

# Alibaba Cloud Express Connect

## Best Practices

Issue: 20190812

# Legal disclaimer

---

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company, or individual in any form or by any means without the prior written consent of Alibaba Cloud.
3. The content of this document may be changed due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, special, or punitive damages, including lost profits arising from the use

or trust in this document, even if Alibaba Cloud has been notified of the possibility of such a loss.

5. By law, all the content of the Alibaba Cloud website, including but not limited to works, products, images, archives, information, materials, website architecture, website graphic layout, and webpage design, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of the Alibaba Cloud website, product programs, or content shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates).
6. Please contact Alibaba Cloud directly if you discover any errors in this document.



## Generic conventions

Table -1: Style conventions

Style	Description	Example
	This warning information indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	 <b>Danger:</b> Resetting will result in the loss of user configuration data.
	This warning information indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	 <b>Warning:</b> Restarting will cause business interruption. About 10 minutes are required to restore business.
	This indicates warning information, supplementary instructions, and other content that the user must understand.	 <b>Notice:</b> Take the necessary precautions to save exported data containing sensitive information.
	This indicates supplemental instructions, best practices, tips, and other content that is good to know for the user.	 <b>Note:</b> You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
<b>Bold</b>	It is used for buttons, menus, page names, and other UI elements.	Click OK.
Courier font	It is used for commands.	Run the <code>cd / d C :/ windows</code> command to enter the Windows system folder.
<i>Italics</i>	It is used for parameters and variables.	<code>bae log list --instanceid Instance_ID</code>
[ ] or [a b]	It indicates that it is an optional value, and only one item can be selected.	<code>ipconfig [-all -t]</code>

Style	Description	Example
<code>{}</code> or <code>{a b}</code>	It indicates that it is a required value, and only one item can be selected.	<code>switch {stand   slave}</code>



# Contents

---

Legal disclaimer.....	I
Generic conventions.....	I
1 Access cloud services through physical connections.....	1
2 Test the network performance of a physical connection.....	3
3 Use Express Connect to connect your network to Alibaba Cloud.....	12
4 Connect a VPC to an on-premises data center.....	16
5 Migrate peering connections to CEN.....	20
5.1 Migrate a VPC in a peering connection to a CEN instance.....	20
5.2 Migrate a VBR in a peering connection to a CEN instance.....	25
5.3 Roll back the migration.....	30
6 Implement redundancy for physical connections.....	31



# 1 Access cloud services through physical connections

---

## AnyTunnel VIP

AnyTunnel VIP belongs to 100.64.0.0/10 of each VPC. DNS, YUM, NTP, OSS, SLS, and other cloud services in VPCs are all using IP addresses that belong 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 data center, you must set the router interface pointing to the 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 data center.



### Note:

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 two route entries on the VBR.

## Configure the route on the VBR

1. Log on to the [Express Connect console](#).
2. In the left-side navigation pane, choose Physical Connections > Virtual Border Routers (VBRs).
3. Find the target VBR and click Manage in the Actions column.
4. On the VBR Details page, click Add Route Entry and configure the route entry. The following configurations are used in this topic:
  - Destination CIDR Block: Enter 100.64.0.0/11 and 100.96.0.0/11 respectively.
  - Next Hop Direction: To VPC.
  - Next Hop: Select the exit for data packets. In this topic, select the router interface on the VBR.
5. 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 }
```

## 2 Test the network performance of a physical connection

---

After a physical connection is established, you need to test the performance of the physical connection to ensure that it can meet your service needs.

### Prerequisites

Before the test, make sure that you have made the following preparations:

- Configurations regarding the physical connection and routes are completed. The on-premises data center is connected to the VPC through a leased line.
- A network access device for the on-premises data center is prepared. The network access device is subjected to a stress test to measure the packets per second (pps) of the on-premises data center. It serves as the client or server in the Netperf or iPerf3 test.

In this topic, the IP address of the network device for the on-premises data center is 192.168.100.1.

- Eight VPC-type ECS instances are created. 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 data center to transmit test configurations and test results.

For the eight ECS instances, the specification is ecs.se1.2xlarge, the image is centos\_7\_2\_64\_40G\_base\_20170222.vhd, and the IP address range is from 172.16.0.2 to 172.16.0.9.

### Build the test environment

#### Install Netperf

Netperf is a tool used for testing network performance and is targeted for TCP or UDP transmission.

Follow these steps to install Netperf on the network access device of the on-premises data center and the eight ECS instances.

**1. 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
```

**2. 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
```

**3. Run the `netperf -h` and `netserver -h` commands to verify if the installation is successful.****Install iPerf3**

iPerf3 is a tool used 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 data center and the eight ECS instances.

**1. Run the following command to download iPerf3:**

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

**2. 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
```

**3. 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 data center 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 physical connection

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

Tool name	Main parameter	Description
netserver (server side: receiving side tool)	-p	The port of the server.
netperf (client side: sending side tool)	-H	The IP address of the network access device of the on-premises data center or the VPC server.
	-p	The port of the network access device of the on-premises data center or the VPC server.
	-l	The running duration.
	-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

1. Start the netserver process on the network access device of the on-premises data center 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
```

2. 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 data center.

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

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

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

Test the outbound direction

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

```
netserver -p 11256
```

2. Start eight netperf processes on the network access device of the on-premises data center and specify different IP addresses:

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

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

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

### Analyze the test result

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

	Socket Size bytes	Message Size bytes	Elapsed Time secs	Messages Okay #	Errors #	Throughput 10 ^ 6bits / sec
63	124928	1	10 . 00	4532554	0	3 .

212992

10 . 00

1099999

0 . 88

The fields in the test result are described in the following table:

Field	Description
Socket Size	The buffer size.
Message Size	The packet size. Unit: Byte
Elapsed Time	The duration of the test. Unit: seconds
Message Okay	The number of packets successfully sent out.
Message Errors	The number of packets that fail to be sent out.
Throughput	The network throughput. Unit: 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 the test. That is,  $\text{pps} = \frac{\text{the number of packets successfully sent out}}{\text{the duration of the test}}$ .

Use iPerf3 to test the bandwidth of the physical connection

Main parameters of iPerf3 are described in the following table:

Tool name	Main parameter	Description
iPerf3	-s	Indicates receiving packets as the server.
	-i	The interval between every two reports. Unit: 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.



Tool name	Main parameter	Description
	-l	Indicates the length of the read/write buffer. We recommend that you set the value to 16 when you test the packet forwarding performance and to 1400 when you test the bandwidth.
	-b	The bandwidth used in the UDP mode. Unit: bit/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 associate an iperf3 process with 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 data center 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
```

