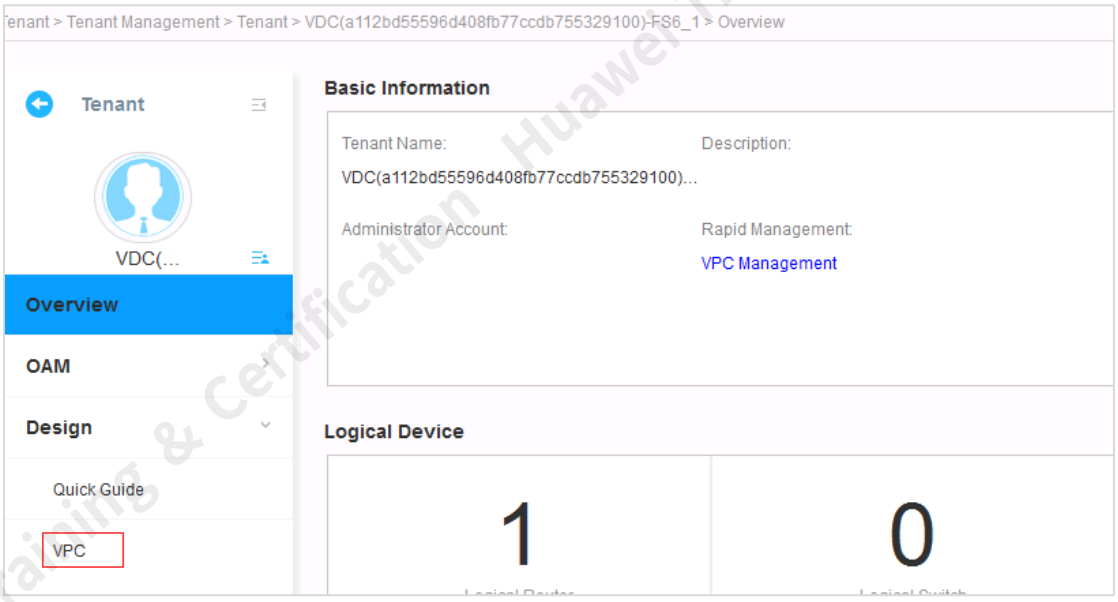

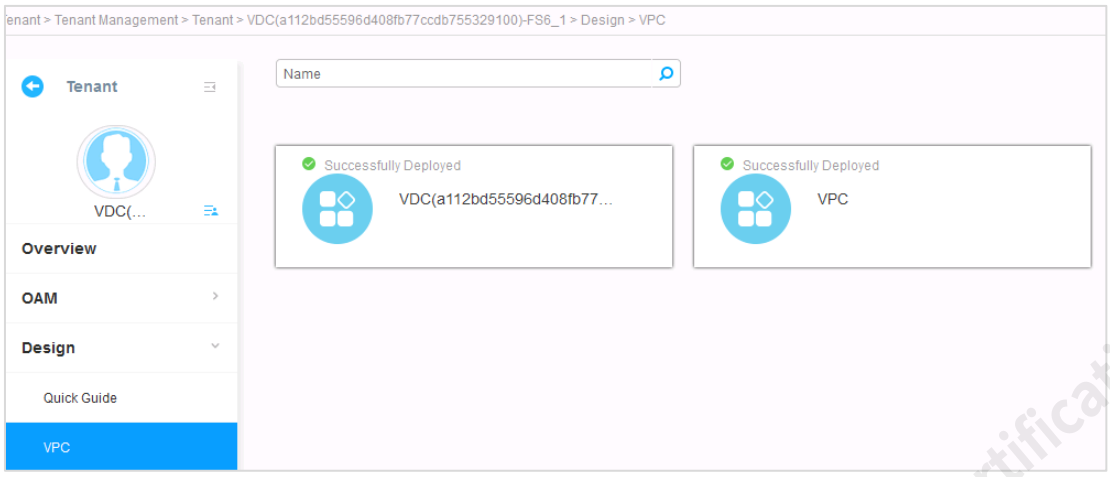

Click VPC to view the network model. No subnet information is available.

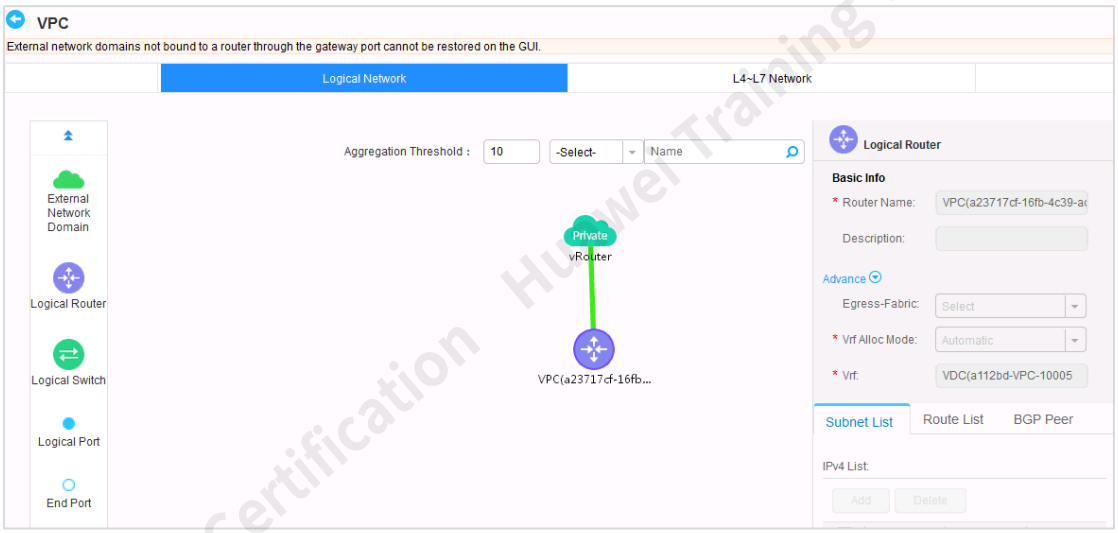

----End

7.2.3 Applying for a Virtual Firewall

If cloud hosts in a VDC need to access the public network (for example, SNAT/EIP services are used), you need to apply for the virtual firewall first.

Step 1 **Create a virtual firewall.**

Choose Console > Networks > Firewalls and click Create.

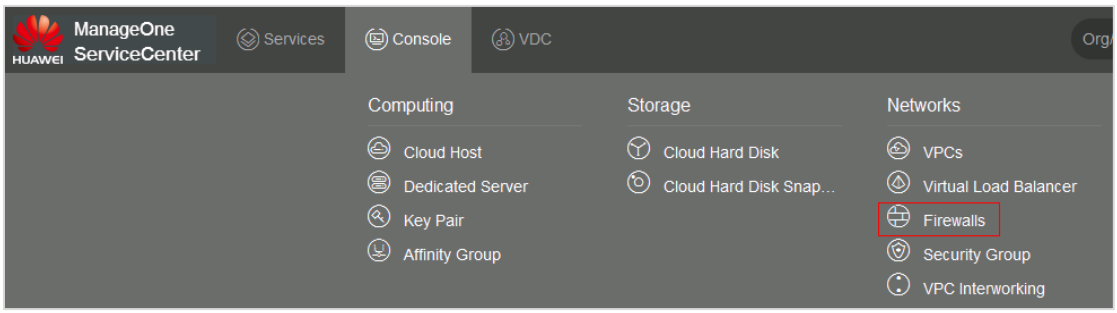

Enter the names of the firewall and the associated VPC, and click **OK**.

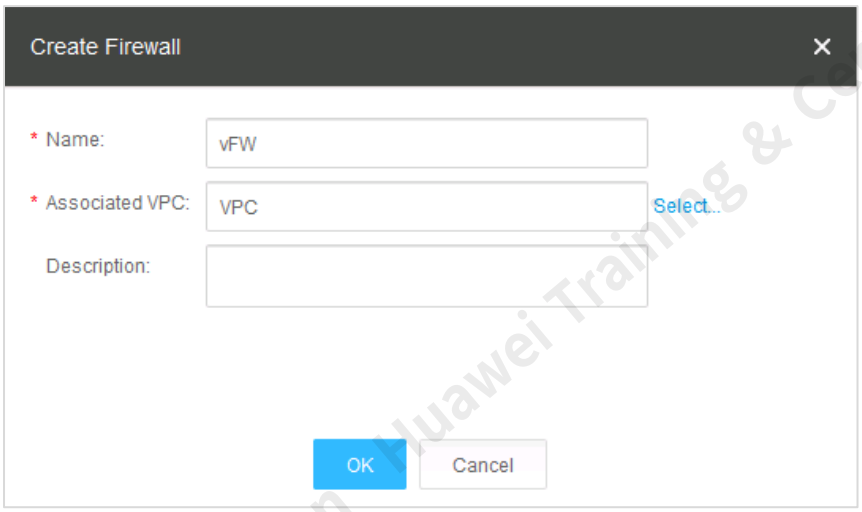

Step 2 **Check the created firewall.**

Check the Agile Controller operation logs and the logs shows that the vsys is successfully created.

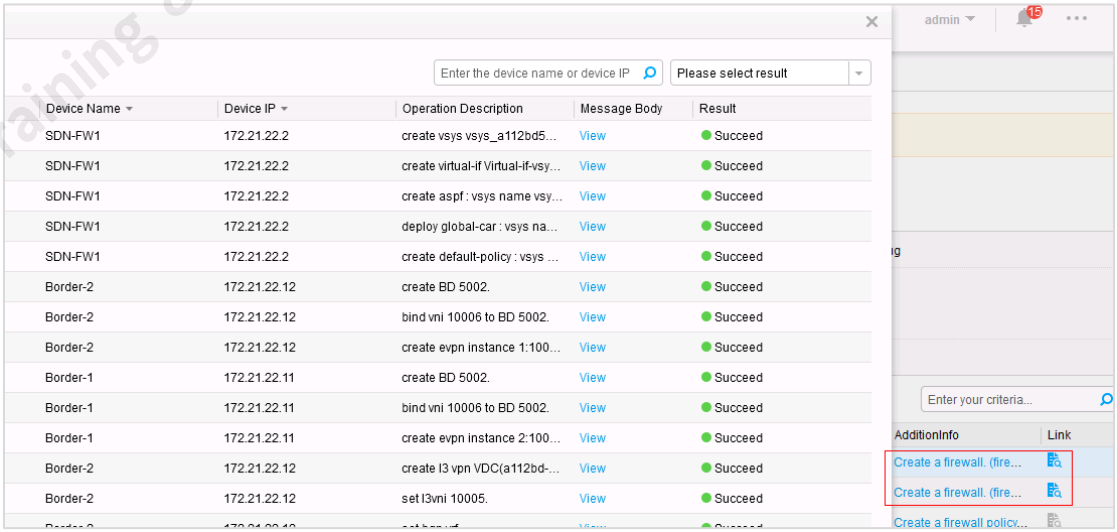

The firewall status on ServiceCenter is normal.

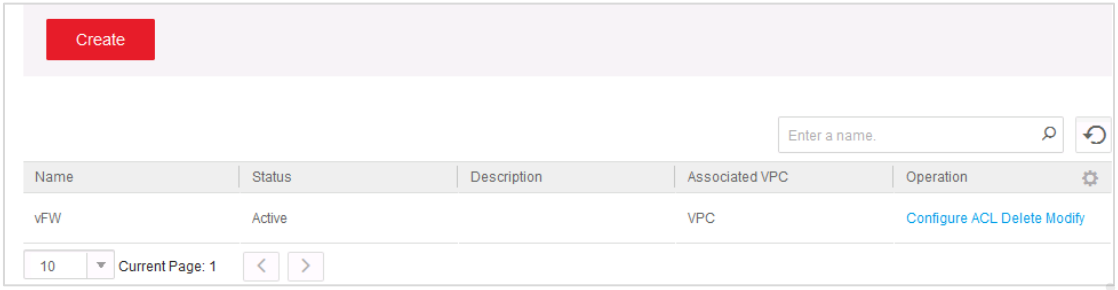

----End

7.2.4 Creating a Routed Network

After applying for a router and a firewall, you need to create different subnets associated with the router. In this experiment, the network segment 192.168.100.0/24 is created and associated with the router, forming a routed network (named route01).

Step 1 **Create a routed network.**

Log in to ServiceCenter as the VDC administrator **vdcadmin**. Choose **VDC** > **Networks** > **Routed/Internal network** on the **VDC** page, and click **Create**.

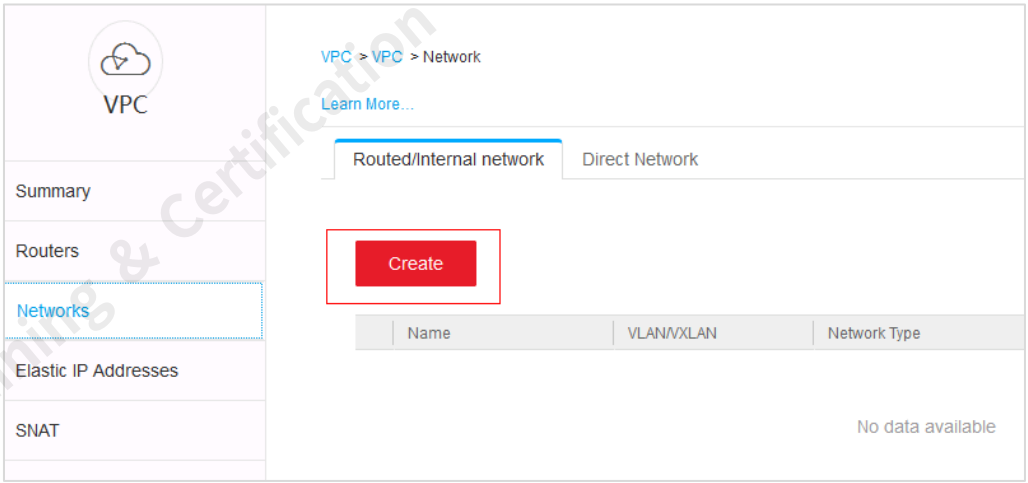

Set the routed network name to **route01**, select **Routed Network** for **Network Type**, and click **Next**.

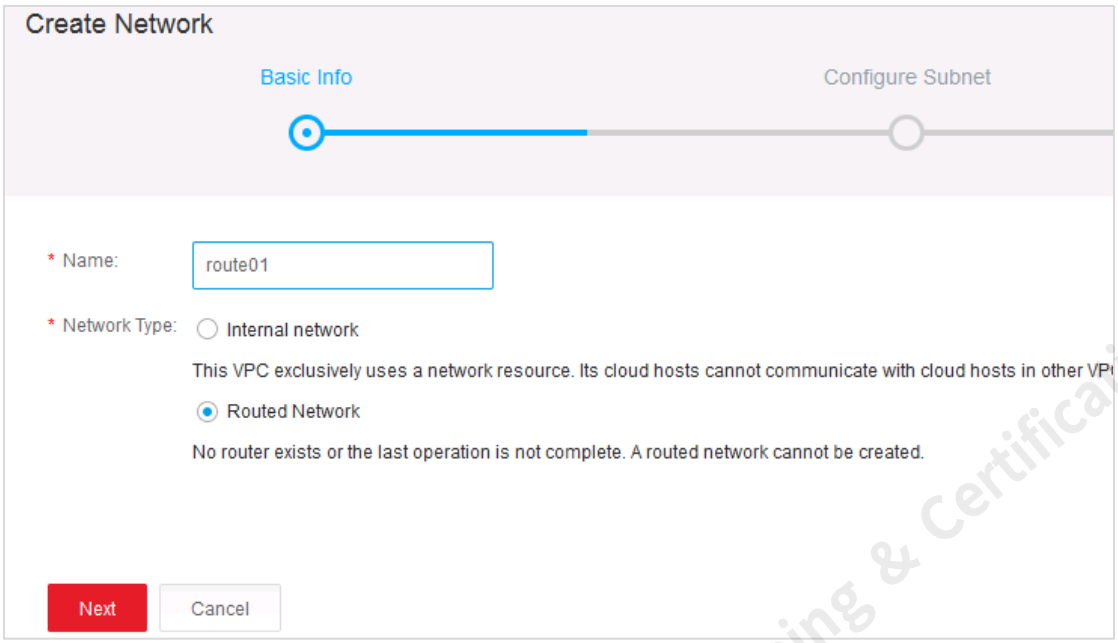

Enter the information about the network and network segment.

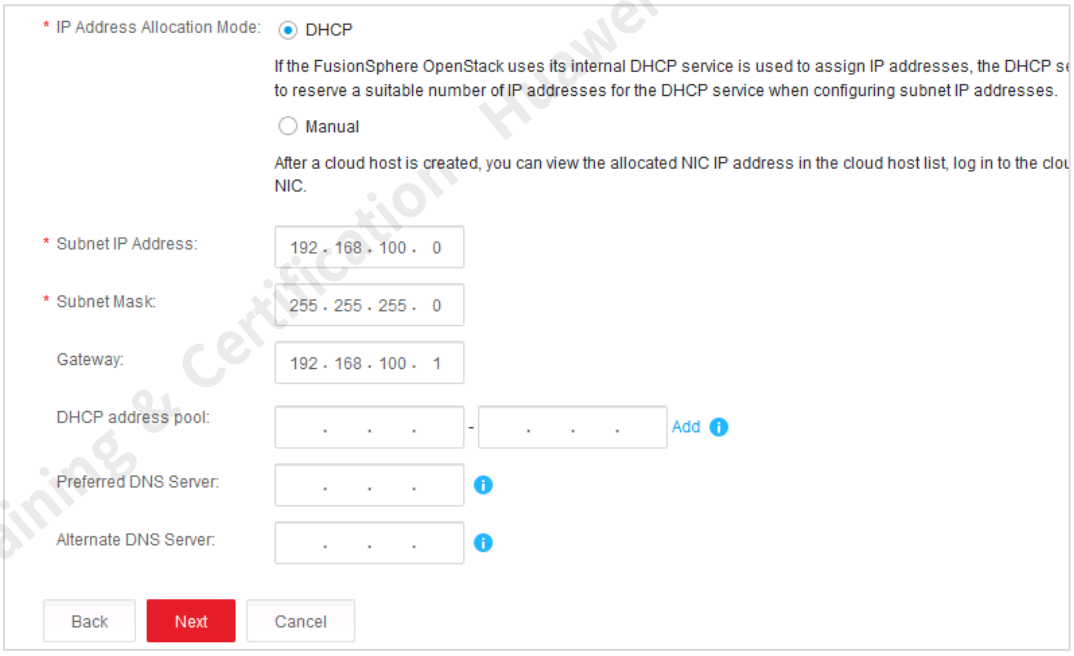

The routed network is created.

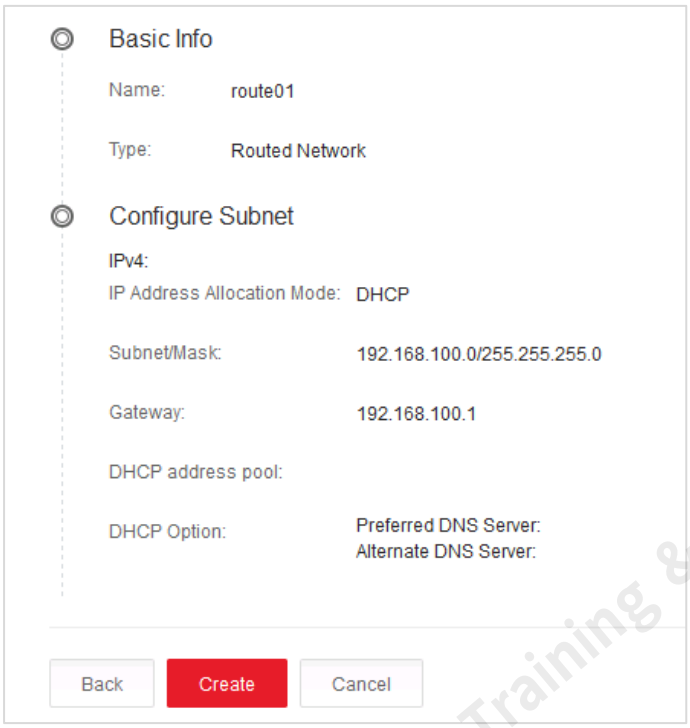

The routed network is successfully created.

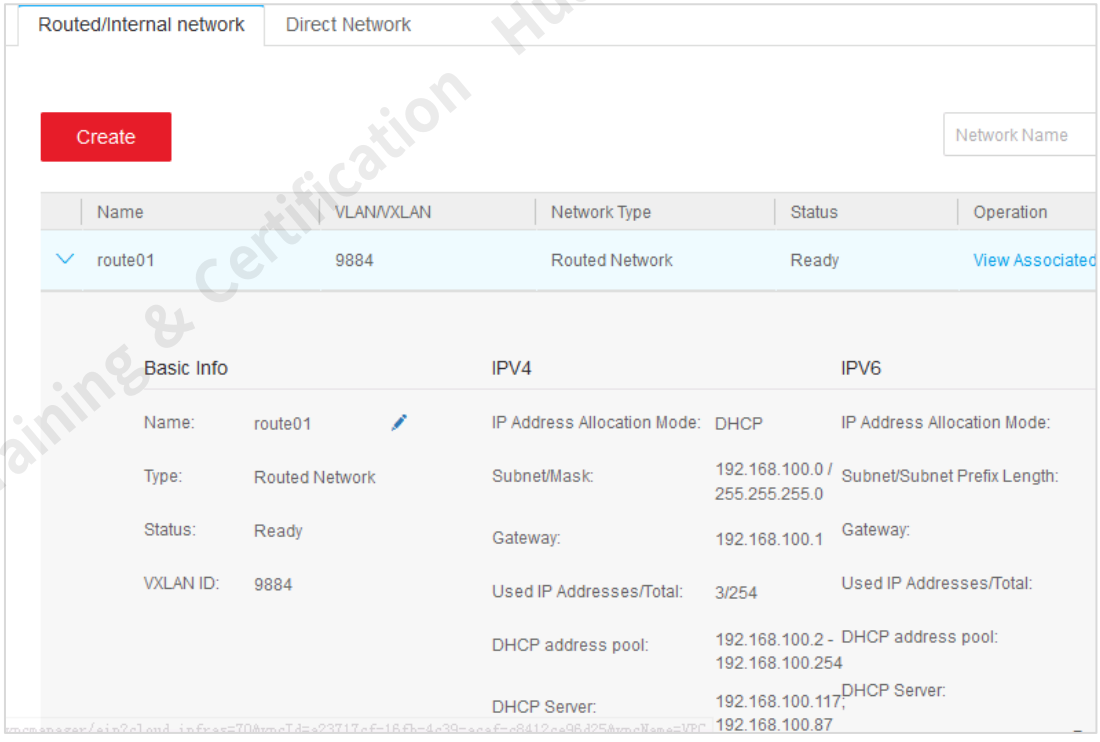

Step 2 **Check the configuration delivered by the cloud-network.**

Check the VPC information of the tenant on the Agile Controller.

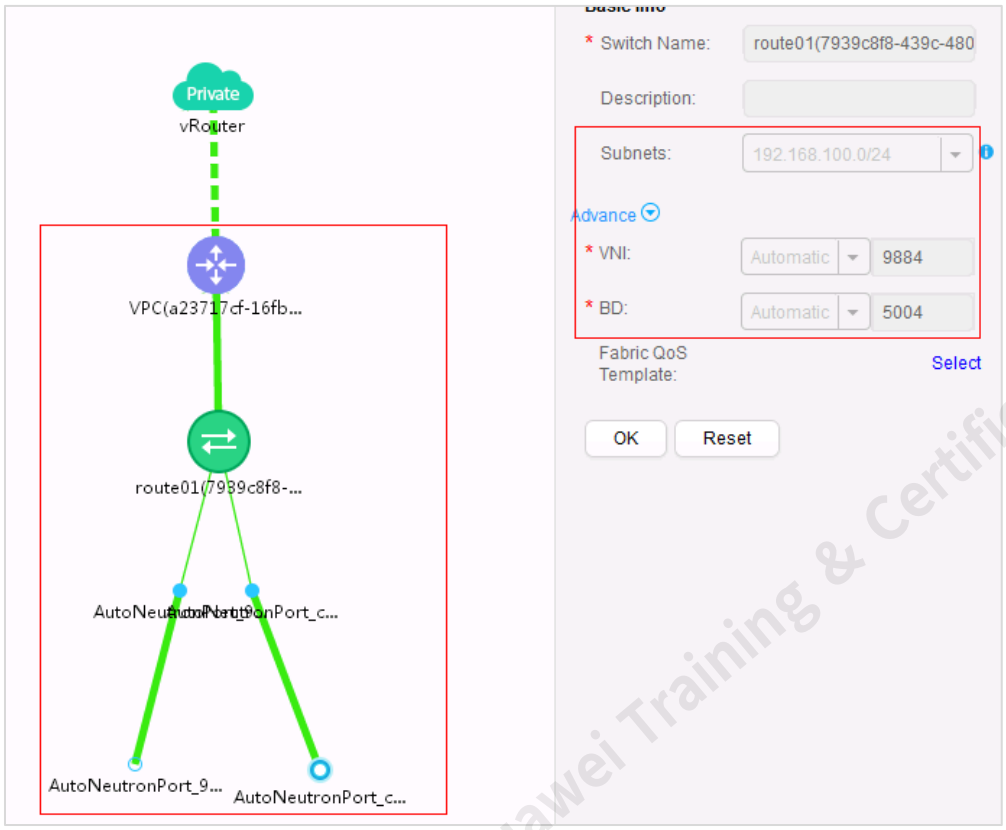

Create a logical switch, and check its subnet information.

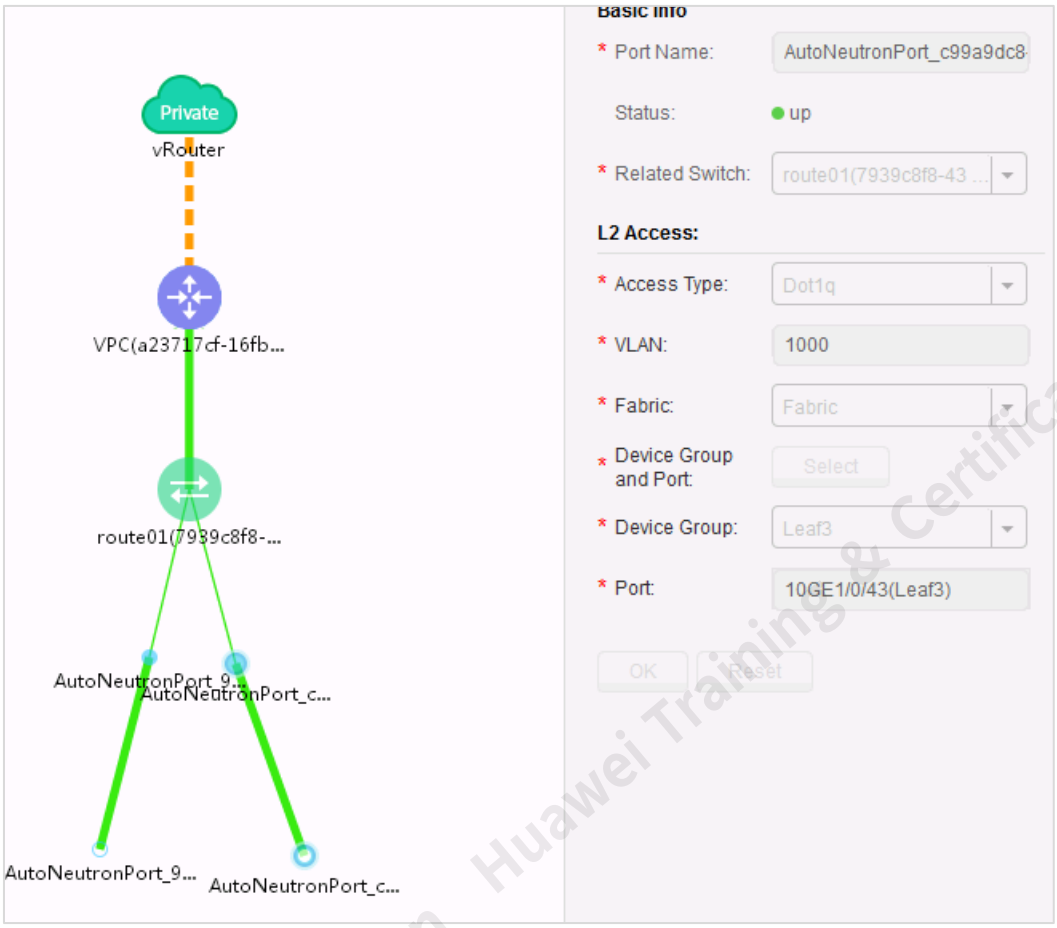

The logical port goes online and configurations are correctly delivered.

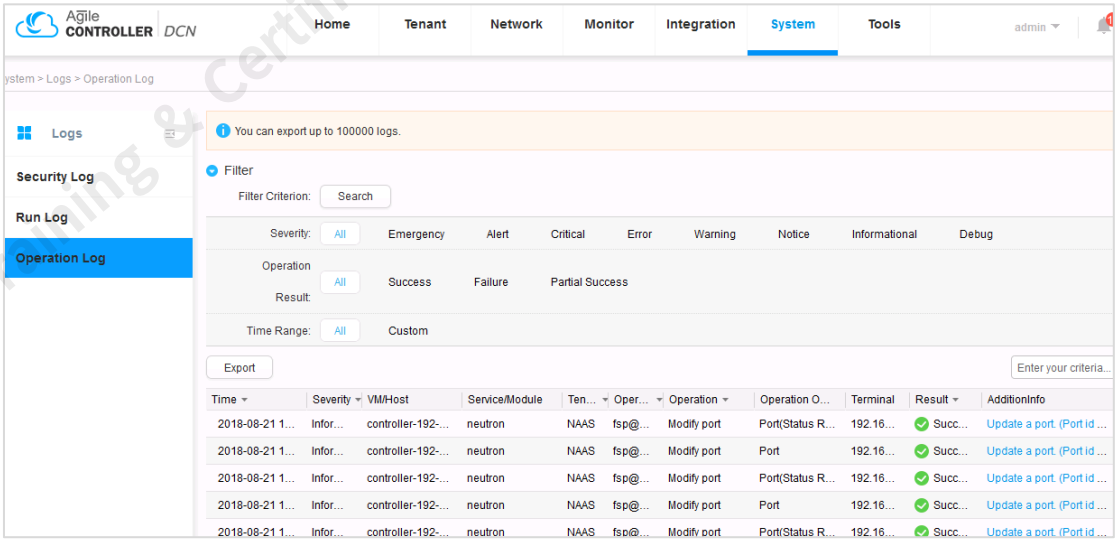

Step 3 **Check the configuration delivered by the switch.**

```
<Leaf-3>display configuration commit changes last 1
Building configuration
  #
+ bridge-domain 5004
+  vxlan vni 9849
+  evpn
+   route-distinguisher 2:9849
+   vpn-target 0:9849 export-extcommunity
+   vpn-target 0:9849 import-extcommunity
  #
+ interface 10GE1/0/45.3 mode l2
+  encapsulation dot1q vid 1001
+  bridge-domain 5004
  #
+ interface 10GE1/0/46.3 mode l2
+  encapsulation dot1q vid 1001
+  bridge-domain 5004
  #
  interface Nve1
+  vni 9849 head-end peer-list protocol bgp
  #
<Leaf-3>
```

Create a VXLAN sub-interface, Bridge-domain5004, and VXLAN VNI 9849 on the port connected to the FSP server.

Check the configuration delivered by the spine node.

```
<Spine-2>display configuration commit changes last 1
Building configuration
  #
+ bridge-domain 5004
+  vxlan vni 9849
+  evpn
+   route-distinguisher 3:9849
+   vpn-target 0:9849 export-extcommunity
+   vpn-target 0:10005 export-extcommunity
+   vpn-target 0:9849 import-extcommunity
  #
+ interface Vbdif5004
+  ip binding vpn-instance VDC(2dba14-VPC-10005
+  ip address 192.168.100.1 255.255.255.0
+  mac-address 0000-5e00-0102
  #
  interface Nve1
+  vni 9849 head-end peer-list protocol bgp
  #
+ ip route-static 192.168.100.0 255.255.255.0 10.125.97.242
  #
<Spine-2>
```

Create a VBridge-domainIF gateway on the spine node and create a default route destined for the firewall interface (10.125.97.242).

----End

7.3 Commissioning the Cloud Host

This section describes how to apply for three cloud hosts (Host01, Host02, and Host03), associate cloud hosts with different networks, and test the connectivity.

7.3.1 Creating the Cloud Host Service

Step 1 **Log in to ServiceCenter as the system super administrator.**

Enter the username **cloud_admin** and password **FusionSphere123**.

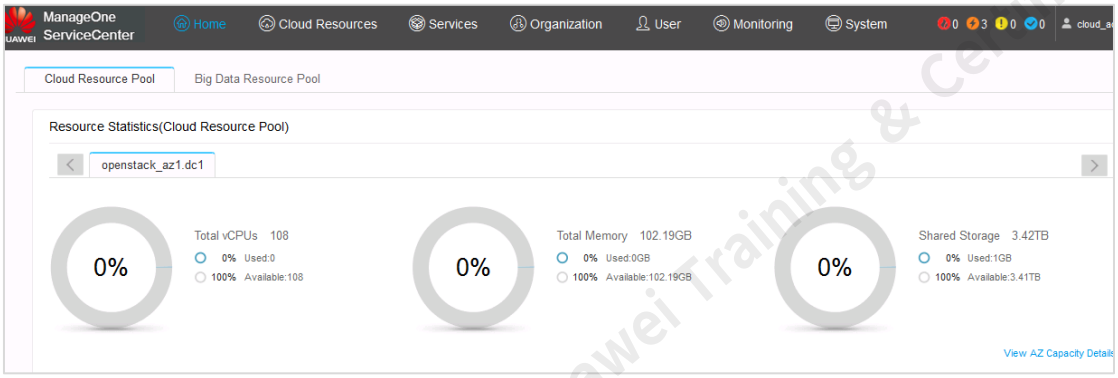

Step 2 **Create the cloud host service.**

Choose Services > Computing > Cloud Host, and click Create Service.

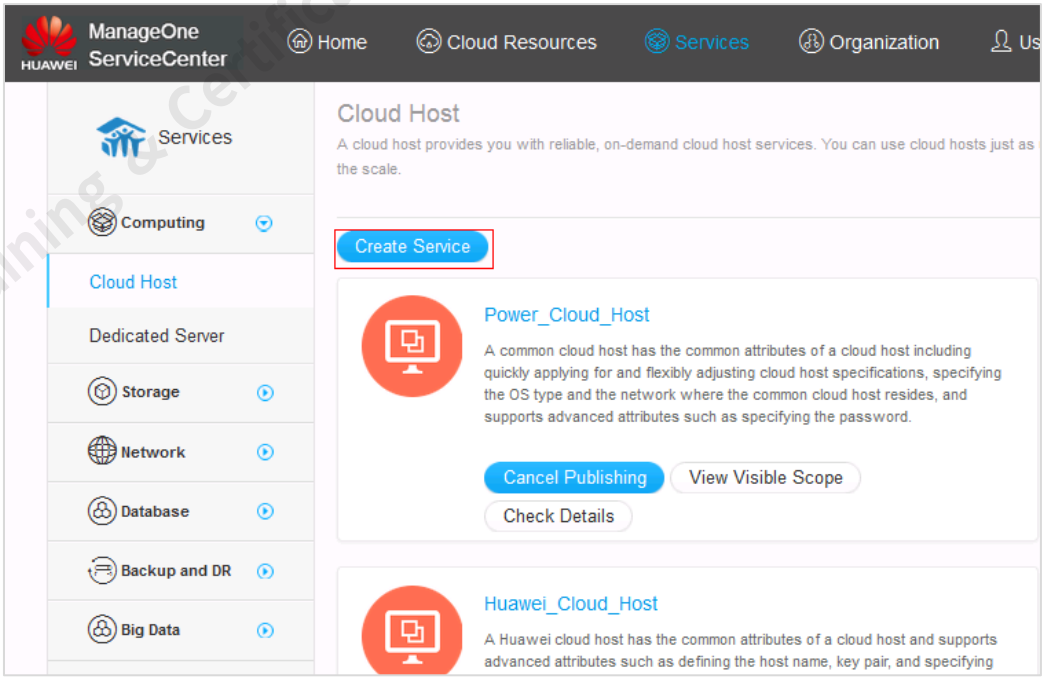

Select the cloud host type and approval type.

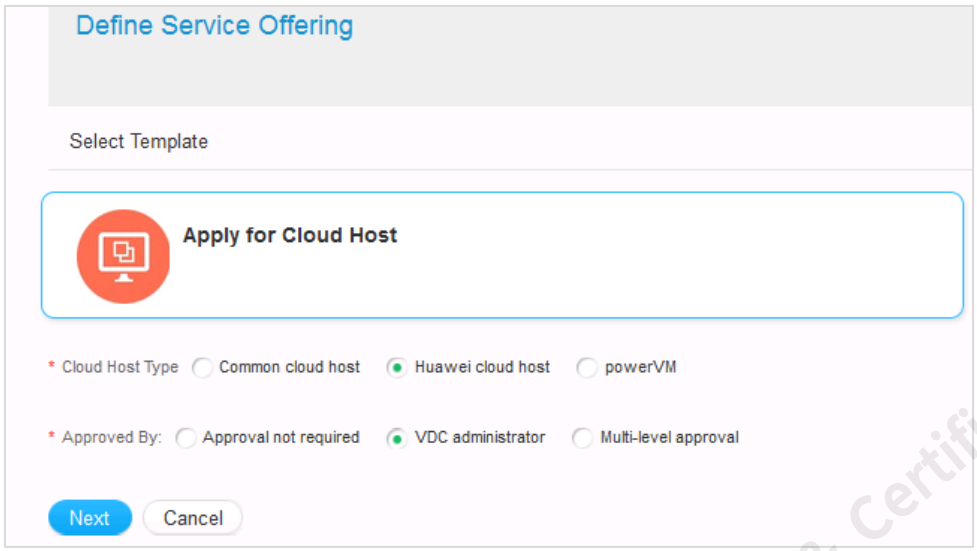

Select a cloud resource pool and an AZ.

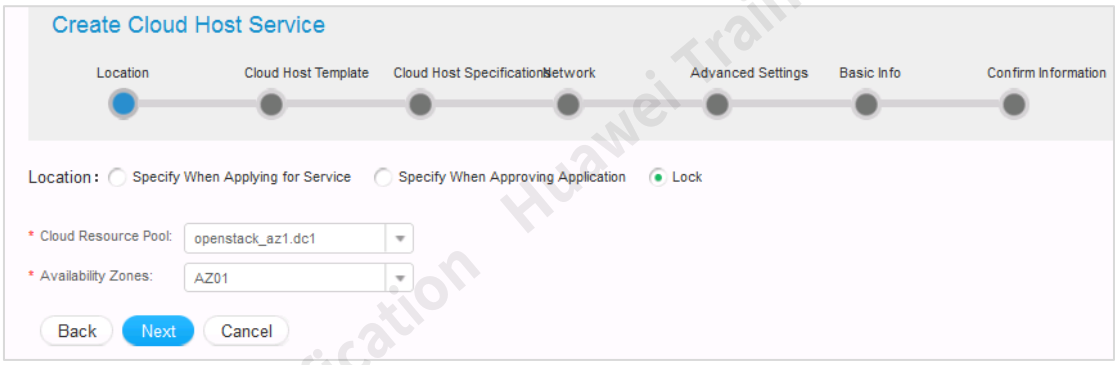

Select an OS.

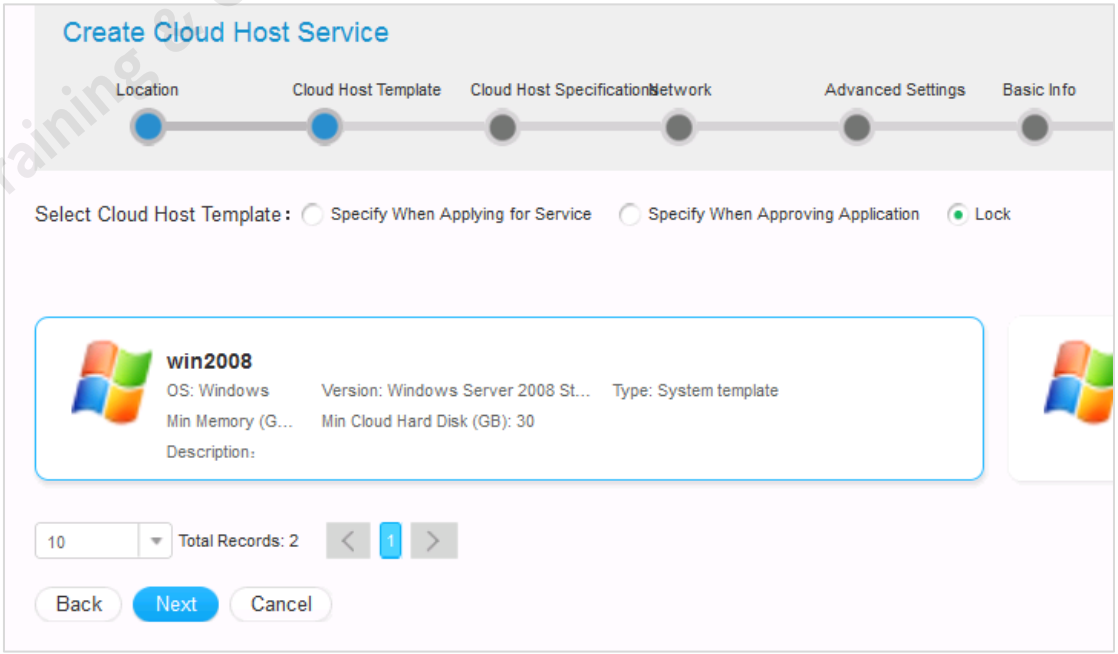

Specify the VM specifications.

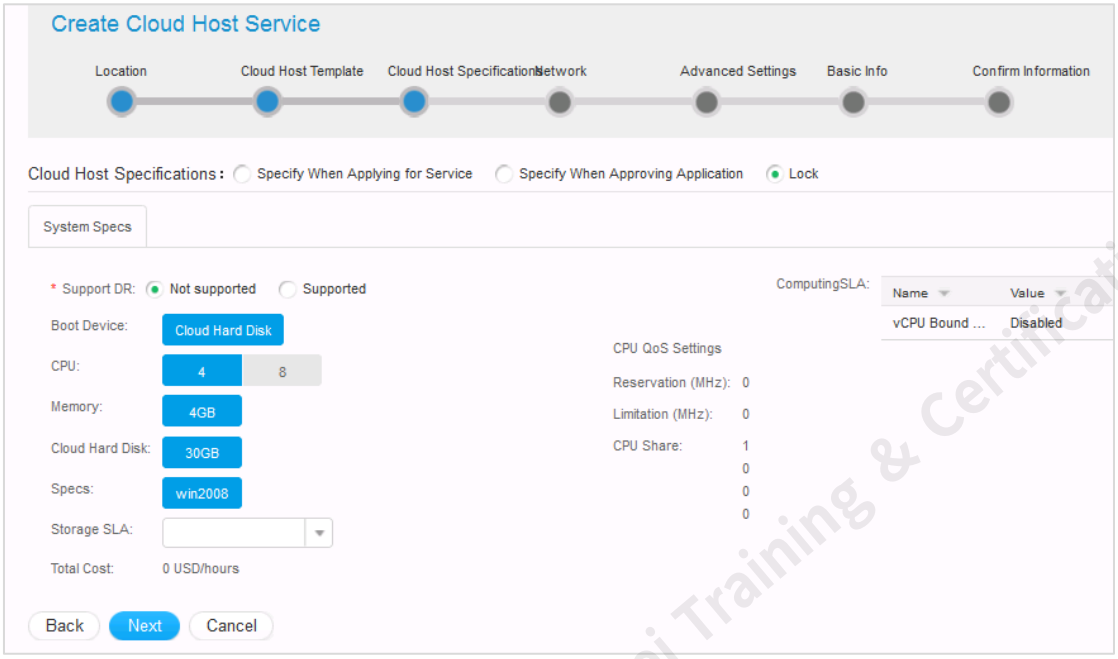

Select Specify When Approving Application for Select Network.

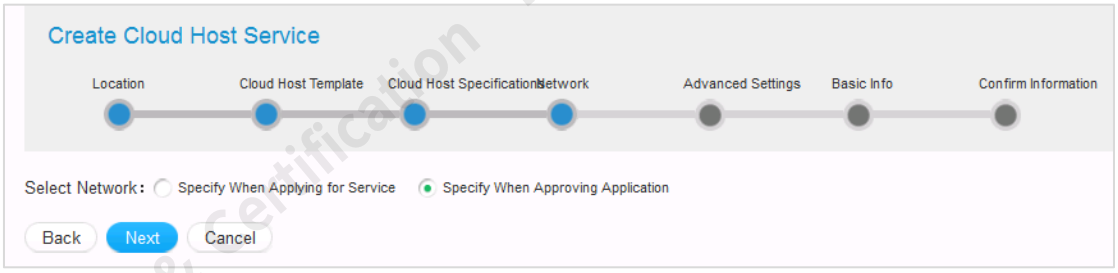

Select Specify When Approving Application for Password Login.

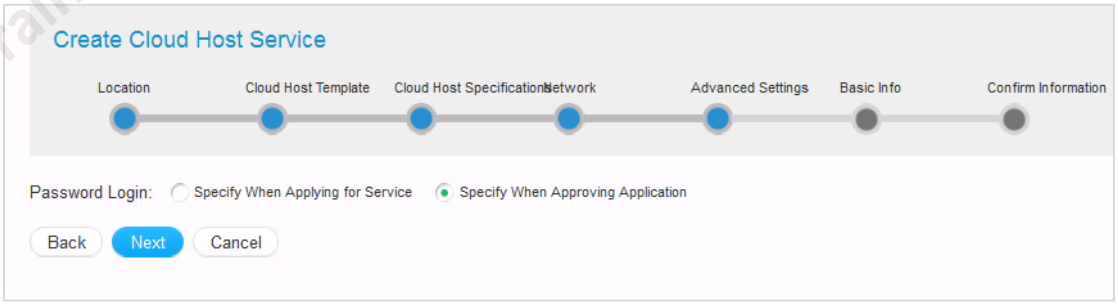

Enter the service name and the cloud host is successfully created.

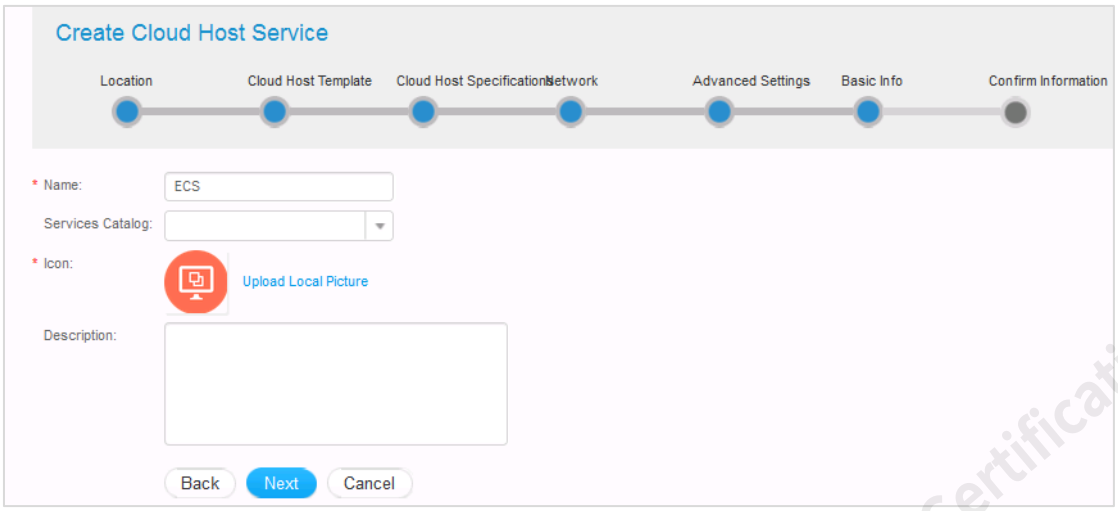

Step 3 **Publish a cloud service.**

After a service is created, you can click **Set Visible Scope** to set the service to be visible to the specified VDC or click **Publish** to publish the service.

After the service is published, VDC users can apply for the service.

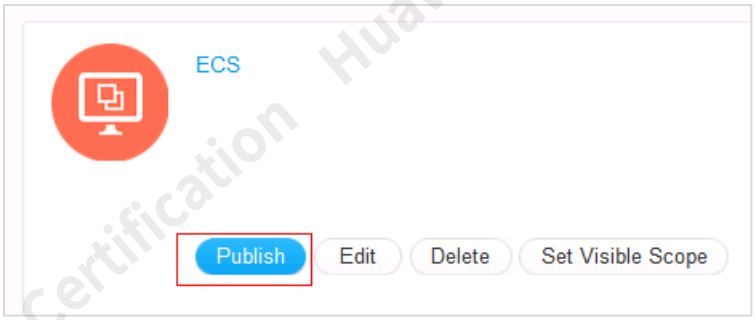

----End

7.3.2 Applying for a Cloud Host

Step 1 **Apply for a cloud host as the VDC user.**

Log in to the MangeOne ServiceCenter using the username **vdcuser01**.

Choose **Services** > **Computing** > **Cloud Host** > **ECS** to enter the page for service application.

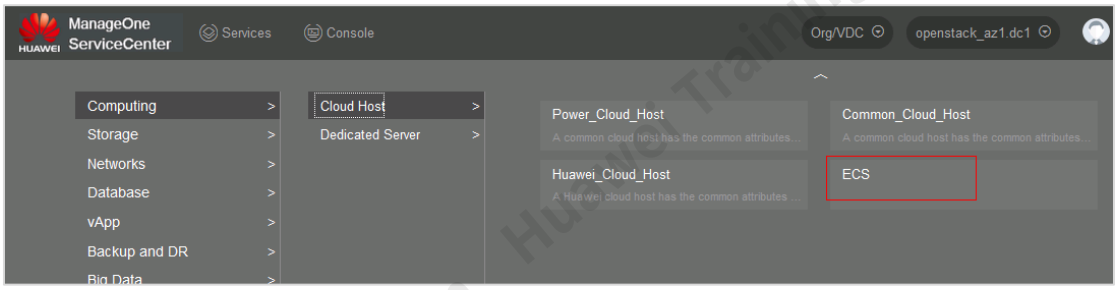

Click Apply for.

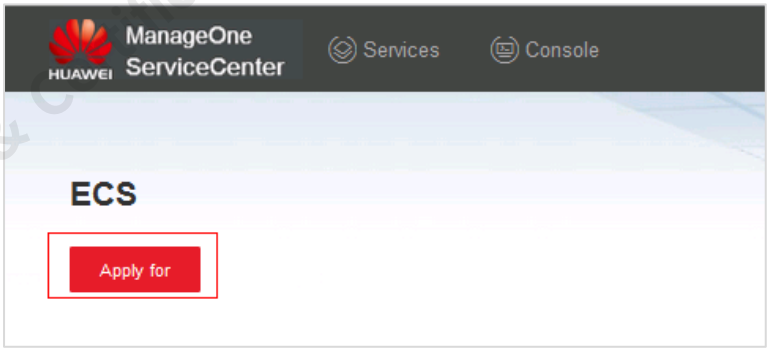

Enter the cloud host name and click **Next**.

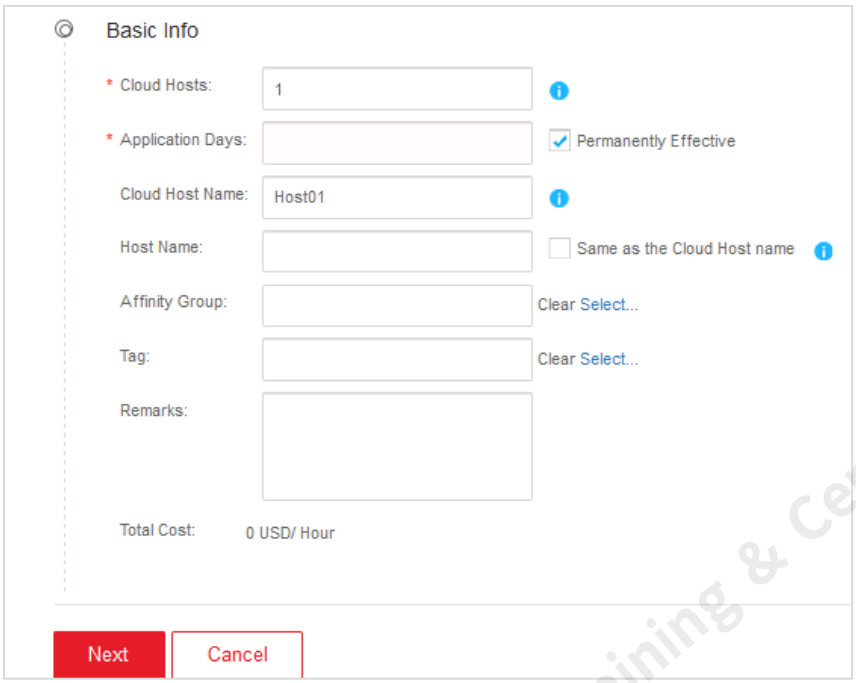

Click **Submit**.

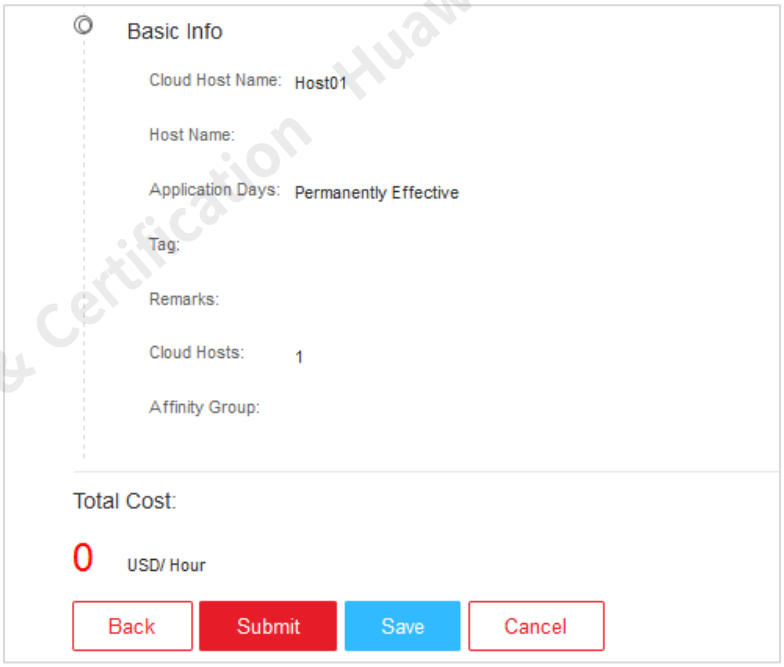

Step 2 **Approve a cloud host.**

Use the username **vdcadmin** to log in to ServiceCenter as the VDC administrator.

Click To Approve.

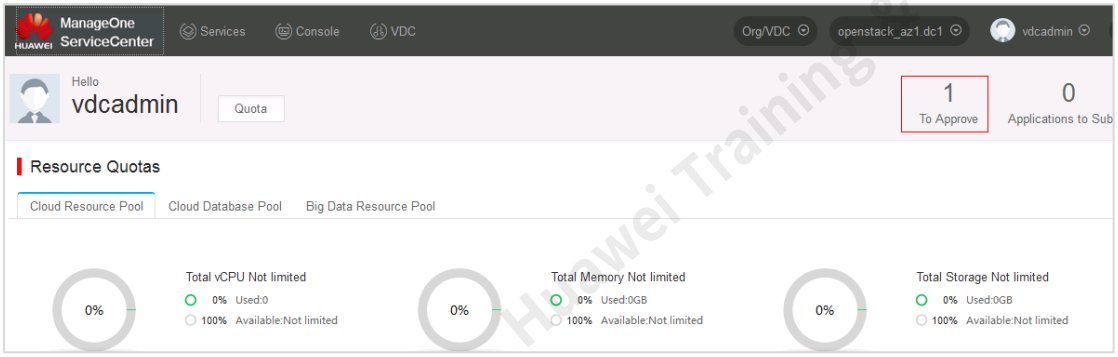

On the **User Center** page, click **Approve**.

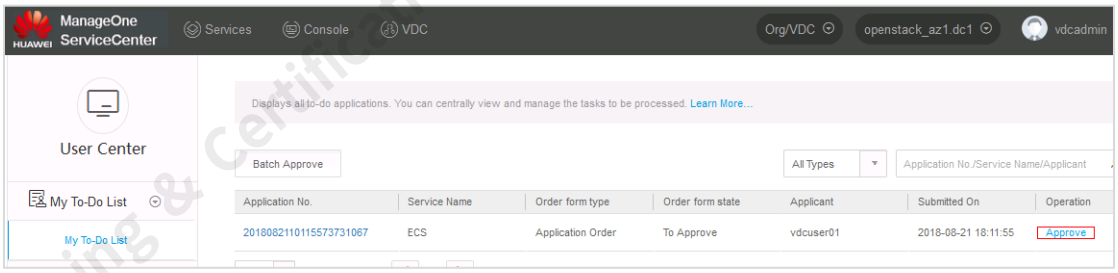

Select the network used by the cloud host.

