Alibaba Cloud

Elastic Compute Service Network

Document Version: 20220706

C-J Alibaba Cloud

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.

- You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloudauthorized 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 because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should 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 this document based on the "status quo", "being defective", and "existing functions" of its products and services. 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 take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, 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 contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, 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 this document 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 and/or its affiliates 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 directly contact Alibaba Cloud for any errors of this document.

Document conventions

Style	Description	Example		
▲ Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	Danger: Resetting will result in the loss of user configuration data.		
O Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	Warning: Restarting will cause business interruption. About 10 minutes are required to restart an instance.		
C) Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	Notice: If the weight is set to 0, the server no longer receives new requests.		
? Note	A note indicates supplemental instructions, best practices, tips, and other content.	Note: You can use Ctrl + A to select all files.		
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click Settings> Network> Set network type.		
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click OK.		
Courier font	Courier font is used for commands	Run the cd /d C:/window command to enter the Windows system folder.		
Italic	Italic formatting is used for parameters and variables.	bae log listinstanceid Instance_ID		
[] or [a b]	This format is used for an optional value, where only one item can be selected.	ipconfig [-all -t]		
{} or {a b}	This format is used for a required value, where only one item can be selected.	switch {active stand}		

Table of Contents

1.Network types	07
2.Instance IP addresses	09
2.1. IP addresses of ECS instances in VPCs	09
2.2. IP addresses of ECS instances in the classic network	10
2.3. Elastic IP addresses	12
2.4. IPv6 addresses	14
2.5. Internal network	15
3.Manage IPv4 addresses	17
3.1. Modify a private IP address	17
3.2. Change the public IP address of an instance	18
3.3. Associate or disassociate an EIP	19
3.4. Convert the public IP address of a VPC-type instance to a	20
3.5. Convert the public IP address of an instance in the classi	22
4.Configure IPv6 addresses	25
4.1. Configure an IPv6 address for a Windows instance	25
4.1.1. Use IPv6 addresses of Windows instances	25
4.1.2. Step 1: Create an IPv6 VPC	26
4.1.3. Step 2: Assign an IPv6 address	26
4.1.4. Step 3: Enable IPv6 public bandwidth	27
4.1.5. Step 4: Configure an IPv6 address	28
4.1.6. Step 5: Add IPv6 security group rules	30
4.1.7. Step 6: Test network connectivity	31
4.1.8. Step 7: (Optional) Delete an IPv6 address	32
4.2. Configure an IPv6 address for a Linux instance	32
4.2.1. Use IPv6 addresses of Linux instances	32
4.2.2. Step 1: Create an IPv6 VPC	33

4.2.3. Step 2: Assign an IPv6 address ₃	33
4.2.4. Step 3: Enable IPv6 public bandwidth 3	34
4.2.5. Step 4: Configure an IPv6 address 3	35
4.2.6. Step 5: Add IPv6 security group rules 4	18
4.2.7. Step 6: Test network connectivity 4	19
4.2.8. Step 7: (Optional) Delete an IPv6 address	19
5.View IP addresses 5	51
6.Elastic Network Interfaces 5	53
6.1. Overview 5	53
6.2. Managed ENIs 5	56
6.3. Create an ENI 5	57
6.4. Bind an ENI 5	59
6.5. Configure a secondary ENI 6	53
6.6. Assign secondary private IP addresses	75
6.7. Unassign secondary private IP addresses8	36
6.8. Modify an ENI 8	37
6.9. Edit the tags of an ENI 8	39
6.10. Unbind an ENI 9	90
6.11. Delete an ENI 9	91
6.12. Have ENI operations automatically performed in response9	92
7.Prefix lists 10)0
7.1. Overview 10)0
7.2. Create a prefix list 10)1
7.3. Clone a prefix list 10)3
7.4. Manage the entries in a prefix list 10)4
7.5. Delete prefix lists 10)6
7.6. Grant RAM users permissions on prefix lists 10)7
7.7. Use prefix lists to simplify management of security group r 10)9

8.Change the VPC of an ECS instance	114
9.Configure NIC multi-queue	117
10.Set the MTU size of an NIC	120
11.Migrate an ECS instance from the classic network to a VPC (n	124
11.1. Migrate ECS instances from the classic network to a VPC	124
12.Connect a classic network to a VPC	138
13.Network FAQ	140

1.Network types

Alibaba Cloud Elastic Compute Service (ECS) supports the classic network and virtual private clouds (VPCs).

VPCs

A VPC is an isolated virtual network built on Alibaba Cloud. VPCs are logically isolated from each other. You can customize the topology and IP addresses within a VPC. VPCs are suitable for users who have high network security requirements and network management capabilities.

For more information about VPCs, see What is a VPC?

Classic network

Services that use the classic network are deployed in the public infrastructure of Alibaba Cloud, and are planned and managed by Alibaba Cloud. The classic network is suitable for users who have high requirements for network availability.

Note If you purchase an ECS instance after 12:00:00 (UT C+8) on June 16, 2017, you cannot select the classic network.

Differences

The following table describes the differences between VPCs and the classic network.

ltem	VPC	Classic network
Layer 2 logical isolation	Supported.	Not supported.
Custom private CIDR block	Supported.	Not supported.
Private IP address planning	Private IP addresses must be unique within a single VPC, but can be duplicate across VPCs.	Private IP addresses must be unique in the classic network.
Instance communication within or between private networks	Instances in the same VPC can communicate with each other. However, instances in different VPCs are isolated from each other.	Instances in the classic network can communicate with each other if they belong to the same region and the same account.
Tunneling	Supported.	Not supported.
Custom router	Supported.	Not supported.
Route table	Supported.	Not supported.
vSwitch	Supported.	Not supported.
SDN	Supported.	Not supported.
Self-managed NAT gateway	Supported.	Not supported.

Network•Network types

ltem	VPC	Classic network
Self-managed VPN	Supported.	Not supported.

2.Instance IP addresses 2.1. IP addresses of ECS instances in VPCs

IP addresses are used to connect to Elastic Compute Service (ECS) instances or services deployed on ECS instances. Two types of IP addresses can be assigned to ECS instances in virtual private clouds (VPCs): private IP addresses and public IP addresses.

Private IP addresses

Each new ECS instance in a VPC is assigned a private IP address based on the CIDR block of the VPC and the CIDR block of the vSwitch to which the instance is connected.

Private IP addresses can be used in the following scenarios:

- Load balancing
- Communication over the internal network between ECS instances within the same VPC
- Communication over the internal network between an ECS instance and other cloud services such as Object Storage Service (OSS) and ApsaraDB RDS within the same VPC

You can use the ECS console to change private IP addresses. For more information, see Modify a private IP address. For more information about internal network communication, see Internal network.

Public IP addresses

ECS instances in VPCs support the following types of public IP addresses:

- PublicIP addresses: the public IP addresses assigned by ECS.
- Elastic IP addresses (EIPs). For more information, see Elastic IP addresses .

The following table describes the major differences between the two types of public IP addresses.

ltem	PublicIP address	EIP	
Scenario	If you want an instance to have a public IP address that will not be retained when the instance is released, you can configure a PublicIP address to be automatically assigned when you create the instance.	If you want to retain a public IP address after its associated instance is released for use with other ECS instances, use an EIP. Each EIP can be associated with or disassociated from different ECS instances. After an instance is released, its associated EIP is retained.	
Description	If you select Assign Public IPv4 Address when you create an ECS instance, a PublicIP address is assigned to the instance.	After an EIP is created, it can be associated with an ECS instance that is not assigned a PublicIP address. For more information, see Apply for an EIP and Associate an EIP with an instance.	

Network-Instance IP addresses

ltem	PublicIP address	EIP	
Maximum number of IP addresses that can be assigned to or associated with a single ECS instance	Only a single PublicIP address can be assigned to each ECS instance.	Multiple EIPs can be associated with a single ECS instance in multi-EIP-to-ENI mode. For more information, see Associate EIPs with secondary ENIs in multi-EIP-to-ENI mode (new applications are not accepted).	
Method to disassociate an IP address	After a PublicIP address is assigned to an ECS instance, the address can only be released along with the instance but cannot be disassociated from the instance. You can convert PublicIP addresses into EIPs as needed. EIPs can be associated with or disassociated from ECS instances at any time. For more information, see Convert the public IP address of a VPC-type instance to an EIP.	EIPs can be disassociated from ECS instances. For more information, see Disassociate an EIP from an instance.	
Method to release an IP address	 When ECS instances are released, their PublicIP addresses are also released. During the lifecycle of an ECS instance, you can release its PublicIP address by setting the public bandwidth of the instance to 0 Mbit/s. For more information about how to modify the public bandwidth of a subscription ECS instance, see Modify the bandwidth configurations of subscription instances. For more information about how to modify the public bandwidth of a pay- as-you-go ECS instance, see Modify the bandwidth configurations of pay-as- you-go instances. 	For more information, see Release a pay- as-you-go EIP.	
Method to view a media access control (MAC) address	ECS instances in VPCs are connected to the Int addresses to internal network interface contr ECS instances in VPCs cannot be discovered fr whether the instances are assigned PublicIP a	ternet by using the mapping of public IP rollers (NICs). Therefore, the public NICs of rom within the instances regardless of ddresses or associated with EIPs.	

Billing

You are charged only for outbound traffic to the Internet. For more information, see Public bandwidth.

2.2. IP addresses of ECS instances in the classic network

IP addresses are used to connect to Elastic Compute Service (ECS) instances or services deployed on ECS instances. Two types of IP addresses can be assigned to ECS instances in the classic network: internal IP addresses and public IP addresses.

Internal IP addresses

Each new ECS instance in the classic network is assigned an internal IP address for communication over the internal network. For more information about internal network communication, see Internal network.

? Note Internal IP addresses of ECS instances in the classic network do not support multicasting or broadcasting.

Internal IP addresses can be used in the following scenarios:

- Load balancing
- Communication over the internal network between ECS instances within the same LAN
- Communication over the internal network between an ECS instance and other cloud services such as Object Storage Service (OSS) and ApsaraDB RDS within the same LAN

The internal IP addresses of ECS instances in the classic network are assigned by the system and cannot be modified. Do not change internal IP addresses within the operating systems. Otherwise, communication within the internal network may be interrupted.

Public IP addresses

If you set the public bandwidth of an instance to a value greater than 0 when you purchase the instance in the classic network, the system assigns a public IP address to the instance.

Public IP addresses can be used in the following scenarios:

- Communication between ECS instances and the Internet
- Communication between an ECS instance and other cloud services that are not within the same LAN

The following section describes the operations related to public IP addresses:

• If you select **Assign Public IP Address** when you create an ECS instance in the classic network, the system assigns a public IP address to the instance. For more information, see **Create an instance by using the wizard**.

(?) Note After the system assigns a public IP address to an ECS instance in the classic network, you cannot release or disassociate the public IP address. If you set the public bandwidth of an instance to 0 Mbit/s when you downgrade the configurations of the instance during renewal or change the bandwidth of the pay-as-you-go instance, the instance cannot be connected to the Internet but the public IP address is retained.

- If you do not select Assign Public IP Address when you create an ECS instance in the classic network, you can assign a public IP address to the instance by modifying the outbound public bandwidth. For more information about how to modify the public bandwidth of a subscription ECS instance, see Modify the bandwidth configurations of subscription instances. For more information about how to modify the public bandwidth of a pay-as-you-go ECS instance, see Modify the bandwidth configurations of pay-as-you-go instances.
- You can modify the public IP address of an instance only within six hours after the instance is created. For more information, see Change the public IP address of an instance.
- When you manually release an ECS instance in the classic network, you can convert the public IP

address of the instance into an elastic IP address (EIP). EIPs can be associated with ECS instances in virtual private clouds (VPCs) for a variety of scenarios such as network migration and flexible bandwidth adjustment. For more information, see Convert the public IP address of an instance in the classic network into an EIP.

Billing

Communication traffic generated by using internal IP addresses is free of charge.

You are charged only for outbound Internet traffic. For more information, see Public bandwidth.

2.3. Elastic IP addresses

An elastic IP address (EIP) is a public IP address that you can purchase and use as an independent resource. EIPs can be retained independently and associated with or disassociated from Elastic Compute Service (ECS) instances in virtual private clouds (VPCs).

Overview

EIPs are NAT IP addresses that are located in the public gateway of Alibaba Cloud. By means of NAT, EIPs are mapped to the primary elastic network interfaces (ENIs) of the ECS instances that are associated with the EIPs. You can associate EIPs with ECS instances that are located in VPCs to enable the instances to communicate over the Internet. Similarly to system-assigned public IP addresses (PublicIP addresses), EIPs remain invisible inside the operating systems of ECS instances.

Advantages

ECS instances located in VPCs support both PublicIP addresses and EIPs. For information about the use scenarios of and differences between these public IP addresses, see IP addresses of ECS instances within VPCs.

The following table describes the advantages of EIPs over PublicIP addresses.

ltem	PublicIP address	EIP
Can the public IP address be independently purchased and used?	No	Yes
Can the public IP address be associated with or disassociated from an ECS instance as needed?	No	Yes
Is the public IP address retained when the associated ECS instance is released?	Νο	Yes
Maximum number of public IP addresses per ECS instance	1	In multi-EIP-to-ENI mode, multiple EIPs can be associated with a single ECS instance. For more information, see Associate EIPs with secondary ENIs in multi-EIP-to-ENI mode (new applications are not accepted).

Billing methods

EIPs support the subscription and pay-as-you-go billing methods. For more information, see Billing overview.

Limits

EIPs can be associated with ECS instances or ENIs. An EIP can be associated only with an ECS instance that meets the following requirements:

- The ECS instance is located in a VPC.
- The ECS instance is located in the same region as the EIP.
- The ECS instance is in the **Running** or **Stopped** state.
- No public IP addresses are associated with the primary ENI of the ECS instance.

Operations

• Create an EIP

You can go to the Elastic IP Addresses page in the VPC console to create an EIP and associate the EIP with an ECS instance that is located in a VPC and is not assigned public IP addresses. For more information, see Apply for an EIP.

Note In multi-EIP-to-ENI mode, multiple EIPs can be associated with a single ECS instance. For more information, see Associate EIPs with secondary ENIs in multi-EIP-to-ENI mode (new applications are not accepted).

- Associate an EIP with an ECS instance
 - You can log on to the ECS console and associate an EIP with an ECS instance that is located in a VPC and is not assigned public IP addresses. For more information, see Associate an EIP with an instance.
 - You can also go to the Elastic IP Addresses page in the VPC console and associate an EIP with an ECS instance that is located in a VPC and is not assigned public IP addresses. For more information, see Associate an EIP with an ECS instance.
- Associate an EIP with a secondary ENI

You can go to the Elastic IP Addresses page in the VPC console and associate an EIP with a secondary ENI. Each ENI is automatically assigned a private IP address. After you associate an EIP with an ENI, the ENI has both a private IP address and a public IP address. For more information, see Overview.

- Disassociate an EIP from an ECS instance
 - If your ECS instance no longer needs an EIP, you can disassociate the EIP from the instance in the ECS console. For more information, see Disassociate an EIP from an instance.
 - You can also disassociate the EIP from the ECS instance on the Elastic IP Addresses page in the VPC console. For more information, see Disassociate an EIP from a cloud resource.
- Release an EIP

Billing of an EIP continues after it is disassociated. If you no longer need the EIP, go to the Elastic IP Addresses page in the VPC console to release the EIP. For more information, see Release a pay-as-you-go EIP.

For more information, see EIP overview.

2.4. IPv6 addresses

IPv4 addresses are widely used, but the limited number of IPv4 addresses restricts the development of the Internet. Compared with IPv4 addresses, IPv6 addresses are more sufficient and allow more types of devices to access the Internet. Elastic Compute Service (ECS) supports both IPv4 and IPv6 addresses.

Comparison between IPv4 and IPv6

ltem	IPv4	IPv6
Address length	32 bits (4 bytes)	128 bits (16 bytes)
Number of addresses	2^32	2^128
Address format	xxx.xxx.xxx.xxx Where xxx is a decimal number that can range from 0 to 255. Each x is a decimal integer, and leading zeros can be omitted. Example: 192.168.1.1	<pre>xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxx xx Where each x is a hexadecimal number, and leading zeros can be omitted. You can use a double colon (::) once in an IPv6 address to indicate a series of zeros. Example: CDCD:0000:0000:0000:8475:1111:390 0:2020=CDCD::8475:1111:3900:2020</pre>
Address Resolution Protocol (ARP)	Uses broadcast ARP Request frames to resolve an IP address to a link layer address.	Uses multicast neighbor solicitation messages to resolve an IP address to a link layer address.
Security	Implements a security mechanism based on applications and cannot provide protections at the IP layer.	Supports packet fragmentation to ensure data confidentiality and integrity and provides security at the IP layer.
LAN connection	Connects to LANs by using network interfaces.	Can work with Ethernet adapters and is supported over virtual Ethernet networks between logical partitions.
Address type	Unicast addressMulticast addressBroadcast address	Unicast addressMulticast addressAnycast address

Limits

- Only ECS instances that reside in virtual private clouds (VPCs) support IPv6 addresses. For information about ECS instance types that support IPv6 addresses, see Instance families.
- IPv6 addresses support communication over the internal network within VPC. To use an IPv6 address to communicate with the Internet, you must purchase a public bandwidth plan for the IPv6 address on the IPv6 Gateway page in the VPC console. For more information, see Purchase a public bandwidth

plan for an IPv6 address.

• Only a single IPv6 address can be assigned to each ECS instance.

2.5. Internal network

If you need to transmit data between two ECS instances within the same region, we recommend that you transmit data over the internal network. ECS instances can also be connected to ApsaraDB RDS, SLB, and OSS over the internal network.

In the internal network, each non-I/O optimized instance has a shared bandwidth of 1 Gbit/s and each I/O optimized instance has a shared bandwidth of 10 or 25 Gbit/s. The internal network is a shared network and bandwidth may fluctuate.

The following rules apply to an ECS instance in the internal network:

• The internal communication is affected by the network type, account, region, or security group of an ECS instance. The following table describes how the internal communication is affected.

Network type	Account	Region	Security group	Internal communication
One VPC	One or more accounts	One region	One security group	By default, Internal communication is enabled. Instances within a security group can also be isolated from each other. For more information, see Network isolation within a basic security group.
			Different security groups	You can implement internal communication by authorizing mutual access between two security groups. For more information, see Security groups for different use casesConfiguration guide for ECS security groups.
Different VPCs	One or more accounts	One region	Different security groups	You can implement internal communication by using Express Connect. For more information, see Common scenarios.
		Different regions		
	One account	One region	One security group	By default, internal communication is enabled.
		,	1	

Classic network Network type	Account	Region	Security group	Internal communication
	Different accounts	One region	Different security groups	You can implement internal communication by authorizing mutual access between two security groups. For more information, see Security groups for different use casesConfiguration guide for ECS security groups.

- The private IP address of an ECS instance within a VPC can be modified. For more information, see Modify a private IP address. The private IP address of an ECS instance within the classic network cannot be modified.
- You can use the ClassicLink feature of VPC to connect an ECS instance in the classic network to the cloud resources deployed in a VPC. For more information, see ClassicLink.

3.Manage IPv4 addresses 3.1. Modify a private IP address

After you create a Virtual Private Cloud-type (VPC-type) Elastic Compute Service (ECS) instance, you can modify the private IP address of the instance directly or by changing the vSwitch of the instance.

Prerequisites

- No secondary private IP addresses are configured for elastic network interfaces (ENIs) of the instance.
- The instance whose private IP address you want to modify is in the Stopped state.

Procedure

- 1.
- 2.
- 3.
- 4. Find the instance for which you want to modify the private IP address and choose More > Network and Security Group > Modify Private IP Address in the Actions column.
- 5. In the **Modify Private IP Address** dialog box, modify the parameters and click **Modify**.
 - If you want to change the vSwitch, make sure that the selected vSwitch resides within the same zone as the instance.
 - If you do not want to change the vSwitch, modify only the private IP address.

Modify Private IP Address				
Instance:	i-b			
Zone:	China East 1 Zone G			
VSwitch:	vs by the same zone as the instance	✓ 4090 private IP addresses available		
Private IP Address:	17 7 The specified private IP address must be unoccupied in the VSwitch network segment. If no private IP address is specified, an idle private IP address will be automatically assigned to the ECS instance.			
		Modify Cancel		

6. Choose More > Instance Status > Start in the Actions column.

Onte The new private IP address takes effect after the ECS instance is restarted.

Related information

• ModifyInstanceVpcAttribute

3.2. Change the public IP address of an instance

If an Elastic Compute Service (ECS) instance in the classic network or in a virtual private cloud (VPC) is assigned a public IP address, you can change the IP address within 6 hours after the instance is created.

Prerequisites

The ECS instance whose public IP address you want to change meets the following requirements:

• The instance is in the Stopped (Stopped) state.

? Note

- The instance is assigned a public IP address.
- The instance is less than 6 hours old.

Context

The following limits apply when you change the public IP address of an ECS instance:

- The public IP address of the instance can be changed up to three times.
- The operations described in this topic are not applicable to an instance that is not assigned a public IP address. In this case, you can use one of the following methods to assign a public IP address (system-assigned public IP address or elastic IP address) to the instance:
 - Apply for and associate an elastic IP address (EIP) with the instance. For more information, see Apply for an EIP of *EIP documentation*.
 - Modify the public bandwidth of the instance to allocate a system-assigned public IP address. For information about how to perform this operation on subscription instances, see Overview of instance configuration changes. For information about how to perform this operation on pay-as-you-go instances, see Modify the bandwidth configurations of pay-as-you-go instances.

Procedure

1.

- 2.
- 3.
- 4. Find the instance whose public IP address you want to change. Choose More > Network and Security Group > Change Public IP Address.

IP Address	Status 👻	Network Type 👻	Configuration	Billing Method 👻		Actions
47. 187(Internet) 172.16.236.212(Private)	Stopped	VPC	2 vCPU 8 GiB (I/O Optimized) ecs.g5.large 5Mbps (Peak Value)	Pay-As-You-Go Released At April 18, 2019, 11:22	Change Instance Type	Manage More -
itch to Subscription Rele	ase Setting	More		Add to Security Group Configure Security Group	Buy Same Type Instance Status Instance Settings	Þ
				Modify Private IP Address Manage Secondary Private IP Address	Password/Key Pair Configuration Change	Þ
				Convert to EIP Change Public IP Address	Disk and Image	up b
					Operations and Troublesh	ooting

- 5. In the **Change Public IP Address** dialog box, click **Start Now**. If the public IP address of the instance is changed, a new public IP address is displayed in the dialog
- If the public IP address of the instance is changed, a new public IP address is displayed in the dial box.
- 6. Click OK.

3.3. Associate or disassociate an EIP

An elastic IP address (EIP) is a public IP address that you can purchase and use as an independent resource. You can create an EIP and associate it with an Elastic Compute Service (ECS) instance in a virtual private cloud (VPC) that has no public IP addresses to make the instance accessible to the Internet.

Prerequisites

To associate an EIP with an instance, make sure that the instance meets the following requirements:

- The instance resides within a VPC.
- The instance resides within the same region as the EIP.
- The instance is in the **Running** or **Stopped** state.
- No public IP addresses are associated with the primary network interface controller (NIC) of the instance.

Associate an EIP with an instance

- 1.
- 2.
- 3.
- 4. Find the instance with which you want to associate an EIP and click **More > Network and Security Group > Associate EIP** in the Actions column.
- 5. In the Associate EIP dialog box, select an EIP from the EIP drop-down list.

If no EIPs are available, click **Create EIP** on the right side of the EIP drop-down list to create an EIP. For more information, see Apply for an EIP.

6. Click OK.

After the EIP is associated with the instance, the EIP is displayed in the **IP Address** column corresponding to the instance on the Instances page. Example: 192.168.XX.XX (EIP).

Disassociate an EIP from an instance

- 1.
- 2.
- 3.
- 4. Find the instance from which you want to disassociate an EIP and click **More > Network and Security Group > Disassociate EIP** in the Actions column.
- 5. In the **Disassociate EIP** message, confirm the message.
- 6. Click OK.

After the EIP is disassociated from the instance, the EIP is not displayed in the **IP Address** column corresponding to the instance on the Instances page.

? Note

- After an EIP is disassociated from an instance, you are still charged for the EIP. If you no longer need the EIP, manually release it. For more information, see Release a pay-as-you-go EIP.
- After an EIP is disassociated from an instance, the instance cannot use the EIP to access the Internet.

3.4. Convert the public IP address of a VPC-type instance to an EIP

This topic describes how to convert the public IP address of an Elastic Compute Service (ECS) instance in a virtual private cloud (VPC) to an elastic IP address (EIP). After the conversion, you can disassociate the EIP from the instance and associate the EIP with the instance again or with another instance at any time. Address conversion does not affect the access from the Internet to your instance or cause transient network outage.

Prerequisites

Before you convert the public IP address of an instance in a VPC to an EIP, make sure that the following requirements are met:

- The instance is assigned a public IP address.
- If the instance is a pay-as-you-go instance, you have no overdue payments within your account.
- If the instance is a subscription instance, the instance has at least 24 hours remaining in its validity period.
- If the instance is a subscription instance, the instance uses the **pay-by-traffic** billing method for network usage. You can change the billing method for network usage from **pay-by-bandwidth** to pay-by-traffic by upgrading or downgrading instance configurations. For more information, see Overview of instance configuration changes.
- If the instance type has been changed, wait until the change takes effect before you convert the address.
- The instance is in the **Running** or **Stopped** state.

Context

After you convert the public IP address of an instance in a VPC to an EIP, the following results occur:

- The billing method for network usage remains unchanged.
- The EIP is separately billed, and separate bills are generated. For more information about EIP billing, see Billing overview. You can go to the Billing Management console. In the left-side navigation pane, click Usage Records. On the Usage Records page, select Elastic IP from the Product drop-down list to export EIP usage records.

This section describes how to convert the public IP address of an instance in a VPC to an EIP by using the ECS console. You can also convert the public IP address by calling the ConvertNatPublicIpToEip operation. To call this operation, use SDKs 4.3.0 or later. For more information, see Overview.

Procedure

- 1.
- 2.
- 3.
- 4.
- Find the instance whose network type is VPC, and then choose More > Network and Security Group > Convert to EIP in the Actions column.
- 6. In the dialog box that appears, confirm the information and click **OK**.
- 7. Refresh the instance list.

Result

After the public IP address is converted to an EIP, the public IP address is marked by (Elastic).

You can click the EIP to go to the Elastic IP Addresses page in the VPC console to manage the EIP.

What's next

After the public IP address is converted to an EIP, you can perform the following operations:

- Disassociate the EIP from the instance and then release the EIP or associate the EIP with a different instance. For more information, see Disassociate an EIP from a cloud resource.
- Add the EIP to an EIP bandwidth plan to reduce costs. For more information, see Associate an EIP with an EIP bandwidth plan, Select a product to gain access to the Internet, and Reduce the costs of data transfer over the Internet.

Related information

• Convert Nat PublicIpT oEip

3.5. Convert the public IP address of an instance in the classic network into an EIP

When you manually release an Elastic Compute Service (ECS) instance that is located in the classic network, you can convert the public IP address of the instance into an elastic IP address (EIP). EIPs can be associated as needed with ECS instances that are located in virtual private clouds (VPCs) for variety of scenarios such as network migration and flexible bandwidth adjustment. You can convert the public IP address of an ECS instance that is located in the classic network into an EIP only when you manually release the instance.

Prerequisites

Before you convert the public IP address of an ECS instance in the classic network into an EIP, make sure that the instance meets the following requirements:

- The instance is assigned a public IP address.
- If the instance is a pay-as-you-go instance, it is in the **Stopped** state and you do not have overdue payments in your account.
- If the instance is a subscription instance, it is in the **Expired** state.

(?) Note If a subscription instance in the classic network has expired, you can manually release the instance. When you release the instance, you can convert its public IP address of the instance into an EIP. If a subscription instance in the classic network has not expired, you can change its billing method to pay-as-you-go and then manually release the instance. When you release the instance, you can convert its public IP address into an EIP. For information about how to change the billing methods of instances, see Change the billing method of an instance from subscription to pay-as-you-go.

- If the instance is a subscription instance, the instance uses the **pay-by-traffic** billing method for network usage. You can change the billing method for network usage from **pay-by-bandwidth** to pay-by-traffic by upgrading or downgrading instance configurations. For more information, see Overview of instance configuration changes.
- If the configurations of the instance were changed, the new configurations have taken effect.

• Snapshots are created to back up data and prevent data loss caused by accidental operations. For more information, see Create a snapshot for a disk.

Context

After the public IP address of an ECS instance in the classic network is converted into an EIP, the following limits apply to the EIP:

- The EIP is billed based on data transfers.
- The public bandwidth of the EIP is the same as that of the ECS instance. You can increase the public bandwidth of the EIP in the VPC console. However, if the public bandwidth of the ECS instance is 0 Mbit/s, the public bandwidth of the EIP is 1 Mbit/s.
- The EIP cannot be associated with an ECS instance in the classic network.
- Unlike ECS instances in VPCs, ECS instances in the classic network have public network interface controllers (NICs). After the public IP address of an instance in the classic network is converted into an EIP, the public NIC is unbound from the instance and the MAC address of the NIC is released.

Procedure

1.

2.

3.

- 4. Find the ECS instance that you want to release and perform one of the following operations to release the instance:
 - To release a subscription instance, click **Release** in the **Actions** column.
 - To release a pay-as-you-go instance, choose **More > Instance Status > Release** in the **Actions** column.
- 5. Select Release Now, select Convert the public IP address of the ECS instance in a classic network to an EIP address. (The EIP addresses that are not bound to ECS instances will be billed.), and then click Next.

Release	\times	
*Release Mode:	Release Now O Scheduled Release	
How to retain Handling Resources:	 disks while the instance is released? Convert the public IP address of the ECS instance in a classic network to an EIP address. (The EIP addresses that are not bound to ECS instances will be billed.) 	
	Next Cancel	

6. Click OK.

Result

After the public IP address of the ECS instance is converted into an EIP, the ECS instance is released. You can view the EIP in the VPC console.

Instance ID/Name	IP Address	Monitor	Bandwidth	Connection Type	Charge Type(All)	Status(All) 77	Shared Bandwidth/Global Acceleration
eip- v3f1e5 -	39. 151	II	1 Mbps Pay By Traffic	BGP	Pay-As-You-Go //16/2019, 14:17:09 Created	Available	Add to Shared Bandwidth Package Add to Global Acceleration

What's next

You can associate this EIP with another ECS instance. For more information, see Associate an EIP with an ECS instance.