2. 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 data center:

```
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

1. 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
```

2. Start eight iPerf3 processes in the client mode on the network access device of the on-premises data center and the value of `-c` is 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 the test result

The following result is 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

```

The fields in the test results are described in the following table:

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

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, for example, `sar -n DEV 1 320`.

### Alibaba Cloud-side speed limit

In addition to limits on the physical connection, the following are limits on the communication between the VPC and the on-premises data center:

- 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 one twelfth of the Express Connect bandwidth. For example, if the bandwidth from the VBR to the VPC is large1, namely 1 Gbit/s, the maximum bandwidth of a single hash stream is 85 Mbit/s.

**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.

## 3 Use Express Connect to connect your network to Alibaba Cloud

---

### Overview

Express Connect allows you to implement a highly reliable intranet communication between your on-premises data center and Alibaba Cloud VPC, and between VPCs that are based in different regions. Specifically, Express Connect provides the following two key features:

- VPC-to-VPC connection

Express Connect supports intranet communication between two VPCs regardless of the regions they are based or the accounts under which they are billed. The connection between two VPCs that are based in the same region is provided free of charge, whereas the connection between two VPCs that are based in different regions incurs fees.

To implement a connection between VPCs, Alibaba Cloud creates a [Router Interface \(RI\)](#) on the VRouter of each of the two VPCs, and uses its own backbone transmission network to achieve secure, reliable, and fast communication between the two VPCs. For more information, see [VPC interconnection](#).

- Physical connection

You can use a leased line to physically connect your on-premises data center to Alibaba Cloud. After that, you can create a VBR and RIs to achieve communication between the on-premises data center and a VPC in Alibaba Cloud. For more information, see [Connect an on-premises data center to a VPC through a physical connection](#).

### Access points

If you use a leased line to connect an on-premises data center to an Alibaba Cloud VPC, you only need to select an access point that is closest in geographic proximity to your on-premises data center. You do not need to build a physical connection between your on-premises data center and Alibaba Cloud VPC. Access points include Alibaba Cloud access points and access points provided by Alibaba Cloud partners.

- Alibaba Cloud access points

You can view all [access points](#) of Alibaba Cloud in the Express Connect console. If an access point is available in the city where your on-premises data center is located, you can directly select this access point for the leased line connection.

Region	China North 1 (Qingdao)	China North 2 (Beijing)	China North 3 (Zhangjiakou)	China North 5 (Hohhot)	China East 1 (Hangzhou)
	China East 2 (Shanghai)	China South 1 (Shenzhen)	Hong Kong	Singapore	Australia (Sydney)
	Malaysia (Kuala Lumpur)	Indonesia (Jakarta)	Japan (Tokyo)	India (Mumbai)	US (Silicon Valley)
	US (Virginia)	Germany (Frankfurt)	UK(London)	UAE (Dubai)	
SP	China Unicom	China Telecom	China Mobile	Others	
Access Point	Hangzhou-Yuhang-A- Ali	Hangzhou-Linan-A- HuatongCloud	Hangzhou-Xiaoshan- A-CU	Hangzhou-Jiangan- B-21vianet	Hangzhou-Deqing-A- CU
Port Specification	1G and below	10G	The fee charged for renting resources changes based on the specification of a port. Apply for a port as required.		
Port Type	100Base-T	1000Base-LX			
Redundant Connection ID	None				

- Access points of Alibaba Cloud partners

Some access points are also provided by Alibaba Cloud partners, and are connected with Alibaba Cloud through dedicated physical lines. This means that if no Alibaba Cloud access point is available, you can select an access point provided by one of our partners. To obtain the information about your selected physical connection, contact the [Alibaba Cloud partner](#).

As a best practice, if no Alibaba Cloud access point or access point provided by an Alibaba Cloud partner is available directly in the city where your on-premises data center is located, we recommend that you select an access point that is nearest to the city where your on-premises data center is based.

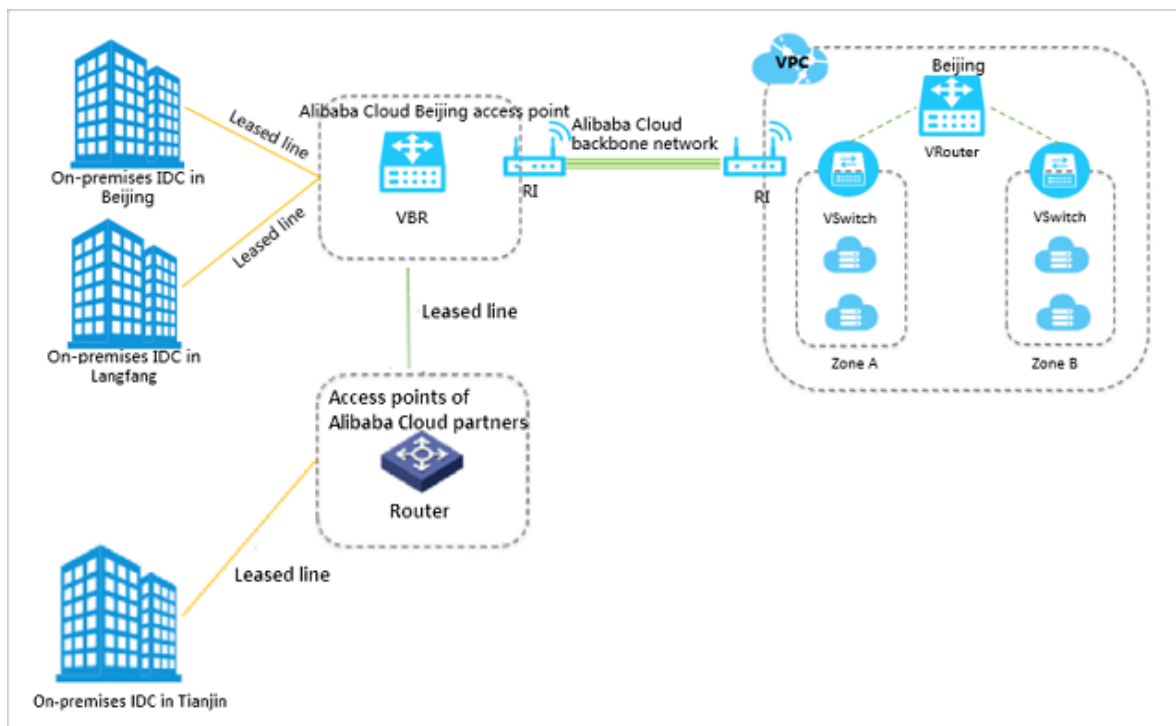
For example, the following figure shows two on-premises data center centers (one in Beijing, one in Langfang) connected to an Alibaba Cloud access point based