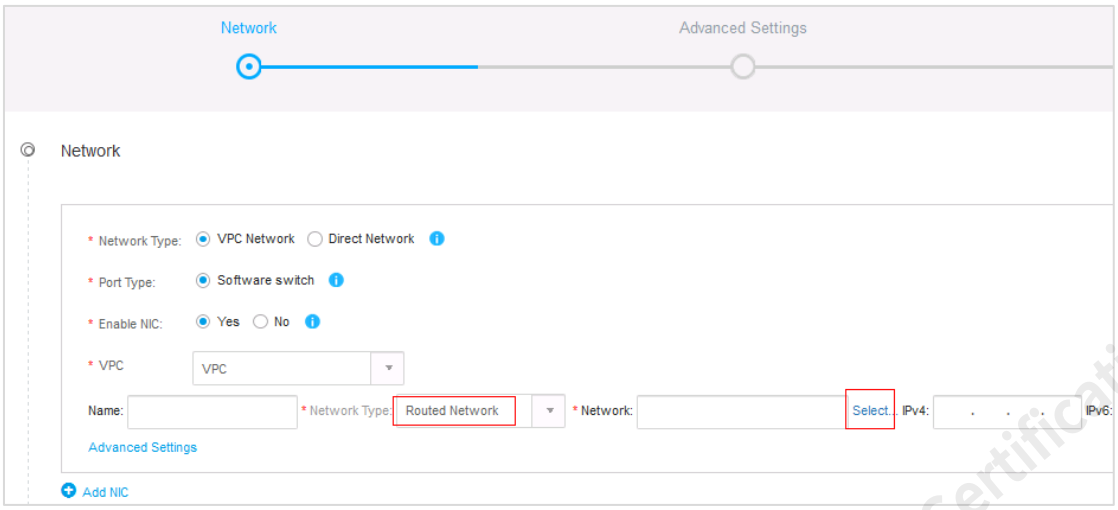

Select the corresponding routed network.

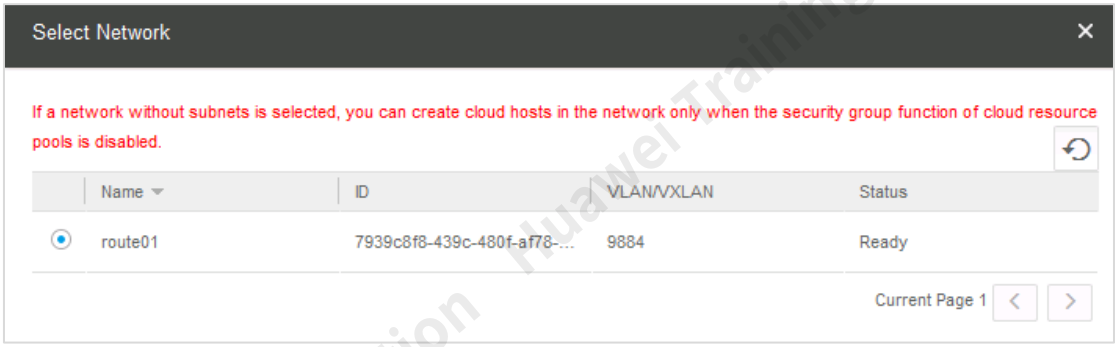

Enter the password and click **Next**.

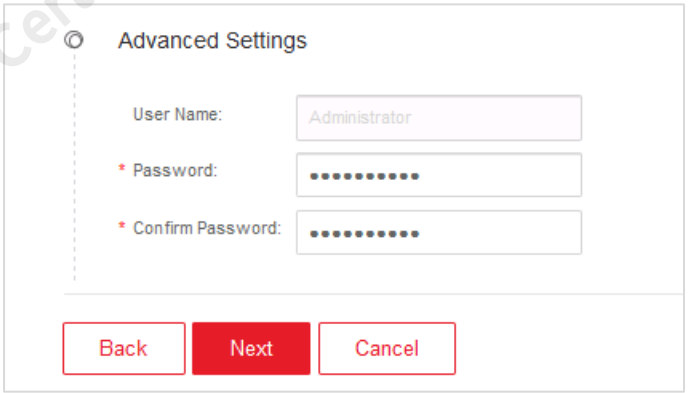

Click **Submit** to submit the approval result.

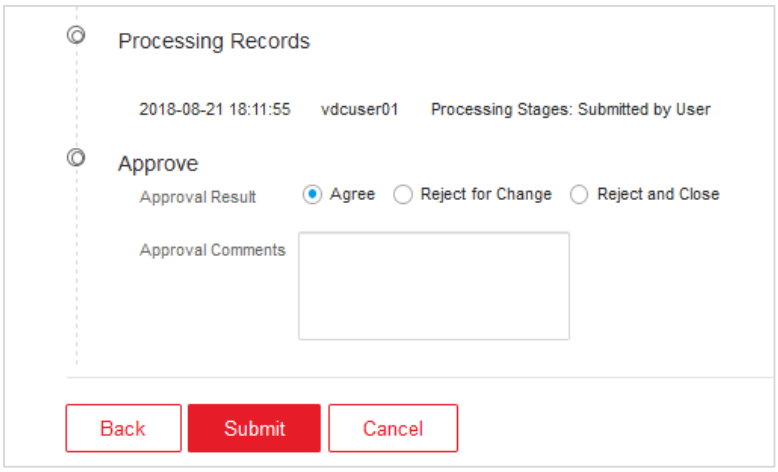

Step 3 **Check the created VM.**

Log in to the FusionCompute using the username **admin** and the password **Huawei@123**. The VM is being created.

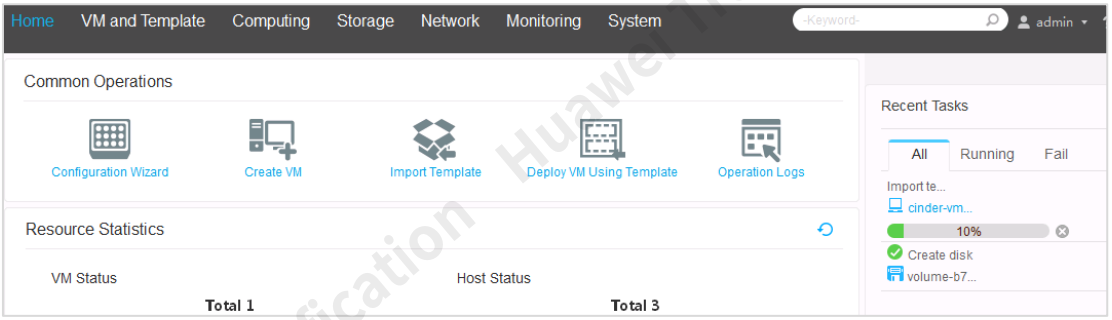

Log in to ServiceCenter using the username **vdcuser01**. On the **User Center** page, click **My Applications**, and the status is **Succeeded**.

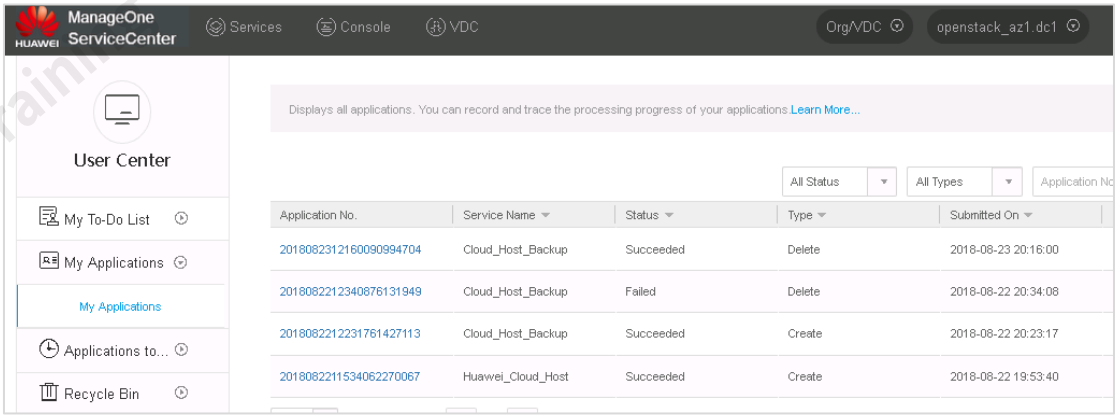

Choose **Console** > **Cloud Host**, and log in to the cloud host using VNC.

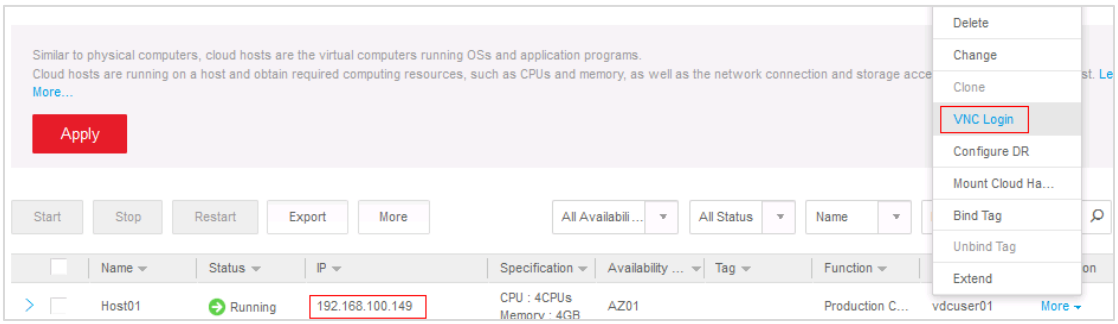

Click Send CtralAltDel.

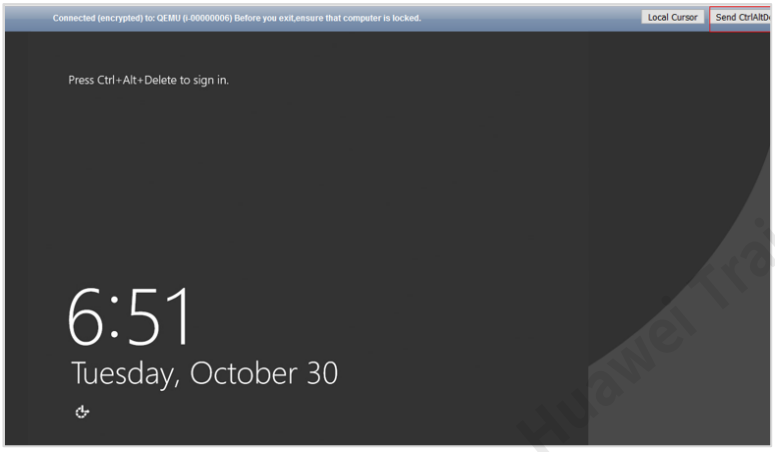

Enter the login password **Huawei@123**.

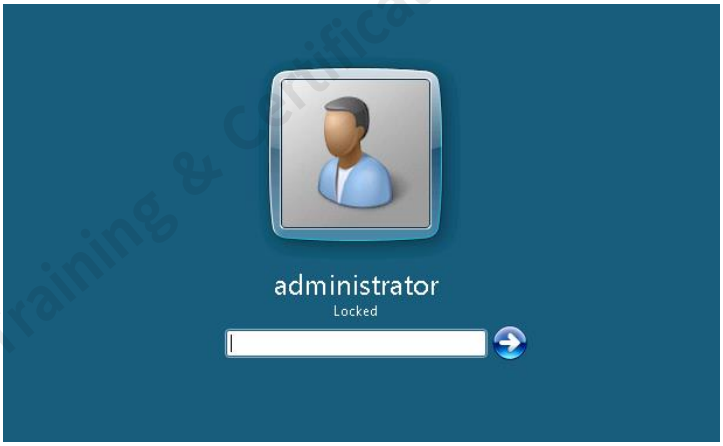

The login is successful.

Test the network of the VM. The gateway is successfully pinged.

```
C:\Users\Administrator>ping 192.168.100.1

Pinging 192.168.100.1 with 32 bytes of data:
Reply from 192.168.100.1: bytes=32 time<1ms TTL=128
Reply from 192.168.100.1: bytes=32 time<1ms TTL=128
Reply from 192.168.100.1: bytes=32 time<1ms TTL=128

Ping statistics for 192.168.100.1:
    Packets: Sent = 3, Received = 3, Lost = 0 (0% loss),
```

----End

7.3.3 Mutual Access between Two Cloud Hosts Using the Same Subnet

In this experiment, the cloud host Host02 uses the same subnet as Host01 so that the two hosts can communicate with each other.

Step 1 **Apply for a cloud host.**

Use the username **vdcadmin** to log in to ServiceCenter and apply for another cloud host.

Choose Services > Cloud Host > ECS.

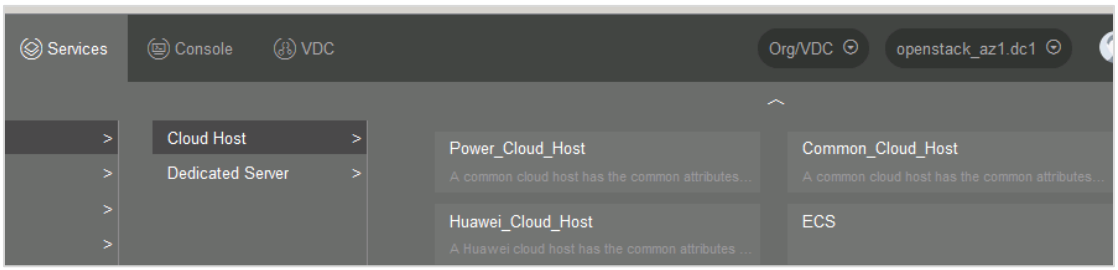

Click Apply for.

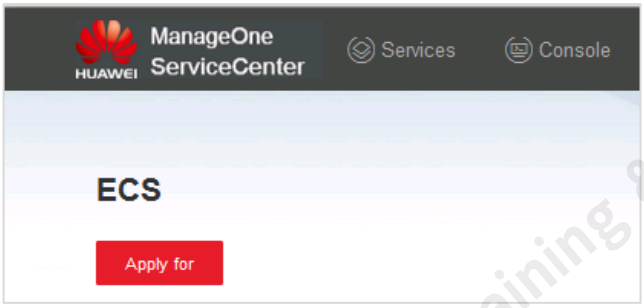

Submit the application following the same steps in section 7.3.2 "**Applying for a Cloud Host**".

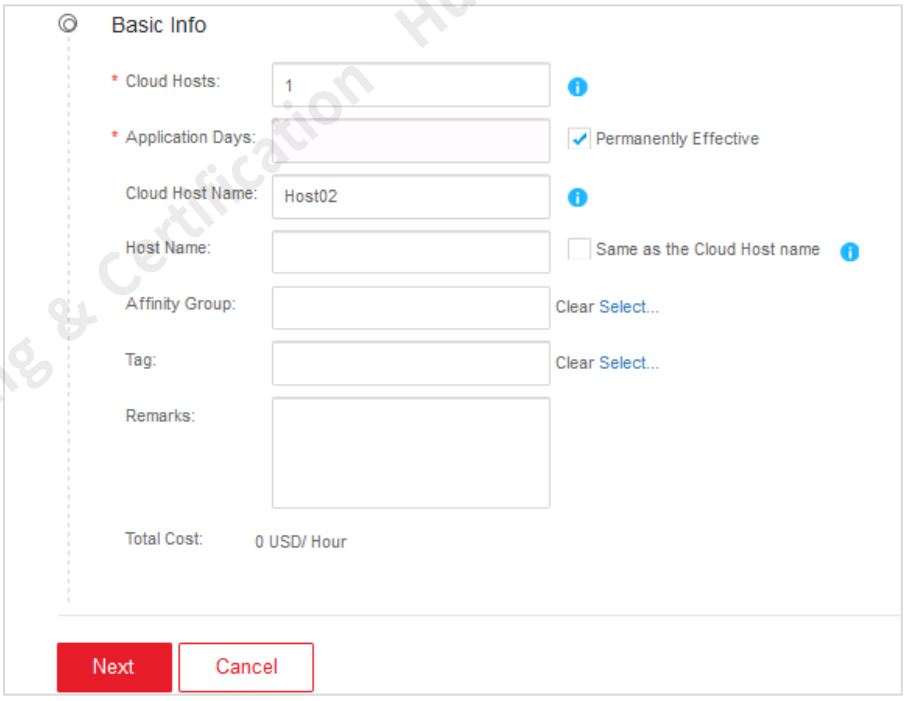

Step 2 **Approve a cloud host.**

Choose User Center > My To-Do List, and click Approve.

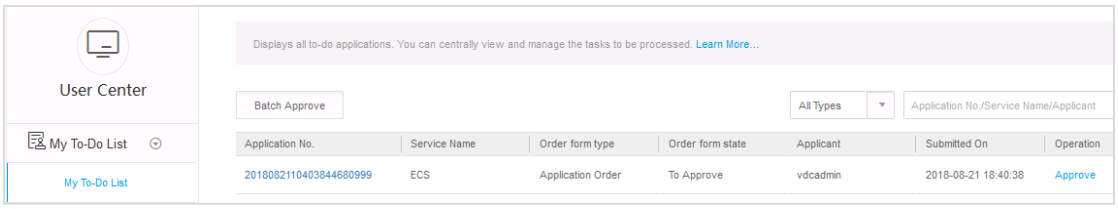

Select the same subnet for Host02 as that of Host01.

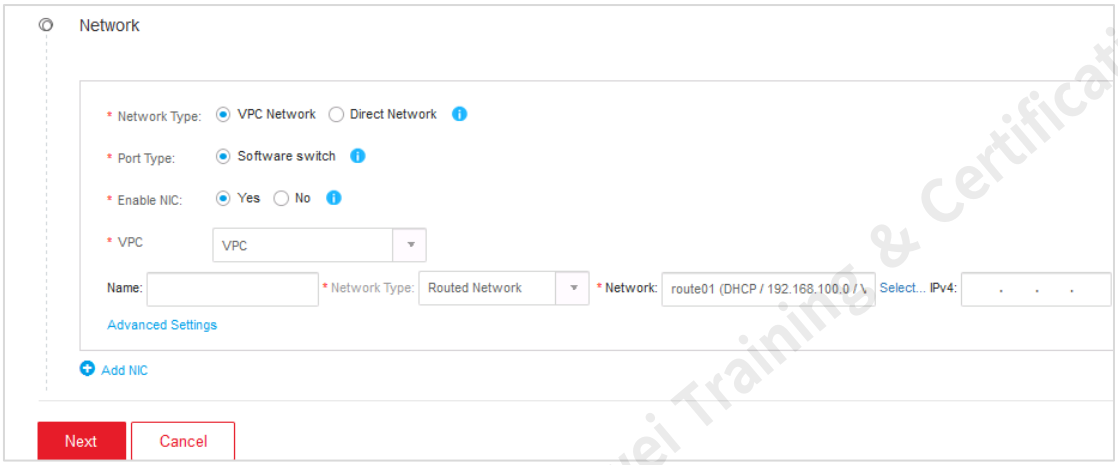

Create a VM following the same steps in section 7.3.2 "**Applying for a Cloud Host**".

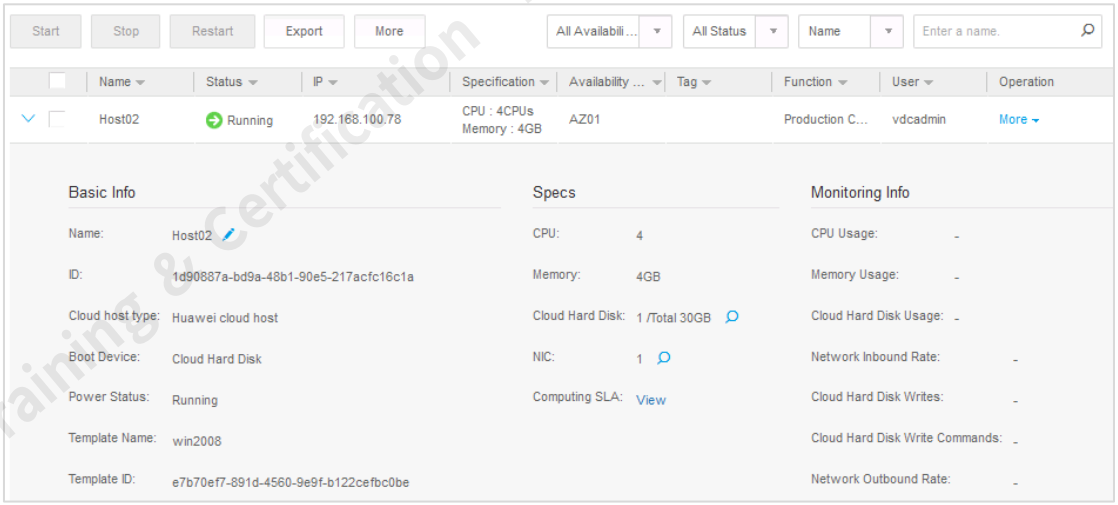

The DHCP address is 192.168.100.78. Log in to Host02 using VNC, and test the connectivity of Host01 (192.168.100.149 in this case) and Host02. (you need to disable the virtual firewall).

```
C:\Users\Administrator>ping 192.168.100.149

Pinging 192.168.100.149 with 32 bytes of data:
Reply from 192.168.100.149: bytes=32 time<1ms TTL=128
Reply from 192.168.100.149: bytes=32 time<1ms TTL=128
Reply from 192.168.100.149: bytes=32 time<1ms TTL=128
Reply from 192.168.100.149: bytes=32 time<1ms TTL=128

Ping statistics for 192.168.100.149:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

After the Host01 firewall is disabled, the test is successful.

----End

7.3.4 Mutual Access between Two Cloud Hosts Using Different Subnets

In this experiment, the Host03 uses a new routed network route02 (192.168.200.0/24) to communicate with cloud hosts in the network segment route01 (192.168.100.0/24).

Step 1 **Create a routed network route02.**

Choose VPC > Networks > Routed/Internal network, and click Create.

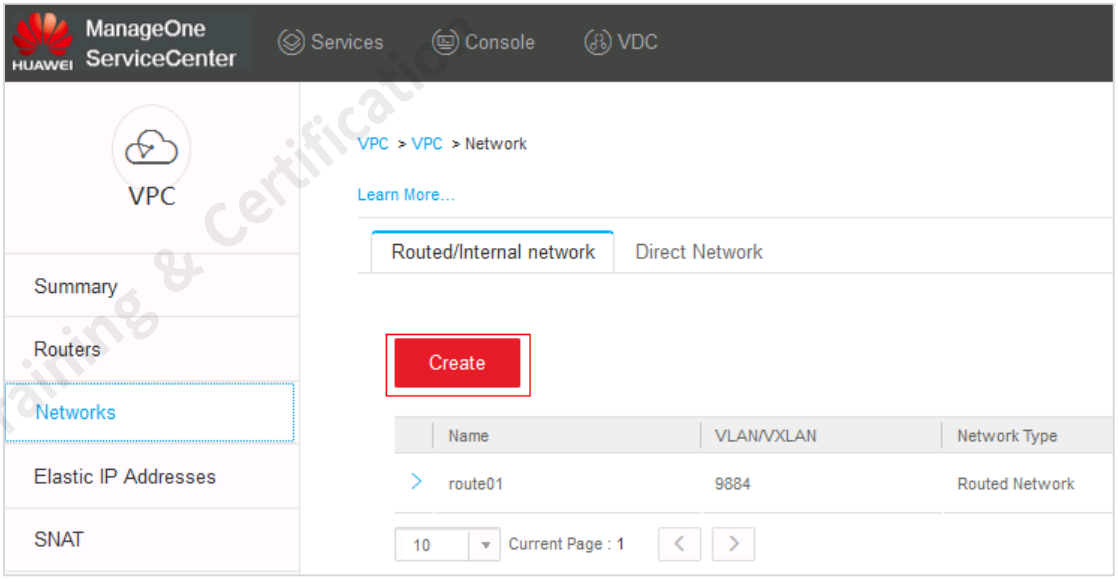

Create a routed network.

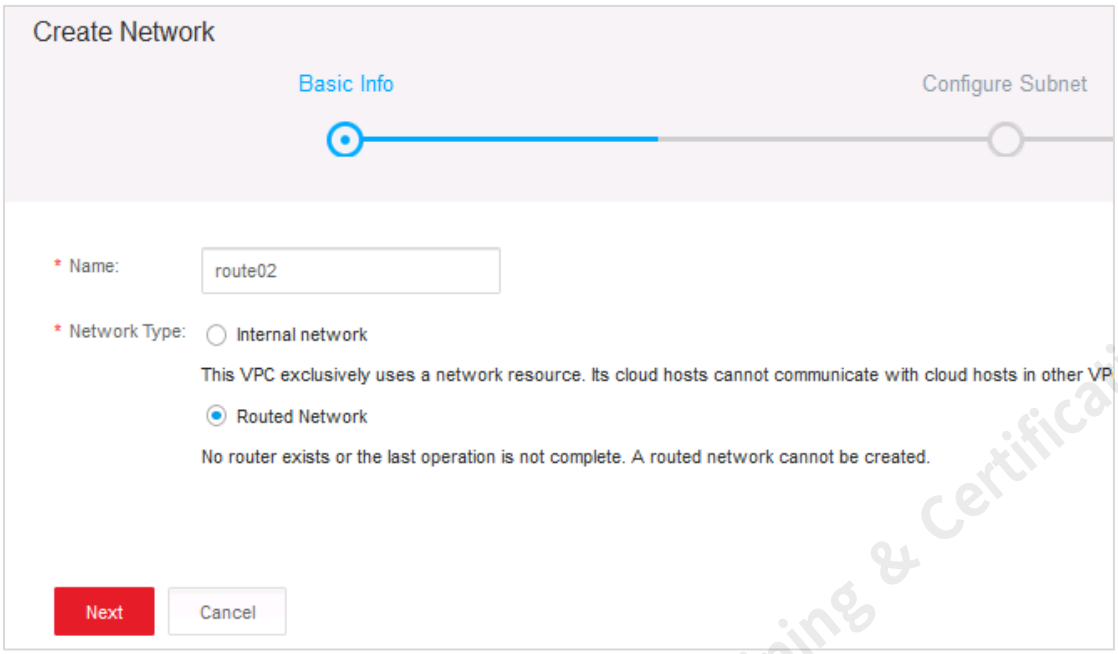

Enter the IP address of the network segment.

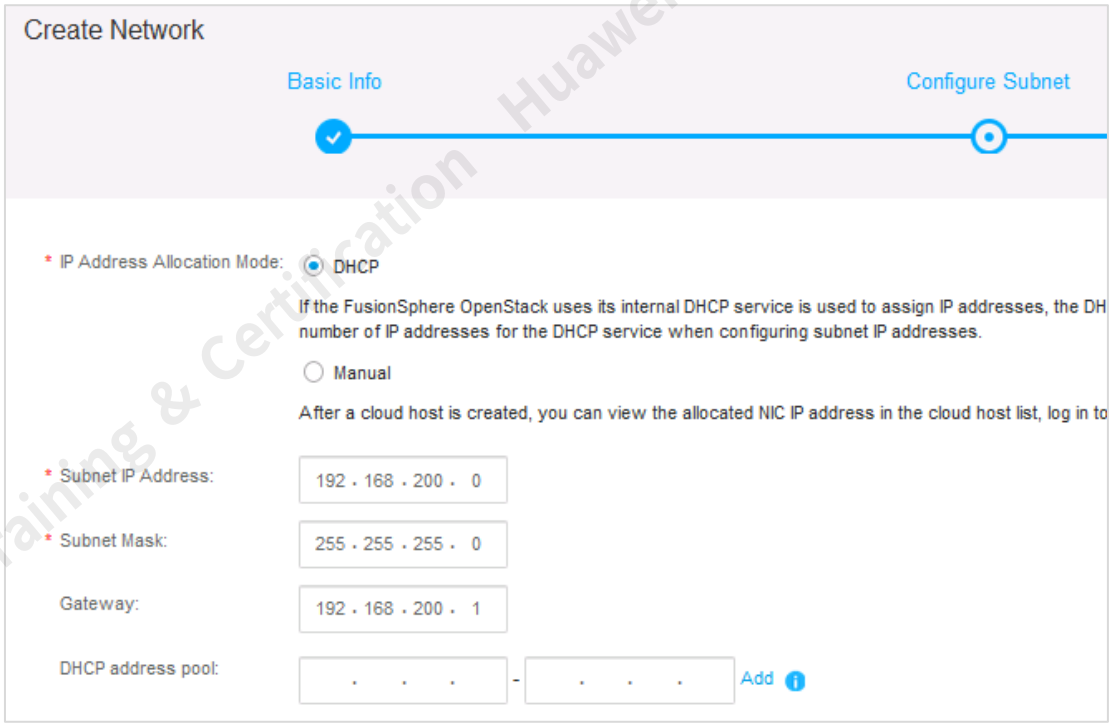

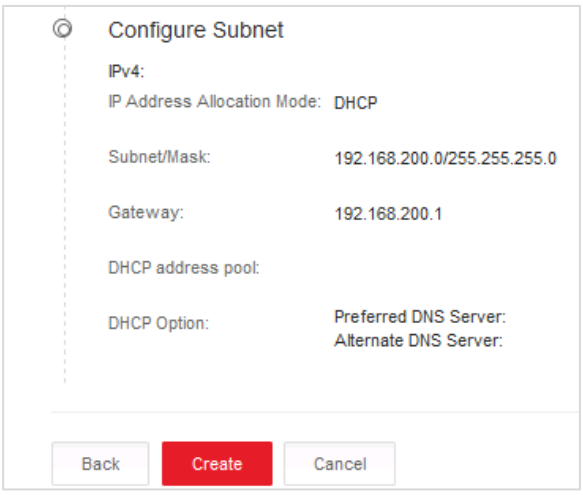

Click **Next** to complete the creation.

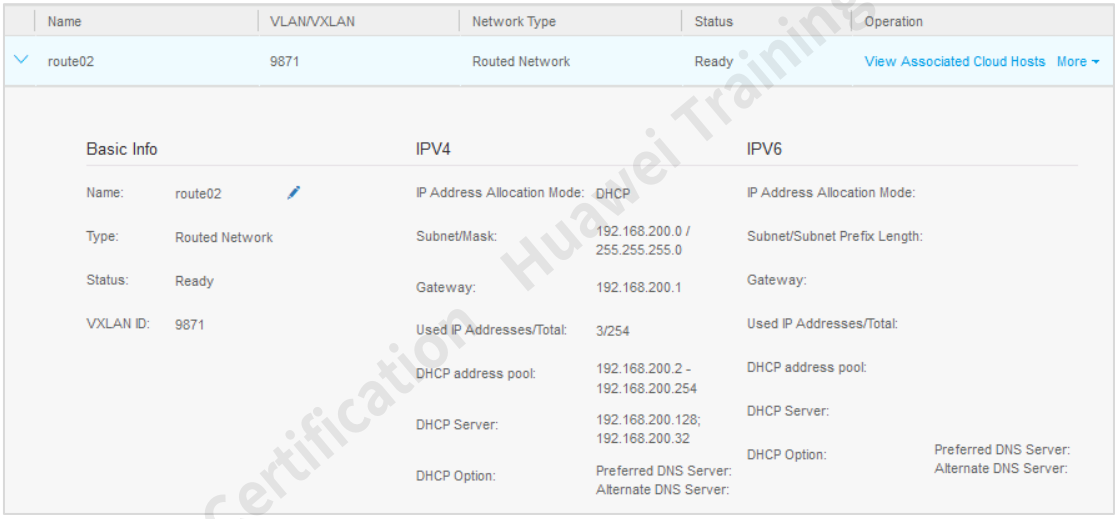

Step 2 **Apply for a cloud host.**

Apply for a cloud host following the same steps in section 7.3.2 "**Applying for a Cloud Host**".

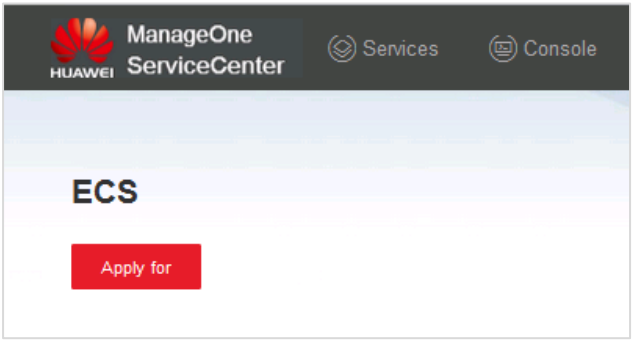

When approving Host03, select **route02** for the network type.

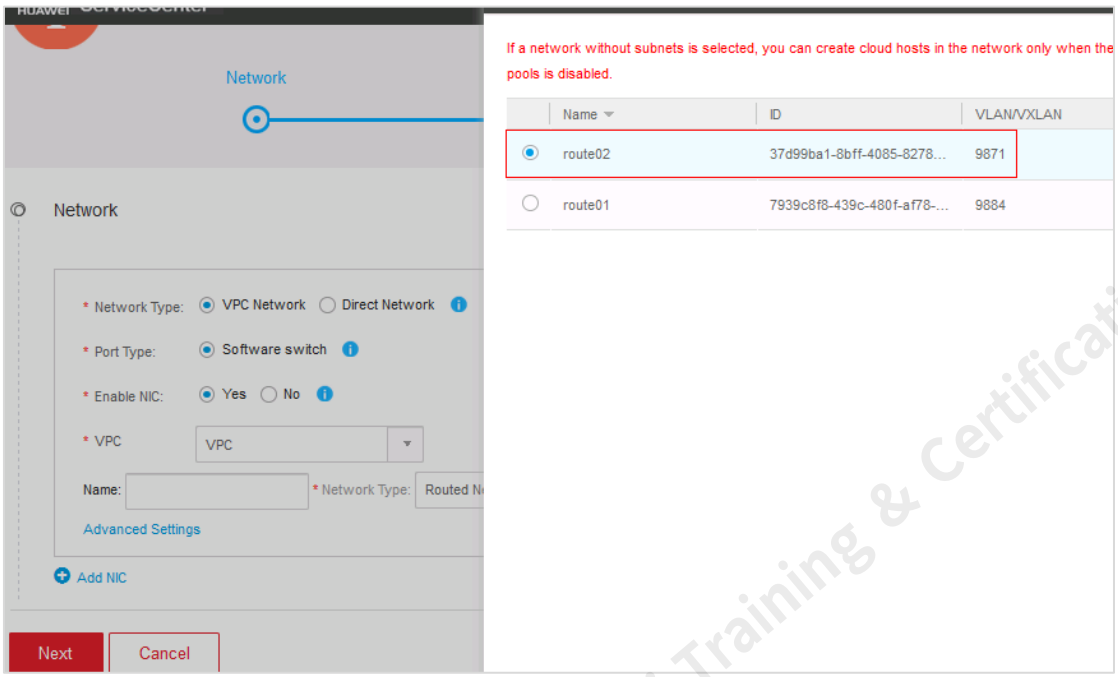

The cloud host is successfully delivered and the DHCP address is 192.168.200.87.

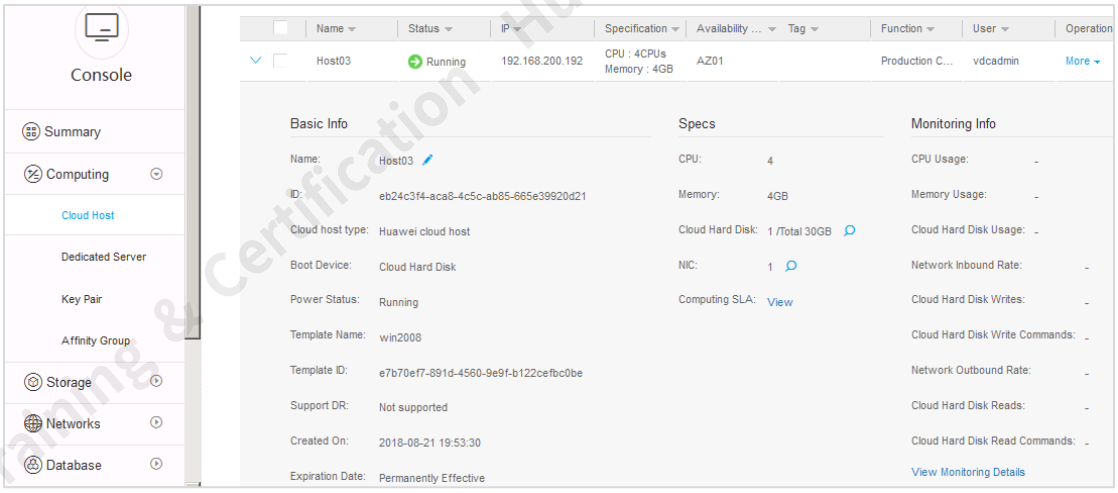

Step 3 **Test the service connectivity.**

User VNC to log in to Host03 to test the gateway connectivity.

```
C:\Users\Administrator>ping 192.168.200.1

Pinging 192.168.200.1 with 32 bytes of data:
Reply from 192.168.200.1: bytes=32 time<1ms TTL=128
Reply from 192.168.200.1: bytes=32 time<1ms TTL=128
Reply from 192.168.200.1: bytes=32 time<1ms TTL=128
Reply from 192.168.200.1: bytes=32 time<1ms TTL=128

Ping statistics for 192.168.200.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

Test the connectivity between Host01 and Host03 (192.168.100.149).

```
C:\Users\Administrator>ping 192.168.100.149

Pinging 192.168.100.149 with 32 bytes of data:
Reply from 192.168.100.149: bytes=32 time<1ms TTL=128
Reply from 192.168.100.149: bytes=32 time<1ms TTL=128
Reply from 192.168.100.149: bytes=32 time<1ms TTL=128
Reply from 192.168.100.149: bytes=32 time<1ms TTL=128

Ping statistics for 192.168.100.149:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

The current VPC network topology is as follows:

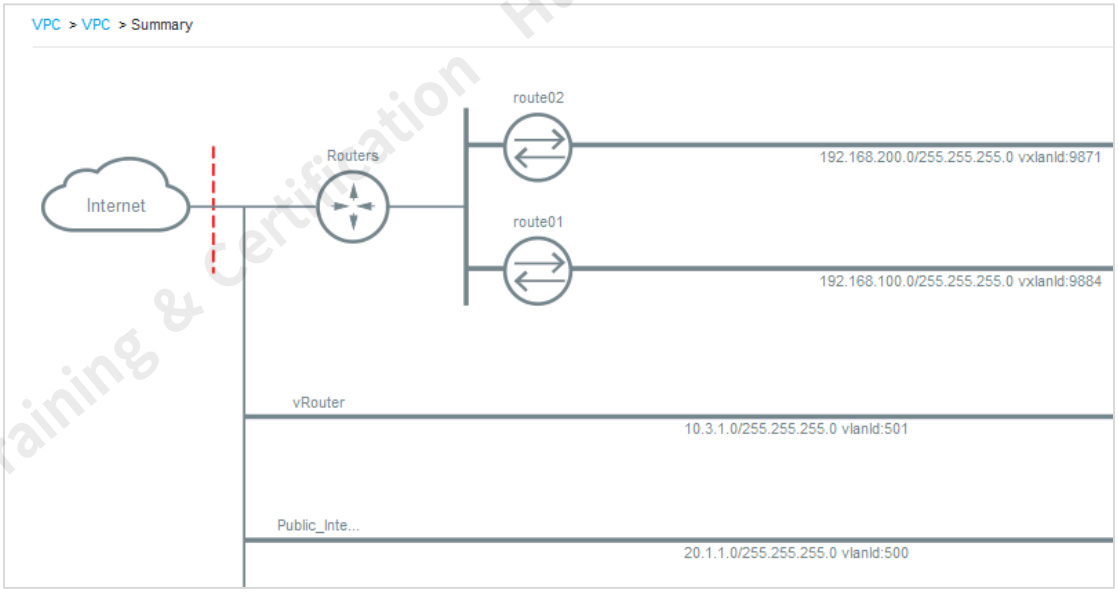

route01 and route02 are routed networks and can communicate with each other through the route. For details about how to perform access control on the cloud host, see section 7.5 "**Security Group**".

On the AC-DCN, the VPC logical network is as follows:

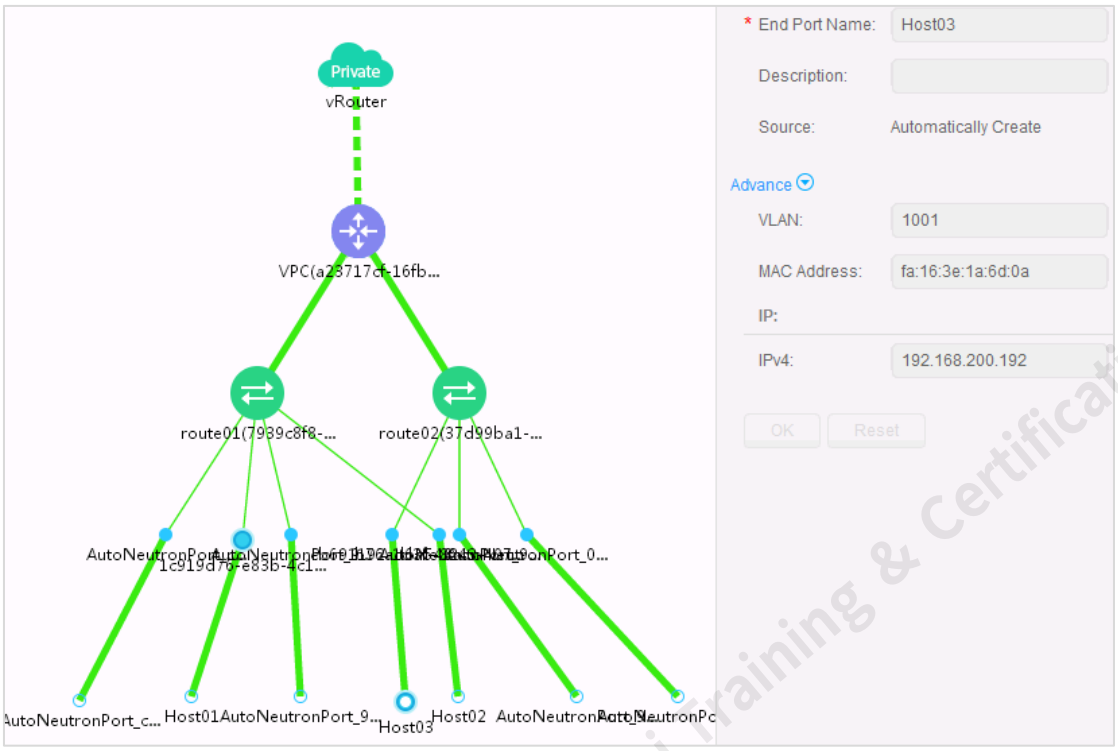

----End

7.4 Commissioning SNAT

The SNAT service helps VDC VMs obtain public network access rights. The physical firewall provides the SNAT service for cloud-network integration. The cloud platform delivers service requirements to the USG firewall through the Agile Controller-DCN.

In this experiment, you can apply for the SNAT service so that the subnet (192.168.100.0/24) where the Host01 is located can obtain the address of the Public_Internet network segment (20.1.1.0/24) and access the public network 30.1.1.1/32.

7.4.1 Applying for SNAT

Step 1 **Log in to ServiceCenter as the VDC administrator.**

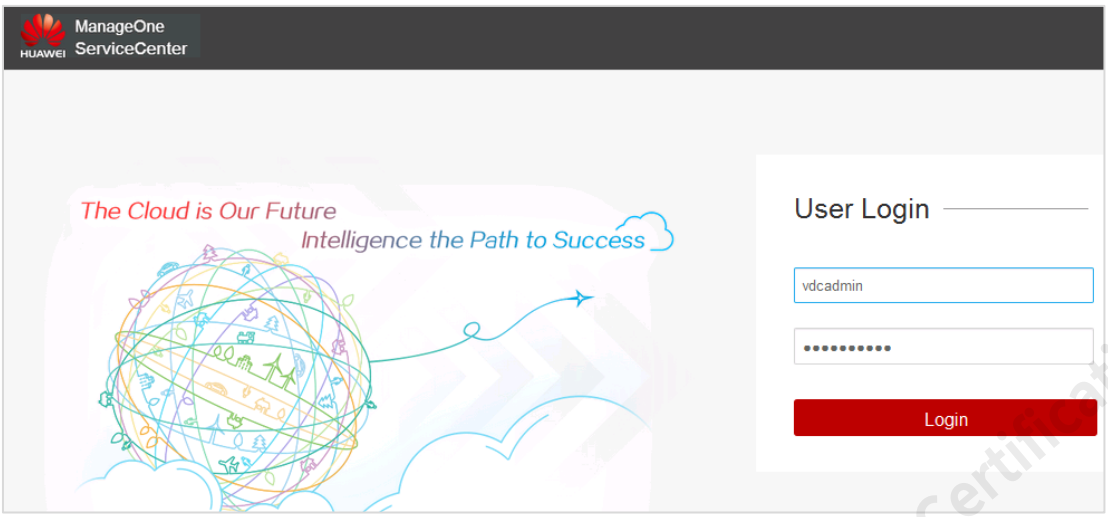

Step 2 **Apply for the SNAT service.**

Choose Console > Networks > VPCs.

Click **SNAT** on the **VPC** page.

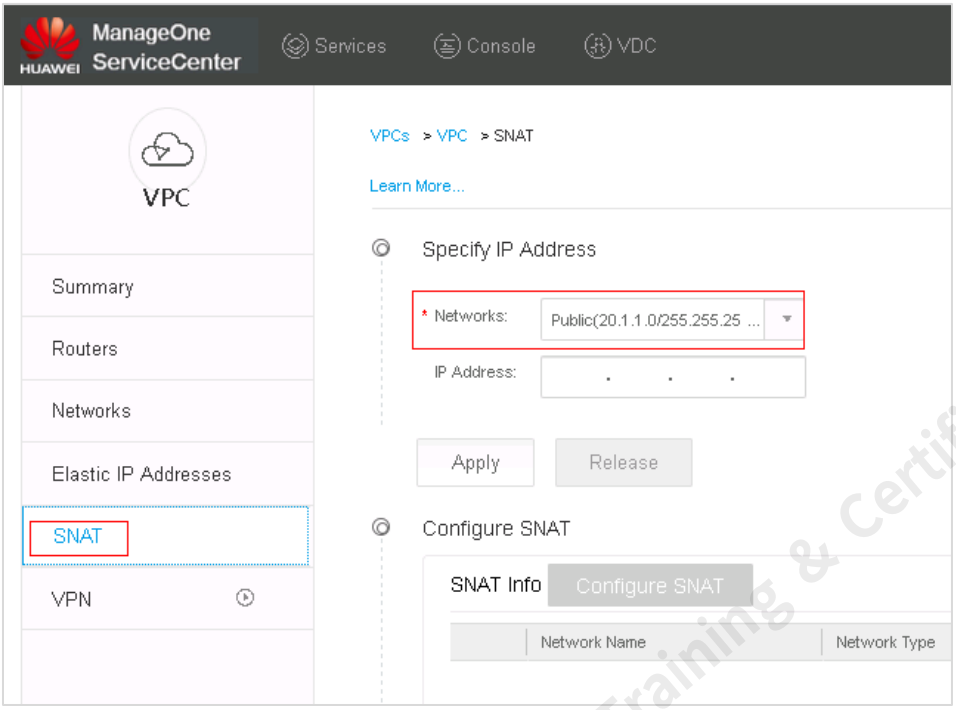

The network is auto-selected. Click **Apply**.

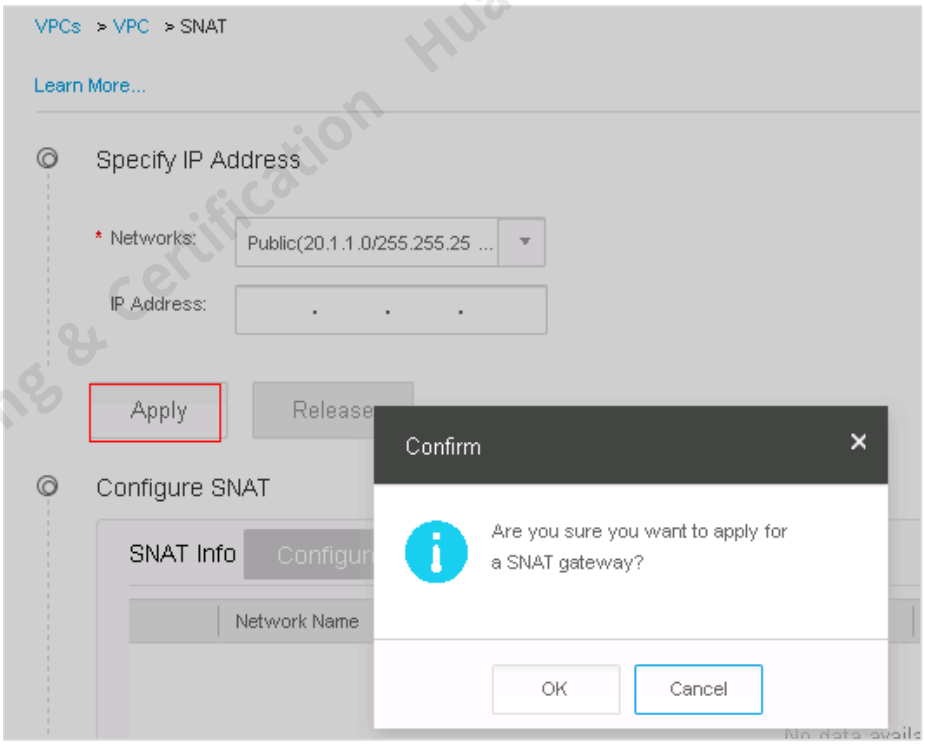

Click Configure SNAT, and select Enabled for Enable SNAT.

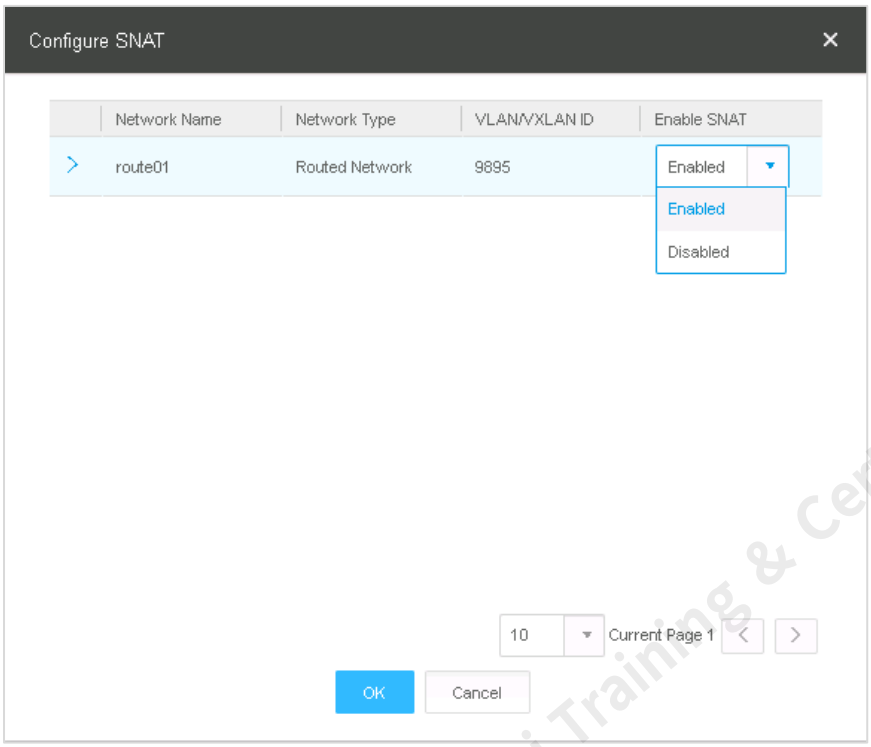

The configuration is successful.

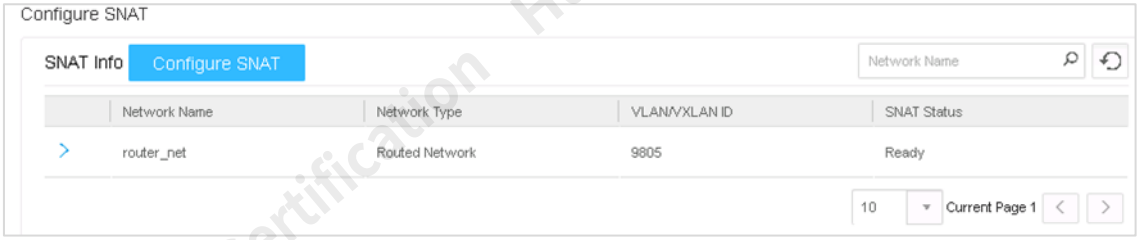

Step 3 **Check the configurations delivered by the Agile Controller-DCN.**

Check the operation logs on the Agile Controller-DCN.

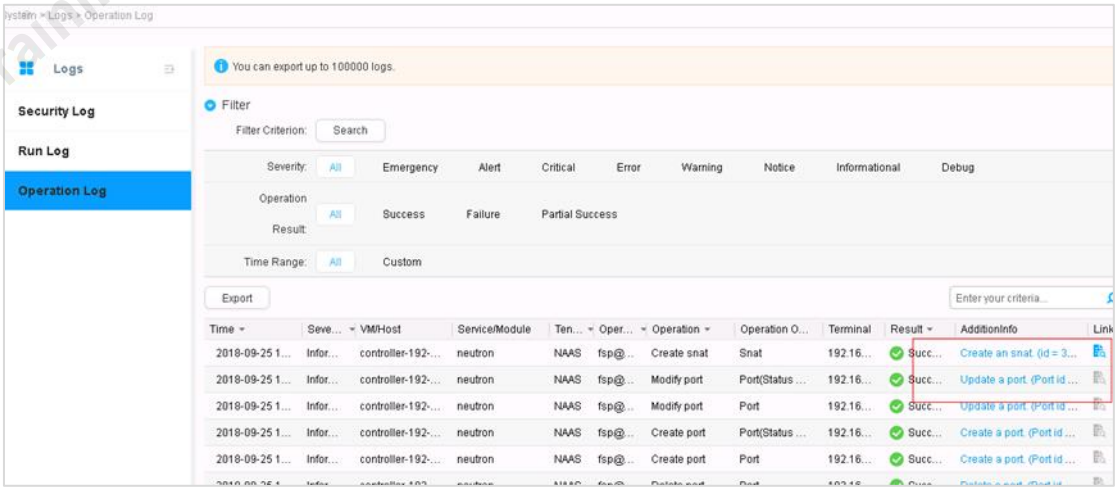

The virtual firewall is successfully created. Check the security policy and the VXLAN configuration of the Fabric network delivered by the Agile Controller-DCN.

Config Record

Enter the device name or device IP | Please select result

Configuration Time	Device Name	Device IP	Operation Description	Message Body	Result
2018-09-25 10:49:57	Border-2	172.21.22.12	create BD 5002.	View	Succeed
2018-09-25 10:49:57	Border-2	172.21.22.12	bind vni 10009 to BD 5002.	View	Succeed
2018-09-25 10:49:57	Border-2	172.21.22.12	create evpn instance 5:100...	View	Succeed
2018-09-25 10:49:57	Border-2	172.21.22.12	enable arp bd 5002.	View	Succeed
2018-09-25 10:49:57	Border-1	172.21.22.11	create BD 5002.	View	Succeed
2018-09-25 10:49:57	Border-1	172.21.22.11	bind vni 10009 to BD 5002.	View	Succeed
2018-09-25 10:49:57	Border-1	172.21.22.11	create evpn instance 6:100...	View	Succeed
2018-09-25 10:49:57	Border-1	172.21.22.11	enable arp bd 5002.	View	Succeed
2018-09-25 10:49:57	Border-2	172.21.22.12	create interface Eth-Trunk2...	View	Succeed
2018-09-25 10:49:57	Border-2	172.21.22.12	set if Eth-Trunk2.202 acces...	View	Succeed
2018-09-25 10:49:57	Border-2	172.21.22.12	bind bd 5002 to if Eth-Trun...	View	Succeed
2018-09-25 10:49:57	Border-2	172.21.22.12	create interface Eth-Trunk3...	View	Succeed
2018-09-25 10:49:57	Border-2	172.21.22.12	set if Eth-Trunk3.202 acces...	View	Succeed
2018-09-25 10:49:57	Border-2	172.21.22.12	bind bd 5002 to if Eth-Trun...	View	Succeed
2018-09-25 10:49:57	Border-1	172.21.22.11	create interface Eth-Trunk2...	View	Succeed
2018-09-25 10:49:57	Border-1	172.21.22.11	set if Eth-Trunk2.202 acces...	View	Succeed
2018-09-25 10:49:57	Border-1	172.21.22.11	bind bd 5002 to if Eth-Trun...	View	Succeed
2018-09-25 10:49:57	Border-1	172.21.22.11	create interface Eth-Trunk3...	View	Succeed
2018-09-25 10:49:57	Border-1	172.21.22.11	set if Eth-Trunk3.202 acces...	View	Succeed

Check whether the configuration of the firewall is the same as that of the elastic IP address applied on ServiceCenter.

