

Elastic Cloud Server

User Guide

Issue 01
Date 2024-04-15



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2024. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Cloud Computing Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.

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

Notice

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

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

Huawei Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road
Qianzhong Avenue
Gui'an New District
Gui Zhou 550029
People's Republic of China

Website: <https://www.huaweicloud.com/intl/en-us/>

Contents

1 Service Overview	1
1.1 What Is ECS?	1
1.2 ECS Advantages	2
1.3 ECS Application Scenarios	4
1.4 Notes and Constraints on Using ECSs	5
1.5 ECS and Other Services	6
1.6 Instances	7
1.6.1 ECS Overview	7
1.6.2 ECS Lifecycle	7
1.6.3 ECS Types	9
1.7 x86 ECS Specifications and Types	10
1.7.1 General-Purpose ECSs	10
1.7.2 General Computing-plus ECSs	15
1.7.3 Memory-optimized ECSs	24
1.7.4 Disk-intensive ECSs	30
1.7.5 Ultra-high I/O ECSs	33
1.7.6 High-Performance Computing ECSs	41
1.7.7 GPU-accelerated ECSs	45
1.7.8 AI-accelerated ECSs	69
1.8 Kunpeng ECS Specifications and Types	73
1.8.1 Kunpeng General Computing ECSs	73
1.8.2 Kunpeng General Computing-plus ECSs	76
1.8.3 Kunpeng Memory-optimized ECSs	80
1.8.4 Kunpeng Ultra-high I/O ECSs	84
1.8.5 Kunpeng AI Inference-accelerated ECSs	89
1.9 Images	93
1.9.1 Image Types	93
1.9.2 Cloud-Init	94
1.10 EVS Disks	95
1.11 Network	96
1.12 Security	97
1.12.1 Fault Recovery	97
1.13 Permissions Management	98

1.14 Region and AZ.....	104
2 Getting Started.....	106
2.1 Creating an ECS.....	106
2.1.1 Overview.....	106
2.1.2 Step 1: Configure Basic Settings.....	106
2.1.3 Step 2: Configure Network.....	108
2.1.4 Step 3: Configure Advanced Settings.....	110
2.1.5 Step 4: Confirm.....	113
2.2 Logging In to an ECS.....	114
2.3 Initializing EVS Data Disks.....	115
2.3.1 Scenarios and Disk Partitions.....	115
2.3.2 Initializing a Windows Data Disk (Windows Server 2008).....	116
2.3.3 Initializing a Windows Data Disk (Windows Server 2019).....	123
2.3.4 Initializing a Linux Data Disk (fdisk).....	131
2.3.5 Initializing a Linux Data Disk (parted).....	137
2.3.6 Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2008).....	143
2.3.7 Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2012).....	151
2.3.8 Initializing a Linux Data Disk Larger Than 2 TiB (parted).....	159
3 Instances.....	166
3.1 Viewing ECS Information.....	166
3.1.1 Viewing ECS Creation Statuses.....	166
3.1.2 Viewing Failed Tasks.....	167
3.1.3 Viewing ECS Details (List View).....	168
3.1.4 Exporting ECS Information.....	168
3.2 Logging In to a Windows ECS.....	169
3.2.1 Login Overview.....	169
3.2.2 Remotely Logging In to a Windows ECS (Using VNC).....	170
3.2.3 Remotely Logging In to a Windows ECS (Using MSTSC).....	171
3.2.4 Remotely Logging In to a Windows ECS (from a Linux Computer).....	176
3.2.5 Remotely Logging In to a Windows ECS (from a Mobile Terminal).....	178
3.2.6 Remotely Logging In to a Windows ECS (from a macOS Server).....	183
3.3 Logging In to a Linux ECS.....	186
3.3.1 Login Overview.....	186
3.3.2 Remotely Logging In to a Linux ECS (Using VNC).....	187
3.3.3 Remotely Logging In to a Linux ECS (Using an SSH Password).....	189
3.3.4 Remotely Logging In to a Linux ECS (from a Mobile Terminal).....	191
3.3.5 Remotely Logging In to a Linux ECS (from a macOS Server).....	203
3.4 Managing ECSs.....	203
3.4.1 Changing ECS Names.....	203
3.4.2 Reinstalling the OS.....	204
3.4.3 Changing the OS.....	205
3.4.4 Managing ECS Groups.....	208

3.4.5 Changing the Time Zone for an ECS.....	210
3.4.6 Automatically Recovering ECSs.....	212
3.4.7 Obtaining ECS Console Logs.....	213
3.5 Modifying ECS Specifications.....	215
3.5.1 General Operations.....	216
3.6 Obtaining Metadata and Passing User Data.....	217
3.6.1 Obtaining Metadata.....	217
3.6.2 Passing User Data to ECSs.....	225
3.7 (Optional) Configuring Mapping Between Hostnames and IP Addresses.....	233
3.8 (Optional) Installing a Driver and Toolkit.....	234
3.8.1 GPU Driver.....	234
3.8.2 Installing a GRID Driver on a GPU-accelerated ECS.....	235
3.8.3 Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS.....	247
3.8.4 Obtaining a Tesla Driver and CUDA Toolkit.....	261
4 Images.....	264
4.1 Overview.....	264
4.2 Creating an Image.....	265
5 EVS Disks.....	267
5.1 Adding a Disk to an ECS.....	267
5.2 Attaching an EVS Disk to an ECS.....	268
5.3 Detaching an EVS Disk from a Running ECS.....	269
5.4 Expanding the Capacity of an EVS Disk.....	271
5.5 Enabling Advanced Disk.....	272
6 CBR.....	274
6.1 Overview.....	274
6.2 Backing Up an ECS.....	278
7 NICs.....	281
7.1 Attaching a Network Interface.....	281
7.2 Detaching a Network Interface.....	283
7.3 Modifying a Private IP Address.....	283
7.4 Managing Virtual IP Addresses.....	284
7.5 Enabling NIC Multi-Queue.....	284
7.6 Dynamically Assigning IPv6 Addresses.....	290
8 EIP.....	309
8.1 Binding an EIP.....	309
8.2 Unbinding an EIP.....	309
8.3 Changing an EIP Bandwidth.....	310
8.4 Enabling Internet Connectivity for an ECS Without an EIP.....	310
9 Security.....	314
9.1 Security Groups.....	314

9.1.1 Overview.....	314
9.1.2 Default Security Group and Rules.....	315
9.1.3 Security Group Configuration Examples.....	316
9.1.4 Configuring Security Group Rules.....	321
9.1.5 Changing a Security Group.....	323
10 Passwords and Key Pairs.....	324
10.1 Passwords.....	324
10.1.1 Application Scenarios for Using Passwords.....	324
10.1.2 Resetting the Password for Logging In to an ECS on the Management Console.....	325
10.1.3 Changing the Login Password on an ECS.....	327
10.1.4 Resetting the Password for Logging In to a Windows ECS.....	328
10.1.5 Resetting the Password for Logging In to a Linux ECS.....	331
10.2 One-Click Password Reset Plug-in.....	334
10.2.1 Installing the One-Click Password Reset Plug-in on an ECS.....	334
10.3 Key Pairs.....	340
10.3.1 Application Scenarios for Using Key Pairs.....	340
10.3.2 (Recommended) Creating a Key Pair on the Management Console.....	341
10.3.3 Creating a Key Pair Using PuTTYgen.....	342
10.3.4 Importing a Key Pair.....	346
10.3.5 Obtaining and Deleting the Password of a Windows ECS.....	347
10.3.5.1 Obtaining the Password for Logging In to a Windows ECS.....	347
10.3.5.2 Deleting the Initial Password for Logging In to a Windows ECS.....	348
11 Permissions Management.....	350
11.1 Creating a User and Granting ECS Permissions.....	350
11.2 ECS Custom Policies.....	351
12 Resources.....	353
12.1 Quota Adjustment.....	353
13 Monitoring.....	354
13.1 Monitoring ECSs.....	354
13.2 Basic ECS Metrics.....	354
13.3 OS Monitoring Metrics Supported by ECSs with the Agent Installed.....	359
13.4 Setting Alarm Rules.....	401
13.5 Viewing ECS Metrics.....	403
14 FAQs.....	404
14.1 ECS Overview.....	404
14.1.1 What Are the Precautions for Using ECSs?.....	404
14.1.2 What Can I Do with ECSs?.....	404
14.2 Creation and Deletion.....	404
14.2.1 Why Does the Failures Area Show an ECS Creation Failure But the ECS List Displays the Created ECS?.....	404
14.2.2 How Quickly Can I Obtain an ECS?.....	405

14.2.3 What Happens After I Click the Delete Button?.....	405
14.2.4 Can a Deleted ECS Be Provisioned Again?.....	405
14.2.5 Can I Forcibly Restart or Stop an ECS?.....	405
14.3 Login and Connection.....	406
14.3.1 What Should I Do If Starting an ECS Remains in "Waiting for cloudResetPwdAgent" State?.....	406
14.3.2 What Should I Do If I Cannot Use MSTSC to Log In to an ECS Running the Windows Server 2012 OS?.....	407
14.3.3 How Can I Change a Remote Login Port?.....	408
14.3.4 What Browser Version Is Required to Remotely Log In to an ECS?.....	411
14.3.5 What Should I Do If the Page Does not Respond After I Log In to an ECS Using VNC and Do Not Perform Any Operation for a Long Period of Time?.....	412
14.3.6 What Should I Do If I Cannot View Data After Logging In to an ECS Using VNC?.....	412
14.3.7 Why Does a Blank Screen Appear After I Attempted to Log In to an ECS Using VNC?.....	412
14.3.8 Why Does an Authentication Failure Occurs After I Attempt to Remotely Log In to a Windows ECS?.....	412
14.3.9 Why Can't I Use the Local Computer to Connect to My Windows ECS?.....	414
14.3.10 How Can I Obtain the Permission to Remotely Log In to a Windows ECS?.....	419
14.3.11 Why Does the System Display No Remote Desktop License Servers Available to Provide a License When I Log In to a Windows ECS?.....	421
14.3.12 Why Does the System Display Error Code 0x112f When I Log In to a Windows ECS?.....	423
14.3.13 Why Does the System Display Error Code 0x1104 When I Log In to a Windows ECS?.....	424
14.3.14 Why Does the System Display Error Code 122.112... When I Log In to a Windows ECS?.....	428
14.3.15 Why Does the System Display Invalid Certificate or Associated Chain When I Log In to a Windows ECS from a Mac?.....	430
14.3.16 Why Is My Remote Session Interrupted by a Protocol Error?.....	434
14.3.17 Why Am I Seeing an Error Message That Says Identity of Remote Computer Cannot be Verified When I Log In to a Windows ECS?.....	436
14.3.18 Why Am I Seeing An Error Message That Says The Two Computers Couldn't Be Connected in the Amount of Time Allotted When I Log In to a Windows ECS?.....	437
14.3.19 Why Am I Seeing an Error Message That Says User Account is not Authorized for Remote Login When I Log In to a Windows ECS?.....	437
14.3.20 Why Does My Remote Desktop Session End Because Another User Logs In When I Log In to a Windows ECS?.....	441
14.3.21 Why Does an ECS Fail to Be Remotely Connected Using RDP and Internal Error Code 4 Is Displayed?.....	444
14.3.22 Why Am I Seeing the Error Message "Module is unknown" When I Remotely Log In to a Linux ECS?.....	445
14.3.23 What Should I Do If Error Message "Permission denied" Is Displayed When I Remotely Log In to a Linux ECS?.....	447
14.3.24 What Should I Do If Error Message "read: Connection reset by peer" Is Displayed When I Remotely Log In to a Linux ECS?.....	449
14.3.25 Why Am I Seeing the Error Message "Access denied" When I Remotely Log In to a Linux ECS?.....	450
14.3.26 What Should I Do If Error Message "Disconnected: No supported authentication methods available" Is Displayed When I Remotely Log In to a Linux ECS?.....	451
14.4 How Do I Handle Error Messages Displayed on the Management Console?.....	451
14.5 Why Does the System Display a Question Mark When I Attempt to Obtain Console Logs?.....	454

14.6 ECS Management.....	454
14.6.1 How Can a Changed Static Hostname Take Effect Permanently?.....	454
14.6.2 Is an ECS Hostname with Suffix .novalocal Normal?.....	457
14.6.3 Why Do the Disks of a Windows ECS Go Offline After I Modify the ECS Specifications?.....	458
14.6.4 Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?.....	461
14.7 OS Management.....	462
14.7.1 Can I Install or Upgrade the OS of an ECS?.....	462
14.7.2 Can I Change the OS of an ECS?.....	462
14.7.3 How Long Does It Take to Change an ECS OS?.....	463
14.7.4 Can I Select Another OS During ECS OS Reinstallation?.....	463
14.7.5 How Long Does It Take to Reinstall an ECS OS?.....	463
14.7.6 How Do I View the GPU Usage of a GPU-accelerated ECS?.....	463
14.8 File Upload/Data Transfer.....	465
14.8.1 How Do I Upload Files to My ECS?.....	465
14.8.2 How Can I Transfer Files from a Local Windows Computer to a Windows ECS?.....	466
14.8.3 How Can I Use WinSCP to Transfer Files from a Local Windows Computer to a Linux ECS?.....	469
14.8.4 How Can I Transfer Files from a Local Mac to a Windows ECS?.....	470
14.8.5 How Can I Use SCP to Transfer Files Between a Local Linux Computer and a Linux ECS?.....	473
14.8.6 How Can I Use SFTP to Transfer Files Between a Local Linux Computer and a Linux ECS?.....	474
14.8.7 How Can I Use FTP to Transfer Files from a Local Windows Computer to a Windows or Linux ECS?	476
14.8.8 How Can I Use FTP to Transfer Files Between a Local Linux Computer and a Linux ECS?.....	477
14.8.9 What Should I Do If the Connection Between the Client and the Server Times Out When I Upload a File Using FTP?.....	478
14.8.10 What Should I Do If Writing Data Failed When I Upload a File Using FTP?.....	478
14.8.11 Why Am I Seeing an FTP Folder Error When I Open a Folder on an FTP Server?.....	480
14.8.12 Why Do I Fail to Connect to a Linux ECS Using WinSCP?.....	481
14.9 ECS Migration.....	482
14.9.1 Can I Migrate an ECS to Another Region, AZ, or Account?.....	483
14.10 Disk Management.....	483
14.10.1 Why Can't I Find My Newly Purchased Data Disk After I Log In to My Windows ECS?.....	483
14.10.2 How Can I Adjust System Disk Partitions?.....	484
14.10.3 How Can I Obtain the Mapping Between Disk Partitions and Disk Devices on a Windows ECS?.....	490
14.10.4 How Can I Obtain the Mapping Between Disk Partitions and Disk Devices on a Linux ECS?.....	493
14.10.5 How Can I Enable Virtual Memory on a Windows ECS?.....	495
14.10.6 How Can I Add the Empty Partition of an Expanded System Disk to the End Root Partition Online?.....	497
14.10.7 How Can I Add the Empty Partition of an Expanded System Disk to the Non-end Root Partition Online?.....	498
14.10.8 Can I Attach Multiple Disks to an ECS?.....	501
14.10.9 What Are the Requirements for Attaching an EVS Disk to an ECS?.....	502
14.10.10 Which ECSs Can Be Attached with SCSI EVS Disks?.....	503
14.10.11 How Do I Obtain My Disk Device Name in the ECS OS Using the Device Identifier Provided on the Console?.....	503

14.10.12 Why Does a Linux ECS with a SCSI Disk Attached Fails to Be Restarted?.....	507
14.10.13 Why Does a Disk Attached to a Windows ECS Go Offline?.....	508
14.10.14 Why Does the Disk Drive Letter Change After the ECS Is Restarted?.....	509
14.10.15 How Can I Obtain Data Disk Information If Tools Are Uninstalled?.....	510
14.11 Passwords and Key Pairs.....	511
14.11.1 How Can I Set the Validity Period of the Image Password?.....	511
14.11.2 Why Does Login to My ECS Using the Reset Password Fail?.....	512
14.11.3 Why Am I Seeing the Message Indicating That the Port Is Used by a One-Click Password Reset Plug-in?.....	515
14.11.4 Why Does the One-Click Password Reset Plug-in Use Too Much VIRT and SHR?.....	518
14.11.5 How Can I Obtain the Key Pair Used by My ECS?.....	519
14.11.6 What Should I Do If a Key Pair Cannot Be Imported?.....	520
14.11.7 Why Does the Login to My Linux ECS Using a Key File Fail?.....	520
14.11.8 Why Does a Key Pair Created Using puttygen.exe Fail to Be Imported on the Management Console?.....	521
14.11.9 What Is the Cloudbase-Init Account in Windows ECSs Used for?.....	523
14.11.10 What Should I Do If Cloud-Init Does Not Work After Python Is Upgraded?.....	524
14.12 Network Configurations.....	525
14.12.1 Can Multiple EIPs Be Bound to an ECS?.....	525
14.12.2 Can an ECS Without an EIP Bound Access the Internet?.....	526
14.12.3 What Should I Do If an EIP Cannot Be Pinged?.....	527
14.12.4 Why Can I Remotely Access an ECS But Cannot Ping It?.....	532
14.12.5 Will NICs Added to an ECS Start Automatically?.....	532
14.12.6 How Can I Check Whether the Network Communication Is Normal Between Two ECSs Equipped with an InfiniBand NIC Driver?.....	532
14.12.7 How Can I Manually Configure an IP Address for an InfiniBand NIC?.....	533
14.12.8 How Can I Handle the Issue that a Windows 7 ECS Equipped with an Intel 82599 NIC Reports an Error in SR-IOV Scenarios?.....	534
14.12.9 Why Can't My Linux ECS Obtain Metadata?.....	535
14.12.10 How Can I Obtain the MAC Address of My ECS?.....	536
14.12.11 How Can I Test Network Performance?.....	538
14.12.12 Why Can't I Use DHCP to Obtain a Private IP Address?.....	547
14.12.13 How Can I View and Modify Kernel Parameters of a Linux ECS?.....	549
14.12.14 How Can I Configure Port Redirection?.....	554
14.12.15 Can the ECSs of Different Accounts Communicate over an Intranet?.....	556
14.12.16 Will ECSs That I Purchased Deployed in the Same Subnet?.....	556
14.13 Resource Monitoring.....	556
14.13.1 Why Is My Windows ECS Running Slowly?.....	556
14.13.2 Why Is My Linux ECS Running Slowly?.....	560
14.14 Database Applications.....	564
14.14.1 Can a Database Be Deployed on an ECS?.....	564
14.14.2 Does an ECS Support Oracle Databases?.....	564

A Change History.....565

1 Service Overview

1.1 What Is ECS?

An Elastic Cloud Server (ECS) is a basic computing unit that consists of vCPUs, memory, OS, and Elastic Volume Service (EVS) disks.

You can create an ECS by specifying its vCPUs, memory, OS, and login mode. After creating an ECS, you can use it on the cloud like using your local PC or physical server. You can also modify its specifications if necessary. ECS lets your applications run in a reliable, secure, efficient computing environment.


- For details about the operating systems supported by an ECS, see [Image Types](#).
- For details about the login authentication modes, see "Logging In to an ECS" in the *Elastic Cloud Server User Guide*.

System Architecture

ECS works with other products and services to provide computing, storage, and network resources.

- You can deploy ECSs across different availability zones (AZs) that are connected over an intranet. If one AZ becomes unavailable, ECSs in other AZs can continue to provide services.
- Virtual Private Cloud (VPC) helps you build your own dedicated network on the cloud. You can set subnets and security groups within your VPC for further isolation. You can also bind an EIP to your ECSs for Internet access.
- With the Image Management Service (IMS), you can use an image to create ECSs. You can also use an existing ECS to create a private image and use the private image to create the same ECSs for rapid service deployment.
- Elastic Volume Service (EVS) provides storage space. Volume Backup Service (VBS) provides data backup and restoration.
- Cloud Eye lets you keep a close eye on the performance and resource utilization of ECSs, ensuring ECS reliability and availability.
- Cloud Backup and Recovery (CBR) allows you to create backups for EVS disks and ECSs and use the backups to restore them.

Figure 1-1 System architecture



Access Methods

You can access ECS through the web-based management console or HTTPS-based application programming interfaces (APIs).

- Accessing ECSs through APIs
Use this method if you intend to integrate ECSs into a third-party system for secondary development. For details, see *Elastic Cloud Server API Reference*.
- Accessing ECSs through the management console
Use this method if you are not required to integrate ECSs with a third-party system.

1.2 ECS Advantages

ECS supports automated scaling of compute resources based on traffic changes and predefined scaling policies. You can customize ECS specifications including vCPUs, memory, and bandwidth to let your applications run in a flexible, efficient environment.

Reliability

- A broad range of EVS disk types
EVS disk types are classified based on I/O performance. Select EVS disks based on service requirements.

For more information about EVS disk specifications and performance, see *Elastic Volume Service User Guide*.

- Distributed architecture

ECS provides scalable, reliable, and high-throughput virtual block storage on a distributed architecture. This ensures that data can be rapidly migrated and restored if any data replica is unavailable, preventing data loss caused by a single hardware fault.

- Backup and restoration

You can set automatic backup policies to back up in-service ECSs and EVS disks. You can also configure policies on the management console or use an API to back up the data of ECSs and EVS disks at a specified time.

Security

- Multi-dimensional protection

A number of security services, such as Web Application Firewall (WAF) and Vulnerability Scan Service (VSS) are available.

- Security evaluation

Cloud security evaluation and security configuration check help you identify security vulnerabilities and threats, reducing or eliminating your loss from viruses or attacks.

- Intelligent process management

You can customize an allowlist to automatically prohibit the execution of unauthorized programs.

- Vulnerability scan

Comprehensive scan services are available, including general web vulnerability scan, third-party application vulnerability scan, port detection, and fingerprint identification.

Hardware and Software

- Professional hardware devices

You can deploy ECSs on professional hardware devices that allow in-depth virtualization optimization, delivering superior virtual server performance.

- Virtual resources accessible anytime, anywhere

You can obtain scalable, dedicated resources from the virtual resource pool anytime, anywhere, so your applications can run in reliable, secure, flexible, and efficient environments. You can use your ECS like the way you are using your local computer.

Scalability

- Automated scaling of computing resources

Dynamic scaling: AS automatically increases or decreases the number of ECSs in an AS group based on monitored data.

Periodic/Scheduled scaling: AS increases or decreases the number of ECSs in an AS group at a regular interval or a specified time based on the predicted load or a pre-set plan.

- Flexible adjustment of ECS specifications
ECS specifications and bandwidth can be flexibly adjusted based on service requirements.

1.3 ECS Application Scenarios

Internet

- No special requirements on CPUs, memory, disk space, or bandwidth
- High security and reliability standards
- Deploying an application on one or only a few ECSs to minimize upfront investment and maintenance costs, such as website development and testing, and small databases

Use general computing ECSs, which provide a balance of computing, memory, and network resources. This ECS type is appropriate for medium-load applications and meets the cloud service needs of both enterprises and individuals.

E-Commerce

- Large amount of memory
- Quick processing of large volumes of data
- Large incoming traffic

Use memory-optimized ECSs, which provide a large memory, ultra-high I/O EVS disks, and the needed bandwidths. This ECS type is suitable for precision marketing, E-Commerce, and mobile apps.

Graphics Rendering

- High-quality graphics and video
- Large amount of memory and rapid processing of large volumes of data
- Fast network with high I/O
- High GPU performance for graphics rendering and engineering drawing

Use GPU-accelerated ECSs, which adopt NVIDIA Tesla M60 hardware virtualization and provide cost-effective graphics acceleration. These ECSs support DirectX and OpenGL, and provide up to 1 GiB of GPU memory and 4096 x 2160 resolution.

Data Analytics

- Capable of processing large volumes of data
- High I/O performance and rapid data switching and processing, such as MapReduce and Hadoop

Use disk-intensive ECSs, which are designed for applications requiring sequential read/write on ultra-large datasets in local storage (such as distributed Hadoop computing) as well as large-scale parallel data processing and log processing. Disk-intensive ECSs use hard disk drives (HDDs) and a default network bandwidth of 10GE, providing high packets per second (PPS) and low network latency. Each disk-intensive ECS supports up to 24 local disks, 48 vCPUs, and 384 GiB of memory.

High-Performance Computing

High computing performance and throughput, such as scientific computing, genetic engineering, games and animation, biopharmaceuticals, and storage systems

Use high-performance computing ECSs for tasks that require large amounts of resources for parallel computing.

1.4 Notes and Constraints on Using ECSs

Notes

- Do not use ECSs as unauthorized servers for any illegal or violation service, such as gambling or cross-border VPN.
- Do not use ECSs for fraudulent transactions, such as click farming on e-commerce websites.
- Do not use ECSs to initiate network attacks, such as DDoS attacks, CC attacks, web attacks, brute force cracking, or to spread viruses and Trojan horses.
- Do not use ECSs for traffic transit.
- Do not use ECSs for web crawling.
- Do not use ECSs to detect other systems like scanning or penetration unless otherwise being authorized.
- Do not deploy any illegal websites or applications on ECSs.
- Do not use ECSs to send spams.

Restricted Operations on ECSs

- Do not uninstall drivers on the ECS hardware.
- Do not install external hardware devices, such as encryption dongles, USB flash drives, external hard disks, or bank USB security keys on ECSs.
- Do not change the MAC address of NICs.
- Do not install virtualization software on ECSs for nested virtualization.
- Do not associate software licenses with the physical server hosting an ECS. Once an ECS is migrated from one physical server to another, the associated licenses may become invalid.
- Do not deploy applications on a single ECS if you require high availability. Set up auto start for your ECSs or deploy applications in cluster or active/standby mode.
- Data on ECSs running core applications needs to be backed up.
- Monitoring needs to be configured for ECSs.
- Do not change the default DNS server address. If you need to configure a public DNS address, configure both a public and a private DNS address on your ECS.
- The system disk can boot from Basic Input Output System (BIOS) or Unified Extensible Firmware Interface (UEFI) according to the boot mode in the image file.

- You can change the OS to convert the boot mode of the ECS.
- You can create a UEFI or BIOS private image and use it to create an ECS.

Precautions for Using Windows ECSs

- Do not stop system processes if you are not sure about the consequences. Otherwise, blue screen of death (BSOD) may occur on the ECS, or the ECS may restart.
- Ensure that there is at least 2 GB of idle memory. Otherwise, BSOD, frame freezing, or service failure may occur.
- Do not modify the registry. Otherwise, the system may fail to start. If the modification is mandatory, back up the registry before modifying it.
- Do not modify ECS clock settings. Otherwise, DHCP lease may fail, leading to the loss of IP addresses.
- Do not delete the CloudResetPwdAgent or CloudResetPwdUpdateAgent process. Otherwise, one-click password reset will become unavailable.
- Do not disable virtual memory. Otherwise, system performance may deteriorate, or system exceptions may occur.
- Do not delete the VMTool program. Otherwise, the ECS may be abnormal.

Precautions for Using Linux ECSs

- Do not modify the `/etc/issue` file. Otherwise, the OS distribution will not be identified.
- Do not delete system directories or files. Otherwise, the system may fail to start or run properly.
- Do not change the permissions for or names of system directories. Otherwise, the system may fail to start or run properly.
- Do not upgrade the kernel of the Linux unless necessary.
- Do not delete the CloudResetPwdAgent or CloudResetPwdUpdateAgent process. Otherwise, one-click password reset will be unavailable.
- Do not change the default `/etc/resolv.conf` of the DNS server. Otherwise, software sources and NTP may be unavailable.
- Do not modify default intranet configurations, such as IP addresses, subnet mask, and gateway address of an ECS. Otherwise, network exceptions may occur.
- Manually specified IP addresses for Linux ECSs are generally static IP addresses. To avoid network exceptions caused by conflicts between NetworkManager and internal network services, do not enable NetworkManager when not required, such as when installing Kubernetes.

1.5 ECS and Other Services

ECS-related Services

- Auto Scaling (AS)
Automatically adjusts ECS resources based on the configured AS policies. This improves resource usage and reduces resource costs.

- (ELB)
Automatically distributes traffic to multiple ECSs. This enhances system service and fault tolerance capabilities.
- Elastic Volume Service (EVS)
Enables you to attach EVS disks to an ECS and expand their capacity.
- Virtual Private Cloud (VPC)
Enables you to configure internal networks and change network configurations by customizing security groups, VPNs, IP address ranges, and bandwidth. This simplifies network management. You can also customize the ECS access rules within a security group and between security groups to improve ECS security.
- Image Management Service (IMS)
Enables you to create ECSs using images. This improves the efficiency of ECS creation.
- Cloud Eye
Allows you to check the status of monitored service objects after you have obtained an ECS. This can be done without requiring additional plug-ins be installed.
- Cloud Backup and Recovery (CBR)
Backs up EVS disks and ECSs for restoration. You can back up all EVS disks (including the system disk and data disks) attached to an ECS and use the backup to restore the ECS data.

1.6 Instances

1.6.1 ECS Overview

An ECS is a basic computing unit that consists of vCPUs, memory, OS, and EVS disks.

After creating an ECS, you can use it like using your local computer or physical server, ensuring secure, reliable, and efficient computing. ECSs support self-service creation, modification, and operation. You can create an ECS by specifying its vCPUs, memory, OS, and login authentication. After the ECS is created, you can modify its specifications as required. This ensures a reliable, secure, efficient computing environment.

The cloud platform provides multiple ECS types for different computing and storage capabilities. One ECS type provides various flavors with different vCPU and memory configurations for you to select.

- For details about ECS types, see [ECS Types](#).
- For details about all ECS statuses in a lifecycle, see [ECS Lifecycle](#).

1.6.2 ECS Lifecycle

The ECS lifecycle refers to the entire journey an ECS goes through, from creation to deletion (or release).

Table 1-1 ECS statuses

Status	Status Attribute	Description
Creating	Intermediate	The ECS is being created.
Starting	Intermediate	The ECS is being started.
Running	Stable	The ECS is running properly.
Stopping	Intermediate	The ECS is being stopped.
Stopped	Stable	The ECS has been stopped.
Restarting	Intermediate	The ECS is being restarted.
Resizing	Intermediate	The ECS has received a resizing request and has started to resize.
Verifying resizing	Intermediate	The ECS is verifying the new size.
Deleting	Intermediate	The ECS is being deleted. If the ECS remains in this state for a long time, exceptions may have occurred. In such a case, contact technical support.
Deleted	Intermediate	The ECS has been deleted. An ECS in this state cannot provide services and will be promptly cleared from the system.
Faulty	Stable	An exception has occurred on the ECS. Contact technical support for assistance.
Reinstalling OS	Intermediate	The ECS has received a request to reinstall the OS and has begun the reinstallation.
Reinstalling OS failed	Stable	The ECS received a request to reinstall the OS, but the reinstallation failed. Contact technical support for assistance.
Changing OS	Intermediate	The ECS received a request to change the OS and has begun implementing the changes.
OS change failed	Stable	The ECS has received a request to change the OS, but due to exceptions, the change attempt failed. Contact technical support for assistance.
Forcibly restarting	Intermediate	The ECS is being forcibly restarted.
Rolling back resizing	Intermediate	The ECS is rolling back a resizing operation.

1.6.3 ECS Types

The cloud platform provides the following ECS types for different application scenarios:

- Kunpeng architecture
 - Kunpeng general computing
 - Kunpeng general computing-plus
 - Kunpeng memory-optimized
- x86 architecture
 - General-purpose
 - General computing-plus
 - Memory-optimized
 - Disk-intensive
 - Ultra-high I/O
 - High-performance computing
 - GPU-accelerated
 - AI-accelerated

ECS Flavor Naming Rules

ECS flavors are named in the "AB.C" format.

Example: *s2.medium*

The format is defined as follows:

- **A** specifies the ECS type. For example, **s** indicates a general-purpose ECS, **c** a general computing-plus ECS, and **m** a memory-optimized ECS. Kunpeng flavor names are started with letter **k**. For example, **kc** indicates Kunpeng general computing-plus.
- **B** specifies the type ID. For example, **2** in **s2** indicates a computing II ECS.
- **C** specifies the flavor size, such as medium, large, or xlarge.

vCPU

ECS supports hyper-threading, which enables two threads to run concurrently on a single CPU core. Each thread is represented as a virtual CPU (vCPU) and a CPU core contains two vCPUs (logical cores).

Hyper-threading is enabled for most ECS flavors by default. If hyper-threading is disabled during the ECS creation or flavor change, the number of vCPUs queried from the ECS is half of the number of vCPUs defined by the ECS flavor.

For example, a 2-core physical CPU contains 4 vCPUs (threads).

Network Bandwidth

The intranet bandwidth and packets per second (PPS) of an ECS are determined by the ECS flavor.

- Assured intranet bandwidth: indicates the guaranteed bandwidth allocated to an ECS when there is a network bandwidth contention in the entire network.
- Maximum intranet bandwidth: indicates the maximum bandwidth that can be allocated to an ECS when the ECS does not compete for network bandwidth (other ECSs on the host do not have high requirements on network bandwidth).
- Maximum intranet PPS: indicates the maximum ECS capability in sending and receiving packets.

NOTE

- For details about how to test packets per second (PPS), see "How Can I Test Network Performance?" in *Elastic Cloud Server User Guide*.
- For how to enable NIC multi-queue, see "Enabling NIC Multi-Queue" in *Elastic Cloud Server User Guide*.
- If the cluster uses 10GE networking, the ECS performance cannot be ensured in extreme scenarios.
- The maximum bandwidth is the total bandwidth allocated to an ECS. If an ECS has multiple NICs, the sum of the maximum bandwidths allocated to all NICs cannot exceed the maximum bandwidth allocated to the ECS.

1.7 x86 ECS Specifications and Types

1.7.1 General-Purpose ECSs

Overview

General-purpose ECSs provide a balance of compute, memory, and networking resources and a baseline level of vCPU performance with the ability to burst above the baseline. These ECSs are suitable for applications with general workloads, such as web servers, enterprise R&D, and small-scale databases.

S7n ECSs use the 3rd generation Intel® Xeon® Scalable processors and 25GE high-speed intelligent NICs to provide high network bandwidth and PPS.

S6 and S6nl ECSs are suitable for applications that require moderate performance generally but burstable high performance occasionally, such as light-workload web servers, development and testing environments, and low- and medium-performance databases.

GS7 ECSs are deployed on a new virtualization platform and have multiple technologies optimized. They use Chinese-developed HYGON C86 7380 CPUs and 25GE high-speed intelligent NICs to provide high network bandwidth, high PPS, and high cost-effectiveness.

GS6 ECSs are deployed on a new virtualization platform and have multiple technologies optimized. They use Chinese-developed HYGON C86 7285 CPUs and 25GE high-speed intelligent NICs to provide high network bandwidth, high PPS, and high cost-effectiveness.

Scenarios

- Applications
General-purpose ECSs are suitable for applications that have no special requirements on CPU performance, memory, disk capacity, or bandwidth, but have high requirements on security and reliability. Example applications include web servers, R&D and testing environments, and small-scale databases. They feature low initial investment and maintenance costs.
GS7 and GS6 ECSs are suitable for lightweight databases, cache servers, and medium- and light-load enterprise applications.
- Application scenarios
Enterprise website deployment, enterprise office environment setup, and enterprise R&D and testing activities
GS7 and GS6 ECSs are suitable for government, enterprise, and finance scenarios.

Specifications

Table 1-2 S7n ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	Virtualization
s7n.small.1	1	1	0.8/0.1	10	1	2	4	KVM
s7n.medium.2	1	2	0.8/0.1	10	1	2	4	KVM
s7n.large.2	2	4	1.5/0.2	15	1	2	8	KVM
s7n.xlarge.2	4	8	2/0.35	25	1	2	16	KVM
s7n.2xlarge.2	8	16	3/0.75	50	2	2	32	KVM
s7n.medium.4	1	4	0.8/0.1	10	1	2	4	KVM
s7n.large.4	2	8	1.5/0.2	15	1	2	8	KVM
s7n.xlarge.4	4	16	2/0.35	25	1	2	16	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	Virtualization
s7n.2xlarge.4	8	32	3/0.75	50	2	2	32	KVM

Table 1-3 S6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
s6.small.1	1	1	0.8/0.1	10	1	2	KVM
s6.xlarge.1	4	4	2/0.35	25	1	2	KVM
s6.2xlarge.1	8	8	3/0.75	50	2	2	KVM
s6.medium.2	1	2	0.8/0.1	10	1	2	KVM
s6.large.2	2	4	1.5/0.2	15	1	2	KVM
s6.xlarge.2	4	8	2/0.35	25	1	2	KVM
s6.2xlarge.2	8	16	3/0.75	50	2	2	KVM
s6.4xlarge.2	16	32	6/1.5	100	4	2	KVM
s6.medium.4	1	4	0.8/0.1	10	1	2	KVM
s6.large.4	2	8	1.5/0.2	15	1	2	KVM
s6.xlarge.4	4	16	2/0.35	25	1	2	KVM

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
s6.2xlarge.4	8	32	3/0.75	50	2	2	KVM
s6.4xlarge.4	16	64	6/1.5	100	4	2	KVM

Table 1-4 S6nl ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
s6nl.small.1	1	1	0.8/0.15	10	1	2	KVM
s6nl.medium.2	1	2	0.8/0.15	10	1	2	KVM
s6nl.large.2	2	4	1.5/0.3	15	1	2	KVM
s6nl.xlarge.2	4	8	2/0.5	25	1	2	KVM
s6nl.2xlarge.2	8	16	3/1	50	2	2	KVM
s6nl.4xlarge.2	16	32	6/1.5	100	4	2	KVM
s6nl.medium.4	1	4	0.8/0.15	10	1	2	KVM
s6nl.large.4	2	8	1.5/0.3	15	1	2	KVM
s6nl.xlarge.4	4	16	2/0.5	25	1	2	KVM
s6nl.2xlarge.4	8	32	3/1	50	2	2	KVM
s6nl.4xlarge.4	16	64	6/1.5	100	4	2	KVM

Table 1-5 GS7 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
gs7.small.1	1	1	0.8/0.1	10	1	2	KVM
gs7.medium.2	1	2	0.8/0.1	10	1	2	KVM
gs7.large.2	2	4	1/0.2	15	1	2	KVM
gs7.xlarge.2	4	8	1.5/0.3	20	1	2	KVM
gs7.2xlarge.2	8	16	2/0.5	35	2	2	KVM
gs7.medium.4	1	4	0.8/0.1	10	1	2	KVM
gs7.large.4	2	8	1/0.2	15	1	2	KVM
gs7.xlarge.4	4	16	1.5/0.3	20	1	2	KVM
gs7.2xlarge.4	8	32	2/0.5	35	2	2	KVM

Table 1-6 GS6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
gs6.small.1	1	1	0.8/0.1	10	1	2	KVM
gs6.medium.2	1	2	0.8/0.1	10	1	2	KVM
gs6.large.2	2	4	1/0.2	15	1	2	KVM
gs6.xlarge.2	4	8	1.5/0.3	20	1	2	KVM

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
gs6.2xlarge.2	8	16	2/0.5	35	2	2	KVM
gs6.medium.4	1	4	0.8/0.1	10	1	2	KVM
gs6.large.4	2	8	1/0.2	15	1	2	KVM
gs6.xlarge.4	4	16	1.5/0.3	20	1	2	KVM
gs6.2xlarge.4	8	32	2/0.5	35	2	2	KVM

1.7.2 General Computing-plus ECSs

Overview

General computing-plus ECSs use dedicated vCPUs to deliver powerful performance. In addition, the ECSs use the latest-generation network acceleration engines and Data Plane Development Kit (DPDK) to provide high network performance.

- C7n ECSs use the third-generation Intel Xeon scalable processors to provide enhanced computing, security, and stability. Each C7n ECS can have a maximum number of 96 vCPUs and a memory speed of 3200 MHz, and provide a secure, trusted cloud environment.
- C6s ECSs use the second-generation Intel® Xeon® Scalable processors that feature high performance, stability, low latency, and cost-effectiveness. They are suitable for Internet, gaming, and rendering scenarios, especially those that require high computing and network stability.
- C6 and C6nl ECSs use the second-generation Intel® Xeon® Scalable processors to provide powerful and stable computing performance. By using 25GE high-speed intelligent NICs, C6 and C6nl ECSs offer ultra-high network bandwidth and PPS.
- GC7 ECSs are deployed on a new virtualization platform and have multiple technologies optimized. They use Chinese-developed HYGON C86 7380 CPUs to provide stable, robust computing performance and use 25GE high-speed intelligent NICs to provide high network bandwidth and high packets per second (PPS).
- GC6 ECSs are deployed on a new virtualization platform and have multiple technologies optimized. They use Chinese-developed HYGON C86 7285 CPUs to provide stable, robust computing performance and use 25GE high-speed intelligent NICs to provide high network bandwidth and high PPS.

Specifications

Table 1-7 C7n ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	Virtualization
c7n.large.2	2	4	4/0.8	40	2	2	16	KVM
c7n.xlarge.2	4	8	8/1.6	80	2	3	32	KVM
c7n.2xlarge.2	8	16	15/3	150	4	4	64	KVM
c7n.3xlarge.2	12	24	17/5	200	4	6	96	KVM
c7n.4xlarge.2	16	32	20/6	280	8	8	128	KVM
c7n.6xlarge.2	24	48	25/9	400	8	8	192	KVM
c7n.8xlarge.2	32	64	30/12	550	16	8	256	KVM
c7n.12xlarge.2	48	96	35/18	750	16	8	256	KVM
c7n.16xlarge.2	64	128	36/24	800	28	8	256	KVM
c7n.24xlarge.2	96	192	40/36	850	32	8	256	KVM
c7n.large.4	2	8	4/0.8	40	2	2	16	KVM
c7n.xlarge.4	4	16	8/1.6	80	2	3	32	KVM
c7n.2xlarge.4	8	32	15/3	150	4	4	64	KVM
c7n.3xlarge.4	12	48	17/5	200	4	6	96	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	Virtualization
c7n.4xlarge.4	16	64	20/6	280	8	8	128	KVM
c7n.6xlarge.4	24	96	25/9	400	8	8	192	KVM
c7n.8xlarge.4	32	128	30/12	550	16	8	256	KVM
c7n.12xlarge.4	48	192	35/18	750	16	8	256	KVM
c7n.16xlarge.4	64	256	36/24	800	28	8	256	KVM
c7n.24xlarge.4	96	384	40/36	850	32	8	256	KVM

Table 1-8 C6s ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
c6s.large.2	2	4	1/1	30	2	2	KVM
c6s.xlarge.2	4	8	2/2	60	2	3	KVM
c6s.2xlarge.2	8	16	4/4	120	4	4	KVM
c6s.3xlarge.2	12	24	5.5/5.5	180	4	6	KVM
c6s.4xlarge.2	16	32	7.5/7.5	240	8	8	KVM
c6s.6xlarge.2	24	48	11/11	350	8	8	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
c6s.8xlarge.2	32	64	15/15	450	16	8	KVM
c6s.12xlarge.2	48	96	22/22	650	16	8	KVM
c6s.16xlarge.2	64	128	30/30	850	32	8	KVM
c6s.large.4	2	8	1/1	30	2	2	KVM
c6s.xlarge.4	4	16	2/2	60	2	3	KVM
c6s.2xlarge.4	8	32	4/4	120	4	4	KVM
c6s.3xlarge.4	12	48	5.5/5.5	180	4	6	KVM
c6s.4xlarge.4	16	64	7.5/7.5	240	8	8	KVM
c6s.6xlarge.4	24	96	11/11	350	8	8	KVM
c6s.8xlarge.4	32	128	15/15	450	16	8	KVM
c6s.12xlarge.4	48	192	22/22	650	16	8	KVM
c6s.16xlarge.4	64	256	30/30	850	32	8	KVM

Table 1-9 C6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
c6.large.2	2	4	4/1.2	40	2	KVM
c6.xlarge.2	4	8	8/2.4	80	2	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
c6.2xlarge.2	8	16	15/4.5	150	4	KVM
c6.3xlarge.2	12	24	17/7	200	4	KVM
c6.4xlarge.2	16	32	20/9	280	8	KVM
c6.6xlarge.2	24	48	25/14	400	8	KVM
c6.8xlarge.2	32	64	30/18	550	16	KVM
c6.12xlarge.2	48	96	35/27	750	16	KVM
c6.16xlarge.2	64	128	40/36	1,000	32	KVM
c6.large.4	2	8	4/1.2	40	2	KVM
c6.xlarge.4	4	16	8/2.4	80	2	KVM
c6.2xlarge.4	8	32	15/4.5	150	4	KVM
c6.3xlarge.4	12	48	17/7	200	4	KVM
c6.4xlarge.4	16	64	20/9	280	8	KVM
c6.6xlarge.4	24	96	25/14	400	8	KVM
c6.8xlarge.4	32	128	30/18	550	16	KVM
c6.12xlarge.4	48	192	35/27	750	16	KVM
c6.16xlarge.4	64	256	40/36	1,000	32	KVM

Table 1-10 C6nl ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
c6nl.large.2	2	4	4/1	32	2	KVM
c6nl.xlarge.2	4	8	8/2	64	2	KVM
c6nl.2xlarge.2	8	16	15/4	120	4	KVM
c6nl.3xlarge.2	12	24	17/6	160	4	KVM
c6nl.4xlarge.2	16	32	20/8	224	8	KVM
c6nl.6xlarge.2	24	48	25/12	320	8	KVM
c6nl.8xlarge.2	32	64	30/16	440	16	KVM
c6nl.16xlarge.2	64	128	40/32	800	32	KVM
c6nl.large.1	2	2	4/1	32	2	KVM
c6nl.xlarge.1	4	4	8/2	64	2	KVM
c6nl.2xlarge.1	8	8	15/4	120	4	KVM
c6nl.3xlarge.1	12	12	17/6	160	4	KVM
c6nl.4xlarge.1	16	16	20/8	224	8	KVM
c6nl.large.4	2	8	4/1	32	2	KVM
c6nl.xlarge.4	4	16	8/2	64	2	KVM
c6nl.2xlarge.4	8	32	15/4	120	4	KVM
c6nl.3xlarge.4	12	48	17/6	160	4	KVM

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
c6nl.4xlarge.4	16	64	20/8	224	8	KVM
c6nl.6xlarge.4	24	96	25/12	320	8	KVM
c6nl.8xlarge.4	32	128	30/16	440	16	KVM
c6nl.16xlarge.4	64	256	40/32	800	32	KVM

Table 1-11 GC7 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
gc7.large.2	2	4	2/0.4	30	2	KVM
gc7.xlarge.2	4	8	3/0.8	50	2	KVM
gc7.2xlarge.2	8	16	4/1.5	80	4	KVM
gc7.3xlarge.2	12	24	5/2	110	4	KVM
gc7.4xlarge.2	16	32	7/3	140	8	KVM
gc7.6xlarge.2	24	48	8/4	200	8	KVM
gc7.8xlarge.2	32	64	10/5	260	16	KVM
gc7.12xlarge.2	48	96	15/8	350	16	KVM
gc7.16xlarge.2	64	128	18/10	400	32	KVM
gc7.large.1	2	2	2/0.4	30	2	KVM

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
gc7.xlarge.1	4	4	3/0.8	50	2	KVM
gc7.2xlarge.1	8	8	4/1.5	80	4	KVM
gc7.3xlarge.1	12	12	5/2	110	4	KVM
gc7.4xlarge.1	16	16	7/3	140	8	KVM
gc7.large.4	2	8	2/0.4	30	2	KVM
gc7.xlarge.4	4	16	3/0.8	50	2	KVM
gc7.2xlarge.4	8	32	4/1.5	80	4	KVM
gc7.3xlarge.4	12	48	5/2	110	4	KVM
gc7.4xlarge.4	16	64	7/3	140	8	KVM
gc7.6xlarge.4	24	96	8/4	200	8	KVM
gc7.8xlarge.4	32	128	10/5	260	16	KVM
gc7.12xlarge.4	48	192	15/8	350	16	KVM
gc7.16xlarge.4	64	256	18/10	400	32	KVM

Table 1-12 GC6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
gc6.large.2	2	4	2/0.4	30	2	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
gc6.xlarge.2	4	8	3/0.8	50	2	KVM
gc6.2xlarge.2	8	16	4/1.5	80	4	KVM
gc6.3xlarge.2	12	24	5/2	110	4	KVM
gc6.4xlarge.2	16	32	7/3	140	8	KVM
gc6.6xlarge.2	24	48	8/4	200	8	KVM
gc6.8xlarge.2	32	64	10/5	260	16	KVM
gc6.12xlarge.2	48	96	15/8	350	16	KVM
gc6.16xlarge.2	64	128	18/10	400	32	KVM
gc6.large.1	2	2	2/0.4	30	2	KVM
gc6.xlarge.1	4	4	3/0.8	50	2	KVM
gc6.2xlarge.1	8	8	4/1.5	80	4	KVM
gc6.3xlarge.1	12	12	5/2	110	4	KVM
gc6.4xlarge.1	16	16	7/3	140	8	KVM
gc6.large.4	2	8	2/0.4	30	2	KVM
gc6.xlarge.4	4	16	3/0.8	50	2	KVM
gc6.2xlarge.4	8	32	4/1.5	80	4	KVM
gc6.3xlarge.4	12	48	5/2	110	4	KVM

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
gc6.4xlarge.4	16	64	7/3	140	8	KVM
gc6.6xlarge.4	24	96	8/4	200	8	KVM
gc6.8xlarge.4	32	128	10/5	260	16	KVM
gc6.12xlarge.4	48	192	15/8	350	16	KVM
gc6.16xlarge.4	64	256	18/10	400	32	KVM

Scenarios

- C7n
Medium- and heavy-load enterprise applications with strict requirements on computing and network performance, such as web applications, e-commerce platforms, short video platforms, online games, and insurance and finance.
- C6s
Internet, gaming, and rendering scenarios, especially those with strict requirements on computing and network stability
 - Gaming: C6s ECSs meet requirements of high performance and stability.
 - Rendering: C6s ECSs provide high-quality rendering at the most optimal cost.
 - Other scenarios: C6s ECSs can be used in gaming acceleration, video bullet screen, website building, and app development.
- C6 and C6nl
Websites and web applications, generalized databases and cache servers, and medium- and heavy-load enterprise applications with strict requirements on computing and network performance
- GC7 and GC6
Government, enterprise, and finance scenarios
Websites and web applications, generalized databases and cache servers, medium- and heavy-load enterprise applications, gaming, and rendering with strict requirements on computing and network performance

1.7.3 Memory-optimized ECSs

Overview

Memory-optimized ECSs have a large memory size and provide high memory performance. They are designed for memory-intensive applications that process

large volumes of data, such as precision marketing, e-commerce, and IoT big data analysis.

- M7n ECSs use the third-generation Intel® Xeon® Scalable processors to provide enhanced computing, security, and stability. Each M7n ECS can have a maximum number of 96 vCPUs and a memory speed of 3,200 MHz, and provide a secure and trusted cloud environment for memory-intensive computing applications.
- M6s ECSs use the second-generation Intel® Xeon® Scalable processors with technologies optimized to offer powerful and stable computing performance. Using 25GE high-speed intelligent NICs, M6s ECSs provide a maximum memory size of 512 GiB based on DDR4 for memory-intensive applications with high requirements on network bandwidth and Packets Per Second (PPS).
- M6 and M6nl ECSs use the second-generation Intel® Xeon® Scalable processors with technologies optimized to offer powerful and stable computing performance. Using 25GE high-speed intelligent NICs, M6 and nM6 ECSs provide a maximum memory size of 512 GiB based on DDR4 for memory-intensive applications with high requirements on network bandwidth and Packets Per Second (PPS).
- GM7 ECSs are deployed on a new virtualization platform, use Chinese-developed HYGON C86 7380 CPUs, and have multiple technologies optimized. They use 25GE high-speed intelligent NICs to provide high network bandwidth and high PPS, and provide up to 512 GiB of DDR4 memory for memory-intensive applications.
- GM6 ECSs are deployed on a new virtualization platform, use Chinese-developed HYGON C86 7285 CPUs, and have multiple technologies optimized. They use 25GE high-speed intelligent NICs to provide high network bandwidth and high PPS, and provide up to 512 GiB of DDR4 memory for memory-intensive applications.

Scenarios

- Applications
Memory-optimized ECSs are suitable for applications that require a large amount of memory such as relational databases, NoSQL databases, and memory data analysis
GM7 and GM6 ECSs are suitable for massively parallel processing (MPP) data warehouse, MapReduce and Hadoop distributed computing, distributed file systems, network file systems, and log or data processing applications.
- Application scenarios
Big data analysis for precision marketing, e-commerce, and IoT, relational databases, NoSQL databases, and memory data analysis
GM7 and GM6 ECSs are suitable for government, enterprise, and finance scenarios.

Specifications

Table 1-13 M7n ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	Virtualization
m7n.large.8	2	16	4/0.8	40	2	2	16	KVM
m7n.xlarge.8	4	32	8/1.6	80	2	3	32	KVM
m7n.2xlarge.8	8	64	15/3	150	4	4	64	KVM
m7n.3xlarge.8	12	96	17/5	200	4	6	96	KVM
m7n.4xlarge.8	16	128	20/6	280	8	8	128	KVM
m7n.6xlarge.8	24	192	25/9	400	8	8	192	KVM
m7n.8xlarge.8	32	256	30/12	550	16	8	256	KVM
m7n.12xlarge.8	48	384	35/18	750	16	8	256	KVM
m7n.16xlarge.8	64	512	36/24	800	28	8	256	KVM
m7n.24xlarge.8	96	768	40/36	850	32	8	256	KVM

Table 1-14 M6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
m6.large.8	2	16	4/1.2	40	2	KVM

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
m6.xlarge.8	4	32	8/2.4	80	2	KVM
m6.2xlarge.8	8	64	15/4.5	150	4	KVM
m6.3xlarge.8	12	96	17/7	200	4	KVM
m6.4xlarge.8	16	128	20/9	280	8	KVM
m6.6xlarge.8	24	192	25/14	400	8	KVM
m6.8xlarge.8	32	256	30/18	550	16	KVM
m6.16xlarge.8	64	512	40/36	1,000	32	KVM

Table 1-15 M6s ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
m6s.large.8	2	16	3/1	30	2	2	KVM
m6s.xlarge.8	4	32	6/2	60	2	3	KVM
m6s.2xlarge.8	8	64	12/4	120	4	4	KVM
m6s.3xlarge.8	12	96	14/5.5	160	4	6	KVM
m6s.4xlarge.8	16	128	16/7.5	220	8	8	KVM
m6s.6xlarge.8	24	192	20/11	320	8	8	KVM
m6s.8xlarge.8	32	256	25/15	450	16	8	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
m6s.16xlarge.8	64	512	34/30	850	32	8	KVM

Table 1-16 M6nl ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
m6nl.large.8	2	16	4/1	32	2	KVM
m6nl.xlarge.8	4	32	8/2	64	2	KVM
m6nl.2xlarge.8	8	64	15/4	120	4	KVM
m6nl.3xlarge.8	12	96	17/6	160	4	KVM
m6nl.4xlarge.8	16	128	20/8	224	8	KVM
m6nl.6xlarge.8	24	192	25/12	320	8	KVM
m6nl.8xlarge.8	32	256	30/16	440	16	KVM
m6nl.16xlarge.8	64	512	40/32	800	32	KVM

Table 1-17 GM7 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
gm7.large.8	2	16	2/0.4	30	2	KVM

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
gm7.xlarge.8	4	32	3/0.8	50	2	KVM
gm7.2xlarge.8	8	64	4/1.5	80	4	KVM
gm7.3xlarge.8	12	96	5/2	110	4	KVM
gm7.4xlarge.8	16	128	7/3	140	8	KVM
gm7.6xlarge.8	24	192	8/4	200	8	KVM
gm7.8xlarge.8	32	256	10/5	260	16	KVM
gm7.16xlarge.8	64	512	18/10	400	32	KVM

Table 1-18 GM6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
gm6.large.8	2	16	2/0.4	30	2	KVM
gm6.xlarge.8	4	32	3/0.8	50	2	KVM
gm6.2xlarge.8	8	64	4/1.5	80	4	KVM
gm6.3xlarge.8	12	96	5/2	110	4	KVM
gm6.4xlarge.8	16	128	7/3	140	8	KVM
gm6.6xlarge.8	24	192	8/4	200	8	KVM

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
gm6.8xlarge.8	32	256	10/5	260	16	KVM
gm6.16xlarge.8	64	512	18/10	400	32	KVM

1.7.4 Disk-intensive ECSs

Overview

Disk-intensive ECSs are delivered with local disks for high storage bandwidth and IOPS. In addition, local disks are more cost-effective in massive data storage scenarios. Disk-intensive ECSs have the following features:

- They use local disks to provide high sequential read/write performance and low latency, improving file read/write performance.
- They provide powerful and stable computing capabilities, ensuring efficient data processing.
- They provide high intranet performance, including high intranet bandwidth and packets per second (PPS), meeting requirements for data exchange between ECSs during peak hours.

D7 ECSs use third-generation Intel® Xeon® Scalable processors to provide enhanced computing, security, and stability. The RAM frequency is up to 3,200 MHz. Equipped with 25GE high-speed intelligent NICs and local SATA disks, D7 ECSs offer ultra-high network bandwidth, PPS, and local storage. The capacity of a single SATA disk is up to 3,600 GiB, and an ECS can be attached with up to 32 such disks.

D6 ECSs, with a vCPU/memory ratio of 1:4, use 2nd Generation Intel® Xeon® Scalable processors to offer powerful and stable computing performance. Equipped with 25GE high-speed intelligent NICs and local SATA disks, D6 ECSs offer ultra-high network bandwidth, PPS, and local storage. The capacity of a single SATA disk is up to 3600 GiB, and an ECS can be attached with up to 36 such disks.

D3 ECSs use Intel® Xeon® Scalable processors to offer powerful and stable computing performance. Equipped with proprietary 25GE high-speed intelligent NICs and local SAS disks, D3 ECSs offer ultra-high network bandwidth, PPS, and local storage.

Scenario

- Applications: Massively parallel processing (MPP) database, MapReduce and Hadoop distributed computing, and big data computing

- Features: Suitable for applications that require large volumes of data to process, high I/O performance, and rapid data switching and processing.
- Application scenarios: Distributed file systems, network file systems, and logs and data processing applications

Specifications

Table 1-19 D7 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	Local Disk (GiB)	Virtualization
d7.xlarge.4	4	16	5/1.7	60	2	3	32	2 × 3,600	KVM
d7.2xlarge.4	8	32	10/3.5	120	4	4	64	4 × 3,600	KVM
d7.4xlarge.4	16	64	20/6.7	240	4	6	96	8 × 3,600	KVM
d7.6xlarge.4	24	96	25/10	350	8	8	128	12 × 3,600	KVM
d7.8xlarge.4	32	128	30/13.5	450	8	8	192	16 × 3,600	KVM
d7.12xlarge.4	48	192	40/20	650	16	8	256	24 × 3,600	KVM
d7.16xlarge.4	64	256	42/27	850	16	8	256	32 × 3,600	KVM

Table 1-20 D6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Local Disks (GiB)	Virtualization
d6.xlarge.4	4	16	5/2	60	2	3	2 × 3,600	KVM
d6.2xlarge.4	8	32	10/4	120	4	4	4 × 3,600	KVM
d6.4xlarge.4	16	64	20/7.5	240	8	8	8 × 3,600	KVM
d6.6xlarge.4	24	96	25/11	350	8	8	12 × 3,600	KVM
d6.8xlarge.4	32	128	30/15	450	16	8	16 × 3,600	KVM
d6.12xlarge.4	48	192	40/22	650	16	8	24 × 3,600	KVM
d6.16xlarge.4	64	256	42/30	850	32	8	32 × 3,600	KVM
d6.18xlarge.4	72	288	44/34	900	32	8	36 × 3,600	KVM

Table 1-21 D3 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Local Disks (GiB)	Virtualization
d3.xlarge.8	4	32	2.5/2.5	50	2	3	2 × 1,675	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Local Disks (GiB)	Virtualization
d3.2xlarge.8	8	64	5/5	100	2	4	4 × 1,675	KVM
d3.4xlarge.8	16	128	10/10	120	4	8	8 × 1,675	KVM
d3.6xlarge.8	24	192	15/15	160	6	8	12 × 1,675	KVM
d3.8xlarge.8	32	256	20/20	200	8	8	16 × 1,675	KVM
d3.12xlarge.8	48	384	32/32	220	16	8	24 × 1,675	KVM
d3.14xlarge.10	56	560	40/40	500	16	8	28 × 1,675	KVM

1.7.5 Ultra-high I/O ECSs

Overview

Ultra-high I/O ECSs use high-performance local NVMe SSDs to provide high storage input/output operations per second (IOPS) and low read/write latency. You can create such ECSs with high-performance local NVMe SSDs attached on the management console.

Scenarios

- Ultra-high I/O ECSs are suitable for high-performance relational databases.
- Ultra-high I/O ECSs are suitable for NoSQL databases (such as Cassandra and MongoDB) and Elasticsearch.

Specifications

Table 1-22 Ir7 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	Local Disks (GiB)	Virtualization
ir7.large.4	2	8	3/0.8	40	2	3	32	2 × 50	KVM
ir7.xlarge.4	4	16	6/1.5	80	2	3	32	2 × 100	KVM
ir7.2xlarge.4	8	32	15/3.1	150	4	4	64	2 × 200	KVM
ir7.4xlarge.4	16	64	20/6.2	300	4	6	96	2 × 400	KVM
ir7.8xlarge.4	32	128	30/12	400	8	8	192	2 × 800	KVM
ir7.16xlarge.4	64	256	40/25	600	16	8	256	2 × 1,600	KVM

Table 1-23 Ir7n ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	Local Disks (GiB)	Virtualization
ir7n.large.4	2	8	3/0.9	40	2	3	32	2 × 50	KVM
ir7n.xlarge.4	4	16	6/1.8	80	2	3	32	2 × 100	KVM
ir7n.2xlarge.4	8	32	15/3.6	150	4	4	64	2 × 200	KVM
ir7n.4xlarge.4	16	64	20/7.3	300	4	6	96	2 × 400	KVM
ir7n.8xlarge.4	32	128	30/14.5	400	8	8	192	2 × 800	KVM
ir7n.16xlarge.4	64	256	40/29	600	16	8	256	2 × 1,600	KVM

Table 1-24 I7n ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	Local Disks (GiB)	Virtualization
i7n.2xlarge.4	8	32	10/3.4	120	4	4	64	1 × 1,600 GiB NVMe	KVM
i7n.4xlarge.4	16	64	15/6.7	200	4	6	96	2 × 1,600 GiB NVMe	KVM
i7n.8xlarge.4	32	128	25/13.5	400	8	8	192	4 × 1,600 GiB NVMe	KVM
i7n.12xlarge.4	48	192	30/20	500	16	8	256	6 × 1,600 GiB NVMe	KVM
i7n.16xlarge.4	64	256	35/27	600	16	8	256	8 × 1,600 GiB NVMe	KVM
i7n.24xlarge.4	96	420	44/20	800	32	8	256	12 × 1,600 GiB NVMe	KVM

Table 1-25 I3 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Local Disks	Virtualization
i3.2xlarge.8	8	64	2.5/2.5	100	4	4	1 × 1,600 GiB NVMe	KVM
i3.4xlarge.8	16	128	5/5	150	4	8	2 × 1,600 GiB NVMe	KVM
i3.8xlarge.8	32	256	10/10	200	8	8	4 × 1,600 GiB NVMe	KVM
i3.12xlarge.8	48	384	15/15	240	8	8	6 × 1,600 GiB NVMe	KVM
i3.15xlarge.8	60	512	25/25	500	16	8	7 × 1,600 GiB NVMe	KVM
i3.16xlarge.8	64	512	25/25	500	16	8	8 × 1,600 GiB NVMe	KVM

Features

Table 1-26 lists the IOPS performance of local disks attached to an Ir7 ECS.

Table 1-26 IOPS performance of local disks used by Ir7 ECSs

Flavor	Maximum IOPS for Random 4 KB Read
ir7.large.4	28,125
ir7.xlarge.4	56,250
ir7.2xlarge.4	112,500

Flavor	Maximum IOPS for Random 4 KB Read
ir7.4xlarge.4	225,000
ir7.8xlarge.4	450,000
ir7.16xlarge.4	900,000

[Table 1-27](#) lists the IOPS performance of local disks attached to an Ir7n ECS.

Table 1-27 IOPS performance of local disks used by Ir7n ECSs

Flavor	Maximum IOPS for Random 4 KB Read
ir7n.large.4	28,125
ir7n.xlarge.4	56,250
ir7n.2xlarge.4	112,500
ir7n.4xlarge.4	225,000
ir7n.8xlarge.4	450,000
ir7n.16xlarge.4	900,000

[Table 1-28](#) and [Table 1-29](#) list the IOPS performance of local disks and specifications of a single local disk attached to an I7n ECS.

Table 1-28 IOPS performance of local disks used by I7n ECSs

Flavor	Maximum IOPS for Random 4 KB Read
i7n.2xlarge.4	900,000
i7n.4xlarge.4	1,800,000
i7n.8xlarge.4	3,600,000
i7n.12xlarge.4	5,400,000
i7n.16xlarge.4	7,200,000
i7n.24xlarge.4	10,800,000

Table 1-29 Specifications of a single local disk attached to an I7n ECS

Metric	Performance
Disk capacity	1.6 TB
IOPS for random 4 KB read	900,000
IOPS for random 4 KB write	250,000
Read throughput	6.2 GiB/s
Write throughput	2.1 GiB/s
Access latency	Within microseconds

Table 1-30 and **Table 1-31** list the IOPS performance of local disks and specifications of a single local disk attached to an I3 ECS.

Table 1-30 IOPS performance of local disks used by I3 ECSs

Flavor	Maximum IOPS for Random 4 KB Read
i3.2xlarge.8	750,000
i3.4xlarge.8	1,500,000
i3.8xlarge.8	3,000,000
i3.12xlarge.8	4,500,000
i3.15xlarge.8	5,250,000
i3.16xlarge.8	6,000,000

Table 1-31 Specifications of a single I3 local disk

Metric	Performance
Disk capacity	1.6 TB
IOPS for random 4 KB read	750,000
IOPS for random 4 KB write	200,000
Read throughput	2.9 GiB/s
Write throughput	1.9 GiB/s
Access latency	Within microseconds

Notes

- Ultra-high I/O ECSs support the following OSs:
 - EulerOS 2.2
 - CentOS 7.2
 - CentOS 7.3
 - Ubuntu Server 16.04
 - SUSE Linux Enterprise Server 12 SP2
 - Fedora 25 64bit
 - OpenSUSE 42.2 64bit

NOTE

EulerOS 2.2 and Ubuntu Server 16.04 are recommended.

- If the host where an ultra-high I/O ECS is deployed is faulty, the ECS cannot be restored through live migration.
 - If the host is faulty or subhealthy, you need to stop the ECS for hardware repair.
 - In case of system maintenance or hardware faults, the ECS will be redeployed (to ensure HA) and cold migrated to another host. The local disk data of the ECS will not be retained.
- Ultra-high I/O ECSs do not support specifications change.
- Ultra-high I/O ECSs do not support local disk snapshots or backups.
- Ultra-high I/O ECSs can use local disks, and can also have EVS disks attached to provide a larger storage size. Note the following when using the two types of storage media:
 - Only an EVS disk, not a local disk, can be used as the system disk of an ultra-high I/O ECS.
 - Both EVS disks and local disks can be used as data disks of an ultra-high I/O ECS.
 - An ultra-high I/O ECS can have a maximum of 60 attached disks (including VBD, SCSI, and local disks).
- Modify the **fstab** file to set automatic disk mounting at ECS start. For details, see "Configuring Automatic Mounting at System Start" in the *Elastic Cloud Server User Guide*.
- The local disk data of an ultra-high I/O ECS if an exception occurs, such as physical server breakdown or local disk damage. If your application does not use the data reliability architecture, it is a good practice to use EVS disks to build your ECS.
- When an ultra-high I/O ECS is deleted, the data on local NVMe SSDs will also be automatically deleted, which can take some time. As a result, an ultra-high I/O ECS takes a longer time than other ECSs to be deleted. Back up the data before deleting such an ECS.
- The data reliability of local disks depends on the reliability of physical servers and hard disks, which are SPOF-prone. It is a good practice to use data redundancy mechanisms at the application layer to ensure data availability. Use EVS disks to store service data that needs to be stored for a long time.

- The device name of a local disk attached to an ultra-high I/O ECS is `/dev/nvme0n1` or `/dev/nvme0n2`.
- Local disks attached to Ir3 ECSs can be split for multiple ECSs to use. If a local disk is damaged, the ECSs that use this disk will be affected.
You are advised to add Ir3 ECSs to an ECS group during the creation process to prevent such failures. For details, see "Managing ECS Groups" in the *Elastic Cloud Server User Guide*.
- The basic resources, including vCPUs, memory, and image of an ultra-high I/O ECS will continue to be billed after the ECS is stopped. To stop the ECS from being billed, delete it and its associated resources.

1.7.6 High-Performance Computing ECSs

Overview

- The processor and memory ratio of an Hc2 (high-performance computing second-generation) ECS is 1:2 or 1:4. Each vCPU corresponds to the hyperthreading of an Intel® Xeon® Scalable processor core. Hc2 ECSs can be used for high-performance computing services. They provide a large number of parallel computing resources and high-performance infrastructure services to meet the requirements of high-performance computing and massive storage and ensure rendering efficiency. Hc2 ECSs are frequently used in the following scenarios:
 - Computing and storage systems for genetic engineering, games, animations, biopharmaceuticals, and scientific computing
 - Public rendering platforms for renderfarms and animation and film bases; other rendering platforms for movies and videos
 - High-performance frontend clusters, web servers, high-performance science and engineering applications, advertisements, video coding, and distributed analysis
- H2 ECSs are designed to meet high-end computational needs, such as molecular modeling and computational fluid dynamics. In addition to the substantial CPU power, the H2 ECSs offer diverse options for low-latency RDMA networking using EDR InfiniBand NICs to support memory-intensive computational requirements.
- H1 ECSs feature large memory capacity. They allow ECSs to interconnect with each other using 100 Gbit/s RDMA InfiniBand NICs and support 56 Gbit shared high I/O storage.
- H3 ECSs use high-performance Intel® Xeon® Scalable processors. Each vCPU corresponds to the hyper thread of an Intel® Xeon® Scalable processor core, providing stable computing capabilities. H3 ECSs are suitable for high-performance computing services. In addition, such ECSs use latest-generation network acceleration engines and DPDK rapid packet processing mechanism to provide high network performance.

Specifications

Table 1-32 High-performance computing ECS specifications

Type	vCPUs	Memory (GiB)	Flavor	Virtualization
Hc2	2	4	hc2.large.2	KVM
	4	8	hc2.xlarge.2	KVM
	8	16	hc2.2xlarge.2	KVM
	16	32	hc2.4xlarge.2	KVM
	32	64	hc2.8xlarge.2	KVM
	2	8	hc2.large.4	KVM
	4	16	hc2.xlarge.4	KVM
	8	32	hc2.2xlarge.4	KVM
	16	64	hc2.4xlarge.4	KVM
	32	128	hc2.8xlarge.4	KVM
Hl1	32	256	hl1.8xlarge.8	KVM
H2	16	128	h2.4xlarge.8	KVM
	16	256	h2.4xlarge.16	KVM
H3	2	4	h3.large.2	KVM
	4	8	h3.xlarge.2	KVM
	8	16	h3.2large.2	KVM
	12	24	h3.3xlarge.2	KVM
	16	32	h3.4large.2	KVM
	24	48	h3.6xlarge.2	KVM
	32	64	h3.8large.2	KVM
	2	8	h3.large.4	KVM
	4	16	h3.xlarge.4	KVM
	8	32	h3.2large.4	KVM
	12	48	h3.3xlarge.4	KVM
	16	64	h3.4large.4	KVM
	24	96	h3.6xlarge.4	KVM
	32	128	h3.8large.4	KVM

Scenarios

- Applications
H2 and H1: High-performance computing (HPC), big data, and Artificial Intelligence (AI)
- Application scenarios
H2 and H1
 - High-performance hardware: The ratio of memory to vCPU is 8:1, and a large number of multi-thread physical CPUs are available to provide high-performance storage I/O and high-throughput network connections.
 - Designed for HPC clusters: Multiple H1 ECSs can be clustered to install scalable, clustered file system, such as Lustre. HPC applications running on H2 ECSs can read and modify the data stored in the ECSs.
 - RDMA network connection: Same as H2 ECSs, H1 ECSs also offer RDMA network using EDR 100 Gbit/s InfiniBand NICs. H1 ECSs can communicate with H2 ECSs with RDMA protocol. In addition, H1 ECSs can access EVS disks over the RDMA protocol, which allows up to 56 Gbit/s storage bandwidth.
- Application scenarios
H2 and H1 ECSs provide computing capabilities for clusters with a large memory, good connectivity between nodes, and high storage I/O. The typical application scenarios include HPC, big data, and AI. In HPC solution, H1 ECSs are perfectly suited for the Lustre parallel distributed file system, generally used for large-scale cluster computing.
For example, in HPC scenario, H2 ECSs can be used as compute nodes, and H1 ECSs can be used as storage nodes.

Features

High-performance computing ECSs have the following features:

- They provide large memory capacity and more processor cores than other types of ECSs.
- They offer up to 32 vCPUs.
- H2 and H1 ECSs use InfiniBand NICs that provide a bandwidth of 100 Gbit/s.
- H1 ECSs can use the following types of EVS disks as system disk and data disk:
 - High I/O (performance-optimized I)
 - Ultra-high I/O (latency-optimized)
- H1 ECSs support 56 Gbit shared high I/O storage.
To support 56 Gbit shared high I/O storage, you only need to attach high I/O (performance-optimized I) or ultra-high I/O (latency-optimized) EVS disks to target H1 ECSs.

Notes on Using H2 ECSs

- H2 ECSs do not support OS reinstallation or change.
- H2 ECSs do not support specifications modification.

- H2 ECSs do not support cold migration, live migration, or HA.
- H2 ECSs support the following OSs:
 - For public images:
 - CentOS 7.3 64bit
 - SUSE Linux Enterprise Server 11 SP4 64bit
 - SUSE Linux Enterprise Server 12 SP2 64bit
 - For private images:
 - CentOS 6.5 64bit
 - CentOS 7.2 64bit
 - CentOS 7.3 64bit
 - SUSE Linux Enterprise Server 11 SP4 64bit
 - SUSE Linux Enterprise Server 12 SP2 64bit
 - Red Hat Enterprise Linux 7.2 64bit
 - Red Hat Enterprise Linux 7.3 64bit
- H2 ECSs use InfiniBand NICs that provide a bandwidth of 100 Gbit/s.
- Each H2 ECS uses one PCIe 3.2 TB SSD card for temporary local storage.
- If an H2 ECS is created using a private image, install an InfiniBand NIC driver on the ECS after the ECS creation following the instructions provided by Mellanox. Download the required version (4.2-1.0.0.0) of InfiniBand NIC driver from the official Mellanox website and install the driver by following the instructions provided by Mellanox.
 - InfiniBand NIC type: **Mellanox Technologies ConnectX-4 Infiniband HBA (MCX455A-ECAT)**
 - Mellanox official website: <http://www.mellanox.com/>
 - NIC driver download path: https://network.nvidia.com/products/infiniband-drivers/linux/mlnx_ofed/
- For SUSE H2 ECSs, if IP over InfiniBand (IPoIB) is required, you must manually configure an IP address for the InfiniBand NIC after installing the InfiniBand driver. For details, see "How Can I Manually Configure an IP Address for an InfiniBand NIC?" in the *Elastic Cloud Server User Guide*.
- After you delete an H2 ECS, the data stored in SSDs is automatically cleared. Therefore, do not store persistence data into SSDs during ECS running.

Notes on Using H1 ECSs

- H1 ECSs only support the attachment of high I/O (performance-optimized I) and ultra-high I/O (latency-optimized) EVS disks.

To support 56 Gbit shared high I/O storage, you only need to attach high I/O (performance-optimized I) or ultra-high I/O (latency-optimized) EVS disks to target H1 ECSs.
- H1 ECSs do not support specifications modification.

- H1 ECSs use InfiniBand NICs that provide a bandwidth of 100 Gbit/s.
- H1 ECSs created using a private image must have the InfiniBand NIC driver installed. Download the required version (4.2-1.0.0.0) of InfiniBand NIC driver from the official Mellanox website and install the driver by following the instructions provided by Mellanox.
 - InfiniBand NIC type: **Mellanox Technologies ConnectX-4 Infiniband HBA (MCX455A-ECAT)**
 - Mellanox official website: <http://www.mellanox.com/>
- For SUSE H1 ECSs, if IP over InfiniBand (IPoIB) is required, you must manually configure an IP address for the InfiniBand NIC after installing the InfiniBand driver. For details, see "How Can I Manually Configure an IP Address for an InfiniBand NIC?" in the *Elastic Cloud Server User Guide*.
- H1 ECSs support the following OSs:
 - For public images:
 - CentOS 7.3 64bit
 - SUSE Linux Enterprise Server 11 SP4 64bit
 - SUSE Linux Enterprise Server 12 SP2 64bit
 - For private images:
 - CentOS 6.5 64bit
 - CentOS 7.2 64bit
 - CentOS 7.3 64bit
 - SUSE Linux Enterprise Server 11 SP4 64bit
 - SUSE Linux Enterprise Server 12 SP2 64bit
 - Red Hat Enterprise Linux 7.2 64bit
 - Red Hat Enterprise Linux 7.3 64bit

1.7.7 GPU-accelerated ECSs

GPU-accelerated ECSs provide outstanding floating-point computing capabilities. They are suitable for applications that require real-time, highly concurrent massive computing.

GPU-accelerated ECSs are classified as G series and P series of ECSs.

- G series: Graphics-accelerated ECSs, which are suitable for 3D animation rendering and CAD
- P series: Computing-accelerated or inference-accelerated ECSs, which are suitable for deep learning, scientific computing, and CAE

GPU-accelerated ECS Types

Available now: All GPU models except the recommended ones. If available ECSs are sold out, use the recommended ones.

- G series
 - [Graphics-accelerated Enhancement G7r](#)
 - [Graphics-accelerated Enhancement G7v](#)
 - [Graphics-accelerated Enhancement G7](#)
 - [Graphics-accelerated Enhancement G6](#)
 - [Graphics-accelerated Enhancement G5r](#)
 - [Graphics-accelerated Enhancement G5](#)
- P series
 - [Computing-accelerated P3v](#)
 - [Computing-accelerated P2s](#) (recommended)
 - [Inference-accelerated Pi3](#)
 - [Inference-accelerated Pi2](#) (recommended)
 - [Inference-accelerated Pi2nl](#)

Helpful links:

- See "Installing a GRID Driver on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.
- See "Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.

Images Supported by GPU-accelerated ECSs

Table 1-33 Images supported by GPU-accelerated ECSs

Category	ECS Type	Supported Image
Graphics-accelerated	G7r	Provided by the cloud desktop
Graphics-accelerated	G7v	<ul style="list-style-type: none">• CentOS 8.2 64bit• CentOS 7.6 64bit• Ubuntu 20.04 Server 64bit• Ubuntu 18.04 Server 64bit• Windows Server 2019 Standard 64bit• Windows Server 2016 Standard 64bit
Graphics-accelerated	G7	<ul style="list-style-type: none">• CentOS 8.2 64bit• CentOS 7.6 64bit• Ubuntu 20.04 Server 64bit• Ubuntu 18.04 Server 64bit• Windows Server 2019 Standard 64bit• Windows Server 2016 Standard 64bit

Category	ECS Type	Supported Image
Graphics-accelerated	G6	<ul style="list-style-type: none">CentOS 8.2 64bitCentOS 7.6 64bitUbuntu 20.04 64bitUbuntu 18.04 64bitWindows Server 2019 Standard 64bitWindows Server 2016 Standard 64bit
Graphics-accelerated	G5r	<ul style="list-style-type: none">Windows Server 2016 Standard 64bitWindows Server 2012 R2 Standard 64bit
Graphics-accelerated	G5	<ul style="list-style-type: none">CentOS 8.2 64bitCentOS 7.6 64bitCentOS 7.5 64bitUbuntu 20.04 64bitUbuntu 18.04 64bitWindows Server 2019 Standard 64bitWindows Server 2016 Standard 64bitWindows Server 2019 Datacenter 64bitWindows Server 2016 Datacenter 64bit
Computing-accelerated	P3v	<ul style="list-style-type: none">CentOS 8.2 64bitCentOS 8.1 64bitCentOS 8.0 64bitCentOS 7.9 64bitCentOS 7.8 64bitCentOS 7.7 64bitCentOS 7.6 64bitUbuntu 20.04 server 64bitUbuntu 18.04 server 64bit
Computing-accelerated	P2s	<ul style="list-style-type: none">Windows Server 2016 Standard 64bit
Inference-accelerated	Pi3	<ul style="list-style-type: none">CentOS 8.2 64bitCentOS 8.1 64bitCentOS 8.0 64bitCentOS 7.9 64bitCentOS 7.8 64bitCentOS 7.7 64bitCentOS 7.6 64bitUbuntu 20.04 server 64bitUbuntu 18.04 server 64bit

Category	ECS Type	Supported Image
Inference-accelerated	Pi2	<ul style="list-style-type: none"> CentOS 7.5 64bit Windows Server 2019 Standard 64bit Windows Server 2016 Standard 64bit
Inference-accelerated	Pi2nl	<ul style="list-style-type: none"> CentOS 7.5 64bit Ubuntu 16.04 Server 64bit Windows Server 2016 Standard 64bit

Graphics-accelerated Enhancement G7r

Overview

G7r ECSs use NVIDIA Quadro RTX A6000 graphics card with up to 48 GiB GDDR6 GPU memory and support DirectX, Shader Model, OpenGL, and Vulkan. These ECSs theoretically can provide 38.7 TFLOPS of single precision (FP32) performance and 309.7 TFLOPS of sparse tensor performance. More tensor cores deliver more powerful performance to meet diverse graphics processing requirements.

Select your desired GPU-accelerated ECS type and specifications.

Specifications

Table 1-34 G7r ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	GPUs	GPU Memory (GiB)	Virtualization
g7r.3xlarge.4	12	48	17/5	200	4	6	1 × NVIDIA RTX A6000-6Q	6	KVM
g7r.4xlarge.4	16	64	20/6	280	8	8	1 × NVIDIA RTX A6000-8Q	8	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	GPUs	GPU Memory (GiB)	Virtualization
g7r.6xlarge.4	24	96	25/9	400	8	8	1 × NVIDIA RTX A6000-12Q	12	KVM
g7r.8xlarge.4	32	128	30/12	550	16	8	1 × NVIDIA RTX A6000-16Q	16	KVM
g7r.12xlarge.4	48	192	35/18	750	16	8	1 × NVIDIA RTX A6000-24Q	24	KVM
g7r.24xlarge.4	96	384	40/36	1,100	32	8	1 × NVIDIA RTX A6000	48	KVM

G7r ECS Features

- CPU: 3rd Generation Intel® Xeon® Scalable 8378A processors (3.0 GHz of base frequency and 3.5 GHz of turbo frequency)
- Graphics acceleration APIs
 - DirectX 12.0, Direct2D, DirectX Video Acceleration (DXVA)
 - Shader Model 5.1
 - OpenGL 4.6
 - Vulkan 1.1
- CUDA, DirectCompute, and OpenCL
- A single card is equipped with 10,752 CUDA cores, 84 second-generation RT cores, and 576 third-generation Tensor cores.
- Graphics applications accelerated
- Heavy-load CPU inference

- Application flow identical to common ECSs
- Automatic scheduling of G7r ECSs to AZs where NVIDIA Quadro RTX A6000 GPUs are used
- One NVENC (encoding) engine and two NVDEC (decoding) engines (including AV1 decoding) embedded

Supported Common Software

G7r ECSs are used in graphics acceleration scenarios, such as video rendering, cloud desktop, and 3D visualization. If the software relies on GPU DirectX and OpenGL hardware acceleration, use G7r ECSs. G7r ECSs support the following commonly used graphics processing software:

- AutoCAD
- 3DS MAX
- MAYA
- Agisoft PhotoScan
- ContextCapture

Notes

- After a G7r ECS is stopped, basic resources (including vCPUs, memory, image, and GPUs) are not billed, but its system disk is billed based on the disk capacity. If other products, such as EVS disks, EIP, and bandwidth are associated with the ECS, these products are billed separately.

NOTE

Resources will be released after a G7r ECS is stopped. If resources are insufficient at the next start, the start may fail. If you want to use such an ECS for a long period of time, do not stop the ECS.

- G7r ECSs created using a public image have had the GRID driver of a specific version installed by default. However, you need to purchase and configure a GRID license by yourself. Ensure that the GRID driver version meets service requirements.

For details about how to configure a GRID license, see "Installing a GRID Driver on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.

- If a G7r ECS is created using a private image, make sure that the GRID driver was installed during the private image creation. If the GRID driver has not been installed, install the driver for graphics acceleration after the ECS is created.

For details, see "Installing a GRID Driver on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.

- GPU-accelerated ECSs differ greatly in general-purpose and heterogeneous computing power. Their specifications can only be changed to other specifications of the same instance type.

Graphics-accelerated Enhancement G7v

Overview

G7v ECSs use NVIDIA A40 GPUs and support DirectX, Shader Model, OpenGL, and Vulkan. Each GPU provides 48 GiB of GPU memory. Theoretically, G7v ECSs can

provide 37.4 TFLOPS of FP32 peak performance and 74.8 TFLOPS (sparsity disabled) or 149.6 TFLOPS (sparsity enabled) of TF32 peak tensor performance. They deliver two times the rendering performance and 1.4 times the graphics processing performance of RTX6000 GPUs to meet professional graphics processing requirements.

Select your desired GPU-accelerated ECS type and specifications.

Specifications

Table 1-35 G7v ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	GPUs	GPU Memory (GiB)	Virtu alization
g7v.2xlarge.8	8	64	15/3	100	4	4	1 × NVIDIA A40-8Q	8	KVM
g7v.4xlarge.8	16	128	20/6	150	8	8	1 × NVIDIA A40-16Q	16	KVM
g7v.6xlarge.8	24	192	25/9	200	8	8	1 × NVIDIA A40-24Q	24	KVM

G7v ECS Features

- CPU: 3rd Generation Intel® Xeon® Scalable 8378A processors (3.0 GHz of base frequency and 3.5 GHz of turbo frequency)
- Graphics acceleration APIs
 - DirectX 12.07, Direct2D, DirectX Video Acceleration (DXVA)
 - Shader Model 5.17
 - OpenGL 4.68
 - Vulkan 1.18
- CUDA, DirectCompute, OpenACC, and OpenCL
- A single card is equipped with 10,752 CUDA cores, 84 second-generation RT cores, and 336 third-generation Tensor cores.

- Graphics applications accelerated
- Heavy-load CPU inference
- Application flow identical to common ECSs
- Automatic scheduling of G7v ECSs to AZs where NVIDIA A40 GPUs are used
- One NVENC (encoding) engine and two NVDEC (decoding) engines (including AV1 decoding) embedded

Supported Common Software

G7v ECSs are used in graphics acceleration scenarios, such as video rendering, cloud desktop, and 3D visualization. If the software relies on GPU DirectX and OpenGL hardware acceleration, use G7v ECSs. G7v ECSs support the following commonly used graphics processing software:

- AutoCAD
- 3DS MAX
- MAYA
- Agisoft PhotoScan
- ContextCapture
- Adobe Premiere Pro
- Solidworks
- Unreal Engine
- Blender
- Vray

Notes

- After a G7v ECS is stopped, basic resources (including vCPUs, memory, image, and GPUs) are not billed, but its system disk is billed based on the disk capacity. If other products, such as EVS disks, EIP, and bandwidth are associated with the ECS, these products are billed separately.

NOTE

Resources will be released after a G7v ECS is stopped. If resources are insufficient at the next start, the start may fail. If you want to use such an ECS for a long period of time, do not stop the ECS.

- G7v ECSs created using a public image have had the GRID driver of a specific version installed by default. However, you need to purchase and configure a GRID license by yourself. Ensure that the GRID driver version meets service requirements.

For details about how to configure a GRID license, see "Installing a GRID Driver on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.

- If a G7v ECS is created using a private image, make sure that the GRID driver was installed during the private image creation. If the GRID driver has not been installed, install the driver for graphics acceleration after the ECS is created.

For details, see "Installing a GRID Driver on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.

- GPU-accelerated ECSs differ greatly in general-purpose and heterogeneous computing power. Their specifications can only be changed to other specifications of the same instance type.

Graphics-accelerated Enhancement G7

Overview

G7 ECSs use NVIDIA A40 GPUs and support DirectX, Shader Model, OpenGL, and Vulkan. Each GPU provides 48 GiB of GPU memory. Theoretically, G7 ECSs provide 37.4 TFLOPS of FP32 peak performance and 74.8 TFLOPS (sparsity disabled) or 149.6 TFLOPS (sparsity enabled) of TF32 peak tensor performance. They deliver two times the rendering performance and 1.4 times the graphics processing performance of RTX6000 GPUs to meet professional graphics processing requirements.

Select your desired GPU-accelerated ECS type and specifications.

Specifications

Table 1-36 G7 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	GPUs	GPU Memory (GiB)	Virtu alization
g7.1 2xlarge.8	48	384	35/18	750	16	8	1 × NVIDIA A-A40	1 × 48	KVM
g7.2 4xlarge.8	96	768	40/36	850	16	8	2 × NVIDIA A-A40	2 × 48	KVM

G7 ECS Features

- CPU: 3rd Generation Intel® Xeon® Scalable 8378A processors (3.0 GHz of base frequency and 3.5 GHz of turbo frequency)
- Graphics acceleration APIs
 - DirectX 12.07, Direct2D, DirectX Video Acceleration (DXVA)
 - Shader Model 5.17
 - OpenGL 4.68
 - Vulkan 1.18
- CUDA, DirectCompute, OpenACC, and OpenCL

- A single card is equipped with 10,752 CUDA cores, 84 second-generation RT cores, and 336 third-generation Tensor cores.
- Graphics applications accelerated
- Heavy-load CPU inference
- Application flow identical to common ECSs
- Automatic scheduling of G7 ECSs to AZs where NVIDIA A40 GPUs are used
- One NVENC (encoding) engine and two NVDEC (decoding) engines (including AV1 decoding) embedded

Supported Common Software

G7 ECSs are used in graphics acceleration scenarios, such as video rendering, cloud desktop, and 3D visualization. If the software relies on GPU DirectX and OpenGL hardware acceleration, use G7 ECSs. G7 ECSs support the following commonly used graphics processing software:

- AutoCAD
- 3DS MAX
- MAYA
- Agisoft PhotoScan
- ContextCapture
- Adobe Premiere Pro
- Solidworks
- Unreal Engine
- Blender
- Vray

Notes

- After a G7 ECS is stopped, basic resources (including vCPUs, memory, image, and GPUs) are not billed, but its system disk is billed based on the disk capacity. If other products, such as EVS disks, EIP, and bandwidth are associated with the ECS, these products are billed separately.

NOTE

Resources will be released after a G7 ECS is stopped. If resources are insufficient at the next start, the start may fail. If you want to use such an ECS for a long period of time, do not stop the ECS.

- G7 ECSs created using a public image have had the GRID driver of a specific version installed by default. However, you need to purchase and configure a GRID license by yourself. Ensure that the GRID driver version meets service requirements.

For details about how to configure a GRID license, see "Installing a GRID Driver on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.

- If a G7 ECS is created using a private image, make sure that the GRID driver was installed during the private image creation. If the GRID driver has not been installed, install the driver for graphics acceleration after the ECS is created.

For details, see "Installing a GRID Driver on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.

- GPU-accelerated ECSs differ greatly in general-purpose and heterogeneous computing power. Their specifications can only be changed to other specifications of the same instance type.

Graphics-accelerated Enhancement G6

Overview

G6 ECSs use NVIDIA Tesla T4 GPUs to support DirectX, OpenGL, and Vulkan and provide 16 GiB of GPU memory. The theoretical Pixel rate is 101.8 Gpixel/s and Texture rate 254.4 GTexel/s, meeting professional graphics processing requirements.

Select your desired GPU-accelerated ECS type and specifications.