in Beijing. Additionally, the figure shows an on-premises data center located in Tianjin that is connected to Alibaba Cloud VPC through a leased line connected to the access point of an Alibaba Cloud partner.



**Note:**

The lines in yellow are leased lines that need to be installed by your service provider.



### Access global resources from one point

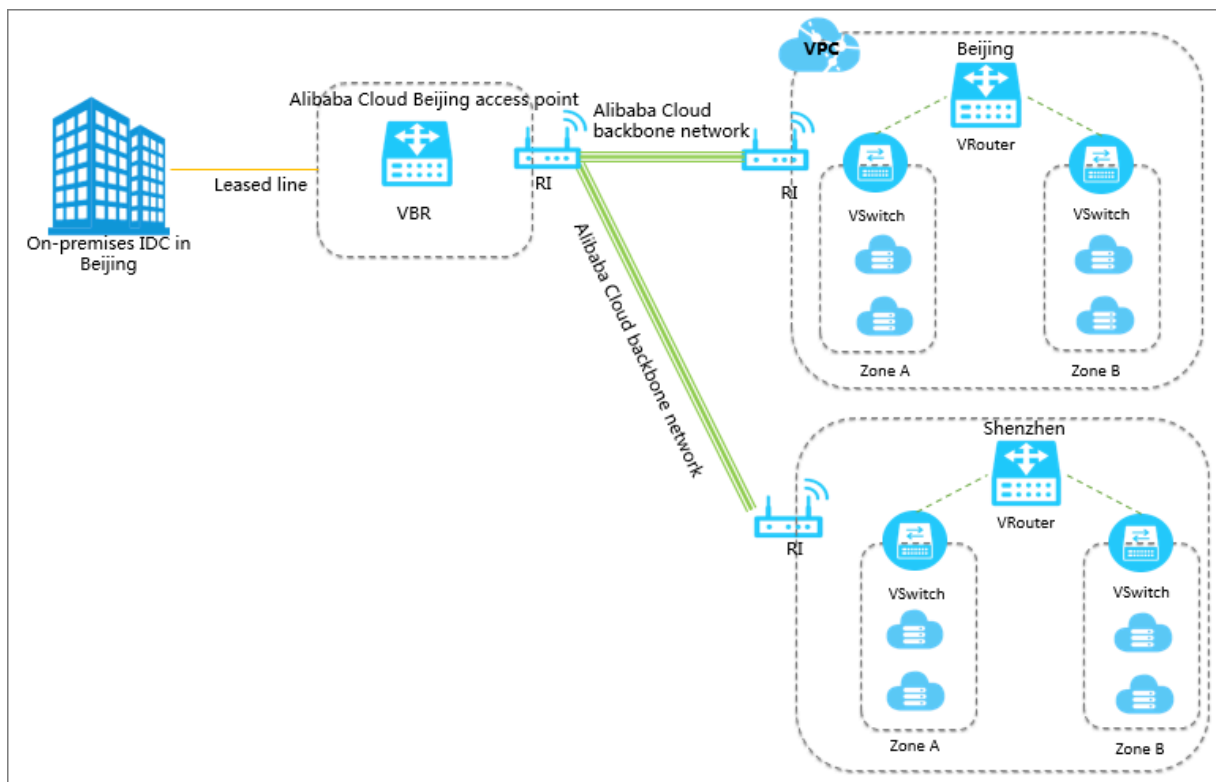
By connecting to any one access point, you can connect your resources to Alibaba Cloud VPCs around the globe through the access point.

For example, you want to connect an on-premises data center in Beijing to a VPC in Beijing and a VPC in Shenzhen through a physical connection. To implement that, you only need to use a leased line to connect the on-premises data center to an Alibaba Cloud Beijing access point, and create two RIs respectively connecting to the two VPCs on the VBR.



**Note:**

In the following figure, only the lines in yellow need to be installed by your service provider.




## 4 Connect a VPC to an on-premises data center

You can connect an on-premises data center to a VPC by using VPN Gateway, a physical connection of Express Connect, or Smart Access Gateway to build a hybrid cloud.

### Overview

You can establish intranet communication between a local data center and Alibaba Cloud to build a hybrid cloud. Then you can seamlessly expand your local IT infrastructure to Alibaba Cloud to cope with service fluctuation and improve application stability by right of the mass computing, storage, network, and CDN resources of Alibaba Cloud.

You can use VPN Gateway, a physical connection of Express Connect and Smart Access Gateway to connect a local data center to a VPC. In addition, you can interconnect global networks by using CEN.





## Solutions

Solution	Description
VPN Gateway	<p>You can use IPsec-VPN to connect a local data center to a VPC. VPN Gateway contains two different gateway instances which form active/standby hot backup. The traffic is automatically distributed to the standby node when the active node fails.</p> <p>The VPN Gateway is based on Internet communication , so its network latency and availability are decided by the Internet. If you do not have a particularly high demand for network latency, we recommend that you use VPN Gateway.</p> <p>For more information, see <a href="#">Establish a connection between a VPC and an on-premises data center</a>.</p>
Physical connection	<p>You can use a leased line of your service provider to establish a physical connection between your on-premises IDC and an Alibaba Cloud access point.</p> <p>Physical connection features good network quality and large bandwidth. Therefore, if your priority is good network quality, we recommend that you select physical connection.</p> <p>For more information, see <a href="#">Connect a local data center to a VPC through a physical connection</a>.</p>

Solution	Description
Redundant physical connections	<p>You can use redundant physical connections to connect your on-premises data center to a VPC. Redundant physical connections provide high-quality and high-reliability intranet communication between your local data center and Alibaba Cloud. Alibaba Cloud supports up to four physical connections to achieve Equal-Cost Multipath Routing (ECMP).</p> <p>For more information, see <a href="#">Create redundant physical connections</a>.</p>
Smart Access Gateway	<p>Smart Access Gateway (SAG) is an all-in-one solution for connecting local branches of an enterprise to the Alibaba Cloud. With Smart Access Gateway, enterprises can access Alibaba Cloud through the Internet using a fully encrypted connection, which is more intelligent, more reliable, and more secure.</p> <p>Smart Access Gateway is an easy-to-configure and low-cost service. If you want to connect multiple local branches of an enterprise to the cloud, we recommend that you select Smart Access Gateway.</p> <p>For more information, see <a href="#">Connect local branches to Alibaba Cloud through Smart Access Gateway</a>.</p>
BGP active/standby links	<p>Function by using both a physical connection and CEN, allowing you to connect an on-premises data center to VPCs in different regions through active/standby links.</p> <p>For more information, see <a href="#">Connect a local data center to Alibaba Cloud by using BGP active/standby links</a>.</p>

Solution	Description
Physical connection + Smart Access Gateway	<p>A solution using Smart Access Gateway as the backup link of the existing physical connection to build a reliable and high-availability hybrid cloud.</p> <p>For more information, see <a href="#">Tutorial for configuring Smart Access Gateway as the backup of a physical connection</a>.</p>

## 5 Migrate peering connections to CEN

### 5.1 Migrate a VPC in a peering connection to a CEN instance

This topic describes how to migrate a VPC that uses a peering connection to a Cloud Enterprise Network (CEN) instance. By using CEN, you can build private network communication channels between VPCs or between VPCs and on-premises data centers. CEN uses automatic route distribution and learning, which can improve the network convergence and the quality and security of cross-network communication, and achieve the interconnection of all network resources.



