CloudPond

Best Practices

Issue	01
Date	2024-05-16

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions

Notice

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

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

Huawei Technologies Co., Ltd.

Address:	Huawei Industrial Base Bantian, Longgang Shenzhen 518129 People's Republic of China
Website:	https://www.huawei.com
Email:	support@huawei.com

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process*. For details about this process, visit the following web page:
https://www.huawei.com/en/psirt/vul-response-process
For vulnerability information, enterprise customers can visit the following web page:
https://securitybulletin.huawei.com/enterprise/en/security-advisory

Contents

1 CloudPond Best Practices of Edge-Cloud Networking....1
1.1 Scenarios.... 1
1.1.1 Local Deployment.... 1
1.1.1.1 Deploying Local Office Services on CloudPond.... 2
1.1.1.2 Deploying ERP Services on CloudPond.... 3
1.1.2 Hybrid Deployment.... 4
1.1.3 Collaboration Deployment....5
1.2 Operations.... 6
1.2.1 Removing the EIP of the Data Plane from a Shared Bandwidth.... 7
1.2.2 Configuring the Data Center Public Network for CloudPond.... 8
1.2.3 Creating a VPC and Assigning a Subnet.... 8
1.2.4 Configuring Intranet Private Lines for CloudPond....9

2 Best Practices of Load Balancing in Customer-Built Data Centers.... 11
2.1 Load Balancing in Customer-Built Data Centers....11
2.1.1 LVS Load Balancing....11
2.1.2 Configuration on the IES Console....12
2.1.3 Configuration Preparations.... 13
2.1.3.1 Installing Tools on LVS Nodes....13
2.1.3.2 Switching the IPVS Kernel Module on LVS Nodes.... 13
2.1.4 Configuration Guide.... 14
2.1.4.1 Configuring Keepalived.... 14
2.1.4.2 Configuring the **/etc/sysctl.conf** file....16
2.1.5 Verification test.... 17
2.1.5.1 LVS Load Balancing Test.... 17
2.1.5.2 LVS Automatically Isolating Faulty Backend Nodes....18
2.1.5.3 LVS Automatically Managing Recovered Nodes.... 18
2.1.5.4 LVS Automatically Switching to the Standby Node When the Primary Node Is Faulty....18
2.1.5.5 LVS Automatically Switching to the Primary Node After the Node Is Recovered....19

1 CloudPond Best Practices of Edge-Cloud Networking

Background

CloudPond extends Huawei Cloud infrastructure and cloud services to customers' local data centers through innovative distributed cloud technologies. It provides different edge-cloud communication solutions for service planning and O&M based on your needs.

Scenario

Based on the deployment areas and purposes, CloudPond can be used for:

- Local deployment: All applications are deployed at the edge. The cloud is only responsible for resource management and monitoring.
- Hybrid deployment: Some applications are deployed at the edge, and the others on the cloud. Generally, applications in the two places do not communicate with each other.
- Collaboration deployment: Applications deployed at the edge and on the cloud need low latency, large data volume, or frequent communication.

1.1 Scenarios

1.1.1 Local Deployment

All your applications are deployed at the edge. Service communications are within your data center. So, configure the minimum bandwidth for edge-cloud communications and pay for the EIPs by traffic. You can configure edge-cloud networking with different reliability levels.

1.1.1.1 Deploying Local Office Services on CloudPond

Background

A customer plans to migrate their OA services from a local virtualization environment to the cloud by deploying CloudPond in their data center. The service status is as follows:

- OA services cannot be migrated to the public cloud because of the service purposes.
- Currently, there is no plan to deploy a large number of services on the public cloud.
- The service purposes and people who use OA services seldom change. So the service performance is predictable, and the services are not required for high reliability (available: 5 working days x 8 hours + some holidays).
- The public network bandwidth of the data center is limited and is shared with services of other domains.

Networking Configuration

The recommended edge-cloud networking is shown in **Figure 1-1**.

Figure 1-1 Networking design

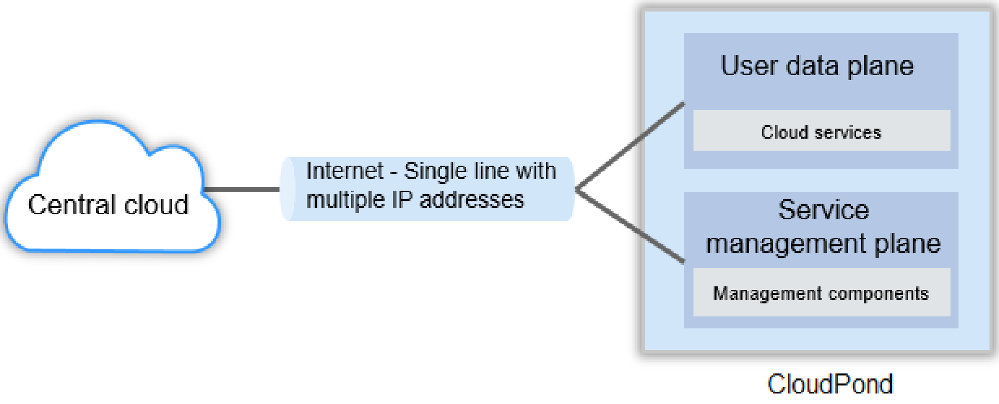

Key Points

- The minimum bandwidth is provided for management plane and data plane networks because all services are deployed at the edge and the edge-cloud traffic is small.
- CloudPond EIPs use the cloud's shared bandwidths (≥100 Mbit/s) because the customer data center's public network only has limited bandwidths shared by services.
- EIP billing by traffic is recommended for the default settings of edge-cloud communications because the cloud data plane does not need to communicate with the edge.
- Internet single-line connection is used because edge OA services are not required for high reliability. If the management plane is disconnected, edge

services can still run properly but cannot be managed by the central cloud, which has a small impact on the entire service system.