Specifications

Table 1-37 G6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	GPUs	GPU Memory (GiB)	Virtualization
g6.xlarge.4	4	16	6/2	200	8	8	1 × T4	16	KVM
g6.4xlarge.4	16	64	15/8	200	8	8	1 × T4	16	KVM
g6.6xlarge.4	24	96	25/15	200	8	8	1 × T4	16	KVM
g6.9xlarge.7	36	252	25/15	200	16	8	1 × T4	16	KVM
g6.10xlarge.7	40	280	25/15	200	16	8	1 × T4	16	KVM
g6.18xlarge.7	72	504	30/30	400	32	16	2 × T4	32	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	GPUs	GPU Memory (GiB)	Virtualization
g6.2 Oxlarge.7	80	560	30/30	400	32	16	2 × T4	32	KVM

G6 ECS Features

- CPU: 2nd Generation Intel® Xeon® Scalable 6266 processors (3.0 GHz of base frequency and 3.4 GHz of turbo frequency)
- Graphics acceleration APIs
 - DirectX 12, Direct2D, and DirectX Video Acceleration (DXVA)
 - OpenGL 4.5
 - Vulkan 1.0
- CUDA and OpenCL
- NVIDIA T4 GPUs
- Graphics applications accelerated
- Heavy-load CPU inference
- Automatic scheduling of G6 ECSs to AZs where NVIDIA T4 GPUs are used
- One NVENC engine and two NVDEC engines embedded

Supported Common Software

G6 ECSs are used in graphics acceleration scenarios, such as video rendering, cloud desktop, and 3D visualization. If the software relies on GPU DirectX and OpenGL hardware acceleration, use G6 ECSs. G6 ECSs support the following commonly used graphics processing software:

- AutoCAD
- 3DS MAX
- MAYA
- Agisoft PhotoScan
- ContextCapture

Notes

- After a G6 ECS is stopped, basic resources (including vCPUs, memory, image, and GPUs) are not billed, but its system disk is billed based on the disk capacity. If other products, such as EVS disks, EIP, and bandwidth are associated with the ECS, these products are billed separately.

 **NOTE**

Resources will be released after a G6 ECS is stopped. If resources are insufficient at the next start, the start may fail. If you want to use such an ECS for a long period of time, do not stop the ECS.

- G6 ECSs created using a public image have had the GRID driver of a specific version installed by default. However, you need to purchase and configure a GRID license by yourself. Ensure that the GRID driver version meets service requirements.

For details about how to configure a GRID license, see "Installing a GRID Driver on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.

- If a G6 ECS is created using a private image, make sure that the GRID driver was installed during the private image creation. If not, install the driver for graphics acceleration after the ECS is created.

For details, see "Installing a GRID Driver on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.

- GPU-accelerated ECSs differ greatly in general-purpose and heterogeneous computing power. Their specifications can only be changed to other specifications of the same instance type.

Graphics-accelerated Enhancement G5r

Overview

G5r ECSs are based on PCI passthrough and exclusively use GPUs for professional graphics acceleration. G5r ECSs equipped with NVIDIA Quadro RTX5000 GPUs support DirectX and OpenGL and provide a maximum GPU memory of 16 GiB for rendering, cloud gaming, and graphics workstations.

Specifications

Table 1-38 G5r ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	GPUs	GPU Memory (GiB)	Virtualization
g5r.8xlarge.2	32	64	10/4	100	4	1 × RTX5000	16	KVM

 **NOTE**

NVIDIA Quadro RTX5000 GPUs use the latest-generation Turing architecture and work with the latest NVIDIA RTX platform.

G5r ECS Features

- CPU: 2nd Generation Intel® Xeon® Scalable 6278 processors (2.6 GHz of base frequency and 3.5 GHz of turbo frequency), or Intel® Xeon® Scalable 6151 processors (3.0 GHz of base frequency and 3.4 GHz of turbo frequency)
- Grating acceleration based on 48 RT cores
- NVIDIA RTX5000 GPUs
- Rendering graphics acceleration
- Deep learning based on 3,072 CUDA cores and 384 Tensor cores
- GPU passthrough
- A maximum GPU memory of 16 GiB

Notes

- After a G5r ECS is stopped, basic resources (including vCPUs, memory, image, and GPUs) are not billed, but its system disk is billed based on the disk capacity. If other products, such as EVS disks, EIP, and bandwidth are associated with the ECS, these products are billed separately.

NOTE

Resources will be released after a G5r ECS is stopped. If resources are insufficient at the next start, the start may fail. If you want to use such an ECS for a long period of time, do not stop the ECS.

- G5r ECSs do not support specifications modification.
- G5r ECSs are in open beta testing. Contact customer service for the test.
- GPU-accelerated ECSs differ greatly in general-purpose and heterogeneous computing power. Their specifications can only be changed to other specifications of the same instance type.

Graphics-accelerated Enhancement G5

Overview

G5 ECSs use NVIDIA GRID vGPUs and provide comprehensive, professional graphics acceleration. They use NVIDIA Tesla V100 GPUs and support DirectX, OpenGL, and Vulkan. These ECSs provide 1, 2, 4, 8, or 16 GiB of GPU memory and up to 4096 x 2160 resolution, meeting requirements from entry-level through professional graphics processing.

Select your desired GPU-accelerated ECS type and specifications.

Specifications

Table 1-39 G5 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	GPUs	GPU Memory (GiB)	Virtualization
g5.xlarge.4	32	128	25/15	200	16	1 × V100	16	KVM

 **NOTE**

V100-xQ indicates that V100 GPUs are virtualized to vGPUs with different specifications and models using GRID. **x** specifies the vGPU memory, and **Q** indicates that the vGPU of this type is designed to work in workstations and desktop scenarios. For more details about GRID vGPUs, see [GRID VIRTUAL GPU User Guide](#).

G5 ECS Features

- CPU: 2nd Generation Intel® Xeon® Scalable 6278 processors (2.6 GHz of base frequency and 3.5 GHz of turbo frequency), or Intel® Xeon® Scalable 6151 processors (3.0 GHz of base frequency and 3.4 GHz of turbo frequency)
- Graphics acceleration APIs
 - DirectX 12, Direct2D, and DirectX Video Acceleration (DXVA)
 - OpenGL 4.5
 - Vulkan 1.0
- CUDA and OpenCL
- Quadro vDWS for professional graphics acceleration
- NVIDIA V100 GPUs
- Graphics applications accelerated
- GPU hardware virtualization (vGPUs)
- Automatic scheduling of G5 ECSs to AZs where NVIDIA V100 GPUs are used
- A maximum specification of 16 GiB of GPU memory and 4096 x 2160 resolution for processing graphics and videos

Supported Common Software

G5 ECSs are used in graphics acceleration scenarios, such as video rendering, cloud desktop, and 3D visualization. If the software relies on GPU DirectX and OpenGL hardware acceleration, use G5 ECSs. G5 ECSs support the following commonly used graphics processing software:

- AutoCAD
- 3DS MAX
- MAYA
- Agisoft PhotoScan
- ContextCapture

Notes

- After a G5 ECS is stopped, basic resources (including vCPUs, memory, image, and GPUs) are not billed, but its system disk is billed based on the disk capacity. If other products, such as EVS disks, EIP, and bandwidth are associated with the ECS, these products are billed separately.

NOTE

Resources will be released after a G5 ECS is stopped. If resources are insufficient at the next start, the start may fail. If you want to use such an ECS for a long period of time, do not stop the ECS.

- When the Windows OS running on a G5 ECS is started, the GRID driver is loaded by default, and vGPUs are used for video output by default. In such a case, the remote login function provided on the management console is not supported. To access such an ECS, use RDP, such as Windows MSTSC. Then, install a third-party VDI tool on the ECS for remote login, such as VNC.
- For G5 ECSs, you need to configure the GRID license after the ECS is created.
- G5 ECSs created using a public image have had the GRID driver of a specific version installed by default. However, you need to purchase and configure a GRID license by yourself. Ensure that the GRID driver version meets service requirements.

For details about how to configure a GRID license, see "Installing a GRID Driver on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.

- If a G5 ECS is created using a private image, make sure that the GRID driver was installed during the private image creation. If not, install the driver for graphics acceleration after the ECS is created.

For details, see "Installing a GRID Driver on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.

- GPU-accelerated ECSs differ greatly in general-purpose and heterogeneous computing power. Their specifications can only be changed to other specifications of the same instance type.

Computing-accelerated P3v

Overview

P3v ECSs use NVIDIA A800 GPUs and provide flexibility and ultra-high-performance computing. P3v ECSs have strengths in AI-based deep learning, scientific computing, Computational Fluid Dynamics (CFD), computing finance, seismic analysis, molecular modeling, and genomics. P3v ECSs provide theoretic 19.5 TFLOPS of FP32 single-precision performance, 156 TFLOPS of TF32 core floating-point performance, and 312 TFLOPS of Bfloat16.

Specifications

Table 1-40 P3v ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	GPUs	GPU Connection	GPU Memory (GiB)	Virtualization
p3v.3xlarge.8	12	96	17/5	200	4	4	1 × NVIDIA A800 80GB	N/A	80	KVM
p3v.24xlarge.8	96	768	40/36	850	32	8	8 × NVIDIA A800 80GB	NVL link	640	KVM

P3v ECS Features

- CPU: 3rd Generation Intel® Xeon® Scalable 6348 processors (2.6 GHz of base frequency and 3.5 GHz of turbo frequency)
- Up to eight NVIDIA A800 GPUs on an ECS
- NVIDIA CUDA parallel computing and common deep learning frameworks, such as TensorFlow, Caffe, PyTorch, and MXNet
- 19.5 TFLOPS of single-precision computing and 9.7 TFLOPS of double-precision computing
- NVIDIA Tensor cores with 156 TFLOPS of single- and double-precision computing for deep learning
- Up to 40 Gbit/s of network bandwidth on a single ECS
- 80 GB HBM2 GPU memory per graphics card, and multiple GPU cards interconnected based on NVLink for up to 2,039 Gbit/s
- Comprehensive basic capabilities
 - User-defined network with flexible subnet division and network access policy configuration
 - Mass storage, elastic expansion, and backup and restoration
 - Elastic scaling
- Flexibility
Similar to other types of ECSs, P3v ECSs can be provisioned in a few minutes.
- Excellent supercomputing ecosystem
The supercomputing ecosystem allows you to build up a flexible, high-performance, cost-effective computing platform. A large number of HPC applications and deep-learning frameworks can run on P3v ECSs.

Supported Common Software

P3v ECSs are used in computing acceleration scenarios, such as deep learning training, inference, scientific computing, molecular modeling, and seismic analysis. If the software is required to support GPU CUDA, use P3v ECSs. P2vs ECSs support the following commonly used software:

- Common deep learning frameworks, such as TensorFlow, Spark, PyTorch, MXNet, and Caffee
- CUDA GPU rendering supported by RedShift for Autodesk 3dsMax and V-Ray for 3ds Max
- Agisoft PhotoScan
- MapD
- More than 2,000 GPU-accelerated applications such as Amber, NAMD, and VASP

Notes

- After a P3v ECS is stopped, basic resources (including vCPUs, memory, image, and GPUs) are not billed, but its system disk is billed based on the disk capacity. If other products, such as EVS disks, EIP, and bandwidth are associated with the ECS, these products are billed separately.

NOTE

- Resources will be released after a P3v ECS is stopped. If resources are insufficient at the next start, the start may fail. If you want to use such an ECS for a long period of time, do not stop the ECS.
- If a P3v ECS is created using a private image, make sure that the Tesla driver was installed during the private image creation. If not, install the driver for computing acceleration after the ECS is created. For details, see "Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.
- GPU-accelerated ECSs differ greatly in general-purpose and heterogeneous computing power. Their specifications can only be changed to other specifications of the same instance type.

Computing-accelerated P2s

Overview

P2s ECSs use NVIDIA Tesla V100 GPUs to provide flexibility, high-performance computing, and cost-effectiveness. P2s ECSs provide outstanding general computing capabilities and have strengths in AI-based deep learning, scientific computing, Computational Fluid Dynamics (CFD), computing finance, seismic analysis, molecular modeling, and genomics.

Specifications

Table 1-41 P2s ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	GPUs	GPU Connection	GPU Memory (GiB)	Virtualization
p2s.2xlarge.8	8	64	10/4	50	4	4	1 × V100	PCIe Gen 3	1 × 32 GiB	KVM
p2s.4xlarge.8	16	128	15/8	100	8	8	2 × V100	PCIe Gen 3	2 × 32 GiB	KVM
p2s.8xlarge.8	32	256	25/15	200	16	8	4 × V100	PCIe Gen 3	4 × 32 GiB	KVM
p2s.16xlarge.8	64	512	30/30	400	32	8	8 × V100	PCIe Gen 3	8 × 32 GiB	KVM

P2s ECS Features

- CPU: 2nd Generation Intel® Xeon® Scalable 6278 processors (2.6 GHz of base frequency and 3.5 GHz of turbo frequency), or Intel® Xeon® Scalable 6151 processors (3.0 GHz of base frequency and 3.4 GHz of turbo frequency)
- Up to eight NVIDIA Tesla V100 GPUs on an ECS
- NVIDIA CUDA parallel computing and common deep learning frameworks, such as TensorFlow, Caffe, PyTorch, and MXNet
- 14 TFLOPS of single-precision computing and 7 TFLOPS of double-precision computing
- NVIDIA Tensor cores with 112 TFLOPS of single- and double-precision computing for deep learning
- Up to 30 Gbit/s of network bandwidth on a single ECS
- 32 GiB of HBM2 GPU memory with a bandwidth of 900 Gbit/s
- Comprehensive basic capabilities
 - User-defined network with flexible subnet division and network access policy configuration
 - Mass storage, elastic expansion, and backup and restoration
 - Elastic scaling
- Flexibility
Similar to other types of ECSs, P2s ECSs can be provisioned in a few minutes.
- Excellent supercomputing ecosystem

The supercomputing ecosystem allows you to build up a flexible, high-performance, cost-effective computing platform. A large number of HPC applications and deep-learning frameworks can run on P2s ECSs.

Supported Common Software

P2s ECSs are used in computing acceleration scenarios, such as deep learning training, inference, scientific computing, molecular modeling, and seismic analysis. If the software is required to support GPU CUDA, use P2s ECSs. P2s ECSs support the following commonly used software:

- Common deep learning frameworks, such as TensorFlow, Caffe, PyTorch, and MXNet
- CUDA GPU rendering supported by RedShift for Autodesk 3dsMax and V-Ray for 3ds Max
- Agisoft PhotoScan
- MapD

Notes

- After a P2s ECS is stopped, basic resources (including vCPUs, memory, image, and GPUs) are not billed, but its system disk is billed based on the disk capacity. If other products, such as EVS disks, EIP, and bandwidth are associated with the ECS, these products are billed separately.

NOTE

Resources will be released after a P2s ECS is stopped. If resources are insufficient at the next start, the start may fail. If you want to use such an ECS for a long period of time, do not stop the ECS.

- If a P2s ECS is created using a private image, make sure that the Tesla driver was installed during the private image creation. If not, install the driver for computing acceleration after the ECS is created. For details, see "Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.
- GPU-accelerated ECSs differ greatly in general-purpose and heterogeneous computing power. Their specifications can only be changed to other specifications of the same instance type.

Inference-accelerated Pi3

Overview

Pi3 ECSs use NVIDIA A30 GPUs dedicated for real-time AI inference. 24 GiB of GPU memory plus up to 933 GB/s of bandwidth allows Pi3 ECSs to be used in AI training scenarios. Its theoretical AI training throughput is three times that of NVIDIA V100 graphics card and six times that of T4 graphics cards on previous-generation Pi2 ECSs. With NVIDIA A30 GPUs, Pi3 ECSs provide 330 peak INT 8 TOPS (with sparsity enabled). Theoretically, Pi3 ECSs provide 10.3 TFLOPS of TF32 single-precision performance and 82 TFLOPS (sparsity disabled) or 165 TFLOPS (sparsity enabled) of TF32 peak tensor performance.

Specifications

Table 1-42 Pi3 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	GPUs	GPU Connection	GPU Memory (GiB)	Virtualization
pi3.6xlarge.4	24	96	25/9	400	8	8	1 × NVIDIA A30	24	-	KVM
pi3.12xlarge.4	48	192	35/18	750	8	8	2 × NVIDIA A30	48	-	KVM

Pi3 ECS Features

- CPU: 3rd Generation Intel® Xeon® Scalable 6348 processors (2.6 GHz of base frequency and 3.5 GHz of turbo frequency)
- Up to two NVIDIA A30 GPUs on an ECS, and multiple GPU cards interconnected based on NVLink
- Up to 10.3 TFLOPS of single-precision computing on a single GPU
- Up to 330 TOPS of INT8 computing on a single GPU
- 24 GiB of HBM2 GPU memory with a bandwidth of 933 GiB/s on a single GPU
- One OFA, one NVJPEG, and four NVDECs embedded

Supported Common Software

Pi3 ECSs are used in GPU-based inference computing scenarios, such as image recognition, speech recognition, and natural language processing. The Pi3 ECSs can also be used for light-load training.

Pi3 ECSs support the following commonly used software:

- Deep learning frameworks, such as TensorFlow, Caffe, PyTorch, MXNet, and Spark
- More than 2,000 GPU-accelerated software applications such as AMBER, NAMD, and OPENFOAM

Notes

- After a Pi3 ECS is stopped, basic resources (including vCPUs, memory, image, and GPUs) are not billed, but its system disk is billed based on the disk capacity. If other products, such as EVS disks, EIP, and bandwidth are associated with the ECS, these products are billed separately.

 NOTE

Resources will be released after a Pi3 ECS is stopped. If resources are insufficient at the next start, the start may fail. If you want to use such an ECS for a long period of time, do not stop the ECS.

- Pi3 ECSs support automatic recovery when the hosts accommodating such ECSs become faulty.
- When creating a Pi3 ECS, make sure that the Tesla driver was installed during the private image creation. If not, install the driver for computing acceleration after the ECS is created. For details, see "Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.
- GPU-accelerated ECSs differ greatly in general-purpose and heterogeneous computing power. Their specifications can only be changed to other specifications of the same instance type.

Inference-accelerated Pi2

Overview

Pi2 ECSs use NVIDIA Tesla T4 GPUs dedicated for real-time AI inference. These ECSs use the T4 INT8 calculator for up to 130 TOPS of INT8 computing. The Pi2 ECSs can also be used for light-load training.

Specifications

Table 1-43 Pi2 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	GPUs	GPU Memory (GiB)	Local Disks	Virtualization
pi2.xlarge.4	4	16	8/2	25	2	1 × T4	1 × 16	N/A	KVM
pi2.2xlarge.4	8	32	10/4	50	4	1 × T4	1 × 16	N/A	KVM
pi2.3xlarge.4	12	48	12/6	80	6	1 × T4	1 × 16	N/A	KVM
pi2.4xlarge.4	16	64	15/8	100	8	2 × T4	2 × 16	N/A	KVM
pi2.8xlarge.4	32	128	25/15	200	16	4 × T4	4 × 16	N/A	KVM

Pi2 ECS Features

- CPU: 2nd Generation Intel® Xeon® Scalable 6278 processors (2.6 GHz of base frequency and 3.5 GHz of turbo frequency), or Intel® Xeon® Scalable 6151 processors (3.0 GHz of base frequency and 3.4 GHz of turbo frequency)
- Up to four NVIDIA Tesla T4 GPUs on an ECS
- GPU hardware passthrough
- Up to 8.1 TFLOPS of single-precision computing on a single GPU
- Up to 130 TOPS of INT8 computing on a single GPU
- 16 GiB of GDDR6 GPU memory with a bandwidth of 320 GiB/s on a single GPU
- One NVENC engine and two NVDEC engines embedded

Supported Common Software

Pi2 ECSs are used in GPU-based inference computing scenarios, such as image recognition, speech recognition, and natural language processing. The Pi2 ECSs can also be used for light-load training.

Pi2 ECSs support the following commonly used software:

- Deep learning frameworks, such as TensorFlow, Caffe, PyTorch, and MXNet

Notes

- After a Pi2 ECS is stopped, basic resources (including vCPUs, memory, image, and GPUs) are not billed, but its system disk is billed based on the disk capacity. If other products, such as EVS disks, EIP, and bandwidth are associated with the ECS, these products are billed separately.

NOTE

Resources will be released after a Pi2 ECS is stopped. If resources are insufficient at the next start, the start may fail. If you want to use such an ECS for a long period of time, do not stop the ECS.

- Pi2 ECSs support automatic recovery when the hosts accommodating such ECSs become faulty.
- By default, Pi2 ECSs created using a public image have the Tesla driver installed.
- If a Pi2 ECS is created using a private image, make sure that the Tesla driver was installed during the private image creation. If not, install the driver for computing acceleration after the ECS is created. For details, see "Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.
- GPU-accelerated ECSs differ greatly in general-purpose and heterogeneous computing power. Their specifications can only be changed to other specifications of the same instance type.

Inference-accelerated Pi2nl

Overview

Pi2nl ECSs use NVIDIA Tesla T4 GPUs dedicated for real-time AI inference. These ECSs use the T4 INT8 calculator for up to 130 TOPS of INT8 computing. The Pi2nl ECSs can also be used for light-workload training.

Specifications

Table 1-44 Pi2nl ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	GPUs	GPU Memory (GiB)	Local Disks	Virtualization
pi2nl.2xlarge.4	8	32	10/4	50	4	1 × T4	1 × 16	N/A	KVM
pi2nl.4xlarge.4	16	64	15/8	100	8	2 × T4	2 × 16	N/A	KVM
pi2nl.8xlarge.4	32	128	25/15	200	16	4 × T4	4 × 16	N/A	KVM

Pi2nl ECS Features

- CPU: 2nd Generation Intel® Xeon® Scalable 6278 processors (2.6 GHz of base frequency and 3.5 GHz of turbo frequency), or Intel® Xeon® Scalable 6151 processors (3.0 GHz of base frequency and 3.4 GHz of turbo frequency)
- Up to four NVIDIA Tesla T4 GPUs on an ECS
- GPU hardware passthrough
- Up to 8.1 TFLOPS of single-precision computing on a single GPU
- Up to 130 TOPS of INT8 computing on a single GPU
- 16 GiB of GDDR6 GPU memory with a bandwidth of 320 GiB/s on a single GPU
- One NVENC engine and two NVDEC engines embedded

Supported Common Software

Pi2nl ECSs are used in GPU-based inference computing scenarios, such as image recognition, speech recognition, and natural language processing. The Pi2nl ECSs can also be used for light-load training.

Pi2 ECSs support the following commonly used software:

- Deep learning frameworks, such as TensorFlow, Caffe, PyTorch, and MXNet

Notes

- After a Pi2nl ECS is stopped, basic resources (including vCPUs, memory, image, and GPUs) are not billed, but its system disk is billed based on the disk capacity. If other products, such as EVS disks, EIP, and bandwidth are associated with the ECS, these products are billed separately.

NOTE

Resources will be released after a Pi2nl ECS is stopped. If resources are insufficient at the next start, the start may fail. If you want to use such an ECS for a long period of time, do not stop the ECS.

- Pi2nl ECSs support automatic recovery when the hosts accommodating such ECSs become faulty.
- By default, Pi2nl ECSs created using a public image have the Tesla driver installed.
- If a Pi2nl ECS is created using a private image, make sure that the Tesla driver was installed during the private image creation. If the Tesla driver has not been installed, install the driver for computing acceleration after the ECS is created. For details, see "Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS" in the *Elastic Cloud Server User Guide*.
- GPU-accelerated ECSs differ greatly in general-purpose and heterogeneous computing power. Their specifications can only be changed to other specifications of the same instance type.

1.7.8 AI-accelerated ECSs

AI-accelerated ECSs, powered by Ascend processors and software stacks, are dedicated for accelerating AI applications.

AI inference-accelerated ECSs use Ascend 310 processors for AI inference acceleration.

AI-accelerated ECS Types

AI inference-accelerated: [Enhanced AI Inference-accelerated Ai2 \(Type II\)](#) and [Enhanced AI Inference-accelerated Ai1s \(Type I\)](#)

Public Images Supported by AI-accelerated ECSs

Table 1-45 Public images

Category	ECS Type	Public Images
Enhanced AI inference-accelerated (type II)	Ai2	Ubuntu Server 18.04 64bit CentOS 7.8 64bit CentOS 7.6 64bit
Enhanced AI inference-accelerated (type I)	Ai1s	Ubuntu Server 18.04 64bit CentOS 7.6 64bit

Enhanced AI Inference-accelerated Ai2 (Type II)

Overview

Ai2 ECSs use Ascend 310P processors for AI acceleration. Featuring low power consumption and high computing power, Ascend 310P processors have

significantly improved energy efficiency ratio (EER), promoting the wide application of AI inference. AI inference-accelerated Ai2 ECSs enable you to easily use the powerful processing capability of Ascend 310P processors.

Ai2 ECSs are ideal for computer vision, smart campus, smart city, smart transportation, smart retail, Internet-based real-time communication, and video encoding and decoding scenarios.

Specifications

Table 1-46 Ai2 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max. / Assured Bandwidth	Max. PPS (10,000)	Ascend 310P	Ascend RAM (GiB)	Max. NIC Queues	Max. NICs	Virtualization
ai2.xlarge.4	4	16	8/1.6	80	1	24	3	3	KVM
ai2.2xlarge.4	8	32	15/3	150	1	24	4	4	KVM
ai2.4xlarge.4	16	64	20/6	280	1	24	8	8	KVM
ai2.8xlarge.4	32	128	30/12	550	2	48	8	8	KVM
ai2.16xlarge.4	64	256	36/24	800	4	96	16	8	KVM
ai2.24xlarge.4	96	384	40/36	850	6	144	32	8	KVM

Features

Ai2 ECSs have the following features:

- 1:4 ratio of vCPUs to memory
- CPU: 3rd Generation Intel® Xeon® Scalable 6348 processors (2.6 GHz of base frequency and 3.5 GHz of turbo frequency), or Intel® Xeon® Scalable 8378A processors (3.0 GHz of base frequency and 3.4 GHz of turbo frequency)

- Each Atlas 300I Pro accelerator card is equipped with one Ascend 310P processor, including eight DaVinci AI cores and eight self-developed CPU cores.
- 140 TeraOPS of integer-precision (INT8) on a single GPU
- 24 GiB of GPU memory with a bandwidth of 204.8 GiB/s on a single GPU
- Built-in hardware video codec engine for HD video decoder (H.264/265) and JPEG image encoding and decoding

Notes

1. Ai2 ECSs support the following public images:
 - Ubuntu Server 18.04 64bit
 - CentOS 7.8 64bit
 - CentOS 7.6 64bit
2. Ai2 ECSs do not support modification of specifications.
3. Ai2 ECSs support automatic recovery when the hosts accommodating such ECSs become faulty.

Enhanced AI Inference-accelerated Ai1s (Type I)

Overview

Ai1s ECSs use Ascend 310 processors for AI acceleration. Ascend 310 processors feature low power consumption, high computing capabilities, and significantly improved energy efficiency ratio (EER). This facilitates the wide application of AI inference. Ai1s ECSs deliver the computing acceleration capabilities of the Ascend 310 processors on the cloud platform.

Ai1s ECSs are based on Atlas 300I accelerator cards. For details, go to [Ascend Community](#).

AI-accelerated ECSs are ideal for computer vision, smart campus, smart city, smart transportation, smart retail, Internet-based real-time communication, and video encoding and decoding scenarios.

Specifications

Table 1-47 Ai1s ECS specifications

Flavor	vCPUs	Memory (GiB)	Max. / Assured Bandwidth	Max. PPS (10,000)	Ascend 310 Processors	Ascend RAM (GiB)	Max. NIC Queues	Max. NICs	Virtualization
ai1s.large.4	2	8	4/1.3	20	1	8	2	2	KVM

Flavor	vCPUs	Memory (GiB)	Max. / Assured Bandwidth	Max. PPS (10,000)	Ascend 310 Processors	Ascend RAM (GiB)	Max. NIC Queues	Max. NICs	Virtualization
ai1s.xlarge.4	4	16	6/2	35	2	16	2	3	KVM
ai1s.2xlarge.4	8	32	10/4	50	4	32	4	4	KVM
ai1s.4xlarge.4	16	64	15/8	100	8	64	8	8	KVM
ai1s.8xlarge.4	32	128	25/15	200	16	128	8	8	KVM
ai1s.6xlarge.4	24	96	25/14	400	4	96	8	8	KVM
ai1s.9xlarge.4	36	144	30/18	550	4	144	16	8	KVM
ai1s.18xlarge.4	72	288	40/36	1000	8	288	32	8	KVM

Features

Ai1s ECSs have the following features:

- vCPU to memory ratio: 1:4
- CPU: 2nd Generation Intel® Xeon® Scalable 6278 processors (2.6 GHz of base frequency and 3.5 GHz of turbo frequency), or Intel® Xeon® Scalable 6151 processors (3.0 GHz of base frequency and 3.4 GHz of turbo frequency)
- Ascend 310 processors, four of which in an Atlas300I accelerator card
- 16 TeraOPS of integer-precision computing (INT8) on one processor
- 8 GiB of GPU memory with a memory bandwidth of 50 GiB/s on one processor
- 5-channel HD video decoder (H.264/H.265) based on built-in hardware video codec engine

Notes

1. Ai1s ECSs support the following public images:
 - Ubuntu Server 18.04 64bit
 - CentOS 7.6 64bit
2. Ai1s ECSs do not support modification of specifications.
3. Ai1s ECSs support automatic recovery when the hosts accommodating such ECSs become faulty.

Using an AI-accelerated ECS

Perform the following steps:

1. Create an ECS. For details, see "Step 1: Configure Basic Settings" in the *Elastic Cloud Server User Guide*.
 - In the **Specifications** field, select AI-accelerated specifications.
 - In the **Image** field, select **Public image** or **Private image**.
 - **Public image:** The CANN 3.1.0 development kit has been included and environment variables have been configured in public images by default. You need to verify the environment availability.
 - **Private image:** You need to install the driver, firmware, and development kit, and configure environment variables by yourself. For details, see the *CANN Software Installation Guide* of the corresponding version in [Ascend Documentation](#).
2. Remotely log in to the ECS.

If your Ai1 ECS runs Linux, use an SSH password to log in to the ECS.
3. Verify the environment availability.

Use a sample for compilation and running. For details, see "Sample Overview" in the *Model Development Learning Map* of the corresponding CANN edition in [Ascend Documentation](#).

The sample shows how to classify images (decode, resize, and infer images) based on the Caffe ResNet-50 network.

1.8 Kunpeng ECS Specifications and Types

1.8.1 Kunpeng General Computing ECSs

Overview

Kunpeng general computing ECSs use Kunpeng 920 processors and 25GE high-speed intelligent NICs to cost-effectively provide a baseline level of vCPU performance with the ability to burst above the baseline, meeting the requirements of migrating infrastructure services to the cloud.

Specifications

Table 1-48 ks1 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
ks1.large.1	2	2	1.5/0.3	15	1	1	KVM
ks1.xlarge.1	4	4	2/0.5	20	1	2	KVM
ks1.2xlarge.1	8	8	3.5/1	35	2	2	KVM
ks1.3xlarge.1	12	12	5/1.5	45	2	2	KVM
ks1.4xlarge.1	16	16	7/2	60	2	2	KVM
ks1.large.2	2	4	1.5/0.3	15	1	1	KVM
ks1.xlarge.2	4	8	2/0.5	20	1	2	KVM
ks1.2xlarge.2	8	16	3.5/1	35	2	2	KVM
ks1.3xlarge.2	12	24	5/1.5	45	2	2	KVM
ks1.4xlarge.2	16	32	7/2	60	2	2	KVM
ks1.large.4	2	8	1.5/0.3	15	1	1	KVM
ks1.xlarge.4	4	16	2/0.5	20	1	2	KVM
ks1.2xlarge.4	8	32	3.5/1	35	2	2	KVM
ks1.3xlarge.4	12	48	5/1.5	45	2	2	KVM
ks1.4xlarge.4	16	64	7/2	60	2	2	KVM

Table 1-49 kS1s ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
ks1s.large.1	2	2	1.5/0.3	15	1	1	KVM
ks1s.xlarge.1	4	4	2/0.5	20	1	2	KVM
ks1s.2xlarge.1	8	8	3.5/1	35	2	2	KVM
ks1s.3xlarge.1	12	12	5/1.5	45	2	2	KVM
ks1s.4xlarge.1	16	16	7/2	60	2	2	KVM
ks1s.large.2	2	4	1.5/0.3	15	1	1	KVM
ks1s.xlarge.2	4	8	2/0.5	20	1	2	KVM
ks1s.2xlarge.2	8	16	3.5/1	35	2	2	KVM
ks1s.3xlarge.2	12	24	5/1.5	45	2	2	KVM
ks1s.4xlarge.2	16	32	7/2	60	2	2	KVM
ks1s.large.4	2	8	1.5/0.3	15	1	1	KVM
ks1s.xlarge.4	4	16	2/0.5	20	1	2	KVM
ks1s.2xlarge.4	8	32	3.5/1	35	2	2	KVM
ks1s.3xlarge.4	12	48	5/1.5	45	2	2	KVM
ks1s.4xlarge.4	16	64	7/2	60	2	2	KVM

Scenarios

Kunpeng general computing ECSs are suitable for applications that require moderate CPU performance generally but burstable high performance

occasionally, such as light-load web servers, enterprise R&D and testing environments, and low- and medium-performance databases.

1.8.2 Kunpeng General Computing-plus ECSs

Overview

Kunpeng general computing-plus ECSs use Kunpeng processors to provide powerful compute and high-performance networks, meeting enterprise requirements for cost-effective, secure, reliable cloud services.

Available now: kC1, kC1s, and kC1m

Table 1-50 Kunpeng general computing-plus ECS features

Type	Compute	Disk Type	Network
kC1	<ul style="list-style-type: none">vCPU to memory ratio: 1:1, 1:2, or 1:4Number of vCPUs: 1 to 60Kunpeng 920 processorBase frequency: 2.6 GHz	<ul style="list-style-type: none">High I/OUltra-high I/O	<ul style="list-style-type: none">Ultra-high packets per second (PPS) throughputAn ECS with higher specifications has better network performance.Maximum PPS: 4,000,000Maximum intranet bandwidth: 30 Gbit/s
kC1s	<ul style="list-style-type: none">vCPU to memory ratio: 1:2 or 1:4Number of vCPUs: 1 to 48Kunpeng 920 processorBase frequency: 2.6 GHz	<ul style="list-style-type: none">Ultra-high I/OHigh I/O	<ul style="list-style-type: none">Ultra-high PPS throughputAn ECS with higher specifications has better network performance.Maximum PPS: 4,000,000Maximum intranet bandwidth: 18 Gbit/s
kC1m	<ul style="list-style-type: none">vCPU to memory ratio: 1:2 or 1:4Number of vCPUs: 2 to 48Kunpeng 920 processorBase frequency: 2.6 GHz	<ul style="list-style-type: none">Ultra-high I/OHigh I/O	<ul style="list-style-type: none">Ultra-high PPS throughputAn ECS with higher specifications has better network performance.Maximum PPS: 3,500,000Maximum intranet bandwidth: 25 Gbit/s

Specifications

Table 1-51 kc1 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
kc1.small.1	1	1	2/0.5	20	1	2	KVM
kc1.large.2	2	4	3/0.8	30	2	2	KVM
kc1.xlarge.2	4	8	5/1.5	50	2	3	KVM
kc1.2xlarge.2	8	16	7/3	80	4	4	KVM
kc1.3xlarge.2	12	24	9/4.5	110	4	5	KVM
kc1.4xlarge.2	16	32	12/6	140	4	6	KVM
kc1.6xlarge.2	24	48	15/8.5	200	8	6	KVM
kc1.8xlarge.2	32	64	18/10	260	8	6	KVM
kc1.12xlarge.2	48	96	25/16	350	16	6	KVM
kc1.15xlarge.2	60	120	30/20	400	16	6	KVM
kc1.large.4	2	8	3/0.8	30	2	2	KVM
kc1.xlarge.4	4	16	5/1.5	50	2	3	KVM
kc1.2xlarge.4	8	32	7/3	80	4	4	KVM
kc1.3xlarge.4	12	48	9/4.5	110	4	5	KVM
kc1.4xlarge.4	16	64	12/6	140	4	6	KVM
kc1.6xlarge.4	24	96	15/8.5	200	8	6	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
kc1.8xlarge.4	32	128	18/10	260	8	6	KVM
kc1.12xlarge.4	48	192	25/16	350	16	6	KVM

Table 1-52 kc1s ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
kc1s.large.2	2	4	2/0.4	25	2	2	KVM
kc1s.xlarge.2	4	8	3/0.8	40	2	3	KVM
kc1s.2xlarge.2	8	16	4/1.5	65	4	4	KVM
kc1s.3xlarge.2	12	24	5/2	90	4	5	KVM
kc1s.4xlarge.2	16	32	7/3	120	4	6	KVM
kc1s.6xlarge.2	24	48	8/4	160	8	6	KVM
kc1s.8xlarge.2	32	64	10/5	200	8	6	KVM
kc1s.12xlarge.2	48	96	15/8	280	16	6	KVM
kc1s.15xlarge.2	60	120	18/10	320	16	6	KVM
kc1s.large.4	2	8	2/0.4	25	2	2	KVM
kc1s.xlarge.4	4	16	3/0.8	40	2	3	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
kc1s.2xlarge.4	8	32	4/1.5	65	4	4	KVM
kc1s.3xlarge.4	12	48	5/2	90	4	5	KVM
kc1s.4xlarge.4	16	64	7/3	120	4	6	KVM
kc1s.6xlarge.4	24	96	8/4	160	8	6	KVM
kc1s.8xlarge.4	32	128	10/5	200	8	6	KVM
kc1s.12xlarge.4	48	192	15/8	280	16	6	KVM

Table 1-53 kc1m ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
kc1m.large.2	2	4	3/0.8	30	2	2	KVM
kc1m.xlarge.2	4	8	5/1.5	50	2	3	KVM
kc1m.2xlarge.2	8	16	7/3	80	4	4	KVM
kc1m.3xlarge.2	12	24	9/4.5	110	4	5	KVM
kc1m.4xlarge.2	16	32	12/6	140	4	6	KVM
kc1m.6xlarge.2	24	48	15/8.5	200	8	6	KVM
kc1m.8xlarge.2	32	64	18/10	260	8	6	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
kc1m.12xlarge.2	48	96	25/16	350	16	6	KVM
kc1m.large.4	2	8	3/0.8	30	2	2	KVM
kc1m.xlarge.4	4	16	5/1.5	50	2	3	KVM
kc1m.2xlarge.4	8	32	7/3	80	4	4	KVM
kc1m.3xlarge.4	12	48	9/4.5	110	4	5	KVM
kc1m.4xlarge.4	16	64	12/6	140	4	6	KVM
kc1m.6xlarge.4	24	96	15/8.5	200	8	6	KVM
kc1m.8xlarge.4	32	128	18/10	260	8	6	KVM
kc1m.12xlarge.4	48	192	25/16	350	16	6	KVM

Scenarios

Kunpeng general computing-plus ECSs are suitable for:

- Governments, enterprises, and the financial industry with strict requirements on security and privacy
- Internet applications with high requirements on network performance
- Big data and HPC with requirements on a large number of vCPUs
- Website setups and e-Commerce requiring cost-effectiveness

1.8.3 Kunpeng Memory-optimized ECSs

Overview

Kunpeng memory-optimized ECSs use Kunpeng 920 processors and 25GE high-speed intelligent NICs to provide up to 480 GiB DDR4-based memory with high network performance for large in-memory datasets.

Available now: kM1, kM1s, and kM1m

Table 1-54 Kunpeng memory-optimized ECS features

Type	Compute	Disk Type	Network
kM1	<ul style="list-style-type: none">vCPU to memory ratio: 1:8Number of vCPUs: 2 to 60Kunpeng 920 processorBase frequency: 2.6 GHz	<ul style="list-style-type: none">High I/OGeneral Purpose SSDUltra-high I/O	<ul style="list-style-type: none">Ultra-high packets per second (PPS) throughputAn ECS with higher specifications has better network performance.Maximum PPS: 4,000,000Maximum intranet bandwidth: 30 Gbit/s
kM1s	<ul style="list-style-type: none">vCPU to memory ratio: 1:8Number of vCPUs: 2 to 60Kunpeng 920 processorBase frequency: 2.6 GHz	<ul style="list-style-type: none">Ultra-high I/OHigh I/O	<ul style="list-style-type: none">Ultra-high PPS throughputAn ECS with higher specifications has better network performance.Maximum PPS: 4,000,000Maximum intranet bandwidth: 18 Gbit/s
kM1m	<ul style="list-style-type: none">vCPU to memory ratio: 1:8Number of vCPUs: 2 to 48Kunpeng 920 processorBase frequency: 2.6 GHz	<ul style="list-style-type: none">Ultra-high I/OHigh I/O	<ul style="list-style-type: none">Ultra-high PPS throughputAn ECS with higher specifications has better network performance.Maximum PPS: 3,500,000Maximum intranet bandwidth: 25 Gbit/s

Specifications

Table 1-55 km1 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
km1.large.8	2	16	3/0.8	30	2	2	KVM
km1.xlarge.8	4	32	5/1.5	50	2	3	KVM
km1.2xlarge.8	8	64	7/3	80	4	4	KVM
km1.3xlarge.8	12	96	9/4.5	110	4	5	KVM
km1.4xlarge.8	16	128	12/6	140	4	6	KVM
km1.6xlarge.8	24	192	15/8	200	8	6	KVM
km1.8xlarge.8	32	256	18/10	260	8	6	KVM
km1.12xlarge.8	48	384	25/16	350	16	8	KVM
km1.15xlarge.8	60	480	30/20	400	16	8	KVM

Table 1-56 km1s ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
km1s.large.8	2	16	2/0.4	30	2	2	KVM
km1s.xlarge.8	4	32	3/0.8	50	2	3	KVM
km1s.2xlarge.8	8	64	4/1.5	80	4	4	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
km1s.3xlarge.8	12	96	5/2	110	4	5	KVM
km1s.4xlarge.8	16	128	7/3	140	4	6	KVM
km1s.6xlarge.8	24	192	8/4	200	8	6	KVM
km1s.8xlarge.8	32	256	10/5	260	8	6	KVM
km1s.12xlarge.8	48	384	15/8	350	16	8	KVM
km1s.15xlarge.8	60	480	18/10	400	16	8	KVM

Table 1-57 km1m ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
km1m.large.8	2	16	3/0.8	30	2	2	KVM
km1m.xlarge.8	4	32	5/1.5	50	2	3	KVM
km1m.2xlarge.8	8	64	7/3	80	4	4	KVM
km1m.3xlarge.8	12	96	9/4.5	110	4	5	KVM
km1m.4xlarge.8	16	128	12/6	140	4	6	KVM
km1m.6xlarge.8	24	192	15/8.5	200	8	6	KVM
km1m.8xlarge.8	32	256	18/10	260	8	6	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
km1m.12xlarge.8	48	384	25/16	350	16	8	KVM

Scenarios

Big data analysis, precision marketing, e-commerce, IoT, and in-memory storage (such as Memcache)

1.8.4 Kunpeng Ultra-high I/O ECSs

Overview

Kunpeng ultra-high I/O ECSs use Kunpeng 920 processors and 25GE high-speed intelligent NICs to provide up to 480 GiB DDR4-based memory with high network performance for large in-memory datasets.

Available now: k11 and k11s

Table 1-58 Kunpeng ultra-high I/O ECS features

Type	Compute	Disk Type	Network
k11s	<ul style="list-style-type: none">vCPU to memory ratio: 1:4Number of vCPUs: 8 to 64Kunpeng 920 processorBase frequency: 2.6 GHz	<ul style="list-style-type: none">Ultra-high I/OHigh I/O	<ul style="list-style-type: none">Ultra-high packets per second (PPS) throughputAn ECS with higher specifications has better network performance.Maximum PPS: 3,200,000Maximum intranet bandwidth: 18 Gbit/s

Type	Compute	Disk Type	Network
ki1	<ul style="list-style-type: none"> vCPU to memory ratio: 1:4 Number of vCPUs: 8 to 64 Kunpeng 920 processor Base frequency: 2.6 GHz 	<ul style="list-style-type: none"> High I/O Ultra-high I/O 	<ul style="list-style-type: none"> Ultra-high packets per second (PPS) throughput An ECS with higher specifications has better network performance. Maximum PPS: 4,000,000 Maximum intranet bandwidth: 30 Gbit/s

Specifications

Table 1-59 ki1s ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NICs	Max. NIC Queues	Local Disks	Virtualization
ki1s.2xlarge.4	8	32	4/1.5	65	4	4	1 × 3200GiB	KVM
ki1s.4xlarge.4	16	64	7/3	120	6	4	2 × 3200GiB	KVM
ki1s.6xlarge.4	24	96	8/4	160	6	8	3 × 3200GiB	KVM
ki1s.8xlarge.4	32	128	10/5	200	6	8	4 × 3200GiB	KVM
ki1s.12xlarge.4	48	192	15/8	280	6	16	6 × 3200GiB	KVM
ki1s.16xlarge.4	64	228	18/10	320	6	16	8 × 3200GiB	KVM

Table 1-60 k1 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NICs	Max. NIC Queues	Local Disks	Virtualization
ki1.2xlarge.4	8	32	7/3	80	4	4	1 × 3,200 GiB	KVM
ki1.4xlarge.4	16	64	12/6	140	6	4	2 × 3,200 GiB	KVM
ki1.6xlarge.4	24	96	15/8.5	200	6	8	3 × 3,200 GiB	KVM
ki1.8xlarge.4	32	128	18/10	260	6	8	4 × 3,200 GiB	KVM
ki1.12xlarge.4	48	192	25/16	350	6	16	6 × 3,200 GiB	KVM
ki1.16xlarge.4	64	228	30/20	400	6	16	8 × 3,200 GiB	KVM

Features

[Table 1-61](#) and [Table 1-62](#) lists the IOPS performance of k1s ECSs and specifications of a single local disk attached to a k1s ECS.

Table 1-61 k1s ECS IOPS performance

Flavor	Maximum IOPS for Random 4 KB Read
ki1s.2xlarge.4	750,000
ki1s.4xlarge.4	1,500,000
ki1s.6xlarge.4	2,250,000
ki1s.8xlarge.4	3,000,000
ki1s.12xlarge.4	4,500,000
ki1s.16xlarge.4	6,000,000

Table 1-62 Specifications of a single NVMe disk attached to a k1s ECS

Metric	Performance
Disk capacity	3.2 TB
IOPS for random 4 KB read	750,000
IOPS for random 4 KB write	200,000
Read throughput	2.9 GiB/s
Write throughput	1.9 GiB/s
Access latency	Within microseconds

Table 1-63 and **Table 1-64** lists the IOPS performance of k1 ECSs and specifications of a single local disk attached to a k1 ECS.

Table 1-63 k1 ECS IOPS performance

Flavor	Maximum IOPS for Random 4 KB Read
ki1.2xlarge.4	750,000
ki1.4xlarge.4	1,500,000
ki1.6xlarge.4	2,250,000
ki1.8xlarge.4	3,000,000
ki1.12xlarge.4	4,500,000
ki1.16xlarge.4	6,000,000

Table 1-64 Specifications of a single NVMe disk attached to a k1 ECS

Metric	Performance
Disk capacity	3.2 TB
IOPS for random 4 KB read	750,000
IOPS for random 4 KB write	200,000
Read throughput	2.9 GiB/s
Write throughput	1.9 GiB/s
Access latency	Within microseconds

Notes

- Kunpeng ultra-high I/O ECSs do not support specifications modification.
- Kunpeng ultra-high I/O ECSs do not support local disk snapshots or backups.
- Kunpeng ultra-high I/O ECSs can use both local disks and EVS disks to store data. In addition, they can have EVS disks attached to provide a larger storage size. Use restrictions on the two types of storage media are as follows:
 - Only an EVS disk, not a local disk, can be used as the system disk of a Kunpeng ultra-high I/O ECS.
 - Both EVS disks and local disks can be used as data disks of a Kunpeng ultra-high I/O ECS.
 - A Kunpeng ultra-high I/O ECS can have a maximum of 60 attached disks (including VBD, SCSI, and local disks). Among the 60 disks, the maximum number of SCSI disks is 30, and the maximum number of VBD disks is 22 (including the system disk).
 - It is a good practice to use World Wide Names (WWNs), but not drive letters, to perform operations on local disks to prevent drive letter drift (low probability) on Linux. Take local disk attachment as an example:
If the local disk WWN is `wwn-0x50014ee2b14249f6`, run the **`mount /dev/disk/by-id/wwn-0x50014ee2b14249f6`** command.

NOTE

How can I view the local disk WWN?

1. Log in to the ECS.
2. Run the following command:

```
ll /dev/disk/by-id
```

- The local disk data of a Kunpeng ultra-high I/O ECS if an exception occurs, such as physical server breakdown or local disk damage. If your application does not use the data reliability architecture, it is a good practice to use EVS disks to build your ECS.
- When a Kunpeng ultra-high I/O ECS is deleted, the data on local NVMe SSDs will also be automatically deleted, which can take some time. As a result, a Kunpeng ultra-high I/O ECS takes a longer time than other ECSs to be deleted. Back up the data before deleting such an ECS.
- The data reliability of local disks depends on the reliability of physical servers and hard disks, which are SPOF-prone. It is a good practice to use data redundancy mechanisms at the application layer to ensure data availability. Use EVS disks to store service data that needs to be stored for a long time.
- The device name of a local disk attached to a Kunpeng ultra-high I/O ECS is `/dev/nvme0n1` or `/dev/nvme0n2`.
- The basic resources, including vCPUs, memory, and image of a Kunpeng ultra-high I/O ECS will continue to be billed after the ECS is stopped. To stop the ECS from being billed, delete it and its associated resources.

Scenarios

Kunpeng ultra-high I/O ECSs can be used for high-performance relational databases, NoSQL databases (such as Cassandra and MongoDB), and ElasticSearch.

1.8.5 Kunpeng AI Inference-accelerated ECSs

Kunpeng AI inference-accelerated ECSs are designed to provide acceleration services for AI services. These ECSs are provided with the Ascend AI Processors and Ascend AI Software Stack.

Kunpeng AI inference-accelerated ECSs use Ascend 310 processors for AI inference acceleration.

Table 1-65 Kunpeng AI Inference-accelerated ECSs

Type	Compute	Disk Type	Network
kAi1s	<ul style="list-style-type: none">vCPU to memory ratio: 1:1 or 1:2Number of vCPUs: 4 to 48Kunpeng 920 processorBase frequency: 2.6 GHz	<ul style="list-style-type: none">High I/OGeneral Purpose SSDUltra-high I/OExtreme SSDGeneral Purpose SSD V2	<ul style="list-style-type: none">Ultra-high packets per second (PPS) throughputAn ECS with higher specifications has better network performance.Maximum PPS: 2,000,000Maximum intranet bandwidth: 12 Gbit/s
kAi2	<ul style="list-style-type: none">vCPU to memory ratio: 1:4Number of vCPUs: 16 to 96Kunpeng 920 processorBase frequency: 2.6 GHz	<ul style="list-style-type: none">High I/OGeneral Purpose SSDUltra-high I/OExtreme SSDGeneral Purpose SSD V2	<ul style="list-style-type: none">Ultra-high packets per second (PPS) throughputAn ECS with higher specifications has better network performance.Maximum PPS: 7,000,000Maximum intranet bandwidth: 38 Gbit/s

Kunpeng Enhanced AI Inference-accelerated kAi1s (Type I)

Overview

Kunpeng AI inference-accelerated kAi1s ECSs use Ascend 310 processors for AI acceleration. Ascend 310 processors feature low power consumption, high computing capabilities, and significantly improved energy efficiency ratio (EER),

facilitating the wide application of AI inference. kAi1s ECSs deliver the computing acceleration capabilities of the Ascend 310 processors on the cloud platform. This helps you quickly and simply use the Ascend 310 processors.

kAi1s ECSs are based on Atlas 300I accelerator cards. For details, see [Ascend Community](#).

kAi1s ECSs are used for general technologies, such as computer vision, speech recognition, and natural language processing to support smart retail, smart campus, robot cloud brain, and safe city scenarios.

Specifications

Table 1-66 kAi1s ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Ascend 310 Processors	Virtualization
kai1s.xlarge.1	4	4	3/0.8	20	2	2	1	KVM
kai1s.2xlarge.1	8	8	4/1.5	40	2	3	2	KVM
kai1s.4xlarge.1	16	16	6/3	80	4	4	4	KVM
kai1s.3xlarge.2	12	24	8/4	100	4	4	4	KVM
kai1s.4xlarge.2	16	32	10/6	140	4	5	6	KVM
kai1s.6xlarge.2	24	48	12/8	200	8	6	8	KVM
kai1s.9xlarge.2	36	72	12/8	200	8	6	12	KVM
kai1s.12xlarge.2	48	96	12/8	200	16	6	12	KVM

Features

kAi1s ECSs have the following features:

- 1:1 or 1:2 ratio of vCPUs to memory
- CPU: Kunpeng 920 (2.6 GHz)
- Ascend 310 processors, four of which in an Atlas 300I accelerator card
- 8 TeraFLOPS of half-precision computing (FP16) on one processor
- 16 TeraOPS of integer-precision computing (INT8) on one processor
- 8 GiB of GPU memory with a memory bandwidth of 50 GiB/s on one processor
- Built-in hardware video codec engine, supporting H.264/H.265

Notes

- kAi1s ECSs support the following OSs:
 - Ubuntu Server 18.04 64bit
 - CentOS 7.6 64-bit
- kAi1s ECSs do not support modification of specifications.
- kAi1s ECSs support automatic recovery when the hosts accommodating such ECSs become faulty.

Kunpeng Enhanced AI Inference-accelerated kAi2 (Type II)

Overview

Kunpeng enhanced AI inference-accelerated kAi2 ECSs use Ascend 310P processors for AI acceleration. The Atlas 300I Pro accelerator card with the 24 GiB of GPU memory and Ascend 310P processor helps boost your AI inference services.

Specifications

Table 1-67 kAi2 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Ascend 310P	GPU Memory (GiB)	Virtualization
kai2.4xlarge.4	16	64	15/5.5	200	8	8	1	1 × 24	KVM
kai2.8xlarge.4	32	128	24/8	250	8	8	2	2 × 24	KVM
kai2.16xlarge.4	64	256	30/16	350	16	8	4	4 × 24	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Ascend 310P	GPU Memory (GiB)	Virtualization
kai2.2 Oxlarge.4	80	320	38/30	700	32	16	5	5 × 24	KVM
kai2.2 4xlarge.4	96	384	38/30	700	32	16	5	5 × 24	KVM

Features

kAi2 ECSs have the following features:

- 1:4 ratio of vCPUs to memory
- CPU: Kunpeng 920 (2.6 GHz)
- Ascend 310P processors, one of which in an Atlas 300I Pro accelerator card
Each processor with eight DaVinci AI cores and eight self-developed CPU cores
- 70 TeraFLOPS of half-precision computing (FP16) on one processor
- 140 TeraOPS of integer-precision computing (INT8) on one processor
- 24 GiB of GPU memory with a memory bandwidth of 240.8 GiB/s on one processor
- Built-in hardware video codec engine, supporting H.264/H.265
- JPEG image codec

Notes

- kAi2 ECSs support the following OSs:
 - HCE 2.0 64bit
 - CentOS 7.6 64bit
 - Ubuntu 18.04 server 64bit
- kAi2 ECSs support automatic recovery when the hosts accommodating such ECSs become faulty.
- kAi2 ECSs support specifications modification but do not support the instance type change.

Using a kAi ECS

Perform the following steps:

1. Create an ECS. For details, see "Step 1: Configure Basic Settings" in the *Elastic Cloud Server User Guide*.
 - In the **Specifications** field, select kAi-accelerated specifications.
 - In the **Image** field, select **Public image** or **Private image**.

- **Public image:** The CANN 3.1.0 development kit has been included and environment variables have been configured in public images by default. You need to verify the environment availability.
 - **Private image:** You need to install the driver, firmware, and development kit, and configure environment variables by yourself. For details, see the *CANN Software Installation Guide* of the corresponding version in [Ascend Documentation](#).
2. Remotely log in to the ECS.
If your ECS runs Linux, use an SSH password to log in to the ECS.
 3. Verify the environment availability.
Use a sample for compilation and running. For details, see "Sample Overview" in the *Model Development Learning Map* of the corresponding CANN edition in [Ascend Documentation](#).
The sample shows how to classify images (decode, resize, and infer images) based on the Caffe ResNet-50 network.

1.9 Images

1.9.1 Image Types

What Is Image?

An image is an ECS template that contains an OS. It may also contain proprietary software and application software, such as database software. You can use images to create ECSs.

Images can be public or private. Public images are provided by the system by default, and private images are manually created. You can use any type of image to create an ECS. You can also create a private image using an existing ECS or external image. This provides you with a simple and fast way to create ECSs tailored to your needs. For example, if you use web services, your image can contain web server configurations, static configurations, and dynamic page code. After you use this image to create an ECS, the web server will run on the created ECS.

Image Types

Image Type	Description
Public image	A public image is a standard, widely used image. It contains an OS and preinstalled public applications and is available to all users. Public images are very stable and their OS and any included software have been officially authorized for use. If a public image does not contain the environments or software you need, you can use a public image to create an ECS and then deploy the required environments or software on it.

Image Type	Description
Private image	<p>A private image contains an OS or service data, preinstalled public applications, and a user's personal applications. Private images are only available to the users who created them.</p> <p>A private image can be a system disk image, data disk image, ISO image, or full-ECS image.</p> <ul style="list-style-type: none">• A system disk image contains an OS and preinstalled software for various services. You can use a system disk image to create ECSs and migrate your services to the cloud.• A data disk image contains only service data. You can use a data disk image to create EVS disks and use them to migrate your service data to the cloud.• An ISO image is created from an external ISO image file. It is a special image that is not available on the ECS console.• A full-ECS image contains an OS, preinstalled software, and service data. A full-ECS image is created using differential backups and the creation takes less time than creating a system or data disk image that has the same disk capacity.
Shared image	<p>A shared image is a private image another user has shared with you.</p> <p>For more information about shared images, see "Sharing Images" in <i>Image Management Service User Guide</i>.</p>

1.9.2 Cloud-Init

Cloud-Init is an open-source cloud initialization program, which initializes some of the customized configurations of a newly created ECS, such as the hostname and user data.

Using Cloud-Init to initialize your ECSs will affect your ECS, IMS, and AS services.

Impact on IMS

To ensure that ECSs created using a private image support custom configurations, you must install Cloud-Init or Cloudbase-Init on the ECSs before using them to create private images.

- For Windows OSs, download and install Cloudbase-Init.
- For Linux OSs, download and install Cloud-Init.

After being installed in an image, Cloud-Init or Cloudbase-Init automatically configures initial attributes for the ECSs created using this image.

For more information, see *Image Management Service User Guide*.

Impact on ECS

- When creating an ECS, if the selected image supports Cloud-Init, you can use the **User Data** function to specify custom configuration, such as ECS login

password to the ECS. Such custom settings will take effect upon ECS initialization.

- If Cloud-Init is supported, ECS login modes will be changed.
- If Cloud-Init is supported, you can view and use metadata to configure and manage running ECSs.

Impact on AS

- When creating an AS configuration, you can use the **User Data** function to specify ECS configurations for initialization. If the AS configuration has taken effect in an AS group, the ECSs newly created in the AS group will automatically initialize their configurations based on the specified ECS configurations.
- For an existing AS configuration, if its private image does not have Cloud-Init or Cloudbase-Init installed, the login mode of the ECSs created in the AS group where the AS configuration takes effect may fail to take effect. To resolve this issue, see "How Does Cloud-Init Affect the AS Service?" in *Auto Scaling User Guide*.

Notes

- When using Cloud-Init, enable DHCP in the VPC to which the ECS belongs.
- When using Cloud-Init, ensure that security group rules for the outbound direction meet the following requirements:
 - **Protocol: TCP**
 - **Port: 80**
 - **Destination: 169.254.0.0/16**

NOTE

If you use the default security group rules for the outbound direction, the metadata can be accessed because the default rules meet the preceding requirements. For details about the default security group rules for the outbound direction, see [Security Group](#).

1.10 EVS Disks

What Is Elastic Volume Service?

Elastic Volume Service (EVS) offers scalable block storage for ECSs. With high reliability, high performance, and rich specifications, EVS disks can be used for distributed file systems, development and test environments, data warehouses, and high-performance computing (HPC) scenarios to meet diverse service requirements.

Disk Types

EVS disk types differ in performance. Choose a disk type based on your requirements.

For more information about EVS disk specifications and performance, see *Elastic Volume Service User Guide*.

Device Types

EVS disks have two device types, Virtual Block Device (VBD) and Small Computer System Interface (SCSI).

- VBD

When you create an EVS disk on the management console, **Device Type** of the EVS disk is VBD by default. VBD EVS disks support only simple SCSI read/write commands.

- SCSI

You can create EVS disks whose **Device Type** is SCSI on the management console. These EVS disks support transparent SCSI command transmission, allowing ECS OS to directly access underlying storage media. SCSI EVS disks support both basic and advanced SCSI commands.

NOTE

For more information about how to use SCSI EVS disks, for example, how to install a driver for SCSI EVS disks, see "Device Types and Usage Instructions" in *Elastic Volume Service User Guide*.

1.11 Network

VPC

Virtual Private Cloud (VPC) allows you to create customized virtual networks in your logically isolated AZ. Such networks are dedicated zones that are logically isolated, providing secure network environments for your ECSs. You can define security groups, virtual private networks (VPNs), IP address segments, and bandwidth for a VPC. This facilitates internal network configuration and management and allows you to change your network in a secure and convenient network manner. You can also customize the ECS access rules within a security group and between security groups to improve ECS security.

For more information about VPC, see *Virtual Private Cloud User Guide*.

Subnet

A subnet is a range of IP addresses in your VPC and provides IP address management and DNS resolution functions for ECSs in it. The IP addresses of all ECSs in a subnet belong to the subnet.

By default, ECSs in all subnets of the same VPC can communicate with each other, while ECSs in different VPCs cannot.

Security Group

A security group is a collection of access control rules for ECSs that have the same security protection requirements and that are mutually trusted. By adding an ECS to a security group, you apply all the rules defined for this security group to this ECS.

Your account automatically comes with a default security group. The default security group allows all outbound data, denies all inbound data, and allows all

data between ECSs in the group. Your ECSs in the security group can communicate with each other without the need to add rules.

Figure 1-2 Default security group

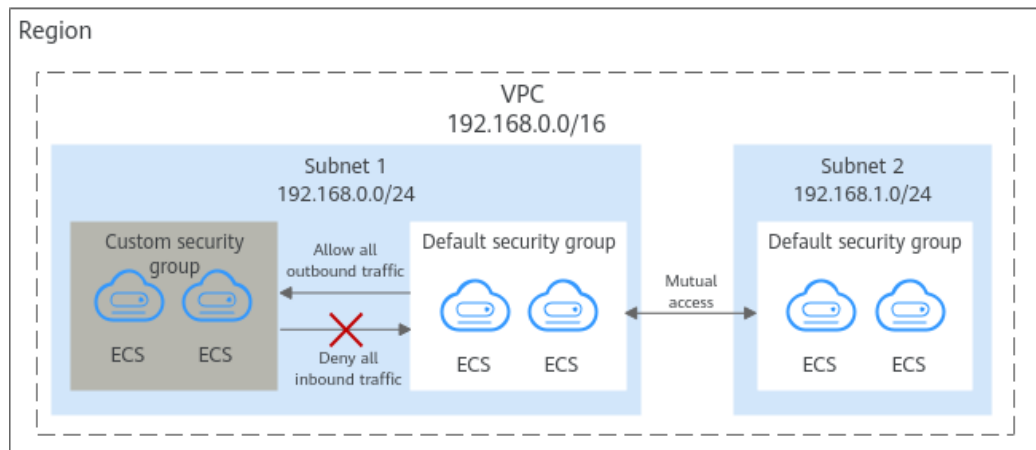


Table 1-68 describes default security group rules.

Table 1-68 Default security group rules

Direction	Protocol	Port/Range	Source/Destination	Description
Outbound	All	All	Destination: 0.0.0.0/0	Allows all outbound traffic.
Inbound	All	All	Source: the current security group (for example, sg-xxxxx)	Allows communications among ECSs within the security group and denies all inbound traffic (incoming data packets).

EIP

The Elastic IP (EIP) service enables your cloud resources to communicate with the Internet using static public IP addresses and scalable bandwidths. EIPs can be bound to or unbound from ECSs, BMSs, virtual IP addresses, or load balancers.

Each EIP can be used by only one cloud resource at a time.

1.12 Security

1.12.1 Fault Recovery

Cloud Backup and Recovery (CBR) lets you back up and restore data in case of a failure. If an ECS or EVS disk is faulty or data is deleted accidentally, you can use data backups to quickly restore data.

What Is CBR?

CBR enables you to back up ECSs and EVS disks with ease. If any exceptions occur, such as virus attack, accidental deletion, or software or hardware fault, you can restore data to any point in the past when the data was backed up.

CBR secures your services by ensuring the security and consistency of your data.

Differences Between Cloud Server Backup and Cloud Disk Backup

You can back up ECS data using Cloud Server Backup or Cloud Disk Backup.

- **Cloud Server Backup (recommended):** Use this backup function if you want to back up the data of all EVS disks (system and data disks) on an ECS. This prevents data inconsistency caused by time difference in creating a backup.
- **Cloud Disk Backup:** Use this backup function if you want to back up the data of one or more EVS disks (system or data disk) on an ECS. This minimizes backup costs on the basis of data security.

Table 1-69 Differences between cloud server backup and cloud disk backup

Item	Cloud Server Backup	Cloud Disk Backup
Resources to be backed up or restored	All disks (system and data disks) on a server	One or more specified disks (system or data disks)
Recommended scenario	An entire cloud server needs to be protected.	Only data disks need to be backed up, because the system disk does not contain users' application data.
Advantages	All disks on a server are backed up at the same time, ensuring data consistency.	Backup cost is reduced without compromising data security.

1.13 Permissions Management

If you need to assign different permissions to employees in your enterprise to access your ECS resources, IAM is a good choice for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you secure access to your resources.

With IAM, you can use your account to create IAM users, and assign permissions to the users to control their access to specific resources. For example, some software developers in your enterprise need to use ECS resources but should not be allowed to delete the resources or perform any other high-risk operations. In this scenario, you can create IAM users for the software developers and grant them only the permissions required for using ECS resources.

If your account does not need individual IAM users for permissions management, skip this section.

IAM is a free service. You pay only for the resources in your account. For more information about IAM, see *Identity and Access Management Service Overview*.

ECS Permissions

By default, new IAM users do not have permissions assigned. You need to add a user to one or more groups, and attach permissions policies or roles to these groups. Users inherit permissions from the groups to which they are added and can perform specified operations on cloud services based on the permissions.

ECS is a project-level service deployed and accessed in specific physical regions. To assign ECS permissions to a user group, specify the scope as region-specific projects and select projects for the permissions to take effect. If you select **All projects**, the permissions will take effect for user groups in all region-specific projects. When accessing ECS, the users need to switch to a region where they have got permissions to use this service.

You can grant users permissions by using roles and policies.

- **Roles:** A type of coarse-grained authorization mechanism that defines permissions related to user responsibilities. This mechanism provides only a limited number of service-level roles for authorization. When using roles to grant permissions, you need to also assign other roles which the permissions depend on to take effect. However, roles are not an ideal choice for fine-grained authorization and secure access control.
- **Policies:** A fine-grained authorization strategy that defines permissions required to perform operations on specific cloud resources under certain conditions. This mechanism allows for more flexible policy-based authorization, meeting requirements for secure access control. For example, you can grant ECS users only the permissions for managing a certain type of ECSs.

Most policies define permissions based on APIs. For the API actions supported by ECS, see "Permissions Policies and Supported Actions" in *Elastic Cloud Server API Reference*.

Table 1-70 lists all the system policies supported by ECS.

Table 1-70 System-defined permissions for ECS

Policy/Role Name	Description	Type	Policy Content
ECS FullAccess	Administrator permissions for ECS. Users granted these permissions can perform all operations on ECSs, including creating, deleting, and viewing ECSs, and modifying ECS specifications.	System-defined policy	ECS FullAccess Policy Content
ECS CommonOperations	Common user permissions for ECS. Users granted these permissions can start, stop, restart, and query ECSs.	System-defined policy	ECS CommonOperations Policy Content

Policy/Role Name	Description	Type	Policy Content
ECS ReadOnlyAccess	Read-only permissions for ECS. Users granted these permissions can only view ECS data.	System-defined policy	ECS ReadOnlyAccess Policy Content
Server Administrator	<p>Full permissions for ECS. This role must be used together with the Tenant Guest role in the same project.</p> <p>If a user needs to create, delete, or change resources of other services, the user must also be granted administrator permissions of the corresponding services in the same project.</p> <p>For example, if a user needs to create a new VPC when creating an ECS, the user must also be granted permissions with the VPC Administrator role.</p>	System role	Server Administrator Policy Content

Table 1-71 lists the common operations supported by each system-defined policy of ECS. Select the policies as required.

Table 1-71 Common operations supported by each system-defined policy

Operation	ECS FullAccess	ECS CommonOperations	ECS ReadOnlyAccess
Creating an ECS	Supported	Not supported	Not supported
Remotely logging in to an ECS on the management console	Supported	Supported	Not supported (VNC login not supported)
Querying an ECS list	Supported	Supported	Supported
Querying ECS details	Supported	Supported	Supported
Modifying ECS details	Supported	Not supported	Not supported
Starting an ECS	Supported	Supported	Not supported
Stopping an ECS	Supported	Supported	Not supported
Restarting an ECS	Supported	Supported	Not supported
Deleting an ECS	Supported	Not supported	Not supported

Operation	ECS FullAccess	ECS CommonOperations	ECS ReadOnlyAccess
Reinstalling an ECS OS	Supported	Not supported	Not supported
Changing an ECS OS	Supported	Not supported	Not supported
Attaching a disk to an ECS	Supported	Not supported	Not supported
Detaching a disk from an ECS	Supported	Not supported	Not supported
Querying a disk list	Supported	Supported	Supported
Attaching a NIC to an ECS	Supported	Not supported	Not supported
Detaching a NIC from an ECS	Supported	Not supported	Not supported
Querying a NIC list	Supported	Supported	Supported
Modifying ECS specifications	Supported	Not supported	Not supported
Querying the ECS flavor list	Supported	Supported	Supported
Querying ECS groups	Supported	Supported	Supported

ECS FullAccess Policy Content

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ecs:*",
        "evs:*:get",
        "evs:*:list",
        "evs:volumes:create",
        "evs:volumes:delete",
        "evs:volumes:attach",
        "evs:volumes:detach",
        "evs:volumes:manage",
        "evs:volumes:update",
        "evs:volumes:use",
        "evs:volumes:uploadImage",
        "evs:snapshots:create",
        "vpc:*:get",
        "vpc:*:list",
        "vpc:networks:create",
        "vpc:networks:update",
        "vpc:subnets:update",
        "vpc:subnets:create",
        "vpc:ports:*",
        "vpc:routers:get",
        "vpc:routers:update",

```

```
        "vpc:securityGroups:*",
        "vpc:securityGroupRules:*",
        "vpc:floatingIps:*",
        "vpc:publicIps:*",
        "ims:images:create",
        "ims:images:delete",
        "ims:images:get",
        "ims:images:list",
        "ims:images:update",
        "ims:images:upload"
    ]
}
]
```

ECS CommonOperations Policy Content

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ecs:*:get*",
        "ecs:*:list*",
        "ecs:*:start",
        "ecs:*:stop",
        "ecs:*:reboot",
        "ecs:blockDevice:use",
        "ecs:cloudServerFpgaImages:relate",
        "ecs:cloudServerFpgaImages:register",
        "ecs:cloudServerFpgaImages:delete",
        "ecs:cloudServerFpgaImages:unrelate",
        "ecs:cloudServers:setAutoRecovery",
        "ecs:cloudServerPasswords:reset",
        "ecs:cloudServerPorts:modify",
        "ecs:cloudServers:vnc",
        "ecs:diskConfigs:use",
        "ecs:securityGroups:use",
        "ecs:serverGroups:manage",
        "ecs:serverFloatingIps:use",
        "ecs:serverKeyPairs:*",
        "ecs:serverPasswords:manage",
        "ecs:servers:createConsole",
        "ecs:servers:createImage",
        "ecs:servers:setMetadata",
        "ecs:servers:setTags",
        "ecs:serverVolumes:use",
        "evs:*:get*",
        "evs:*:list*",
        "evs:snapshots:create",
        "evs:volumes:uploadImage",
        "evs:volumes:delete",
        "evs:volumes:update",
        "evs:volumes:attach",
        "evs:volumes:detach",
        "evs:volumes:manage",
        "evs:volumes:use",
        "vpc:*:get*",
        "vpc:*:list*",
        "vpc:floatingIps:create",
        "vpc:floatingIps:update",
        "vpc:floatingIps:delete",
        "vpc:publicIps:update",
        "vpc:publicIps:delete",
        "ims:images:create",
        "ims:images:delete",
        "ims:images:get",
        "ims:images:list",

```



```
        "ims:images:update",  
        "ims:images:upload"  
    ]  
  }  
]  
}
```

ECS ReadOnlyAccess Policy Content

```
{  
  "Version": "1.1",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "ecs:*:get*",  
        "ecs:*:list*",  
        "ecs:serverGroups:manage",  
        "ecs:serverVolumes:use",  
        "evs:*:get*",  
        "evs:*:list*",  
        "vpc:*:get*",  
        "vpc:*:list*",  
        "ims:*:get*",  
        "ims:*:list*"  
      ]  
    }  
  ]  
}
```

Server Administrator Policy Content

```
{  
  "Version": "1.1",  
  "Statement": [  
    {  
      "Action": [  
        "ecs:*:*",  
        "evs:*:get",  
        "evs:*:list",  
        "evs:volumes:create",  
        "evs:volumes:delete",  
        "evs:volumes:attach",  
        "evs:volumes:detach",  
        "evs:volumes:manage",  
        "evs:volumes:update",  
        "evs:volumes:uploadImage",  
        "evs:snapshots:create",  
        "vpc:*:get",  
        "vpc:*:list",  
        "vpc:networks:create",  
        "vpc:networks:update",  
        "vpc:subnets:update",  
        "vpc:subnets:create",  
        "vpc:routers:get",  
        "vpc:routers:update",  
        "vpc:ports:*",  
        "vpc:privateIps:*",  
        "vpc:securityGroups:*",  
        "vpc:securityGroupRules:*",  
        "vpc:floatingIps:*",  
        "vpc:publicIps:*",  
        "vpc:bandwidths:*",  
        "vpc:firewalls:*",  
        "ims:images:create",  
        "ims:images:delete",  
        "ims:images:get",  
        "ims:images:list",  
        "ims:images:update",  
      ]  
    }  
  ]  
}
```

```
    "ims:images:upload"  
  ],  
  "Effect": "Allow"  
}  
]  
}
```

1.14 Region and AZ


Concept

A region and availability zone (AZ) identify the location of a data center. You can create resources in a specific region and AZ.

- A region is a physical data center, which is completely isolated to improve fault tolerance and stability. The region that is selected during resource creation cannot be changed after the resource is created.
- An AZ is a physical location where resources use independent power supplies and networks. A region contains one or more AZs that are physically isolated but interconnected through internal networks. Because AZs are isolated from each other, any fault that occurs in one AZ will not affect others.

[Figure 1-3](#) shows the relationship between regions and AZs.

Figure 1-3 Regions and AZs



Selecting a Region

Select a region closest to your target users for lower network latency and quick access.

Selecting an AZ

When deploying resources, consider your applications' requirements on disaster recovery (DR) and network latency.

- For high DR capability, deploy resources in different AZs within the same region.
- For lower network latency, deploy resources in the same AZ.

Regions and Endpoints

Before you use an API to call resources, specify its region and endpoint. For more details, see [Regions and Endpoints](#).

2 Getting Started

2.1 Creating an ECS

2.1.1 Overview

Scenarios

ECSs are more cost-effective than physical servers. Within minutes, you can obtain ECS resources from the cloud service platform. ECS resources are flexible and on-demand. This section describes how to create an ECS on the management console.

Creation process:

- [Step 1: Configure Basic Settings](#)
- [Step 2: Configure Network](#)
- [Step 3: Configure Advanced Settings](#)
- [Step 4: Confirm](#)

2.1.2 Step 1: Configure Basic Settings


Accessing the ECS Creation Page

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Click **Create ECS**.

The page for creating ECSs is displayed.

Basic Settings

1. Confirm the region.

If the region is incorrect, click  in the upper left corner of the page to select your region.

2. Select an AZ.

An AZ is a physical location that uses independent power supply and networks. AZs in the same region can communicate with each other over an intranet.

- To enhance application availability, create ECSs in different AZs.
- To shorten network latency, create ECSs in the same AZ.

NOTE

Random AZ allocation is available on the console when you create an ECS. The system will use a hash algorithm to select an AZ as the default AZ based on your universally unique identifier (UUID).

The available ECS types and flavors vary depending on AZs. To view all supported ECS types and flavors on the cloud service platform, set **AZ** to **Random**. Then, the system automatically allocates an AZ according to your selected ECS flavor.

For example, S3 ECSs are available only in AZ 1; S2 ECSs are available in AZ 2 and AZ 3 and have been sold out in AZ 1. If you set **AZ** to **Random**, you can view both S3 and S2 ECSs. If you create an S3 ECS, the system automatically allocates it to AZ 1. If you create an S2 ECS, the system randomly allocates it to AZ 2 or AZ 3.

3. Set **Specifications**.

The cloud platform provides various ECS types for different application scenarios. You can choose from existing ECS types and flavors in the list. Alternatively, you can enter a flavor or specify vCPUs and memory size to search for the flavor suited to your needs.

NOTE

- Before selecting an ECS type, learn about the introduction and notes on each type of ECSs. For details, see [ECS Types](#).

4. Select an image.

- Public image

A public image is a standard, widely used image. It contains an OS and preinstalled public applications and is available to all users. You can configure the runtime environment or software in the public image as needed.

- Private image

A private image is an image available only to the user who created it. It contains an OS, preinstalled public applications, and the user's private applications. Using a customized private image, you can create ECSs tailored to your needs in batches.

 NOTE

- If a full-ECS image is in **Normal** state and the system displays message "Available in AZx", the full-ECS image can be used to create ECSs in this AZ only. Additionally, the SCSI and sharing attribute settings of the system and data disks cannot be modified during ECS creation.
- If a full-ECS image is in **Normal** state but the system does not display message "Available in AZx", the full-ECS image can be used to create ECSs in the entire region. Additionally, the SCSI and sharing attribute settings of data disks can be modified during ECS creation.
- To ensure that NIC multi-queue is enabled on an ECS created using a private image, configure NIC multi-queue when creating such a private image. NIC multi-queue routes NIC interrupt requests among multiple vCPUs for higher network packets per second (PPS) and bandwidth.

For details, see "How Do I Set NIC Multi-Queue Feature of an Image?"

– Shared image

A shared image is a private image shared by another user.

5. Set **System Disk** and **Data Disk** if required.

– System disk

For details about the disk types supported by ECS, see [EVS Disks](#).

– Data disk

You can create multiple data disks for an ECS and enable required functions for each data disk. When creating an ECS, you can add up to 23 data disks with customized sizes to it. After the ECS is created, you can add up to 23 VBD disks or 59 SCSI disks to it.

Click **Show**  and set the following functions if required:

- **SCSI**: indicates that the device type of the data disk is SCSI if you select this option. For more information about SCSI disks and the ECSs that can be attached with SCSI disks, see .
- **Share**: indicates that the EVS disk is sharable if you select this option. Such an EVS disk can be attached to multiple ECSs.

6. Click **Next: Configure Network**.

2.1.3 Step 2: Configure Network

Network Settings

1. Set **Network** by selecting an available VPC and subnet from the drop-down list and specifying a private IP address assignment mode.

VPC provides a dedicated network for your ECS. A VPC can contain subnets for further isolation. You can configure security groups per subnet to control access to cloud resources.

You can select an existing VPC or create a new one.

For more information about VPC, see *Virtual Private Cloud User Guide*.

 NOTE

- Ensure that DHCP is enabled in the VPC to which the ECS belongs.
 - When you use VPC for the first time, the system automatically creates a VPC for you, including the security group and NIC.
2. (Optional) Add an extension NIC. You can add multiple extension NICs to an ECS and specify IP addresses for them (including primary NICs).

 NOTE

If you specify an IP address for a NIC when creating multiple ECSs in a batch:

- This IP address serves as the start IP address.
 - Ensure that the IP addresses required by the NICs are within the subnet, consecutive, and available.
 - The subnet with the specified IP address cannot overlap with other subnets.
- By default, the system assigns IPv4 addresses. If you select **Automatically-assigned IPv6 address**, NICs support both IPv4 and IPv6 addresses, and the system assigns IPv4 and IPv6 dual-stack addresses. In a VPC, an ECS uses an IPv6 address to access the dual-stack intranet. To access the Internet, you must enable **IPv6 Bandwidth** and select a shared bandwidth. Then, the ECS accesses the IPv6 Internet through the IPv6 address.

After creating an ECS, manually configure it, including allowing it to dynamically obtain an IPv6 address and enabling IPv6. For details, see [Dynamically Assigning IPv6 Addresses](#).

 NOTE

- IPv6 can be enabled only during ECS creation, and the configuration cannot be modified. If **IPv6 Bandwidth** is not enabled when you create an ECS, you can enable it after the ECS is created.
 - Dedicated bandwidth is not supported.
3. Set **Security Group** by selecting an available security group from the drop-down list or creating a new one.

A security group controls ECS access within or between security groups by defining access rules. This enhances ECS security.

When creating an ECS, you can select multiple (recommended not more than five) security groups. In such a case, the access rules of all the selected security groups apply on the ECS.

 NOTE

Before initializing an ECS, ensure that the security group rules for the outbound direction meet the following requirements:

- **Protocol:** TCP
- **Port Range:** 80
- **Remote End:** 169.254.0.0/16

If you use the default security group rules for the outbound direction, the preceding requirements are met, and the ECS can be initialized. The default security group rules for the outbound direction are as follows:

- **Protocol:** ANY
- **Port Range:** ANY
- **Remote End:** 0.0.0.0/16

4. Set EIP.

An EIP is a static public IP address bound to an ECS in a VPC. Using the EIP, the ECS provides services externally.

The following options are provided:

- Auto assign

The system automatically assigns an EIP for the ECS. The EIP provides a dedicated bandwidth that is configurable.

- Specify

An existing EIP is assigned for the ECS. When using an existing EIP, you are not allowed to create ECSs in a batch.

- Do not use

Without an EIP, the ECS cannot access the Internet and is used in the private network or cluster only.

5. Set **Bandwidth Size**.

Select the bandwidth based on service requirements. The unit is Mbit/s.

2.1.4 Step 3: Configure Advanced Settings

Advanced Settings

1. Set **ECS Name**.

The name can be customized but can contain only letters, digits, underscores (_), hyphens (-), and periods (.).

If you want to create multiple ECSs at a time, the system automatically sequences these ECSs.

If multiple ECSs are created at the same time, the system automatically adds a hyphen followed by a four-digit incremental number to the end of each ECS name. For example, if you enter **ecs**, the ECSs will be named **ecs-0001**, **ecs-0002**, ... If you create multiple ECSs again, the values in the new ECS names increase from the existing maximum value. For example, the existing ECS with the maximum number in name is **ecs-0010**. If you enter **ecs**, the names of the new ECSs will be **ecs-0011**, **ecs-0012**, ... When the value reaches **9999**, it will start from **0001**.

Allow duplicate name: allows ECS names to be duplicate. If you select **Allow duplicate name** and create multiple ECSs in a batch, the created ECSs will have the same name.

The **ECS Name** set in this step will be the initial host name in the ECS OS.

 **NOTE**

Consecutive periods (.) or hyphens (-) will be replaced with the first character to prevent unknown issues.

2. Set **Login Mode**.

- Password

A username and its initial password are used for ECS login authentication.

The initial password of user **root** is used for authenticating Linux ECSs, while that of user **Administrator** is used for authenticating Windows ECSs.

The passwords must meet the requirements described in [Table 2-1](#).

Table 2-1 Password complexity requirements

Parameter	Requirement
Password	<ul style="list-style-type: none">• Consists of 8 to 26 characters. For specific requirements on the password length, see the information displayed on the console.• Contains at least three of the following character types:<ul style="list-style-type: none">- Uppercase letters- Lowercase letters- Digits- Special characters for Windows: \$!@%-_+=[]:./,?- Special characters for Linux: !@%-_+=[]:./^,{}?• Cannot contain the username or the username spelled backwards.• Cannot contain more than two consecutive characters in the same sequence as they appear in the username. (This requirement applies only to Windows ECSs.)• Cannot start with a slash (/) for Windows ECSs.

 **NOTE**

The system does not periodically change the ECS password. It is recommended that you change your password regularly for security.

- Image password

If you use a Linux private image to create an ECS, you can use the image password for login authentication. Make sure that a password has been set for the selected private image.

- Set password later
The password for logging in to the ECS is not configured during the ECS creation. After the ECS is created, select **Reset Password** in the **Operation** column, set a password for the ECS as prompted, and log in to the ECS.

3. Set **Cloud Backup and Recovery**.

Cloud Backup and Recovery (CBR) provides backup protection for EVS disks and ECSs, and uses backups to restore the EVS disks and ECSs. After you set **Cloud Backup and Recovery**, the system binds the target ECS to the cloud backup vault and associates the ECS with the selected backup policy to periodically back up the ECS.

The following options are provided:

- Create new
 - i. Set the name of the cloud backup vault, which consists of 1 to 64 characters, containing only letters, digits, underscores (_), and hyphens (-). For example, **vault-f61e**. The default naming rule is **vault_xxxx**.
 - ii. Enter the vault capacity, which is required for backing up the ECS. The vault capacity cannot be smaller than that of the ECS to be backed up. Its value ranges from the total capacity of the ECS to 10,485,760 in the unit of GB.
 - iii. Select a backup policy from the drop-down list, or log in to the CBR console and configure a desired one.
- Specify
 - i. Select an existing cloud backup vault from the drop-down list.
 - ii. Select a backup policy from the drop-down list, or log in to the CBR console and configure a desired one.
- Do not use
Skip this configuration if CBR is not required. If you need to enable CBR after creating an ECS, log in to the CBR console, locate the target vault, and bind the ECS to the vault.

4. Set **ECS Group (Optional)**.

An ECS group applies the anti-affinity policy to the ECSs in it so that the ECSs are automatically allocated to different hosts. This configuration is optional. For instructions about how to create an ECS group, see [Managing ECS Groups](#).

 **NOTE**

An existing ECS attached with a local disk cannot be added to an ECS group. To use ECS group functions, select an ECS group when creating an ECS.

5. Configure automatic recovery.

If a physical server accommodating ECSs breaks down, the ECSs with automatic recovery enabled will automatically be migrated to a functional server to minimize the impact on your services. During this process, the ECSs will restart. If automatic recovery is not enabled, the affected users must wait for the system administrator to recover ECSs once the physical server becomes faulty. For more information about automatic recovery, see [Automatically Recovering ECSs](#).

 **NOTE**

An ECS with any of the following resources cannot be automatically recovered:

- Local disk
 - Passthrough FPGA
 - Passthrough InfiniBand NIC
6. To use functions listed in **Advanced Options**, select **Configure now**. Otherwise, do not select it.
- User Data
You can specify the user data. The user data will be automatically passed to the ECS when the ECS starts for the first time. This configuration is optional.
For example, if you activate user **root** permission by passing a script file to an ECS, you can log in to the ECS as user **root**.
For detailed operations, see [Passing User Data to ECSs](#).
 - Agency
This configuration is optional. When your ECS resources need to be shared with other accounts, or your ECS is delegated to professional personnel or team for management, the tenant administrator creates an agency in IAM and grants the ECS management permissions to the personnel or team. The delegated account can log in to the cloud system and switch to your account to manage resources. You do not need to share security credentials (such as passwords) with other accounts, ensuring the security of your account.
If you have created an agency in IAM, you can select the agency from the drop-down list and obtain specified operation permissions. For instructions about how to create an agency, see *Identity and Access Management User Guide*.
7. Click **Next: Confirm**.

2.1.5 Step 4: Confirm

Confirming the Order

1. On the **Confirm** page, view details about the ECS configuration.
2. Set the number of ECSs to be created.
3. Confirm the configuration and click **Apply Now**.

Follow-up Procedure

- If your ECS is created using a private image that is created using an external image file, and the ECS has no one-click password reset plug-in installed, it is a good practice to install the password reset plug-in after you log in to the ECS. If your login password is forgotten or expires, you can use the one-click password reset function to set a new password for the ECS.
For details, see "Installing the One-Click Password Reset Plug-in".
- After an ECS with automatic recovery enabled is created, you can check whether this function has been enabled by viewing **Failures**. For details, see [Viewing Failed Tasks](#).

2.2 Logging In to an ECS

Logging In to a Windows ECS in China

You can log in to a Windows ECS using either VNC or MSTSC provided on the management console.

- Management console (VNC)
For details, see [Remotely Logging In to a Windows ECS \(Using VNC\)](#).
- Remote desktop connection (MSTSC)
For details, see [Remotely Logging In to a Windows ECS \(Using MSTSC\)](#).

Logging In to a Windows ECS Outside China

You can log in to a Windows ECS using either VNC or MSTSC provided on the management console.

- Management console (VNC)
For details, see [Remotely Logging In to a Windows ECS \(Using VNC\)](#).
- Remote desktop connection (MSTSC)
For details, see [Remotely Logging In to a Windows ECS \(Using MSTSC\)](#).

Logging In to a Linux ECS in China

The method of logging in to an ECS varies depending on the login authentication configured during ECS creation.

- To log in to a password-authenticated ECS for the first time, use either of the following methods:
 - For logins using VNC on the management console, the login username is **root**.
For details about how to log in to the ECS using VNC, see [Remotely Logging In to a Linux ECS \(Using VNC\)](#).
 - For logins using an SSH password, the login username is **root** and the ECS must have an EIP bound.
For details, see [Remotely Logging In to a Linux ECS \(Using an SSH Password\)](#).

Logging In to a Linux ECS Outside China

Select a login method and log in to the ECS.

- VNC
For details, see [Remotely Logging In to a Linux ECS \(Using VNC\)](#).

Follow-up Procedure

- If you have added a data disk during ECS creation, you must initialize the data disk after logging in to the ECS.

For details, see [Scenarios and Disk Partitions](#).

- Certain ECSs require the installation of a driver after you log in to them. For details about available ECS types and functions, see [ECS Types](#). For details about restrictions on using different types of ECSs, see their notes.

2.3 Initializing EVS Data Disks

2.3.1 Scenarios and Disk Partitions

If you have added a data disk during ECS creation, you must initialize the data disk after logging in to the ECS.

Scenarios

After a disk is attached to a server, you need to log in to the server to initialize the disk, that is, format the disk. You must initialize a disk before accessing it.

- System disk
A system disk does not require manual initialization because it is automatically created and initialized upon server creation. The default partition style is master boot record (MBR).
- Data disk
 - If a data disk is created along with a server, it will be automatically attached to the server.
 - If a data disk is created separately, you need to manually attach it to a server.

In both cases, you must initialize the data disk before using it. Choose an appropriate partition style based on your service plan.

Partitioning Operation Guide

[Table 2-2](#) lists the common disk partition styles. In Linux, different disk partition styles require different partitioning tools.

Table 2-2 Disk partition styles

Disk Partition Style	Maximum Disk Capacity Supported	Maximum Number of Partitions Supported	Linux Partitioning Tool
Master Boot Record (MBR)	2 TiB	<ul style="list-style-type: none">• 4 primary partitions• 3 primary partitions and 1 extended partition <p>With MBR, you can create several primary partitions and one extended partition. The extended partition must be divided into logical partitions before use. For example, if 6 partitions need to be created, you can create them in the following two ways:</p> <ul style="list-style-type: none">• 3 primary partitions and 1 extended partition, with the extended partition divided into 3 logical partitions• 1 primary partition and 1 extended partition, with the extended partition divided into 5 logical partitions	<ul style="list-style-type: none">• fdisk• parted
GUID Partition Table (GPT)	18 EiB 1 EiB = 1048576 TiB	Unlimited Disk partitions created using GPT are not categorized.	parted

2.3.2 Initializing a Windows Data Disk (Windows Server 2008)

Scenarios

This section uses Windows Server 2008 R2 Enterprise 64bit to describe how to initialize a data disk attached to a server running Windows.

The maximum disk capacity supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Therefore, use the GPT partition style if your disk capacity is larger than 2 TiB. For details, see [Initializing a Windows Data Disk Larger Than 2 TiB \(Windows Server 2008\)](#). To learn more about disk partition styles, see [Scenarios and Disk Partitions](#).

The method for initializing a disk varies slightly depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

NOTICE

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

Prerequisites

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
 - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.

Procedure

Step 1 On the desktop of the server, right-click **Computer** and choose **Manage** from the shortcut menu.

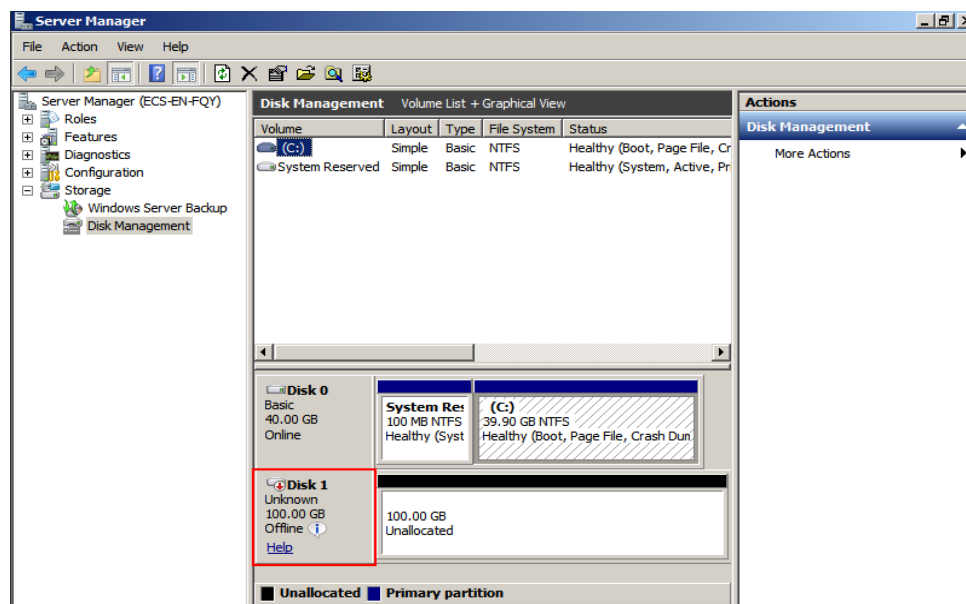
The **Server Manager** window is displayed.