#### Warning:

After you migrate a VPC to a CEN instance, do not freeze or delete the same-region peering connections that belong to the China (Hangzhou) or China (Shanghai) regions.

#### Prerequisites

If you want to use an existing CEN instance, make sure that the overlapping routing function is enabled.



#### Note:

If the overlapping routing function is not enabled for the target CEN instance, enable the function first.

**CEN**[Get Started](#)[Documentation](#)

**Basic Settings**

ID: cer-  
Name: test\_11 [Edit](#)  
Description: - [Edit](#)

Status: **Ready**  

Overlapping Routing: Disabled [Enable](#)

Function

**Networks**

Bandwidth PackagesRegion ConnectionsRoutesPrivateZone

[Attach Network](#)

[Refresh](#)

Instance ID/Name	Region	Network Type	Account ID	Attach Time	Status	Actions
vp- vp- c0- 64-	China (Hangzhou)	VPC	5- 8	02/18/2019, 13:56:00	● Attached	<a href="#">Detach</a>

Contact Us

## Procedure

To migrate a VPC in a peering connection to a CEN instance, follow these steps:



### Note:

Make sure that you have made the necessary preparations before migration.

1. Log on to the [CEN console](#).
2. On the Instances page, find the target CEN instance and click the instance ID.
3. On the Networks tab, click Attach Network and add the VPC to be migrated. For more information, see [Networks](#).

Attach Network

Your Account Different Account

Note: You cannot attach networks that are already attached to the CEN instance.

• Network Type ?  
VPC

• Region ?  
China (Hangzhou)

• Networks ?  
-/vpc-

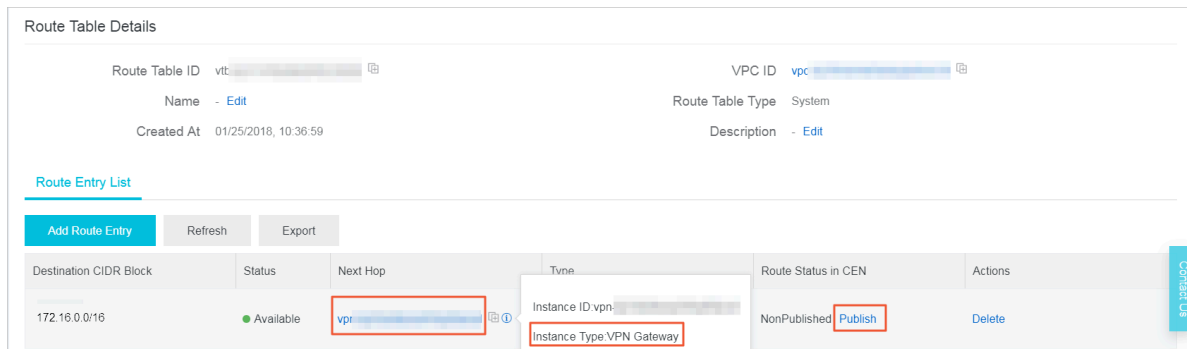
OK Cancel

Contact Us

4. If you need the VPC to communicate with other resources that belong to different regions, you need to buy a bandwidth package and set an intranet communication bandwidth value.

For more information, see [Set a cross-region connection bandwidth](#).

5. If you have added routes destined for ECS instances, VPN Gateways, or High-Availability Virtual IP Addresses (HaVips) in the VPC, you need to publish these routes to the CEN instance.



6. Log on to the [CEN console](#), click the ID of the target CEN instance, and on the Routes tab, check the routes. Make sure that the routes do not conflict with each other after you add the VPC to the CEN instance.

The static route configured for the peering connection takes precedence over the dynamic route of the CEN instance. Specifically, if a static route is configured for the peering connection, no CEN route that is more detailed than or the same as the static route is allowed to be learnt by the CEN instance. In this case, we

recommend that you divide a large route segment into smaller route segments and delete these routes after CEN learns the routes to ensure smooth migration.

For example, the CEN route 172.16.1.0/24 in the following figure is more detailed than the route 172.16.0.0/16 configured for the peering connection, which constitutes a route conflict.

The screenshot shows the AWS VPC console 'Routes' tab for a VPC in China (Shanghai). The table lists four routes. The third route, with destination CIDR block 172.16.0.0/16, is of type 'Custom' and is active, with 'ExpressConnect' as the next hop. The fourth route, with destination CIDR block 172.16.1.0/24, is of type 'CEN' and has a status of 'Rejected', with 'China (Qingdao)' as the next hop. A red box highlights these two routes, indicating a conflict.

Destination CIDR Block	Publish Status	Type	Status	Next Hop
10.0.0.0/8	(value, select, Published (Published NonPublished (NonPublished) other {}))	CEN	Active	China (Beijing)
100.64.0.0/10	(value, select, Published (Published NonPublished (NonPublished) other {}))	System	Active	—
172.16.0.0/16	(value, select, Published (Published NonPublished (NonPublished) other {}))	Custom	Active	ExpressConnect
172.16.1.0/24	(value, select, Published (Published NonPublished (NonPublished) other {}))	CEN	Rejected	China (Qingdao)

- You can directly delete the route of the peering connection through the VPC console. Then, the CEN route takes effect automatically. However, this method causes intermittent disconnections.

The duration of disconnections is in proportion to the number of CEN routes. Therefore, we recommend that you use the following method to smoothly migrate the VPC for important services.

- You can divide the peering connection route 172.16.0.0/16 into two smaller route segments, 172.16.1.0/25 and 172.16.1.128/25, which are smaller than the CEN route 172.16.1.0/24.
  - a. Log on to the [VPC console](#) and find the route table to which the target peering connection route belongs.
  - b. Click Add Route Entry. Add two route entries that are respectively destined for 172.16.1.0/25 and 172.16.1.128/25 with the Express Connect route interface as the next hop type.