Related Operations

Removing the EIP of the Data Plane from a Shared Bandwidth

1.1.1.2 Deploying ERP Services on CloudPond

Background

A customer uses CloudPond to avoid data security risks of their purchased SaaS ERP software. The service status is as follows:

- ERP is mainly used inside the company. So, it needs to be deployed locally for low latency.
- ERP needs to be used outside the company during business trips.
- ERP requires high reliability. It cannot be interrupted except during holidays for both internal and external use.
- The network and security system of the customer data center meets the ERP deployment requirements.

Networking Configuration

The recommended edge-cloud networking is shown in **Figure 1-2**.

Figure 1-2 Solution design

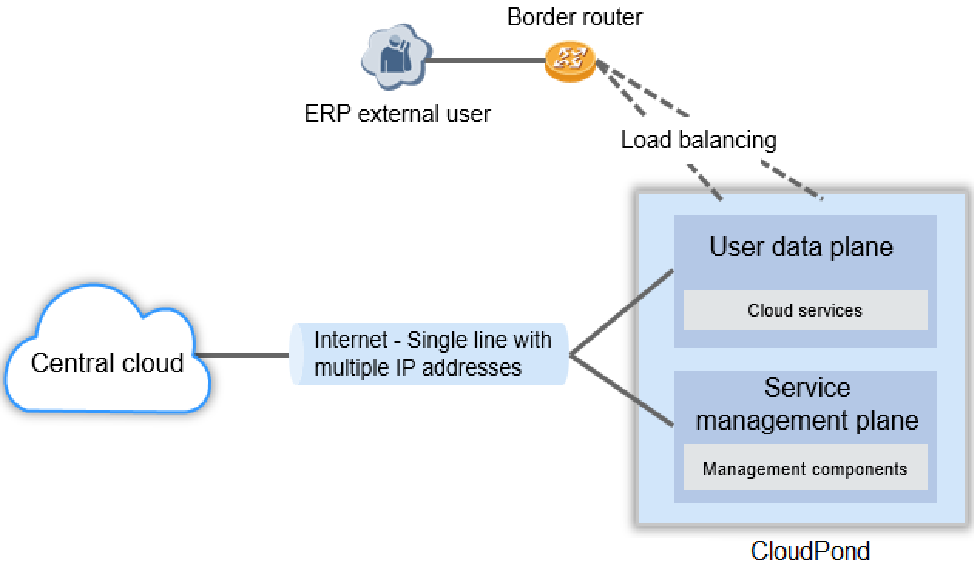

Key Points

- The minimum bandwidth is provided for management plane and data plane networks because all services are deployed at the edge and the edge-cloud traffic is small.

- It is recommended that the public network of the data center is expanded for ERP use only because the network of the data center may be used by users outside the company.
- EIP billing by traffic is recommended for the default settings of edge-cloud communications because the cloud data plane does not need to communicate with the edge.
- It is recommended that load balancing is built for public network access, meeting the high reliability requirements of ERP services. The internal network of the data center does not need load balancing due to its maturity.
- The security system of the customer data center should cover ERP applications. If the security system cannot cover them, CloudPond can provide the security capabilities.

Related Operations

1. **Removing the EIP of the Data Plane from a Shared Bandwidth**
2. **Configuring the Data Center Public Network for CloudPond**
3. **Building Load Balancing for Your Data Center**

1.1.2 Hybrid Deployment

Applications deployed at the edge and on the cloud do not need to communicate with each other. So, configure the minimum bandwidth for edge-cloud communications and pay for the EIPs by traffic. Services are processed in different VPCs based on their types. Edge-cloud networking with different reliability levels can be configured.

Background

A customer has deployed their company website and WeChat mini program on Huawei Cloud. They plan to migrate their OA services to the cloud by using CloudPond deployed in their data center. The service status is as follows:

- OA services cannot be migrated to the public cloud because of the service purposes. Services (websites and WeChat mini programs) deployed on the public cloud are running stably and do not need to be migrated to CloudPond.
- The service purposes and people who use OA services seldom change. So the service performance is predictable, and the services are not required for high reliability (available: 5 working days x 8 hours + some holidays).
- The public network bandwidth of the data center is limited and is shared with services of other domains.

Networking Configuration

The recommended edge-cloud networking is shown in **Figure 1-3**.

Figure 1-3 Solution design

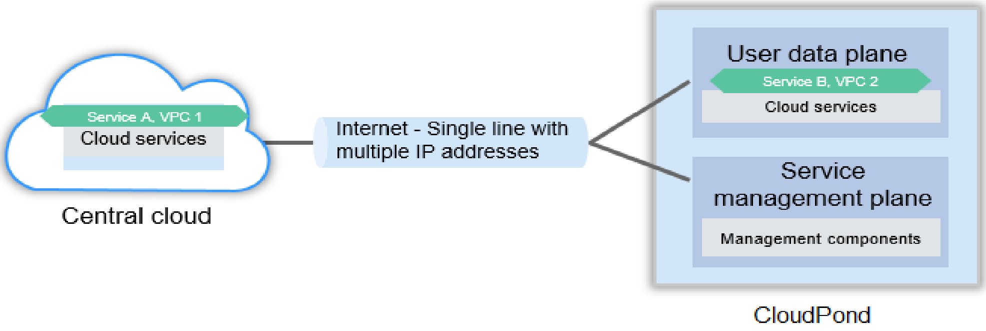

Key Points

- Services deployed at the edge and on the cloud do not need to communicate with each other, so they are located in different VPCs. (Services on the cloud are for external publicity, and services at the edge are used by employees.)
- CloudPond EIPs use the cloud's shared bandwidths because the customer data center 's public network only has limited bandwidths that are shared with services of other domains.
- EIP billing by traffic is recommended for the default settings of edge-cloud communications because the cloud data plane does not need to communicate with the OA services to be deployed.
- Internet single-line connection is used because edge OA services are not required for high reliability. If the management plane is disconnected, edge services can still run properly but cannot be managed by the central cloud, which has a small impact on the entire service system.