4.Configure IPv6 addresses 4.1. Configure an IPv6 address for a Windows instance

4.1.1. Use IPv6 addresses of Windows instances

This topic describes how to use IPv6 addresses of Windows Elastic Compute Service (ECS) instances.

Prerequisites

The Windows instances are deployed in virtual private clouds (VPCs). Before you assign IPv6 addresses, familiarize yourself with the basics of IPv6 addresses. For more information, see IPv6 addresses.

Procedure

The following figure shows how to use IPv6 addresses.



To use IPv6 addresses, perform the following operations:

1. Create an IPv6 VPC. For more information, see Step 1: Create an IPv6 VPC.

Before you assign an IPv6 address to an ECS instance, you must create an IPv6 VPC.

2. Assign an IPv6 address to an ECS instance. For more information, see Step 2: Assign an IPv6 address.

By default, when you create an ECS instance, a private IPv4 address is assigned instead of an IPv6 address. You must assign an IPv6 address to the ECS instance by using the ECS console or by calling an API operation.

3. Enable IPv6 public bandwidth. For more information, see Step 3: Enable IPv6 public bandwidth.

When an IPv6 address is assigned to an ECS instance, only communication within the VPC is allowed. If you want to allow traffic to or from the IPv6 address over the Internet, you must enable public bandwidth for the IPv6 address.

4. Configure the IPv6 address. For more information, see Step 4: Configure an IPv6 address.

After an IPv6 address is assigned to an ECS instance, you must log on to the ECS instance and configure the IPv6 address in the operating system before you can use the IPv6 address. We recommend that you use the automatic configuration tool to configure the IPv6 address.

5. Add security group rules. For more information, see Step 5: Add IPv6 security group rules.

You can use security group rules to allow or deny access to or from the Internet or internal network for ECS instances within a security group. For information about common cases, see Security groups for different use casesConfiguration guide for ECS security groups.

6. Test network connectivity. For more information, see Step 6: Test network connectivity.

You can log on to the ECS instance to test the network connectivity to ensure that the configured IPv6 address can access the Internet.

7. (Optional) Delete the IPv6 address. For more information, see Step 7: (Optional) Delete an IPv6 address.

You can delete IPv6 addresses that are no longer needed. After you delete the IPv6 address of an instance, the instance can still use its IPv4 address.

4.1.2. Step 1: Create an IPv6 VPC

This topic describes how to create an IPv6 virtual private cloud (VPC). Only VPC-type Elastic Compute Service (ECS) instances support IPv6 addresses. To configure an IPv6 address for a Windows instance, you must first create an IPv6 VPC.

Background information

By default, VPCs use the IPv4 addressing protocol. You can enable the IPv6 addressing protocol based on your business requirements. For more information about IPv6 addresses, see IPv6 addresses.

Procedure

- If no VPCs are created, you can enable IPv6 when you create a VPC. For more information, see Enable IPv6 for a VPC.
- If a VPC is created, you can enable IPv6 for the existing VPC. For more information, see Enable an IPv6 CIDR block for a VPC network.

4.1.3. Step 2: Assign an IPv6 address

This topic describes how to assign an IPv6 address to a Windows Elastic Compute Service (ECS) instance. You can assign an IPv6 address to a Windows instance when you create the instance. You can also assign an IPv6 address to an existing Windows instance.

Context

By default, when you create an ECS instance, a private IPv4 address is assigned instead of an IPv6 address.

Assign an IPv6 address when you create an ECS instance

Prerequisites: The virtual private cloud (VPC) and vSwitch of the instance are assigned IPv6 CIDR blocks. For more information, see Create a VPC with an IPv6 CIDR block.

Procedure: Create an ECS instance. For more information, see Create an instance by using the wizard. Take note of the following items when you configure parameters:

- In the **Basic Configurations** step, filter out **IPv6-supported** instance types, and then select an instance type.
- In the **Networking** step, select a VPC and a vSwitch for which IPv6 is enabled, and then select **Assign IPv6 Address Free of Charge**.
- In the Preview step, confirm that the IPv6 address is selected.

Assign an IPv6 address to an existing instance

Prerequisites:

• The instance supports IPv6. For more information about the instance types that support IPv6, see

Instance family.

- The network type of the instance is VPC. The VPC and vSwitch of the instance are assigned IPv6 CIDR blocks. For more information, see Create a VPC with an IPv6 CIDR block.
- The instance is in the Running or Stopped state.

To assign an IPv6 address to an existing instance, perform the following operations:

- 1.
- 2.
- 3.
- 4. Select the ECS instance to which you want to assign an IPv6 address and click **More** in the **Actions** column.
- 5. Choose Network and Security Group > Manage Secondary Private IP Address.
- 6. Click Assign New IP next to IPv6 Address.
- 7. Use one of the following methods to assign an IPv6 address:
 - Auto-assign: The system assigns a new IPv6 address.
 - Specify Address: You must specify an IPv6 address.
- 8. Click OK.

Related information

Assignlpv6Addresses

4.1.4. Step 3: Enable IPv6 public bandwidth

This topic describes how to enable IPv6 public bandwidth. If you want to use an IPv6 address to communicate over the Internet, you must enable IPv6 public bandwidth for the IPv6 address.

Context

By default, the IPv6 addresses assigned when Elastic Compute Service (ECS) instances are created are used only for communication over the internal network within virtual private clouds (VPCs).

Procedure

- 1. Log on to the VPC console.
- 2. In the left-side navigation pane, choose Access to Internet > IPv6 Gateway.
- 3. On the IPv6 Gateway page, select a region, find the IPv6 gateway for which you want to enable public bandwidth, and then click **Manage** in the Actions column.
- 4. On the IPv6 Gateway Details page, click the IPv6 Internet Bandwidth tab.
- 5. Find the IPv6 address that you want to use for communication over the Internet, and click **Create** IPv6 Internet Bandwidth.
- 6. Set the billing method for network usage and the maximum bandwidth, and click **Buy Now** to make the payment.

Two billing methods for network usage are available: pay-by-bandwidth and pay-by-traffic. For more information, see Billing.

4.1.5. Step 4: Configure an IPv6 address

This topic describes how to configure an IPv6 address for a Windows instance. IPv6 addresses can be manually configured, or automatically configured by the system. We recommend that you use an appropriate tool to automatically configure IPv6 addresses.

Automatically configure an IPv6 address

The ecs-util-ipv6 tool can be used to configure IPv6 addresses for instances that are assigned IPv6 addresses, or clear IPv6 configurations for instances that are not assigned IPv6 addresses.

The ecs-util-ipv6 tool has the following limits:

- The ecs-util-ipv6 tool applies only to Elastic Compute Service (ECS) instances in virtual private clouds (VPCs) and relies on instance metadata. Before you use this tool, make sure that the network service or the relevant outbound IP port (100.100.100.200:80) is enabled. For more information, see Overview of ECS instance metadata.
- When the ecs-util-ipv6 tool is running, network interface controllers (NICs) and the network service are restarted. As a result, your network connection may be unavailable for a period of time. Proceed with caution.

You can download the ecs-util-ipv6 tool from the following links:

- Windows Server 2003/2008/2012/2016 (32-bit)
- Windows Server 2003/2008/2012/2016 (64-bit)

Download the tool to your operating system and run the following command to run the tool as an administrator:

ecs-utils-ipv6.exe

If your ECS instance is assigned an IPv6 address, the tool automatically configures the IPv6 address. Otherwise, the tool clears the existing IPv6 address configurations.

If you want to configure IPv6 addresses for multiple instances at a time, we recommend that you use Cloud Assistant or user data to configure scripts that automate the process of configuring IPv6 addresses. For more information, see Overview and Overview of ECS instance user data. In the following script, Windows 64-bit and PowerShell are used:

```
#powershell
$install_dir="C:\Windows\system32"
$install_path = "$install_dir\ecs-utils-ipv6.exe"
if(-not (Test-Path -Path $install_path)){
    # download the tool
    $tool_url = 'http://ecs-image-utils.oss-cn-hangzhou.aliyuncs.com/ipv6/win/64/ecs-utils-
ipv6.exe'
    Invoke-WebRequest -uri $tool_url -OutFile $install_path
    Unblock-File $install_path
}
# run the tool
Start-Process -FilePath "$install_path" -ArgumentList "--noenterkey" -NoNewWindow
```

Manually configure an IPv6 address

Perform the following steps to manually configure an IPv6 address:

- 1. Check whet her IPv6 is enabled for your instance.
 - i. Connect to the instance. For more information, see Connect to a Windows instance by using a password.
 - ii. Start Windows CMD.
 - iii. Run the ipconfig command.
 - If the IPv6 address information is returned, IPv6 is enabled for your instance.
 - If no IPv6 address information is returned, IPv6 is not enabled for your instance. Perform the following steps to enable IPv6.
- 2. Enable IPv6.
 - Perform the following steps to enable IPv6 on Windows Server 2008, 2012, and 2016 operating systems:
 - a. Connect to the instance.
 - b. Choose Control Panel > Network and Internet > Network and Sharing Center.
 - c. In the View your active networks section of the Network and Sharing Center dialog box, click the current network connection name. In the Wi-Fi Status dialog box, click **Wireless Properties**.
 - d. Check whether the IPv6 option is selected. If the IPv6 option is not selected, select it and click **OK**.
 - Perform the following steps to enable IPv6 on Windows Server 2003:
 - a. Connect to the instance. For more information, see Connect to a Windows instance by using a password.
 - b. Choose Control Panel > Network and Sharing Center > Network Connections.
 - c. In the View your active networks section of the Network and Sharing Center dialog box, click the current network connection name. In the Wi-Fi Status dialog box, click **Wireless Properties**.
 - d. Perform the following operations as needed:
 - If IPv6 is displayed, select Internet Protocol Version 6 (TCP/IPv6), and then click OK.
 - If IPv6 is not displayed, install IPv6, select Internet Protocol Version 6 (TCP/IPv6), and then click OK. Perform the following steps to install IPv6:
 - a. In the Local Area Connection Properties dialog box, click **Install**. In the Select Network Component Type dialog box, choose **Protocol > Add**.
 - b. In the Select Network Protocol dialog box, choose Microsoft TCP/IP Version 6 > OK.
- 3. Check the IPv6 address assigned to the instance.

You can view the IPv6 address assigned to the instance by using the ECS console or instance metadata.

- ECS Console: For more information, see Step 2: Assign an IPv6 address.
- Instance metadata: You can view the IPv6 address by using the following metadata. For more information, see Overview of ECS instance metadata.
 - IPv6 address: network/interfaces/macs/[mac]/ipv6s

- IPv6 gateway: network/interfaces/macs/[mac]/ipv6-gateway
- IPv6 vSwitch CIDR block: network/interfaces/macs/[mac]/vswitch-ipv6-cidr-block
- 4. Manually configure IPv6 addresses.
 - Perform the following steps to enable IPv6 on Windows Server 2008, 2012, and 2016 operating systems:
 - a. Connect to the instance. For more information, see Connect to a Windows instance by using a password.
 - b. Choose **Control Panel > Network and Internet >** Network and Sharing Center.
 - c. In the View your active networks section of the Network and Sharing Center dialog box, click the current network connection name. In the Wi-Fi Status dialog box, click **Wireless Properties**.
 - d. Choose Internet Protocol Version 6 (TCP/IPv6) > Properties.
 - e. Select **Use the Following IPv6 Address**, enter the IPv6 address, subnet prefix length, and IPv6 gateway, and then click **OK**.
 - f. (Optional) Assign multiple IPv6 addresses: In the Internet Protocol Version 6 (TCP/IP) Properties dialog box, click Advanced. In the Advanced Settings dialog box, click Add to add multiple IPv6 addresses, and then click OK.
 - Perform the following steps to enable IPv6 on Windows Server 2003:
 - a. Connect to the instance. For more information, see Connect to a Windows instance by using a password.
 - b. Choose **Control Panel > Network Connections** to view the current network connection name. In this example, **Local Area Connection 2** is displayed.
 - c. Start Windows CMD.
 - d. Add IPv6 addresses.
 - Run the following command to add a single IPv6 address:

netsh interface ipv6 add address "Local Area Connection 2" <IPv6 address>

Run the following command to add multiple IPv6 addresses:

netsh interface ipv6 add address "Local Area Connection 2" <IPv6 address 1> netsh interface ipv6 add address "Local Area Connection 2" <IPv6 address 2>

e. Run the following command to add a default route:

netsh interface ipv6 add route ::/0 "Local Area Connection 2" <IPv6 gateway>

4.1.6. Step 5: Add IPv6 security group rules

This topic describes how to add IPv6 security group rules to an Elastic Compute Service (ECS) instance. IPv4 and IPv6 addresses are independent of each other. If the current security group rules do not apply to your IPv6 services, you must configure security group rules for the ECS instances to regulate communication with IPv6 addresses.

Procedure

1.

2.

3.

- 4. Find the security group and click Add Rules in the Actions column.
- 5. Click Add Rule.
- 6. Configure security group rules.

Enter the authorized IPv6 CIDR block in the Authorization Object field. For example, enter ::/0 to authorize all IPv6 addresses.

For more information about configuration operations and common scenarios of security group rules, see Add a security group rule and Security groups for different use casesConfiguration guide for ECS security groups.

4.1.7. Step 6: Test network connectivity

This topic describes how to test network connectivity. You can log on to the Elastic Compute Service (ECS) instance to test network connectivity to ensure that the configured IPv6 address can access the Internet.

Context

In this example, the **ping** -6 command is used to test the connectivity between the ECS instance and the Alibaba Cloud China site (aligun.com).

Procedure

1. Connect to an ECS instance that has an IPv6 address configured.

For more information, see Connect to a Windows instance by using a password.

2. Start the command-line tool.

You can use one of the following methods to start the command-line tool:

- In the lower-left corner of the Windows desktop, right-click the Start icon and open Windows PowerShell.
- On the desktop, press *Win+ R* to open the **Run** dialog box and enter cmd to start the command-line tool.
- 3. Run the following command to test network connectivity.

ping -6 aliyun.com

The following figure shows an example of the test result.

```
C:\Users\Administrator>ping -6 aliyun.com

Pinging aliyun.com [2401:b180: ...:f] with 32 bytes of data:

Reply from 2401:b180: ...f: time=7ms

Ping statistics for 2401:b180:1:50::f:

Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),

Approximate round trip times in milli-seconds:

Minimum = 7ms, Maximum = 7ms, Average = 7ms

C:\Users\Administrator>_
```

4.1.8. Step 7: (Optional) Delete an IPv6 address

This topic describes how to delete an IPv6 address when your Elastic Compute Service (ECS) instance no longer requires the IPv6 address. After you delete the IPv6 address of an instance, the instance can still use its IPv4 address.

Prerequisites

The instance is in the **Running** or **Stopped** state.

Procedure

- 1.
- 2.
- 3.
- 4. Select the ECS instance to which you want to assign an IPv6 address and click **More** in the **Actions** column.
- 5. Choose Network and Security Group > Manage Secondary Private IP Address.
- 6. Click Unassign next to IPv6 Address.
- 7. Click OK.

4.2. Configure an IPv6 address for a Linux instance

4.2.1. Use IPv6 addresses of Linux instances

This topic describes how to use IPv6 addresses of Linux Elastic Compute Service (ECS) instances.

Prerequisites

The Linux instances are deployed in virtual private clouds (VPCs). Before you assign IPv6 addresses, familiarize yourself with the basics of IPv6 addresses. For more information, see IPv6 addresses.

Procedure

The following figure shows the procedure to use IPv6 addresses.



To use IPv6 addresses, perform the following operations:

1. Create an IPv6 VPC. For more information, see Step 1: Create an IPv6 VPC.

Before you assign an IPv6 address to an ECS instance, you must create an IPv6 VPC.

2. Assign an IPv6 address to an ECS instance. For more information, see Step 2: Assign an IPv6 address.

By default, when you create an ECS instance, a private IPv4 address is assigned instead of an IPv6 address. You must manually assign an IPv6 address to the ECS instance.

3. Enable IPv6 public bandwidth. For more information, see Step 3: Enable IPv6 public bandwidth.

When an IPv6 address is assigned to an ECS instance, only communication within the VPC in which the instance is deployed is allowed. If you want to allow traffic to or from the IPv6 address over the Internet, you must enable IPv6 public bandwidth.

4. Configure the IPv6 address. For more information, see Step 4: Configure an IPv6 address.

IPv6 addresses can be manually configured, or automatically configured by the system. We recommend that you use an appropriate tool to automatically configure IPv6 addresses.

5. Add security group rules. For more information, see Step 5: Add IPv6 security group rules.

You can use security group rules to allow or deny access to or from the Internet or internal network for ECS instances within a security group. For information about common scenarios, see Security groups for different use casesConfiguration guide for ECS security groups.

6. Test network connectivity. For more information, see Step 6: Test network connectivity.

You can log on to the ECS instance to test the network connectivity to ensure that the configured IPv6 address can access the Internet.

7. (Optional) Delete the IPv6 address. For more information, see Step 7: (Optional) Delete an IPv6 address.

You can delete IPv6 addresses that are no longer needed. After you delete the IPv6 address of an instance, the instance can still use its IPv4 address.

4.2.2. Step 1: Create an IPv6 VPC

This topic describes how to create an IPv6 VPC. Only Elastic Compute Service (ECS) instances in virtual private clouds (VPCs) support IPv6 addresses. To assign an IPv6 address to a Linux instance, you must first create an IPv6 VPC.

Background information

By default, VPCs use the IPv4 addressing protocol. You can enable the IPv6 addressing protocol based on your business requirements. For more information about IPv6 addresses, see IPv6 addresses.

Procedure

- If no VPCs are created, you can enable IPv6 when you create a VPC. For more information, see Enable IPv6 for a VPC.
- If a VPC is created, you can enable IPv6 for the existing VPC. For more information, see Enable an IPv6 CIDR block for a VPC network.

4.2.3. Step 2: Assign an IPv6 address

This topic describes how to assign an IPv6 address to a Linux instance. You can assign an IPv6 address to a Linux instance when you create the instance, or assign an IPv6 address to an existing Linux instance.

Context

By default, when you create an ECS instance, a private IPv4 address is assigned instead of an IPv6 address.

Assign an IPv6 address when you create an ECS instance

Prerequisites: The virtual private cloud (VPC) and vSwitch of the instance are assigned IPv6 CIDR blocks. For more information, see Create a VPC with an IPv6 CIDR block.

Procedure: Create an ECS instance. For more information, see Create an instance by using the wizard. Take note of the following items when you configure parameters:

- In the **Basic Configurations** step, filter out **IPv6-supported** instance types, and then select an instance type.
- In the **Networking** step, select a VPC and a vSwitch for which IPv6 is enabled, and then select **Assign IPv6 Address Free of Charge**.
- In the **Preview** step, confirm that the IPv6 address is selected.

Assign an IPv6 address to an existing instance

Prerequisites:

- The instance supports IPv6. For more information about the instance types that support IPv6, see Instance family.
- The network type of the instance is VPC. The VPC and vSwitch of the instance are assigned IPv6 CIDR blocks. For more information, see Create a VPC with an IPv6 CIDR block.
- The instance is in the Running or Stopped state.

To assign an IPv6 address to an existing instance, perform the following operations:

- 1.
- 2.
- 3.
- 4. Select the ECS instance to which you want to assign an IPv6 address and click **More** in the **Actions** column.
- 5. Choose Network and Security Group > Manage Secondary Private IP Address.
- 6. Click Assign New IP next to IPv6 Address.
- 7. Use one of the following methods to assign an IPv6 address:
 - Auto-assign: The system assigns a new IPv6 address.
 - Specify Address: You must specify an IPv6 address.
- 8. Click OK.

Related information

AssignIpv6Addresses

4.2.4. Step 3: Enable IPv6 public bandwidth

This topic describes how to enable IPv6 public bandwidth. If you want to use an IPv6 address to communicate over the Internet, you must enable IPv6 public bandwidth.

Context

By default, the IPv6 addresses assigned when Elastic Compute Service (ECS) instances are created are used only for communication over the internal network within virtual private clouds (VPCs).

Procedure

- 1. Log on to the VPC console.
- 2. In the left-side navigation pane, choose Access to Internet > IPv6 Gateway.
- 3. On the IPv6 Gateway page, select a region, find the IPv6 gateway for which you want to enable public bandwidth, and then click **Manage** in the Actions column.
- 4. On the IPv6 Gateway Details page, click the IPv6 Internet Bandwidth tab.
- 5. Find the IPv6 address that you want to use for communication over the Internet, and click **Create** IPv6 Internet Bandwidth.
- 6. Set the billing method for network usage and the maximum bandwidth, and click **Buy Now** to make the payment.

Two billing methods for network usage are available: pay-by-bandwidth and pay-by-traffic. For more information, see Billing.

4.2.5. Step 4: Configure an IPv6 address

This topic describes how to configure IPv6 addresses for Linux Elastic Compute Service (ECS) instances. IPv6 addresses can be manually or automatically configured. We recommend that you use an appropriate tool to automatically configure IPv6 addresses.

Automatically configure IPv6 addresses

The ecs-util-ipv6 tool can be used to configure IPv6 addresses for instances that are already assigned IPv6 addresses and clear IPv6 configurations for instances that are not assigned IPv6 addresses.

Series	Distribution	Download URL
Red Hat Enterprise Linux (RHEL)	 CentOS 5, CentOS 6, CentOS 7, and CentOS 8 Red Hat 5, Red Hat 6, and Red Hat 7 Anolis OS Fedora Alibaba Cloud Linux 2/3 	Download URL
Debian	 Ubuntu 14, Ubuntu 16, Ubuntu 18, and Ubuntu 20 Debian 8, Debian 9, Debian 10, and Debian 11 	Download URL
SUSE Linux Enterprise Server (SLES)	 SUSE 11, SUSE 12, and SUSE 15 openSUSE 15 and openSUSE 42 	Download URL
FreeBSD	FreeBSD 11	Download URL

The following table lists the download URLs for the ecs-util-ipv6 tool.

The ecs-util-ipv6 tool has the following limits:

• The ecs-util-ipv6 tool applies only to ECS instances located in virtual private clouds (VPCs) and depends on instance metadata. Before you use this tool, make sure that the network service is not

disabled and that outbound access to 100.100.100.200 is allows on port 80. For more information, see Overview of ECS instance metadata.

• When the ecs-util-ipv6 tool runs, network interface controllers (NICs) and the network service are restarted. This may cause a brief network interruption. Proceed with caution.

Download an appropriate version of the ecs-util-ipv6 tool to your instance. Run the following commands to grant the execute permissions on the tool and run the tool as an administrator:

```
chmod +x ./ecs-utils-ipv6
./ecs-utils-ipv6
```

Notice If your instance uses a Ubuntu 14 public image, you must restart the instance after you run the preceding commands to make the configurations take effect. For more information, see Restart an instance.

If your instance is already assigned an IPv6 address, the IPv6 address is automatically configured. Otherwise, the existing IPv6 address configurations are automatically cleared.

Command parameters:

```
ecs-utils-ipv6 --help  # show usage
ecs-utils-ipv6 --version  # show version
ecs-utils-ipv6  # auto config all dev ipv6
ecs-utils-ipv6 --static [dev] [ip6s] [prefix_len] [gw6] # config dev static ipv6
e.g. ecs-utils-ipv6 --static eth0
        ecs-utils-ipv6 --static eth0 xxx::x1 64 xxx::x0
        ecs-utils-ipv6 --static eth0 "xxx::x1 xxx:x2 xxx:x3" 64 xxx::x0
ecs-utils-ipv6 --enable  # enable ipv6
ecs-utils-ipv6 --disable  # disable ipv6
```

You can enable, disable, or manually configure an IPv6 address or have an IPv6 address automatically configured. By default, an IPv6 address is automatically configured.

If you want to configure IPv6 addresses for multiple instances at a time, we recommend that you use Cloud Assistant or user data to configure a script that automates the configuration process. For more information, see Cloud Assistant and Prepare user data. The following sample script is a Bash shell script that is used to configure IPv6 addresses for an instance that runs a Red Hat Enterprise Linux (RHEL) operating system:
```
#!/bin/sh
install dir=/usr/sbin
install path="$install dir"/ecs-utils-ipv6
if [ ! -f "$install path" ]; then
    tool url="http://ecs-image-utils.oss-cn-hangzhou.aliyuncs.com/ipv6/rhel/ecs-utils-ipv6"
   # download the tool
   if ! wget "$tool url" -0 "$install path"; then
       echo "[Error] download tool failed, code $?"
       exit "$?"
   fi
fi
# chmod the tool
if ! chmod +x "$install path"; then
   echo "[Error] chmod tool failed, code $?"
   exit "$?"
fi
# run the tool
"$install_path"
```

Manually configure IPv6 addresses for an instance that runs an Alibaba Cloud Linux 2 or 3 operating system

Perform the following steps to manually configure IPv6 addresses for an Alibaba Cloud Linux instance.

- 1. Connect to the instance. For more information, see Connect to a Linux instance by using a password.
- 2. Check whet her IPv6 is enabled for your instance.

(?) Note IPv6 is disabled in Alibaba Cloud Linux 2 images of the aliyun_2_1903_64_20G_alib ase_20190829.vhd version and earlier. By default, IPv6 is enabled in Alibaba Cloud Linux 2 images of the aliyun_2_1903_x64_20G_alibase_20200221.vhd version and later.

Runthe ip addr | grep inet6 Or ifconfig | grep inet6 command.

- If the command output contains inet6 information, IPv6 is enabled for your instance. You can skip Step 3 and proceed to Step 4.
- If the command output does not contain <u>inet6</u> information, IPv6 is not enabled for your instance. Perform the following steps to enable IPv6.
- 3. Enable IPv6.

IPv6 can be temporarily or permanently enabled. If IPv6 is temporarily enabled, it becomes disabled when your instance is stopped or restarted. If IPv6 is permanently enabled, it remains enabled regardless of the instance state. Temporarily or permanently enable IPv6 based on your business requirements.

- Temporarily enable IPv6.
 - a. Go to the /etc/systemd/network/ directory.

cd /etc/systemd/network/

b. Run the ls command to check the .network file in the preceding directory.

```
In this example, the 50-dhcp.network file is used.
```

c. Modify the 50-dhcp.network file.

```
vi /etc/systemd/network/50-dhcp.network
```

d. Press the /key to enter the insert mode.

```
Change the information below [Network] to DHCP=yes .
```

```
⑦ Note
[Match]
Name=eth*
[Network]
DHCP=yes
```

After you make the change, press the *Esc* key, enter :wq, and then press the Enter key to save the change and exit.

- e. Enable IPv6 for all NICs or a specific NIC based on your business requirements.
 - Enable IPv6 for all NICs.

```
sudo sysctl -w net.ipv6.conf.all.disable_ipv6=0
sudo sysctl -w net.ipv6.conf.default.disable_ipv6=0
```

• Enable IPv6 for a specific NIC. In this example, the eth0 NIC is used.

sudo sysctl -w net.ipv6.conf.eth0.disable_ipv6=0

Restart the systemd-networkd service to make the configurations take effect.

sudo systemctl restart systemd-networkd

• Permanently enable IPv6.

a. Modify the /etc/sysctl.conf file.

vi /etc/sysctl.conf

- b. Press the /key to enter the edit mode. Use one of the following methods to modify the file:
 - Delete the following configurations:

```
net.ipv6.conf.all.disable_ipv6 = 1
net.ipv6.conf.default.disable_ipv6 = 1
net.ipv6.conf.lo.disable ipv6 = 1
```

To enable IPv6 for all NICs, make the following modification:

```
net.ipv6.conf.all.disable_ipv6 = 0
net.ipv6.conf.default.disable_ipv6 = 0
net.ipv6.conf.lo.disable ipv6 = 0
```

To enable IPv6 for a specific NIC, set the disable_ipv6 parameter that corresponds to the NIC to 0. In this example, the eth0 NIC is used.

```
net.ipv6.conf.eth0.disable_ipv6 = 0
```

After you make the modification, press the *Esc* key, enter :wq, and then press the Enter key to save the modification and exit.

c. Check whether the configurations in the /etc/sysctl.conf file are consistent with those in the /etc/sysctl.conf file in the initial RAM file system (initramfs).

diff -u /etc/sysctl.conf <(lsinitrd -f /etc/sysctl.conf)</pre>

Note Alibaba Cloud Linux 2 has an initramfs configured. If the configurations in the /etc/sysct1.conf file in the initramfs are inconsistent with those in the /etc/sysct1.conf file, the system may accept the file in the initramfs.

d. If the two files are inconsistent, run the following command to generate a new initramfs:

```
sudo dracut -v -f
```

e. Restart the instance.

reboot

f. Run the ifconfig command to check whether IPv6 is enabled.

If IPv6 is enabled, the following network configurations are displayed:

inet6 <Unicast address starting with fe80::>
inet6 <IPv6 address of ECS instance>

- 4. Manually configure IPv6 addresses.
 - i. Run the vi /etc/sysconfig/network-scripts/ifcfg-<NIC identifier> command to open the configuration file of an NIC. Example command: vi /etc/sysconfig/network-scripts/ifcfg-eth0. Replace eth0 in the command with the actual identifier of your NIC. Add the following configurations to the file based on your business requirements:

DHCPV6C=yes IPV6INIT=yes

Save the modification and exit.

ii. Restart the instance.

reboot

Manually configure IPv6 addresses for an instance that runs an operating system of another series

Perform the following steps to manually configure IPv6 addresses for an instance that runs an operating system of another series, such as CentOS, Debian, Ubuntu, or Fedora.

- 1. Connect to the instance. For more information, see Connect to a Linux instance by using a password.
- 2. Check whet her IPv6 is enabled for your instance.

? Note By default, IPv6 is enabled on Cent OS 8, Debian 10 and later, Ubuntu 18 and later.

Runthe ip addr | grep inet6 Or ifconfig | grep inet6 command.

- If the command output contain inet6 information, IPv6 is enabled for your instance. You can skip Step 3 and proceed to Step 4.
- If the command output does not contain <u>inet6</u> information, IPv6 is not enabled for your instance. Perform the following steps to enable IPv6.
- 3. Enable IPv6.