**Route Table**

**Route Table Details**

Route Table ID: vtb-m-...  
 Name: - Edit  
 Created At: 04/29/2019, 16:28:12

VPC ID: vpc-m-...  
 Route Table Type: System  
 Description: - Edit

[Route Entry List](#) [Associated VSwitches](#)

[Add Route Entry](#) [Refresh](#) [Export](#)

Destination CIDR Block	Status	Next Hop	Type	Route Status in CEN	Actions
172.16.1.0/25 test2	Available	ri-m-...	Custom	-	Delete
172.16.1.128/25 test1	Available	ri-m-...	Custom	-	Delete

- c. In the VPC route table, find the target peering connection route 172.16.0.0/16 and click Delete to delete this route.

**Route Table**

**Route Table Details**

Route Table ID: vtb-m-...  
 Name: - Edit  
 Created At: 04/29/2019, 16:28:12

VPC ID: vpc-m-...  
 Route Table Type: System  
 Description: - Edit

[Route Entry List](#) [Associated VSwitches](#)

[Add Route Entry](#) [Refresh](#) [Export](#)

Destination CIDR Block	Status	Next Hop	Type	Route Status in CEN	Actions
172.16.1.0/25 test2	Available	ri-m-...	Custom	-	Delete
172.16.1.128/25 test1	Available	ri-m-...	Custom	-	Delete
172.16.0.0/16	Available	ri-m-...	Custom	-	Delete

- d. Click Refresh to check if the CEN route has taken effect.

**Route Table**

**Route Table Details**

Route Table ID: vtb-m-...  
 Name: - Edit  
 Created At: 04/29/2019, 16:28:12

VPC ID: vpc-m-...  
 Route Table Type: System  
 Description: - Edit

[Route Entry List](#) [Associated VSwitches](#)

[Add Route Entry](#) [Refresh](#) [Export](#)

Destination CIDR Block	Status	Next Hop	Type	Route Status in CEN	Actions
172.16.1.0/25 test2	Available	ri-m-...	Custom	-	Delete
172.16.1.128/25 test1	Available	ri-m-...	Custom	-	Delete
172.16.0.0/24	Available	vpc-m-...	Cloud Enterprise Network	-	Delete

- e. After the CEN route takes effect, delete the added two route entries 172.16.1.0/25 and 172.16.1.128/25 to complete the smooth migration.



## 5.2 Migrate a VBR in a peering connection to a CEN instance

This topic describes how to migrate a Virtual Border Router (VBR) that uses a peering connection to a Cloud Enterprise Network (CEN) instance. By using CEN, you can build private network communication channels between VPCs or between VPCs and on-premises data centers. CEN uses automatic route distribution and learning, which can improve the network convergence and the quality and security of cross-network communication, and achieve the interconnection of all network resources.

### Prerequisites

If you want to use an existing CEN instance, make sure that the overlapping routing function is enabled.



#### Note:

If the overlapping routing function is not enabled for the target CEN instance, enable the function first.

**CEN** [Get Started](#) [Documentation](#)

**Basic Settings**

ID: cer-...-ty  
 Name: test\_11 [Edit](#)  
 Description: - [Edit](#)

Status: Ready  
 Overlapping Routing: Disabled [Enable](#)  
 Function

**Networks** | Bandwidth Packages | Region Connections | Routes | PrivateZone

[Attach Network](#) [Refresh](#)

Instance ID/Name	Region	Network Type	Account ID	Attach Time	Status	Actions
vp-...-13	China (Hangzhou)	VPC	5-...-8	02/18/2019, 13:56:00	<span style="color: green;">●</span> Attached	<a href="#">Detach</a>

[Contact Us](#)

### Procedure

To migrate a VBR in a peering connection to a CEN instance, follow these steps:



#### Note:

Make sure that you have made the necessary preparations before migration.

1. If you have enabled the health check function for the VBR, we recommend that you first disable the health check function in the Express Connect console.
2. Log on to the [CEN console](#).

3. On the Instances page, find the target CEN instance and click the instance ID.
4. On the Networks tab, click Attach Network and add the VBR and VPC to be migrated. For more information, see [Networks](#).

**Attach Network**

**Your Account** Different Account

**Note:** You cannot attach networks that are already attached to the CEN instance.

• **Network Type** ?  
Virtual Border Router (VBR)

• **Region** ?  
China (Beijing)

• **Networks** ?  
vbr-2z...

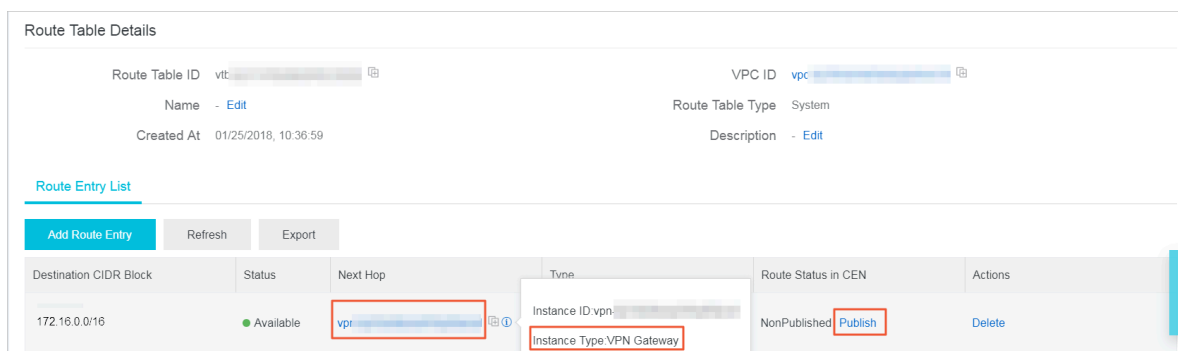
OK Cancel

Contact Us

5. If you need the VPC to communicate with other resources that belong to different regions, you need to buy a bandwidth package and set an intranet communication bandwidth value.

For more information, see [Set a cross-region connection bandwidth](#).

6. If you have added routes destined for ECS instances, VPN Gateways, or High-Availability Virtual IP Addresses (HaVips) in the VPC, you need to publish these routes to the CEN instance.



7. If an on-premises data center needs to access cloud resources, such as OSS and PrivateZone, perform the configurations through the CEN console.

For more information, see [Set PrivateZone access](#).

8. Log on to the CEN console, click the ID of the target CEN instance, and on the Routes tab, check the routes. Make sure that the routes do not conflict with each other after you add the VBR and VPC to the CEN instance.