Step 2 In the navigation tree, choose **Storage > Disk Management**.

The **Disk Management** window is displayed.


- If **Figure 2-1** is displayed, the new disk is offline. Go to **Step 3**.
- If **Figure 2-4** is displayed, the **Initialize Disk** window is prompted. Go to **Step 5**.

Figure 2-1 Disk Management



Step 3 Disks are displayed in the right pane. In the **Disk 1** area, right-click **Offline** and choose **Online** from the shortcut menu to online the disk.

Figure 2-2 Online the disk

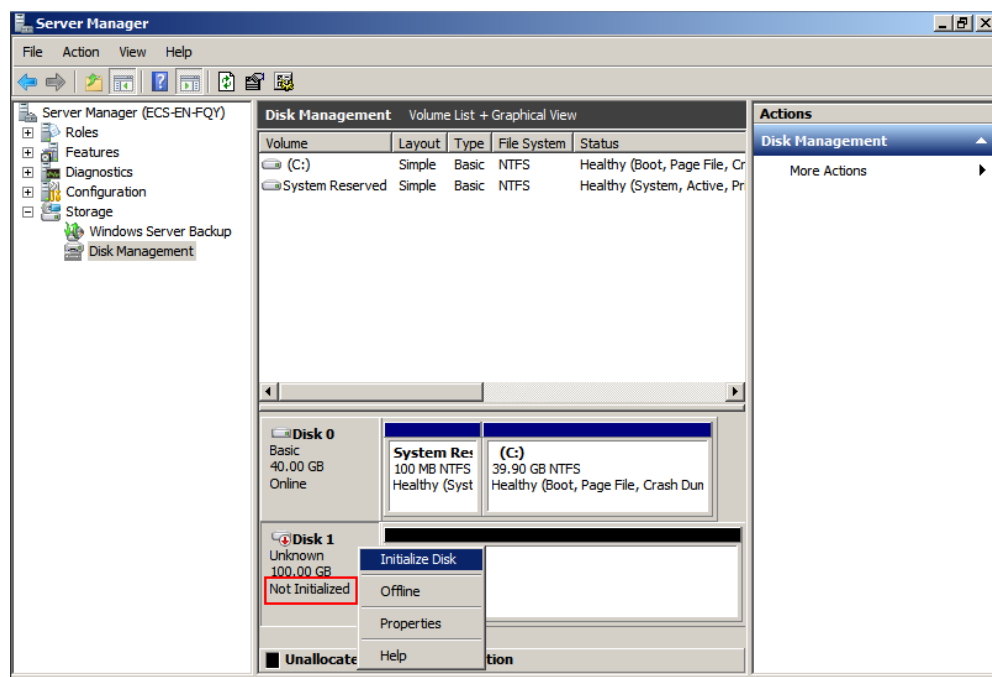


NOTE

If the disk is offline, you need to bring it online before initializing it.


- Step 4** After making the disk online, the disk status changes from **Offline** to **Not Initialized**. Right-click the disk status and choose **Initialize Disk** from the shortcut menu.

Figure 2-3 Initialize Disk



Step 5 In the **Initialize Disk** dialog box, select the target disk, click **MBR (Master Boot Record)** or **GPT (GUID Partition Table)**, and click **OK**.

Figure 2-4 Unallocated space



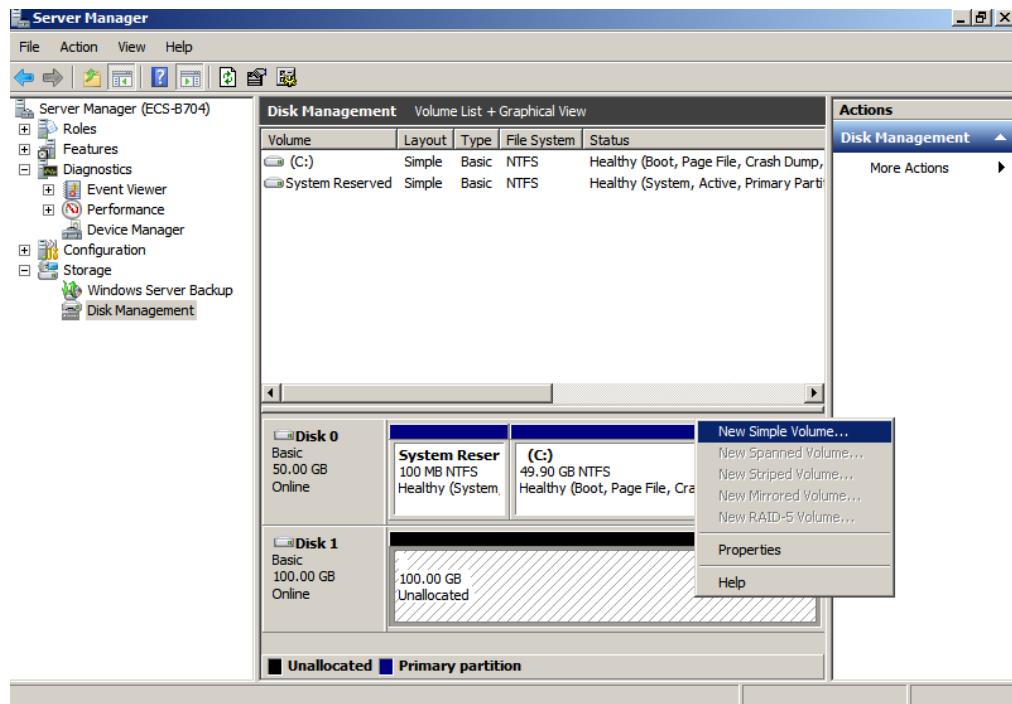
NOTICE

The maximum disk size supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Because an EVS data disk currently supports up to 32 TiB, use GPT if your disk size is larger than 2 TiB.

If the partition style is changed after the disk has been used, all data on the disk will be lost, so take care to select an appropriate partition style when initializing the disk. If you must change the partition style to GPT after a disk has been used, it is recommended that you back up the disk data before the change.


Step 6 Right-click at the unallocated space and choose **New Simple Volume** from the shortcut menu.

Figure 2-5 New Simple Volume




Step 7 On the displayed **New Simple Volume Wizard** window, click **Next**.

Figure 2-6 New Simple Volume Wizard




Step 8 Specify the volume size and click **Next**. The default value is the maximum size.

Figure 2-7 Specify Volume Size



Step 9 Assign the drive letter and click **Next**.

Figure 2-8 Assign Drive Letter or Path



Step 10 On the displayed **Format Partition** page, click **Format this volume with the following settings**, set parameters based on the requirements, and select **Perform a quick format**. Then, click **Next**.

Figure 2-9 Format Partition


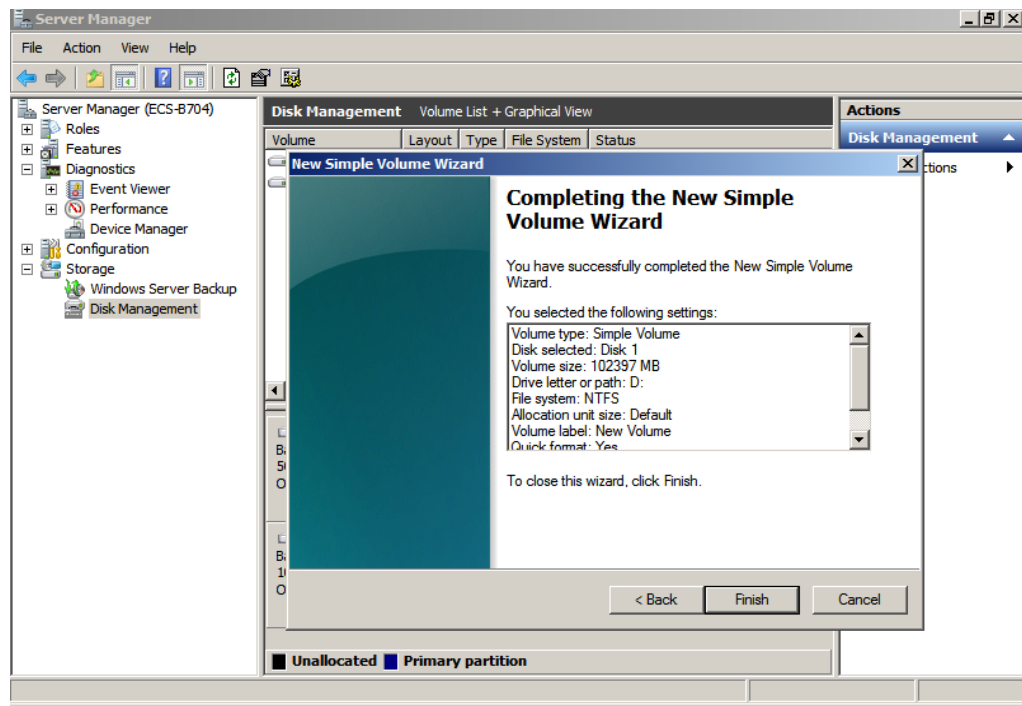


Figure 2-10 Completing the partition creation




NOTICE

The partition sizes supported by file systems vary. Choose an appropriate file system format based on your service requirements.

- Step 11** Click **Finish**. Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished successfully.

Figure 2-11 Disk initialization succeeded



----End

2.3.3 Initializing a Windows Data Disk (Windows Server 2019)

Scenarios

This section uses Windows Server 2019 Standard 64bit to describe how to initialize a data disk attached to a server running Windows.

The maximum disk capacity supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Therefore, use the GPT partition style if your disk capacity is larger than 2 TiB. For details, see [Initializing a Windows Data Disk Larger Than 2 TiB \(Windows Server 2008\)](#). To learn more about disk partition styles, see [Scenarios and Disk Partitions](#).

The method for initializing a disk varies slightly depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

NOTICE

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

Prerequisites

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
 - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.

Procedure

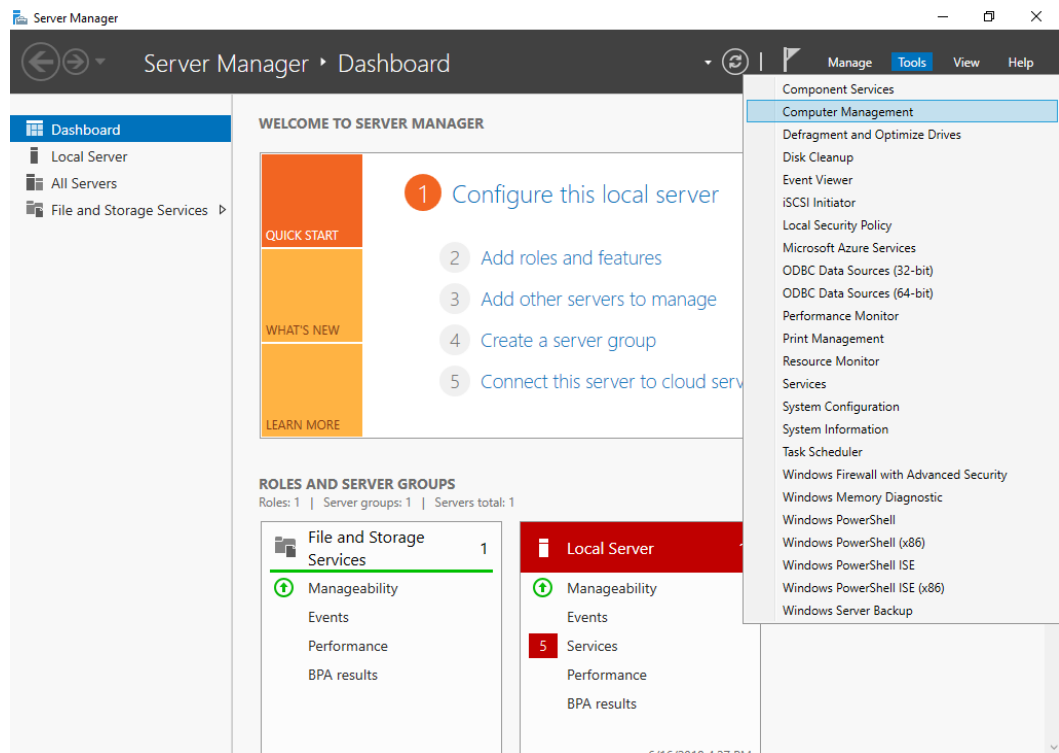
Step 1 On the desktop of the server, click the start icon in the lower left corner.

The **Windows Server** window is displayed.

Step 2 Click **Server Manager**.

The **Server Manager** window is displayed.


Figure 2-12 Server Manager



Step 3 In the upper right corner, choose **Tools > Computer Management**.

The **Computer Management** window is displayed.

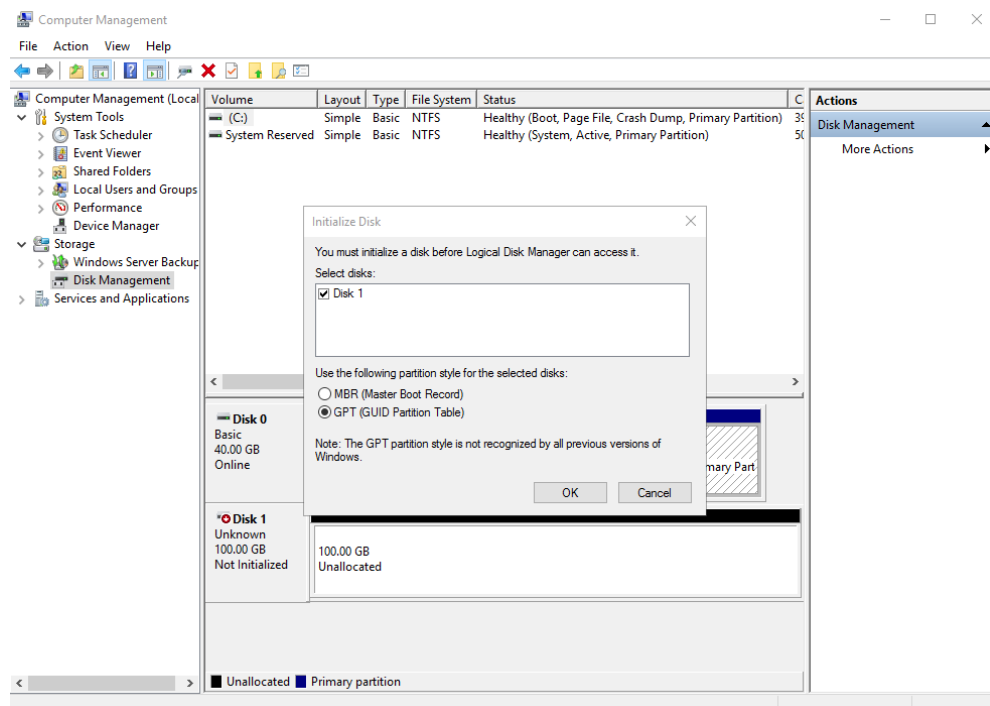
Figure 2-13 Computer Management



Step 4 Choose **Storage > Disk Management**.

Disks are displayed in the right pane. If there is a disk that is not initialized, the system will prompt you with the **Initialize Disk** dialog box.

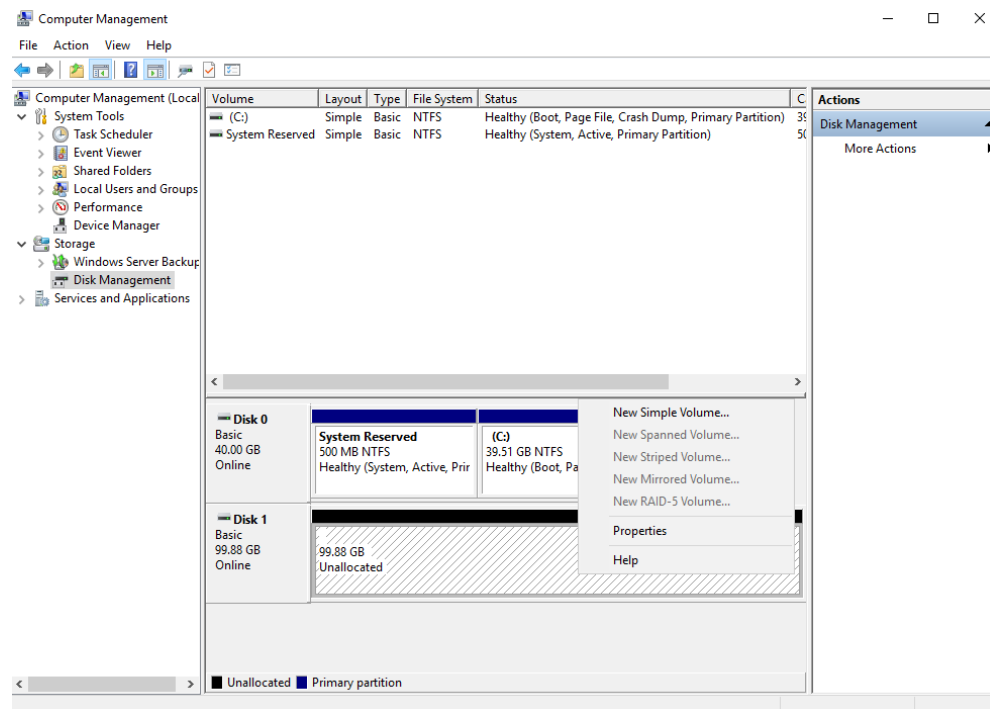
Figure 2-14 Disk list



Step 5 In the **Initialize Disk** dialog box, the to-be-initialized disk is selected. Select a disk partition style and click **OK**. In this example, **GPT (GUID Partition Table)** is selected.

The **Computer Management** window is displayed.

Figure 2-15 Computer Management



NOTICE


The maximum disk size supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Because an EVS data disk currently supports up to 32 TiB, use GPT if your disk size is larger than 2 TiB.

If the partition style is changed after the disk has been used, all data on the disk will be lost, so take care to select an appropriate partition style when initializing the disk. If you must change the partition style to GPT after a disk has been used, it is recommended that you back up the disk data before the change.

Step 6 Right-click at the unallocated disk space and choose **New Simple Volume** from the shortcut menu.

The **New Simple Volume Wizard** window is displayed.

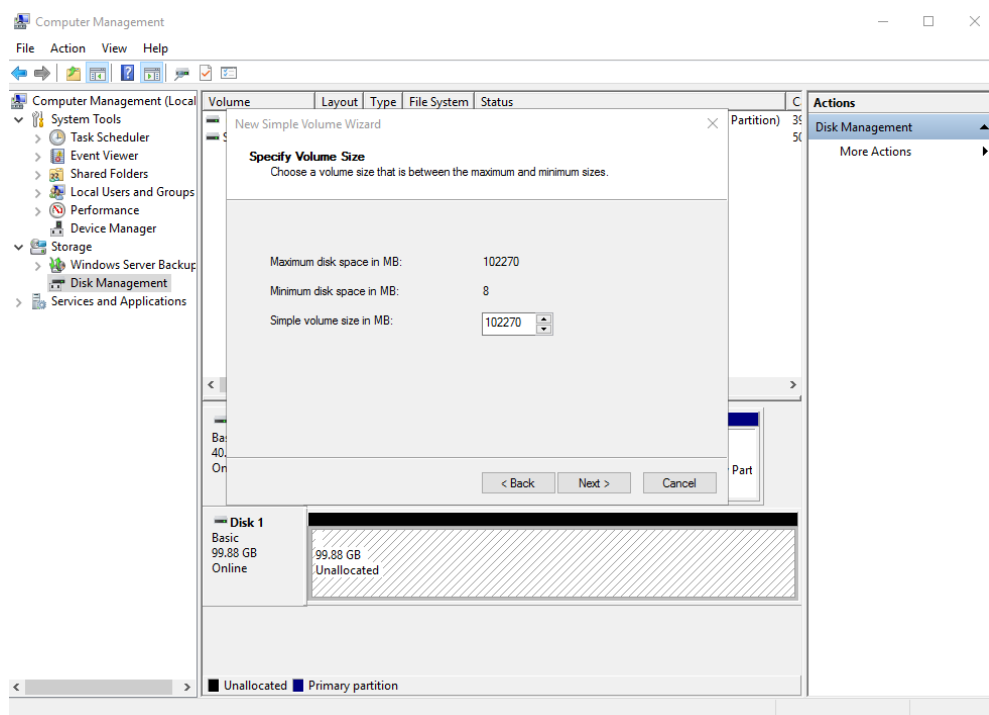
Figure 2-16 New Simple Volume Wizard



Step 7 Follow the prompts and click **Next**.

The **Specify Volume Size** page is displayed.

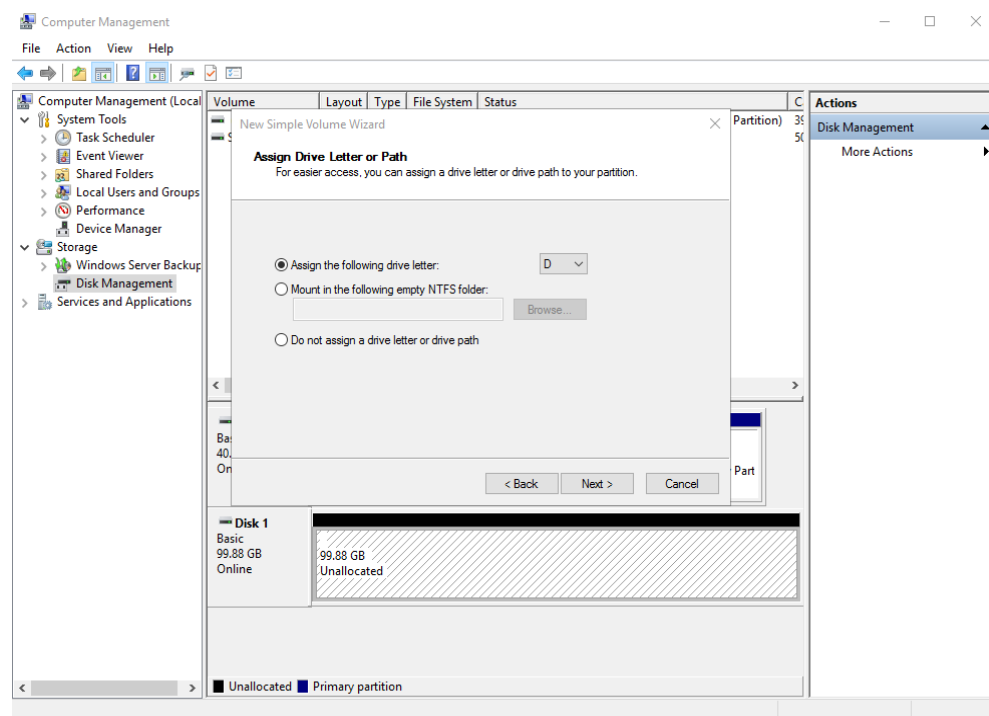
Figure 2-17 Specify Volume Size



Step 8 Specify the volume size and click **Next**. The system selects the maximum volume size by default. You can specify the volume size as required. In this example, the default setting is used.

The **Assign Drive Letter or Path** page is displayed.

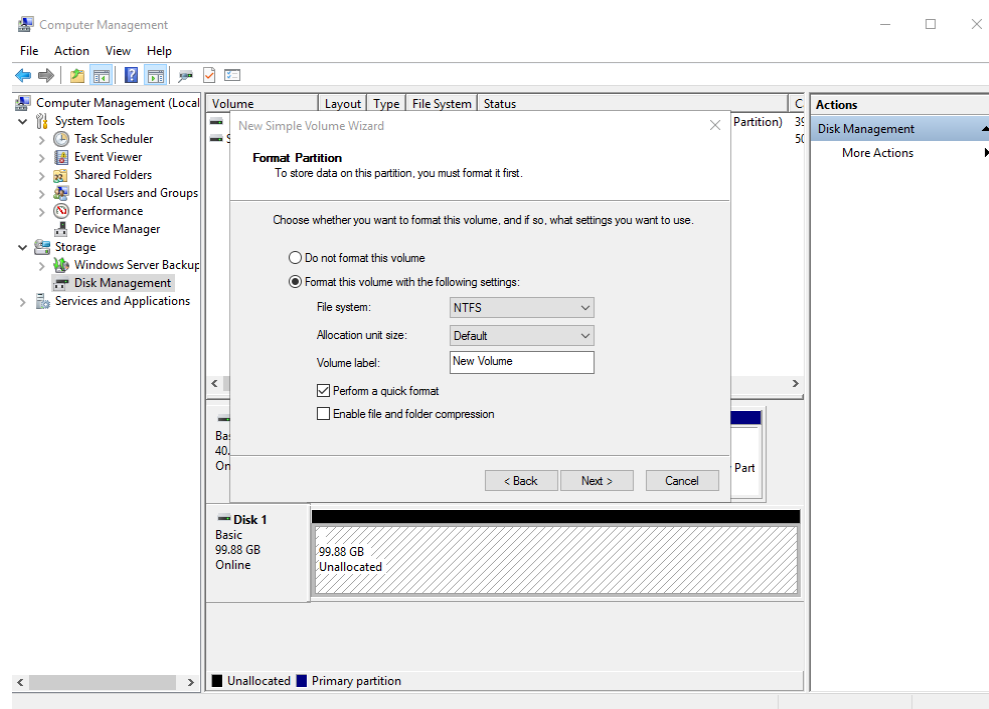
Figure 2-18 Assign Drive Letter or Path



Step 9 Assign a drive letter or path to your partition and click **Next**. The system assigns drive letter D by default. In this example, the default setting is used.

The **Format Partition** page is displayed.


Figure 2-19 Format Partition



Step 10 Specify format settings and click **Next**. The system selects the NTFS file system by default. You can specify the file system type as required. In this example, the default setting is used.

The **Completing the New Simple Volume Wizard** page is displayed.

Figure 2-20 Completing the New Simple Volume Wizard




NOTICE

The partition sizes supported by file systems vary. Choose an appropriate file system format based on your service requirements.

Step 11 Click **Finish**.

Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished successfully.

Figure 2-21 Disk initialized





Step 12 After the volume is created, click  on the task bar and check whether a new volume appears in **This PC**. In this example, New Volume (D:) is the new volume. If New Volume (D:) appears, the disk is successfully initialized and no further action is required.

Figure 2-22 This PC



----End

2.3.4 Initializing a Linux Data Disk (fdisk)

Scenarios

This section uses CentOS 7.4 64bit to describe how to initialize a data disk attached to a server running Linux and use fdisk to partition the data disk.

The maximum partition size that MBR supports is 2 TiB and that GPT supports is 18 EiB. If the disk size you need to partition is greater than 2 TiB, partition the disk using GPT.

The fdisk partitioning tool is suitable only for MBR partitions, and the parted partitioning tool is suitable for both MBR and GPT partitions. For more information, see [Scenarios and Disk Partitions](#).

The method for initializing a disk varies slightly depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

NOTICE

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

Prerequisites

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
 - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.

Creating and Mounting a Partition

The following example shows you how a new primary partition can be created on a new data disk that has been attached to a server. The primary partition will be created using `fdisk`, and MBR will be used. Furthermore, the partition will be formatted using the file system, mounted on `/mnt/sdc`, and configured to mount automatically at startup.

Step 1 Query what block devices are available on the server.

fdisk -l

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# fdisk -l
```

```
Disk /dev/vda: 42.9 GiB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000bcb4e
```

Device	Boot	Start	End	Blocks	Id	System
/dev/vda1	*	2048	83886079	41942016	83	Linux

```
Disk /dev/vdb: 107.4 GiB, 107374182400 bytes, 209715200 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

In the command output, this server contains two disks. `/dev/vda` and `/dev/vdb`. `/dev/vda` is the system disk, and `/dev/vdb` is the new data disk.

Step 2 Launch `fdisk` to partition the new data disk.

fdisk *New data disk*

In this example, run the following command:

fdisk /dev/vdb

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# fdisk /dev/vdb
```

```
Welcome to fdisk (util-linux 2.23.2).
```

```
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.
```

```
Device does not contain a recognized partition table
Building a new DOS disklabel with disk identifier 0x38717fc1.
```

```
Command (m for help):
```

Step 3 Enter `n` and press **Enter** to create a new partition.

Information similar to the following is displayed:

```
Command (m for help): n
```

```
Partition type:
```

```
p primary (0 primary, 0 extended, 4 free)
e extended
```

There are two types of disk partitions:

- Choosing **p** creates a primary partition.
- Choosing **e** creates an extended partition.

 **NOTE**

If MBR is used, a maximum of four primary partitions, or three primary partitions plus one extended partition can be created. The extended partition must be divided into logical partitions before use.

Disk partitions created using GPT are not categorized.

Step 4 Enter **p** and press **Enter** to create a primary partition in this example.

Information similar to the following is displayed:

```
Select (default p): p
Partition number (1-4, default 1):
```

Partition number indicates the serial number of the primary partition. The value ranges from **1** to **4**.

Step 5 Enter the serial number of the primary partition and press **Enter**. Primary partition number **1** is used in this example. One usually starts with partition number **1** when partitioning an empty disk.

Information similar to the following is displayed:

```
Partition number (1-4, default 1): 1
First sector (2048-209715199, default 2048):
```

First sector indicates the start sector. The value ranges from **2048** to **209715199**, and the default value is **2048**.

Step 6 Select the default start sector **2048** and press **Enter**.

The system displays the start and end sectors of the partition's available space. You can customize the value within this range or use the default value. The start sector must be smaller than the partition's end sector.

Information similar to the following is displayed:

```
First sector (2048-209715199, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-209715199, default 209715199):
```

Last sector indicates the end sector. The value ranges from **2048** to **209715199**, and the default value is **209715199**.

Step 7 Select the default end sector **209715199** and press **Enter**.

The system displays the start and end sectors of the partition's available space. You can customize the value within this range or use the default value. The start sector must be smaller than the partition's end sector.

Information similar to the following is displayed:

```
Last sector, +sectors or +size{K,M,G} (2048-209715199, default 209715199):
Using default value 209715199
Partition 1 of type Linux and of size 100 GiB is set
Command (m for help):
```

A primary partition has been created for the new data disk.

Step 8 Enter **p** and press **Enter** to print the partition details.

Information similar to the following is displayed:

```
Command (m for help): p

Disk /dev/vdb: 107.4 GiB, 107374182400 bytes, 209715200 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x38717fc1

   Device Boot      Start         End      Blocks   Id  System
/dev/vdb1            2048    209715199    104856576   83   Linux
```

Command (m for help):

Details about the **/dev/vdb1** partition are displayed.

Step 9 Enter **w** and press **Enter** to write the changes to the partition table.

Information similar to the following is displayed:

```
Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

The partition is created.

NOTE

In case that you want to discard the changes made before, you can exit fdisk by entering **q**.

Step 10 Synchronize the new partition table to the OS.

partprobe

Step 11 Format the new partition with a desired file system format.

mkfs -t *File system format* /dev/vdb1

In this example, the **ext4** format is used for the new partition.

mkfs -t ext4 /dev/vdb1

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# mkfs -t ext4 /dev/vdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
6553600 inodes, 26214144 blocks
1310707 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2174746624
800 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000, 23887872

Allocating group tables: done
Writing inode tables: done
```



```
Creating journal (32768 blocks): done  
Writing superblocks and filesystem accounting information: done
```

The formatting takes a period of time. Observe the system running status and do not exit.

NOTICE

The partition sizes supported by file systems vary. Choose an appropriate file system format based on your service requirements.

Step 12 Create a mount point.

mkdir *Mount point*

In this example, the **/mnt/sdc** mount point is created.

mkdir /mnt/sdc

 **NOTE**

The **/mnt** directory exists on all Linux systems. If the mount point cannot be created, it may be that the **/mnt** directory has been accidentally deleted. You can run **mkdir -p /mnt/sdc** to create the mount point.

Step 13 Mount the new partition on the created mount point.

mount *Disk partition Mount point*

In this example, the **/dev/vdb1** partition is mounted on **/mnt/sdc**.

mount /dev/vdb1 /mnt/sdc

Step 14 Check the mount result.

df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH  
Filesystem Type Size Used Avail Use% Mounted on  
/dev/vda1 ext4 43G 1.9G 39G 5% /  
devtmpfs devtmpfs 2.0G 0 2.0G 0% /dev  
tmpfs tmpfs 2.0G 0 2.0G 0% /dev/shm  
tmpfs tmpfs 2.0G 9.1M 2.0G 1% /run  
tmpfs tmpfs 2.0G 0 2.0G 0% /sys/fs/cgroup  
tmpfs tmpfs 398M 0 398M 0% /run/user/0  
/dev/vdb1 ext4 106G 63M 101G 1% /mnt/sdc
```

You should now see that partition **/dev/vdb1** is mounted on **/mnt/sdc**.

 **NOTE**

After the server is restarted, the disk will not be automatically mounted. You can modify the **/etc/fstab** file to configure automount at startup. For details, see [Configuring Automatic Mounting at System Start](#).

----End

Configuring Automatic Mounting at System Start

The **fstab** file controls what disks are automatically mounted at startup. You can use **fstab** to configure your data disks to mount automatically. This operation will not affect the existing data.

The example here uses UUIDs to identify disks in the **fstab** file. You are advised not to use device names to identify disks in the file because device names are assigned dynamically and may change (for example, from **/dev/vdb1** to **/dev/vdb2**) after a server stop or start. This can even prevent the server from booting up.

NOTE

UUIDs are the unique character strings for identifying partitions in Linux.

Step 1 Query the partition UUID.

blkid *Disk partition*

In this example, the UUID of the **/dev/vdb1** partition is queried.

blkid /dev/vdb1

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# blkid /dev/vdb1
/dev/vdb1: UUID="0b3040e2-1367-4abb-841d-ddb0b92693df" TYPE="ext4"
```

Carefully record the UUID, as you will need it for the following step.

Step 2 Open the **fstab** file using the vi editor.

vi /etc/fstab

Step 3 Press **i** to enter editing mode.

Step 4 Move the cursor to the end of the file and press **Enter**. Then, add the following information:

```
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df /mnt/sdc ext4 defaults 0 2
```

The preceding information is used for reference only. The line starting with **UUID** is the information added. Edit this line from left to right to match the following format:

- **UUID**: The UUID obtained in [Step 1](#).
- **Mount point**: The directory on which the partition is mounted. You can query the mount point using **df -TH**.
- **Filesystem**: The file system format of the partition. You can query the file system format using **df -TH**.
- **Mount option**: The partition mount option. Usually, this parameter is set to **defaults**.
- **Dump**: The Linux dump backup option.
 - **0**: Linux dump backup is not used. Usually, dump backup is not used, and you can set this parameter to **0**.
 - **1**: Linux dump backup is used.
- **fsck**: The fsck option, which means whether to use fsck to check the disk during startup.

- **0**: not use fsck.
- If the mount point is the root partition (/), this parameter must be set to **1**.

If this parameter is set to **1** for the root partition, this parameter for other partitions must start with **2** because the system checks the partitions in the ascending order of the values.

Step 5 Press **Esc**, enter **:wq**, and press **Enter**.

The system saves the configurations and exits the vi editor.

Step 6 Verify that the disk is auto-mounted at startup.

1. Unmount the partition.

umount *Disk partition*

In this example, run the following command:

```
umount /dev/vdb1
```

2. Reload all the content in the **/etc/fstab** file.

```
mount -a
```

3. Query the file system mounting information.

```
mount | grep Mount point
```

In this example, run the following command:

```
mount | grep /mnt/sdc
```

If information similar to the following is displayed, automatic mounting has been configured:

```
root@ecs-test-0001 ~]# mount | grep /mnt/sdc  
/dev/vdb1 on /mnt/sdc type ext4 (rw,relatime,data=ordered)
```

----End

2.3.5 Initializing a Linux Data Disk (parted)

Scenarios

This section uses CentOS 7.4 64bit to describe how to initialize a data disk attached to a server running Linux and use parted to partition the data disk.

The maximum partition size that MBR supports is 2 TiB and that GPT supports is 18 EiB. If the disk size you need to partition is greater than 2 TiB, partition the disk using GPT.

The fdisk partitioning tool is suitable only for MBR partitions, and the parted partitioning tool is suitable for both MBR and GPT partitions. For more information, see [Scenarios and Disk Partitions](#).

The method for initializing a disk varies slightly depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

NOTICE

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

Prerequisites

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
 - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.

Creating and Mounting a Partition

The following example shows you how a new partition can be created on a new data disk that has been attached to a server. The partition will be created using parted, and GPT will be used. Furthermore, the partition will be formatted using the ext4 file system, mounted on **/mnt/sdc**, and configured to mount automatically at startup.

Step 1 Query information about the new data disk.

lsblk

Information similar to the following is displayed:

```
root@ecs-test-0001 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
└─vda1 253:1 0 40G 0 part /
vdb 253:16 0 100G 0 disk
```

In the command output, this server contains two disks. **/dev/vda** and **/dev/vdb**. **/dev/vda** is the system disk, and **/dev/vdb** is the new data disk.

Step 2 Launch parted to partition the new data disk.

parted *New data disk*

In this example, run the following command:

parted /dev/vdb

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted)
```

Step 3 Enter **p** and press **Enter** to view the current disk partition style.

Information similar to the following is displayed:

```
(parted) p
Error: /dev/vdb: unrecognised disk label
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 107GiB
Sector size (logical/physical): 512B/512B
Partition Table: unknown
Disk Flags:
(parted)
```

In the command output, the **Partition Table** value is **unknown**, indicating that no partition style is set for the new disk.

Step 4 Set the disk partition style.

mklabel *Disk partition style*

This command lets you control whether to use MBR or GPT for your partition table. In this example, GPT is used.

mklabel gpt

NOTICE

The maximum disk size supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Because an EVS data disk currently supports up to 32 TiB, use GPT if your disk size is larger than 2 TiB.

If the partition style is changed after the disk has been used, all data on the disk will be lost, so take care to select an appropriate partition style when initializing the disk. If you must change the partition style to GPT after a disk has been used, it is recommended that you back up the disk data before the change.

Step 5 Enter **p** and press **Enter** to view the disk partition style.

Information similar to the following is displayed:

```
(parted) mklabel gpt
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 107GiB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start  End  Size  File system  Name  Flags
(parted)
```

In the command output, the **Partition Table** value is **gpt**, indicating that the disk partition style is GPT.

Step 6 Enter **unit s** and press **Enter** to set the measurement unit of the disk to sector.

Step 7 Create a new partition.

mkpart *Partition name Start sector End sector*

In this example, run the following command:

mkpart test 2048s 100%

In this example, one partition is created for the new data disk, starting on **2048** and using **100%** of the rest of the disk. The two values are used for reference only. You can determine the number of partitions and the partition size based on your service requirements.

Information similar to the following is displayed:

```
(parted) mkpart opt 2048s 100%
(parted)
```

Step 8 Enter **p** and press **Enter** to print the partition details.

Information similar to the following is displayed:

```
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 209715200s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start End      Size      File system Name  Flags
 1     2048s 209713151s 209711104s          test

(parted)
```

Step 9 Enter **q** and press **Enter** to exit parted.

Information similar to the following is displayed:

```
(parted) q
Information: You may need to update /etc/fstab.
```

You can configure automatic mounting by updating the **/etc/fstab** file. Before doing so, format the partition with a desired file system and mount the partition on the mount point.

Step 10 View the disk partition information.

lsblk

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda   253:0  0  40G  0 disk
├─vda1 253:1  0  40G  0 part /
vdb   253:16  0 100G  0 disk
├─vdb1 253:17  0 100G  0 part
```

In the command output, **/dev/vdb1** is the partition you created.

Step 11 Format the new partition with a desired file system format.

mkfs -t *File system format* **/dev/vdb1**

In this example, the **ext4** format is used for the new partition.

mkfs -t ext4 /dev/vdb1

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# mkfs -t ext4 /dev/vdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
6553600 inodes, 26213888 blocks
1310694 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2174746624
800 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000, 23887872

Allocating group tables: done
```

```
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

The formatting takes a period of time. Observe the system running status and do not exit.

NOTICE

The partition sizes supported by file systems vary. Choose an appropriate file system format based on your service requirements.

Step 12 Create a mount point.

mkdir *Mount point*

In this example, the **/mnt/sdc** mount point is created.

mkdir /mnt/sdc

 **NOTE**

The **/mnt** directory exists on all Linux systems. If the mount point cannot be created, it may be that the **/mnt** directory has been accidentally deleted. You can run **mkdir -p /mnt/sdc** to create the mount point.

Step 13 Mount the new partition on the created mount point.

mount *Disk partition Mount point*

In this example, the **/dev/vdb1** partition is mounted on **/mnt/sdc**.

mount /dev/vdb1 /mnt/sdc

Step 14 Check the mount result.

df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/vda1       ext4      43G  1.9G  39G   5% /
devtmpfs        devtmpfs  2.0G   0  2.0G   0% /dev
tmpfs           tmpfs     2.0G   0  2.0G   0% /dev/shm
tmpfs           tmpfs     2.0G  9.0M  2.0G   1% /run
tmpfs           tmpfs     2.0G   0  2.0G   0% /sys/fs/cgroup
tmpfs           tmpfs     398M   0  398M   0% /run/user/0
/dev/vdb1       ext4     106G  63M  101G   1% /mnt/sdc
```

You should now see that partition **/dev/vdb1** is mounted on **/mnt/sdc**.

 **NOTE**

After the server is restarted, the disk will not be automatically mounted. You can modify the **/etc/fstab** file to configure automount at startup. For details, see [Configuring Automatic Mounting at System Start](#).

----End

Configuring Automatic Mounting at System Start

The **fstab** file controls what disks are automatically mounted at startup. You can configure the **fstab** file of a that has data. This operation will not affect the existing data.

The following example uses UUIDs to identify disks in the **fstab** file. You are advised not to use device names (like **/dev/vdb1**) to identify disks in the file because device names are assigned dynamically and may change (for example, from **/dev/vdb1** to **/dev/vdb2**) after a stop or start. This can even prevent your from booting up.

NOTE

UUIDs are the unique character strings for identifying partitions in Linux.

Step 1 Query the partition UUID.

blkid *Disk partition*

In this example, the UUID of the **/dev/vdb1** partition is queried.

blkid /dev/vdb1

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# blkid /dev/vdb1
/dev/vdb1: UUID="0b3040e2-1367-4abb-841d-ddb0b92693df" TYPE="ext4"
```

Carefully record the UUID, as you will need it for the following step.

Step 2 Open the **fstab** file using the vi editor.

vi /etc/fstab

Step 3 Press **i** to enter editing mode.

Step 4 Move the cursor to the end of the file and press **Enter**. Then, add the following information:

```
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df /mnt/sdc ext4 defaults 0 2
```

The preceding information is used for reference only. The line starting with **UUID** is the information added. Edit this line from left to right to match the following format:

- **UUID**: The UUID obtained in [Step 1](#).
- **Mount point**: The directory on which the partition is mounted. You can query the mount point using **df -TH**.
- **Filesystem**: The file system format of the partition. You can query the file system format using **df -TH**.
- **Mount option**: The partition mount option. Usually, this parameter is set to **defaults**.
- **Dump**: The Linux dump backup option.
 - **0**: Linux dump backup is not used. Usually, dump backup is not used, and you can set this parameter to **0**.
 - **1**: Linux dump backup is used.
- **fsck**: The fsck option, which means whether to use fsck to check the disk during startup.

- **0**: not use fsck.
- If the mount point is the root partition (/), this parameter must be set to **1**.

If this parameter is set to **1** for the root partition, this parameter for other partitions must start with **2** because the system checks the partitions in the ascending order of the values.

Step 5 Press **Esc**, enter **:wq**, and press **Enter**.

The system saves the configurations and exits the vi editor.

Step 6 Verify that the disk is auto-mounted at startup.

1. Unmount the partition.

umount *Disk partition*

In this example, run the following command:

```
umount /dev/vdb1
```

2. Reload all the content in the **/etc/fstab** file.

```
mount -a
```

3. Query the file system mounting information.

```
mount | grep Mount point
```

In this example, run the following command:

```
mount | grep /mnt/sdc
```

If information similar to the following is displayed, automatic mounting has been configured:

```
root@ecs-test-0001 ~]# mount | grep /mnt/sdc  
/dev/vdb1 on /mnt/sdc type ext4 (rw,relatime,data=ordered)
```

----End

2.3.6 Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2008)

Scenarios

This section uses Windows Server 2008 R2 Standard 64bit to describe how to initialize a data disk whose capacity is larger than 2 TiB. In the following operations, the capacity of the example disk is 3 TiB.

The maximum disk capacity supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Therefore, use the GPT partition style if your disk capacity is larger than 2 TiB. For details, see [Initializing a Windows Data Disk Larger Than 2 TiB \(Windows Server 2008\)](#). To learn more about disk partition styles, see [Scenarios and Disk Partitions](#).

The method for initializing a disk varies slightly depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

NOTICE

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

Prerequisites

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
 - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.

Procedure

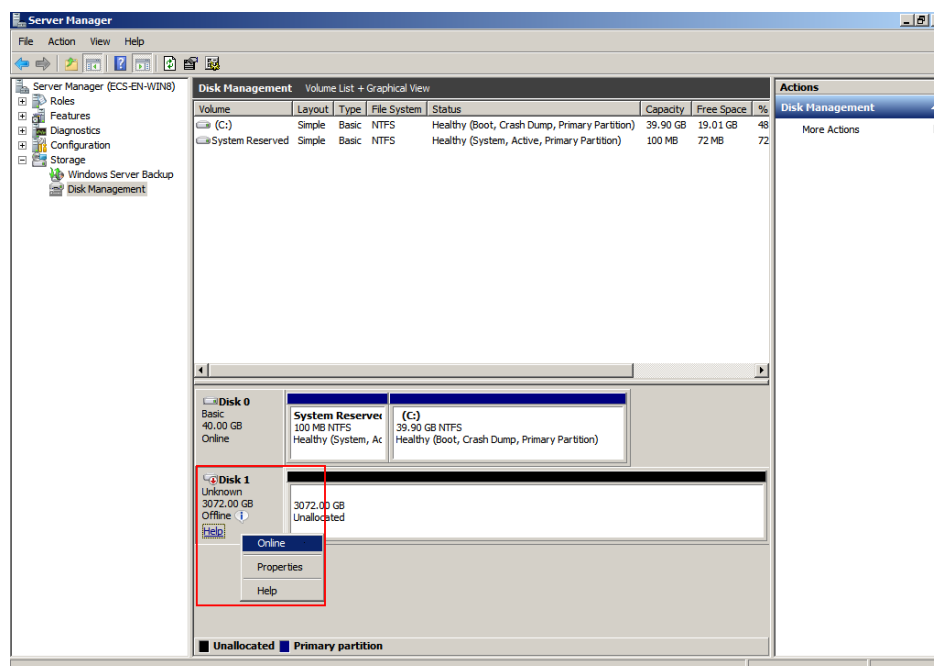
Step 1 On the desktop of the server, click **Start**.

The **Start** window is displayed.

Step 2 Right-click **Computer** and choose **Manage** from the short-cut menu.

The **Server Manager** window is displayed.

Figure 2-23 Server Manager (Windows Server 2008)

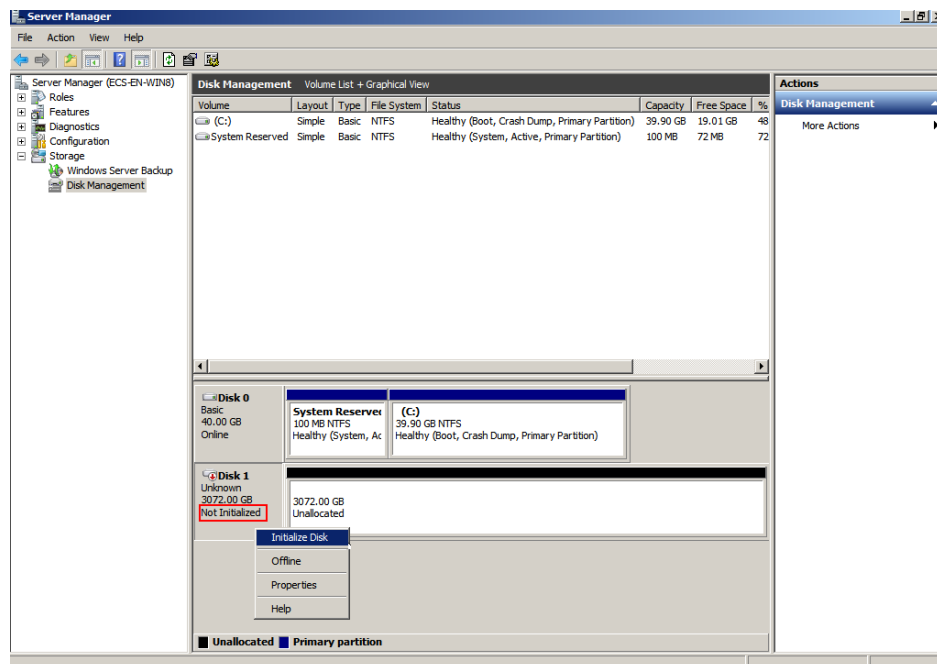


Step 3 Disks are listed in the right pane. If the new disk is offline, bring it online before initializing it.

In the **Disk 1** area, right-click and choose **Online** from the shortcut menu.

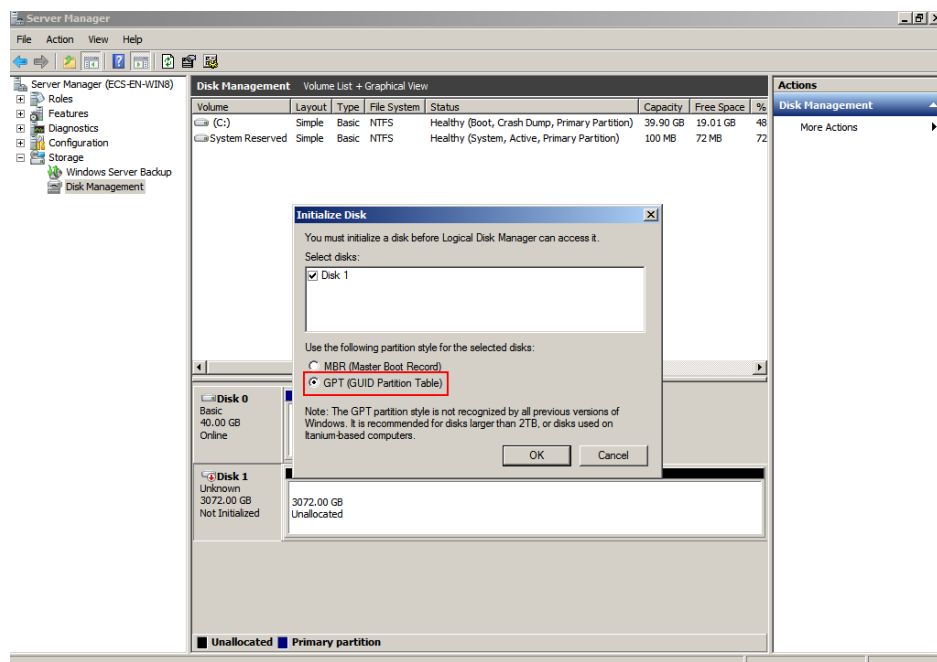
When the status of Disk 1 changes from **Offline** to **Not Initialized**, the disk has been brought online.

Figure 2-24 Bring online succeeded (Windows Server 2008)




Step 4 In the **Disk 1** area, right-click and choose **Initialize Disk** from the shortcut menu. The **Initialize Disk** dialog box is displayed.

Figure 2-25 Initialize Disk (Windows Server 2008)



Step 5 In the **Initialize Disk** dialog box, the to-be-initialized disk is selected. In this example, the disk capacity is larger than 2 TiB. Therefore, select **GPT (GUID Partition Table)** and click **OK**.

The **Server Manager** window is displayed.

Figure 2-26 Server Manager (Windows Server 2008)**NOTICE**


The maximum disk size supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Because an EVS data disk currently supports up to 32 TiB, use GPT if your disk size is larger than 2 TiB.

If the partition style is changed after the disk has been used, all data on the disk will be lost, so take care to select an appropriate partition style when initializing the disk. If you must change the partition style to GPT after a disk has been used, it is recommended that you back up the disk data before the change.

Step 6 Right-click at the unallocated disk space and choose **New Simple Volume** from the shortcut menu.

The **New Simple Volume Wizard** window is displayed.


Figure 2-27 New Simple Volume Wizard (Windows Server 2008)



Step 7 Follow the prompts and click **Next**.

The **Specify Volume Size** page is displayed.

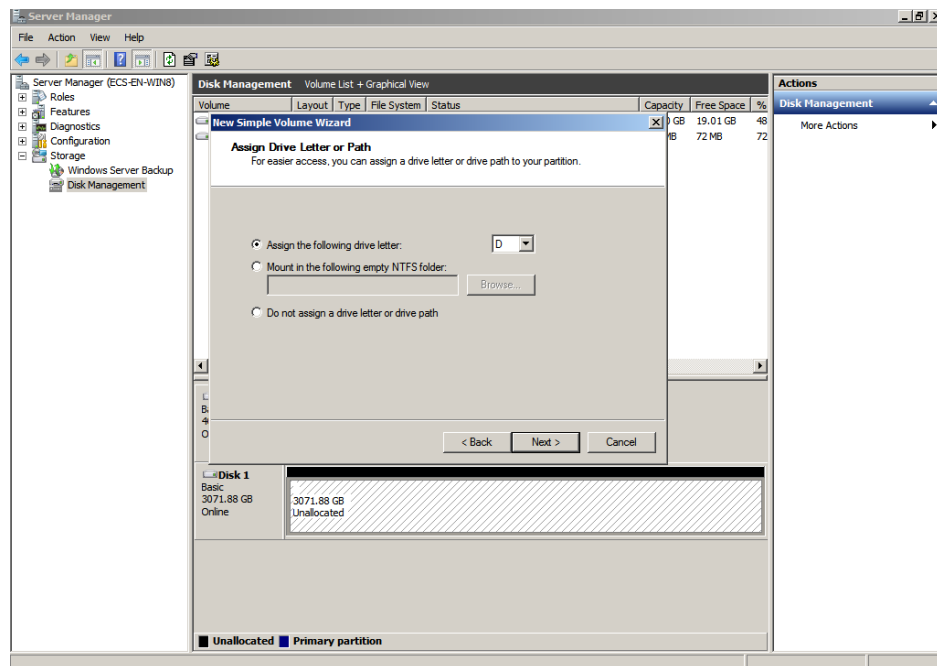
Figure 2-28 Specify Volume Size (Windows Server 2008)



Step 8 Specify the volume size and click **Next**. The system selects the maximum volume size by default. You can specify the volume size as required. In this example, the default setting is used.

The **Assign Drive Letter or Path** page is displayed.

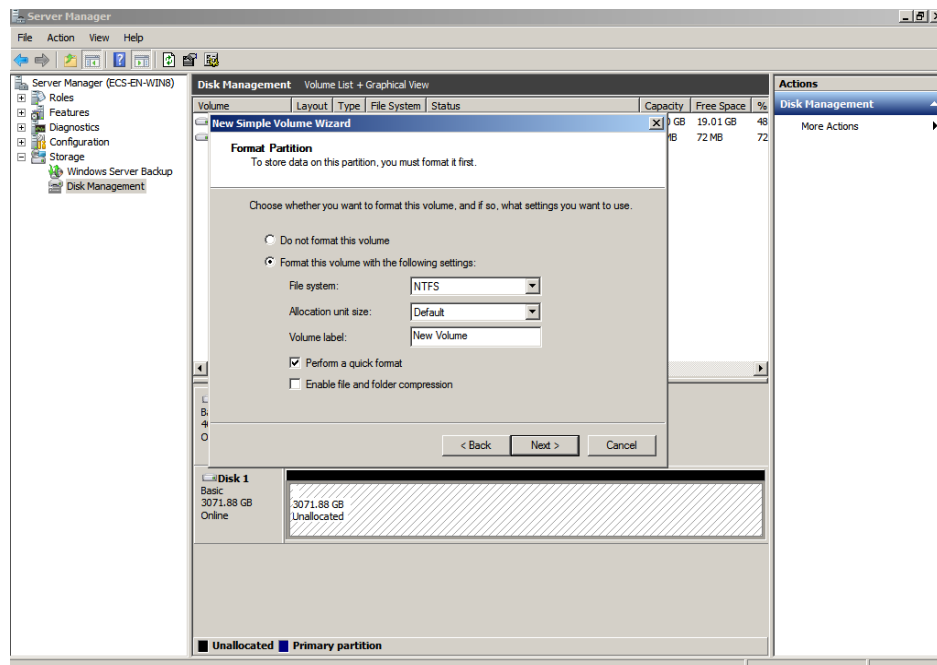
Figure 2-29 Assign Drive Letter or Path (Windows Server 2008)



Step 9 Assign a drive letter or path to your partition and click **Next**. The system assigns drive letter D by default. In this example, the default setting is used.

The **Format Partition** page is displayed.

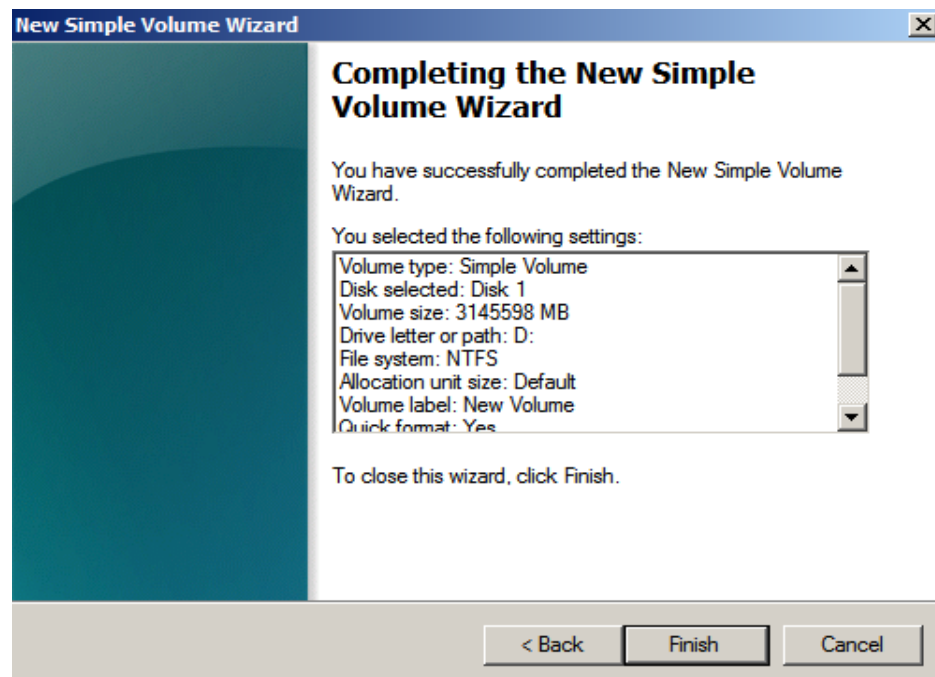
Figure 2-30 Format Partition (Windows Server 2008)



Step 10 Specify format settings and click **Next**. The system selects the NTFS file system by default. You can specify the file system type as required. In this example, the default setting is used.

The **Completing the New Simple Volume Wizard** page is displayed.

Figure 2-31 Completing the New Simple Volume Wizard




NOTICE

The partition sizes supported by file systems vary. Choose an appropriate file system format based on your service requirements.

Step 11 Click **Finish**.

Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished successfully.


Figure 2-32 Disk initialization succeeded (Windows Server 2008)



Step 12 After the volume is created, click  and check whether a new volume appears in **Computer**. In this example, New Volume (D:) is the new volume.

If New Volume (D:) appears, the disk is successfully initialized and no further action is required.

Figure 2-33 Computer (Windows Server 2008)



----End

2.3.7 Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2012)

Scenarios

This section uses Windows Server 2012 R2 Standard 64bit to describe how to initialize a data disk whose capacity is larger than 2 TiB. In the following operations, the capacity of the example disk is 3 TiB.

The maximum disk capacity supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Therefore, use the GPT partition style if your disk capacity is larger than 2 TiB. For details, see [Initializing a Windows Data Disk Larger Than 2 TiB \(Windows Server 2008\)](#). To learn more about disk partition styles, see [Scenarios and Disk Partitions](#).

The method for initializing a disk varies slightly depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

NOTICE

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

Prerequisites

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
 - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.

Procedure


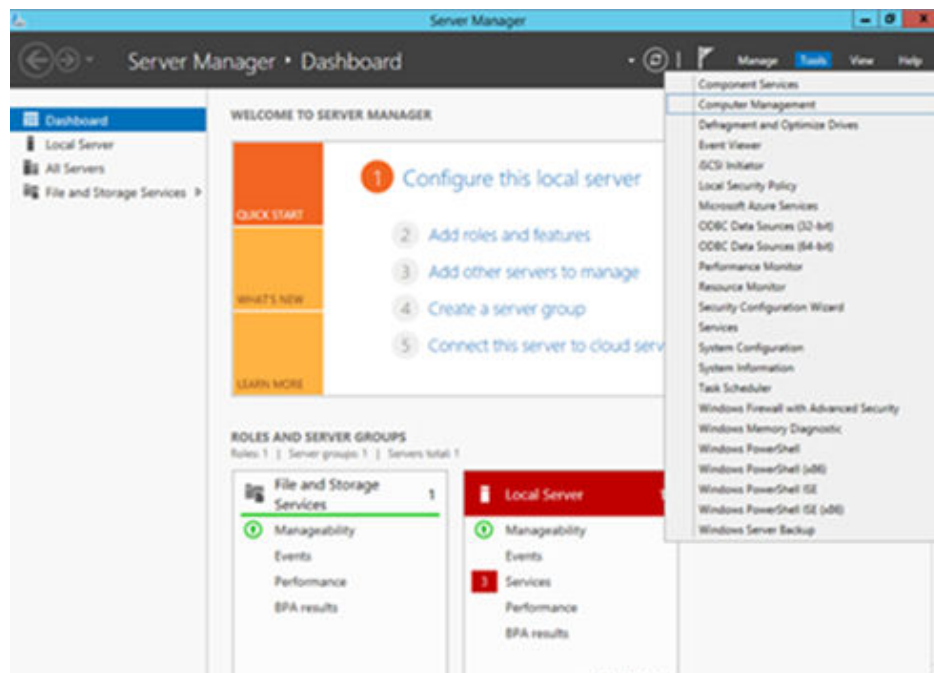
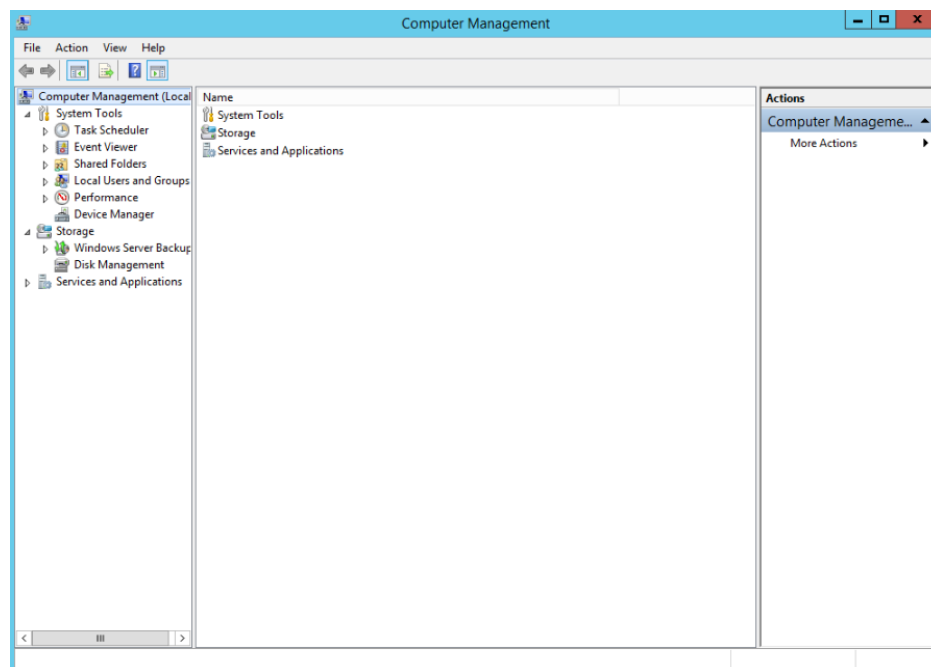
- Step 1** On the desktop of the server, click  in the lower area. The **Server Manager** window is displayed.

Figure 2-34 Server Manager (Windows Server 2012)




Step 2 In the upper right corner, choose **Tools > Computer Management**.
The **Computer Management** window is displayed.

Figure 2-35 Computer Management window (Windows Server 2012)



Step 3 Choose **Storage > Disk Management**.
Disks are displayed in the right pane.

Figure 2-36 Disk Management list (Windows Server 2012)




Step 4 (Optional) If the new disk is offline, bring it online before initializing it.

In the **Disk 1** area, right-click and choose **Online** from the shortcut menu.

When the status of Disk 1 changes from **Offline** to **Not Initialized**, the disk has been brought online.


Figure 2-37 Bring online succeeded (Windows Server 2012)



Step 5 (Optional) In the **Disk 1** area, right-click and choose **Initialize Disk** from the shortcut menu.

The **Initialize Disk** dialog box is displayed.


Figure 2-38 Initialize Disk (Windows Server 2012)



Step 6 In the **Initialize Disk** dialog box, the to-be-initialized disk is selected. In this example, the disk capacity is larger than 2 TiB. Therefore, select **GPT (GUID Partition Table)** and click **OK**.

The **Computer Management** window is displayed.

Figure 2-39 Computer Management (Windows Server 2012)



NOTICE


The maximum disk size supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Because an EVS data disk currently supports up to 32 TiB, use GPT if your disk size is larger than 2 TiB.

If the partition style is changed after the disk has been used, all data on the disk will be lost, so take care to select an appropriate partition style when initializing the disk. If you must change the partition style to GPT after a disk has been used, it is recommended that you back up the disk data before the change.

Step 7 Right-click at the unallocated disk space and choose **New Simple Volume** from the shortcut menu.

The **New Simple Volume Wizard** window is displayed.

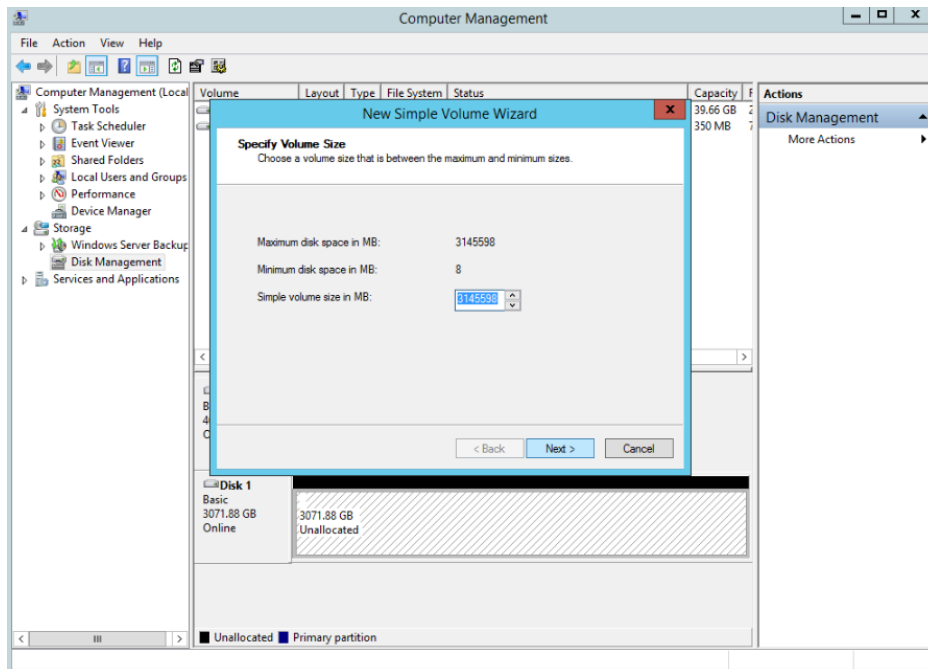
Figure 2-40 New Simple Volume Wizard (Windows Server 2012)



Step 8 Follow the prompts and click **Next**.

The **Specify Volume Size** page is displayed.


Figure 2-41 Specify Volume Size (Windows Server 2012)



Step 9 Specify the volume size and click **Next**. The system selects the maximum volume size by default. You can specify the volume size as required. In this example, the default setting is used.

The **Assign Drive Letter or Path** page is displayed.


Figure 2-42 Assign Drive Letter or Path (Windows Server 2012)



Step 10 Assign a drive letter or path to your partition and click **Next**. The system assigns drive letter D by default. In this example, the default setting is used.

The **Format Partition** page is displayed.


Figure 2-43 Format Partition (Windows Server 2012)



Step 11 Specify format settings and click **Next**. The system selects the NTFS file system by default. You can specify the file system type as required. In this example, the default setting is used.

The **Completing the New Simple Volume Wizard** page is displayed.

Figure 2-44 Completing the New Simple Volume Wizard (Windows Server 2012)



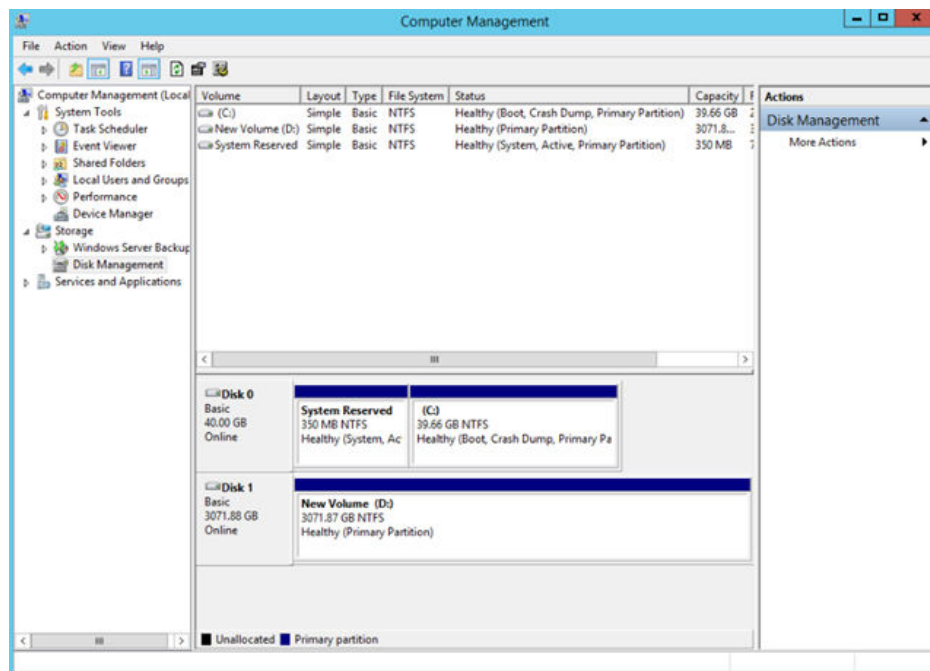
NOTICE

The partition sizes supported by file systems vary. Choose an appropriate file system format based on your service requirements.


Step 12 Click **Finish**.

Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished successfully.

Figure 2-45 Disk initialization succeeded (Windows Server 2012)

**Step 13** After the volume is created, click  and check whether a new volume appears in **This PC**. In this example, New Volume (D:) is the new volume.

If New Volume (D:) appears, the disk is successfully initialized and no further action is required.

Figure 2-46 This PC (Windows Server 2012)

----End

2.3.8 Initializing a Linux Data Disk Larger Than 2 TiB (parted)

Scenarios

This section uses CentOS 7.4 64bit to describe how to use parted to initialize a data disk whose capacity is larger than 2 TiB. In the following operations, the capacity of the example disk is 3 TiB.

The maximum partition size that MBR supports is 2 TiB and that GPT supports is 18 EiB. If the disk size you need to partition is greater than 2 TiB, partition the disk using GPT.

The fdisk partitioning tool is suitable only for MBR partitions, and the parted partitioning tool is suitable for both MBR and GPT partitions. For more information, see [Scenarios and Disk Partitions](#).

The method for initializing a disk varies slightly depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

NOTICE

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

Prerequisites

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.

- For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.

Creating and Mounting a Partition

The following example shows you how a new partition can be created on a new data disk that has been attached to a server. The partition will be created using parted, and GPT will be used. Furthermore, the partition will be formatted using the ext4 file system, mounted on **/mnt/sdc**, and configured to mount automatically at startup.

Step 1 Query information about the new data disk.

lsblk

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda   253:0   0 40G  0 disk
├─vda1 253:1   0  1G  0 part /boot
└─vda2 253:2   0 39G  0 part /
vdb   253:16  0  3T  0 disk
```

In the command output, this server contains two disks. **/dev/vda** and **/dev/vdb**. **/dev/vda** is the system disk, and **/dev/vdb** is the new data disk.

Step 2 Launch parted to partition the new data disk.

parted *New data disk*

In this example, run the following command:

parted /dev/vdb

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted)
```

Step 3 Enter **p** and press **Enter** to view the current disk partition style.

Information similar to the following is displayed:

```
(parted) p
Error: /dev/vdb: unrecognised disk label
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 3299GiB
Sector size (logical/physical): 512B/512B
Partition Table: unknown
Disk Flags:
(parted)
```

In the command output, the **Partition Table** value is **unknown**, indicating that no partition style is set for the new disk.

Step 4 Set the disk partition style.

mklabel *Disk partition style*

The disk partition style can be MBR or GPT. If the disk capacity is greater than 2 TiB, use GPT.

mklabel gpt

NOTICE

The maximum disk size supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Because an EVS data disk currently supports up to 32 TiB, use GPT if your disk size is larger than 2 TiB.

If the partition style is changed after the disk has been used, all data on the disk will be lost, so take care to select an appropriate partition style when initializing the disk. If you must change the partition style to GPT after a disk has been used, it is recommended that you back up the disk data before the change.

Step 5 Enter **p** and press **Enter** to view the disk partition style.

Information similar to the following is displayed:

```
(parted) mklabel gpt
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 3299GiB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start  End  Size  File system  Name  Flags
(parted)
```

Step 6 Enter **unit s** and press **Enter** to set the measurement unit of the disk to sector.

Step 7 Create a new partition.

mkpart *Partition name Start sector End sector*

In this example, run the following command:

mkpart opt 2048s 100%

In this example, one partition is created for the new data disk, starting on **2048** and using **100%** of the rest of the disk. The two values are used for reference only. You can determine the number of partitions and the partition size based on your service requirements.

Information similar to the following is displayed:

```
(parted) mkpart opt 2048s 100%
Warning: The resulting partition is not properly aligned for best performance.
Ignore/Cancel? Ignore
```

If the preceding warning message is displayed, enter **Ignore** to ignore the performance warning.

Step 8 Enter **p** and press **Enter** to print the partition details.

Information similar to the following is displayed:

```
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 6442450944s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
```

```
Number Start End Size File system Name Flags
1 2048s 6442448895s 6442446848s opt
```

Details about the **dev/vdb1** partition are displayed.

Step 9 Enter **q** and press **Enter** to exit parted.

Step 10 View the disk partition information.

lsblk

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
├─vda1 253:1 0 1G 0 part /boot
├─vda2 253:2 0 39G 0 part /
vdb 253:16 0 3T 0 disk
├─vdb1 253:17 0 3T 0 part
```

In the command output, **/dev/vdb1** is the partition you created.

Step 11 Format the new partition with a desired file system format.

mkfs -t *File system format* **/dev/vdb1**

In this example, the **ext4** format is used for the new partition.

mkfs -t ext4 /dev/vdb1

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# mkfs -t ext4 /dev/vdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
201326592 inodes, 805305856 blocks
40265292 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2952790016
24576 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
    102400000, 214990848, 512000000, 550731776, 644972544

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

The formatting takes a period of time. Observe the system running status and do not exit.

NOTICE

The partition sizes supported by file systems vary. Choose an appropriate file system format based on your service requirements.

Step 12 Create a mount point.

mkdir *Mount point*

In this example, the **/mnt/sdc** mount point is created.

mkdir /mnt/sdc

 **NOTE**

The **/mnt** directory exists on all Linux systems. If the mount point cannot be created, it may be that the **/mnt** directory has been accidentally deleted. You can run **mkdir -p /mnt/sdc** to create the mount point.

Step 13 Mount the new partition on the created mount point.

mount *Disk partition Mount point*

In this example, the **/dev/vdb1** partition is mounted on **/mnt/sdc**.

mount /dev/vdb1 /mnt/sdc

Step 14 Check the mount result.

df -TH

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/vda2       ext4      42G   1.5G  38G   4% /
devtmpfs        devtmpfs  2.0G   0   2.0G   0% /dev
tmpfs           tmpfs     2.0G   0   2.0G   0% /dev/shm
tmpfs           tmpfs     2.0G   8.9M  2.0G   1% /run
tmpfs           tmpfs     2.0G   0   2.0G   0% /sys/fs/cgroup
/dev/vda1       ext4      1.1G   153M  801M  17% /boot
tmpfs           tmpfs     398M   0   398M   0% /run/user/0
/dev/vdb1       ext4      3.3T   93M   3.1T   1% /mnt/sdc
```

You should now see that partition **/dev/vdb1** is mounted on **/mnt/sdc**.

----End

Configuring Automatic Mounting at System Start

The **fstab** file controls what disks are automatically mounted at startup. You can configure the **fstab** file of a that has data. This operation will not affect the existing data.

The following example uses UUIDs to identify disks in the **fstab** file. You are advised not to use device names (like **/dev/vdb1**) to identify disks in the file because device names are assigned dynamically and may change (for example, from **/dev/vdb1** to **/dev/vdb2**) after a stop or start. This can even prevent your from booting up.

 **NOTE**

UUIDs are the unique character strings for identifying partitions in Linux.

Step 1 Query the partition UUID.

blkid *Disk partition*

In this example, the UUID of the `/dev/vdb1` partition is queried.

```
blkid /dev/vdb1
```

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# blkid /dev/vdb1
/dev/vdb1: UUID="0b3040e2-1367-4abb-841d-ddb0b92693df" TYPE="ext4"
```

Carefully record the UUID, as you will need it for the following step.

Step 2 Open the `fstab` file using the vi editor.

```
vi /etc/fstab
```

Step 3 Press `i` to enter editing mode.

Step 4 Move the cursor to the end of the file and press **Enter**. Then, add the following information:

```
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df /mnt/sdc          ext4 defaults        0 2
```

The preceding information is used for reference only. The line starting with **UUID** is the information added. Edit this line from left to right to match the following format:

- **UUID:** The UUID obtained in [Step 1](#).
- **Mount point:** The directory on which the partition is mounted. You can query the mount point using `df -TH`.
- **Filesystem:** The file system format of the partition. You can query the file system format using `df -TH`.
- **Mount option:** The partition mount option. Usually, this parameter is set to **defaults**.
- **Dump:** The Linux dump backup option.
 - **0:** Linux dump backup is not used. Usually, dump backup is not used, and you can set this parameter to **0**.
 - **1:** Linux dump backup is used.
- **fsck:** The fsck option, which means whether to use fsck to check the disk during startup.
 - **0:** not use fsck.
 - If the mount point is the root partition (`/`), this parameter must be set to **1**.

If this parameter is set to **1** for the root partition, this parameter for other partitions must start with **2** because the system checks the partitions in the ascending order of the values.

Step 5 Press **Esc**, enter `:wq`, and press **Enter**.

The system saves the configurations and exits the vi editor.

Step 6 Verify that the disk is auto-mounted at startup.

1. Unmount the partition.

```
umount Disk partition
```

In this example, run the following command:

```
umount /dev/vdb1
```

2. Reload all the content in the **/etc/fstab** file.

mount -a

3. Query the file system mounting information.

mount | grep Mount point

In this example, run the following command:

mount | grep /mnt/sdc

If information similar to the following is displayed, automatic mounting has been configured:

```
root@ecs-test-0001 ~]# mount | grep /mnt/sdc  
/dev/vdb1 on /mnt/sdc type ext4 (rw,relatime,data=ordered)
```

----End

3 Instances

3.1 Viewing ECS Information

3.1.1 Viewing ECS Creation Statuses

Scenarios

After submitting the request for creating an ECS, you can view the creation status. This section describes how to view the creation status of an ECS.

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. After creating an ECS, view the creation status above the ECS list beside the common operations (**Start**, **Stop**, **Restart**, and **More**).
4. Click the number displayed above **Creating** and view task details.

NOTE

- An ECS that is being created is in one of the following states:
 - **Creating**: The ECS is being created.
 - **Faulty**: Creating the ECS failed. In such a case, the system automatically rolls back the task and displays an error code on the GUI, for example, **Ecs.0013 Insufficient EIP quota**.
 - **Running**: The request of creating the ECS has been processed, and the ECS is running properly. An ECS in this state can provide services for you.
See [How Do I Handle Error Messages Displayed on the Management Console?](#) for troubleshooting.
- If you find that the task status area shows an ECS creation failure but the ECS has been created successfully and displayed in the ECS list, see [Why Does the Failures Area Show an ECS Creation Failure But the ECS List Displays the Created ECS?](#)

3.1.2 Viewing Failed Tasks

Scenarios

You can view the details of failed task (if any) in the **Failures** area, including the task names and statuses. This section describes how to view failures.

Failure Types

Table 3-1 lists the types of failures that can be recorded in the **Failures** area.

Table 3-1 Failure types

Failure Type	Description
Creation failures	A task failed. For a failed task, the system rolls back the task and displays an error code, for example, Ecs.0013 Insufficient EIP quota . For details about how to handle the failure, see How Do I Handle Error Messages Displayed on the Management Console?
Operation failures	<ul style="list-style-type: none">• Modifying ECS specifications If an ECS specifications modification failed, this operation is recorded in Failures.• Automatic recovery enabled during ECS creation Automatic recovery is enabled during ECS creation. After the ECS is created, if the system fails to enable automatic recovery, this operation is recorded in Failures.

Procedure

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. View **Failures** on the right side of common operations.

 **NOTE**

- If **Failures** is not displayed on the management console, the following tasks have been successfully executed:
- The ECS specifications are modified.
 - Automatic recovery is enabled during ECS creation.
4. Click the number displayed in the **Failures** area to view task details.
 - **Creation Failures**: show the failed ECS creation tasks.
 - **Operation Failures**: show the tasks with failed operations and error codes, which help you troubleshoot the faults.

3.1.3 Viewing ECS Details (List View)

Scenarios

After obtaining ECSs, you can view and manage them on the management console. This section describes how to view detailed ECS configurations, including its name, image, system disk, data disks, VPC, NIC, security group, and EIP.

To view the private IP address of an ECS, view it on the **Elastic Cloud Server** page.

Procedure


1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
The **Elastic Cloud Server** page is displayed. On this page, you can view your ECSs and the basic information about the ECSs, such as their specifications, images, and IP addresses.
3. In the search box above the ECS list, select a filter (such as ECS name, ID, or private IP address), enter the corresponding information, and press **Enter**.
4. Click the name of the target ECS.
The page providing details about the ECS is displayed.
5. View the ECS details.
You can click the tabs and perform operations. For details, see [Changing a Security Group](#), [Attaching a Network Interface](#), and [Binding an EIP](#).

3.1.4 Exporting ECS Information


Scenarios

The information of all ECSs under your account can be exported in a CSV file to a local directory. The file includes the IDs, private IP addresses, and EIPs of your ECSs.

Procedure

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. In the upper right corner above the ECS list, click .
The system will automatically export all ECSs in the current region under your account to a local directory.

NOTE

- To export certain ECSs, select the target ECSs and click  in the upper right corner of the page.
4. In the default download path, view the exported ECS information.

3.2 Logging In to a Windows ECS

3.2.1 Login Overview

Constraints

- Only a running ECS can be logged in.
- The username for logging in to a Windows ECS is **Administrator**.
- Certain G-series ECSs do not support remote login from the console. If you need to remotely log in to the ECSs, install the VNC server on them. For details, see [GPU-accelerated ECSs](#). You are suggested to log in to the ECSs using MSTSC.

Login Modes

You can choose from a variety of login modes based on your local OS type.

Table 3-2 Windows login modes

ECS OS	Local OS	Connection Method	Requirement
Windows	Windows	Use MSTSC. Click Start on the local computer. In the Search programs and files text box, enter mstsc to open the Remote Desktop Connection dialog box. For details, see Remotely Logging In to a Windows ECS (Using MSTSC) .	The target ECS has had an EIP bound.
	Linux	Install a remote connection tool, for example, rdesktop. For details, see Remotely Logging In to a Windows ECS (from a Linux Computer) .	
	macOS	Install a remote connection tool, for example, Microsoft Remote Desktop on the macOS. For details, see Remotely Logging In to a Windows ECS (from a macOS Server) .	

ECS OS	Local OS	Connection Method	Requirement
	Mobile terminal	Install a remote connection tool, for example, Microsoft Remote Desktop. For details, see Remotely Logging In to a Windows ECS (from a Mobile Terminal) .	
	Windows	Through the management console. For details, see Remotely Logging In to a Windows ECS (Using VNC) .	No EIP is required.

3.2.2 Remotely Logging In to a Windows ECS (Using VNC)

Scenarios

This section describes how to use VNC provided on the management console to log in to an ECS.


Constraints

- The remote login function is implemented using customized ports. Therefore, before attempting to log in remotely, ensure that the port to be used is not blocked by the firewall. For example, if the remote login link is xxx:8002, ensure that port 8002 is not blocked by the firewall.
- If the client OS uses a local proxy and the firewall port cannot be configured on the local proxy, disable the proxy mode and then try logging in remotely.

Logging In to a Windows ECS

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Obtain the password for logging in to the ECS.
Before logging in to the ECS, you must have the login password.
 - If your ECS uses password authentication, log in to the ECS using the password you configured during the ECS creation.
4. In the **Operation** column of the target ECS, click **Remote Login**.
5. (Optional) When the system displays "Press CTRL+ALT+DELETE to log on", click **Send CtrlAltDel** in the upper part of the remote login page to log in to the ECS.

Figure 3-1 Send CtrlAltDel



6. Enter the ECS password as prompted.

3.2.3 Remotely Logging In to a Windows ECS (Using MSTSC)

Scenarios

This section describes how to use the remote login tool MSTSC to log in to a Windows ECS from a local computer.

Prerequisites

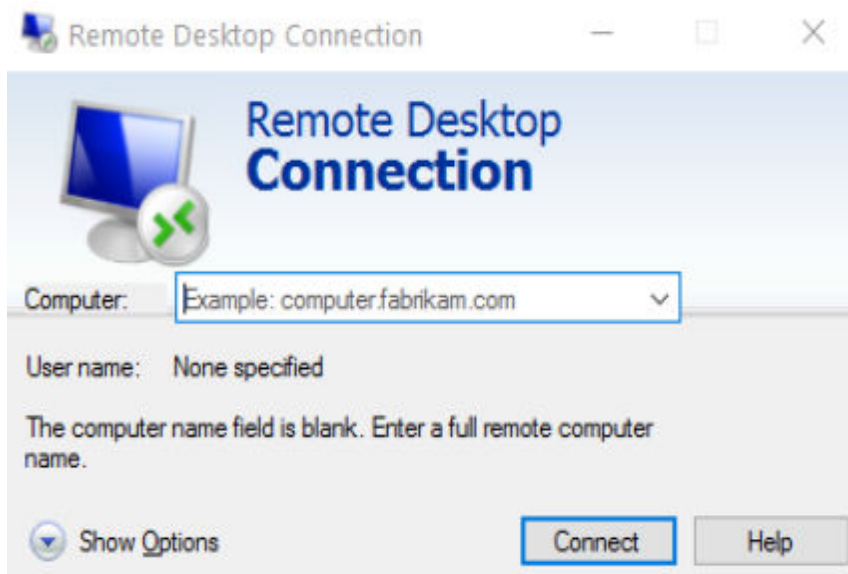
- The target ECS is running.
- You have bound an EIP to the ECS. For details, see [Binding an EIP](#).
- Access to port 3389 is allowed in the inbound direction of the security group to which the ECS belongs. For details, see [Configuring Security Group Rules](#).
- The network connection between the login tool and the target ECS is normal. For example, the default port 3389 is not blocked by the firewall.
- Remote Desktop Protocol (RDP) needs to be enabled on the target ECS. For ECSs created using public images, RDP has been enabled by default. For instructions about how to enable RDP, see [Enabling RDP](#).

Logging In to a Windows ECS Using MSTSC

If your local server runs Windows, you can use the remote desktop connection tool MSTSC delivered with the Windows OS to log in to a Windows ECS.

1. Click the start menu on the local server.
2. In the **Search programs and files** text box, enter **mstsc**.
3. In the **Remote Desktop Connection** dialog box, click **Show Options**.

Figure 3-2 Show Options

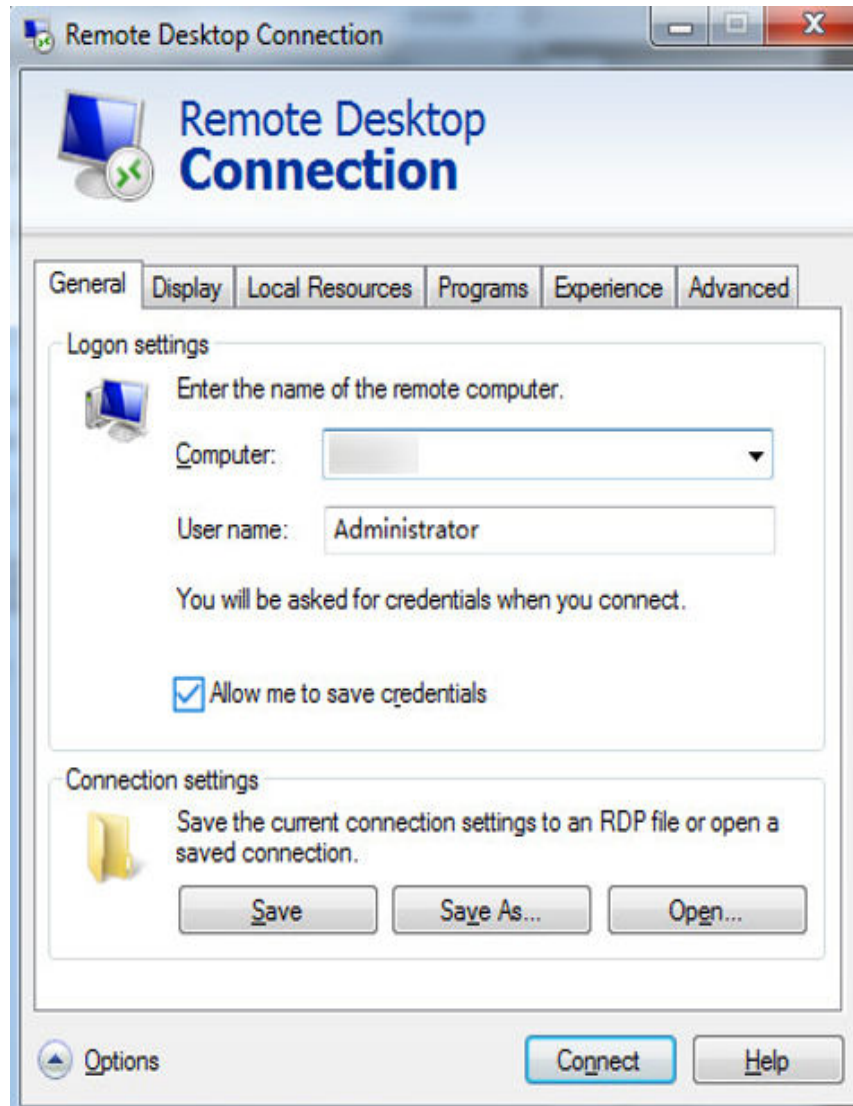


4. Enter the EIP and username (**Administrator** by default) of the target ECS.

 **NOTE**

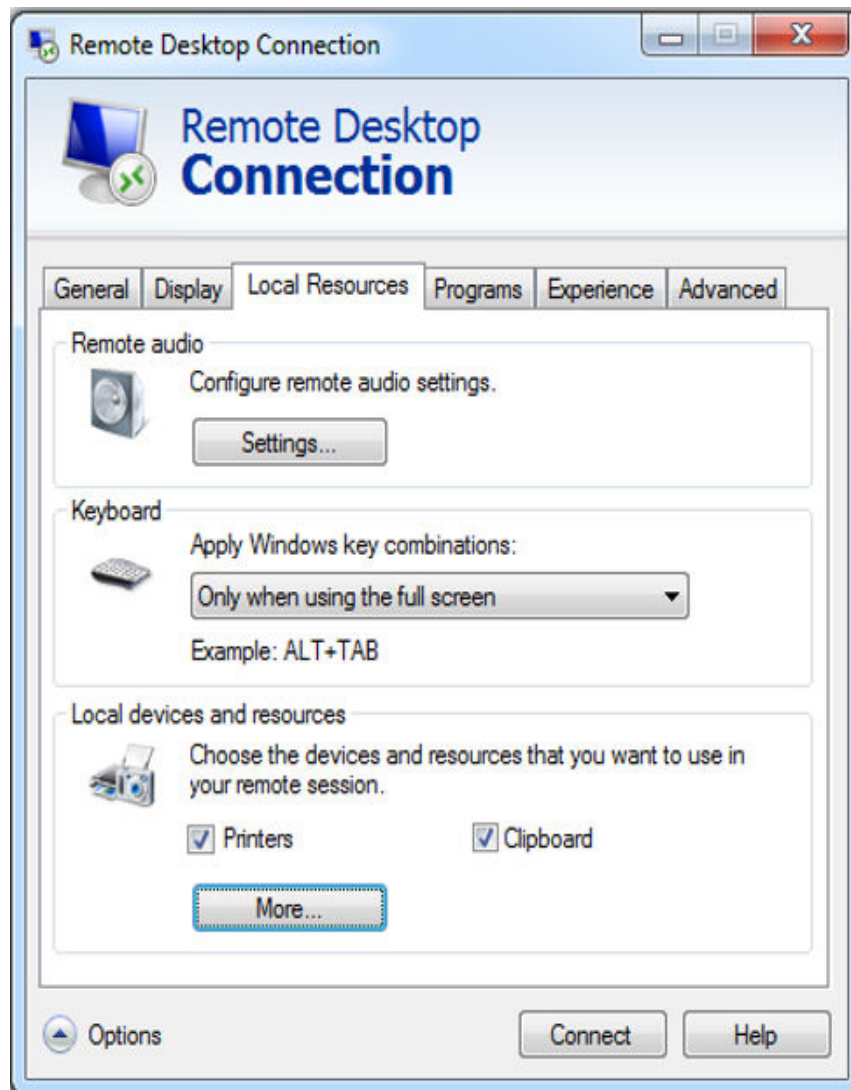
If you do not want to enter the username and password in follow-up logins, select **Allow me to save credentials**.