Operating system Procedure

Operating system	Procedure	
	 i. Run the vi /etc/modprobe.d/disable_ipv6.conf command and change the value of options ipv6 disable to 0. Then, save the modification and exit. ii. Run the vi /etc/sysconfig/network command and change the value of NETWORKING_IPV6 to yes. Then, save the modification and exit. 	
	iii. Run the following commands:	
	modprobe ipv6 -r modprobe ipv6	
	iv. Run the lsmod grep ipv6 command. If IPv6 is enabled, the following information is displayed in the command output:	
	ipv6 xxxxx 8	
CentOS 6	Note The parameter value in the third column cannot be 0. Otherwise, you must configure IPv6 again.	
	v. Run the vi /etc/sysctl.conf command and make the following modification:	
	<pre>#net.ipv6.conf.all.disable_ipv6 = 1 #net.ipv6.conf.default.disable_ipv6 = 1 #net.ipv6.conf.lo.disable_ipv6 = 1 net.ipv6.conf.all.disable_ipv6 = 0 net.ipv6.conf.default.disable_ipv6 = 0 net.ipv6.conf.lo.disable_ipv6 = 0</pre>	
	Save the modification and exit.	
	vi. Run the sysctl -p command to make the modification take effect.	

Operating system	Procedure	
	 i. Run the vi /etc/modprobe.d/disable_ipv6.conf command and change the value of options ipv6 disable to 0 . Then, save the modification and exit. ii. Run the vi /etc/sysconfig/network command and change the value of NETWORKING_IPV6 to yes . Then, save the modification and exit. 	
	iii. Run the vi /etc/sysctl.conf command and make the following modification:	
CentOS 7	<pre>#net.ipv6.conf.all.disable_ipv6 = 1 #net.ipv6.conf.default.disable_ipv6 = 1 #net.ipv6.conf.lo.disable_ipv6 = 1 net.ipv6.conf.all.disable_ipv6 = 0 net.ipv6.conf.default.disable_ipv6 = 0</pre>	
	Save the modification and exit. iv. Run the sysctl -p command to make the modification take effect.	
	 Run the vi /etc/default/grub command and delete ipv6.disab le=1 . Then, save the modification and exit. 	
	ii. Run the vi /boot/grub/grub.cfg command and delete ipv6.dis able=1 . Then, save the modification and exit.	
	iii. Restart the instance.	
	<pre>iv. Run the vi /etc/sysctl.conf command and make the following modification:</pre>	
Debian 8 and Debian 9	<pre>#net.ipv6.conf.all.disable_ipv6 = 1 #net.ipv6.conf.default.disable_ipv6 = 1 #net.ipv6.conf.lo.disable_ipv6 = 1 net.ipv6.conf.all.disable_ipv6 = 0 net.ipv6.conf.default.disable_ipv6 = 0 net.ipv6.conf.lo.disable_ipv6 = 0</pre>	
	Save the modification and exit. v. Run the sysctl -p command to make the modification take effect.	

Operating system Procedure	Procedure	
 Obuntu 14 and Ubuntu 16 Run the vi /etc/sysctl.conf command and make the follow modification: 	ving	
<pre>⑦ Note In openSUSE 42, after IPv6 is enabled, IPv6 addresses are automatically obtained and do not need to be manually configured.</pre> #net.ipv6.conf.all.disable_ipv6 = 1 #net.ipv6.conf.lo.disable_ipv6 = 1 net.ipv6.conf.all.disable_ipv6 = 0 net.ipv6.conf.default.disable_ipv6 = 0 net.ipv6.conf.lo.disable_ipv6 = 0 Save the modification and exit. ii. Run the sysctl -p command to make the modification take exists.	ffect.	
i. Run the vi /etc/rc.conf command and add ipv6_activate _interfaces="YES" . Then, save the modification and exit. ii. Run the /etc/netstart restart command to restart the network service.	 i. Run the vi /etc/rc.conf command and add ipv6_activate_allinterfaces="YES" . Then, save the modification and exit. ii. Run the /etc/netstart restart command to restart the network service. 	
SUSE 11 and SUSE 12 i. Run the vi /etc/modprobe.d/50-ipv6.conf command and d install ipv6 /bin/true command and d or command and or command and d or command and d or command and d or	elete kit.	
Image: Substant state Note In SUSE 11 and 12, Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state Image: Substant state </td <td>ving</td>	ving	
addresses are automatically obtained and do not need to be manually#net.ipv6.conf.all.disable_ipv6 = 1 #net.ipv6.conf.lo.disable_ipv6 = 1 met.ipv6.conf.all.disable_ipv6 = 1 net.ipv6.conf.all.disable_ipv6 = 0 net.ipv6.conf.all.disable_ipv6 = 0 net.ipv6.conf.lo.disable_ipv6 = 0		
Save the modification and exit.	ffect.	

4. View the IPv6 addresses assigned to your instance.

You can view the IPv6 address assigned to your instance by using the ECS console or instance metadata.

- ECS console: For more information about how to view the IPv6 address assigned to an instance in the ECS console, see Step 2: Assign an IPv6 address.
- Instance metadata: You can view the IPv6 address by using the following metadata items. For more information, see Overview of ECS instance metadata.
 - IPv6 address: network/interfaces/macs/[mac]/ipv6s
 - IPv6 gateway: network/interfaces/macs/[mac]/ipv6-gateway

• vSwitch IPv6 CIDR block: network/interfaces/macs/[mac]/vswitch-ipv6-cidr-block

5. Manually configure IPv6 addresses.

Operating system	Procedure		
	In this example, centos_6_10_x64_20G_alibase_20201120.vhd and centos_7_9_x64_20G_alibase_20211227.vhd are used. i. Run the vi /etc/sysconfig/network-scripts/ifcfg- <nic identif<="" td=""></nic>		
 CentOS 6 and CentOS 7 	command: vi /etc/sysconfig/network-scripts/ifcfg-eth0. Replace eth0		
• Red Hat 6 and Red Hat 7	in the command with the actual identifier of your NIC. Add the following configurations to the file based on your business requirements:		
AlmalinuxRocky Linux	DHCPV6C=yes IPV6INIT=yes		
	Save the modification and exit.		
	ii. Run the reboot command to restart the instance.		

Operating system	Procedure	
CentOS 8	In this example, centos_8_5_x64_20G_alibase_20211228.vhd is used. Check whether the NIC configuration file contains the IPV6INIT=yes and DHCPV6C=yes configurations. If yes, proceed to Step 2. If not, add the configurations to the file. vi /etc/sysconfig/network-scripts/ifcfg-eth0 In this example, eth0 is used. You must replace it with the actual identifier of your NIC. Save the modification and exit. Disable the ability of cloud-init to modify the NIC files in the /etc/sys config/network-scripts/ directory. Note Assigned IPv6 addresses do not need to be manually configured, but they may be lost when the instance is restarted. 	
	 a. Run the vi /etc/cloud/cloud.cfg command to open the NIC configuration file 	
	vi /etc/cloud/cloud.cfg	
	b. Add the following information before Example datasource con fig .	
	network: config: disabled	
	Save the modification and exit. iii. Run the reboot command to restart the instance.	

Operating system	Procedure	
Ubuntu 14	 In this example, ubuntu_14_0405_64_20G_alibase_20170824.vhd is used. i. Run the vi /etc/sysctl.conf command to open the NIC configuration file. Make the following modification: 	
	<pre>#net.ipv6.conf.all.disable_ipv6 = 1 #net.ipv6.conf.default.disable_ipv6 = 1 #net.ipv6.conf.lo.disable_ipv6 = 1 net.ipv6.conf.all.disable_ipv6 = 0 net.ipv6.conf.default.disable_ipv6 = 0 net.ipv6.conf.lo.disable_ipv6 = 0</pre>	
	Save the modification and exit. ii. Run the vi /etc/network/interfaces command to open the NIC configuration file. Add the following information to the file based on your business requirements:	
	iface eth0 inet6 dhcp	
	In this example, eth0 is used. You must replace it with the actual identifier of your NIC. Save the modification and exit. iii. Run the reboot command to restart the instance.	
	In this example, debian_9_13_x64_20G_alibase_20211227.vhd, debian_10_11_x64_20G_alibase_20211227.vhd, debian_11_2_x64_20G_alibase_20211227.vhd, and ubuntu 16 04 x64 20G alibase 20211028.vhd are used.	
 Ubuntu 16 Debian 8, Debian 9, Debian 10, and Debian 11 	i. Run the vi /etc/network/interfaces command to open the NIC configuration file. Add the following information to the file based on your business requirements:	
	iface eth0 inet6 dhcp	
	 In this example, eth0 is used. You must replace it with the actual identifier of your NIC. Save the modification and exit. ii. Run the reboot command to restart the instance. 	

Operating system	Procedure	
Ubuntu 18 and Ubuntu 20	In this example, ubuntu_18_04_x64_20G_alibase_20211227.vhd and ubuntu_20_04_x64_20G_alibase_20211227.vhd are used. i. Disable the ability of cloud-init to modify the NIC files in the /etc/sys config/network-scripts/ directory. ⑦ Note Assigned IPv6 addresses do not need to be manually configured, but they may be lost when the instance is restarted. Therefore, the ability of cloud-init to modify NIC files must be disabled. a. Run the vi /etc/cloud/cloud.cfg command to open the NIC configuration file. vi /etc/cloud/cloud.cfg b. Add the following information before Example datasource con fig . network: config: disabled Save the modification and exit. i. Run the reboot command to restart the instance.	
 Anolis OS 7.9 and Anolis OS 8.4 CentOS Stream Fedora 	In this example, anolisos_7_9_x64_20G_rhck_alibase_20220110.vhd, anolisos_8_4_x64_20G_rhck_alibase_20211008.vhd, centos_stream_8_x64_20G_alibase_20211227.vhd, and fedora_34_1_x64_20G_alibase_20211028.vhd are used. i. Check whether the NIC configuration file contains the IPV6INIT=yes and DHCPV6C=yes configurations. If yes, no further operations are required. If not, add the configurations to the file. vi /etc/sysconfig/network-scripts/ifcfg-eth0 In this example, eth0 is used. You must replace it with the actual identifier of your NIC. Save the modification and exit. ii. Run the reboot command to restart the instance.	

Operating system	Procedure	
	 In this example, freebsd_11_4_x64_30G_alibase_20210319.vhd is used. i. Run the vi /etc/rc.conf command to open the NIC configuration file. Add the following information to the file based on your business requirements: 	
	<pre>ipv6_enable="YES" ipv6_ifconfig_vtnet0="<ipv6 address=""> <subnet length="" prefix="">"</subnet></ipv6></pre>	
	ii. Make the following modification in the file. Save the modification and exit.	
	<pre>ip6addrctl_enable="YES" ipv6_activate_all_interfaces="YES" ipv6_network_interfaces="auto"</pre>	
	After the modification is complete, the configuration file contains the following content:	
FreeBSD 11	<pre>hostname="Aliyun" sshd_enable="YES" dumpdev="NO" ipv6_enable="YES" ip6addrctl_enable="YES" ip6addrctl_policy="ipv4_prefer" ipv6_activate_all_interfaces="YES" ipv6_network_interfaces="auto" ifconfig_lo0="inet 127.0.0.1 netmask 255.0.0.0" ifconfig_vtnet0="inet 192.168.XX.XX netmask 255.255.255.0" ipv6_ifconfig_vtnet0="2001:XXXX:4:4:4:4:4:4 prefixlen 64" defaultrouter="192.168.XX.XX" hostname="freebsd" ii. Run the reboot command to restart the instance.</pre>	

4.2.6. Step 5: Add IPv6 security group rules

This topic describes how to add IPv6 security group rules to an Elastic Compute Service (ECS) instance. IPv4 and IPv6 addresses are independent of each other. If the current security group rules do not apply to your IPv6 services, you must configure security group rules for the ECS instances to regulate communication with IPv6 addresses.

Procedure

- 1.
- 2.

3.

- 4. Find the security group and click Add Rules in the Actions column.
- 5. Click Add Rule.
- 6. Configure security group rules.

Enter the authorized IPv6 CIDR block in the Authorization Object field. For example, enter ::/0 to authorize all IPv6 addresses.

For more information about configuration operations and common scenarios of security group rules, see Add a security group rule and Security groups for different use casesConfiguration guide for ECS security groups.

4.2.7. Step 6: Test network connectivity

This topic describes how to test network connectivity. You can log on to the Elastic Compute Service (ECS) instance to test network connectivity to ensure that the configured IPv6 address can access the Internet.

Context

In this example, the ping -6 (or ping6) command is used to test the connectivity between the ECS instance and the Alibaba Cloud China site (aliyun.com).

Procedure

1. Connect to an ECS instance that has an IPv6 address configured.

For more information, see Connect to a Linux instance by using a password.

2. Run the following command to test network connectivity:

```
ping -6 aliyun.com
```

The following figure shows an example of the test result.

[root@test ~]# ping -6 aliyun.com		
PING aliyun.com(2401:b180:1:50::f (2401:b180 ::f)) 56 data bytes		
64 bytes from 2401:b180:1:50::f (2401:b180: ::f): icmp_seq=1 ttl=251 time=2.32 ms		
64 bytes from 2401:b180:1:50::f (2401:b180:] ::f): icmp_seq=2 ttl=251 time=2.28 ms		
64 bytes from 2401:b180:1:50::f (2401:b180: ■ ■::f): icmp_seq=3 ttl=251 time=2.28 ms		
64 bytes from 2401:b180:1:50::f (2401:b180:] ::f): icmp_seq=4 ttl=251 time=2.34 ms		
64 bytes from 2401:b180:1:50::f (2401:b180:] =::f): icmp_seq=5 ttl=251 time=2.33 ms		
64 bytes from 2401:b180:1:50::f (2401:b180: 📲::f): icmp_seq=6 ttl=251 time=2.31 ms		
^c		
aliyun.com ping statistics		
6 packets transmitted, 6 received, 0% packet loss, time 5006ms		
rtt min/avg/max/mdev = 2.284/2.315/2.344/0.035 ms		

4.2.8. Step 7: (Optional) Delete an IPv6 address

This topic describes how to delete an IPv6 address when your Linux instance no longer requires the IPv6 address. After you delete the IPv6 address of an instance, the instance can still use its IPv4 address.

Prerequisites

The instance is in the **Running** or **Stopped** state.

Procedure

- 1.
- 2.
- 3.
- 4. Select the ECS instance to which you want to assign an IPv6 address and click **More** in the **Actions** column.
- 5. Choose Network and Security Group > Manage Secondary Private IP Address.
- 6. Click Unassign next to IPv6 Address.
- 7. Click OK.

5.View IP addresses

This topic describes how to view the IP addresses of Elastic Compute Service (ECS) instances in the ECS console, including public IP addresses, elastic IP addresses (EIPs), primary private IP addresses, secondary private IP addresses, and IPv6 addresses.

View IP addresses on the Instances page

You can view the public IP addresses, EIPs, and primary private IP addresses of instances on the Instances page.

- 1.
- 2.
- 3.
- 4. On the **Instances** page, find the instance whose IP addresses you want to view, and view the IP addresses in the IP Address column.

Examples:

- Public IP address: 47.98.XX.XX(Public)
- EIP: 121.41.XX.XX(Elastic)
- Primary private IP address: 172.16.XX.XX(Private)

View IP addresses on the Instance Details page

You can view the public IP addresses, EIPs, primary private IP addresses, secondary private IP addresses, and IPv6 addresses of instances on the Instance Details page.

1.

2.

3.

- 4. On the **Instances** page, click the name of an instance in the **Instance ID/Name** column corresponding to the instance.
- 5. On the Instance Details page, view the IP addresses of the instance.

Examples:

- Basic Information section:
 - Public IP: 47.98.XX.XX
 - EIP: 121.41.XX.XX

• Network Information section:

- Primary Private IP Address: 192.168.XX.XX
- IPv6 Address: 2408:XXXX:325:a216:95f1:3dd9:6640:8b9e
- Secondary Private IP Address: 192.168.XX.XX,192.168.XX.XX

View IP addresses on the Network Interfaces page

You can view the EIPs, primary private IP addresses, and secondary private IP addresses of instances on the Network Interfaces page.

1.

- 2.
- 3.
- 4. On the **Network Interfaces** page, find an elastic network interface (ENI) bound to an instance and view the IP addresses in the Public IP Address and Private IP Address columns.

Examples:

- Public IP Address column: The associated EIP (if available) is displayed.
- Private IP Address column:
 - Primary private IP address: 192.168.XX.XX(Primary)
 - Secondary private IP address: 192.168.XX.XX(Secondary)

⑦ Note Multiple secondary private IP addresses for each ENI are displayed on separate lines.

6.Elastic Network Interfaces 6.1. Overview

An elastic network interface (ENI) is a virtual network interface controller (NIC) that can be bound to an Elastic Compute Service (ECS) instance of the Virtual Private Cloud (VPC) type. You can use ENIs to deploy high availability clusters and perform low-cost failover and fine-grained network management.

Attributes

An ENI is a virtual network interface that must be bound to an instance of the VPC type before you can use the ENI. The following table describes the attributes of an ENI.

Attribute	Description
ENI type	 ENIs consist of primary and secondary ENIs. Primary ENIs: created together with the instance. The lifecycle of a primary ENI is the same as the instance to which the primary ENI is bound. You cannot unbind a primary ENI from the instance to which the primary ENI is bound. Secondary ENIs: can be separately created. You can bind a secondary ENI to an instance or unbind a secondary ENI from an instance.
VPC	Only instances of the VPC type support ENIs. An ENI must reside within the same VPC as the instance to which the ENI is bound.
Zone	The vSwitch to which the ENI belongs must reside within the same zone as the instance to which the ENI is bound.
Security group	An ENI must be added to at least one security group. The security group controls the inbound and outbound traffic of the ENI.
EIP	An ENI can be associated with one or more elastic IP addresses (EIPs).
Primary private IP address	The primary private IP address is an IP address specified by the user or assigned by the system during ENI creation. The primary private IP address must be an idle IP address within the CIDR block of the vSwitch.
Secondary private IP address	The secondary private IP address must be an idle IP address within the CIDR block of the vSwitch. You can assign or revoke the secondary private IP address.
MAC address	A media access control (MAC) address is a globally unique identifier of an ENI.

Features

An ENI is an independent virtual NIC that can be migrated among multiple instances to support the flexible scaling and migration of services. When you create an ENI together with an instance, the ENI is automatically bound to the instance. You can also separately create a secondary ENI and bind it to an instance.

ENIs have the following features:

• In addition to the primary ENI that is created together with an instance, you can also bind multiple secondary ENIs to the instance. The ECS instance and the secondary ENIs that you want to bind to

the instance must reside within the same zone and VPC, but can belong to different vSwitches and security groups.

- Each ENI can be assigned multiple secondary private IP addresses based on the instance type of the instance to which the ENI is bound.
- When you unbind a secondary ENI from an instance and bind the ENI to another instance, the attributes of the ENI remain unchanged and the network traffic is redirected to the new instance.
- ENIs support hot-plug and can be migrated among instances. When you unbind an ENI from an instance and bind the ENI to another instance, services on the instances are not affected, and you do not need to restart the instances.

Limits

- The following limits apply to the resources supported by a single ENI:
 - Primary private IP address: one.
 - Secondary private IP address: one or more. The number of secondary private IP addresses is determined based on the instance type of the instance to which the ENI is bound. For more information, see Instance family.
 - EIP: one or more. The number of EIPs is determined based on how the EIPs are associated with the ENI. For more information, see Associate an EIP with an ECS instance.
 - MAC address: one.
 - Security group: one to five. At least one security group is required.
- A limited number of ENIs can be created for one account in each region. For more information, see the "ENI limits" section in Limits.
- The ENI and the instance to which the ENI is bound must reside within the same zone and VPC, but can belong to different vSwitches and security groups.
- The number of secondary ENIs that can be bound to an ECS instance is determined based on the instance type.
- Only I/O optimized instance types support ENIs.
- ECS instances of the classic network type do not support ENIs.
- The instance bandwidth is determined based on the instance type. You cannot increase the bandwidth of an ECS instance by binding multiple secondary ENIs to the instance.

Use scenarios

ENIs are suitable for the following scenarios:

• Deployment of high availability clusters

Multiple ENIs can be bound to a single ECS instance within a high availability architecture.

• Low-cost failover

You can unbind an ENI from a failed ECS instance and bind the ENI to another instance to redirect traffic to the backup instance. This allows quick recovery of services.

• Fine-grained network management

You can configure multiple ENIs for an instance. For example, you can use some ENIs for internal management and other ENIs for Internet business access to isolate management data from business data. You can also configure specific security group rules for each ENI based on the source IP addresses, protocols, and ports to achieve access control.

• Configuration of multiple private IP addresses for a single instance

You can assign multiple secondary private IP addresses to an ENI. If multiple applications are managed on your instance, you can assign an independent IP address for each application to improve the utilization of your instance.

• Configuration of multiple public IP addresses for a single instance

Only a single public IP address can be assigned to an ECS instance that has no ENIs bound. To assign multiple public IP addresses to an instance, you can associate EIPs with one or more ENIs of the instance. In NAT mode, each private IP address of an ENI can have EIPs associated.

Operations in the ECS console

The following table describes the operations that you can perform in the ECS console to manage ENIs.

Operation	Description	References
Create an ENI	You can create an ENI together with an instance or separately create an ENI.	Create an ENI
Bind an ENI	When you create an ENI together with an instance, the ENI is automatically bound to the instance. You can also separately create an ENI and bind it to an instance. An ENI can be bound only to a single ECS instance at a time. However, an ECS instance can have multiple ENIs bound to it.	Bind an ENI
Configure an ENI	For instances whose images cannot identify secondary ENIs, you must log on to the instance to configure the ENIs.	Configure a secondary ENI
	Note If an instance runs an image of CentOS 7.3 64-bit, CentOS 6.8 64-bit, or Windows Server 2008 R2 or later, you do not need to configure ENIs.	
Assign or revoke secondary private IP addresses	You can assign or revoke multiple secondary private IP addresses to or from an ENI.	 Assign secondary private IP addresses Unassign secondary private IP addresses
Modify an ENI	You can modify the security groups to which the primary and secondary ENIs belong. You can also modify the names and descriptions of secondary ENIs.	Modify an ENI
Unbind an ENI	You can unbind an ENI from an instance.	Unbind an ENI
Delete an ENI	You can delete an ENI after you unbind it from an instance.	Delete an ENI

API operations

The following table describes the API operations that you can call to manage ENIs.	

API	Description				
CreateNetworkInterface	Creates a secondary ENI.				
DeleteNetworkInterface	Deletes a secondary ENI.				
DescribeNetworkInterfa ces	Queries the details of one or more ENIs.				
AttachNetworkInterface	Binds a secondary ENI to an instance.				
AssignPrivatelpAddresse s	Assigns one or more secondary private IP addresses to an ENI.				
UnassignPrivatelpAddre sses	Revokes one or more secondary private IP addresses from an ENI.				
DetachNetworkInterface	Unbinds a secondary ENI from an instance.				
ModifyNetworkInterface Attribute	Modifies the name, description, and security group of a secondary ENI.				
DescribeInstances	Queries the information about ENIs that are bound to instances.				

6.2. Managed ENIs

When you create elastic network interfaces (ENIs) for specific Alibaba Cloud services such as Container Service for Kubernetes (ACK) and NAT Gateway, you can configure the ENIs to be managed by the services. ENIs managed by Alibaba Cloud services are called managed ENIs. Managed ENIs help prevent accidental resource deletion and ensure service availability. This topic describes the managed ENI feature and permissions on API operations used to query or manage managed ENIs.

Introduction

The managed ENI feature allows Alibaba Cloud services to have control on ENIs. When you use the Elastic Compute Service (ECS) console or the console of another Alibaba Cloud service to access managed ENIs, you can view the information of the ENIs but cannot manage them.

⑦ Note Procedure to create a managed ENI:

After you use Alibaba Cloud Security Token Service (STS) to grant permissions to an Alibaba Cloud service, the service calls the CreateNetworkInterface operation provided by the ECS API to create an ENI. For more information about STS, see What is STS?.

You can call the DescribeNetworkInterfaces operation and check the ServiceManaged and Description values in the response to determine whether an ENI is a managed ENI.

Note If an ENI is a managed one, the <u>ServiceManaged</u> value for it is true and the <u>Description</u> value is the name of the Alibaba Cloud service that manages the ENI.

Permissions on API operations used to query or manage managed ENIs

When you use OpenAPI to access managed ENIs, you can call API operations only to query managed ENIs. If you attempt to call an API operation to manage a managed ENI, you are prompted that the ENI is a managed ENI and cannot be manually managed and the InvalidOperation.EniServiceManaged error code is returned. The following table describes whether your Alibaba Cloud account or Alibaba Cloud services that create managed ENIs have permissions to call the API operations to query or manage the managed ENIs.

API operation	Description	Can be called by your Alibaba Cloud account for a managed ENI	Can be called by the Alibaba Cloud service that creates a managed ENI for the ENI	
DescribeNetworkInterfaces	Queries ENIs.	Yes	Yes	
DeleteNetworkInterface	Deletes an ENI.	No	Yes	
ModifyNetworkInterfaceAttribu te	Modifies the attributes such as the name, description, and security group of an ENI.	No	Yes	
AttachNetworkInterface	Binds an ENI.	No	Yes	
DetachNetworkInterface	Unbinds an ENI.	No	Yes	
AssignPrivatelpAddresses	Assigns one or more secondary private IP addresses to an ENI.	No	Yes	
UnassignPrivatelpAddresses	Unassigns one or more secondary private IP addresses from an ENI.	No	Yes	
Assignlpv6Addresses	Assigns one or more IPv6 address to an ENI.	No	Yes	
Unassignlpv6Addresses	Unassigns one or more IPv6 addresses from an ENI.	No	Yes	

6.3. Create an ENI

You can use elastic network interfaces (ENIs) to deploy high-availability clusters and perform low-cost failover and fine-grained network management. This topic describes how to separately create an ENI in the Elastic Compute Service (ECS) console.

Prerequisites

- A virtual private cloud (VPC) is created in a specified region and a vSwitch is created in the VPC. For more information, see Create and manage a VPC and Work with vSwitches.
- A security group is created in the specified VPC. For more information, see Create a security group.

Context

You can use one of the following methods to create an ENI. This topic describes how to separately create an ENI.

• Create an ENI when you create an instance.

When you create an instance, you can bind only one primary and one secondary ENIs to the instance. If the secondary ENI is not unbound from the instance, the secondary ENI is released together with the instance. For more information, see Bind an ENI when you create an instance.

• Separately create an ENI.

ENIs that are separately created are secondary ENIs and can be bound to instances.

After ENIs are created, deleted, bound to instances, or unbound from instances, ENI operation events are triggered. You can use CloudMonitor to set notifications of ENI operation events to obtain the ENI operation results such as whether ENIs are created. For more information, see ENI operation event notifications.

Procedure

- 1.
- 2.
- 3.
- 4. Click Create ENI.
- 5. In the Create ENI dialog box, configure the parameters described in the following table.

Parameter	Description					
ENI Name	Enter a name in the ENI Name field.					
VPC	Select the VPC where the instance is deployed. After an ENI is created, its VPC cannot be changed.					
	Note An ENI can be bound to only an instance that is in the same VPC as the ENI.					
VSwitch	Select a vSwitch that is in the same zone as the instance. After an ENI is created, you cannot change its vSwitch.					
	Note An ENI can be bound to only an instance that is in the same zone as the ENI. The instance and the ENI can connect to different vSwitches.					
Primary Private IP	(Optional) Enter the primary private IPv4 address of the ENI. The IPv4 address must be an idle IP address within the CIDR block of the vSwitch. If you do not specify an IPv4 address, an idle private IPv4 address is automatically assigned to your ENI after the ENI is created.					

Parameter	Description					
Secondary Private IP Addresses	 (Optional) Set the secondary private IP addresses of the ENI. Not set: No secondary private IP addresses are assigned to the ENI. Auto: You can enter an integer from 1 to 9 as the number of secondary private IP addresses that you want to assign to the ENI. The system automatically assigns the corresponding number of idle IP addresses in the vSwitch to the ENI. Manual: You can manually add secondary private IP addresses to the ENI. You can specify up to nine secondary private IP addresses. 					
Security Group	Select security groups in the specified VPC. You can specify one to five security groups. Image: The security groups and advanced security groups cannot be selected at the same time.					
Description	(Optional) Enter the description of the ENI for future management.					
Resource Group	(Optional) Select a resource group to which to add the ENI. For more information about resource groups, see Resource groups .					
Tag(Optional) Select one or more tags to be added to the ENI. For more information about tags, see Overview.						

6. Click OK.

Result

If the ENI is created, Available is displayed in the Status column in the ENI list.

What's next

After you separately create an ENI, you can bind it to an instance. For more information, see Bind an ENI.

Related information

• CreateNetworkInterface

6.4. Bind an ENI

You can use elastic network interfaces (ENIs) to deploy high-availability clusters and perform low-cost failover and fine-grained network management. This topic describes how to bind an ENI when you create an Elastic Compute Service (ECS) instance and how to create an ENI separately and bind it to an existing instance.

Prerequisites

- If you want to bind an ENI when you create an instance, make sure that the required preparations are made. For more information, see Create an instance by using the wizard.
- If you want to create an ENI separately and bind it to an existing instance, make sure that the

following requirements are met:

• The instance to which you want to bind an ENI is of an I/O optimized instance type and is in the **Stopped** or **Running** state.

Note For specific instance types, secondary ENIs can be bound to instances only when the instances are in the **Stopped** state. For more information about instance types, see **Instance types for which instances must be in the Stopped state**.

• The maximum number of ENIs that can be bound to the instance is not reached.

(?) Note An ENI can be bound to only a single ECS instance at a time. However, an ECS instance can have multiple ENIs. For information about the maximum number of ENIs that can be bound per instance for different instance types, see Instance family.

• If the instance was last started, restarted, or reactivated before April 1, 2018 and has remained in the Running state since then, you must restart the instance before you can bind ENIs to it.

? Note

- An ENI is created and is in the Available state. For more information, see Create an ENI.
- The instance and the ENI reside in the same virtual private cloud (VPC).
- The instance and the ENI reside in the same zone.

Context

You can use one of the following methods to bind an ENI to an instance:

• Bind an ENI when you create an instance.

When you create an instance in the ECS console, you can bind only a primary ENI and a secondary ENI to the instance. For more information, see Bind an ENI when you create an instance.

(?) Note For specific instance types, you cannot bind secondary ENIs when you create instances. To bind secondary ENIs to instances of these instance types, wait until the instances are created. For more information, see Instance family.

• Create an ENI Separately and bind it to an existing instance.

