Express Connect

Best Practices

Best Practices

AnyTunnel VIP

AnyTunnel VIP provides a quick way for VPC to access cloud services. AnyTunnel VIP is a public VIP provided for VPC and can be accessed by all VPCs.

AnyTunnel VIP is 100.64.0.0/10, which is the same as that in a classic network. An IP belonging to 100.64.0.0/10, such as 100.64.1.1, can exist both in a class network and a VPC.

DNS, OSS, Log Service, and other cloud services are all using IPs belonging to 100.64.0.0/10. If you need to access these cloud services from the peer end of the leased line, namely your on-premises IDC, you must set the router interface pointing to VPC as the next hop of the route destined for 100.64.0.0/10 after you create the VBR. You also need to set the router interface pointing to Alibaba Cloud as the next hop of the route destined for 100.64.0.0/10 on the gateway device of the on-premises IDC. Because 100.64.0.0/10 is a reserved CIDR block of VPC, you need to split it into 100.64.0.0/11 and 100.96.0.0/11 and configure the two CIDR blocks on the VBR.

Configure routes on the VBR of the leased line

Set the router interface pointing to VPC as the next hop of the route destined for 100.64.0.0/11

Log on to the Express Connect console.

In the left-side navigation pane, select **Virtual Border Router**.

On the **My VBRs** page, select the target VBR and click **Manage**.

On the **VBR Details** page, click **Add Route Entry** and configure the route entry. The following configurations are used in this tutorial:

Destination CIDR Block: 100.64.0.0/11

Next Hop Direction: **To VPC**

Next Hop: Select the exit for data packets. In this tutorial, select the router interface on the VBR.

Click **OK** to complete the configuration.

Set the router interface pointing to VPC as the next hop of the route destined for 100.96.0.0/11

Go back to the **VBR Details** page, click **Add Route Entry** and configure the route entry. The following configurations are used in this tutorial:

Destination CIDR Block: 100.96.0.0/11

Next Hop Direction: **To VPC**

Next Hop: Select the exit for data packets. In this tutorial, select the router interface on the VBR.

Click **OK** to complete the configuration.

Configure the route on the customer-side access device of the leased line

Add a static route pointing to Alibaba Cloud on the customer-side access device of the leased line:

ip route 100.64.0.0/10 {Alibaba Cloud-side IP address}

After completing the leased line access, you need to test the performance of the leased line to ensure that it can meet your service needs.

Prerequisites

Before the test, ensure you have made the following preparations:

Complete configurations regarding leased line access and routes. The on-premises IDC must be connected to the VPC through a leased line.

Prepare a network access device for the on-premises IDC: The network access device is subjected to stress testing to test the packets per second (pps) of the on-premises IDC. It serves as the client or server in the netperf or iperf3 test.

In this tutorial, the IP address of the on-premises IDC is 192.168.100.1.

Prepare eight VPC ECS instances: The ECS instances serve as clients or servers in the netperf

or iperf3 test. They are connected to the network access device of the on-premises IDC to transmit test configurations and test results.

In this tutorial, eight ECS instances are used, of which the specification is ecs.se1.2xlarge, the image is centos_7_2_64_40G_base_20170222.vhd, and the IP address range is 172.16.0.2 – 172.16.0.9.

Build the test environment

Install netperf

Netperf is a tool for testing network performance and is based on TCP or UDP transmission.

Follow these steps to install netperf on the network access device of the on-premises IDC and the eight ECS instances.

Run the following command to download netperf:

```
wget -c "https://codeload.github.com/HewlettPackard/netperf/tar.gz/netperf-2.5.0" -O netperf-2.5.0.tar.gz
```

Run the following command to install netperf:

```
tar -zxvf netperf-2.5.0.tar.gz
cd netperf-netperf-2.5.0
./configure && make && make install && cd ..
```

Run the netperf -h and netserver -h commands to verify if the installation is successful.

Install iperf3

Iperf3 is a tool for testing network performance and can test the maximum TCP or UDP bandwidth.

Follow these steps to install iperf3 on the network access device of the on-premises IDC and the eight ECS instances.