Related Operations

1. **Creating a VPC and Assigning a Subnet**
2. **Removing the EIP of the Data Plane from a Shared Bandwidth**

1.1.3 Collaboration Deployment

Applications deployed at the edge and on the cloud need to communicate with each other for data transmission, backup, API invoking, permission authentication, and so on. EIP and shared bandwidth configurations need to be adjusted to meet specific service requirements. Edge-cloud networking with different reliability levels can be configured.

Background

A customer plans to build a data analysis platform based on industry AI models to provide more accurate sales and production prediction for their core services. They decide to deploy CloudPond in their own data center and use it together with the central cloud. The service status is as follows:

- Data needs to be collected and processed at scattered points, and then stored in a data warehouse for analysis. Huawei Cloud IoT will be used at each point for public cloud access.
- Some analyzed raw data requires high security and cannot be stored outside the customer data center for a long time according to industry laws and regulations.
- There is a private line between the data center and the central cloud, which can be used for heavy-traffic data transmission.

Networking Configuration

The recommended edge-cloud networking is shown in **Figure 1-4**.

Figure 1-4 Solution design

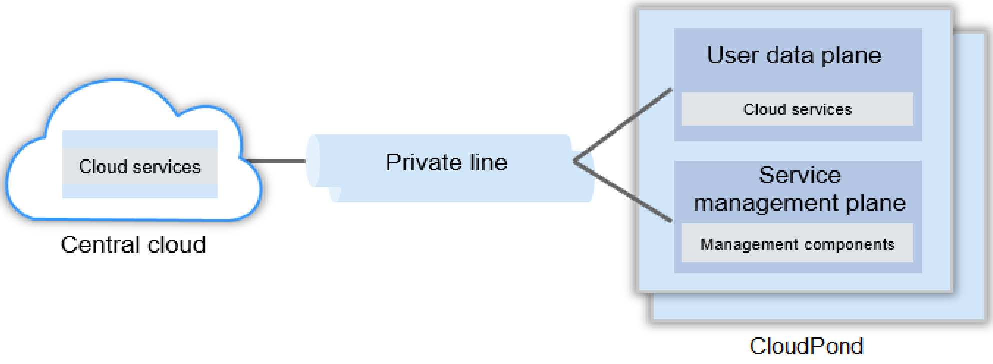

Key Points

- Services that do not need to store data for a long time, such as data collection, data processing, and AI modeling, are deployed on the public cloud. Services such as databases, data analysis, data warehouses, and BI are deployed on CloudPond. This ensures security compliance and reduces overall costs.
- The bandwidth of the private line between the customer data center and Huawei Cloud can be used directly without configuring public IP addresses or shared bandwidths.
- The edge and the cloud communicate with each other in a VPC for frequent data transmission and security.
- The edge and the cloud are interconnected by private lines for high service reliability. Multiple CloudPond sites are deployed locally to ensure the high reliability of core data.

Related Operations

Configuring Intranet Private Lines for CloudPond

1.2 Operations

1.2.1 Removing the EIP of the Data Plane from a Shared Bandwidth

Procedure

1. Log in to the management console using a Huawei Cloud account with CloudPond purchased.
2. Select a Huawei Cloud region where CloudPond is deployed.
3. In the service list, choose **Networking** > **Virtual Private Cloud**.
4. In the navigation pane on the left, choose **Elastic IP and Bandwidth** > **Shared Bandwidths**.
5. Locate the target bandwidth and click its name.
6. On the shared bandwidth details page, select the EIP to be removed. Click **Remove**.

Figure 1-5 Removing an EIP

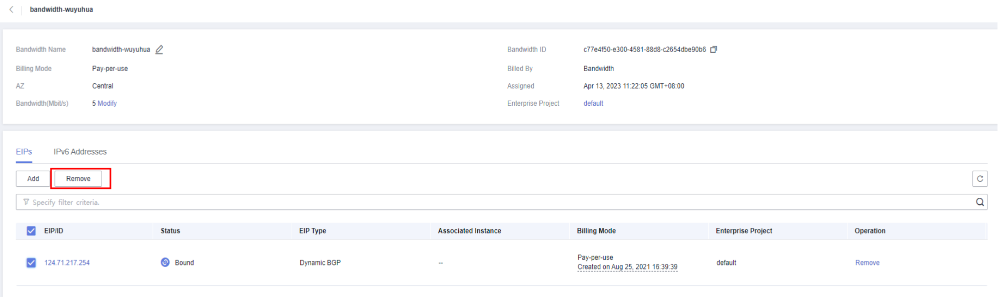

bandwidth-wuyuhua

Bandwidth Name	bandwidth-wuyuhua	Bandwidth ID	c77e4f50-e300-4581-88d8-c2654dbe90b6
Billing Mode	Pay-per-use	Billed By	Bandwidth
AZ	Central	Assigned	Apr 13, 2023 11:22:05 GMT+08:00
Bandwidth(Mbit/s)	5 Modify	Enterprise Project	default

☑ EIP/ID	Status	EIP Type	Associated Instance	Billing Mode	Enterprise Project	Operation
☑ 124.71.217.254	Bound	Dynamic BGP	--	Pay-per-use Created on Aug 25, 2021 16:39:39	default	Remove