You can perform this operation on the Instances, Network Interfaces, or Security Groups page in the ECS console. For more information, see Bind an ENI to an existing instance on the Instances page, Bind an ENI to an existing instance on the Network Interfaces page, and Bind an ENI to an existing instance on the Security Groups page.

After ENIs are created, deleted, bound to instances, or unbound from instances, ENI operation events are triggered. You can use CloudMonitor to set notifications of ENI operation events to obtain the ENI operation results such as whether ENIs are created. For more information, see ENI operation event notifications.

Bind an ENI when you create an instance

For information about how to create an instance, see Create an instance by using the wizard. If you want to bind an ENI when you create an instance, take note of the following configurations:

- Basic configurations:
 - Region: ENIs are supported in all regions.
 - Instance Type: Select an I/O optimized instance type that allows ENIs to be bound during instance creation. For more information, see Instance families.
 - Image: ENIs can be automatically identified without additional configurations if you select one of the following images:
 - Alibaba Cloud Linux 3.2104 64-bit
 - Cent OS 8.0 64-bit, Cent OS 8.1 64-bit, and Cent OS 8.2 64-bit
 - Cent OS 7.3 64-bit, Cent OS 7.4 64-bit, and Cent OS 7.5 64-bit
 - Cent OS 6.8 64-bit and Cent OS 6.9 64-bit
 - Debian 10.5 64-bit and Debian 10.6 64-bit
 - Windows Server 2008 R2 and later

If you select an image that is not in the preceding list, you must configure ENIs so that they can be identified after the instance is created. For more information, see Configure a secondary ENI.

- Networking configurations:
 - Network Type: You must select VPC. Then, select a VPC and a vSwitch for the instance.
 - Elastic Network Interface: Click Add ENI to create an ENI. An ENI and the instance to which it is bound must reside within the same zone but do not need to connect to the same vSwitch.

After the instance is created, you can go to its Instance Details page and view the state of ENIs on the ENIs tab. If an ENI is bound to the instance, InUse is displayed in the Status/Creation Time column corresponding to the ENI.

Bind an ENI to an existing instance on the Instances page

- 1.
- 2.
- 3.
- 4. Find the instance to which you want to bind an ENI and choose **More > Network and Security Group > Bind Secondary ENI** in the **Actions** column.
- 5. In the Bind Secondary ENI dialog box, select a secondary ENI and click OK. After you bind the ENI to the instance, you can go to the Instance Details page of the instance and view the state of the ENI on the ENIs tab. If the ENI is bound to the instance, InUse is displayed in the Status/Creation Time column corresponding to the ENI.

Bind an ENI to an existing instance on the Network Interfaces page

- 1.
- 2.
- 3.
- 4. Find an available secondary ENI and click **Bind to Instance** in the Actions column.
- In the Bind to Instance dialog box, select an instance and click OK. Refresh the ENI list. If the ENI is bound to the instance, InUse is displayed in the Status/Creation Time column corresponding to the ENI.

Bind an ENI to an existing instance on the Security Groups page

- 1.
- 2.
- 3.
- 4. Find the security group to which an available ENI belongs and click **Manage ENIs** in the **Actions** column.
- 5. On the ENIs in Security Group page, find the ENI and click Bind to Instance in the Actions column.
- 6. In the Bind to Instance dialog box, select an instance and click OK. After you bind the ENI to the instance, you can choose the Network & Security > ENIs in the left-side navigation pane and view the state of the ENI on the Network Interfaces page. If the ENI is bound to the instance, InUse is displayed in the Status/Creation Time column corresponding to the ENI.

What to do next

For instances that use specific images, you may need to manually configure ENIs so that they can be identified by the instances. For more information, see Configure a secondary ENI.

Instance types for which instances must be in the Stopped state

For specific instance types, ENIs can be bound to or unbound from instances only when the instances are in the **Stopped** state. For more information, see **Stop** an instance.

Instance family	Instance type
s6, shared standard instance family	ecs.s6-c1m1.small, ecs.s6-c1m2.large, ecs.s6-c1m2.small, ecs.s6- c1m4.large, and ecs.s6-c1m4.small
t6, burstable instance family	ecs.t6-c1m1.large, ecs.t6-c1m2.large, ecs.t6-c1m4.large, ecs.t6- c2m1.large, and ecs.t6-c4m1.large
t5, burstable instance family	ecs.t5-c1m1.large, ecs.t5-c1m2.large, ecs.t5-c1m4.large, ecs.t5- lc1m1.small, ecs.t5-lc1m2.large, ecs.t5-lc1m2.small, ecs.t5- lc1m4.large, and ecs.t5-lc2m1.nano
Previous-generation shared instance families xn4, n4, mn4, and e4	 ecs.xn4.small ecs.n4.small and ecs.n4.large ecs.mn4.small and ecs.mn4.large ecs.e4.small and ecs.e4.large

The following table describes the instance types. For more information, see Instance family.

Related information

References

- Assign secondary private IP addresses
- RunInstances
- CreateInstance

• AttachNetworkInterface

6.5. Configure a secondary ENI

After secondary elastic network interfaces (ENIs) are bound to Elastic Compute Service (ECS) instances, some images used by these instances may not recognize the secondary ENIs and configure routes for the secondary ENIs. If this occurs, the secondary ENIs cannot be used on the instances. This topic describes how to configure secondary ENIs from within instances to have their IP addresses recognized and how to configure routes for the secondary ENIs.

Prerequisites

- A secondary ENI is bound to an ECS instance.
- You are connected to the ECS instance. For more information, see Connection methodsGuidelines on instance connection.

Context

If automatic configuration tools have been pre-installed in images that instances use, secondary ENIs that are bound to the instances can be automatically configured by the tools. You can use the secondary ENIs without manually configuring them. Skip the topic if your instance uses an image of one of the following versions:

- Alibaba Cloud Linux 3.2104 64-bit
- Cent OS 8.0 64-bit, Cent OS 8.1 64-bit, and Cent OS 8.2 64-bit
- Cent OS 7.3 64-bit, Cent OS 7.4 64-bit, and Cent OS 7.5 64-bit
- CentOS 6.8 64-bit and CentOS 6.9 64-bit
- Debian 10.5 64-bit and Debian 10.6 64-bit
- Windows Server 2008 R2 and later

Procedure

1. Check whether the IP address of a secondary ENI bound to an instance can be recognized.

For more information, see the Check whether the IP addresses of ENIs can be recognized section of this topic. If the IP address of the secondary ENI is recognized, skip the following steps. If the IP address of the secondary ENI is not recognized, proceed with the following steps to configure the secondary ENI.

2. Obtain the information of the secondary ENI.

When you configure a secondary ENI, the primary private IP address and media access control (MAC) address may be required. Prepare the information for subsequent configurations. For more information, see the Obtain the information of an ENI section of this topic.

In the examples provided in this topic, the sample values listed in the following table are used. In actual scenarios, replace them with the attribute values of your secondary ENI.

Secondary ENI attribute	Sample value
ENI name	eth1
MAC address	00:16:3e:0f:**:**

Secondary ENI attribute	Sample value
Primary private IP address	192.168.**.*2
Subnet mask	255.255.255.0
Gateway address	192.168.**.253

3. Configure the secondary ENI to have its IP address recognized.

The operations required to configure secondary ENIs vary based on the operating systems of instances to which the secondary ENIs are bound.

Operating system	References			
 Alibaba Cloud Linux 2 (Instances that run this operating system use the network-scripts network service) CentOS RedHat 	Configure a secondary ENI for an instance that runs an Alibaba Cloud Linux 2, CentOS 6, CentOS 7, or Red Hat operating system (network-scripts)			
Alibaba Cloud Linux 2 (Instances that run this operating system use the systemd-networkd network service) Note For more information, see Methods and impacts of switching the network service for instances that run Alibaba Cloud Linux 2.	Configure a secondary ENI for an instance that runs an Alibaba Cloud Linux 2 operating system (systemd-networkd)			
 Ubuntu Debian	Configure a secondary ENI for an instance that runs a Ubuntu or Debian operating system			
SUSEOpenSUSE	Configure a secondary ENI for an instance that runs a SUSE or openSUSE operating system			

4. Check whether routes are configured for the secondary ENI.

You can run the route -n command to check the route information. If no routes are configured for the secondary ENI or if the existing routes do not meet your requirements, manually configure routes for the secondary ENI. The following sections provide examples on configuring routes for a secondary ENI that is bound to an instance that runs one of the following operating systems:

- Configure routes for a secondary ENI that is bound to an instance that runs an Alibaba Cloud Linux 2 or Cent OS 7 operating system
- Configure routes for a secondary ENI that is bound to an instance that runs a CentOS 8 operating system

Check whether the IP addresses of ENIs can be recognized

Run the following command to check whether the IP addresses of ENIs can be recognized:

ip address show

Sample command outputs:

• The following command output shows that the IP address of the eth0 primary ENI is recognized but the IP address of the eth1 secondary ENI is not recognized. You can perform operations described in this topic to configure the secondary ENI.

[root@ecs ~]# ip address show
1: lo: <loopback,up,lower_up> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000</loopback,up,lower_up>
link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
inet 127. 🔰 /8 scope host lo
valid_lft forever preferred_lft forever
inet6 ::1 in scope host
valid_lft forever preferred_lft forever
2: eth0: <broadcast,multicast,up,lower_up> mtu 1500 qdisc mq state UP group default qlen 1000</broadcast,multicast,up,lower_up>
link/ether 00:16:3e:16 brd ff:ff:ff:ff:ff:ff
inet 192.168.
valid lft 315359900sec preferred_lft 315359900sec
inet6 fe80::216:3eff: scope link
valid lft forever preferred lft forever
3: eth1: <broadcast,multicast> mtu 1500 qdisc noop state DOWN group default qlen 1000</broadcast,multicast>
link/ether 00:16:3e:0f is brd ff:ff:ff:ff:ff:ff
[root@ecs ~]#

• The following command output shows that the IP addresses of both the eth0 primary ENI and the eth1 secondary ENI are recognized. You do not need to configure the secondary ENI.



Note In the preceding command outputs, 00:16:3e:16:**:** is the MAC address of the primary ENI and 00:16:3e:0f:**:** is the MAC address of the secondary ENI.

Obtain the information of an ENI

You can obtain the information of an ENI from instance metadata, by using the ECS console, or by calling an API operation. You can use one of the following methods to obtain the information of an ENI:

- Obtain the information of an ENI from instance metadata.
 - Obtain the MAC addresses of ENIs that are bound to an instance.

curl http://100.100.200/latest/meta-data/network/interfaces/macs/

? Note The MAC addresses of ENIs are required to obtain the primary private IP addresses, subnet masks, and gateway addresses of the ENIs.

• Obtain the primary private IP address of the specified ENI.

curl http://100.100.100.200/latest/meta-data/network/interfaces/macs/00:16:3e:19:**:**/
primary-ip-address

• Obtain the subnet mask of the specified ENI.

curl http://100.100.100.200/latest/meta-data/network/interfaces/macs/00:16:3e:19:**:**/
netmask

• Obtain the gateway address of the specified ENI.

```
curl http://100.100.100.200/latest/meta-data/network/interfaces/macs/00:16:3e:19:**:**/
gateway
```

The following figure shows the sample command output. In the command output,00:16:3e:16:**:** is the MAC address of the primary ENI and00:16:3e:0f:**:**is the MAC address of thesecondary ENI.

(?) Note After you run the <u>ip address show</u> command, you can determine which is the primary ENI and which is the secondary ENI based on the order in which the MAC addresses are displayed in the command output.

[root@ecs ~]# curl http://100.100.100.200/latest/meta-data/network/interfaces/macs/
00:16:3e:0f
00:16:3e:16 curl http://100.100.100.200/latest/meta-data/network/interfaces/macs/00:16:3e:0f //rimary-ip-address
192.168Curl http://100.100.100.200/latest/meta-data/network/interfaces/macs/00:16:3e:0f //netmask
255.255cul http://100.100.100.200/latest/meta-data/network/interfaces/macs/00:16:3e:0f 💶 🖉 / gateway
192.168. [root@ecs ~]#
255.255fcurl http://100.100.100.200/latest/meta-data/network/interfaces/macs/00:16:3e:0f/gateway 192.168[root@ecs ~]#

• Obtain the information of an ENI by using the ECS console.

i.

ii.

iii. On the **Network Interfaces** page, find the ENIs whose information you want to query and view their primary private IP addresses and MAC addresses in the Private IP Address and Type/MAC Address(All) columns.

The following figure shows an example of the Network Interfaces page in the ECS console.

Network Interfaces									
Create ENI Name	Create ENI Name Search by ENI name Q Tag			Tag			C		
ID/Name	Tag VSwitc	/VPC Zone	Security Group ID	Bound Instance	Public IP Address	Private IP Address	Type/MAC Address(Secondary) 꼬	Status/Creation Time(All) 모	Actions
eni. La forganizzazione da ylumnana arti inna a ili)2	vsw bp' vpc bp'	Hangzho Zone K	sg	i-t, f ray,		192.16 (Primary)	Secondary 00:16 5	InUse November 19, 2021, 16:17	Modify Unbind Delete Manage Secondary Private IP Address Check Security Group Rules
elle by linear could be placed an one	vswi bp1 vpci bp1 even bp1 bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 bp1 bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp1 even bp bp bp bp bp bp bp bp bp bp bp bp bp	Hangzh Zone K	ou sg- or			192.16 (Primary)	Secondary 00:16:	Available November 17, 2021, 18:23	Modify Bind to Instance Delete Manage Secondary Private IP Address
eri- ty-filolof-syffic (27) ylanoinai mi	 vsw bp1 vpc bp1 	Hangzhe Zone I	sg -tar ta			192.16 (Primary) 192.16 (Secondary) 192.16 (Secondary)	Secondary 00:16: The second	Available August 25, 2021, 16:48	Modify Bind to Instance Delete Manage Secondary Private IP Address

• Obtain the information of an ENI by running commands in Alibaba Cloud CLI to call the DescribeNetworkInterfaces operation.

```
aliyun ecs DescribeNetworkInterfaces \
--output cols=MacAddress,PrivateIpAddress rows=NetworkInterfaceSets.NetworkInterfaceSet[]
\
--RegionId 'cn-hangzhou' \
--InstanceId 'i-bpla5qj0bzhwz7g****'
```

The following figure shows the sample command output. In the command output,00:16:3e:16:**:** is the MAC address of the primary ENI and00:16:3e:0f:**:**is the MAC address of thesecondary ENI.

? Note After you run the <u>ip address show</u> command, you can determine which is the primary ENI and which is the secondary ENI based on the order in which the MAC addresses are displayed in the command output.



Configure a secondary ENI for an instance that runs an Alibaba Cloud Linux 2, CentOS 6, CentOS 7, or Red Hat operating system (network-scripts)

If your instance runs an Alibaba Cloud Linux 2, CentOS 6, CentOS 7, or Red Hat operating system and uses the network-scripts network service, you can use the multi-nic-util tool to have ENIs bound to the instance automatically configured. You can also manually modify the ENI configuration files to configure the ENIs.

• Use the multi-nic-util tool to have a secondary ENI automatically configured.

(?) Note If you want to use the multi-nic-util tool to have secondary ENIs automatically configured for CentOS instances, note that the multi-nic-util tool is supported only on some versions of CentOS images. If your instance uses a CentOS 6 image, make sure that the instance uses CentOS 6.8 or later. If your instance uses a CentOS 7.3 or later. If the multi-nic-util tool is not supported on the image version that your instance uses, you must manually modify the configuration files of secondary ENIs to configure the secondary ENIs.

i. Download and install the multi-nic-util tool.

```
wget https://image-offline.oss-cn-hangzhou.aliyuncs.com/multi-nic-util/multi-nic-util
-0.6.tgz && \
tar -zxvf multi-nic-util-0.6.tgz && \
cd multi-nic-util-0.6 && \
bash install.sh
```

ii. Restart the ENI service.

systemctl restart eni.service

• Manually modify the configuration file of a secondary ENI to configure the secondary ENI.

i. Open the configuration file of the secondary ENI.

```
vi /etc/sysconfig/network-scripts/ifcfg-eth1
```

ii. Add the information of the secondary ENI to the configuration file. Then, save and close the configuration file.

The following section provides an example of the ENI information to add to the configuration file:

```
DEVICE=eth1 # Specify the ENI that you want to configure.
BOOTPROTO=dhcp
ONBOOT=yes
TYPE=Ethernet
USERCTL=yes
PEERDNS=no
IPV6INIT=no
PERSISTENT_DHCLIENT=yes
HWADDR=00:16:3e:0f:**:** # Use the obtained MAC address of the ENI.
DEFROUTE=no # Specify that the ENI is not the default route. To prevent the default
route of the ECS instance from being changed when you run the ifup command to start t
he secondary ENI, do not set eth1 as the default route.
```

- iii. Restart the network service.
 - Versions earlier than CentOS 7, such as CentOS 6:

service network restart

Cent OS 7 or later and Alibaba Cloud Linux 2:

systemctl restart network

? Note

- After you configure the secondary ENI, you can configure routes for the ENI. For more information, see the Configure routes for a secondary ENI that is bound to an instance that runs an Alibaba Cloud Linux 2 or CentOS 7 operating system section of this topic.
- If you want to create custom images from the instance whose ENIs are configured, you must first run the /etc/eni_utils/eni-cleanup command to remove network configurations from /etc/udev/rules.d/70-persistent-net.rules and /etc/sysconfig/network-scripts/.

Configure a secondary ENI for an instance that runs an Alibaba Cloud Linux 2 operating system (systemd-networkd)

If your instance runs an Alibaba Cloud Linux 2 operating system and uses the systemd-networkd network service, you must manually modify the ENI configuration file to configure an ENI.

1. Open the configuration file of the secondary ENI.

vi /etc/systemd/network/60-eth1.network

2. Add the information of the secondary ENI to the configuration file. Then, save and close the configuration file.

You can assign a dynamic or static IP address to the secondary ENI. You can use one of the following methods based on your requirements. The following section provides an example of the ENI information to add to the configuration file:

• Assign a dynamic IP address to the secondary ENI by using the Dynamic Host Configuration Protocol (DHCP).

```
[Match]
Name=eth1 # Specify the ENI that you want to configure.
[Network]
DHCP=yes
[DHCP]
UseDNS=yes
```

• Assign a static IP address to the secondary ENI.

```
[Match]
Name=eth1 # Specify the ENI that you want to configure.
[Network]
Address=192.168.**.*2/24 # Specify the static IP address and subnet mask to be assign
ed.
```

```
Note In the preceding example, 192.168.**.*2 is the primary private IP address and the /24 subnet mask is 255.255.255.0.
```

3. Restart the network service.

```
systemctl restart systemd-networkd
```

Configure a secondary ENI for an instance that runs a Ubuntu or Debian operating system

If your instance runs a Ubuntu or Debian operating system, you must modify the configuration file of the secondary ENI based on your image version.

- Perform the following operations on an instance that runs Ubuntu 14.04, Ubuntu 16.04, or Debian:
 - i. Open the configuration file of the secondary ENI.

```
vi /etc/network/interfaces
```

ii. Add the information of the secondary ENI to the configuration file. Then, save and close the configuration file.

The following section provides an example of the ENI information to add to the configuration file.

```
auto eth0
iface eth0 inet dhcp
auto eth1 # Specify the ENI that you want to configure.
iface eth1 inet dhcp
```

Note The eth0 primary ENI is configured in the same configuration file as the eth1 secondary ENI. You must add the information of the primary ENI to the configuration file.

- iii. Restart the network service.
 - Versions earlier than Ubuntu 16.04, such as Ubuntu 14.04:

service networking restart

Ubuntu 16.04 and Debian:

systemctl restart networking

The configurations of the secondary ENI can take effect regardless of whether the following alert notification appears. You can run the <u>ip address show</u> command to check whether the IP address of the secondary ENI can be recognized.

root@ecs:~# service networking restart Job for networking.service failed because the control process exited with error code. See "systemctl status networking.service" and "journalctl -xe" for details.

- Perform the following operations on an instance that runs Ubuntu 18.04:
 - i. Open the configuration file of the secondary ENI.

vi /etc/netplan/eth1-netcfg.yaml

ii. Add the information of the secondary ENI to the configuration file. Then, save and close the configuration file.

? Note When you modify the configuration file, take note of the following items:

- The configuration file is in the YAML format. You must follow the YAML syntax rules when you configure the file.
- Tabs cannot be used for indentation in YAML files. Use spaces instead.
- We recommend that you copy information from the default /etc/netplan/99-netcfg. yaml configuration file to prevent format issues.

The following section provides an example of the ENI information to add to the configuration file:

```
network:
  version: 2
  renderer: networkd
  ethernets:
    eth1:
    dhcp4: yes
    dhcp6: no
```

iii. Apply the added configurations.

netplan apply

Configure a secondary ENI for an instance that runs a SUSE or openSUSE operating system

If your instance run a SUSE or openSUSE operating system, you must manually modify the ENI configuration file to configure the secondary ENI.

1. Open the configuration file of the secondary ENI.

- vi /etc/sysconfig/network/ifcfg-eth1
- 2. Add the information of the secondary ENI to the configuration file. Then, save and close the configuration file.

In the following example, a dynamic IP address is assigned to the secondary ENI by using DHCP.

BOOTPROTO='dhcp4' STARTMODE='auto' USERCONTROL='no'

- 3. Restart the network service.
 - Versions earlier than SUSE Linux Enterprise Server 12:

service network restart

• SUSE Linux Enterprise Server 12 or later:

systemctl restart network

Configure routes for a secondary ENI that is bound to an instance that runs an Alibaba Cloud Linux 2 or CentOS 7 operating system

If you manually configure secondary ENIs but do not configure routes for the secondary ENIs or if routes configured by the multi-nic-util tool do not meet your requirements, perform the following steps to configure routes:

1. View the route information.

route -n

Sample command outputs:

• The following command output shows only the route information of the eth0 primary ENI, which indicates that no routes are configured for the eth1 secondary ENI.

[root@ecs ~]# route -n										
Kernel IP routing table										
Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface			
0.0.0	192.168.	0.0.0.0	UG	0	0	0	eth0			
169.254.	0.0.0.0	255.255.	U	1002	0	0	eth0			
192.168.	0.0.0	255.255.	U	0	0	0	eth0			

• The following command output shows the route information of both the eth0 primary ENI and the eth1 secondary ENI. If the configured routes do not meet your requirements, you can modify the route configurations.

root@ecs ~]# route -n							
Kernel IP routing table							
Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
0.0.0	192.168.	0.0.0.0	UG	0	0	0	eth0
0.0.0	192.168.	0.0.0.0	UG	1001	0	0	eth1
169.254	0.0.0	255.255	U	1002	0	0	eth0
169.254	0.0.0.0	255.255	U	1003	0	0	eth1
192.168	0.0.0.0	255.255	U	0	0	0	eth0
192.168	0.0.0.0	255.255	U	0	0	0	eth1

2. Plan the default route based on your requirements.

In this example, the sample values listed in the following table are used.

Secondary ENI attribute	Sample value
ENI name	eth1
Primary private IP address	192.168.**.*2
Gateway address	192.168.**.253
metric	1001

3. Configure the default route.

You can run the following commands to add the default route for the eth1 secondary ENI, create a route table, and then attach a routing policy to the table. In this example, *table 1001* is created as the route table. We recommend that you keep the name of the route table the same as the metric value in the default route of the ENI. *192.168.**.253* is the gateway address and *192.168.*.*2* is the primary private IP address of the eth1 secondary ENI.

ip -4 route add default via 192.168.**.253 dev ethl metric 1001 && $\$ ip -4 route add default via 192.168.**.253 dev ethl table 1001 && $\$ ip -4 rule add from 192.168.**.*2 lookup 1001

4. View the created route table and routing policy.

```
ip route list table 1001 && \
ip rule list
```

The following figure shows that the route table and routing policy are created.

[root@ecs ~]# ip route list table 1001 & > ip rule list	& \				
default via 192.168253 dev eth1					
0: trom all lookup local					
32765: from 192.168.47 12 lookup 1001					
32766: from all lookup main					
32767: from all lookup default					
[root@ecs ~]#					

5. Configure routes to automatically update on instance startup.

After you perform the preceding steps to configure routes for the eth1 secondary ENI, you must perform the following steps to configure the routes to automatically update on instance startup. Otherwise, the routes become invalid after the instance is restarted.

i. Open the /etc/rc.local file.

vim /etc/rc.local

ii. Add the configuration information of the routes to the */etc/rc.local* file. Then, save and close the file.

```
ip -4 route add default via 192.168.**.253 dev eth1 metric 1001
ip -4 route add default via 192.168.**.253 dev eth1 table 1001
ip -4 rule add from 192.168.**.*2 lookup 1001
```

iii. Grant execution permissions on the /etc/rc.local file.

```
chmod +x /etc/rc.local
```
Configure routes for a secondary ENI that is bound to an instance that runs a CentOS 8 operating system

If routes configured by the system do not meet your requirements, perform the following steps to configure routes:

1. View the route information.

route -n

The following figure shows the route information of both the eth0 primary ENI and the eth1 secondary ENI. If the configured routes do not meet your requirements, you can modify the route configurations.

I rootkees alt r	outo _n						
[lootwees ~]# loute -n							
Kernel IP routi	Kernel IP routing table						
Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
0.0.0	192.168.	0.0.0	UG	100	0	0	eth0
0.0.0	192.168.	0.0.0	UG	101	0	0	eth1
192.168.	0.0.0	255.255.	U	100	0	0	eth0
192.168.	0.0.0	255.255.	U	101	0	0	eth1
[root@ecs ~]#							

2. Plan the default route based on your requirements.

In this example, the sample values listed in the following table are used.

Secondary ENI attribute	Sample value
ENI name	eth1
Primary private IP address	192.168.**.*2
Gateway address	192.168.**.253
table	1001

- 3. Create a script for configuring routes.
 - i. Create the */root/route.sh* file and open it.

ii. Add the configuration information of the routes to the */root/route.sh* file. Then, save and close the file.

The following section shows how to create a route table and attach a routing policy to the route table for the eth1 secondary ENI. In this example, *table 1001* is created as the route table. *192.168.**.253* is the gateway address and *192.168.*.*2* is the primary private IP address of the eth1 secondary ENI.

```
#!/bin/bash
i=0
while true; do
       /usr/sbin/ip -4 route add default via 192.168.**.253 dev eth1 table 1001
       if [ $? -eq 0 ]; then
               break
    fi
       sleep 3
       let i++
       if [ $i -gt 10 ]; then
             exit -1
       fi
done
i=0
while true; do
       /usr/sbin/ip -4 rule add from 192.168.**.*2 lookup 1001
       if [ $? -eq 0 ]; then
               break
    fi
       sleep 3
       let i++
       if [ $i -gt 10 ]; then
              exit -1
       fi
done
```

4. Configure the default route.

sh /root/route.sh

5. View the created route table and routing policy.

```
ip route list table 1001 && \
ip rule list
```

The following figure shows that the route table and routing policy are created.

6. Configure routes to automatically update on instance startup.

After you perform the preceding steps to configure routes for the eth1 secondary ENI, you must perform the following steps to configure the routes to automatically update on instance startup. Otherwise, the routes become invalid after the instance is restarted.

i. Open the /etc/rc.local file.

vim /etc/rc.local

ii. Add the configuration information of the routes to the */etc/rc.local* file. Then, save and close the file.

sh /root/route.sh

iii. Grant execution permissions on the /etc/rc.local file.

chmod +x /etc/rc.local

Related information

DescribeNetworkInterfaces

6.6. Assign secondary private IP addresses

You can assign one or more secondary private IP addresses to a primary or secondary elastic network interface (ENI). This topic describes how to assign secondary private IP addresses and configure secondary private IP addresses in an Elastic Compute Service (ECS) instance.

Context

Secondary private IP addresses are suitable for the following scenarios:

- Scenarios that involve multiple applications: If your instance hosts multiple applications, you can
 assign secondary private IP addresses to the applications so that each application can use a separate
 IP address for outbound connections. This way, a single instance can provide multiple services
 optimally.
- Failover scenarios: If an instance fails, you can unbind ENIs from the instance and bind the ENIs to another instance so that traffic destined for the secondary private IP addresses of the failed instance is diverted to the normal instance. This ensures service continuity.

When you assign secondary private IP addresses, take note of the following limits:

- Limits on security groups: A limited number of private IP addresses can be contained in a security group of the Virtual Private Cloud (VPC) type. For more information, see the "Security group limits" section in Limits.
- Limits on ENIs: The maximum number of private IP addresses that can be assigned to an ENI varies based on the state of the ENI.
 - For an ENI in the Available state, up to 10 private IP addresses can be assigned.
 - For an ENI in the **Bound** state, the maximum number of private IP addresses that can be assigned is subject to the instance type of the associated instance. For more information, see Instance family.

Procedure

- 1. Make sure that the following prerequisites are satisfied:
 - The instance to which an ENI is bound is in the **Running** (Running) state. For more information, see **Start** an instance.
 - A secondary ENI is bound to the instance. For more information, see Bind an ENI.

2. In the ECS console, assign secondary private IP addresses to an ENI.

You can assign secondary private IP addresses to an ENI on the Network Interfaces, Instances, or Security Groups page. For more information, see the following sections:

- Assign secondary private IP addresses on the Network Interfaces page
- Assign secondary private IP addresses on the Instances page
- Assign secondary private IP addresses on the Security Groups page
- 3. In the instance to which the ENI is bound, configure the assigned secondary private IP addresses.

This topic describes how to configure secondary private IPv4 addresses. Operations vary based on the operating system type and the IP address type. For more information, see the following sections:

- Configure secondary private IPv4 addresses in a Windows instance
- Configure secondary private IPv4 addresses in a Linux instance that runs a Red Hat Enterprise Linux (RHEL) operating system
- Configure secondary private IPv4 addresses in a Linux instance that runs a Debian operating system
- Configure secondary private IPv4 addresses in a Linux instance that runs a SUSE Linux Enterprise Server (SLES) operating system

(?) Note For more information about how to configure secondary private IPv6 addresses, see Step 4: Configure an IPv6 address for Windows instances and Step 4: Configure an IPv6 address for Linux instances.

Assign secondary private IP addresses on the Network Interfaces page

1.

2.

3.