Run the following command to download iperf3:

```
yum install git -y
git clone https://github.com/esnet/iperf
```

Run the following command to install iperf3:

```
cd iperf
./configure && make && make install && cd ..
cd src
ADD_PATH="$(pwd)"
PATH="${ADD_PATH}:${PATH}"
export PATH
```

Run the iperf3 -h command to verify if the installation is successful.

Enable the multiple queue feature

Run the following command on the network access device of the on-premises IDC to enable the multiple queue feature (assume the interface connected to the leased line is eth0):

```
ethtool -L eth0 combined 4
echo "ff" > /sys/class/net/eth0/queues/rx-0/rps_cpus
echo "ff" > /sys/class/net/eth0/queues/rx-1/rps_cpus
echo "ff" > /sys/class/net/eth0/queues/rx-2/rps_cpus
echo "ff" > /sys/class/net/eth0/queues/rx-3/rps_cpus
```

Use netperf to test the packet forwarding performance of the leased line

After being installed, netperf creates two command line tools: netserver (server side) and netperf (client side). Main parameters of the two tools are shown in the following table.

Tool name	Main parameter	Parameter description
netserver (server side: receiving side tool)	-p	The port of the server.
netperf	-H	The IP address of the network access device of on-premises IDC or the VPC server.
	-p	The port of the network access device of the on-premises IDC or the VPC server.
	-l	The duration of running.
	-t	The protocol used for sending packets: TCP_STREAM or UDP_STREAM.

		We recommend UDP_STREAM.
	-m	The data packet size. - We recommend that you set the value to 1 when testing pps. - We recommend that you set the value to 1400 when testing bps (bit per second).

Test the inbound direction

Start the netserver process on the network access device of the on-premises IDC and specify different ports:

```
netserver -p 11256
netserver -p 11257
netserver -p 11258
netserver -p 11259
netserver -p 11260
netserver -p 11261
netserver -p 11262
netserver -p 11263
```

Start the netperf process on the eight ECS instances in the VPC and specify different ports connecting to the network access device of the on-premises IDC.

```
netperf -H 192.168.100.1 -p 11256 -t UDP_STREAM -l 300 -- -m 1 #first ECS instance
netperf -H 192.168.100.1 -p 11257 -t UDP_STREAM -l 300 -- -m 1 #second ECS instance
netperf -H 192.168.100.1 -p 11258 -t UDP_STREAM -l 300 -- -m 1 #third ECS instance
netperf -H 192.168.100.1 -p 11259 -t UDP_STREAM -l 300 -- -m 1 #fourth ECS instance
netperf -H 192.168.100.1 -p 11260 -t UDP_STREAM -l 300 -- -m 1 #fifth ECS instance
netperf -H 192.168.100.1 -p 11261 -t UDP_STREAM -l 300 -- -m 1 #sixth ECS instance
netperf -H 192.168.100.1 -p 11262 -t UDP_STREAM -l 300 -- -m 1 #seventh ECS instance
netperf -H 192.168.100.1 -p 11263 -t UDP_STREAM -l 300 -- -m 1 #eighth ECS instance
```

3. If you want to test bps, change the preceding commands to:

```
netperf -H 192.168.100.1 -p 11256 -t UDP_STREAM -l 300 -- -m 1400 #first ECS instance
netperf -H 192.168.100.1 -p 11257 -t UDP_STREAM -l 300 -- -m 1400 #second ECS instance
netperf -H 192.168.100.1 -p 11258 -t UDP_STREAM -l 300 -- -m 1400 #third ECS instance
netperf -H 192.168.100.1 -p 11259 -t UDP_STREAM -l 300 -- -m 1400 #fourth ECS instance
```

```
netperf -H 192.168.100.1 -p 11260 -t UDP_STREAM -l 300 -- -m 1400 #fifth ECS instance
netperf -H 192.168.100.1 -p 11261 -t UDP_STREAM -l 300 -- -m 1400 #sixth ECS instance
netperf -H 192.168.100.1 -p 11262 -t UDP_STREAM -l 300 -- -m 1400 #seventh ECS instance
netperf -H 192.168.100.1 -p 11263 -t UDP_STREAM -l 300 -- -m 1400 #eighth ECS instance
```