7. In the **EIP Bandwidth After Removal** area, select **Traffic** for **Billing Mode**.

Figure 1-6 Billing by traffic

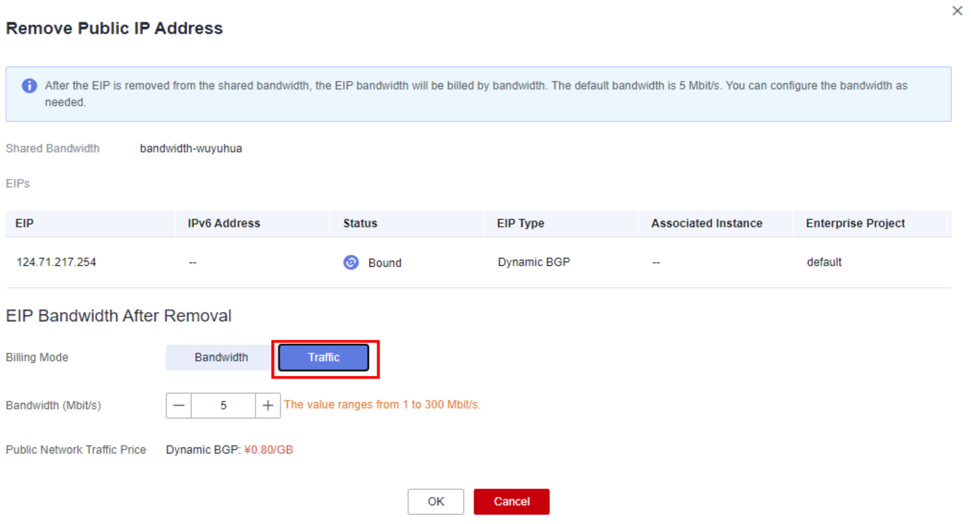

Remove Public IP Address

After the EIP is removed from the shared bandwidth, the EIP bandwidth will be billed by bandwidth. The default bandwidth is 5 Mbit/s. You can configure the bandwidth as needed.

Shared Bandwidth bandwidth-wuyuhua

EIPs

EIP	IPv6 Address	Status	EIP Type	Associated Instance	Enterprise Project
124.71.217.254	--	Bound	Dynamic BGP	--	default

EIP Bandwidth After Removal

Billing Mode: Bandwidth / Traffic

Bandwidth (Mbit/s): 5 The value ranges from 1 to 300 Mbit/s.

Public Network Traffic Price: Dynamic BGP: ¥0.80/GB

NOTE

Billing by traffic is the most cost effective when the total edge-cloud communication traffic is low.

1.2.2 Configuring the Data Center Public Network for CloudPond

Procedure

1. Log in to the management console using a Huawei Cloud account with CloudPond purchased.
2. In the upper right corner of the management console, choose **More** > **Service Tickets** > **Create Service Ticket**.
3. On the **Create Service Ticket** page, enter **CloudPond** in the search box of **My Service/Product**. Click **CloudPond**.

Figure 1-7 Searching for a product/service

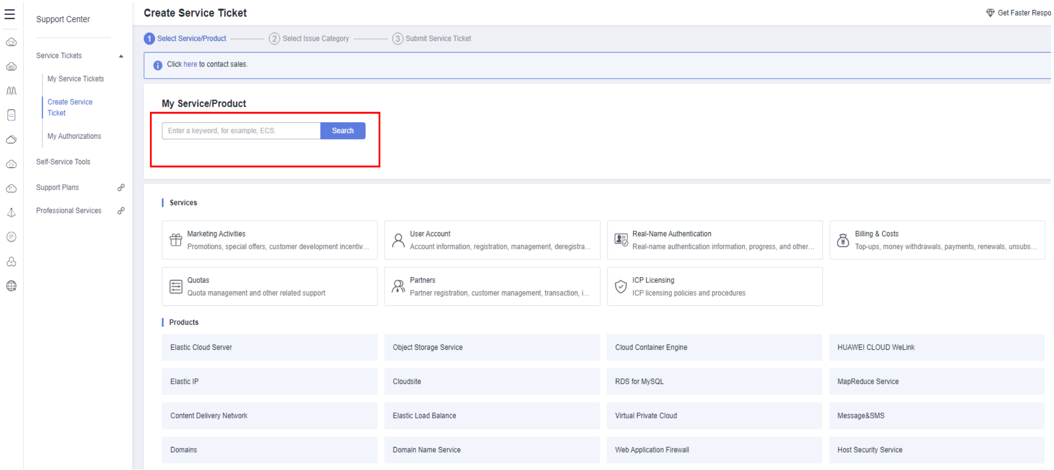

4. On the **Issue Categories** page, click **Other**. Click **Create Now**.
5. Describe the public network resources purchased for your data center (including the public IP address and bandwidth for ERP services). Leave your contact information, and wait for the service personnel to contact you.
6. After receiving the service ticket, Huawei Cloud customer service personnel will verify the information and arrange public network resources.
7. After the configuration is complete, you can view the shared bandwidth and EIP to be assigned on the **Shared Bandwidths** page of **Elastic IP** service.
8. The EIP can be used by ERP VMs.

1.2.3 Creating a VPC and Assigning a Subnet

Procedure

1. Log in to the management console using a Huawei Cloud account with CloudPond purchased.
2. Select a Huawei Cloud region where CloudPond is deployed.
3. In the service list, choose **Networking** > **Virtual Private Cloud**. The **Network Console** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud** > **My VPCs**.

5. In the upper right corner of the **Virtual Private Cloud** page, click **Create VPC**.
6. Select a region. Specify the name and network segment, and select an enterprise project based on your needs.
7. Select an AZ for the default subnet based on your needs. If the subnet is used on the cloud, select **Central** for **AZ**. If the subnet is used on CloudPond, select **Edge** for **AZ**. See **Figure 1-8**.