The static route configured for the peering connection takes precedence over the dynamic route of the CEN instance. Specifically, if a static route is configured for the peering connection, no CEN route that is more detailed than or the same as the static route is allowed to be learnt by the CEN instance. In this case, we recommend that you divide a large route segment for the peering connection into

smaller route segments and delete these routes after CEN learns the routes to ensure smooth migration.

For example, the CEN route 192.168.1.0/24 in the following figure is more detailed than the route 192.168.0.0/16 configured for the peering connection, which constitutes a route conflict.

Networks Bandwidth Packages Region Connections <b>Routes</b> AnyTunnel PrivateZone				
Networks	China (Hangzhou) vbr-		Refresh	
Destination CIDR Block	Publish Status	Type	Status	Next Hop
10.0.0.0/8	{value, select, Published (Published NonPublished (NonPublished) other {-}) }	CEN	Active	China (Beijing)
100.64.0.0/10	{value, select, Published (Published NonPublished (NonPublished) other {-}) }	System	Active	—
172.16.0.0/24	{value, select, Published (Published NonPublished (NonPublished) other {-}) } Withdraw	System	Active	—
192.168.0.0/16	{value, select, Published (Published NonPublished (NonPublished) other {-}) }	Custom	Active	ExpressConnect
192.168.1.0/24	{value, select, Published (Published NonPublished (NonPublished) other {-}) }	CEN	Rejected	China (Qingdao)

- You can directly delete the route of the peering connection. Then, the CEN route takes effect automatically. However, this method causes intermittent disconnections.

The duration of disconnections is in proportion to the number of CEN routes. Therefore, we recommend that you use the following method to smoothly migrate the VPC for important services.

- You can divide the peering connection route 192.168.0.0/16 into two smaller route segments, 192.168.1.0/25 and 192.168.1.128/25, which are smaller than the CEN route 192.168.1.0/24.
  - a. Log on to the Express Connect console, find the target VBR, click the VBR ID, and then click the Routes tab.
  - b. Click Add Route. Add two routes that are respectively destined for 192.168.1.0/25 and 192.168.1.128/25 with the next hop type of VPCs.

**vbr-2**

Basic Information

VBR vbr-2

Access Point Beijing-Daxing-A

Status ● Active

Name

Created At Mar 6, 2018, 19:16:34

CEN cen- Unbind

Physical Connection Interfaces **Routes** Advertised BGP Subnets BGP Groups BGP Peers CEN Authorization Peering Connections

Add Route Refresh

Route Table ID	Destination Subnet	Status	Next Hop Instance	Next Hop Type	Route Type	CEN Publication Status	Actions
vtb-2	192.168.1.128/25	<span style="color: green;">●</span> Available	vpc-m	VPC	Custom	-	Delete
vtb-2	192.168.1.0/25	<span style="color: green;">●</span> Available	vpc-m	VPC	Custom	-	Delete

- c. For BGP routing, you need to advertise the CIDR blocks related to 192.168.1.0/25 and 192.168.1.128/25.

**vbr-2**

Basic Information

VBR vbr-2

Access Point Beijing-Daxing-A

Status ● Active

Name

Created At Mar 6, 2018, 19:16:34

CEN cen-7 Unbind

Physical Connection Interfaces Routes **Advertised BGP Subnets** BGP Groups BGP Peers CEN Authorization Peering Connections

Advertise BGP Subnet Refresh

Advertised Subnet	Actions
192.168.1.0/25	Delete
192.168.1.128/25	Delete

- d. Delete the peering connection route 192.168.0.0/16.

**vbr-2**

Basic Information

VBR vbr-2

Access Point Beijing-Daxing-A

Status ● Active

Name

Created At Mar 6, 2018, 19:16:34

CEN cen-7 Unbind

Physical Connection Interfaces **Routes** Advertised BGP Subnets BGP Groups BGP Peers CEN Authorization Peering Connections

Add Route Refresh

Route Table ID	Destination Subnet	Status	Next Hop Instance	Next Hop Type	Route Type	CEN Publication Status	Actions
vtb-9	192.168.1.128/25	<span style="color: green;">●</span> Available	vpc-m	VPC	Custom	-	Delete
vtb-9	192.168.1.0/25	<span style="color: green;">●</span> Available	vpc-m	VPC	Custom	-	Delete
vtb-9	192.168.0.0/16	<span style="color: green;">●</span> Available	vpc-m	VPC	Custom	-	Delete

- e. Click Refresh and check whether the CEN route has taken effect.

Basic Information

VBR: vbr-2  
Access Point: Beijing-Daxing-A  
Status: Active

Name: [redacted]  
Created At: Mar 6, 2018, 19:16:34  
CEN: cen-2 [redacted] Unbind

Physical Connection Interfaces | **Routes** | Advertised BGP Subnets | BGP Groups | BGP Peers | CEN Authorization | Peering Connections

Add Route Refresh

Route Table ID	Destination Subnet	Status	Next Hop Instance	Next Hop Type	Route Type	CEN Publishment Status	Actions
vtb-2-9	192.168.1.128/25	Available	vpc-m5	VPC	Custom	-	Delete
vtb-2-9	192.168.1.0/25	Available	vpc-m5	VPC	Custom	-	Delete
vtb-2-9	10.0.0.0/24	Available	vpc-m5	VPC	Custom	-	Delete
vtb-2-9	10.0.0.0/8	Available	pc-2	Physical Connection Interface	Custom	-	Delete
vtb-2-9	192.168.1.0/24	Available	vpc-m5	VPC	CEN	-	Delete

- f. Delete the two routes 192.168.1.0/25 and 192.168.1.128/25 in the VBR route table, and delete the advertised BGP routes.