Test the outbound direction

Start the netserver process on the eight ECS instances in the VPC and specify the port as follows:

```
netserver -p 11256
```

Start eight netperf processes on the network access device of the on-premises IDC and specify different IP addresses.

```
netperf -H 172.16.0.2 -p 11256 -t UDP_STREAM -l 300 -- -m 1 #first ECS instance
netperf -H 172.16.0.3 -p 11256 -t UDP_STREAM -l 300 -- -m 1 #second ECS instance
netperf -H 172.16.0.4 -p 11256 -t UDP_STREAM -l 300 -- -m 1 #third ECS instance
netperf -H 172.16.0.5 -p 11256 -t UDP_STREAM -l 300 -- -m 1 #fourth ECS instance
netperf -H 172.16.0.6 -p 11256 -t UDP_STREAM -l 300 -- -m 1 #fifth ECS instance
netperf -H 172.16.0.7 -p 11256 -t UDP_STREAM -l 300 -- -m 1 #sixth ECS instance
netperf -H 172.16.0.8 -p 11256 -t UDP_STREAM -l 300 -- -m 1 #seventh ECS instance
netperf -H 172.16.0.9 -p 11256 -t UDP_STREAM -l 300 -- -m 1 #eighth ECS instance
```

3. If you want to test bps, change the preceding commands to:

```
netperf -H 192.168.100.1 -p 11256 -t UDP_STREAM -l 300 -- -m 1400 #first ECS instance
netperf -H 192.168.100.1 -p 11257 -t UDP_STREAM -l 300 -- -m 1400 #second ECS instance
netperf -H 192.168.100.1 -p 11258 -t UDP_STREAM -l 300 -- -m 1400 #third ECS instance
netperf -H 192.168.100.1 -p 11259 -t UDP_STREAM -l 300 -- -m 1400 #fourth ECS instance
netperf -H 192.168.100.1 -p 11260 -t UDP_STREAM -l 300 -- -m 1400 #fifth ECS instance
netperf -H 192.168.100.1 -p 11261 -t UDP_STREAM -l 300 -- -m 1400 #sixth ECS instance
netperf -H 192.168.100.1 -p 11262 -t UDP_STREAM -l 300 -- -m 1400 #seventh ECS instance
netperf -H 192.168.100.1 -p 11263 -t UDP_STREAM -l 300 -- -m 1400 #eighth ECS instance
```

Analyze test results

The following results are displayed when netperf processes on the client side are completed.

```
  Socket  Message  Elapsed      Messages
Size Size Time Okay Errors Throughput
bytes bytes secs # # 10^6bits/sec

124928 1 10.00 4532554 0 3.63
```

```
212992 10.00 1099999 0.88
```

Descriptions of fields in the test results are shown in the following table:

Field	Description
Socket Size	Buffer size
Message Size	Packet size (Byte)
Elapsed Time	The duration of test (s)
Message Okay	The number of packets successfully sent out
Message Errors	The number of packets fail to be sent out
Throughput	Network throughput (Mbit/s)

You can obtain the pps of the tested link if you divide the number of packets successfully sent out by the duration of test. That is, pps = the number of packets successfully sent out / the duration of test

Use iperf3 to test the bandwidth of the leased line

Main parameters of iperf3 are shown in the following table.

Tool name	Main parameter	Description
iperf3	-s	Indicates receiving packets as the server.
	-i	Interval between every two reports, in seconds.
	-p	The listening port of the server.
	-u	Indicates using the UDP protocol to send packets. If this parameter is not specified, the TCP protocol is used.
	-l	Indicates the length of the read-write buffer. We recommend that you set the value to 16 when testing the packet forwarding performance and to 1400 when testing the bandwidth.
	-b	The bandwidth used by the UDP mode, in bits/s.
	-t	Set the duration of transmission. In the specified time period, iperf repeatedly sends packets of specified

		length. The default value is 10 seconds.
	-A	CPU affinity. You can bind an iperf3 process to the logic CPU of the corresponding number to avoid cross-CPU scheduling of the iperf3 process.