Figure 3-3 Remote Desktop Connection




5. (Optional) To use local server resources in a remote session, configure parameters on the **Local Resources** tab.
To copy data from the local server to your ECS, select **Clipboard**.

Figure 3-4 Clipboard




To copy files from the local server to your ECS, click **More** and select your desired disks.

Figure 3-5 Drives



6. (Optional) Click the **Display** tab and then adjust the size of the remote desktop.

Figure 3-6 Adjusting the size of the desktop

7. Click **Connect** and enter the login password as prompted to log in to the ECS. To ensure system security, change the login password after you log in to the ECS for the first time.
8. (Optional) Copy local files to the Windows ECS using clipboard. If the file size is greater than 2 GB, an error will occur. To resolve this issue, see [troubleshooting cases](#).

Enabling RDP

For your first login, use VNC to log in and enable RDP for your ECS. Then, use MSTSC to log in.

NOTE


By default, RDP has been enabled on the ECSs created using a public image.

1. Log in to the Windows ECS using VNC. For details, see [Remotely Logging In to a Windows ECS \(Using VNC\)](#).

2. Click **Start** in the task bar and choose **Control Panel > System and Security > System > Remote settings**.

The **System Properties** dialog box is displayed.

Figure 3-7 System Properties



3. Click the **Remote** tab and select **Allow remote connections to this computer**.
4. Click **OK**.

3.2.4 Remotely Logging In to a Windows ECS (from a Linux Computer)

Scenarios

This section describes how to log in to a Windows ECS from a Linux computer.

Prerequisites

- The target ECS is running.
- The ECS must have an EIP bound.
- Access to port 3389 is allowed in the inbound direction of the security group to which the ECS belongs.
- Data can be exchanged between the login tool and the target ECS. For example, the default port 3389 is not blocked by the firewall.

- RDP has been enabled on the target ECS. By default, RDP has been enabled on the ECSs created using a public image. For instructions about how to enable RDP, see [Enabling RDP](#).

Procedure

To log in to a Windows ECS from a local Linux computer, use a remote access tool, such as rdesktop.

1. Run the following command to check whether rdesktop has been installed on the ECS:

rdesktop

If the message "command not found" is displayed, rdesktop is not installed. In such a case, obtain the rdesktop installation package at the [official rdesktop website](#).

2. Run the following command to log in to the ECS:

```
rdesktop -u Username -p Password -g Resolution EIP
```

For example, run **rdesktop -u administrator -p password -g 1024*720 121.xx.xx.xx**.

Table 3-3 Parameters in the remote login command

Parameter	Description
-u	Username, which defaults to Administrator for Windows ECSs
-p	Password for logging in to the Windows ECS
-f	Full screen by default, which can be switched using Ctrl+Alt+Enter
-g	Resolution, which uses an asterisk (*) to separate numbers. This parameter is optional. If it is not specified, the remote desktop is displayed in full screen by default, for example, 1024*720 .
EIP	EIP of the Windows ECS to be remotely logged in. Replace it with the EIP bound to your Windows ECS.

Enabling RDP

For your first login, use VNC to log in and enable RDP for your ECS. Then, use MSTSC to log in.

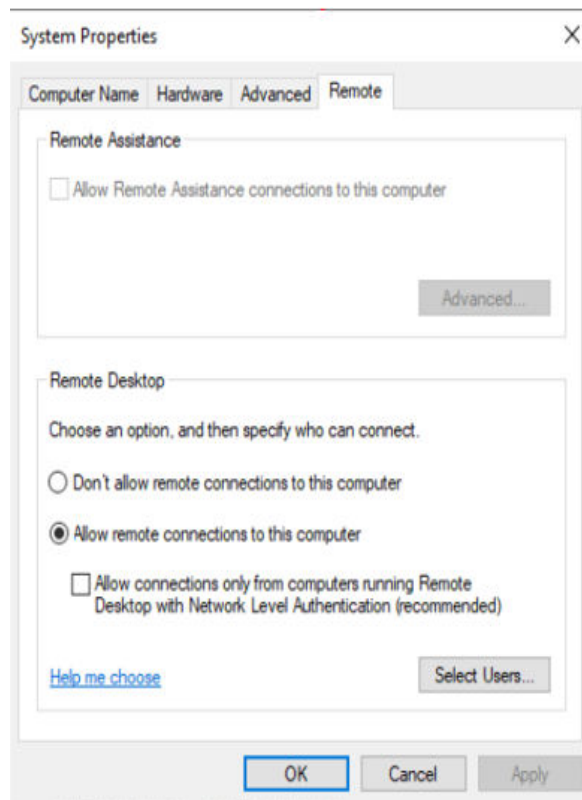
NOTE

By default, RDP has been enabled on the ECSs created using a public image.

1. Log in to the Windows ECS using VNC.
For details, see [Remotely Logging In to a Windows ECS \(Using VNC\)](#).
2. Click **Start** in the task bar and choose **Control Panel > System and Security > System > Remote settings**.

The **System Properties** dialog box is displayed.

Figure 3-8 System Properties



3. Click the **Remote** tab and select **Allow remote connections to this computer**.
4. Click **OK**.

3.2.5 Remotely Logging In to a Windows ECS (from a Mobile Terminal)

Scenarios

This section describes how to log in to an ECS running Windows Server 2012 R2 DataCenter 64bit from a mobile terminal via the Microsoft Remote Desktop client.

Prerequisites

- The target ECS is running.
- You have obtained the username and password for logging in to the ECS.
- You have bound an EIP to the ECS. For details, see [Binding an EIP](#).
- Access to port 3389 is allowed in the inbound direction of the security group to which the ECS belongs. For details, see [Configuring Security Group Rules](#).
- [Microsoft Remote Desktop](#) has been installed on the mobile terminal.

Procedure


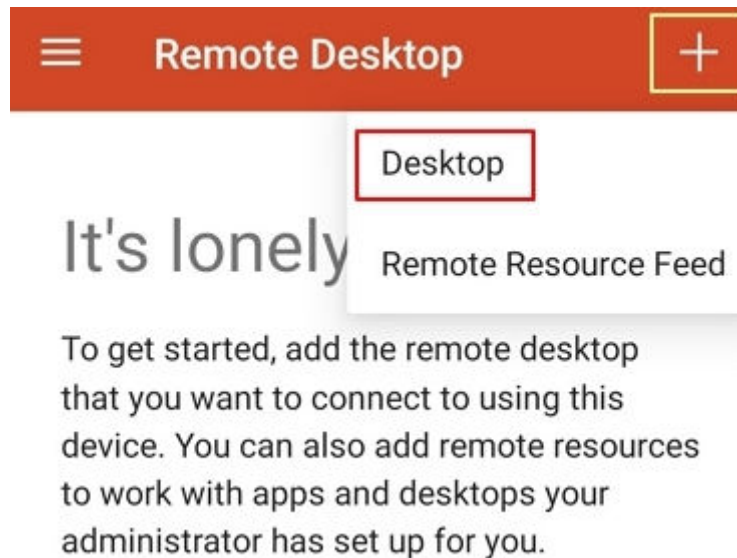
1. Start the Microsoft Remote Desktop client.
2. In the upper right corner of the **Remote Desktop** page, tap  and select **Desktop**.

Figure 3-9 Remote Desktop



3. On the **Add desktop** page, set login information and tap **SAVE**.
 - **PC name:** Enter the EIP bound to the target Windows ECS.
 - Perform the following operations to set **User name:**
 - i. Tap **User name** and select **Add user account** from the drop-down list.
The **Add user account** dialog box is displayed.
 - ii. Enter the username **administrator** and password for logging in to the Windows ECS and click **SAVE**.

Figure 3-10 Setting the login information

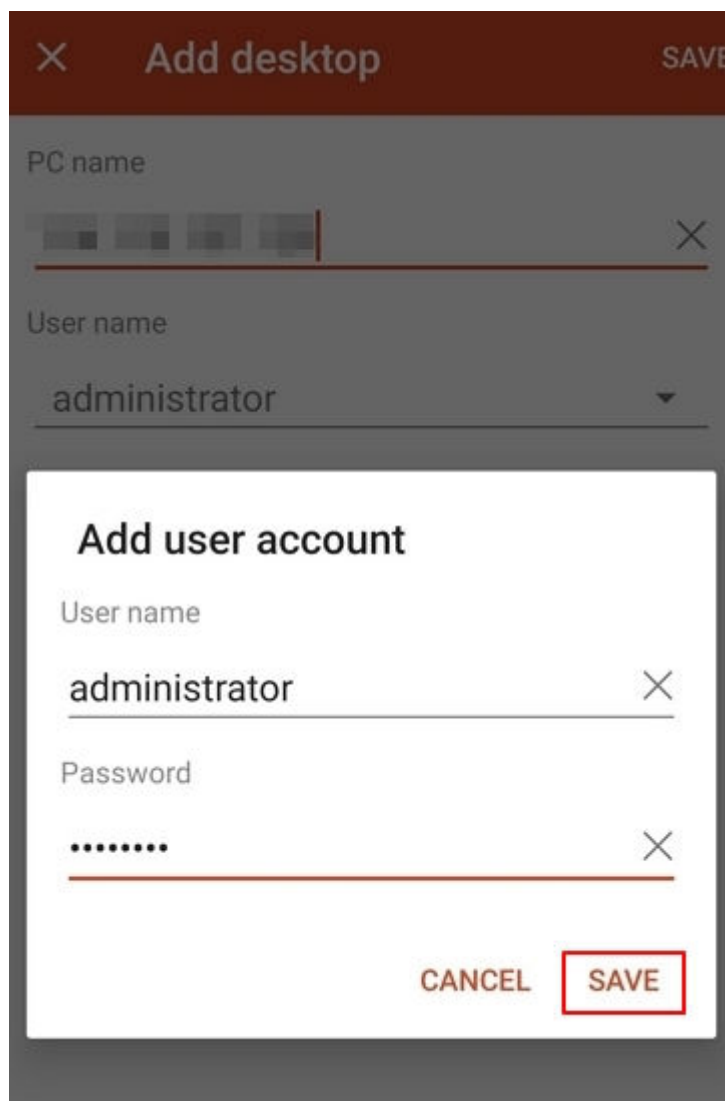



Figure 3-11 Saving the settings




4. On the **Remote Desktop** page, tap the icon of the target Windows ECS.

Figure 3-12 Logging in to the Windows ECS

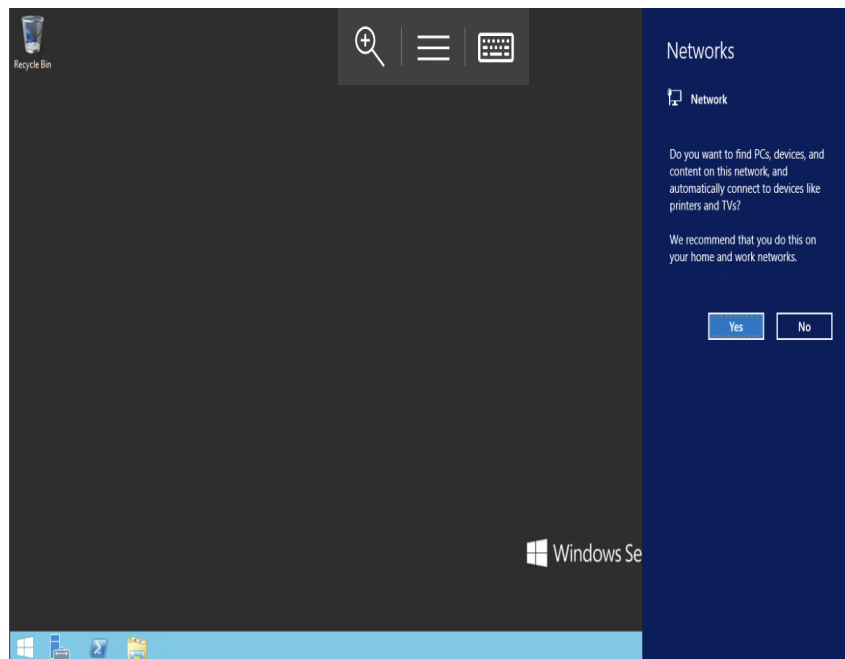


5. Confirm the information and tap **CONNECT**.

Figure 3-13 CONNECT



You have logged in to the Windows ECS.

Figure 3-14 Successful login

3.2.6 Remotely Logging In to a Windows ECS (from a macOS Server)

Scenarios

This section describes how to use a remote login tool to log in to a Windows ECS from a macOS server. In this section, the remote login tool Microsoft Remote Desktop for Mac and the ECS running Windows Server 2012 R2 Data Center 64bit are used as an example.

Prerequisites


- The target ECS is running.
- You have obtained the username and password for logging in to the ECS.
- You have bound an EIP to the ECS. For details, see [Binding an EIP](#).
- Access to port 3389 is allowed in the inbound direction of the security group to which the ECS belongs. For details, see [Configuring Security Group Rules](#).
- The remote access tool supported by Mac, such as Microsoft Remote Desktop for Mac has been installed. For details, see [Download Microsoft Remote Desktop for Mac](#).

Microsoft stopped providing the link for downloading the Remote Desktop client. You can download the beta version by visiting [Microsoft Remote Desktop Beta](#).

Procedure

1. Start Microsoft Remote Desktop.
2. Click **Add Desktop**.

Figure 3-15 Add Desktop



3. On the **Add PC** page, set login information.
 - **PC name:** Enter the EIP bound to the target Windows ECS.
 - **User account:** Select **Add user account** from the drop-down list. The **Add user account** dialog box is displayed.
 - i. Enter the username **administrator** and password for logging in to the Windows ECS and click **Add**.

Figure 3-16 Add user account

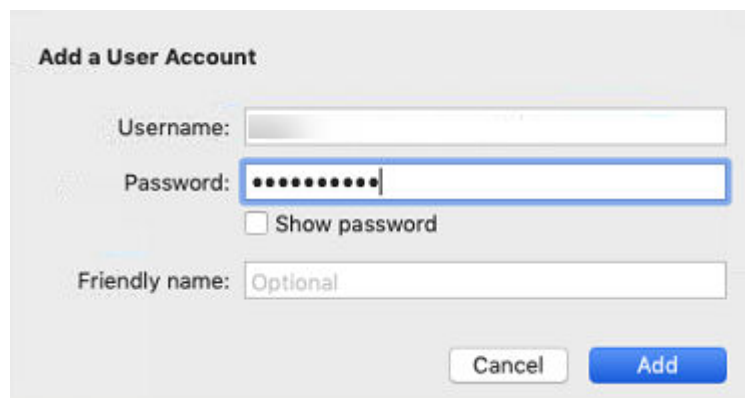
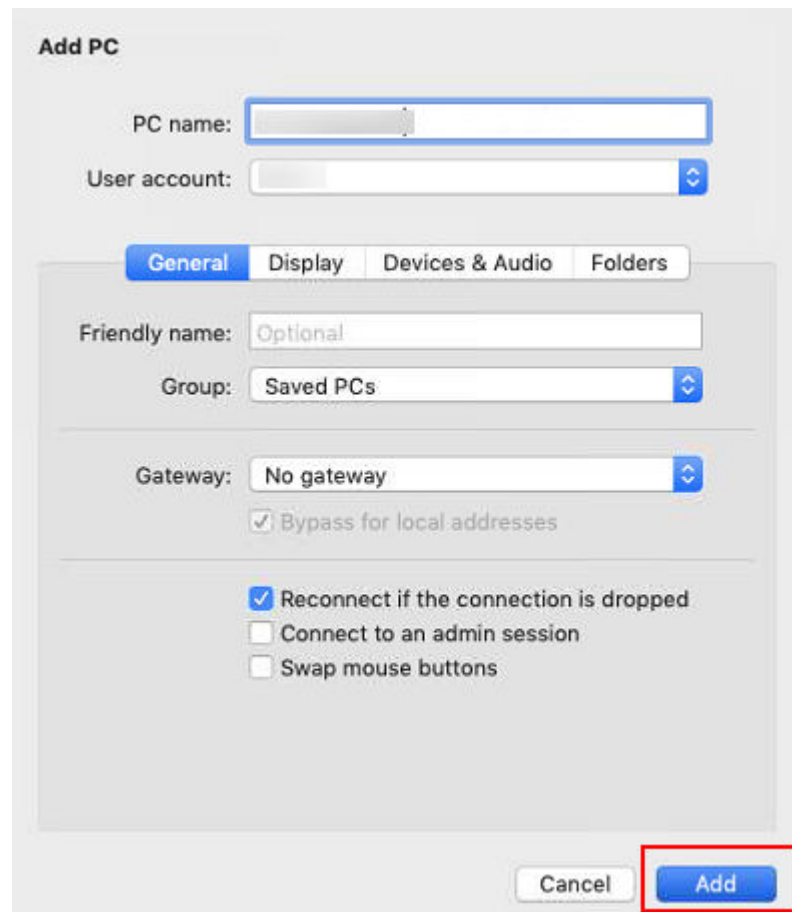
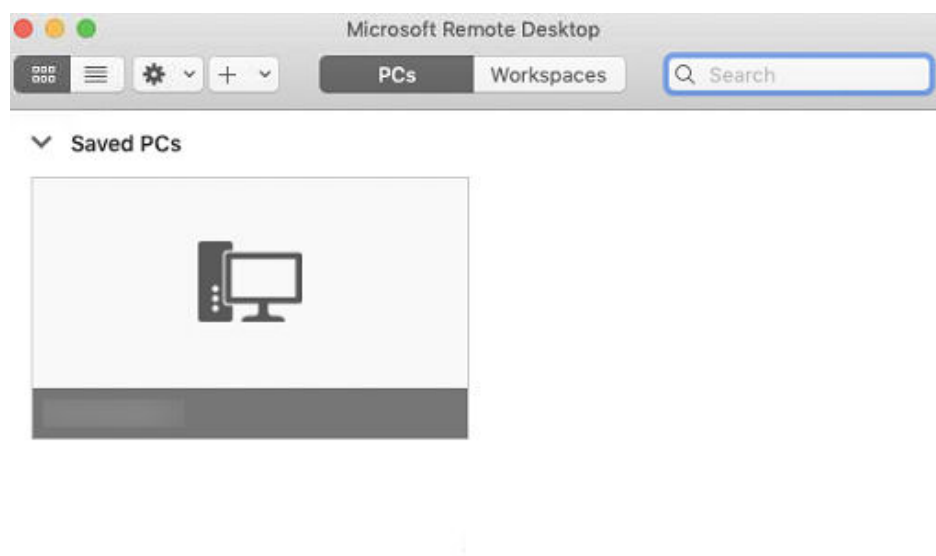


Figure 3-17 Add PC



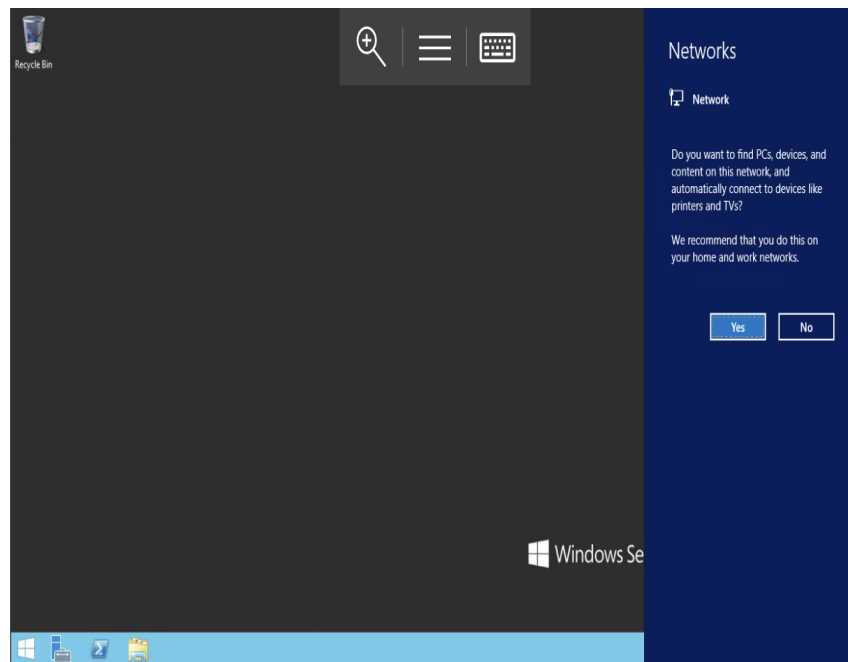
4. On the **Remote Desktop** page, double-click the icon of the target Windows ECS.

Figure 3-18 Double-click for login



5. Confirm the information and click **Continue**.
You have logged in to the Windows ECS.

Figure 3-19 Successful login



3.3 Logging In to a Linux ECS

3.3.1 Login Overview

Constraints

- Only a running ECS can be logged in.
- The username for logging in to a Linux ECS is **root**.

Login Modes

You can choose from a variety of login modes based on your local OS type.

Table 3-4 Linux ECS login modes

ECS OS	Local OS	Connection Method	Requirement
Linux	Windows	Use a remote login tool, such as PuTTY or Xshell. <ul style="list-style-type: none"> • Password-authenticated: Logging In to a Linux ECS from a Local Windows Server 	The target ECS has an EIP bound.
	Linux	Run commands. <ul style="list-style-type: none"> • Password-authenticated: Logging In to a Linux ECS from a Local Linux Server 	

ECS OS	Local OS	Connection Method	Requirement
	Mobile terminal	Use an SSH client tool, such as Termius or JuiceSSH, to log in to the ECS. Remotely Logging In to a Linux ECS (from a Mobile Terminal)	
	macOS	Use the terminal included in the macOS. Remotely Logging In to a Linux ECS (from a macOS Server)	
	Windows	Use the remote login function available on the management console. For details, see Remotely Logging In to a Linux ECS (Using VNC) .	No EIP is required.

3.3.2 Remotely Logging In to a Linux ECS (Using VNC)

Scenarios

This section describes how to use VNC provided on the management console to log in to an ECS.

For instructions about how to copy and paste data on VNC pages after the ECS login, see [Follow-up Procedure](#).

Constraints

- The remote login function is implemented using customized ports. Therefore, before attempting to log in remotely, ensure that the port to be used is not blocked by the firewall. For example, if the remote login link is xxx:8002, ensure that port 8002 is not blocked by the firewall.
- If the client OS uses a local proxy and the firewall port cannot be configured on the local proxy, disable the proxy mode and then try logging in remotely.

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Obtain the password for logging in to the ECS.
Before logging in to the ECS, you must have the login password.
 - If your ECS uses password authentication, log in to the ECS using the password you configured during the ECS creation.
4. In the **Operation** column of the target ECS, click **Remote Login**.

- (Optional) When the system displays "Press CTRL+ALT+DELETE to log on", click **Send CtrlAltDel** in the upper part of the remote login page to log in to the ECS.

 **NOTE**

Do not press **CTRL+ALT+DELETE** on the physical keyboard because this operation does not take effect.

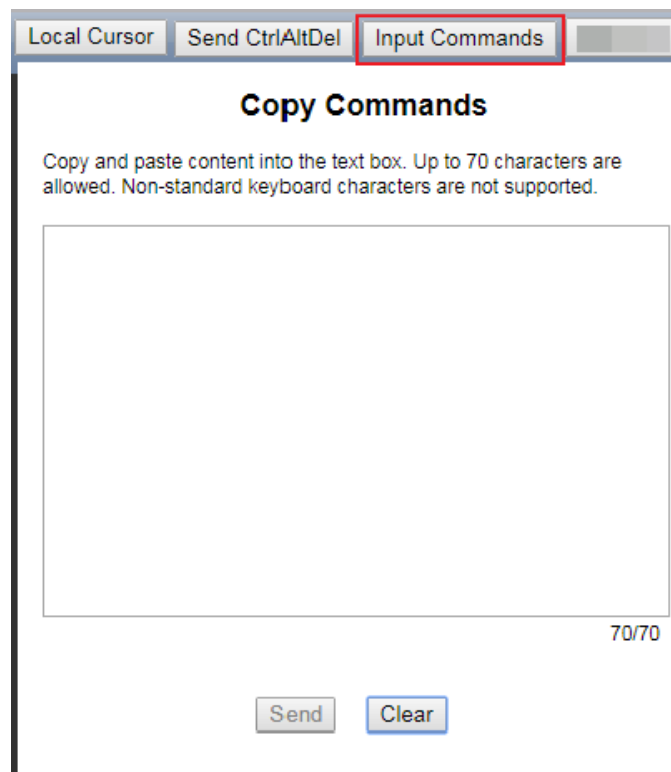
- Enter the ECS password as prompted.

Follow-up Procedure

Local commands can be copied to an ECS. To do so, perform the following operations:

- Log in to the ECS using VNC.
- Click **Input Commands** in the upper right corner of the page.

Figure 3-20 Input Commands



- Press **Ctrl+C** to copy data from the local computer.
- Press **Ctrl+V** to paste the local data to the **Paste & Send** window.
- Click **Send**.
Send the copied data to the CLI.

 NOTE

There is a low probability that data is lost when you use Input Commands on the VNC page of a GUI-based Linux ECS. This is because the number of ECS vCPUs fails to meet GUI requirements. In such a case, it is a good practice to send a maximum of 5 characters at a time or switch from GUI to CLI (also called text interface), and then use the command input function.

3.3.3 Remotely Logging In to a Linux ECS (Using an SSH Password)

Scenarios

This section describes how to remotely log in to a Linux ECS using an SSH password from a Windows and a Linux server, respectively.

NOTICE

Logging in to a Linux ECS using SSH password authentication is disabled by default. If you require password authentication, configure it after logging in to the ECS. To ensure system security, reset the common user password for logging in to the Linux ECS after configuring SSH password authentication.

Prerequisites

- The target ECS is running.
- You have bound an EIP to the ECS. For details, see [Binding an EIP](#).
- Access to port 22 is allowed in the inbound direction of the security group to which the ECS belongs. For details, see [Configuring Security Group Rules](#).
- The network connection between the login tool (PuTTY) and the target ECS is normal. For example, the default port 22 is not blocked by the firewall.

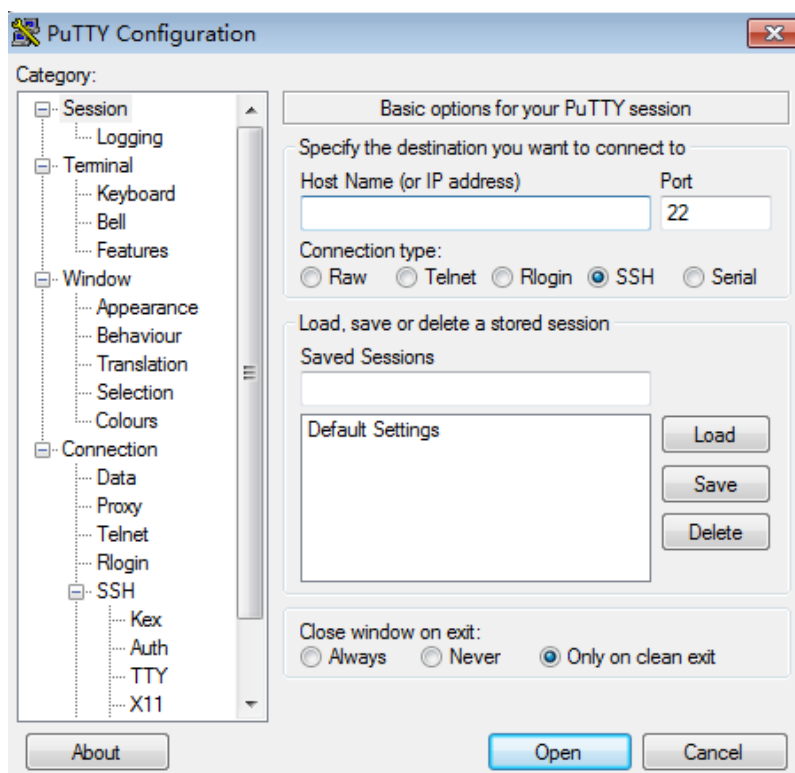
Logging In to a Linux ECS from a Local Windows Server

To log in to a Linux ECS from a local Windows server, perform the operations below.

The following operations use PuTTY as an example to log in to the ECS.

1. Visit the following website and download PuTTY and PuTTYgen:
<https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html>
2. Run PuTTY.
3. Choose **Session**.
 - a. **Host Name (or IP address)**: Enter the EIP bound to the ECS.
 - b. **Port**: Enter **22**.
 - c. **Connection type**: Click **SSH**.
 - d. **Saved Sessions**: Enter the task name, which can be clicked for remote connection when you use PuTTY next time.

Figure 3-21 Session



4. Choose **Window**. Then, select **UTF-8** for **Received data assumed to be in which character set:** in **Translation**.
5. Click **Open**.
If you log in to the ECS for the first time, PuTTY displays a security warning dialog box, asking you whether to accept the ECS security certificate. Click **Yes** to save the certificate to your local registry.
6. After the SSH connection to the ECS is set up, enter the username and password as prompted to log in to the ECS.

Logging In to a Linux ECS from a Local Linux Server

To log in to a Linux ECS from a local Linux server, perform the operations below.

1. On the Linux CLI, run the following command to log in to the ECS:

```
ssh xx.xx.xx.xx
```

NOTE

xx.xx.xx.xx indicates the EIP bound to the ECS.

2. Verify the SSH fingerprint of the ECS and enter **yes**.
The authenticity of host '**xx.xx.xx.xx** (**xx.xx.xx.xx**)' can't be established.
ECDSA key fingerprint is SHA256:rnKuzrUSYS03MCoaXXXXXXXXXXXXXXXXXXXXXXXXXXXX.
ECDSA key fingerprint is MD5:cf:64:5b:5e:74:30:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx.
Are you sure you want to continue connecting (yes/no)? **yes**
Warning: Permanently added '**xx.xx.xx.xx**' (ECDSA) to the list of known hosts.
3. Enter the password for logging in to ECS.

3.3.4 Remotely Logging In to a Linux ECS (from a Mobile Terminal)

Scenarios

This section describes how to access a Linux ECS from a mobile terminal.

- For instructions about how to log in to a Linux ECS from an iOS terminal through iTerminal-SSH Telnet, see [Logging In to a Linux ECS from an iOS Terminal](#).
- For instructions about how to log in to a Linux ECS from an Android terminal through JuiceSSH, see [Logging In to a Linux ECS from an Android Terminal](#).

Prerequisites

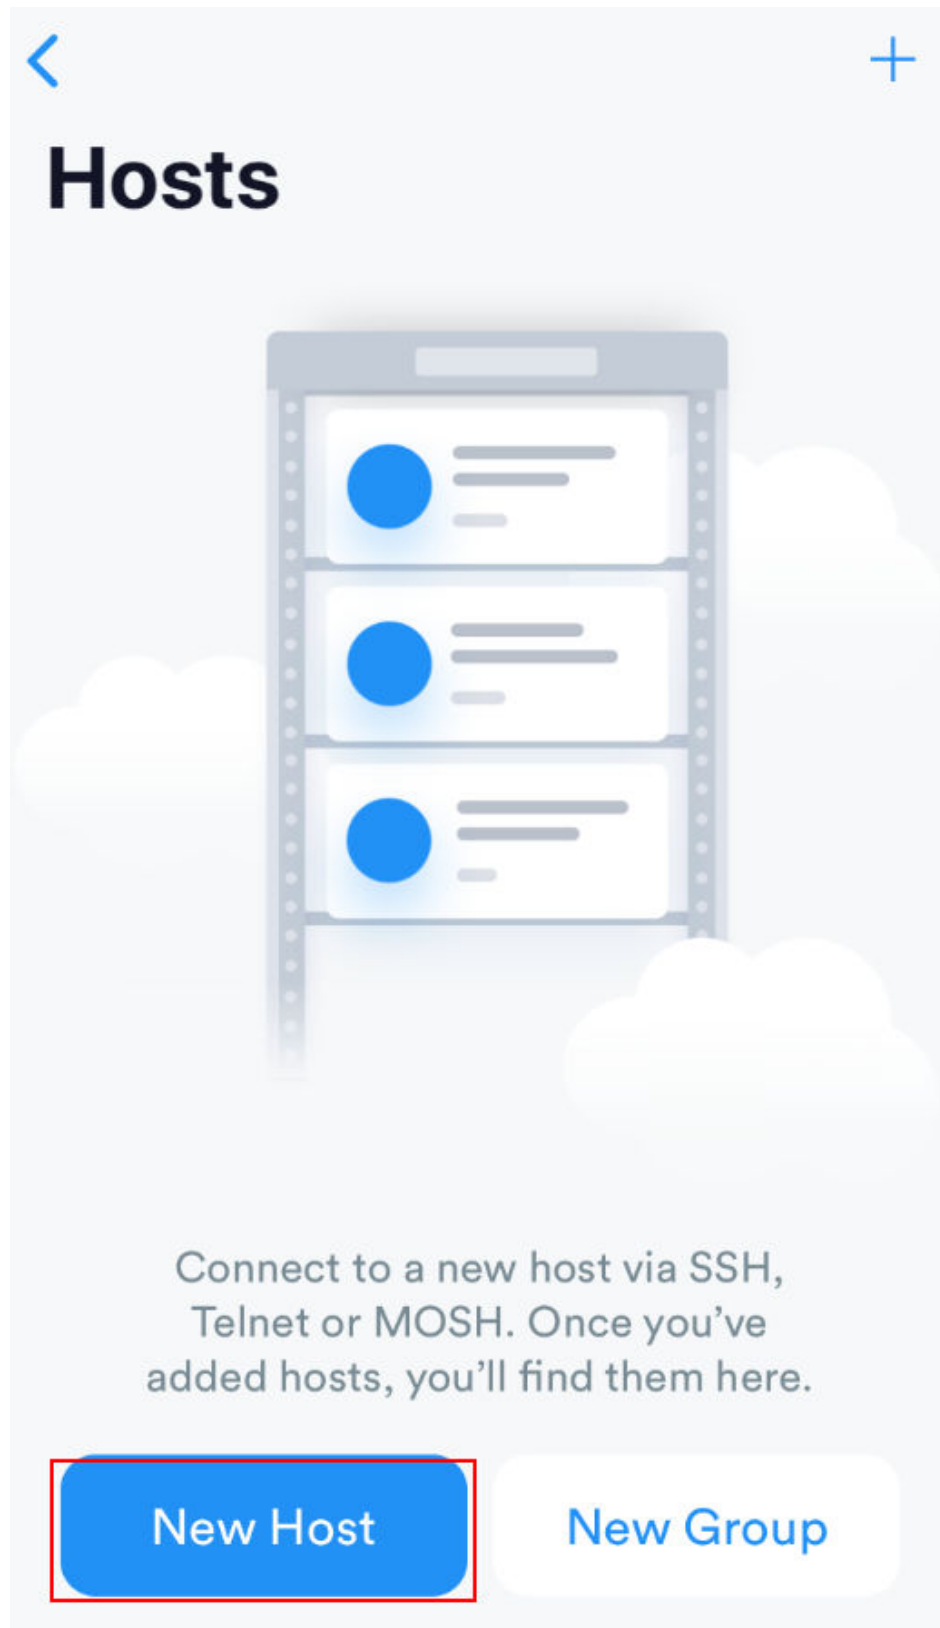
- The target ECS is running.
- You have obtained the username and password for logging in to the ECS.
- You have bound an EIP to the ECS. For details, see [Binding an EIP](#).
- Access to port 22 is allowed in the inbound direction of the security group to which the ECS belongs. For details, see [Configuring Security Group Rules](#).

Logging In to a Linux ECS from an iOS Terminal

Before performing the operation, make sure that you have installed an SSH client tool, for example, Termius, on the iOS terminal. In this example, the Linux ECS runs CentOS 7.6, and it is authenticated using a username and password.

1. Start Termius and tap **New Host**.

Figure 3-22 New Host



2. On the **New Host** page, set the following parameters:
 - **Alias**: Enter the hostname. In this example, set this parameter to **ecs01**.

- **Hostname:** Enter the EIP bound to the target ECS.
- **Use SSH:** Enable it.
- **Host:** Enter the EIP bound to the target ECS.
- **Port:** Enter port number **22**.
- **Username:** Enter **root**.
- **Password:** Enter the login password.

Figure 3-23 Setting parameters

Cancel New Host Save

1 Alias

2 Hostname

Group >

Tags >

Backspace as CTRL+H

SSH / MOSH

3 Use SSH

Use Mosh (Beta)


4 Port 22
Default

5 Username root

6 Password ●●●●●●●●


3. Tap **Save** in the upper right corner of the page to save the login settings. On the **Hosts** page, tap the name of the connection.

Figure 3-24 Login information



If the following page is displayed, you have connected to the Linux ECS.

Figure 3-25 Connected

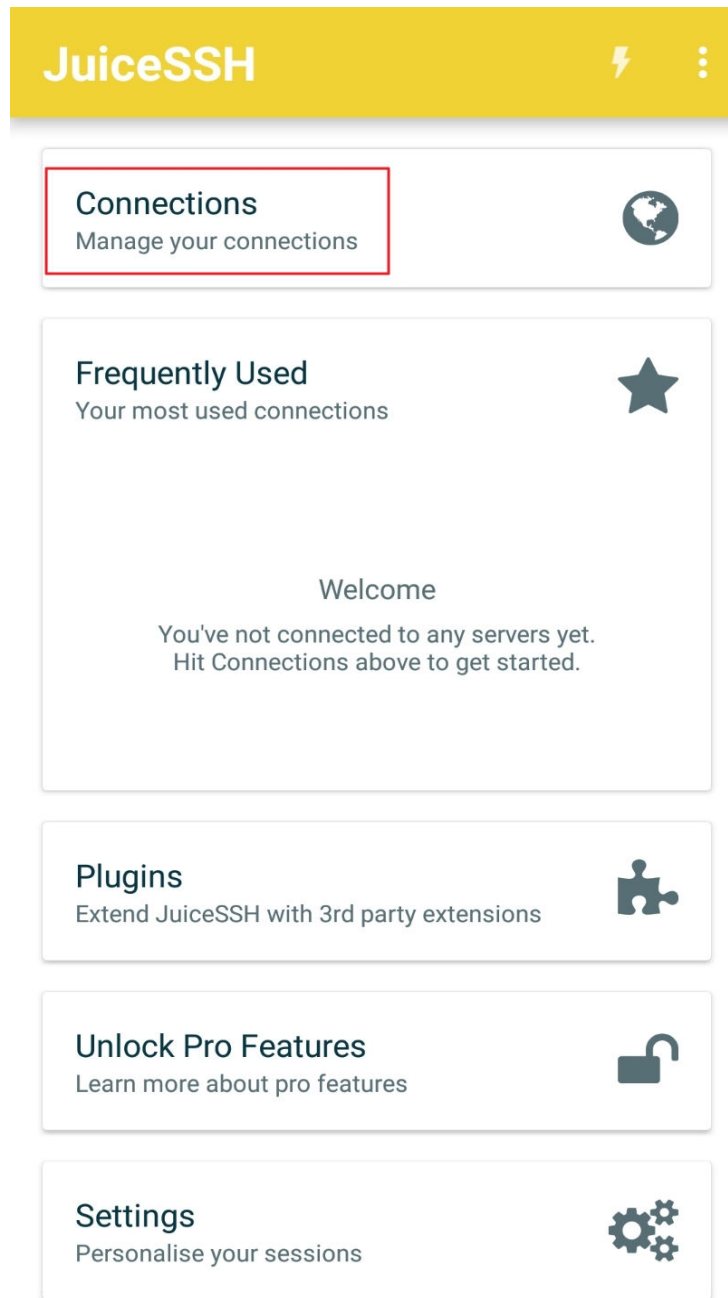


Logging In to a Linux ECS from an Android Terminal

Before performing the operation, make sure that you have installed JuiceSSH on the Android terminal. In this example, the Linux ECS runs CentOS 7.6, and it is authenticated using a username and password.


1. Start JuiceSSH and tap **Connections**.

Figure 3-26 Starting JuiceSSH




2. On the **Connections** page, tap .

Figure 3-27 Connections



No Connections

You do not currently have any connections configured. Use the button below to get started.



3. On the **New Connection** page, configure basic and advanced settings and save the settings. The parameters are as follows:
 - **Nickname:** Set the name of the login session. In this example, set this parameter to **linux_test**.
 - **Type:** Retain the default value **SSH**.
 - **Address:** Enter the EIP bound to the target Linux ECS.
 - Perform the following operations to set **Identity**:
 - i. Tap **Identity** and choose **New** from the drop-down list.


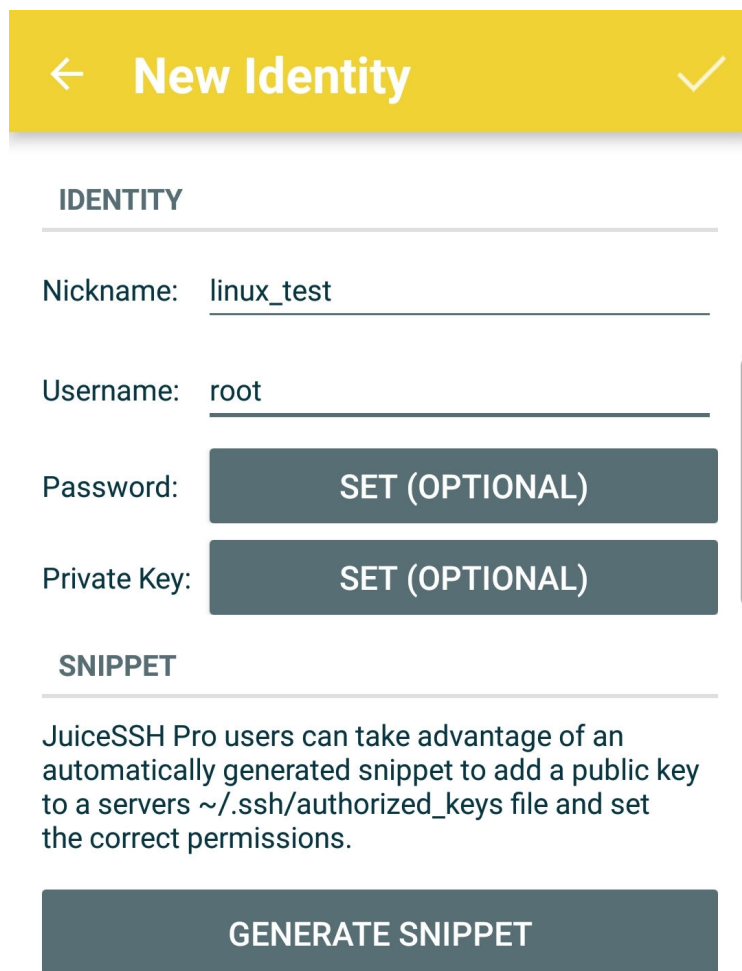
- ii. On the **New Identity** page, set the following parameters and tap .
 - **Nickname:** Set an identity name as required to facilitate subsequent management. This parameter is optional. In this example, set it to **linux_test**.
 - **Username:** Enter **root**.
 - **Password:** Tap **SET (OPTIONAL)**, enter the login password, and tap **OK**.

Figure 3-28 New Identity



← **New Identity** ✓

IDENTITY

Nickname: linux_test

Username: root

Password: **SET (OPTIONAL)**

Private Key: **SET (OPTIONAL)**

SNIPPET

JuiceSSH Pro users can take advantage of an automatically generated snippet to add a public key to a servers `~/.ssh/authorized_keys` file and set the correct permissions.

GENERATE SNIPPET

- **Port:** Enter port number **22**.

Figure 3-29 Port

← **New Connection** ✓

BASIC SETTINGS

Nickname:

Type:

Address:

Identity:

ADVANCED SETTINGS

Port:

Connect Via:

Run Snippet:


Backspace:

GROUPS

ADD TO GROUP

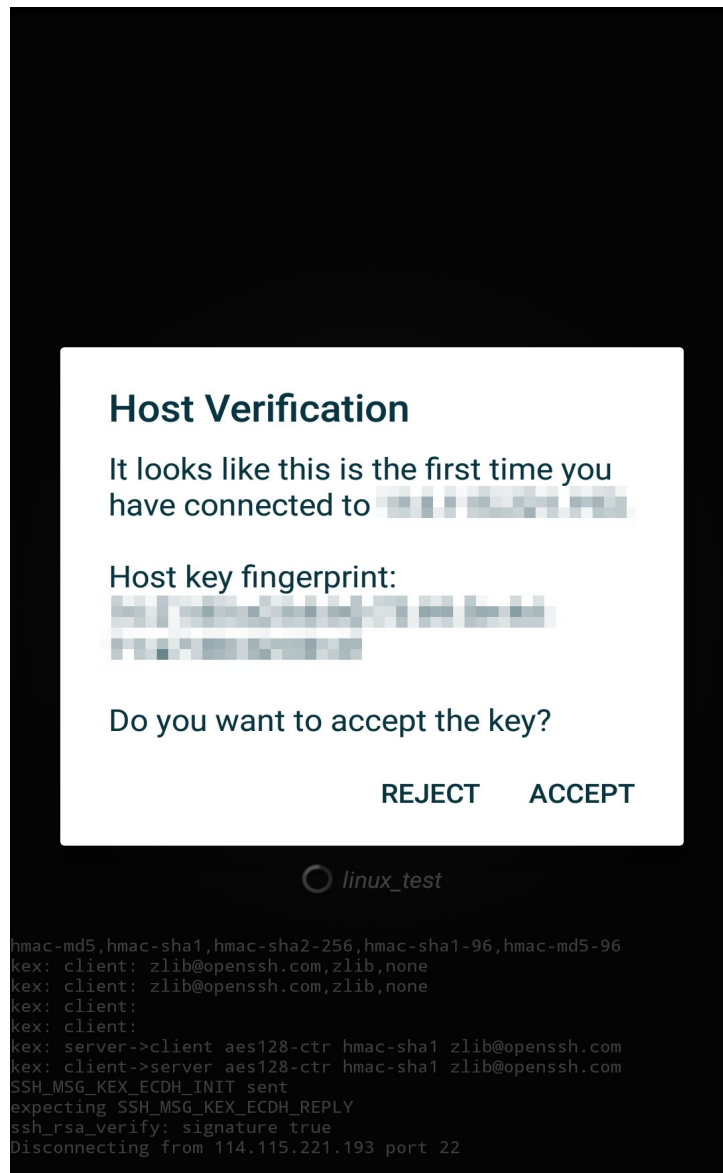
4. On the **Connections** page, tap the created connection.

Figure 3-30 Connections



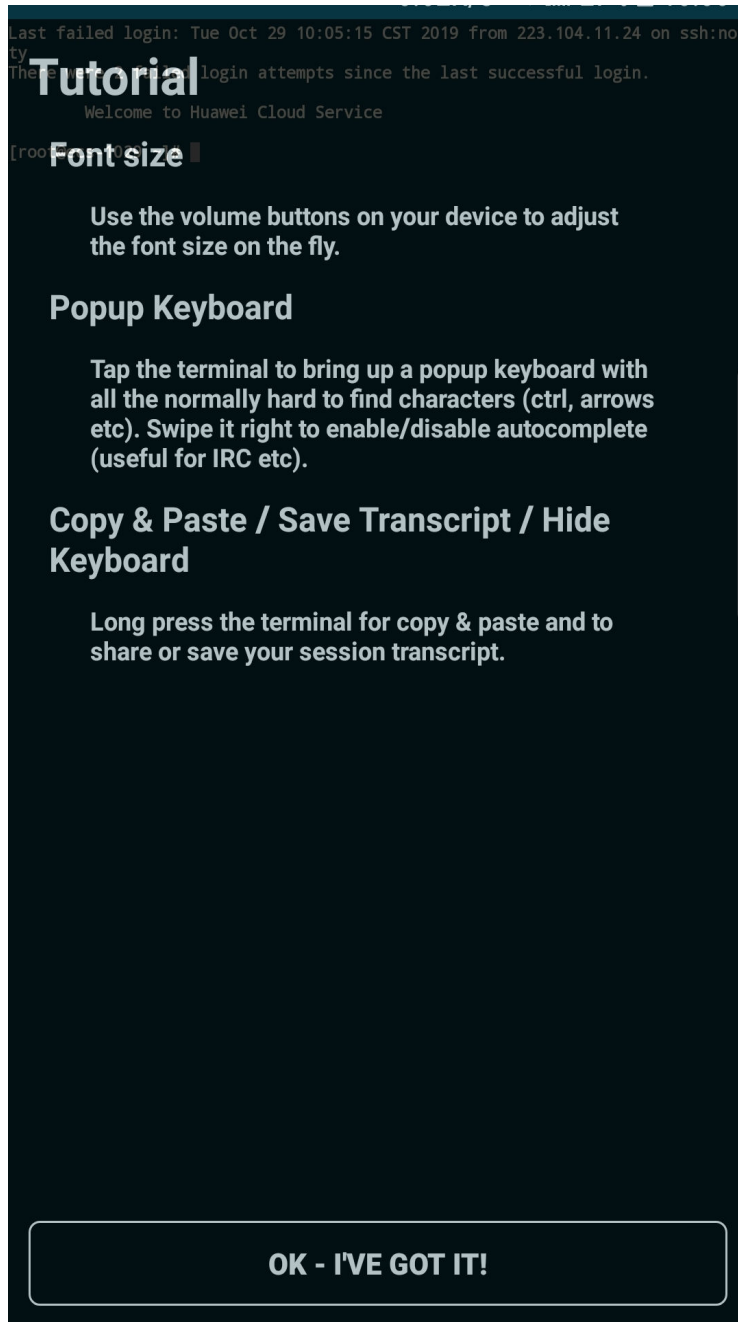
5. Confirm the information that is displayed and tap **ACCEPT**.

Figure 3-31 Confirming the information




6. (Optional) When you log in to the ECS for the first time, JuiceSSH displays a tutorial for you, including setting the font size and popping up the keyboard. Confirm the information and click **OK - I'VE GOT IT**.

Figure 3-32 Tutorial



You have logged in to the Linux ECS.

Figure 3-33 Successful login



3.3.5 Remotely Logging In to a Linux ECS (from a macOS Server)

Scenario

This section describes how to log in to a Linux ECS from a macOS server.

Prerequisites

- The target ECS is running.
- You have obtained the username and password for logging in to the ECS.
- You have bound an EIP to the ECS. For details, see [Binding an EIP](#).
- Port 22 is allowed in the inbound direction of the security group to which the ECS belongs. For details, see [Configuring Security Group Rules](#).

Procedure

You can log in to the Linux ECS through the terminal included in the macOS.

- Using an SSH password
 - a. Open the terminal of the macOS and run the following command to log in to the ECS:

```
ssh Username@EIP
```

3.4 Managing ECSs


3.4.1 Changing ECS Names

Scenarios


The name of a created ECS can be changed to meet your service requirements.

Multiple ECS names can be changed in a batch. After the change, the ECS names are the same.

Changing the Name of a Single ECS

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. Click the name of the target ECS.
4. On the page providing details about the ECS, click  next to the ECS name. Then, change the name as prompted.

Allow duplicate name: allows ECS names to be duplicate. If **Allow duplicate name** is not selected and the new name you configure is the same as an existing ECS name, the system displays a message indicating that the name has been used and you need to change it to another name.

5. Click  next to the ECS name.

6. Click **OK**.

Changing the Names of Multiple ECSs in a Batch

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Select the target ECSs.
4. Click **More** above the ECS list and select **Change ECS Name** from the drop-down list.
5. Enter a new name.
6. Click **OK**.

If you change ECS names in a batch, the new ECS names are the same, for example, all are **ecs-test**.

3.4.2 Reinstalling the OS

Scenarios

If the OS of an ECS fails to start or requires optimization, reinstall the OS.

Notes

- After the OS is reinstalled, the IP and MAC addresses of the ECS remain unchanged.
- Reinstalling the OS clears the data in all partitions of the EVS system disk, including the system partition. Therefore, back up data before reinstalling the OS.
- Reinstalling the OS does not affect data disks.
- Do not perform any operations on the ECS immediately after its OS is reinstalled. Wait for several minutes until the system successfully injects the password. Otherwise, the injection may fail, and the ECS cannot be logged in to.

Constraints

- The EVS disk quotas must be greater than 0.
- If the target ECS is created using a private image, ensure that the private image is available.
- H2 ECSs do not support OS reinstallation.

Prerequisites

- The target ECS has a system disk attached.

Procedure

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. Locate the row containing the target ECS and choose **More > Manage Image/Disk/Backup > Reinstall OS** in the **Operation** column.

Before reinstalling the OS, stop the ECS first or select Automatically stop the ECSs and reinstall their OSs.

4. Select the login mode.
5. Click **OK**.
6. On the **Reinstall OS** page, confirm the settings, and click **OK**.

After the request is submitted, the status **Reinstalling** is displayed. When this status disappears, the reinstallation is complete.

NOTE

A temporary ECS is created during the reinstallation process. After reinstallation, this ECS will be automatically deleted. Do not perform any operation on the temporary ECS during the reinstallation process.

Follow-up Procedure

If the reinstallation fails, perform steps [2](#) to [6](#) again to retry the OS installation.

If the second reinstallation attempt is unsuccessful, contact the administrator for manual recovery.

3.4.3 Changing the OS

Scenarios

Changing an ECS OS will change the system disk attached to the ECS. After the change, the system disk ID of the ECS will be changed, and the original system disk will be deleted.

If the OS running on an ECS cannot meet service requirements, change the ECS OS.

The cloud platform supports changing between image types (public images, private images, and shared images) and between OSs. You can change your OS by changing your ECS image.

Constraints

- The OS change takes about 10 to 20 minutes. During this process, the ECS status is **Changing OS**.
- Do not perform any operations on the ECS before the system injects the password. Otherwise, the login will fail.
- The target ECS must have a system disk attached.
- The EVS disk quota must be greater than 0.
- The system disk type cannot be changed.
- H2 ECSs do not support OS change.

Notes

- After the OS is changed, the original OS is not retained, and the original system disk is deleted, including the data in all partitions of the system disk.
- Back up data before changing the OS.

- Changing the OS does not affect data in data disks.
- After the OS is changed, your service running environment must be deployed in the new OS again.
- After the OS is changed, the ECS will be automatically started.
- After the OS is changed, the system disk type of the ECS cannot be changed.
- After the OS is changed, the IP and MAC addresses of the ECS remain unchanged.
- After the OS is changed, customized configurations, such as DNS and hostname of the original OS will be reset and require reconfiguration.
- It takes about 10 to 20 minutes to change the OS. During this process, the ECS is in **Changing OS** state.
- The system disk capacity of an ECS with OS changed may change because the system disk capacity specified by the image of the changed OS may be changed.

Notes on Change Between Windows and Linux

When you change the OS from Windows to Linux or from Linux to Windows, note the following:

- To change Windows to Linux, install an NTFS partition tool, such as NTFS-3G for data reads and writes on the Windows ECS.
- To change Linux to Windows, install software, such as Ext2Read or Ext2Fsd to identify ext3 or ext4.

NOTE

If there are LVM partitions on the Linux ECS, these partitions may fail after the OS is changed to Windows. Therefore, a change from Linux to Windows is not recommended.

Prerequisites

- The data is backed up.
- If you plan to use a private image to change the OS, ensure that a private image is available. For details about how to create a private image, see *Image Management Service User Guide*.
 - If the image of a specified ECS is required, make sure that a private image has been created using this ECS.
 - If a local image file is required, make sure that the image file has been imported to the cloud platform and registered as a private image.
 - If a private image from another region is required, make sure that the image has been copied.
 - If a private image from another user account is required, make sure that the image has been shared with you.

Procedure

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.

3. Locate the row containing the target ECS and choose **More > Manage Image/Disk > Change OS** in the **Operation** column.

Before changing the OS, stop the ECS first or select **Automatically stop the ECSs and change their OSs** in the displayed dialog box.

 **NOTE**

ECS OSs can be changed in a batch. To do so, perform the following operations:

1. Select the target ECSs.
 2. Click **More** above the ECS list and select **Change OS** from the drop-down list.
4. Select the target image.
For more details, see [Creating an ECS](#).
 5. Configure the login mode.
When using a private image to change the OS, you can use the private image password.
 6. Click **OK**.
 7. On the **Change OS** page, confirm the specifications, and click **Submit**.
After the application is submitted, the status **Changing OS** is displayed. When this status disappears, the OS change is complete.

 **NOTE**

A temporary ECS is created during the OS change process. After the process is complete, this ECS will be automatically deleted.

Follow-up Procedure

- If the OSs before and after the OS change are both Linux, and automatic mounting upon system startup has been enabled for data disks, the data disk partition mounting information will be lost after the OS is changed. In such a case, you need to update the **/etc/fstab** configuration.
 - a. Write the new partition information into **/etc/fstab**.
It is a good practice to back up the **/etc/fstab** file before writing data into it.
To enable automatic partition mounting upon system startup, see [Initializing a Linux Data Disk \(fdisk\)](#).
 - b. Mount the partition so that you can use the data disk.
mount *Disk partition Device name*
 - c. Check the mount result.
df -TH
- If the OS change is unsuccessful, perform steps **2** to **7** again to retry the OS change.
- If the second OS change attempt is unsuccessful, contact the administrator for manual recovery.

3.4.4 Managing ECS Groups

Scenarios

An ECS group logically groups ECSs. ECSs in an ECS group comply with the same policy associated with the ECS group.

Currently, only the anti-affinity policy is supported.

This policy enables ECSs in the same ECS group to run on different hosts for improved reliability, high availability, and disaster recovery.

You can perform the following operations on an ECS group:

- [Creating an ECS Group](#)
- [Adding an ECS to an ECS Group](#)
 - Add an ECS to an ECS group during ECS creation.
For details, see [Step 3: Configure Advanced Settings](#).
 - Add an existing ECS to an ECS group.
- [Removing an ECS from an ECS Group](#)
- [Deleting an ECS Group](#)

Creating an ECS Group

Create an ECS group and associate the same policy to all group members. ECS groups are independent from each other.

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. In the navigation pane on the left, choose **ECS Group**.
4. On the **ECS Group** page, click **Create ECS Group**.
5. Enter the name of an ECS group.
6. Select a policy for the ECS group.
7. Click **OK**.

Adding an ECS to an ECS Group

To improve service reliability, you can add ECSs to an ECS group so that these ECSs in this group can run on different hosts.

 **NOTE**

- ECSs of specific types must be stopped before being added to an ECS group. Stop these ECSs as prompted when adding them to an ECS group.
 - After an ECS is added to an ECS group, the system reallocates a host to run this ECS to ensure that ECSs in this group are running on different hosts. When the ECS is being restarted, the startup may fail due to insufficient resources. In such a case, remove the ECS from the ECS group and try to restart the ECS again.
 - ECSs that have local disks attached can be added to an ECS group only during the creation process. Once created, they can no longer be added to any ECS groups.
 - ECSs that have local disks, GPU cards, or FPGA cards attached can be added to an ECS group only during the creation process. Once created, they can no longer be added to any ECS groups.
1. Log in to the management console.
 2. Under **Computing**, choose **Elastic Cloud Server**.
 3. In the navigation pane on the left, choose **ECS Group**.
 4. Locate the row that contains the target ECS group and click **Add ECS** in the **Operation** column.
 5. On the **Add ECS** page, select an ECS to be added.
 6. Click **OK**. The ECS is added to the ECS group.

Removing an ECS from an ECS Group

After an ECS is removed from an ECS group, the ECS does not comply with the anti-affinity policy anymore.

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. In the navigation pane on the left, choose **ECS Group**.
4. Expand the ECS group information and view the ECSs in the ECS group.
5. Locate the ECS to be removed and click **Remove** in the **Operation** column.
6. In the displayed dialog box, click **Yes**.
The ECS is removed from the ECS group.

Deleting an ECS Group

After an ECS group is deleted, the policy does not apply to the ECSs in the ECS group anymore.

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. In the navigation pane on the left, choose **ECS Group**.
4. Locate the ECS group to be deleted and click **Delete** in the **Operation** column.
5. In the displayed dialog box, click **Yes**.

3.4.5 Changing the Time Zone for an ECS

Scenarios

The default time zone for an ECS is the one you selected when creating the image that was used to create the ECS. This section describes how to change the time zone for an ECS to the local one or to another time zone in your network.

After you log in to your ECS, if you find that the time on the ECS is different from the local time, you can change the time zone for the ECS so that the time on the ECS is the same as the local time.

For Linux ECSs

The process of changing the time zone for a Linux ECS depends on the OS. In this section, the CentOS 6.x 64bit OS is used to demonstrate how to change the time zone for a Linux ECS.

1. Log in to the ECS.
2. Run the following command to switch to user **root**:

```
su - root
```

3. Run the following command to obtain the time zones supported by the ECS:

```
ls /usr/share/zoneinfo/
```

In the terminal display, the **/usr/share/zoneinfo** directory contains a hierarchy of time zone data files. Use the directory structure to obtain your desired time zone file.

The directory structure shown in **/usr/share/zoneinfo** includes both time zones and directories. The directories contain time zone files for specific cities. Locate the time zone for the city in which the ECS is located.

4. Set the target time zone.
 - a. Run the following command to open the **/etc/sysconfig/clock** file:

```
vim /etc/sysconfig/clock
```
 - b. Locate the **ZONE** entry and change its value to the name of the desired time zone file.
5. Press **Esc**. Then, run the following command to save and exit the **/etc/sysconfig/clock** file:

```
:wq
```
6. Run the following command to check whether the **/etc/localtime** file is available on the ECS:

```
ls /etc/localtime
```

 - If the file is available, go to step [7](#).
 - If the file is not available, go to step [8](#).
7. Run the following command to delete the existing **/etc/localtime** file:

```
rm /etc/localtime
```
8. Run the following command to create a symbolic link between **/etc/localtime** and your time zone file so that the ECS can find this time zone file when it references the local time:

ln -sf /usr/share/zoneinfo/Asia/city1/etc/localtime

- Run the following command to restart the ECS so that all services and applications running on the ECS use the new time zone:

reboot

- Log in to the ECS again and run the following command as user **root** to check whether the time zone has been changed:

ls -lh /etc/localtime

The following information is displayed:


```
# ls -lh /etc/localtime  
lrwxrwxrwx 1 root root 33 Nov 27 11:01 /etc/localtime -> /usr/share/zoneinfo/Asia/city1
```

For Windows ECSs

- Log in to the ECS.
- Click the time display on the far right side of the task bar located at the bottom of your screen. In the dialog box that is displayed, click **Change date and time settings**.

The **Date and Time** page is displayed.

Figure 3-34 Date and Time



- Click **Change time zone**.

The **Time Zone Settings** page is displayed.

4. In the **Set the time zone** pane, choose the target time zone from the **Time zone** drop-down list.
5. Click **OK**.

3.4.6 Automatically Recovering ECSs

Scenarios

ECSs run on physical servers. Although there are multiple mechanisms to ensure system reliability, error tolerance, and high availability, server hardware might be damaged or power failure might occur. If physical servers cannot be powered on or restarted due to damage, CPU and memory data will be lost, and the ECSs cannot recover through live migration.

The cloud platform provides automatic recovery to restart ECSs through cold migration, ensuring high availability and top-performing dynamic migration capability of ECSs. You can enable automatic recovery during or after ECS creation. If a physical server accommodating ECSs breaks down, the ECSs with automatic recovery enabled will automatically be migrated to a functional server to minimize the impact on your services. During this process, the ECSs will restart.

Notes

- Automatic recovery does not ensure user data consistency.
- An ECS can be automatically recovered only if the physical server on which it is deployed becomes faulty. This function does not take effect if the fault is caused by the ECS itself.
- An ECS can be automatically recovered only after the physical server on which it is deployed is shut down. If the physical server is not shut down due to a fault, for example, a memory fault, automatic recovery fails to take effect.
- An ECS can be automatically recovered only once within 12 hours if the server on which it is deployed becomes faulty.
- ECS automatic recovery may fail in the following scenarios:
 - No physical server is available for migration due to a system fault.
 - The target physical server does not have sufficient temporary capacity.
- An ECS with any of the following resources cannot be automatically recovered:
 - Local disk
 - Passthrough FPGA card
 - Passthrough InfiniBand NIC

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Click the name of the target ECS.
The page providing details about the ECS is displayed.

4. Set **Auto Recovery** to **Enable** or **Disable**.

Automatic recovery is enabled by default.

- If a physical server accommodating ECSs breaks down, the ECSs with automatic recovery enabled will automatically be migrated to a functional server to minimize the impact on your services. During this process, the ECSs will restart.
- If **Auto Recovery** is disabled, you must wait for the system administrator to recover ECSs when hardware becomes faulty.

3.4.7 Obtaining ECS Console Logs

Scenarios

When an ECS cannot start or run properly, you can download and view ECS console logs for troubleshooting, for example, checking whether the kernel and service configuration are correct.

The ECS console logs record ECS operations, such as ECS starting, stopping, restarting, or forcibly restarting. Through the management console, you can obtain the ECS logs within one hour.

Notes

- The system does not record the logs for forcible ECS stopping.
- The system supports viewing console logs for the ECSs running the following OSs:
 - Red Hat Enterprise Linux 6.x series
 - Red Hat Enterprise Linux 7.x series
 - CentOS 6.x series
 - CentOS 7.x series
 - Ubuntu 14.x series
 - Ubuntu 16.x series
 - SUSE 11.x series
 - SUSE 12.x series
 - OpenSUSE 13.x series
 - OpenSUSE 42.x series
 - Debian 16.x series
 - Fedora series
 - FreeBSD series
 - CoreOS series
- The ECSs running Windows do not support console logs.
- The system can save up to 100 KB log files.

Procedure

Step 1 Log in to the ECS.


Step 2 Check and modify the grub file.

The configuration method varies depending on the OS.

NOTE

To prevent impact on the start of the recovery mode, you are advised to modify only the item used for the default start.

- For CentOS and Red Hat 6, perform the following steps:
 - a. Run the following command to open the configuration file:
vi /boot/grub/menu.lst
 - b. Locate the row that contains **linux**, **linux16**, or **kernel** (depending on the system), add **console=ttyS0** to its end, and delete parameter **rhgb quiet**. If **console=ttyS0** already exists, you do not need to add it. Save the change and exit.
- For CentOS 7, Red Hat 7, and Ubuntu 14, perform the following steps:
 - a. Run the following command to open the configuration file:
vi /boot/grub2/grub.cfg
 - b. Locate the row that contains **linux**, **linux16**, or **kernel** (depending on the system), add **console=ttyS0** to its end, and delete parameter **rhgb quiet**. If **console=ttyS0** already exists, you do not need to add it. Save the change and exit.
- For SUSE Linux 11, perform the following steps:
 - a. Run the following command to open the configuration file:
vi /boot/grub/menu.1st
 - b. Locate the row that contains **linux**, **linux16**, or **kernel** (depending on the system) and add **console=ttyS0** to its end. If **console=ttyS0** already exists, you do not need to add it. Save the change and exit.
- For SUSE Linux 12, openSUSE 13, and openSUSE 42, perform the following steps:
 - a. Run the following command to open the configuration file:
vi /boot/grub2/grub.cfg
 - b. Locate the row that contains **linux**, **linux16**, or **kernel** (depending on the system) and add **console=ttyS0** to its end. If **console=ttyS0** already exists, you do not need to add it. Save the change and exit.
- For Debian and Ubuntu 16, perform the following steps:
 - a. Run the following command to open the configuration file:
vi /boot/grub/grub.cfg
 - b. Locate the row that contains **linux**, **linux16**, or **kernel** (depending on the system) and add **console=ttyS0** to its end. If **console=ttyS0** already exists, you do not need to add it. Save the change and exit.
- For Fedora, perform the following steps:
 - a. Run the following command to open the configuration file:
vi /boot/grub2/grub.cfg
 - b. Locate the row that contains **linux**, **linux16**, or **kernel** (depending on the system) and add **console=ttyS0** to its end. If **console=ttyS0** already exists, you do not need to add it. Save the change and exit.

- For FreeBSD, perform the following steps:
 - a. Run the following command to open the configuration file:
vi /boot/loader.conf
 - b. Add **console="comconsole"**. If **console="comconsole"** already exists, you do not need to add it. Save the change and exit.
- For CoreOS, perform the following steps:
 - a. Run the following command to check whether **ttyS0** has been configured:
cat /proc/cmdline | grep ttyS0
 - If yes, **ttyS0** has been configured.
 - If no, **ttyS0** has not been configured. Go to [Step 2.b](#).
 - b. Run the following command to open the configuration file to be edited:
vi /usr/share/oem/grub.cfg
 **NOTE**

If the **/usr/share/oem/grub.cfg** configuration file does not exist, manually create the file.
 - c. Add **set linux_append="console=ttyS0"**. If **set linux_append="console=ttyS0"** already exists, you do not need to add it. Save the change and exit.

Step 3 On the **Elastic Cloud Server** page, click **Restart**.

Step 4 Obtain ECS console logs.

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. On the **Elastic Cloud Server** page, click the name of the target ECS.
4. On the page providing details about the ECS, click the **Console Logs** tab.
5. Choose the number of lines to be displayed for a log from the **Displayed Lines** drop-down list.
6. Click **Query**.

View details of the displayed log.

 **NOTE**

After you click **Query**, the system will not automatically update the displayed log. To view the latest log, click **Query** again.

7. (Optional) Click **Download** to download the information of the displayed log. Downloaded log files are in .txt format.

----End

3.5 Modifying ECS Specifications

3.5.1 General Operations

Scenarios

If ECS specifications do not meet service requirements, you can modify the ECS specifications, including vCPUs and memory. Certain ECSs allow you to change their types when you modify their specifications.

Notes

- Downgrading ECS specifications (vCPU or memory) will reduce performance.
- Certain ECS types do not support specifications modification currently. For details about available ECS types and functions, see [ECS Types](#). For details about restrictions on using different types of ECSs, see their notes.
- When the disk status is **Expanding**, you are not allowed to modify the specifications of the ECS where the disk is attached.
- Before modifying the specifications of a Windows ECS, modify the SAN policy by following the instructions provided in [Why Does a Disk Attached to a Windows ECS Go Offline?](#) to prevent disks from going offline after the specifications are modified.

Step 1: Modify Specifications

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Click **More** in the **Operation** column and select **Modify Specifications**.
The **Modify ECS Specifications** page is displayed.
4. Select the new ECS type, vCPUs, and memory as prompted.
Before modifying the specifications, stop the ECS or select **Automatically stop the ECSs and then modify specifications**.
5. Click **Next**.
6. Confirm the settings, read and select the disclaimer, and then click **Submit Application**.

Step 2: Check Disk Attachment

After specifications are modified, disk attachment may fail. Therefore, check disk attachment after specifications modification. If disks are properly attached, the specifications modification is successful.

- Windows ECS
For details, see [Why Do the Disks of a Windows ECS Go Offline After I Modify the ECS Specifications?](#)
- Linux ECS
For details, see [Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?](#)

3.6 Obtaining Metadata and Passing User Data

3.6.1 Obtaining Metadata

Scenarios

ECS metadata includes basic information of an ECS on the cloud platform, such as the ECS ID, hostname, and network information. ECS metadata can be obtained using either OpenStack or EC2 compatible APIs, as shown in [Table 3-5](#). The following describes the URI and methods of using the supported ECS metadata.

Notes

If the metadata contains sensitive data, take appropriate measures to protect the sensitive data, for example, controlling access permissions and encrypting the data.

Perform the following configuration on the firewall:

- Windows

If you need to assign permissions only to the administrator to access custom data, enable the firewall as an administrator and run the following commands in PowerShell:

```
PS C:\>$RejectPrincipal = New-Object -TypeName  
System.Security.Principal.NTAccount ("Everyone")
```

```
PS C:\>$RejectPrincipalSID =  
$RejectPrincipal.Translate([System.Security.Principal.SecurityIdentifier]).V  
alue
```

```
PS C:\>$ExceptPrincipal = New-Object -TypeName  
System.Security.Principal.NTAccount ("Administrator")
```

```
PS C:\>$ExceptPrincipalSID =  
$ExceptPrincipal.Translate([System.Security.Principal.SecurityIdentifier]).  
Value
```

```
PS C:\>$PrincipalSDDL = "O:LSD:(D;;CC;;;$ExceptPrincipalSID)  
(A;;CC;;;$RejectPrincipalSID)"
```

```
PS C:\>New-NetFirewallRule -DisplayName "Reject metadata service for $  
($RejectPrincipal.Value), exception: $($ExceptPrincipal.Value)" -Action  
block -Direction out -Protocol TCP -RemoteAddress 169.254.169.254 -  
LocalUser $PrincipalSDDL
```

- Linux

If you need to assign permissions only to user **root** to access custom data, run the following command as user **root**:

```
iptables --append OUTPUT --proto tcp --destination 169.254.169.254 --  
match owner ! --uid-owner root --jump REJECT
```

ECS Metadata Types

[Table 3-5](#) does not contain the following metadata items: ami-id, ami-launch-index, ami-manifest-path, block-device-mapping/, instance-action, instance-id,

reservation-id, ramdisk-id, and kernel-id. These metadata items are meaningless and are not recommended.

Table 3-5 ECS metadata types

Metadata Type	Metadata Item	Description
OpenStack	/meta_data.json	Displays ECS metadata. For the key fields in the ECS metadata, see Table 3-6 .
OpenStack	/password	Displays the password for logging in to an ECS. This metadata is used by Cloudbase-Init to store ciphertext passwords during initialization of key-pair-authenticated Windows ECSs.
OpenStack	/user_data	Displays ECS user data. This metadata allows you to specify scripts and configuration files for initializing ECSs. For details, see Passing User Data to ECSs . For password-authenticated Linux ECSs, this metadata is used to save password injection scripts.
OpenStack	/network_data.json	Displays ECS network information.
OpenStack	/securitykey	Obtains temporary AKs and SKs. Before enabling an ECS to obtain a temporary AK and SK, make sure that the op_svc_ecs account has been authorized on IAM and that the desired ECS resources have been authorized for management.
EC2-compatible	/meta-data/hostname	Displays the name of the host accommodating an ECS. To remove the suffix .novalocal from an ECS, see: Is an ECS Hostname with Suffix .novalocal Normal?
EC2-compatible	/meta-data/local-hostname	The meaning of this field is the same as that of hostname.
EC2-compatible	/meta-data/public-hostname	The meaning of this field is the same as that of hostname.

Metadata Type	Metadata Item	Description
EC2-compatible	/meta-data/ instance-type	Displays an ECS flavor.
EC2-compatible	/meta-data/ local-ipv4	Displays the fixed IP address of an ECS. If there are multiple NICs, only the IP address of the primary NIC is displayed.
EC2-compatible	/meta-data/ placement/ availability-zone	Displays the AZ accommodating an ECS.
EC2-compatible	/meta-data/ public-ipv4	Displays the EIP bound to the ECS. If there are multiple NICs, only the EIP of the primary NIC is displayed.
EC2-compatible	/meta-data/ public-keys/0/ openssh-key	Displays the public key of an ECS.
EC2-compatible	/user-data	Displays ECS user data.
EC2-compatible	/meta-data/ security-groups	Displays the security group of an ECS.

Table 3-6 Metadata key fields

Parameter	Type	Description
uuid	String	Specifies an ECS ID.
availability_zone	String	Specifies the AZ where an ECS locates.
meta	Dict	Specifies the metadata information, including the image name, image ID, and VPC ID.
hostname	String	Specifies the name of the host accommodating an ECS. To remove the suffix .novalocal from an ECS, see: Is an ECS Hostname with Suffix .novalocal Normal?

Prerequisites

- The target ECS has been logged in.
- Security group rules in the outbound direction meet the following requirements:

- **Protocol:** TCP
- **Port:** 80
- **Destination:** 169.254.0.0/16

NOTE

If you use the default security group rules for the outbound direction, the metadata can be accessed because the default rules meet the preceding requirements. For details about the default security group rules for the outbound direction, see [Default Security Group and Rules](#).

Metadata (OpenStack Metadata API)

This API is used to query ECS metadata.

- URI
`/169.254.169.254/openstack/latest/meta_data.json`
- Usage method
Supports GET requests.
- Example

To use cURL to view Linux ECS metadata, run the following command:

```
curl http://169.254.169.254/openstack/latest/meta_data.json
```

To use Invoke-RestMethod to view Windows ECS metadata, run the following command:

```
Invoke-RestMethod http://169.254.169.254/openstack/latest/meta_data.json | ConvertTo-Json
```

```
{
  "random_seed": "rEocCViRS+dNwlYdGlxJHU+p+00poeUsAdBFkbPbYQTmpNwpoEb43k9z+96TyrekNKS
+ilYDdRny4kKGoNPEVBCC05Hg1TcDbIAPfJwgJS1okqEtlcofUhKmlL3K0fto
+5KXEDU3GNuGwyZXjdVb9HQWU+E1jztAJjqsahnU+g/tawABTVySLBKlAT8fMGax1mTgGArucn/
WzDcy19DGioKPE7F8lLtSQ4Ww3VCLK5VYB/h0x+4r7IVHrPmYX/
bi1Yhm3Dc4rRYNaTjdOV5gUOsbo3oAeQkmKwQ/
NO0N8qw5Ya4l8ZUW4tMav4mOsRySOOB35v0bvaJc6p
+50DTbWNeX5A2MLiEhTP3vsPrmvk4LRF7CLz2J2TGIM14OoVBw7LARwmv9cz532zHki/c8tLhRzLmOTXh/
wL36zFW10DeuReUGmxth7IGNmRMQKV6+mil78jm/KMPpgAdK3vwYF/
GcelOFJD2HghMUUCeMbwYnvijLTejuBpwhJMNiHA/NvlEsxJDxqBCoss/Jfe+yCmUFyxovJ
+L8oNkTzkmCNzw3Ra0hiKchGhqK3BleToV/kVx5DdF081xrEA
+qyoM6CVyfJtEoz1zRRyoo9bJ65Eg6JJd8dj1UCVsDqRY1pljgzE/
Mzsw6AaaCVhaMJL7u7YMVdyKzA6z65Xtvujz0Vo=",
  "uuid": "ca9e8b7c-f2be-4b6d-a639-f10b4d994d04",
  "availability_zone": "lt-test-1c",
  "hostname": "ecs-ddd4.novalocal",
  "launch_index": 0,
  "meta": {
    "metering.image_id": "3a64bd37-955e-40cd-ab9e-129db56bc05d",
    "metering.imagetype": "gold",
    "metering.resourcespeccode": "s3.medium.2.linux",
    "image_name": "CentOS 7.6 64bit",
    "metering.resourcetype": "1",
    "vpc_id": "3b6c201f-aeb3-4bce-b841-64756e66cb49",
    "os_bit": "64",
    "cascaded.instance_extrainfo": "pcibridge:1",
    "os_type": "Linux",
    "charging_mode": "0"
  },
  "project_id": "6e8b0c94265645f39c5abbe63c4113c6",
  "name": "ecs-ddd4"
}
```


Invoke-RestMethod http://169.254.169.254/openstack/latest/network_data.json | ConvertTo-Json

```
{
  "services": [{
    "type": "dns",
    "address": "xxx.xx.x.x"
  },
  {
    "type": "dns",
    "address": "100.125.21.250"
  }
],
  "qos": {
    "instance_min_bandwidth": 100,
    "instance_max_bandwidth": 500
  },
  "networks": [{
    "network_id": "67dc10ce-441f-4592-9a80-cc709f6436e7",
    "type": "ipv4_dhcp",
    "link": "tap68a9272d-71",
    "id": "network0"
  }
],
  "links": [{
    "vif_id": "68a9272d-7152-4ae7-a138-3ef53af669e7",
    "ethernet_mac_address": "fa:16:3e:f7:c1:47",
    "mtu": null,
    "type": "cascading",
    "id": "tap68a9272d-71"
  }
]
}
```

Security Key (OpenStack Metadata API)

This API is used to obtain temporary AKs and SKs.

NOTE

- If an ECS needs to obtain a temporary AK and SK, go to the ECS details page, and configure **Agency** for the ECS in the **Management Information** area so that the ECS is authorized on IAM.
- The validity period of a temporary AK and SK is one hour. The temporary AK and SK are updated 10 minutes ahead of the expiration time. During the 10 minutes, both the new and old temporary AKs and SKs can be used.
- When using temporary AKs and SKs, add '**X-Security-Token**':{**securitytoken**} in the message header. **securitytoken** is the value returned when a call is made to the API.
- URI
/openstack/latest/securitykey
- Usage method
Supports GET requests.
- Examples
Linux:
curl http://169.254.169.254/openstack/latest/securitykey
Windows:
Invoke-RestMethod http://169.254.169.254/openstack/latest/securitykey

User Data (EC2 Compatible API)

This API is used to query ECS user data. The value is configured only when you create an ECS. It cannot be changed after the configuration.

- URI
`/169.254.169.254/latest/user-data`
- Usage method
Supports GET requests.
- Example
Linux:
curl http://169.254.169.254/latest/user-data
Windows:
Invoke-RestMethod http://169.254.169.254/latest/user-data
ICAgICAgDQoiQSBjbG91ZCBkb2VzIG5vdCBrbm93IHdoeSBpdCBtb3ZlcyBpbiBqdXN0IHN1Y2ggYSBkaXJlY3Rpb24gYW5kIGF0IHN1Y2ggYSBzZGVlZC4uLkl0IGZlZWxzIGFuIGltcHVsc2lvbi4uLnRoXMGaXMgdGh1IHBSYWNlIHhrvIGdvlG5vdy4gQnV0IHRobzSBza3kga25vd3MgdGh1IHJlYXNvbnMgYW5kIHRobzSBwYXR0ZXJucyBiZWphbmQgYWxslGNsb3Vkc3VwYW5kIHlvdSB3aWxslGtub3csIHRobybygd2hlbiB5b3UgbGlmdCB5b3Vyc2VsZiBoaWdoIGVub3VnaCB0byBzZWUgYmV5b25kIGhvcml6b25zLlINCg0KLVJpY2hcmQgQmFjaA=

Hostname (EC2 Compatible API)

This API is used to query the name of the host accommodating an ECS. The **.novalocal** suffix will be added later.

- URI
`/169.254.169.254/latest/meta-data/hostname`
- Usage method
Supports GET requests.
- Example
Linux:
curl http://169.254.169.254/latest/meta-data/hostname
Windows:
Invoke-RestMethod http://169.254.169.254/latest/meta-data/hostname
vm-test.novalocal

Instance Type (EC2 Compatible API)

This API is used to query an ECS flavor.

- URI
`/169.254.169.254/latest/meta-data/instance-type`
- Usage method
Supports GET requests.
- Example
Linux:
curl http://169.254.169.254/latest/meta-data/instance-type
Windows:
Invoke-RestMethod http://169.254.169.254/latest/meta-data/instance-type
s3.medium.2

Local IPv4 (EC2 Compatible API)

This API is used to query the fixed IP address of an ECS. If there are multiple NICs, only the IP address of the primary NIC is displayed.

- URI
/169.254.169.254/latest/meta-data/local-ipv4
- Usage method
Supports GET requests.
- Example
Linux:
curl http://169.254.169.254/latest/meta-data/local-ipv4
Windows:
Invoke-RestMethod http://169.254.169.254/latest/meta-data/local-ipv4
192.1.1.2

Availability Zone (EC2 Compatible API)

This API is used to query the AZ accommodating an ECS.

- URI
/169.254.169.254/latest/meta-data/placement/availability-zone
- Usage method
Supports GET requests.
- Example
Linux:
curl http://169.254.169.254/latest/meta-data/placement/availability-zone
Windows:
Invoke-RestMethod http://169.254.169.254/latest/meta-data/placement/availability-zone
az1.dc1

Public IPv4 (EC2 Compatible API)

This API is used to query the EIP bound to an ECS. If there are multiple NICs, only the EIP of the primary NIC is displayed.

- URI
/169.254.169.254/latest/meta-data/public-ipv4
- Usage method
Supports GET requests.
- Example
Linux:
curl http://169.254.169.254/latest/meta-data/public-ipv4
Windows:
Invoke-RestMethod http://169.254.169.254/latest/meta-data/public-ipv4
46.1.1.2

Public Keys (EC2 Compatible API)

This API is used to query the public key of an ECS.

- URI
/169.254.169.254/latest/meta-data/public-keys/0/openssh-key
- Usage method
Supports GET requests.

- Example

Linux:

```
curl http://169.254.169.254/latest/meta-data/public-keys/0/openssh-key
```

Windows:

```
Invoke-RestMethod http://169.254.169.254/latest/meta-data/public-keys/0/openssh-key
```

```
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQDI5Fw5k8Fgzajn1zJwLoV3+wMP+6CyvsSilc/hioggSnYu/AD0Yqm8vVO0kWLun1rFbdO+QUZKyVr/OPUjQSw4SRh4qsTKf/+eFoWTjplFvd1WCBZzS/WRenxlwR00KkcZHSJro763+wYcwKieb4eKRxaQoQvoFgVjLBULXAJH4eKoKTVNtMXAvPP9aMy2SLgsJNtMb9ArfziAiblQynq7UIflnN3VclzPeiWrqtzjyOp6CPUXnL0lVPTvbLe8sUteBsJZwLL6K4i+Y0lf3ryqnmQgC21yW4Dzu+kwk8FVT2MgWkCwiZd8gQ/+ujzrJFyMfUOBikIOBfuUENIUhABGenerated-by-Nova
```

Helpful Links

[Why Can't My Linux ECS Obtain Metadata?](#)

3.6.2 Passing User Data to ECSs

Scenarios

Specify **User Data** to pass user data to ECSs to:

- Simplify ECS configuration.
- Initialize the ECS OS configuration.
- Upload your scripts to ECSs during ECS creation.
- Perform other tasks using scripts.

Use Restrictions

- Linux
 - The image that is used to create ECSs must have Cloud-Init installed.
 - The user data to be specified must be less than or equal to 32 KB.
 - If user data is uploaded as text, the data can contain only ASCII characters. If user data is uploaded using a file, the file can contain any characters and the file size cannot exceed 32 KB.
 - The image that is used to create ECSs must be a public image, a private image created based on a public image, or a private image with Cloud-Init installed.
 - The format of the customized scripts must be supported by Linux ECSs.
 - DHCP must be enabled on the VPC network, and port 80 must be enabled for the security group in the outbound direction.

- When the password login mode is selected, user data cannot be passed.
- Windows
 - The image that is used to create ECSs must have Cloudbase-Init installed.
 - The user data to be specified must be less than or equal to 32 KB.
 - If user data is uploaded as text, the data can contain only ASCII characters. If user data is uploaded using a file, the file can contain any characters and the file size cannot exceed 32 KB.
 - The image that is used to create ECSs must be a public image, a private image created based on a public image, or a private image with Cloudbase-Init installed.
 - DHCP must be enabled on the VPC network, and port 80 must be enabled for the security group in the outbound direction.

Passing User Data

1. Create a user data script, the format of which complies with user data script specifications. For details, see [Helpful Links](#).
2. When creating an ECS, set **Advanced Options** to **Configure now**, and paste the content of the user data script to the **User Data** text box or upload the user data file.

NOTE

You can pass user data to an ECS as text or as a file.

Text: Copy the content of the user data script to the text box.

File: Save the user data script to a text file and then upload the file.

3. The created ECS automatically runs Cloud-Init/Cloudbase-Init and reads the user data script upon startup.

User Data Scripts of Linux ECSs

Customized user data scripts of Linux ECSs are based on the open-source Cloud-Init architecture. This architecture uses ECS metadata as the data source for automatically configuring the ECSs. The customized script types are compatible with open-source Cloud-Init. For details about Cloud-Init, see <http://cloudinit.readthedocs.io/en/latest/topics/format.html>.

- Script execution time: A customized user data script is executed after the status of the target ECS changes to **Running** and before `/etc/init` is executed.

NOTE

By default, the scripts are executed as user **root**.

- Script type: Both user-data scripts and Cloud-Config data scripts are supported.

Table 3-7 Linux ECS script types

-	User-Data Script	Cloud-Config Data Script
Description	Scripts, such as Shell and Python scripts, are used for custom configurations.	Methods pre-defined in Cloud-Init, such as the yum repository, are used for configuring certain ECS applications.
Format	The first line must start with #! (for example, #!/bin/bash or #!/usr/bin/env python) and no spaces are allowed at the beginning. When a script is started for the first time, it will be executed at the rc.local-like level, indicating a low priority in the boot sequence.	The first line must be #cloud-config , and no space is allowed in front of it.
Constraint	Before Base64 encoding, the size of the script, including the first line, cannot exceed 32 KB.	Before Base64 encoding, the size of the script, including the first line, cannot exceed 32 KB.
Frequency	The script is executed only once when the ECS is started for the first time.	The execution frequency varies according to the applications configured on the ECS.

- How can I view the customized user data passed to a Linux ECS?
 - a. Log in to the ECS.
 - b. Run the following command to view the customized user data as user **root**:
curl http://169.254.169.254/openstack/latest/user_data
- Script usage examples
This section describes how to inject scripts in different formats into Linux ECSs and view script execution results.

Example 1: Inject a user-data script.

When creating an ECS, set **User Data** to **As text** and enter the customized user data script.

```
#!/bin/bash
echo "Hello, the time is now $(date -R)" | tee /root/output.txt
```

After the ECS is created, start it and run the **cat [file]** command to check the script execution result.

```
[root@XXXXXXXX ~]# cat /root/output.txt
Hello, the time is now Mon, 16 Jul 2016 16:03:18+0800
```

Example 2: Inject a Cloud-Config data script.

When creating an ECS, set **User Data** to **As text** and enter the customized user data script.

```
#cloud-config
bootcmd:
- echo 192.168.1.130 us.archive.ubuntu.com >> /etc/hosts
```

After the ECS is created, start it and run the **cat /etc/hosts** command to check the script execution result.

Figure 3-36 Viewing operating results

```
localhost localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
192.168.1.130 us.archive.ubuntu.com
```

User Data Scripts of Windows ECSs

Customized user data scripts of Windows ECSs are based on the open-source Cloudbase-Init architecture. This architecture uses ECS metadata as the data source for initializing and automatically configuring the ECSs. The customized script types are compatible with open-source Cloudbase-Init. For details about Cloudbase-Init, see <https://cloudbase-init.readthedocs.io/en/latest/userdata.html>.

- Script type: Both batch-processing program scripts and PowerShell scripts are supported.

Table 3-8 Windows ECS script types

-	Batch-Processing Program Script	PowerShell Script
Format	The script must be started with rem cmd , which is the first line of the script. No space is allowed at the beginning of the first line.	The script must be started with #ps1 , which is the first line of the script. No space is allowed at the beginning of the first line.
Constraint	Before Base64 encoding, the size of the script, including the first line, cannot exceed 32 KB.	Before Base64 encoding, the size of the script, including the first line, cannot exceed 32 KB.

- How can I view the customized user data passed into a Windows ECS?
 - Log in to the ECS.
 - Access the following URL in the address box of the browser and view the user data:

http://169.254.169.254/openstack/latest/user_data

- Script usage examples

This section describes how to inject scripts in different formats into Windows ECSs and view script execution results.