Figure 1-8 Selecting an AZ

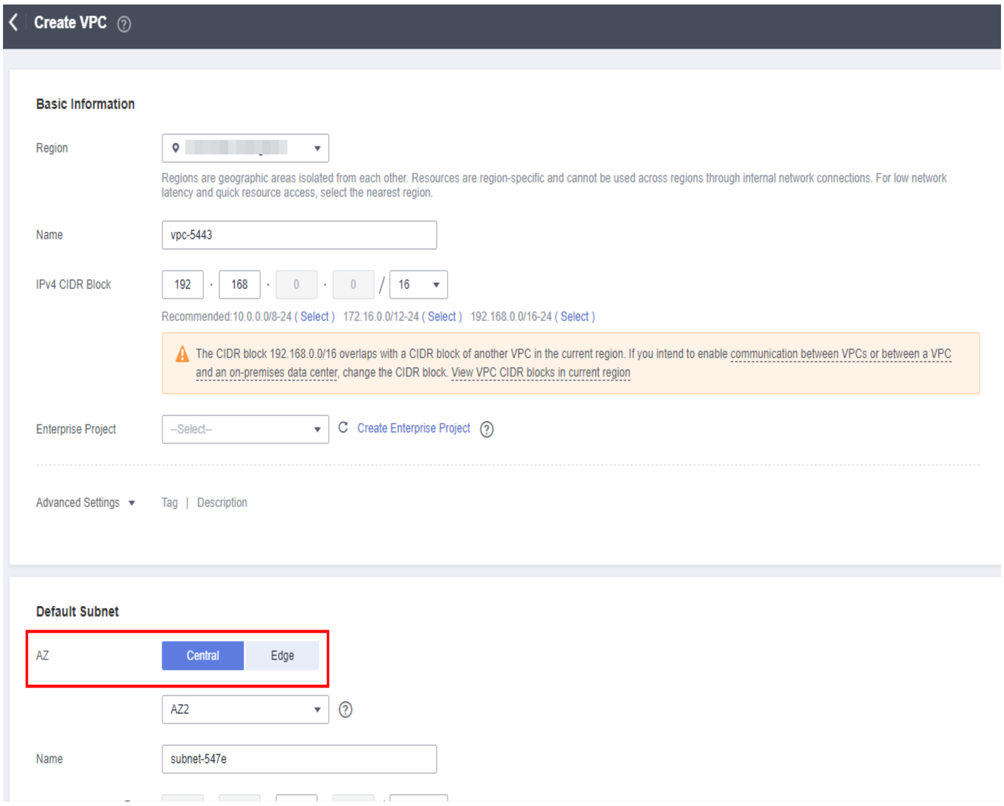

Create VPC

Basic Information

Region [redacted]

Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.

Name vpc-5443

IPv4 CIDR Block 192 · 168 · 0 · 0 / 16

Recommended:10.0.0.0/8-24 (Select) 172.16.0.0/12-24 (Select) 192.168.0.0/16-24 (Select)

The CIDR block 192.168.0.0/16 overlaps with a CIDR block of another VPC in the current region. If you intend to enable communication between VPCs or between a VPC and an on-premises data center, change the CIDR block. View VPC CIDR blocks in current region

Enterprise Project --Select-- Create Enterprise Project

Advanced Settings Tag | Description

Default Subnet

AZ Central Edge

AZ2

Name subnet-547e

8. Click **Create Now**. The new VPC will be displayed in the VPC list.

1.2.4 Configuring Intranet Private Lines for CloudPond

Procedure

1. Log in to the management console using a Huawei Cloud account with CloudPond purchased.
2. In the upper right corner of the management console, choose **More** > **Service Tickets** > **Create Service Ticket**.
3. On the **Create Service Ticket** page, enter **CloudPond** in the search box of **My Service/Product**. Click **CloudPond**.
4. On the **Issue Categories** page, click **Other**. Click **Create Now**.
5. Describe the detailed information about the Direct Connect in use, including the physical connection, virtual gateways, and virtual interfaces. If you do not have a Direct Connect, establish one before performing this step.

📖 NOTE

For details, see **Establishing a Direct Connection**.

6. After receiving the service ticket, Huawei Cloud customer service personnel will verify the information and arrange the connection.
7. The edge and the cloud can communicate with each other using the Direct Connection.

2 Best Practices of Load Balancing in Customer-Built Data Centers

2.1 Load Balancing in Customer-Built Data Centers

2.1.1 LVS Load Balancing

LVS FULLNAT mode: DNAT is performed on packets to change VIP VPORT to RIP RPORT, and SNAT is performed to change CIP CPORT to LIP LPORT. The packet is forwarded to the RealServer for processing. NAT is performed in the opposite direction to send packets returned from the RealServer to LVS through LIP. No configuration is required for back-end network devices and the RealServer.

NOTE

LVS (Linux virtual server) is a virtual linux cluster server. A user requests an LVS VIP. LVS forwards the request to a backend server based on the forwarding mode and algorithm. The backend server receives the request and returns it to the user. **Figure 2-1** shows the LVS networking.

Figure 2-1 LVS networking

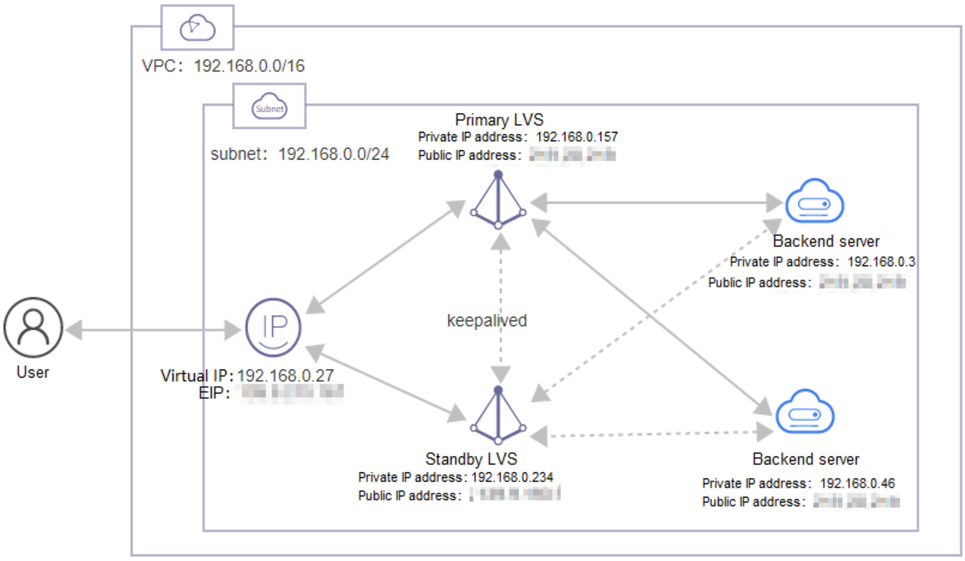
VPC：192.168.0.0/16
subnet：192.168.0.0/24
Primary LVS
Private IP address：192.168.0.157
Public IP address：[redacted]
Backend server
Private IP address：192.168.0.3
Public IP address：[redacted]
keepalived
User
Virtual IP：192.168.0.27
EIP：[redacted]
Standby LVS
Private IP address：192.168.0.234
Public IP address：[redacted]
Backend server
Private IP address：192.168.0.46
Public IP address：[redacted]