Test the inbound direction

1. Start the iperf3 process in the server mode on the network access device of the on-premises IDC and specify different ports as follows:

```
iperf3 -s -i 1 -p 16001
iperf3 -s -i 1 -p 16002
iperf3 -s -i 1 -p 16003
iperf3 -s -i 1 -p 16004
iperf3 -s -i 1 -p 16005
iperf3 -s -i 1 -p 16006
iperf3 -s -i 1 -p 16007
iperf3 -s -i 1 -p 16008
```

1. Start the iperf3 process in the client mode on the eight ECS instances in the VPC and specify different ports connecting to the network access device of the on-premises IDC.

```
iperf3 -u -l 16 -b 100m -t 120 -c 192.168.100.1 -i 1 -p 16001 -A 1
iperf3 -u -l 16 -b 100m -t 120 -c 192.168.100.1 -i 1 -p 16002 -A 2
iperf3 -u -l 16 -b 100m -t 120 -c 192.168.100.1 -i 1 -p 16003 -A 3
iperf3 -u -l 16 -b 100m -t 120 -c 192.168.100.1 -i 1 -p 16004 -A 4
iperf3 -u -l 16 -b 100m -t 120 -c 192.168.100.1 -i 1 -p 16005 -A 5
iperf3 -u -l 16 -b 100m -t 120 -c 192.168.100.1 -i 1 -p 16006 -A 6
iperf3 -u -l 16 -b 100m -t 120 -c 192.168.100.1 -i 1 -p 16007 -A 7
iperf3 -u -l 16 -b 100m -t 120 -c 192.168.100.1 -i 1 -p 16008 -A 8
```

Test the outbound direction

Start the iperf3 process in the server mode on each ECS instance in the VPC and specify the port:

```
iperf3 -s -i 1 -p 16001
```

Start eight iperf3 processes in the client mode on the network access device of the on-premises IDC and the value of -cis the IP address of each ECS instance.

```
iperf3 -u -l 16 -b 100m -t 120 -c 172.16.0.2 -i 1 -p 16001 -A 1
iperf3 -u -l 16 -b 100m -t 120 -c 172.16.0.3 -i 1 -p 16001 -A 2
iperf3 -u -l 16 -b 100m -t 120 -c 172.16.0.4 -i 1 -p 16001 -A 3
iperf3 -u -l 16 -b 100m -t 120 -c 172.16.0.5 -i 1 -p 16001 -A 4
iperf3 -u -l 16 -b 100m -t 120 -c 172.16.0.6 -i 1 -p 16001 -A 5
iperf3 -u -l 16 -b 100m -t 120 -c 172.16.0.7 -i 1 -p 16001 -A 6
iperf3 -u -l 16 -b 100m -t 120 -c 172.16.0.8 -i 1 -p 16001 -A 7
iperf3 -u -l 16 -b 100m -t 120 -c 172.16.0.9 -i 1 -p 16001 -A 8
```

Analyze test results

The following results are displayed when iperf3 processes on the client side are completed.

```
[ ID]  Interval       Transfer    Bandwidth      Jitter    Lost/Total Datagrams
[ 4] 0.00-10.00 sec 237 MBytes 199 Mbits/sec 0.027 ms 500/30352 (1.6%)
[ 4] Sent 30352 datagrams
```

Descriptions of fields in the rest results are shown in the following table:

Field	Description
Transfer	The total number of data transmitted
Bandwidth	Bandwidth
Jitter	Jitter
Lost/Total Datagrams	The number of dropped packets / The total number of packets (packet loss)

PPS = The number of packets received by the peer end/Duration

> **Note:** We recommend that you run the sar command on the server side to count the packets actually received and use the obtained value as the actual result, such as sar -n DEV 1 320.

Alibaba side speed limit

In addition to limits on the leased line, the following are limits on the communication between the VPC and the on-premises IDC:

The maximum read-write speed of OSS is 5 Gbit/s.

To improve the reliability, the speed of a single hash stream from the VPC to the VBR is limited to “Express Connect bandwidth / 12”. For example, if the bandwidth from the VBR to the VPC is large1, namely 1 Gbps, the maximum bandwidth of a single hash stream is 83 Mbps.