- g. In the CEN console, configure health checks for the migrated VBR. For more information, see [Configure health check](#).

## 5.3 Roll back the migration

This topic describes how to roll back your migration by modifying the routes.

Rollback solutions depend on the migration methods you have adopted. The available rollback solutions are as follows:

- **Migration with intermittent disconnections:** Re-add the deleted static route of the peering connection. All the routes that are more detailed than or equals the re-added peering connection route are automatically deleted.
- **Smooth migration:** Re-add the deleted detailed routes directly.



### Note:

If the migrated Virtual Border Router (VBR) is configured with BGP routes, you need to re-advertise the related CIDR blocks.

## 6 Implement redundancy for physical connections

---

To make sure that traffic is forwarded to the standby physical connection when the active physical connection fails, you must configure health checks for the associated VBR-to-VPC peering connections and configure route weights.

### Prerequisites

Before you configure health checks and route weights, make sure that the following operations are completed:

- Two physical connection interfaces are applied for and the connection between the on-premises data center and Alibaba Cloud is established.
- Two VBR-to-VPC peering connections are created. For more information, see [Apply for a physical connection interface](#) and [Interconnect a VPC and a VBR](#).
- Static routes are configured between the Virtual Border Routers (VBRs) and the on-premises data center. No BGP is used.

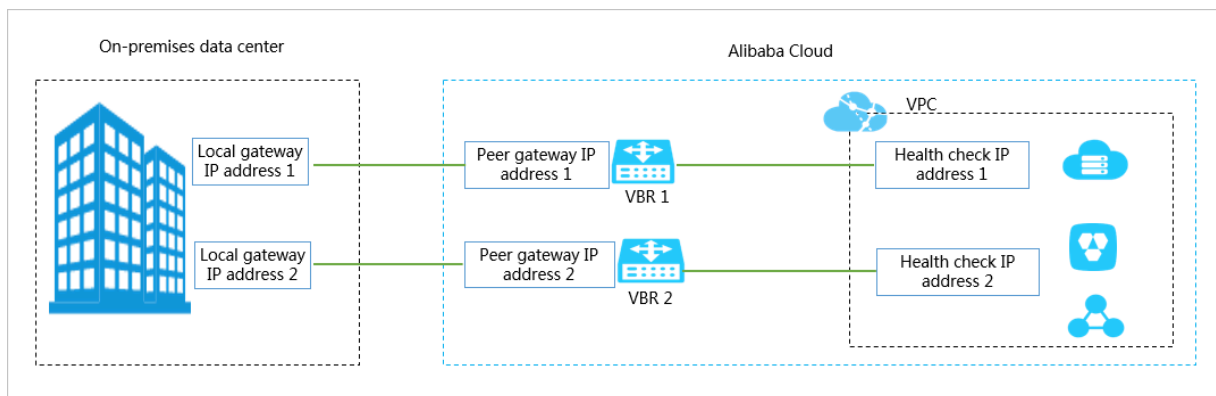
### Context

Alibaba Cloud sends a `ping` packet once every two seconds from the health check IP address to the customer-side IP address of the on-premises data center. If no response is received for the `ping` packet for eight consecutive times on one physical connection, traffic is switched to the other physical connection.



#### Note:

If Control Plane Policing (Copp) (such as Cisco devices) or Local Attack Defense Policy (Huawei devices) is configured on the on-premises data center, health check packets may be discarded and the health check link shocks. Therefore, we recommend that you cancel the speed limitation on the network device of the on-premises data center.




The network configurations are as follows:

Configuration	CIDR block
The connected VPC	192.168.0.0/16
The on-premises data center	172.16.0.0/16
The connection between one VBR and the on-premises data center	<ul style="list-style-type: none"> <li>VBR gateway IP address: 10.10.10.1</li> <li>Gateway IP address of the on-premises data center: 10.10.10.2</li> <li>Subnet mask: 255.255.255.252</li> </ul>
The connection between the other VBR and the on-premises data center	<ul style="list-style-type: none"> <li>VBR gateway IP address: 10.10.11.1</li> <li>Gateway IP address of the on-premises data center: 10.10.11.2</li> <li>Subnet mask: 255.255.255.252</li> </ul>
Health checks for one VBR-to-VPC peering connection	<ul style="list-style-type: none"> <li>Source IP address: 192.168.10.1</li> <li>Destination IP address: 10.10.10.2</li> </ul>
Health checks for the other VBR-to-VPC peering connection	<ul style="list-style-type: none"> <li>Source IP address: 192.168.10.2</li> <li>Destination IP address: 10.10.11.2</li> </ul>

### Step 1 Configure health checks

You must configure health checks for the two peering connections. To do so, follow these steps:

1. Log on to the [Express Connect](#) console.
2. In the left-side navigation pane, choose VPC Peering Connections > VBR-to-VPC.
3. Find the target peering connection and choose  > Health Check in the Actions column.



4. On the Health Check page, click Configure.
5. On the Edit VBR page, configure health checks.

The following table shows the required parameters.

Parameter	Description
Source IP	Any idle private IP address in the connected VPC.
Destination IP	<p>The interface IP address of the network device of the on-premises data center.</p> <p>If you need to perform ICMP health checks from the on-premises data center to the VPC, enter the health check IP address of the VPC as the destination IP address and configure a route that points to the new health check destination.</p>

Edit VBR

\* Source IP

192.168.10.1

Enter an unused VSwitch IP address.

\* Destination IP

10.10.10.2

Enter an interface IP address of the network equipment on the customer's data center side.

Send Packet Every (Seconds)

2

Packets Detected

8

OK

Cancel

Contact Us

6. Click OK.

7. Repeat the preceding steps to configure health checks for the other peering connection.

**Note:**

The source IP address of the health check for the other peering connection cannot be the same as that for the first peering connection.

## Step 2 Configure routes

In this example, load balancing routing is configured.

1. In the left-side navigation pane, choose Route Tables.
2. Find the target VPC and click the ID of the corresponding route table.