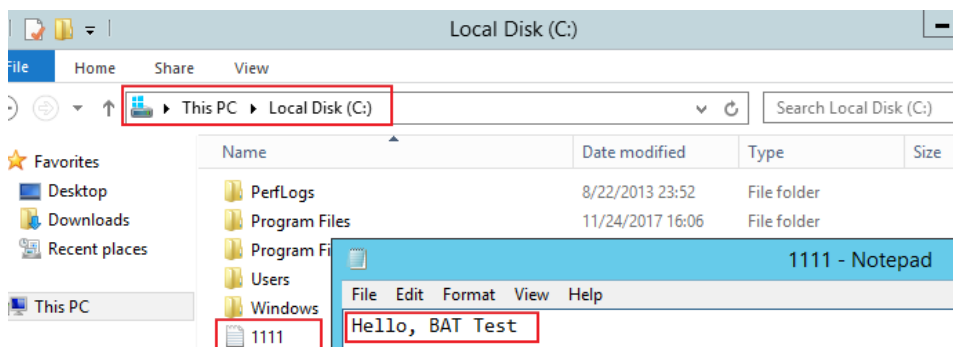
Example 1: Inject a batch-processing program script.

When creating an ECS, set **User Data** to **As text** and enter the customized user data script.

```
rem cmd
echo "Hello, BAT Test" > C:\1111.txt
```


After the ECS is created, start it and check the script execution result. In this example, a text file named **1111** is added to disk C:\.

Figure 3-37 Creating text file (Batch)



To view the user data passed to the Windows ECS, log in at http://169.254.169.254/openstack/latest/user_data.

Figure 3-38 Viewing user data (Batch)




Example 2: Inject a PowerShell script.

When creating an ECS, set **User Data** to **As text** and enter the customized user data script.

```
#ps1
echo "Hello, Powershell Test" > C:\\aaaa.txt
```

After the ECS is created, start it and check the script execution result. In this example, a text file named **aaaa** is added to disk C:\.

Figure 3-39 Creating text file (PowerShell)



To view the user data passed to the Windows ECS, log in at http://169.254.169.254/openstack/latest/user_data.

Figure 3-40 Viewing user data (PowerShell)

Case 1

This case illustrates how to pass user data to simplify Linux ECS configuration.

In this example, vim is configured to enable syntax highlighting, display line numbers, and set the tab stop to **4**. The `.vimrc` configuration file is created and injected into the `/root/.vimrc` directory during ECS creation. After the ECS is created, vim is automatically configured based on your requirements. This improves ECS configuration efficiency, especially in batch ECS creation scenarios.

User data example:

```
#cloud-config
write_files:
  - path: /root/.vimrc
    content: |
      syntax on
      set tabstop=4
      set number
```

Case 2

This case illustrates how to use the user data passing function to set the password for logging in to a Linux ECS.

NOTE

The new password must meet the password complexity requirements listed in [Table 3-9](#).

Table 3-9 Password complexity requirements

Parameter	Requirement
Password	<ul style="list-style-type: none">• Consists of 8 to 26 characters. For specific requirements on the password length, see the information displayed on the console.• Contains at least three of the following character types:<ul style="list-style-type: none">– Uppercase letters– Lowercase letters– Digits– Special characters for Windows: \$!@%-_=[:./,?– Special characters for Linux: !@%-_=[:./^,}{?• Cannot contain the username or the username spelled backwards.• Cannot contain more than two consecutive characters in the same sequence as they appear in the username. (This requirement applies only to Windows ECSs.)• Cannot start with a slash (/) for Windows ECSs.

User data example:

Using a ciphertext password (recommended)

```
#!/bin/bash  
echo 'root:$6$V6azyeLwcD3CHlpY$BN3VVq18fmCkj66B4zdHLWevqcxlig' | chpasswd -e;
```

In the preceding command output, **\$6\$V6azyeLwcD3CHlpY\$BN3VVq18fmCkj66B4zdHLWevqcxlig** is the ciphertext password, which can be generated as follows:

1. Run the following command to generate an encrypted ciphertext value:

```
python -c "import crypt, getpass, pwd;print crypt.mksalt()"
```

The following information is displayed:

```
$6$V6azyeLwcD3CHlpY
```

2. Run the following command to generate a ciphertext password based on the salt value:

```
python -c "import crypt, getpass, pwd;print crypt.crypt('Cloud.1234!','$6$V6azyeLwcD3CHlpY')"
```

The following information is displayed:

```
$6$V6azyeLwcD3CHlpY$BN3VVq18fmCkj66B4zdHLWevqcxlig
```

After the ECS is created, you can use the password to log in to it.

Case 3

This case illustrates how to use the user data passing function to reset the password for logging in to a Linux ECS.

In this example, the password of user **root** is reset to *********.

 NOTE

The new password must meet the password complexity requirements listed in [Table 3-10](#).

Table 3-10 Password complexity requirements

Parameter	Requirement
Password	<ul style="list-style-type: none">• Consists of 8 to 26 characters. For specific requirements on the password length, see the information displayed on the console.• Contains at least three of the following character types:<ul style="list-style-type: none">– Uppercase letters– Lowercase letters– Digits– Special characters for Windows: \$!@%-_=[:./,;?– Special characters for Linux: !@%-_=[:./^,{}?• Cannot contain the username or the username spelled backwards.• Cannot contain more than two consecutive characters in the same sequence as they appear in the username. (This requirement applies only to Windows ECSs.)• Cannot start with a slash (/) for Windows ECSs.

User data example (Retain the indentation in the following script):

```
#cloud-config
chpasswd:
  list: |
    root:*****
  expire: False
```

After the ECS is created, you can use the reset password to log in to it. To ensure system security, change the password of user **root** after logging in to the ECS for the first time.

Case 4

This case illustrates how to use the user data passing function to create a user on a Windows ECS and configure the password for the user.

In this example, the user's username is **abc**, its password is *********, and the user is added to the **administrators** user group.

 NOTE

The new password must meet the password complexity requirements listed in [Table 3-10](#).

User data example:

```
rem cmd
net user abc ***** /add
net localgroup administrators abc /add
```

After the ECS is created, you can use the created username and password to log in to it.

Case 5

This case illustrates how to use the user data passing function to update system software packages for a Linux ECS and enable the HTTPd service. After the user data is passed to an ECS, you can use the HTTPd service.

User data example:

```
#!/bin/bash
yum update -y
service httpd start
chkconfig httpd on
```

Helpful Links

For more information about user data passing cases, visit the official Cloud-init/Cloudbase-init website:

- <https://cloudinit.readthedocs.io/en/latest/>
- <https://cloudbase-init.readthedocs.io/en/latest/>

3.7 (Optional) Configuring Mapping Between Hostnames and IP Addresses

ECSs in the same VPC can communicate with each other using hostnames. In such a case, you are required to configure the mapping between hostnames and IP addresses. The communication using hostnames is more convenient than that using IP addresses.

Constraints

This method applies only to Linux ECSs.

Procedure

For example, there are two ECSs in a VPC, ecs-01 and ecs-02. Perform the following operations to enable communication using hostnames between ecs-01 and ecs-02:

Step 1 Log in to ecs-01 and ecs-02 and obtain their private IP addresses.

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. On the **Elastic Cloud Server** page, obtain the private IP address in the **IP Address** column.

For example, the obtained private IP addresses are as follows:

ecs-01: 192.168.0.1

ecs-02: 192.168.0.2

Step 2 Obtain the hostnames for the two ECSs.

1. Log in to an ECS.
2. Run the following command to view the ECS hostname:

```
sudo hostname
```

For example, the obtained hostnames are as follows:

```
ecs-01: hostname01
```

```
ecs-02: hostname02
```

Step 3 Create a mapping between the hostnames and IP addresses and add information about other ECSs in the same VPC.

1. Log in to ecs-01.
2. Run the following command to switch to user **root**:

```
sudo su -
```

3. Run the following command to edit the hosts configuration file:

```
vi /etc/hosts
```

4. Press **i** to enter editing mode.
5. Add the statement in the following format to set up the mapping:

```
Private IP address hostname
```

For example, add the following statement:

```
192.168.0.1 hostname01
```

```
192.168.0.2 hostname02
```

6. Press **Esc** to exit editing mode.
7. Run the following command to save the configuration and exit:

```
:wq
```

8. Log in to ecs-02.
9. Repeat [Step 3.2](#) to [Step 3.7](#).

Step 4 Check whether the ECSs can communicate with each other using hostnames.

Log in to an ECS in the same VPC, run the following command to ping the added host, and check whether the operation is successful:

```
ping Hostname
```

```
----End
```

3.8 (Optional) Installing a Driver and Toolkit

3.8.1 GPU Driver

Overview

Before using a GPU-accelerated ECS, make sure that a GPU driver has been installed on the ECS for GPU acceleration.

GPU-accelerated ECSs support GRID and Tesla drivers.

- To use graphics acceleration, such as OpenGL, DirectX, or Vulkan, install a GRID driver and separately purchase and configure a GRID license. The GRID driver with a vDWS license also supports CUDA for both computing and graphics acceleration.
 - A graphics-accelerated (G series) ECS created using a public image has had a GRID driver of a specified version installed by default, but the GRID license must be configured separately. Before using such an ECS, check whether the desired driver has been installed on it and whether the version of the installed driver meets service requirements.
 - To install a GRID driver on a GPU-accelerated ECS created using a private image, see [Installing a GRID Driver on a GPU-accelerated ECS](#).
- To use computing acceleration, install a Tesla driver.
 - A computing-accelerated (P series) ECS created using a public image has had a Tesla driver of a specified version installed by default.
 - To install a Tesla driver on a GPU-accelerated ECS created using a private image, see [Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS](#).

Table 3-11 Acceleration supported by GPU drivers

Driver	License	CUDA	OpenGL	DirectX	Vulkan	Application Scenario	Description
GRID	Required	Supported	Supported	Supported	Supported	3D rendering, graphics workstation, and game acceleration	The GRID driver must meet the requirements for accelerating graphics and image applications.
Tesla	Not required	Supported	Not supported	Not supported	Not supported	Scientific computing, deep learning training, and inference	The Tesla driver is downloaded free of charge and usually used with NVIDIA CUDA SDKs to accelerate general computing applications.

3.8.2 Installing a GRID Driver on a GPU-accelerated ECS

Scenarios

To use graphics acceleration, such as OpenGL, DirectX, or Vulkan, install a GRID driver and separately purchase and configure a GRID license. The GRID driver with a vDWS license also supports CUDA for both computing and graphics acceleration.

- A graphics-accelerated (G series) ECS created using a public image has had a GRID driver of a specified version installed by default, but the GRID license must be configured separately.
- If a GPU-accelerated ECS is created using a private image, install a GRID driver and separately configure a GRID license.

This section describes how to install a GRID driver, apply for a GRID license, and configure the license server.

Process of installing a GRID driver:

1. [Configuring a GRID License](#)
2. [Downloading GRID Driver and Software License Packages](#)
3. [Deploying and Configuring the License Server](#)
4. [Installing the GRID Driver and Configuring the License](#)

 **NOTE**

- NVIDIA allows you to apply for a 90-day trial license.
- For details about GPU-accelerated ECSs with different specifications and application scenarios, see [GPU-accelerated ECSs](#).

Configuring a GRID License

- Configure an official license.
To obtain an official license, contact NVIDIA or their NVIDIA agent in your local country or region.
- Apply for a trial license.
Log in at the [official NVIDIA website](#) and enter desired information.
For details about how to sign up for an account and apply for a trial license, see [official NVIDIA help page](#).

 **NOTE**

The method of using a trial license is the same as that of using an official license. You can use an official license to activate an account with a trial license to prevent repetitive registration. The trial license has a validity period of 90 days. After the trial license expires, it cannot be used anymore. Configure an official license then.

Figure 3-41 Applying for a trial license

START YOUR 90-DAY TRIAL

Please register with your corporate email address.
Personal email addresses or extensions will not be approved.
If already registered, [click here](#).
If you need assistance, please review [FAQ](#).

* First name * Last name
* Email address * Phone
* Company * Industry
* Job role * Location
* Street 1 Street 2
* City * State/Province
* Postal Code

* Certified Server * NVIDIA GPUs
Certified Server Other * VDI Hypervisor
* VDI Remoting Client * VDI Seats
* Primary Application

Send me the latest enterprise news, announcements, and more from NVIDIA. I can unsubscribe at any time.

* Required Fields

By registering, you agree to [NVIDIA Account Terms and Conditions](#) & [Privacy Policy](#).

Downloading GRID Driver and Software License Packages

1. Obtain the driver installation package required for an OS. For details, see [Table 3-12](#).

For more information about the GRID driver, see [NVIDIA vGPU Software Documentation](#).

NOTE

For a GPU passthrough ECS, select a GRID driver version as required.

For a GPU virtualization ECS, select a driver version based on the following table.


Table 3-12 GRID driver versions supported by GPU-accelerated ECSs

ECS Type	GPU Attachment	OS	Driver Version	CPU Architecture
G7v	GPU virtualization	<ul style="list-style-type: none"> CentOS 8.2 64bit CentOS 7.6 64bit Ubuntu 20.04 Server 64bit Ubuntu 18.04 Server 64bit Windows Server 2019 Standard 64bit Windows Server 2016 Standard 64bit 	GRID 13.0	x86_64
G7	GPU passthrough	<ul style="list-style-type: none"> CentOS 8.2 64bit CentOS 7.6 64bit Ubuntu 20.04 Server 64bit Ubuntu 18.04 Server 64bit Windows Server 2019 Standard 64bit Windows Server 2016 Standard 64bit 	Select a version as needed.	x86_64
G6	GPU passthrough	<ul style="list-style-type: none"> CentOS 8.2 64bit CentOS 7.6 64bit Ubuntu 20.04 64bit Ubuntu 18.04 64bit Windows Server 2019 Standard 64bit Windows Server 2016 Standard 64bit 	Select a version as needed.	x86_64

ECS Type	GPU Attachment	OS	Driver Version	CPU Architecture
G5.xlarge.4	GPU passthrough	<ul style="list-style-type: none">CentOS 8.2 64bitCentOS 7.6 64bitCentOS 7.5 64bitUbuntu 20.04 64bitUbuntu 18.04 64bitWindows Server 2019 Standard 64bitWindows Server 2016 Standard 64bitWindows Server 2019 Datacenter 64bitWindows Server 2016 Datacenter 64bit	Select a version as needed.	x86_64
P12	GPU passthrough	<ul style="list-style-type: none">CentOS 7.5 64bitWindows Server 2019 Standard 64bitWindows Server 2016 Standard 64bit	Select a version as needed.	x86_64
Pi2nl	GPU passthrough	<ul style="list-style-type: none">CentOS 7.5 64bitUbuntu 16.04 Server 64bitWindows Server 2016 Standard 64bit	Select a version as needed.	x86_64


- After the registration, log in at the [official NVIDIA website](#) and enter the account.
- Check whether NVIDIA is used for the first time.
 - If yes, go to step [4](#).
 - If no, go to step [6](#).
- Obtain the Product Activation Key (PAK) from the email indicating successful registration with NVIDIA.

Figure 3-42 PAK




5. Enter the PAK obtained in step 4 on the **Redeem Product Activation Keys** page and click **Redeem**.

Figure 3-43 Redeem Product Activation Keys



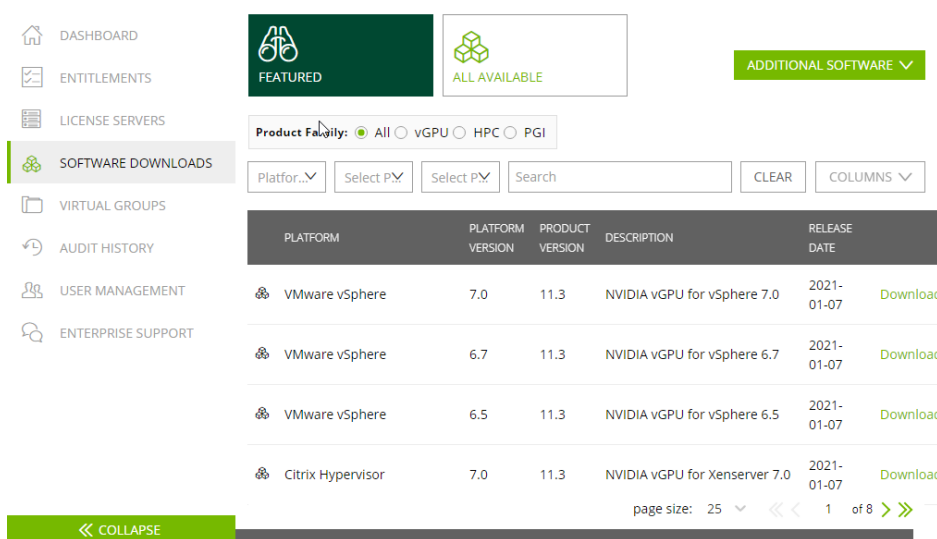
6. Specify **Username** and **Password** and click **LOGIN**.

Figure 3-44 Logging in to the official NVIDIA website




7. Log in at the official NVIDIA website as prompted and select **SOFTWARE DOWNLOADS**.

Figure 3-45 SOFTWARE DOWNLOADS page



8. Download the GRID driver of the required version. For details, see [Table 3-12](#).
9. Decompress the GRID driver installation package and install the driver that matches your ECS OS.
10. On the **SOFTWARE DOWNLOADS** page, click **ADDITIONAL SOFTWARE** to download the license software package.

Figure 3-46 ADDITIONAL SOFTWARE

Deploying and Configuring the License Server

The following uses an ECS running CentOS 7.5 as an example to describe how to deploy and configure the license server on the ECS.

NOTE

- The target ECS must have at least 2 vCPUs and 4 GiB of memory.
- Ensure that the MAC address of the target ECS has been recorded.
- If the license server is used in the production environment, deploy it in high availability mode. For details, see [official NVIDIA documentation for license server high availability](#).

1. Configure the network.

- If the license server is to be accessed using the VPC:

- Single-user: Ensure that the license server and the GPU-accelerated ECS with the GRID driver installed are in the same VPC subnet.
- Multi-user: If the license server and the GPU-accelerated ECS with the GRID driver installed use different accounts, perform the following operations:

For example, the license server uses account A and needs to use a license from account B. Additionally, the GPU-accelerated ECS under account B has had the GRID driver installed.

- 1) Use account A to create a VPC endpoint service and set **Backend Resource Type** to **ECS**.
- 2) Add account B's domain ID to the whitelist of the VPC endpoint service under account A.
- 3) Under account B, create a VPC endpoint to access the VPCEP service. Then, use the node IP address or private network domain name to authorize account A to use the license.

For details, see "Configuring a VPC Endpoint for Communication Across VPCs of Different Accounts" in *VPC Endpoint User Guide*.

- If the license server is to be accessed using a public IP address, configure the security group to which license server belongs and add inbound rules for TCP 7070 and TCP 8080.
2. Install the license server.
 - a. Run the following command to decompress the installation package. The **Installer.zip** in the command indicates the name of the software package obtained in [10](#).
unzip Installer.zip
 - b. Run the following command to assign execution permissions to the installer:
chmod +x setup.bin
 - c. Run the installer as user **root**:
sudo ./setup.bin -i console
 - d. In the Introduction section, press **Enter** to continue.

```
=====
Introduction
-----

InstallAnywhere will guide you through the installation of License Server.

It is strongly recommended that you quit all programs before continuing with
this installation.

Respond to each prompt to proceed to the next step in the installation. If
you want to change something on a previous step, type 'back'.

You may cancel this installation at any time by typing 'quit'.

PRESS <ENTER> TO CONTINUE: █
```

- e. In the License Agreement section, press **Enter** to turn to last pages and accept the license agreement.
Enter **Y** and press **Enter**.
- ```
DO YOU ACCEPT THE TERMS OF THIS LICENSE AGREEMENT? (Y/N): Y █
```
- f. In the Choose Install Folder section, press **Enter** to retain the default path for installing the License Server software.
  - g. In the Choose Local Tomcat Server Path section, enter the Tomcat's local path in the `"/var/lib/Tomcat version"` format, for example, `/var/lib/tomcat8`.
  - h. In the Choose Firewall Options section, confirm the port to be enabled in the firewall and press **Enter**.

```
Choose Firewall Options

The license server listens on port 7070. This port must be opened in the
firewall for other machines to obtain licenses from this server.

The license server's management interface listens on port 8080. Leave this
port closed to prevent unauthorized access to the management interface.

->1- License server (port 7070)
 2- Management interface (port 8080)

ENTER A COMMA-SEPARATED LIST OF NUMBERS REPRESENTING THE DESIRED CHOICES, OR
PRESS <ENTER> TO ACCEPT THE DEFAULT: █
```

- i. In the Pre-Installation Summary section, confirm the information and press **Enter** to start the installation.

```
Pre-Installation Summary

Please Review the Following Before Continuing:

Product Name:
 License Server

Install Folder:
 /opt/flexnet1s/nvidia

Link Folder:
 /root/NVIDIA Corporation/License Server

Disk Space Information (for Installation Target):
 Required: 105,216,774 Bytes
 Available: 35,501,248,512 Bytes

PRESS <ENTER> TO CONTINUE: █
```

- j. In the Install Complete section, press **Enter** to end the installation.

```
Install Complete

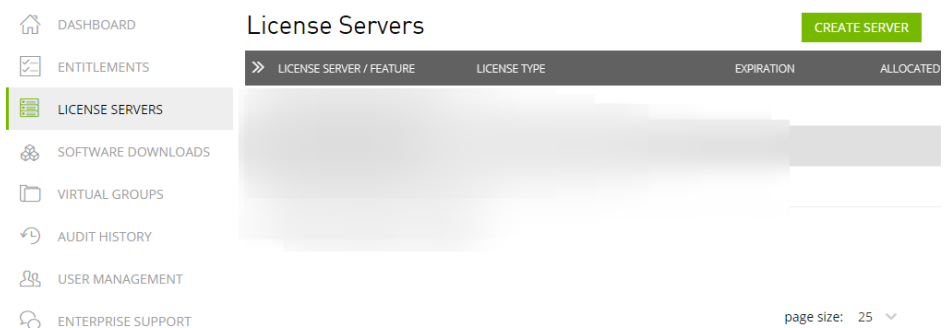
License Server has been successfully installed to:

 /opt/flexnet1s/nvidia

PRESS <ENTER> TO EXIT THE INSTALLER:
```

- 3. Obtain the license file.
  - a. Log in to the [NVIDIA website](#) on a new tab and select **LICENSE SERVERS**.

**Figure 3-47** LICENSE SERVERS



- b. Click **CREATE SERVER**.

- c. On the displayed **Create License Server** page, configure parameters.

**Figure 3-48** Create License Server

**Create License Server** X

**Server Name**  
Name this license server

**Feature**  
Select a feature

**Licenses**  
1 ADD

**Description**  
Provide a short description

**MAC Address**  
MAC Address (XX:XX:XX:XX:XX:XX or XX-XX-XX-XX-XX-XX)

Failover server configuration is optional.  
If configuring, you must provide a name AND MAC address

**Failover License Server**  
Failover License Server

**Failover MAC Address**  
Failover MAC Address

**Added Features**


| FEATURE                         | COUNT |
|---------------------------------|-------|
| No features have been added yet |       |

CANCEL RESET CREATE LICENSE SERVER


**Table 3-13** Parameters for creating a license server

| Parameter   | Description                                                                                                                                                                                                                                                                        |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Server Name | License server name, which can be customized.                                                                                                                                                                                                                                      |
| Description | License description information.                                                                                                                                                                                                                                                   |
| MAC Address | MAC address of the ECS where the license server is deployed.<br>You can log in to the ECS and run <b>ipconfig -a</b> to query the MAC address.                                                                                                                                     |
| Feature     | Select a feature, enter the number of required licenses in the <b>Licenses</b> text box, and click <b>ADD</b> .<br>In active/standby deployment, enter the name of the standby server in <b>Failover License Server</b> and enter the MAC address in <b>Failover MAC Address</b> . |

- d. Click **CREATE LICENSE SERVER**.
- e. Download the license file.

**Figure 3-49** Downloading the license file

4. In the web browser, access the homepage of the license server management page using the link configured during the installation.  
Default URL: `http://IP address of the EIP.8080/licserver`
5. In the navigation pane on the left, click **License Server** > **License Management**.
6. Select the .bin license file to be uploaded and click **Upload**.

**Figure 3-50** Uploading a license file


## Installing the GRID Driver and Configuring the License

1. Install the GRID driver of a desired version, for example, on a GPU-accelerated Windows ECS.

### NOTE

Microsoft remote login protocols do not support GPU 3D hardware acceleration. To use this function, install third-party desktop protocol-compliant software, such as VNC, PCoIP, or NICE DCV, and access the ECS through the client.

2. Open the NVIDIA control panel on the Windows control panel.
3. Enter the IP address and port number of the deployed license server in the level-1 license server, and then click **Apply**. If the message indicating that you have obtained a GRID license is displayed, the installation is successful. Additionally, the MAC address of the GPU-accelerated ECS with the GRID driver installed is displayed on the **Licensed Clients** page of the license server management console.

**Figure 3-51** License server management console



## 3.8.3 Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS

### Scenarios

Before using a GPU-accelerated ECS, make sure that the desired Tesla driver and CUDA toolkit have been installed on the ECS for computing acceleration.

- A computing-accelerated (P series) ECS created using a public image has had a Tesla driver of a specified version installed by default.
- After a GPU-accelerated ECS is created using a private image, it must have a Tesla driver installed. Otherwise, computing acceleration will not take effect.

This section describes how to install a Tesla driver and CUDA toolkit on a GPU-accelerated ECS.

### Notes

- The ECS must have an EIP bound.
- Check whether the CUDA toolkit and Tesla driver have been installed on the ECS.

#### NOTE

- If the CUDA toolkit has not been installed, download it from the official NVIDIA website and install it. A Tesla driver matching the CUDA version will be automatically installed then. However, if there are specific requirements or dependencies on the Tesla driver version, download the matching Tesla driver from the official NVIDIA website first and then install the driver before installing the CUDA toolkit.
- If a Tesla driver has been installed on the ECS, check the driver version. Before installing a new driver version, uninstall the original Tesla driver to prevent an installation failure due to driver conflicts.

Installation process:

- [Obtaining a Tesla Driver and CUDA Toolkit](#)
- Installing a Tesla Driver
  - [Installing a Tesla Driver on a Linux ECS](#)
  - [Installing a Tesla Driver on a Windows ECS](#)
- Installing a CUDA Toolkit
  - [Installing the CUDA Toolkit on a Linux ECS](#)
  - [Installing the CUDA Toolkit on a Windows ECS](#)

### Installing a Tesla Driver on a Linux ECS

The following uses Ubuntu 16.04 64bit as an example to describe how to install the Tesla driver matching CUDA 10.1 on a GPU-accelerated ECS.

 NOTE

The Linux kernel version is compatible with the driver version. If installing the driver failed, check the driver installation log, which is generally stored in `/var/log/nvidia-installer.log`. If the log shows that the failure was caused by a driver compilation error, for example, the `get_user_pages` parameter setting is incorrect, the kernel version is incompatible with the driver version. In such a case, select the desired kernel version and driver version and reinstall them. It is recommended that the release time of the kernel version and driver version be the same.

1. Log in to the ECS.
2. Update the system software based on the OS.
  - Ubuntu  
Update the software installation source: **`apt-get -y update`**  
Install necessary programs: **`apt-get install gcc g++ make`**
  - CentOS  
Update the software installation source: **`yum -y update --exclude=kernel* --exclude=centos-release* --exclude=initcripts*`**  
Install the desired program: **`yum install -y kernel-devel `uname -r` gcc gcc-c++`**
3. Download the NVIDIA driver package.  
Select a driver version at [NVIDIA Driver Downloads](#) based on the ECS type.  
Click **SEARCH**.

**Figure 3-52** Selecting a NVIDIA driver version

### NVIDIA Driver Downloads


Advanced Driver Search

|                                    |                                           |
|------------------------------------|-------------------------------------------|
| Product Type:                      | Operating System:                         |
| <input type="text" value="Tesla"/> | <input type="text" value="Linux 64-bit"/> |
| Product Series:                    | CUDA Toolkit:                             |
| <input type="text"/>               | <input type="text" value="10.1"/>         |
| Product:                           | Language:                                 |
| <input type="text"/>               | <input type="text" value="English (US)"/> |
|                                    | Recommended/Beta:                         |
|                                    | <input type="text" value="All"/> ?        |

4. Select a driver version as required. The following uses Tesla 418.67 as an example.

Figure 3-53 Selecting a driver version

### NVIDIA Driver Downloads



5. Click the driver to be downloaded. On the **TESLA DRIVER FOR LINUX X64** page that is displayed, click **DOWNLOAD**.
6. Copy the download link.

Figure 3-54 Copying the download link


### Download

By clicking the "Agree & Download" button below, you are confirming that you have read and agree to be bound by the [License For Customer Use of NVIDIA Software](#) for use of the driver. The driver will begin downloading immediately after clicking on the "Agree & Download" button below. NVIDIA recommends users update to the latest driver version. Please review [NVIDIA Product Security](#) for more information.



7. Run the following command on the ECS to download the driver:  
**wget Copied link**  
For example, **wget http://us.download.nvidia.com/tesla/418.67/NVIDIA-Linux-x86\_64-418.67.run**

Figure 3-55 Obtaining the installation package




8. Run the following command to install the driver:

```
sh NVIDIA-Linux-x86_64-418.67.run
```


9. (Optional) If the following information is displayed after the command for installing the driver is executed, disable the Nouveau driver.

**Figure 3-56** Disabling the Nouveau driver



- a. Run the following command to check whether the Nouveau driver has been installed:  
**lsmod | grep nouveau**
    - If the command output contains information about the Nouveau driver, the Nouveau driver has been installed and must be disabled. Then, go to step [9.b](#).
    - If the command output does not contain information about the Nouveau driver, the Nouveau driver has been disabled. Then, go to step [10](#).
  - b. Edit the **blacklist.conf** file.  
If the **/etc/modprobe.d/blacklist.conf** file is unavailable, create it.  
**vi /etc/modprobe.d/blacklist.conf**  
Add the following statement to the end of the file:

```
blacklist nouveau
options nouveau modeset=0
```
  - c. Run the following command to back up and create an initramfs application:
    - Ubuntu  
**sudo update-initramfs -u**
    - CentOS:  
**mv /boot/initramfs-\$(uname -r).img /boot/initramfs-\$(uname -r).img.bak**  
**dracut -v /boot/initramfs-\$(uname -r).img \$(uname -r)**
  - d. Restart the ECS:  
**reboot**
10. Select **OK** for three consecutive times as prompted to complete the driver installation.

**Figure 3-57** Completing the NVIDIA driver installation

11. Run the following command to set systemd:  
**systemctl set-default multi-user.target**
12. Run the **reboot** command to restart the ECS.
13. Log in to the ECS and run the **nvidia-smi** command. If the command output contains the installed driver version, the driver has been installed.

**Figure 3-58** Viewing the NVIDIA driver version

```
root@ecs-474b:~# nvidia-smi
Thu Mar 26 20:05:17 2020

+-----+
| NVIDIA-SMI 418.67 Driver Version: 418.67 CUDA Version: 10.1 |
+-----+
| GPU Name Persistence-M| Bus-Id Disp.A | Volatile Uncorr. ECC |
| Fan Temp Perf Pwr:Usage/Cap| Memory-Usage | GPU-Util Compute M. |
+-----+-----+
| 0 Tesla Off | 00000000:21:01.0 Off | |
| N/A 52C P0 29W / 70W | 0MiB / 15079MiB | 0% Default |
+-----+-----+

+-----+
| Processes: GPU Memory |
| GPU PID Type Process name Usage |
+-----+-----+
| No running processes found |
+-----+

root@ecs-474b:~#
```

## Installing a Tesla Driver on a Windows ECS

The following uses Windows Server 2016 Standard 64bit as an example to describe how to install a Tesla driver on a GPU-accelerated ECS.

1. Log in to the ECS.
2. Download the NVIDIA driver package.  
Select a driver version at [NVIDIA Driver Downloads](#) based on the ECS type.

**Figure 3-59** Selecting a driver type (Windows)

Advanced Driver Search

|                                                     |                                                                       |
|-----------------------------------------------------|-----------------------------------------------------------------------|
| Product Type:<br><input type="text" value="Tesla"/> | Operating System:<br><input type="text" value="Windows Server 2016"/> |
| Product Series:<br><input type="text"/>             | CUDA Toolkit:<br><input type="text" value="10.1"/>                    |
| Product:<br><input type="text"/>                    | Language:<br><input type="text" value="English (US)"/>                |
|                                                     | Recommended/Beta:<br><input type="text" value="All"/> ?               |

3. Select a driver version as required. The following uses Tesla 425.25 as an example.

**Figure 3-60** Selecting a driver version (Windows)

Advanced Driver Search

Product Type:

Operating System:

Product Series:

CUDA Toolkit:

Product:

Language:

Recommended/Beta:  ?

| Name                                                   | Version | Release Date      | CUDA Toolkit |
|--------------------------------------------------------|---------|-------------------|--------------|
| <input type="checkbox"/> Tesla Driver for Windows WHQL | 426.50  | February 28, 2020 | 10.1         |
| <input type="checkbox"/> Tesla Driver for Windows WHQL | 426.32  | December 9, 2019  | 10.1         |
| <input type="checkbox"/> Tesla Driver for Windows WHQL | 426.23  | October 3, 2019   | 10.1         |
| <input type="checkbox"/> Tesla Driver for Windows WHQL | 426.00  | August 14, 2019   | 10.1         |
| <input type="checkbox"/> Tesla Driver for Windows WHQL | 425.25  | May 7, 2019       | 10.1         |
| <input type="checkbox"/> Tesla Driver for Windows WHQL | 419.69  | March 25, 2019    | 10.1         |

4. Click the driver to be downloaded. On the **TESLA DRIVER FOR WINDOWS** page that is displayed, click **DOWNLOAD**.
5. Click **AGREE & DOWNLOAD** to download the installation package.

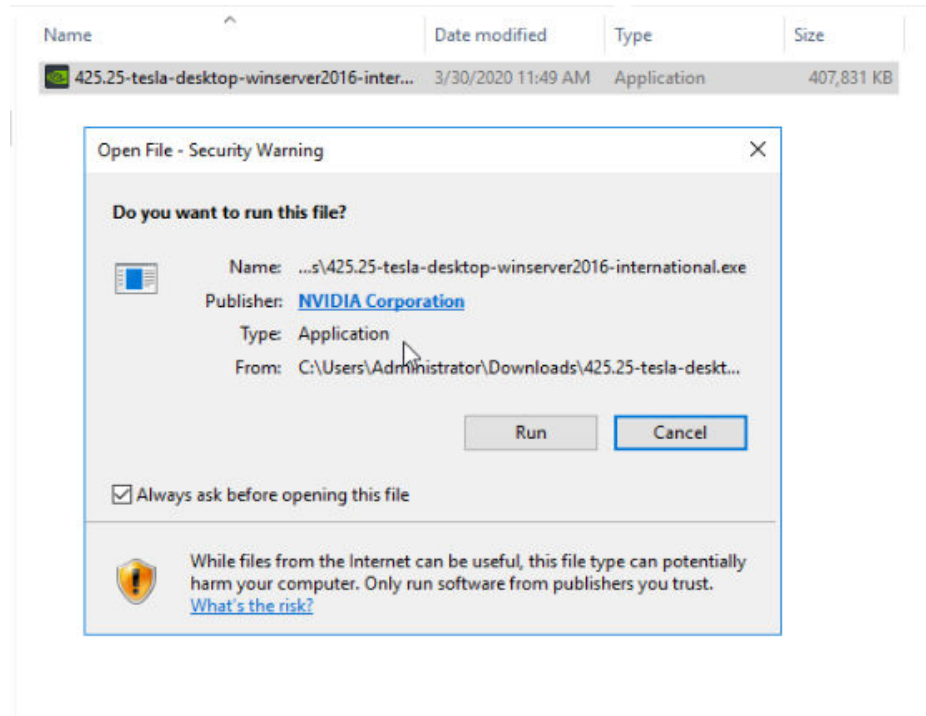
**Figure 3-61** Downloading the driver installation package

### Download

By clicking the "Agree & Download" button below, you are confirming that you have read and agree to be bound by the [License For Customer Use of NVIDIA Software](#) for use of the driver. The driver will begin downloading immediately after clicking on the "Agree & Download" button below. NVIDIA recommends users update to the latest driver version. Please review [NVIDIA Product Security](#) for more information.


6. Double-click the driver and click **Run**.

**Figure 3-62** Running the NVIDIA driver installation program



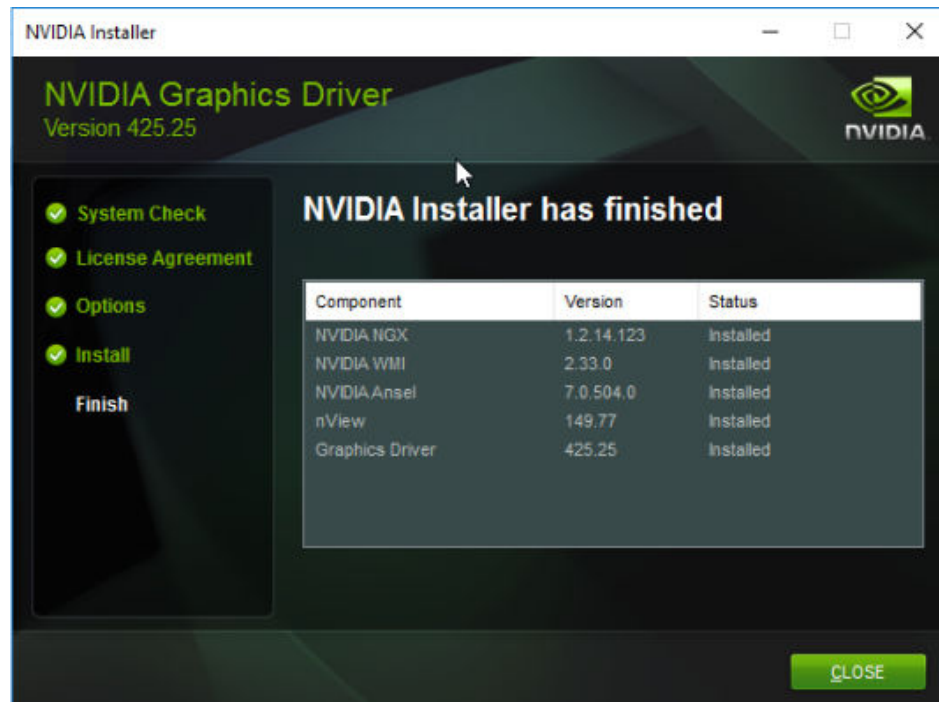
- 7. Select an installation path and click **OK**.

**Figure 3-63** Selecting an installation path



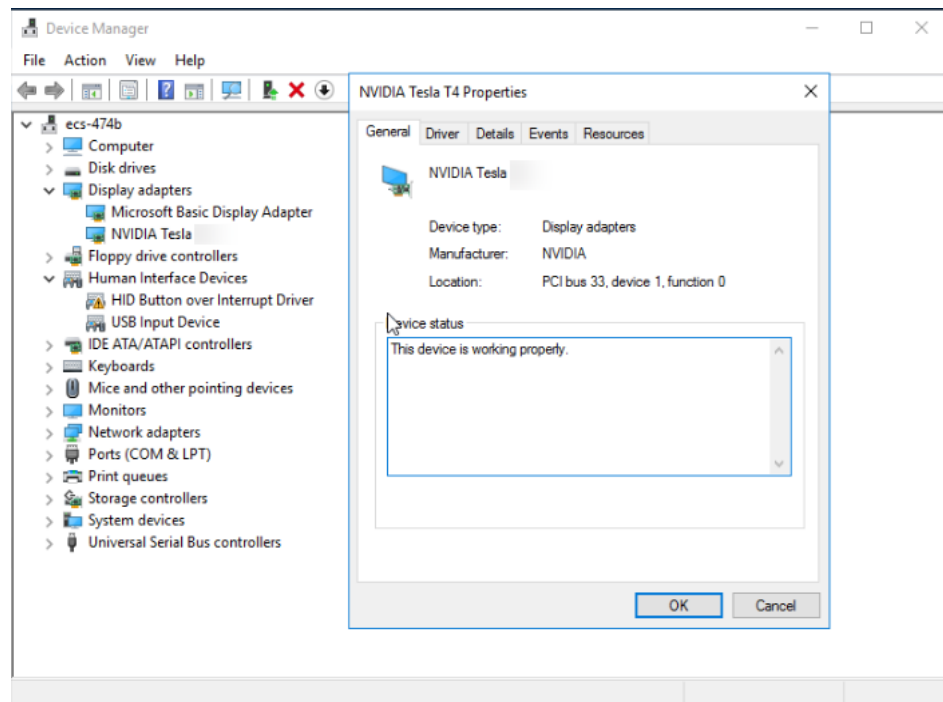
- 8. Install the NVIDIA program as prompted.

Figure 3-64 Completing the driver installation



9. Restart the ECS.
10. Check whether the NVIDIA driver has been installed.
  - a. Switch to **Device Manager** and click **Display adapters**.

Figure 3-65 Display adapters



- b. Open the **cmd** window on the ECS and run the following commands:  
**cd C:\Program Files\NVIDIA Corporation\NVSMI**



### nvidia-smi

If the command output contains the installed driver version, the driver has been installed.

**Figure 3-66** Viewing the NVIDIA driver version


```
C:\Program Files\NVIDIA Corporation\NVSMI>nvidia-smi
2020
+-----+
| NVIDIA-SMI 425.25 Driver Version: 425.25 CUDA Version: 10.1 |
+-----+-----+-----+-----+-----+-----+
| GPU Name TCC/WDDM | Bus-Id Disp.A | Volatile Uncorr. ECC |
| Fan Temp Perf Pwr:Usage/Cap| Memory-Usage | GPU-Util Compute M. |
+-----+-----+-----+-----+-----+-----+
| 0 Tesla TCC | 00000000:21:01.0 Off | | 0% Default |
|N/A 33C P8 11W / 70W | 0MiB / 15205MiB | | |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| Processes: GPU Memory |
| GPU PID Type Process name Usage |
+-----+-----+-----+-----+-----+-----+
| No running processes found | |
+-----+-----+-----+-----+-----+
C:\Program Files\NVIDIA Corporation\NVSMI>
```

## Installing the CUDA Toolkit on a Linux ECS

The following uses Ubuntu 16.04 64bit as an example to describe how to install the CUDA 10.1 toolkit on a GPU-accelerated ECS.


1. Log in to the ECS.
2. Update the system software based on the OS.
  - Ubuntu  
Update the software installation source: **apt-get -y update**  
Install necessary programs: **apt-get install gcc g++ make**
  - CentOS  
Update the software installation source: **yum -y update --exclude=kernel\* --exclude=centos-release\* --exclude=initscripts\***  
Install the desired program: **yum install -y kernel-devel`uname -r` gcc gcc-c++**
3. On the CUDA download page, set parameters according to the information shown in [Obtaining a Tesla Driver and CUDA Toolkit](#).

**Figure 3-67** Selecting a CUDA version



- 4. Find the link for downloading CUDA 10.1 and copy the link.

Figure 3-68 Copying the link for downloading CUDA




- 5. Run the following command on the ECS to download CUDA:

**wget Copied link**

For example, **wget https://developer.nvidia.com/compute/cuda/10.1/Prod/local\_installers/cuda\_10.1.105\_418.39\_linux.run**


Figure 3-69 Downloading CUDA



- 6. Install CUDA.

Follow the instructions provided on the official NVIDIA website.

Figure 3-70 Installing CUDA



- 7. Run the following command to install CUDA:

**sh cuda\_10.1.243\_418.87.00\_linux.run**

8. Select **accept** on the installation page and press **Enter**.

**Figure 3-71** Installing CUDA\_1

```
End User License Agreement

Preface

The Software License Agreement in Chapter 1 and the Supplement
in Chapter 2 contain license terms and conditions that govern
the use of NVIDIA software. By accepting this agreement, you
agree to comply with all the terms and conditions applicable
to the product(s) included herein.

NVIDIA Driver

Description

This package contains the operating system driver and

Do you accept the above EULA? (accept/decline/quit):
accept
```

9. Select **Install** and press **Enter** to start the installation.

**Figure 3-72** Installing CUDA\_2

```
CUDA Installer
- [X] Driver
 [X] 418.39
+ [X] CUDA Toolkit 10.1
 [X] CUDA Samples 10.1
 [X] CUDA Demo Suite 10.1
 [X] CUDA Documentation 10.1
 Install
 Options

Up/Down: Move | Left/Right: Expand | 'Enter': Select | 'A': Advanced options
```

**Figure 3-73** Completing the installation

```
=====
= Summary =
=====
Driver: Installed
Toolkit: Installed in /usr/local/cuda-10.1/
Samples: Installed in /root/, but missing recommended libraries

Please make sure that
- PATH includes /usr/local/cuda-10.1/bin
- LD_LIBRARY_PATH includes /usr/local/cuda-10.1/lib64, or, add /usr/local/cuda-10.1/lib64 to /etc/ld.so.conf and run ldconfig
as root

To uninstall the CUDA Toolkit, run cuda-uninstaller in /usr/local/cuda-10.1/bin
To uninstall the NVIDIA Driver, run nvidia-uninstall

Please see CUDA_Installation_Guide_Linux.pdf in /usr/local/cuda-10.1/doc/pdf for detailed information on setting up CUDA.
Logfile is /var/log/cuda-installer.log
root@ecs-474b:~#
```

10. Run the following command to switch to `/usr/local/cuda-10.1/samples/1_Utilities/deviceQuery`:  
**cd /usr/local/cuda-10.1/samples/1\_Utilities/deviceQuery**
  11. Run the **make** command to automatically compile the deviceQuery program.
  12. Run the following command to check whether CUDA has been installed:  
**./deviceQuery**
- If the command output contains the CUDA version, CUDA has been installed.

**Figure 3-74** deviceQuery common output

```
root@ecs-474b:/usr/local/cuda-10.1/samples/1_Utilities/deviceQuery# ./deviceQuery
./deviceQuery Starting...

CUDA Device Query (Runtime API) version (CUDA RT static linking)
Detected 1 CUDA Capable device(s)
Device 0: "Tesla "
 CUDA Driver Version / Runtime Version 10.1 / 10.1
 CUDA Capability Major/Minor version number: 7.5
 Total amount of global memory: 15080 MBytes (15812263936 bytes)
 (40) Multiprocessors, (64) CUDA Cores/MP: 2560 CUDA Cores
 GPU Max Clock rate: 1590 Mhz (1.59 GHz)
 Memory Clock rate: 5001 Mhz
 Memory Bus Width: 256-bit
 L2 Cache Size: 4194304 bytes
 Maximum Texture Dimension Size (x,y,z) 1D=(131072), 2D=(131072, 65536), 3D=(16384, 16384, 16384)
 Maximum Layered 1D Texture Size, (num) layers 1D=(32768), 2048 layers
 Maximum Layered 2D Texture Size, (num) layers 2D=(32768, 32768), 2048 layers
 Total amount of constant memory: 65536 bytes
 Total amount of shared memory per block: 49152 bytes
 Total number of registers available per block: 65536
 Warp size: 32
 Maximum number of threads per multiprocessor: 1024
 Maximum number of threads per block: 1024
 Max dimension size of a thread block (x,y,z): (1024, 1024, 64)
 Max dimension size of a grid size (x,y,z): (2147483647, 65535, 65535)
 Maximum memory pitch: 2147483647 bytes
 Texture alignment: 512 bytes
 Concurrent copy and kernel execution: Yes with 3 copy engine(s)
 Run time limit on kernels: No
 Integrated GPU sharing Host Memory: No
 Support host page-locked memory mapping: Yes
 Alignment requirement for Surfaces: Yes
 Device has ECC support: Enabled
 Device supports Unified Addressing (UVA): Yes
 Device supports Compute Preemption: Yes
 Supports Cooperative Kernel Launch: Yes
 Supports MultiDevice Co-op Kernel Launch: Yes
 Device PCI Domain ID / Bus ID / location ID: 0 / 33 / 1
 Compute Mode:
 < Default (multiple host threads can use ::cudaSetDevice() with device simultaneously) >
deviceQuery, CUDA Driver = CUDART, CUDA Driver Version = 10.1, CUDA Runtime Version = 10.1, NumDevs = 1
Result = PASS
root@ecs-474b:/usr/local/cuda-10.1/samples/1_Utilities/deviceQuery#
```

13. Check the CUDA version.  
**/usr/local/cuda/bin/nvcc -V**

**Figure 3-75** Checking the CUDA version

```
root@ecs-474b:/usr/local/cuda-10.1/samples/1_Utilities/deviceQuery# ./deviceQuery
[roo@ecs-474b deviceQuery]# /usr/local/cuda/bin/nvcc -V
nvcc: NVIDIA (R) Cuda compiler driver
Copyright (c) 2005-2019 NVIDIA Corporation
Built on Fri Feb 8 19:08:17 PST 2019
Cuda compilation tools, release 10.1, V10.1.105
[roo@ecs-474b deviceQuery]#
```


14. Run the following command to enable the persistent mode:  
**sudo nvidia-smi -pm 1**
- Enabling the persistent mode optimizes the GPU performance on Linux ECSs.

## Installing the CUDA Toolkit on a Windows ECS

The following uses Windows Server 2016 Standard 64bit as an example to describe how to install the CUDA 10.1 toolkit on a GPU-accelerated ECS.


1. Log in to the ECS.
2. On the CUDA download page, set parameters according to the information shown in [Downloading a CUDA Toolkit](#).

**Figure 3-76** Selecting a CUDA version




3. Find the link for downloading CUDA 10.1.

**Figure 3-77** Finding the link for downloading CUDA




4. Click **Download** to download the CUDA toolkit.
5. Double-click the installation file and click **Run** to install the CUDA toolkit.

**Figure 3-78** Installing CUDA




6. On the **CUDA Setup Package** page, select an installation path and click **OK**.

**Figure 3-79** Selecting an installation path



7. Install the CUDA toolkit as prompted.

**Figure 3-80** Completing the installation

8. Check whether CUDA has been installed  
Open the **cmd** window and run the following command:  
**nvcc -V**  
If the command output contains the CUDA version, CUDA has been installed.

**Figure 3-81** Successful installation

```
C:\Users\Administrator>nvcc -V
nvcc: NVIDIA (R) Cuda compiler driver
Copyright (c) 2005-2019 NVIDIA Corporation
Built on Fri Feb 8 19:08:26 Pacific Standard Time 2019
Cuda compilation tools, release 10.1, V10.1.105

C:\Users\Administrator>
```

## 3.8.4 Obtaining a Tesla Driver and CUDA Toolkit

### Scenarios

Before using a GPU-accelerated ECS, make sure that the desired Tesla driver and CUDA toolkit have been installed on the ECS. Otherwise, computing acceleration will not take effect. This section describes how to obtain a Tesla driver and CUDA toolkit. Select a driver version based on your ECS type.

For instructions about how to install the Tesla driver and CUDA toolkit, see [Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS](#).

### Downloading a Tesla Driver

[Download a driver](#) based on your ECS type.

**Table 3-14** Mapping between Tesla drivers and ECS types

| ECS Type | Driver | Product Series | Product |
|----------|--------|----------------|---------|
| Pi2      | Tesla  | T              | T4      |
| Pi2nl    | Tesla  | T              | T4      |
| G7       | Tesla  | A              | A40     |
| G6       | Tesla  | T              | T4      |
| G5       | Tesla  | V              | V100    |

## Downloading a CUDA Toolkit

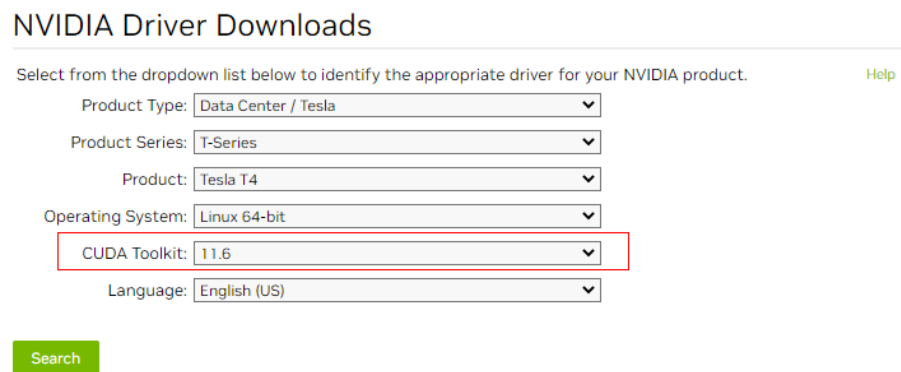
Download the [CUDA software package](#) and select the corresponding CUDA Toolkit software package based on the instance type and driver version.

### NOTE

[NVIDIA Driver Downloads](#) provides the mapping between the driver version and CUDA Toolkit. If the versions do not match, the driver may be unavailable.

The following uses Tesla T4 as an example to describe how to download the driver package and CUDA Toolkit.

1. Select the Linux operating system and the CUDA Toolkit 11.6 version.

**Figure 3-82** Selecting the CUDA Toolkit version

The screenshot shows the 'NVIDIA Driver Downloads' page. It features a series of dropdown menus for selecting product type, series, product, operating system, CUDA Toolkit version, and language. The 'CUDA Toolkit' dropdown is highlighted with a red box, showing '11.6' selected. A green 'Search' button is located below the dropdowns.

NVIDIA Driver Downloads

Select from the dropdown list below to identify the appropriate driver for your NVIDIA product. [Help](#)

Product Type: Data Center / Tesla

Product Series: T-Series

Product: Tesla T4

Operating System: Linux 64-bit

CUDA Toolkit: 11.6

Language: English (US)

[Search](#)

2. Select a CUDA Toolkit 11.6 package to download.



**Figure 3-83** Downloading a CUDA Toolkit 11.6 package

**Archived Releases**

[CUDA Toolkit 11.7.1 \(August 2022\), Versioned Online Documentation](#)  
[CUDA Toolkit 11.7.0 \(May 2022\), Versioned Online Documentation](#)  
[CUDA Toolkit 11.6.2 \(March 2022\), Versioned Online Documentation](#)  
[CUDA Toolkit 11.6.1 \(February 2022\), Versioned Online Documentation](#)  
[CUDA Toolkit 11.6.0 \(January 2022\), Versioned Online Documentation](#)  
[CUDA Toolkit 11.5.2 \(February 2022\), Versioned Online Documentation](#)  
[CUDA Toolkit 11.5.1 \(November 2021\), Versioned Online Documentation](#)  
[CUDA Toolkit 11.5.0 \(October 2021\), Versioned Online Documentation](#)  
[CUDA Toolkit 11.4.4 \(February 2022\), Versioned Online Documentation](#)  
[CUDA Toolkit 11.4.3 \(November 2021\), Versioned Online Documentation](#)  
[CUDA Toolkit 11.4.2 \(September 2021\), Versioned Online Documentation](#)  
[CUDA Toolkit 11.4.1 \(August 2021\), Versioned Online Documentation](#)  
[CUDA Toolkit 11.4.0 \(June 2021\), Versioned Online Documentation](#)

# 4 Images

---

## 4.1 Overview

### Image

An image is an ECS or BMS template that contains an OS or service data. It may also contain proprietary software and application software, such as database software. Images are classified into public, private, and shared images.

Image Management Service (IMS) allows you to easily create and manage images. You can create an ECS using a public image, private image, or shared image. You can also use an existing ECS or external image file to create a private image.

### Public Image

A public image is a standard, widely used image that contains a common OS, such as Ubuntu, CentOS, or Debian, and preinstalled public applications. This image is available to all users. Select your desired public image. Alternatively, create a private image based on a public image to copy an existing ECS or rapidly create ECSs in a batch. You can customize a public image by configuring the application environment or software.

### Private Image

A private image contains an OS or service data, preinstalled public applications, and private applications. It is available only to the user who created it.

**Table 4-1** Private image types

| Image Type        | Description                                                                                                                                          |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| System disk image | Contains an OS and application software for running services. You can use a system disk image to create ECSs and migrate your services to the cloud. |

| Image Type      | Description                                                                                                                                       |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| Data disk image | Contains only service data. You can use a data disk image to create EVS disks and migrate your service data to the cloud.                         |
| Full-ECS image  | Contains an OS, application software, and data for running services. A full-ECS image contains the system disk and all data disks attached to it. |
| ISO image       | Created from an external ISO image file. It is a special image that can only be used to create temporary ECSs.                                    |

If you plan to use a private image to change the OS, ensure that the private image is available. For instructions about how to create a private image, see *Image Management Service User Guide*.

- If the image of a specified ECS is required, make sure that a private image has been created using this ECS.
- If a local image file is required, make sure that the image file has been imported to the cloud platform and registered as a private image.
- If a private image from another region is required, make sure that the image has been copied.
- If a private image from another user account is required, make sure that the image has been shared with you.

## Shared Image

A shared image is a private image shared by another user and can be used as your own private image.

- Images can be shared within a region only.
- Each image can be shared to a maximum of 128 tenants.
- You can stop sharing images anytime without notifying the recipient.
- You can delete shared image anytime without notifying the recipient.
- Full-ECS images cannot be shared.

## 4.2 Creating an Image

### Scenarios

You can use an existing ECS to create a system disk image, data disk image, and full-ECS image.

- **System disk image:** contains an OS and application software for running services. You can use a system disk image to create ECSs and migrate your services to the cloud.
- **Data disk image:** contains only service data. You can create a data disk image from an ECS data disk. You can also use a data disk image to create EVS disks and migrate your service data to the cloud.

- Full-ECS image: contains all the data of an ECS, including the data on the data disks attached to the ECS. A full-ECS image can be used to rapidly create ECSs with service data.
- ISO image: is created from an external ISO image file. It is a special image that can only be used to create temporary ECSs.

You can use a private image to change the OS. For instructions about how to create a private image, see *Image Management Service User Guide*.

## Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. In the ECS list, choose **More > Manage Image > Create Image** in the **Operation** column.
4. Configure the following information:  
**Table 4-2** and **Table 4-3** list the parameters in the **Image Type and Source** and **Image Information** areas, respectively.

**Table 4-2** Image type and source

| Parameter  | Description                                                              |
|------------|--------------------------------------------------------------------------|
| Image Type | Select <b>System disk image</b> .                                        |
| Source     | Click the <b>ECS</b> tab and select an ECS with required configurations. |

**Table 4-3** Image information

| Parameter   | Description                                  |
|-------------|----------------------------------------------|
| Name        | Set a name for the image.                    |
| Description | (Optional) Enter a description of the image. |

### NOTE

The parameters may vary depending on enterprises and organizations.

5. Click **Create Now**.

# 5 EVS Disks

---

## 5.1 Adding a Disk to an ECS

### Scenarios

The disks attached to an ECS include one system disk and one or more data disks. The system disk is automatically created and attached when the ECS is created. You do not need to add it again. The data disks can be added in either of the following ways:

- If you add data disks when creating an ECS, the data disks will be automatically attached to the ECS.
- If you create data disks after an ECS is created, the data disks need to be manually attached to the ECS.

This section describes how to add a data disk after creating an ECS.

### Procedure

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. Locate the row containing the target ECS and choose **More > Manage Image/Disk/Backup > Add Disk** in the **Operation** column.

The page for adding a disk is displayed.

4. Set parameters for the new EVS disk as prompted.

For instructions about how to set EVS disk parameters, see "Create an EVS Disk" in *Elastic Volume Service User Guide*.

#### NOTE

- By default, the billing mode of the new disk is the same as that of the ECS.
- By default, the new disk is in the same region as the ECS.
- By default, the new disk is in the same AZ as the ECS, and the AZ of the disk cannot be changed.
- After the new disk is created, it is attached to the ECS by default.

5. Click **Create Now**.

The system automatically switches back to the **Disks** tab on the ECS management console. Then, you can view the information of the new disk.

## Follow-up Procedure

The system automatically attaches the new disk to the ECS, but the disk can be used only after it is initialized. To do so, log in to the ECS and initialize the disk.

For details about how to initialize a data disk, see [Scenarios and Disk Partitions](#).

## 5.2 Attaching an EVS Disk to an ECS

### Scenarios

If the existing disks of an ECS fail to meet service requirements, for example, due to insufficient disk space or poor disk performance, you can attach more available EVS disks to the ECS, or create more disks (in **Storage > Elastic Volume Service**) and attach them to the ECS.

### Prerequisites

- EVS disks are available.  
For instructions about how to create an EVS disk, see "Creating an EVS Disk" in *Elastic Volume Service User Guide*.

### Procedure

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID for search.
4. Click the name of the target ECS.  
The page providing details about the ECS is displayed.
5. Click the **Disks** tab. Then, click **Attach Disk**.  
The **Attach Disk** dialog box is displayed.
6. Select the target disk and specify the disk as the system disk or data disk
  - For KVM ECSs, you can specify a disk as a system disk or data disk but cannot specify a device name for the disk.
  - For Xen ECSs, you can specify the device name of a disk, such as **/dev/vdb**.

#### NOTE

- If no EVS disks are available, click **Create Disk** in the lower part of the list.
  - For details about restrictions on attaching a disk, see [What Are the Requirements for Attaching an EVS Disk to an ECS?](#)
7. Click **OK**.  
After the disk is attached, you can view the information about it on the **Disks** tab.

## Follow-up Procedure

If the attached disk is newly created, the disk can be used only after it is initialized.

For details about how to initialize a data disk, see [Scenarios and Disk Partitions](#).

## 5.3 Detaching an EVS Disk from a Running ECS

### Scenarios

You can detach EVS disks from an ECS.

- System disks (mounted to **/dev/sda** or **/dev/vda**) can only be detached offline. They must be stopped before being detached.
- Data disks (mounted to points other than **dev/sda**) can be detached online if the attached ECS is running certain OSs. You can detach these data disks without stopping the ECS.

This section describes how to detach a disk from a running ECS.

### Constraints

- The EVS disk to be detached must be mounted to a point other than **/dev/sda** or **/dev/vda**.  
EVS disks mounted to **/dev/sda** or **/dev/vda** are system disks and cannot be detached from running ECSs.
- Before detaching an EVS disk from a running Windows ECS, make sure that UVP VMTools have been installed on the ECS and that the tools are running properly.
- Before detaching an EVS disk from a running Windows ECS, ensure that no programs are reading data from or writing data to the disk. Otherwise, data will be lost.
- SCSI EVS disks cannot be detached from running Windows ECSs.
- Before detaching an EVS disk from a running Linux ECS, you must log in to the ECS and run the **umount** command to cancel the association between the disk and the file system. In addition, ensure that no programs are reading data from or writing data to the disk. Otherwise, detaching the disk will fail.

### Notes

- On a Windows ECS, if the disk is in non-offline state, the system forcibly detaches the EVS disk. If this occurs, the system may generate a xenvbd alarm. You can ignore this alarm.

 NOTE

To view the status of an EVS disk, perform the following operations:

1. Click **Start** in the task bar. In the displayed **Start** menu, right-click **Computer** and choose **Manage** from the shortcut menu.

The **Server Manager** page is displayed.

2. In the navigation pane on the left, choose **Storage > Disk Management**.

The EVS disk list is displayed in the right pane.

3. View the status of each EVS disk.

- Do not detach an EVS disk from an ECS that is being started, stopped, or restarted.
- Do not detach an EVS disk from a running ECS whose OS does not support this feature. OSs supporting EVS disk detachment from a running ECS are listed in [OSs Supporting EVS Disk Detachment from a Running ECS](#).
- For a running Linux ECS, the drive letter may be changed after an EVS disk is detached from it and then attached to it again. This is a normal case due to the drive letter allocation mechanism of the Linux system.
- For a running Linux ECS, the drive letter may be changed after an EVS disk is detached from it and the ECS is restarted. This is a normal case due to the drive letter allocation mechanism of the Linux system.

## OSs Supporting EVS Disk Detachment from a Running ECS

OSs supporting EVS disk detachment from a running ECS include two parts:

- For the first part, see **Overview > Concept > OSs for Public Images Supported by IMS** in *Image Management Service User Guide*.
- [Table 5-1](#) lists the second part of supported OSs.

**Table 5-1** OSs supporting EVS disk detachment from a running ECS

| OS     | Version                                   |
|--------|-------------------------------------------|
| CentOS | 7.3 64bit                                 |
|        | 7.2 64bit                                 |
|        | 6.8 64bit                                 |
|        | 6.7 64bit                                 |
| Debian | 8.6.0 64bit                               |
|        | 8.5.0 64bit                               |
| Fedora | 25 64bit                                  |
|        | 24 64bit                                  |
| SUSE   | SUSE Linux Enterprise Server 12 SP2 64bit |
|        | SUSE Linux Enterprise Server 12 SP1 64bit |
|        | SUSE Linux Enterprise Server 11 SP4 64bit |



| OS                                                              | Version                                 |
|-----------------------------------------------------------------|-----------------------------------------|
|                                                                 | SUSE Linux Enterprise Server 12 64bit   |
| OpenSUSE                                                        | 42.2 64bit                              |
|                                                                 | 42.1 64bit                              |
| Oracle Linux Server release                                     | 7.3 64bit                               |
|                                                                 | 7.2 64bit                               |
|                                                                 | 6.8 64bit                               |
|                                                                 | 6.7 64bit                               |
| Ubuntu Server                                                   | 16.04 64bit                             |
|                                                                 | 14.04 64bit                             |
|                                                                 | 14.04.4 64bit                           |
| Windows (SCSI EVS disks cannot be detached from a running ECS.) | Windows Server 2008 R2 Enterprise 64bit |
|                                                                 | Windows Server 2012 R2 Standard 64bit   |
|                                                                 | Windows Server 2016 R2 Standard 64bit   |
| Red Hat Linux Enterprise                                        | 7.3 64bit                               |
|                                                                 | 6.8 64bit                               |

 **NOTE**

Online detachment is not supported by the ECSs running OSs not listed in the preceding table. For such ECSs, stop the ECSs before detaching disks from them to prevent any possible problems from occurring.

## Procedure

1. On the **Elastic Cloud Server** page, click the name of the ECS from which the EVS disk is to be detached. The page providing details about the ECS is displayed.
2. Click the **Disks** tab. Locate the row containing the EVS disk to be detached and click **Detach**.

## 5.4 Expanding the Capacity of an EVS Disk

### Scenarios

You can expand the disk capacity if the disk space is insufficient. The capacities of both system disks and data disks can be expanded.

## Procedure

The capacity of an EVS disk can be expanded in either of the following ways:

- Create an EVS disk and attach it to an ECS.
- Expand the capacity of an existing EVS disk. The capacities of both system disks and data disks can be expanded.

You can expand the disk capacities when the EVS disks are in the **In-use** or **Available** state.

- Expanding an **In-use** EVS disk means expanding the capacity of an EVS disk that has been attached to an ECS. Only certain OSs support the expansion of **In-use** EVS disks. For details, see "Expanding an In-use EVS Disk" in *Elastic Volume Service User Guide*.
- Expanding an **Available** EVS disk means expanding the capacity of an EVS disk that has not been attached to any ECS. For details, see "Expanding an Available EVS Disk" in *Elastic Volume Service User Guide*.

### NOTE

After the disk capacity is expanded, only the storage capacity of the EVS disk is expanded. To use the added storage space, you also need to log in to the ECS and extend the partition and file system.

## 5.5 Enabling Advanced Disk

### Scenarios

- Disk functions have been upgraded on the platform. Newly created ECSs can have up to 60 attached disks. However, an existing ECS can still have a maximum of 24 attached disks (40 for certain ECSs). To allow such ECSs to have up to 60 attached disks, enable advanced disk.
- After advanced disk is enabled, you can view the mapping between device names and disks. For details, see [How Do I Obtain My Disk Device Name in the ECS OS Using the Device Identifier Provided on the Console?](#)

This section describes how to enable advanced disk on an ECS.

### Procedure

1. Log in to management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Click the name of the target ECS. The page providing details about the ECS is displayed.
4. Click the **Disks** tab.
5. View the current number of disks that can be attached to the ECS and enable advanced disk as prompted.

The **Enable Advanced Disk** dialog box is displayed.

6. Click **OK**.
7. Stop and then start the target ECS.

This operation allows advanced disk to take effect.

8. Switch to the page providing details about the ECS again, click the **Disks** tab, and check whether the number of disks that can be attached to the ECS has been changed.
  - If yes, advanced disk has been enabled.
  - If no, enabling advanced disk failed. In such a case, contact the administrator.

# 6 CBR

---

## 6.1 Overview

### What Is CBR?

Cloud Backup and Recovery (CBR) enables you to back up cloud servers and disks with ease. In case of a virus attack, accidental deletion, or software or hardware fault, you can restore data to any point in the past when the data was backed up.

CBR protects your services by ensuring the security and consistency of your data.

### What Are the Differences Between Backup and Image?

You can use the cloud server backup function to create ECSs and the cloud disk backup function to create EVS disks.

An image can be a system disk image, data disk image, or full-ECS image.

| Backup Type         | Backup Object                                      | Application Scenario                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Differences and Advantages                                                                                                                                                                |
|---------------------|----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cloud server backup | All disks (system and data disks) on an ECS        | <ul style="list-style-type: none"> <li>● <b>Hacker attacks and viruses</b><br/>You can use cloud server backup to restore data to the latest backup point at which the ECS has not been affected by hacker attacks and viruses.</li> <li>● <b>Accidental data deletion</b><br/>You can use cloud server backup to restore data to the backup point prior to the accidental deletion.</li> <li>● <b>Application update errors</b><br/>You can use cloud server backup to restore data to the backup point prior to the application update.</li> <li>● <b>System breakdown</b><br/>You can use cloud server backup to restore an ECS to the backup point in time prior to system breakdown.</li> </ul> | <p>All disks on an ECS are backed up at the same time, ensuring data consistency.</p> <p>In addition, you can configure backup policies for automatic backup.</p>                         |
| Cloud disk backup   | One or more specified disks (system or data disks) | <ul style="list-style-type: none"> <li>● <b>Only data disks need to be backed up, because the system disk does not contain users' application data.</b><br/>You can use cloud disk backup to back up and restore data if an EVS disk is faulty or encounters a logical error, for example, accidental deletion, hacker attacks, and virus infection.</li> <li>● <b>Use backups as baseline data.</b><br/>After a backup policy has been set, the EVS disk data can be automatically backed up based on the policy. You can use the backups created on a timely basis as the baseline data to create new EVS disks or to restore the backup data to EVS disks.</li> </ul>                             | <p>Backup data is stored in OBS, instead of disks. This ensures data restoration upon disk data loss or corruption.</p> <p>Backup cost is reduced without compromising data security.</p> |

| Backup Type       | Backup Object                               | Application Scenario                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Differences and Advantages                                                                                                                                 |
|-------------------|---------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| System disk image | System disk                                 | <ul style="list-style-type: none"><li>• <b>Rapid system recovery</b><br/>You can create a system disk image for the system disk of an ECS before OS change, application software upgrade, or service data migration. If an exception occurs during the migration, you can use the system disk image to change ECS OS or create a new ECS.</li><li>• <b>Rapid deployment of multiple services</b><br/>You can use a system disk image to quickly create multiple ECSs with the same OS, thereby quickly deploying services these ECSs.</li></ul>                | A system disk image can help an ECS with OS damaged to quickly change its OS.                                                                              |
| Data disk image   | Specific data disk                          | <b>Rapid data replication</b><br>You can use a data disk image to create multiple EVS disks containing the same initial data, and then attach these disks to ECSs to provide data resources for multiple services.                                                                                                                                                                                                                                                                                                                                             | A data disk image can replicate all data on a disk and create new EVS disks. The EVS disks can be attached to other ECSs for data replication and sharing. |
| Full-ECS image    | All disks (system and data disks) on an ECS | <ul style="list-style-type: none"><li>• <b>Rapid system recovery</b><br/>You can create a full-ECS image for the system disk and data disks of an ECS before OS change, application software upgrade, or service data migration. If an exception occurs during the migration, you can use the full-ECS image to change ECS OS or create a new ECS.</li><li>• <b>Rapid deployment of multiple services</b><br/>You can use a full-ECS image to quickly create multiple ECSs with the same OS and data, thereby quickly deploying services these ECSs.</li></ul> | A full-ECS image facilitates service migration.                                                                                                            |

## CBR Architecture

CBR consists of backups, vaults, and policies.

- **Backup**

A backup is a copy of a particular chunk of data and is usually stored elsewhere so that it may be used to restore the original data in the event of data loss. CBR supports the following backup types:

- Cloud server backup: This type of backup uses the consistency snapshot technology for disks to protect data of ECSs and BMSs. The backups of servers without deployed databases are common server backups, and those of servers with deployed databases are application-consistent backups.
- Cloud disk backup: This type of backup provides snapshot-based data protection for EVS disks.

- **Vault**

CBR uses vaults to store backups. Before creating a backup, you need to create at least one vault and associate the resource you want to back up with the vault. Then the backup of the resource is stored in the associated vault.

Vaults can be classified into two types: backup vaults and replication vaults. Backup vaults store backups, whereas replication vaults store replicas of backups.

The backups of different types of resources must be stored in different types of vaults.

- **Policy**

Policies are divided into backup policies and replication policies.

- Backup policies: To perform automatic backups, configure a backup policy by setting the execution times of backup tasks, the backup cycle, and retention rules, and then apply the policy to a vault.
- Replication policies: To automatically replicate backups or vaults, configure a replication policy by setting the execution times of replication tasks, the replication cycle, and retention rules, and then apply the policy to a vault. Replicas of backups must be stored in replication vaults.

## Backup Mechanism

A full backup is performed only for the first backup and backs up all used data blocks.

For example, if the size of a disk is 100 GB and the used space is 40 GB, the 40 GB of data is backed up.

An incremental backup backs up only the data changed since the last backup, which is storage- and time-efficient.

When a backup is deleted, only the data blocks that are not depended on by other backups are deleted, so that other backups can still be used for restoration. Both a full backup and an incremental backup can restore data to the state at a given backup point in time.

When creating a backup of a disk, CBR also creates a snapshot for it. Every time a new disk backup is created, CBR deletes the old snapshot and keeps only the latest snapshot.

CBR stores backup data in OBS, enhancing backup data security.

## Backup Options

CBR supports one-off backup and periodic backup. A one-off backup task is manually created by users and is executed only once. Periodic backup tasks are automatically executed based on a user-defined backup policy.

**Table 6-1** One-off backup and periodic backup

| Item                   | One-Off Backup                                                                                                                                                                                         | Periodic Backup                                                                                                                             |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| Backup policy          | Not required                                                                                                                                                                                           | Required                                                                                                                                    |
| Number of backup tasks | One manual backup task                                                                                                                                                                                 | Periodic tasks driven by a backup policy                                                                                                    |
| Backup name            | User-defined backup name, which is <b>manualbk_XXXX</b> by default                                                                                                                                     | System-assigned backup name, which is <b>autobk_XXXX</b> by default                                                                         |
| Backup mode            | Full backup for the first time and incremental backup subsequently, by default                                                                                                                         | Full backup for the first time and incremental backup subsequently, by default                                                              |
| Application scenario   | Executed before patching or upgrading the OS or upgrading an application on a resource. A one-off backup can be used to restore the resource to the original state if the patching or upgrading fails. | Executed for routine maintenance of a resource. The latest backup can be used for restoration if an unexpected failure or data loss occurs. |

## 6.2 Backing Up an ECS

### Scenarios

CBR enhances data integrity and service continuity. For example, if an ECS or EVS disk is faulty or a misoperation causes data loss, you can use data backups to quickly restore data. This section describes how to back up ECSs and EVS disks.

For more information, [CBR Architecture](#), [Backup Mechanism](#), and [Backup Options](#).

You can back up ECS data using Cloud Server Backup or Cloud Disk Backup.



- **Cloud Server Backup (recommended):** Use this backup function if you want to back up the data of all EVS disks (system and data disks) on an ECS. This prevents data inconsistency caused by time difference in creating a backup.
- **Cloud Disk Backup:** Use this backup function if you want to back up the data of one or more EVS disks (system or data disk) on an ECS. This minimizes backup costs on the basis of data security.

## ECS Backup Procedure


1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. In the ECS list, locate the target ECS and choose **More > Manage Image/Disk/Backup > Create Backup**.
  - If the ECS has been associated with a vault, configure the backup information as prompted.
    - **Server List:** The ECS to be backed up is selected by default.
    - **Name:** Customize your backup name.
    - **Description:** Supplementary information about the backup.
    - **Full Backup:** If this option is selected, the system will perform full backup for the ECS to be associated. The storage capacity used by the backup increases accordingly.
  - If the ECS is not associated with a vault, buy a vault first and then configure the backup information as prompted.

For details, see *Cloud Backup and Recovery User Guide*.
4. Click **OK**. The system automatically creates a backup for the ECS.

On the **Backups** tab page, if the status of the backup is **Available**, the backup task is successful.

The ECS can be restarted if the backup progress of an ECS exceeds 10%. However, to ensure data integrity, restart it after the backup is complete.

## EVS Disk Backup Procedure

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. In the ECS list, locate the target ECS and choose **More > Manage Image/Disk > Create Backup**.
  - If the ECS has been associated with a vault, configure the backup information as prompted.
    - **Server List:** The ECS to be backed up is selected by default. Click  to view the disks attached to the ECSs. Select the disks to be backed up.
    - **Name:** Customize your backup name.
    - **Description:** Supplementary information about the backup.

- **Full Backup:** If this option is selected, the system will perform full backup for the disks to be associated. The storage capacity used by the backup increases accordingly.
- If the ECS is not associated with a vault, buy a vault first and then configure the backup information as prompted.

For details, see *Cloud Backup and Recovery User Guide*.

4. Click **OK**. The system automatically creates a backup for the disk.

On the **Backups** tab of the CBR console, if the status of the backup is **Available**, the backup task is successful.

If some files are deleted from the disk during the backup, the deleted files may fail to be backed up. Therefore, to ensure data integrity, delete the target data after the backup is complete.

# 7 NICs

---

## 7.1 Attaching a Network Interface

### Scenarios

If your ECS requires multiple network interfaces, you can attach them to your ECS.

### Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Click the name of the target ECS.  
The page providing details about the ECS is displayed.
4. On the **Network Interfaces** tab, click **Attach Network Interface**.
5. Select either of the following methods to attach the network interface.
  - Use an existing network interface.
    - i. (Optional) Search for the network interface by name, ID, or private IP address.
    - ii. In the network interface list, select the target one.
  - Create a new network interface.  
Set the subnet and security group for the network interface to be attached.
    - **Subnet:** specifies the subnet which the network interface belongs to.
    - **Private IP Address:** If you want to add a network interface with a specified IP address, enter an IP address into the **Private IP Address** field.
    - **Security Group:** You can select multiple security groups. In such a case, the access rules of all the selected security groups will apply to the ECS.
6. Click **OK**.

## Follow-up Procedure

Some OSs cannot identify newly added network interfaces. In this case, you must manually activate the network interfaces. Ubuntu is used as an example in the following network interface activation procedure. Required operations may vary among systems. For additional information, see the documentation for your OS.

1. Locate the row containing the target ECS and click **Remote Login** in the **Operation** column.  
Log in to the ECS.
2. Run the following command to view the network interface name:  
**ifconfig -a**  
In this example, the network interface name is **eth2**.
3. Run the following command to switch to the target directory:  
**cd /etc/network**
4. Run the following command to open the **interfaces** file:  
**vi interfaces**
5. Add the following information to the **interfaces** file:  
**auto eth2**  
**iface eth2 inet dhcp**
6. Run the following command to save and exit the **interfaces** file:  
**:wq**
7. Run either the **ifup eth2** command or the **/etc/init.d/networking restart** command to make the newly added network interface take effect.  
*X* in the preceding command indicates the network interface name and SN, for example, **ifup eth2**.
8. Run the following command to check whether the network interface name obtained in step 2 is displayed in the command output:  
**ifconfig**  
For example, check whether **eth2** is displayed in the command output.
  - If yes, the newly added network interface has been activated, and no further action is required.
  - If no, the newly added network interface failed to be activated. Go to step 9.
9. Log in to the management console. Locate the row containing the target ECS, click **More** in the **Operation** column, and select **Restart**.
10. Run the following command to check whether the network interface name obtained in step 2 is displayed in the command output:
  - If yes, no further action is required.
  - If no, contact the administrator.

## 7.2 Detaching a Network Interface

### Scenarios

An ECS can have up to 12 network interfaces, including one primary network interface that cannot be deleted and extension network interfaces. This section describes how to detach a network interface.

### Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. On the **Elastic Cloud Server** page, click the name of the target ECS.  
The page providing details about the ECS is displayed.
4. On the **Network Interfaces** tab, locate the target network interface and click **Detach**.

#### NOTE

You are not allowed to delete the primary ECS network interface. By default, the primary ECS network interface is the first network interface displayed in the network interface list.

5. In the displayed dialog box, click Yes.

#### NOTE

Certain ECSs do not support network interface detachment when they are running. For details, see the GUI display. To detach a network interface from such an ECS, stop the ECS first.

## 7.3 Modifying a Private IP Address

### Scenarios

You can modify the private IP address of the primary NIC. If you want to modify the private IP address of an extension NIC, delete the NIC and attach a new NIC.

### Constraints

- The ECS must be stopped.
- If a virtual IP address or DNAT rule has been configured for the NIC, cancel the configuration before modifying the private IP address.
- If the NIC has an IPv6 address, its private IP address (IPv4 or IPv6 address) cannot be modified.
- Before changing the private IP address of an ELB backend server, delete the backend server group.

## Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Click the name of the target ECS.  
The ECS details page is displayed.
4. Click the **Network Interfaces** tab. Locate the row containing the primary network interface and click **Modify Private IP**.  
The **Modify Private IP** dialog box is displayed.
5. Change the subnet and private IP address of the primary NIC as required.

### NOTE

Subnets can be changed only within the same VPC.

If the target private IP address is not specified, the system will automatically assign one to the primary NIC.

## 7.4 Managing Virtual IP Addresses

### Scenarios

A virtual IP address provides the second IP address for one or more ECS NICs, improving high availability between the ECSs.

### Binding a Virtual IP Address

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. On the **Elastic Cloud Server** page, click the name of the target ECS.  
The page providing details about the ECS is displayed.
4. On the **Network Interfaces** tab, locate the target virtual IP address and click **Manage Virtual IP Address**.
5. On the **IP Addresses** tab of the displayed page, locate the row containing the target virtual IP address and select **Bind to EIP** or **Bind to Server** in the **Operation** column.  
Multiple ECSs deployed to work in active/standby mode can be bound with a virtual IP address to improve DR performance.
6. Click **OK**.

## 7.5 Enabling NIC Multi-Queue

### Scenarios

Single-core CPU performance cannot meet the requirement of processing NIC interruptions incurred with the increase of network I/O bandwidth. NIC multi-queue enables multiple CPUs to process ECS NIC interruptions, thereby improving packets per second (PPS) and I/O performance.

The ECS described in this section is assumed to comply with the requirements on specifications and virtualization type.

- If the ECS was created using a public image listed in [Support of NIC Multi-Queue](#), NIC multi-queue has been enabled on the ECS by default. Therefore, you do not need to perform the operations described in this section.
- If the ECS was created using a private image and the OS of the external image file is listed in [Support of NIC Multi-Queue](#), perform the following operations to enable NIC multi-queue:
  - a. [Importing the External Image File to the IMS Console](#)
  - b. [Setting NIC Multi-Queue for the Image](#)
  - c. [Creating an ECS Using a Private Image](#)
  - d. [Enabling NIC Multi-Queue](#)

## Support of NIC Multi-Queue

NIC multi-queue can be enabled on an ECS only when the ECS specifications, virtualization type, and image OS meet the requirements described in this section.

- For details about the ECS specifications that support NIC multi-queue, see [x86 ECS Specifications and Types](#).

### NOTE

If the number of NIC queues is greater than 1, NIC multi-queue is supported.

- The virtualization type must be KVM.
- The Linux public images listed in [Table 7-2](#) support NIC multi-queue.

### NOTE

- Windows public images have not supported NIC multi-queue. If you enable NIC multi-queue in a Windows public image, starting an ECS created using such an image may be slow.
- It is a good practice to upgrade the kernel version of the Linux ECS to 2.6.35 or later. Otherwise, NIC multi-queue is not supported.

Run the `uname -r` command to obtain the kernel version. If the kernel version is earlier than 2.6.35, contact technical support to upgrade the kernel.

**Table 7-1** Support of NIC multi-queue for Linux ECSs

| Image                                              | Support of NIC Multi-Queue | NIC Multi-Queue Enabled by Default |
|----------------------------------------------------|----------------------------|------------------------------------|
| Ubuntu 14.04/16.04/18.04/20.04 server 64bit        | Yes                        | Yes                                |
| OpenSUSE 42.2/15.* 64bit                           | Yes                        | Yes                                |
| SUSE Enterprise 12 SP1/SP2 64bit                   | Yes                        | Yes                                |
| CentOS 6.8/6.9/7.*/8.* 64bit                       | Yes                        | Yes                                |
| Debian 8.0.0/8.8.0/8.9.0/9.0.0/10.0.0/10.2.0 64bit | Yes                        | Yes                                |

| Image                     | Support of NIC Multi-Queue | NIC Multi-Queue Enabled by Default |
|---------------------------|----------------------------|------------------------------------|
| Fedora 24/25/30 64bit     | Yes                        | Yes                                |
| EulerOS 2.2/2.3/2.5 64bit | Yes                        | Yes                                |

**Table 7-2** Support of NIC multi-queue for KVM ECSs

| OS      | Image                                                       | Status                         |
|---------|-------------------------------------------------------------|--------------------------------|
| Windows | Windows Server 2008 Web R2 64-bit                           | Supported using private images |
|         | Windows Server 2008 R2 Standard/DataCenter/Enterprise 64bit | Supported using private images |
|         | Windows Server 2012 R2 Standard/DataCenter 64bit            | Supported using private images |
|         | Windows Server 2016 Standard/DataCenter 64bit               | Supported using private images |
| Linux   | Ubuntu 14.04/16.04 server 64bit                             | Supported                      |
|         | OpenSUSE 42.2 64bit                                         | Supported                      |
|         | SUSE Enterprise 12 SP1/SP2 64bit                            | Supported                      |
|         | CentOS 6.8/6.9/7.0/7.1/7.2/7.3/7.4/7.5/7.6 64bit            | Supported                      |
|         | Debian 8.0.0/8.8.0/8.9.0/9.0.0 64bit                        | Supported                      |
|         | Fedora 24/25 64bit                                          | Supported                      |
|         | EulerOS 2.2 64bit                                           | Supported                      |

## Importing the External Image File to the IMS Console

For details, see "Registering an Image File as a Private Image" in *Image Management Service User Guide*.

## Setting NIC Multi-Queue for the Image

Windows OSs have not commercially supported NIC multi-queue. If you enable NIC multi-queue in a Windows image, starting an ECS created using such an image may be slow.

Use one of the following methods to set the NIC multi-queue attribute:

### Method 1:



1. Log in to the management console.
2. Under **Computing**, click **Image Management Service**.
3. Click the **Private Images** tab, locate the row containing the target image, click **Modify** in the **Operation** column.
4. Set the NIC multi-queue attribute of the image.

**Method 2:**

1. Log in to the management console.
2. Under **Computing**, click **Image Management Service**.
3. Click the **Private Images** tab. In the image list, click the name of the target image to switch to the page providing details about the image.
4. Click **Modify** in the upper right corner. In the displayed **Modify Image** dialog box, set the NIC multi-queue attribute.

**Method 3: Add `hw_vif_multiqueue_enabled` to an image through the API.**

1. For instructions about how to obtain the token, see **Calling APIs > Authentication** in *Image Management Service API Reference*.
2. For instructions about how to call an API to update image information, see "Updating Image Information (Native OpenStack API)" in *Image Management Service API Reference*.
3. Add **X-Auth-Token** to the request header.  
The value of **X-Auth-Token** is the token obtained in step 1.
4. Add **Content-Type** to the request header.

The value of **Content-Type** is **application/openstack-images-v2.1-json-patch**.

The request URI is in the following format:


PATCH /v2/images/{image\_id}

The request body is as follows:

```
[
 {
 "op": "add",
 "path": "/hw_vif_multiqueue_enabled",
 "value": "true"
 }
]
```

**Figure 7-1** shows an example request body for modifying the NIC multi-queue attribute.

**Figure 7-1** Example request body



## Creating an ECS Using a Private Image

Create an ECS using a registered private image. For details, see [Creating an ECS](#). Note the following when setting the parameters:

- **Region:** Select the region where the private image is located.
- **Image:** Select **Private image** and then the desired image from the drop-down list.

## Enabling NIC Multi-Queue

KVM Windows ECSs use private images to support NIC multi-queue. For details, see "How Do I Set NIC Multi-queue Feature of an Image?" in *Image Management Service User Guide*.

This section uses a Linux ECS running CentOS 7.4 as an example to describe how to enable NIC multi-queue.

### Step 1 Enable NIC multi-queue.

1. Log in to the ECS.
2. Run the following command to obtain the number of queues supported by the NIC and the number of queues with NIC multi-queue enabled:

```
ethtool -l NIC
```

3. Run the following command to configure the number of queues used by the NIC:

```
ethtool -L NIC combined Number of queues
```

An example is provided as follows:

```
[root@localhost ~]# ethtool -l eth0 #View the number of queues used by NIC eth0.
Channel parameters for eth0:
Pre-set maximums:
RX: 0
TX: 0
Other: 0
Combined: 4 #Indicates that a maximum of four queues can be enabled for the NIC.
Current hardware settings:
RX: 0
TX: 0
Other: 0
Combined: 1 #Indicates that one queue has been enabled.

[root@localhost ~]# ethtool -L eth0 combined 4 #Enable four queues on NIC eth0.
```

### Step 2 (Optional) Enable irqbalance so that the system automatically allocates NIC interrupts on multiple vCPUs.

1. Run the following command to enable irqbalance:  
**service irqbalance start**
2. Run the following command to view the irqbalance status:  
**service irqbalance status**

If the **Active** value in the command output contains **active (running)**, irqbalance has been enabled.

**Figure 7-2** Enabled irqbalance

```
root@localhost ~# service irqbalance status
Redirecting to /bin/systemctl status irqbalance.service
● irqbalance.service - irqbalance daemon
 Loaded: loaded (/usr/lib/systemd/system/irqbalance.service; enabled; vendor preset: enabled)
 Active: active (running) since Wed 2018-08-15 10:27:30 CST; 4h 5min ago
 Main PID: 858 (irqbalance)
 CGroup: /system.slice/irqbalance.service
 └─858 /usr/sbin/irqbalance --foreground

Aug 15 10:27:30 localhost.localdomain systemd[1]: Started irqbalance daemon.
Aug 15 10:27:30 localhost.localdomain systemd[1]: Starting irqbalance daemon...
```

**Step 3** (Optional) Enable interrupt binding.

Enabling irqbalance allows the system to automatically allocate NIC interrupts, improving network performance. If the improved network performance still fails to meet your requirements, manually configure interrupt affinity on the ECS.

To do so, perform the following operations:

Configure the following script so that one ECS vCPU serves the interrupt requests initialized by one queue. One queue corresponds to one interrupt, and one interrupt binds to one vCPU.

```
#!/bin/bash
service irqbalance stop

eth_dirs=$(ls -d /sys/class/net/eth*)
if [$? -ne 0];then
 echo "Failed to find eth* , sleep 30" >> $ecs_network_log
 sleep 30
 eth_dirs=$(ls -d /sys/class/net/eth*)
fi

for eth in $eth_dirs
do
 cur_eth=$(basename $eth)
 cpu_count=`cat /proc/cpuinfo| grep "processor"| wc -l`
 virtio_name=$(ls -l /sys/class/net/"$cur_eth"/device/driver/ | grep pci |awk '{print $9}')

 affinity_cpu=0
 virtio_input="$virtio_name"-input"
 irqs_in=$(grep "$virtio_input" /proc/interrupts | awk -F ":" '{print $1}')
 for irq in ${irqs_in[*]}
 do
 echo $((affinity_cpu%cpu_count)) > /proc/irq/"$irq"/smp_affinity_list
 affinity_cpu=$((affinity_cpu+2))
 done

 affinity_cpu=1
 virtio_output="$virtio_name"-output"
 irqs_out=$(grep "$virtio_output" /proc/interrupts | awk -F ":" '{print $1}')
 for irq in ${irqs_out[*]}
 do
 echo $((affinity_cpu%cpu_count)) > /proc/irq/"$irq"/smp_affinity_list
 affinity_cpu=$((affinity_cpu+2))
 done
done
```

**Step 4** (Optional) Enable XPS and RPS.

XPS allows the system with NIC multi-queue enabled to select a queue by vCPU when sending a data packet.

```
#!/bin/bash
enable XPS feature
cpu_count=$(grep -c processor /proc/cpuinfo)
dec2hex(){
 echo $(printf "%x" $1)
}
```

```
eth_dirs=$(ls -d /sys/class/net/eth*)
if [$? -ne 0];then
 echo "Failed to find eth* , sleep 30" >> $ecs_network_log
 sleep 30
 eth_dirs=$(ls -d /sys/class/net/eth*)
fi
for eth in $eth_dirs
do
 cpu_id=1
 cur_eth=$(basename $eth)
 cur_q_num=$(ethtool -l $cur_eth | grep -iA5 current | grep -i combined | awk {'print $2'})
 for((i=0;i<cur_q_num;i++))
 do
 if [$i -eq $cpu_count];then
 cpu_id=1
 fi
 xps_file="/sys/class/net/${cur_eth}/queues/tx-$i/xps_cpus"
 rps_file="/sys/class/net/${cur_eth}/queues/rx-$i/rps_cpus"
 cpuset=$(dec2hex "$cpu_id")
 echo $cpuset > $xps_file
 echo $cpuset > $rps_file
 let cpu_id=cpu_id*2
 done
done
```

----End

## 7.6 Dynamically Assigning IPv6 Addresses

### Scenarios

IPv6 addresses are used to deal with IPv4 address exhaustion. If an ECS uses an IPv4 address, the ECS can run in dual-stack mode after IPv6 is enabled for it. Then, the ECS will have two IP addresses to access the intranet and Internet: an IPv4 address and an IPv6 address.