- 4. On the **Network Interfaces** page, find the ENI to which you want to assign secondary private IP addresses and click **Manage Secondary Private IP Address** in the **Actions** column.
- 5. In the Manage Secondary Private IP Address dialog box, click Assign New IP to assign IP addresses based on your business needs.
 - To automatically assign IP addresses, accept the default Auto-assign value. Then, the system randomly assigns IP addresses from within the private CIDR blocks in the IPv4 Private CIDR Block and IPv6 Private CIDR Block values of the ENI.
 - To manually assign IP addresses, enter specific IP addresses from within the private CIDR blocks in the IPv4 Private CIDR Block and IPv6 Private CIDR Block values of the ENI.

dresses.	the ENI. If you leave IP address fie	INI. You can also click Assign New IP adds empty, the system automatically	to assign
Instance:	i-tay 128-bit Societana ang si		
ENI ID:	eni- hari katu k atu katu		
Primary Private IP Address:	192.		
IPv4 Private CIDR Block:	192		
IPv6 Private CIDR Block:	2408:4(4		
IPv4 Addresses:	The current ENI supports up to including 1 primary private IP a addresses.	9 6 private IPv4 addresses, address and 5 secondary private IP	
	Auto-assign	Unassign	
	Assign New IP		
IPv6 Address:	The current ENI supports up to	1 IPv6 addresses.	
	2408:4005:340:c800: Auto-assig	gn Unassign	
IPv6 Address: have made the follow Automatically assign	addresses. Auto-assign Assign New IP The current ENI supports up to 2408:4005:340:c800: Auto-assign wing changes: ted 1 IPv4 addresses.	unassign o 1 IPv6 addresses. gn Unassign	

6. Click **OK**.

Assign secondary private IP addresses on the Instances page

When you assign secondary private IP addresses for an instance on the Instances page, the IP addresses are assigned to the primary ENI of the instance.

- 1.
- 2.
- 3.
- 4. On the Instances page, find the instance for which you want to assign secondary private IP addresses and choose More > Network and Security Group > Manage Secondary Private IP Address in the Actions column.
- 5. In the Manage Secondary Private IP Address dialog box, click Assign New IP to assign IP addresses based on your business needs.
 - To automatically assign IP addresses, accept the default Auto-assign value. Then, the system randomly assigns IP addresses from within the private CIDR blocks in the IPv4 Private CIDR Block and IPv6 Private CIDR Block values of the ENI.
 - To manually assign IP addresses, enter specific IP addresses from within the private CIDR blocks in the IPv4 Private CIDR Block and IPv6 Private CIDR Block values of the ENI.

You can assign new IP ad the IP address field empt	Idresses to the instance or modify and delete y and click Assign New IP, the system automa	existing IP addresses. (If you leav atically assigns an IP address.)
Instance:	ity manufactures.com	
Primary ENI ID:	eni-	
Primary Private IP Address :	192.196.00	
IPv4 Private CIDR Block:	192.	
IPv6 Private CIDR Block:	2408	
IPv4 Addresses:	The current ENI supports up to 6 privat including 1 primary private IP address a addresses.	e IPv4 addresses, and 5 secondary private IP
	Auto-assign	Unassign
	Assign New IP	
IPv6 Address:	The current ENI supports up to 1 IPv6 a	ddresses.
IPv6 Address:	The current ENI supports up to 1 IPv6 a 2408:4005:340:c800: Auto-assign	ddresses.

6. Click OK.

Assign secondary private IP addresses on the Security Groups page

- 1.
- 2.
- 3.
- 4. Find a security group of the ENI to which you want to assign secondary private IP addresses and click **Manage ENIs** in the **Actions** column.
- 5. On the ENIs in Security Group page, find the ENI to which you want to assign secondary private IP addresses and click Manage Secondary Private IP Address in the Actions column.
- 6. In the Manage Secondary Private IP Address dialog box, click Assign New IP to assign IP addresses based on your business needs.
 - To automatically assign IP addresses, accept the default Auto-assign value. Then, the system randomly assigns IP addresses from within the private CIDR blocks in the IPv4 Private CIDR Block and IPv6 Private CIDR Block values of the ENI.
 - To manually assign IP addresses, enter specific IP addresses from within the private CIDR blocks in the IPv4 Private CIDR Block and IPv6 Private CIDR Block values of the ENI.

/ou can modify or unassi assign new IP addresses t P addresses.	gn the existing IP addresses of this ENI. You o to the ENI. If you leave IP address fields empt	can also click Assign New IP to ty, the system automatically assigr
Instance:	i-ty-Weddheithendaryd	
ENI ID:	eni- by' beilagiturg Sichera t	
Primary Private IP Address :	192.	
IPv4 Private CIDR Block:	192	
IPv6 Private CIDR Block:	2408:40 4	
IPv4 Addresses:	The current ENI supports up to 6 privat including 1 primary private IP address a addresses.	te IPv4 addresses, and 5 secondary private IP
	Auto-assign	Unassign
	Assign New IP	
IPv6 Address:	The current ENI supports up to 1 IPv6 a	addresses.
	2408:4005:340:c800: Auto-assign	Unassign

7. Click **OK**.

Configure secondary private IPv4 addresses in a Windows instance

1.

- 2. View the subnet mask and default gateway of the instance.
 - i. Open Command Prompt or Windows PowerShell.
 - ii. Run the <code>ipconfig</code> command to view the subnet mask and default gateway of the instance.

A command output similar to the following one is displayed.

PS C:\Users\Administrator> ipconfig
Tindows IP 配置
以太网道翻器 以太网:
连接特定的 DNS 后缀
以太网话配器 vBthernet (nat):
连接特定的 DNS 后缀

- 3. Open Network and Sharing Center.
- 4. Click Change adapter settings.
- 5. Double-click the network connection in use. In this example, the network connection named **Ethernet** is used. Click **Properties** in the **Ethernet Status** dialog box.
- 6. In the Ethernet Properties dialog box, double-click Internet Protocol Version 4 (TCP/IPv4).

- 7. In the Internet Protocol Version 4 (TCP/IPv4) Properties dialog box, select Use the following IP address and click Advanced...
- 8. In the Advanced TCP/IP Settings dialog box, configure IP addresses.
 - i. In the IP addresses section, click Add... and enter one of the assigned IP addresses in the IP address field and the obtained subnet mask in the Subnet mask field.

You can repeat this step to add multiple IP addresses to the same adapter.

Advanced T	CP/IP Settings				
IP Settings DNS WINS					
IP addresses					
IP address	Subnet mask				
172	255.255.0.0				
172 3	255.255.0.0				
<u>A</u> dd	Edit Remo <u>v</u> e				
Gateway	Metric				
A <u>d</u>	Ediţ Remove				

- ii. In the **Default gateways** section, click **Add**... and enter the obtained default gateway in the **Default gateway** field.
- 9. Click **OK**.
- 10. Run the *ipconfig* command to check the configuration result.

The following figure shows the output of an example command used to configure two secondary private IP addresses.

PS C:\Users\Administrator> ipconfig
Windows IP 酌置
이 수정권하고 이 수정·
IPv6 地址
本地链接 IPv6 地址 fe80::4080:146f:%9
IPv4 地址
十四進約
1794 地址
子网摘码 255 255
野认网关
192. 168 53
以太网话武器 vBthernet (nat):
连接特定的 DMS 后缀 :
本地链接 IPv6 地址 : fe80::9428:27b
IPv4 地址
子网摘码
默认网关

? Note

If a Windows instance cannot access the Internet after you configure secondary private IP addresses for the instance, troubleshoot the problem by following the instructions in After I configure a secondary private IP address for a Windows instance, the instance cannot connect to the Internet. Why?

Configure secondary private IPv4 addresses in a Linux instance that runs a Red Hat Enterprise Linux (RHEL) operating system

Before you perform the following procedure, take note of the following items:

- This procedure applies to the following operating systems: Alibaba Cloud Linux 2, Alibaba Cloud Linux 3, Cent OS 6, Cent OS 7, Cent OS 8, Red Hat 6, Red Hat 7, Red Hat 8, Anolis 7, Anolis 8, Fedora 33, and Fedora 34.
- In the following example, the *eth0* primary ENI is used. If you are working with a secondary ENI, modify the ENI ID.

1.

2. Run the ifconfig command to view the subnet mask and run the route -n command to view the default gateway.

The following figure shows the output of an example command.



In the preceding command output, 255.255.**.** that corresponds to netmask is the IPv4 subnet mask and 192.**.**.253 that corresponds to Gateway is the default gateway.

Once If the Linux distribution used by the instance does not support the ifconfig command, run the ip a or ip addr show command instead.

- 3. Modify the network configuration file.
 - To configure a single private IPv4 address, run the vi /etc/sysconfig/network-scripts/ifcfg-e th0:0 command to add the corresponding configuration items.

Example:

```
DEVICE=eth0:0
TYPE=Ethernet
BOOTPROTO=static
ONBOOT=yes
IPADDR=<Assigned secondary private IPv4 address 1>
NETMASK=<IPv4 subnet mask>
```

• To configure multiple private IPv4 addresses, increment the sequence number in the DEVICE value

and continue to add configuration items.

For example, run the vi /etc/sysconfig/network-scripts/ifcfg-eth0:1 command to add the following configuration items:

DEVICE=eth0:1 TYPE=Ethernet BOOTPROTO=static ONBOOT=yes IPADDR=<Assigned secondary private IPv4 address 2> NETMASK=<IPv4 subnet mask>

4. Run the corresponding command based on the operating system for the configuration to take effect.

Operating system	Command	
 Alibaba Cloud Linux 2 Cent OS 7 Red Hat 7 Anolis 7 	<pre>Run one of the following commands to restart the network service: service network restart systemctl restart network</pre>	
CentOS 6Red Hat 6	Run the service network restart command to restart the network service.	
 Alibaba Cloud Linux 3 Cent OS 8 Red Hat 8 Anolis 8 Fedora 33/34 	 Perform the following operations: Run the systemctl restart NetworkMana ger command to restart the network service. Run the nmcli device reapply eth0 command to restart the eth0 ENI, or run the reboot command to restart the instance. 	

5. Run the ifconfig command to check the configuration result.

The following figure shows the output of an example command used to configure two secondary private IP addresses.



Configure secondary private IPv4 addresses in a Linux instance that runs a Debian operating system

Before you perform the following procedure, take note of the following items:

- This procedure applies to the following operating systems: Ubuntu 18, Ubuntu 20, Ubuntu 14, Ubuntu16, Debian 8, Debian9, and Debian10.
- In the following example, the *eth0* primary ENI is used. If you are working with a secondary ENI, modify the ENI ID.

1.

2. Run the ifconfig command to view the subnet mask and run the route -n command to view the default gateway.

The following figure shows the output of an example command.

[root@ecs ~]#	ifconfig						
eth0: flags=42	L63 <up,broadcast,r< td=""><td>UNNING, MULTICAST</td><td>> mtu</td><td>1500</td><td></td><td></td><td></td></up,broadcast,r<>	UNNING, MULTICAST	> mtu	1500			
inet :	192.168 net	mask 255.255.	🔳 bro	oadcast	192.16	8.	300
inet6 fe80::216:3eff: prefixlen 64 scopeid 0x20 <link/>							
ether 00:16:3e:1b txqueuelen 1000 (Ethernet)							
RX packets 52842 bytes 77066906 (73.4 MiB)							
RX er	RX errors Ø dropped Ø overruns Ø frame Ø						
тх рас	TX packets 4983 bytes 560229 (547.0 KiB)						
TX er	rors 0 dropped 0	overruns 0 carr	ier Ø	collisi	ions 0		
lo: flags=73<	JP,LOOPBACK,RUNNIN	G> mtu 65536					
inet :	127.0 netmask	255.0					
inet6	::1 prefixlen 12	8 scopeid 0x10<	host≻				
loop	txqueuelen 1000	(Local Loopback)					
RX pa	BX packets 0 bytes 0 (0.0 B)						
BX errors 0 dropped 0 overruns 0 frame 0							
TX packets 0 bytes 0 (0.0 B)							
TX er	rors 0 dropped 0	overruns 0 carr	ier Ø	collisi	ions 0		
[root@ecs ~]#	route -n						
Kernel IP rou	ting table						
Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
0.0.0.0	192. 253	0.0.0.0	UG	0	0	0	eth0
169.254	0.0.0	255.255.	U	1002	0	0	eth0
192.168.	0.0.0.0	255.255.	U	0	0	0	eth0

In the preceding command output, 255.255.**.** that corresponds to netmask is the IPv4 subnet mask and 192.**.**.253 that corresponds to Gateway is the default gateway.

Once If the Linux distribution used by the instance does not support the ifconfig command, run the ip a or ip addr show command instead.

- 3. Configure secondary private IP addresses based on the operating system of your instance.
 - Debian series: Ubunt u 18 and Ubunt u 20
 - a. Disable the network configuration feature of cloud-init for the instance. Run the vim /etc/ cloud/cloud.cfg.d/99-disable-network-config.cfg command and add the corresponding configuration items.

Example:

network: {config: disabled}

b. Run the vim /etc/netplan/50-cloud-init.yaml command to open the network configuration file and change the IP addresses that are configured by using Host Configuration Protocol (DHCP) to static IP addresses.

Example:

```
network:
version: 2
ethernets:
eth0:
match:
macaddress: 00:16:3e:36:**:**
addresses:
- <Primary private IPv4 address>/<Subnet mask bit>
- <Assigned secondary private IPv4 address 1>/<Subnet mask bit>
- <Assigned secondary private IPv4 address 2>/<Subnet mask bit>
gateway4: <Default gateway>
```

(?) Note In the example, <Subnet mask bit> is replaced with the subnet mask bit corresponding to the subnet mask. For example, the 255.255.255.0 subnet mask corresponds to the /24 subnet mask bit.

- c. Run the netplan apply command to restart the network service.
- d. Run the ip -a command to check the configuration result.

The following figure shows the output of an example command used to configure two secondary private IP addresses.

root@ecs:~# ip a
1: lo: <loopback,up,lower up=""> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000</loopback,up,lower>
link/loopback 00:00:00:00 🗰 brd 00:00:00:00:
inet 127.0. 3/8 scope host lo
valid lft forever preferred lft forever
inet6 ::1/128 scope host
valid lft forever preferred lft forever
2: eth0: <broadcast,multicast,up,lower_up> mtu 1500 qdisc mq state UP group default qlen 1000</broadcast,multicast,up,lower_up>
link/ether 00:16:3e:0f: brd ff:ff:ff:ff:
inet 192.168. 📭 💷/24 brd 192.168.45.255 scope global eth0
valid_lft forever preferred_lft forever
inet 192.168. // // // // // // // // // // // // //
valid lft forever preferred_lft forever
inet 192.168.
valid_lft forever preferred_lft forever
inet6 2408:4005:325:a206:
valid_lft 127787sec preferred_lft 84587sec
inet6 fe80::216:3eff 🚺 3/64 scope link
valid 1ft forever preferred 1ft forever

- Debian series: Ubuntu 14, Ubuntu 16, Debian 8, Debian 9, and Debian 10.
 - a. Run the vi /etc/network/interfaces command to open the network configuration file and add the corresponding configuration items.

Example:

```
auto eth0:0
iface eth0:0 inet static
address <Assigned secondary private IPv4 address 1>
netmask <IPv4 mask>
auto eth0:1
iface eth0:1 inet static
address <Assigned secondary private IPv4 address 2>
netmask <IPv4 mask>
```

- b. Run the reboot command to restart the instance.
- c. Run the ifconfig command to check the configuration result.

The following figure shows the output of an example command used to configure two secondary private IP addresses.

root@ecs:	"# ifconfig
eth0	Link encap:Ethernet HWaddr 00:16:3e:0f:
	inet addr:192.168 🖬 🗰 Bcast:192.168. 🖬 🖬 Mask:255.255.
	UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
	RX packets:4146 errors:0 dropped:0 overruns:0 frame:0
	TX packets:2373 errors:0 dropped:0 overruns:0 carrier:0
	collisions:0 txqueuelen:1000
	RX bytes:4046275 (4.0 MB) TX bytes:299061 (299.0 KB)
eth0:0	Link encap:Ethernet HWaddr 00:16:3e:0f:🖿 💷
	inet addr:192.168. Bcast:192.168. Mask:255.255.
	UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
eth0:1	Link encap:Ethernet HWaddr 00:16:3e:0f:
	inet addr:192.168. 🖬 🗰 Bcast:192.168. 🖬 🖬 Mask:255.255.
	UP BROADCAST RUNNING MULTICAST MIU:1500 Metric:1

Configure secondary private IPv4 addresses in a Linux instance that runs a SUSE Linux Enterprise Server (SLES) operating system

Before you perform the following procedure, take note of the following items:

- This procedure applies to the following operating systems: SUSE 11, SUSE 12, SUSE 15, OpenSUSE 15, and OpenSUSE 42.
- In the following example, the *eth0* primary ENI is used. If you are working with a secondary ENI, modify the ENI ID.
 - 1.
 - 2. Run the ifconfig command to view the subnet mask and run the route -n command to view the default gateway.

The following figure shows the output of an example command.



In the preceding command output, 255.255.**.** that corresponds to netmask is the IPv4 subnet mask and 192.**.**.253 that corresponds to Gateway is the default gateway.

Once If the Linux distribution used by the instance does not support the ifconfig command, run the ip a or ip addr show command instead.

3. Run the vi /etc/sysconfig/network/ifcfg-eth0 command to open the network configuration file and add the following configuration items:

```
IPADDR_0=<Assigned secondary private IPv4 address 1>
NETMASK_0=<IPv4 subnet mask>
LABEL_0='0'
IPADDR_1=<Assigned secondary private IPv4 address 2>
NETMASK_1=<IPv4 subnet mask>
LABEL_1='1'
```

- 4. Run the service network restart or systemctl restart network command to restart the network service.
- 5. Run the ifconfig command to check the configuration result.

The following figure shows the output of an example command used to configure two secondary private IP addresses.

ec	s:~ # The addition and a supervision and a super
1:	lo: <loopback,up,lower_up> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1</loopback,up,lower_up>
	link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
	inet 127.0. In brd 127.255. State scope host lo
	valid_lft_forever_preferred_lft_forever
2:	eth0: <broadcast,multicast,up,lower_up> mtu 1500 qdisc mq state UP group default qlen 1000</broadcast,multicast,up,lower_up>
	link/ether 00:16:3e:0f: 🖛 🔤 brd ff:ff:ff:ff: 🖬 💵
	inet 192.168. 7/24 brd 192.168. 5 scope global eth0
	valid Ift torever preferred Ift forever
	inet 192.168. 9/24 brd 192.168. scope global secondary eth0:0
	valid lft forever preferred lft forever
	inet 192.168. 0/24 brd 192.168. scope global secondary eth0:1
	valid_lft_forever_preferred_lft_forever
3:	eth1: <broadcast,multicast> mtu 1500 qdisc noop state DOWN group default qlen 1000</broadcast,multicast>
Γ	link/ether 00:16:3e:12: brd ff:ff:ff:ff:ff:ff

Related information

AssignPrivatelpAddresses

6.7. Unassign secondary private IP addresses

When an elastic network interface (ENI) no longer needs one or more secondary private IP addresses, you can unassign the addresses from the ENI.

Prerequisites

The following requirements are met:

- One or more secondary private IP address are assigned to the ENI from which you want to unassign secondary private IP addresses.
- The ENI is in the Available (Available) or Bound (InUse) state.
- When you unassign secondary private IP addresses from a primary ENI, the Elastic Compute Service (ECS) instance to which the primary ENI is attached is in the **Running** (Running) or **Stopped** (Stopped) state.

Unassign secondary private IP addresses on the Network Interfaces page

- 1.
- 2.
- 3.
- 4. On the **Network Interfaces** page, find the ENI from which you want to unassign secondary private IP addresses and click **Manage Secondary Private IP Address** in the **Actions** column.
- 5. In the **Manage Secondary Private IP Address** dialog box, find the secondary private IP addresses that you want to unassign and click **Unassign** for each of the IP addresses.
- 6. Click OK.

Unassign secondary private IP addresses on the Instances page

When you unassign secondary private IP addresses for an instance on the Instances page of the ECS console, the secondary private IP addresses are unassigned from the primary ENI of the instance.

- 1.
- 2.
- 3.
- 4. On the **Instances** page, find the instance for which you want to unassign secondary private IP addresses and choose **More > Network and Security Group > Manage Secondary Private IP Address** in the **Actions** column.
- 5. In the **Manage Secondary Private IP Address** dialog box, find the secondary private IP addresses that you want to unassign and click **Unassign** for each of the IP addresses.
- 6. Click OK.

Unassign secondary private IP addresses on the Security Groups page

- 1.
- 2.
- 3.
- 4. Find the security group to which the ENI that you want to manage belongs and click **Manage ENIs** in the **Actions** column.
- 5. On the ENIs in Security Group page, find the ENI from which you want to unassign secondary private IP addresses and click Manage Secondary Private IP Address in the Actions column.
- 6. In the **Manage Secondary Private IP Address** dialog box, find the secondary private IP addresses that you want to unassign and click **Unassign** for each of the IP addresses.
- 7. Click OK.

Related information

• UnassignPrivatelpAddresses

6.8. Modify an ENI

This topic describes how to modify a primary or secondary elastic network interface (ENI). You can change the security group of a primary ENI by moving its bound Elastic Compute Service (ECS) instance to a different security group. You can modify the attributes of a secondary ENI, such as the name, security group, and description.

Prerequisites

Before you add an primary ENI to a new security group, make sure that the primary ENI belongs to the same virtual private cloud (VPC) and the same zone as the new security group. For more information, see Overview.

Context

If you want to change the security group of an ENI, take note of the following limits on the ENI and its bound ECS instance:

- An ECS instance cannot belong to both a basic and an advanced security group at the same time.
- An ENI cannot belong to both a basic and an advanced security group at the same time.
- An ENI can only be bound to an ECS instance when they belong to the same type of security groups.

For more information about security groups, see Overview.

Modify a primary ENI on the Security Groups page

The primary ENI and the secondary ENIs of an ECS instance can belong to different security groups. If you move the ECS instance to a different security group, the primary ENI is also moved to the security group, but the secondary ENIs are not. Perform the following steps to change the security group to which the primary ENI belongs on the Security Groups page.

- 1.
- 2.
- 3.
- 4. Find the security group that you want to change and click **Manage Instances** in the **Actions** column.
- 5. On the **Instances in Security Group** page, change the security group to which the primary ENI belongs:
 - Perform the following steps to add the primary ENI to a new security group:
 - a. In the upper-right corner of the Instances in Security Group page, click Add Instance.
 - b. In the Add Instance dialog box, select the ID of the instance to which the primary ENI is bound. Click OK.

The primary ENI of the selected instance is also added to the new security group.

- Perform the following steps to remove the primary ENI from its current security group:
 - a. On the **Instances in Security Group** page, select one or more instances and click **Remove from Security Group**.
 - b. In the Remove ECS Instances from Security Group message, click OK.

The primary ENIs of the selected instances are also removed from the current security group. The primary ENI and the ECS instance must belong to at least one security group.

After the modification, choose **Network & Security > ENIs** in the left-side navigation pane. If the security group of a primary ENI is changed, the new security group is displayed in the Security Group

ID column corresponding to the ENI.

Modify a secondary ENI on the Network Interfaces page

Perform the following steps to modify the name, security group, and description of the secondary ENI:

- 1.
- 2.
- 3.
- 4. Find the secondary ENI that you want to modify and click **Modify** in the Actions column.
- 5. In the **Modify** dialog box, modify the following ENI attributes:
 - ENI Name: Specify a new ENI name based on the naming conventions displayed under this field.
 - **Security Group**: Select a new security group for the ENI, or remove the ENI from a security group. The ENI must belong to at least one security group.
 - **Description**: Modify the description as prompted.
- 6. Click OK.

Refresh the ENI list. If the secondary ENI is modified, you can you can find its new attributes in the corresponding columns.

Modify a secondary ENI on the Security Groups page

Perform the following steps to modify the name, security group, or description of a secondary ENI:

- 1.
- 2.
- 3.
- 4. Find the security group of the secondary ENI that you want to modify and click **Manage Instances** in the **Actions** column.
- 5. On the ENIs in Security Group page, find the secondary ENI that you want to modify and click Manage Secondary Private IP Address in the Actions column.
- 6. In the **Modify** dialog box, modify the following ENI attributes:
 - ENI Name: Specify a new ENI name based on the naming conventions displayed under this field.
 - **Security Group**: Select a new security group for the ENI, or remove the ENI from a security group. The ENI must belong to at least one security group.
 - **Description**: Modify the description as prompted.
- 7. Click OK.

After the modification, choose **Network & Security > ENIs** in the left-side navigation pane. If the secondary ENI is modified, you can you can find its new attributes in the corresponding columns.

Related information

References

• ModifyNetworkInterfaceAttribute

6.9. Edit the tags of an ENI

Tags can be used to identify resources with the same characteristics (such as elastic network interfaces that belong to the same organization or that serve the same purpose) for easy search and management. This topic describes how to edit the tags of an existing elastic network interface (ENI).

Context

For information about how to use tags, the resources that support tags, and the limits on tags, see Overview and the "Tag limits" section of the Limits topic.

Procedure

1.

2.

3.

4. Find the ENI whose tags you want to edit, move the pointer over the 💿 icon in the Tag column,

and then click Edit Tags.

5. In the Edit Tags dialog box, click Available Tags to select existing tags or click Create to create tags. Then, click OK.

What's next

After tags are added to your ENIs, you can filter the ENIs by tag to perform different operations. For example, you can assign secondary private IP addresses to ENIs that have a set of tags and detach ENIs that have a different set of tags.

6.10. Unbind an ENI

This topic describes how to unbind a secondary elastic network interface (ENI) from your Elastic Compute Service (ECS) instance.

Prerequisites

Before you unbind an ENI, make sure that the following requirements are met:

- The ENI to be unbound is a secondary ENI. Primary ENIs cannot be unbound from instances.
- The instance from which you want to unbind an ENI is in the **Stopped** or **Running** state.

(2) Note ENIs can be unbound from instances of specific instance types only when the instances are in the Stopped state. For more information about instance types, see the "Instance types for which instances must be in the Stopped state" section in Instance types for which instances must be in the Stopped state.

Context

After ENIs are created, deleted, bound to instances, or unbound from instances, ENI operation events are triggered. You can use CloudMonitor to set notifications of ENI operation events to obtain the ENI operation results such as whether ENIs are created. For more information, see ENI operation event notifications.

Unbind an ENI on the Network Interfaces page

1.

2.

- 3.
- 4. Find the ENI in the InUse state that you want to unbind and click Unbind in the Actions column.
- In the Unbind message, confirm the information and click OK.
 If the state of the ENI changes to Available after you refresh the ENI list, the ENI is unbound from the instance.

Unbind an ENI on the Security Groups page

1.

2.

- 3.
- 4. Find the security group that contains the ENI that you want to unbind and click **Manage ENIs** in the **Actions** column.
- 5. On the ENIs in Security Group page, find the ENI in the InUse state and click Unbind in the Actions column.
- 6. In the **Unbind** message, confirm the information and click **OK**. If the state of the ENI changes to **Available** after you refresh the ENI list, the ENI is unbound from the instance.

What to do next

You can perform the following operations on an ENI that is in the **available** state:

- Bind an ENI
- Delete an ENI
- Modify an ENI

Related information

• DetachNetworkInterface

6.11. Delete an ENI

If an elastic network interface (ENI) is no longer needed, you can delete it. This topic describes how to delete an ENI. Secondary ENIs can be deleted, whereas primary ENIs cannot.

Prerequisites

- The ENI that you want to delete is in the Available state.
- If the ENI that you want to delete was bound to an Elastic Compute Service (ECS) instance, the ENI has been unbound from the instance. For more information, see Unbind an ENI.

⑦ Note When an instance is released, the ENIs bound to the instance are deleted.

Context

When an ENI is deleted, the following results occur:

• All private IP addresses of the ENI are automatically released.

• The ENI is removed from all security groups to which it was added.

Delete an ENI on the Network Interfaces page

1.

- 2.
- 3.
- 4. On the Network Interfaces page, find an ENI that is in the **Available** state and click **Delete** in the **Actions** column.
- In the message that appears, click OK.
 On the Network Interfaces page, refresh the ENI list. If the ENI is deleted, it is no longer displayed.

Delete an ENI on the Security Groups page

1.

2.

3.

- 4. On the Security Groups page, find the security group to which an ENI belongs and click Manage ENIs in the Actions column.
- 5. On the ENIs in Security Group page, find the ENI in the Available state and click Delete in the Actions column.
- 6. In the message that appears, click **OK**. On the Network Interfaces page, refresh the ENI list. If the ENI is deleted, it is no longer displayed.

Related information

• DeleteNetworkInterface

6.12. Have ENI operations automatically performed in response to ENI operation events

After elastic network interfaces (ENIs) are created, deleted, bound, or unbound, ENI operation events are triggered. You can use Alibaba Cloud CloudMonitor and Message Service (MNS) to have ENI operations automatically performed in response to notifications of ENI operation events. This topic describes how to have your ENI operations automatically performed in response to ENI operation events by using Java code.

Context

CloudMonitor can work with MNS to configure event notifications. When CloudMonitor that works with MNS is used to monitor ENI operation events, CloudMonitor can synchronize event notifications to MNS. You can obtain the event information from MNS queues and then have subsequent operations automatically performed based on the notifications. For best practice, you can perform the following steps:

1. Create an MNS queue and a CloudMonitor alert rule to monitor and collect an ENI operation event. For more information, see the Preparations section of this topic. 2. Configure ENI operations to be automatically performed. In this topic, Java code is used to poll every message in the MNS queue and return the information of the ENI operation event in the messages. For more information, see the Use Java code to have an ENI operation event automatically handled section of this topic.

Onte The ENI operation event feature is in invitational preview. To use this feature, submit a ticket.

Preparations

1. Create an MNS queue.

This topic elaborates only steps that are required by this best practice. For more information about how to create MNS queues, see Create a queue.

- i. Log on to the MNS console.
- ii. In the left-side navigation pane, click Queues.
- iii. In the top navigation bar, select a region.

The queue must be created in the same region as the ENI that you want to monitor. For example, if the ENI that you want to monitor is created in the China (Hangzhou) region, you must create an MNS queue in the China (Hangzhou) region.

- iv. Click Create Queue.
- v. In the Create Queue panel, configure the parameters for the queue.

In this example, enter eni-operate-completed-event in the Name field and use the default settings for other parameters.

vi. Click OK.

The following figure shows the created queue.

Name Jr	Available Messages ? ↓	Scheduled Messages ? ↓	Logging Feature @ √	Created At 🗤	Actions
eni-operate-completed-event Messane Retention Period: 4 Davs	0	0	 Not Enabled 	Dec 17, 2021	Details More 🗸

2. Create a CloudMonitor alert rule.

This topic elaborates only steps that are required by this best practice. For more information about how to create CloudMonitor alert rules, see Configure event notifications.