Table 2-1 IP information

Server Name	Private IP Address	Public IP Address	OS Type
LVS-MASTER	192.168.0.157	x.x.x.x	EulerOS 2.8
LVS-SLAVE	192.168.0.234	x.x.x.x	EulerOS 2.8
CDN-1	192.168.0.3	x.x.x.x	EulerOS 2.8
CDN-2	192.168.0.46	x.x.x.x	EulerOS 2.8

NOTE

- EIPs are not required in actual services. You can delete EIPs bound to each VM after the environment is deployed.
- The IP addresses used in the configuration commands in the following sections are for reference only. Replace them with the actual IP addresses in your service environment.

2.1.2 Configuration on the IES Console

1. Apply for a virtual IP address in the ECS subnet. Bind the virtual IP address to an EIP. Bind the virtual IP to the primary and standby LVS servers. See **Figure 2-2**.

Figure 2-2 Binding to the primary and standby LVS servers

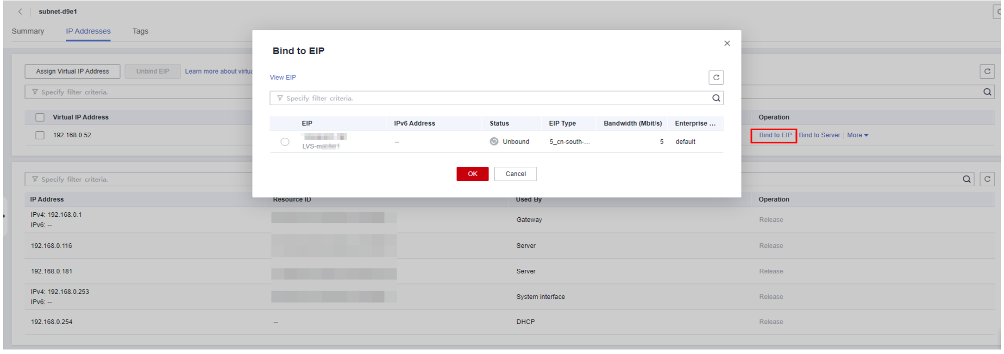

subnet-d9e1

Summary | IP Addresses | Tags

Assign Virtual IP Address | Unbind EIP | Learn more about virtu

Virtual IP Address	Operation
192.168.0.52	Bind to EIP | Bind to Server | More

Bind to EIP

View EIP

	EIP	IPv6 Address	Status	EIP Type	Bandwidth (Mbit/s)	Enterprise ...
○	[redacted] LVS-m[redacted]	--	Unbound	5_cn-south-...	5	default

OK | Cancel

IP Address	Resource ID	Used By	Operation
IPv4: 192.168.0.1 IPv6: --	[redacted]	Gateway	Release
192.168.0.116	[redacted]	Server	Release
192.168.0.181	[redacted]	Server	Release
IPv4: 192.168.0.253 IPv6: --	[redacted]	System interface	Release
192.168.0.254	--	DHCP	Release

2. **Figure 2-3** shows the EIP bound to the virtual IP address.

Figure 2-3 Binding an EIP

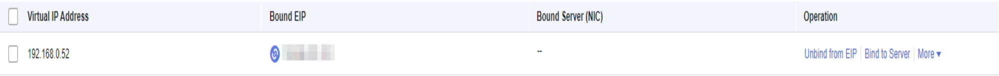

Virtual IP Address	Bound EIP	Bound Server (NIC)	Operation
192.168.0.52	[redacted]	--	Unbind from EIP | Bind to Server | More

2.1.3 Configuration Preparations

2.1.3.1 Installing Tools on LVS Nodes

1. Run the following commands in sequence to install tools on LVS nodes:

   yum insatll keepalived

   yum install ipvsadm

   yum install ipvs-fnat

   yum install ipvs-fnat-tools

   NOTE

   Run these commands on both primary and standby LVS nodes.

2.1.3.2 Switching the IPVS Kernel Module on LVS Nodes

Overview

This chapter describes how to switch the IPVS kernel module so that the FNAT mode can be used.

Precautions

1. To switch the IPVS kernel module, you need to have the root permission.
2. Keepalived must be disabled before you switch the kernel.
3. During the switchover, the traffic is interrupted and IPVS tables in all kernels are deleted.

Verification

1. Run the following command to switch the IPVS mode to the FullNAT mode:

   ipvs-switch fnat

```
[root@localhost keepalived]# ipvs-switch fnat
Clear ip_vs tables
rmmod ip_vs modules
doing depmod
modprobe ip_vs modules
switch to fnat mode successfully
```

NOTE

Run these commands on both primary and standby LVS nodes.