In some cases, an ECS cannot dynamically acquire an IPv6 address even if it meets all the requirements in [Constraints](#). You need to configure the ECS to dynamically acquire IPv6 addresses. For public images:

- By default, dynamic IPv6 address assignment is enabled for Windows public images. You do not need to configure it. The operations in [Windows Server 2012](#) and [Windows Server 2008](#) are for your reference only.
- Before enabling dynamic IPv6 address assignment for a Linux public image, check whether IPv6 has been enabled and then whether dynamic IPv6 address assignment has been enabled. Currently, IPv6 is enabled for all Linux public images.

### Constraints

- Ensure that IPv6 has been enabled on the subnet where the ECS works. If IPv6 is not enabled on the subnet, enable it by referring to [Enabling IPv6 for an ECS](#). IPv6 cannot be disabled once it is enabled.
- Ensure that **Self-assigned IPv6 address** is selected during ECS creation.
- After the ECS is started, its hot-swappable NICs cannot automatically acquire IPv6 addresses.
- Only ECSs can work in dual-stack mode and BMSs cannot.

- Only one IPv6 address can be bound to a NIC.

## Procedure

- Windows: Windows Server 2012/2008 is used as an example to describe how to enable dynamic assignment of IPv6 addresses in Windows.
- Linux: Dynamic assignment of IPv6 addresses can be enabled automatically (recommended) or manually.

If a private image created from a CentOS 6.x or Debian ECS with automatic IPv6 address assignment enabled is used to create an ECS in an environment that does not support IPv6, the ECS may start slow because of IPv6 address assignment timeout. You can set the timeout duration for assigning IPv6 addresses by referring to [Setting the Timeout Duration for IPv6 Address Assignment](#).

**Table 7-3** Enabling dynamic assignment of IPv6 addresses for different OSs

| OS                  | Automatically/<br>Manually Enabling | Reference                                                                           |
|---------------------|-------------------------------------|-------------------------------------------------------------------------------------|
| Windows Server 2012 | Automatically                       | <a href="#">Windows Server 2012</a>                                                 |
| Windows Server 2008 | Automatically                       | <a href="#">Windows Server 2008</a>                                                 |
| Linux               | Automatically<br>(recommended)      | <a href="#">Linux (Automatically Enabling Dynamic Assignment of IPv6 Addresses)</a> |
| Linux               | Manually                            | <a href="#">Linux (Manually Enabling Dynamic Assignment of IPv6 Addresses)</a>      |

## Enabling IPv6 for an ECS

Enabling IPv6 on a Subnet

### NOTE

After IPv6 is enabled on the subnet where the ECS works, an IPv6 CIDR block is automatically assigned to the subnet. IPv6 cannot be disabled once it is enabled.

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Click the target ECS to go to the detail page.
4. In the **ECS Information** area, click the VPC name.
5. Click the number in the **Subnets** column.  
The **Subnets** page is displayed.
6. In the subnet list, locate the target subnet and click its name.

The subnet details page is displayed.

7. In the **Subnet Information** area, click **Enable** for **IPv6 CIDR Block**.
8. Click **Yes**.

Enabling IPv6 on a Network Interface

1. Go back to the ECS details page.
2. On the **Network Interfaces** tab, click **Enable IPv6** in the upper right corner of the folded panel.

#### NOTE

You can disable IPv6 on this page if it is no longer used. After IPv6 is disabled, the IPv6 address of the network interface is empty.

If you want to enable IPv6 again after it is disabled, you need to restart the ECS, log in to the ECS and manually clear the IPv6 cache, and request an IPv6 address again.

3. Click **OK**.

## Windows Server 2012

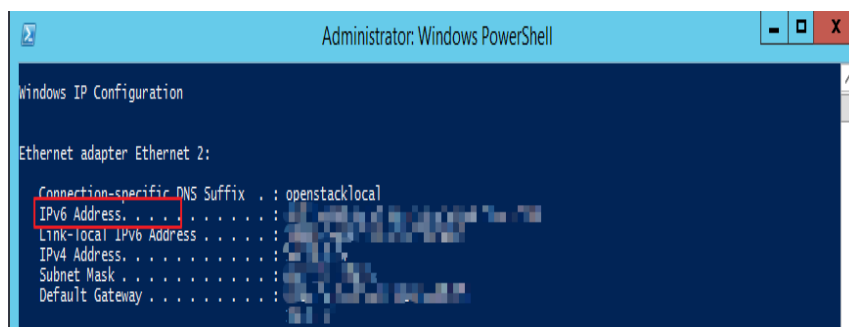
**Step 1** Check whether IPv6 is enabled for the ECS.

Run the following command in the CMD window to check it:

### ipconfig

- If an IPv6 address and a link-local IPv6 address are displayed, IPv6 is enabled and dynamic IPv6 assignment is also enabled.

**Figure 7-3** Querying the IPv6 address



```
Administrator: Windows PowerShell

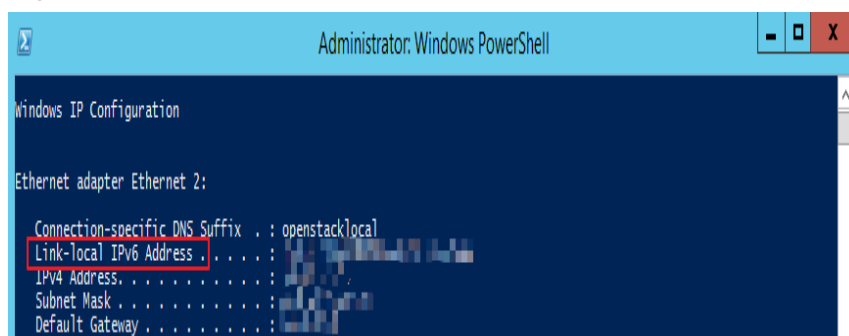
Windows IP Configuration

Ethernet adapter Ethernet 2:

 Connection-specific DNS Suffix . : openstacklocal
 IPv6 Address. :
 Link-local IPv6 Address :
 IPv4 Address. :
 Subnet Mask :
 Default Gateway :
```

- If only a link-local IPv6 address is displayed, IPv6 is enabled but dynamic IPv6 assignment is not enabled. Go to [Step 2](#).

**Figure 7-4** Link-local IPv6 address



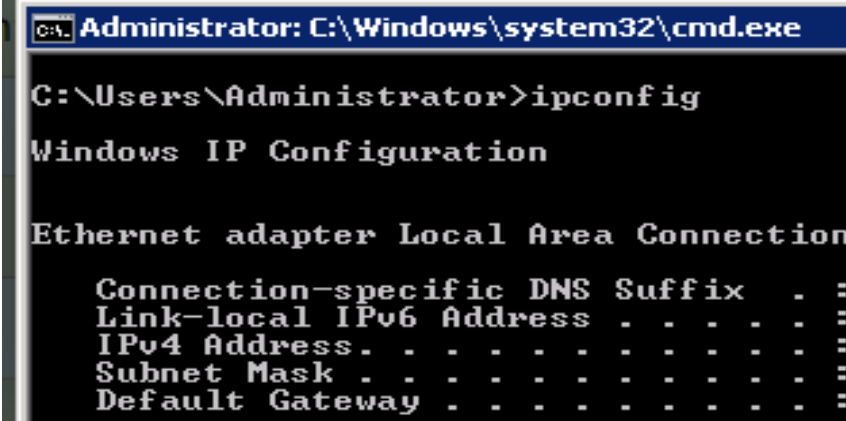
```
Administrator: Windows PowerShell

Windows IP Configuration

Ethernet adapter Ethernet 2:

 Connection-specific DNS Suffix . : openstacklocal
 Link-local IPv6 Address :
 IPv4 Address. :
 Subnet Mask :
 Default Gateway :
```

- If neither an IPv6 address nor link-local IPv6 address is displayed, IPv6 is disabled. Go to [Step 3](#).

**Figure 7-5** IPv6 disabled

```
C:\Users\Administrator>ipconfig

Windows IP Configuration

Ethernet adapter Local Area Connection:


 Connection-specific DNS Suffix . :
 Link-local IPv6 Address :
 IPv4 Address. :
 Subnet Mask :
 Default Gateway :
```

**NOTE**

By default, dynamic IPv6 address assignment is enabled for Windows public images, as shown in [Figure 7-3](#). No additional configuration is required.

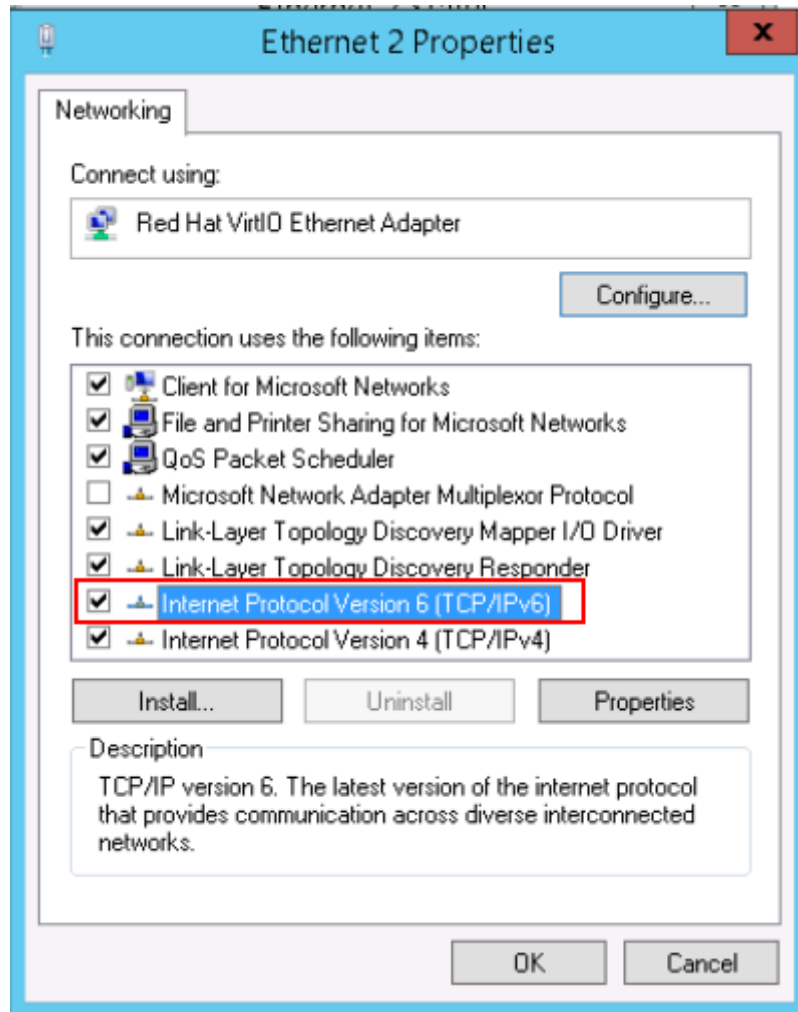
**Step 2** Enable dynamic IPv6 address assignment.

1. Choose **Start > Control Panel**.
2. Click **Network and Sharing Center**.
3. Click the Ethernet connection.

**Figure 7-6** Ethernet connection

4. In the **Ethernet Status** dialog box, click **Properties** in the lower left corner.
5. Select **Internet Protocol Version 6 (TCP/IPv6)** and click **OK**.

**Figure 7-7** Configuring dynamic IPv6 address assignment




6. Perform [Step 1](#) to check whether dynamic IPv6 address assignment is enabled.

**Step 3** Enable and configure IPv6.

1. In the **Internet Protocol Version 6 (TCP/IPv6) Properties** dialog box, configure an IPv6 address and a DNS server address.
  - **IPv6 address:** IPv6 address allocated during ECS creation. Obtain the value from the ECS list on the console.
  - **Subnet prefix length: 64**
  - **Preferred DNS server: 240c::6666** (recommended)



**Figure 7-8** Configuring an IPv6 address and a DNS server address

2. (Optional) Run the following command depending on your ECS OS.  
For Windows Server 2012, run the following command in PowerShell or CMD:  
**Set-NetIPv6Protocol -RandomizeIdentifiers disabled**
3. Perform [Step 1](#) to check whether dynamic IPv6 address assignment is enabled.

----End

## Windows Server 2008


**Step 1** Check whether IPv6 is enabled for the ECS.

Run the following command in the CMD window to check it:

### ipconfig


- If an IPv6 address and a link-local IPv6 address are displayed, IPv6 is enabled and dynamic IPv6 assignment is also enabled.

**Figure 7-9** Querying the IPv6 address




- If only a link-local IPv6 address is displayed, IPv6 is enabled but dynamic IPv6 assignment is not enabled. Go to [Step 2](#).

**Figure 7-10** Link-local IPv6 address



- If neither an IPv6 address nor link-local IPv6 address is displayed, IPv6 is disabled. Go to [Step 3](#).

**Figure 7-11** IPv6 disabled



**NOTE**


By default, dynamic IPv6 address assignment is enabled for Windows public images, as shown in [Figure 7-9](#). No additional configuration is required.

**Step 2** Enable dynamic IPv6 address assignment.

1. Choose **Start > Control Panel**.

2. Click **Network and Sharing Center**.
3. Click **Change adapter settings**.
4. Right-click the local network connection and choose **Properties**.
5. Select **Internet Protocol Version 6 (TCP/IPv6)** and click **OK**.

**Figure 7-12** Configuring dynamic IPv6 address assignment

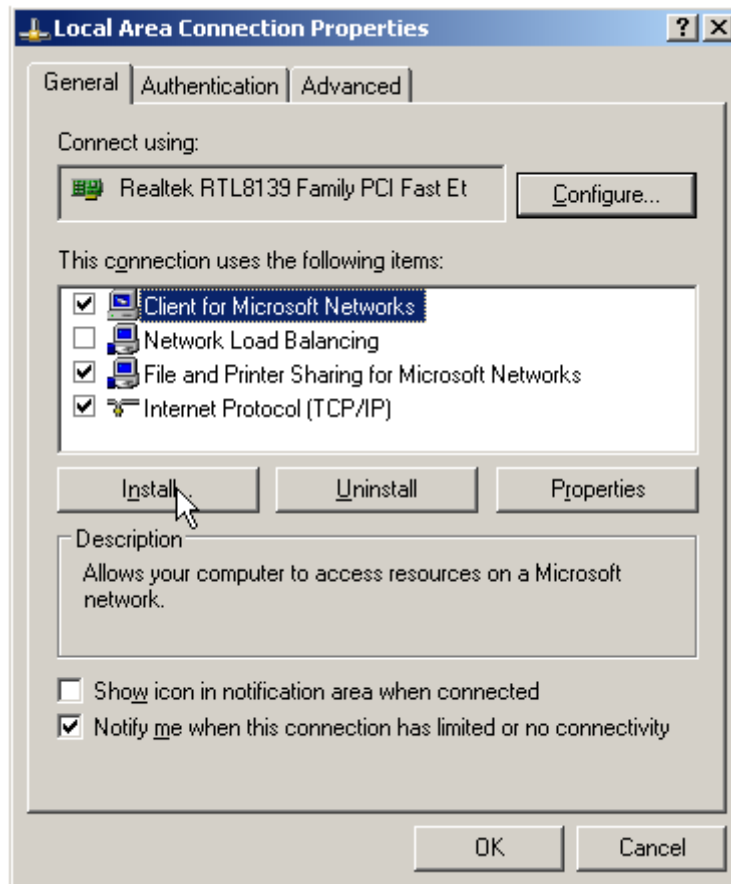


6. Perform **Step 1** to check whether dynamic IPv6 address assignment is enabled.

**Step 3** Enable and configure IPv6.


1. Choose **Start > Control Panel > Network Connection > Local Connection**.
2. Select **Properties**, select the following options, and click **Install**.

Figure 7-13 Enabling and configuring IPv6



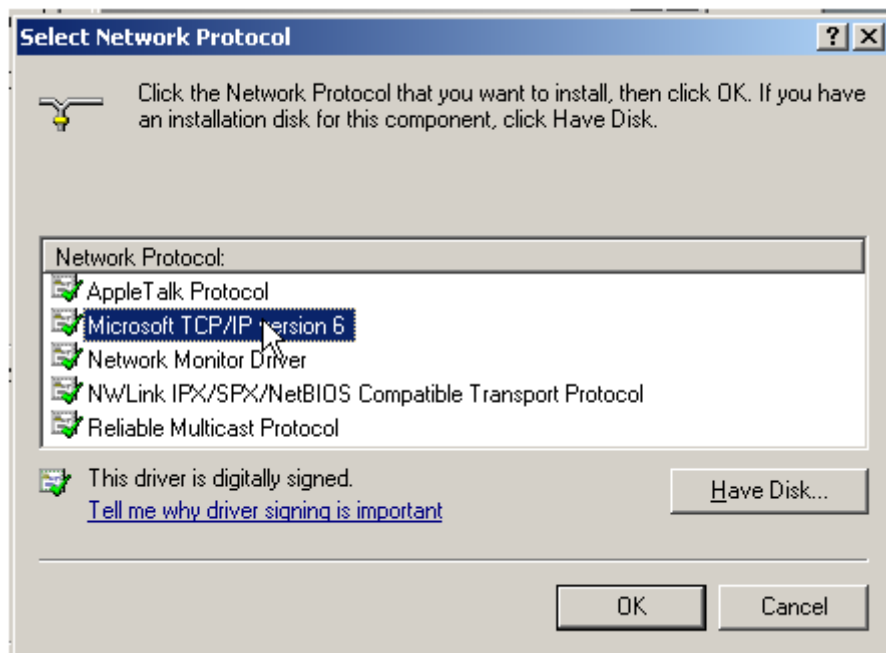
3. Select **Protocol** and click **Add**.

Figure 7-14 Adding the protocol



4. Select **Microsoft TCP/IP Version 6** and click **OK**.

Figure 7-15 Network protocols



- (Optional) Run the following commands depending on your ECS OS.  
For Windows Server 2008, run the following command in PowerShell or CMD:  
**netsh interface ipv6 set global randomizeidentifiers=disable**  
Disable the local connection and then enable it again.  
To disable the local connection, choose **Start > Control Panel > Network and Internet > Network and Sharing Center > Change Adapter Options**. Right-click the local connection and choose **Disable** from the shortcut menu.  
To enable the local connection, choose **Start > Control Panel > Network and Internet > Network and Sharing Center > Change Adapter Options**. Right-click the local connection and choose **Enable** from the shortcut menu.
- Perform [Step 1](#) to check whether dynamic IPv6 address assignment is enabled.

----End

## Linux (Automatically Enabling Dynamic Assignment of IPv6 Addresses)

The **ipv6-setup-xxx** tool can be used to enable Linux OSs to automatically acquire IPv6 addresses. *xxx* indicates a tool, which can be *rhel* or *debian*.

You can also enable dynamic IPv6 address assignment by following the instructions in [Linux \(Manually Enabling Dynamic Assignment of IPv6 Addresses\)](#).

**CAUTION**

- When you run **ipv6-setup-xxx**, the network service will be automatically restarted. As a result, the network is temporarily disconnected.
- If a private image created from a CentOS 6.x or Debian ECS with automatic IPv6 address assignment enabled is used to create an ECS in an environment that does not support IPv6, the ECS may start slow because of IPv6 address assignment timeout. Set the timeout duration for assigning IPv6 addresses to 30s by referring to [Setting the Timeout Duration for IPv6 Address Assignment](#) and try to create a new private image again.

**Step 1** Run the following command to check whether IPv6 is enabled for the ECS:

**ip addr**

- If only an IPv4 address is displayed, IPv6 is disabled. Enable it by referring to [Setting the Timeout Duration for IPv6 Address Assignment](#).

**Figure 7-16** IPv6 disabled

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
link/ether fa:16:3e: b8:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff
inet 192.168.1.100 brd 192.168.1.255 scope global noprefixroute dynamic eth0
 valid_lft 1193sec preferred_lft 1193sec
```

- If a link-local address (starting with fe80) is displayed, IPv6 is enabled but dynamic assignment of IPv6 addresses is not enabled.

**Figure 7-17** IPv6 enabled

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
link/ether fa:16:3e:75:af:4c brd ff:ff:ff:ff:ff:ff
inet 192.168.1.100 brd 192.168.1.255 scope global noprefixroute dynamic eth0
 valid_lft 76391sec preferred_lft 76391sec
inet6 fe80::f816:3eff:fe00:0000/64 scope link
 valid_lft forever preferred_lft forever
```

- If the following address is displayed, IPv6 is enabled and an IPv6 address has been assigned:

**Figure 7-18** IPv6 enabled and an IPv6 address assigned

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
link/ether fa:16:3e:75:af:4c brd ff:ff:ff:ff:ff:ff
inet 192.168.1.100 brd 192.168.1.255 scope global noprefixroute dynamic eth0
 valid_lft 86395sec preferred_lft 86395sec
inet6 2407:c080:802:0000:0000:0000:0000:0000/128 scope global dynamic
 valid_lft 7496sec preferred_lft 7196sec
inet6 fe80::f816:3eff:fe00:0000/64 scope link noprefixroute
 valid_lft forever preferred_lft forever
```

**NOTE**

IPv6 is enabled for Linux public images by default, as shown in [Figure 7-17](#).

**Step 2** Enable IPv6 for the ECS.

1. Run the following command to check whether IPv6 is enabled for the kernel:  
**sysctl -a | grep ipv6**
  - If a command output is displayed, IPv6 is enabled.
  - If no information is displayed, IPv6 is disabled. Go to [Step 2.2](#) to load the IPv6 module.

2. Run the following command to load the IPv6 module:  
**modprobe ipv6**
3. Add the following content to the `/etc/sysctl.conf` file:  
**net.ipv6.conf.all.disable\_ipv6=0**
4. Save the configuration and exit. Then, run the following command to load the configuration:  
**sysctl -p**

**Step 3** Enable dynamic IPv6 address assignment for the ECS.

1. Download **ipv6-setup-rhel** or **ipv6-setup-debian** with a required version and upload it to the target ECS.  
**ipv6-setup-xxx** modifies the configuration file of a NIC to enable dynamic IPv6 address assignment or adds such a configuration file for a NIC, and then restarts the NIC or network service.  
Contact the administrator to obtain the download paths of **ipv6-setup-rhel** and **ipv6-setup-debian**.
2. Run the following command to make **ipv6-setup-xxx** executable:  
**chmod +x ipv6-setup-xxx**
3. Run the following command to enable dynamic IPv6 address assignment for a NIC:

```
./ipv6-setup-xxx --dev [dev]
```

Example:

```
./ipv6-setup-xxx --dev eth0
```

 **NOTE**

- To enable dynamic IPv6 address assignment for all NICs, run the `./ipv6-setup-xxx` command.
- To learn how to use **ipv6-setup-xxx**, run the `./ipv6-setup-xxx --help` command.

----End

## Linux (Manually Enabling Dynamic Assignment of IPv6 Addresses)

---

 **CAUTION**

If a private image created from a CentOS 6.x or Debian ECS with automatic IPv6 address assignment enabled is used to create an ECS in an environment that does not support IPv6, the ECS may start slow because of IPv6 address assignment timeout. Set the timeout duration for assigning IPv6 addresses to 30s by referring to [Setting the Timeout Duration for IPv6 Address Assignment](#) and try to create a new private image again.

---

**Step 1** Run the following command to check whether IPv6 is enabled for the ECS:

```
ip addr
```

- If only an IPv4 address is displayed, IPv6 is disabled. Enable it by referring to [Step 2](#).

**Figure 7-19** IPv6 disabled

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
link/ether fa:16:3e: brd ff:ff:ff:ff:ff:ff
inet brd scope global noprefixroute dynamic eth0
valid_lft 1193sec preferred_lft 1193sec
```

- If a link-local address (starting with fe80) is displayed, IPv6 is enabled but dynamic assignment of IPv6 addresses is not enabled.

**Figure 7-20** IPv6 enabled

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
link/ether fa:16:3e: brd ff:ff:ff:ff:ff:ff
inet brd scope global noprefixroute dynamic eth0
valid_lft 76391sec preferred_lft 76391sec
inet6 fe80::f816: /64 scope link
valid_lft forever preferred_lft forever
```

- If the following address is displayed, IPv6 is enabled and an IPv6 address has been assigned:

**Figure 7-21** IPv6 enabled and an IPv6 address assigned

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
link/ether fa:16:3e:75:af:4c brd ff:ff:ff:ff:ff:ff
inet brd scope global noprefixroute dynamic eth0
valid_lft 86395sec preferred_lft 86395sec
inet6 2407:c080:802: /128 scope global dynamic
valid_lft 7496sec preferred_lft 7196sec
inet6 fe80::f816:3eff: /64 scope link noprefixroute
valid_lft forever preferred_lft forever
```

**NOTE**

IPv6 is enabled for Linux public images by default, as shown in [Figure 7-20](#).

**Step 2** Enable IPv6 for the ECS.

1. Run the following command to check whether IPv6 is enabled for the kernel:  
**sysctl -a | grep ipv6**
  - If a command output is displayed, IPv6 is enabled.
  - If no information is displayed, IPv6 is disabled. Go to [Step 2.2](#) to load the IPv6 module.
2. Run the following command to load the IPv6 module:  
**modprobe ipv6**
3. Add the following content to the `/etc/sysctl.conf` file:  
**net.ipv6.conf.all.disable\_ipv6=0**
4. Save the configuration and exit. Then, run the following command to load the configuration:  
**sysctl -p**

**Step 3** Enable dynamic IPv6 address assignment for the ECS.

- Ubuntu 18.04/20.04
  - a. Run the following command to access `/etc/netplan/`:  
**cd /etc/netplan/**
  - b. Run the following command to list the configuration file:  
**ls**



**Figure 7-22** Configuration file name

```
root@ecs- :/etc/netplan# ls
01-netcfg.yaml 01-network-manager-all.yaml
```

- c. Run the following command to edit the configuration file:  
**vi 01-network-manager-all.yaml**
- d. Append the following content to the configuration file (pay attention to the yaml syntax and text indentation):  
ethernets:  
  eth0:  
    dhcp6: true

**Figure 7-23** Edited configuration file

```
Let NetworkManager manage all devices on this system
network:
 version: 2
 renderer: NetworkManager
 ethernets:
 eth0:
 dhcp6: true
```

Save the changes and exit.

- e. Run the following command to make the changes take effect:  
**sudo netplan apply**
- Ubuntu 22.04
  - a. Run the following command to access `/etc/netplan/`:  
**cd /etc/netplan**
  - b. Run the following command to list the configuration file:  
**ls**

**Figure 7-24** Configuration file name

```
root@ecs-485b:/etc/netplan# ls
01-netcfg.yaml
```

- c. Run the following command to edit the configuration file:  
**vi 01-netcfg.yaml**
- d. Append the following content to the configuration file **01-netcfg.yaml** (pay attention to the yaml syntax and text indentation):  
ethernets:  
  eth0:  
    dhcp6: true

Figure 7-25 Edited configuration file

```
network:
 version: 2
 renderer: NetworkManager
 ethernets:
 eth0:
 dhcp4: true
 dhcp6: true
 eth1:
 dhcp4: true
 eth2:
 dhcp4: true
 eth3:
 dhcp4: true
 eth4:
 dhcp4: true
```

Save the changes and exit.

- e. Run the following command to make the changes take effect:  
**sudo netplan apply**
- f. Run the following command to edit `/etc/NetworkManager/NetworkManager.conf`:  
**vi /etc/NetworkManager/NetworkManager.conf**
- g. Append the following content to the configuration file **NetworkManager.conf** (pay attention to the file format and indentation):

```
[main]
plugins=ifupdown,keyfile
dhcp=dhclient

[ifupdown]
managed=true

[device]
wifi.scan-rand-mac-address=no
```

Figure 7-26 Modification result

```
[main]
plugins=ifupdown,keyfile
dhcp=dhclient

[ifupdown]
managed=true

[device]
wifi.scan-rand-mac-address=no
```

- h. Run the following command for the configuration to take effect:  
**systemctl restart NetworkManager**
- Debian
    - a. Add the following content to the `/etc/network/interfaces` file:

```
auto lo
iface lo inet loopback
```

```
auto eth0
iface eth0 inet dhcp
iface eth0 inet6 dhcp
pre-up sleep 3
```

- b. Add configurations for each NIC to the `/etc/network/interfaces` file. The following uses eth1 as an example:

```
auto eth1
iface eth1 inet dhcp
iface eth1 inet6 dhcp
pre-up sleep 3
```

- c. Run the following command to restart the network service:

```
service networking restart
```

#### NOTE

If no IPv6 address is assigned after the NICs are brought down and up, you can run this command to restart the network.

- d. Perform [Step 1](#) to check whether dynamic IPv6 address assignment is enabled.
- CentOS, EulerOS, or Fedora
    - a. Open the configuration file `/etc/sysconfig/network-scripts/ifcfg-eth0` of the primary NIC.

Add the following configuration items to the file:

```
IPV6INIT=yes
DHCPV6C=yes
```

- b. Edit the `/etc/sysconfig/network` file to add or modify the following line:  
`NETWORKING_IPV6=yes`
- c. For an ECS running CentOS 6, you need to edit the configuration files of its extension NICs. For example, if the extension NIC is eth1, you need to edit `/etc/sysconfig/network-scripts/ifcfg-eth1`.

Add the following configuration items to the file:

```
IPV6INIT=yes
DHCPV6C=yes
```

In CentOS 6.3, dhcpv6-client requests are filtered by **ip6tables** by default. So, you also need to add a rule allowing the dhcpv6-client request to the **ip6tables** file.

- i. Run the following command to add the rule to **ip6tables**:  
**ip6tables -A INPUT -m state --state NEW -m udp -p udp --dport 546 -d fe80::/64 -j ACCEPT**
- ii. Run the following command to save the rule in **ip6tables**:  
**service ip6tables save**

#### Figure 7-27 Example command

```
root@ecs-cd02 log]# ip6tables -A INPUT -m state --state NEW -m udp -p udp --dport 546 -d fe80::/64 -j ACCEPT
nf_conntrack version 0.5.0 (7964 buckets, 31856 max)
root@ecs-cd02 log]# service ip6tables save
ip6tables: Saving firewall rules to /etc/sysconfig/ip6tablef OK]
```

- d. (Optional) For CentOS 7/CentOS 8, change the IPv6 link-local address mode of extension NICs to EUI64.
  - i. Run the following command to query the NIC information:  
**nmcli con**

**Figure 7-28** Querying NIC information

```
[root@ecs-166b ~]# nmcli con
NAME UUID TYPE DEVICE
System eth0 5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03 ethernet eth0
Wired connection 1 9c92fad9-6ecb-3e6c-eb4d-8a47c6f50c04 ethernet eth1
Wired connection 1 3a73717e-65ab-93e8-b518-24f5af32dc0d ethernet eth2
```

- ii. Run the following command to change the IPv6 link-local address mode of eth1 to EUI64:

```
nmcli con modify "Wired connection 1" ipv6.addr-gen-mode eui64
```

**NOTE**

The NIC information varies depending on the CentOS series. In the command, *Wired connection 1* needs to be replaced with the value in the **NAME** column of the queried NIC information.

- iii. Run the following commands to bring eth1 down and up:

```
ifdown eth1
```

```
ifup eth1
```

- e. Restart the network service.

- i. For CentOS 6, run the following command to restart the network service:

```
service network restart
```

- ii. For CentOS 7/EulerOS/Fedora, run the following command to restart the network service:

```
systemctl restart NetworkManager
```

- f. Perform [Step 1](#) to check whether dynamic IPv6 address assignment is enabled.

- SUSE, openSUSE, or CoreOS

SUSE 11 SP4 does not support dynamic IPv6 address assignment.

No additional configuration is required for SUSE 12 SP1 or SUSE 12 SP2.

No additional configuration is required for openSUSE 13.2 or openSUSE 42.2.

No additional configuration is required for CoreOS 10.10.5.

----End

## Setting the Timeout Duration for IPv6 Address Assignment

After automatic IPv6 address assignment is configured on an ECS running CentOS 6.x or Debian, the ECS will be created as a private image. When this image is used to create an ECS in an environment that IPv6 is unavailable, the ECS may start slow because acquiring an IPv6 address times out. Before creating the private image, you can set the timeout duration for acquiring IPv6 addresses to 30s as follows:

- CentOS 6.x:

- a. Run the following command to edit the **dhclient.conf** file:

```
vi /etc/dhcp/dhclient.conf
```

- b. Press **i** to enter editing mode and add the timeout attribute to the file.  
timeout 30;

- c. Enter **:wq** to save the settings and exit.
- Debian 7.5:
  - a. Run the following command to edit the **networking** file:  
**vi /etc/init.d/networking**
  - b. Press **i** to enter editing mode and add the timeout attribute.

**Figure 7-29** Modification 1

```
115 case "$1" in
116 start)
117 if init_is_upstart; then
118 exit 1
119 fi
120 process_options
121 check_ifstate
122
123 if ["$CONFIGURE_INTERFACES" = no]
124 then
125 log_action_msg "Not configuring network interfaces, see /etc/default/networking"
126 exit 0
127 fi
128 set -f
129 exclusions=$(process_exclusions)
130 log_action_begin_msg "Configuring network interfaces"
131 if /usr/bin/timeout 30 ifup -a $exclusions $verbose && ifup_hotplug $exclusions $verbose
132 then
133 log_action_end_msg $?
134 else
135 log_action_end_msg $?
136 fi
137 ;;
138
139 stop)
140 if init_is_upstart; then
141 exit 0
142 fi
143 check_network_file_systems
144 check_network_swap
145
146 log_action_begin_msg "Deconfiguring network interfaces"
147 if /usr/bin/timeout 30 ifdown -a --exclude=lo $verbose; then
148 log_action_end_msg $?
149 fi
150 ;;
151 *)
152 echo "Usage: networking {start|stop|force-reload|restart}"
153 ;;
154 *)
155 ;;
156 *)
157 ;;
158 *)
159 ;;
160 *)
161 ;;
162 *)
163 ;;
164 *)
165 ;;
166 *)
167 ;;
168 *)
169 ;;
170 *)
171 ;;
172 *)
173 ;;
174 *)
175 ;;
176 *)
177 ;;
178 *)
179 ;;
180 *)
181 ;;
182 *)
183 ;;
```

**Figure 7-30** Modification 2

```
154 reload)
155 process_options
156
157 log_action_begin_msg "Reloading network interfaces configuration"
158 state=$(cat /run/network/ifstate)
159 if /usr/bin/timeout 30 ifdown -a --exclude=lo $verbose || true
160 then
161 if /usr/bin/timeout 30 ifup --exclude=lo $state $verbose; then
162 log_action_end_msg $?
163 else
164 log_action_end_msg $?
165 fi
166 fi
167 ;;
168
169 force-reload|restart)
170 if init_is_upstart; then
171 exit 1
172 fi
173 process_options
174
175 log_warning_msg "Running $0 $1 is deprecated because it may not re-enable some interfaces"
176 log_action_begin_msg "Reconfiguring network interfaces"
177 if /usr/bin/timeout 30 ifdown -a --exclude=lo $verbose || true
178 then
179 set -f
180 exclusions=$(process_exclusions)
181 if /usr/bin/timeout 30 ifup -a --exclude=lo $exclusions $verbose && ifup_hotplug $exclusions $verbose
182 then
183 log_action_end_msg $?
184 else
185 log_action_end_msg $?
186 fi
187 fi
188 ;;
189 *)
190 ;;
191 *)
192 ;;
193 *)
194 ;;
195 *)
196 ;;
197 *)
198 ;;
199 *)
200 ;;
201 *)
202 ;;
203 *)
204 ;;
205 *)
206 ;;
207 *)
208 ;;
209 *)
210 ;;
211 *)
212 ;;
213 *)
214 ;;
215 *)
216 ;;
217 *)
218 ;;
219 *)
220 ;;
221 *)
222 ;;
223 *)
224 ;;
225 *)
226 ;;
227 *)
228 ;;
229 *)
230 ;;
231 *)
232 ;;
233 *)
234 ;;
235 *)
236 ;;
237 *)
238 ;;
239 *)
240 ;;
241 *)
242 ;;
243 *)
244 ;;
245 *)
246 ;;
247 *)
248 ;;
249 *)
250 ;;
251 *)
252 ;;
253 *)
254 ;;
255 *)
256 ;;
257 *)
258 ;;
259 *)
260 ;;
261 *)
262 ;;
263 *)
264 ;;
265 *)
266 ;;
267 *)
268 ;;
269 *)
270 ;;
271 *)
272 ;;
273 *)
274 ;;
275 *)
276 ;;
277 *)
278 ;;
279 *)
280 ;;
281 *)
282 ;;
283 *)
284 ;;
285 *)
286 ;;
287 *)
288 ;;
289 *)
290 ;;
291 *)
292 ;;
293 *)
294 ;;
295 *)
296 ;;
297 *)
298 ;;
299 *)
300 ;;
301 *)
302 ;;
303 *)
304 ;;
305 *)
306 ;;
307 *)
308 ;;
309 *)
310 ;;
311 *)
312 ;;
313 *)
314 ;;
315 *)
316 ;;
317 *)
318 ;;
319 *)
320 ;;
321 *)
322 ;;
323 *)
324 ;;
325 *)
326 ;;
327 *)
328 ;;
329 *)
330 ;;
331 *)
332 ;;
333 *)
334 ;;
335 *)
336 ;;
337 *)
338 ;;
339 *)
340 ;;
341 *)
342 ;;
343 *)
344 ;;
345 *)
346 ;;
347 *)
348 ;;
349 *)
350 ;;
351 *)
352 ;;
353 *)
354 ;;
355 *)
356 ;;
357 *)
358 ;;
359 *)
360 ;;
361 *)
362 ;;
363 *)
364 ;;
365 *)
366 ;;
367 *)
368 ;;
369 *)
370 ;;
371 *)
372 ;;
373 *)
374 ;;
375 *)
376 ;;
377 *)
378 ;;
379 *)
380 ;;
381 *)
382 ;;
383 *)
384 ;;
385 *)
386 ;;
387 *)
388 ;;
389 *)
390 ;;
391 *)
392 ;;
393 *)
394 ;;
395 *)
396 ;;
397 *)
398 ;;
399 *)
400 ;;
401 *)
402 ;;
403 *)
404 ;;
405 *)
406 ;;
407 *)
408 ;;
409 *)
410 ;;
411 *)
412 ;;
413 *)
414 ;;
415 *)
416 ;;
417 *)
418 ;;
419 *)
420 ;;
421 *)
422 ;;
423 *)
424 ;;
425 *)
426 ;;
427 *)
428 ;;
429 *)
430 ;;
431 *)
432 ;;
433 *)
434 ;;
435 *)
436 ;;
437 *)
438 ;;
439 *)
440 ;;
441 *)
442 ;;
443 *)
444 ;;
445 *)
446 ;;
447 *)
448 ;;
449 *)
450 ;;
451 *)
452 ;;
453 *)
454 ;;
455 *)
456 ;;
457 *)
458 ;;
459 *)
460 ;;
461 *)
462 ;;
463 *)
464 ;;
465 *)
466 ;;
467 *)
468 ;;
469 *)
470 ;;
471 *)
472 ;;
473 *)
474 ;;
475 *)
476 ;;
477 *)
478 ;;
479 *)
480 ;;
481 *)
482 ;;
483 *)
484 ;;
485 *)
486 ;;
487 *)
488 ;;
489 *)
490 ;;
491 *)
492 ;;
493 *)
494 ;;
495 *)
496 ;;
497 *)
498 ;;
499 *)
500 ;;
501 *)
502 ;;
503 *)
504 ;;
505 *)
506 ;;
507 *)
508 ;;
509 *)
510 ;;
511 *)
512 ;;
513 *)
514 ;;
515 *)
516 ;;
517 *)
518 ;;
519 *)
520 ;;
521 *)
522 ;;
523 *)
524 ;;
525 *)
526 ;;
527 *)
528 ;;
529 *)
530 ;;
531 *)
532 ;;
533 *)
534 ;;
535 *)
536 ;;
537 *)
538 ;;
539 *)
540 ;;
541 *)
542 ;;
543 *)
544 ;;
545 *)
546 ;;
547 *)
548 ;;
549 *)
550 ;;
551 *)
552 ;;
553 *)
554 ;;
555 *)
556 ;;
557 *)
558 ;;
559 *)
560 ;;
561 *)
562 ;;
563 *)
564 ;;
565 *)
566 ;;
567 *)
568 ;;
569 *)
570 ;;
571 *)
572 ;;
573 *)
574 ;;
575 *)
576 ;;
577 *)
578 ;;
579 *)
580 ;;
581 *)
582 ;;
583 *)
584 ;;
585 *)
586 ;;
587 *)
588 ;;
589 *)
590 ;;
591 *)
592 ;;
593 *)
594 ;;
595 *)
596 ;;
597 *)
598 ;;
599 *)
600 ;;
601 *)
602 ;;
603 *)
604 ;;
605 *)
606 ;;
607 *)
608 ;;
609 *)
610 ;;
611 *)
612 ;;
613 *)
614 ;;
615 *)
616 ;;
617 *)
618 ;;
619 *)
620 ;;
621 *)
622 ;;
623 *)
624 ;;
625 *)
626 ;;
627 *)
628 ;;
629 *)
630 ;;
631 *)
632 ;;
633 *)
634 ;;
635 *)
636 ;;
637 *)
638 ;;
639 *)
640 ;;
641 *)
642 ;;
643 *)
644 ;;
645 *)
646 ;;
647 *)
648 ;;
649 *)
650 ;;
651 *)
652 ;;
653 *)
654 ;;
655 *)
656 ;;
657 *)
658 ;;
659 *)
660 ;;
661 *)
662 ;;
663 *)
664 ;;
665 *)
666 ;;
667 *)
668 ;;
669 *)
670 ;;
671 *)
672 ;;
673 *)
674 ;;
675 *)
676 ;;
677 *)
678 ;;
679 *)
680 ;;
681 *)
682 ;;
683 *)
684 ;;
685 *)
686 ;;
687 *)
688 ;;
689 *)
690 ;;
691 *)
692 ;;
693 *)
694 ;;
695 *)
696 ;;
697 *)
698 ;;
699 *)
700 ;;
701 *)
702 ;;
703 *)
704 ;;
705 *)
706 ;;
707 *)
708 ;;
709 *)
710 ;;
711 *)
712 ;;
713 *)
714 ;;
715 *)
716 ;;
717 *)
718 ;;
719 *)
720 ;;
721 *)
722 ;;
723 *)
724 ;;
725 *)
726 ;;
727 *)
728 ;;
729 *)
730 ;;
731 *)
732 ;;
733 *)
734 ;;
735 *)
736 ;;
737 *)
738 ;;
739 *)
740 ;;
741 *)
742 ;;
743 *)
744 ;;
745 *)
746 ;;
747 *)
748 ;;
749 *)
750 ;;
751 *)
752 ;;
753 *)
754 ;;
755 *)
756 ;;
757 *)
758 ;;
759 *)
760 ;;
761 *)
762 ;;
763 *)
764 ;;
765 *)
766 ;;
767 *)
768 ;;
769 *)
770 ;;
771 *)
772 ;;
773 *)
774 ;;
775 *)
776 ;;
777 *)
778 ;;
779 *)
780 ;;
781 *)
782 ;;
783 *)
784 ;;
785 *)
786 ;;
787 *)
788 ;;
789 *)
790 ;;
791 *)
792 ;;
793 *)
794 ;;
795 *)
796 ;;
797 *)
798 ;;
799 *)
800 ;;
801 *)
802 ;;
803 *)
804 ;;
805 *)
806 ;;
807 *)
808 ;;
809 *)
810 ;;
811 *)
812 ;;
813 *)
814 ;;
815 *)
816 ;;
817 *)
818 ;;
819 *)
820 ;;
821 *)
822 ;;
823 *)
824 ;;
825 *)
826 ;;
827 *)
828 ;;
829 *)
830 ;;
831 *)
832 ;;
833 *)
834 ;;
835 *)
836 ;;
837 *)
838 ;;
839 *)
840 ;;
841 *)
842 ;;
843 *)
844 ;;
845 *)
846 ;;
847 *)
848 ;;
849 *)
850 ;;
851 *)
852 ;;
853 *)
854 ;;
855 *)
856 ;;
857 *)
858 ;;
859 *)
860 ;;
861 *)
862 ;;
863 *)
864 ;;
865 *)
866 ;;
867 *)
868 ;;
869 *)
870 ;;
871 *)
872 ;;
873 *)
874 ;;
875 *)
876 ;;
877 *)
878 ;;
879 *)
880 ;;
881 *)
882 ;;
883 *)
884 ;;
885 *)
886 ;;
887 *)
888 ;;
889 *)
890 ;;
891 *)
892 ;;
893 *)
894 ;;
895 *)
896 ;;
897 *)
898 ;;
899 *)
900 ;;
901 *)
902 ;;
903 *)
904 ;;
905 *)
906 ;;
907 *)
908 ;;
909 *)
910 ;;
911 *)
912 ;;
913 *)
914 ;;
915 *)
916 ;;
917 *)
918 ;;
919 *)
920 ;;
921 *)
922 ;;
923 *)
924 ;;
925 *)
926 ;;
927 *)
928 ;;
929 *)
930 ;;
931 *)
932 ;;
933 *)
934 ;;
935 *)
936 ;;
937 *)
938 ;;
939 *)
940 ;;
941 *)
942 ;;
943 *)
944 ;;
945 *)
946 ;;
947 *)
948 ;;
949 *)
950 ;;
951 *)
952 ;;
953 *)
954 ;;
955 *)
956 ;;
957 *)
958 ;;
959 *)
960 ;;
961 *)
962 ;;
963 *)
964 ;;
965 *)
966 ;;
967 *)
968 ;;
969 *)
970 ;;
971 *)
972 ;;
973 *)
974 ;;
975 *)
976 ;;
977 *)
978 ;;
979 *)
980 ;;
981 *)
982 ;;
983 *)
984 ;;
985 *)
986 ;;
987 *)
988 ;;
989 *)
990 ;;
991 *)
992 ;;
993 *)
994 ;;
995 *)
996 ;;
997 *)
998 ;;
999 *)
1000 ;;
```

- Debian 8.2.0/8.8.0
  - a. Run the following command to edit the **network-pre.conf** file:  
**vi /lib/systemd/system/networking.service.d/network-pre.conf**
  - b. Press **i** to enter editing mode and add the timeout attribute to the file.  
[Service]  
TimeoutStartSec=30
- Debian 9.0
  - a. Run the following command to edit the **networking.service** file:  
**vi /etc/systemd/system/network-online.target.wants/networking.service**

- b. Press **i** to enter editing mode and change **TimeoutStartSec=5min** to **TimeoutStartSec=30**.

# 8 EIP

## 8.1 Binding an EIP

### Scenarios

You can assign an EIP and bind it to an ECS to enable the ECS to access the Internet.

### Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. In the ECS list, select the ECS to which an EIP is to be bound, and choose **More > Manage Network > Bind EIP** in the **Operation** column.
4. In the displayed dialog box, select an EIP

#### NOTE

If no EIP is available in the current region, the EIP list is empty. In such a case, allocate an EIP and then bind it.

5. Click **OK**.  
After the EIP is bound, view it in the ECS list on the **Elastic Cloud Server** page.

## 8.2 Unbinding an EIP

### Scenarios

This section describes how to unbind an EIP from an ECS.

### Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.

3. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network > Unbind EIP**.
4. Confirm the EIP to be unbound and click **OK**.

 **NOTE**

Unreleased EIPs will continue to be billed. To stop the EIPs from being billed, release them.

## 8.3 Changing an EIP Bandwidth

### Scenarios

If an EIP has been bound to the ECS, the ECS can access the Internet using the bandwidth associated with the EIP. This section describes how to adjust the bandwidth of an ECS.

### Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network > Modify Bandwidth**.
4. Change the bandwidth name and size as prompted.

## 8.4 Enabling Internet Connectivity for an ECS Without an EIP

### Scenarios

To ensure platform security and conserve EIPs, EIPs are assigned only to specified ECSs. ECSs without EIPs cannot access the Internet directly. If these ECSs need to access the Internet (for example, to perform a software upgrade or install a patch), you can select an ECS with an EIP bound to function as a proxy ECS, providing an access channel for these ECSs.

 **NOTE**

NAT Gateway is recommended, which provides both the SNAT and DNAT functions for your ECSs in a VPC and allows the ECSs to access or provide services accessible from the Internet. For details, see *NAT Gateway User Guide*.

### Prerequisites


- A proxy ECS with an EIP bound is available.
- The IP address of the proxy ECS is in the same network and same security group as the ECSs that need to access the Internet.



## Linux Proxy ECS

In this example, the proxy ECS runs CentOS 6.5.

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. In the search box above the upper right corner of the ECS list, enter the proxy ECS name for search.
4. Click the name of the proxy ECS. The page providing details about the ECS is displayed.

5. On the **Network Interfaces** tab, click . Then, disable **Source/Destination Check**.

By default, the source/destination check function is enabled. When this function is enabled, the system checks whether source IP addresses contained in the packets sent by ECSs are correct. If the IP addresses are incorrect, the system does not allow the ECSs to send the packets. This mechanism prevents packet spoofing, thereby improving system security. However, this mechanism prevents the packet sender from receiving returned packets. Therefore, disable the source/destination check.

6. Log in to the proxy ECS.  
For more details, see [Login Overview](#).
7. Run the following command to check whether the proxy ECS can access the Internet:

```
ping www.baidu.com
```

The proxy ECS can access the Internet if information similar to the following is displayed:

**Figure 8-1** Checking connectivity

```
[root@ecs-f4f0 ~]# ping www.baidu.com
PING www.a.shifen.com (61.135.169.121) 56(84) bytes of data:
64 bytes from 61.135.169.121 (61.135.169.121): icmp_seq=1 ttl=47 time=2.77 ms
64 bytes from 61.135.169.121 (61.135.169.121): icmp_seq=2 ttl=47 time=2.65 ms
64 bytes from 61.135.169.121 (61.135.169.121): icmp_seq=3 ttl=47 time=2.61 ms
64 bytes from 61.135.169.121 (61.135.169.121): icmp_seq=4 ttl=47 time=2.83 ms
64 bytes from 61.135.169.121 (61.135.169.121): icmp_seq=5 ttl=47 time=2.69 ms
64 bytes from 61.135.169.121 (61.135.169.121): icmp_seq=6 ttl=47 time=2.63 ms
```

8. Run the following command to check whether IP forwarding is enabled on the proxy ECS:

```
cat /proc/sys/net/ipv4/ip_forward
```

- If **0** (disabled) is displayed, go to **9**.
- If **1** (enabled) is displayed, go to **14**.

9. Run the following command to open the IP forwarding configuration file in the vi editor:

```
vi /etc/sysctl.conf
```

10. Press **i** to enter editing mode.
11. Set the **net.ipv4.ip\_forward** value to **1**.  
Set the **net.ipv4.ip\_forward** value to **1**.

 NOTE

If the `sysctl.conf` file does not contain the `net.ipv4.ip_forward` parameter, run the following command to add it:

```
echo net.ipv4.ip_forward=1 >> /etc/sysctl.conf
```

12. Press **Esc**, type `:wq`, and press **Enter**.

The system saves the configurations and exits the vi editor.

13. Run the following command to make the modification take effect:

```
sysctl -p /etc/sysctl.conf
```

14. Run the following commands to configure default **iptables** rules:

```
iptables -P INPUT ACCEPT
```

```
iptables -P OUTPUT ACCEPT
```

```
iptables -P FORWARD ACCEPT
```

---

 CAUTION

Running **iptables -P INPUT ACCEPT** will set default INPUT policy to ACCEPT, which poses security risks. You are advised to set security group rules to restrict inbound access.

- 
15. Run the following command to configure source network address translation (SNAT) to enable ECSs in the same network segment to access the Internet through the proxy ECS:

```
iptables -t nat -A POSTROUTING -o eth0 -s subnet/netmask-bits -j SNAT --to nat-instance-ip
```

For example, if the proxy ECS is in network 192.168.125.0, the subnet mask has 24 bits, and the private IP address is 192.168.125.4, run the following command:

```
iptables -t nat -A POSTROUTING -o eth0 -s 192.168.125.0/24 -j SNAT --to 192.168.125.4
```

 NOTE

To retain the preceding configuration even after the ECS is restarted, run the `vi /etc/rc.local` command to edit the `rc.local` file. Specifically, copy the rule described in step 15 into `rc.local`, press **Esc** to exit Insert mode, and enter `:wq` to save the settings and exit.

16. Run the following commands to save the iptables configuration and make it start up automatically upon ECS startup:

```
service iptables save
```

```
chkconfig iptables on
```

17. Run the following command to check whether SNAT has been configured:

```
iptables -t nat --list
```

SNAT has been configured if information similar to [Figure 8-2](#) is displayed.

**Figure 8-2** Successful SNAT configuration

```
[root@host- ~]# iptables -t nat --list
Chain PREROUTING (policy ACCEPT)
target prot opt source destination
Chain POSTROUTING (policy ACCEPT)
target prot opt source destination
SNAT all -- 192.168.125.0/24 anywhere to:192.168.125.4
SNAT all -- 0.0.0.0/0 anywhere to:192.168.125.4
```

18. Add a route.
  - a. Log in to the management console.
  - b. Under **Network**, click **Virtual Private Cloud**.
  - c. Choose **Route Tables** in the left navigation pane. In the route table list, click a target route table. On the displayed page, click **Add Route**.
  - d. Set route information on the displayed page.
    - **Destination:** indicates the destination network segment. The default value is **0.0.0.0/0**.
    - **Next Hop:** indicates the private IP address of the proxy ECS. You can obtain the private IP address of the ECS on the **Elastic Cloud Server** page.
19. To delete the added iptables rules, run the following command:  
**iptables -t nat -D POSTROUTING -o eth0 -s subnet/netmask-bits -j SNAT --to nat-instance-ip**

For example, if the proxy ECS is in network segment 192.168.125.0, the subnet mask has 24 bits, and the private IP address is 192.168.125.4, run the following command:

```
iptables -t nat -D POSTROUTING -o eth0 -s 192.168.125.0/24 -j SNAT --to 192.168.125.4
```

# 9 Security

---

## 9.1 Security Groups

### 9.1.1 Overview

#### Security Group

A security group is a collection of access control rules for ECSs that have the same security protection requirements and that are mutually trusted. After a security group is created, you can create various access rules for the security group, these rules will apply to all ECSs added to this security group.

You can also customize a security group or use the default one. The system provides a default security group for you, which permits all outbound traffic and denies inbound traffic. ECSs in a security group are accessible to each other. For details about the default security group, see [Default Security Group and Rules](#).

#### NOTE

If two ECSs are in the same security group but in different VPCs, the security group does not take effect. You can use a VPC peering connection to connect the two VPCs first.

#### Security Group Rules

After a security group is created, you can add rules to the security group. A rule applies either to inbound traffic (ingress) or outbound traffic (egress). After ECSs are added to the security group, they are protected by the rules of that group.

Each security group has default rules. For details, see [Default Security Group and Rules](#). You can also customize security group rules. For details, see [Configuring Security Group Rules](#).

#### Security Group Constraints

- By default, you can add up to 50 security group rules to a security group.

- By default, you can add an ECS or extension NIC to up to five security groups. In such a case, the rules of all the selected security groups are aggregated to take effect.


## 9.1.2 Default Security Group and Rules

Note the following when using default security group rules:

- Inbound rules control incoming traffic to instances in the default security group. The instances can only communicate with each other but cannot be accessed from external networks.
- Outbound rules allow all traffic from the instances in the default security group to external networks.

Figure 9-1 shows the default security group.

Figure 9-1 Default security group



### NOTE

- You cannot delete the default security group, but you can modify existing rules or add rules to the group.
- The default security group is automatically created to simplify the process of creating an instance for the first time. The default security group denies all external requests. To log in to an instance, add a security group rule by referring to [Remotely Logging In to an ECS from a Local Server](#).

Table 9-1 describes the rules in the default security group.

Table 9-1 Default security group rules

| Direction | Protocol | Port/Range | Source/Destination     | Description                  |
|-----------|----------|------------|------------------------|------------------------------|
| Outbound  | All      | All        | Destination: 0.0.0.0/0 | Allows all outbound traffic. |

| Direction | Protocol | Port/Range | Source/Destination                                         | Description                                                                                                        |
|-----------|----------|------------|------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| Inbound   | All      | All        | Source: the current security group (for example, sg-xxxxx) | Allows communications among ECSs within the security group and denies all inbound traffic (incoming data packets). |

### 9.1.3 Security Group Configuration Examples

When you create instances, such as cloud servers, containers, and databases, in a VPC subnet, you can use the default security group or create a security group. You can add inbound and outbound rules to the default or your security group to control traffic from and to the instances in the security group. Here are some common security group configuration examples:

- [Remotely Logging In to an ECS from a Local Server](#)
- [Remotely Connecting to an ECS from a Local Server to Upload or Download FTP Files](#)
- [Setting Up a Website on an ECS to Provide Services Externally](#)
- [Using ping Command to Verify Network Connectivity](#)
- [Enabling Communications Between Instances in Different Security Groups](#)
- [Allowing External Instances to Access the Database Deployed on an ECS](#)
- [Allowing ECSs to Access Specific External Websites](#)

#### Precautions

Note the following before configuring security group rules:

- Instances associated with different security groups are isolated from each other by default.
- Generally, a security group denies all external requests by default. You need to add inbound rules to allow specific traffic to the instances in the security group.
- By default, outbound security group rules allow all requests from the instances in the security group to access external resources. If outbound rules are deleted, the instances in the security group cannot communicate with external resources. To allow outbound traffic, you need to add outbound rules by referring to [Table 9-2](#).

**Table 9-2** Default outbound rules in a security group

| Direction | Protocol & Port | Destination | Description                                                                                    |
|-----------|-----------------|-------------|------------------------------------------------------------------------------------------------|
| Outbound  | All             | 0.0.0.0/0   | This rule allows the instances in the security group to access any IPv4 address over any port. |

## Remotely Logging In to an ECS from a Local Server

A security group denies all external requests by default. To remotely log in to an ECS in a security group from a local server, add an inbound rule based on the OS running on the ECS.

- To remotely log in to a Linux ECS using SSH, enable port 22. For details, see [Table 9-3](#).
- To remotely log in to a Windows ECS using RDP, enable port 3389. For details, see [Table 9-4](#).

**Table 9-3** Remotely logging in to a Linux ECS using SSH

| Direction | Protocol & Port | Source                |
|-----------|-----------------|-----------------------|
| Inbound   | TCP: 22         | IP address: 0.0.0.0/0 |

**Table 9-4** Remotely logging in to a Windows ECS using RDP

| Direction | Protocol & Port | Source                |
|-----------|-----------------|-----------------------|
| Inbound   | TCP: 3389       | IP address: 0.0.0.0/0 |

### NOTICE

If the source is set to 0.0.0.0/0, any IP address can be used to remotely log in to the ECS. To ensure security, set the source to a specific IP address based on service requirements. For details about the configuration example, see [Table 9-5](#).

**Table 9-5** Remotely logging in to an ECS using a specified IP address

| ECS Type  | Direction | Protocol & Port | Source                     |
|-----------|-----------|-----------------|----------------------------|
| Linux ECS | Inbound   | TCP: 22         | IP address: 192.168.0.0/24 |

| ECS Type    | Direction | Protocol & Port | Source                   |
|-------------|-----------|-----------------|--------------------------|
| Windows ECS | Inbound   | TCP: 3389       | IP address: 10.10.0.0/24 |

## Remotely Connecting to an ECS from a Local Server to Upload or Download FTP Files

By default, a security group denies all external requests. If you need to remotely connect to an ECS from a local server to upload or download files, you need to enable FTP ports 20 and 21.

**Table 9-6** Remotely connecting to an ECS from a local server to upload or download files

| Direction | Protocol & Port | Source                |
|-----------|-----------------|-----------------------|
| Inbound   | TCP: 20-21      | IP address: 0.0.0.0/0 |

### NOTICE

You must first install the FTP server program on the ECSs and check whether ports 20 and 21 are working properly.

## Setting Up a Website on an ECS to Provide Services Externally

A security group denies all external requests by default. If you have set up a website on an ECS that can be accessed externally, you need to add an inbound rule to the ECS security group to allow access over specific ports, such as HTTP (80) and HTTPS (443).

**Table 9-7** Setting up a website on an ECS to provide services externally

| Direction | Protocol & Port | Source                |
|-----------|-----------------|-----------------------|
| Inbound   | TCP: 80         | IP address: 0.0.0.0/0 |
| Inbound   | TCP: 443        | IP address: 0.0.0.0/0 |

## Using ping Command to Verify Network Connectivity

Ping works by sending an Internet Control Message Protocol (ICMP) Echo Request. To ping an ECS from your PC to verify the network connectivity, you need to add an inbound rule to the security group of the ECS to allow ICMP traffic.



**Table 9-8** Using **ping** command to verify network connectivity

| Direction | Protocol & Port | Source                |
|-----------|-----------------|-----------------------|
| Inbound   | ICMP: All       | IP address: 0.0.0.0/0 |

## Enabling Communications Between Instances in Different Security Groups

Instances in the same VPC but associated with different security groups cannot communicate with each other. If you want ECSs in security group **sg-A** to access MySQL databases in security group **sg-B**, you need to add an inbound rule to security group **sg-B** to allow access from ECSs in security group **sg-A**.

**Table 9-9** Enabling communications between instances in different security groups

| Direction | Protocol & Port | Source               |
|-----------|-----------------|----------------------|
| Inbound   | TCP: 3306       | Security group: sg-A |

## Allowing External Instances to Access the Database Deployed on an ECS

A security group denies all external requests by default. If you have deployed a database on an ECS and want the database to be accessed from external instances on a private network, you need to add an inbound rule to the security group of the ECS to allow access over corresponding ports. Here are some common ports for databases:

- MySQL: port 3306
- Oracle: port 1521
- MS SQL: port 1433
- PostgreSQL: port 5432
- Redis: port 6379

**Table 9-10** Allowing external instances to access the database deployed on an ECS

| Direction | Protocol & Port | Source                     | Description                                                                                             |
|-----------|-----------------|----------------------------|---------------------------------------------------------------------------------------------------------|
| Inbound   | TCP: 3306       | Security group: sg-A       | This rule allows the ECSs in security group <b>sg-A</b> to access the MySQL database service.           |
| Inbound   | TCP: 1521       | Security group: sg-B       | This rule allows the ECSs in security group <b>sg-B</b> to access the Oracle database service.          |
| Inbound   | TCP: 1433       | IP address: 172.16.3.21/32 | This rule allows the ECS whose private IP address is 172.16.3.21 to access the MS SQL database service. |

| Direction | Protocol & Port | Source                        | Description                                                                                                                   |
|-----------|-----------------|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| Inbound   | TCP: 5432       | IP address:<br>192.168.0.0/24 | This rule allows ECSs whose private IP addresses are in the 192.168.0.0/24 network to access the PostgreSQL database service. |

**NOTICE**

In this example, the source is for reference only. Set the source address based on your requirements.

## Allowing ECSs to Access Specific External Websites

By default, a security group allows all outbound traffic. [Table 9-12](#) lists the default rules. If you want to allow ECSs to access specific websites, configure the security group as follows:

1. Add outbound rules to allow traffic over specific ports to specific IP addresses.

**Table 9-11** Allowing ECSs to access specific external websites

| Direction | Protocol & Port | Destination                  | Description                                                                                                                                       |
|-----------|-----------------|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| Outbound  | TCP: 80         | IP address:<br>132.15.XX.XX  | This rule allows ECSs in the security group to access the external website at <a href="http://132.15.XX.XX:80">http://132.15.XX.XX:80</a> .       |
| Outbound  | TCP: 443        | IP address:<br>145.117.XX.XX | This rule allows ECSs in the security group to access the external website at <a href="https://145.117.XX.XX:443">https://145.117.XX.XX:443</a> . |

2. Delete the original outbound rules that allow all traffic.

**Table 9-12** Default outbound rules in a security group

| Direction | Protocol & Port | Destination | Description                                                                                    |
|-----------|-----------------|-------------|------------------------------------------------------------------------------------------------|
| Outbound  | All             | 0.0.0.0/0   | This rule allows the instances in the security group to access any IPv4 address over any port. |

## 9.1.4 Configuring Security Group Rules

### Scenarios

Similar to firewall, a security group is a logical group used to control network access. You can define access rules for a security group to protect the ECSs that are added to this security group.

- **Inbound:** Inbound rules allow external network traffic to be sent to the ECSs in the security group.
- **Outbound:** Outbound rules allow network traffic from the ECSs in the security group to be sent out of the security group.

For details about the default security group rules, see *Virtual Private Cloud User Guide*. For details about configuration examples for security group rules, see [Security Group Configuration Examples](#).

### Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. On the **Elastic Cloud Server** page, click the name of the target ECS.  
The page providing details about the ECS is displayed.
4. Click the **Security Groups** tab, expand the information of the security group, and view security group rules.
5. Click the security group ID.  
The system automatically switches to the security group details page.
6. Configure required parameters.  
You can click + to add more inbound rules.

**Table 9-13** Inbound rule parameter description

| Parameter       | Description                                                                                                                                                                                          | Example Value   |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Protocol & Port | The network protocol used to match traffic in a security group rule. The value can be <b>All</b> , <b>TCP</b> , <b>UDP</b> , <b>GRE</b> , and <b>ICMP</b> .                                          | TCP             |
|                 | Destination port used to match traffic in a security group rule. The value can be from 1 to 65535.<br>Inbound rules control incoming traffic over specific ports to instances in the security group. | 22, or<br>22-30 |

| Parameter   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                             | Example Value |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| Source      | <p>Source of the security group rule. The value can be an IP address or a security group to allow access from IP addresses or instances in the security group.</p> <ul style="list-style-type: none"> <li>IPv4 address: xxx.xxx.xxx.xxx/32</li> <li>Subnet: xxx.xxx.xxx.0/24</li> <li>Any IP address: 0.0.0.0/0</li> </ul> <p>If the source is a security group, this rule will apply to all instances associated with the selected security group.</p> | 0.0.0.0/0     |
| Description | <p>Supplementary information about the security group rule. This parameter is optional.</p> <p>The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (&lt; or &gt;).</p>                                                                                                                                                                                                                        | N/A           |

- Configure required parameters.  
You can click + to add more outbound rules.

**Table 9-14** Outbound rule parameter description

| Parameter                | Description                                                                                                                                                                                                                                                                                                                                                                                                              | Example Value |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| Protocol/<br>Application | The network protocol. Currently, the value can be <b>All</b> , <b>TCP</b> , <b>UDP</b> , <b>ICMP</b> , or more.                                                                                                                                                                                                                                                                                                          | TCP           |
| Port & Destination       | <p><b>Port:</b> The port or port range over which the traffic can leave your ECS. The value ranges from 1 to 65535.</p>                                                                                                                                                                                                                                                                                                  | 22, or 22-30  |
|                          | <p><b>Destination:</b> The destination of the security group rule. The value can be a single IP address or a security group to allow access to IP addresses or instances in the security group. For example:</p> <ul style="list-style-type: none"> <li>xxx.xxx.xxx.xxx/32 (IPv4 address)</li> <li>xxx.xxx.xxx.0/24 (IP address range)</li> <li>0.0.0.0/0 (all IP addresses)</li> <li>sg-abc (security group)</li> </ul> | 0.0.0.0/0     |

| Parameter   | Description                                                                                                                                                                                                     | Example Value |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| Description | Supplementary information about the security group rule. This parameter is optional.<br>The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >). | N/A           |

- Click **OK** to complete the security rule configuration.

## 9.1.5 Changing a Security Group

### Scenarios

To change the security group associated with an ECS network interface, perform the operations described in this section.

### Constraints

- Changing the security group will overwrite the original security group settings.
- Using multiple security groups may deteriorate ECS network performance. You are advised to select no more than five security groups.

### Procedure

- Log in to the management console.
- Under **Computing**, click **Elastic Cloud Server**.
- In the ECS list, locate the row that contains the target ECS. Click **More** in the **Operation** column and select **Manage Network > Change Security Group**. The **Change Security Group** dialog box is displayed.
- Select the target NIC and security groups.  
To create a security group, click **Create Security Group**.

#### NOTE

Using multiple security groups may deteriorate ECS network performance. You are suggested to select no more than five security groups.

- Click **OK**.

# 10 Passwords and Key Pairs

## 10.1 Passwords

### 10.1.1 Application Scenarios for Using Passwords

The password for logging in to your ECS is important and please keep it secure. You can reset the password if it is forgotten or expires.

**Table 10-1** provides guidance on how to reset your password in different scenarios.

**Table 10-1** Resetting a password

| Reference                                                                                 | Prerequisites                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">Resetting the Password for Logging In to an ECS on the Management Console</a> | The password reset plug-in has been installed.<br><b>NOTE</b> <ul style="list-style-type: none"><li>If your ECS was created using a public image, the password reset plug-in was installed on the ECS by default.</li><li>The reference is for Windows or Linux ECSs.</li><li>One-click password reset is available only in Chinese regions.</li></ul> |
| <a href="#">Changing the Login Password on an ECS</a>                                     | N/A<br><b>NOTE</b><br>The reference is for Windows or Linux ECSs.                                                                                                                                                                                                                                                                                      |
| <a href="#">Resetting the Password for Logging In to a Windows ECS</a>                    | The password reset plug-in has not been installed.                                                                                                                                                                                                                                                                                                     |
| <a href="#">Resetting the Password for Logging In to a Linux ECS</a>                      | The password reset plug-in has not been installed.                                                                                                                                                                                                                                                                                                     |

## Background

**Table 10-2** shows the ECS password complexity requirements.

**Table 10-2** Password complexity requirements

| Parameter | Requirement                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Password  | <ul style="list-style-type: none"><li>• Consists of 8 to 26 characters. For specific requirements on the password length, see the information displayed on the console.</li><li>• Contains at least three of the following character types:<ul style="list-style-type: none"><li>– Uppercase letters</li><li>– Lowercase letters</li><li>– Digits</li><li>– Special characters for Windows: \$!@%-_=[:./,?</li><li>– Special characters for Linux: !@%-_=[:./^,{}?</li></ul></li><li>• Cannot contain the username or the username spelled backwards.</li><li>• Cannot contain more than two consecutive characters in the same sequence as they appear in the username. (This requirement applies only to Windows ECSs.)</li><li>• Cannot start with a slash (/) for Windows ECSs.</li></ul> |

## 10.1.2 Resetting the Password for Logging In to an ECS on the Management Console

### Scenarios

If you did not set a password when creating an ECS, or the password expires or is forgotten, reset the password by following the instructions provided in this section.

### Prerequisites

- You have installed the password reset plug-in before your ECS password expires or is forgotten.
  - If your ECS was created using a public image, the password reset plug-in was installed on the ECS by default.
  - If your ECS was created using a private image and has no password reset plug-in installed, see [Resetting the Password for Logging In to a Windows ECS](#) or [Resetting the Password for Logging In to a Linux ECS](#).
- One-click password reset is available only in Chinese regions.
- Do not delete the CloudResetPwdAgent or CloudResetPwdUpdateAgent process. Otherwise, one-click password reset will not be available.
- One-click password reset can be used on the ECSs created using SUSE 11 SP4 only if their memory capacity is greater than or equal to 4 GiB.

- DHCP is enabled in the VPC to which the ECS belongs.
- The ECS network connectivity is normal.

## Procedure

Perform the following operations to change the login password of one or multiple ECSs in a batch on the management console.

### NOTE

If you reset the password when the ECS is running, the new password takes effect only after the ECS is restarted. You can manually restart the ECS after resetting the password, or select **Auto Restart** when resetting the password.

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Locate the row containing the target ECS, click **More** in the **Operation** column, and select **Reset Password** from the drop-down list.

### NOTE

The ECSs with the one-click password reset plug-in installed support password reset in a batch. To do so, perform the following operations:

1. Select the target ECSs.
2. Click **More** above the ECS list and select **Reset Password** from the drop-down list.
3. Set a new password as prompted.  
After the resetting, the login passwords of these ECSs are the same.
4. Set and confirm a new password as prompted.

The new password must meet the complexity requirements listed in [Table 10-3](#).



**Table 10-3** Password complexity requirements

| Parameter | Requirement                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Password  | <ul style="list-style-type: none"><li>• Consists of 8 to 26 characters. For specific requirements on the password length, see the information displayed on the console.</li><li>• Contains at least three of the following character types:<ul style="list-style-type: none"><li>- Uppercase letters</li><li>- Lowercase letters</li><li>- Digits</li><li>- Special characters for Windows: \$!@%-_+=+[]:./,?</li><li>- Special characters for Linux: !@%-_+=+[]:/^,{}?</li></ul></li><li>• Cannot contain the username or the username spelled backwards.</li><li>• Cannot contain more than two consecutive characters in the same sequence as they appear in the username. (This requirement applies only to Windows ECSs.)</li><li>• Cannot start with a slash (/) for Windows ECSs.</li></ul> |

5. Click **OK**.

It takes about 10 minutes for the system to reset the password. Do not repeatedly perform this operation.

- If the ECS is running when you reset the password, you need to manually restart the ECS for the new password to take effect.
- If the ECS is stopped, the new password will take effect after you start the ECS.

## Helpful Links

- [Why Does Login to My ECS Using the Reset Password Fail?](#)
- [Why Am I Seeing the Message Indicating That the Port Is Used by a One-Click Password Reset Plug-in?](#)

## 10.1.3 Changing the Login Password on an ECS

### Scenarios

This section describes how to change the password for logging in to an ECS when the password is about to expire, the password is forgotten, or you are logging in to the ECS for the first time. It is a good practice to change the initial password upon the first login.

You are advised to reset the ECS login password on the management console by referring to [Resetting the Password for Logging In to an ECS on the Management Console](#).

## Prerequisites

The ECS can be logged in.

## Background

[Table 10-4](#) shows the ECS password complexity requirements.

**Table 10-4** Password complexity requirements

| Parameter | Requirement                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Password  | <ul style="list-style-type: none"><li>• Consists of 8 to 26 characters. For specific requirements on the password length, see the information displayed on the console.</li><li>• Contains at least three of the following character types:<ul style="list-style-type: none"><li>– Uppercase letters</li><li>– Lowercase letters</li><li>– Digits</li><li>– Special characters for Windows: \$!@%-_+=[:./,;?</li><li>– Special characters for Linux: !@%-_+=[:./^,{}?</li></ul></li><li>• Cannot contain the username or the username spelled backwards.</li><li>• Cannot contain more than two consecutive characters in the same sequence as they appear in the username. (This requirement applies only to Windows ECSs.)</li><li>• Cannot start with a slash (/) for Windows ECSs.</li></ul> |

## Windows

1. Log in to the ECS.  
For details, see [Login Overview](#).
2. Press **Win+R** to start the **Run** dialog box.
3. Enter **cmd** to open the command-line interface (CLI) window.
4. Run the following command to change the password (the new password must meet the requirements described in [Table 10-4](#)):

```
net user Administrator New password
```

## 10.1.4 Resetting the Password for Logging In to a Windows ECS

### Scenarios

If the password reset plug-in is not installed on a Windows ECS or a password reset does not take effect, you can reset the password following the instructions provided in this section.

The method described in this section can only be used to change the password of a local Windows account, but not the password of a domain account.

For details about the operations performed on Linux ECS, see [Resetting the Password for Logging In to a Linux ECS](#).

#### NOTE

- If your ECS has password reset plug-in installed, follow the instructions provided in [Resetting the Password for Logging In to an ECS on the Management Console](#) to obtain a new password.
- If your ECS was created using a public image, the password reset plug-in was installed on the ECS by default. To reset the password, see [Resetting the Password for Logging In to an ECS on the Management Console](#).
- Before you perform operations in this section, we recommend you to back up the system disk to prevent data loss.

## Prerequisites

- A temporary Linux ECS which runs Ubuntu 14.04 or later and locates in the same AZ as the target ECS is available.
- You have bound an EIP to the temporary ECS and configured the apt-get source.
- You have used either of the following methods to install **ntfs-3g** and **chntpw** software packages on the temporary ECS:

Method 1:

Run the following command to install the **ntfs-3g** and **chntpw** software packages:

```
sudo apt-get install ntfs-3g chntpw
```

Method 2:

Download the ntfs-3g and chntpw software packages of the version required by the temporary ECS OS.

## Procedure

1. Stop the original ECS and detach the system disk.
  - a. Log in to the management console.
  - b. Under **Computing**, click **Elastic Cloud Server**.
  - c. Stop the original Windows ECS, switch to the page providing details about the ECS, and click the **Disks** tab.

#### NOTE

Do not forcibly stop the Windows ECS. Otherwise, password reset may fail.

- d. Locate the row containing the system disk to be detached and click **Detach** to detach the system disk from the ECS.
2. Attach the system disk to the temporary ECS.
    - a. On the page providing details about the temporary ECS, click the **Disks** tab.
    - b. Click **Attach Disk**. In the displayed dialog box, select the system disk detached in step [1.d](#) and attach it to the temporary ECS.

- c. Remotely log in to the temporary ECS.
- d. Run the following command to view the directory of the system disk detached from the original Windows ECS now attached to the temporary ECS:

```
fdisk -l
```

- e. Run the following command to mount the file system of the detached system disk to the temporary ECS:

```
mount -t ntfs-3g /dev/Result obtained in step 2.d /mnt/
```

For example, if the result obtained in step 2.d is **xvde2**, run the following command:

```
mount -t ntfs-3g /dev/xvde2 /mnt/
```

If the following error information is displayed after the preceding command is executed, the NTFS file systems may be inconsistent. In such a case, rectify the file system inconsistency.

```
The disk contains an unclean file system (0, 0).
Metadata kept in Windows cache, refused to mount.
Failed to mount '/dev/xvde2': Operation not permitted
The NTFS partition is in an unsafe state. Please resume and shutdown
Windows fully (no hibernation or fast restarting), or mount the volume
read-only with the 'ro' mount option.
```

Back up the disk data, run the following command to rectify the NTFS file system inconsistency, and attach the system disk:

```
ntfsfix /dev/Result obtained in step 2.d
```

For example, if the result obtained in step 2.d is **xvde2**, run the following command:

```
ntfsfix /dev/xvde2
```

3. Change the password of the specified user and clear the original password.

- a. Run the following command to back up the SAM file:

```
cp /mnt/Windows/System32/config/SAM /mnt/Windows/System32/
config/SAM.bak
```

- b. Run the following command to change the password of the specified user:

```
chntpw -u Administrator /mnt/Windows/System32/config/SAM
```

- c. Enter **1**, **q**, and **y** as prompted, and press **Enter**.

The password has been reset if the following information is displayed:

```
Select: [q] > 1
Password cleared!
Select: [q] > q
Hives that have changed:
#Name
0<SAM>
Write hive files? (y/n) [n] : y
0<SAM> - OK
```

4. Stop the temporary ECS, detach the system disk, and attach the system disk to the original Windows ECS.
  - a. Stop the temporary ECS, switch to the page providing details about the ECS, and click the **Disks** tab.
  - b. Click **Detach** to detach the data disk temporarily attached in step 2.b.

- c. On the page providing details about the original Windows ECS, click the **Disks** tab.
  - d. Click **Attach Disk**. In the displayed dialog box, select the data disk detached in step 4.b and device name `/dev/sda`.
5. Start the original Windows ECS and set a new login password.
- a. Click **Start** to start the original Windows ECS. After the status becomes **Running**, click **Remote Login** in the **Operation** column.
  - b. Click **Start**. Enter **CMD** in the search box and press **Enter**.
  - c. Run the following command to change the password (the new password must meet the requirements described in [Table 10-5](#)):  
**net user Administrator New password**

**Table 10-5** Password complexity requirements

| Parameter | Requirement                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Password  | <ul style="list-style-type: none"><li>• Consists of 8 to 26 characters. For specific requirements on the password length, see the information displayed on the console.</li><li>• Contains at least three of the following character types:<ul style="list-style-type: none"><li>- Uppercase letters</li><li>- Lowercase letters</li><li>- Digits</li><li>- Special characters for Windows: <code>\$!@%-_=[ ]:./,?</code></li><li>- Special characters for Linux: <code>!@%-_=[ ]:/^,{}?</code></li></ul></li><li>• Cannot contain the username or the username spelled backwards.</li><li>• Cannot contain more than two consecutive characters in the same sequence as they appear in the username. (This requirement applies only to Windows ECSs.)</li><li>• Cannot start with a slash (/) for Windows ECSs.</li></ul> |

## 10.1.5 Resetting the Password for Logging In to a Linux ECS

### Scenarios

If the password reset plug-in is not installed on a Linux ECS or a password reset does not take effect, you can reset the password following the instructions provided in this section.

For details about the operations performed on Windows ECSs, see [Resetting the Password for Logging In to a Windows ECS](#).

 NOTE

- If your ECS has password reset plug-in installed, follow the instructions provided in [Resetting the Password for Logging In to an ECS on the Management Console](#) to obtain a new password.
- If your ECS was created using a public image, the password reset plug-in was installed on the ECS by default. To reset the password, see [Resetting the Password for Logging In to an ECS on the Management Console](#).
- Before you perform operations in this section, we recommend you to back up the system disk to prevent data loss.

## Prerequisites

- A temporary Linux ECS which locates in the same AZ as the target ECS is available.
- You have bound an EIP to the temporary ECS.

## Procedure

1. Download the script for resetting the password and upload the script to the temporary ECS.  
Contact the administrator to obtain the password reset script. Use a connection tool, such as WinSCP, to upload the obtained **changepasswd.sh** script to the temporary ECS.  
To download WinSCP, log in at <https://winscp.net/>.
2. Stop the original Linux ECS, detach the system disk from it, and attach the system disk to the temporary ECS.
  - a. Log in to the management console.
  - b. Under **Computing**, click **Elastic Cloud Server**.
  - c. Stop the original ECS, switch to the page providing details about the ECS, and click the **Disks** tab.

 NOTE

Do not forcibly stop the original ECS. Otherwise, password reset may fail.

- d. Locate the row containing the system disk to be detached and click **Detach** to detach the system disk from the ECS.
3. Attach the system disk to the temporary ECS.
    - a. On the page providing details about the temporary ECS, click the **Disks** tab.
    - b. Click **Attach Disk**. In the displayed dialog box, select the system disk detached in step [2.d](#) and attach it to the temporary ECS.
  4. Log in to the temporary ECS remotely and reset the password.
    - a. Locate the row containing the temporary ECS and click **Remote Login** in the **Operation** column.
    - b. Run the following command to view the directory of the system disk detached from the original Linux ECS now attached to the temporary ECS:  
**fdisk -l**

**Figure 10-1** Viewing the directory of the system disk

```
root@ecs-~:~# fdisk -l
Disk /dev/vda: 40 GiB, 42949672960 bytes, 83886080 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0x43591807

Device Boot Start End Sectors Size Id Type
/dev/vda1 * 2048 83884031 83881984 40G 83 Linux

Disk /dev/vdb: 40 GiB, 42949672960 bytes, 83886080 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0x5e9a7bb5

Device Boot Start End Sectors Size Id Type
/dev/vdb1 * 2048 83886079 83884032 40G 83 Linux
```

- c. Run the following commands in the directory where the **changepasswd.sh** script is stored to run the script for resetting the password:

```
chmod +x changepasswd.sh
```

```
./changepasswd.sh
```

When you run the password reset script, if the system displays a message indicating that there is no command related to logical volume manager (LVM), such as the message "no lvs command", install an LVM tool on the temporary ECS. The LVM2 tool is recommended, which can be installed by running the **yum install lvm2** command.

 **NOTE**

If the original ECS and the temporary ECS both run CentOS 7, a mount failure may occur during script execution. To resolve this issue, replace **mount \$dev \$mountPath** with **mount -o nouuid \$dev \$mountPath** in the script.

- d. Enter the new password and the directory obtained in step 4.b as prompted.

If the following information is displayed, the password has been changed:  
set password success.

5. (Optional) Enable remote root login for non-root users.

```
vi /etc/ssh/sshd_config
```

Modify the following settings:

- Change **PasswordAuthentication no** to **PasswordAuthentication yes**.  
Alternatively, uncomment **PasswordAuthentication yes**.
- Change **PermitRootLogin no** to **PermitRootLogin yes**.  
Alternatively, uncomment **PermitRootLogin yes**.
- Change the value of **AllowUsers** to **root**.

Search for **AllowUsers** in the file. If **AllowUsers** is missing, add **AllowUsers root** at the end of the file.

6. Stop the temporary ECS, detach the system disk, attach the system disk to the original Linux ECS, and restart the original Linux ECS.
  - a. Stop the temporary ECS, switch to the page providing details about the ECS, and click the **Disks** tab.
  - b. Click **Detach** to detach the data disk attached in step 2.
  - c. On the page providing details about the original Linux ECS, click the **Disks** tab.
  - d. Click **Attach Disk**. In the displayed dialog box, select the data disk detached in 6.b.
7. Restart the original Linux ECS.

## 10.2 One-Click Password Reset Plug-in

### 10.2.1 Installing the One-Click Password Reset Plug-in on an ECS

You can reset the password for logging in to an ECS with just a few clicks if you forgot the password or the password expires.

After you have created an ECS, it is a good practice to log in to it and install the password reset plug-in.

#### NOTE

One-click password reset is available only in Chinese regions.

The password reset plug-in has been installed on the ECSs created using a public image by default. To check whether the plug-in has been installed, see [Step 1](#).

#### Notes

1. The password reset plug-in is not installed by default. You can determine whether to install it.
2. After the installation, do not uninstall the plug-in by yourself. Otherwise, the ECS password cannot be reset.
3. After you reinstall or change the OS of an ECS, the one-click password reset function will become invalid. If you want to continue using this function, reinstall the password reset plug-in.
4. After you replace the system disk of an ECS, the one-click password reset function will become invalid. If you want to continue using this function, reinstall the password reset plug-in.
5. The password reset plug-in cannot be installed on a CoreOS ECS.
6. To reset the password, the one-click password reset plug-in must be installed before the ECS password is lost or expires.
7. The one-click password reset plug-in can be installed only after an EIP is bound to the ECS.



## Prerequisites

- The available space in drive C of a Windows ECS is greater than 300 MB, and data can be written to it.
- The available space in the root directory of a Linux ECS is greater than 300 MB, and data can be written to it.
- One-click password reset can be used on the ECSs created using SUSE 11 SP4 only if their memory capacity is greater than or equal to 4 GiB.
- DHCP is enabled in the VPC to which the ECS belongs.
- The ECS network connectivity is normal.
- Set the NIC to DHCP so that the ECS can dynamically obtain an IP address.
- The ECS security group rule in the outbound direction meets the following requirements:
  - **Protocol: TCP**
  - **Port Range: 80**
  - **Remote End: 169.254.0.0/16**

If you use the default security group rules for the outbound direction, the preceding requirements are met, and the ECS can be initialized. The default security group rules for the outbound direction are as follows:

- **Protocol: ANY**
- **Port Range: ANY**
- **Remote End: 0.0.0.0/16**

## Installing the Password Reset Plug-in on a Linux ECS

**Step 1** Use either of the following methods to check whether the password reset plug-in has been installed on the ECS:

Method 1: Use the management console for query.

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Locate the row containing the target ECS, click **More** in the **Operation** column, and select **Reset Password** from the drop-down list.
  - If a dialog box is displayed, asking you to enter the new password, the password reset plug-in has been installed. No further action is required.
  - If a dialog box is displayed, asking you to download a password reset script, the password reset plug-in has not been installed. Then, install it.

Method 2: Use the ECS for query.

1. Log in to the ECS as user **root**.
2. Run the following command to check whether CloudResetPwdAgent has been installed:

```
ls -lh /Cloud*
```

**Figure 10-2** Checking whether the plug-in has been installed

```
[root@ecs-test ~]# ls -lh /Cloud*
total 20K
drwx----- 2 root root 4.0K Jun 13 14:13 bin
drwxr-xr-x 2 root root 4.0K Jun 13 11:53 conf
drwx----- 3 root root 4.0K Jun 13 11:53 depend
drwx----- 2 root root 4.0K Jun 13 11:53 lib
drwx----- 2 root root 4.0K Jun 13 14:13 logs
[root@ecs-test ~]#
[root@ecs-test ~]#
```

Check whether the obtained information is similar to that shown in [Figure 10-2](#).

- If yes, the plug-in has been installed.
- If no, the plug-in has not been installed. Then, install it.

**Step 2** Download the one-click password reset plug-in **CloudResetPwdAgent.zip**.

The plug-in is stored in the OBS bucket. Contact the administrator to obtain the path of the OBS bucket.

For example, the download path is [http://xxx-cloud-reset-pwd.obs.com/linux/64/reset\\_pwd\\_agent/CloudResetPwdAgent.zip](http://xxx-cloud-reset-pwd.obs.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip).

Then, run the following command:

```
wget http://xxx-cloud-reset-pwd.obs.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip
```

**Step 3** Run the following command to decompress **CloudResetPwdAgent.zip**:

There is no special requirement for the directory that stores the decompressed **CloudResetPwdAgent.zip**. Use any directory.

```
unzip -o -d Decompressed directory CloudResetPwdAgent.zip
```

An example is provided as follows:

If the plug-in is decompressed to **/home/linux/test**, run the following command:

```
unzip -o -d /home/linux/test CloudResetPwdAgent.zip
```

**Step 4** Install the one-click password reset plug-in.

1. Run the following command to open the **CloudResetPwdAgent.Linux** file:  

```
cd {Plug-in decompressed directory}/CloudResetPwdAgent/CloudResetPwdAgent.Linux
```

For example:

If the plug-in is decompressed to **/home/linux/test**, run the following command:

```
cd /home/linux/test/CloudResetPwdAgent/CloudResetPwdAgent.Linux
```

2. Run the following command to add the execute permission for the **setup.sh** file:  

```
chmod +x setup.sh
```

3. Run the following command to install the plug-in:

```
sudo sh setup.sh
```

If "cloudResetPwdAgent install successfully." is displayed and "Failed to start service cloudResetPwdAgent" is not displayed, the installation is successful.

#### NOTE

- You can also check whether the password reset plug-in has been installed using the methods provided in [Step 1](#).
- If the installation failed, check whether the installation environment meets requirements and install the plug-in again.

- Step 5** Modify the file permission of the password reset plug-in.

```
chmod 700 /CloudrResetPwdAgent/bin/cloudResetPwdAgent.script
```

```
chmod 700 /CloudrResetPwdAgent/bin/wrapper
```

```
chmod 600 /CloudrResetPwdAgent/lib/*
```

----End

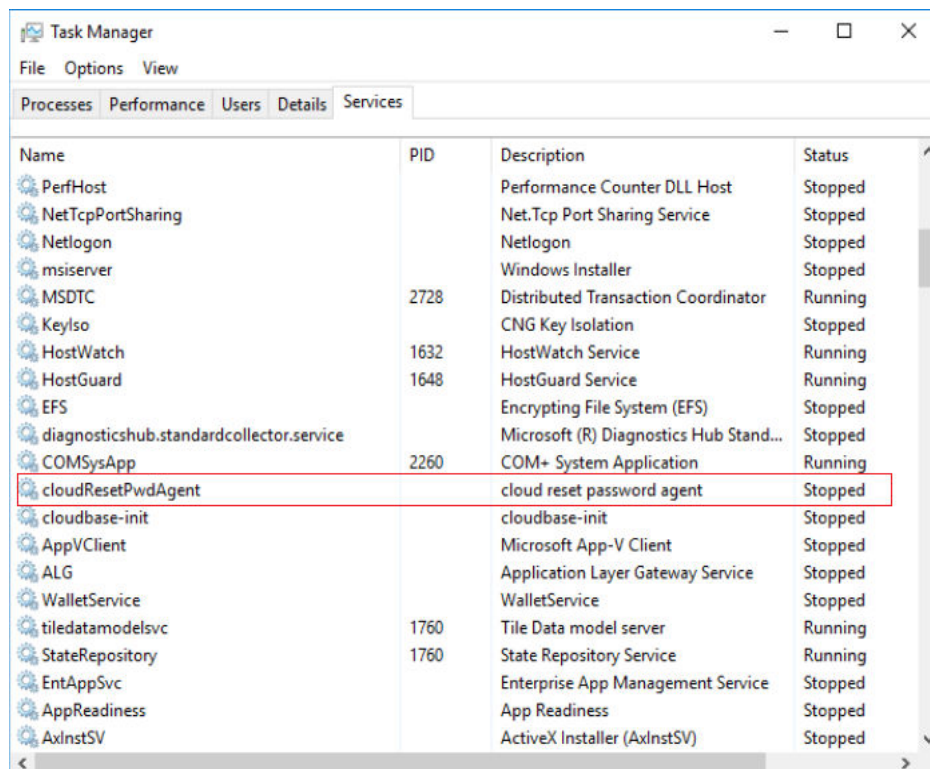
## Installing the Password Reset Plug-in on a Windows ECS

- Step 1** Log in to the ECS.