Hash stream: The data stream that is defined by the combination of the source IP address, source port, transport layer protocol, destination IP address, and destination port. For example, “192.168.1.1 10000 TCP 121.14.88.76 80” forms a hash stream. A terminal whose IP address is 192.168.1.1 is connected to port 80 of a terminal whose IP address is 121.14.88.76 through port 10000 by using the TCP protocol.

You can use a leased line that is already connected to an access point of Alibaba Cloud to connect multiple VPCs.

Note: Now a leased can be connected to 5 VPCs at most. You can open a ticket to increase the quota.

Scenario

A company has opened account A on Alibaba Cloud and created VPC-A. Account A already opened a leased line that connects the on-premises IDC of the company to VPC-A. A subsidiary of the company has open account B on Alibaba Cloud and VPC-B is under account B. The subsidiary wants to connect VPC-B to the on-premises IDC.

Because there is already a leased line under account A that connects the on-premises IDC to the access point of Alibaba Could, VPC-B under account B of the subsidiary can reuse the leased line and VBR of account A. The company only needs to create a router interface for the VBR under account A and the VPC under account B respectively and connect the two interfaces, as shown in the following figure.

This tutorial uses the case as the example to illustrate how to reuse a leased line to connect multiple VPCs. In this tutorial, VPC and leased line configurations are as follows:

Account A	Account B
Account ID: 12345678	Account ID: 87654321
VPC - Name: VPC-A - Region: China North 2 (Beijing) - VPC ID: vpc-12345678 - CIDR block: 10.10.0.0/16	VPC - Name: VPC-B - Region: China East 1(Hangzhou) - VPC ID: vpc-87654321 - CIDR block: 192.168.0.0/16
Physical Connection - VBR name: VPC-Beijing - VBR ID: vbr-12345678 - Leased line ID: pc-AAA - VLAN ID: 1000	-

Step 1: Create router interfaces

Create a router interface on the VBR under account A and the VPC under account B respectively, so that the VRouter of the VPC and the VBR can send messages to each other through the router interfaces. For more information, see Router interfaces.

Note: The router interface on the VBR must act as the initiator.

Create the initiator router interface

Follow these steps to create a router interface for the VBR:

Use account A to log on to the Express Connect console.

In the left-side navigation pane, select **Router Interface**. Click **Create Router Interface**.

Configure the router interface. This tutorial uses the following configurations.

Scenario: **Physical Access**.

Router Creation: **Create Initiator**.

Router Type: **VRouter**

Local Region: **China North 2 (Beijing)**.

Access Point: **Beijing-Daxing-A**.

VBR ID: vbr-12345678.

Peer Region: **China East 1 (Hangzhou)**.

Peer Router Type: **VRouter**.

Specification: **Large.1(1Gb)**.

Click **Buy Now** to complete the creation.

Go back to the **Router Interface** page after about one minute and select the target region. Then you can see the newly created router interface under account A. In this tutorial, the ID

of the router interface under account A is ri-AAA.

Create the receiver router interface

Follow these steps to create the receiver router interface:

Use account B to log on to the Express Connect console.

In the left-side navigation pane, select **VPC Connection > Router Interface**.

Click **Create Router Interface**.

Configure the router interface. This tutorial uses the following configurations.

Billing Method: **Pay-As-You-Go**.

Scenario: **Physical Access**.

Router Creation: **Create Receiver**.

Router Type: **VRouter**

Local Region: **China East 1 (Hangzhou)**.

VPC ID: vpc-87654321.

Peer Region: **China North 2 (Beijing)**.

Peer Access Point: **Beijing-Daxing-A**.

Peer Router Type: **VRouter**

Click **Buy Now**.

Go back to the **Router Interface** page after about one minute and select the target region. Then you can see the newly created router interface under account B. In this tutorial, the ID of the router interface under account B is ri-BBB.

Step 2: Initiate a connection

After creating router interfaces, you need to add peer router interfaces and initiate a connection. Only the initiator router interface can initiate a connection.