```
vsys enable
resource-class r0
#
#
vsys name vsys_2dba1437_VPC_10004 1
 assign vlan 3002
 assign global-ip 20.1.1.141 20.1.1.141 free
#
ip vpn-instance default
 ipv4-family
#
ip vpn-instance vsys_2dba1437_VPC_10004
 ipv4-family
#
```

----End

7.4.2 Configuring the Firewall Policy

As the default policy of the firewall is denying all traffic, you need to release the traffic from the cloud host to the public network.

In this example, the public IP address is 30.1.1.1/32. Click **ACL Management**.

Click **Add Rule**, enter the source IP address and the destination IP address, and click **Create**.

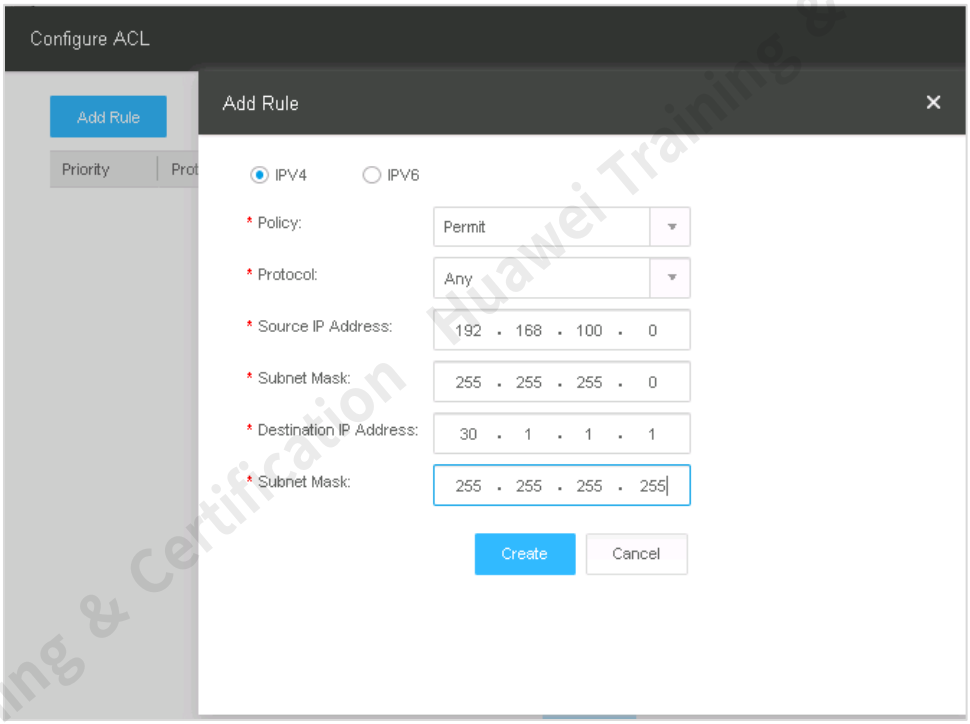

The creation is successful.

Test the connectivity of the public network. The public network can be connected.

```
PS C:\Users\Administrator> ping 30.1.1.1

Pinging 30.1.1.1 with 32 bytes of data:
Reply from 30.1.1.1: bytes=32 time<1ms TTL=128
Reply from 30.1.1.1: bytes=32 time<1ms TTL=128
Reply from 30.1.1.1: bytes=32 time<1ms TTL=128
Reply from 30.1.1.1: bytes=32 time<1ms TTL=128

Ping statistics for 30.1.1.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

7.5 Security Group

A security group implements access control for the cloud hosts, enhancing the security of the cloud hosts. The VDC administrator can define various access control rules for a security group, and these rules take effect for all cloud hosts added to this security group.

7.5.1 Creating a Security Group

Use the username **vdcadmin** to log in to ServiceCenter.

Choose Console > Networks > Default security group, and click Create.

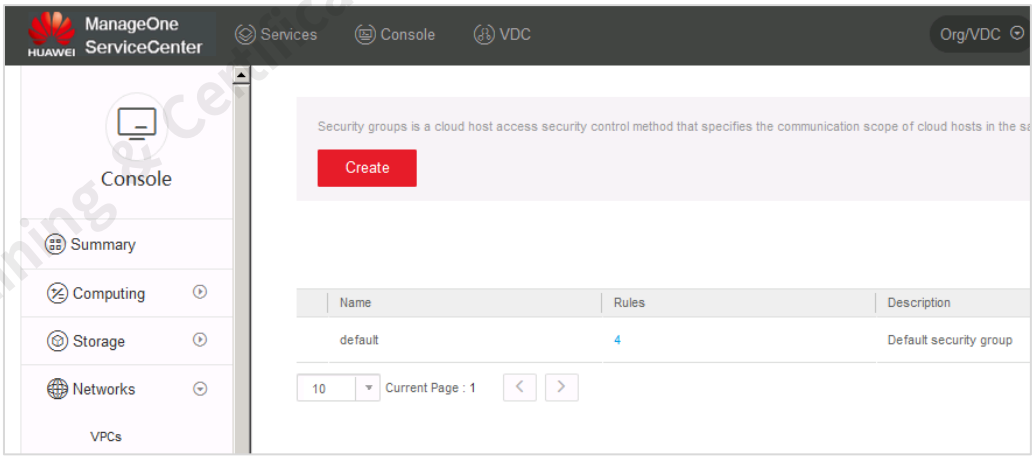

Enter the security group name **SG01**.

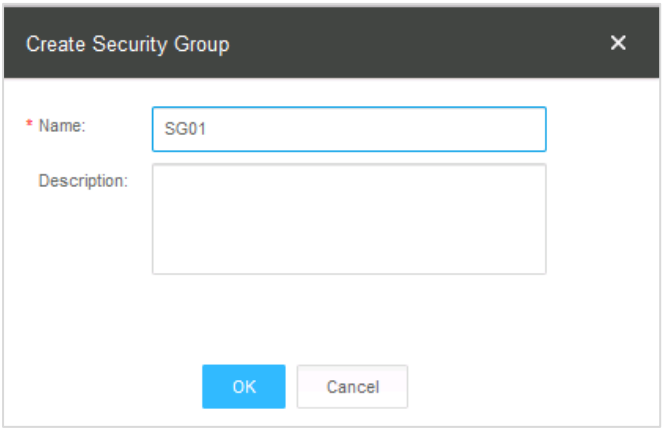

7.5.2 Adding a Cloud Host to a Security Group

Choose **Console** > **Computing** > **Cloud Host**, and check the basic information of the cloud host.

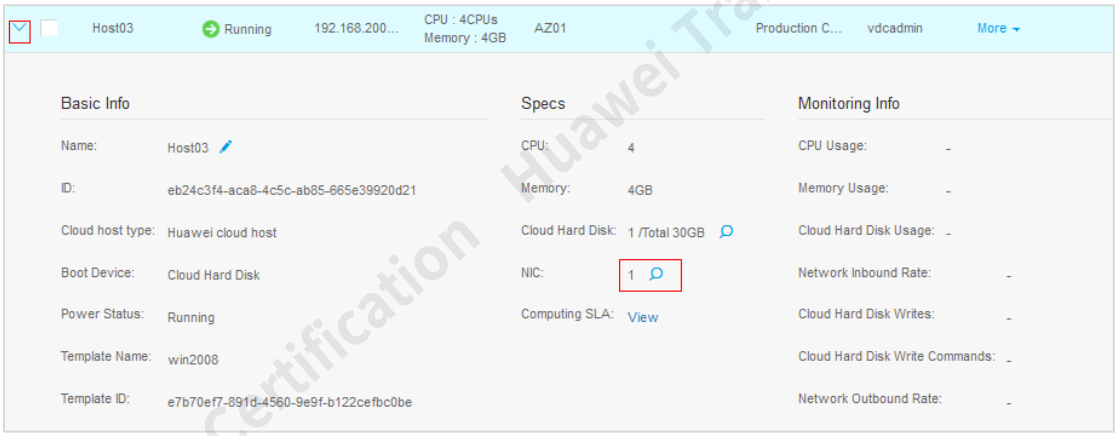

On the **NIC-Host03** page, click **More** and select **Join Security Group**.

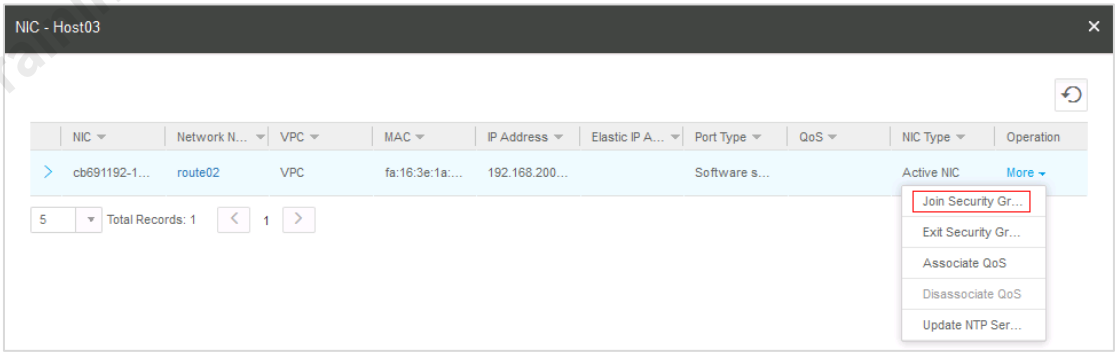

Select SG01 and click **OK**.

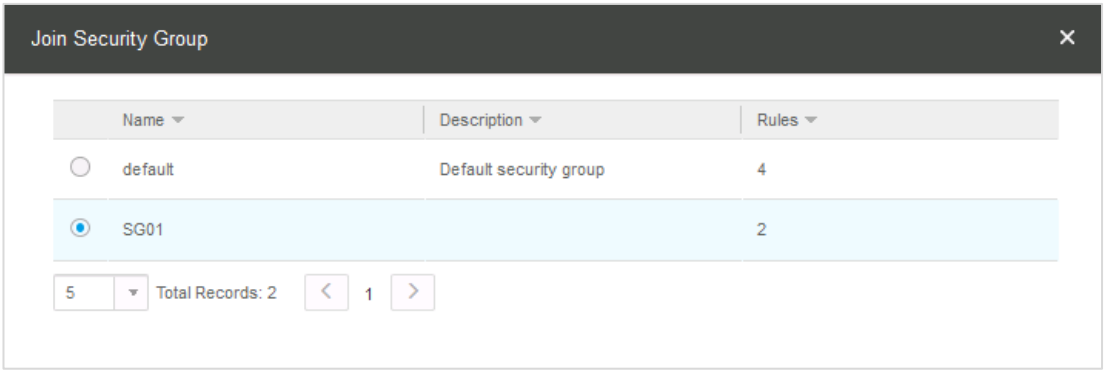

Host03 is successfully added into the security group. Configure the Host03 to exit the default security group.

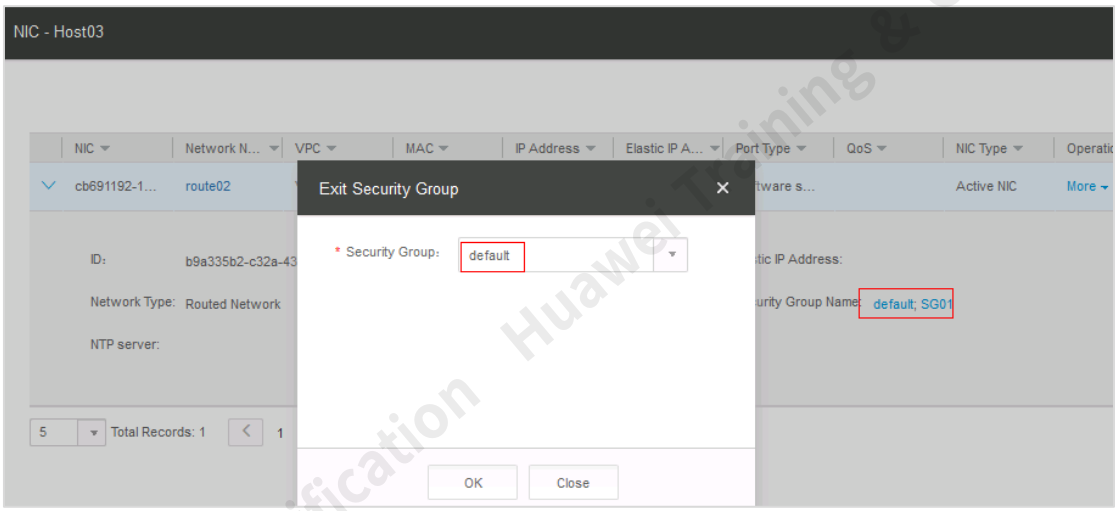

Host01 and Host03 cannot communicate with each other.

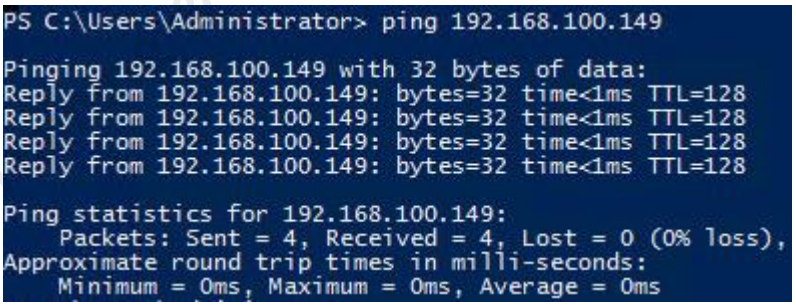

7.5.3 Adding a Security Group Rule

By default, **Default security group** has four rules (two for IPv4 and two for IPv6).

- Cloud hosts added to **Default security group** can access cloud hosts in any other subnets or security groups.

- Cloud hosts added to **Default security group** do not allow access from cloud hosts in any other subnets or security groups.

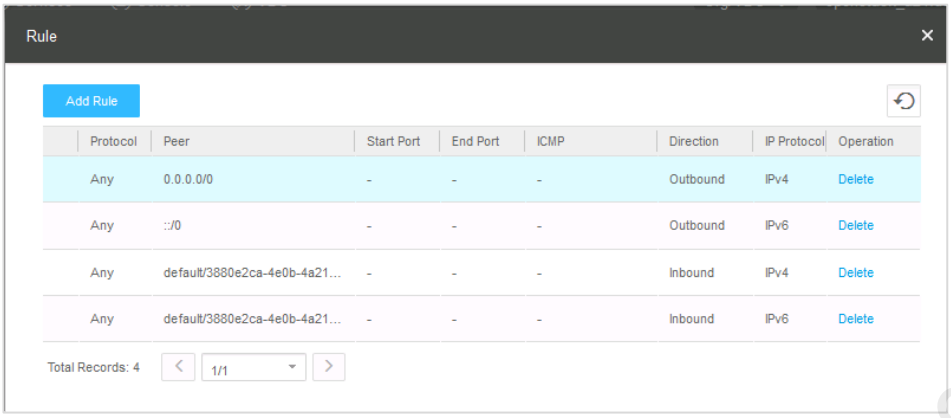

Click Add Rule.

Allow the traffic of Host03 to access **Default security group**.

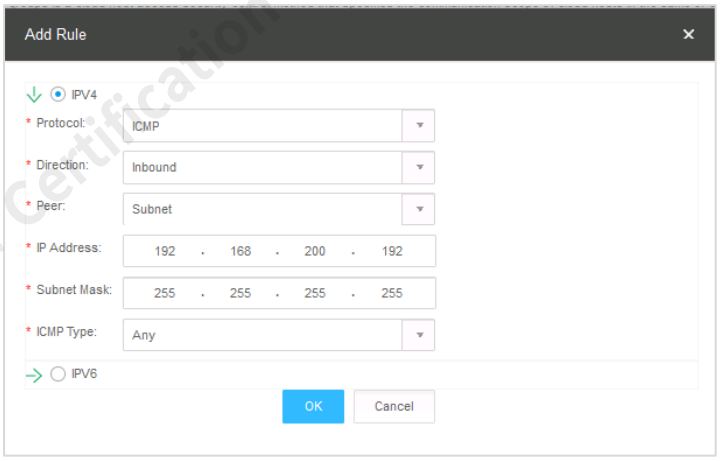

Host03 and Host01 can communicate with each other.

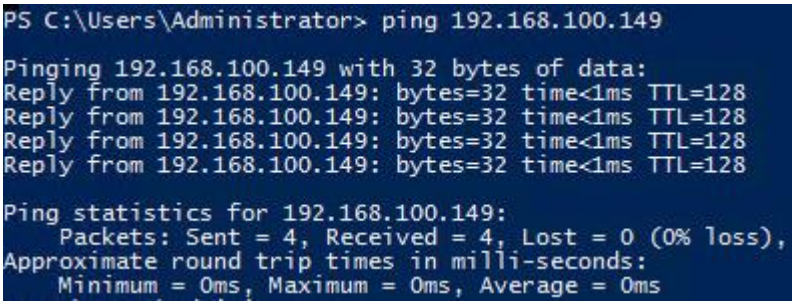

8 Experiment on Cloud Host Backup in Cloud Data Centers

8.1 Objectives

- To have a good command of the networking of the Huawei FusionCloud backup solution
- To master the FusionCloud backup principles and operations

8.2 Networking and Service Description

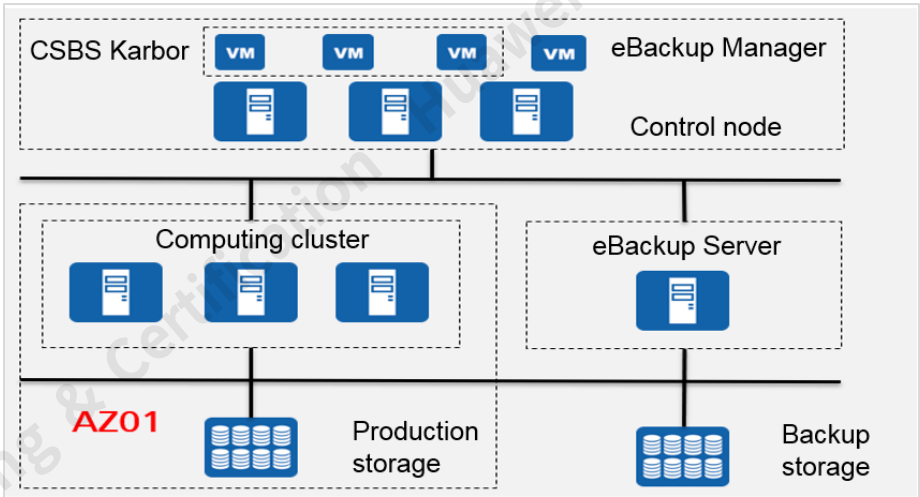

This experiment applies to the FusionCloud backup scenarios. The backup components include CSBS Karbor, eBackup Manager, and eBackup Server. Karbor provides orchestration and scheduling capabilities for CSBS. eBackup Manager receives service requests such as backup and restoration requests from Karbor. eBackup Server directly receives requests from users and executes backup tasks.

AZ01 is an AZ of FusionSphere OpenStack. It contains a computing cluster that consists of three CNAs and FusionStorage production storage. The backup storage is the NAS provided by OceanStor V3.

This experiment is conducted to implement backup and recovery of a cloud host in AZ01 by using BCManager eBackup.

8.3 Planning Interconnection Data

During the experiment, the backup server and OpenStack need to be interconnected. The following table lists interconnection parameters.

Category	Parameter Name	Value
-	Protocol	HTTPS
Authentication configuration	Type	Keystone
	URL	identity.az1.dc1.domainname.com:443
	Username	Karbor
	Password	CloudService@123!
Nova configuration	URL	compute.az1.dc1.domainname.com:443
Cinder configuration	URL	volume.az1.dc1.domainname.com:443
Neutron configuration	URL	network.az1.dc1.domainname.com:443

8.4 Configuring eBackup Manager

Step 1 **Conduct pre-configuration.**

Install the eBackup Server, and configure the network plane, OpenStack domain name information, and system time.

Step 2 **Perform OpenStack interconnection.**

Log in to the management page of FusionInsight Manager, choose **Settings** > **Configuration** > **OpenStack Settings**, and set parameters based on the plan.

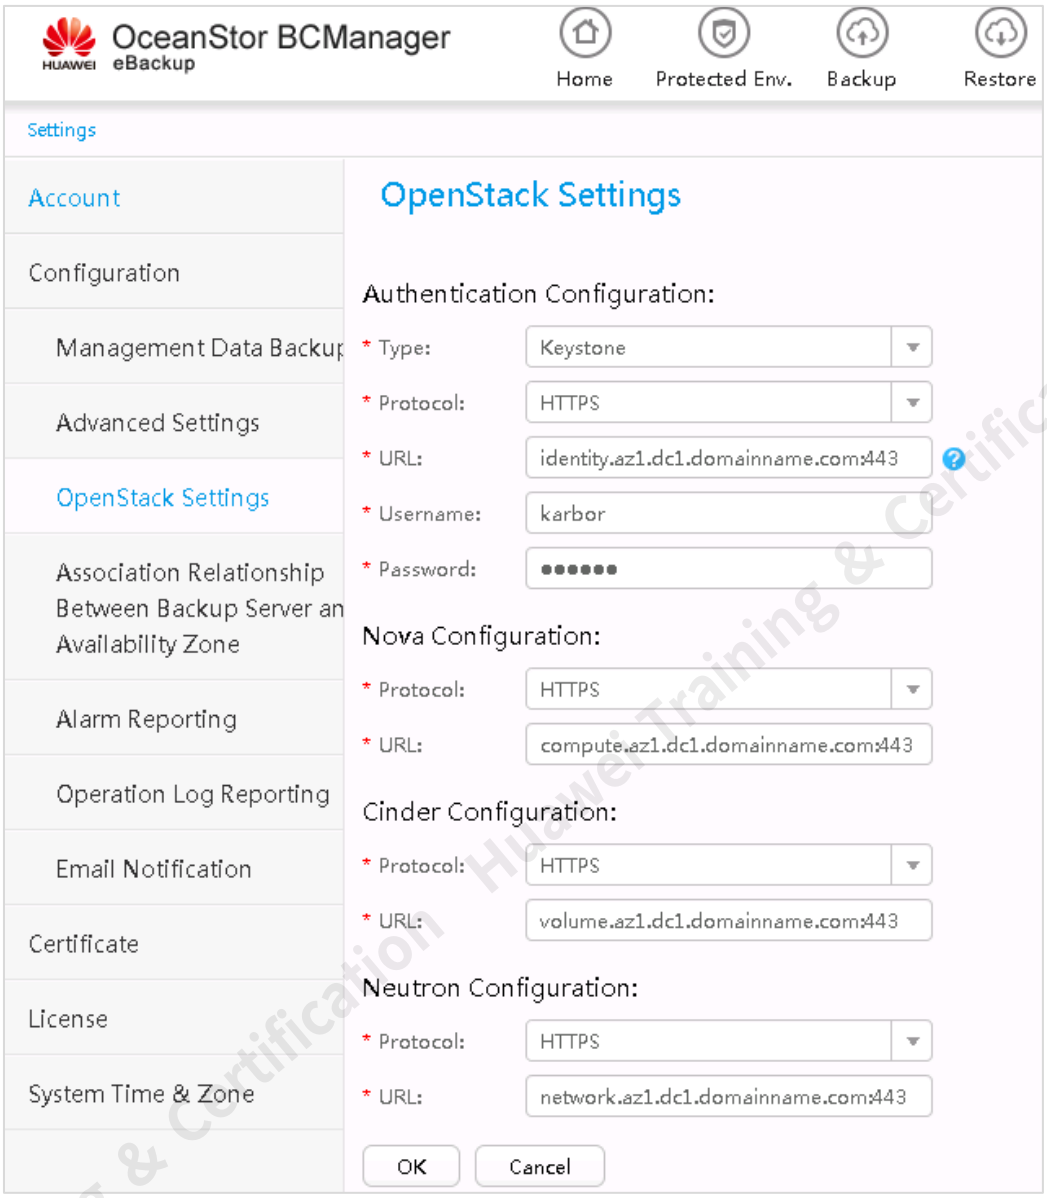

Step 3 **Configure the association relationship between Manager and the AZ.**

Ensure that the association between the server and AZ is the same as that between the server and OpenStack, facilitating backup service provisioning from eBackup to AZs.

Choose Settings > Configuration > Association Relationship Between Backup Server and Availability Zone > Add Association Relationship.

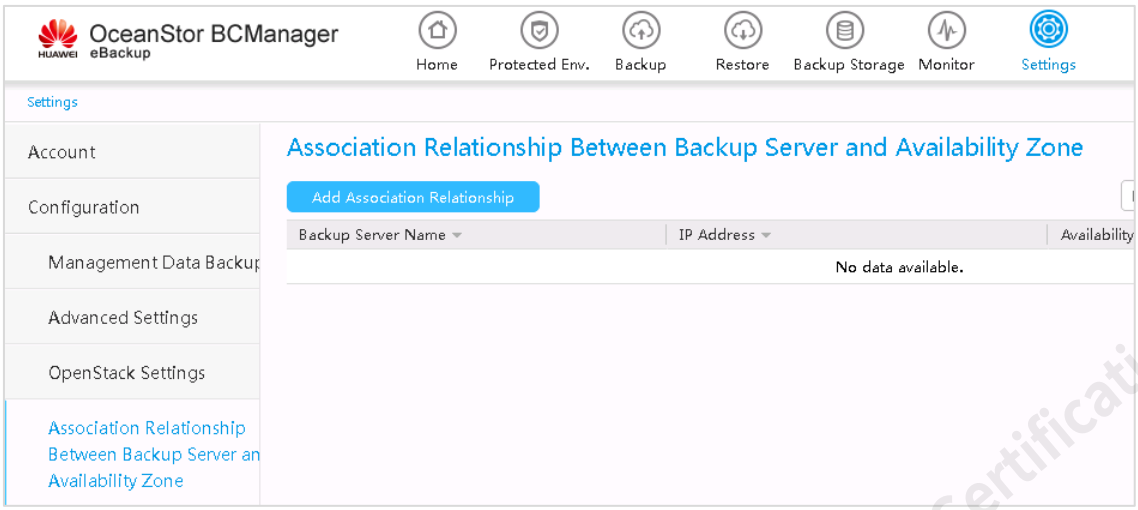

In the **Add Association Relationship** dialog box, enter the default username (**NBIUser**) and password (**Huawei@CLOUD8!**). For **Availability Zone**, select one or more AZs of OpenStack. For **IP Address**, enter the IP address of the backup management plane of the server.

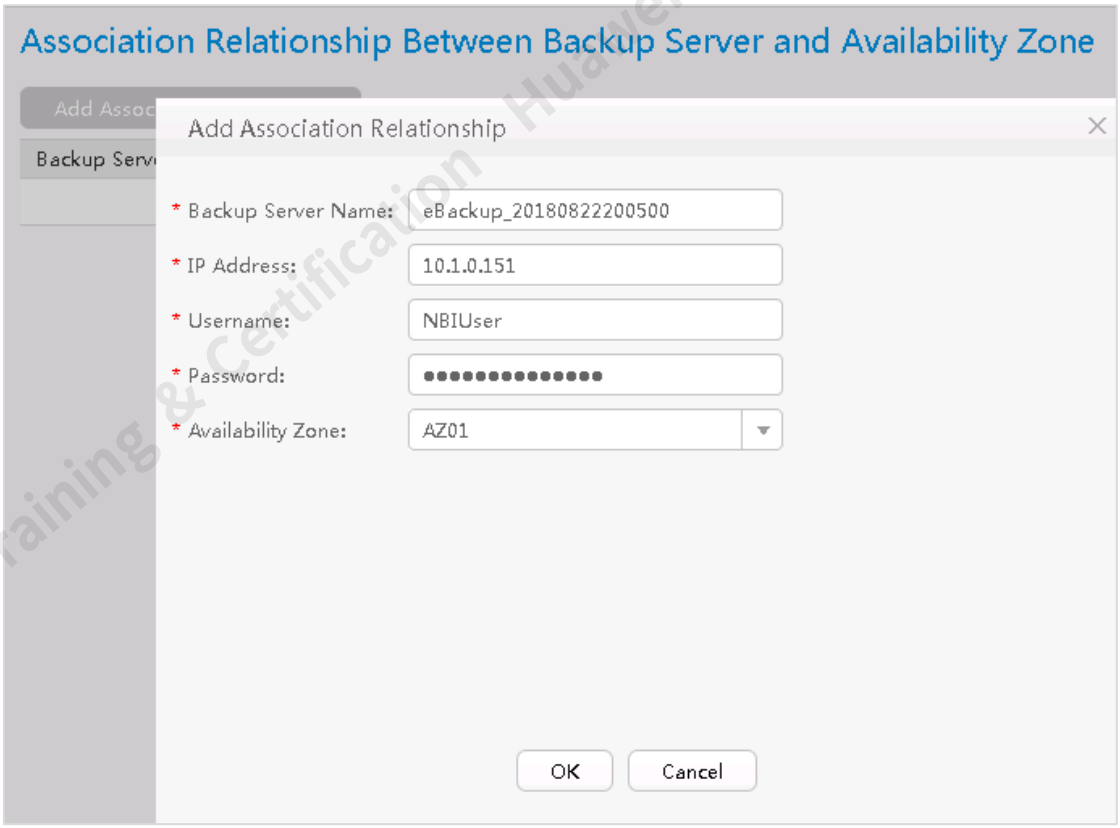

The association relationship is added.

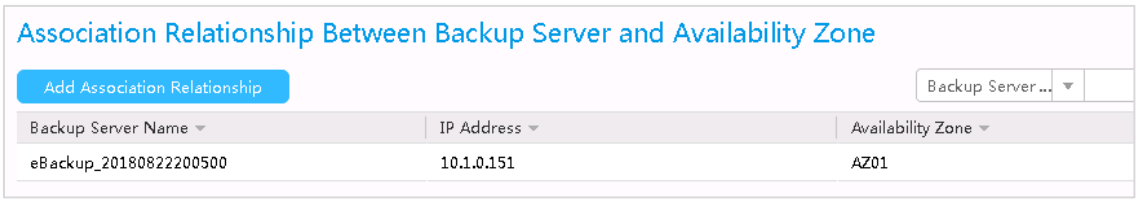

----End

8.5 Configuring Production Storage

If the production storage type planned on the eBackup management system is FusionStorage, add eBackup Server to the FusionStorage cluster (eBackup Manager does not need to be added to the cluster) and use the LAN-Free backup networking mode.

8.5.1 Configuring eBackup Server

Complete eBackup Server installation and network plane configuration.

Check the system dependency package and run the **rpm -qa |grep sysstat** command to check whether the sysstat component has been installed.

Step 1 **Log in to the eBackup server as user root.**

Log in to the eBackup server and switch to the **root** user.

Step 2 **Configure the FusionStorage management plane so that it can pass through the firewall.**

Run the **cd /opt/huawei-data-protection/ebackup/bin** command to go to the directory where the **iptables** script resides.

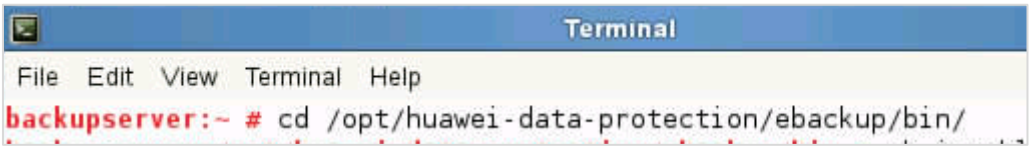

Run the sh iptablesHelper.sh accept ProtectedEnvironmentManagementPlane Management plane IP address of FusionStage (network segment) Production- and management-plane IP address of eBackup (network segment) command to add the management plane networks of FusionStorage to firewall rules of the eBackup server.

```
backupserver:/opt/huawei-data-protection/ebackup/bin # sh iptablesHelper.sh acce
pt ProtectedEnvrionmentManagementPlane 10.1.0.0/24 10.1.0.0/24
Usage:
    iptablesHelper.sh init
```

Step 3 **Configure the production storage plane so that it can pass through the firewall.**

Run the sh iptablesHelper.sh accept StoragePlane Storage plane IP address of FusionStage (network segment) Production- and storage-plane IP address of eBackup (network segment) command to add the storage plane networks of FusionStorage to firewall rules of the eBackup server.

```
backupserver:/opt/huawei-data-protection/ebackup/bin # sh iptablesHelper.sh acce
pt StoragePlane 10.11.0.0/24 10.11.0.0/24
Run command: iptables -I INPUT -s 10.11.0.0/24 -d 10.11.0.0/24 -p icmp -j ACCEPT
```

----End

8.5.2 Installing FusionStorage Agent on eBackup Server

Install FusionStorage Agent on the eBackup Server to invoke production storage.

Step 1 **Add eBackup Server as the server.**

Log in to FSM, click **Hardware**, and add eBackup Server.

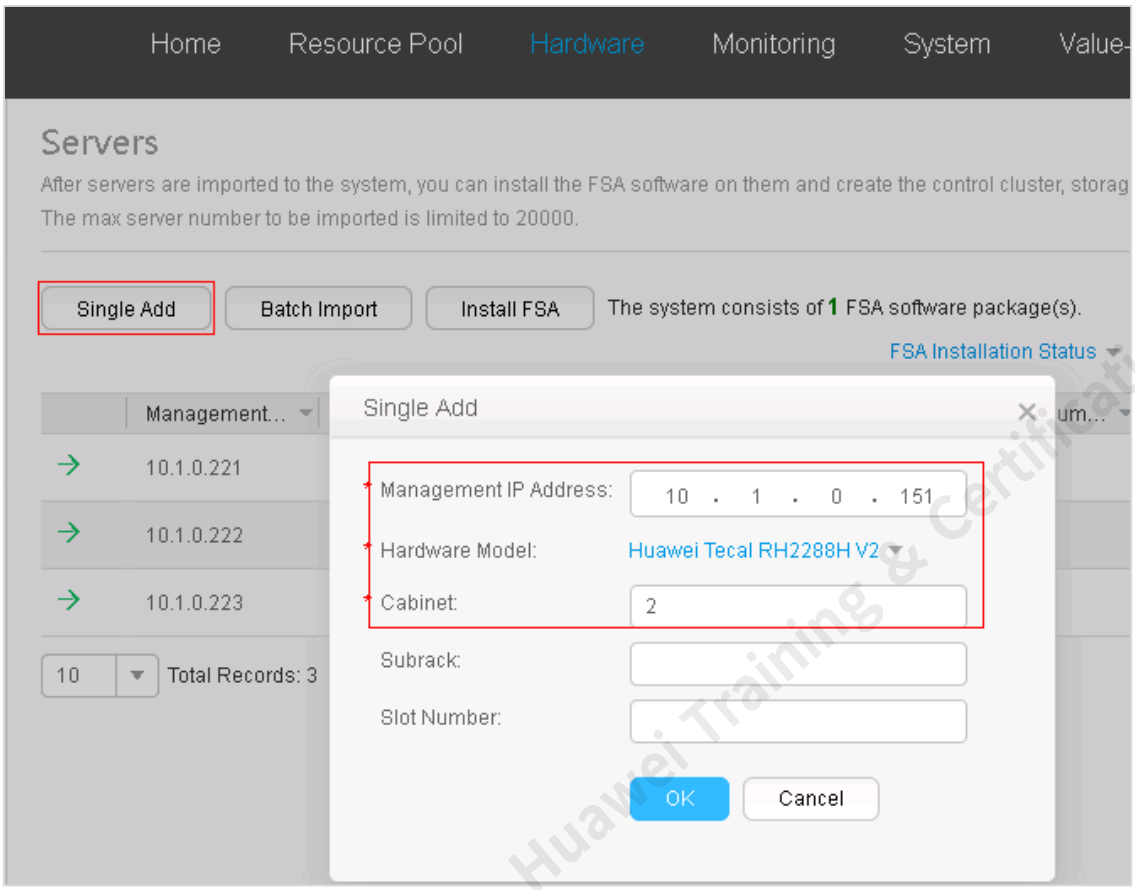

Step 2 **Install FusionStorage Agent.**

Install FSA by referring to section 5.3.2 "**Installing FSA**". The default username for logging in to eBackup is **hcp** and the default password is **Huawei@CLOUD8**.

Step 3 **Create block storage clients.**

Create a block storage client for the eBackup server. Choose **Resource Pool** > **Block Storage Clients**.

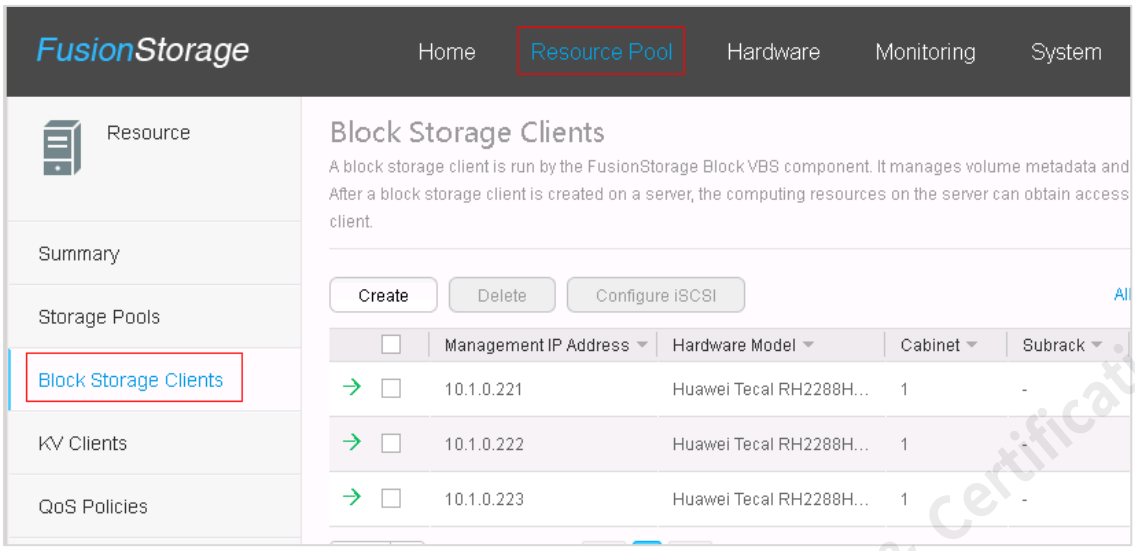

Create block storage clients.

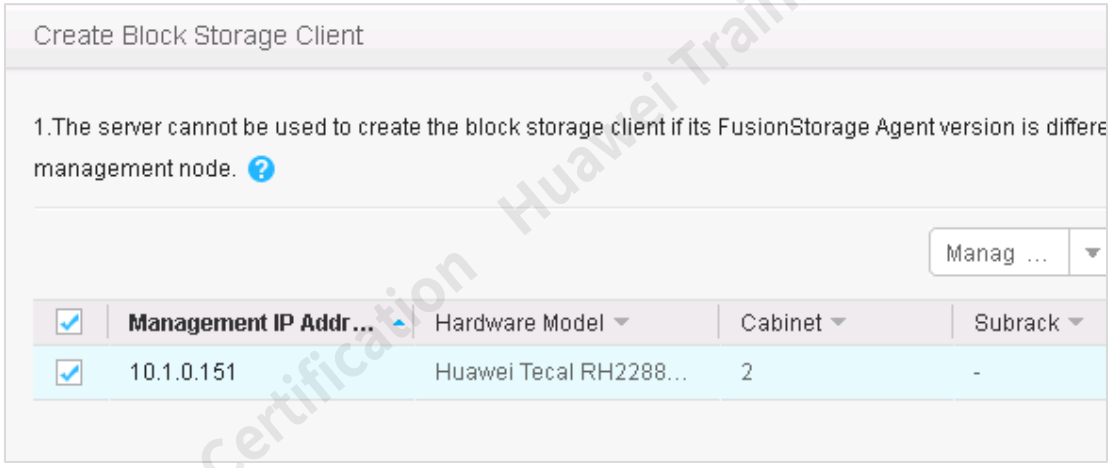

The installation is complete.

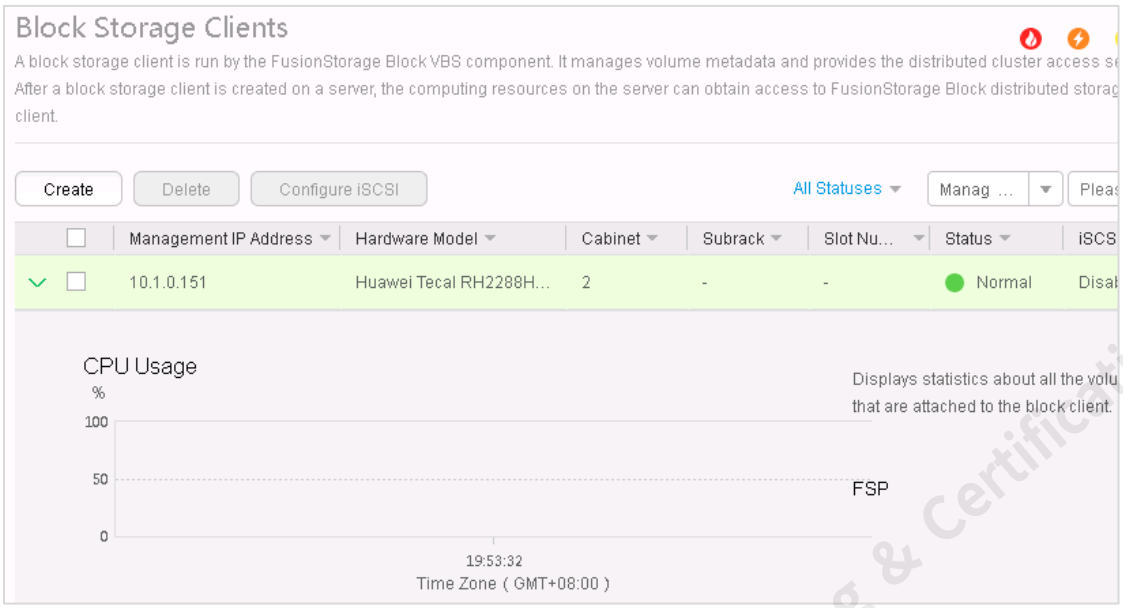

----End

8.6 Configuring Backup Storage

eBackup allows you to use NAS and S3 storage as backup storage in FusionCloud backup scenarios. You need to create storage units for level-1 backup storage (NAS) and level-2 backup storage (NAS and S3). In this experiment, NAS of OceanStor V3 is used as the backup storage.

8.6.1 Configuring OceanStor V3 NAS

A disk domain and a storage pool have been created.

Step 1 **Create a file system.**

Choose Provisioning > File System.

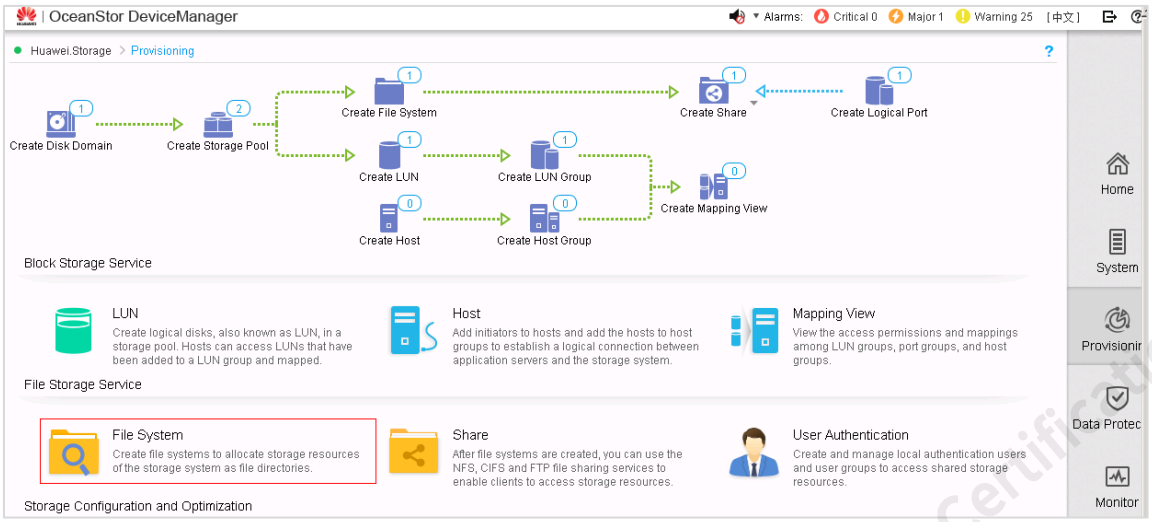

In the **Create File System** dialog box, select a proper application scenario based on the storage type.

Step 2 **Configure the NFS logical port.**

Logical ports are created based on Ethernet ports, bond ports, or VLANs. Connect physical cables and then create logical ports.

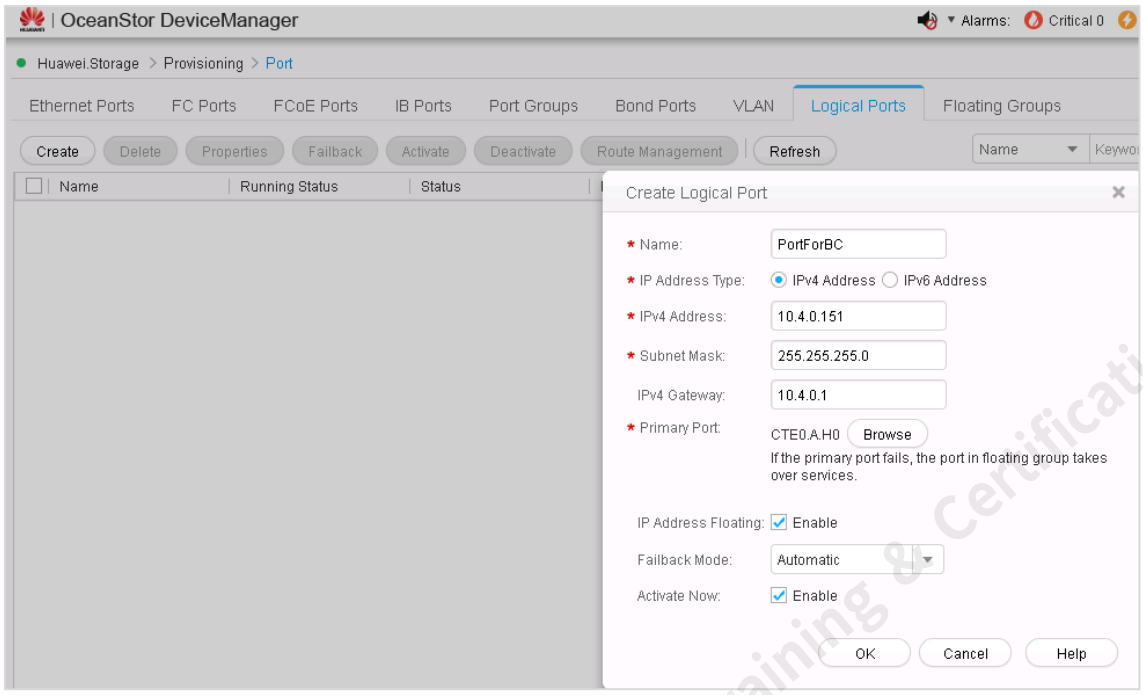

Step 3 **Create an NFS.**

Choose **Provisioning** > **Share** > **NFS (Linux/UNIX/MAC)**. Click **Create** and then select a file system in the displayed dialog box.

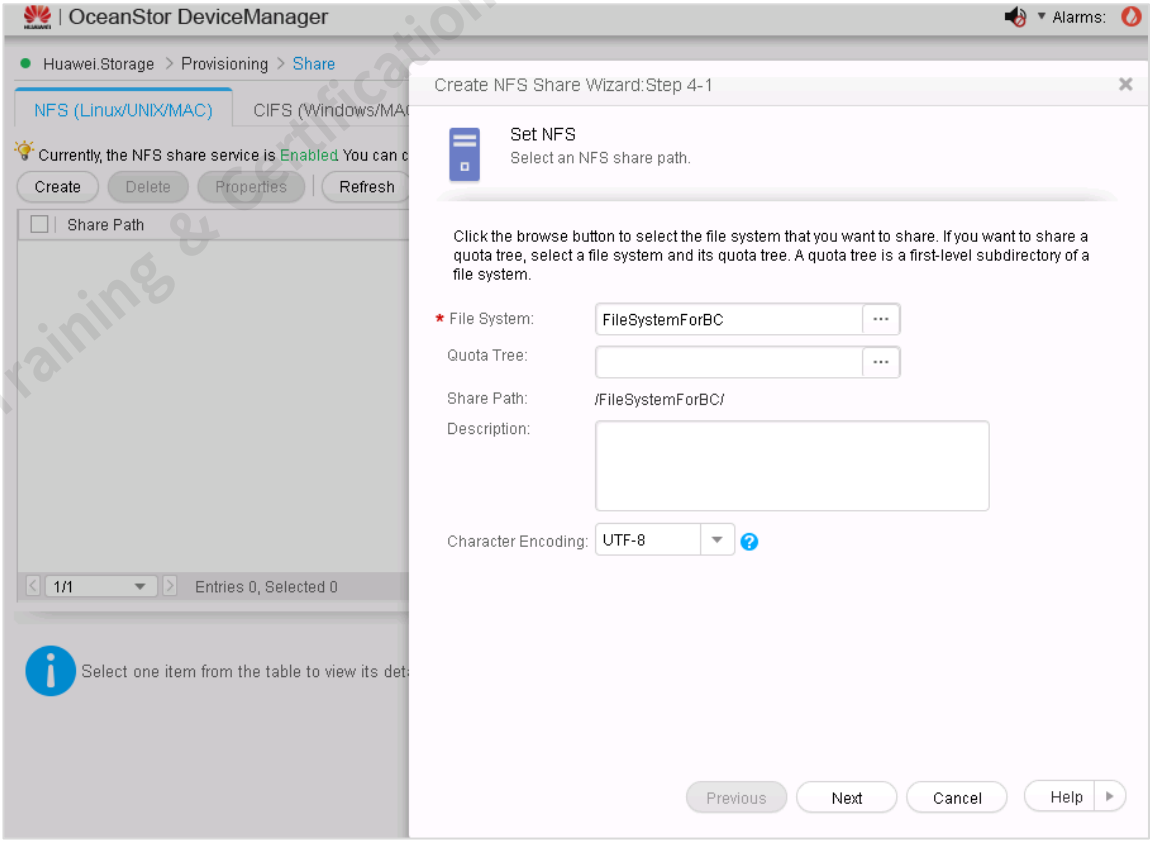

Set the access permission.

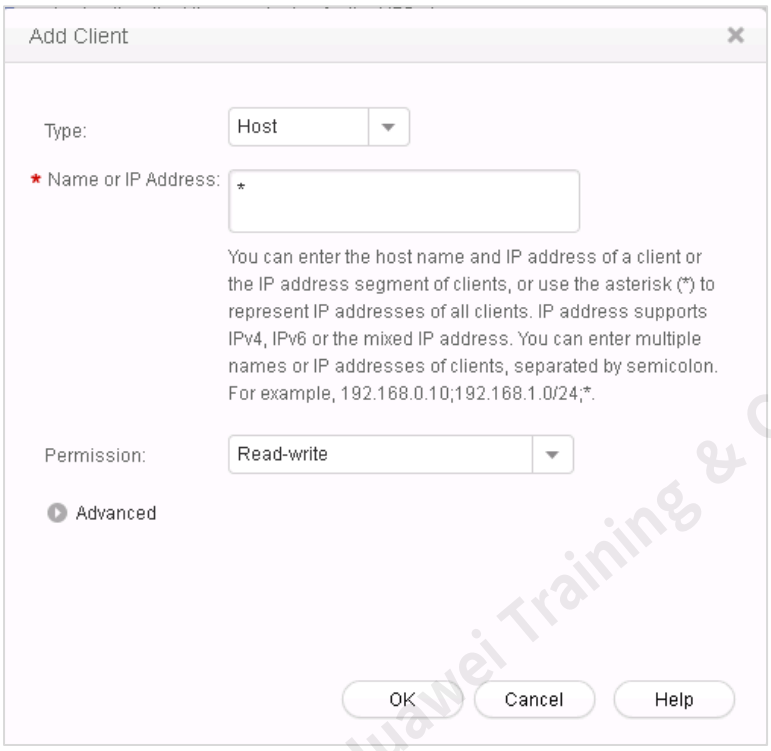

----End

8.6.2 Configuring Backup Storage

In FusionCloud backup scenarios, you only need to create storage units. The corresponding storage pools and repositories are automatically created by the cloud platform based on the backup task.

Log in to eBackup Server. Choose **Backup Storage** > **Storage Unit**.

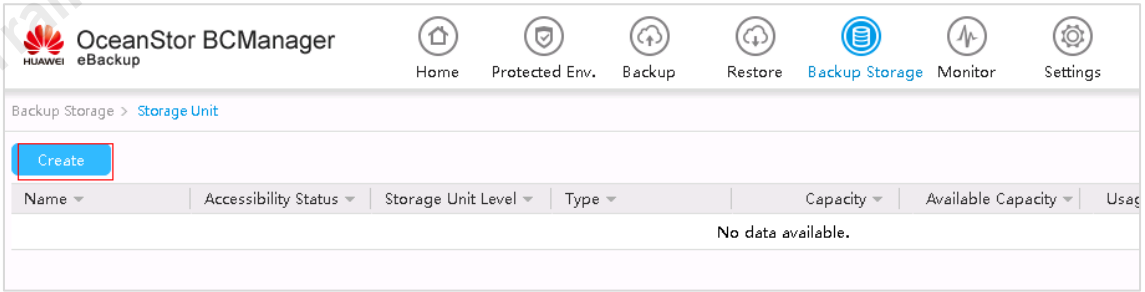

Enter the NFS path.

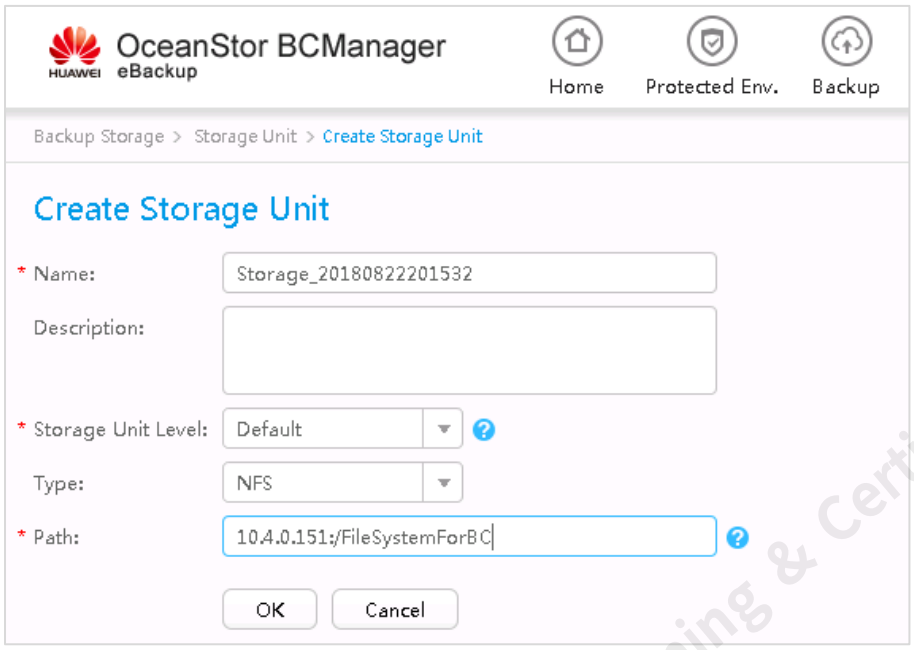

Ensure that the storage unit is accessible.

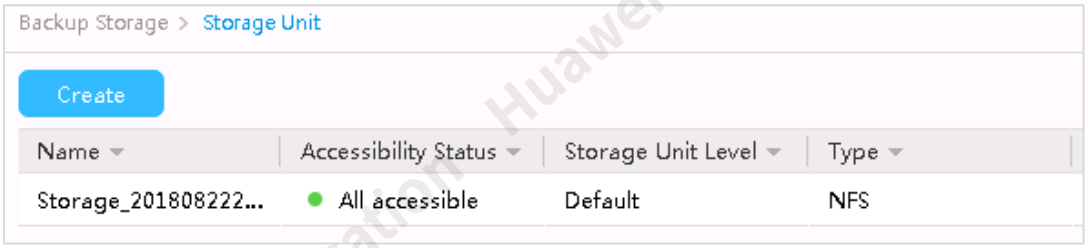

8.7 Configuring CSBS Karbor

Karbor has been connected with OpenStack and the certificate has been imported.