- i. Log on to the CloudMonitor console.
- ii. In the left-side navigation pane, click **Event Monitoring**.
- iii. On the Event Monitoring page, click the Alert Rules tab. On the Alert Rules tab, click Create Event Alert.

- iv. In the **Create / Modify Event Alert** panel, configure the following parameters for the alert rule:
 - Alert Rule Name: Specify a name for the rule. In this example, enter eni-event-test-rule.
 - Event Type: Select System Event.
 - Product Type: Select ECS.
 - Event Type: Select Status Notification.
 - Event Level: Select INFO.
 - Event Name: Select NetworkInterface:NetworkInterfaceOperateCompleted.
 - Resource Range: Use the default setting. You can modify the configuration based on your needs.
 - Alert Type: In this example, select only MNS queue and set Region and Queue based on the MNS queue that you created in the previous step. You can select multiple alert types based on your needs.
- v. Click OK.

The following figure shows the created alert rule.

Rule Name	Enable	Rule Description	Resource Range	Target	Actions
eni-event-test-rule	Enabled	ECS INFO NetworkInterface:NetworkInterfaceOperateCompleted	All Resources	MNS queue Asia Pacific SE 2 (Sydney) eni-operate- completed-event	Modify test Disable Delete

Use Java code to have an ENI operation event automatically handled

In this topic, Java Development Kit (JDK) 1.8 is used. In actual scenarios, you can modify and test the code based on your development tools and programming languages.

MNS provides a complete sample code for the queue operation. For more information, see Release notes of the SDK for Java.

1. Manually create a file named *.aliyun-mns.properties* under the user directory on your computer and add the endpoint of the queue and the AccessKey pair to the file.

? Note

- In Linux, the user directory is */home/<username>/*. In Windows, the user directory is *C:\Us ers\<username>*.
- The file must be named *.aliyun-mns.properties*, where *.properties* is the file name extension.

mns.accountendpoint=<Public endpoint of the queue>
mns.accesskeyid=<yourAccessKeyId>
mns.accesskeysecret=<yourAccessKeySecret>

Variables:

- *<Public endpoint of the queue>*: For information about how to obtain the public endpoint of a queue, see the "View the endpoints of a queue" in View the endpoints of a queue.
- *<yourAccessKeyId>*: the AccessKey ID of your Alibaba Cloud account. For information about how to obtain an AccessKey pair, see Obtain an AccessKey pair.
- <yourAccessKeySecret>: the AccessKey secret of your Alibaba Cloud account.
- 2. Create a Maven project by using a Java development tool.

Make sure that you create a Maven project by using a Java development tool such as Eclipse and Intellij IDEA before you configure ENI operations to be automatically performed.

3. Add Maven dependencies into the <dependencies></dependencies> section of the *pom.xml* file to install Alibaba Cloud SDKs.

? Note If you want to install the new versions of SDKs for multiple programming languages, see SDK Center.

Add the following Maven dependencies in sequence to install Alibaba Cloud SDKs.

• Install an SDK core library.

```
<dependency>
    <groupId>com.aliyun</groupId>
    <artifactId>aliyun-java-sdk-core</artifactId>
    <version>4.5.18</version>
</dependency>
```

• Install ECS SDKs.

```
<dependency>
    <groupId>com.aliyun</groupId>
    <artifactId>aliyun-java-sdk-ecs</artifactId>
    <version>4.23.10</version>
</dependency>
```

• Install MNS SDKs.

```
<dependency>
    <groupId>com.aliyun.mns</groupId>
    <artifactId>aliyun-sdk-mns</artifactId>
    <version>1.1.9</version>
</dependency>
```

• Install the Fastjson dependency.

```
<dependency>
    <groupId>com.alibaba</groupId>
    <artifactId>fastjson</artifactId>
    <version>1.2.73</version>
</dependency>
```

4. Create an entity class to define the structure of the ENI operation event.

In this example, create an entity class named EniEventMessage . For information about the structure of an ENI operation event notification, see ENI operation event notifications.

? Note In this topic, the package information is ignored in the sample code. In actual scenarios, add the package information to the code.

```
import java.util.Map;
public class EniEventMessage {
    private String resourceId;
    private String product;
```

```
private String ver;
private String instanceName;
private String regionId;
private String eventTime;
private String name;
private String ruleName;
private String id;
private String status;
private Map<String, String> content;
public String getResourceId() {
    return resourceId;
}
public void setResourceId(String resourceId) {
   this.resourceId = resourceId;
}
public String getProduct() {
   return product;
}
public void setProduct(String product) {
   this.product = product;
}
public String getVer() {
   return ver;
}
public void setVer(String ver) {
   this.ver = ver;
}
public String getInstanceName() {
   return instanceName;
}
public void setInstanceName(String instanceName) {
   this.instanceName = instanceName;
}
public String getRegionId() {
   return regionId;
}
public void setRegionId(String regionId) {
   this.regionId = regionId;
}
public String getEventTime() {
   return eventTime;
}
public void setEventTime(String eventTime) {
    this.eventTime = eventTime;
}
public String getName() {
   return name;
}
public void setName(String name) {
   this.name = name;
}
public String getRuleName() {
   return ruleName;
}
public void setRuleName(String ruleName) {
```

}

```
public vota bechatemane (betting tatemane) (
   this.ruleName = ruleName;
}
public String getId() {
   return id;
public void setId(String id) {
   this.id = id;
public String getStatus() {
   return status;
1
public void setStatus(String status) {
   this.status = status;
public Map<String, String> getContent() {
   return content;
}
public void setContent(Map<String, String> content) {
   this.content = content;
```

Create a code file to have operations triggered by the ENI operation event automatically performed.

In this example, the Java class name of the automated operations is ComsumerDemo. In this Java class, the information of the MNS queue and ENI operation event can be obtained, and then the obtained queue information can be cleared.

```
import com.aliyun.mns.client.CloudAccount;
import com.aliyun.mns.client.CloudQueue;
import com.aliyun.mns.client.MNSClient;
import com.aliyun.mns.common.ClientException;
import com.aliyun.mns.common.ServiceException;
import com.aliyun.mns.common.utils.ServiceSettings;
import com.aliyun.mns.model.Message;
import java.util.Map;
public class ComsumerDemo {
    public static void main(String[] args) {
        // Connect to the MNS queue.
        CloudAccount account = new CloudAccount(
                ServiceSettings.getMNSAccessKevId(),
                ServiceSettings.getMNSAccessKeySecret(),
                ServiceSettings.getMNSAccountEndpoint());
        MNSClient client = account.getMNSClient();
        try{
            // Obtain the information of the ni-operate-completed-event queue.
            CloudQueue queue = client.getQueueRef("eni-operate-completed-event");
            // Simulate polling of messages in the queue to have operations automatical
ly performed.
            for (int time = 0; time < 100; time++)
            {
                // Obtain the messages in the queue.
                Message popMsg = queue.popMessage();
                if (popMsg != null) {
```

TT (Pobusa .- untt)(System.out.println("message handle: " + popMsg.getReceiptHandle()); System.out.println("message body: " + popMsg.getMessageBodyAsString ()); System.out.println("message id: " + popMsg.getMessageId()); System.out.println("message dequeue count:" + popMsg.getDequeueCoun t()); // Deserialize the message bodies. EniEventMessage messageBody = com.alibaba.fastjson.JSON.parseObject (popMsg.getMessageBodyAsString(), EniEventMessage.class); // Obtain the content of ENI operation events contained in the mess age bodies. Map<String, String> messageContent = messageBody.getContent(); // Obtain the ID information of the ENI. String eniId = messageContent.get("eniId"); // Obtain the state of the ENI. String eniStatus = messageContent.get("eniStatus"); // Obtain the result of the operation performed on the ENI. String result = messageContent.get("result"); // Obtain the name of the operation performed on the ENI. String operation = messageContent.get("operation"); $\ensuremath{{\prime}}\xspace$ // Obtain the ID of the request to the ENI operation. String requestId = messageContent.get("requestId"); // Obtain the relevant information. System.out.println("ENI ID: " + eniId); System.out.println("ENI status: " + eniStatus); System.out.println("result: " + result); System.out.println("operation: " + operation); System.out.println("requestId: " + requestId); // Clear the information that corresponds to the queue. queue.deleteMessage(popMsg.getReceiptHandle()); System.out.println("delete message successfully.\n"); } } } catch (ClientException ce) System.out.println("Something wrong with the network connection between cli ent and MNS service." + "Please check your network and DNS availablity."); ce.printStackTrace(); } catch (ServiceException se) { if (se.getErrorCode().equals("QueueNotExist")) { System.out.println("Queue is not exist.Please create queue before use") ; } else if (se.getErrorCode().equals("TimeExpired")) System.out.println("The request is time expired. Please check your loca l machine timeclock"); } se.printStackTrace(); } catch (Exception e) { System.out.println("Unknown exception happened!");

```
e.printStackTrace();
}
client.close();
}
```

What's next

You can run the main function in ComsumerDemo to obtain the notification of an ENI operation event in an MNS queue. Then, you can have operations automatically performed in response to the ENI operation event. If an ENI operation event is triggered by the successful deletion of an ENI, a command output similar to the following one is returned.



7.Prefix lists 7.1. Overview

A prefix list is a set of one or more network prefixes (CIDR blocks). You can reference prefix lists to configure network rules for other network resources. You can add frequently used CIDR blocks to prefix lists to eliminate the need to repeatedly add multiple rules for CIDR blocks when you configure network rules. This improves O&M efficiency. Prefix lists can be referenced when you configure security group rules.

Concepts

Concept	Description
Maximum number of entries	The maximum number of CIDR blocks in a prefix list. Each entry consists of a CIDR block and the description for the CIDR block.
Address family	The address family of entries in prefix lists. Prefix lists support the IPv4 or IPv6 address type. Entries in a prefix list must belong to the same address family.

Concept	Description
	CIDR is an addressing scheme for the Internet that allows for IP addresses to be assigned in a more efficient manner than the traditional scheme based on classes A, B, and C. CIDR notation is used to denote IP addresses and IP ranges. It consists of an IP address and a forward slash followed by a decimal number that denotes how many bits are in the network prefix.
	• Example 1: Convert a CIDR block into an IP address range
CIDR block	For example, you can convert the 10.0.0.0/8 CIDR block into a 32-bit binary IP address of 00001010.00000000.00000000.00000000. In this CIDR block, /8 represents an 8-bit network ID. The first 8 bits of the 32-bit binary IP address are fixed, and the corresponding IP addresses are from 00001010.00000000.00000000.00000000 to 00001010.1111111111111111111111111. After you convert the preceding IP addresses into IP addresses in the decimal format, the 10.0.0.0/8 CIDR block indicates the IP addresses from 10.0.0.0 to 10.255.255.255 with a subnet mask of 255.0.0.0.
	• Example 2: Convert an IP address range into a CIDR block
	For example, you have a range of IP addresses from 192.168.0.0 to 192.168.31.255. You can convert the last two parts of the first and last IP addresses to binary numbers from 00000000.00000000 to 00011111.11111111. The first 19 (8 × 2 + 3) bits are fixed. After you convert the IP addresses to IP addresses in the CIDR format, the corresponding CIDR block is 192.168.0.0/19.
Associated resource	Other resources that reference prefix lists.

Limits

Use scenarios

You can maintain CIDR blocks in prefix lists and reference prefix lists in the rules of other resources. When you modify the entry information in a prefix list, the modifications take effect on all rules that reference the prefix list. This eliminates the need to modify multiple entries and improves O&M efficiency.

For example, when you add a rule to a security group, you can reference a prefix list to apply the rule to all CIDR blocks in the prefix list.

7.2. Create a prefix list

A prefix list is a set of one or more network prefixes (CIDR blocks). You can reference prefix lists to configure network rules for other network resources. This topic describes how to create a prefix list.

Prerequisites

If RAM users are used, the RAM users are granted permissions on prefix lists. For more information, see Grant RAM users permissions on prefix lists.

Procedure

- 1.
- 2.
- 3.
- 4. On the Prefix List page, click Create Prefix List.
- 5. In the **Create Prefix List** dialog box, configure the parameters described in the following table for the prefix list.

Create Prefix List ×			×		
* Name	doc_demo				
Description				1	
* Address Family	● IPv4 ○ IPv6 The select	ted address family cannot be changed	after the prefix list is created.		
* Max Entries	Max Entries 10 Entries Each prefix list can contain 1 to 200 entries. The maximum number of entries that the			that the	
	prefix list supports cannot be ch	nged after the prefix list is created.			
	124000-007500	en meneraleten n	CHARD, DARCH,		
Entries 🕐	+ Add Entries				
	CIDR Block 🕥	Description		Actions	
	192.168.0.0/24 ×	web		Delete	
	192.168.1.0/24 ×	sql		Delete	
				Delete	
			Create	Cancel	
Parameter		Description			
Parameter Name		Description Enter a name for the	prefix list.		
Parameter Name Description		Description Enter a name for the Enter a description for describe the intende	prefix list. or the prefix list. We r d purpose of the pref	ecommend t fix list.	hat you
Parameter Name Description		Description Enter a name for the Enter a description for describe the intende Select IPv4 or IPv6 its address family.	prefix list. or the prefix list. We r d purpose of the pref After you create a pre you can specify only	ecommend t fix list. efix list, you c IPv4 CIDR blo	hat you cannot modif
Parameter Name Description Address Far	nily	Description Enter a name for the Enter a description for describe the intender Select IPv4 or IPv6 its address family. o If you select IPv4, entries of the preference	prefix list. or the prefix list. We r d purpose of the pref After you create a pre you can specify only 'ix list.	ecommend t Fix list. Efix list, you c IPv4 CIDR blo	hat you cannot modif cks in the

Parameter	Description
Max Entries	Specify the maximum number of entries in the prefix list. After you create a prefix list, you cannot modify the maximum number of entries in the prefix list. Valid values: 1 to 200.
	Specify the information of CIDR blocks in the prefix list. You can click Add Entries to add a CIDR block and enter a description for the CIDR block. The following limits apply to optries in a prefix list:
	 The total number of entries cannot exceed the value set for Max
	Entries.
	 You can enter multiple CIDR blocks at a time. Separate multiple CIDR blocks with spaces or commas (,).
	 The address type of a CIDR block in each entry is determined by the Address Family parameter. You cannot combine IPv4 and IPv6 CIDR blocks in a single prefix list.
Entries	 CIDR blocks within entries in a prefix list must be unique. For example, you cannot specify 192.168.1.0/24 twice in the entries of the prefix list.
	• You can specify an IP address. The system automatically converts the IP address to a CIDR block.
	For example, if you specify 192.168.1.100, the system automatically converts the IP address to 192.168.1.100/32.
	 If an IPv6 CIDR block is specified, the system automatically converts it to the zero compression format.
	For example, if you specify 2001:0DB8:0000:0000:0000:0000:0000/32, the system converts it into 2001:db8::/32.

6. Click Create.

What's next

After the prefix list is created, you can perform the following operations:

- Maintain the prefix list. For information about how to modify the name of a prefix list or entries in a prefix list, see Manage the entries in a prefix list.
- Reference the prefix list. For information about how to reference a prefix list in security group rules, see Add a security group rule.

Related information

• CreatePrefixList

7.3. Clone a prefix list

If you want to create identical prefix lists in multiple regions, you can use the prefix list clone feature to create identical prefix lists in multiple regions in a quick manner.

Prerequisites

If RAM users are used, the RAM users are granted permissions on prefix lists. For more information, see Grant RAM users permissions on prefix lists.

Procedure

1.

- 2.
- 3.
- 4. On the **Prefix List** page, find the prefix list that you want to clone and click **Clone** in the **Actions** column.
- 5. In the Clone Prefix List dialog box, modify the Region, Name, and Description parameters.
- 6. Click Clone.

After the prefix list is cloned, you can switch to the intended region to view the cloned prefix list.

7.4. Manage the entries in a prefix list

Each entry consists of a CIDR block and a description for the CIDR block. You can add, modify, or delete entries for a prefix list on the Prefix List Details page.

Prerequisites

If RAM users are used, the RAM users are granted permissions on prefix lists. For more information, see Grant RAM users permissions on prefix lists.

Add entries to a prefix list

On the Prefix List Details page, you can add an entry to a prefix list.

1.

2.

- 3.
- 4. On the **Prefix List** page, find the prefix list to which you want to add an entry and click the ID of the prefix list.
- 5. On the Prefix List Details page, add an entry to the prefix list.
 - i. On the Entries tab, click Add.

ii. Enter a CIDR block in the CIDR Block field and enter a description for the CIDR block in the Description field.

The following limits apply to entries in a prefix list:

- The total number of entries cannot exceed the value set for Max Entries.
- You can enter multiple CIDR blocks at a time. Separate multiple CIDR blocks with spaces or commas (,).
- The address type of a CIDR block in each entry is determined by the Address Family parameter. You cannot combine IPv4 and IPv6 CIDR blocks in a single prefix list.
- CIDR blocks within entries in a prefix list must be unique. For example, you cannot specify 192.168.1.0/24 twice in the entries of the prefix list.
- You can specify an IP address. The system automatically converts the IP address to a CIDR block.

For example, if you specify 192.168.1.100, the system automatically converts the IP address to 192.168.1.100/32.

 If an IPv6 CIDR block is specified, the system automatically converts it to the zero compression format.

For example, if you specify 2001:0DB8:0000:0000:0000:0000:0000/32, the system converts it into 2001:db8::/32.

iii. Click **Save** in the Actions column.

If you want to add multiple entries to the prefix list, repeat the preceding operations.

Modify entries in a prefix list

On the Prefix List Details page, you can modify entries in a prefix list.

1.

2.

- 3. On the **Prefix List** page, find the prefix list in which you want to modify entries and click the ID of the prefix list.
- 4. On the Prefix List Details page, find the entry that you want to modify and click **Modify** in the Actions column.
- 5. Modify the CIDR block and its description.

The following limits apply to entries in a prefix list:

- You can enter multiple CIDR blocks at a time. Separate multiple CIDR blocks with spaces or commas (,).
- The address type of a CIDR block in each entry is determined by the **Address Family** parameter. You cannot combine IPv4 and IPv6 CIDR blocks in a single prefix list.
- CIDR blocks within entries in a prefix list must be unique. For example, you cannot specify 192.168.1.0/24 twice in the entries of the prefix list.
- You can specify an IP address. The system automatically converts the IP address to a CIDR block.

For example, if you specify 192.168.1.100, the system automatically converts the IP address to 192.168.1.100/32.

• If an IPv6 CIDR block is specified, the system automatically converts it to the zero compression format.

For example, if you specify 2001:0DB8:0000:0000:0000:0000:0000/32, the system converts it into 2001:db8::/32.

6. Click **Save** in the Actions column.

Delete entries from a prefix list

On the Prefix List page, you can delete entries from a prefix list.

1.

2.

- 3. On the **Prefix List** page, find the prefix list from which you want to delete entries and click the ID of the prefix list.
- 4. On the Prefix List Details page, use one of the following methods to delete entries:
 - Delete a single entry: Find the entry that you want to delete and click **Delete** in the **Actions** column.
 - Delete multiple entries: Select the entries that you want to delete and click **Delete** in the lower part of the page.

Related information

• ModifyPrefixList

7.5. Delete prefix lists

You can delete prefix lists that are no longer needed.

Prerequisites

- No resources are associated with the prefix lists that you want to delete. If the prefix lists are referenced by other resources, the prefix lists cannot be deleted. You must remove all references to the prefix lists from the resources on the configuration pages of the resources before you can delete the prefix lists.
- If RAM users are used, the RAM users are granted permissions on prefix lists. For more information, see Grant RAM users permissions on prefix lists.

Procedure

1.

2.

3.

- 4. On the Prefix List page, use one of the following methods to delete prefix lists:
 - Delete a single prefix list: Find the prefix list that you want to delete and click **Delete** in the **Actions** column.
 - Delete multiple prefix lists: Select the prefix lists that you want to delete and click **Batch Delete** in the lower part of the page.

Related information

• Delet ePref ixList

7.6. Grant RAM users permissions on prefix lists

You can use RAM users to avoid sharing the AccessKey pair of your Alibaba Cloud account with other users. You can grant permissions to RAM users based on the principle of least privilege to minimize security risks for your enterprise. This topic describes how to grant RAM users the permissions on prefix lists.

Context

This topic describes how to grant RAM users the permissions on prefix lists. If you want to use other resources in the Resource Access Management (RAM) console, you must attach policies that correspond to the resources to RAM users. For example, you can click **System Policy** and click **AliyunECSReadOnlyAccess** to grant the read-only permissions on Elastic Compute Service (ECS).

Procedure

- 1. Log on to the RAM console by using your Alibaba Cloud account.
- 2. Create the policies on prefix lists. For more information, see Create a custom policy.

* Policy Name
PrefixListPolicy
Note
Configuration Mode
O Visualized
Script
Policy Document
Import an existing system policy
1 {
4 "Action": [
5 "ecs:CreatePrefixList",
6 "ecs:ModifyPrefixList",
7 "ecs:DescribePrefixLists",
8 "ecs:DescribePrefixListAssociations",
9 "ecs:DescribePrefixListAttributes",
10 "ecs:DeletePrefixList"
11],
12 "Resource": "*",
13 "Effect": "Allow"
15],
OK Return

Create the PrefixListPolicy policy. The following code shows the content of the policy:

```
{
   "Statement": [
       {
           "Action": [
               "ecs:CreatePrefixList",
               "ecs:ModifyPrefixList",
               "ecs:DescribePrefixLists",
               "ecs:DescribePrefixListAssociations",
                "ecs:DescribePrefixListAttributes",
               "ecs:DeletePrefixList"
           ],
            "Resource": "*",
            "Effect": "Allow"
       }
   ],
   "Version": "1"
}
```

Note The preceding code shows only authentication rules for prefix lists. For more information about ECS-related authentication rules, see Authentication rules.

3. Grant RAM users the permissions on prefix lists. For more information, see Grant permissions to a RAM user.
| * Authorized Scope | | | | |
|------------------------------|-----------------|---|-------------------------|--------|
| Alibaba Cloud Account | | | | |
| O Specific Resource Group | | | | |
| Enter a resource group name. | | | | \sim |
| * Principal | | | | |
| (page and an other states) | L.com X | | | |
| * Select Policy | | | | |
| System Policy Custom Policy | + Create Policy | | Selected (2) | Clear |
| PrefixListPolicy | | G | AliyunECSReadOnlyAccess | × |
| Authorization Policy Name | Description | | PrefixListPolicy | × |
| PrefixListPolicy | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | E |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| OK Cancel | | | | |

After the permissions are granted, you can use the RAM users to manage prefix lists.

7.7. Use prefix lists to simplify management of security group rules

This topic describes how to use prefix lists to simplify management of security group rules.

Context

A prefix list is a set of one or more network prefixes (CIDR blocks). You can reference prefix lists to configure security group rules. When entries in a prefix list are modified, all security group rules that reference this prefix list are also updated. You can put frequently-used IP addresses in a prefix list and reference the prefix list in security group rules instead of referencing the IP addresses individually. This way, you can consolidate security group rules that share the same attributes except for the authorization object into a single rule that uses a prefix list as the authorization object, and reduce the burdens of managing security group rules. For more information about prefix lists, see Overview.

Use scenarios

Assume that you have planned multiple security domains for your resources in the cloud to ensure resource security. Each security domain corresponds to a security group. A public resource such as an office network off the cloud requires access to your resources in multiple security domains. This public resource has multiple variable CIDR blocks.

If you do not use the prefix list feature, you must configure multiple rules that reference the CIDR blocks of the public resource as authorization objects in multiple security groups to allow access from the public resource. The configured security group rules must share the same attributes except for the authorization object. If the CIDR blocks of the public resource change, you must modify the corresponding rules of security groups in multiple security domains. The greater the number of security groups and CIDR blocks, the more difficult to manage the security group rules.

If you use the prefix list feature, you can create a prefix list from the CIDR blocks of the public resource and configure a rule that references the prefix list as the authorization object in multiple security groups to allow access from the public resource. If the CIDR blocks of the public resource change, you need only to modify the corresponding entries in the prefix list, and the associated security group rules are also updated. This eliminates the need to modify the security group rules one by one and simplifies management of security group rules.

If you have resources in multiple Alibaba Cloud regions, you can use the clone feature to clone prefix lists across regions.

Procedure

This section describes how to add or modify security group rules by using a prefix list to deny or allow access from specific IP addresses. In the examples, two IP addresses are used.

Note If you are using a RAM user, grant permissions on prefix lists to the user. For more information, see **Grant RAM users permissions on prefix lists**.

1.

- 2. Create a prefix list.
 - i.

ii.

iii. Click Create Prefix List.

iv. In the Create Prefix List dialog box, configure parameters and click Create.

In this example, two IPv4 entries are added. Examples values of parameters in the Create Prefix List dialog box:

- Name: RemoteLogon
- Description: Allow access from the CIDR blocks in the prefix list to Elastic Compute Service (ECS) instances in a security group.
- Address Family: IPv4
- Max Entries: 2

(?) Note The rule quotas of resources (such as security groups) that are associated with a prefix list are calculated based on the maximum number of entries in the prefix list, instead of the actual number of entries. Set a proper value for Max Entries.

 Entries: Click Add Entries and enter 192.168.1.0/24 for an entry. Click Add Entries again and enter 192.168.2.0/24 for another entry.

Each of preceding CIDR blocks is a set of consecutive IP addresses.

- 192.168.1.0/24 : 192.168.1.0 **to** 192.168.1.255 .
- 192.168.2.0/24 : 192.168.2.0 to 192.168.2.255 .
- 3. Add security group rules that reference the prefix list.

Repeat the following steps to add a rule that references the RemoteLogon prefix list to allow access to remote connection ports in multiple security groups:

i.

- ii. Find the security group to which you want to add a rule and click Add Rules in the Actions column.
- iii. On the Inbound tab, click Add Rule.

(?) Note In this example, a security group of the Virtual Private Cloud (VPC) type is used. For a security group of the classic network type, select a tab based on whether the CIDR blocks are public ones.

iv. Configure parameters to add a rule and click Save.

In this example, a rule is added to allow SSH access and Remote Desktop Protocol (RDP) access to the ECS instances in the security group. Example values of parameters in the rule entry:

- Action: Allow
- Priority: 1
- Protocol Type: Custom TCP
- Port Range: SSH (22) and RDP (3389)
- Authorization Object: the RemoteLogon prefix list

After the rule is added, the CIDR blocks contained in the RemoteLogon prefix list are allowed to connect to the instances within the security group.

4. Modify the entries in the prefix list.

After the rules are added to security groups, if you want to deny access from specific IP addresses

contained in the RemoteLogon prefix list, modify the corresponding entries in the prefix list, instead of modifying the security group rules one by one. For example, assume that you use the instances whose private IP addresses are 192.168.1.1 and 192.168.2.1 as jump servers and you want to allow access to the security groups only from the jump servers. Perform the following steps to modify the entries in the prefix list:

i.

- ii. Find the RemoteLogon prefix list and click View Details in the Actions column.
- iii. Click the Entries tab.
- iv. Click Modify in the Actions column corresponding to one entry.
- v. Change the CIDR block and click **Save**.

Repeat the preceding steps to modify the other entry. The CIDR block in one entry is changed to 192.168.1.1/32, and the CIDR block in the other entry is changed to 192.168.2.1/32.

After the entries are modified, the modifications immediately take effect. The security group rules that use the RemoteLogon prefix list are updated to allow access only from 192.168.1. 1/32 and 192.168.2.1/32.

More use scenarios

This section compares the numbers of operations required for security group rules that reference individual IP addresses and for security group rules that reference prefix lists to show the advantages of prefix lists in improving efficiency. Assume that you have 50 security groups. The following table describes the numbers of operations required in different scenarios when a prefix list is used and when a prefix list is not used.

Scenario	Security group rule that references an individual IP address	Security group rule that references a prefix list
Deny access from five IP addresses	If you remove five allow rules that reference five IP addresses one by one from each of the 50 security groups, 250 remove operations are required. Event if you batch remove the five allow rules from each security group, 50 remove operations are still required.	If you remove five entries that contain five IP addresses one by one from the prefix list, five remove operations are required. If you batch remove the five entries from the prefix list, only a single remove operation is required.
Modify rules or entries to allow access from five IP addresses	If you modify five rules in each of the 50 security groups to allow access from five IP addresses, 250 modify operations are required.	If you modify five entries to include five IP addresses in the prefix list, five modify operations are required.
Add rules or entries to allow access from five IP addresses	If you add five rules in each of the 50 security groups to allow access from five IP addresses. In this case, 250 add operations are required. Event if you add a single rule that references five IP addresses to each security group, and 50 add operations are still required.	If you add five entries that include five IP addresses to the prefix list, five add operations are required. If you add a single entry that includes the five IP addresses to the prefix list, only a single add operation is required.

Scenario	Security group rule that references an individual IP address	Security group rule that references a prefix list
Modify rules or entries to allow access from five IP addresses, and add rules or entries to access from another five IP addresses	If you modify five rules in each of the 50 security groups to allow access from five IP addresses and then add five rules to each security group to allow access from another five IP addresses, a total of 500 operations are required. Even if you modify five rules that reference five IP addresses and add a single rule that references another five IP addresses in each security group, a total of 300 operations are still required.	If you modify five entries to include five IP addresses and add five entries to include the other five IP addresses in the prefix list, a total of 10 operations are required. If you modify five entries and add a single entry that includes another five IP addresses in the prefix list, a total of six operations are required.

8.Change the VPC of an ECS instance

This topic describes how to migrate a VPC-type ECS instance from the current VPC to another. If you select a VPC that does not meet your business needs when you create an ECS instance or if you want to replan the network of an instance, you can use this feature to change the VPC of the instance.

Prerequisites

• The instance is in the Stopped state. For more information, see Stop an instance.

? Note If the instance has not been restarted after it is created, you must restart it before you stop it.

- The instance is not added as a backend server of a Server Load Balancer (SLB) instance. For more information, see Remove a backend server.
- Secondary elastic network interfaces (ENIs) bound to the instance are unbound. Multiple secondary private IP addresses assigned to the ENIs are revoked. For more information, see Unbind an ENI and Unassign secondary private IP addresses.
- The destination VPC, vSwitch, and security groups are created and available.