Add peer router interface for the VPC under account B

Use account B to log on to the Express Connect console.

In the left-side navigation pane, select **Router Interface**.

Click the region where the target router interface is located and find the target router interface.

Click **Add** in the **Peer Router Interface** column or click **More** > **Edit Peer Interface** in the **Actions** column.

In the displayed dialog box, select **Other Account** and enter the account ID (12345678), VBR ID (vbr-AAA), and router interface ID (ri-AAA) of account A.

Add peer router interface for the router interface on the VBR under account A and initiate a connection

Follow these steps to add peer router interface for the router interface on the VBR under account A and initiate a connection:

Use account A to log on to the Express Connect console.

In the left-side navigation pane, select **Router Interface**.

Click the region where the target router interface is located and find the target router interface.

Click **Add** in the **Peer Router Interface** column or click **More** > **Edit Peer Interface** in the **Actions** column.

In the displayed dialog box, select **Other Account** and enter the account ID (87654321), VBR ID (vbr-BBB), and router interface ID (ri-BBB) of account B.

Find the router interface of the VBR, and click **Initiate a Connection**.

The connection is established successfully when the status of the router interfaces ri-AAA and ri-BBB changes to **Active**.

Step 3: Configure routes

After creating the router interfaces, you need to configure routes so that the on-premises IDC can communicate with the VPC.

Configure routes on the VBR

Follow these steps to forward the traffic, destined for the on-premises IDC (CIDR block: 172.16.0.0/12), from the VBR to the leased line:

Use account A to log on to the Express Connect console.

In the left-side navigation pane, click **Virtual Border Router**.

Find the target VBR and click **Manage**. Then click **Add Route Entry** on the page of VBR details.

Configure the route. In this tutorial, the route configurations are as follows:

Destination CIDR Block: The CIDR block of the on-premises IDC. In this tutorial, enter 172.16.0.0/12.

Next Hop Direction: Select **To VPC**.

Next Hop: Select the existing leased line.

Follow these steps to forward the traffic, destined for the VPC (CIDR block: 192.168.0.0/16), from the VBR to the VPC:

Use account A to log on to the Express Connect console.

In the left-side navigation pane, click **Virtual Border Router**.

Find the target VBR and click **Manage**. Then click **Add Route Entry** on the page of VBR details.

Configure the route. In this tutorial, the route configurations are as follows:

Destination CIDR Block: The CIDR block of the peer VPC. In this tutorial, enter 192.168.0.0/16.

Next Hop Direction: Select **To VPC**.

Next Hop: Select the router interface of the VBR. In this tutorial, select ri-BBB.

Configure the route on the VPC

Follow these steps to forward the traffic, destined for the on-premises IDC (CIDR block: 172.16.0.0/12), from the VPC to the VBR:

Use account B to log on to the Express Connect console.

In the left-side navigation pane, click **Router Interface**. Find the target router interface and click **Router Configuration**.

Configure the route. In this tutorial, the route configurations are as follows:

Destination CIDR Block: The CIDR block of the on-premises IDC. In this tutorial, enter 172.16.0.0/12.

Next Hop Type: Select **Router Interface**.

Next Hop: Select the router interface of VPC-B. In this tutorial, select ri-AAA.

Configure the route on the on-premises IDC

Till now, the route configuration on Alibaba Cloud has been completed. You still need to add a route entry for the VPC CIDR block in the physical access device of the customer. The destination CIDR block is the Alibaba Cloud-side IP address. For example:

```
ip route 172.16.0.0/12 10.100.0.1
```

You can also configure BGP dynamic routing to direct traffic to the VBR:

Create BGP peer groups. For more information, see Manage BGP peer groups.

Add BGP peers to the BGP groups, see Manage BGP peers.

Advertise BGP network, see Advertise BGP network.

Note: Make sure the destination CIDR block of the BGP routing is the same as that of the static route. In this tutorial, it is 192.168.0.0/16.

Till now, all configurations have been completed.

Step 5: Performance test

After the two networks are connected with each other, test the speed of the leased line to ensure it can meet service needs. For more information, see Test the network performance of a physical connection.