Step 1 **Log in to the Karbor VM.**

The initial password for user **root** is **CloudService@123!**.

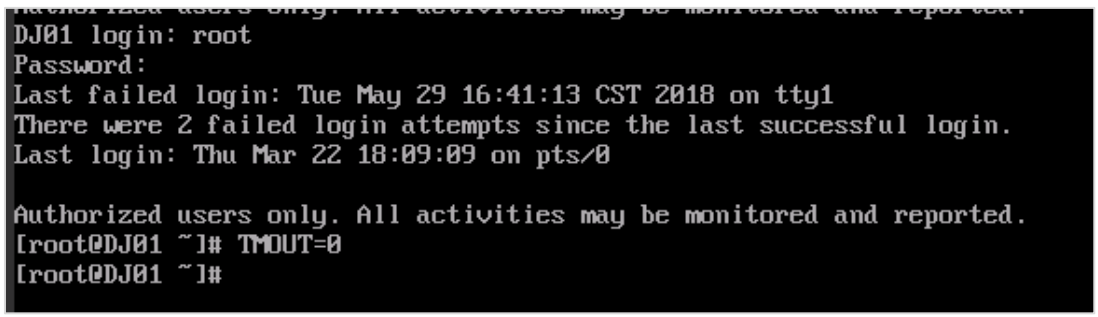

Step 2 **Configure AZs.**

After you have successfully configured an AZ on any CSBS Karbor node, the system will automatically synchronize the configuration of the rest nodes.

set_az_backup_cp --op operation --name az_name --local_backup backup --remote_copy copy --remote_copy_targets re_cptargets --local_restore_targets lo_retargets --remote_restore_targets re_retargets

```
[root@DJ01 ~]# set_az_backup_cp --op add --name AZ01 --local_backup supported --remote_copy supported --remote_copy_targets "['AZ01']" --local_restore_targets "['AZ01']" --remote_restore_targets "['AZ01']"
Operation succeeded.
```

Step 3 **Query the current AZs.**

```
[root@DJ01 ~]# set_az_backup_cp --op list
+-----------+--------------+-------------+--------------------+---------------------+----------------------+
|    name   | local_backup | remote_copy | remote_copy_targets | local_restore_targets | remote_restore_targets |
+-----------+--------------+-------------+--------------------+---------------------+----------------------+
|    fc01   |  supported   |  supported  |       ['fc01']      |        ['fc01']       |        ['fc01']       |
| manage-az |  supported   |  supported  |    ['manage-az']    |     ['manage-az']     |     ['manage-az']     |
|    AZ01   |  supported   |  supported  |       ['AZ01']      |        ['AZ01']       |        ['AZ01']       |
+-----------+--------------+-------------+--------------------+---------------------+----------------------+
Operation succeeded.
```

----End

8.8 Applying for Cloud Host Backup

8.8.1 Applying for Cloud Host Backup

Step 1 **Log in to ServiceCenter.**

Step 2 **Apply for cloud host backup.**

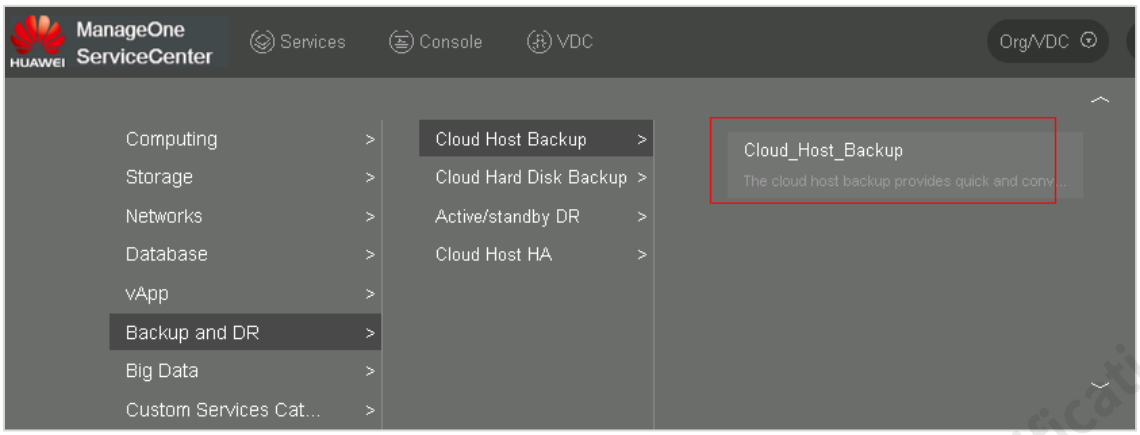

Select an AZ and a cloud host to be backed up.

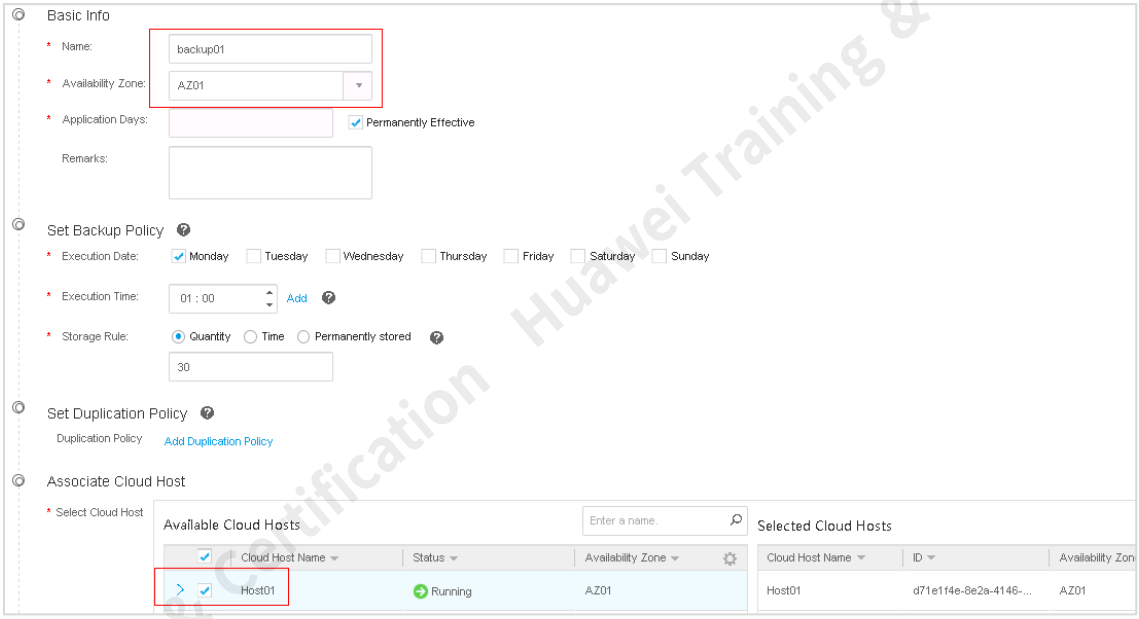

Step 3 **Check the status of applying for backup.**

In **User Center**, the backup is under implementation.

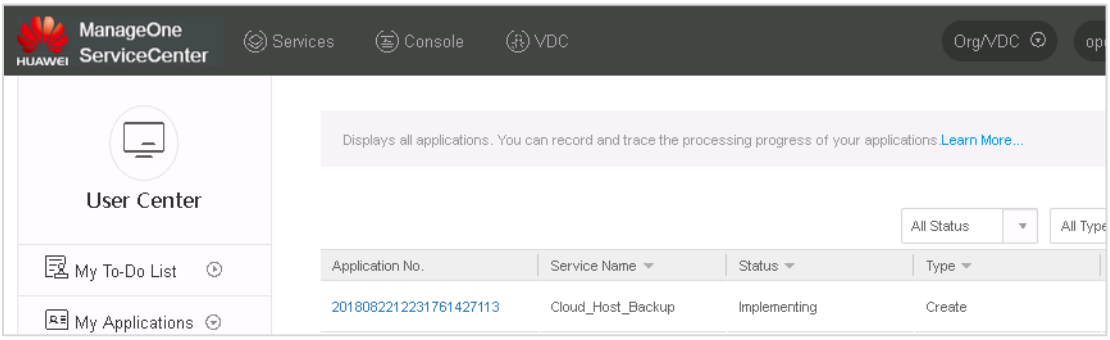

----End

8.8.2 Using Cloud Host Backup

Step 1 **Check the cloud host backup service.**

Choose **Console** > **Backup and DR** > **Cloud Host Backup**. Click **Manage Services**. No backup copies are available at present.

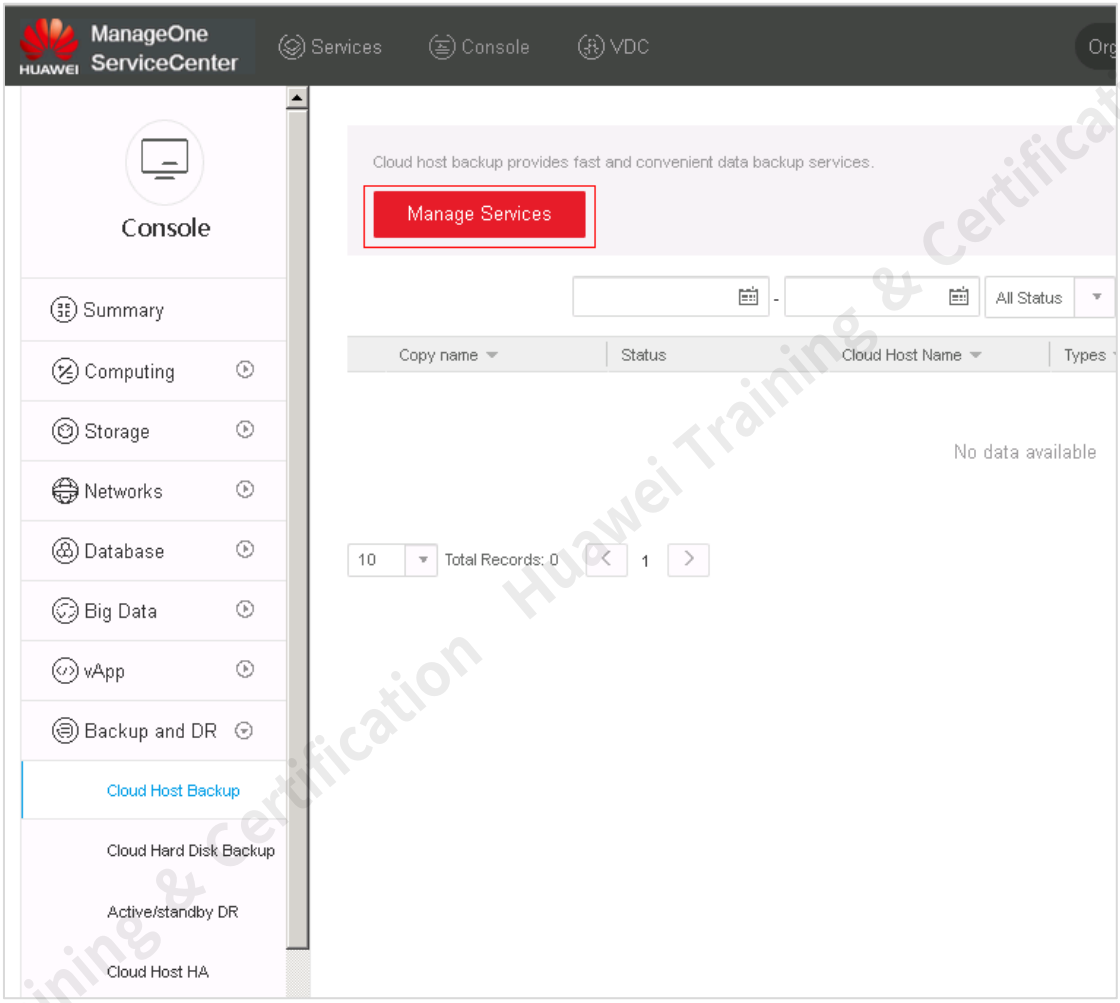

Step 2 **Manually back up the cloud host.**

Check the current backup policy. Click **More**, select **Manually Backup**.

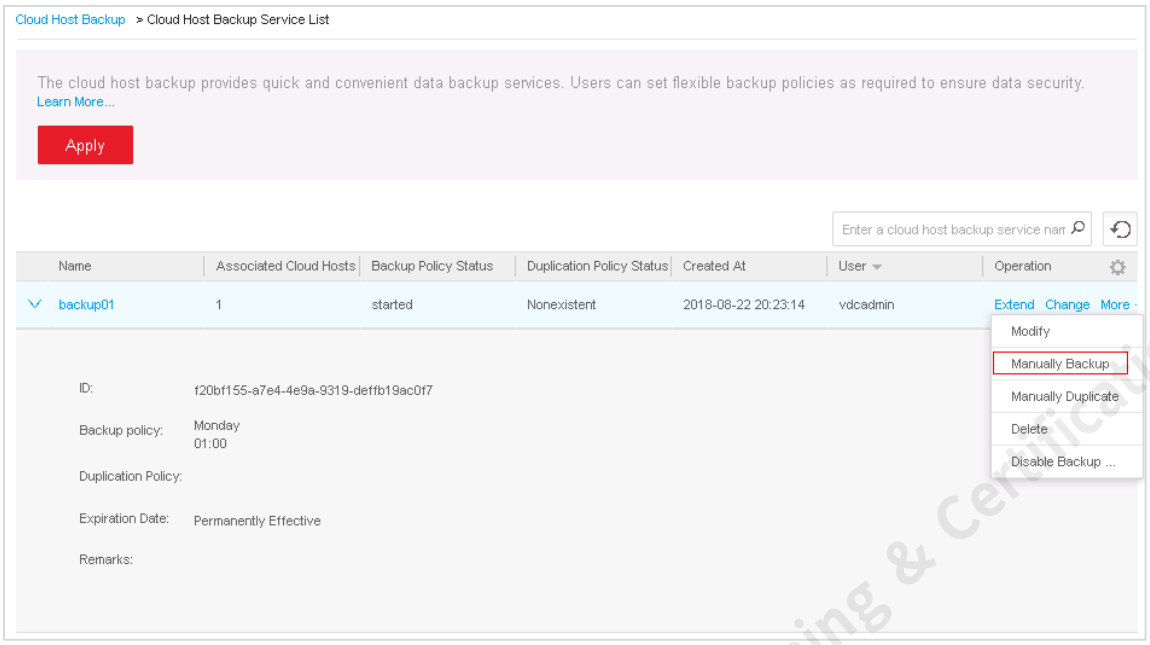

Confirm the execution of manual backup.

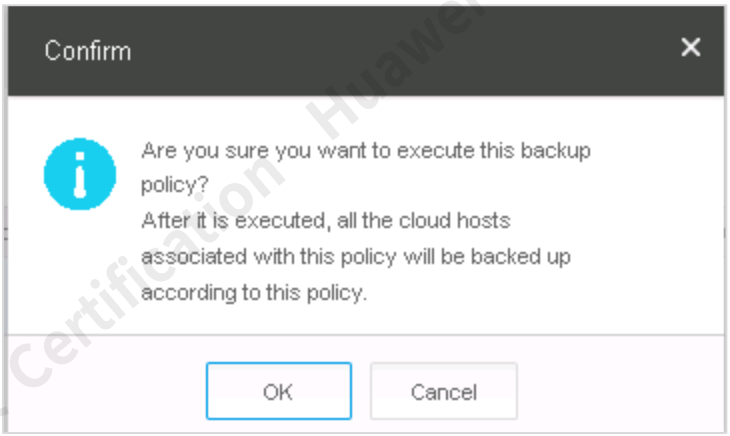

Click **Copy**. The status is **Executing**.

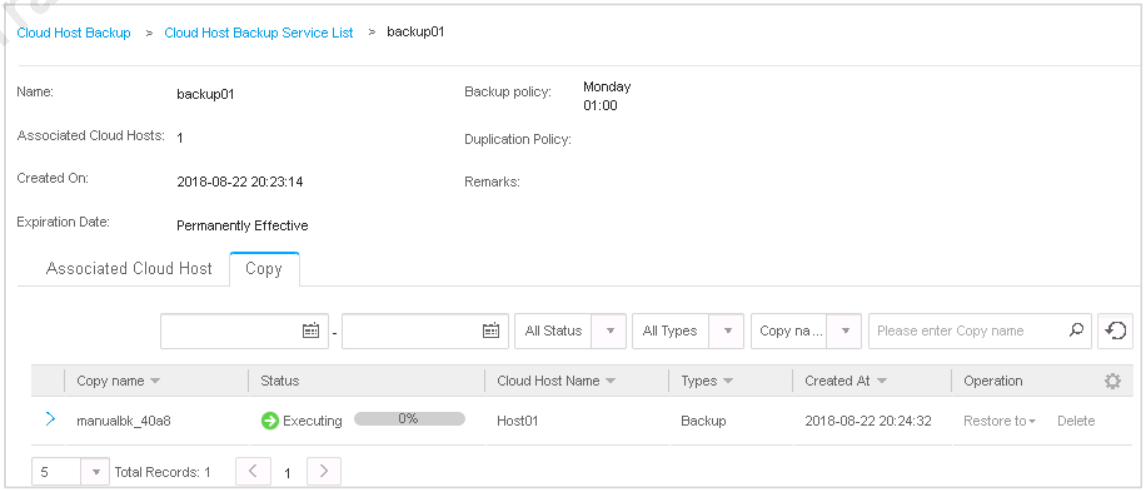

Step 3 **Restore the backup data.**

On the cloud host backup page, select a cloud host copy. Select **Original cloud host** from the **Restore to** drop-down list. The backup server automatically starts backup.

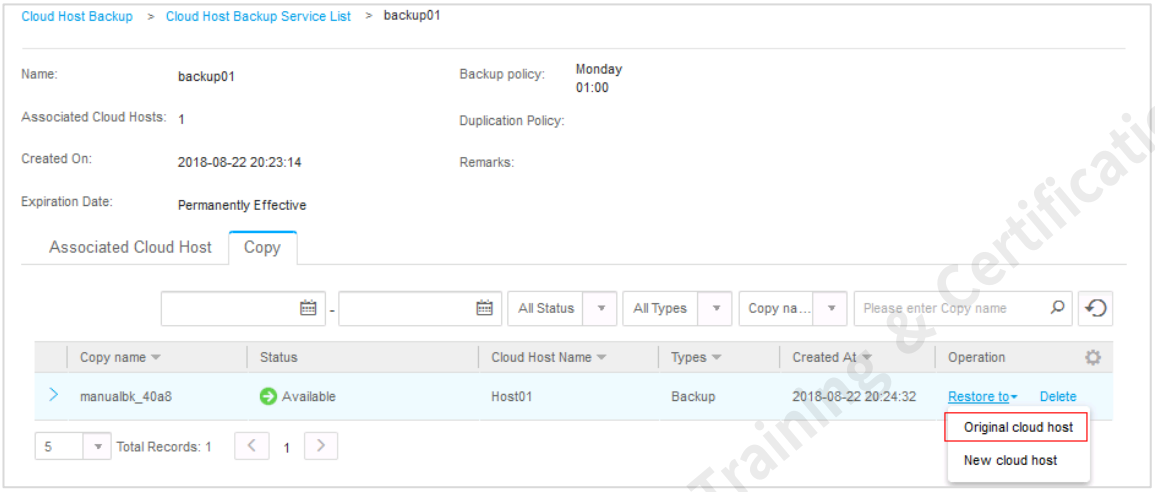

----End

9 O&M Experiment in Cloud Data Centers

9.1 Performing eSight O&M

9.1.1 About This Experiment

Adding devices to eSight is the prerequisite for network monitoring. In terms of the vendor and device type, devices managed by eSight include Huawei devices are non-Huawei devices (such as H3C and Cisco devices). During deployment, capacity expansion, or the first use of eSight, you can add devices to eSight in batches or by using automatic device discovery.

SNMP-based device discovery helps discover devices in batches in a timely manner or periodically. It is recommended for further device management, including querying device details, monitoring device faults, and monitoring device performance. This section describes how to automatically discover devices using SNMPv2c.

9.1.2 Objectives

- To have a good command of SNMP configurations of switches and routers
- To understand how to manually create an NE
- To understand how to use SNMP to automatically discover devices

9.1.3 Configuring Experiment Tasks

9.1.3.1 Configuration Roadmap

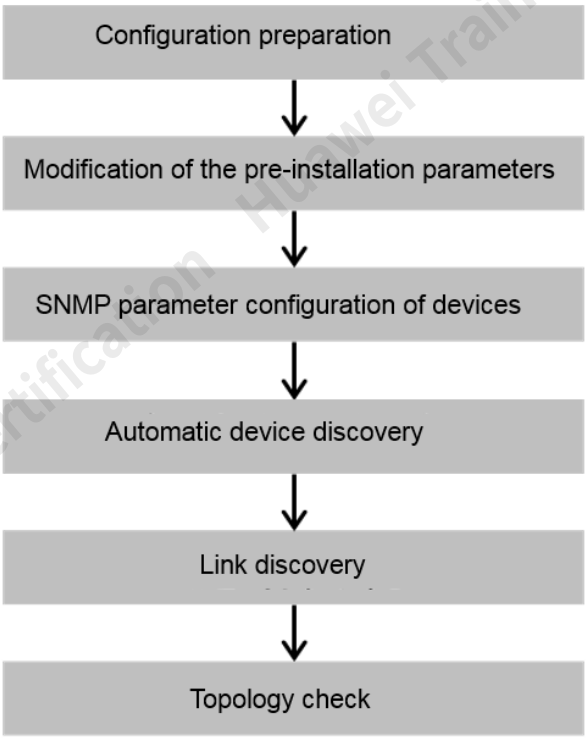

9.1.3.2 Configuration Procedure

Step 1 **Complete device pre-configuration.**

For details about device pre-configuration, see the first experiment. Prepare the lab environment.

Step 2 **Configure SNMP parameters for the devices.**

Configurations on FW1 and FW2 are as follows:

```
<FW1>system-view
[FW1]interface GigabitEthernet  1/0/1
[FW1-GigabitEthernet0/0/0]ip address 172.21.22.2 16
[FW1-GigabitEthernet0/0/0]service-manage snmp permit
[FW1]quit
[FW1]snmp-agent
[FW1]snmp-agent community read cipher Huawei@123
[FW1]snmp-agent community write cipher Admin@123
[FW1]snmp-agent sys-info version v2c
[FW1]snmp-agent target-host trap address udp-domain 172.21.22.100 params
securityname cipher Huawei@123
[FW1]snmp-agent trap source GigabitEthernet1/0/1
[FW1]snmp-agent trap enable
```

Configurations on Leaf-1, Leaf-2, Leaf-3, Spine-1, and Spine-2 are as follows:

```
<Leaf-1>system-view immediately
[Leaf-1]interface  MEth  0/0/0
[Leaf1-MEth0/0/0]ip address 172.21.22.6 16
[Leaf1-MEth0/0/0]quit
[Leaf-1]snmp-agent
[Leaf-1]snmp-agent sys-info version v2c
[Leaf-1]snmp-agent community read Huawei@123
[Leaf-1]snmp-agent community write Admin@123
[Leaf-1]snmp-agent target-host trap address udp-domain 172.21.22.100 params
securityname cipher Huawei@123
[Leaf-1]snmp-agent trap source Meth 0/0/0
[Leaf-1]snmp-agent trap enable
```

Server configurations are as follows:

Log in to the iBMC WebUI, choose **Configure** > **System Settings**. Select **v2C**. Enter **Huawei@123** into the box on the right of **Read-Only Community** and **Admin@123** into the box on the right of **Read/Write Community**.

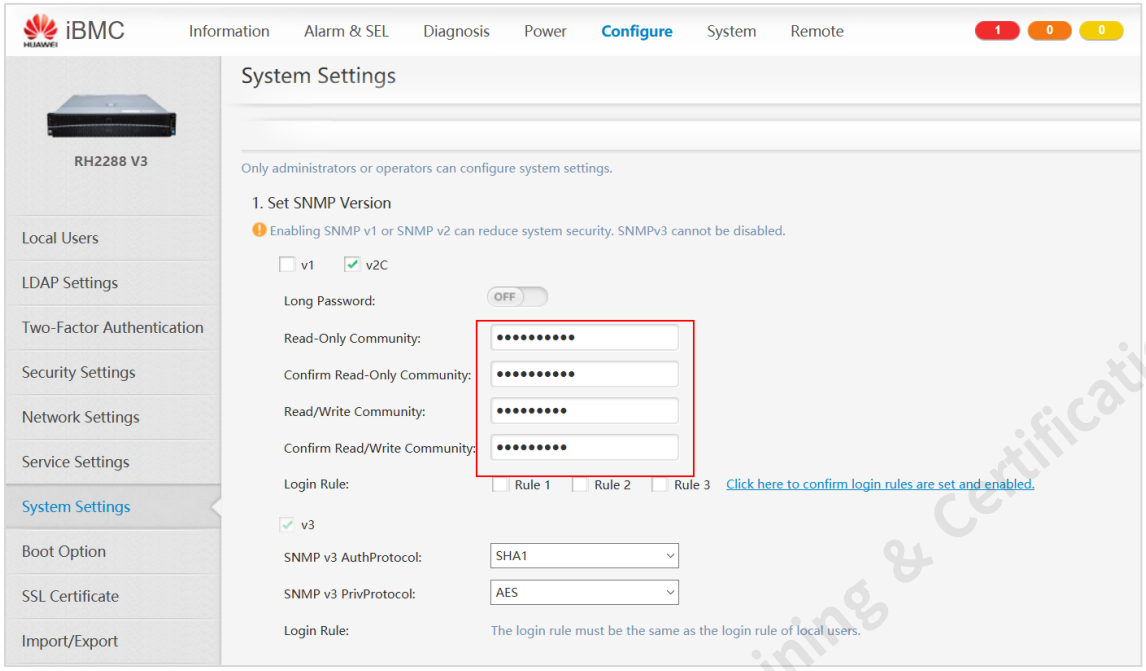

Choose **Settings** > **Alarm Settings** > **Trap IP Address Management** > **Add**. In the displayed **Modify Server IP Address** dialog box, set parameters and then click **OK**.

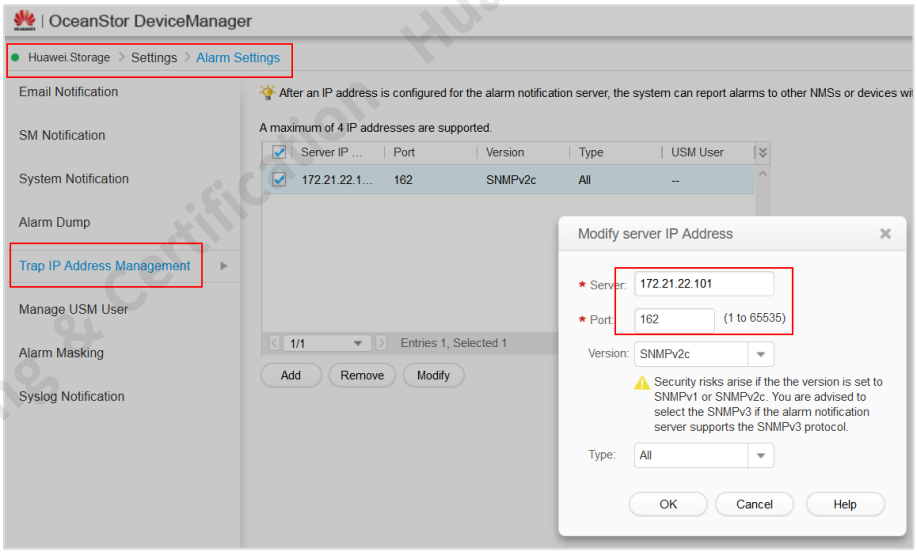

Step 3 **Configure Telnet parameters for the devices.**

Configurations on FW1 and FW2 are as follows:

```
<FW1>system-view
[FW1]telnet server enable
[FW1]user-interface vty 0 4
[FW1-ui-vty0-4]authentication-mode password
[FW1-ui-vty0-4]set authentication password cipher Huawei@123
[FW1-ui-vty0-4]user privilege level 15
```

```
[FW1-ui-vty0-4]protocol inbound telnet
```

Configurations on Leaf-1, Leaf-2, Leaf-3, Spine-1, and Spine-2 are as follows:

```
<Leaf-1>system-view immediately
[Leaf-1]undo telnet server disable
[Leaf-1]user-interface vty 0 4
[Leaf-1-ui-vty0-4]authentication-mode password
[Leaf-1-ui-vty0-4]set authentication password cipher Huawei@123
[Leaf-1-ui-vty0-4]user privilege level 3
[Leaf-1-ui-vty0-4]protocol inbound telnet
```

Step 4 **Manually add an NE.**

Create an SNMP template on the eSight by performing the following operations: On the main menu, choose **Resource** > **Common**.

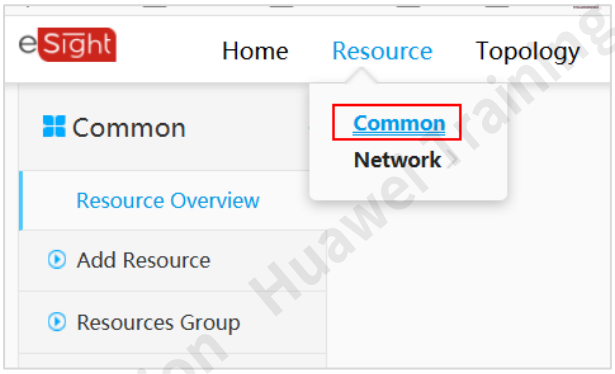

In the navigation tree on the left, choose **Add Resource** > **Protocol Template** > **SNMP Template**, and click **Create**.

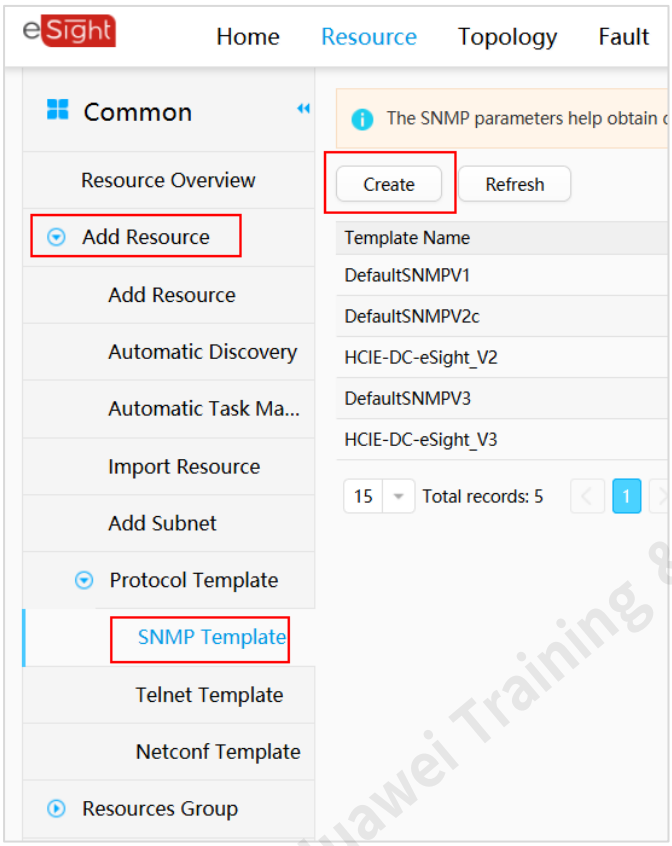

Create an SNMP parameter template by referring to the protocol parameter planning in the data preparation, set **Read community** and **Write community** to **Huawei@123** and **Admin@123**, respectively, and click OK. The template is created.

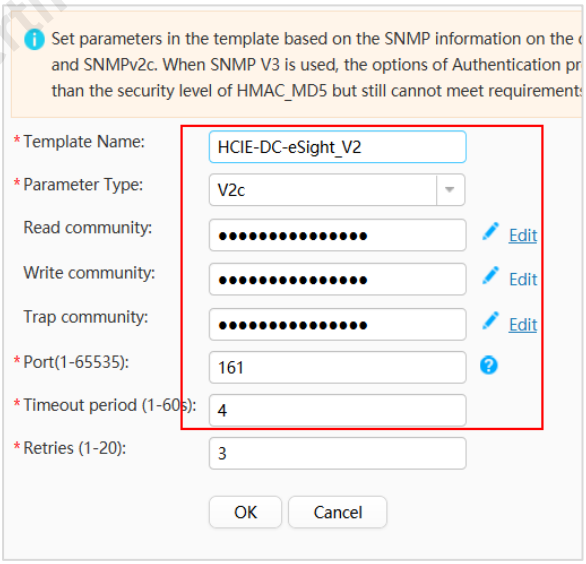

Create a Telnet template on the eSight by performing the following operations: On the main menu, choose **Resource** > **Common**.

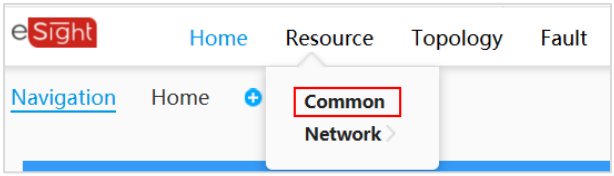

In the navigation tree on the left, choose **Add Resource** > **Protocol Template** > **Telnet Template**, and click **Create**.

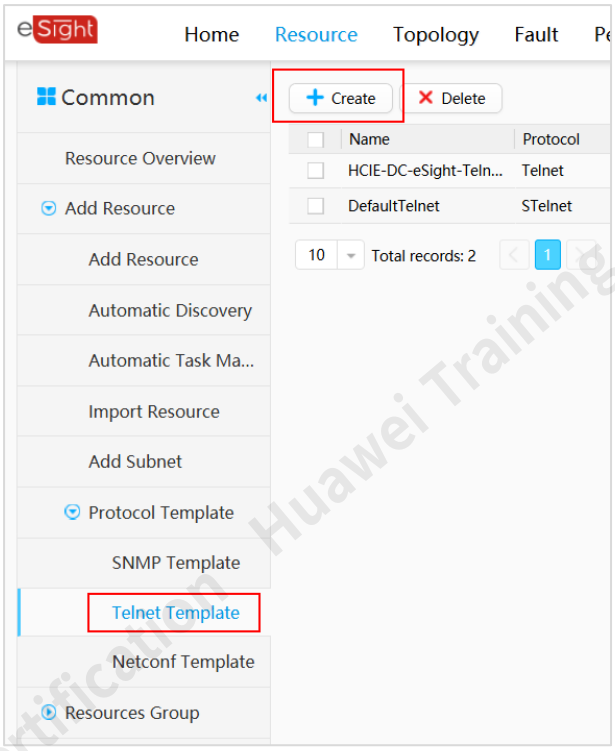

Create a Telnet parameter template. Set **Template name** to **HCIP-DC-eSight-Telnet**, **Protocol** to **Telnet**. Set the authentication mode to password authentication, and enter the password **Huawei@123**. Click **OK**.

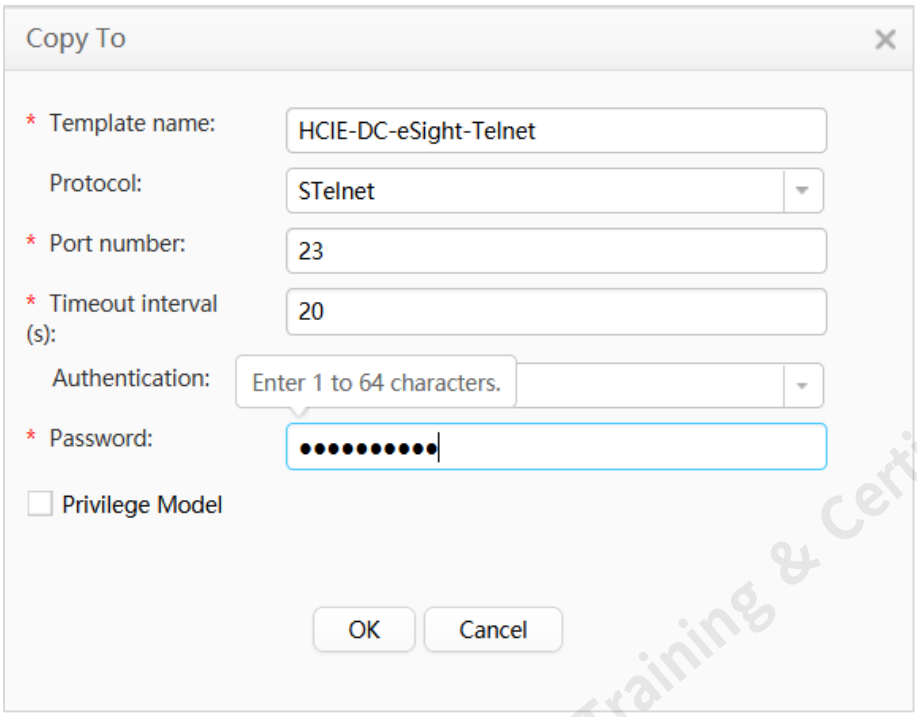

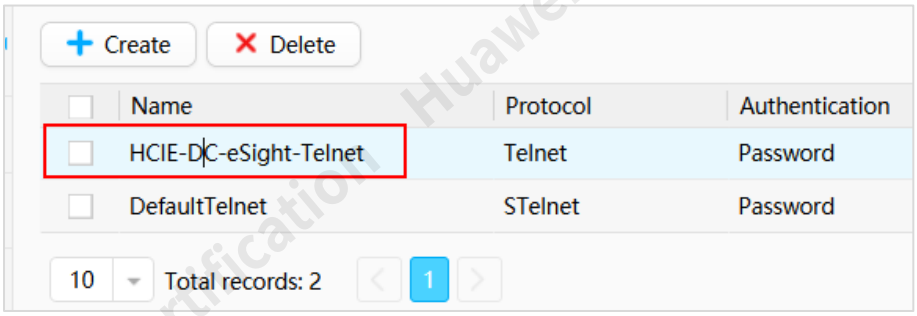

Step 5 **Use the automatic discovery function to add NEs.**

On the menu bar, choose Resource > Add Resource > Automatic Discovery.

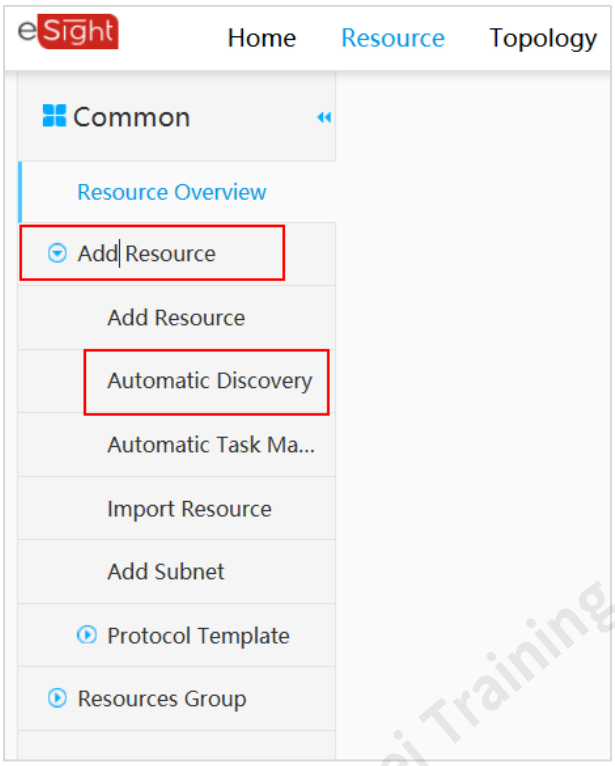

In Basic Settings, set Start IP address to 172.21.22.1, End IP address to 172.21.22.27, and Add to subnet to Root.

Set the task execution frequency and select the **Automatically match template** check box. Click **OK** and then **Discover Devices**.

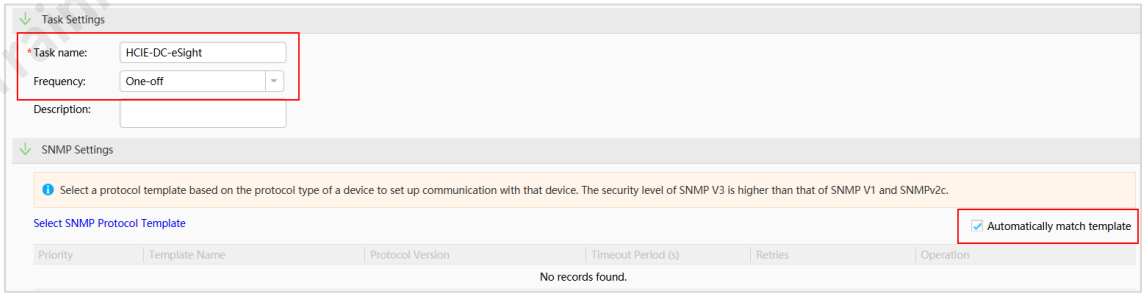

According to the result, eight devices are discovered successfully, and no devices fail to be discovered. Device addition fails if the device has been added.

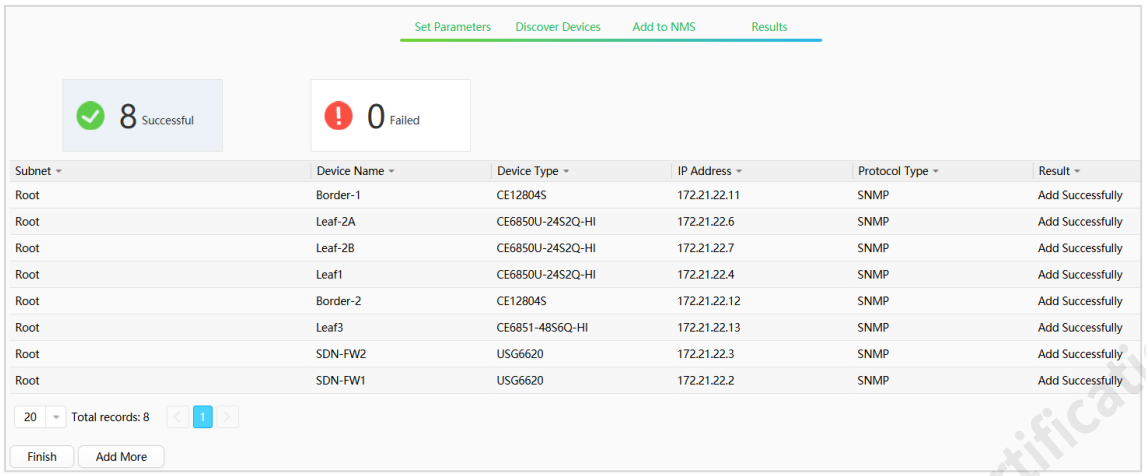

Set Telnet parameters for discovered devices. On the main menu, choose **Resource** > **Common** > **Network Device**. Select **Set Telnet Parameters** from the **Set Protocol** drop-down list. Select a parameter template and click **OK**.

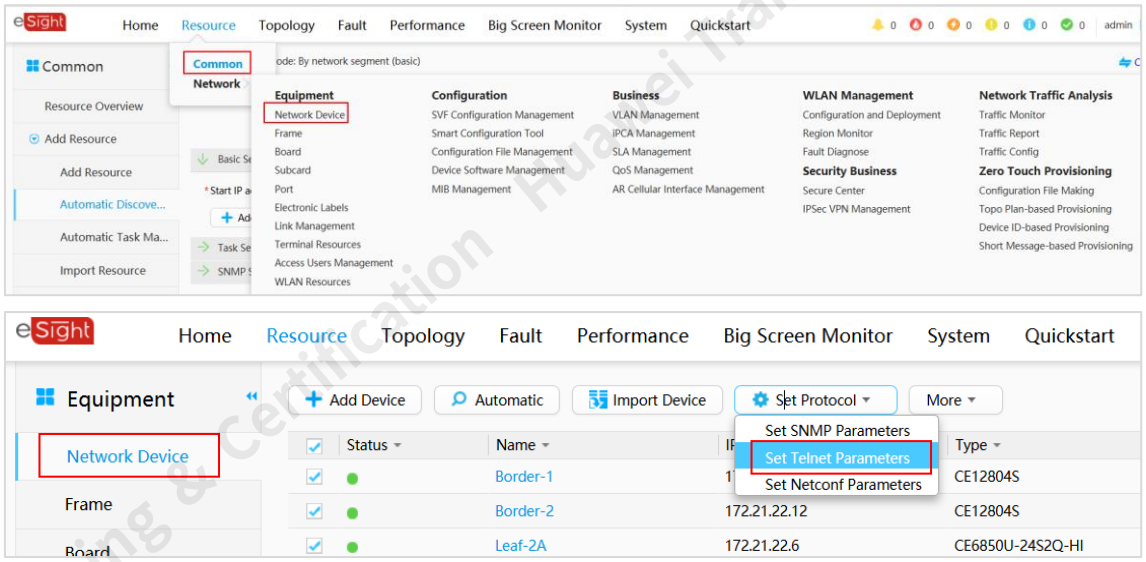

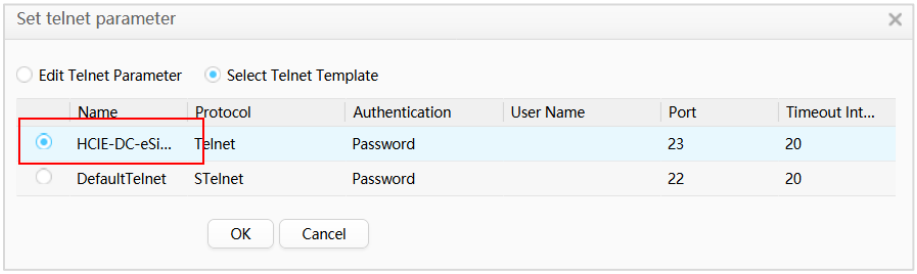

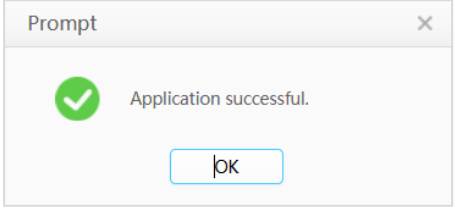

----End

9.1.4 Verifying the Result

9.1.4.1 Checking the Topology

Step 1 **Discover links.**

Choose **Resource** > **Network** > **Link Management** from the main menu.

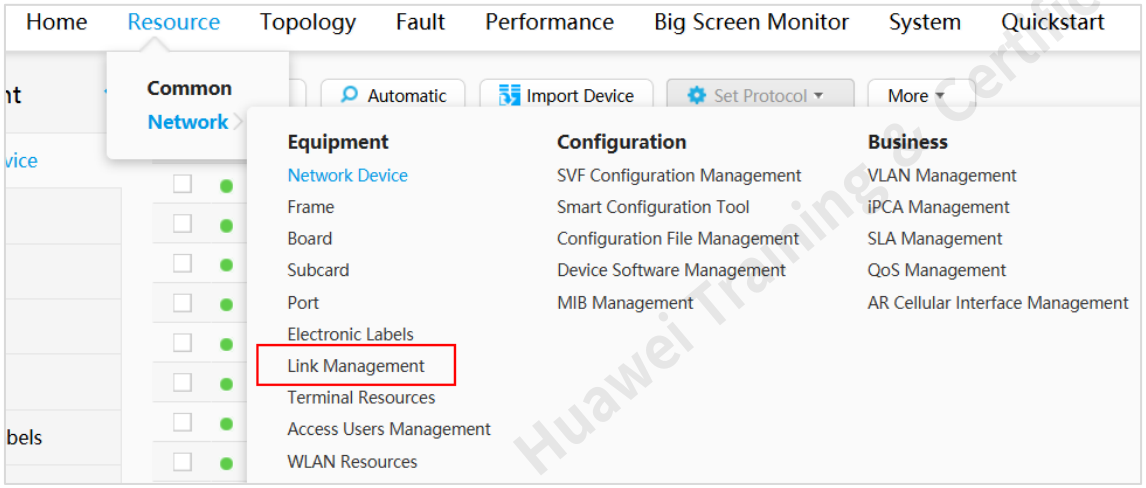

Click **Discover Link**, select the subnet where the device resides, and start to discover links.

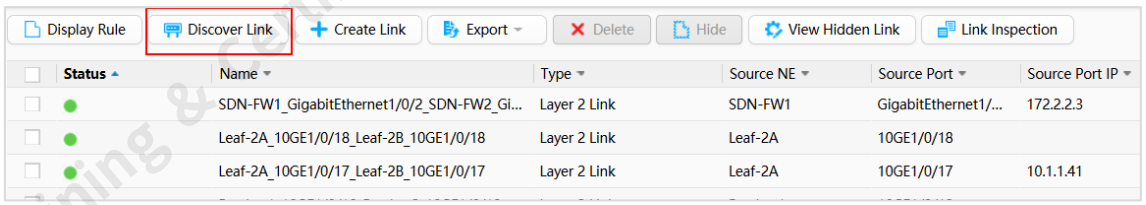

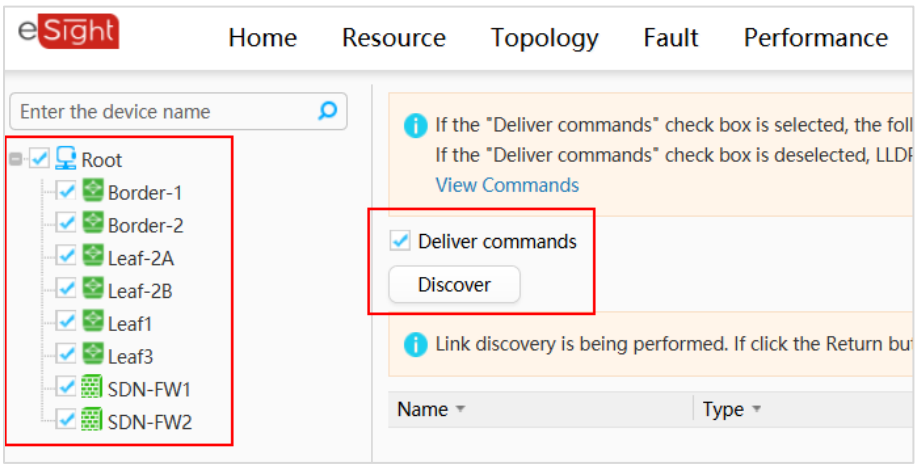

The link discovery is complete.

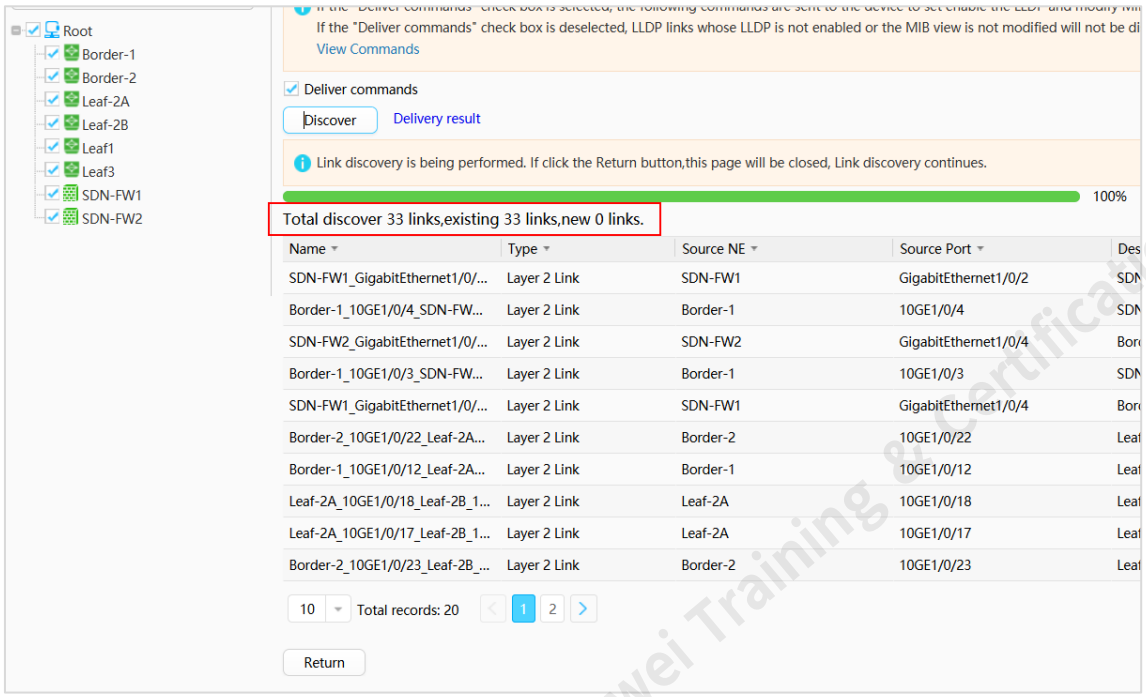

Step 2 **Adjust the network topology.**

Choose **Topology** > **Topology Management** from the main menu. Open the subnet topology named eSight-lab.

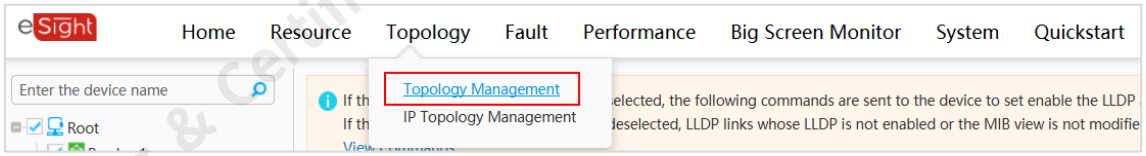

Step 3 **Adjust the topology structure. Click to save the topology.**

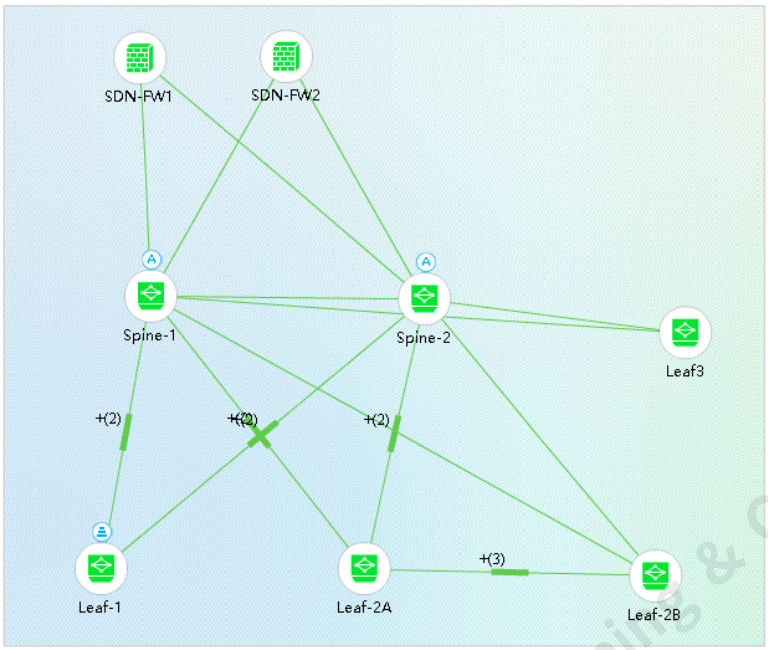

----End

9.2 Performing an O&M Experiment on the ManageOne OperationCenter System

9.2.1 Objectives

- To understand how to install and deploy ManageOne OperationCenter
- To understand how to configure interconnections between ManageOne OperationCenter and other components
- To have a good command of ManageOne OperationCenter service commissioning

9.2.2 Planning Networking

- ManageOne OperationCenter is deployed on a VM.
- The OpenStack External_API network segment is used and the IP address 192.168.0.180 is used.
- The OpenStack NTP server with the IP address of 10.1.0.9 is used as the NTP server.

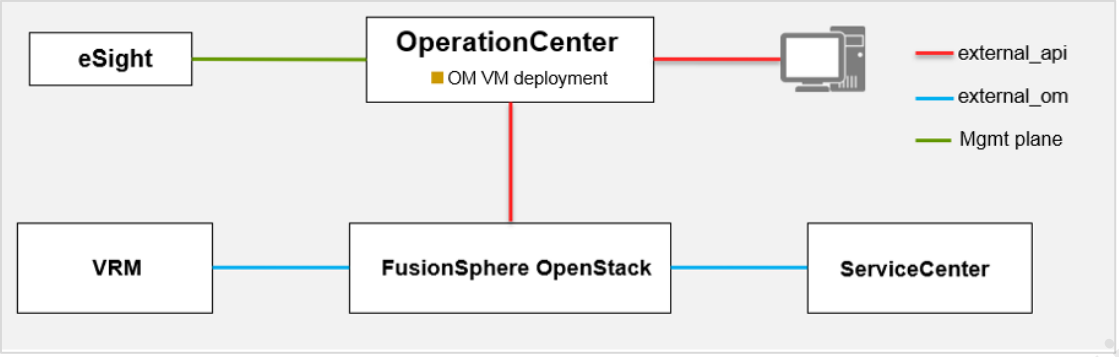

ManageOne OperationCenter VMs are deployed on FusionSphere OpenStack, interconnect with FusionSphere OpenStack and ManageOne ServiceCenter through external_api, and interconnect with the eSight through the management plane.