Scenarios

- You want to replan the VPCs of your ECS instances because the original VPCs are unable to keep up with the growing needs of your business.
- In the early business stage, only one VPC is planned. Risks in data operations exist because different projects and usage environments share this VPC. You want to use different VPCs for different projects and environments.
- Your ECS instances are deployed in the default VPCs of different accounts. Connections between instances across Alibaba Cloud accounts cannot be implemented due to IP address conflicts. In this case, you must change the VPCs of the ECS instances and resolve the address conflicts before you can interconnect the instances across Alibaba Cloud accounts.

Limits

- The instance cannot be used in other cloud services. For example, the instance cannot be in the process of being migrated or having its VPC changed, or the databases deployed in the instance cannot be managed by Data Transmission Service (DTS).
- After the VPC is changed, the new vSwitch of the instance must be in the same zone as the original vSwitch.
- You must select one to five security groups for an instance and the security groups must be of the same type (basic security group or advanced security group).
- Instances of some instance families cannot be migrated to a VPC if advanced VPC features are enabled on the VPC. For more information, see Instance families that do not support advanced VPC features.
- You can change the VPCs of up to 20 instances at a time.
- After you change the VPC of an instance, the instance can no longer communicate with other instances in the original VPC. For information about how to communicate with other instances, see

What is Express Connect?

- The cut-through mode or multi-EIP to ENI mode cannot be enabled for the instance.
- The instance cannot be associated with a high-availability virtual IP address (HAVIP).
- The vSwitch of the instance cannot be associated with a custom route table.
- Global Accelerator (GA) cannot be activated for the instance.

Procedure

- 1.
- 2.
- 3.
- 5.
- 4.
- 5. Change the VPC of one or more ECS instances at a time.
 - Change the VPC of a single instance

Find the instance for which you want to change the VPC, and choose **More > Network and Security Group > Change VPC** in the **Actions** column.

• Change the VPCs of multiple ECS instances at a time

Select the instances for which you want to change the VPCs, and choose **More > Network and Security Group > Change VPC** in the lower part of the page.

6. In the Change VPC wizard, follow the instructions to change the VPC of the ECS instance.

Prepare	2 Select VPC	③	Configure Primary Private IP Addre:	(4) Res
he secondary elastic network	interfaces (ENIs) must be deta	ached and all secondary private	IP addresses assigned to the ENI mus	st be
evoked from the instance. The instance must be removed	I from the SLB backend server	s in advance. VPC	Primary Private IP Addre	ess
evoked from the instance. The instance must be removed nstance ID/Name -/i-0xi§ whto2e	I from the SLB backend server Zone Virginia Zone B	s in advance. VPC vpc-0	Primary Private IP Addre	ess

- i. In the Prepare step, check the network information and precautions and click Next.
- ii. In the Select VPC step, configure the Destination VPC, Destination VSwitch, and Destination Security Group parameters and click Next.
- iii. (Optional)In the **Configure Primary Private IP Address** step, specify a primary private IP address for the selected instance to use in the destination VPC.
 - The primary private IP address must be within the CIDR block of the destination vSwitch.
 - If you do not manually set the primary private IP address, the system assigns one.
- iv. Click OK.

After you change the VPC of the instance, you can click the instance ID to view the new VPC and vSwitch in the **Network Information** section on the **Instance Details** page.

Related information

• ModifyInstanceVpcAttribute

9.Configure NIC multi-queue

Network interface controller (NIC) multi-queue enables an Elastic Compute Service (ECS) instance to use multiple NIC queues to improve network performance. Performance bottlenecks may occur when a single vCPU is used to process NIC interrupts on an instance. To solve this issue, you can use NIC multi-queue to distribute NIC interrupts across different vCPUs. This way, network performance is improved.

Prerequisites

The following requirements are met:

• The instance type of your instance supports the NIC multi-queue feature. For more information about the instance types that support NIC multi-queue, see Instance family. If the number of NIC queues is greater than one, NIC multi-queue is supported.

Note In an instance of the re6p persistent memory-optimized instance family, if NIC interrupts are not distributed across different vCPUs, we recommend that you upgrade the ecs_mq configuration script to the latest version.

• The image of your instance supports the NIC multi-queue feature and has this feature disabled by default. The following public images provided by Alibaba Cloud support the NIC multi-queue feature. The support for this feature is irrelevant to the bit sizes of the operating systems contained in images.

? Note

- Even if your operating system is included in the list, public images of earlier versions may not support the NIC multi-queue feature. We recommend that you use the latest public images. If the image of your instance has the NIC multi-queue feature enabled by default, skip this topic.
- The following procedure applies only to Linux instances. You do not need to configure the NIC multi-queue feature for ECS instances that run Windows 2012 or later because this feature is automatically configured.

Public image	NIC multi-queue supported	NIC multi-queue enabled
Cent OS 6.8/6.9/7.2/7.3/7.4/8.*	Yes	Yes
Ubuntu 14.04/16.04/18.04/20.04	Yes	Yes
Debian 8.9/9.2/10.*	Yes	Yes
SUSE Linux Enterprise Server 12 SP1/12 SP2/15 SP1/15 SP2	Yes	Yes
Red Hat Enterprise Linux 6.9/7.4/7.5	Yes	No
OpenSUSE 42.3/15.*	Yes	No
Alibaba Cloud Linux 2.1903	Yes	Yes

Public image	NIC multi-queue supported	NIC multi-queue enabled
Windows 2012 or later	Yes	Yes

Context

NIC multi-queue is a technology that can fix Quality of Service (QoS) issues of I/O bandwidth. The NIC multi-queue driver uses interrupts to bind each queue to different vCPUs to solve processing bottlenecks of single vCPUs when network I/O bandwidth increases and improve the packet forwarding rate and bandwidth performance. Under identical packet forwarding rate and network bandwidth conditions, two queues have performance 50% to 100% higher than that of a single queue. Performance can be improved even greater with the use of four queues.

Automatic configuration

1.

2. Download the ecs mq automatic configuration script package.

wget https://ecs-image-tools.oss-cn-hangzhou.aliyuncs.com/ecs_mq/ecs_mq_latest.tgz

3. Decompress the script package.

tar -xzf ecs_mq_latest.tgz

4. Change the working path.

cd ecs_mq/

5. Run the extracted script.

The command format varies based on image versions. For example, use bash install.sh centos 7 for Cent OS 7.6 images.

bash install.sh <Operating system name> <Major version number of the operating system>

6. Start the service.

systemctl start ecs_mq

Manual configuration

In the following example, a CentOS 7.6 image is used. The primary elastic network interface (ENI) is named eth0, and the secondary ENI is named eth1. This section describes how to manually configure NIC multi-queue.

1. Run the ethtool -1 eth0 command to check whether the primary ENI supports NIC multi-queue.

```
[root@localhost ~]# ethtool -l eth0
Channel parameters for eth0:
Pre-set maximums:
RX: 0
TX: 0
Other: 0
Combined: 2 # Indicates that a maximum of two queues can be configured.
Current hardware settings:
RX: 0
TX: 0
Other: 0
Combined: 1 # Indicates that one queue is in effect.
```

? Note If the return values of the two Combined fields are the same, NIC multi-queue is enabled for the primary ENI.

2. Run the ethtool -L eth0 combined 2 command to enable NIC multi-queue.

This command configures the eth0 primary ENI to use two queues.

[root@localhost ~]# ethtool -L eth0 combined 2

3. Configure NIC multi-queue for the secondary ENI.

```
# Check whether the ethl secondary ENI supports NIC multi-queue.
[root@localhost ~]# ethtool -1 eth1
Channel parameters for eth1:
Pre-set maximums:
RX: 0
TX: 0
Other: 0
Combined: 4 # Indicates that a maximum of four queues can be configured.
Current hardware settings:
RX: 0
TX: 0
Other: 0
Combined: 1 # Indicates that one queue is in effect.
# Configure the ethl secondary ENI to use four queues.
[root@localhost ~]# ethtool -L eth1 combined 4
```

10.Set the MTU size of an NIC

A maximum transmission unit (MTU) is the largest size of a packet that can be transmitted without the need to fragment the packet. This topic describes how to set the MTU size of a network interface controller (NIC) on an Elastic Compute Service (ECS) instance by using the ecs-utils-jumbof rame script.

Context

To limit the maximum sizes of transmitted packets and prevent excess latency resulted when large packets are fragmented, we recommend that you keep the local and network MTU sizes consistent.

Notice The ecs-utils-jumbof rame script does not contain commands that can be used to restart NICs or network services. Different types of network services and NIC drivers may be used in different systems, and transient connections may occur. If your business does not allow transient connections, exercise caution when you perform this operation.

Limits

- You can run the ecs-utils-jumbof rame script only on ECS instances of the Virtual Private Cloud (VPC) type. The ecs-utils-jumbof rame script uses instance metadata services to check whether instance types support the Jumbo Frame feature. Before you run the script on an ECS instance, access <IP address of the ECS instance>: <Port number> in your browser to check whether the port used to exchange data between networks and instance metadata services is enabled on the ECS instance. Example: 100.100.100.200:80. If the port is not enabled, enable it.
- Only seventh-generation and later instances support the Jumbo Frame feature. The ecs-utilsjumboframe script determines the maximum MTU size based on instance types. You can also use this script to set a maximum MTU size of 1,500 for NICs on ECS instances of other instance types.
- The Jumbo Frame feature is applicable only to direct TCP communication between seventhgeneration or later instances of the VPC type, and does not support the communication through intermediate nodes such as Server Load Balancer (SLB) instances.

Set the MTU size of an NIC on a Linux instance

Use a Linux instance that uses one of the following images:

- CentOS 6.x, CentOS 7.x, and CentOS 8.x
- Debian 9.x and Debian 10.x
- SUSE Linux Enterprise Server 15
- Ubuntu 16.x, Ubuntu 18.x, and Ubuntu 20.x
- 1. Connect to the instance. For more information, see Connection methods.
- 2. Download the ecs-utils-jumboframe-linux.sh script.

wget https://ecs-image-tools.oss-cn-hangzhou.aliyuncs.com/jumboframe/ecs-utils-jumbofra me-linux.sh

3. Grant the execute permissions on the script.

```
chmod +x ./ecs-utils-jumboframe-linux.sh
```

ip addr show

You can run ./ecs-utils-jumboframe-linux.sh -h to view the help.

[root@i ______~]# chmod +x ./ecs-utils-jumboframe-linux.sh
[root@i ______~]# ./ecs-utils-jumboframe-linux.sh -h
Usage: ecs-utils-jumboframe-linux.sh [-h] [-p] IFNAME MTU
 __h Print this help.
 __p Also configure new MTU value persistently.

4. Query the list of NICs on the instance and the MTU sizes of the NICs.

- 5. Set the MTU size of a NIC.
 - Set a temporary MTU size for the NIC. To test the Jumbo Frame feature, you can use the following command to set a temporary MTU size for the NIC. The MTU is restored to the default size when the instance is restarted. Format:

./ecs-utils-jumboframe-linux.sh <NIC name> <New MTU size>

For example, you can run the following command to set a temporary MTU size of 8500 for the eth0 NIC:

./ecs-utils-jumboframe-linux.sh eth0 8500

 Set a permanent MTU size (a MTU size that remains unchanged when the instance is restarted) for the NIC. To enable the Jumbo Frame feature, you can use the following command to set a permanent MTU size and specify the -p mode. After the MTU size is set, it remains unchanged when the instance is restarted. Format:

./ecs-utils-jumboframe-linux.sh -p <NIC name> <New MTU size>

For example, you can run the following command to set a permanent MTU size of 8500 for the eth0 NIC:

./ecs-utils-jumboframe-linux.sh -p eth0 8500

After the new MTU size is set for the eth0 NIC, you can run the <u>ip addr show</u> command to query information of the NIC. The MTU size of the NIC is 8500 in the command output.



Set the MTU size of an NIC on a Windows instance

Make sure that the Windows instance meets the following requirements:

- The instance uses a Windows Server 2012 or later image.
- The instance has PowerShell installed.
 - 1. Connect to the instance.
 - 2. Download the *ecs-utils-jumboframe-windows.ps1* script.

Visit Download link in your browser.

- 3. Start PowerShell.
- 4. Switch to the directory where the script is located.

In this example, the script is placed on the desktop and located in the following directory:

cd C:\Users\Administrator\Desktop

You can run Get-Help .\ecs-utils-jumboframe-windows.ps1 to view the help.



5. Query the list of NICs.

Get-NetAdapte	er			
PS C:\Users\Admini	strator\Desktop> Get=NetAdapter			
Name	InterfaceDescription	ifIndex Status	MacAddress	LinkSpeed
以太网	Red Hat VirtIO Ethernet Adapter	3 Vp	•	10 Gbps

6. Set the MTU size of a NIC.

? Note You can set only permanent MTU sizes on Windows instances. After a permanent MTU size is set for a NIC on a Windows instance, the MTU size remains unchanged when the instance is restarted.

Format:

.\ecs-utils-jumboframe-windows.ps1 -NetworkAdapter <"NIC name"> -NewMTUValue <New MTU s
ize>

For example, you can run the following command to set a permanent MTU size of 8500 for an NIC named Ethernet:

.\ecs-utils-jumboframe-windows.ps1 -NetworkAdapter "Ethernet" -NewMTUValue 8500

 After the new MTU size is set for the NIC, you can run the
 Get-NetAdapterAdvancedProperty -Name

 "Ethernet" -RegistryKeyword "MTU"
 command to query information of the NIC. The MTU size of the NIC is 8500 in the command output.

PS C:\Users\Administrator	\Desktop> Get=NetAdapterAdvance	dProperty -Name "以太网"-Regis	tryKeyword "MTU"	
Name	DisplayName	DisplayValue	RegistryKeyword	RegistryValue
以太网	Init.MTUSize	8500	MTU	{8500}

11.Migrate an ECS instance from the classic network to a VPC (new version)

11.1. Migrate ECS instances from the classic network to a VPC

Compared with Elastic Compute Service (ECS) instances located in the classic network, ECS instances located in virtual private clouds (VPCs) are more secure and support additional features such as associating elastic IP addresses (EIPs). This topic describes how to use a migration plan to migrate one or more ECS instances from the classic network to a VPC.

Prerequisites

The ECS instances that you want to migrate from the classic network to a VPC meet the following requirements:

- The instances do not have local disks attached. If the instances have local disks attached, submit a ticket to seek advice from Alibaba Cloud on how to migrate the instances.
- The instances have a public bandwidth greater than 0 Mbit/s. If an instance has a public IP address and a public bandwidth of 0 Mbit/s, you must upgrade the public bandwidth before you can migrate the instance. For more information, see Modify public bandwidth.
- The instances are located in one of the following regions that support the migration plan feature: China (Qingdao), China (Beijing), China (Hangzhou), China (Shanghai), China (Shenzhen), China (Hong Kong), US (Silicon Valley), and Singapore (Singapore).

(?) Note Some instances located in Hangzhou Zone C cannot be migrated from the classic network to VPCs.

Impacts of migrating an ECS instance from the classic network to a VPC

ltem	Description
	It takes about 15 minutes from when the instance in the classic network is stopped until when it is migrated and started in the VPC.
Amount of time required to migrate the instance	Note After the computing and network resources of an instance are migrated, the instance is started in the VPC. If the instance is migrated across zones, the system continues to migrate disk data of the instance after the instance is started. Typically, it takes about 4 hours to migrate 100 GiB of disk data. During the migration, the I/O performance of disks degrades and snapshot- and disk-related features are not supported.

ltem	Description
Instance state	During migration, the instance is stopped and then started again. We recommend that you schedule to migrate your instance during off-peak hours.
Network type	After the instance is migrated, its network type changes from classic network to VPC. For information about VPCs, see What is a VPC?.
Software authorization code	After the instance is migrated, its software authorization codes may change.
	• The public IP address of the instance remains unchanged.
IP address	Notice ECS instances located in VPCs do not have public network interface controllers (NICs), and use NAT devices to access the Internet. You can find only internal IP addresses inside the instances. If your applications require a public IP address visible in the instance operating system, reconsider whether to migrate your instance from the classic network to a VPC.
	• You can specify whether to retain the internal IP address of the instance when you create a migration plan to migrate the instance. You can also modify the internal IP address of the instance after the instance is migrated. For more information, see Modify a private IP address.
Disk name	 Some ECS instances have their underlying virtualization technology upgraded when they are migrated from the classic network to VPCs. This may cause the disk names on the instances to change. On Linux instances, disks names follow a naming convention of vd?, such as vda, vdb, and vdc. If a disk name is in the vd? format before the instance is migrated, the disk name remains unchanged after the instance is migrated. If a disk name is in the xvd? format before the instance is migrated, the disk name is converted to the vd? format such as vda, vdb, or vdc after the instance is migrated. Alibaba Cloud updates the <i>/etc/fstab</i> file for Linux instances. However, you must check whether applications are dependent on the original disk names.
Fee	 You are not charged for the migration. After a subscription instance is migrated from the classic network to a VPC, a new billing cycle immediately starts and the unit price of the instance type changes. An instance located in a VPC is more cost-effective than an instance with the same configurations located in the classic network. Orders for instance renewal and configuration changes that have not taken effect or are unpaid are canceled. You can renew the instance and change its configurations again.

ltem	Description
Others	 The ID, username, and logon password of the instance remain unchanged. If the ECS instance has been added to the vServer group of a Server Load Balancer (SLB) instance before the ECS instance is migrated, the ECS instance is not automatically associated with the SLB instance after the ECS instance is migrated. You must manually add the ECS instance to the vServer group of the SLB instance. For more information, see Create a vServer group.

Preparations

1. Create snapshots for the disks on the ECS instances to be migrated to back up data.

For more information, see Create a snapshot of a disk.

2. (Optional) If an ECS instance to be migrated is associated with an Alibaba Cloud database service, you must enable the hybrid access mode for the database service beforehand.

In hybrid access mode, Alibaba Cloud database services are accessible to ECS instances regardless of whether the instances are located in the classic network or in VPCs. For more information, see Overview of the hybrid access mode of ApsaraDB.

3. (Optional) If an ECS instance to be migrated is associated with an Alibaba Cloud database service (such as ApsaraDB RDS) that provides the whitelist feature, you must add the CIDR block of the destination vSwitch to the corresponding whitelists of the database service beforehand.

For more information, see Configure a whitelist.

- 4. (Optional) To ensure that services can be rapidly restored after migration, we recommend that you configure application services to run on instance startup and monitor service availability.
- 5. Disable or uninstall server security software on the ECS instances to be migrated.

Note The device drivers of ECS instances are updated when the instances are migrated. You must disable or uninstall security software such as Safedog, Huweishen, and Yunsuo on the instances beforehand.

- 6. Reserve at least 500 MiB of free space on the system disk of each ECS instance to be migrated.
- 7. Make sure that the destination vSwitch has sufficient internal IP addresses available. The number of the available internal IP addresses must be greater than that of ECS instances to be migrated.

Step 1: Create a migration plan

1.

2.

- 3.
- 4. Click Create Migration Plan.
- 5. In the **Configure Migration Plan** step, configure parameters in different sections and then click **Next**.
 - i. Configure parameters in the Destination Zone and VPC section.

Destination Zone and VPC	
Plan Name	
Default migration plan	
* Select a destination zone	
Shanghai Zone B 🗸	No zone found 🞱
* Destination VPC or Create a VPC @	
(Default) Automatically create a VPC, CIDR block: 10.0.0.0/8 $$	С

Parameter	Description
Plan Name	Enter a name for the migration plan.
	Select a destination zone from the drop-down list. The available zones are automatically planned and displayed based on resource availability. If you want to specify a zone that is not in the drop-down list, submit a ticket.
Select a destination zone	Note Only a single zone can be specified in each migration plan. If you want to migrate multiple ECS instances to different zones, you must create multiple migration plans.
	Select a destination VPC from the drop-down list. The CIDR block of the selected destination VPC determines whether to retain the internal IP addresses of the ECS instances from the classic network.
	If you want to retain the internal IP addresses of the ECS instances, you must select a VPC that is associated with the 10.0.0.0/8 CIDR block. You can select the default option or a VPC that you created.
	If you have not created VPCs that are associated with the 10.0.0.0/8 CIDR block, select (Default) Automatically create a VPC, CIDR block: 10.0.0.0/8 for the system to create a VPC that is associated with the 10.0.0.0/8 CIDR block.
Destination VPC or	If you have created a VPC that is associated with the 10.0.0/8 CIDR block, select the VPC.
	If you do not want to retain the internal IP addresses of the ECS instances, you must select a VPC that is associated with a CIDR block other than 10.0.0.0/8.

ii. Configure parameters in the Instance Network Properties section.

Instance Network Properties
* Destination Security Group
● (Default) Clone Security Groups of Classic Network-type Instances Learn more about cloning ⊘
Specify Security Groups
* Mar Address Retention Dolisy
* Mac Address Netendon Policy
(Default) Private Mac Address
Public Mac Address

Parameter	Description			
Destination Security Group	 Specify destination security groups for the ECS instances from the classic network. Valid values: (Default) Clone Security Groups of Classic Network-type Instances: The security groups of the ECS instances are automatically cloned from the classic network to the destination VPC. The rules of the new security groups (clone security groups) in the VPC are identical to those in the original security groups in the classic network. If you set Destination VPC or Create a VPC to (Default) Automatically create a VPC, CIDR block: 10.0.0.0/8, Destination Security Group is automatically set to (Default) Clone Security Groups of Classic Network-type Instances and cannot be modified. Note If rules of a security group contains other security groups as sources or destinations for traffic, the security group cannot be cloned. Specify Security Groups: Select one or more existing security groups from the drop-down list. Note Improper security group settings affect the connectivity of ECS instances. Make sure that your security group rules meet your connectivity requirements. 			
Mac Address Retention Policy	 Specify which MAC address to retain for the ECS instances from the classic network. In the classic network, if an ECS instance is assigned a public IP address, the instance has a public MAC address and a private MAC address. In a VPC, each ECS instance has only a private MAC address and can have its internal IP address mapped by a NAT device to a public IP address for Internet access. You can select (Default) Private Mac Address or Public Mac Address based on your needs. If your business system is associated with a MAC address (example: if your software is associated with a MAC address for registration), select the corresponding option to retain the associated MAC address. If your business system is not associated with a MAC address, select (Default) Private Mac Address or Public Mac Address or Public Mac Address or Public Mac Address. 			

iii. Configure parameters in the Instance Network Connectivity section.

(Default) Yes vSwitch Creation Policy O Automatic Manual The system automatically allocates a vSwitch with a 16-bit CDR block to the destination VPC based on internal IP addresses of classic network-type instance. If the vSwitch cannot be created because the specified CDR block is occupied or	
vSwitch Creation Policy Automatic Manual The system automatically allocates a vSwitch with a 16-bit CIDR block to the destination VPC based on internal IP addresses of classic network-type instances. If the vSwitch cannot be created because the specified CIDR block is occupied or	
The system submatically allocates a dynitch with a 16-bit CIDR block to the destination VPC based on internal IP addresses of 16-bit network-type instances. If the v3witch earnot be received because the specified CIDR block is occupied or	
because the maximum number of vSwitches within the VPC has been reached, select Manual.	
* Ensure interconnections between the migrated instances and the classic network-type instances specified in the plan over the internal n	etwork @
(Default) No	
Ves	

Parameter	Description
5 f C 5 r	Specify whether to retain the internal IP addresses of the ECS instances from the classic network. If you specify to retain the internal IP addresses of the ECS instances, you must specify how to create a vSwitch. If you specify not to retain the internal IP addresses of the ECS instances, you must select a vSwitch from the drop-down list.
	 (Default) Yes: retains the internal IP addresses of the ECS instances from the classic network. If (Default) Yes is selected, you must continue to specify vSwitch Creation Policy.
	If vSwitch Creation Policy is set to Automatic, a vSwitch is automatically created and associated with a CIDR block based on the internal IP addresses of the ECS instances. Make sure that the CIDR block that corresponds to the internal IP addresses of the ECS instances is not used by other instances. Otherwise, the vSwitch cannot be created.
	 Note If you set Destination VPC or Create a VPC to (Default) Automatically create a VPC, CIDR block: 10.0.0.0/8, Retain Internal IP Address is automatically set to (Default) Yes, and vSwitch Creation Policy is automatically set to Automatic and cannot be modified.
Retain Internal IP Address	If vSwitch Creation Policy is set to Manual, you must manually create a vSwitch in the specified destination zone based on the internal IP addresses of the ECS instances from the classic network.
	Note You can set vSwitch Creation Policy to Manual only when you select a user-created VPC that is associated with the 10.0.0.0/8 CIDR block for Destination VPC or Create a VPC .
	 No: does not retain the internal IP addresses of the ECS instances from the classic network. You must select a vSwitch from the drop-down list.
	Note If you cannot find the vSwitches that you created in the drop-down list, it may be because that the vSwitches are not located in the specified destination zone. Create a vSwitch in the destination zone. For more information, see Work with vSwitches.

Description
 Specify whether to allow mutual access over the internal network between migrated and unmigrated instances that are included in this migration plan. Configure this parameter based on the value of Retain Internal IP Address. Retain Internal IP Address set to (Default) Yes: If you do not want to allow mutual access over the internal network between migrated and unmigrated instances that are included in this migration plan, select (Default) No. If you want to allow mutual access over the internal network between migrated and unmigrated instances that are included in this migration plan, select (Default) No. If you want to allow mutual access over the internal network between migrated and unmigrated instances that are included in this migration plan, select Yes. Then, in the Select Instances step, select all ECS instances in the classic network that require mutual access over the internal network. You can schedule different migration times for these instances to control the order in which to migrate them.
Note ECS instances in the classic network that are not included in this migration plan cannot communicate with the ECS instances that are migrated to the specified VPC. After this migration plan is created, ECS instances cannot be added to or removed from it.
 Retain Internal IP Address set to No: If you do not want to allow mutual access over the internal network between migrated and unmigrated instances that are included in this migration plan, proceed to the Select Instances step. If you want to allow mutual access over the internal network between migrated and unmigrated instances that are included in this migration plan, configure ClassicLink to link these instances to the specified VPC before you migrate them. For more information, see Connect a classic network to a VPC.

6. In the Select Instances step, select ECS instance and click Next.

If you set **Retain Internal IP Address** to **(Default) Yes** and specify to allow mutual access over the internal network between migrated and unmigrated instances that are included in the migration plan, you must select all ECS instances in the classic network that require mutual access over the internal network. You can schedule different migration times for these instances to control the order in which to migrate them.

? Note ECS instances in the classic network that are not included in this migration plan cannot communicate with the ECS instances that are migrated to the specified VPC. After this migration plan is created, ECS instances cannot be added to or removed from it.

In the following figure, the section tagged ① shows the instances that you want to migrate first, and the section tagged ② shows the instances that you want to migrate afterward.

← Create Migration Plan									
For more information about instance migration and how to flexibly configure migration plans, Learn More									
Configure Migration Plan 2 Select Instances									
Plan Name	Plan Name Default migration plan Destination Zone Shanghai Zone 8								
Destination VPC	(Default) Automatically create a VPC	CIDR block: 10.0.0.0/8		Mac Address	Retention Policy	Retain private Mac address			
Change Internal IP Addresses No (Do Not Change)				Ensure Network Connectivity		No			
vSwitch Creation Policy Automatic				Destination Security Group		Clone security groups of classic network-type instances			
All Automatic V Search by ins	tance ID, name, or keyword	Tag filtering ∨	Configurations	Network Town	ID Address				
Kesource ID/Name	lags Status	Zone	Configurations	Network lype	IP Address				
ecs_c2v001	🗣 🕑 Running	Shanghai Zone B	ecs.xn4.small	Classic Network	10 I 10 (Public) 10 5(Internal)	1			
ecs_c2v002	🗞 🕑 Running	Shanghai Zone B	ecs.xn4.small	Classic Network	1(157 (Public) 1(55(Internal)	2			
Selected instances: 2									

7. In the Scheduled Migration step, set migration times for the instances and click Verify.

The instances are stopped and then started again during the migration. We recommend that you schedule to migrate your instances during off-peak hours. An individual migration time can be specified for each instance.

- To set a migration time for only a single instance at a time, click **Schedule Migration Time** in the **Actions** column corresponding to the instance.
- To set a migration time for multiple instances at a time, select the instances and click **Batch** Schedule Migration Time.
- To set the migration time for all of the instances at a time, click Set Unified Migration Time.

♥ Notice

- Set a late migration time for ECS instances that need to remain in the classic network but require mutual access with migrated ECS instances over the internal network. Before the migration time arrives, determine whether to migrate the ECS instances from the classic network.
- The following limits apply to the migration time that can be set for each instance:
 - The migration time cannot be earlier than the current time.
 - The migration time cannot be later than the expiration time of the instance.

After the migration plan is created, the three replicas that correspond to each of some disks are checked first. The amount of time the check takes is determined based on the disk size and the number of disks that are waiting in the queue. Set migration times as prompted.

- 8. In the **Verify** dialog box, read the migration considerations and verify whether your migration plan meets the specified requirements.
 - If your migration plan meets the specified requirements, select options and click **Confirm and Create**.
 - If your migration plan does not meet the requirements, error messages are displayed. You can troubleshoot the errors based on the error messages and modify parameters to create the migration plan again.

Step 2: Migrate the ECS instances

After the migration plan is created, the system migrates the specified ECS instances from the classic network to the destination VPC at the specified times.

Migrate Instances							
For more information and a second	about instance migration and how to fle	exibly configure migration plans, Learn More			×		
Create Migration Plan	Automatic V Search by migrat	tion plan properties or Q			С		
Migration Plan ID/Name	Plan Status 👅	Instances to Migrate	Created At	Actions			
mp-uf6 Default migration plan	Enabling	1/2	Apr 1, 2021, 17:00:00	Manage Cancel			
-							

During the migration, the system performs the following operations:

- 1. Stop the ECS instances to be migrated.
- 2. Migrate the computing and network resources of the ECS instances.
- 3. Start the migrated ECS instances.
- 4. Continue to migrate the disk data of the ECS instances.
- 5. Complete the migration.

? Note For a cross-zone migration, after the computing and network resources are migrated and the instances are started, the system continues to migrate disk data. Typically, it takes about 4 hours to migrate 100 GiB of disk data. During the migration, the I/O performance of disks degrades and snapshot- and disk-related features are not supported.

Step 3: Check the migration results

1.

- 2. Find the migrated ECS instances and click the ID of each of these instances.
- 3. On the Instance Details page, check whether the network type of the instance is VPC.