3. On the Route Table page, click Add Route Entry and configure the load balancing route.

Configure the load balancing route according to the following information:

- **Destination CIDR Block:** Enter the destination CIDR block to which traffic is forwarded.
- **Next Hop Type:** Select Router Interface (To VBR), which means forwarding the traffic with a destination IP address that falls into the destination CIDR block to the router interface associated with the VBR.

Select Load Balancing Routing for the routing type, select the two VBRs connected with the VPC as the next hop, and set weights for the two VBRs. Value

range of the weights of the VBRs: 1 to 255. Default value: 100. The weights of the two VBRs must be the same so that traffic can be evenly distributed to them.

**Add Route Entry**

**Name** ?  
CIDR\_block\_on-premises\_data\_center 34/128 ✓

**Destination CIDR Block**  
172 . 16 . 0 . 0 / 16 ✓

**Next Hop Type**  
Router Interface (To VBR) ✓

General Routing Active/Standby Routing **Load Balancing Routing**

vtb-b [redacted] Weight 100 + -  
vtb-b [redacted] Weight 100 + -

Add Next Hop

**i** Load balancing routing requires 2-8 router interfaces for the next hop. You must specify a weight value between 1 and 255 for each added router interface. The values you specify to these weights must be identical. Therefore, the system will distribute the traffic evenly among these router interfaces.

OK Cancel

4. Click Add Route Entry and add a route from one VBR to the on-premises data center.

Configure the route according to the following information:

- **Destination CIDR Block:** Enter the destination CIDR block.
- **Next Hop Type:** Select Router Interface (To VBR), which means forwarding the traffic with a destination IP address that falls into the destination CIDR block to the router interface associated with the VBR.

Select General Routing for the routing type and select one VBR as the next hop.

- Click Add Route Entry and configure a route from the other VBR to the on-premises data center.

Configure the route according to the following information:

- **Destination CIDR Block:** Enter the destination CIDR block.
- **Next Hop Type:** Select Router Interface (To VBR), which means forwarding the traffic with a destination IP address that falls into the destination CIDR block to the router interface associated with the VBR.

Select General Routing for the routing type and select the other VBR as the next hop.

The following figure shows the configured routes.

Route Entry List Associated VSwitches					
<div> Add Route Entry Refresh Export </div>					
Destination CIDR Block	Status	Next Hop	Type	Actions	
10.10.10.0/24 From_one_VBR_to_data_center	Available	ri-uf...	Custom	Delete	
172.16.0.0/16 CIDR_block_on-premises_data_center	Available	ri-uf... ri-uf...	Custom	Delete	
10.10.11.0/24 From_the_other_VBR_to_data_center	Available	ri-uf...	Custom	Delete	

### Step 3 Configure static routes on the network device of the on-premises data center

If no BGP is used, the following static routes need to be configured for between the on-premises data center and the VBRs on the network device of the on-premises data center:

- A route entry with the health check source IP address of one peering connection as the destination IP address and the IP address of the corresponding VBR (Alibaba Cloud-side IP address) as the next hop.
- A route entry with the health check source IP address of the other peering connection as the destination IP address and the IP address of the corresponding VBR as the next hop.

### Step 6 Test the network connectivity

Ping an instance in the VPC when one physical connection fails to check if the redundant physical connection works.

## Advertise BGP CIDR blocks

If you have configured BGP for your on-premises data center and the VBRs, the VBRs need to advertise BGP CIDR blocks.

1. Log on to the [Express Connect](#) console.
2. In the left-side navigation pane, choose Physical Connections > Virtual Border Routers (VBRs).
3. Find one of the two VBRs and click the VBR ID. On the Routes tab, click Add Route.
4. On the Add Route page, configure a route pointing to the health check source IP address.

Configure the route according to the following information:

- **Destination Subnet:** Enter the source IP address of health checks. In this example, enter 192.168.10.1/32.
- **Next Hop Type:** Select VPC. Then, select the connected VPC as the next hop.

**Add Route**

\* Destination Subnet

192.168.10.1/32

Next Hop Type

☒ VPC ☐ Physical Connection Interface

\* Next Hop

vpc-uf614ufdjl3n871accgzx/ri-uf69ejgh2atodylnx8nl9

OK Cancel

Contact Us

5. Click the Advertised BGP Subnets tab and click Advertise BGP Subnet.

6. On the Advertise BGP Subnet page, enter the source IP address of health checks.

Advertise BGP Subnet

\* Advertised Subnet

192.168.10.1/32

7. Repeat the preceding steps to advertise BGP CIDR blocks for the health check source IP address of the other VBR.