2.1.4 Configuration Guide

2.1.4.1 Configuring Keepalived

Procedure

1. Run the following command in the background to modify the **/etc/keepalived/keepalived.conf** file and configure Keepalived for the primary LVS node:

   vi /etc/keepalived/keepalived.conf

```
! Configuration File for keepalived

global_defs {
    router_id master_node
}

local_address_group laddr_gl {
192.168.0.157 # IP address of the primary LVS node
}

vrrp_instance VI_1 {
    state MASTER
    interface eth0
    virtual_router_id   51
    priority 100
    advert_int   1
    authentication {
        auth_type PASS
        auth_pass 123456
    }
    virtual_ipaddress {
        192.168.0.27
}
notify_master "/etc/keepalived/notify.sh master"
notify_backup "/etc/keepalived/notify.sh backup"
}

Vitual IP port
virtual_server 192.168.0.27 80 {
    delay_loop 6
    lb_algo rr
    lb_kind FNAT
    protocol TCP
    laddr_group_name laddr_gl
Backend server node
    real_server 192.168.0.3 80 {
        weight 1
        TCP_CHECK {
            connect_timeout 15
            connect_port 80
            nb_get_retry 5
            delay_before_retry 3
        }
    }
    real_server 192.168.0.46 80 {
        weight 1
        TCP_CHECK {
            connect_timeout 15
            connect_port 80
            nb_get_retry 5
            delay_before_retry 3
        }
    }
}
```

2. Run the following command in the background to modify the **/etc/keepalived/keepalived.conf** file and configure Keepalived for the standby LVS node:

   vi /etc/keepalived/keepalived.conf

```
! Configuration File for keepalived

global_defs {
    router_id master_node
}

local_address_group laddr_gl {
192.168.0.234 #Local IP address
}

vrrp_instance VI_1 {
    state BACKUP
    interface eth0
    virtual_router_id   51
priority 99 #The value must be lower than that of the primary node.
    advert_int   1
    authentication {
        auth_type PASS
        auth_pass 123456
    }
    virtual_ipaddress {
        192.168.0.27
}
notify_master "/etc/keepalived/notify.sh master"
notify_backup "/etc/keepalived/notify.sh backup"
}

virtual_server 192.168.0.27 80 {
    delay_loop 6
    lb_algo rr
    lb_kind FNAT
    protocol TCP
    laddr_group_name laddr_gl
    real_server 192.168.0.3 80 {
        weight 1
        TCP_CHECK {
            connect_timeout 15
            connect_port 80
            nb_get_retry 5
            delay_before_retry 3
        }
    }
    real_server 192.168.0.46 80 {
        weight 1
        TCP_CHECK {
            connect_timeout 15
            connect_port 80
            nb_get_retry 5
            delay_before_retry 3
        }
    }
}
```

NOTE

Perform steps 3 to 6 on both the primary and standby LVS nodes.

3. Run the following command in the background to modify the **/etc/keepalived/notify.sh** file:

   vi /etc/keepalived/notify.sh

```
#!/bin/bash

if [ $# -eq 0 ]
then
   echo "input master/backup"
   exit 1
fi

case $1 in
master)
       sed -i "s/#laddr_group_name/laddr_group_name/" /etc/keepalived/keepalived.conf 2>/dev/null
       systemctl reload keepalived.service
;;
backup)
       if [ `grep "#laddr_group_name" /etc/keepalived/keepalived.conf|wc -l` -eq 0 ]
       then
           sed -i "s/laddr_group_name/#laddr_group_name/" /etc/keepalived/keepalived.conf 2>/dev/null
           systemctl reload keepalived.service
       fi
;;
esac
```

   chmod +x /etc/keepalived/notify.sh

4. Run the following command to enable keepalived:

   systemctl start keepalived

5. Check whether the configuration is modified:

   ipvsadm -Ln

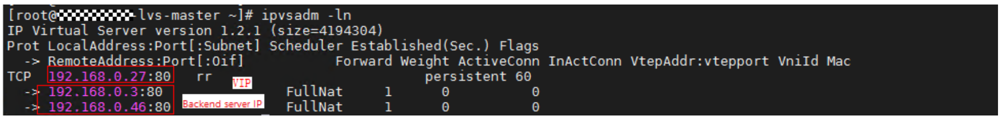

```
[root@[redacted]-lvs-master ~]# ipvsadm -ln
IP Virtual Server version 1.2.1 (size=4194304)
Prot LocalAddress:Port[:Subnet] Scheduler Established(Sec.) Flags
  -> RemoteAddress:Port[:Oif]           Forward Weight ActiveConn InActConn VtepAddr:vtepport VniId Mac
TCP  192.168.0.27:80    rr                      persistent 60
  -> 192.168.0.3:80                 FullNat    1      0          0
  -> 192.168.0.46:80                FullNat    1      0          0
```

NOTE

If the number of servers returned by ipvsadm is different from the actual number, check whether the listener port allows traffic to the associated backend server.

6. Run the following command to configure automatic startup:

   systemctl enable keepalived.service

2.1.4.2 Configuring the /etc/sysctl.conf file

Procedure

1. Run the following command to modify the kernel parameters:

   vi /etc/sysctl.conf

```
net.ipv4.ip_forward = 1
```

2. Run the following command to make the parameter settings to take effect:

   sysctl -p

NOTE

Run these commands on both primary and standby LVS nodes.

2.1.5 Verification test

2.1.5.1 LVS Load Balancing Test

Objectives

LVS can evenly forward requests to two backend server nodes.

Prerequisites

1. The LVS environment of the test cluster has been configured.
2. The web test service has been deployed on the primary and standby backend nodes.

   Install the httpd service on the [root@centos01 ~]# yum -y install httpd <!--Install the httpd service.-->

   [root@centos01 ~]# echo "www.test01.com" > /var/www/html/index.html

   [root@centos01 ~]# chmod 777 /var/www/html/index.html <!--Start the httpd service.-->

   [root@centos01 ~]# systemctl start httpd <!--Start the httpd service.-->

   Set the [root@centos01 ~]# systemctl enable httpd<!--Enable httpd to automatically start upon system boot.-->

   Install the httpd service on the [root@centos02 ~]# yum -y install httpd <!--Install the httpd service.-->

   [root@centos02 ~]# echo "www.test02.com" > /var/www/html/index.html <!--Create a website home page and write the test data. Write different content to the two backend nodes for clear display of test results.-->

   [root@centos02 ~]# chmod 777 /var/www/html/index.html <!--Start the httpd service.-->

   [root@centos02 ~]# systemctl start httpd <!--Start the httpd service.-->

   Set the [root@centos02 ~]# systemctl enable httpd<!--Enable httpd to automatically start upon system boot.-->

Test Methods

1. Run the **curl 192.168.0.27** (VIP) command for multiple times on nodes that can communicate with each other through the VIP.
2. View the response result. The web application content of the two backend nodes is randomly displayed (www.test01.com or www.test02.com).
3. Use the browser to access an EIP (http://x.x.x.x) bound to the VIP for multiple times.
4. View the response result. The web application content of the two backend nodes is randomly displayed (www.test01.com or www.test02.com).