If the instance is migrated to the specified VPC, the network type of the instance changes to VPC.

Basic Information		Diagnose Instance Health 🚥 Start Restart Stop	Configure Security Group Rule Reset Password
ecs_c2v001 🗹 🥑 Running			
Instance ID i-uf	Connect	Region China (Shanghai)	
Public IP 10	Convert to EIP	Zone Shanghai Zone B	
Security Group sg-uf64	Add to Security Group	Hostname ecs001	Modify Hostname
Tags -	Edit Tags	Created At Apr 1, 2021, 16:46:00	
Description -	Modify Instance Description	Expires At Expires May 1, 2021, 23:59:59	Renew
CPU and Memory 1Cores 1 GiB		Cloud Disk 1	Reinitialize Disks
Operating System Alibaba Cloud Linux 2.1903 LTS 64-bit	Replace System Disk	Snapshot O	
Type ecs.xn4.small	Upgrade/Downgrade	Image ID aliyun_2_1903_x64_20G_alibase_20210120.vhd	Create Custom Image
Instance Family Shared Performance Basic		Current Bandwidth 1Mbps	Modify Bandwidth
Network Information			Bind Secondary ENI Change VPC
Network Type VPC		RDMA IP	
ENIs eni-uf6a		EIP ID -	
VPC vpc-uf6 7 🖸		VSwitch vsw-uf6xz∞ 3f ⊡	

4. Check the internal network and business runtime environments.

Scenario	Migration plan	What to do next
Migrate all ECS instances from the classic network to a VPC	 Set Destination VPC or Create a VPC to (Default) Automatically create a VPC, CIDR block: 10.0.0.0/8. Set Ensure interconnections between the migrated instances and the classic network-type instances to (Default) No. 	Check whether your business system runs normally.
Migrate some ECS instances to a VPC and retain other ECS instances in the classic network	 Set Destination VPC or Create a VPC to (Default) Automatically create a VPC, CIDR block: 10.0.0.0/8. Set Ensure interconnections between the migrated instances and the classic network-type instances to Yes. 	Check whether your business system runs normally.

Scenario	Migration plan	What to do next
Other scenario	Set Destination VPC or Create a VPC to a VPC that is associated with a CIDR block other than 10.0.0.0/8.	 i. Check network connectivity. ii. In this scenario, Retain Internal IP Address cannot be set to No. If your business is connected by using internal IP addresses, you must configure new internal IP addresses. iii. Check whether your business system runs normally.

Post-migration considerations

1. If an ECS instance runs a Linux operating system and is assigned a different internal IP address after the instance is migrated, you must modify the */etc/hosts* file of the instance.

::1 1	localhos	t	localhos	st.localdomain	localhos	t6	localhost6.lo	caldomain6
127.0.0.1	1	localhos	t	localhost.local	Ldomain	localhos	t4 local	nost4.localdomain4
72.16.		ecs	ecs					

- i. Run the vi /etc/hosts command to open the hosts file.
- ii. Press the /key to enter the edit mode.
- iii. Change the original internal IP address to the new internal IP address for the instance.
- iv. Press the Esc key to exit the edit mode.
- v. Enter :wq and press the Enter key.
- 2. If you have set Retain Internal IP Address to No in the migration plan, remove the internal IP addresses that are no longer used from the whitelists of other cloud services after the migration,

such as AparaDB RDS, SLB, Object Storage Service (OSS), and Container Service for Kubernetes.

3. If an instance is migrated across zones, its connectivity with other Alibaba Cloud services such as ApsaraDB RDS, ApsaraDB for Redis, and ApsaraDB for MongoDB may be affected. Adjust application configurations in a timely manner. For example, you can migrate the corresponding RDS instances to the same zone as the ECS instance to ensure connectivity.

For more information, see Migrate an ApsaraDB RDS for MySQL instance across zones in the same region.

- 4. If you have not restarted an instance or upgraded its kernel for an extended period of time, problems may occur after the instance is migrated. For example, a file system check (fsck) may be performed, configuration changes may become invalid, or the instance may be unable to start.
- 5. (Optional) Software authorization codes change because NICs are deleted.

If software is associated with a MAC address on your ECS instance and the software vendor approves the migration certificate issued by Alibaba Cloud, you can re-authorize the instance to use the software. If an error occurs, you must modify the configurations or roll back the instance.

6. (Optional) If you have not restarted an ECS instance for an extended period of time or if you have not restarted an instance after its kernel is upgraded, the system checks the file systems of the

instance and updates the configurations of the instance when the instance is restarted. If your ECS instance cannot be started, Submit a ticket in a timely manner to contact Alibaba Cloud.

FAQ

• Why am I unable to open websites, use services, or read data from or write data to databases on an instance after the instance is migrated from the classic network to a VPC?

This may be because traffic is not allowed on the required communication ports in the new security groups of your instance. We recommend that you clone the original security groups. For more information, see Clone a security group.

• After an instance is migrated, some software cannot be used and I am prompted that the authorization code is expired or invalid or that no authorization code exists for the software. Why?

This issue may occur due to one of the following reasons:

- The software vendor has not approved the migration certificate issued by Alibaba Cloud. We recommend that you contact the software vendor or channel partner to submit a verification form for re-authorization.
- The software was associated with a MAC address to register to your instance. Some software is
 registered to a valid environment by associating MAC addresses. After an ECS instance is migrated
 to a VPC, only the public or private MAC address of the instance is retained. If the MAC address with
 which a piece of software was associated to register is deleted, an authorization error occurs. We
 recommend that you contact the software vendor to check whether the software was associated
 with a MAC address to register to your instance. If yes, you must re-associate the MAC address of
 the instance with the software. For more information, see Overview.
- Why am I no longer able to use the FTP service on an instance after the instance is migrated?

After your ECS instance is migrated, its public NIC is deleted and the FTP service becomes unavailable. We recommend that you perform the following operations:

- i. Convert the system-assigned public IP address of an instance that is located in a VPC to an EIP.
- ii. Associate an EIP with a secondary ENI in cut-through mode.

(?) Note Some retired instance types and entry-level instance types of the previous generation do not support ENIs. If the instance type of your instance does not support ENIs, upgrade the instance to an instance type that supports ENIs before you perform the preceding operations. For more information, see Overview of instance configuration changes.

• I cannot find data disks on some Windows instances after the instances are migrated. What do I do?

After some Windows instances are migrated, the disks attached to them go offline. We recommend that you perform the following steps to configure the disks to automatically go back online. For more information, see Methods for processing offline disks on Windows ECS instances.

- i. Log on to the ECS console.
- ii. In the left-side navigation pane, choose Maintenance & Monitoring > Cloud Assistant.
- iii. Click **Create or Run Command** to create and run a Cloud Assistant command.

In the Create Command panel, configure parameters described in the following table. For the parameters that are not described in the table, accept the default values. For more information, see Use the immediate execution feature.

Parameter	Value
Command Type	PowerShell
Command	<pre>@("san policy=onlineall") diskpart</pre>
Select Instances	One or more Windows instances.

- iv. Click Execute and Save.
- Why am I unable to transfer files to or from an instance over FTP after the instance is migrated from the classic network to a VPC?

ECS instances in the classic network have both public and private NICs, whereas ECS instances in VPCs have only ENIs, which are private NICs. If your applications are configured to recognize only public IP addresses, you must reconfigure the applications.

Most FTP clients access FTP servers in passive mode. In passive mode, FTP servers must communicate their IP addresses to FTP clients. In VPCs, public IP addresses cannot be recognized and FTP servers send their internal IP addresses to FTP clients. When the clients use the internal IP addresses to access the servers, errors occur.

When you use an ECS instance located in a VPC as an FTP server, we recommend that you communicate the public IP address of the instance to the FTP server program. The procedures to communicate the public IP addresses of ECS instances vary based on the types of FTP server programs. Select a procedure that is suitable for your FTP server program. In the following example, vsftpd is used. Open the configuration file of vsftpd and add the following content to the file:

```
listen_ipv6=NO
pasv_address=<PublicIP>
```

(?) Note Replace *<PublicIP>* with the system-assigned public IP address or EIP of your instance. If an EIP is associated with the instance, we recommend that you use the EIP.

References

- Change the network type from classic network to VPC
- Change the network type of an ApsaraDB RDS for MySQL instance
- Configure the hybrid access solution for an ApsaraDB RDS for MySQL instance

12.Connect a classic network to a VPC

This topic describes how to connect a classic network to a VPC. You can set up a ClassicLink connection so that ECS instances of the classic network type can access cloud resources in a VPC through the intranet.

Prerequisites

Make sure that you are aware of the limits of ClassicLink. For more information, see Overview.

Procedure

- 1. Log on to the VPC console.
- 2. Select the region of the target VPC, and click the ID of the target VPC.
- 3. On the VPC Details page, click Enable ClassicLink. In the displayed dialog box, click OK.

4.

5.

6.

- 7. Find the target ECS instance of the classic network type, and then choose **More > Network and Security Group > Connect to VPC**.
- 8. In the displayed dialog box, select the target VPC and click **OK**, and then click the security group configuration link.

Connect to VPC	\times
Connected VPC: : 🇆 tanVPC / vpc-bp1c0408b97edxpot1j3u 🗸 ClassicLin	ık
When you connect to a VPC, you must configure the security group rules to connectivity.	ensure
Go to the instance security group list and add ClassicLink rules	

9. Click Add ClassicLink Rules and configure the security rule according to the following information. Then, click OK.

Configuration	Description
Classic Security Group	Display the classic network security group.

Configuration	Description
Select VPC Security Group	Select a security group to use. Up to five security groups can be selected.
Mode	 Select one of the following modes: <i>Classic <=> VPC</i>: The connected resources can access each other (recommended). <i>Classic => VPC</i>: Authorize the classic ECS instance to access cloud resources in the connected VPC. <i>VPC => Classic</i>: Authorize the cloud resources in the connected VPC to access the classic ECS instance.
Protocol Type and Port Range	Select the protocol and port used for the communication. The port must be in the form of xx/xx. For example, if port 80 is used, enter <i>80/80</i> .
Priority	Set the priority for the rule. A smaller number represents a higher priority, for example, <i>1</i> .
Description	Enter a description for the security rule.

10. Return to the ECS console. On the Instance List page, click the Column Filter icon in the upperright corner, and then select the Connection Status check box. Then, click OK.

	Column Filter ×	1
* Select instance attributes or enter keywords		Advanced Search 💆 🗘
T Filters : Network Type: Classic × Clear	Ø Operating System Ø Tags Ø Monitoring Ø Zone Ø	
Instance ID/Name Tags Monitoring Zone II	Ø IP Address Ø Status Ø Network Type Ø Configuration	r Connection Status Actions
1 has a more (7a 1 fam)	VPC Details Instance Type Family 🗭 Billing Method Automatic Renewal	
IZatov mZ Hangzhou Zone I	0 SSH Key Pair 2 🗹 Connection Status 🛛 RAM Role 👘 Cluster ID	Connected Manage Connect 20:40 vpc- bp1c0 vp0t1j3u Change Instance Type More ▼
	Stopped By Dedicated Host Deployment Sets	
Start Stop Restart Reset Password Renew		Total: 1 item(s), Per Page: 20 \bullet item(s) << 1 \rightarrow »
	3 (x	

If **Connection Status** is **Connected**, ECS instances of the classic network are connected to the VPC network.

13.Network FAQ

This topic provides answers to frequently asked questions about networks used by Elastic Compute Service (ECS) instances.

- FAQ about network performance
 - What is the packet loss rate when instances within different regions communicate over the Internet?
 - How is the network latency for instances within the same region that communicate over the internal network?
 - How is the performance of connections guaranteed for instances for which the maximum number of connections is not specified?
 - What do I do if the performance of an ECS instance is unstable when a UDP PPS test or TCP bandwidth test is performed on the instance?
- FAQ about public bandwidth
 - What are the inbound and outbound bandwidths of ECS instances?
 - I purchased a public bandwidth of 5 Mbit/s for an ECS instance. How is this bandwidth used as outbound or inbound bandwidth of the instance?
 - Is public bandwidth exclusive to each ECS instance, or is public bandwidth shared among multiple instances?
 - How am I billed for the network usage of ECS instances?
 - Why is 200 Kbit/s of inbound traffic already consumed on a new ECS instance?
 - How do I view the Internet traffic bills of an ECS instance?
 - Why is the bandwidth usage of my ECS instance displayed in the CloudMonitor console different from that displayed in the ECS console?
 - My ECS instance has been stopped. Why am I still being charged for its outbound traffic on a payas-you-go basis?
- FAQ about IP addresses
 - How do I query the IP addresses of ECS instances?
 - How do I disable the public NIC of an ECS instance?
- FAQ about network access and traffic direction
 - When I attempt to access a website on an ECS instance, a message similar to "Sorry, your access has been blocked because the requested URL may pose a security threat to the website" appears. Why?
 - After I configure a secondary private IP address for a Windows instance, the instance cannot connect to the Internet. Why?
 - An abnormal logon has been detected on one of my ECS instances. What do I do?
 - What is traffic scrubbing?
 - How do I cancel traffic scrubbing for an ECS instance?
 - How do I request reverse lookup for an ECS instance?
 - Can an IP address point to multiple reverse lookup domain names?
- FAQ about public IP addresses
 - o Can I channe the nublic IDvA address of an instance after the instance has been created?

- carrientinge the public in v+ address of an instance after the instance has been created;
- Why am I unable to find the option to change the public IP address of an ECS instance in the ECS console?
- Can I change the private IP address of an instance?
- If no public IPv4 address was assigned to an ECS instance when the instance was being created, how do I assign a public IP address to the instance?
- FAQ about network basics
 - What is a BGP data center?
 - What are WAN and LAN?
 - What is CIDR?
 - How do I express a subnet mask?
 - How do I plan subnets?
- FAQ about quotas
 - How do I view resource quotas?

What is the packet loss rate when instances within different regions communicate over the Internet?

When instances within different regions communicate by using Cloud Enterprise Network (CEN), these instances use Alibaba Cloud backbone networks to transmit data. Alibaba Cloud aims to provide network services with an hourly packet loss rate at the 99th percentile of less than 0.0001%.

How is the network latency for instances within the same region that communicate over the internal network?

You can achieve minimal latency when instances within the same zone and region communicate with each other over the internal network. The one-way latency at the 99th percentile is less than 180 us for communication between instances within the same zone.

How is the performance of connections guaranteed for instances for which the maximum number of connections is not specified?

If an instance family does not have the maximum number of connections specified, this instance family does not ensure that a specific maximum number of connections can be established to a single instance. We recommend that you perform business stress tests on instances to choose appropriate instance types.

Note After a connection is established, the connection counts towards the number of connections before its aging period ends. The displayed number of connections may be greater than the number of connections actually in use.

What do I do if the performance of an ECS instance is unstable when a UDP PPS test or TCP bandwidth test is performed on the instance?

When a network performance test is performed on an ECS instance, the test result may be affected by a number of factors. These factors include the common performance tuning methods such as non-uniform memory access (NUMA) topology adaptation, binding vCPUs for tasks, and binding vCPUs for interrupts.

For example, during a single-stream TCP bandwidth test, if a receive task such as a netserver process and a network interface controller (NIC) receive queue interrupt are bound to the same vCPU, the NIC triggers an interrupt to interrupt the receive task when the NIC receives data frames. If the receive task is frequently interrupted, the test result may not meet your expectations. In this case, you can bind the receive task and the NIC receive queue interrupt to different vCPUs and obtain a better test result by using the performance advantages of multiple vCPUs.

What are the inbound and outbound bandwidths of ECS instances?

Bandwidth type	Description
Inbound bandwidth	 The bandwidth for inbound traffic for an ECS instance, including the following traffic: Traffic generated when you download external resources to the ECS instance Traffic generated when you upload resources to the ECS instance by using an FTP client
Outbound bandwidth	 The bandwidth for outbound traffic for an ECS instance, including the following traffic: Traffic generated when the ECS instance provides external access Traffic generated when you download resources from the ECS instance by using an FTP client

I purchased a public bandwidth of 5 Mbit/s for an ECS instance. How is this bandwidth used as outbound or inbound bandwidth of the instance?

The 5 Mbit/s that you purchased is used as the outbound bandwidth for the instance. The inbound bandwidth of this instance is capped at 10 Mbit/s.

- Out bound bandwidth is consumed when data is transferred from the ECS instance. The maximum out bound bandwidth of an ECS instance is capped at 100 Mbit/s or 200 Mbit/s regardless of whether the instance resides in a virtual private cloud (VPC) or the classic network. The maximum available out bound bandwidth depends on the billing method of the instance.
- Inbound bandwidth is consumed when data is transferred to the ECS instance. The maximum inbound bandwidth is determined by the outbound bandwidth:
 - If the outbound bandwidth is less than 10 Mbit/s, the maximum inbound bandwidth is 10 Mbit/s.
 - If the outbound bandwidth is greater than 10 Mbit/s, the maximum inbound bandwidth is the same as the purchased outbound bandwidth.

Notice When the **pay-by-traffic** billing method for network usage is used, the maximum inbound and outbound bandwidths are used as the upper limits of bandwidths instead of guaranteed performance specifications. In scenarios where demand outstrips resource supplies, these maximum bandwidths may be limited. If you want guaranteed bandwidths for your instance, use the **pay-by-bandwidth** billing method for network usage.

Is public bandwidth exclusive to each ECS instance, or is public bandwidth shared among multiple instances?

The public bandwidth of each instance is exclusive to the instance.

How am I billed for the network usage of ECS instances?

For more information about billing for the network usage of ECS instances, see Public bandwidth.

Why is 200 Kbit/s of inbound traffic already consumed on a new ECS instance?

This traffic was generated by Address Resolution Protocol (ARP) broadcast packets. Each ECS instance is assigned to a large CIDR block. When the gateway receives an ARP request packet for an ECS instance, the gateway broadcasts this packet to all ECS instances within the same CIDR block. The new instance also receives the packet. If the request is not destined for the new instance, the instance does not reply with an ARP reply packet.

How do I view the Internet traffic bills of an ECS instance?

To view the Internet traffic bills of an ECS instance, perform the following steps:

1.

- 2. In the top navigation bar, choose Expenses > User Center.
- 3. In the left-side navigation pane, choose **Spending Summary > Spending Summary**.
- Click the Bills tab. Specify a billing cycle. Set Product Name to Elastic Compute Service, Product Detail to Elastic Computing (Pay-As-You-Go), and Subscription Type to Pay-As-You-Go.
- 5. Click **Export Billing Overview (CSV)**. In the Export Billing Overview (CSV) dialog box, enter the CAPT CHA and click **OK**.
- 6. On the **Export Record** page, wait until the status of the exported file changes to **Exported** and click **Download** in the Actions column.
- 7. Open the exported CSV file to view the Internet traffic bills of the ECS instance.

Why is the bandwidth usage of my ECS instance displayed in the CloudMonitor console different from that displayed in the ECS console?

ECS instances function as backend servers of Server Load Balancer (SLB) instances and use the Layer 7 HTTP forwarding model. In this forwarding model, SLB instances forward client requests to ECS instances, and the ECS instances use their own outbound bandwidth to return responses to the corresponding users. The bandwidth consumed by these responses is not displayed in the ECS console, but the traffic generated by the responses counts towards the outbound traffic of the SLB instances and is displayed in the CloudMonitor console. Therefore, the bandwidth usage of your ECS instance displayed in the CloudMonitor console is different from that displayed in the ECS console.

My ECS instance has been stopped. Why am I still being charged for its outbound traffic on a pay-as-you-go basis?

- Problem description: Your instance is in the **Stopped** state in the ECS console but is in the **Cleaning** state in the Anti-DDoS Basic console. You are charged for outbound traffic from the instance on a pay-as-you-go basis every hour.
- Cause: HTTP flood protection is enabled for the ECS instance. When HTTP flood protection is enabled, the security mechanism sends probe packets to potential attack sources. Therefore, a large

volume of outbound traffic is generated.

• Solution: Disable HTTP flood protection for the ECS instance.

How do I query the IP addresses of ECS instances?

• Linux instance

Run the *ifconfig* command to view NIC information. You can view the IP addresses, subnet masks, gateways, Domain Name System (DNS) servers, and MAC addresses in the command output.

• Windows instance

In Command Prompt, run the ipconfig /all command to view NIC information. You can view the IP addresses, subnet masks, gateways, DNS servers, and MAC addresses in the command output.

How do I disable the public NIC of an ECS instance?

- Linux instance
 - i. Run the <code>ifconfig</code> command to view the name of the public NIC of the instance.
 - ii. Run the ifdown command to disable the public NIC. For example, if the name of the public NIC
 is eth1 , enter ifdown eth1 .

Onte You can run the if up command to re-enable the NIC. For example, if the name of the public NIC is eth1, enter if up eth1.

- Windows instance
 - i. In Command Prompt, run the <code>ipconfig</code> command to view information about the public NIC.
 - ii. Open the **Control Panel** and click View network status and tasks in the Network and Internet section. In the **Network and Sharing Center** window, click **Change adapter settings** in the left-side navigation pane to disable the public NIC.

When I attempt to access a website on an ECS instance, a message similar to "Sorry, your access has been blocked because the requested URL may pose a security threat to the website" appears. Why?

- Problem description: When you attempt to access a website built on an ECS instance, you are prompted with a message similar to "Sorry, your access has been blocked because the requested URL may pose a security threat to the website."
- Cause: Web Application Firewall (WAF) has identified your access request to the URL as an attack and blocked your access.
- Solution: Add the source public IP address that you use to access the website to the WAF whitelist. For more information, see Avoid Anti-DDoS Basic false positives by using a whitelist.

After I configure a secondary private IP address for a Windows instance, the instance cannot connect to the Internet. Why?

- Problem description: After you configure a secondary private IP address for a Windows instance, the instance cannot connect to the Internet.
- Cause: In Windows 2008 and later, the longest prefix match algorithm is used to select next hop IP
addresses based on destination IP addresses of outbound traffic. This may lead to network connection failures.

• Solution: Run the **Netsh** command with skipassource set to true to configure a secondary private IP address for the Windows instance.

Netsh command:

Netsh int ipv4 add address <Interface> <IP Addr> [<Netmask>] [skipassource=true]

The following table describes the parameters in the Netsh command.

Parameter	Description	Example value
<interface></interface>	The network interface with which to associate the secondary private IP address	'Ethernet'
<ip addr=""></ip>	The secondary private IP address	192.168.0.100
<netmask></netmask>	The mask of the secondary private IP address	255.255.255.0

Sample command:

```
Netsh int ipv4 add address 'Ethernet' 192.168.0.100 255.255.255.0 skipassource=true
```

An abnormal logon has been detected on one of my ECS instances. What do I do?

Perform the following operations to solve the problem:

- 1. Check the logon time to see whether the logon was performed by yourself or another administrator.
- 2. If the logon was not performed by yourself or another administrator, it is an unauthorized logon. Perform the following steps:
 - i. Reset the password. For more information, see Reset the logon password of an instance.
 - ii. Check whether the ECS instance is infected.
 - iii. Configure security groups to allow access only from specific IP addresses. For more information, see Security groups for different use cases.

What is traffic scrubbing?

The traffic scrubbing service monitors inbound traffic to ECS instances in real time and identifies abnormal traffic such as DDoS attacks. By default, Anti-DDoS Basic is enabled on ECS instances to provide traffic scrubbing. When ECS instances are under attack, the traffic scrubbing service detects the attack and scrubs malicious traffic without affecting ECS instance services. When suspicious traffic is detected, suspicious traffic is redirected from the destination network to a scrubbing device. The device identifies and removes malicious traffic and then returns legitimate traffic to the network to be forwarded to the ECS instances.

How do I cancel traffic scrubbing for an ECS instance?

1. Log on to the Alibaba Cloud Security Anti-DDoS Basic console.

- 2. Click the ECS tab. In the ECS instance list, find the IP address of an ECS instance that is in the Cleaning state and click **View Details**.
- 3. Click Cancel cleaning.

C Home Products -				۹	A 💴 Billing M:	anagement English 🥘		
III Anti-DD	os Service	Anti-DDoS Basic						
DTplus	Anti-DDoS Basic	Asia Pacific SE 1 (Singapore)	Asia Pacific SOU 1 (Mumbai)	North China 1 North	China 2 China North	h 3 (Zhangjiakou) East China 2		
- Security - Anti-O	1 (Virginia)	China North 5 (Huhehaote)	cn-hongkong Asia Pac	fic SE 3 (Kuala Lumpur)	Middle East 1 (Dubai)	Asia Pacific SE 2 (Sydney)		
Overview Secu	rity Report	Central 1 (Frankfurt) Asia	Pacific NE 1 (104)0) US We	st 1 (allcon valley)				
🔅 Anti-DDoS Sei	ECS Server Loa	I Balancer EIP		_				
	Instance IP addre	B Pisase enter the instanc	e IP address Search					
Server Guard Insta	nce List Instance IP address	name Security Inform	sation(All) + Cleaning Trig	ger Value Black Ho	le Trigger Value(M)	Operation		
SSL Certificates Open	ation Log	Normal	BPS: 500M F	PS: 100000 5100		View details		
Mobile Security								
Domains & Websites	dig signs southly	(DException(ck	eaning) BPS: 500M F	IPS: 100000 5100	Cance	I cleaning View details		

How do I request reverse lookup for an ECS instance?

Reverse lookup is used in mail services to reject mail from IP addresses mapped to unregistered domain names. Most spammers use dynamic IP addresses or IP addresses mapped to unregistered domain names to send unwanted mail and avoid being tracked. When reverse lookup is enabled on a mail server, the server rejects mail sent from dynamic IP addresses or unregistered domain names to reduce the amount of spams received.

You can submit a ticket to request reverse lookup for your ECS instance. To make your ticket easier to process, we recommend that you specify the region, public IP address, and registered domain name of your ECS instance in the ticket.

After your request is approved, you can run the **dig** command to check whether reverse lookup takes effect on your instance. Example:

```
dig -x 121.196.255.** +trace +nodnssec
```

If reverse lookup takes effect on your instance, a command output similar to the following one is displayed:

1.255.196.121.in-addr.arpa. 3600 IN PTR ops.alidns.com.

Can an IP address point to multiple reverse lookup domain names?

No, each IP address can point only to a single reverse lookup domain name. For example, you cannot configure the IP address 121.196.255.** to resolve to multiple domain names such as mail.abc.com, mail.ospf.com, and mail.zebra.com.

Can I change the public IPv4 address of an instance after the instance has been created?

You can change the public IPv4 address of an instance within 6 hours after the instance is created. For more information, see Change the public IP address of an instance.

After 6 hours, the instance network type determines whether the public IP address of the instance can be changed.

• For an instance in a VPC, you can change the public IP address of the instance by converting the IP

address into an elastic IP address (EIP). Then, to assign a new public IP address, you can disassociate the EIP from the instance and associate a new EIP with the instance or upgrade the public bandwidth of the instance. For more information, see Convert the public IP address of a VPC-type instance to an EIP.

• The public IP addresses of instances in the classic network cannot be changed. However, you can convert the public IP address of an instance into an EIP when you release the instance. For more information, see Convert the public IP address of an instance in the classic network into an EIP.

Why am I unable to find the option to change the public IP address of an ECS instance in the ECS console?

- Within 6 hours after a pay-as-you-go instance is created:
- More than 6 hours after the instance is created: You cannot change the public IP address, and the **Change Public IP Address** option is not displayed.

Can I change the private IP address of an instance?

- You can change the private IP addresses of instances in VPCs. For more information, see Modify a private IP address.
- You cannot change the private IP addresses of instances in the classic network.

If no public IPv4 address was assigned to an ECS instance when the instance was being created, how do I assign a public IP address to the instance?

- Apply for and associate an elastic IP address (EIP) with the instance. For more information, see Apply for an EIP of *EIP documentation*.
- Modify the public bandwidth of the instance to allocate a system-assigned public IP address. For information about how to perform this operation on subscription instances, see Overview of instance configuration changes. For information about how to perform this operation on pay-as-you-go instances, see Modify the bandwidth configurations of pay-as-you-go instances.

What is a BGP data center?

Border Gateway Protocol (BGP) is used to connect autonomous systems (AS) over the Internet. The main purpose of BGP is to control route propagation and select the optimal routes.

China Netcom, China Telecom, China Railcom, and some large privately owned IDC service providers all have autonomous system numbers (ASNs). Most major network carriers in China use BGP to implement multi-line connections between their ASNs.

To implement multi-line interconnection in this manner, an IDC must obtain a CIDR block and an ASN from the China Internet Network Information Center (CNNIC) or Asia-Pacific Network Information Center (APNIC), and then broadcast this CIDR block to the networks of other carriers by using BGP. After BGP is used to connect different networks, the backbone routers of the network carriers determine the optimal routes to the CIDR block of the IDC to ensure high-speed access for users of different network carriers.

What are WAN and LAN?

• A wide area network (WAN) is also known as an external or public network. A WAN is a telecommunications network that connects smaller networks such as LANs and metro area networks (MANs). Each WAN extends over a large geographical area that can range in size from as small as a

city or as large as an entire continent to provide telecommunications services and form an international telecommunications network. WAN is not the same as the Internet.

• A LAN is also known as an internal network. A LAN is a network that interconnects computers within a small area. Users can manage files, share application software and printers, schedule work for work groups, and communicate with each other such as by sending emails or faxes within a LAN. A LAN is a closed network that can be as small as consisting of two computers in an office or as large as consisting of thousands of computers in a company. In Alibaba Cloud, ECS instances of the same network type within the same region can communicate with each other over the internal network. ECS instances within different regions are isolated from each other.

What is CIDR?

CIDR is an addressing scheme for the Internet that allows for IP addresses to be assigned in a more efficient manner than the traditional scheme based on classes A, B, and C. CIDR notation is used to denote IP addresses and IP ranges. It consists of an IP address and a forward slash followed by a decimal number that denotes how many bits are in the network prefix.

• Example 1: Convert a CIDR block into an IP address range

• Example 2: Convert an IP address range into a CIDR block

How do I express a subnet mask?

You can use one of the following methods to express a subnet mask:

• Use dotted decimal notation.

The default subnet mask of a Class A network is 255.0.0.0.

• Append a forward slash (/) and a number ranging from 1 to 32 to the end of an IP address to define a subnet mask. This number indicates the length of the network identification bit in the subnet mask.

Example: 192.168.0.3/24.

How do I plan subnets?

For more information about the best practices for planning subnets, see Plan networks.

How can I view the resource quota?

For more information about how to view the limits and quotas of resources, see 使用限制.