9.2.3 Preparing for the Installation

- Complete ManageOne OperationCenter network planning.
- Obtain the centralized installation package of the OS and ManageOne OperationCenter.
- Complete FusionSphere OpenStack installation and deployment.
- Confirm that VM watchdog has been disabled. To check the VM watchdog setting, log in to the FusionSphere OpenStack web client.

Scenario Scale	Small-Scale Deployment	Standard Deployment	Large-Scale Deployment
Configuration requirements	CPU: > 2 GHz CPU core: > 8 cores Memory: > 16 GB Disk capacity: 500 GB	CPU: > 2 GHz CPU core: > 16 cores Memory: > 32 GB Disk capacity: 1500 GB	CPU: > 2 GHz CPU core: > 16 cores Memory: > 64 GB Disk capacity: 1500 GB
Management scale	0 to 3000 VMs and 0 to 500 physical devices	3000 to 10,000 VMs and 0 to 3000 physical devices	10,000 to 50,000 VMs and 0 to 12,000 physical devices

This document describes ManageOne OperationCenter installation and deployment in small-scale scenarios. If the number of data centers exceeds 10, large-scale deployment is recommended.

9.2.4 Configuring the Interconnection

Step 1 **Configure interconnection of OperationCenter with other components to collect alarm and monitoring information.**

In this experiment, ManageOne OperationCenter needs to interconnect with FusionSphere OpenStack, FusionCompute, ServiceCenter, eSight, and BCManager.

The following is the interconnection data list.

Management System to Connect	Interconnection Purpose	Interconnection User Description
Local SNMP	The system can receive the alarm data reported by third-party systems over SNMP.	SNMP security users
FusionSphere OpenStack	Collect and manage information about devices, alarms, and performance from the FusionSphere OpenStack virtualization platform to comprehensively analyze the running status of physical devices in a DC.	-
ServiceCenter	Obtain VDC information to comprehensively analyze the running status of a DC.	The interconnection user must be **OCRest** and the default password is **Huawei@CLOUD8!!**.
eSight	Collect and manage information about devices, alarms, and performance from eSight to comprehensively analyze the running status of physical devices in a DC.	The interconnection user must be an administrator user but cannot be user **admin**, and the password of the administrator user cannot be an initial password of the system.

Configure the local SNMP.

OperationCenter uses SNMP to receive alarms reported by third-party management systems in SNMP trap mode.

The following table lists SNMP data.

Parameter	Description	Reference Value
System name	Local SNMP name	SNMP01
Port number	Default value: **10163**	10163

Parameter	Description	Reference Value
Version	SNMP version	V3
Security user name	V3 USM security username used when OperationCenter connects as the SNMP server to a third-party system	ocuser
Authentication protocol	Authentication protocol used when the SNMP server connects to a third-party system	SHA
Authentication password	-	Huawe@123
Encryption protocol	Information encryption protocol used when the SNMP server connects to a third-party system	AES-128
Encryption password	Encryption password for the V3 USM security username	Huawei@1234
Check read community name		Yes
Read community name	Read community reported from the local system to the peer system	Huawei@123

On the OperationCenter page, choose System > Connection Management > Data Report Configuration.

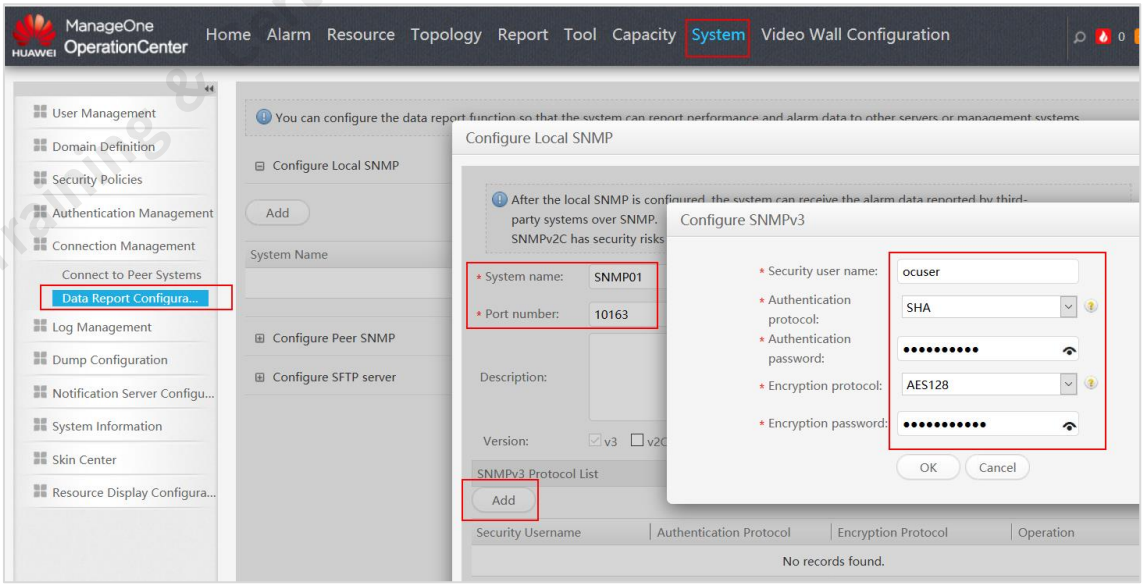

Configure the local SNMP according to the data plan.

Before using OperationCenter, you need to define the location information to identify the physical area to which the monitored object belongs.

On the OperationCenter management portal, choose **Domain Definition** > **Physical Location Management**.

Click **Basic Physical Location** and then **Create**. In the **Create** dialog box, add **Zhejiang** as the level-1 location information.

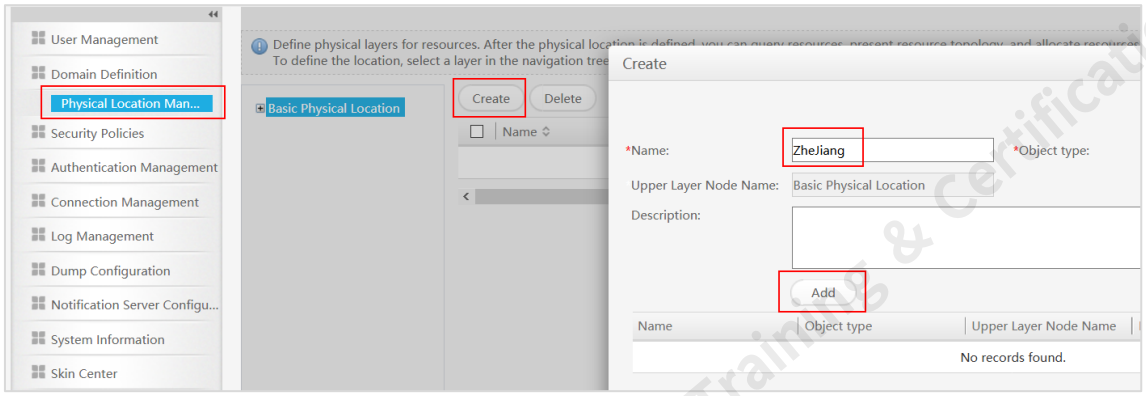

Click **Zhejiang** and in the **Create** dialog box, add **HangZhou** as the level-2 location information.

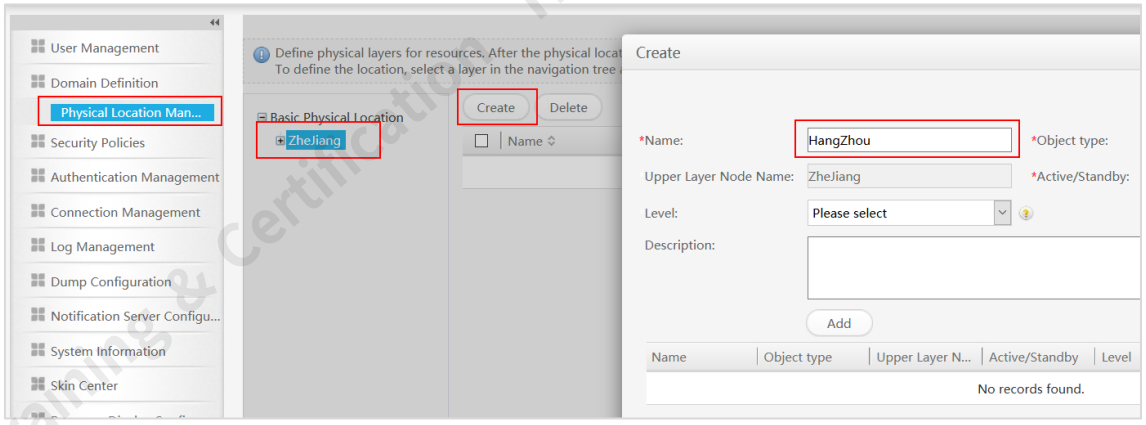

Click **OK**.

Step 2 **Connect OperationCenter with eSight.**

By connecting with eSight, OperationCenter can collect and manage information about devices, alarms, and performance form eSight to comprehensively analyze the running status of physical devices in a DC.

The following table lists interconnection parameters.

Parameter	Description	Reference Value
System name		eSight
IP address and port ID		172.21.22.100 / 32102
Username and password	Username used for interconnection between eSight and OperationCenter. When OperationCenter interconnects with eSight V300R007C00, you need to create an **Open API user** on eSight and ensure that the user belongs to **Open API user group** and **Administrators**. If the belonging roles of the user do not include **Administrators**, the interconnection is successful. However, the local user created on OperationCenter is not automatically synchronized to eSight.	esightuser/Huawei@123
Version		V300R007
IP address or port for collecting performance data in SFTP mode	IP address of the eSight server	172.21.22.100 / 32067
Username and password for collecting performance data in SFTP mode		sftpuser / Huawei@123
Path for saving performance data collected in SFTP mode		/PerformanceFiles

Configurations on eSight:

Create an interconnection account.

Enter https://172.21.22.100:31943/ in to the browser. Log in to eSight. Choose **System** > **User Management** > **User**. Click **Create** to create an interconnection user.

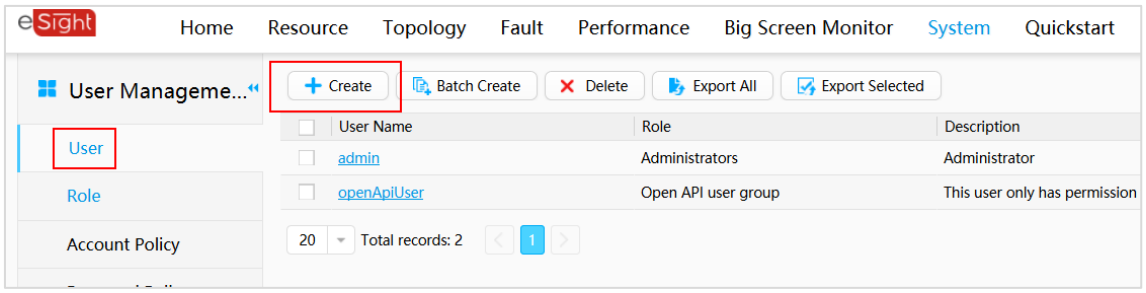

Enter the username and password, select **OpenAPI user**, and click **Next**.

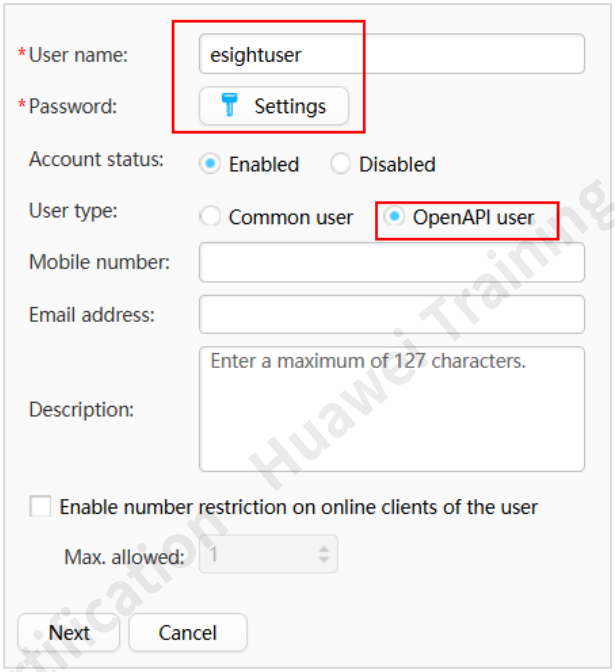

Select the Open API user group and Administrators check boxes.

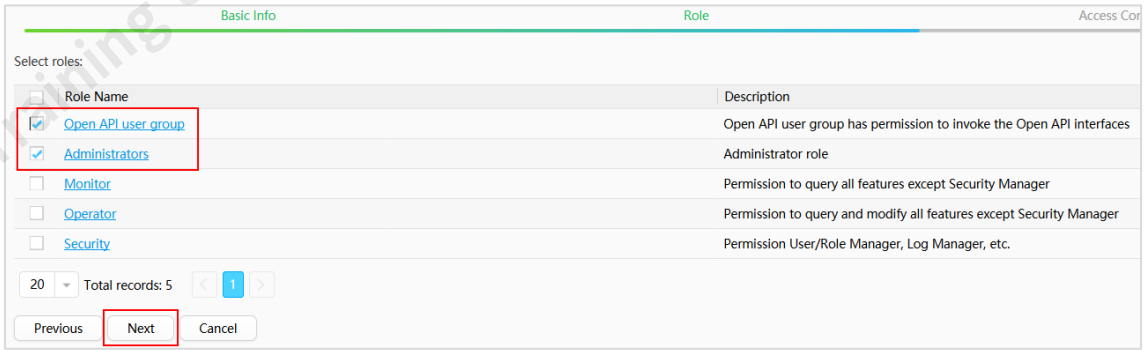

Click **Next**. The interconnection account is created.

Use Remote Desk Connection to log in to the eSight server.

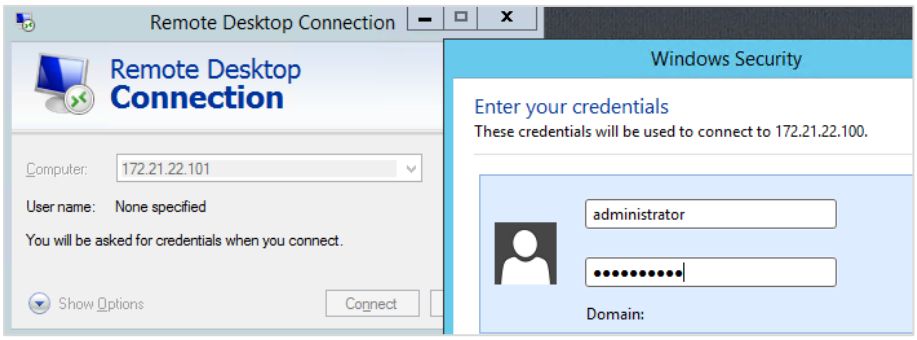

Double-click the eSight console icon and click **Stop**.

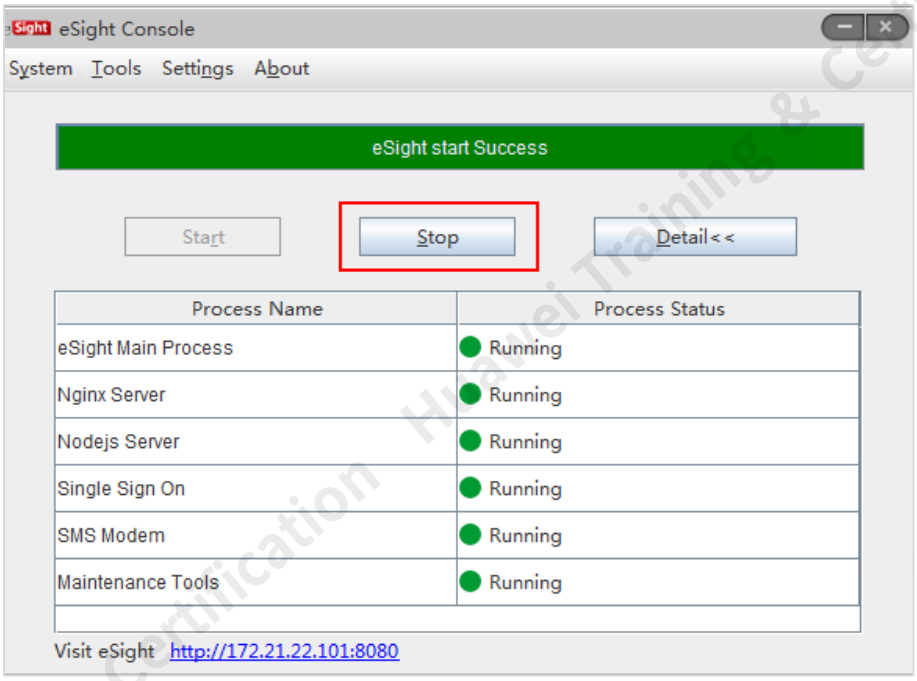

After the eSight is stopped, open the **eSight installation directory\AppBase\etc\oms.nbi\pm** directory on the eSight server and add the **oc_north.xml** configuration file. The content of the file is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<plugin>
    <export all="true" device="*" />
</plugin>
```

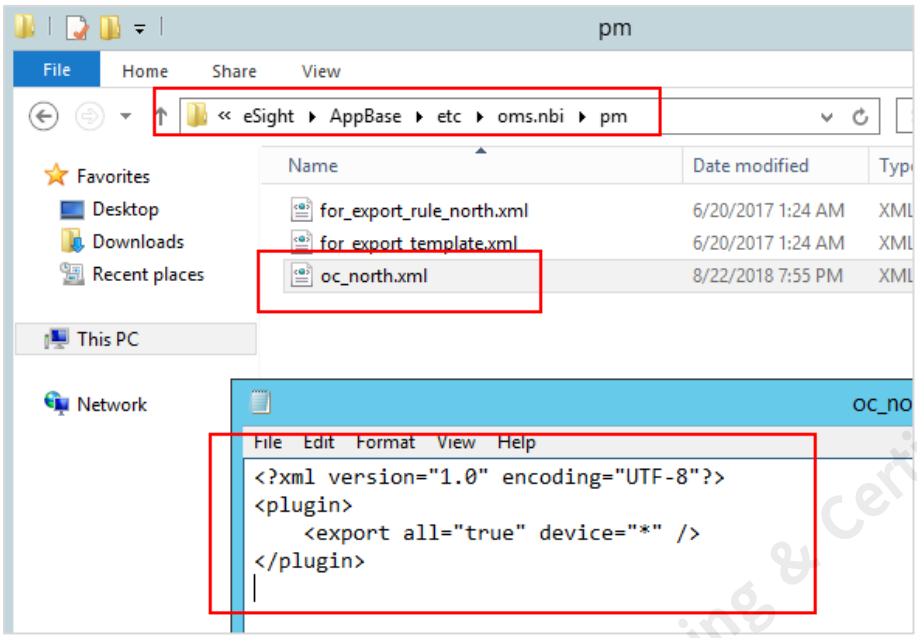

Restart the eSight service.

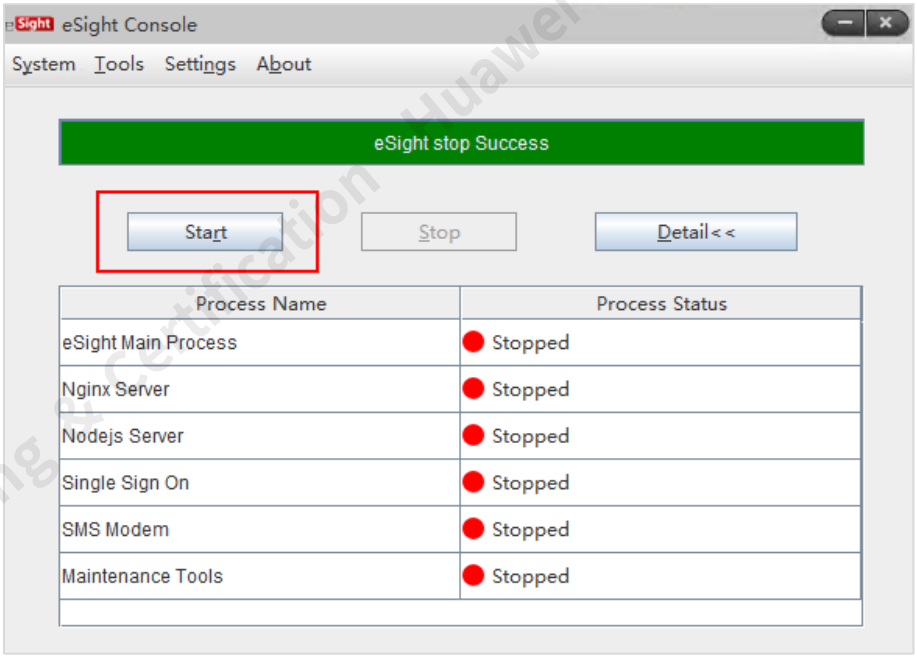

Configure the SNMP, whitelist, and SFTP performance collection information of eSight.

Change the SNMP service IP address to the service IP address of eSight.

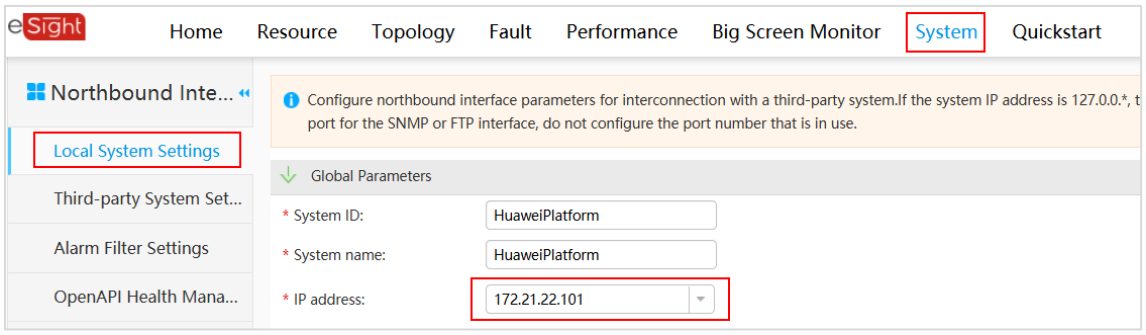

Click **Apply**.

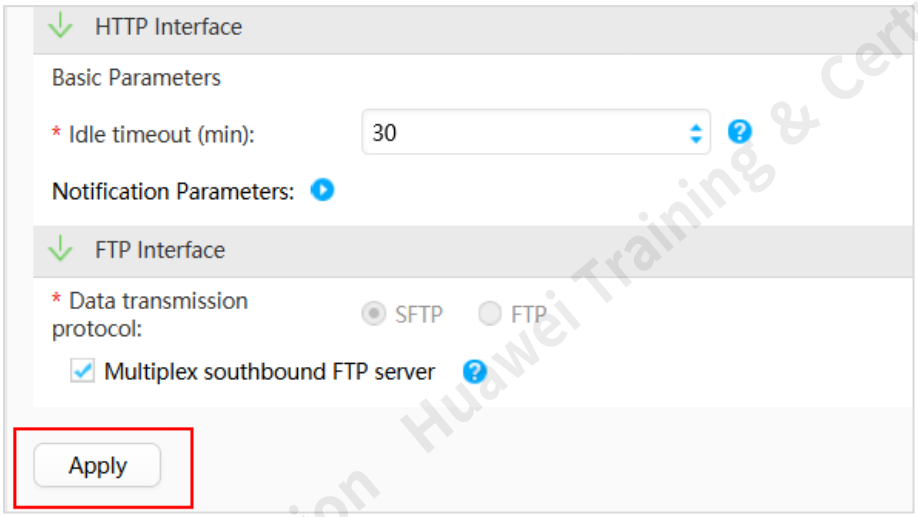

Choose Third-party System Settings > Create.

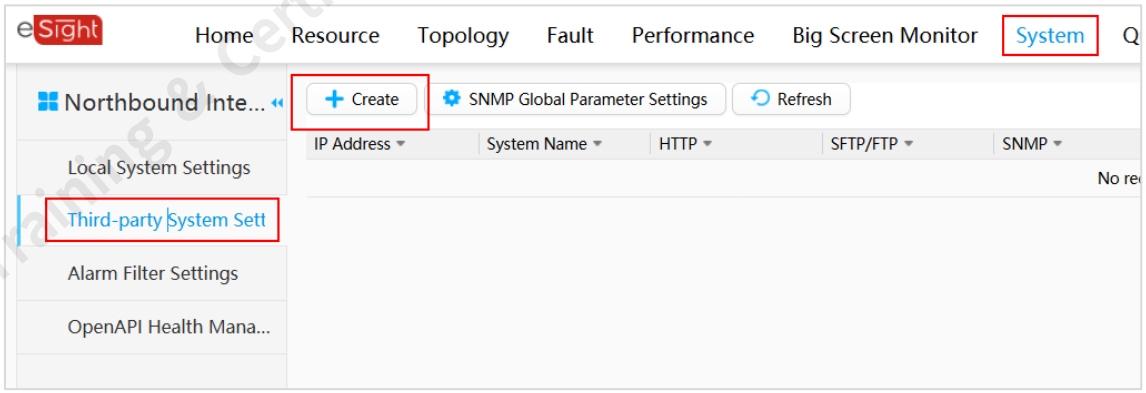

Set the IP address to the floating IP address of OperationCenter. You are advised to select all the protocol types. Ensure that the protocol template parameters are set to the same as the local SNMP. Set SFTP parameters based on the planning.

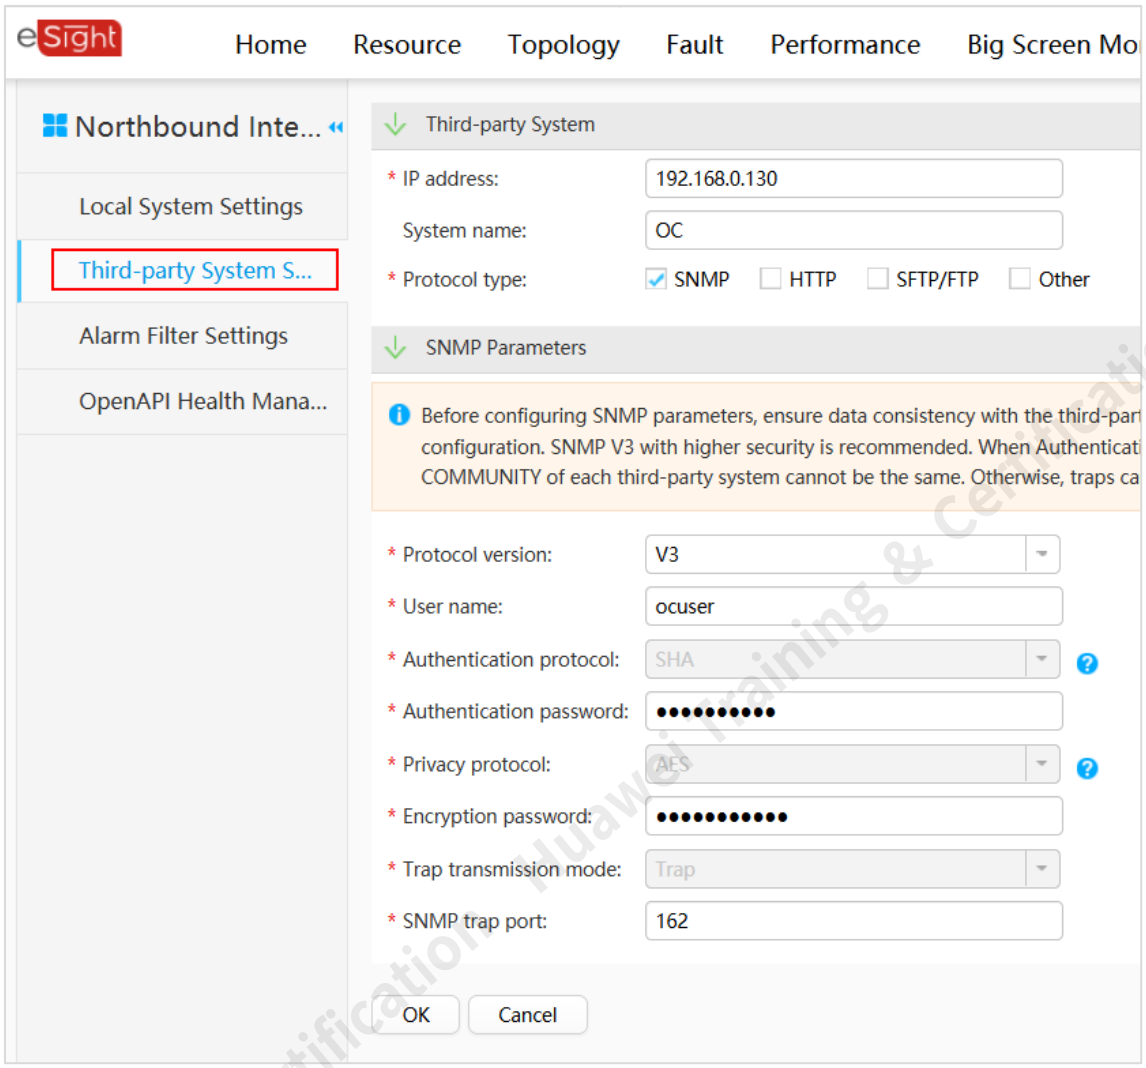

When configuring alarm reporting from eSight to OperationCenter, disabling event reporting.

Click Alarm Filter Setting. Select Event under Alarm/Event Types.

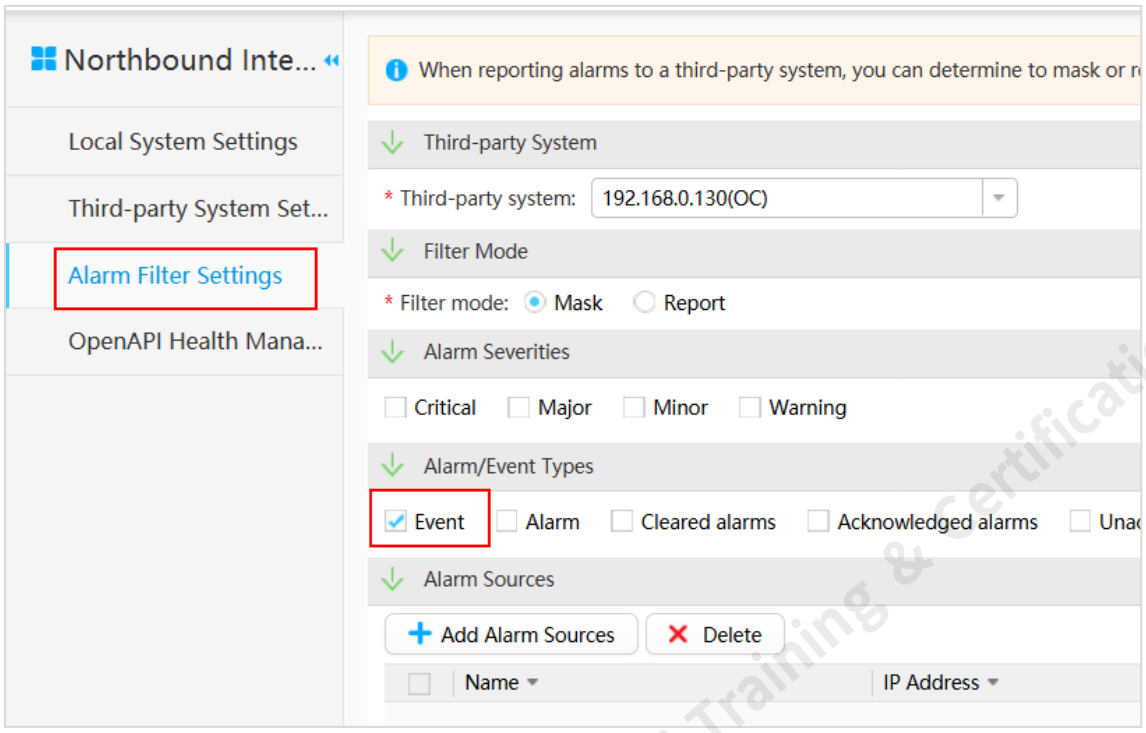

On the OperationCenter page, choose System > Connection Management > Connect to Peer System > Add.

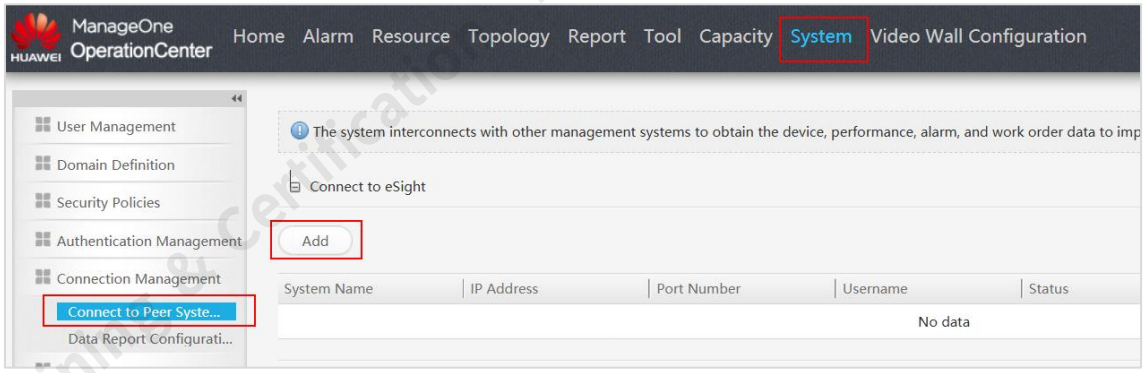

In the **Connect to eSight** dialog box, set parameters as planned.

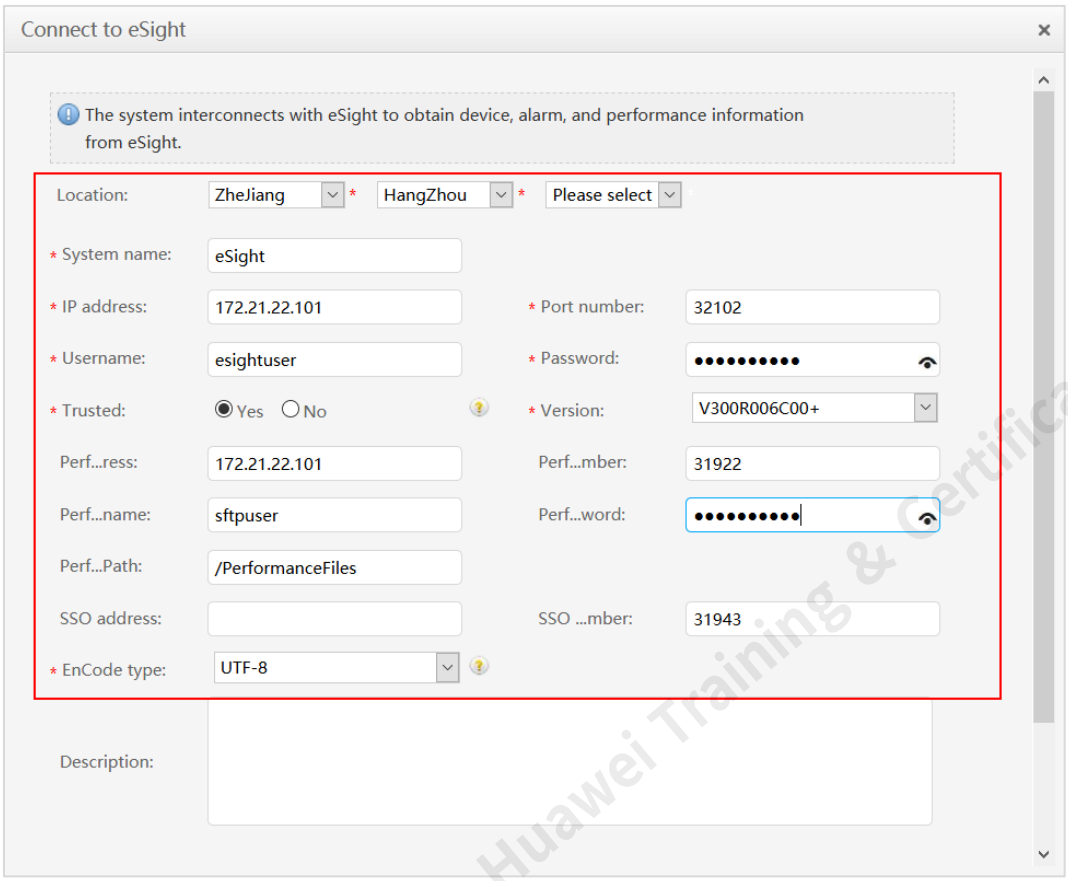

Click **Test**. Import the certificate.

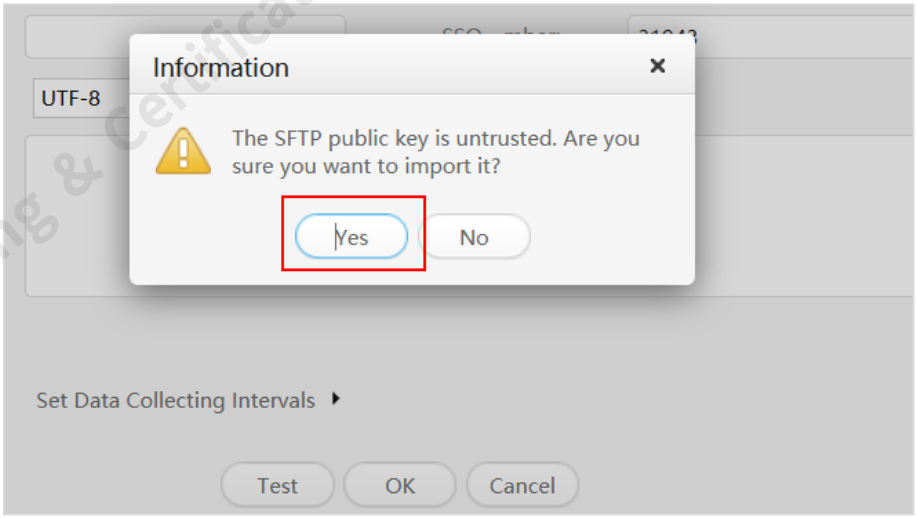

Click **Test** again.

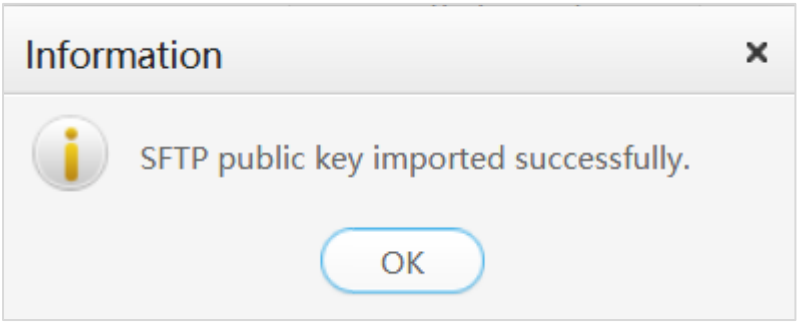

Click **OK**. eSight is connected.

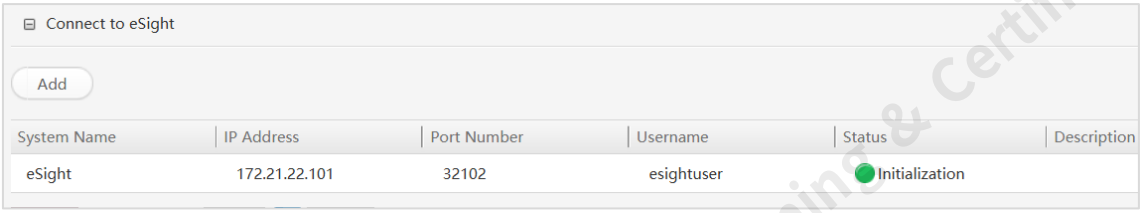

Step 3 **Connect OperationCenter with FusionSphere OpenStack.**

By connecting with the FusionSphere OpenStack system, OperationCenter can collect and manage information about devices, alarms, and performance from the FusionSphere OpenStack system to comprehensively analyze the running status of physical devices in a DC.

The following table lists interconnection parameters.

Parameter	Description	Reference Value
IP Address	IP address of the external API network plane of the host running FusionSphere OpenStack	192.168.0.2
Domain	FusionSphere OpenStack domain name	az1.dc1.domainname.com
Location	Location of the cloud platform	Zhejiang > HangZhou
URL address	On the installation and deployment page of FusionSphere OpenStack, choose **Configuration** > **System** > **Domain Name**, and click **Change** to query the domain name configured for Keystone.	https://identity.az1.dc1.domainname.com:443/identity
Default Project	Default tenant who accesses FusionSphere OpenStack resources. The tenant has the rights to access all resources.	admin
Username	Project administrator	cloud_admin

Parameter	Description	Reference Value
Password		FusionSphere123
Version	FusionSphere OpenStack version	V100R006C10

Interconnect procedure

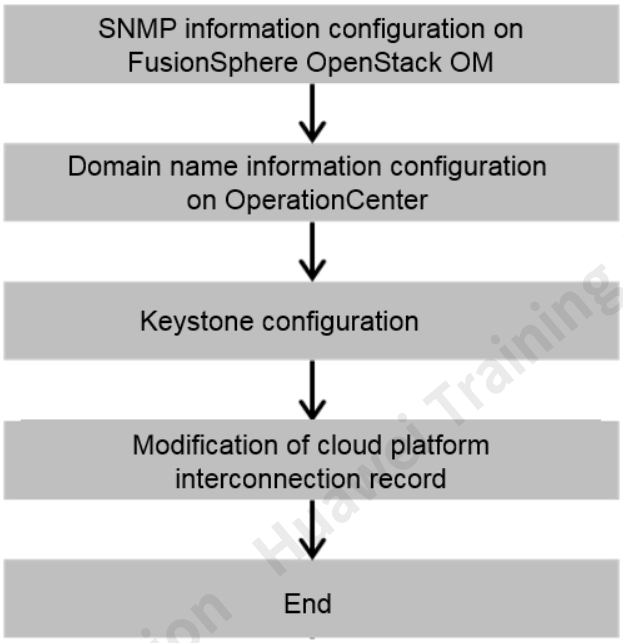

Log in to FusionSphere OpenStack OM, choose Monitoring > Alarm Settings > SNMP Station > Add SNMP Station.

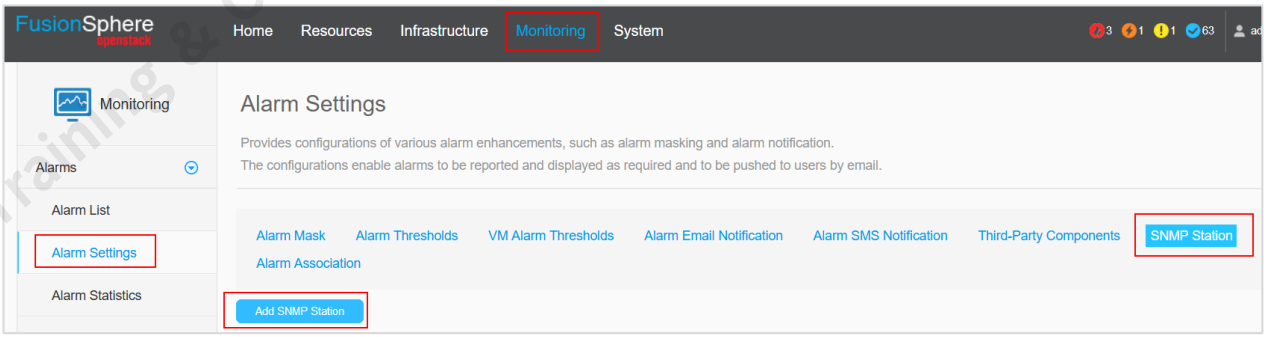

Set the SNMP station information based on the OperationCenter planning information.

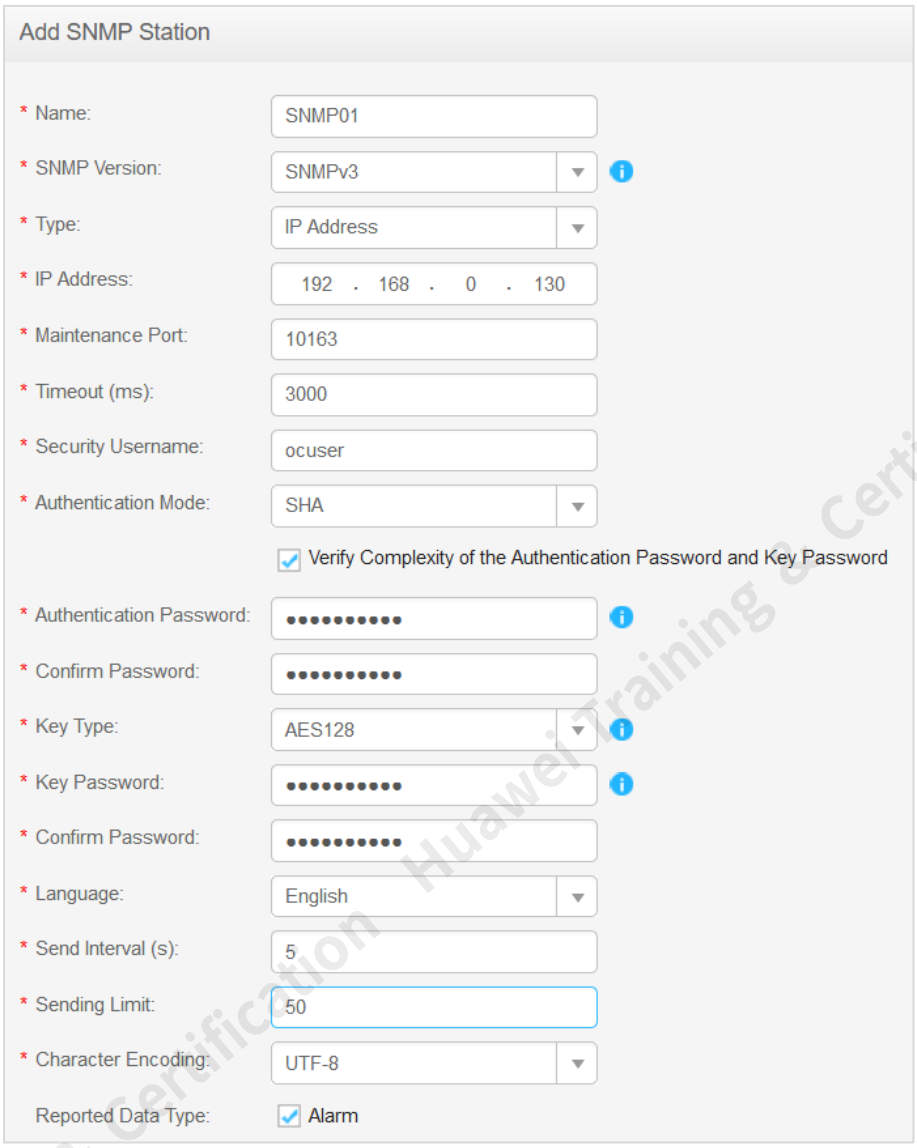

Configure interconnection information on OperationCenter.

On the OperationCenter page, choose **System** > **Authentication Management** > **Configure Keystone**. Configure IP address and domain information.

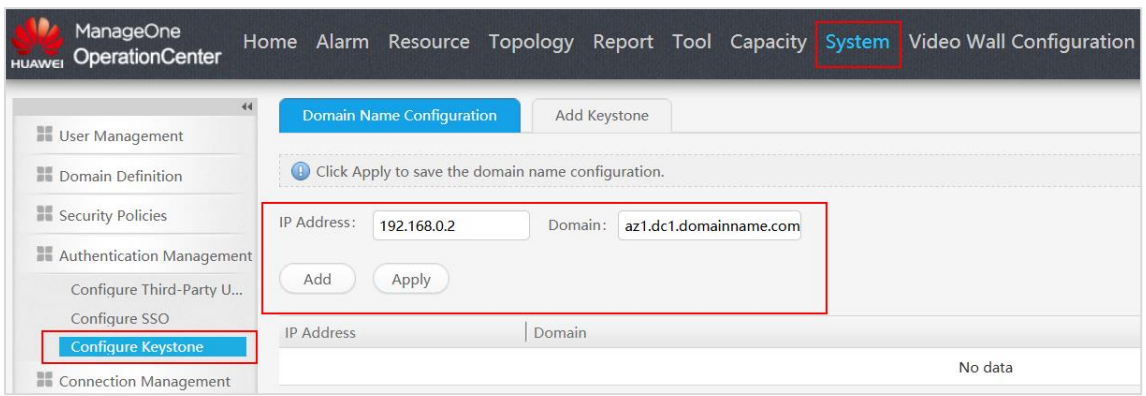

- Add Keystone.

Click **Add** and then **Apply**.

Add Keystone, fill in the information based on the planned data, click **Test**, and click **Yes** to confirm certificate importing.

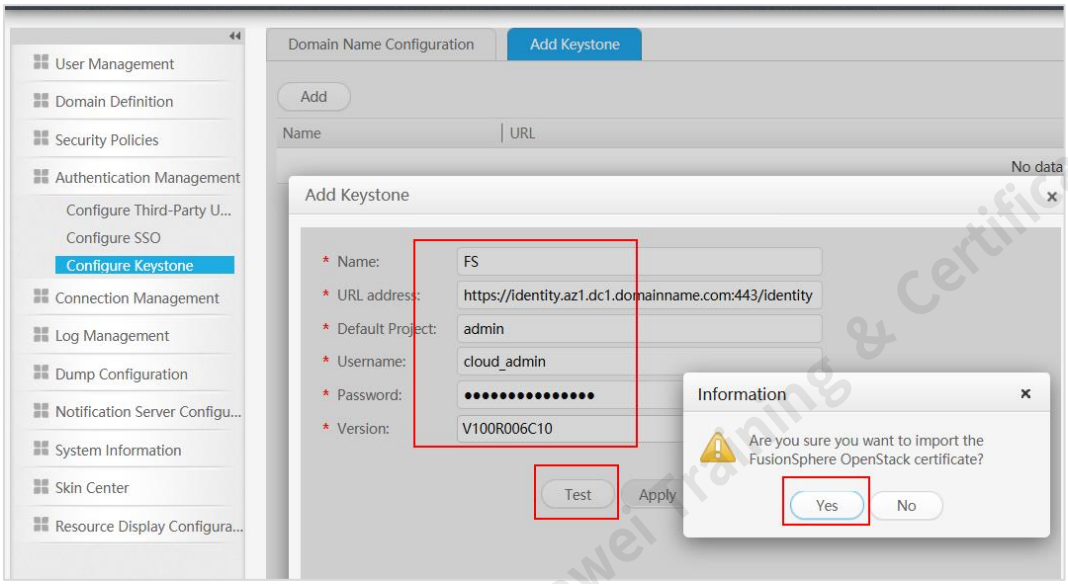

The connectivity test is successful. Click **Apply**. Keystone is added. (You can view the interconnection parameters.)

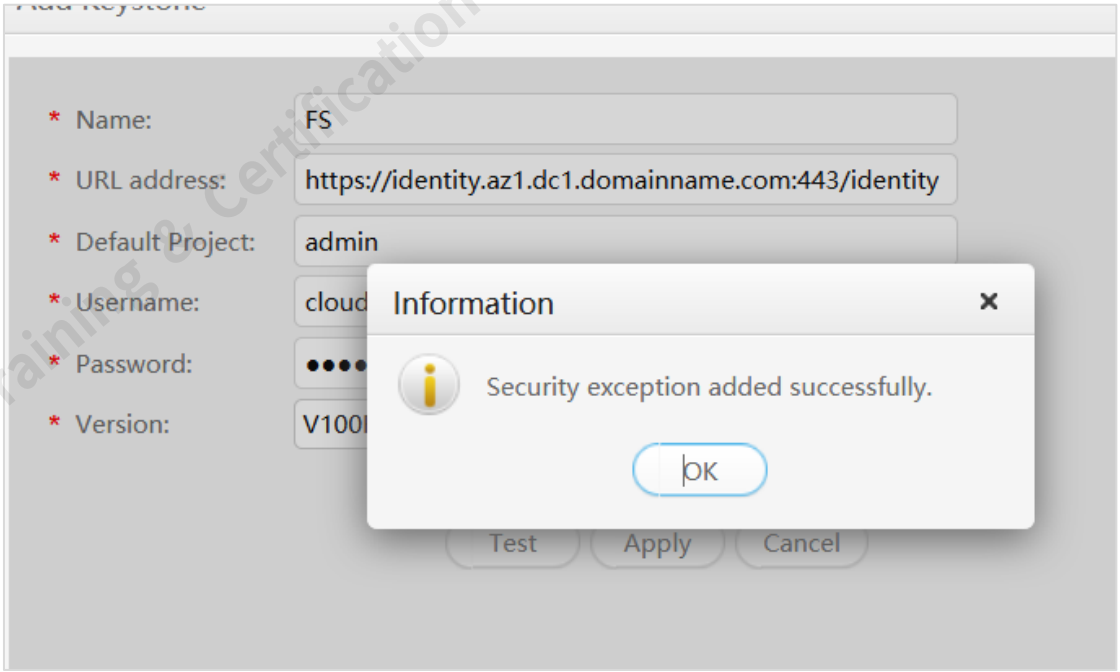

Modify FusionSphere OpenStack interconnection records.

On the OperationCenter page, choose **System** > **Connection Management** > **Connect to Peer System**. Click the button.

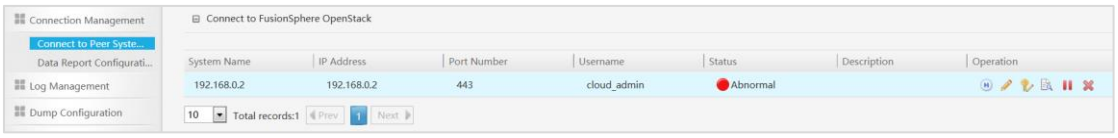

Enter the location information.

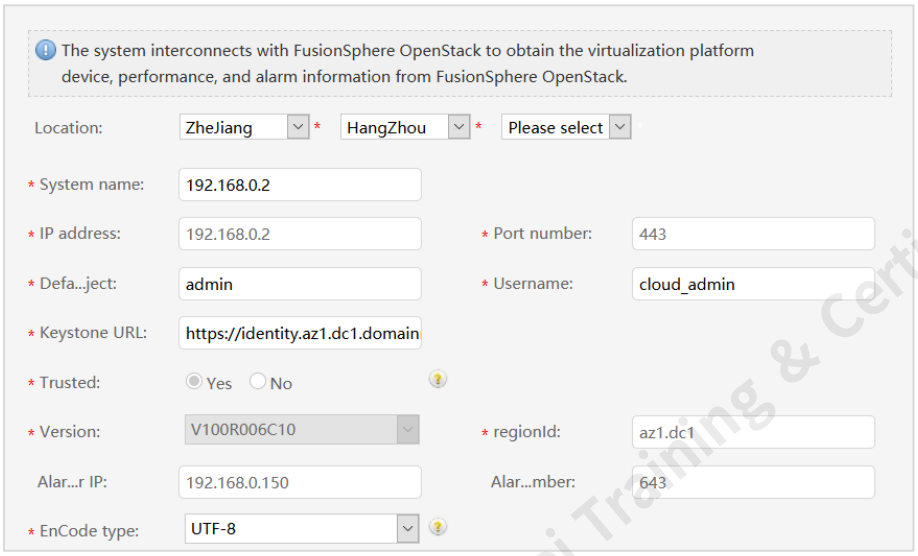

Scroll down the page, click **Test**, and click **OK** in the displayed Information box.

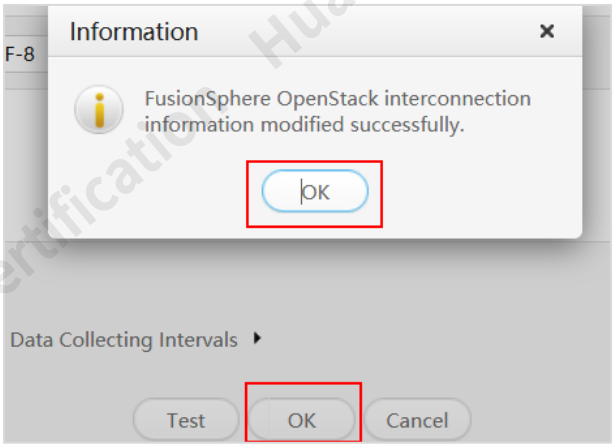

OperationCenter and FusionSphere OpenStack are interconnected.

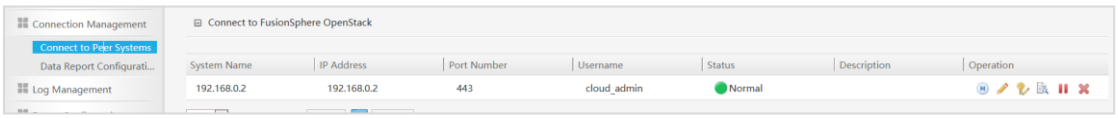

Choose **Resource** > **View Resource** to view cloud platform resource information.

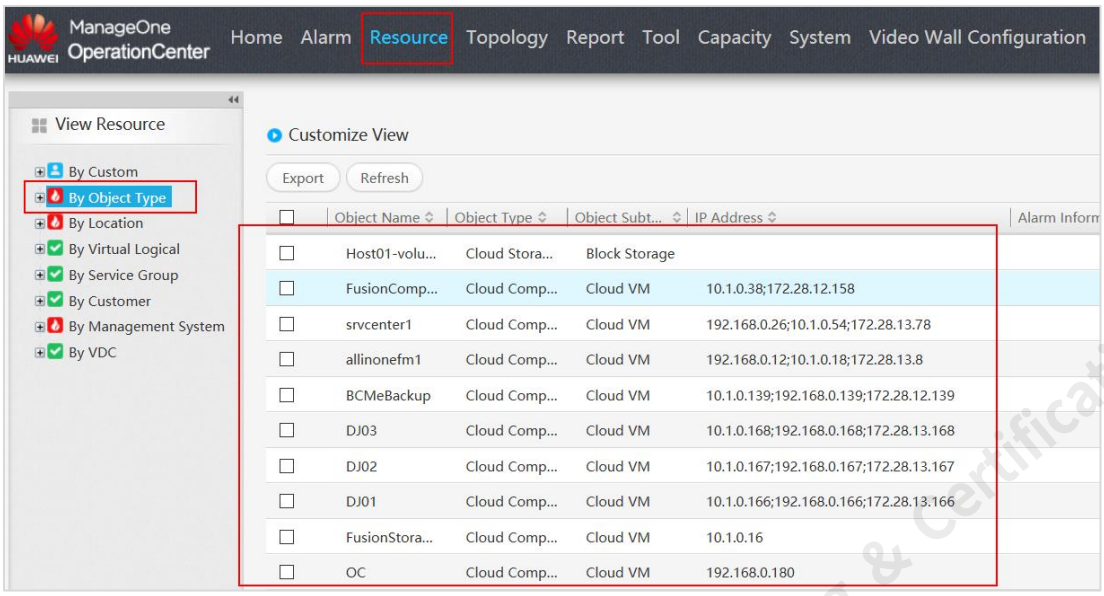

Step 4 **Connect OperationCenter with ServiceCenter.**

The following table lists interconnection parameters.

Parameter	Description	Reference Value
System name		SC
IP address/Port number		192.168.0.170 / 643
Tena...mber		543
Username/Password	**Username** and **Password** have been configured on ManageOne OperationCenter.	OCRest / Huawei@CLOUD8!!

Log in to ServiceCenter.

Enter https://192.168.0.170:643/ into the browser. The default username is **cloud_admin** and password is **FusionSphere123**.

Configure the SNMP information

Log in to the ServiceCenter web client and choose **System** > **System Configuration** > **SNMP Station**.

Ensure that the SNMP information is the same as that on OperationCenter.

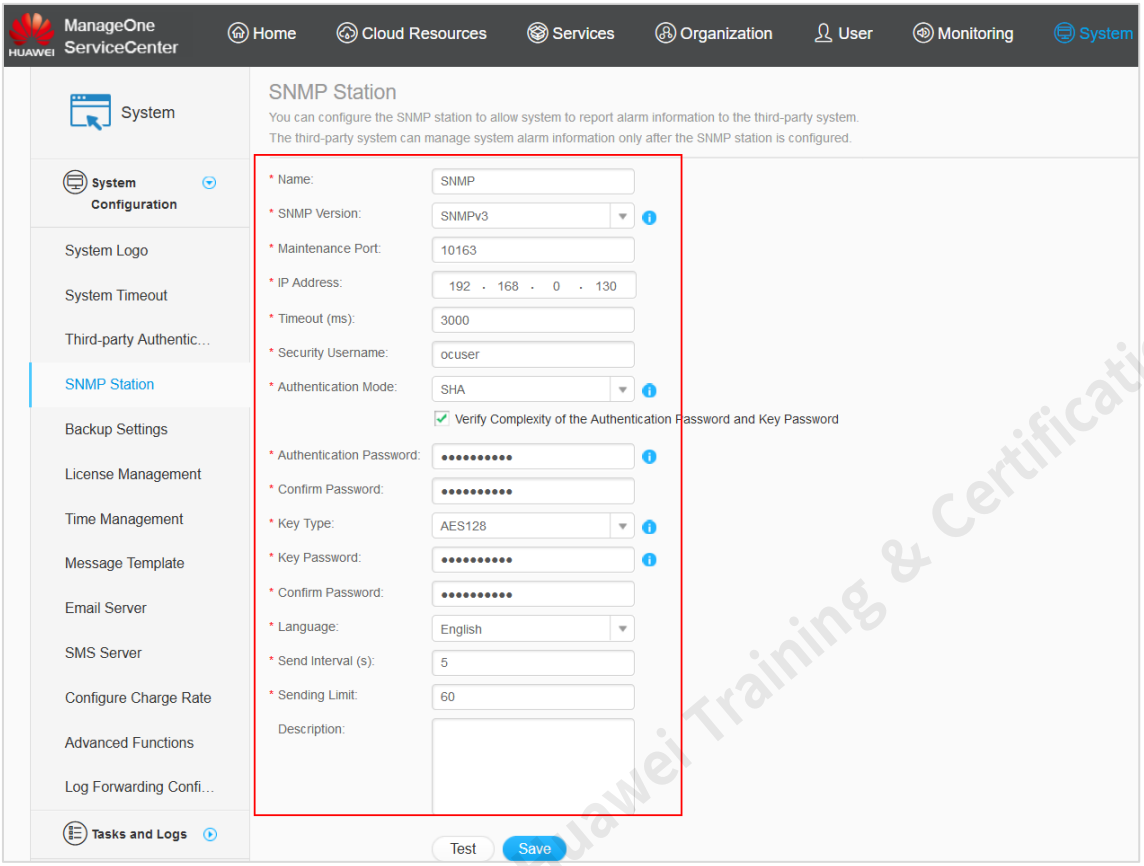

Set the parameters as shown in the preceding figure and click **Save**.

Unlock the interconnection user **OCRest**. Log in to ServiceCenter using VNC on the CPS as user **galaxmanager**. Use the password **Huawei@CLOUD8**.

Run the **unlockSysAccount OCRest** command to unlock the interconnection user.