- Step 2** Check whether the password reset plug-in CloudResetPwdAgent has been installed on the ECS. To check this, perform the following operations:

Start the **Task Manager** and check whether **cloudResetPwdAgent** is displayed on the **Services** tab. As shown in the [Figure 10-3](#), the password reset plug-in has been installed on the ECS.

**Figure 10-3** Successful plug-in installation



- If yes, no further action is required.
- If no, go to [Step 3](#).

**Step 3** Download package **CloudResetPwdAgent.zip**.

There is no special requirement for the directory that stores **CloudResetPwdAgent.zip**. Use any directory.

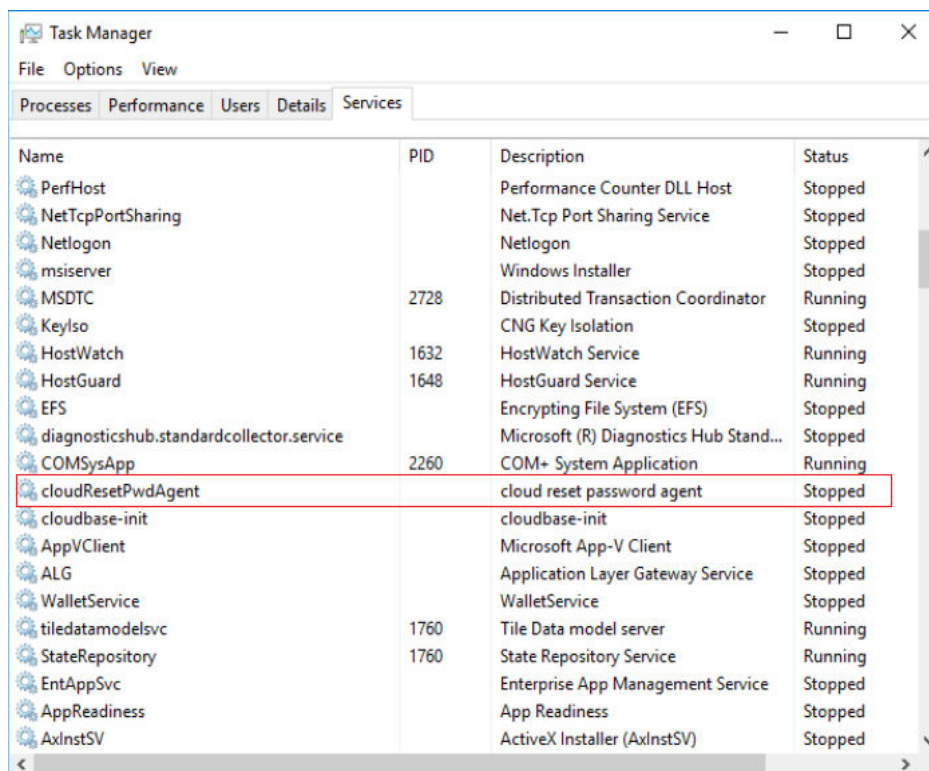
This package is stored in an OBS bucket. Contact the administrator to obtain the path of the OBS bucket.

**Step 4** Decompress **CloudResetPwdAgent.zip**.

There is no special requirement for the directory that stores the decompressed **CloudResetPwdAgent.zip**. Use any directory.

**Step 5** Install the plug-in.

1. Double-click **setup.bat** in **CloudResetPwdAgent.Windows**.  
The password reset plug-in starts to be installed.
2. View the **Task Manager** and check whether the installation was successful.  
If **cloudResetPwdAgent** is displayed in the **Task Manager**, as shown in [Figure 10-4](#), the installation was successful. Otherwise, the installation failed.

**Figure 10-4** Successful plug-in installation**NOTE**

If the installation failed, check whether the installation environment meets requirements and install the plug-in again.

----End

## Follow-up Procedure

- After the one-click password reset plug-in is installed, you can add it to the startup items to ensure that the plug-in automatically starts upon ECS startup. For details, see "What Do I Do If the One-Click Password Resetting Plug-In Failed to Start?" in *Image Management Service User Guide*.
- After installing the one-click password reset plug-in, do not delete the `CloudResetPwdAgent` process. Otherwise, one-click password reset will not be available.
- After the one-click password reset plug-ins are installed, you can view `Dwrapper.key`, which uniquely identifies the JVM process.
- The one-click password reset plug-in has been upgraded automatically. New ECSs work in PIPE mode by default, preventing their plug-ins from using service ports. Existing ECSs still work in AUTO mode, in which the plug-in randomly selects idle ports with the smallest port numbers ranging from 31000 to 32999. If the used port conflicts with a service port, see [Why Am I Seeing the Message Indicating That the Port Is Used by a One-Click Password Reset Plug-in?](#) for troubleshooting.

## Uninstalling the Plug-in

If you do not need the password reset function anymore, perform the following operations to uninstall the plug-in:

- Linux ECS
  - a. Log in to the ECS.
  - b. Run the following commands to switch to the `bin` directory and delete `cloudResetPwdAgent`:

```
cd /CloudrResetPwdAgent/bin
sudo ./cloudResetPwdAgent.script remove
```
  - c. Run the following command to delete the plug-in:

```
sudo rm -rf /CloudrResetPwdAgent
```

Check whether `CloudResetPwdUpdateAgent` exists. If it exists, run the following command to delete it:

```
sudo rm -rf /CloudResetPwdUpdateAgent
```
- Windows ECS
  - a. Uninstall and delete `CloudResetPwdAgent`.
    - i. Switch to the `C:\CloudResetPwdAgent\bin` folder.
    - ii. Double-click `UninstallApp-NT.bat`.
    - iii. Delete the file in `C:\CloudResetPwdAgent`.
  - b. (Optional) Uninstall and delete `CloudResetPwdUpdateAgent`.

The plug-in varies depending on the Windows version. Check whether `CloudResetPwdUpdateAgent` exists. If it exists, perform the following operations to uninstall and delete it. If it does not exist, skip this step.

    - i. Go to the `C:\CloudResetPwdUpdateAgent` folder.
    - ii. Double-click `UninstallApp-NT.bat`.

- iii. Delete the file in **C:\CloudResetPwdUpdateAgent**.  
If the deletion fails, delete **CloudResetPwdUpdateAgent** from Task Manager first and then delete the file in **C:\CloudResetPwdUpdateAgent**.

## 10.3 Key Pairs

### 10.3.1 Application Scenarios for Using Key Pairs

#### Key Pairs

Key pairs are a set of security credentials for identity authentication when you remotely log in to ECSs.

A key pair consists of a public key and a private key. Key Pair Service (KPS) stores the public key and you store the private key. If you have imported a public key into a Linux ECS, you can use the corresponding private key to log in to the ECS without a password. You do not need to worry about password interception, cracking, or leakage.

#### Scenarios

When purchasing an ECS, you are advised to select the key pair login mode. For Windows ECSs, key pairs are required to decrypt the passwords so that you can use the decrypted password to log in.

- Logging in to a Linux ECS  
You can directly use a key pair to log in.
  - When you are creating the ECS, select the key pair login mode. For details, see "Set Login Mode" in [Step 3: Configure Advanced Settings](#).
  - After the ECS is created, bind a key pair to the ECS by referring to "Binding a Key Pair" in the *Data Encryption Workshop User Guide*.
- Logging in to a Windows ECS  
You can use the key pair to obtain a password for login. The password is randomly generated and is more secure.  
For details, see [Obtaining the Password for Logging In to a Windows ECS](#).

#### Creating a Key Pair

You can create a key pair or use an existing one for remote login authentication.

- Creating a key pair  
You can create a key pair using either of the following method:
  - Follow the instructions in [\(Recommended\) Creating a Key Pair on the Management Console](#). The public key is automatically stored in the system, and the private key is stored locally.
  - Follow the instructions in [Creating a Key Pair Using PuTTYgen](#). Both the public and private keys are stored locally.

After the key pair is created, import the key pair following the instructions provided in [Importing a Key Pair](#) so that you can use it.

- Using an existing key pair

If an existing key pair (created using PuTTYgen, for example) is available, you can import the public key by referring to [Importing a Key Pair](#) on the management console to let the system maintain your public key.

#### NOTE

If the public key of the existing key pair is stored by clicking **Save public key** on PuTTY Key Generator, the public key cannot be imported to the management console.

If you want to use this existing key pair for remote login, see [Why Does a Key Pair Created Using puttygen.exe Fail to Be Imported on the Management Console?](#)

## Constraints

- Key pairs can be used to remotely log in to Linux ECSs only.
- SSH-2 key pairs created on the console support only the RSA-2048 cryptographic algorithms.
- Key pairs can be used only for ECSs in the same region.
- Imported key pairs support the following cryptographic algorithms:
  - RSA-1024
  - RSA-2048
  - RSA-4096
- Store your private key in a secure place because you need to use it to prove your identity when logging in to your ECS. The private key can be downloaded once only.

## 10.3.2 (Recommended) Creating a Key Pair on the Management Console

### Scenarios

You can use the management console to create a key pair. ECS stores the public key and you store the private key.

### Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. In the navigation pane on the left, choose **Key Pair**.
4. On the displayed page, click **Create Key Pair**.
5. Enter a key pair name.  
A key pair name consists of two parts: KeyPair and four random digits (KeyPair-xxxx).
6. Click **OK**.
7. Manually or automatically download a .pem private key file with the name that you specify as the key name. Store it in a secure place and click **OK**.

 NOTE

This is the only chance for you to save the private key file. Keep it secure. You'll need to provide the key pair name when you create an ECS, and the corresponding private key each time you connect to the ECS through SSH.

## 10.3.3 Creating a Key Pair Using PuTTYgen

### Scenarios

You can use PuTTYgen to create a key pair and store the public key and private key locally.

 NOTE

Key pairs created using puttygen.exe must be imported by referring to [Importing a Key Pair](#) before they are used.

### Procedure

1. Download and install PuTTY and PuTTYgen.  
<https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html>


 NOTE

PuTTYgen is a key generator, which is used to create a key pair that consists of a public key and a private key for PuTTY.

2. Obtain the public and private keys.
  - a. Double-click **puttygen.exe** to open **PuTTY Key Generator**.




**Figure 10-5** PuTTY Key Generator



- b. Click **Generate**.

The key generator automatically generates a key pair that consists of a public key and a private key. The content shown in the red box in [Figure 10-6](#) is the public key.

**Figure 10-6** Generating the public and private keys

3. Copy the public key to a .txt file and save it to a local directory.

**NOTE**

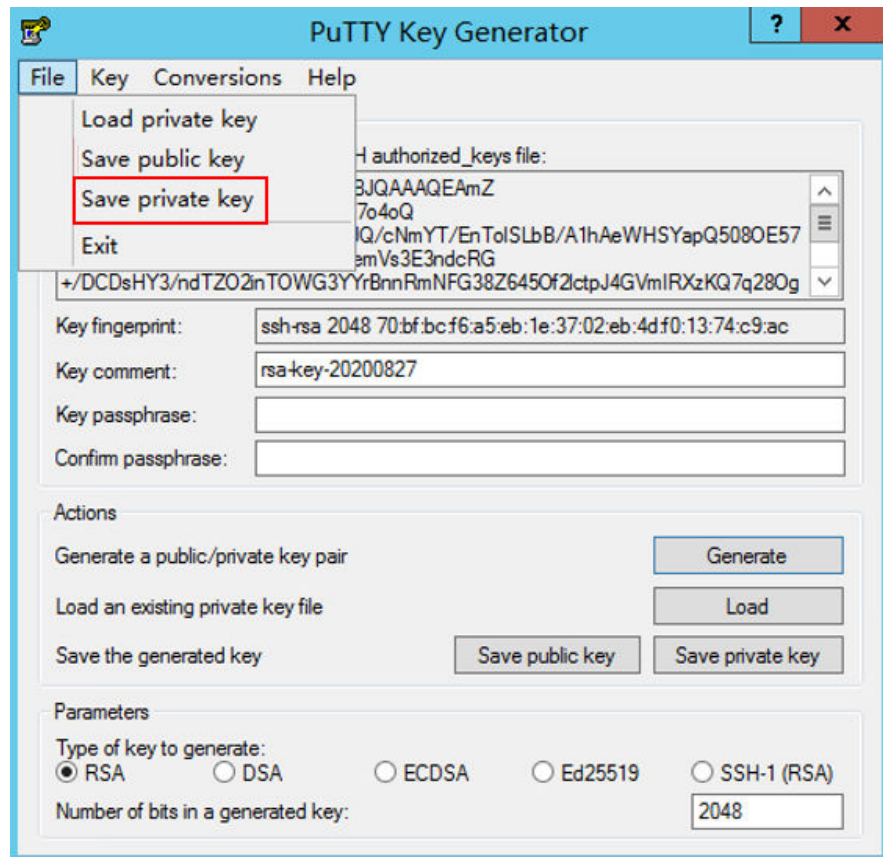
Do not save the public key by clicking **Save public key** because this operation will change the format of the public key content and cause the public key to fail to be imported to the management console.

4. Save the private key and keep it secure. The private key can be downloaded only once.

The format in which to save your private key file varies depending on application scenarios.

- When using PuTTY to log in to a Linux ECS:
  - Save the private key file in the **.ppk** format.
  - i. On the **PuTTY Key Generator** page, choose **File > Save private key**.

Figure 10-7 Saving a private key



- ii. Save the converted private key file, such as **kp-123.ppk**, locally.
- When using Xshell to log in to a Linux ECS or obtaining the password for logging in to a Windows ECS:


Save the private key file in the **.pem** format.

- i. Choose **Conversions > Export OpenSSH key**.

#### NOTE

If you use this private file to obtain the password for logging in to a Windows ECS, do not specify **Key passphrase** for **Export OpenSSH key** so that you can obtain the password successfully.

Figure 10-8 Saving a private key



- ii. Save the private key, for example, **kp-123.pem**, locally.
5. After you have saved the key pair, import your public key to the ECS by referring to [Importing a Key Pair](#).

## 10.3.4 Importing a Key Pair

### Scenarios

You need to import a key pair in either of the following scenarios:

- Create a key pair using PuTTYgen and import the public key to the ECS.
- Import the public key of an existing key pair to the ECS to let the system maintain your public key.

#### NOTE

If the public key of the existing key pair is stored by clicking **Save public key** on PuTTY Key Generator, the public key cannot be imported to the management console.

If you want to use this existing key pair for remote login, see [Why Does a Key Pair Created Using puttygen.exe Fail to Be Imported on the Management Console?](#)

### Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.

3. In the navigation pane on the left, choose **Key Pair**.
4. On the **Key Pair Service** page, click **Import Key Pair**.
5. Use either of the following methods to import the key pair:
  - Selecting a file
    - i. In the **Import Key Pair** dialog box of the management console, click **Select File** and select the locally stored public key file (for example, the .txt file saved in 3 in [Creating a Key Pair Using PuTTYgen](#)).
  - Copying the public key content
    - i. Copy the public key content from the locally stored .txt file into the **Public Key Content** text box.
    - ii. Click **OK**.

 **NOTE**

Make sure that the file to be imported is a public key file.

- ii. Click **OK**.  
After the public key is imported, you can change its name.

## Helpful Links

- [What Should I Do If a Key Pair Cannot Be Imported?](#)
- [Why Does a Key Pair Created Using puttygen.exe Fail to Be Imported on the Management Console?](#)

## 10.3.5 Obtaining and Deleting the Password of a Windows ECS

### 10.3.5.1 Obtaining the Password for Logging In to a Windows ECS

#### Scenarios

Password authentication is required to log in to a Windows ECS. You must use the key file used when you created the ECS to obtain the administrator password generated during ECS creation. The administrator user is **Administrator** or the user configured using Cloudbase-Init. This password is randomly generated, offering high security.

You can obtain the initial password for logging in to a Windows ECS through the management console or APIs. For details, see this section.

#### Obtaining the Password Through the Management Console

1. Obtain the private key file (.pem file) used when you created the ECS.
2. Log in to the management console.
3. Under **Computing**, click **Elastic Cloud Server**.
4. On the **Elastic Cloud Server** page, select the target ECS.
5. In the **Operation** column, click **More** and select **Get Password**.

6. Use either of the following methods to obtain the password through the key file:
  - Click **Select File** and upload the key file from a local directory.
  - Copy the key file content to the text field.
7. Click **Get Password** to obtain a random password.

## Obtaining the Password Through APIs

1. Obtain the private key file (.pem file) used when you created the ECS.
2. Set up the API calling environment.
3. Call APIs. For details, see "Before You Start" in *Elastic Cloud Server API Reference*.
4. Obtain the ciphertext password.

Call the password obtaining APIs to obtain the ciphertext password of the public key encrypted using RSA. The API URI is in the format "GET /v2/{tenant\_id}/servers/{server\_id}/os-server-password".

### NOTE

For details, see "Obtaining the Password for Logging In to an ECS" in the *ECS API Reference*.

5. Decrypt the ciphertext password.

Use the private key file used when you created the ECS to decrypt the ciphertext password obtained in step 4.

  - a. Run the following command to convert the ciphertext password format to ".key -nocrypt" using OpenSSL:  
**openssl pkcs8 -topk8 -inform PEM -outform DER -in rsa\_pem.key -out pkcs8\_der.key -nocrypt**
  - b. Invoke the Java class library **org.bouncycastle.jce.provider.BouncyCastleProvider** and use the key file to edit the code decryption ciphertext.

### 10.3.5.2 Deleting the Initial Password for Logging In to a Windows ECS

#### Scenarios

After you obtain the initial password, it is a good practice to delete it to ensure system security.

Deleting the initial password does not affect ECS operation or login. Once deleted, the password cannot be retrieved. Before you delete a password, it is a good practice to record it.

#### Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. On the **Elastic Cloud Server** page, select the target ECS.
4. In the **Operation** column, click **More** and select **Delete Password**.

The system displays a message, asking you whether you want to delete the password.

5. Click **OK** to delete the password.

# 11 Permissions Management

---

## 11.1 Creating a User and Granting ECS Permissions

Use IAM to implement fine-grained permissions control over your ECSs. With IAM, you can:

- Create IAM users for personnel based on your enterprise's organizational structure. Each IAM user will have their own security credentials for accessing ECS resources.
- Grant only the permissions required for users to perform a specific task.
- Delegate access to other accounts or cloud services for efficient O&M.

If your account does not need individual IAM users, skip this section.

This section describes the procedure for granting permissions (see [Process Flow](#)).

### Prerequisites

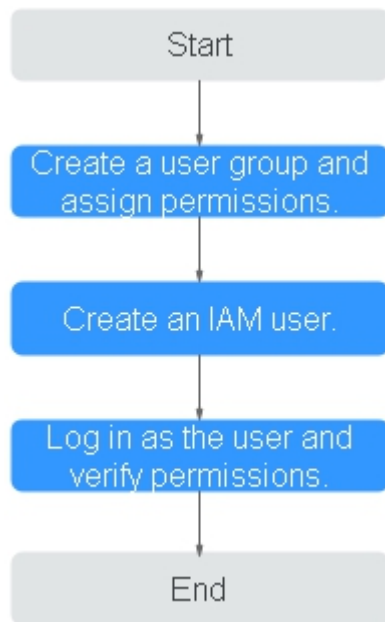
Before assigning permissions to user groups, you should learn about system policies supported by ECS and select the policies based on service requirements.

For more information about system policies supported by ECS, see [Permissions Management](#).



## Process Flow

**Figure 11-1** Process for granting ECS permissions



1. Create a user group and assign permissions.  
Create a user group on the IAM console and click **Authorize** in the **Operation** column to assign the *ECS ReadOnlyAccess* permissions to the group.
2. Create a user and add the user to the user group.  
Create a user on the IAM console and add it to the user group created in 1 by choosing **Authorize** in the **Operation** column.
3. Log in to the management console as the created user.  
In the authorized region, perform the following operations:
  - Choose **Compute > Elastic Cloud Server** in **Service List**. On the ECS console, click **Create ECS**. If the creation attempt failed, the **ECSReadOnlyAccess** policy has already taken effect.
  - Choose any service other than ECS in **Service List**. If a message appears indicating that you have insufficient permissions to access the service, the **ECSReadOnlyAccess** policy has already taken effect.

## 11.2 ECS Custom Policies

Custom policies can be created to supplement the system-defined policies of ECS. For the actions that can be added to custom policies, see "Permissions Policies and Supported Actions" in "Elastic Cloud Server API Reference".

You can create custom policies in either of the following ways:

- Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy syntax.
- JSON: Edit JSON policies from scratch or based on an existing policy.

For details, see "Creating a Custom Policy" in *Identity and Access Management User Guide*. The following provides examples of common ECS custom policies.

## Example Custom Policies

- Example 1: Allowing users to stop and delete multiple ECSs at a time

```
{
 "Version": "1.1",
 "Statement": [
 {
 "Effect": "Allow",
 "Action": [
 "ecs:servers:stop",
 "ecs:servers:get"
]
 }
]
}
```

- Example 2: Denying ECS deletion

A policy with only "Deny" permissions must be used in conjunction with other policies to take effect. If the permissions assigned to a user contain both "Allow" and "Deny", the "Deny" permissions take precedence over the "Allow" permissions.

The following method can be used if you need to assign permissions of the **ECSFullAccess** policy to a user but you want to prevent the user from deleting ECSs. Create a custom policy for denying ECS deletion, and attach both policies to the group to which the user belongs. Then, the user can perform all operations on ECSs except deleting ECSs. The following is an example of a deny policy:

```
{
 "Version": "1.1",
 "Statement": [
 {
 "Effect": "Deny",
 "Action": [
 "ecs:cloudServers:delete"
]
 }
]
}
```

# 12 Resources

---


## 12.1 Quota Adjustment

### What Is Quota?

Quotas can limit the number or amount of resources available to users, such as the maximum number of ECSs or EVS disks that can be created.

If the existing resource quota cannot meet your service requirements, you can apply for a higher quota.

### How Do I View My Quotas?

1. Log in to the management console.
2. In the upper right corner of the page, click  .  
The **Service Quota** page is displayed.
3. View the used and total quota of each type of resources on the displayed page.  
If a quota cannot meet service requirements, apply for a higher quota.

### How Do I Apply for a Higher Quota?

The system does not support online quota adjustment. If you need to adjust a quota, contact the operations administrator.

Before contacting the operations administrator, make sure that the following information has been obtained:

- Account name, which can be obtained by performing the following operations:  
Log in to the management console using the cloud account, click the username in the upper right corner, select **My Credentials** from the drop-down list, and obtain the account name on the **My Credentials** page.
- Quota information, which includes service name, quota type, and required quota

# 13 Monitoring

---

## 13.1 Monitoring ECSs

Monitoring is key for ensuring ECS performance, reliability, and availability. Using monitored data, you can determine ECS resource utilization. The cloud platform provides Cloud Eye to help you obtain the running statuses of your ECSs. You can use Cloud Eye to automatically monitor ECSs in real time and manage alarms and notifications to keep track of ECS performance metrics.

This section covers the following content:

- Viewing basic ECS metrics
- Viewing OS metrics (Agent installed on ECS)
- Customizing ECS alarm rules
- Viewing ECS running statuses for routine monitoring

## 13.2 Basic ECS Metrics

### Description

This section describes basic monitoring metrics reported by ECS to Cloud Eye. You can use Cloud Eye to view these metrics and alarms generated for ECSs.

### Namespace

SYS.ECS

### Basic ECS Metrics

Basic ECS metrics vary depending on ECS OSs and types. For details, see [Table 13-1](#). ✓ indicates that the metric is supported, and x indicates that the metric is not supported.

**Table 13-1** Basic ECS metrics

| Metric                | Windows   | Linux         |
|-----------------------|-----------|---------------|
| CPU Usage             | Supported | Supported     |
| Memory Usage          | Supported | Not supported |
| Disk Usage            | Supported | Not supported |
| Disk Read Bandwidth   | Supported | Supported     |
| Disk Write Bandwidth  | Supported | Supported     |
| Disk Read IOPS        | Supported | Supported     |
| Disk Write IOPS       | Supported | Supported     |
| Inband Incoming Rate  | Supported | Not supported |
| Inband Outgoing Rate  | Supported | Not supported |
| Outband Incoming Rate | Supported | Supported     |
| Outband Outgoing Rate | Supported | Supported     |

**Table 13-2** describes these basic ECS metrics.

The monitoring intervals for the following ECSs with raw monitoring metrics are as follows:

- Xen ECSs: 4 minutes
- KVM ECSs: 5 minutes

**Table 13-2** Basic metric description

| Metric ID | Parameter | Description                                                                                     | Value Range | Monitored Object & Dimension | Monitoring Interval (Raw Metrics and KVM Only) |
|-----------|-----------|-------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------------------------|
| cpu_util  | CPU Usage | CPU usage of an ECS<br>Unit: Percent<br>Formula: CPU usage of an ECS/Number of vCPUs in the ECS | $\geq 0$    | ECS                          | 5 minutes                                      |

| Metric ID            | Parameter           | Description                                                                                                                                                                                                                                                  | Value Range | Monitored Object & Dimension | Monitoring Interval (Raw Metrics and KVM Only) |
|----------------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------------------------|
| mem_util             | Memory Usage        | Memory usage of an ECS<br>This metric is unavailable if the image has no UVP VMTools installed.<br>Unit: Percent<br>Formula: Used memory of an ECS/ Total memory of the ECS<br><b>NOTE</b><br>The memory usage of QingTian ECSs cannot be monitored.         | $\geq 0$    | ECS                          | 5 minutes                                      |
| disk_util_inband     | Disk Usage          | Disk usage of an ECS<br>This metric is unavailable if the image has no UVP VMTools installed.<br>Unit: Percent<br>Formula: Used capacity of an ECS-attached disk/Total capacity of the ECS-attached disk                                                     | $\geq 0$    | ECS                          | 5 minutes                                      |
| disk_read_bytes_rate | Disk Read Bandwidth | Number of bytes read from an ECS-attached disk per second<br>Unit: byte/s<br>Formula: Total number of bytes read from an ECS-attached disk/ Monitoring interval<br>$\text{byte\_out} = (\text{rd\_bytes} - \text{last\_rd\_bytes}) / \text{Time difference}$ | $\geq 0$    | ECS                          | 5 minutes                                      |

| Metric ID                | Parameter            | Description                                                                                                                                                                                                                                        | Value Range | Monitored Object & Dimension | Monitoring Interval (Raw Metrics and KVM Only) |
|--------------------------|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------------------------|
| disk_write_bytes_rate    | Disk Write Bandwidth | Number of bytes written to an ECS-attached disk per second<br>Unit: byte/s<br>Formula: Total number of bytes written to an ECS-attached disk/ Monitoring interval                                                                                  | ≥ 0         | ECS                          | 5 minutes                                      |
| disk_read_requests_rate  | Disk Read IOPS       | Number of read requests sent to an ECS-attached disk per second<br>Unit: request/s<br>Formula: Total number of read requests sent to an ECS-attached disk/ Monitoring interval<br>$req\_out = (rd\_req - last\_rd\_req) / \text{Time difference}$  | ≥ 0         | ECS                          | 5 minutes                                      |
| disk_write_requests_rate | Disk Write IOPS      | Number of write requests sent to an ECS-attached disk per second<br>Unit: request/s<br>Formula: Total number of write requests sent to an ECS-attached disk/ Monitoring interval<br>$req\_in = (wr\_req - last\_wr\_req) / \text{Time difference}$ | ≥ 0         | ECS                          | 5 minutes                                      |

| Metric ID                             | Parameter             | Description                                                                                                                                                                                                          | Value Range | Monitored Object & Dimension | Monitoring Interval (Raw Metrics and KVM Only) |
|---------------------------------------|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------------------------|
| network_incoming_bytes_rate_inband    | Inband Incoming Rate  | Number of incoming bytes on an ECS per second<br>Unit: byte/s<br>Formula: Total number of inband incoming bytes on an ECS/Monitoring interval                                                                        | $\geq 0$    | ECS                          | 5 minutes                                      |
| network_outgoing_bytes_rate_inband    | Inband Outgoing Rate  | Number of outgoing bytes on an ECS per second<br>Unit: byte/s<br>Formula: Total number of inband outgoing bytes on an ECS/Monitoring interval                                                                        | $\geq 0$    | ECS                          | 5 minutes                                      |
| network_incoming_bytes_aggregate_rate | Outband Incoming Rate | Number of incoming bytes on an ECS per second on the hypervisor<br>Unit: byte/s<br>Formula: Total number of outband incoming bytes on an ECS/Monitoring interval<br>This metric is unavailable if SR-IOV is enabled. | $\geq 0$    | ECS                          | 5 minutes                                      |



| Metric ID                              | Parameter             | Description                                                                                                                                                                                                          | Value Range | Monitored Object & Dimension | Monitoring Interval (Raw Metrics and KVM Only) |
|----------------------------------------|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------------------------|
| network_outgoing_bytes_agg_regate_rate | Outband Outgoing Rate | Number of outgoing bytes on an ECS per second on the hypervisor<br>Unit: byte/s<br>Formula: Total number of outband outgoing bytes on an ECS/Monitoring interval<br>This metric is unavailable if SR-IOV is enabled. | $\geq 0$    | ECS                          | 5 minutes                                      |

## Dimensions

| Key         | Value                 |
|-------------|-----------------------|
| instance_id | Specifies the ECS ID. |

## 13.3 OS Monitoring Metrics Supported by ECSs with the Agent Installed

### Description

OS monitoring provides system-level, proactive, and fine-grained monitoring. It requires the Agent to be installed on the ECSs to be monitored. This section describes OS monitoring metrics reported to Cloud Eye.

OS monitoring supports metrics about CPU, CPU load, memory, disk, disk I/O, file system, NIC, NTP, and TCP.

After the Agent is installed, you can view monitoring metrics of ECSs running different OSs. Monitoring data is collected every 1 minute.

### Namespace

AGT.ECS

## OS Metrics: CPU

Table 13-3 CPU metrics

| Metric         | Parameter              | Description                                                                                                                                                                                                                                                                                                                                               | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|----------------|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| cpu_usage      | (Agent) CPU Usage      | <p>CPU usage of the monitored object</p> <p>Unit: percent</p> <ul style="list-style-type: none"> <li>Linux: Check metric value changes in file / <b>proc/stat</b> in a collection period. Run the <b>top</b> command to check the <b>%Cpu(s)</b> value.</li> <li>Windows: Obtain the metric value using the Windows API <b>GetSystemTimes</b>.</li> </ul> | 0-100       | ECS                          | 1 minute                     |
| cpu_usage_idle | (Agent) Idle CPU Usage | <p>Percentage of time that CPU is idle</p> <p>Unit: percent</p> <ul style="list-style-type: none"> <li>Linux: Check metric value changes in file / <b>proc/stat</b> in a collection period.</li> <li>Windows: Obtain the metric value using the Windows API <b>GetSystemTimes</b>.</li> </ul>                                                             | 0-100       | ECS                          | 1 minute                     |

| Metric           | Parameter                      | Description                                                                                                                                                                                                                                                                                                                                                          | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|------------------|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| cpu_usage_user   | (Agent) User Space CPU Usage   | Percentage of time that the CPU is used by user space<br>Unit: percent <ul style="list-style-type: none"><li>Linux: Check metric value changes in file / <b>proc/stat</b> in a collection period. Run the <b>top</b> command to check the <b>%Cpu(s) us</b> value.</li><li>Windows: Obtain the metric value using the Windows API <b>GetSystemTimes</b>.</li></ul>   | 0-100       | ECS                          | 1 minute                     |
| cpu_usage_system | (Agent) Kernel Space CPU Usage | Percentage of time that the CPU is used by kernel space<br>Unit: percent <ul style="list-style-type: none"><li>Linux: Check metric value changes in file / <b>proc/stat</b> in a collection period. Run the <b>top</b> command to check the <b>%Cpu(s) sy</b> value.</li><li>Windows: Obtain the metric value using the Windows API <b>GetSystemTimes</b>.</li></ul> | 0-100       | ECS                          | 1 minute                     |

| Metric          | Parameter                       | Description                                                                                                                                                                                                                                                                                                                                                                                                               | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-----------------|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| cpu_usage_other | (Agent) Other Process CPU Usage | <p>Percentage of time that the CPU is used by other processes</p> <p>Unit: percent</p> <ul style="list-style-type: none"> <li>Linux: <b>Other Process CPU Usage = 1 - Idle CPU Usage - Kernel Space CPU Usage - User Space CPU Usage</b></li> <li>Windows: <b>Other Process CPU Usage = 1 - Idle CPU Usage - Kernel Space CPU Usage - User Space CPU Usage</b></li> </ul>                                                 | 0-100       | ECS                          | 1 minute                     |
| cpu_usage_nice  | (Agent) Nice Process CPU Usage  | <p>Percentage of time that the CPU is in user mode with low-priority processes which can easily be interrupted by higher-priority processes</p> <p>Unit: percent</p> <ul style="list-style-type: none"> <li>Linux: Check metric value changes in file / <b>proc/stat</b> in a collection period. Run the <b>top</b> command to check the <b>%Cpu(s) ni</b> value.</li> <li>Windows is not supported currently.</li> </ul> | 0-100       | ECS                          | 1 minute                     |

| Metric            | Parameter                           | Description                                                                                                                                                                                                                                                                                                                                  | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-------------------|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| cpu_usage_iowait  | (Agent) iowait Process CPU Usage    | Percentage of time that the CPU is waiting for I/O operations to complete<br>Unit: percent <ul style="list-style-type: none"><li>Linux: Check metric value changes in file / <b>proc/stat</b> in a collection period. Run the <b>top</b> command to check the <b>%Cpu(s) wa</b> value.</li><li>Windows is not supported currently.</li></ul> | 0-100       | ECS                          | 1 minute                     |
| cpu_usage_irq     | (Agent) CPU Interrupt Time          | Percentage of time that the CPU is servicing interrupts<br>Unit: percent <ul style="list-style-type: none"><li>Linux: Check metric value changes in file / <b>proc/stat</b> in a collection period. Run the <b>top</b> command to check the <b>%Cpu(s) hi</b> value.</li><li>Windows is not supported currently.</li></ul>                   | 0-100       | ECS                          | 1 minute                     |
| cpu_usage_softirq | (Agent) CPU Software Interrupt Time | Percentage of time that the CPU is servicing software interrupts<br>Unit: percent <ul style="list-style-type: none"><li>Linux: Check metric value changes in file / <b>proc/stat</b> in a collection period. Run the <b>top</b> command to check the <b>%Cpu(s) si</b> value.</li><li>Windows is not supported currently.</li></ul>          | 0-100       | ECS                          | 1 minute                     |

## OS Metric: CPU Load

Table 13-4 CPU load metrics

| Metric        | Parameter                      | Description                                                                                                                                                                                                              | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|---------------|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| load_averge1  | (Agent) 1-Minute Load Average  | CPU load averaged from the last 1 minute<br>Linux: Obtain the metric value from the number of logic CPUs in <b>load1/</b> in file <b>/proc/loadavg</b> . Run the <b>top</b> command to check the <b>load1</b> value.     | $\geq 0$    | ECS                          | 1 minute                     |
| load_averge5  | (Agent) 5-Minute Load Average  | CPU load averaged from the last 5 minutes<br>Linux: Obtain the metric value from the number of logic CPUs in <b>load5/</b> in file <b>/proc/loadavg</b> . Run the <b>top</b> command to check the <b>load5</b> value.    | $\geq 0$    | ECS                          | 1 minute                     |
| load_averge15 | (Agent) 15-Minute Load Average | CPU load averaged from the last 15 minutes<br>Linux: Obtain the metric value from the number of logic CPUs in <b>load15/</b> in file <b>/proc/loadavg</b> . Run the <b>top</b> command to check the <b>load15</b> value. | $\geq 0$    | ECS                          | 1 minute                     |

 NOTE

The Windows OS does not support the CPU load metrics.

## OS Metric: Memory

Table 13-5 Memory metrics

| Metric        | Parameter                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|---------------|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| mem_available | (Agent) Available Memory | <p>Amount of memory that is available and can be given instantly to processes</p> <p>Unit: GB</p> <ul style="list-style-type: none"> <li>Linux: Obtain the metric value from <b>/proc/meminfo</b>. <ul style="list-style-type: none"> <li>If <b>MemAvailable</b> is displayed in <b>/proc/meminfo</b>, obtain the value.</li> <li>If <b>MemAvailable</b> is not displayed in <b>/proc/meminfo</b>, <b>MemAvailable = MemFree + Buffers+Cached</b></li> </ul> </li> <li>Windows: The metric value is calculated by available memory minuses used memory. The value is obtained by calling the Windows API GlobalMemoryStatusEx.</li> </ul> | ≥ 0         | ECS                          | 1 minute                     |

| Metric          | Parameter            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-----------------|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| mem_usedPercent | (Agent) Memory Usage | <p>Memory usage of the monitored object</p> <p>Unit: percent</p> <ul style="list-style-type: none"><li>Linux: Obtain the metric value from the <b>/proc/meminfo</b> file: <b>(MemTotal - MemAvailable)/ MemTotal</b><ul style="list-style-type: none"><li>If <b>MemAvailable</b> is displayed in <b>/proc/meminfo</b>, <b>MemUsedPercent = (MemTotal - MemAvailable)/ MemTotal</b></li><li>If <b>MemAvailable</b> is not displayed in <b>/proc/meminfo</b>, <b>MemUsedPercent = (MemTotal - MemFree - Buffers - Cached)/ MemTotal</b></li></ul></li><li>Windows: The calculation formula is as follows: Used memory size/Total memory size*100%.</li></ul> | 0-100       | ECS                          | 1 minute                     |
| mem_free        | (Agent) Idle Memory  | <p>Amount of memory that is not being used</p> <p>Unit: GB</p> <ul style="list-style-type: none"><li>Linux: Obtain the metric value from <b>/proc/meminfo</b>.</li><li>Windows is not supported currently.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                       | ≥ 0         | ECS                          | 1 minute                     |



| Metric           | Parameter                  | Description                                                                                                                                                                                                                                                                                     | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|------------------|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| mem_buffers      | (Agent) Buffer             | Amount of memory that is being used for buffers<br>Unit: GB <ul style="list-style-type: none"><li>Linux: Obtain the metric value from <b>/proc/meminfo</b>. Run the <b>top</b> command to check the <b>KiB Mem:buffers</b> value.</li><li>Windows is not supported currently.</li></ul>         | $\geq 0$    | ECS                          | 1 minute                     |
| mem_cached       | (Agent) Cache              | Amount of memory that is being used for file caches<br>Unit: GB <ul style="list-style-type: none"><li>Linux: Obtain the metric value from <b>/proc/meminfo</b>. Run the <b>top</b> command to check the <b>KiB Swap:cached Mem</b> value.</li><li>Windows is not supported currently.</li></ul> | $\geq 0$    | ECS                          | 1 minute                     |
| total_open_files | (Agent) Total File Handles | Total handles used by all processes<br>Unit: count <ul style="list-style-type: none"><li>Linux: Use the <b>/proc/{pid}/fd</b> file to summarize the handles used by all processes.</li><li>Windows is not supported currently.</li></ul>                                                        | $\geq 0$    | ECS                          | 1 minute                     |

## OS Metric: Disk

### NOTE

- Currently, only physical disks are monitored. The NFS-attached disks cannot be monitored.
- By default, Docker-related mount points are shielded. The prefix of the mount point is as follows:  
`/var/lib/docker;/mnt/paas/kubernetes;/var/lib/mesos`

**Table 13-6** Disk metrics

| Metric    | Parameter                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-----------|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| disk_free | (Agent) Available Disk Space | Free space on the disks<br>Unit: GB <ul style="list-style-type: none"><li>• Linux: Run the <b>df -h</b> command to check the value in the <b>Avail</b> column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li><li>• Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li></ul> | $\geq 0$    | ECS - Mount point            | 1 minute                     |

| Metric     | Parameter                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|------------|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| disk_total | (Agent) Disk Storage Capacity | <p>Total space on the disks, including used and free<br/>Unit: GB</p> <ul style="list-style-type: none"> <li>Linux: Run the <b>df -h</b> command to check the value in the <b>Size</b> column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li> <li>Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li> </ul> | ≥ 0         | ECS - Mount point            | 1 minute                     |

| Metric    | Parameter               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-----------|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| disk_used | (Agent) Used Disk Space | Used space on the disks<br>Unit: GB <ul style="list-style-type: none"><li>Linux: Run the <b>df -h</b> command to check the value in the <b>Used</b> column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li><li>Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li></ul> | $\geq 0$    | ECS - Mount point            | 1 minute                     |

| Metric           | Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|------------------|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| disk_usedPercent | (Agent) Disk Usage | <p>Percentage of total disk space that is used, which is calculated as follows:<br/><b>Disk Usage = Used Disk Space/Disk Storage Capacity</b></p> <p>Unit: percent</p> <ul style="list-style-type: none"><li>Linux: It is calculated as follows: Used/Size. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li><li>Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li></ul> | 0-100       | ECS - Mount point            | 1 minute                     |

## OS Metric: Disk I/O

**Table 13-7** Disk I/O metrics

| Metric                   | Parameter               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Value Range | Monitored Object & Dimension                                                           | Monitoring Period (Raw Data) |
|--------------------------|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|----------------------------------------------------------------------------------------|------------------------------|
| disk_agt_read_bytes_rate | (Agent) Disks Read Rate | <p>Number of bytes read from the monitored disk per second</p> <p>Unit: byte/s</p> <ul style="list-style-type: none"><li>Linux:<br/>The disk read rate is calculated based on the data changes in the sixth column of the corresponding device in file <b>/proc/diskstats</b> in a collection period.</li><li>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li><li>Windows:<ul style="list-style-type: none"><li>Use Win32_PerformanceData_PerfDisk_LogicalDisk object in the WMI to obtain disk I/O data.</li><li>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li></ul></li></ul> | ≥ 0 bytes/s | <ul style="list-style-type: none"><li>EC S - Disk</li><li>EC S - Mount point</li></ul> | 1 minute                     |

| Metric | Parameter | Description                                                                                                                                                                         | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|--------|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
|        |           | <ul style="list-style-type: none"> <li>- When the CPU usage is high, monitoring data obtaining timeout may occur and result in the failure of obtaining monitoring data.</li> </ul> |             |                              |                              |

| Metric                      | Parameter                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Value Range    | Monitored Object & Dimension                                                              | Monitoring Period (Raw Data) |
|-----------------------------|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|-------------------------------------------------------------------------------------------|------------------------------|
| disk_agt_read_requests_rate | (Agent) Disks Read Requests | <p>Number of read requests sent to the monitored disk per second</p> <p>Unit: request/s</p> <ul style="list-style-type: none"> <li>Linux:                     <p>The disk read requests are calculated based on the data changes in the fourth column of the corresponding device in file <b>/proc/diskstats</b> in a collection period.</p> <p>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</p> </li> <li>Windows:                     <ul style="list-style-type: none"> <li>Use Win32_PerformanceData_PerfDisk_LogicalDisk object in the WMI to obtain disk I/O data.</li> <li>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li> <li>When the CPU usage is high, monitoring data obtaining timeout</li> </ul> </li> </ul> | ≥ 0 requests/s | <ul style="list-style-type: none"> <li>EC S - Disk</li> <li>EC S - Mount point</li> </ul> | 1 minute                     |



| Metric | Parameter | Description                                                       | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|--------|-----------|-------------------------------------------------------------------|-------------|------------------------------|------------------------------|
|        |           | may occur and result in the failure of obtaining monitoring data. |             |                              |                              |

| Metric                    | Parameter                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Value Range | Monitored Object & Dimension                                                           | Monitoring Period (Raw Data) |
|---------------------------|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|----------------------------------------------------------------------------------------|------------------------------|
| disk_agt_write_bytes_rate | (Agent) Disks Write Rate | <p>Number of bytes written to the monitored disk per second</p> <p>Unit: byte/s</p> <ul style="list-style-type: none"><li>Linux:<br/>The disk write rate is calculated based on the data changes in the tenth column of the corresponding device in file <b>/proc/diskstats</b> in a collection period.</li></ul> <p>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</p> <ul style="list-style-type: none"><li>Windows:<ul style="list-style-type: none"><li>Use Win32_PerformanceData_PerfDisk_LogicalDisk object in the WMI to obtain disk I/O data.</li><li>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li><li>When the CPU usage is high, monitoring data obtaining timeout</li></ul></li></ul> | ≥ 0 bytes/s | <ul style="list-style-type: none"><li>EC S - Disk</li><li>EC S - Mount point</li></ul> | 1 minute                     |

| Metric | Parameter | Description                                                       | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|--------|-----------|-------------------------------------------------------------------|-------------|------------------------------|------------------------------|
|        |           | may occur and result in the failure of obtaining monitoring data. |             |                              |                              |

| Metric                       | Parameter                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Value Range    | Monitored Object & Dimension                                                              | Monitoring Period (Raw Data) |
|------------------------------|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|-------------------------------------------------------------------------------------------|------------------------------|
| disk_agt_write_requests_rate | (Agent) Disks Write Requests | <p>Number of write requests sent to the monitored disk per second</p> <p>Unit: request/s</p> <ul style="list-style-type: none"> <li>Linux:                     <p>The disk write requests are calculated based on the data changes in the eighth column of the corresponding device in file <b>/proc/diskstats</b> in a collection period.</p> <p>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</p> </li> <li>Windows:                     <ul style="list-style-type: none"> <li>Use Win32_PerformanceData_PerfDisk_LogicalDisk object in the WMI to obtain disk I/O data.</li> <li>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li> <li>When the CPU usage is high, monitoring data obtaining timeout</li> </ul> </li> </ul> | ≥ 0 requests/s | <ul style="list-style-type: none"> <li>EC S - Disk</li> <li>EC S - Mount point</li> </ul> | 1 minute                     |

| Metric        | Parameter                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Value Range  | Monitored Object & Dimension                                                              | Monitoring Period (Raw Data) |
|---------------|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|-------------------------------------------------------------------------------------------|------------------------------|
|               |                                   | may occur and result in the failure of obtaining monitoring data.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |              |                                                                                           |                              |
| disk_readTime | (Agent) Average Read Request Time | <p>Average amount of time that read requests have waited on the disks<br/>Unit: ms/count</p> <ul style="list-style-type: none"> <li>Linux:<br/>The average read request time is calculated based on the data changes in the seventh column of the corresponding device in file <b>/proc/diskstats</b> in a collection period.<br/>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li> <li>Windows is not supported currently.</li> </ul> | ≥ 0 ms/Count | <ul style="list-style-type: none"> <li>EC S - Disk</li> <li>EC S - Mount point</li> </ul> | 1 minute                     |

| Metric         | Parameter                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Value Range  | Monitored Object & Dimension                                                              | Monitoring Period (Raw Data) |
|----------------|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|-------------------------------------------------------------------------------------------|------------------------------|
| disk_writeTime | (Agent) Average Write Request Time | <p>Average amount of time that write requests have waited on the disks</p> <p>Unit: ms/count</p> <ul style="list-style-type: none"> <li>Linux: The average write request time is calculated based on the data changes in the eleventh column of the corresponding device in file <b>/proc/diskstats</b> in a collection period.</li> <li>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li> <li>Windows is not supported currently.</li> </ul> | ≥ 0 ms/Count | <ul style="list-style-type: none"> <li>EC S - Disk</li> <li>EC S - Mount point</li> </ul> | 1 minute                     |

| Metric       | Parameter              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Value Range | Monitored Object & Dimension                                                              | Monitoring Period (Raw Data) |
|--------------|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-------------------------------------------------------------------------------------------|------------------------------|
| disk_ioUtils | (Agent) Disk I/O Usage | <p>Percentage of the time that the disk has had I/O requests queued to the total disk operation time</p> <p>Unit: percent</p> <ul style="list-style-type: none"> <li>Linux: The disk I/O usage is calculated based on the data changes in the thirteenth column of the corresponding device in file <b>/proc/diskstats</b> in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li> <li>Windows is not supported currently.</li> </ul> | 0-100       | <ul style="list-style-type: none"> <li>EC S - Disk</li> <li>EC S - Mount point</li> </ul> | 1 minute                     |

| Metric            | Parameter                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Value Range | Monitored Object & Dimension                                                           | Monitoring Period (Raw Data) |
|-------------------|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|----------------------------------------------------------------------------------------|------------------------------|
| disk_queue_length | (Agent) Disk Queue Length | <p>Average number of read or write requests queued up for completion for the monitored disk in the monitoring period</p> <p>Unit: count</p> <ul style="list-style-type: none"><li>Linux:<br/>The average disk queue length is calculated based on the data changes in the fourteenth column of the corresponding device in file <b>/proc/diskstats</b> in a collection period.</li></ul> <p>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</p> <ul style="list-style-type: none"><li>Windows is not supported currently.</li></ul> | $\geq 0$    | <ul style="list-style-type: none"><li>EC S - Disk</li><li>EC S - Mount point</li></ul> | 1 minute                     |



| Metric                         | Parameter                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Value Range  | Monitored Object & Dimension                                                              | Monitoring Period (Raw Data) |
|--------------------------------|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|-------------------------------------------------------------------------------------------|------------------------------|
| disk_write_bytes_per_operation | (Agent) Average Disk Write Size | <p>Average number of bytes in an I/O write for the monitored disk in the monitoring period</p> <p>Unit: byte/op</p> <ul style="list-style-type: none"> <li>Linux: The average disk write size is calculated based on the data changes in the tenth column of the corresponding device to divide that of the eighth column in file <b>/proc/diskstats</b> in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li> <li>Windows is not supported currently.</li> </ul> | ≥ 0 bytes/op | <ul style="list-style-type: none"> <li>EC S - Disk</li> <li>EC S - Mount point</li> </ul> | 1 minute                     |

| Metric                        | Parameter                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Value Range  | Monitored Object & Dimension                                                              | Monitoring Period (Raw Data) |
|-------------------------------|-----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|-------------------------------------------------------------------------------------------|------------------------------|
| disk_read_bytes_per_operation | (Agent)<br>Average Disk Read Size | <p>Average number of bytes in an I/O read for the monitored disk in the monitoring period</p> <p>Unit: byte/op</p> <ul style="list-style-type: none"> <li>Linux: The average disk read size is calculated based on the data changes in the sixth column of the corresponding device to divide that of the fourth column in file <b>/proc/diskstats</b> in a collection period.</li> </ul> <p>The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</p> <ul style="list-style-type: none"> <li>Windows is not supported currently.</li> </ul> | ≥ 0 bytes/op | <ul style="list-style-type: none"> <li>EC S - Disk</li> <li>EC S - Mount point</li> </ul> | 1 minute                     |

| Metric        | Parameter                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Value Range | Monitored Object & Dimension                                                              | Monitoring Period (Raw Data) |
|---------------|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-------------------------------------------------------------------------------------------|------------------------------|
| disk_io_svctm | (Agent) Disk I/O Service Time | <p>Average time in an I/O read or write for the monitored disk in the monitoring period</p> <p>Unit: ms/op</p> <ul style="list-style-type: none"> <li>Linux: The average disk I/O service time is calculated based on the data changes in the thirteenth column of the corresponding device to divide the sum of data changes in the fourth and eighth columns in file <b>/proc/diskstats</b> in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</li> <li>Windows is not supported currently.</li> </ul> | ≥ 0         | <ul style="list-style-type: none"> <li>EC S - Disk</li> <li>EC S - Mount point</li> </ul> | 1 minute                     |

| Metric                   | Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|--------------------------|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| disk_device_used_percent | Block Device Usage | <p>Percentage of the physical disk usage of the monitored object.<br/>Calculation formula: Used storage space of all mounted disk partitions/ Total disk storage space</p> <ul style="list-style-type: none"> <li>Collection method for Linux ECSs: Obtain the disk usage of each mount point, calculate the total disk storage space based on the disk sector size and the number of sectors, and then you can calculate the used storage space in total.</li> <li>Windows ECSs do not support this metric.</li> </ul> | 0-100%      | ECS - Disk                   | 1 minute                     |

## OS Metric: File System

Table 13-8 File system metrics

| Metric          | Parameter                             | Description                                                                                                                                                                                                                                               | Value Range                                                                                                    | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-----------------|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------|------------------------------|------------------------------|
| disk_fs_rwstate | (Agent) File System Read/Write Status | <p>Read and write status of the mounted file system of the monitored object<br/>Possible values are <b>0</b> (read and write) and <b>1</b> (read only).</p> <p>Linux: Check file system information in the fourth column in file <b>/proc/mounts</b>.</p> | <ul style="list-style-type: none"> <li><b>0</b>: readable and writable</li> <li><b>1</b>: read-only</li> </ul> | ECS - Mount point            | 1                            |

| Metric                 | Parameter                              | Description                                                                                                                                                                                                                                                                                                               | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|------------------------|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| disk_inodesTotal       | (Agent) Disk inode Total               | Total number of index nodes on the disk<br>Linux: Run the <b>df -i</b> command to check the value in the <b>Inodes</b> column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).                | $\geq 0$    | ECS - Mount point            | 1 minute                     |
| disk_inodesUsed        | (Agent) Total inode Used               | Number of used index nodes on the disk<br>Linux: Run the <b>df -i</b> command to check the value in the <b>IUsed</b> column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).                  | $\geq 0$    | ECS - Mount point            | 1 minute                     |
| disk_inodesUsedPercent | (Agent) Percentage of Total inode Used | Number of used index nodes on the disk<br>Unit: percent<br>Linux: Run the <b>df -i</b> command to check the value in the <b>IUse%</b> column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). | 0-100       | ECS - Mount point            | 1 minute                     |

 NOTE

The Windows OS does not support the file system metrics.

## OS Metric: NIC

Table 13-9 NIC metrics

| Metric      | Parameter                  | Description                                                                                                                                                                                                                                                                                 | Value Range    | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-------------|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|------------------------------|------------------------------|
| net_bitRecv | (Agent) Outbound Bandwidth | Number of bits received by this NIC per second<br>Unit: bit/s<br><ul style="list-style-type: none"><li>Linux: Check metric value changes in file / <b>proc/net/dev</b> in a collection period.</li><li>Windows: Use the MibIfRow object in the WMI to obtain network metric data.</li></ul> | $\geq 0$ bit/s | ECS                          | 1 minute                     |
| net_bitSent | (Agent) Inbound Bandwidth  | Number of bits sent by this NIC per second<br>Unit: bit/s<br><ul style="list-style-type: none"><li>Linux: Check metric value changes in file / <b>proc/net/dev</b> in a collection period.</li><li>Windows: Use the MibIfRow object in the WMI to obtain network metric data.</li></ul>     | $\geq 0$ bit/s | ECS                          | 1 minute                     |

| Metric         | Parameter                       | Description                                                                                                                                                                                                                                                                                   | Value Range  | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|----------------|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|------------------------------|------------------------------|
| net_packetRecv | (Agent) NIC Packet Receive Rate | Number of packets received by this NIC per second<br>Unit: count/s <ul style="list-style-type: none"><li>Linux: Check metric value changes in file / <b>proc/net/dev</b> in a collection period.</li><li>Windows: Use the MibIfRow object in the WMI to obtain network metric data.</li></ul> | ≥ 0 Counts/s | ECS                          | 1 minute                     |
| net_packetSent | (Agent) NIC Packet Send Rate    | Number of packets sent by this NIC per second<br>Unit: count/s <ul style="list-style-type: none"><li>Linux: Check metric value changes in file / <b>proc/net/dev</b> in a collection period.</li><li>Windows: Use the MibIfRow object in the WMI to obtain network metric data.</li></ul>     | ≥ 0 Counts/s | ECS                          | 1 minute                     |
| net_errin      | (Agent) Receive Error Rate      | Percentage of receive errors detected by this NIC per second<br>Unit: percent <ul style="list-style-type: none"><li>Linux: Check metric value changes in file / <b>proc/net/dev</b> in a collection period.</li><li>Windows is not supported currently.</li></ul>                             | 0-100        | ECS                          | 1 minute                     |

| Metric      | Parameter                            | Description                                                                                                                                                                                                                                                                      | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-------------|--------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| net_error   | (Agent) Transmit Error Rate          | Percentage of transmit errors detected by this NIC per second<br>Unit: percent <ul style="list-style-type: none"><li>Linux: Check metric value changes in file / <b>proc/net/dev</b> in a collection period.</li><li>Windows is not supported currently.</li></ul>               | 0-100       | ECS                          | 1 minute                     |
| net_dropin  | (Agent) Received Packet Drop Rate    | Percentage of packets received by this NIC which were dropped per second<br>Unit: percent <ul style="list-style-type: none"><li>Linux: Check metric value changes in file / <b>proc/net/dev</b> in a collection period.</li><li>Windows is not supported currently.</li></ul>    | 0-100       | ECS                          | 1 minute                     |
| net_dropout | (Agent) Transmitted Packet Drop Rate | Percentage of packets transmitted by this NIC which were dropped per second<br>Unit: percent <ul style="list-style-type: none"><li>Linux: Check metric value changes in file / <b>proc/net/dev</b> in a collection period.</li><li>Windows is not supported currently.</li></ul> | 0-100       | ECS                          | 1 minute                     |



## OS Metric: NTP

**Table 13-10** NTP metrics

| Metric     | Parameter          | Description                                                                                                                             | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|------------|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| ntp_offset | (Agent) NTP Offset | NTP offset of the monitored object<br>Unit: ms<br>Collection method for Linux ECSs: Run <b>chronyc sources -v</b> to obtain the offset. | $\geq 0$ ms | ECS                          | 1 minute                     |

## OS Metric: TCP

**Table 13-11** TCP metrics

| Metric        | Parameter         | Description                                                                                                                                                                                                                                                                                                                    | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|---------------|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| net_tcp_total | (Agent) TCP TOTAL | Total number of TCP connections in all states<br>Unit: count <ul style="list-style-type: none"><li>Linux: Obtain TCP connections in all states from the <b>/proc/net/tcp</b> file, and then collect the number of connections in each state.</li><li>Windows: Obtain the metric value using WindowsAPI GetTcpTable2.</li></ul> | $\geq 0$    | ECS                          | 1 minute                     |

| Metric              | Parameter                | Description                                                                                                                                                                                                                                                                                                                                       | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|---------------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| net_tcp_established | (Agent) TCP ESTABLISHED  | Number of TCP connections in ESTABLISHED state<br>Unit: count <ul style="list-style-type: none"><li>Linux: Obtain TCP connections in all states from the <b>/proc/net/tcp</b> file, and then collect the number of connections in each state.</li><li>Windows: Obtain the metric value using WindowsAPI GetTcpTable2.</li></ul>                   | $\geq 0$    | ECS                          | 1 minute                     |
| net_tcp_sys_sent    | (Agent) TCP SYS_SENT     | Number of TCP connections that are being requested by the client<br>Unit: count <ul style="list-style-type: none"><li>Linux: Obtain TCP connections in all states from the <b>/proc/net/tcp</b> file, and then collect the number of connections in each state.</li><li>Windows: Obtain the metric value using WindowsAPI GetTcpTable2.</li></ul> | $\geq 0$    | ECS                          | 1 minute                     |
| net_tcp_sys_rcv     | (Agent) TCP SYS_RECEIVED | Number of pending TCP connections received by the server<br>Unit: count <ul style="list-style-type: none"><li>Linux: Obtain TCP connections in all states from the <b>/proc/net/tcp</b> file, and then collect the number of connections in each state.</li><li>Windows: Obtain the metric value using WindowsAPI GetTcpTable2.</li></ul>         | $\geq 0$    | ECS                          | 1 minute                     |

| Metric            | Parameter             | Description                                                                                                                                                                                                                                                                                                                                                                                                   | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-------------------|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| net_tcp_fin_wait1 | (Agent) TCP FIN_WAIT1 | <p>Number of TCP connections waiting for ACK packets when the connections are being actively closed by the client</p> <p>Unit: count</p> <ul style="list-style-type: none"> <li>Linux: Obtain TCP connections in all states from the <b>/proc/net/tcp</b> file, and then collect the number of connections in each state.</li> <li>Windows: Obtain the metric value using WindowsAPI GetTcpTable2.</li> </ul> | ≥ 0         | ECS                          | 1 minute                     |
| net_tcp_fin_wait2 | (Agent) TCP FIN_WAIT2 | <p>Number of TCP connections in the FIN_WAIT2 state</p> <p>Unit: count</p> <ul style="list-style-type: none"> <li>Linux: Obtain TCP connections in all states from the <b>/proc/net/tcp</b> file, and then collect the number of connections in each state.</li> <li>Windows: Obtain the metric value using WindowsAPI GetTcpTable2.</li> </ul>                                                               | ≥ 0         | ECS                          | 1 minute                     |
| net_tcp_time_wait | (Agent) TCP TIME_WAIT | <p>Number of TCP connections in TIME_WAIT state</p> <p>Unit: count</p> <ul style="list-style-type: none"> <li>Linux: Obtain TCP connections in all states from the <b>/proc/net/tcp</b> file, and then collect the number of connections in each state.</li> <li>Windows: Obtain the metric value using WindowsAPI GetTcpTable2.</li> </ul>                                                                   | ≥ 0         | ECS                          | 1 minute                     |

| Metric             | Parameter              | Description                                                                                                                                                                                                                                                                                                                                                                                      | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|--------------------|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| net_tcp_close      | (Agent) TCP CLOSE      | Number of closed TCP connections<br>Unit: count <ul style="list-style-type: none"><li>Linux: Obtain TCP connections in all states from the <b>/proc/net/tcp</b> file, and then collect the number of connections in each state.</li><li>Windows: Obtain the metric value using WindowsAPI GetTcpTable2.</li></ul>                                                                                | $\geq 0$    | ECS                          | 1 minute                     |
| net_tcp_close_wait | (Agent) TCP CLOSE_WAIT | Number of TCP connections in CLOSE_WAIT TCP state<br>Unit: count <ul style="list-style-type: none"><li>Linux: Obtain TCP connections in all states from the <b>/proc/net/tcp</b> file, and then collect the number of connections in each state.</li><li>Windows: Obtain the metric value using WindowsAPI GetTcpTable2.</li></ul>                                                               | $\geq 0$    | ECS                          | 1 minute                     |
| net_tcp_last_ack   | (Agent) TCP LAST_ACK   | Number of TCP connections waiting for ACK packets when the connections are being passively closed by the client<br>Unit: count <ul style="list-style-type: none"><li>Linux: Obtain TCP connections in all states from the <b>/proc/net/tcp</b> file, and then collect the number of connections in each state.</li><li>Windows: Obtain the metric value using WindowsAPI GetTcpTable2.</li></ul> | $\geq 0$    | ECS                          | 1 minute                     |

| Metric          | Parameter                       | Description                                                                                                                                                                                                                                                                                                                                                                         | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-----------------|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| net_tcp_listen  | (Agent) TCP LISTEN              | Number of TCP connections in the LISTEN state<br>Unit: count <ul style="list-style-type: none"><li>Linux: Obtain TCP connections in all states from the <b>/proc/net/tcp</b> file, and then collect the number of connections in each state.</li><li>Windows: Obtain the metric value using WindowsAPI GetTcpTable2.</li></ul>                                                      | $\geq 0$    | ECS                          | 1 minute                     |
| net_tcp_closing | (Agent) TCP CLOSING             | Number of TCP connections to be automatically closed by the server and the client at the same time<br>Unit: count <ul style="list-style-type: none"><li>Linux: Obtain TCP connections in all states from the <b>/proc/net/tcp</b> file, and then collect the number of connections in each state.</li><li>Windows: Obtain the metric value using WindowsAPI GetTcpTable2.</li></ul> | $\geq 0$    | ECS                          | 1 minute                     |
| net_tcp_retrans | (Agent) TCP Retransmission Rate | Percentage of packets that are resent<br>Unit: percent <ul style="list-style-type: none"><li>Linux: Obtain the metric value from the <b>/proc/net/snmp</b> file. The value is the ratio of the number of sent packets to the number of retransmitted packages in a collection period.</li><li>Windows: Obtain the metric value using WindowsAPI GetTcpStatistics.</li></ul>         | 0-100%      | ECS                          | 1 minute                     |

## OS Metric: GPU

**Table 13-12** GPU metrics

| Metric            | Parameter         | Description                                                                                                                                                                                                                                                                                                 | Value Range                                                                                                                            | Monitored Object & Dimension                                          | Monitoring Period (Raw Data) |
|-------------------|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------|------------------------------|
| gpu_status        | GPU Health Status | Overall measurement of the GPU health<br>Unit: none <ul style="list-style-type: none"><li>Linux: Obtain the metric value using the <b>libnvidia-ml.so.1</b> library file of the graphics card.</li><li>Windows: Obtain the metric value using the <b>nvml.dll</b> library of the graphics card.</li></ul>   | <ul style="list-style-type: none"><li>0: The GPU is healthy.</li><li>1: The GPU is subhealthy.</li><li>2: The GPU is faulty.</li></ul> | <ul style="list-style-type: none"><li>ECS</li><li>ECS - GPU</li></ul> | 1 minute                     |
| gpu_usage_encoder | Encoding Usage    | Encoding capability usage on the GPU<br>Unit: percent <ul style="list-style-type: none"><li>Linux: Obtain the metric value using the <b>libnvidia-ml.so.1</b> library file of the graphics card.</li><li>Windows: Obtain the metric value using the <b>nvml.dll</b> library of the graphics card.</li></ul> | 0-100%                                                                                                                                 | <ul style="list-style-type: none"><li>ECS</li><li>ECS - GPU</li></ul> | 1 minute                     |

| Metric                   | Parameter                       | Description                                                                                                                                                                                                                                                                                                                                                                                         | Value Range | Monitored Object & Dimension                                          | Monitoring Period (Raw Data) |
|--------------------------|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-----------------------------------------------------------------------|------------------------------|
| gpu_usage_decoder        | Decoding Usage                  | <p>Decoding capability usage on the GPU</p> <p>Unit: percent</p> <ul style="list-style-type: none"><li>Linux: Obtain the metric value using the <b>libnvidia-ml.so.1</b> library file of the graphics card.</li><li>Windows: Obtain the metric value using the <b>nvml.dll</b> library of the graphics card.</li></ul>                                                                              | 0-100%      | <ul style="list-style-type: none"><li>ECS</li><li>ECS - GPU</li></ul> | 1 minute                     |
| gpu_volatile_correctable | Volatile Correctable ECC Errors | <p>Number of correctable ECC errors since the GPU is reset. The value is reset to <b>0</b> each time the GPU is reset.</p> <p>Unit: count</p> <ul style="list-style-type: none"><li>Linux: Obtain the metric value using the <b>libnvidia-ml.so.1</b> library file of the graphics card.</li><li>Windows: Obtain the metric value using the <b>nvml.dll</b> library of the graphics card.</li></ul> | $\geq 0$    | <ul style="list-style-type: none"><li>ECS</li><li>ECS - GPU</li></ul> | 1 minute                     |

| Metric                      | Parameter                          | Description                                                                                                                                                                                                                                                                                                                                                                                | Value Range | Monitored Object & Dimension                                          | Monitoring Period (Raw Data) |
|-----------------------------|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-----------------------------------------------------------------------|------------------------------|
| gpu_volatile_uncorrectable  | Volatile Uncorrectable ECC Errors  | Number of uncorrectable ECC errors since the GPU is reset. The value is reset to <b>0</b> each time the GPU is reset.<br>Unit: count <ul style="list-style-type: none"><li>Linux: Obtain the metric value using the <b>libnvidia-ml.so.1</b> library file of the graphics card.</li><li>Windows: Obtain the metric value using the <b>nvml.dll</b> library of the graphics card.</li></ul> | $\geq 0$    | <ul style="list-style-type: none"><li>ECS</li><li>ECS - GPU</li></ul> | 1 minute                     |
| gpu_aggregate_correctable   | Aggregate Correctable ECC Errors   | Aggregate correctable ECC errors on the GPU<br>Unit: count <ul style="list-style-type: none"><li>Linux: Obtain the metric value using the <b>libnvidia-ml.so.1</b> library file of the graphics card.</li><li>Windows: Obtain the metric value using the <b>nvml.dll</b> library of the graphics card.</li></ul>                                                                           | $\geq 0$    | <ul style="list-style-type: none"><li>ECS</li><li>ECS - GPU</li></ul> | 1 minute                     |
| gpu_aggregate_uncorrectable | Aggregate Uncorrectable ECC Errors | Aggregate uncorrectable ECC Errors on the GPU<br>Unit: count <ul style="list-style-type: none"><li>Linux: Obtain the metric value using the <b>libnvidia-ml.so.1</b> library file of the graphics card.</li><li>Windows: Obtain the metric value using the <b>nvml.dll</b> library of the graphics card.</li></ul>                                                                         | $\geq 0$    | <ul style="list-style-type: none"><li>ECS</li><li>ECS - GPU</li></ul> | 1 minute                     |



| Metric                      | Parameter                      | Description                                                                                                                                                                                                                                                                                                                                                                                           | Value Range | Monitored Object & Dimension                                          | Monitoring Period (Raw Data) |
|-----------------------------|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-----------------------------------------------------------------------|------------------------------|
| gpu_retired_page_single_bit | Retired Page Single Bit Errors | <p>Number of retired page single bit errors, which indicates the number of single-bit pages blocked by the graphics card</p> <p>Unit: count</p> <ul style="list-style-type: none"><li>Linux: Obtain the metric value using the <b>libnvidia-ml.so.1</b> library file of the graphics card.</li><li>Windows: Obtain the metric value using the <b>nvml.dll</b> library of the graphics card.</li></ul> | $\geq 0$    | <ul style="list-style-type: none"><li>ECS</li><li>ECS - GPU</li></ul> | 1 minute                     |
| gpu_retired_page_double_bit | Retired Page Double Bit Errors | <p>Number of retired page double bit errors, which indicates the number of double-bit pages blocked by the graphics card</p> <p>Unit: count</p> <ul style="list-style-type: none"><li>Linux: Obtain the metric value using the <b>libnvidia-ml.so.1</b> library file of the graphics card.</li><li>Windows: Obtain the metric value using the <b>nvml.dll</b> library of the graphics card.</li></ul> | $\geq 0$    | <ul style="list-style-type: none"><li>ECS</li><li>ECS - GPU</li></ul> | 1 minute                     |

| Metric                | Parameter                  | Description                                                                                                                                                                                                                                                                                                               | Value Range                                                                                                                                                                                                                                             | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-----------------------|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|------------------------------|
| gpu_performance_state | (Agent) Performance Status | <p>GPU performance of the monitored object</p> <p>Unit: none</p> <ul style="list-style-type: none"> <li>Linux: Obtain the metric value using the <b>libnvidia-ml.so.1</b> library file of the graphics card.</li> <li>Windows: Obtain the metric value using the <b>nvml.dll</b> library of the graphics card.</li> </ul> | <p>P0-P15, P32</p> <ul style="list-style-type: none"> <li><b>P0:</b> indicates the maximum performance status.</li> <li><b>P15:</b> indicates the minimum performance status.</li> <li><b>P32:</b> indicates the unknown performance status.</li> </ul> | ECS - GPU                    | 1 minute                     |

| Metric           | Parameter                | Description                                                                                                                                                                                                                                                                                                     | Value Range | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|------------------|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|------------------------------|------------------------------|
| gpu_usage_memory | (Agent) GPU Memory Usage | GPU memory usage of the monitored object<br>Unit: percent <ul style="list-style-type: none"><li>Linux: Obtain the metric value using the <b>libnvidia-ml.so.1</b> library file of the graphics card.</li><li>Windows: Obtain the metric value using the <b>nvml.dll</b> library of the graphics card.</li></ul> | 0-100       | ECS - GPU                    | 1 minute                     |
| gpu_usage_gpu    | (Agent) GPU Usage        | GPU usage of the monitored object<br>Unit: percent <ul style="list-style-type: none"><li>Linux: Obtain the metric value using the <b>libnvidia-ml.so.1</b> library file of the graphics card.</li><li>Windows: Obtain the metric value using the <b>nvml.dll</b> library of the graphics card.</li></ul>        | 0-100       | ECS - GPU                    | 1 minute                     |

## Dimensions

| Dimension | Key         | Value                 |
|-----------|-------------|-----------------------|
| ECS       | instance_id | Specifies the ECS ID. |

## 13.4 Setting Alarm Rules

### Scenarios

Setting ECS alarm rules allows you to customize the monitored objects and notification policies so that you can closely monitor your ECSs.

This section describes how to set ECS alarm rules, including alarm rule names, monitoring objects, monitoring metrics, alarm thresholds, monitoring intervals, and notifications.

## Procedure

1. Log in to the management console.
2. Under **Management & Deployment**, choose **Cloud Eye**.
3. In the navigation pane on the left, choose **Alarm Management > Alarm Rules**.
4. On the **Alarm Rules** page, click **Create Alarm Rule** to create an alarm rule, or modify an existing alarm rule.

The following uses modifying an existing alarm rule as an example.

- a. Click the target alarm rule.
- b. Click **Modify** in the upper right corner of the page.
- c. In the **Modify Alarm Rule** dialog box shown in [Figure 13-1](#), set the parameters as prompted.

**Figure 13-1** Modifying an alarm rule

The screenshot shows a 'Modify Alarm Rule' dialog box with the following fields and values:

- Name: as-alarm-d3o6
- Description: autoScaling
- Service: Auto Scaling
- Monitored Object: AS Group, 406487d6-a3d4-...
- Metric: CPU Usage
- Threshold: Max., >, 80%
- Occurrences: 1
- Monitoring Interval: 5 minutes
- Send Notification:  Yes,  No

Buttons: OK, Cancel

- d. Click **Modify**.  
After an alarm rule is modified, the system automatically notifies you of an alarm when the alarm complying with the alarm rule is generated.

### NOTE

For more information about ECS alarm rules, see *Cloud Eye User Guide*.

## 13.5 Viewing ECS Metrics

### Scenarios

The cloud platform provides Cloud Eye, which monitors the running statuses of your ECSs. You can obtain the monitoring metrics of each ECS on the management console.

There is a short time delay between transmission and display of monitoring data. The status of an ECS displayed on Cloud Eye is the status obtained 5 to 10 minutes before. If an ECS is just created, wait for 5 to 10 minutes to view the real-time monitoring data.

### Prerequisites

- The ECS is running properly.  
Cloud Eye does not display the monitoring data for a stopped, faulty, or deleted ECS. After such an ECS restarts or recovers, the monitoring data is available in Cloud Eye.

#### NOTE

Cloud Eye discontinues monitoring ECSs that remain in **Stopped** or **Faulty** state for 24 hours and removes them from the monitoring list. However, the alarm rules for such ECSs are not automatically deleted.

- Alarm rules have been configured in Cloud Eye for the target ECS.  
The monitoring data is unavailable for the ECSs without alarm rules configured in Cloud Eye. For details, see [Setting Alarm Rules](#).
- The target ECS has been properly running for at least 10 minutes.  
The monitoring data and graphics are available for a new ECS after the ECS runs for at least 10 minutes.

### Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID for search.
4. Click the name of the target ECS. The page providing details about the ECS is displayed.
5. Click the **Monitoring** tab to view the monitoring data.
6. In the ECS monitoring area, select a duration to view the monitoring data.  
You can view the monitoring data of the ECS in the last 1 hour, last 3 hours, last 12 hours, last 1 day, or last 7 days.

# 14 FAQs

---

## 14.1 ECS Overview

### 14.1.1 What Are the Precautions for Using ECSs?

- Do not upgrade ECS kernel or OS versions. If you want to upgrade the main OS version, for example, from CentOS 7.2 to Cent OS 7.3, use the provided OS changing function.
- Do not uninstall the performance optimization software pre-installed on your ECSs.
- Do not change NIC MAC addresses. Otherwise, the network connection will fail.

### 14.1.2 What Can I Do with ECSs?

You can use ECSs just like traditional physical servers. On an ECS, you can deploy any service application, such as an email system, web system, and Enterprise Resource Planning (ERP) system. After creating an ECS, you can use it like using your local computer or physical server.

## 14.2 Creation and Deletion

### 14.2.1 Why Does the Failures Area Show an ECS Creation Failure But the ECS List Displays the Created ECS?

#### Symptom

After you created an ECS bound with an EIP on the management console, the ECS creation was successful but binding the EIP failed due to insufficient EIPs. Although the **Failures** area showed that the ECS creation failed, the ECS was displayed in the ECS list. The results of the ECS creation task were inconsistent.

## Root Cause

- The ECS list displays created ECSs.
- The **Failures** area shows the ECS creation status, including the statuses of subtasks, such as creating ECS resources and binding an EIP. Only when all subtasks are successful, the ECS is created.

If the ECS is created but EIP binding failed, the task failed. However, the ECS you created is temporarily displayed in the list. After the system rolls back, the ECS is removed from the list.

## 14.2.2 How Quickly Can I Obtain an ECS?

Obtaining an ECS can take as little as a few minutes.

The time it takes to obtain an ECS depends on ECS specifications, available resources (such as EVS disks and EIPs), and system load.

### NOTE

If it takes a long time to obtain your ECS, contact the administrator.

## 14.2.3 What Happens After I Click the Delete Button?

After you click **Delete**, the selected ECSs will be deleted. You can also choose to delete the EVS disks and EIPs together with the selected ECSs. If you do not delete them, they will be retained. If necessary, you can manually delete them later.

To delete selected ECSs, perform the following operations:

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Select the ECSs to be deleted.
4. Above the ECS list, choose **Delete**.

## 14.2.4 Can a Deleted ECS Be Provisioned Again?

No. ECSs in the **Deleted** state cannot provide services and are soon removed from the system.

A deleted ECS is retained in the ECS list on the management console only for a short period of time before it is permanently removed from the system. You can create a new ECS with the same specifications again.

## 14.2.5 Can I Forcibly Restart or Stop an ECS?

Yes. If an ECS remains in the **Restarting** or **Stopping** state for over 30 minutes after it is restarted, you can forcibly restart or stop the ECS as follows:

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Select the target ECS and click **Restart** or **Stop**.

A dialog box is displayed to confirm whether you want to restart or stop the ECS.

4. Select **Forcibly restart the preceding ECSs** or **Forcibly stop the preceding ECSs**.
5. Click **OK**.

## 14.3 Login and Connection

### 14.3.1 What Should I Do If Starting an ECS Remains in "Waiting for cloudResetPwdAgent" State?

#### Symptom

During ECS startup, it remains in "Waiting for cloudResetPwdAgent" state for 20-30 seconds.

Figure 14-1 Starting cloudResetPwdAgent

```
Starting rpcbind: [OK]
Starting NFS statd: [OK]
Starting cloudResetPwdAgent...
Waiting for cloudResetPwdAgent.....
```

#### Possible Causes

This issue is caused by the intranet DNS and user-defined DNS configurations.

#### Solution

1. Log in to the ECS as user **root**.
2. Run the following command to modify the `/etc/cloud/cloud.cfg` configuration file:  
`vi /etc/cloud/cloud.cfg`
3. Add the following statement to the configuration file:  
**manage\_etc\_hosts: true**

Figure 14-2 Editing the configuration file

```
users:
- name: root
 lock_passwd: False

disable_root: 0
ssh_pwauth: 1

datasource_list: ['OpenStack']
manage_etc_hosts: true

datasource:
 OpenStack:
 # timeout: the timeout value for a request at metadata service
 timeout : 50
 # The length in seconds to wait before giving up on the metadata
 # service. The actual total wait could be up to
 # len(resolvable_metadata_urls)*timeout
 max_wait : 120
```



## 14.3.2 What Should I Do If I Cannot Use MSTSC to Log In to an ECS Running the Windows Server 2012 OS?

### Symptom

An ECS running the Windows Server 2012 OS has password authentication configured during ECS creation. When a user used the initial password and MSTSC to log in to the ECS, the login failed and the system displayed the message "You must change your password before logging on for the first time. Please update your password or contact your system administrator or technical support."

### Possible Causes

The local computer used by the user is running the Windows 10 OS.

Due to limitations, the Windows 10 OS does not support remote logins to an ECS running the Windows Server 2012 OS using the initial password.

### Solutions

- Solution 1  
Use a local computer running the Windows 7 OS to remotely log in to the ECS running the Windows Server 2012 OS.
- Solution 2  
Retain the original local computer and change the initial login password.
  - a. Use VNC to log in to the ECS running the Windows Server 2012 OS for the first time.
  - b. Change the login password as prompted.
  - c. Use the changed password and MSTSC to log in to the ECS again.
- Solution 3:  
Retain the original local computer and initial login password.
  - a. Choose **Start**. In the **Search programs and files** text box, enter **mstsc** and press **Enter**.  
The **Remote Desktop Connection** page is displayed.
  - b. Enter the EIP and click **Connect**. Then, use username **administrator** and the login password configured during ECS creation for connection.  
The connection fails, and the system displays the message "You must change your password before logging on for the first time. Please update your password or contact your system administrator or technical support."
  - c. Click **Options** in the lower left corner of the **Remote Desktop Connection** page.
  - d. On the **General** tab, click **Save As** in the **Connection settings** pane and save the remote desktop file in .rdp format.
  - e. Open the .rdp file saved in **d**.
  - f. Add the following statement to the last line of the .rdp file and save the file.

**enablecredsspssupport:i:0**

- g. Double-click the edited .rdp file to set up the remote desktop connection.
- h. Click **Connect** to connect to the ECS running the Windows Server 2012 OS again.

### 14.3.3 How Can I Change a Remote Login Port?

#### Scenarios

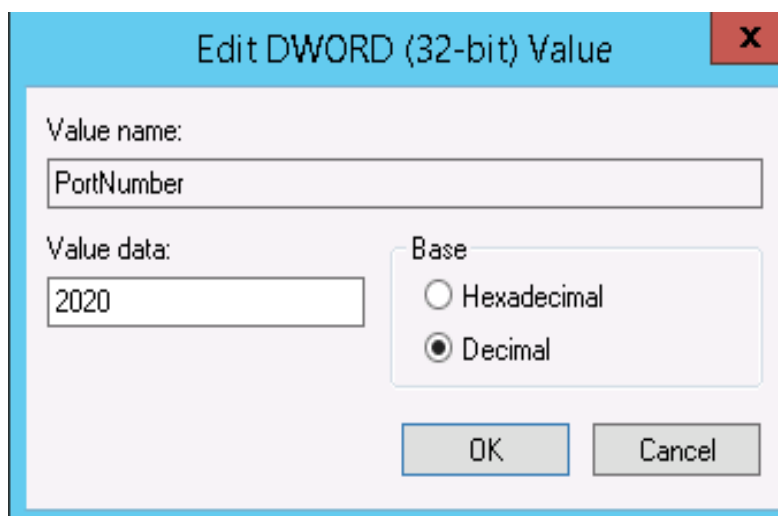
This section describes how to change a port for remote logins.

#### Windows

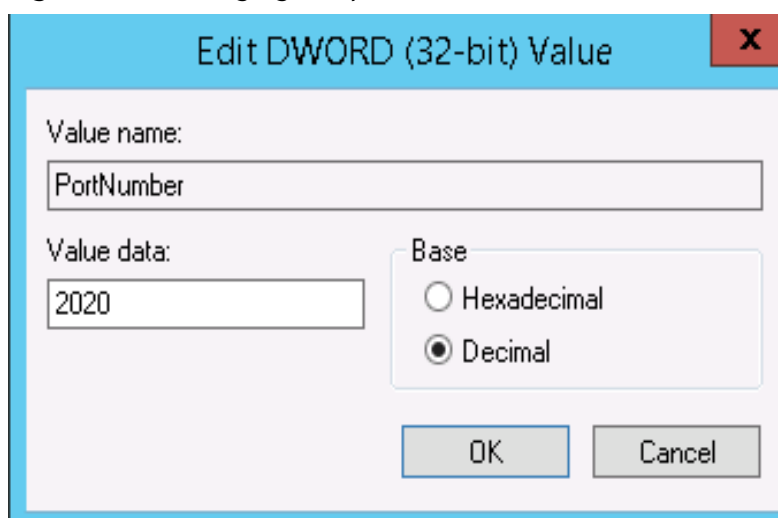
The following procedure uses an ECS running Windows Server 2012 as an example. The default login port of a Windows ECS is 3389. To change it to port 2020, for example, do as follows:

1. Modify the security group rule.
  - a. Log in to the management console.
  - b. Under **Computing**, click **Elastic Cloud Server**.
  - c. On the ECS list, click the name of an ECS for which you want to modify the security group rule.
  - d. On the ECS details page, click the security group in the **Security Groups** area to go to the security group details page.
  - e. On the **Inbound Rules** tab, click **Add Rule**. In the displayed dialog box, set **Protocol & Port** as follows:
    - **Protocols**: TCP (Custom ports)
    - **Port**: 2020

For details, see "Adding a Security Group Rule" in the *Virtual Private Cloud User Guide*.
2. Log in to the ECS.
3. In the **Run** dialog box, enter **regedit** to access the registry editor.
4. In **Registry Editor**, choose **HKEY\_LOCAL\_MACHINE > SYSTEM > CurrentControlSet > Control > Terminal Server > Wds > rdpwd > Tds > tcp** and double-click **PortNumber**.
  - a. In the dialog box that is displayed, set **Base** to **Decimal**.
  - b. Change the value in **Value data** to the new port number, which is **2020** in this example.

**Figure 14-3** Changing the port number to 2020

5. In **Registry Editor**, choose **HKEY\_LOCAL\_MACHINE > SYSTEM > CurrentControlSet > Control > Terminal Server > WinStations > RDP-Tcp** and double-click **PortNumber**.
  - a. In the dialog box that is displayed, set **Base** to **Decimal**.
  - b. Change the value in **Value data** to the new port number, which is **2020** in this example.

**Figure 14-4** Changing the port number to 2020

6. (Skip this step if the firewall is disabled.) Modify the inbound rules of the firewall.  
Choose **Control Panel > Windows Firewall > Advanced Settings > Inbound Rules > New Rule**.
  - **Rule Type:** Port
  - Protocol in **Protocol and Ports:** TCP
  - Port in **Protocol and Ports:** **Specific local ports, 2020** in this example
  - **Action:** **Allow the connection**
  - **Profile:** Default settings

– **Name: RDP-2020**

After the configuration, refresh the page to view the new rule.

7. Open the Windows search box, enter **services**, and select **Services**.

**Figure 14-5** Selecting Services




8. In the **Services** window, restart **Remote Desktop Services** or the ECS.
9. Use "IP address:Port" to remotely access the ECS.

## Linux

The following procedure uses an ECS running CentOS 7.3 as an example. The default login port of a Linux ECS is 22. To change it to port 2020, for example, do as follows:

1. Modify the security group rule.
  - a. Log in to the management console.
  - b. Under **Computing**, click **Elastic Cloud Server**.
  - c. On the ECS list, click the name of an ECS for which you want to modify the security group rule.
  - d. On the ECS details page, click the security group in the **Security Groups** area to go to the security group details page.
  - e. On the **Inbound Rules** tab, click **Add Rule**. In the displayed dialog box, set **Protocol & Port** as follows:
    - **Protocols:** TCP (Custom ports)
    - **Port:** 2020For details, see "Adding a Security Group Rule" in the *Virtual Private Cloud User Guide*.
2. Log in to the ECS.
3. Run the following command to edit the sshd configuration file:  
**vi /etc/ssh/sshd\_config**
4. Delete the comment tag (#) from the **#port 22** line and change **22** to **2020**.

**Figure 14-6** Changing the port number to 2020



5. Press **Esc** to exit Insert mode and enter **:wq!** to save the settings and exit.
6. Run either of the following commands to restart sshd:  
**service sshd restart**  
Or  
**systemctl restart sshd**
7. Skip this step if the firewall is disabled. Configure the firewall.  
The firewall varies depending on the CentOS version. CentOS 7 uses firewalld, and CentOS 6 uses iptables. The following operations use CentOS 7 as an example.  
Run the **firewall-cmd --state** command to check the firewall status.
  - (Recommended) Method 1: Add information about a new port to firewalld.
    - i. Run the following commands to add a rule for port 2020:  
**firewall-cmd --zone=public --add-port=2020/tcp --permanent**  
**firewall-cmd --reload**
    - ii. View the added port. The TCP connection of port 2020 will have been added.  
**firewall-cmd --list-all**
    - iii. Restart firewalld.  
**systemctl restart firewalld.service**
  - Method 2: Disable the firewall and the function of automatically enabling the firewall upon ECS startup.  
**systemctl stop firewalld**  
**systemctl disable firewalld**
8. Run the following command to check whether the port is open:  
**telnet *EIP port***  
For example: **telnet xx.xx.xx.xx 2020**

### 14.3.4 What Browser Version Is Required to Remotely Log In to an ECS?

When you use a browser to remotely log in to an ECS, ensure that the browser version meets the requirements listed in [Table 14-1](#).

**Table 14-1** Browser version requirements

| Browser           | Version   |
|-------------------|-----------|
| Google Chrome     | 31.0-75.0 |
| Mozilla Firefox   | 27.0-62.0 |
| Internet Explorer | 10.0-11.0 |

### 14.3.5 What Should I Do If the Page Does not Respond After I Log In to an ECS Using VNC and Do Not Perform Any Operation for a Long Period of Time?

If your computer is running Windows 7 and you logged in to the ECS using Internet Explorer 10 or 11, click **AltGr** twice on the VNC page to activate the page.

### 14.3.6 What Should I Do If I Cannot View Data After Logging In to an ECS Using VNC?

After you log in to an ECS using VNC and view data, for example, play videos or run the **cat** command to view large files, VNC may become unavailable due to the high memory usage of the browser. In such a case, use another browser and log in to the ECS again.

### 14.3.7 Why Does a Blank Screen Appear After I Attempted to Log In to an ECS Using VNC?

The blank screen means that another user has logged in to this ECS using VNC, so you were logged out.

Only one user can be logged in to an ECS using VNC at a time. If you are already logged in and another user logs in to the same ECS, you will be automatically logged out. You can log back in, but that will kick the other user out.


### 14.3.8 Why Does an Authentication Failure Occurs After I Attempt to Remotely Log In to a Windows ECS?

#### Symptom

When a local computer running Windows attempts to access a Windows ECS using RDP (for example, MSTSC), an identity authentication failure occurs and the desired function is not supported.

- If the error message contains only the information that an identity authentication failure occurs and that the desired function is not supported, rectify the fault by following the instructions provided in [Solution](#).
- If the error message shows that the fault was caused by "CredSSP Encryption Oracle Remediation", as shown in [Figure 14-7](#), the fault may be caused by a security patch released by Microsoft in March 2018. This patch may affect RDP-based CredSSP connections. As a result, setting up RDP-based connections to ECSs failed. Rectify the fault by following the instructions provided in [official Microsoft document](#).

**Figure 14-7** Failed to set up a remote desktop connection




## Solution

Modify the remote desktop connection settings on the Windows ECS:

1. Log in to the ECS.
2. Click **Start** in the lower left corner, right-click **Computer**, and choose **Properties** from the shortcut menu.
3. In the left navigation pane, choose **Remote settings**.
4. Click the **Remote** tab. In the **Remote Desktop** pane, select **Allow connections from computers running any version of Remote Desktop (less secure)**.

Figure 14-8 Remote settings



5. Click **OK**.

### 14.3.9 Why Can't I Use the Local Computer to Connect to My Windows ECS?

#### Symptom

An error message is displayed indicating that your local computer cannot connect to the remote computer.

Figure 14-9 Cannot connect to the remote computer





## Possible Causes

- Port 3389 of the security group on the ECS is disabled. For details, see [Checking Port Configuration on the ECS](#).
- The firewall on the ECS is disabled. For details, see [Checking Whether the Firewall Is Correctly Configured](#).
- The remote desktop connection is not correctly configured. For details, see [Checking Remote Desktop Connection Settings](#).
- Remote Desktop Services are not started. For solution, see [Checking Remote Desktop Services](#).
- Remote Desktop Session Host is not correctly configured. For details, see [Checking Remote Desktop Session Host Configuration](#).

## Checking Port Configuration on the ECS

Check whether port 3389 (used by default) on the ECS is accessible.

Ensure that port 3389 has been added in the inbound rule.


On the page providing details about the ECS, click the **Security Groups** tab and view port 3389 in the inbound rule of the security group.

## Checking Whether the Firewall Is Correctly Configured

Check whether the firewall is enabled on the ECS.


1. Log in to the ECS using VNC available on the management console.
2. Click the Windows icon in the lower left corner of the desktop and choose **Control Panel > Windows Firewall**.

**Figure 14-10** Windows Firewall



3. Click **Turn Windows Firewall on or off**.  
View and set the firewall status.


Figure 14-11 Checking firewall status



To enable Windows firewall, perform the following steps:

4. Click **Advanced settings**.
5. Check **Inbound Rules** and ensure that the following rules are enabled:
  - Remote Desktop - User Mode (TCP-In), Public
  - Remote Desktop - User Mode (TCP-In), Domain, Private

Figure 14-12 Inbound Rules



If the port configured in the inbound rule of the firewall is different from that configured on the remote server, the remote login will fail. If this occurs, add the port configured on the remote server in the inbound rule of the firewall.