2.1.5.2 LVS Automatically Isolating Faulty Backend Nodes

Objectives

When a node in the backend service cluster is faulty, LVS automatically isolates the faulty node.

Prerequisites

The LVS environment of the test cluster has been configured.

Test Methods

1. Stop a backend node, for example, CDN-1.
2. Run this command to check whether CDN-1 node is removed from the cluster.

   ipvsadm -L –n

   The following information is displayed:

   ```
   Prot LocalAddress:Port Scheduler Flags
     -> RemoteAddress:Port           Forward Weight ActiveConn InActConn
   TCP  192.168.0.184:80 wrr persistent 50
     -> 192.168.0.146:80             Route   3      0          0
   ```

2.1.5.3 LVS Automatically Managing Recovered Nodes

Objectives

After the faulty node in the backend service cluster is recovered, LVS automatically manages the recovered node.

Prerequisites

The LVS environment of the test cluster has been configured.

Test Methods

1. Restart the stopped node, for example, CDN-1.
2. Run this command to check whether CDN-1 node is managed again.

   ipvsadm -L –n

   The following information is displayed:

   ```
   Prot LocalAddress:Port Scheduler Flags
     -> RemoteAddress:Port           Forward Weight ActiveConn InActConn
   TCP  192.168.0.184:80 wrr persistent 50
     -> 192.168.0.146:80             Route   3      0          0
     -> 192.168.0.154:80             Route   3      0          0
   ```

2.1.5.4 LVS Automatically Switching to the Standby Node When the Primary Node Is Faulty

Objectives

When the keepalived service on the primary LVS node is faulty (for example, the keepalived service is disabled), the standby LVS node takes over services.

Prerequisites

The LVS primary and standby cluster have been correctly configured.

Test Methods

1. The VIP is on the primary node by default. Run the following command to check the network configuration of the primary node:

   ip addr

   The following information is displayed:

   ```
   [root@lvs-master ~]# ip addr
   eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
       link/ether 52:54:00:68:dc:b6 brd ff:ff:ff:ff:ff:ff
       inet 192.168.0.157/27 brd 182.148.15.255 scope global eth0
       inet 192.168.0.27/32 scope global eth0
       inet6 fe80::5054:ff:fe68:dcb6/64 scope link
          valid_lft forever preferred_lft forever
   ```

2. Run the following command to stop the keepalived service on the primary node.

   service keepalived stop

3. After the VIP is assigned to the standby node, run the following command on the primary node to verify that the VIP information does not exist.

   ip addr

   The following information is displayed:

   ```
   root@lvs-master ~]# ip addr
   eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
       link/ether 52:54:00:68:dc:b6 brd ff:ff:ff:ff:ff:ff
       inet 192.168.0.157/27 brd 182.148.15.255 scope global eth0
       inet 192.168.0.184/27 brd 182.148.15.255 scope global secondary eth0:0
       inet6 fe80::5054:ff:fe68:dcb6/64 scope link
          valid_lft forever preferred_lft forever
   ```

   Run the following command on the standby node to check whether the VIP information exists:

   ip addr

   The following information is displayed:

   ```
   [root@lvs-slave ~]# ip addr
   eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
       link/ether 52:54:00:68:dc:b6 brd ff:ff:ff:ff:ff:ff
       inet 192.168.0.234/27 brd 182.148.15.255 scope global eth0
       inet 192.168.0.27/32 scope global eth0
       inet 192.168.0.184/27 brd 182.148.15.255 scope global secondary eth0:0
       inet6 fe80::5054:ff:fe68:dcb6/64 scope link
          valid_lft forever preferred_lft forever
   ```

2.1.5.5 LVS Automatically Switching to the Primary Node After the Node Is Recovered

Objectives

When the keepalived service on the primary LVS node is faulty (for example, the keepalived service is disabled), the standby LVS node takes over services.

Prerequisites

The LVS primary and standby cluster have been correctly configured.

Test Methods

1. After the standby takes over services, the VIP is on the standby node by default. Run the following command on the standby node:

   ip addr

   The following information is displayed:

   ```
   [root@lvs-slave ~]# ip addr
   eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
       link/ether 52:54:00:68:dc:b6 brd ff:ff:ff:ff:ff:ff
       inet 192.168.0.234/27 brd 182.148.15.255 scope global eth0
       inet 192.168.0.27/32 scope global eth0
       inet6 fe80::5054:ff:fe68:dcb6/64 scope link
          valid_lft forever preferred_lft forever
   ```

2. Run the following command on the primary node to enable the keepalived service:

   service keepalived start

3. The VIP is assigned to the primary node. Run the following command on the primary node:

   ip addr

   The following information is displayed:

   ```
   [root@lvs-master ~]# ip addr
   eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
       link/ether 52:54:00:68:dc:b6 brd ff:ff:ff:ff:ff:ff
       inet 192.168.0.157/27 brd 182.148.15.255 scope global eth0
       inet 192.168.0.27/32 scope global eth0
       inet6 fe80::5054:ff:fe68:dcb6/64 scope link
          valid_lft forever preferred_lft forever
   ```

   Run the following command on the standby node to verify that the VIP information does not exist:

   ip addr

   The following information is displayed:

   ```
   [root@lvs-slave ~]# ip addr
   eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
       link/ether 52:54:00:68:dc:b6 brd ff:ff:ff:ff:ff:ff
       inet 192.168.0.234/27 brd 182.148.15.255 scope global eth0
       inet6 fe80::5054:ff:fe68:dcb6/64 scope link
          valid_lft forever preferred_lft forever
   ```