```
srvcenter login: galaxmanager
Password:
srvcenter:~ $ TMOUT=0
srvcenter:~ $ unlockSysAccount OCRest
[2017-09-29 11:34:45] [ NOTICE] This running is a new time of "unlockSysAccount"
 ...
This shell script unlock system account.
[2017-09-29 11:34:45] [ INFO] unlock system account OCRest begin.
[2017-09-29 11:34:49] [ INFO] unlock system account OCRest success.
[2017-09-29 11:34:49] [ INFO] unlock system account OCRest finish.
srvcenter:~ $
```

On the OperationCenter page, choose System > Connection Management > Connect to Peer System.

In the right pane, click **Connect to ServiceCenter**, and click **Add**.

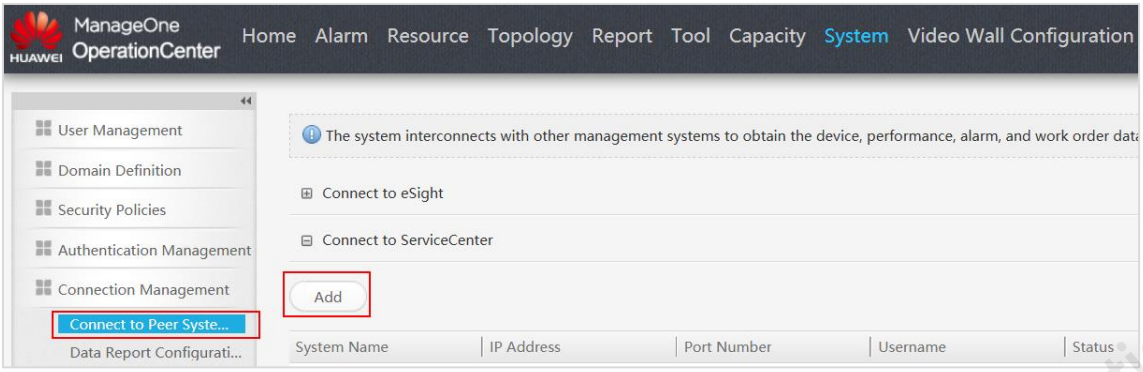

Set parameters as planned. (The username is **OCRest** and the default password is **Huawei@CLOUD8!!**.)

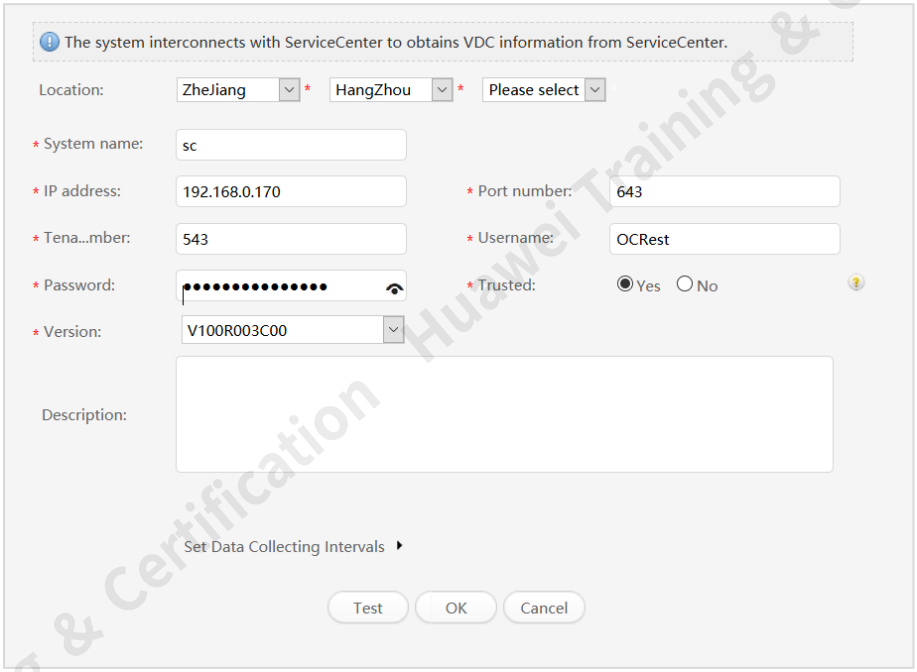

Click **Test**, import the certificate, and perform the test again.

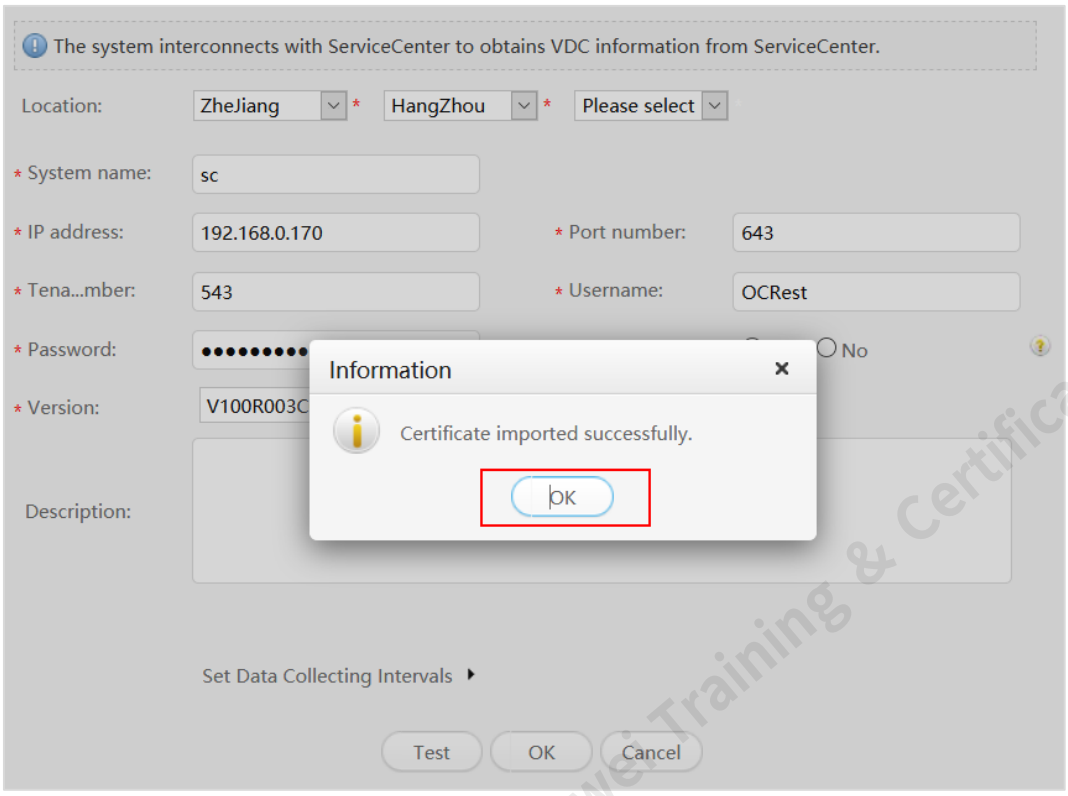

The connection is successful.

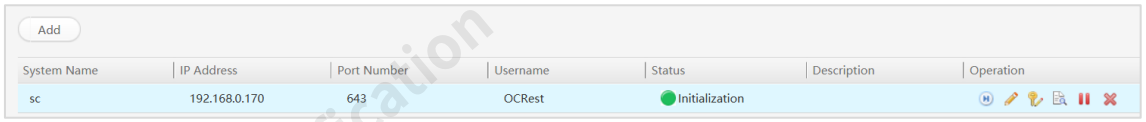

Step 5 **Connect OperationCenter with eBackup Manager.**

Enable the eBackup Manager port.

The port **7** is used to test the connectivity between the two hosts of ManageOne OperationCenter and eBackup.

Log in to Manager as user **hcp** in SSH mode. The default password is **Huawei@CLOUD8!**.

```
hcp@linux-NVjHWa:~> TMOUT=0
```

Run the **su root** command and enter the password (**Huawei@CLOUD8!**) of user **root** as prompted to switch to user **root**:

Run the **TMOUT=0** command to disable automatic logout upon timeout.

Run the following command to enable port **7**:

iptables -I INPUT -s source IP address -d destination IP address -p tcp --dport 7 -j ACCEPT

```
hcp@linux-NVjHWa:~> su root
Password:
linux-NVjHWa:/home/hcp # iptables -I INPUT -s 10.1.0.139 -d 192.168.0.180 -p tcp
 --dport 7 -j ACCEPT
linux-NVjHWa:/home/hcp #
```

In the preceding information, the source IP address is the IP address of the ManageOne OC host, and the destination IP address is the management plane IP address of Manager.

Run the **vi /etc/sysconfig/iptables** command to open the **/etc/sysconfig/iptables** file.

Add firewall rules to the first line of the ***filter** field in the file. (Note: Enter **i** to enter the modification mode and press **ESC** to exit. Enter **:wq** to save and exit the file.)

-A INPUT -s source IP address -d destination IP address -p tcp -m tcp --dport 7 -j ACCEPT

In the preceding information, the source IP address is the network segment of the ManageOne OC host, and the destination IP address is the management plane network segment of Manager.

```
-A INPUT -s 10.1.0.139 -d 192.168.0.180 -p tcp -m tcp --dport 7 -j ACCEPT
```

Run the **cd /opt/huawei-data-protection/ebackup/bin** command to go to the directory where the script resides.

Run the following command to configure firewall rules:

sh iptablesHelper.sh accept HCPManagementPlane Source IP address Destination IP address

In the preceding information, the source IP address is the IP address or network segment of the ManageOne OperationCenter host, and the destination IP address is the management plane IP address or network segment of Manager.

During the operation, enter the role required to configure the firewall rule as prompted. Enter the number corresponding to Backup Manager. (See the following figure.)

```
linux-NVjHWa:/home/hcp # cd /opt/huawei-data-protection/ebackup/bin
linux-NVjHWa:/opt/huawei-data-protection/ebackup/bin # sh iptablesHelper.sh acce
pt HCPManagementPlane 10.1.0.139 192.168.0.180
Please input the role of current node, we will configure for this node type.
0 -- Backup server
1 -- Backup Proxy
2 -- Backup Manager
3 -- Backup Workflow Server
Please input 0, 1, 2 or 3 to select the node type
2
Run command: iptables -I INPUT -s 10.1.0.139 -d 192.168.0.180 -p icmp -j ACCEPT
Run command: iptables -I INPUT -s 10.1.0.139 -d 192.168.0.180 -p tcp --dport 22
-j ACCEPT
Run command: iptables -I INPUT -s 10.1.0.139 -d 192.168.0.180 -p tcp --dport 808
8 -j ACCEPT
Run command: iptables -I INPUT -s 10.1.0.139 -d 192.168.0.180 -p tcp --dport 808
0 -j ACCEPT
Run command: iptables -I INPUT -s 10.1.0.139 -d 192.168.0.180 -p udp --dport 123
 -j ACCEPT
Succeed to save iptables
linux-NVjHWa:/opt/huawei-data-protection/ebackup/bin #
```

Configuration on OperationCenter

Disable the northbound two-way authentication of the OperationCenter server.

Use PuTTY to remotely log in to the server where OperationCenter resides as user **appuser**.

The default password of user **appuser** is **Changeme_123**.

Run the **sudo su – root** command to switch to user **root**.

Enter the password of user **root**. (Note: When the OS and OperationCenter are installed in separated mode, the default password for user **root** to log in to OperationCenter is **UVP_2012lab**. When the OS and OperationCenter are installed in integrated mode, the default password for user **root** to log in to OperationCenter is **Changeme_123**.

Run the **TMOUT=0** command to prevent PuTTY from exiting due to timeout.

Run the **:service oc stop** command to stop the OperationCenter system.

Run the **cd /opt/OperationCenter/bin** command to switch the path.

Run the **sh restClientAuth.sh** command to set the northbound two-way authentication switch of OperationCenter. (The switch status can be **on** or **off**.)

Run the **service oc start** command to start the OperationCenter system.

```
appuser@OCHost1:~> TMOUT=0
appuser@OCHost1:~> su - root
Password:
OCHost1:~ # service oc stop
Stopping...
Stopped Successfully.
OCHost1:~ # cd /opt/OperationCenter/bin
OCHost1:/opt/OperationCenter/bin # sh restClientAuth.sh off
Checking OperationCenter status...
Turn off rest client auth successfully.
OCHost1:/opt/OperationCenter/bin # service oc start
Starting...
The OperationCenter is starting, which takes about 15 minutes. Please wait.
Started Successfully.
```

After the command is executed, run the **service oc status** command to check the OperationCenter system service status, and perform other operations only after all services are normal (no service is in the starting state).

```
OCHost1:~ # service oc status
GaussDB(single)                 ..............  running
Search Engine component         ..............  running
Cas component                   ..............  running
Report Server                   ..............  running
OperationCenter component       ..............  running
NTP                             ..............  running
HA                              ..............  running
OCHost1:~ #
```

Create an interconnection user.

Log in to ManageOne OperationCenter. You must create an account for connecting to the eBackup server on the ManageOne OperationCenter platform first. (Use the **Open API login** mode.)

Log in to OperationCenter, and choose **System** > **User Management** > **User**.

Click **Create**.

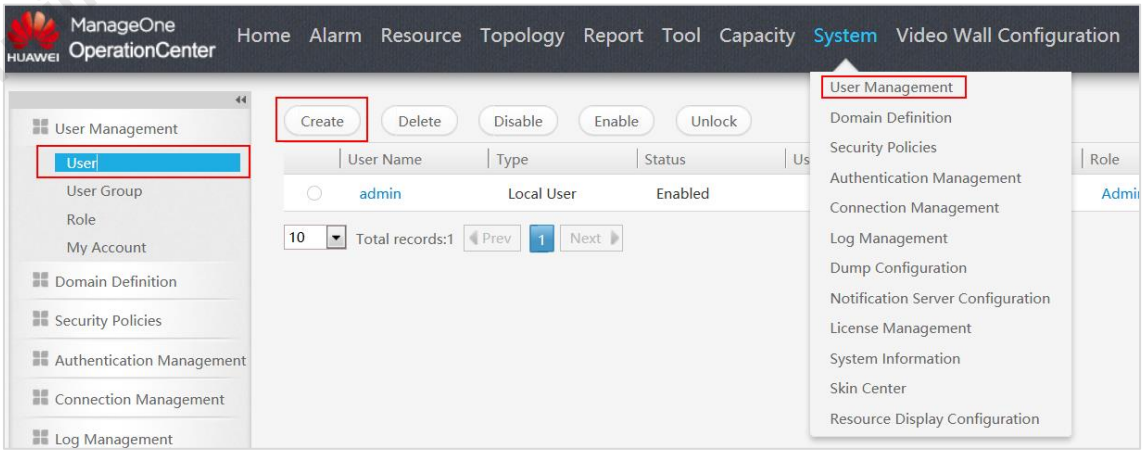

Set parameters as prompted.

Login mode: Open API login

User type: Local user

Username: Name of the user to be created, for example, eBackup

Password: Password of the user to be created, for example, Huawei@123

Confirm password: Enter the password of the user to be created again, for example, Huawei@123

Use the default values for other parameters.

Click **Next**.

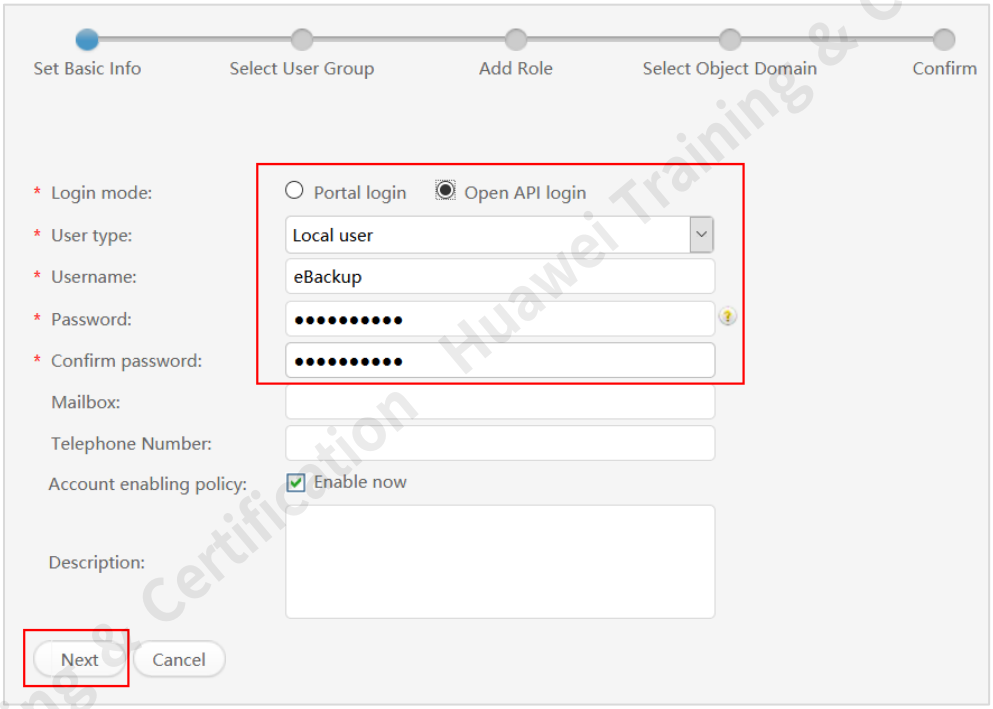

Select System management group. Click Next.

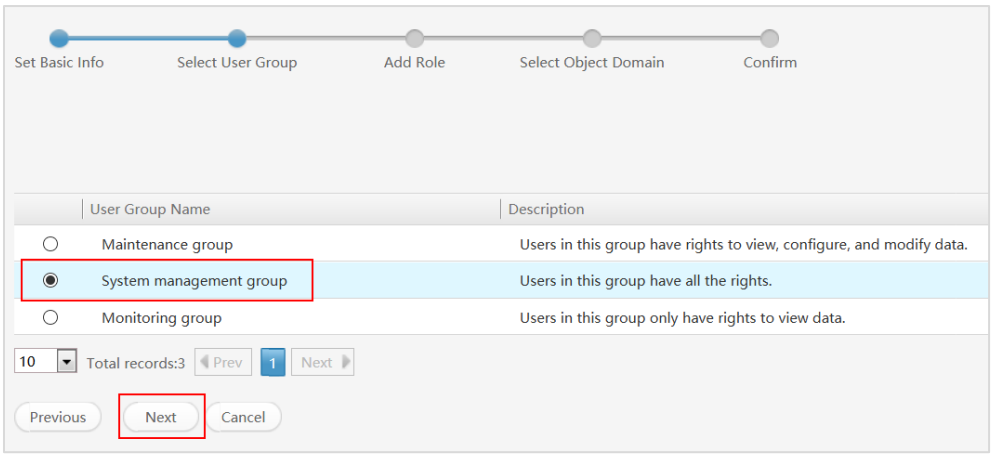

Confirm that all the information is correct and **Finish**. In the displayed **Information** box, click **OK**.

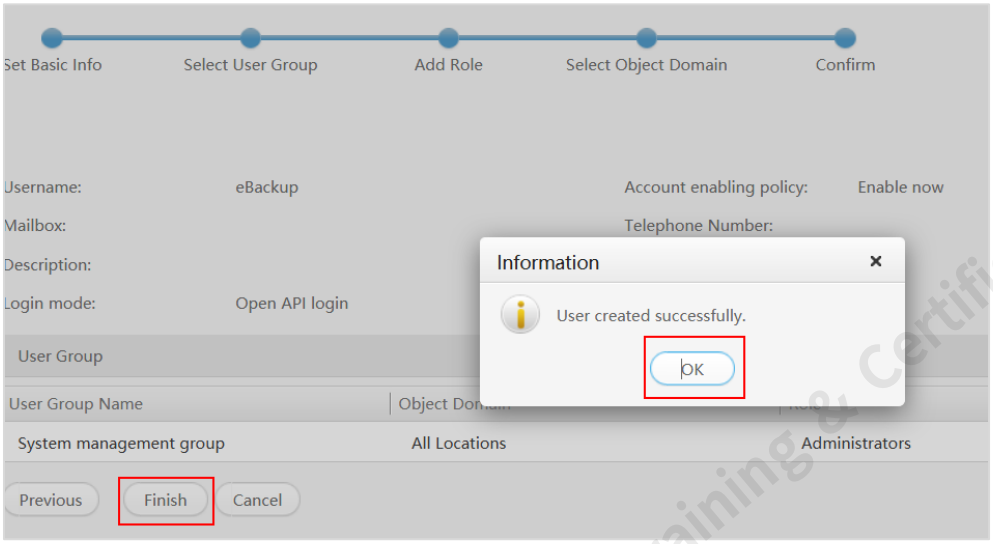

Set the security policy. In the navigation tree on the left, choose **Northbound Interconn...** and click **Add** on the right pane.

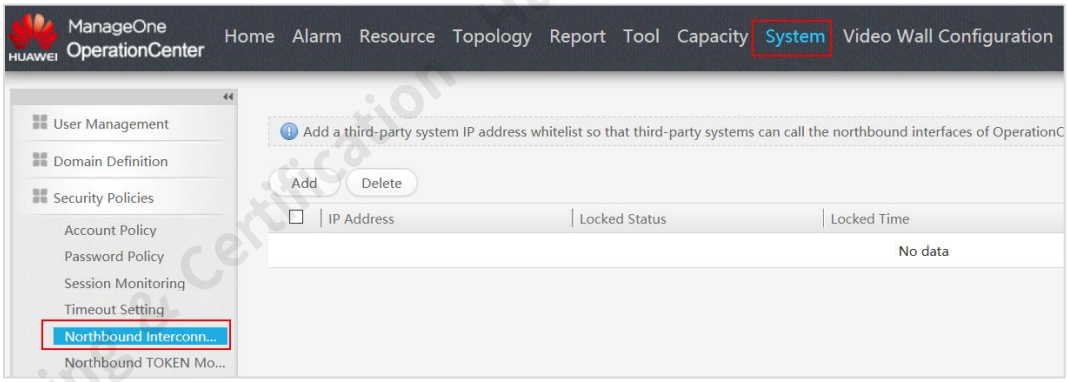

Add the backup management plane IP address of the eBackup server to the IP address whitelist of ManageOne OperationCenter. Click **OK**.

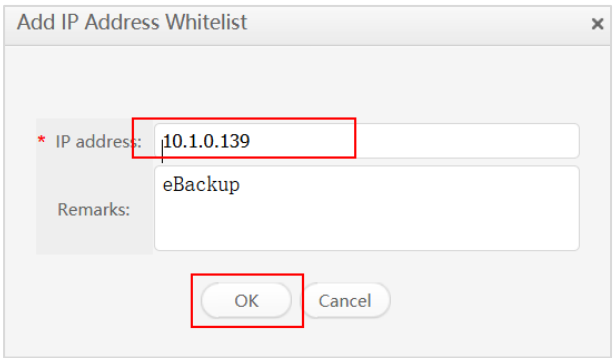

The IP address is added.

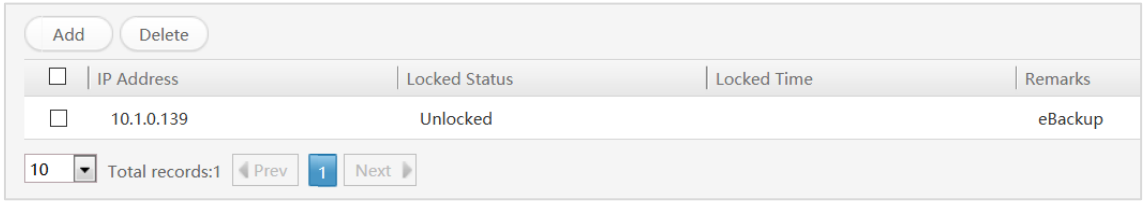

Configure Keystone. The configuration is described in Step 3 and is not described in this step.

BCManager interconnection

On the OperationCenter page, choose System > Connection Management > Connect to Peer System.

Click **Connect to BCManager** on the right pane to expand the **Connect to BCManager** area.

Click **Add**.

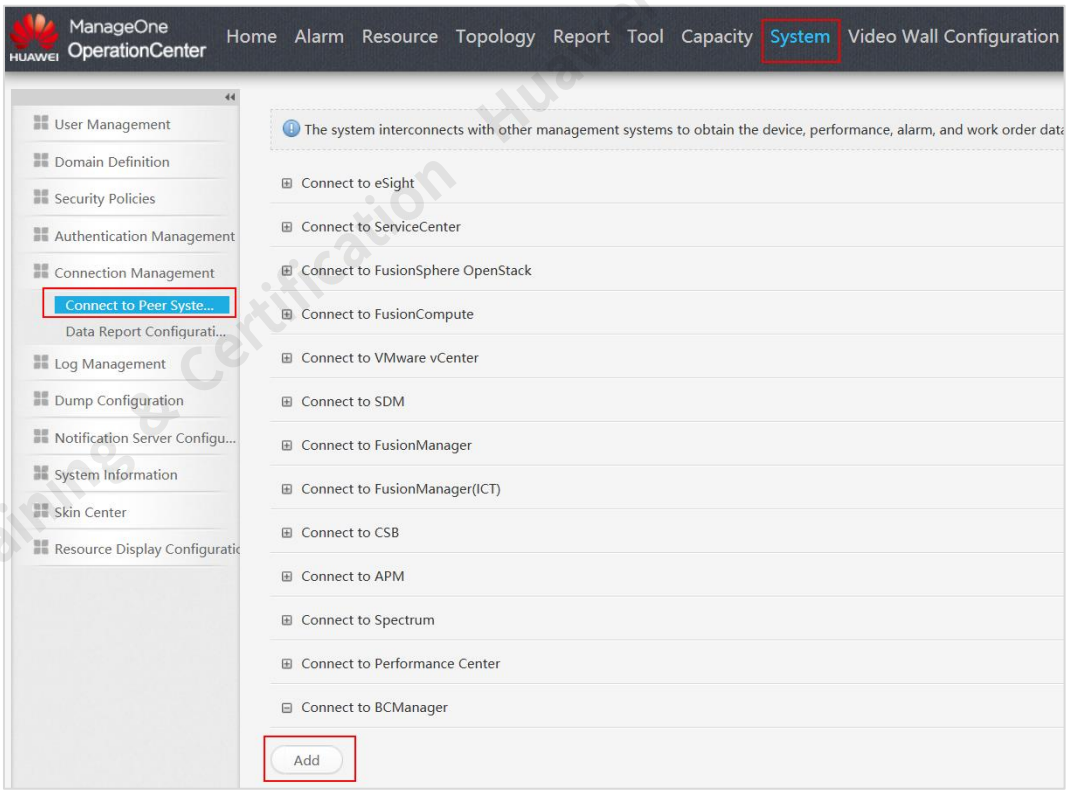

Click **Test** to obtain and import the CA certificate of ManageOne OperationCenter.

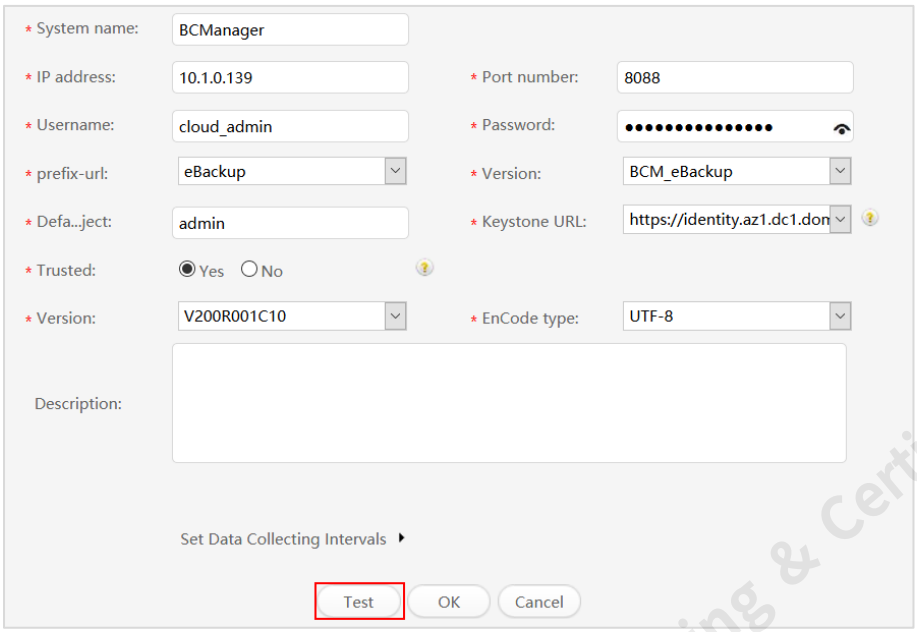

The **Connection established** information box is displayed.

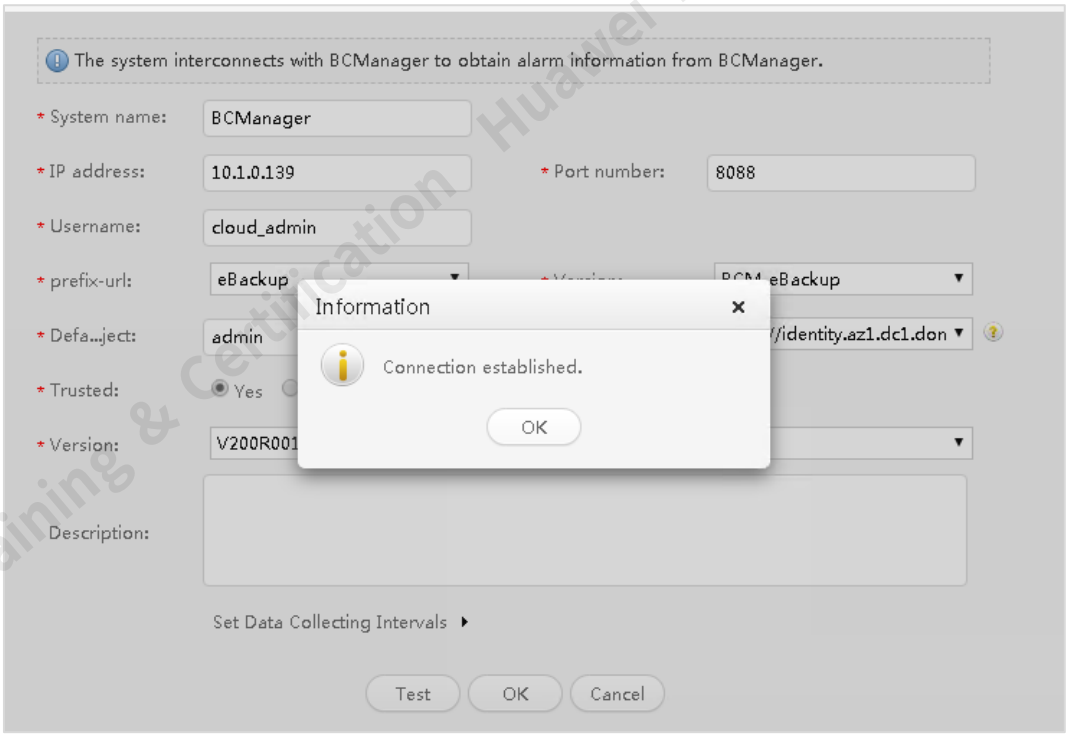

Click **OK**. BCManager is connected.

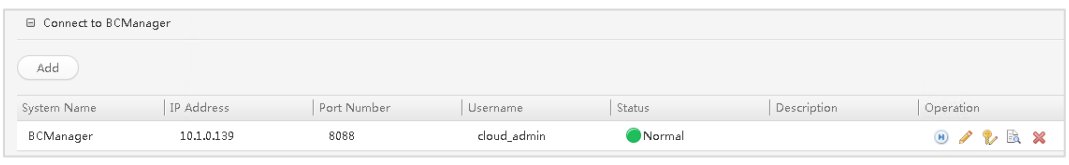

Visit https://192.168.0.170:31943/ and log in to OperationCenter.

Choose **Resource** > **View Resource**. The tree topology on the left displays monitored objects by interconnected system.

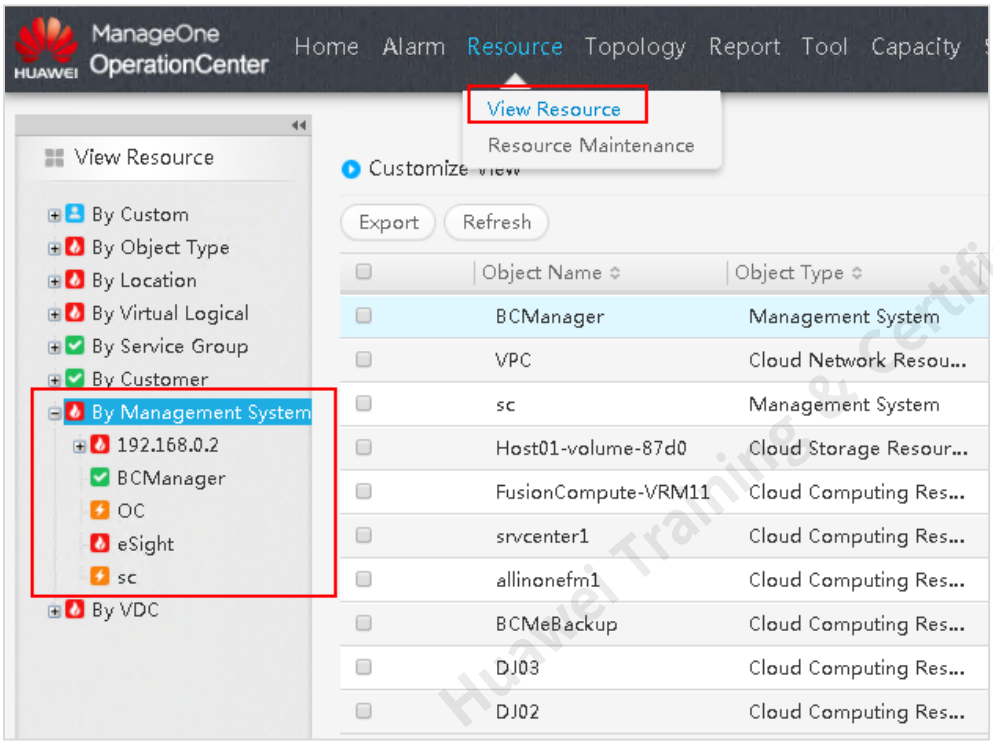

Network device monitoring results are uploaded to OperationCenter through eSight.

Choose **Topology** > **Physical Network Topology**. The physical topology that you set is displayed.

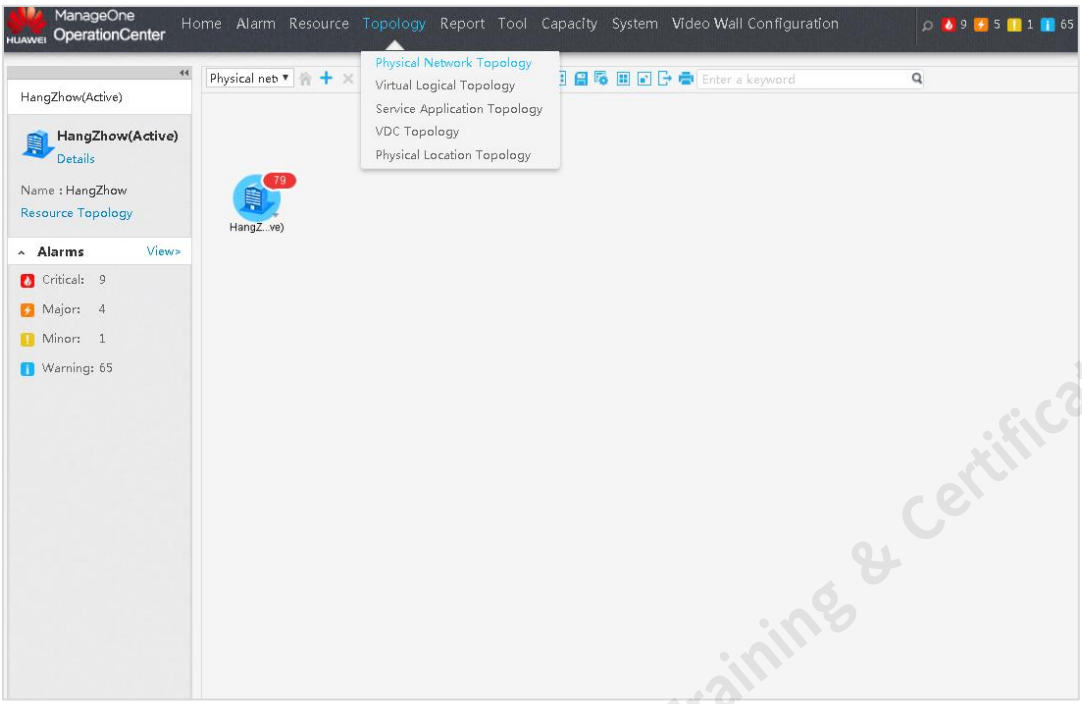

Choose **Topology** > **Virtual Logical Topology** to view the virtual topologies in the cloud platform.

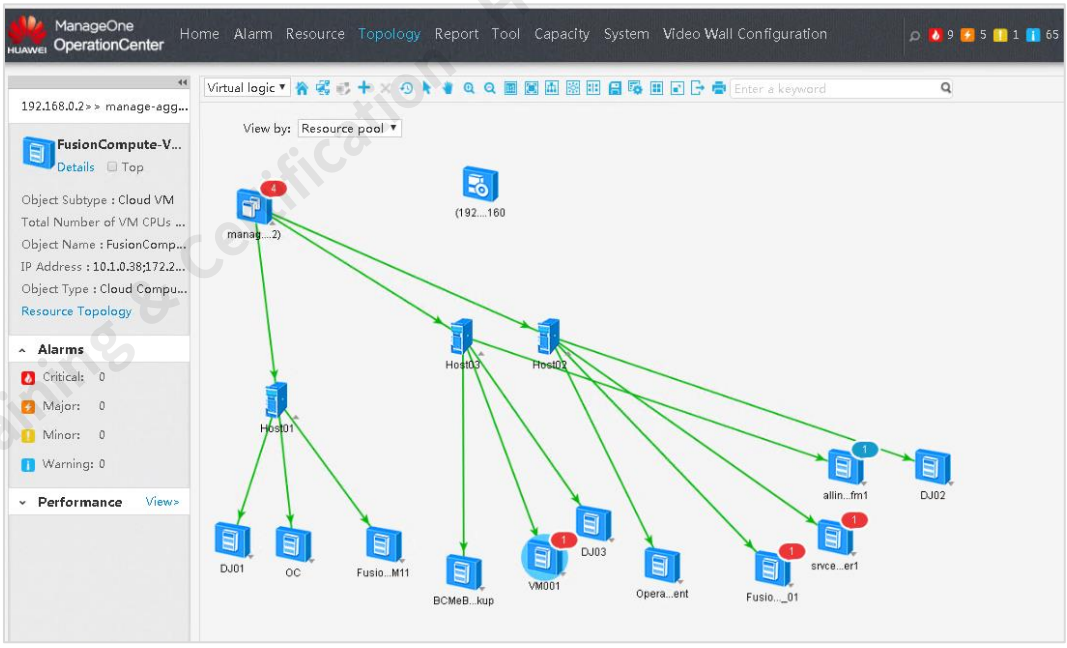

----End

9.2.5 Installing Analysis Tools

Analysis tools are used together with OperationCenter. Tools can be used to calculate and provide scientific data for administrators based on the system running status. Similar to OperationCenter installation, you need to register images, create flavors, and create VMs on FusionSphere OpenStack OM.

Step 1 **Register the installation image.**

Use the administrator username (**admin**) and password (**Huawei@123**) to log in to FusionSphere OpenStack OM. The URL is **https://192.168.0.150:643/center/src/index.html#/res/overview/image**. On the FusionSphere OpenStack OM homepage, choose **Resources** > **Images**.

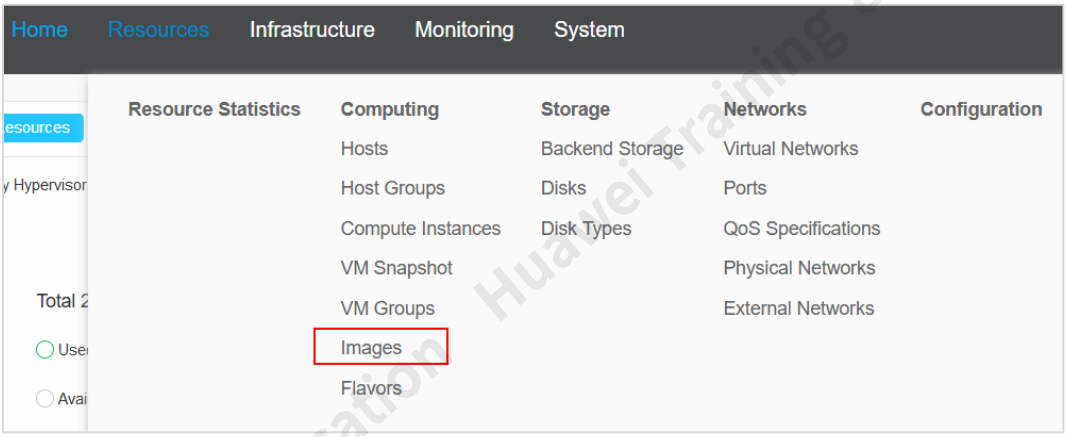

Click **Register** to access the page for setting image parameters and set them following configuration in the following figure. Set the OS version to **SUSE 11SPC3 64bit**.

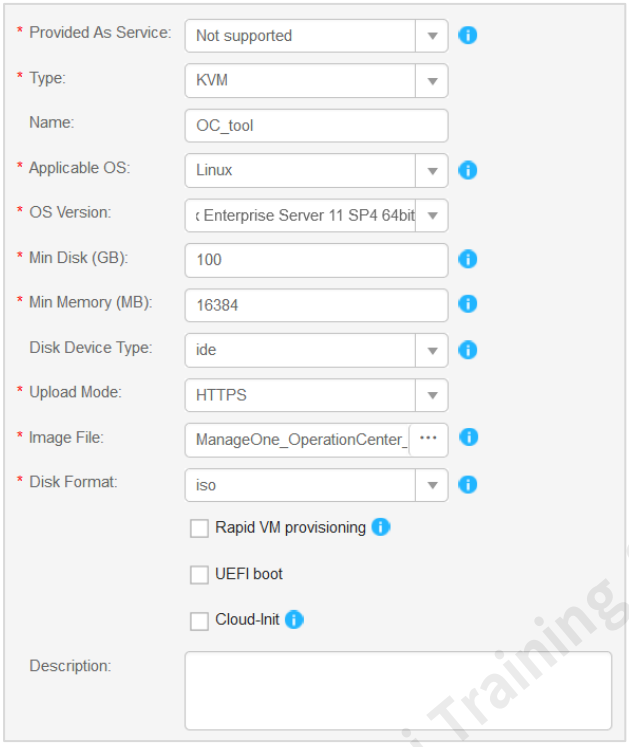

Click **Register** to complete the registration. The registration process may take an extended period of time. You can open a new page to complete the task of creating flavors.

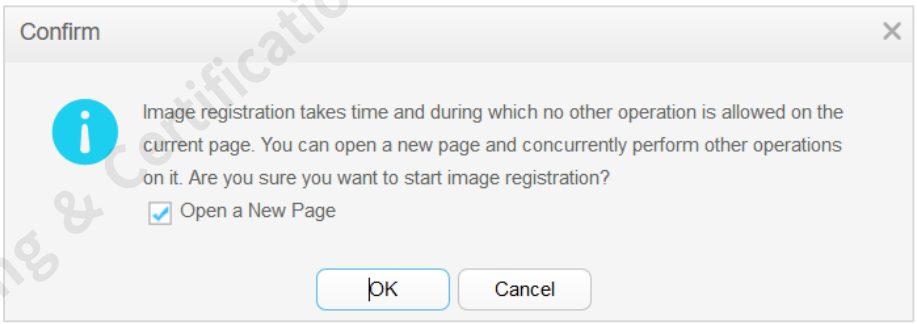

Step 2 **Creates flavors.**

On the new page, choose **Resources** > **Flavors**.

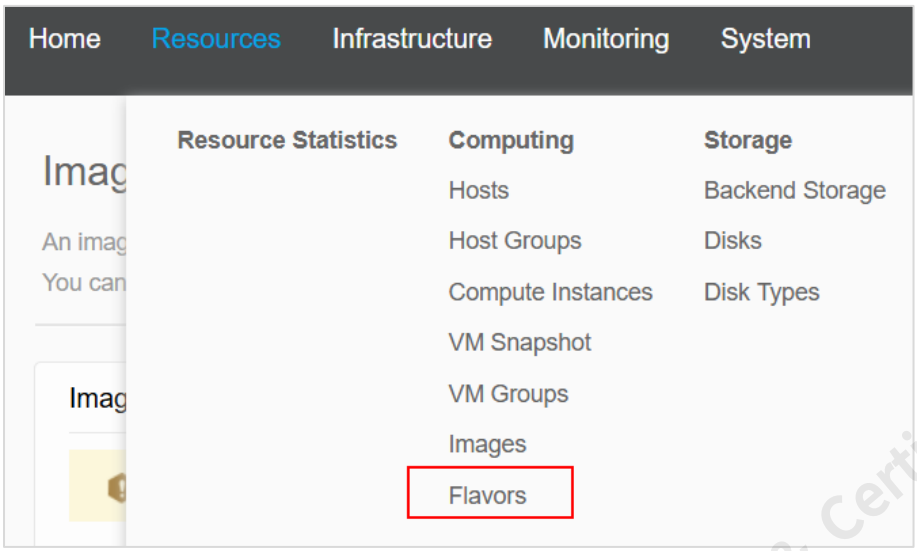

On the page that is displayed, click **Create Flavor**, set parameters in the displayed dialog box, and click **OK**.

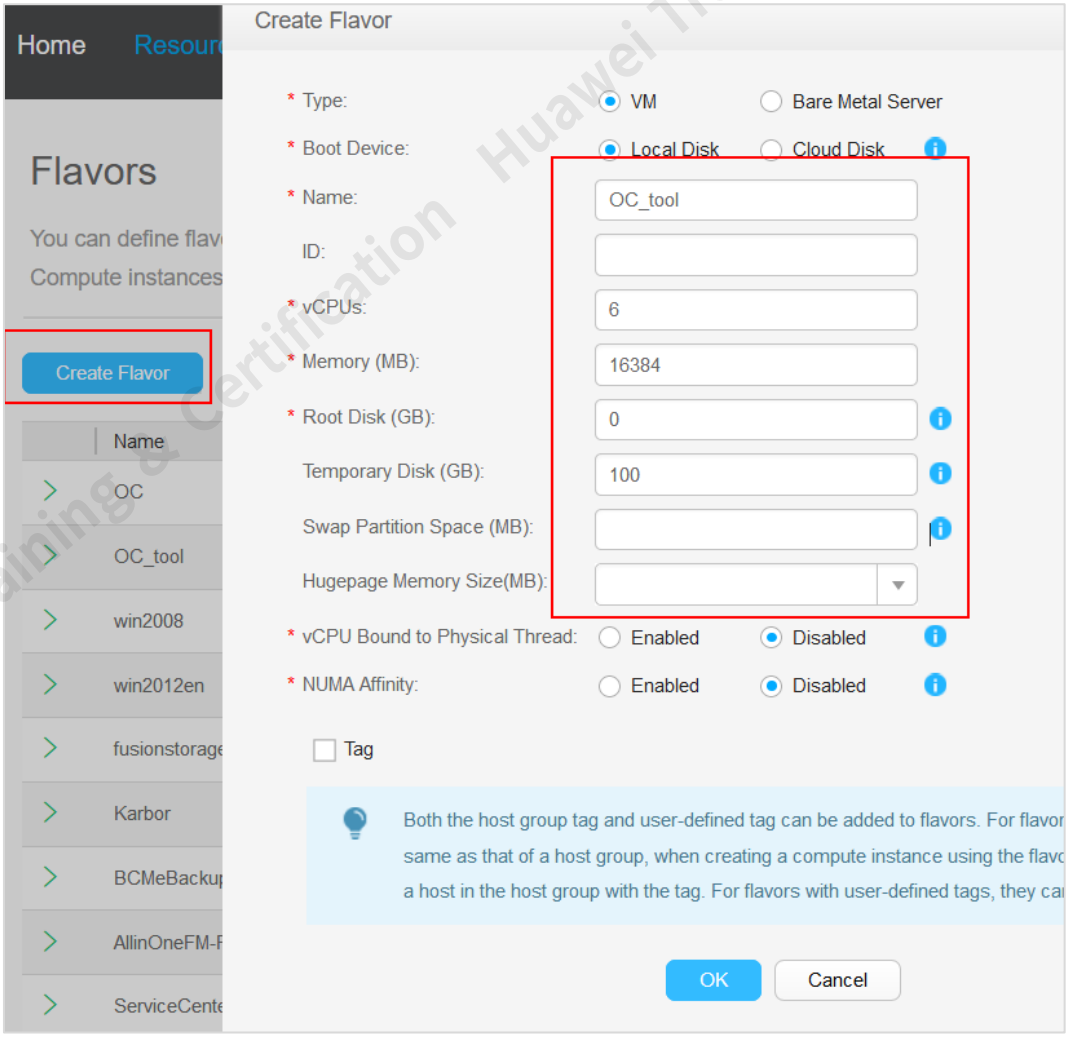

Step 3 **Create VMs.**

On the FusionSphere OpenStack OM page, choose **Resources** > **Compute Instances**.

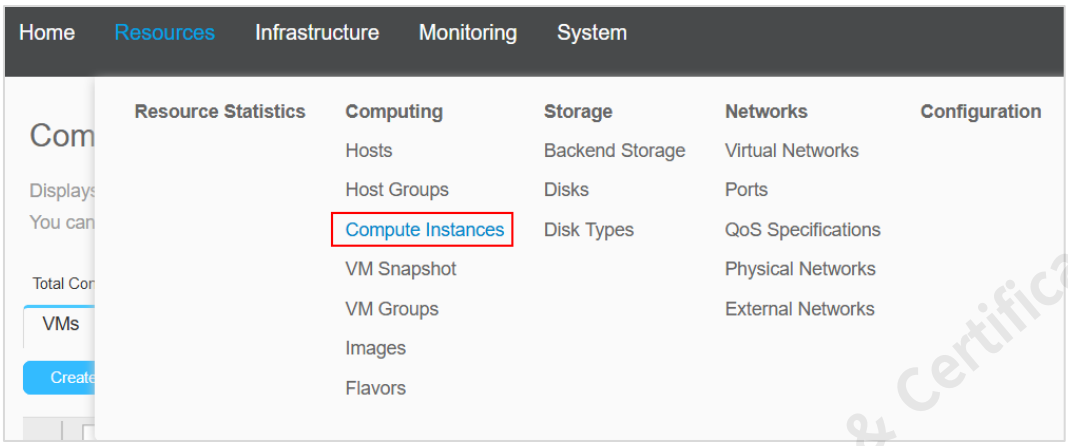

Click **Create** to start VM creation.

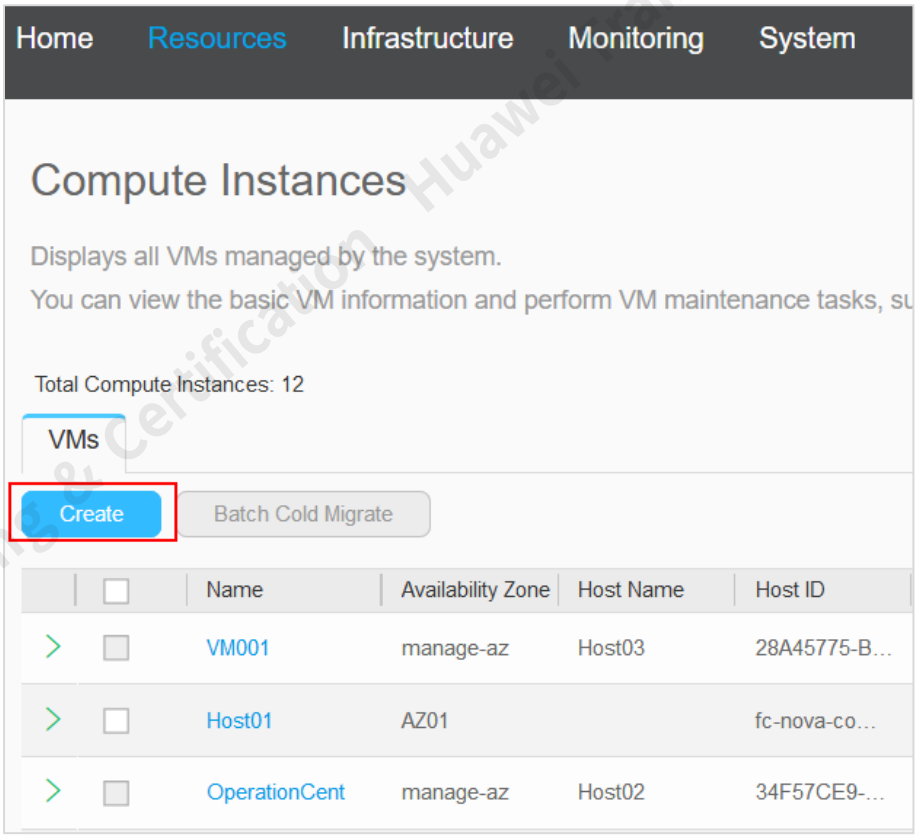

During VM creation, you need to select AZs, images, flavors, and networks. The operations are as follows:

Select an AZ.

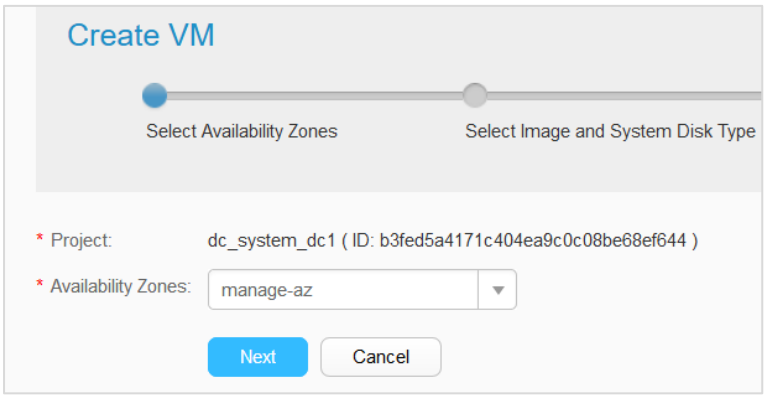

Select an image. (The image is registered in previous work.)

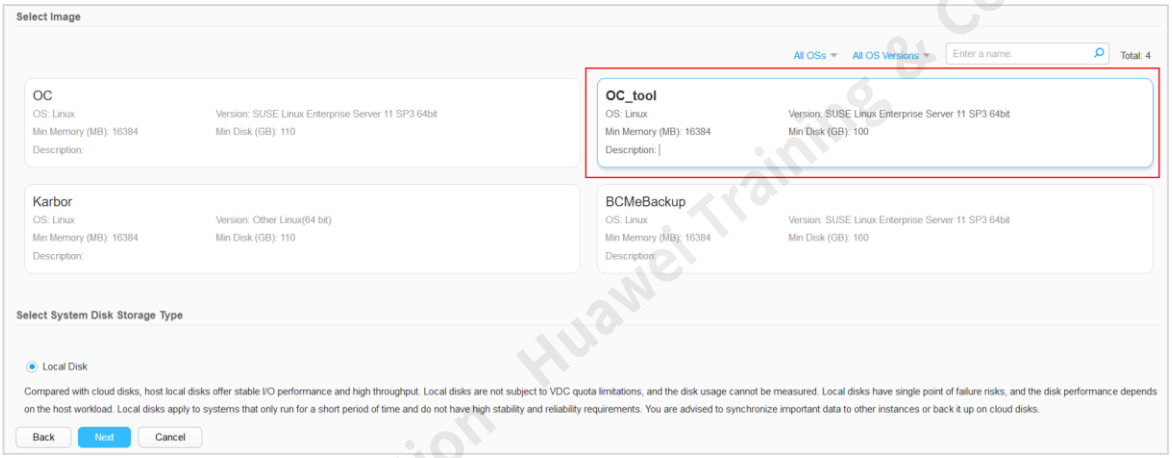

Select a flavor. (The flavor is created in previous work.)

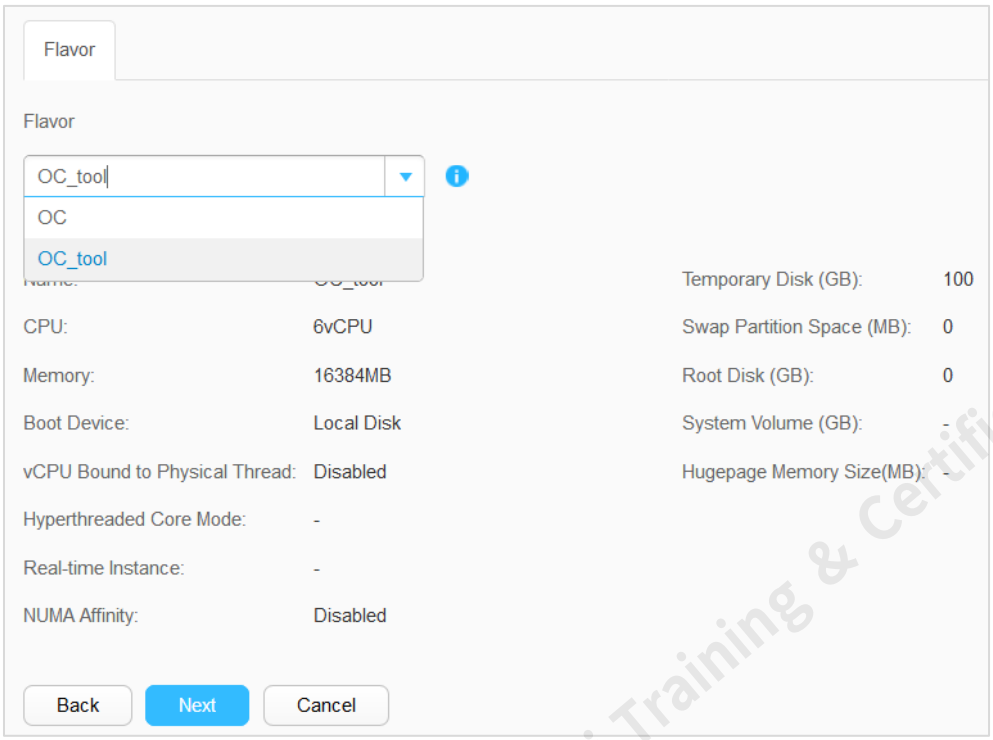

Select a network.

If the security group is enabled at the bottom layer of CPS, the floating IP address is mandatory. If the security group is not enabled, the floating IP address cannot be set. Otherwise, the installation will fail.

Set the VM name and select a host. (You can view the remaining resources of the host and select the host with sufficient resources to complete the installation. The flavor parameters are specific resource requirements.)

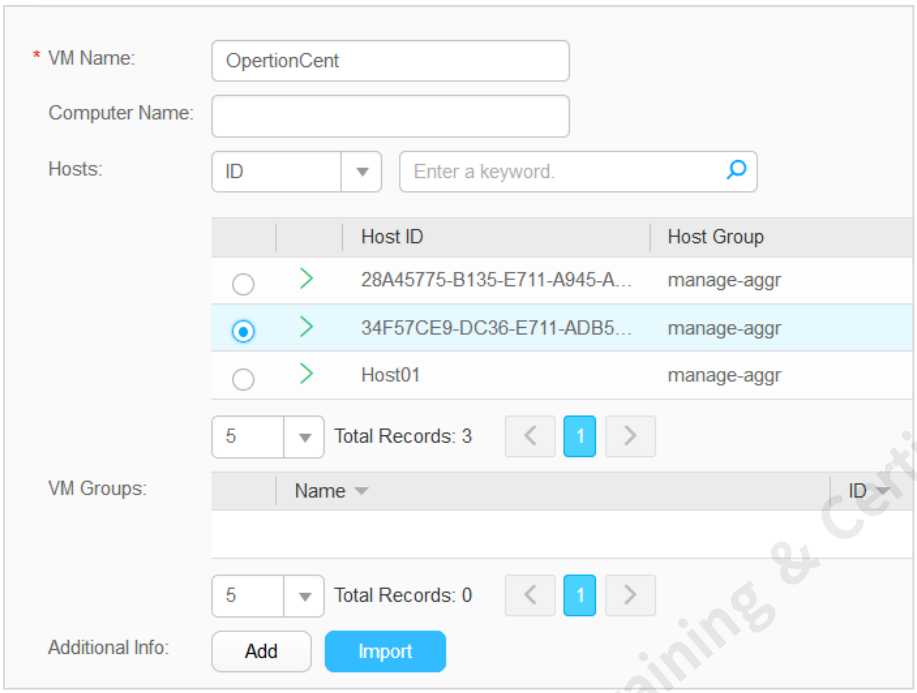

Import additional information about the VM.

In **Additional Information**, click **Import** to copy the following script to the dialog box:

```
{
  "mold":{
  "type":"enum",
  "value":["Dccare"]
  },
"language":{
"type":"enum",
"value":["zh","en"]
},
"mode":{
"type":"enum",
"value":["Single","Double"]
},
"scale":{
"type":"enum",
"value":["Mild","Moderate"]
},
"localIP":{
"type":"ip",
"value":""
},
"gateway":{
"type":"ip",
"value":""
},
"mask":{
"type":"ip",
"value":"255.255.255.0"
```

```
}
}
```

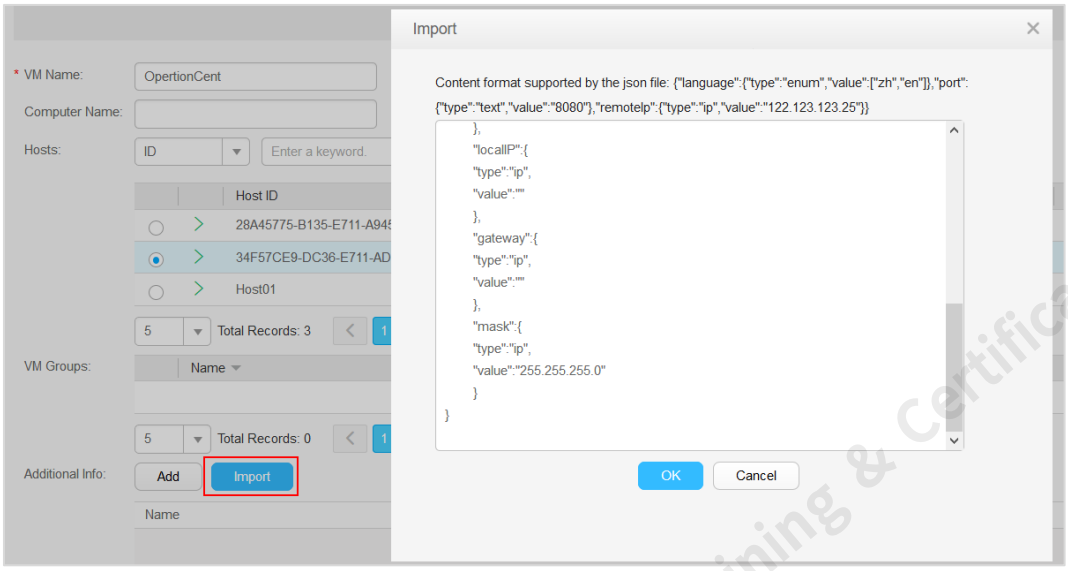

Modify the rest parameters according to the planning.

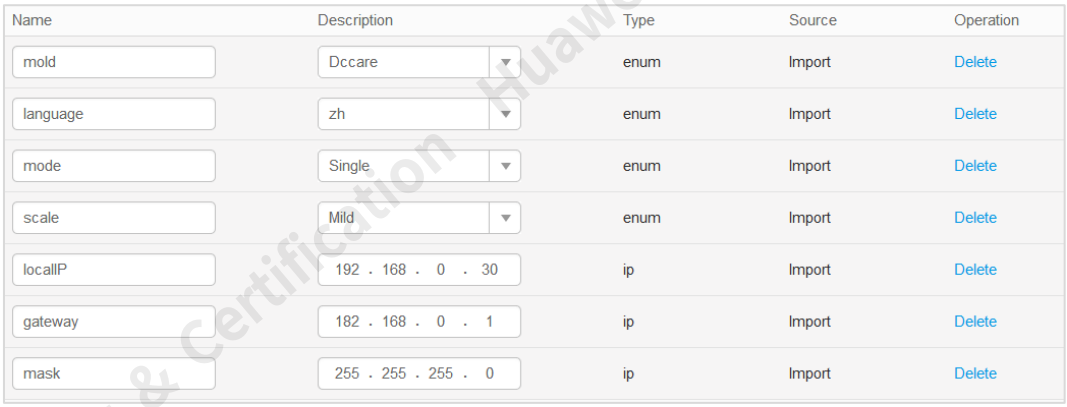

Parameter	How to Set
mold	Name of the analysis tool. Select **mold** to **Dccare**.
language	Installation language of the analysis tool: zh: Chinese en: English Set **language** to **zh**.
mode	Installation mode of the analysis tool: Single: standalone mode Double: active/standby mode

Parameter	How to Set
	Set **mode** to **Single**.
scale	Installation scenario of the analysis tool: Mild: small scale Moderate: standard deployment For the mapping between VM configurations and scales, see the software and hardware configuration requirement table. For example, set **scale** to **Mild**.
localIP	Enter the local IP address of the analysis tool, for example, **192.168.0.30**.
gateway	Enter the gateway address, for example, **192.168.0.1**.
mask	Enter the subnet mask. The default value is **255.255.255.0**.

Click **Next** to confirm the VM information.

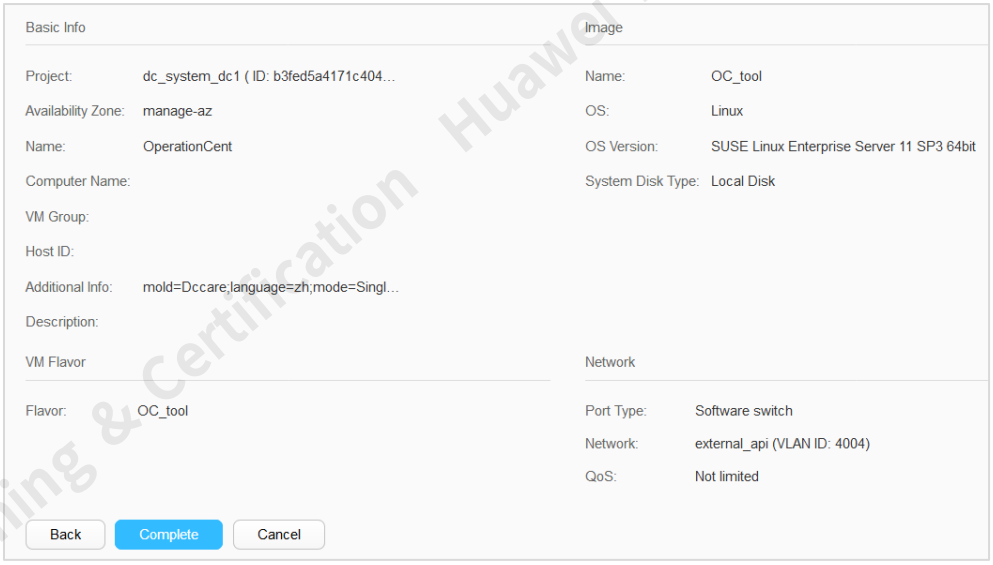

Click **Complete**. You can view the progress in the task center.

----End

9.2.6 Installing UVP VMTools

OperationCenter and analysis tools are deployed on FusionSphere OpenStack. Therefore, you need to install UVP VMTools on the VMs where OperationCenter and analysis tools reside. The installation process is the same. In this experiment, installing the analysis tool is used an example.

Prerequisites

- OperationCenter and analysis tools have been installed.
- The installation package of UVP VMTools has been downloaded.
- (FusionSphere OpenStack V100R006C10RC1_Tool.tar.gz)
- Decompress the preceding tool package to obtain the **uvp-vmtools-2.5.0-106010x86_64.rpm** file.

Step 1 **Copy the rpm file to a VM.**

Use the WinSCP tool to log in to the analysis tool VM. (User name: **appuser**, Password: **Changeme_123**, Host name: 1**92.168.0.30**)

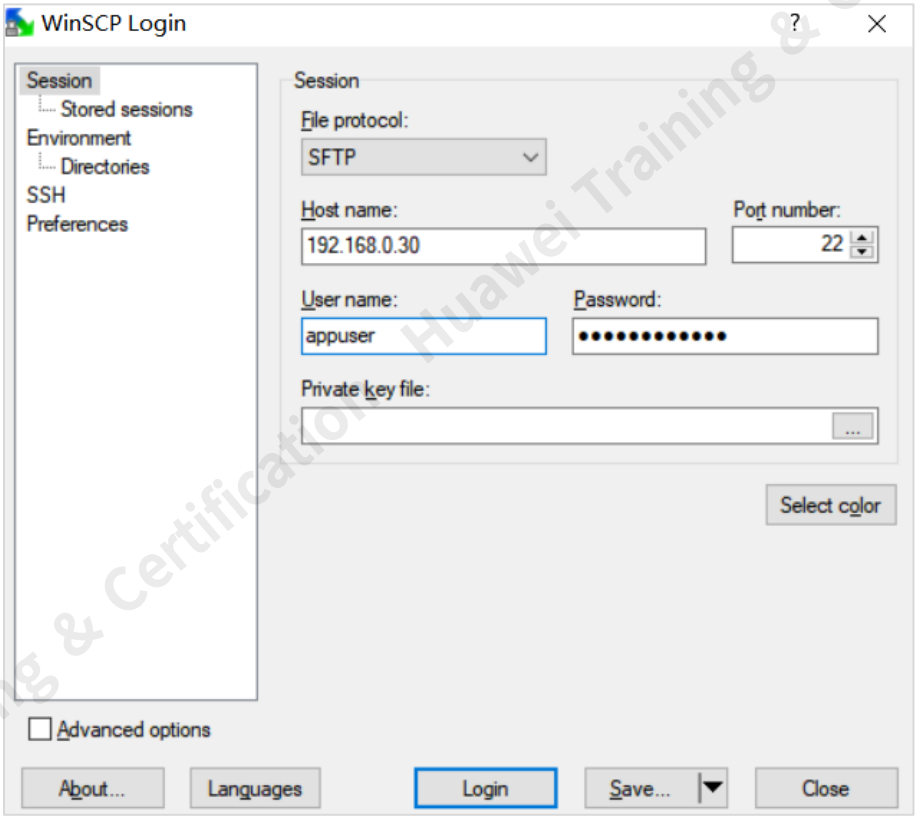

After logging in to the WinCSP, copy the tool to the VM.

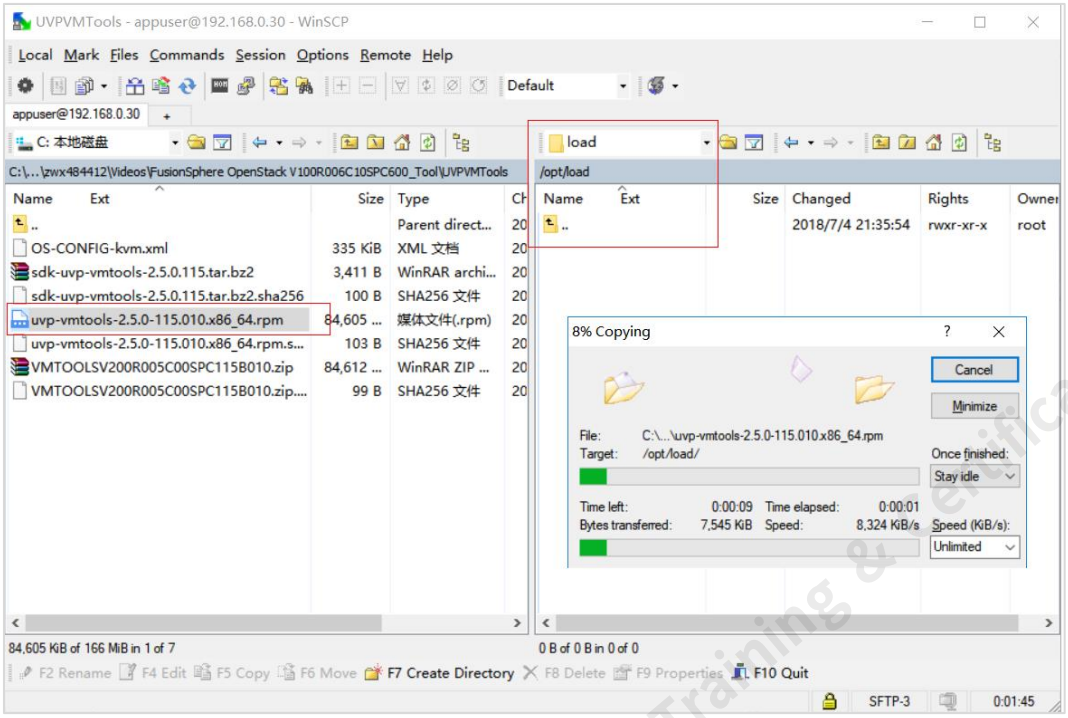

Step 2 **Install UVP VMTools.**

Use PuTTY to log in to the analysis tool. The user name is **appuser**. The password is **Changeme_123**.

Run the following command to switch to user **root**:

```
sudo su - root
```

Enter the password of user **root**. The default password of user **root** is **Changeme_123**.

Run the following command to create an empty **UVP_VMTools** directory under the **/opt/load** directory:

```
appuser@AnalyzerHost1:~> sudo su - root
root's password:
AnalyzerHost1:~ # cd /opt/load
AnalyzerHost1:/opt/load # mkdir UVP_VMTools
AnalyzerHost1:/opt/load #
```

Copy the installation package to **UVP_VMTools**.

```
AnalyzerHost1:/opt # cd /opt/load/
AnalyzerHost1:/opt/load # ls
UVP_VMTools  uvp-vmtools-2.5.0-115.010.x86_64.rpm
AnalyzerHost1:/opt/load # mv uvp-vmtools-2.5.0-115.010.x86_64.rpm
/opt/load/UVP_VMTools/
```

Decompress files.

```
AnalyzerHost1:/opt/load # cd UVP_VMTools/
AnalyzerHost1:/opt/load/UVP_VMTools # ls
uvp-vmtools-2.5.0-115.010.x86_64.rpm
AnalyzerHost1:/opt/load/UVP_VMTools # rpm2cpio uvp-vmtools-2.5.0-
115.010.x86_64.rpm | cp
cp            cpan2dist      cpanp-run-perl  cpufreq-info
cpan          cpanp          cpio            cpufreq-set
AnalyzerHost1:/opt/load/UVP_VMTools # rpm2cpio uvp-vmtools-2.5.0-
115.010.x86_64.rpm | cpio -imd
350461 blocks
AnalyzerHost1:/opt/load/UVP_VMTools #
```

After the decompression, two files are generated in the **/opt/load/UVP_VMTools/opt/patch/programfiles/vmtools** directory.

- vmtools-linux.iso: an image file used to install VMTools on a Linux VM.
- vmtools-windows.iso: an image file used to install VMTools on a Windows VM.

The analysis tool is deployed on a Linux VM. Therefore, you only need to install the **vmtools-linux.iso** file.

Create an empty **uvp_mount** folder in the **/opt/load** directory.

Mount vmtools-linux.iso to the uvp_mount directory.

```
AnalyzerHost1:/opt/load # mkdir uvp_mount
AnalyzerHost1:/opt/load/UVP_VMTools # cd opt/patch/programfiles/vmtools/
AnalyzerHost1:/opt/load/UVP_VMTools/opt/patch/programfiles/vmtools # mount
vmtools-linux.iso -o loop /opt/load/uvp_mount
mount: block device
/opt/load/UVP_VMTools/opt/patch/programfiles/vmtools/vmtools-linux.iso is
write-protected, mounting read-only
```

Copy the tool file to the **/opt/load/UVP_VMTools** directory and enter the directory where the tool file is decompressed.

```
AnalyzerHost1:/opt/load/UVP_VMTools/opt/patch/programfiles/vmtools # cd
/opt/load/uvp_mount/
AnalyzerHost1:/opt/load/uvp_mount # cp vmtools-2.5.0.115.tar.bz2
/opt/load/UVP_VMTools/
AnalyzerHost1:/opt/load/uvp_mount # cd /opt/load/UVP_VMTools/
AnalyzerHost1:/opt/load/UVP_VMTools # tar -jxvf vmtools-2.5.0.115.tar.bz2
```

After the decompression, go to the **vmtools** folder, start installing the tool, and view the returned information.

```
AnalyzerHost1:/opt/load/UVP_VMTools # cd vmtools/
AnalyzerHost1:/opt/load/UVP_VMTools/vmtools # sh install
```

If the returned information indicates that the installation is successful, run the **reboot** command so that the installation takes effect upon VM restart.

```
AnalyzerHost1:/opt/load/UVP_VMTools/vmtools # sh install
Start Installation :
  Install kernel modules.
  Install UVP VMTools agent service.
  Change system configurations.
Update kernel initrd image.
The UVP VMTools is installed successfully.
Reboot the system for the installation to take effect.
AnalyzerHost1:/opt/load/UVP_VMTools/vmtools # reboot

Broadcast message from root (pts/0) (Thu Jul  5 11:08:39 2018):

The system is going down for reboot NOW!
AnalyzerHost1:/opt/load/UVP_VMTools/vmtools #
```

----End

9.2.7 Backing Up and Recovering OperationCenter

Prerequisites:

- OperationCenter has been installed.
- The username and password for logging in to the OperationCenter maintenance tool portal have been obtained. (The username and default password are **sys** and **Changeme_123**, respectively.)

Step 1 **Go to the OperationCenter maintenance tool portal.**

Enter the portal address in the address box of the browser: **http://192.168.0.180:8088**, enter the maintenance tool, and enter the username and password.

Username: **sys**

Password: Changeme_123

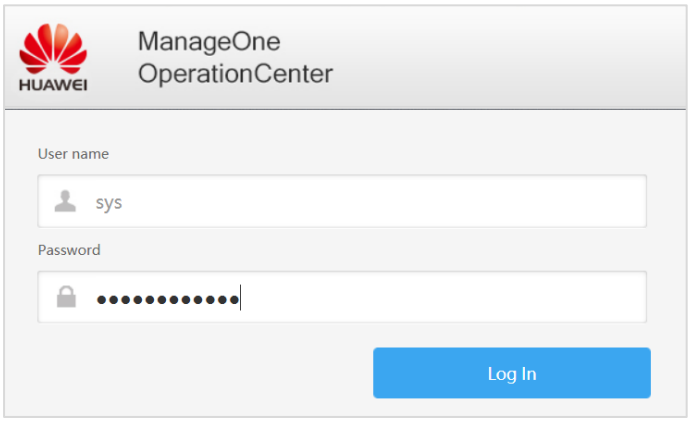

Change the password at initial login. Change the password to **Huawei@123**.

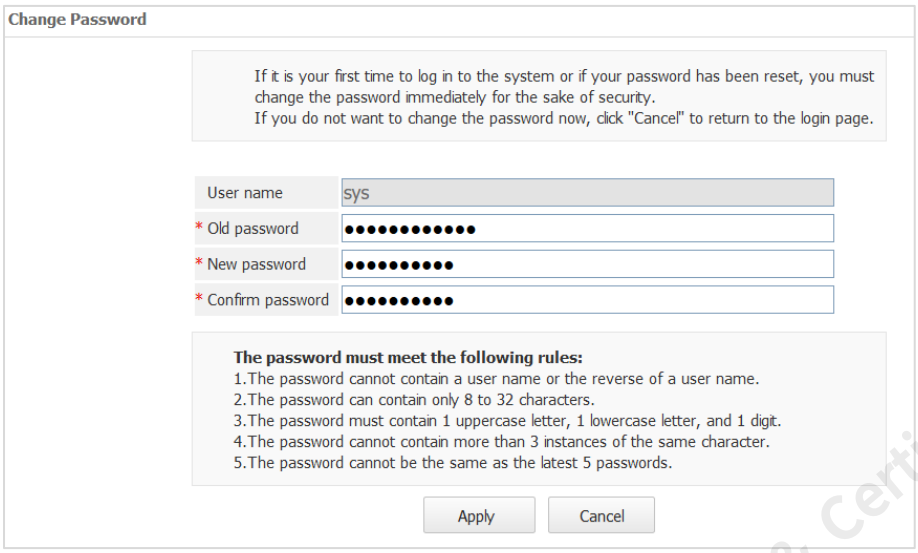

Click **Apply**.

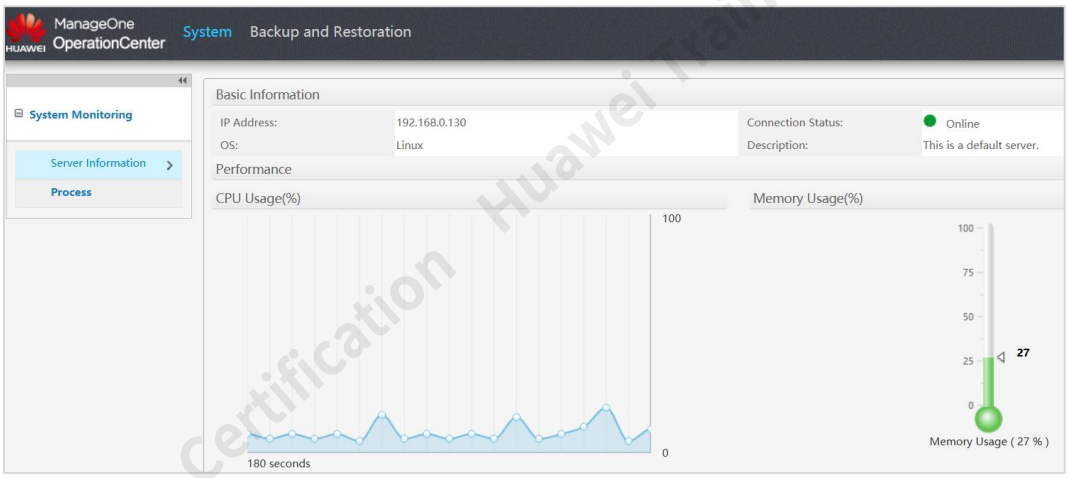

Step 2 **Manually back up OperationCenter.**

On the tool menu, choose **Backup and Restoration** > **Backup Policy Setting** to set backup policies.

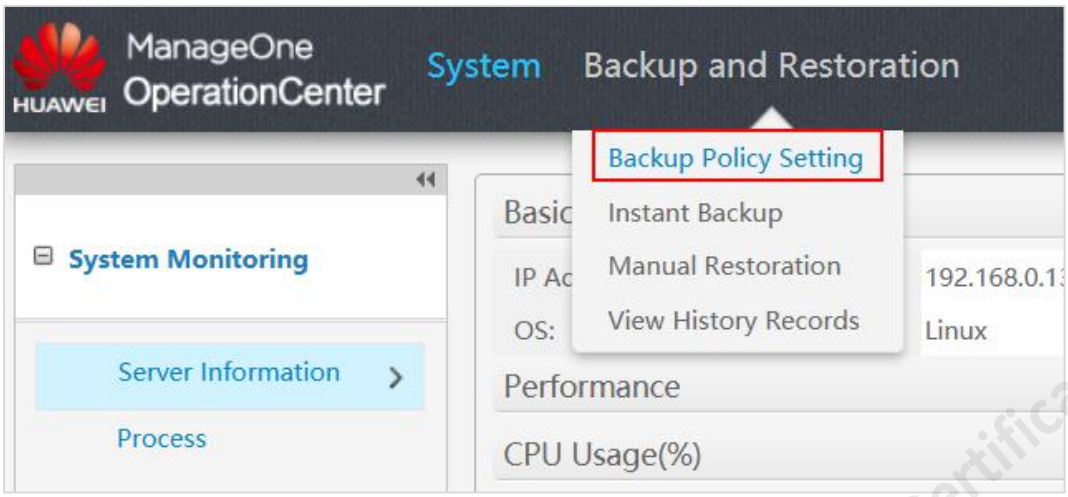

Under **Backup Path**, select **Backup to the local server**. Click **Apply** to save the setting.

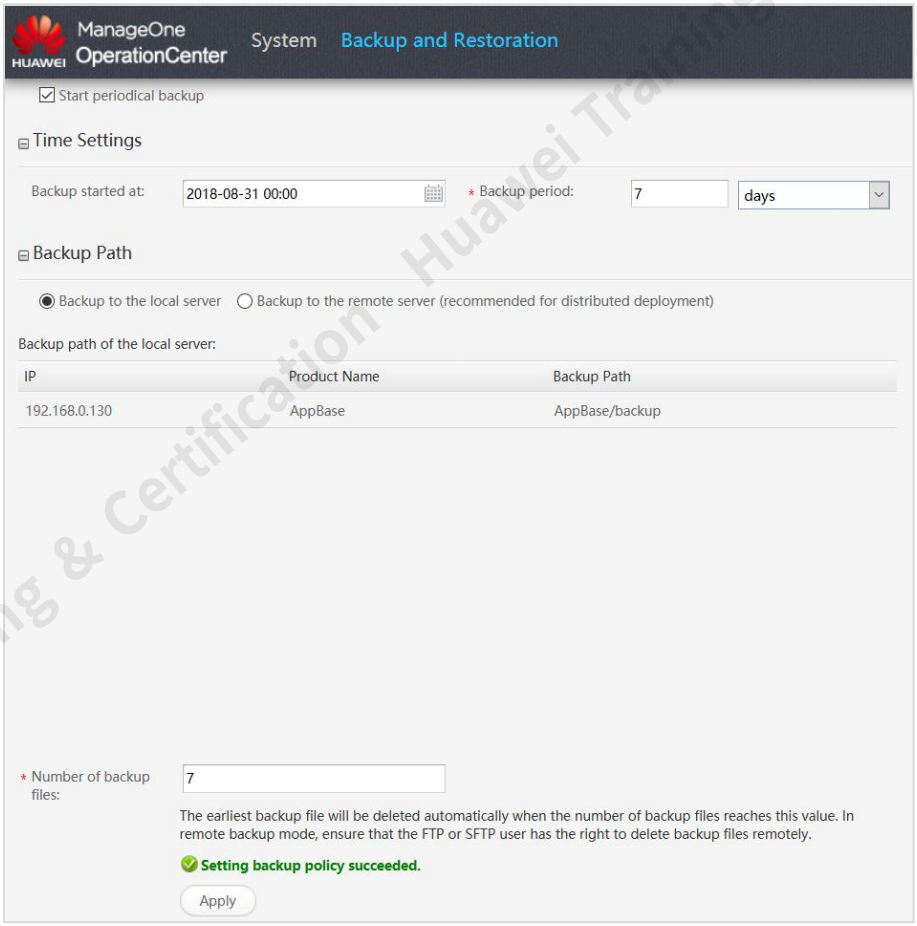

Click Immediate Backup to manually back up OperationCenter.

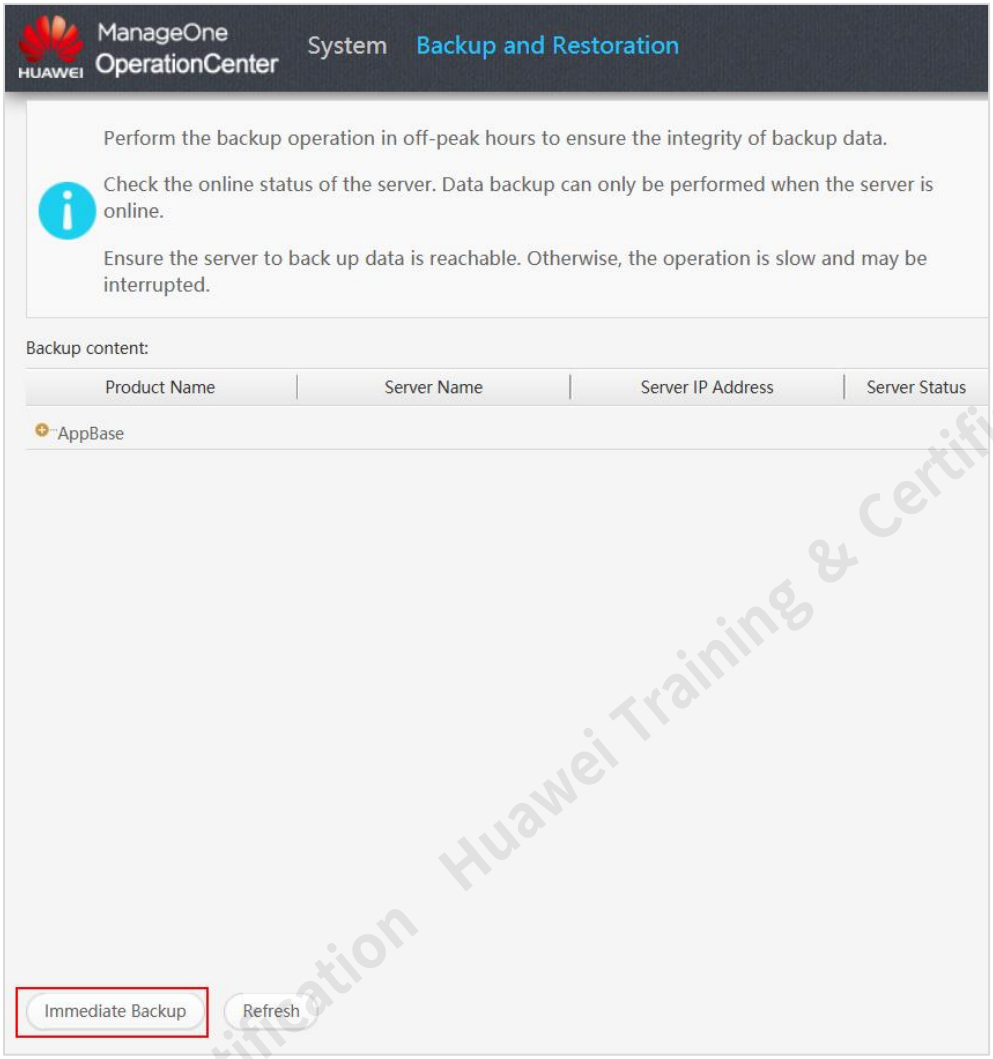

After the backup is complete, click **Complete**.

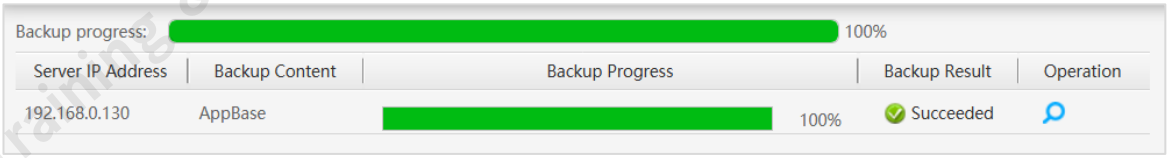

Step 3 **Restore OperationCenter.**

Use PuTTY to log in to the OperationCenter server, switch to user **root**, and stop all services.

```
appuser@OSHost:~> sudo su - root
root's password:
OSHost:~ # TMOUT=0
OSHost:~ # service oc stopforM
Stop for maintaining...
Stopped for maintaining Successfully.
OSHost:~ #
```

Enter the maintenance tool portal, choose **Backup and Restoration** > **Manual Restoration**.

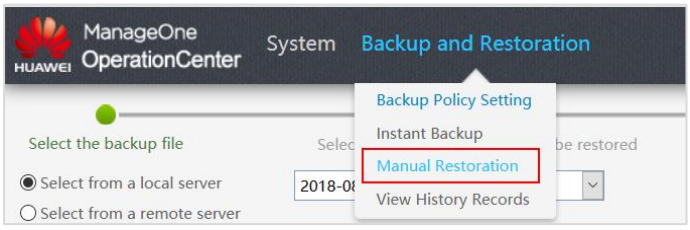

In the following displayed page, click **Next**.

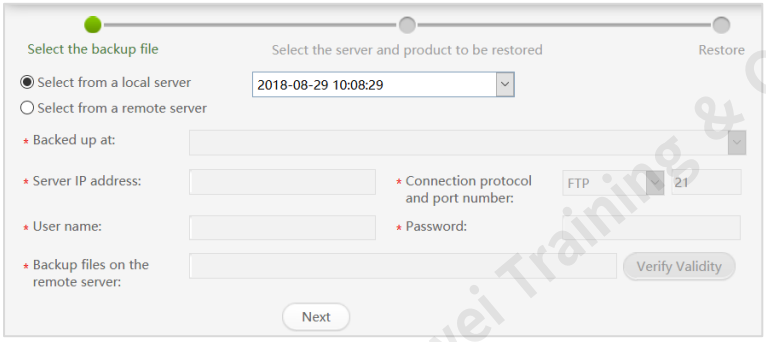

On the **Select the server and product to be restore** page, select **By server**. Then, select the data source and server to be restored. Click **Next**.

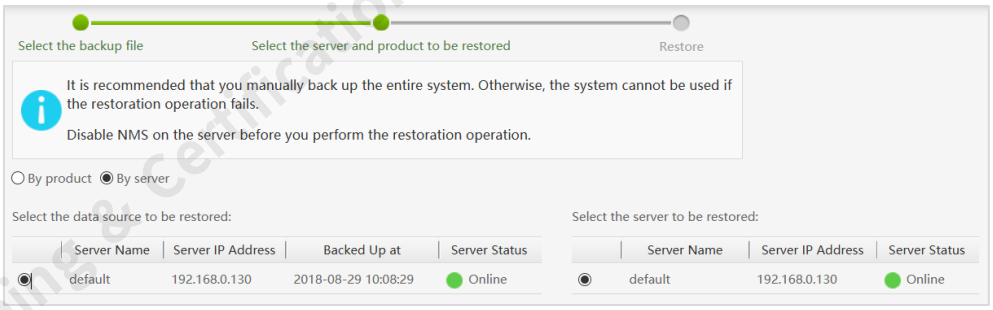

Click **Restore**.

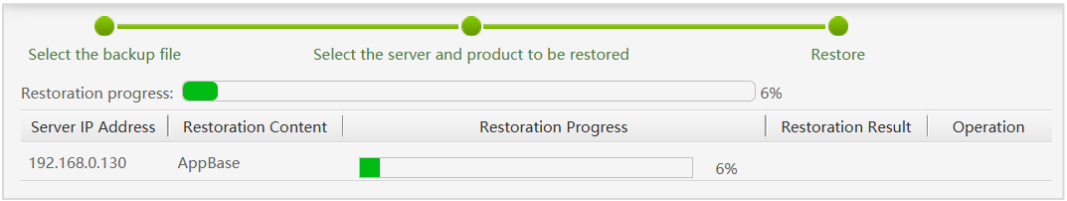

After the task is complete, restart the OperationCenter service. The restoration is complete.

```
appuser@OSHost:~> sudo su - root
```

```
root's password:
OSHost:~ # service oc start
Starting...
The OperationCenter is starting, which takes about 15 minutes. Please wait.
Started Successfully.
```

----End

9.2.8 Adding Analysis Tools

Step 1 **Create an API account on OperationCenter.**

On the OperationCenter page, choose **System** > **User Management** > **User**. Click **Open API login** and set parameters.

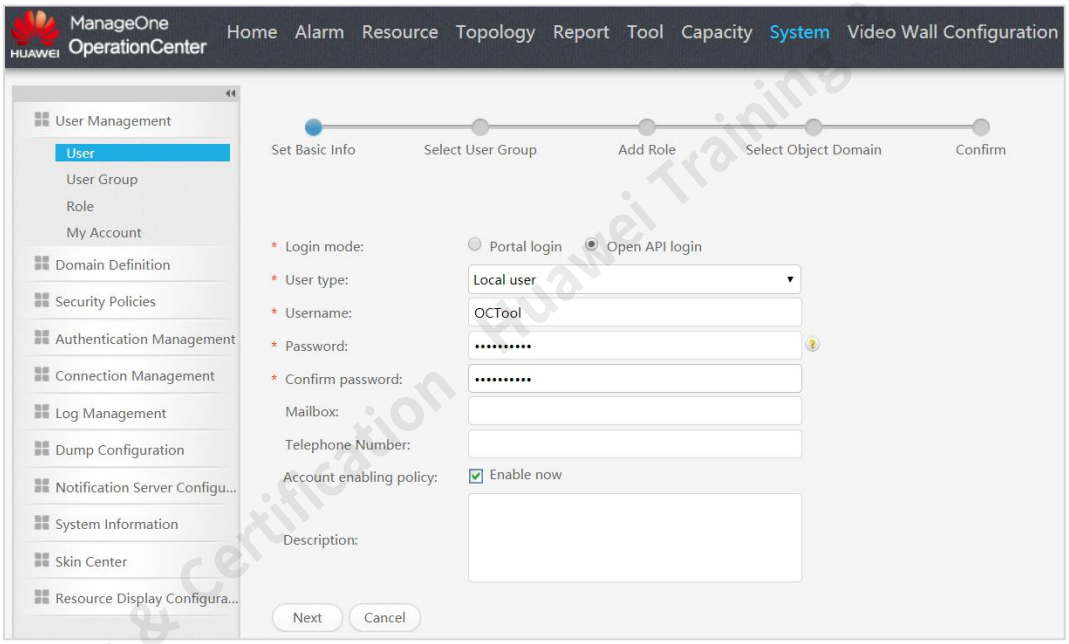

Click Next. On the Select User Group page, select System management group. Click Next.

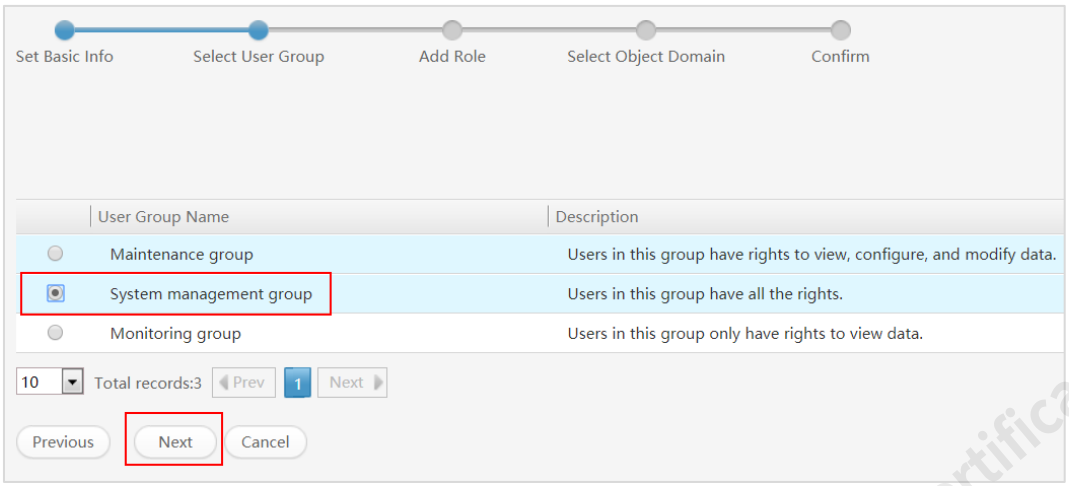

Click **Confirm**.

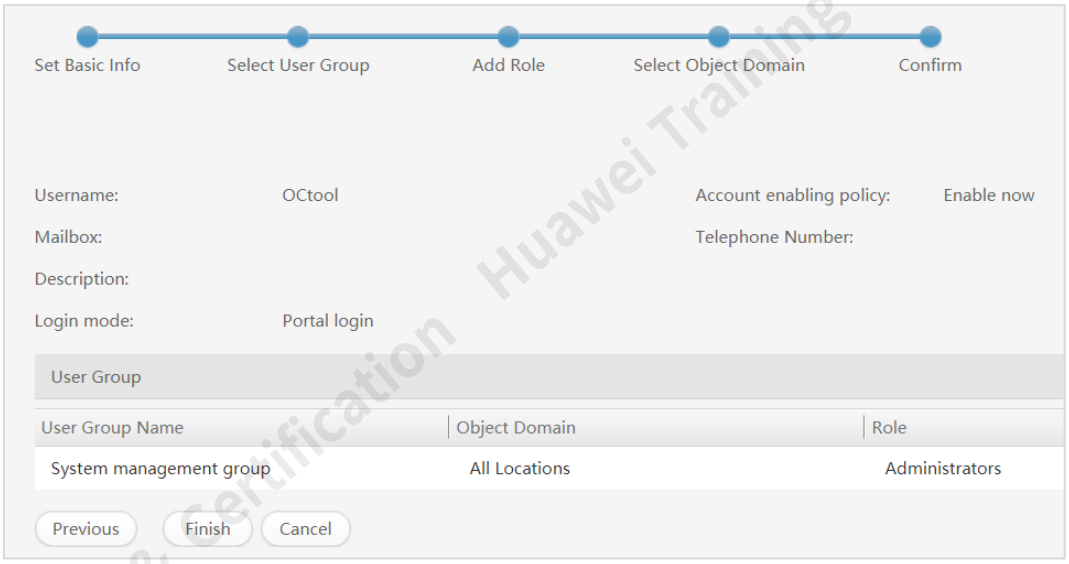

Step 2 **Add analysis tools.**

On the OperationCenter homepage, choose **Tool** > **Tool Management** > **Register**. In the displayed **Register Tool** dialog box, enter information about the analysis tool. Enter the IP address and click to **Obtain Tool ID**.

Set both **Host IP address** and **Floating IP address** to the management IP address because the tool is deployed in standalone mode.

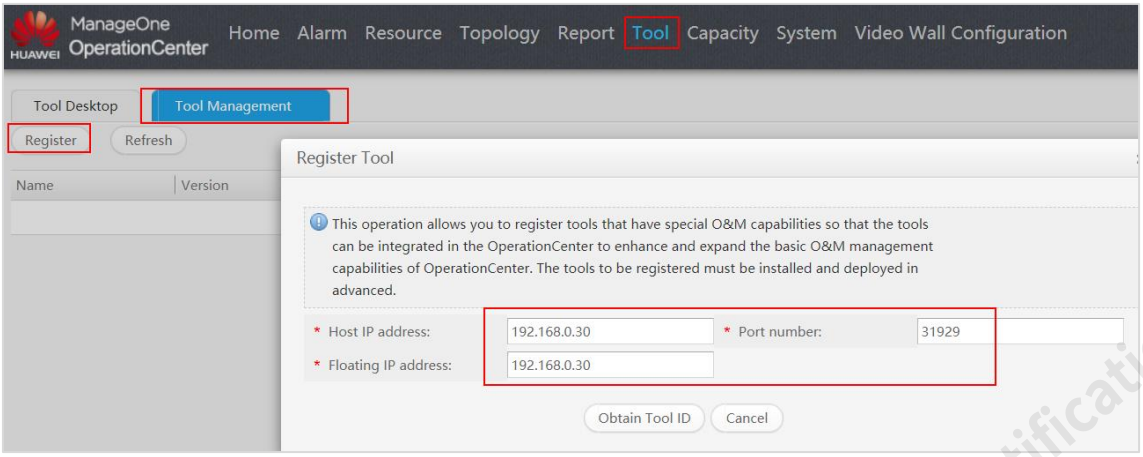

Import the tool certificate. (The username is **appuser** and the initial password is **Changeme_123**.) Then, load the certificate. In the displayed **Register Tool** dialog box, set **Tool ID** to **dccare**.

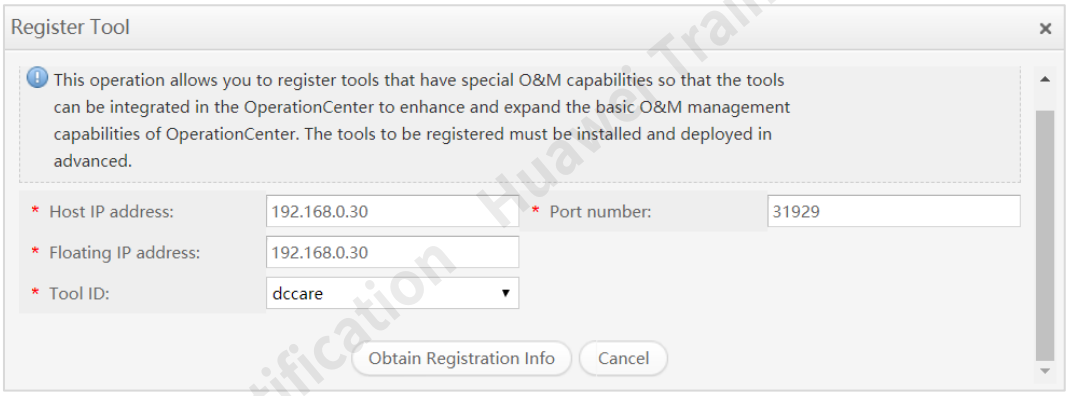

Click **Obtain Tool ID** again. In the **Registration Package Basic Information** dialog box, set the location information, username, and password. The username and password in this step are the same as the API user account created in the previous step.

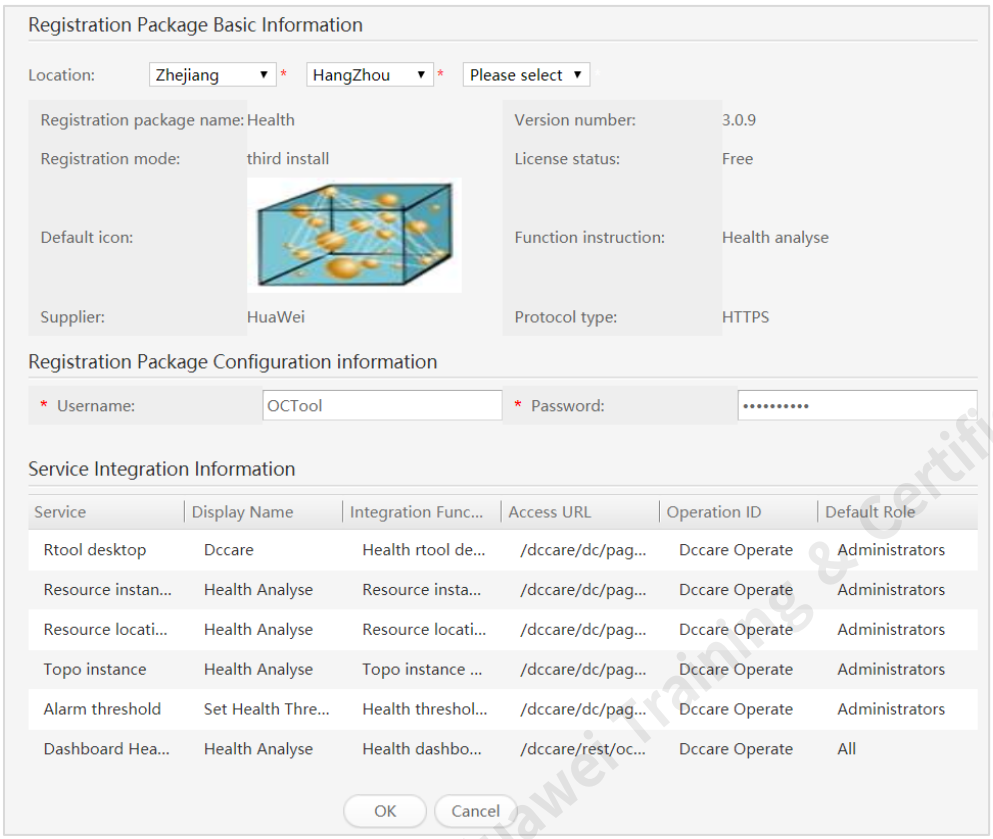

Click **OK**.

Step 3 **Set the threshold.**

On the OperationCenter homepage, choose **Alarm** > **Set Threshold**.

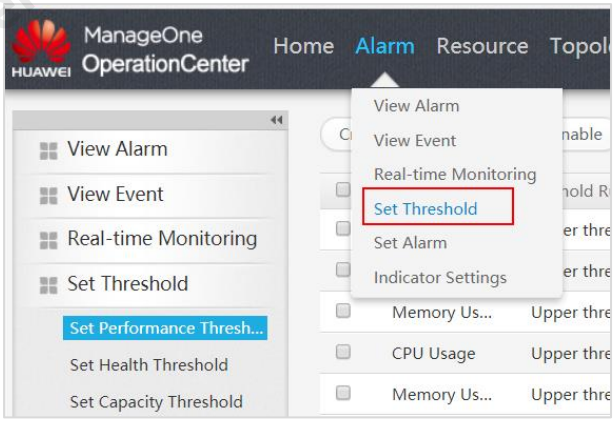

In the displayed page, choose **Set Performance Threshold**. Click **Add Indicator** and select the threshold to be monitored.

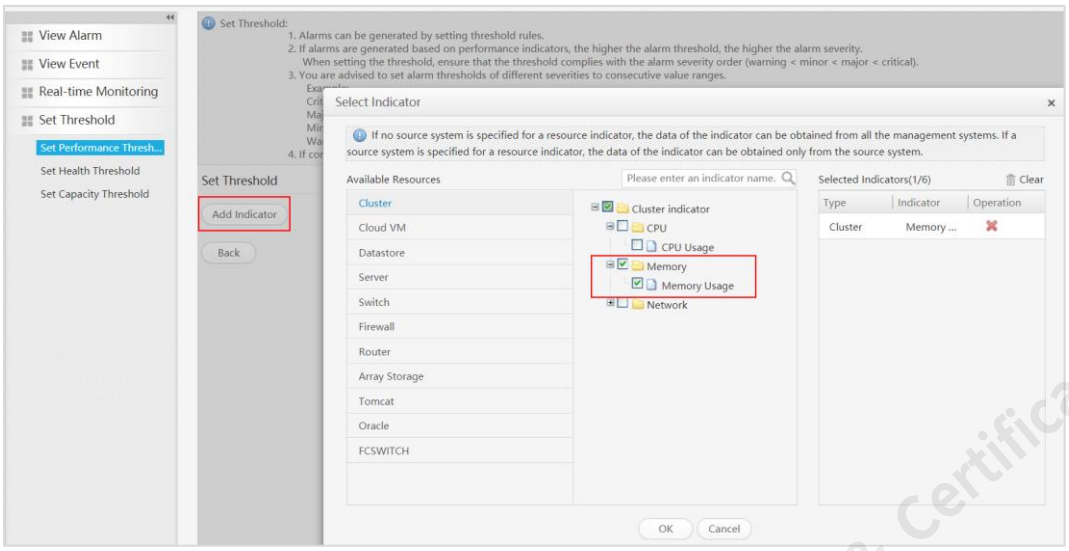

Click **OK**. In the displayed **Set Lower Limit** page, set threshold parameters.

Click OK. If any parameters are incorrectly set, an error message will be prompted. Reconfigure parameters as required.

Enter the analysis tool.

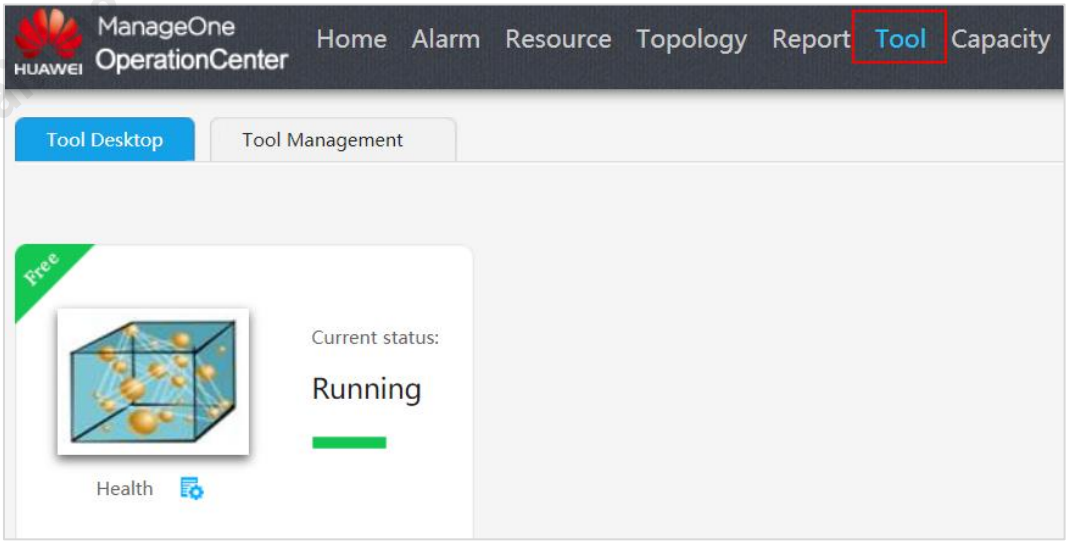

----End

9.3 Agile Controller-DCN Network Service O&M and Monitoring

Agile Controller-DCN supports network service O&M and monitoring after service provisioning.

9.3.1 Three-level Topology Visibility

In the cloud-network synergy scenario, multiple roles such as the tenant administrator, system administrator, and network administrator exist, which correspond to different network models. Agile Controller-DCN provides three-level network topology visibility, that is, visibility of the application network, logical network, and physical network.

Step 1 **View the application network topology.**

Choose Network > TOPO > Application Network.

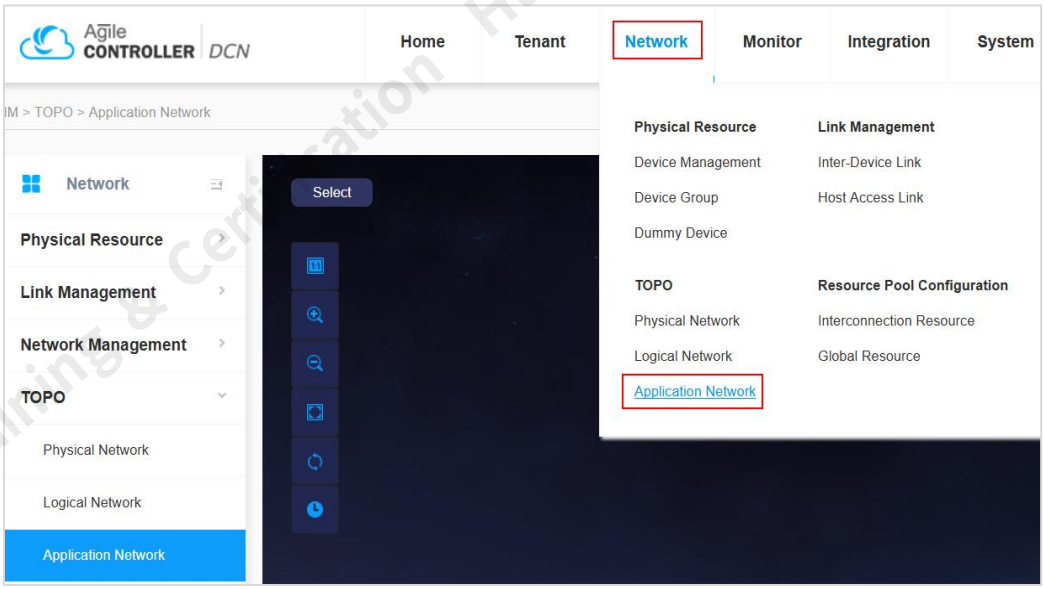

Select the desired tenant.

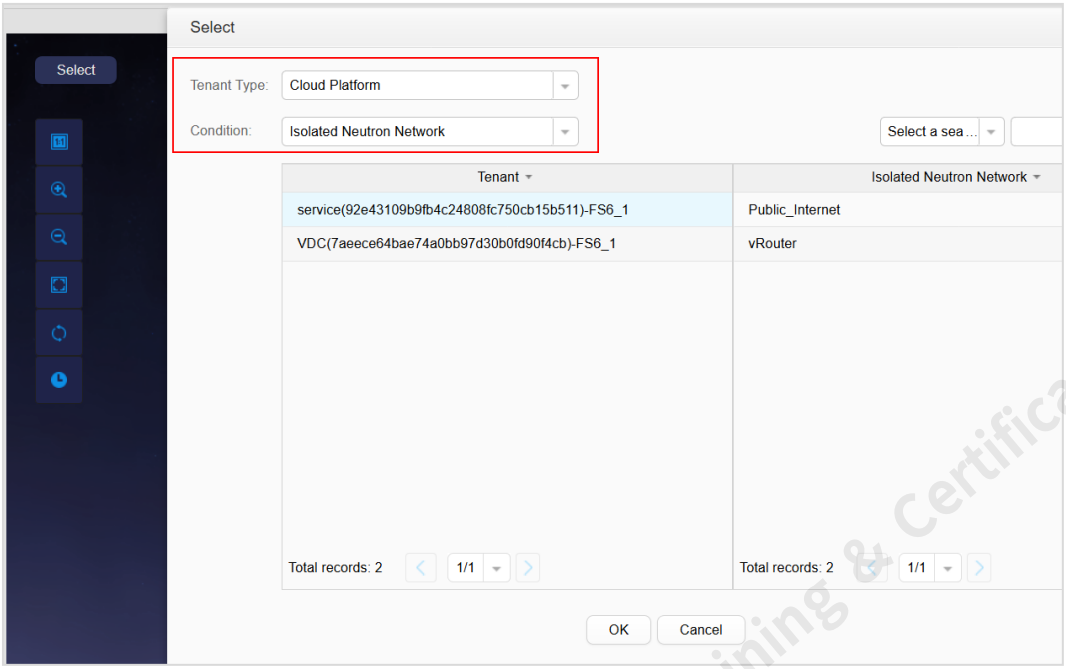

You can view the mapping between the application network and logical network.

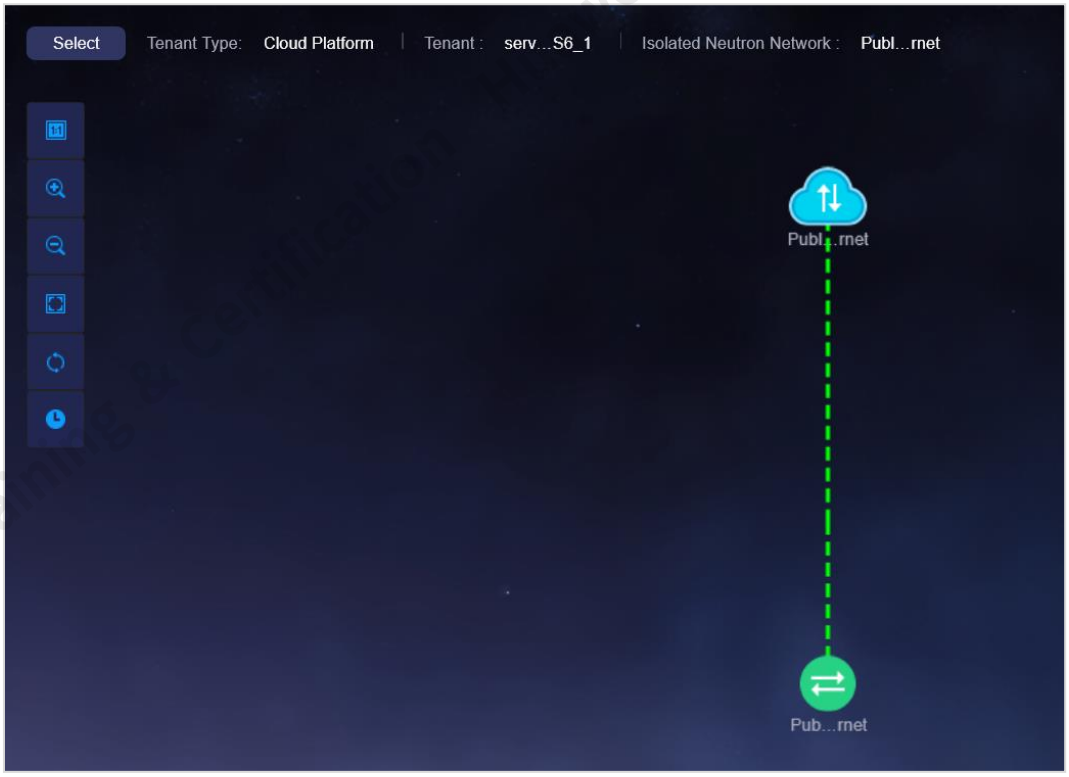

Step 2 **View the logical network topology.**

Click a logical network or click **View More** in the application network topology.

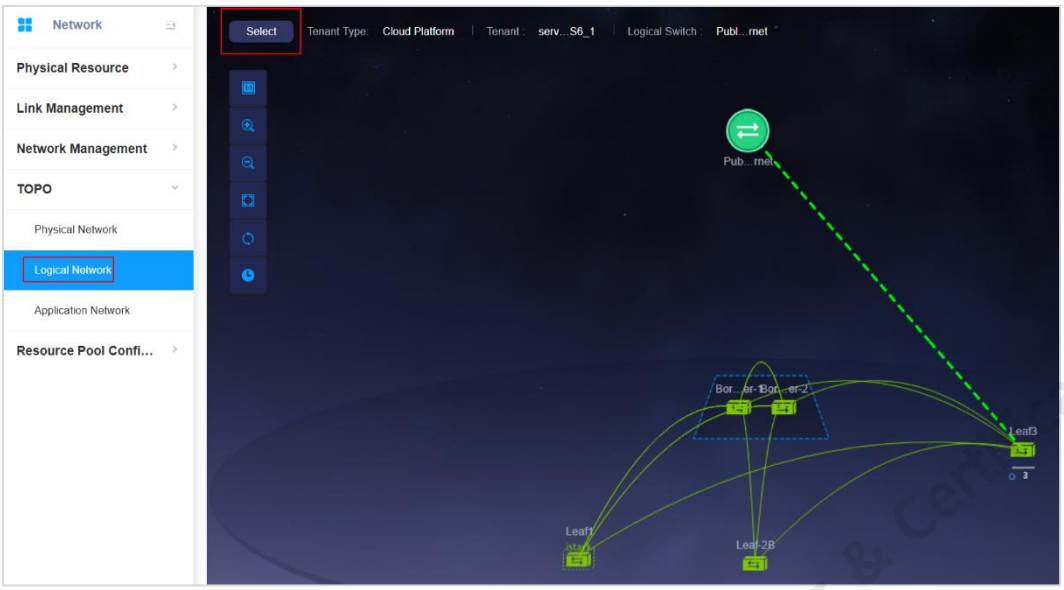

Right-click to view the mapping, port, and subnet information.

The external network plane Public_Internet of FusionSphere corresponds to the logical port created on Leaf-3 of the physical network.

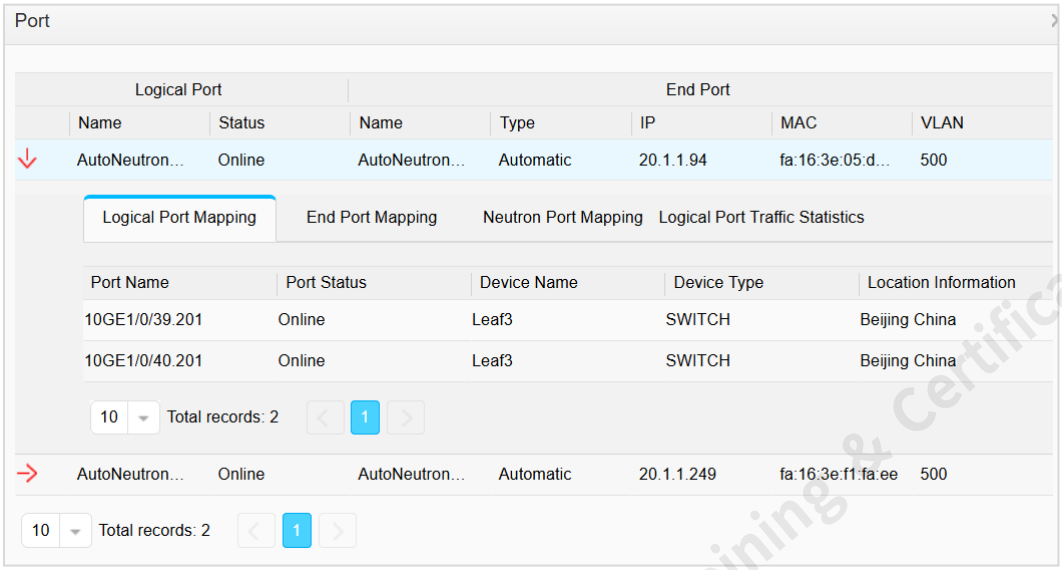

Step 3 **View the physical network topology.**

Choose Network > TOPO > Physical Network.

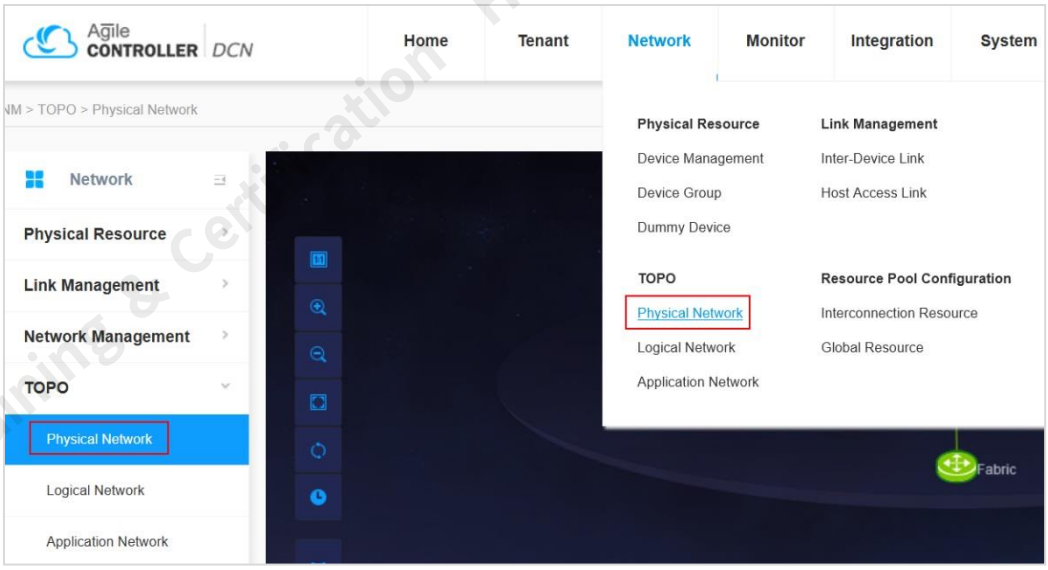

View the physical network topology of the current Fabric.

The digit under the device icon indicates the logical port. You can right-click to view the detailed port information.

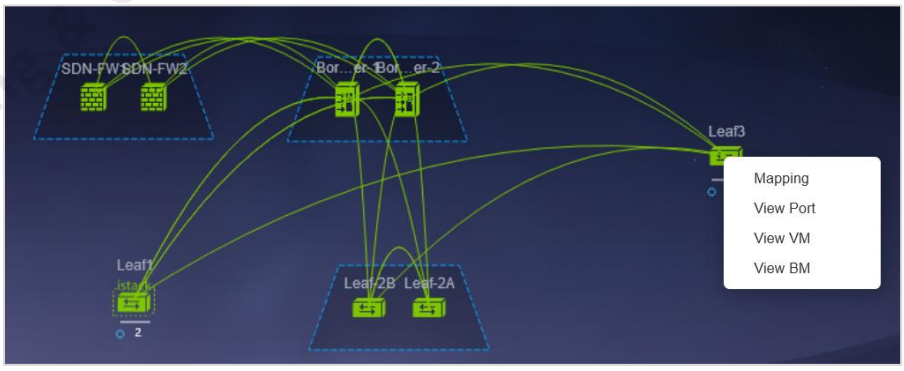

Place the cursor on the device icon to view the basic information.

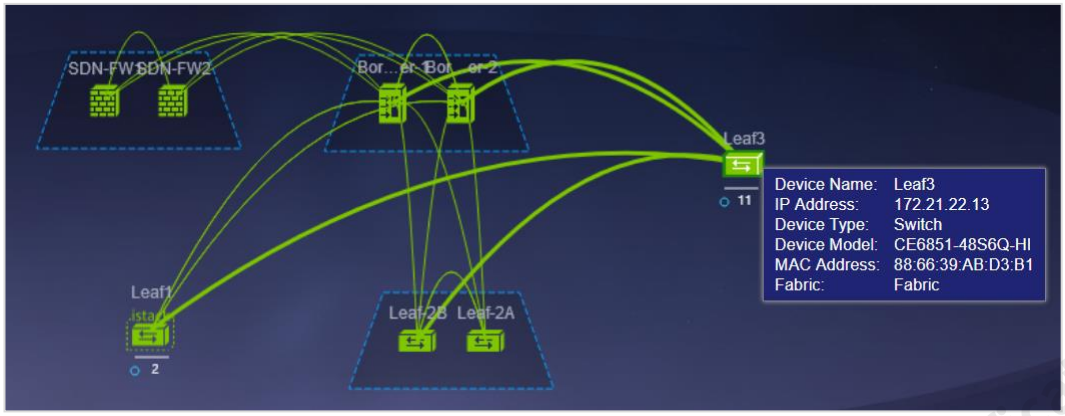

View the Leaf-1 mapping. It is found that the Leaf-1 is a logical switch, which corresponds to two logical networks route01 and route0102. A tenant VDC that is associated with Leaf-1 is generated.

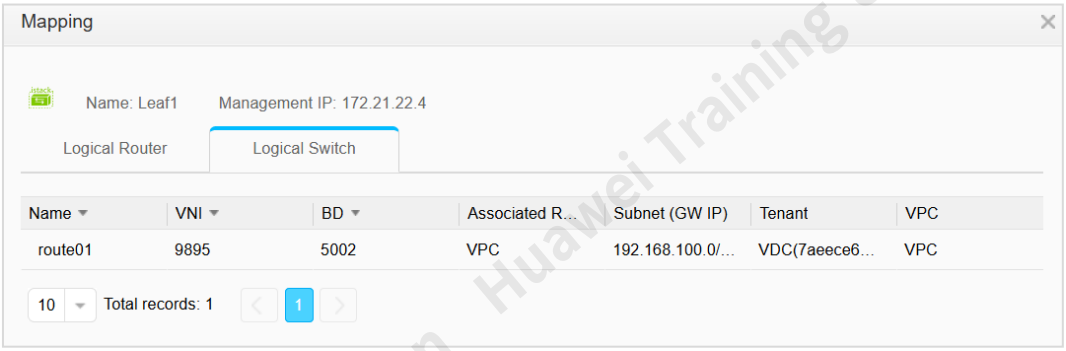

View the VMs that are connected to Leaf-1.

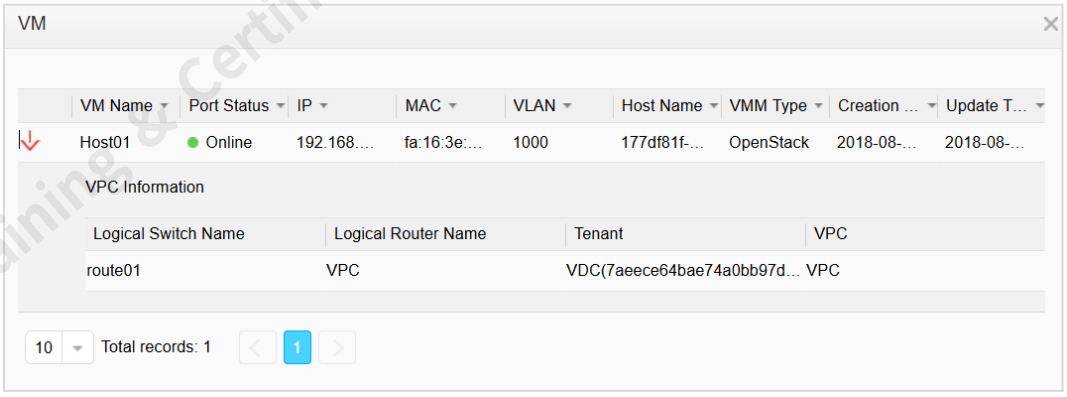

Agile Controller-DCN can identify ECSs connected to different Leaf nodes and present the corresponding VLAN, tenant, and logical network information.

----End

9.3.2 Network Path Detection

This function detects the actual physical paths between VMs or devices and checks whether the service flows are interrupted. Ping is used to check the end-to-end reachability between VMs. Network path detection is used to test the physical connections between NVEs.

Two network path detection modes are supported:

- Single path detection: checks the actual physical path of a service flow between VMs and checks whether the service flow is interrupted.
- Multi path detection: checks multiple actual physical paths between NVE devices and whether service flows are interrupted.

9.3.2.1 Single Path Detection

Step 1 **Enable path detection.**

On the Agile Controller-DCN main menu, choose System > System Settings > Southbound Protocol. Choose OpenFlow > OpenFlow Configuration and set Enabling TLS to OFF.

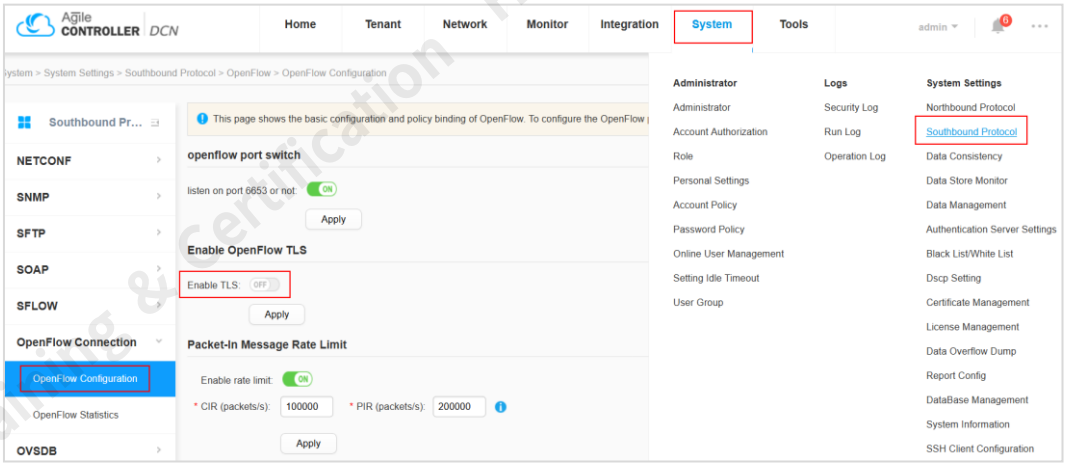

Choose Monitor > Path Detection Settings and set Path Detection Settings to ON.

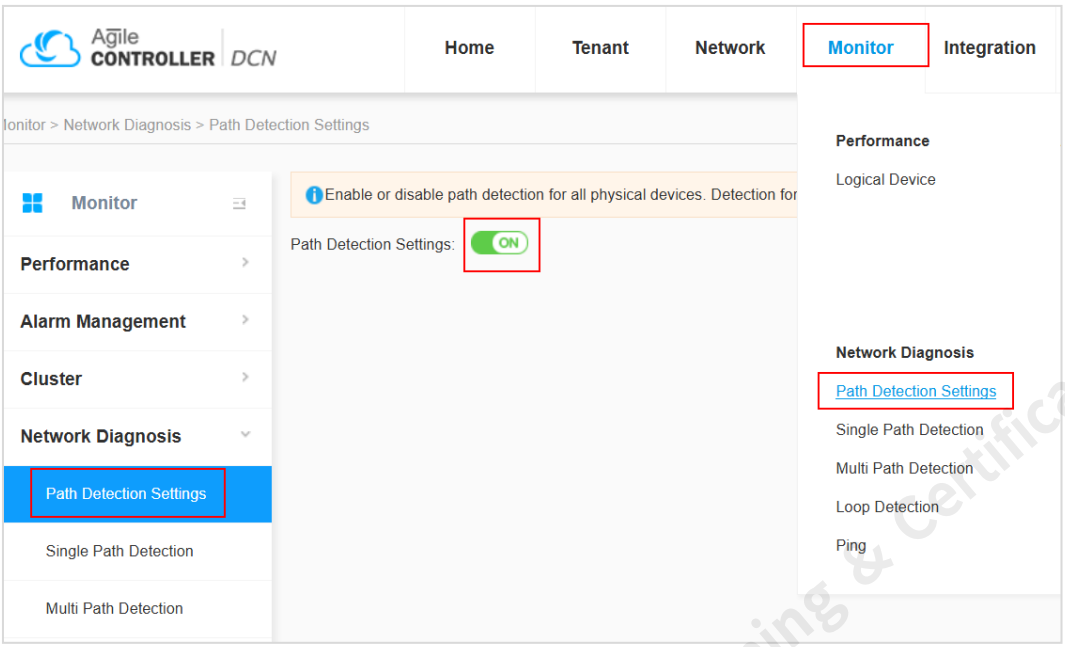

Select **Single Path Detection** and enter the source and destination IP addresses of the VM.

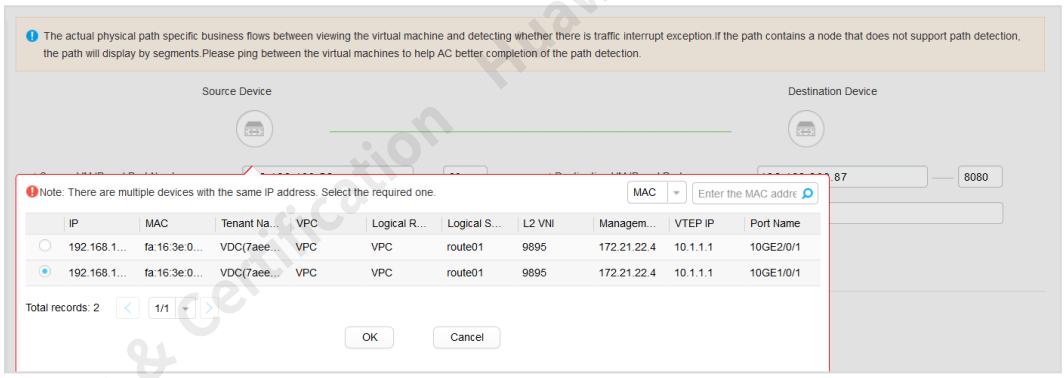

Set Protocol Type and Detection Time, and click Start.

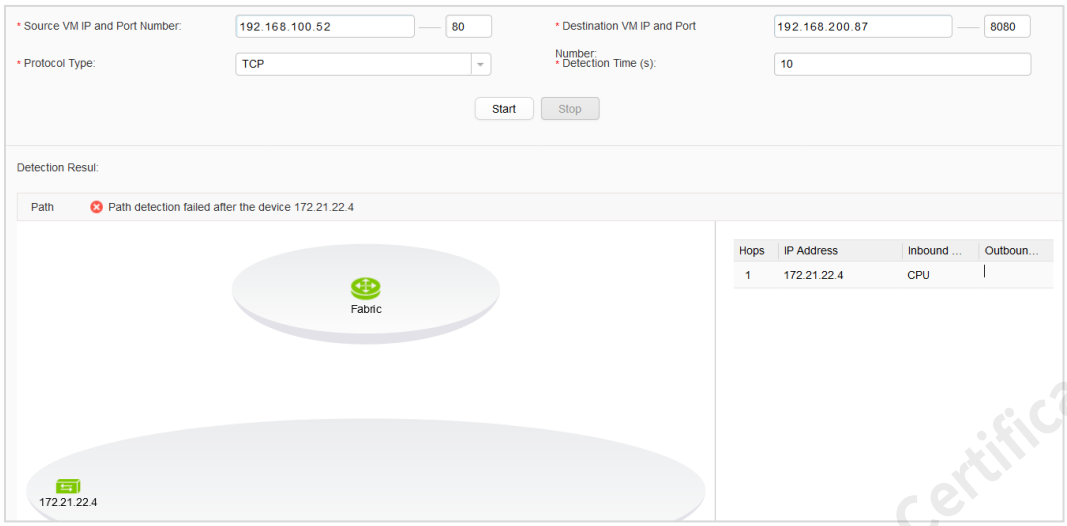

After the detection is complete, traffic between each hop from the Leaf node to the Spine node is displayed. The single-path detection result shows that the ingress port of the first hop is the CPU. The egress port of the last hop is not displayed because it cannot be detected.

----End

9.3.2.2 Multi Path Detection

Select **Multi Path Detection**, enter the source and destination VTEP IP addresses, number of detection packets, and detection time, and click **Start**.

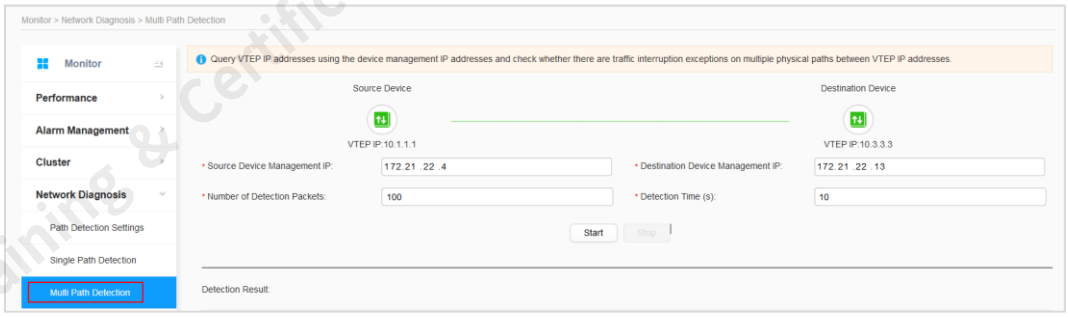

Detection completed

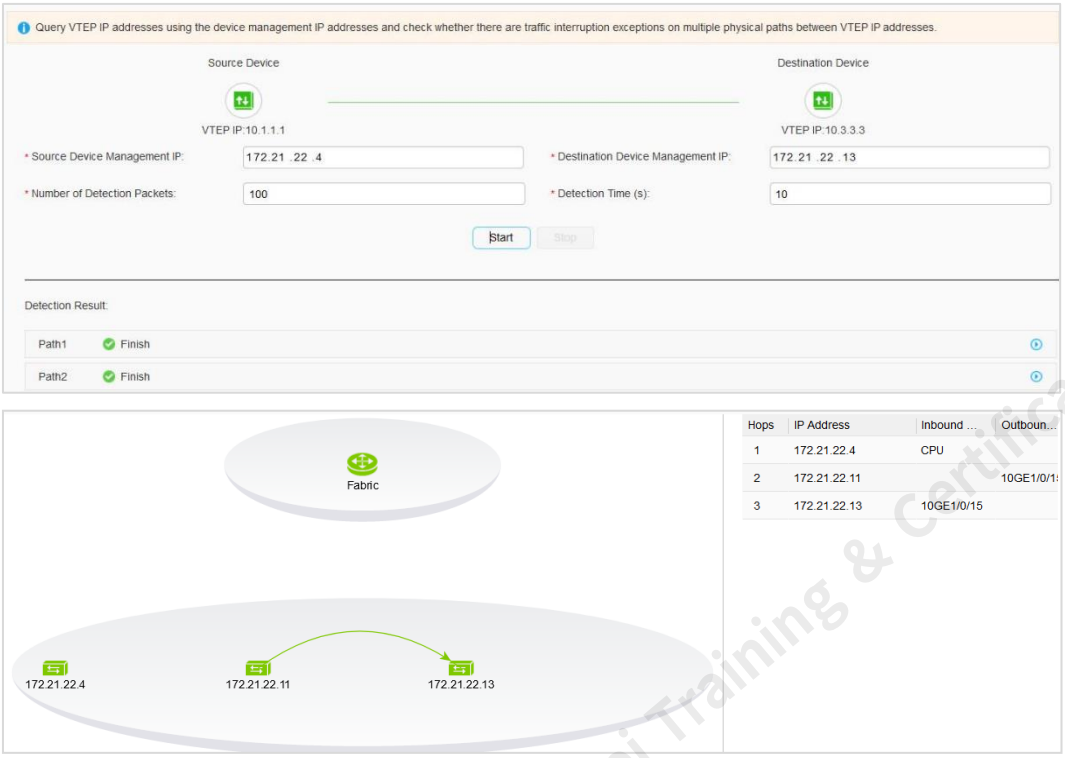

The VTEP IP addresses of two devices can be identified. Four paths are detected. You can expand to view the detailed information about the paths.

9.3.3 Loop Detection

VXLAN and VLAN loops may exist on a Fabric network. If suspected loops are detected, the Agile Controller-DCN samples ARP packets based on alarms reported by the devices and displays all suspected loops in a list.

- If the Agile Controller-DCN collects the same packet within a specific period of time, a real loop exists. The loop information is displayed on the loop detection page and rectification suggestions are provided. The loop detection result displays only the local loopback interfaces.
- If a device interface sends a large number of normal packets, the Agile Controller-DCN may fail to collect the same packet, and therefore cannot determine whether a loop exists. In this case, users can log in to the device and manually confirm whether a loop exists based on the suspected loop information.

VXLAN loops are classified into the following scenarios:

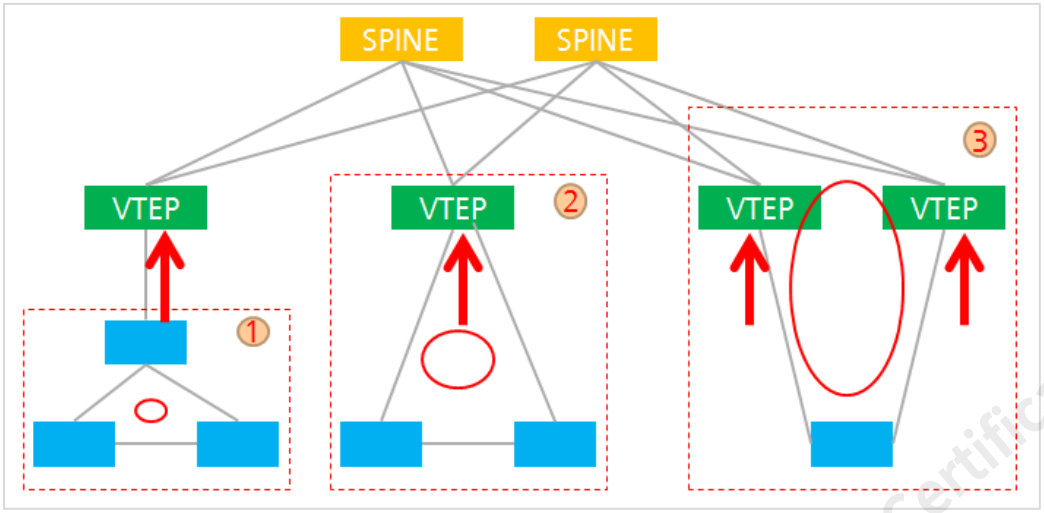

Step 1 **Enable loop alarm on devices.**

Log in to all network devices and enable the loop alarm function.

```
[Spine-1]mac-address flapping periodical trap enable
```

Step 2 **Enable loop detection on the Agile Controller-DCN.**

Choose Southbound Protocol > SNMP > SNMP Configuration, enable Trap Service, set Listening Port, and click Apply.

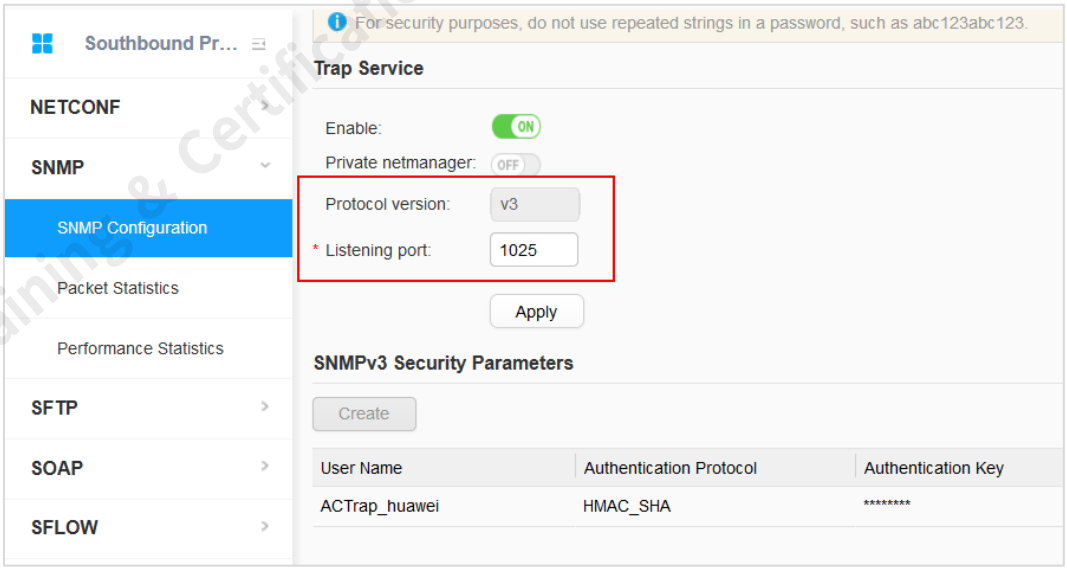

Step 3 **Perform loop detection.**

On the Agile Controller-DCN main menu, choose **Monitor** > **Network Diagnosis** > **Loop Detection** and authorize the Agile Controller-DCN to change the sFlow sampling configuration of the device.

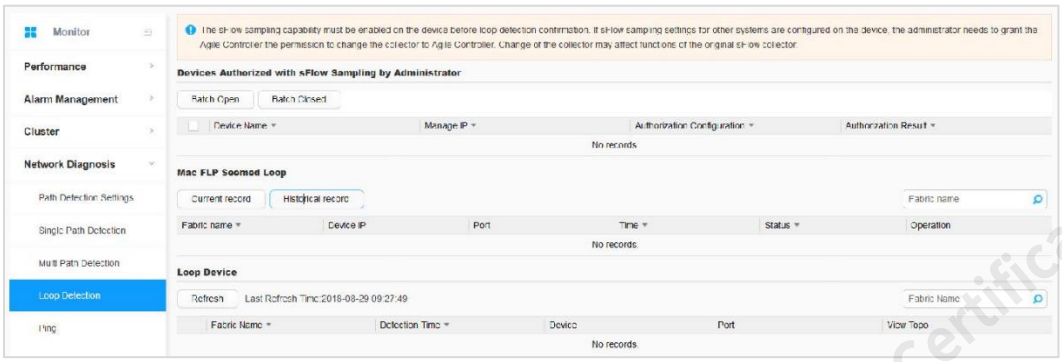

On the **Mac FLP Seemed Loop** tab, you can view the current and historical records of suspected loops and manually delete the current record.

----End

9.4 Agile Controller-DCN Data Inconsistency Discovery

When a device is managed by both the AC-DCN and network administrator, verifying the data consistency between the Agile Controller-DCN and forwarder is important. You need to check data consistency in the inconsistency discovery and inconsistency correct processes. Agile Controller-DCN sends a configuration query request to the forwarder and generates the configuration difference data. Users can correct the configuration inconsistency by synchronizing Agile Controller-DCN configuration to the forwarder or vice visa.

Step 1 **Configure the SFTP southbound server.**

On the Agile Controller-DCN main menu, choose **System** > **System Settings** > **Southbound Protocol** > **SFTP Southbound Server Settings**, and enable the SFTP service.

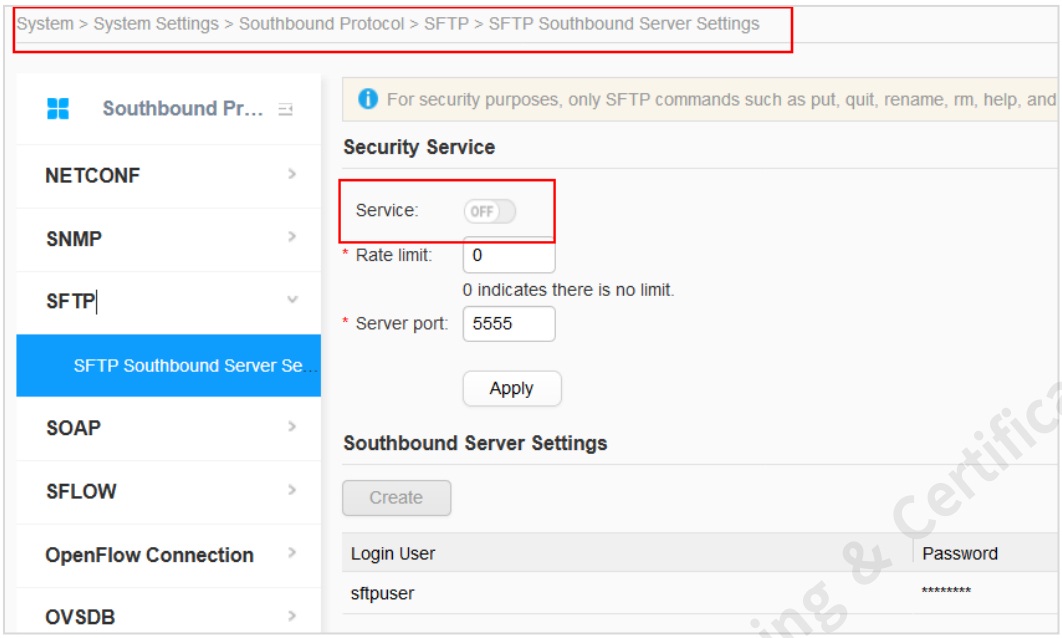

Set **Rate Limit** (maximum rate) and **Server Port**, and click **Apply**.

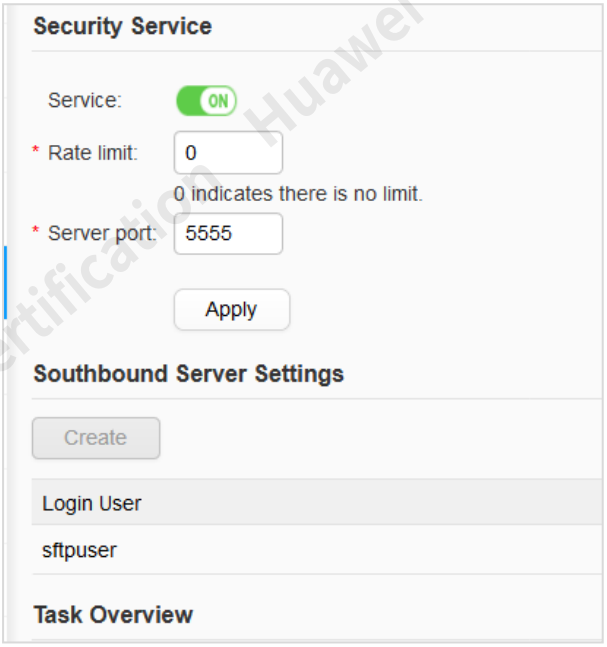

On the **Southbound Server Settings** tab, click **Create**. Set **Login user** and **Password** for the SFTP server. The Agile Controller-DCN automatically delivers the configuration.

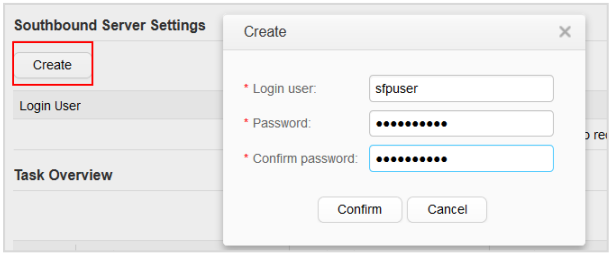

Step 2 **Configure the difference discovery policy.**

From the Agile Controller-DCN main menu, choose **System** > **System Settings** > **Data Consistency**.

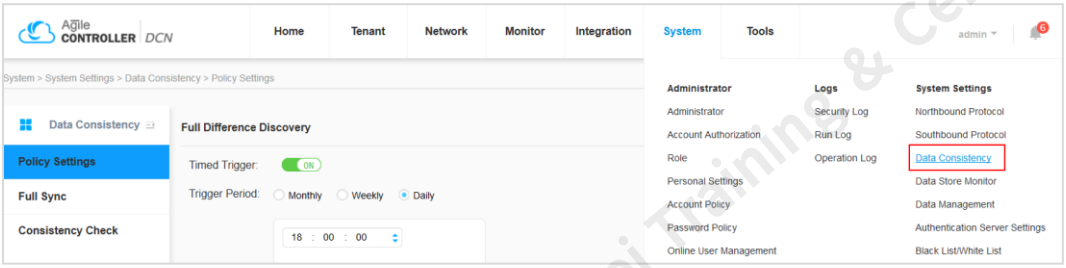

Enable Timed Trigger and set Trigger Period. In the Incremental Difference Discovery area, set Trigger Mode to Automatic and set Minimum Trigger Period. Click Apply.

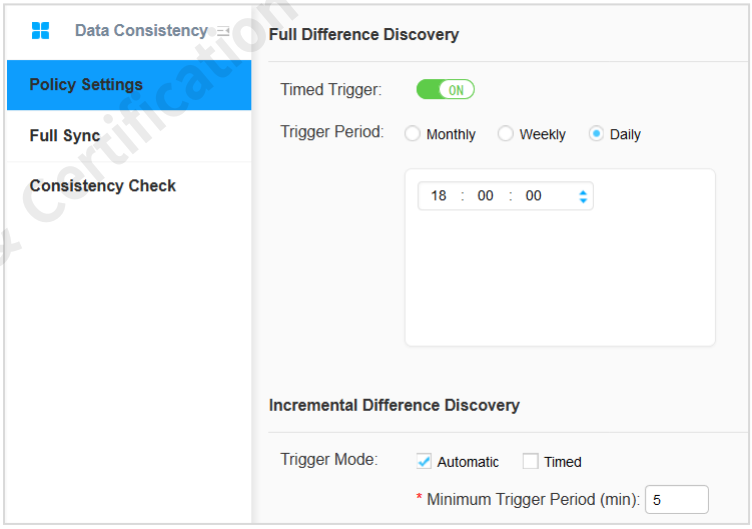

Step 3 **Perform full configuration synchronization for all forwarders.**

Select all devices and click **Full Sync**.

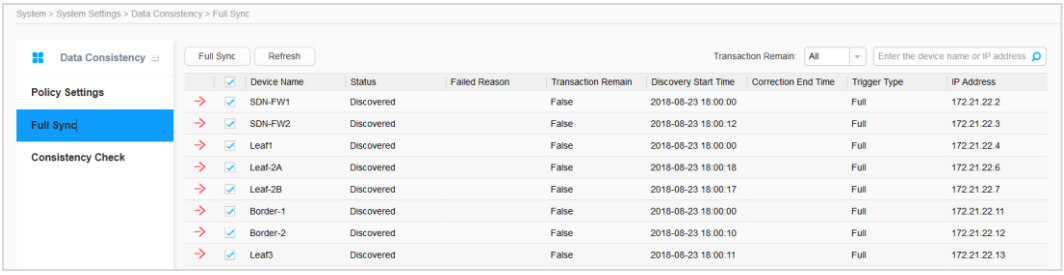

Devices automatically discover and synchronize the configuration.

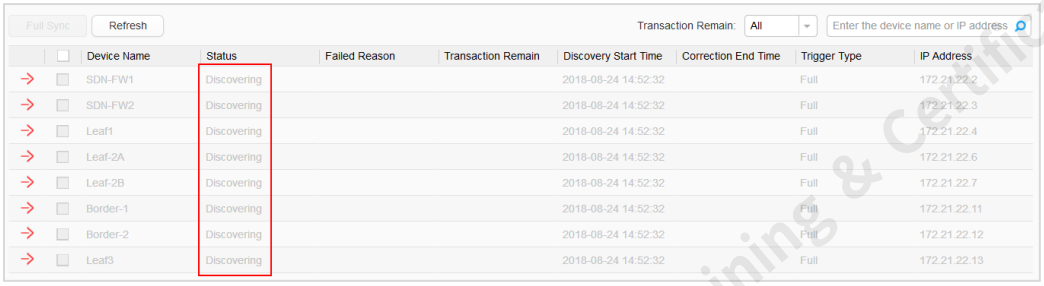

Click **View** to check the configuration differences.

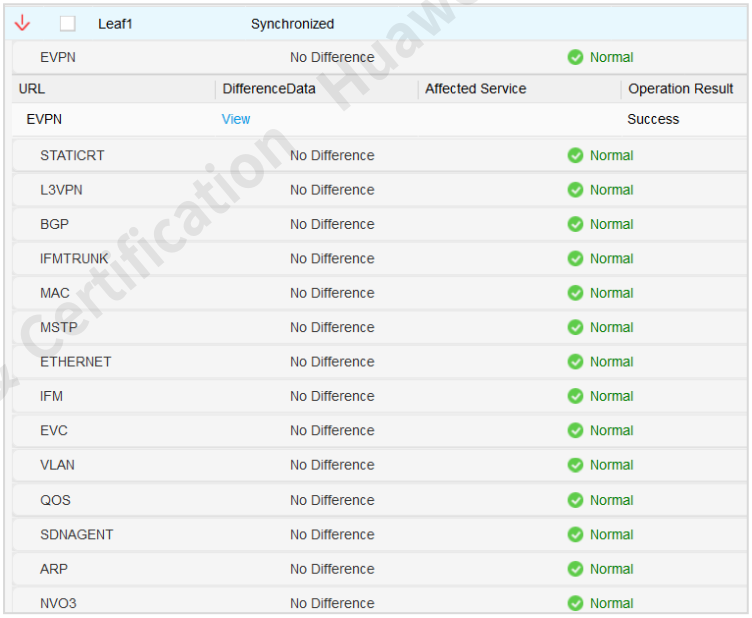

Step 4 **Perform inconsistency discovery.**

Click **Consistency Check** on the left pane and click **Inconsistency Discovery** in the right pane.

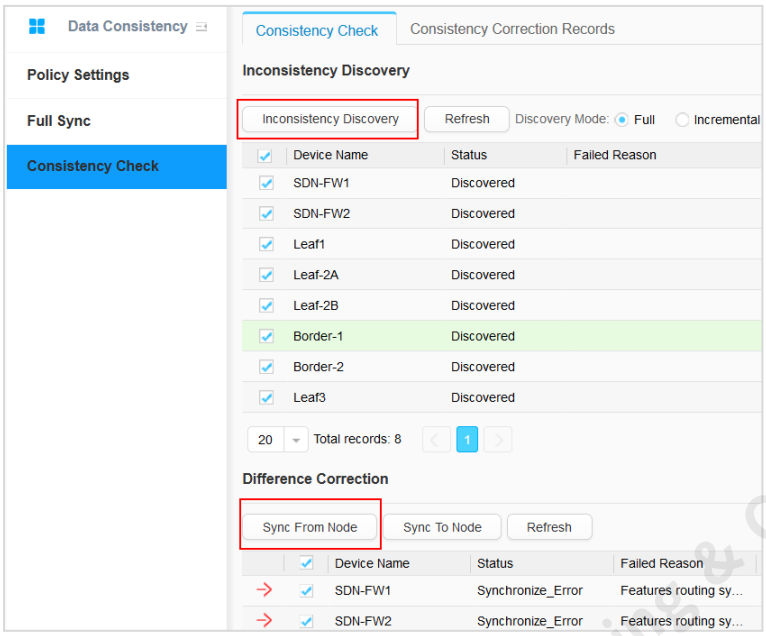

The consistency check includes difference discovery and difference correct. Difference discovery is used to detect the configuration difference between the forwarder and Agile Controller-DCN. Difference correct is used to eliminate the configuration difference between the forwarder and Agile Controller-DCN.

In the difference correct page, click a device and click **View**.

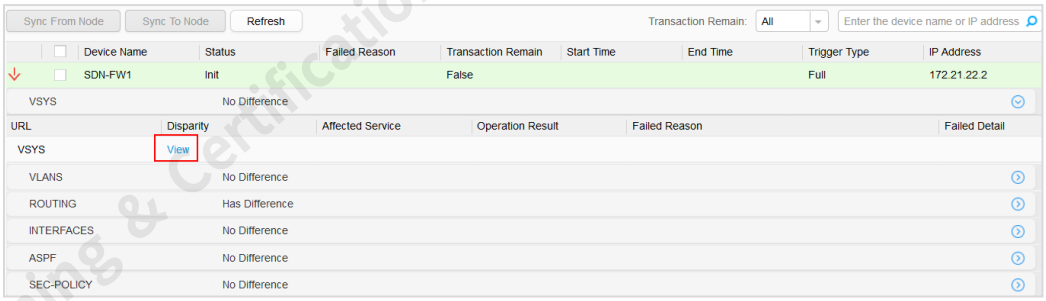

The difference page is displayed. Click  to view all differences.

You can choose to correct configuration based on the configuration of the forwarder or Agile Controller-DCN.

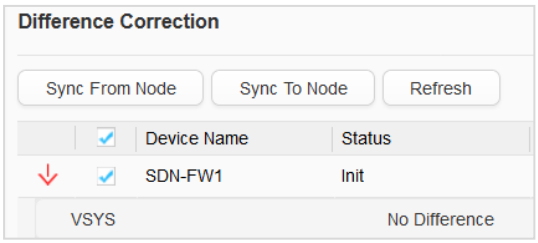

----End

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