**NOTE**

The default port is 3389. If you use another port, add that port in the inbound rule of the firewall.


After you perform the preceding operations, try to remotely log in to the ECS again.

## Checking Remote Desktop Connection Settings

Modify the remote desktop connection settings on the Windows ECS:

1. Log in to the ECS.
2. Click **Start** in the lower left corner, right-click **Computer**, and choose **Properties** from the shortcut menu.
3. In the left navigation pane, choose **Remote settings**.
4. Click the **Remote** tab. In the **Remote Desktop** pane, select **Allow connections from computers running any version of Remote Desktop (less secure)**.

**Figure 14-13** Remote settings




5. Click **OK**.

## Checking Remote Desktop Services


1. Open the Windows search box, enter **services**, and select **Services**.

2. In the **Services** window, restart **Remote Desktop Services**. Ensure that **Remote Desktop Services** is in the **Running** status.


**Figure 14-14** Remote Desktop Services

## Checking Remote Desktop Session Host Configuration

1. Open the **cmd** window and enter **gpedit.msc**.
2. Click **OK** to start Local Group Policy Editor.
3. Choose **Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services**.
4. Choose **Remote Desktop Session Host > Security > Require use of specific security layer for remote (RDP) connections**.

**Figure 14-15** Require use of specific security layer for remote (RDP) connections


5. Set **Require use of specific security layer for remote (RDP) connections** to **Enabled** and **Security layer** to **RDP**.

**Figure 14-16** Setting security layer to RDP

## 14.3.10 How Can I Obtain the Permission to Remotely Log In to a Windows ECS?

### Symptom

When you connect a remote desktop to a Windows ECS, the system prompts that you need to be granted the right to sign in through Remote Desktop Services.


**Figure 14-17** Remote login right missing.

### Solution

1. Open the **cmd** window and enter **gpedit.msc**.
2. Click **OK** to start Local Group Policy Editor.
3. Choose **Computer Configuration > Windows Settings > Security Settings > Local Policies > User Rights Assignment**.

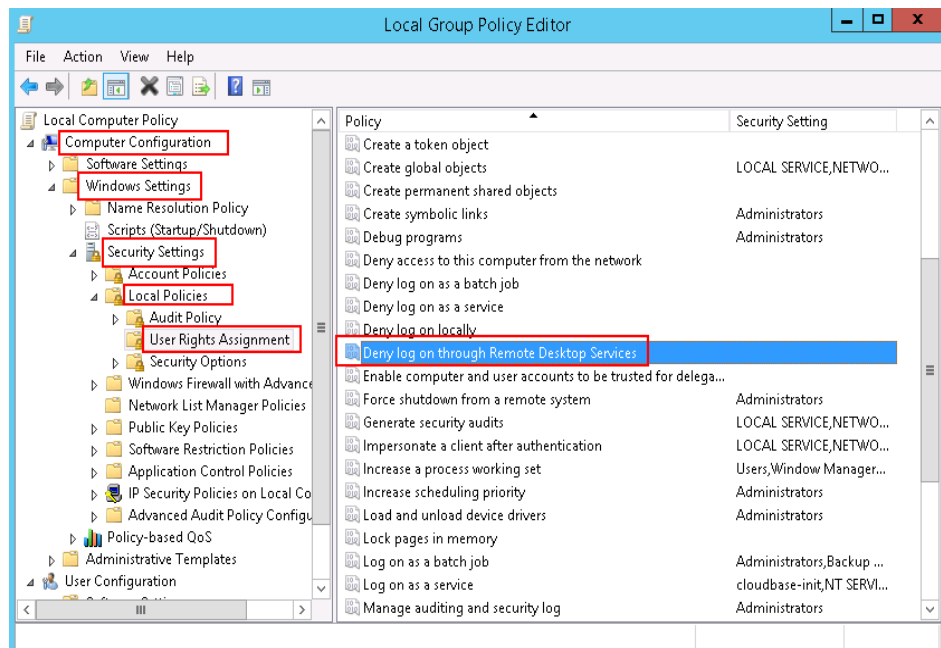
- a. Locate and double-click **Allow log on through Remote Desktop Services**. Ensure that **Administrators** and **Remote Desktop Users** have been added.

**Figure 14-18** Allow log on through Remote Desktop Services properties



- b. Locate and double-click **Deny log on through Remote Desktop Services**. If the administrator account exists, delete it.

**Figure 14-19** Deny log on through Remote Desktop Services properties




## 14.3.11 Why Does the System Display No Remote Desktop License Servers Available to Provide a License When I Log In to a Windows ECS?

### Symptom

An error message is displayed indicating that there are no Remote Desktop License Servers available to provide a license and asks you to contact the administrator.

**Figure 14-20** No Remote Desktop License Servers available to provide a license



### Possible Causes

You have installed the Remote Desktop Session Host.

The grace period for Remote Desktop Services is 120 days. If you do not pay for it when the period expires, the service will stop. Windows allows a maximum of two users (including the local user) in remote desktop connections. To allow the access of more users, install the Remote Desktop Session Host and configure the desired number of authorized users. However, installing the Remote Desktop Session Host will automatically revoke the original two free connections. This leads to the preceding fault if desired number of authorized users has not been configured.


### Precautions

- The operations described in this section apply to the ECSs running a Windows Server 2008 or Windows Server 2012.
- The ECS must be restarted during the operation, which may interrupt services. Back up data before restarting the ECS.

### Windows Server 2008


1. Log in to the Windows ECS using VNC available on the management console.
2. Open **Server Manager**, right-click **Remote Desktop Services** under **Roles**, and choose **Remove Roles** from the shortcut menu.

Figure 14-21 Deleting roles



3. In the displayed dialog box, deselect **Remote Desktop Session Host** and keep clicking **Next** till you finish the operation.

Figure 14-22 Deselecting Remote Desktop Session Host




4. Click **Delete**.
5. Restart the ECS.



## Windows Server 2012

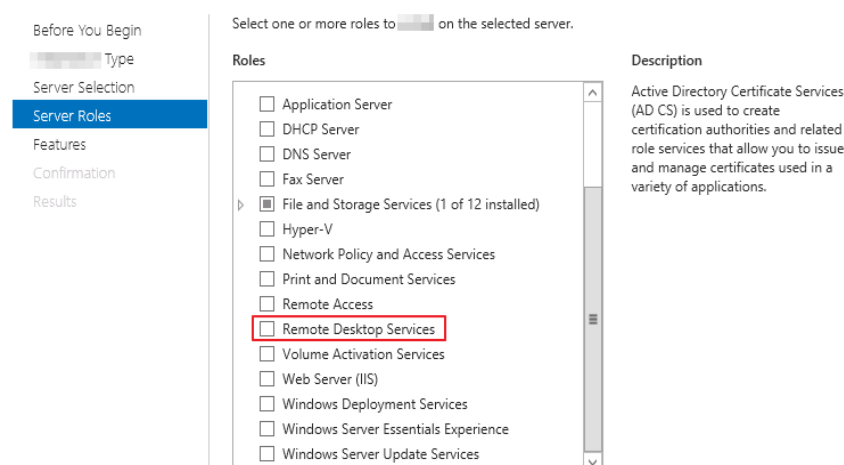
1. Log in to the Windows ECS using VNC available on the management console.
2. Open **Server Manager**, choose **Manage > Remove Roles and Features**, and click **Next**.

Figure 14-23 Deleting roles and features



3. Select the destination server and click **Next**.
4. Deselect **Remote Desktop Services**.

Figure 14-24 Deselecting Remote Desktop Services




5. Click **Delete**.
6. Restart the ECS.

### 14.3.12 Why Does the System Display Error Code 0x112f When I Log In to a Windows ECS?

#### Symptom

When you log in to a Windows ECS, the system displays error code 0x112f, as shown in [Figure 14-25](#).

**Figure 14-25** Error message (code: 0x112f)



## Possible Causes

The ECS memory is insufficient.

## Solution

- Method 1 (recommended)  
Modify the ECS specifications to increase the vCPUs and memory size. For instructions about how to modify ECS specifications, see [General Operations](#).
- Method 2  
Enable virtual memory on the ECS to obtain its idle memory.  
For instructions about how to enable virtual memory, see [How Can I Enable Virtual Memory on a Windows ECS?](#)

### NOTE


This method will deteriorate the disk I/O performance, so use this method only when necessary.

## 14.3.13 Why Does the System Display Error Code 0x1104 When I Log In to a Windows ECS?

### Symptom

The system displays an error message indicating that a protocol error (code: 0x1104) is detected when you use MSTSC to access an ECS running Windows Server 2008.

**Figure 14-26** Protocol error (code: 0x1104)



## Possible Causes

- Port 3389 of the security group on the ECS is disabled.
- The firewall on the ECS is disabled.
- Port 3389 on the ECS is used by other processes.
- The Remote Desktop Session Host is incorrectly configured.


## Solution

**Step 1** Check security group settings.

Check whether port 3389 is allowed in inbound direction. If it is allowed, go to [Step 2](#).


**Step 2** Check whether the firewall is disabled:

1. Log in to the Windows ECS.
2. Click the Windows icon in the lower left corner of the desktop and choose **Control Panel > Windows Firewall**.



3. Click **Turn Windows Firewall on or off**.

View and set the firewall status.

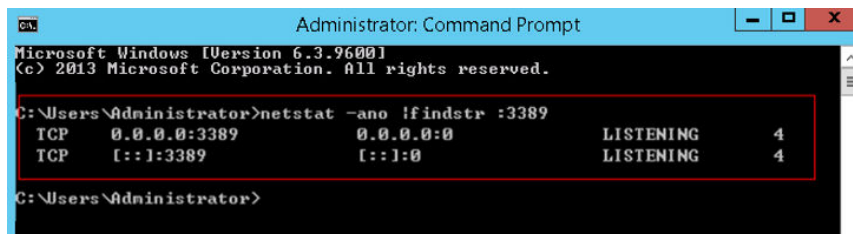


If the firewall is enabled, go to [Step 3](#).

**Step 3** Log in to the ECS using VNC and check the port.

1. Open the **cmd** window and run the following command:  
**netstat -ano |findstr: 3389**

**Figure 14-27** Checking port 3389



```
Administrator: Command Prompt
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>netstat -ano |findstr: 3389
TCP 0.0.0.0:3389 0.0.0.0:0 LISTENING 4
TCP [::]:3389 [::]:0 LISTENING 4

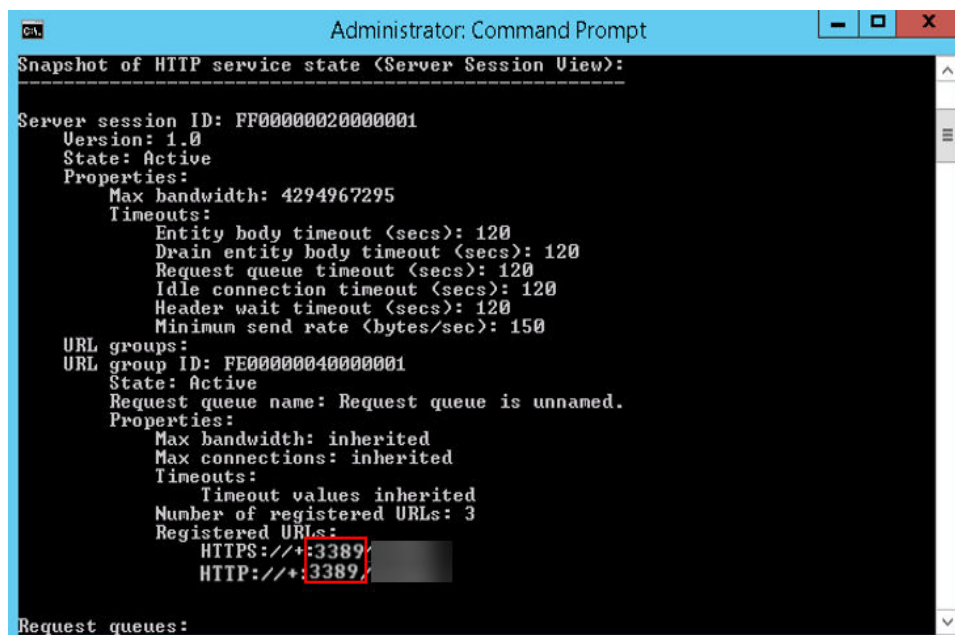
C:\Users\Administrator>
```

As shown in **Figure 14-27**, port 3389 is used by the process with ID of 4.

2. Open Task Manager and find the process with ID of 4 is the System process.
3. Generally, the IIS and SQL Server run as the System process. Run the following HTTP command for further check.

**netsh http show servicestate**

**Figure 14-28** Checking System process



```
Administrator: Command Prompt
Snapshot of HTTP service state (Server Session View):

Server session ID: FF00000020000001
Version: 1.0
State: Active
Properties:
 Max bandwidth: 4294967295
 Timeouts:
 Entity body timeout (secs): 120
 Drain entity body timeout (secs): 120
 Request queue timeout (secs): 120
 Idle connection timeout (secs): 120
 Header wait timeout (secs): 120
 Minimum send rate (bytes/sec): 150
 URL groups:
 URL group ID: FE00000040000001
 State: Active
 Request queue name: Request queue is unnamed.
 Properties:
 Max bandwidth: inherited
 Max connections: inherited
 Timeouts:
 Timeout values inherited
 Number of registered URLs: 3
 Registered URLs:
 HTTPS://+:3389
 HTTP://+:3389

Request queues:
```

4. If port 3389 is used by HTTP protocols, it indicates that the port is used by IIS.
5. Enter **http://127.0.0.1:3389** in the address box of the browser and press **Enter**. Check whether the website can be visited normally.
6. Change the port used by IIS and restart IIS.


**Step 4** If no error occurs during the preceding steps, go to step **Step 5** to check whether error 0x1104 is caused by the configuration of Remote Desktop Session Host.

**Step 5** Check the remote desktop session host configuration.

1. Log in to the ECS using VNC.


2. Open the **cmd** window and enter **gpedit.msc**.
3. Click **OK** to start Local Group Policy Editor.
4. Choose **Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services**.

Figure 14-29 Remote Desktop Services




5. **Remote Desktop Session Host > Security.**

Figure 14-30 Remote (RDP) Connection requires the use of the specified security layer




6. Set **Require use of specific security layer for remote (RDP) connections** to **Enabled** and **Security layer** to **RDP**.

Figure 14-31 Setting security layer



7. Click **OK**.
8. After the configuration is complete, open the **cmd** window.
9. Run the following command to update the group policy:  
**gpupdate**

Figure 14-32 Updating the group policy



----End

## 14.3.14 Why Does the System Display Error Code 122.112... When I Log In to a Windows ECS?

### Symptom

The system displays error 122.112... when you use RDC to locally access an ECS running Windows Server 2012. The ECS is frequently disconnected and the Windows login process is unexpectedly interrupted.

### Possible Causes

1. System resources are insufficient or unavailable.
2. The services cannot be started.

### Solution

**Step 1** Check system logs.




1. Log in to the ECS using VNC.
2. Click  to start the service manager and choose **Administrative Tools > Event Viewer > Windows Logs > System > Filter Current Logs**.

Figure 14-33 Event viewer



3. In the **Event Level** pane, select event levels.

Figure 14-34 Filtering logs



4. Search for login logs.

**Step 2** Check the usage of host resources.

1. Choose **Start > Task Manager > Performance**.
2. Check usage of CPU and memory.

**Step 3** Check whether the purchased Windows ECS is with 1 vCPU and 1 GB of memory.

If it is, change the flavor or stop unnecessary processes.

----End

### 14.3.15 Why Does the System Display Invalid Certificate or Associated Chain When I Log In to a Windows ECS from a Mac?


#### Symptom

When you use Microsoft Remote Desktop for Mac to remotely access a Windows ECS, the system displays invalid certificate or associated chain.



**Figure 14-35** Microsoft Remote Desktop for Mac

Due to the particularity of the Mac system, you need to perform internal configurations on Mac and the Windows ECS to ensure successful remote connection. When you log in to the Windows ECS using Microsoft Remote Desktop for Mac, the system displays an error message indicating that the certificate or associated chain is invalid.


**Figure 14-36** Invalid certificate or associated chain

## Possible Causes


The group policy setting is incorrect on the ECS.

## Procedure


1. On the menu bar in the upper left corner, choose **RDC > Preferences** to open the preference setting page of the Microsoft Remote Desktop.

**Figure 14-37** Preferences setting

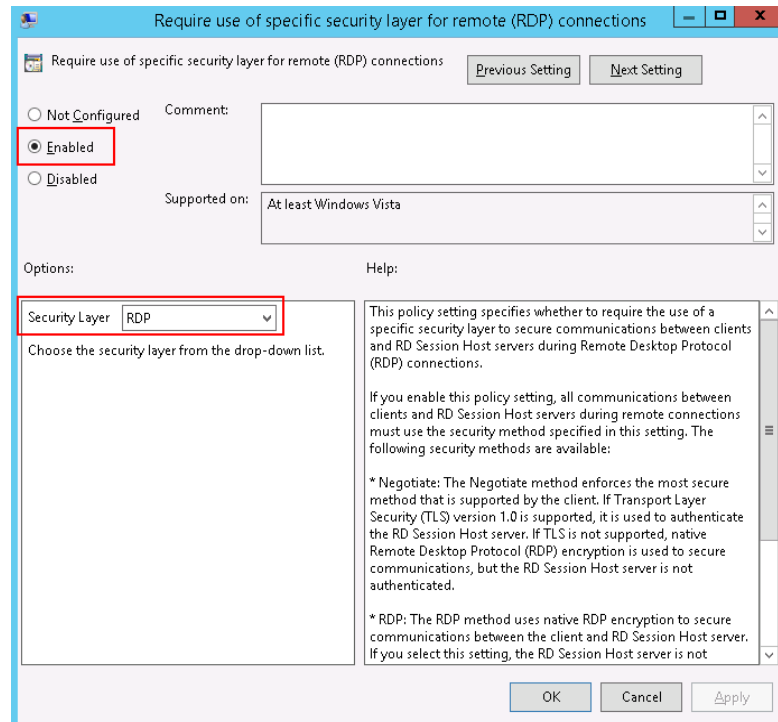
2. Select **Security** and modify the parameter settings according the following figure.

**Figure 14-38** Security setting


3. Remotely connect to the Windows ECS again. If the error message **Invalid certificate or associated chain** is still displayed, go to 4.
4. Log in to the Windows ECS using VNC.
5. Press **Win+R** to start the **Open** text box.
6. Enter **gpedit.msc** to access the Local Group Policy Editor.
7. In the left navigation pane, choose **Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Security**.

**Figure 14-39** Remote Desktop Session Host

8. Modify the following parameters as prompted:
  - Enable **Require use of specific security layer for remote (RDP) connections**.

**Figure 14-40** Require use of specific security layer for remote (RDP) connections

- Disable **Require user authentication for remote connections by using Network Level Authentication**.

**Figure 14-41** Remote connection authentication


9. Close the group policy editor and restart the ECS.

## 14.3.16 Why Is My Remote Session Interrupted by a Protocol Error?

### Symptom

An error message is displayed indicating that the remote session will be disconnected because of a protocol error.

**Figure 14-42** Protocol error



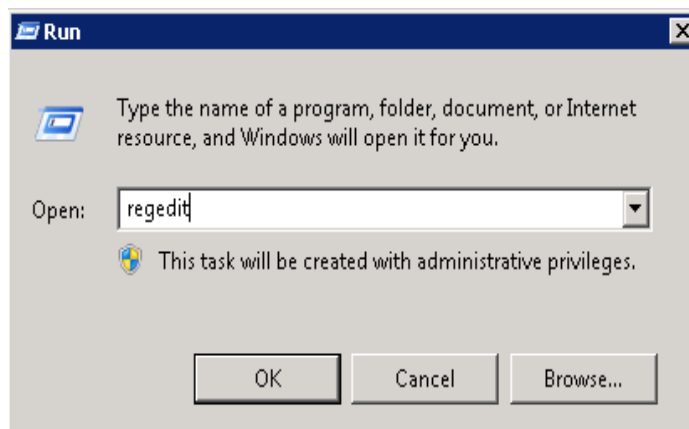
### Possible Causes

The registry subkey Certificate is damaged.

### Solution

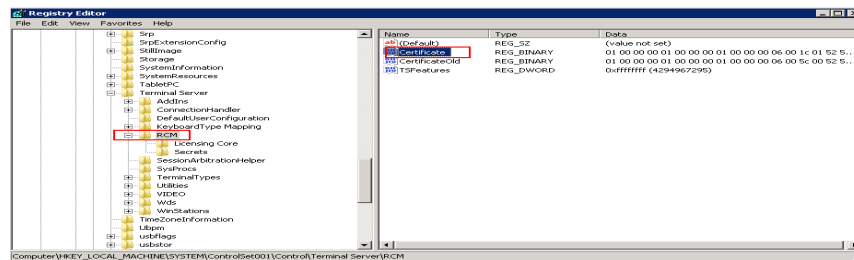
1. In the **Run** dialog box, enter **regedit** and click **OK** to open the registry editor.

**Figure 14-43** Opening the registry editor




2. Choose **HKEY\_LOCAL\_MACHINE > SYSTEM > ControlSet001 > Control > Terminal Server > RCM**.
3. Delete **Certificate**.

Figure 14-44 Deleting Certificate



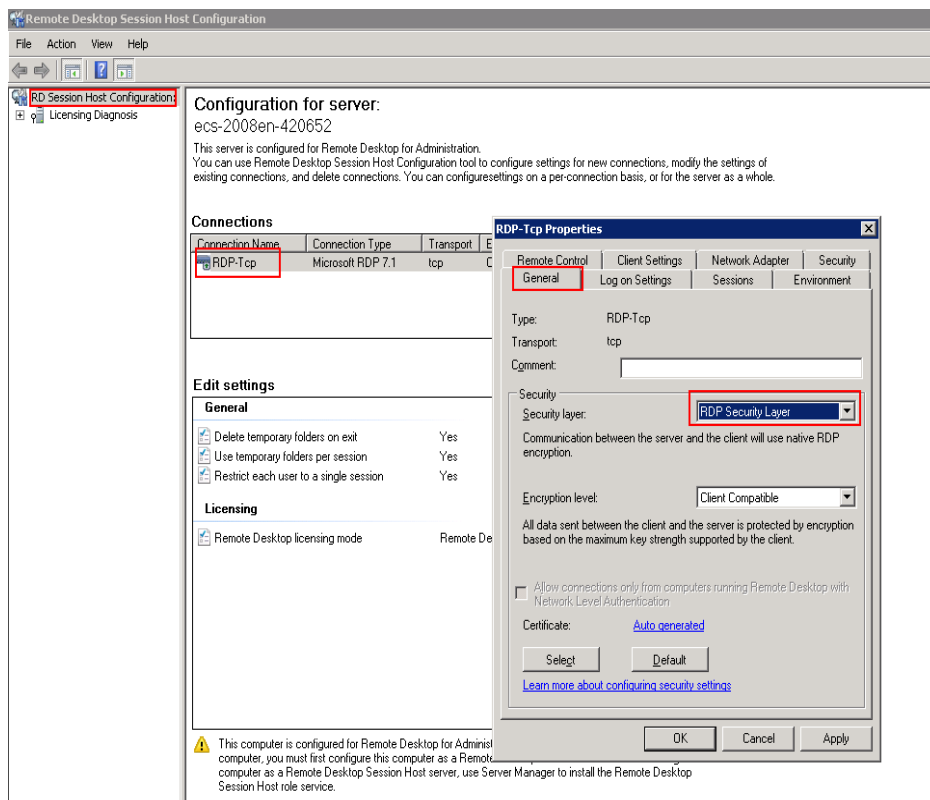
4. Restart the ECS.
5. Choose **Start > Administrative Tools > Remote Desktop Services > Remote Desktop Session Host Configuration**.

Figure 14-45 Opening Remote Desktop Session Host Configuration



6. Right-click **RDP-Tcp** and choose **Properties**. In the displayed dialog box, click **General** and set **Security layer** to **RDP Security Layer**.

Figure 14-46 RDP-Tcp properties




### 14.3.17 Why Am I Seeing an Error Message That Says Identity of Remote Computer Cannot be Verified When I Log In to a Windows ECS?

#### Symptom

An error message is displayed indicating that the identity of the remote computer cannot be verified. You are required to enter the password and log in again.

Figure 14-47 Protocol error



## Possible Causes

Security software installed on the ECS prevents logins from unknown IP addresses.

## Solution

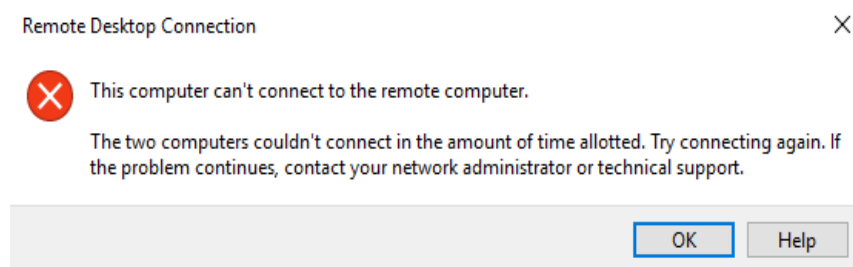
- Uninstall the security software.
- Open the security software and enable the default login mode.

## 14.3.18 Why Am I Seeing An Error Message That Says The Two Computers Couldn't Be Connected in the Amount of Time Allotted When I Log In to a Windows ECS?

### Symptom

An error message is displayed indicating that the computer cannot connect to the remote computer in the amount of time allotted.

**Figure 14-48** Error message




## Solution

1. On the local computer, click on the **Start** icon, type **cmd** into the box, and run the command as an administrator.
2. Run the **netsh winsock reset** command.
3. Restart the local computer as prompted and reconnect to the ECS.

## 14.3.19 Why Am I Seeing an Error Message That Says User Account is not Authorized for Remote Login When I Log In to a Windows ECS?

### Symptom

An error message is displayed indicating that the connection is denied because the user account is not authorized for remote login.

**Figure 14-49** Error message


## Possible Causes

The remote desktop connection permissions have been incorrectly configured.

## Solution

**Step 1** Check remote desktop permissions on the ECS.


1. In the **Run** dialog box, enter **secpol.msc** and click **OK** to open **Local Security Policy**.
2. Choose **Local Policies > User Rights Assignment > Allow log on through Remote Desktop Services**.

**Figure 14-50** Local security policy

3. Check whether there are user groups or users that have been granted the remote login permission.  
If not, add required users or groups.



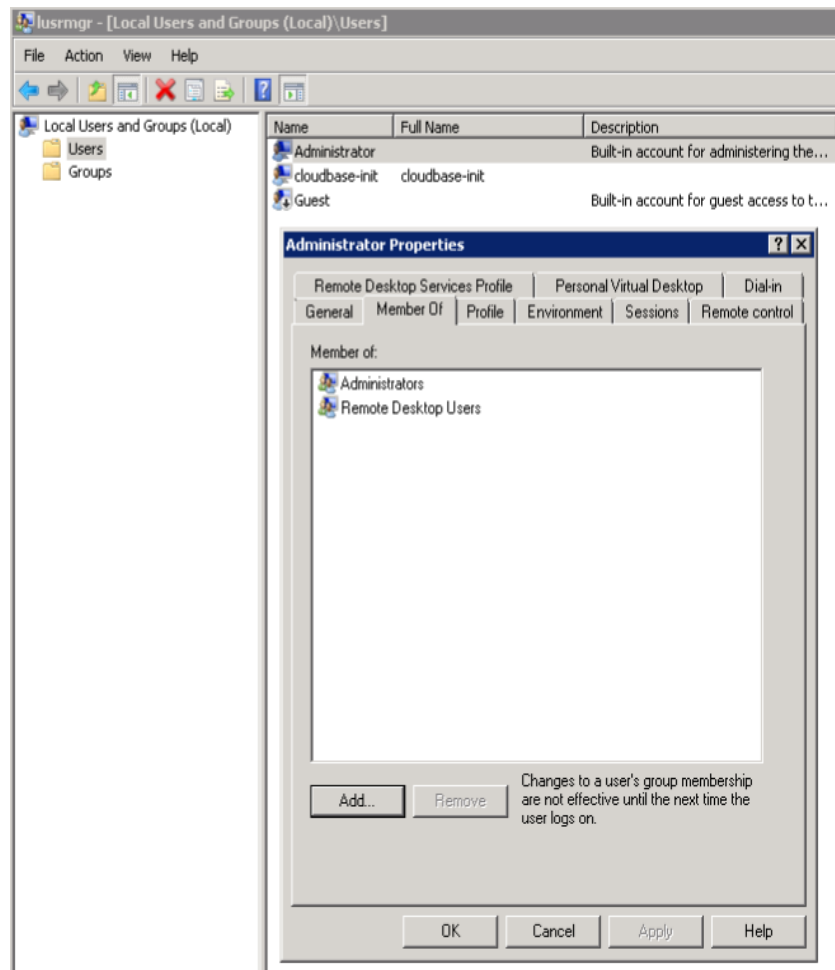
**Figure 14-51** Allow log on through Remote Desktop Services properties



**Step 2** Check the target user group.

1. Open the **Run** dialog box, enter **lusrmgr.msc**, and click **OK** to open **Local Users and Groups**.
2. Double-click **Users** on the left.
3. Double-click the name of the user to whom the login error message was displayed.
4. In the displayed dialog box, click the **Member Of** tab. Ensure that the user belongs to the user group that is assigned with the remote login permission in [Step 2.2](#).


**Figure 14-52** Checking the target user group



**Step 3** Check the remote desktop session host configuration.

1. In the **Run** dialog box, enter **tsconfig.msc** and click **OK** to open **Remote Desktop Session Host Configuration**.
2. Double-click **RDP-Tcp** or other connections added by a user under **Connections** and click the **Security** tab.

Figure 14-53 Security




3. Check whether there are user groups or users that have been granted the remote login permission under **Group or user names**.  
If not, add required users or groups.
  4. Restart the ECS or run the following commands in the CLI to restart the Remote Desktop Services:  
**net stop TermService**  
**net start TermService**
- End

## 14.3.20 Why Does My Remote Desktop Session End Because Another User Logs In When I Log In to a Windows ECS?

### Symptom

An error message is displayed indicating that your remote desktop session has ended because another user has connected to the remote computer.


Figure 14-54 Ended remote desktop session



## Windows Server 2008


1. Choose **Start > Administrative Tools > Remote Desktop Services > Remote Desktop Session Host Configuration**.

Figure 14-55 Remote Desktop Session Host Configuration



2. Double-click **Restrict each user to a single session** and deselect **Restrict each user to a single session**, and click **OK**.

Figure 14-56 Modifying the configuration




## Windows Server 2012

1. Choose **Start > Run**. In the **Run** dialog box, enter **gpedit.msc** and click **OK** to start Local Group Policy Editor.


2. Choose **Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Connections**.

Figure 14-57 Connections



3. Double-click **Restrict Remote Desktop Services users to a single Remote Desktop Services session**, change the value to **Disabled**, and click **OK**.

Figure 14-58 Modifying the configuration



4. Run **gpupdate/force** to update the group policy.

## 14.3.21 Why Does an ECS Fail to Be Remotely Connected Using RDP and Internal Error Code 4 Is Displayed?

### Symptom

An internal error is displayed when you log in to a Windows ECS and you fail to connect to the ECS remotely. Generally, this problem occurs because the Remote Desktop Services is busy.

### Possible Causes

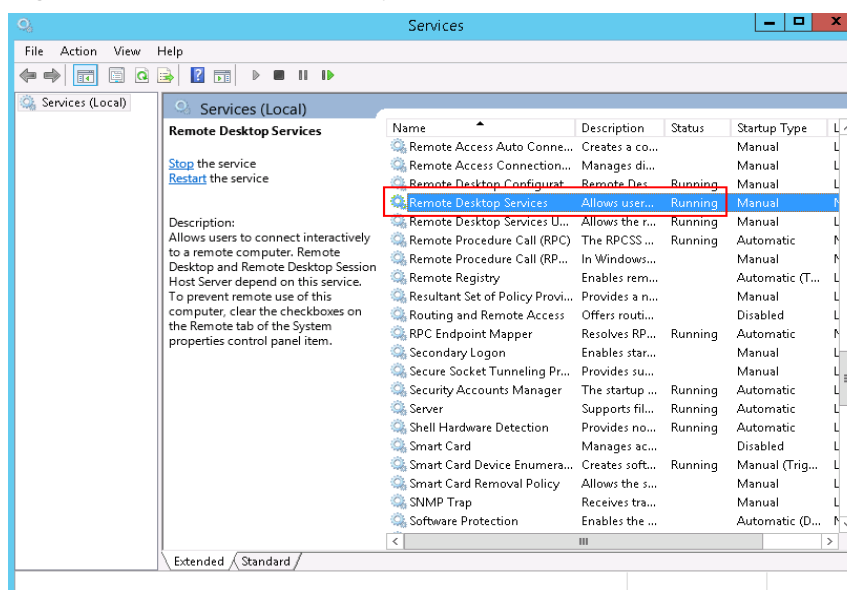
The Remote Desktop Services is busy.

The remote desktop is disconnected after login but is not logged out. To prevent this problem, log out of the ECS if you do not need to remotely connect to it.

### Solution

1. Use VNC provided by the management console to remotely log in to the ECS.
2. Open the Windows search box, enter **services**, and select **Services**.
3. In the **Services** window, restart **Remote Desktop Services**. Ensure that **Remote Desktop Services** is in the **Running** status.

Figure 14-59 Remote Desktop Services



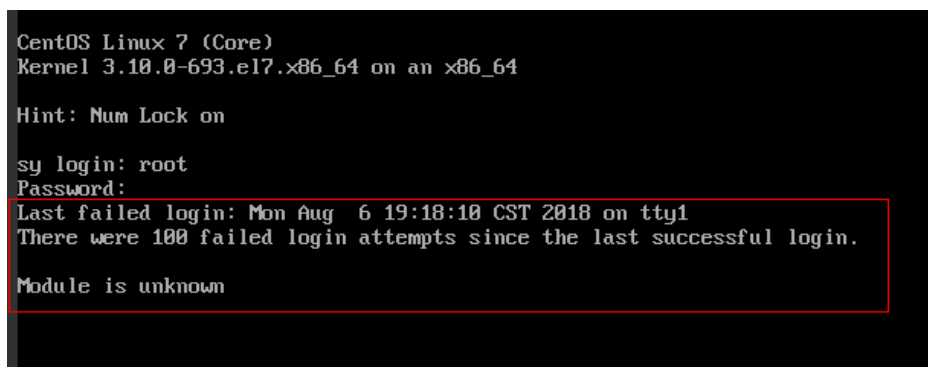
4. Remotely connect to the ECS again.  
If the connection still fails, run the `cmd` command on the local server as the administrator, run the `netsh winsock reset` command to restore the default network connection configurations, and then retry the remote connection.

## 14.3.22 Why Am I Seeing the Error Message "Module is unknown" When I Remotely Log In to a Linux ECS?

### Symptom

When you attempt to remotely log in to a Linux ECS, the system displays the error message "Module is unknown".

**Figure 14-60** Module is unknown



```
CentOS Linux 7 (Core)
Kernel 3.10.0-693.el7.x86_64 on an x86_64

Hint: Num Lock on

sy login: root
Password:
Last failed login: Mon Aug 6 19:18:10 CST 2018 on tty1
There were 100 failed login attempts since the last successful login.

Module is unknown
```

### NOTE

- To resolve this issue, restart the ECS and enter the rescue mode.
- Restarting the ECS may interrupt services. Exercise caution when performing this operation.

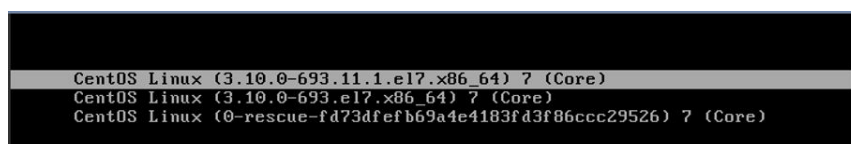
### Root Cause

The file in the `/etc/pam.d/` directory was modified by mistake.

### Solution

1. Enter the single-user mode.  
The following uses CentOS 7 as an example:
  - a. Restart the ECS and click **Remote Login**.
  - b. Click **Ctrl+Alt+Del** in the upper part of the remote login panel to restart the ECS.
  - c. Press the up arrow key to prevent automatic system startup. When the kernels are displayed, press **e** to enter the editing mode.

**Figure 14-61** Entering the kernel editing mode



```
CentOS Linux (3.10.0-693.11.1.el7.x86_64) ? (Core)
CentOS Linux (3.10.0-693.el7.x86_64) ? (Core)
CentOS Linux (0-rescue-fd73dfefb69a4e4183fd3f86ccc29526) ? (Core)
```

**NOTE**

The grub file is encrypted by Euler images by default. Before entering the edit mode, you need to contact the administrator to obtain username and password.

- d. Locate the row containing **linux16** and delete the parameters you do not require.
- e. Change **ro** to **rw** for mounting the root partition with read-write permissions.
- f. Add **rd.break** and press **Ctrl+X**.

**Figure 14-62** Before the modification

```
insmod ext2
set root='hd0,msdos1'
if [x${feature_platform_search_hint} = xy]; then
 search --no-floppy --fs-uuid --set=root --hint='hd0,msdos1' b13ee9c\
8-0ef0-4159-9b90-fc47bde0d464
else
 search --no-floppy --fs-uuid --set=root b13ee9c8-0ef0-4159-9b90-fc47\
bde0d464
fi
linux16 /boot/vmlinuz-3.10.0-327.62.59.83.h162.x86_64 root=UUID=b13ee9\
c8-0ef0-4159-9b90-fc47bde0d464 ro crash_kexec_post_notifiers softlockup_panic=\
1 panic=3 reserve_kbox_mem=16M nmi_watchdog=1 rd.shell=0 net.ifnames=0 spectre\
_v2=off nopti noibrs noibpb crashkernel=auto LANG=en_US.UTF-8
initrd16 /boot/initramfs-3.10.0-327.62.59.83.h162.x86_64.img
```

**Figure 14-63** After the modification

```
insmod ext2
set root='hd0,msdos1'
if [x${feature_platform_search_hint} = xy]; then
 search --no-floppy --fs-uuid --set=root --hint='hd0,msdos1' b13ee9c\
8-0ef0-4159-9b90-fc47bde0d464
else
 search --no-floppy --fs-uuid --set=root b13ee9c8-0ef0-4159-9b90-fc47\
bde0d464
fi
linux16 /boot/vmlinuz-3.10.0-327.62.59.83.h162.x86_64 root=UUID=b13ee9\
c8-0ef0-4159-9b90-fc47bde0d464 rw rd.break
initrd16 /boot/initramfs-3.10.0-327.62.59.83.h162.x86_64.img
```

- g. Run the following command to go to the **/sysroot** directory:  
**# chroot /sysroot**
2. Run the following command to view the system log for error files:  
**grep Module /var/log/messages**

**Figure 14-64** System log

```
Aug 6 18:00:09 sy login: pam_succeed_if(login:auth): requirement "uid >= 1000" not met by user "root"
Aug 6 18:00:11 sy login: FAILED LOGIN 1 FROM tty1 FOR root, Authentication failure
Aug 6 18:00:15 sy login: pam_unix(login:session): session opened for user root by LOGIN(uid=0)
Aug 6 18:00:15 sy login: Module is unknown
Aug 6 18:10:41 sy login: PAM unable to dlopen(/lib/security/pam_limits.so): /lib/security/pam_limits.so: cannot open shared obj
ect file: No such file or directory
Aug 6 18:10:41 sy login: PAM adding faulty module: /lib/security/pam_limits.so
Aug 6 18:10:44 sy login: pam_unix(login:session): session opened for user root by LOGIN(uid=0)
Aug 6 18:10:44 sy login: Module is unknown
```

3. Comment out or modify the error line in the error files displayed in the system log.  
**vi /etc/pam.d/login**



**Figure 14-65** Modifying the error information

```
session required pam_selinux.so open
session required pam_namespace.so
session optional pam_keyinit.so force revoke
session include system-auth
session include postlogin
-session optional pam_ck_connector.so
session required /lib/security/pam_limits.so
```

- Restart the ECS and try to log in to it again.

**NOTE**

- To view the modification records and check whether the modification is caused by unintended actions, run the following command:

```
vi /root/.bash_history
```

Search for the keyword **vi** or **login**.

- Do not modify the files in the `/etc/pam.d/` directory. Run the following command for details about pam:

```
man pam.d
```

## 14.3.23 What Should I Do If Error Message "Permission denied" Is Displayed When I Remotely Log In to a Linux ECS?

### Symptom

When I attempted to remotely log in to a Linux ECS, the system displayed error Message "Permission denied".

**Figure 14-66** Permission denied

```
CentOS Linux 7 (Core)
Kernel 3.10.0-693.11.1.el7.x86_64 on an x86_64

ecs-ams-03 login: :
Password:

Permission denied
_
```

**NOTE**

- To resolve this issue, you are required to restart the ECS and enter the rescue mode.
- Restarting the ECS may interrupt services. Exercise caution when performing this operation.

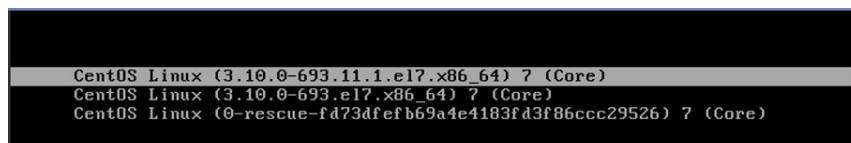
### Root Cause

The **nofile** parameter in `/etc/security/limits.conf` is used to set the maximum number of files that can be opened in the system. If the value is greater than the **fs.nr\_open** value (**1048576** by default) set in **PermissionDenied.png**, a login verification error will occur, leading to "Permission denied".

## Solution

1. Enter the single-user mode.  
The following uses CentOS 7 as an example:
  - a. Restart the ECS and click **Remote Login**.
  - b. Click **Ctrl+Alt+Del** in the upper part of the remote login panel to restart the ECS.
  - c. Press the up arrow key to prevent automatic system startup. When the kernels are displayed, press **e** to enter the editing mode.

**Figure 14-67** Entering the kernel editing mode



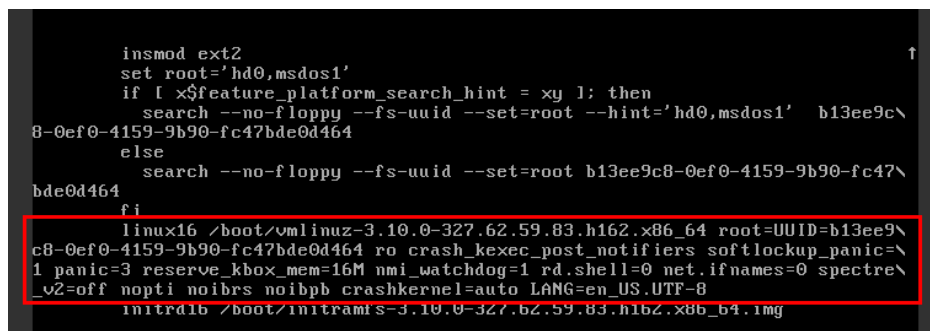
```
CentOS Linux (3.10.0-693.11.1.el7.x86_64) ? (Core)
CentOS Linux (3.10.0-693.el7.x86_64) ? (Core)
CentOS Linux (0-rescue-fd73dfefb69a4e4183fd3f86ccc29526) ? (Core)
```

### NOTE

The grub file is encrypted by Euler images by default. Before entering the edit mode, you need to contact the administrator to obtain username and password.

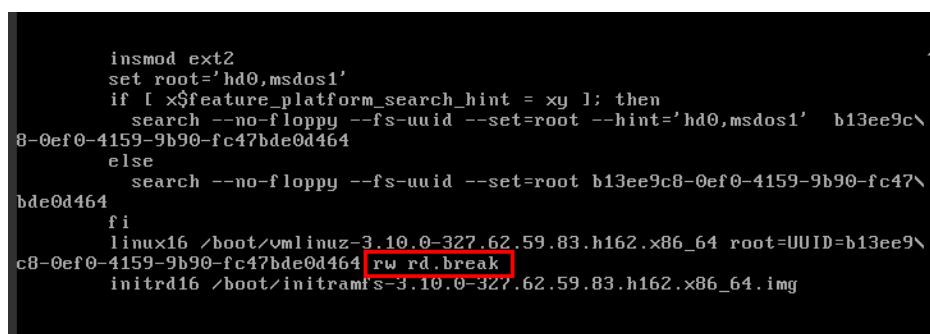
- d. Locate the row containing **linux16** and delete the parameters you do not require.
- e. Change **ro** to **rw** for mounting the root partition with read-write permissions.
- f. Add **rd.break** and press **Ctrl+X**.

**Figure 14-68** Before the modification




```
insmod ext2
set root='hd0,msdos1'
if [x${feature_platform_search_hint} = xy]; then
 search --no-floppy --fs-uuid --set=root --hint='hd0,msdos1' b13ee9c\
8-0ef0-4159-9b90-fc47bde0d464
else
 search --no-floppy --fs-uuid --set=root b13ee9c8-0ef0-4159-9b90-fc47\
bde0d464
fi
linux16 /boot/vmlinuz-3.10.0-327.62.59.83.h162.x86_64 root=UUID=b13ee9\
c8-0ef0-4159-9b90-fc47bde0d464 ro crash_kexec_post_notifiers softlockup_panic=\
1 panic=3 reserve_kbox_mem=16M nmi_watchdog=1 rd.shell=0 net.ifnames=0 spectre\
_v2=off nopti noibrs noibpb crashkernel=auto LANG=en_US.UTF-8
initrd16 /boot/initramfs-3.10.0-327.62.59.83.h162.x86_64.img
```

**Figure 14-69** After the modification



```
insmod ext2
set root='hd0,msdos1'
if [x${feature_platform_search_hint} = xy]; then
 search --no-floppy --fs-uuid --set=root --hint='hd0,msdos1' b13ee9c\
8-0ef0-4159-9b90-fc47bde0d464
else
 search --no-floppy --fs-uuid --set=root b13ee9c8-0ef0-4159-9b90-fc47\
bde0d464
fi
linux16 /boot/vmlinuz-3.10.0-327.62.59.83.h162.x86_64 root=UUID=b13ee9\
c8-0ef0-4159-9b90-fc47bde0d464 rw rd.break
initrd16 /boot/initramfs-3.10.0-327.62.59.83.h162.x86_64.img
```

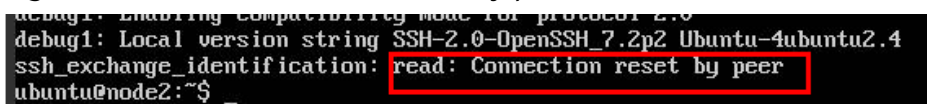
- g. Run the following command to go to the `/sysroot` directory:  
`# chroot /sysroot`
2. Run the following command to view the `fs.nr_open` value:  
`sysctl fs.nr_open`
3. Change the `nofile` value in `/etc/security/limits.conf` so that the value is smaller than the `fs.nr_open` value obtained in 2.  
`vi /etc/security/limits.conf`  
 **NOTE**  
`limits.conf` is the `pam_limits.so` configuration file of Linux Pluggable Authentication Module (PAM). For more details, run the following command:  
`man limits.conf`
4. Restart the ECS and try to log in to it again.

## 14.3.24 What Should I Do If Error Message "read: Connection reset by peer" Is Displayed When I Remotely Log In to a Linux ECS?

### Symptom

When I attempted to remotely log in to a Linux ECS, the system displayed error message "read: Connection reset by peer".

**Figure 14-70** read: Connection reset by peer



```
debug1: Local version string SSH-2.0-OpenSSH_7.2p2 Ubuntu-4ubuntu2.4
ssh_exchange_identification: read: Connection reset by peer
ubuntu@node2:~$
```

### Possible Causes

- The remote login port is not permitted in the security group.
- The firewall is enabled on the ECS, but the remote login port is blocked by the firewall.

### Solution

Perform the following operations for troubleshooting:

- **Check security group rules.**
  - Inbound: Add the remote login port. The default port 22 is used as an example.
  - Outbound: Outbound rules allow network traffic to be out of specified ports.
- **Add a port to the ECS firewall exception.**

The following uses Ubuntu as an example:

  - a. Run the following command to view the firewall status:  
`sudo ufw status`

The following information is displayed:

```
Status: active
```

- b. Add a port to the firewall exception, taking the default port 22 as an example.

```
ufw allow 22
```

```
Rule added
```

```
Rule added (v6)
```

- c. Run following command to check the firewall status again:

```
sudo ufw status
```

```
Status: active
```

| To      | Action | From          |
|---------|--------|---------------|
| --      | ----   | ----          |
| 22      | ALLOW  | Anywhere      |
| 22 (v6) | ALLOW  | Anywhere (v6) |

Try to remotely log in to the ECS again.

## 14.3.25 Why Am I Seeing the Error Message "Access denied" When I Remotely Log In to a Linux ECS?

### Symptom

When you attempt to remotely log in to a Linux ECS, the system displays the error message "Access denied".

### Possible Causes

- Incorrect username or password.
- A policy that denies logins from user **root** is enabled on the SSH server.

### Solution

- If the username or password is incorrect, Check the username and password.
- If a policy that denies logins from user **root** is enabled on the SSH server,
  - a. Edit the `/etc/ssh/sshd_config` file and check the following settings to ensure that the SSH logins from user **root** are allowed:

```
PermitRootLogin yes
```
  - b. Restart SSH.
    - CentOS 6

```
service sshd restart
```
    - CentOS 7

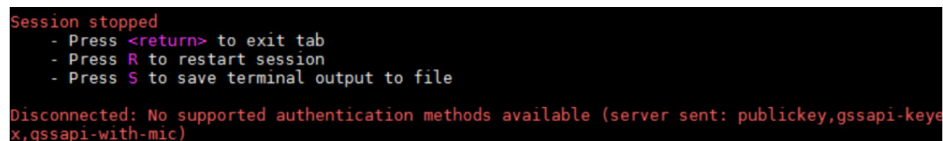
```
systemctl restart sshd
```

## 14.3.26 What Should I Do If Error Message "Disconnected: No supported authentication methods available" Is Displayed When I Remotely Log In to a Linux ECS?

### Symptom

When I attempted to remotely log in to a Linux ECS, the system displayed error message "Disconnected: No supported authentication methods available".

**Figure 14-71** No supported authentication methods available



```
Session stopped
- Press <return> to exit tab
- Press R to restart session
- Press S to save terminal output to file
Disconnected: No supported authentication methods available (server sent: publickey,gssapi-keyex,gssapi-with-mic)
```

### Possible Causes

A policy that denies password-authenticated logins is enabled on the SSH server.

### Solution

1. Open the `/etc/ssh/sshd_config` file and check the following settings:  
`vi /etc/ssh/sshd_config`
1. Modify the following settings:  
Change **PasswordAuthentication no** to **PasswordAuthentication yes**.  
Alternatively, delete the comment tag (`#`) before **PasswordAuthentication yes**.
2. Restart SSH.
  - CentOS 6  
`service sshd restart`
  - CentOS 7  
`systemctl restart sshd`

## 14.4 How Do I Handle Error Messages Displayed on the Management Console?

### Symptom

This section helps you resolve the following issues:

- An error message was displayed on the management console after you performed ECS-related operations.
- An error code was displayed after you used an ECS API (see *Elastic Cloud Server API Reference*).

## Background

After you perform ECS-related operations on the management console, the system displays the request status on the **Elastic Cloud Server** page. You can determine the request execution status based on the information displayed in the request status.

- If the operation request is executed, the system automatically clears the task prompt.
- If an error occurs during the request execution, the system displays an error code and its description in the taskbar.

## Solution

If an error occurs, check the error code and perform the corresponding operations listed in [Table 14-2](#).

**Table 14-2** Error codes and solution suggestions

| Error Code | Message Displayed on the Management Console                                                                            | Solution Suggestion                                                                                                                                                                                           |
|------------|------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Ecs.0000   | Request error. Try again later or contact the administrator.                                                           | Adjust the request structure as requested in the <i>Elastic Cloud Server API Reference</i> .                                                                                                                  |
| Ecs.0001   | The maximum number of ECSs or EVS disks has been reached. Contact the administrator and request an ECS quota increase. | Contact the administrator and request an ECS quota increase.<br><b>NOTE</b><br>Before requesting for increasing your ECS quota, consider the number of to-be-added ECSs, vCPUs, and memory capacity required. |
| Ecs.0003   | You do not have the permission. Contact the administrator to check your account information.                           | Contact the administrator to check your account information.                                                                                                                                                  |
| Ecs.0005   | System error. Try again later or contact the administrator.                                                            | Adjust the request structure as directed in <i>Elastic Cloud Server API Reference</i> .                                                                                                                       |
| Ecs.0010   | The private IP address is in use. Select an available IP address for ECS creation.                                     | Use an idle IP address for ECS creation.                                                                                                                                                                      |

| Error Code        | Message Displayed on the Management Console                                                                                              | Solution Suggestion                                                                                                                  |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| Ecs.0011          | Invalid password.<br>Change the password to make it meet the password complexity requirements, and perform the required operation again. | Input a password that meets password complexity requirements. Then, initial the request again.                                       |
| Ecs.0012          | Insufficient IP addresses in the subnet. Release IP addresses in the subnet or select another subnet for ECS creation.                   | Release IP addresses in the subnet or select another subnet for ECS creation.                                                        |
| Ecs.0013          | Contact the administrator and request an EIP quota increase.                                                                             | Contact the administrator and request an EIP quota increase.                                                                         |
| Ecs.0015          | The disk of this type is not supported by the ECS.                                                                                       | Select a proper disk and attach it to the ECS.                                                                                       |
| Ecs.0100          | Invalid ECS status.<br>Change the status and try again.                                                                                  | Change the ECS status to the desired one and try again.                                                                              |
| Ecs.0103          | The disk is unavailable.                                                                                                                 | Change the ECS status to the desired one and try again. If the EVS disk is faulty, contact the administrator for troubleshooting.    |
| Ecs.0104          | The number of disks to be attached to an ECS exceeds the number allowed.                                                                 | Detach EVS disks from the ECS before attaching new ones.                                                                             |
| Ecs.0105          | No system disk found.                                                                                                                    | Attach the system disk to the ECS and perform the desired operation again.                                                           |
| Ecs.0107          | The number of shared disks to be attached to an ECS exceeds the maximum limit.                                                           | Detach EVS disks from the ECS before attaching new ones.                                                                             |
| Other error codes | Other error messages                                                                                                                     | Initiate the request again. If the error persists, record the returned error code and contact the administrator for troubleshooting. |

## 14.5 Why Does the System Display a Question Mark When I Attempt to Obtain Console Logs?

### Symptom

The system displays a question mark (?) when I attempt to obtain the console logs of an ECS.

### Possible Causes

The image based on which the ECS was created supports viewing console logs. However, this function is not enabled on the ECS.

### Solution

Enable management console log obtaining on the ECS.

For details, see step 1 in section [Obtaining ECS Console Logs](#).

## 14.6 ECS Management

### 14.6.1 How Can a Changed Static Hostname Take Effect Permanently?

#### Symptom

The static hostname of a Linux ECS is user defined and injected using Cloud-Init during the ECS creation. Although the hostname can be changed by running the **hostname** command, the changed hostname is restored after the ECS is restarted.

#### Changing the Hostname on the ECS

To make the changed hostname still take effect even after the ECS is stopped or restarted, save the changed hostname into configuration files.

The changed hostname is assumed to be **new\_hostname**.

1. Modify the **/etc/hostname** configuration file.
  - a. Run the following command to edit the configuration file:  
**sudo vim /etc/hostname**
  - b. Change the hostname to the new one.
  - c. Run the following command to save and exit the configuration file:  
**:wq**
2. Modify the **/etc/sysconfig/network** configuration file.



- a. Run the following command to edit the configuration file:

```
sudo vim /etc/sysconfig/network
```

- b. Change the **HOSTNAME** value to the new hostname.

```
HOSTNAME=Changed hostname
```

 **NOTE**

If there is no **HOSTNAME** in the configuration file, manually add this parameter and set it to the changed hostname.

For example:

```
HOSTNAME=new_hostname
```

- c. Run the following command to save and exit the configuration file:

```
:wq
```

3. Modify the **/etc/cloud/cloud.cfg** configuration file.

- a. Run the following command to edit the configuration file:

```
sudo vim /etc/cloud/cloud.cfg
```

- b. Use either of the following methods to modify the configuration file:

- Method 1: Change the **preserve\_hostname** parameter value or add the **preserve\_hostname** parameter to the configuration file.

If **preserve\_hostname: false** is already available in the **/etc/cloud/cloud.cfg** configuration file, change it to **preserve\_hostname: true**. If **preserve\_hostname** is unavailable in the **/etc/cloud/cloud.cfg** configuration file, add **preserve\_hostname: true** before **cloud\_init\_modules**.

If you use method 1, the changed hostname still takes effect after the ECS is stopped or restarted. However, if the ECS is used to create a private image and the image is used to create a new ECS, the hostname of the new ECS is the hostname (**new\_hostname**) used by the private image, and user-defined hostnames cannot be injected using Cloud-Init.

- Method 2 (recommended): Delete or comment out - **update\_hostname**.

If you use method 2, the changed hostname still takes effect after the ECS is stopped or restarted. If the ECS is used to create a private image and the image is used to create a new ECS, the changed hostname permanently takes effect, and user-defined hostnames (such as **new\_new\_hostname**) can be injected using Cloud-Init.

4. Run the following command to restart the ECS:

```
sudo reboot
```

5. Run the following command to check whether the hostname has been changed:

```
sudo hostname
```

If the changed hostname is displayed in the command output, the hostname has been changed and the new name permanently takes effect.

## Modifying the Mapping Between the ECS Hostname and IP Address (Modifying the hosts File)

If you want to use the changed hostname as the preferred localhost and localhost.localdomain, update the mapping between the hostname and IP address after the hostname is changed and then save the configuration to the corresponding Cloud-Init configuration file so that the new hostname takes effect permanently.

The changed hostname is assumed to be **new\_hostname**.

1. Modify the **/etc/hostname** configuration file.
  - a. Run the following command to edit the configuration file:  
**sudo vim /etc/hostname**
  - b. Change the hostname to the new one.
  - c. Run the following command to save and exit the configuration file:  
**:wq**

2. Modify the **/etc/sysconfig/network** configuration file.
  - a. Run the following command to edit the configuration file:  
**sudo vim /etc/sysconfig/network**
  - b. Change the **HOSTNAME** value to the new hostname.  
**HOSTNAME=Changed hostname**

### NOTE

If there is no **HOSTNAME** in the configuration file, manually add this parameter and set it to the changed hostname.

For example:

```
HOSTNAME=new_hostname
```

3. Modify the **/etc/cloud/cloud.cfg** configuration file.
  - a. Run the following command to edit the configuration file:  
**sudo vim /etc/cloud/cloud.cfg**
  - b. Use either of the following methods to modify the configuration file:
    - Method 1: Change the **preserve\_hostname** parameter value or add the **preserve\_hostname** parameter to the configuration file.  
If **preserve\_hostname: false** is already available in the **/etc/cloud/cloud.cfg** configuration file, change it to **preserve\_hostname: true**.  
If **preserve\_hostname** is unavailable in the **/etc/cloud/cloud.cfg** configuration file, add **preserve\_hostname: true** before **cloud\_init\_modules**.

If you use method 1, the changed hostname still takes effect after the ECS is stopped or restarted. However, if the ECS is used to create a private image and the image is used to create a new ECS, the hostname of the new ECS is the hostname (**new\_hostname**) used by the private image, and user-defined hostnames cannot be injected using Cloud-Init.

- Method 2 (recommended): Delete or comment out - **update\_hostname**.  
If you use method 2, the changed hostname still takes effect after the ECS is stopped or restarted. If the ECS is used to create a private image and the image is used to create a new ECS, the changed hostname permanently takes effect, and user-defined hostnames (such as **new\_new\_hostname**) can be injected using Cloud-Init.
4. Update the mapping between the hostname and IP address in **/etc/hosts** to an entry starting with 127.0.0.1. Use **new\_hostname** as your preferred **localhost** and **localhost.localdomain**.
    - a. Run the following command to edit **/etc/hosts**:  
**sudo vim /etc/hosts**
    - b. Modify the entry starting with 127.0.0.1 and replace **localhost** and **localhost.localdomain** with **new\_hostname**.

```
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
127.0.0.1 new_hostname new_hostname
```
    - c. Run the following command to save and exit the configuration file:  
**:wq**
  5. Modify the **/etc/cloud/cloud.cfg** configuration file.
    - a. Run the following command to edit the configuration file:  
**sudo vim /etc/cloud/cloud.cfg**
    - b. Set **manage\_etc\_hosts** to **manage\_etc\_hosts: false**.

```
manage_etc_hosts: false
```
    - c. Run the following command to save and exit the configuration file:  
**:wq**
  6. Run the following command to restart the ECS:  
**sudo reboot**
  7. Run the following commands to check whether the changes to **hostname** and **hosts** take effect permanently:  
**sudo hostname**  
**sudo cat /etc/hosts**  
If the changed hostname (**new\_hostname**) and **hosts** are displayed in the command output, the changes take effect permanently.

## 14.6.2 Is an ECS Hostname with Suffix .novalocal Normal?

### Symptom

Hostnames of ECSs created based on some types of images have the suffix **.novalocal**, whereas others do not.

For example, the hostname is set to **abc** during ECS creation. [Table 14-3](#) lists the hostnames (obtained by running the **hostname** command) of ECSs created using different images and those displayed after the ECSs are restarted.

**Table 14-3** Hostnames of ECSs created from different images

| Image      | Hostname Before ECS Restart | Hostname After ECS Restart |
|------------|-----------------------------|----------------------------|
| CentOS 6.8 | abc                         | abc.novalocal              |
| CentOS 7.3 | abc.novalocal               | abc.novalocal              |
| Ubuntu 16  | abc                         | abc                        |

## Troubleshooting

This is a normal phenomenon.

The static hostname of a Linux ECS is user defined and injected using Cloud-Init during the ECS creation. According to the test results, Cloud-Init adapts to OSs differently. As a result, hostnames of some ECSs have suffix **.novalocal**, whereas others do not.

If you do not want to have the obtained hostnames contain suffix **.novalocal**, change the hostnames by referring to [How Can a Changed Static Hostname Take Effect Permanently?](#)

### 14.6.3 Why Do the Disks of a Windows ECS Go Offline After I Modify the ECS Specifications?

#### Scenarios

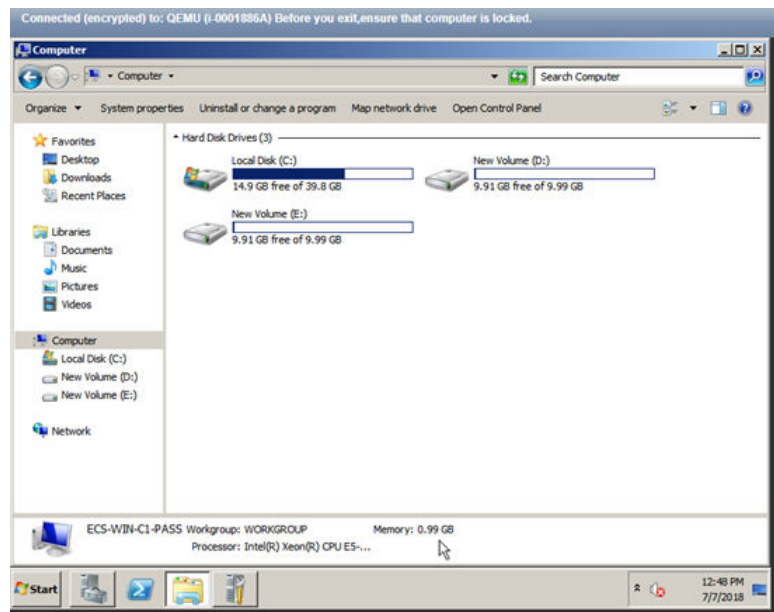
After you modify specifications of a Windows ECS, the disks may go offline. You need to check the number of disks after you modify the specifications.

#### Procedure


1. Check whether the number of disks displayed on the **Computer** page after you modified ECS specifications is the same as the number of disks before you modified ECS specifications.
  - If the numbers are the same, the status of the disks is properly. No further action is required.
  - If the numbers are different, the disks are offline. In this case, go to step [2](#).

For example:

An ECS running Windows Server 2008 has one system disk and two data disks attached before you modified the specifications.

**Figure 14-72** Disks before modifying ECS specifications

After the specifications are modified, check the number of disks.


**Figure 14-73** Disks after modifying ECS specifications

Only one system disk is displayed. The data disks are offline after you modify the specifications.

2. Bring the disks online.
  - a. Click **Start** in the task bar. In the displayed **Start** menu, right-click **Computer** and choose **Manage** from the shortcut menu.  
The **Server Manager** page is displayed.
  - b. In the left navigation pane, choose **Storage > Disk Management**.  
The **Disk Management** page is displayed.

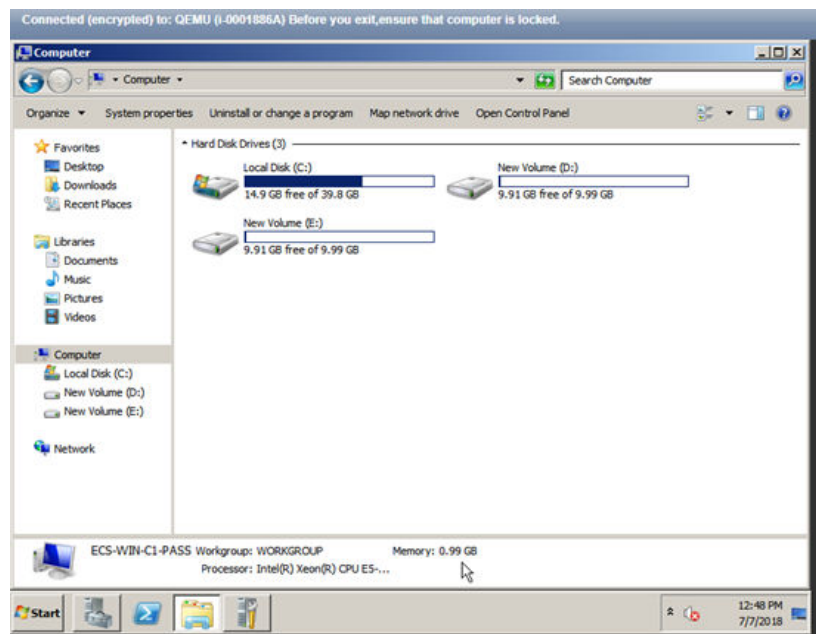
- c. In the left pane, the disk list is displayed. Right-click the offline disk and choose **Online** from the shortcut menu to bring it online.

**Figure 14-74** Bringing the disk online



3. On the **Computer** page, check whether the number of disks after you modified ECS specifications is the same as the number of disks before you modified the ECS specifications.
  - If the numbers are the same, no further action is required.
  - If the numbers are different, contact the administrator.

**Figure 14-75** Disks after you bring the disks online



## 14.6.4 Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?

### Scenarios

After you modify specifications of a Linux ECS, disk attachment may fail. You need to check the disk attachment after you modify the specifications.

### Procedure

1. Log in to the ECS as user **root**.
2. Run the following command to view the disks attached before specifications modification:

```
fdisk -l | grep 'Disk /dev/'
```

Figure 14-76 Viewing disks attached before specifications modification

```
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]# fdisk -l |grep 'Disk /dev/'
Disk /dev/vda: 42.9 GB, 42949672960 bytes, 83886080 sectors
Disk /dev/vdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Disk /dev/vdc: 10.7 GB, 10737418240 bytes, 20971520 sectors
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]#
```

As shown in [Figure 14-76](#), the ECS has three disks attached: **/dev/vda**, **/dev/vdb**, and **/dev/vdc**.

3. Run the following command to view disks attached after specifications modification:

```
df -h | grep '/dev/'
```

Figure 14-77 Viewing disks attached after specifications modification

```
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]# df -h | grep '/dev/'
/dev/vda2 39G 1.4G 35G 4% /
/dev/vda1 976M 146M 764M 16% /boot
```

As shown in [Figure 14-77](#), only one disk **/dev/vda** is attached to the ECS.

4. Check whether the number of disks obtained in step [3](#) is the same as that obtained in step [2](#).
  - If the numbers are the same, the disk attachment is successful. No further action is required.
  - If the numbers are different, the disk attachment failed. In this case, go to step [5](#).
5. Run the **mount** command to attach the affected disks.

For example, run the following command:

```
mount /dev/vdb1 /mnt/vdb1
```

In the preceding command, **/dev/vdb1** is the disk to be attached, and **/mnt/vdb1** is the path for disk attachment.

---

#### NOTICE

Ensure that **/mnt/vdb1** is empty. Otherwise, the attachment will fail.

---

- Run the following commands to check whether the numbers of disks before and after specifications modifications are the same:

```
fdisk -l | grep 'Disk /dev/'
```

```
df -h | grep '/dev/'
```

- If the numbers are the same, no further action is required.
- If the numbers are different, contact the administrator.

Figure 14-78 Checking the number of disks attached

```
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]# mount /dev/vdb1 /mnt/vdb1
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]# mount /dev/vdc1 /mnt/vdc1
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]#
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]#
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]#
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]# fdisk -l |grep 'Disk /dev/'
Disk /dev/vda: 42.9 GB, 42949672960 bytes, 83886080 sectors
Disk /dev/vdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Disk /dev/vdc: 10.7 GB, 10737418240 bytes, 20971520 sectors
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]# df -h | grep '/dev/'
/dev/vda2 39G 1.4G 35G 4% /
/dev/vda1 976M 146M 764M 16% /boot
/dev/vdb1 9.8G 23M 9.2G 1% /mnt/vdb1
/dev/vdc1 9.8G 23M 9.2G 1% /mnt/vdc1
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]#
```

As shown in [Figure 14-78](#), the numbers of disks before and after specifications modifications are the same. The disks are `/dev/vda`, `/dev/vdb`, and `/dev/vdc`.

## 14.7 OS Management

### 14.7.1 Can I Install or Upgrade the OS of an ECS?

You can install or upgrade ECS OSs provided on the cloud platform.

- When you create an ECS, you can select a public image or a private image created from a public image to install the ECS OS. Select an OS image based on the programming language in the actual application scenario.
- You can change your ECS OS through the management console, for example, you can upgrade CentOS 7.2 to CentOS 7.3.

### 14.7.2 Can I Change the OS of an ECS?

Yes, you can change the OS of an ECS.

If the OS running on an ECS cannot meet service requirements, for example, a higher OS version is required, you can change the ECS OS.

The cloud platform allows you to change the image type (public images, private images, and shared images) and OS. You can change the OS by changing the ECS image.

For instructions about how to change an ECS OS, see [Changing the OS](#).



### 14.7.3 How Long Does It Take to Change an ECS OS?

Generally, the process of changing the OS of an ECS takes about 1 to 2 minutes to complete. On the ECS console, stop the ECS and choose **More > Manage Image/Disk > Change OS** in the **Operation** column.

During this process, the ECS is in **Changing OS** state.

### 14.7.4 Can I Select Another OS During ECS OS Reinstallation?

No. You can use only the original image of the ECS to reinstall the OS. To use a new system image, see [Changing the OS](#).

### 14.7.5 How Long Does It Take to Reinstall an ECS OS?

Generally, the process of reinstalling the OS of an ECS takes about 1 to 2 minutes to complete. On the ECS console, stop the ECS and choose **More > Manage Image/Disk/Backup > Reinstall OS** in the **Operation** column.

During this process, the ECS is in **Reinstalling OS** state.

### 14.7.6 How Do I View the GPU Usage of a GPU-accelerated ECS?

#### Symptom

The GPU usage of GPU-accelerated ECSs running Windows Server 2012 and Windows Server 2016 cannot be viewed in Task Manager.

This section provides two methods for you to view the GPU usage. One is to run a command in the command-line interface, and the other is to install the GPU-Z tool.

#### Prerequisites

The NVIDIA driver has been installed on the GPU-accelerated ECS.

#### Method 1

1. Log in to the GPU-accelerated ECS.
2. Start the **Run** dialog box. Enter **cmd** and press **Enter**.
3. Run the following commands to check the GPU usage:

```
cd C:\Program Files\NVIDIA Corporation\NVSMI
nvidia-smi
```

To continuously observe the GPU usage, run the following command:

```
nvidia-smi -l 1
```

**Figure 14-79 GPU usage**

```
C:\Users\Administrator>
C:\Users\Administrator>cd C:\Program Files\NVIDIA Corporation\NvSMI
C:\Program Files\NVIDIA Corporation\NvSMI>nvidia-smi -l 1
Wed Mar 17 15:51:19 2021
```

| NVIDIA-SMI 452.39 |          | Driver Version: 452.39 |                  | CUDA Version: 11.0 |          |                   |
|-------------------|----------|------------------------|------------------|--------------------|----------|-------------------|
| GPU               | Name     | TCC/WDDM               | Bus-Id           | Disp.A             | Volatile | Uncorr. ECC       |
| Fan               | Temp     | Perf                   | Pwr:Usage/Cap    | Memory-Usage       | GPU-Util | Compute M. MIG M. |
| 0                 | Tesla T4 | WDDM                   | 00000000:21:01.0 | Off                | 0%       | Default           |
| N/A               | 33C      | P8                     | 14W / 70W        | 238MiB / 15360MiB  |          | N/A               |

```
Processes:
```

| GPU | GI  | CI  | PID  | Type | Process name                 | GPU Memory Usage |
|-----|-----|-----|------|------|------------------------------|------------------|
| ID  | ID  | ID  |      |      |                              |                  |
| 0   | N/A | N/A | 980  | C+G  | Insufficient Permissions     | N/A              |
| 0   | N/A | N/A | 3788 | C+G  | ...w5nlh2txyewy\SearchUI.exe | N/A              |
| 0   | N/A | N/A | 3896 | C+G  | ...y\ShellExperienceHost.exe | N/A              |

**NOTE**

NVIDIA GPUs can work in Tesla Compute Cluster (TCC) or Windows Display Driver Model (WDDM) mode.

- In TCC mode, the GPU is completely used for computing.
- In WDDM mode, the GPU supports both compute and graphics workloads.

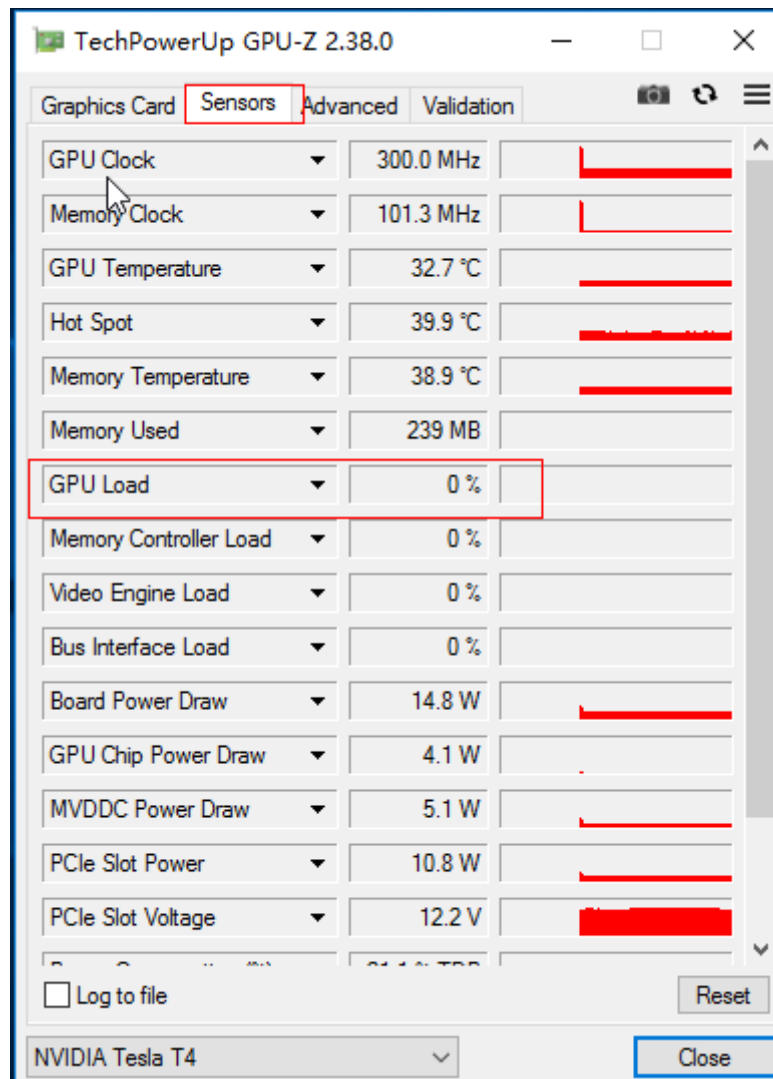
The WDDM mode can be used only when GRID drivers are installed on GPU-accelerated ECSs.

[Learn more](#) about TCC and WDDM.

**Method 2**

1. Log in to the GPU-accelerated ECS.
2. [Download GPU-Z](#) and install it.
3. Open GPU-Z and click **Sensors** to view the GPU usage.

Figure 14-80 GPU usage



## 14.8 File Upload/Data Transfer

### 14.8.1 How Do I Upload Files to My ECS?

#### Windows

- File transfer tool  
Install a file transfer tool, such as FileZilla on both the local computer and the Windows ECS and use it to transfer files. For details, see [How Can I Use FTP to Transfer Files from a Local Windows Computer to a Windows or Linux ECS?](#)
- (Recommended) Local disk mapping  
Use MSTSC to transfer files. This method does not support resumable transmission. Do not use this method to transfer large files.  
For details, see [How Can I Transfer Files from a Local Windows Computer to a Windows ECS?](#)

- FTP site  
Transfer files through an FTP site. Before transferring files from a local computer to a Windows ECS, set up an FTP site on the ECS and install FileZilla on the local computer.  
For details, see [How Can I Use FTP to Transfer Files from a Local Windows Computer to a Windows or Linux ECS?](#)
- From a local Mac  
If your local computer runs macOS, use Microsoft Remote Desktop for Mac to transfer files to the Windows ECS. For details, see [How Can I Transfer Files from a Local Mac to a Windows ECS?](#)

## Linux

- From a local Windows computer  
Use WinSCP to transfer the files to the Linux ECS. For details, see [How Can I Use WinSCP to Transfer Files from a Local Windows Computer to a Linux ECS?](#)  
Before transferring files from a local computer to a Linux ECS, set up an FTP site on the ECS and install FileZilla on the local computer. For details, see [How Can I Use FTP to Transfer Files from a Local Windows Computer to a Windows or Linux ECS?](#)
- From a local Linux computer  
Use SCP to transfer the files to the Linux ECS. For details, see [How Can I Use SCP to Transfer Files Between a Local Linux Computer and a Linux ECS?](#)  
Use SFTP to transfer the files to the Linux ECS. For details, see [How Can I Use SFTP to Transfer Files Between a Local Linux Computer and a Linux ECS?](#)  
Use FTP to transfer the files to the Linux ECS. For details, see [How Can I Use FTP to Transfer Files Between a Local Linux Computer and a Linux ECS?](#)

## Does an ECS Support FTP-based File Transferring by Default?

No. You need to install and configure FTP so that the ECS supports FTP-based file transfer.

## 14.8.2 How Can I Transfer Files from a Local Windows Computer to a Windows ECS?

### Scenarios

You want to transfer files from a local Windows computer to a Windows ECS through an MSTSC-based remote desktop connection.

### Prerequisites


- The target ECS is running.
- An EIP has been bound to the ECS. For details, see [Binding an EIP](#).
- Access to port 3389 is allowed in the inbound direction of the security group to which the ECS belongs. For details, see [Configuring Security Group Rules](#).

## Solution

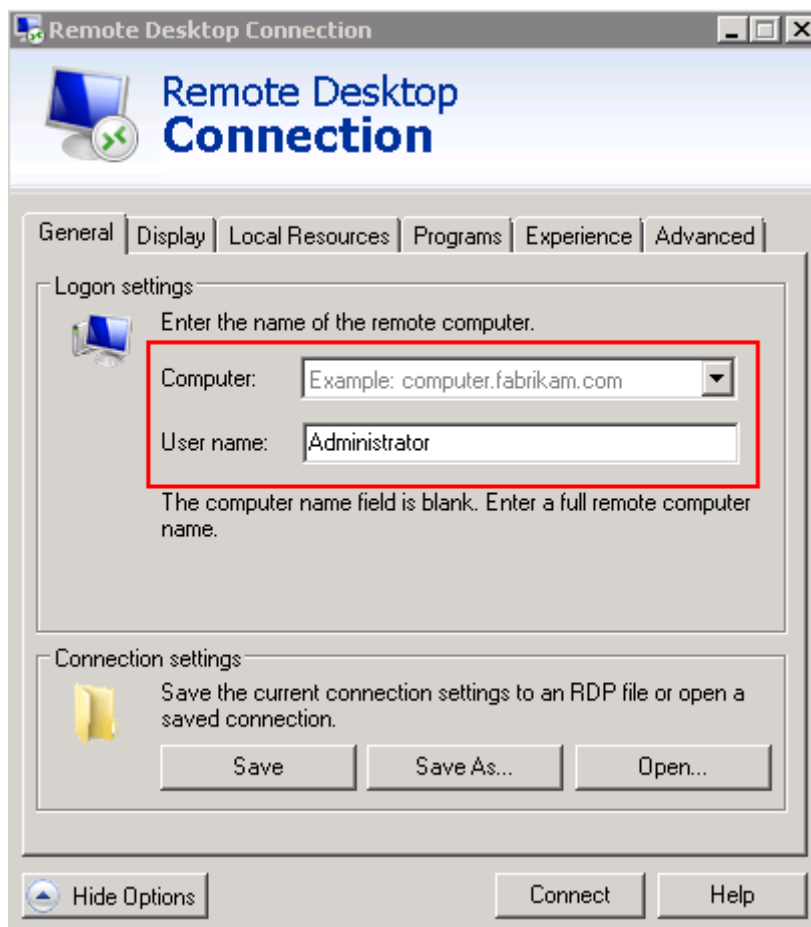
1. On the local Windows computer, click **Start**. In the **Search programs and files** text box, enter **mstsc**.

The **Remote Desktop Connection** window is displayed.

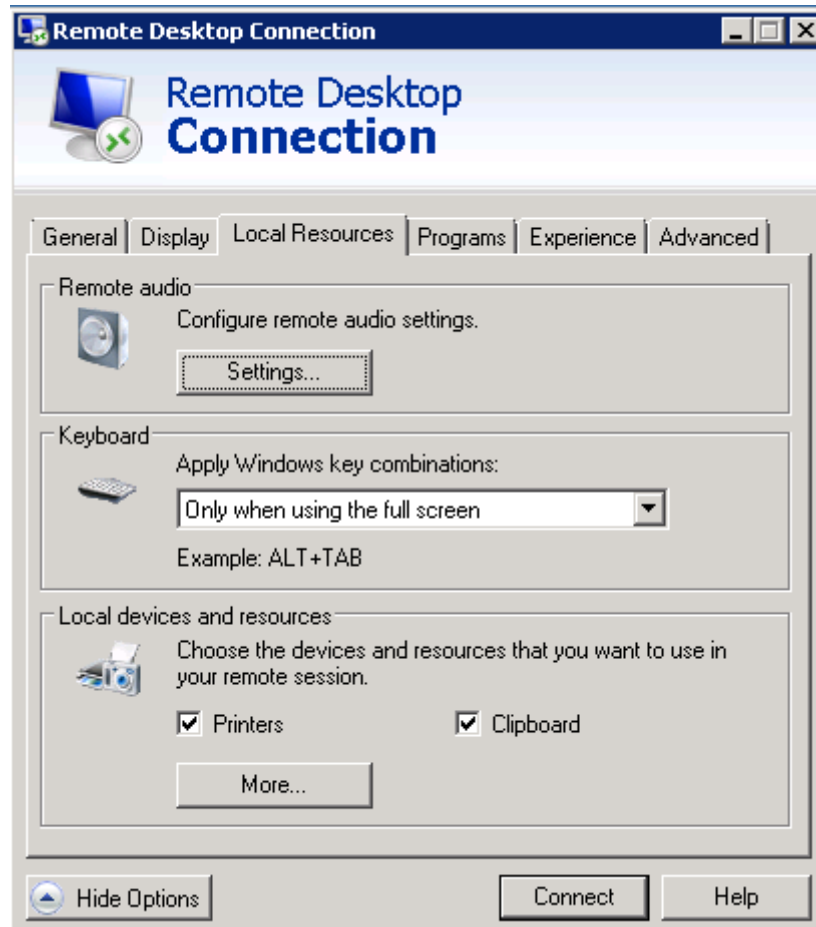
2. Click **Options**.



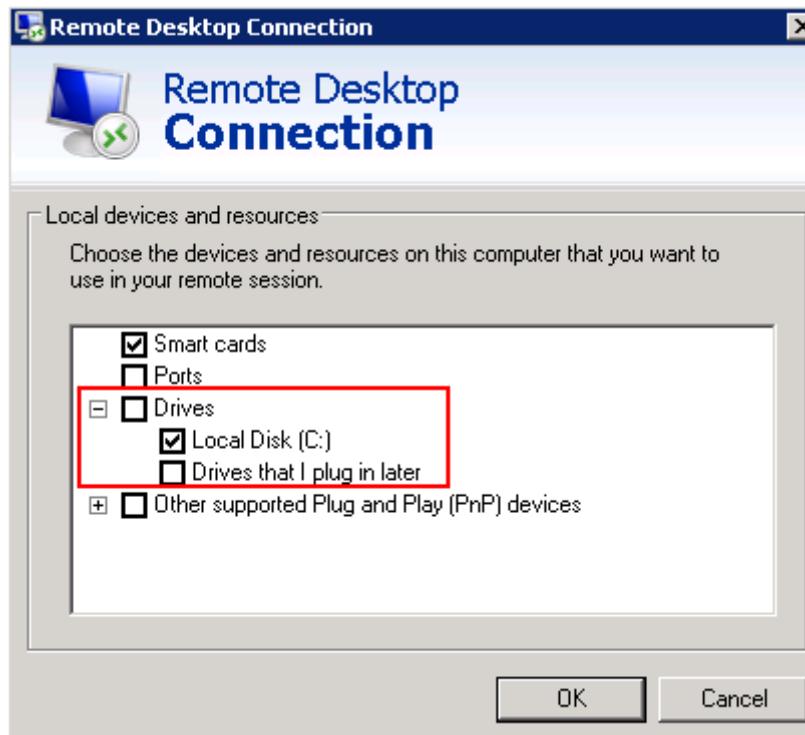
3. On the **General** tab, enter the EIP bound to the ECS and username **Administrator** for logging in to the ECS.



4. Click the **Local Resources** tab and verify that **Clipboard** is selected in the **Local devices and resources** pane.



5. Click **More**.
6. In the **Drives** pane, select the local disk where the file to be transferred to the Windows ECS is located.



7. Click **OK** and log in to the Windows ECS.
8. Choose **Start > Computer**.  
The local disk is displayed on the Windows ECS.
9. Double-click the local disk to access it and copy the file to be transferred to the Windows ECS.

### 14.8.3 How Can I Use WinSCP to Transfer Files from a Local Windows Computer to a Linux ECS?

#### Scenarios

WinSCP can be used to securely copy-paste files across local and remote computers. Compared with FTP, WinSCP allows you to use a username and password to access the destination server without any additional configuration on the server.


To transfer a file from a local Windows computer to a Linux ECS, WinSCP is commonly used. This section describes how to transfer files from a local Windows computer to a Linux ECS using WinSCP. In this example, the ECS running CentOS 7.2 is used as an example.

#### Prerequisites

- The target ECS is running.
- An EIP has been bound to the ECS. For details, see [Binding an EIP](#).
- Access to port 22 is allowed in the inbound direction of the security group to which the ECS belongs. For details, see [Configuring Security Group Rules](#).

## Solution

1. [Download WinSCP.](#)
2. Install WinSCP.
3. Start WinSCP.



Set parameters as follows:

- **File protocol:** Set this to **SFTP** or **SCP**.
  - **Host name:** Enter the EIP bound to the ECS. Log in to the management console to obtain the EIP.
  - **Port number:** **22** by default.
  - **User Name:** Enter the username for logging in to the ECS.
    - If the ECS is logged in using a password, the username is **root** for a public image.
  - **Password:** the password set when you created the ECS.
4. Click **Login**.
  5. Drag a file from the local computer on the left to the remotely logged in ECS on the right to transfer the file.

## 14.8.4 How Can I Transfer Files from a Local Mac to a Windows ECS?

### Scenarios

This section describes how to use Microsoft Remote Desktop for Mac to transfer files from a local Mac to a Windows ECS.



## Prerequisites

- The remote access tool supported by Mac has been installed on the local Mac. This section uses Microsoft Remote Desktop for Mac as an example. [Download Microsoft Remote Desktop for Mac.](#)
- The target Windows ECS has had an EIP bound.
- When you log in to the ECS for the first time, ensure that RDP has been enabled on it. To do so, use VNC to log in to the ECS, enable RDP, and access the ECS using MSTSC.

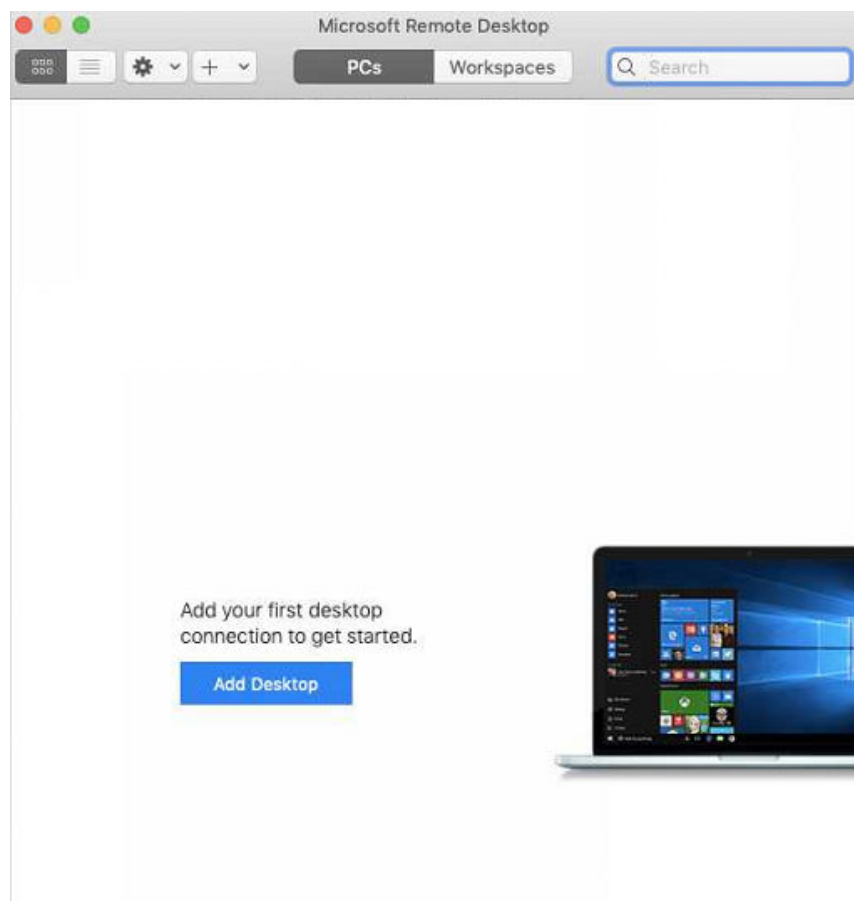
### NOTE

By default, RDP has been enabled on the ECSs created using a public image.

## Procedure

1. Start Microsoft Remote Desktop.
2. Click **Add Desktop**.

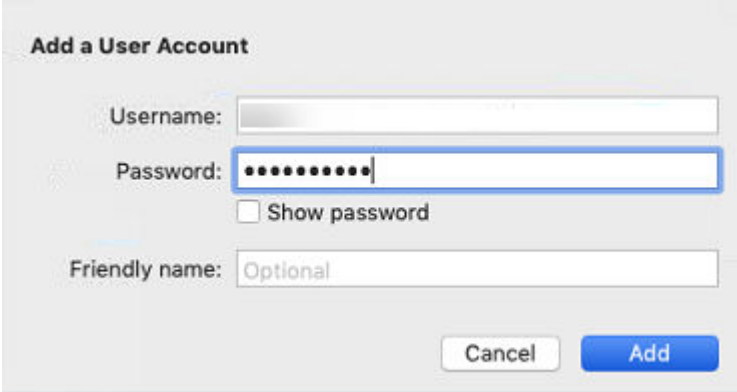
**Figure 14-81** Add Desktop



3. Set login parameters.
  - **PC name:** Enter the EIP bound to the target Windows ECS.
  - **User account:** Select **Add User Account** from the drop-down list. The **Add a User Account** dialog box is displayed.

- i. Enter the username **administrator** and password for logging in to the Windows ECS and click **Add**.

**Figure 14-82** Add user account



**Add a User Account**

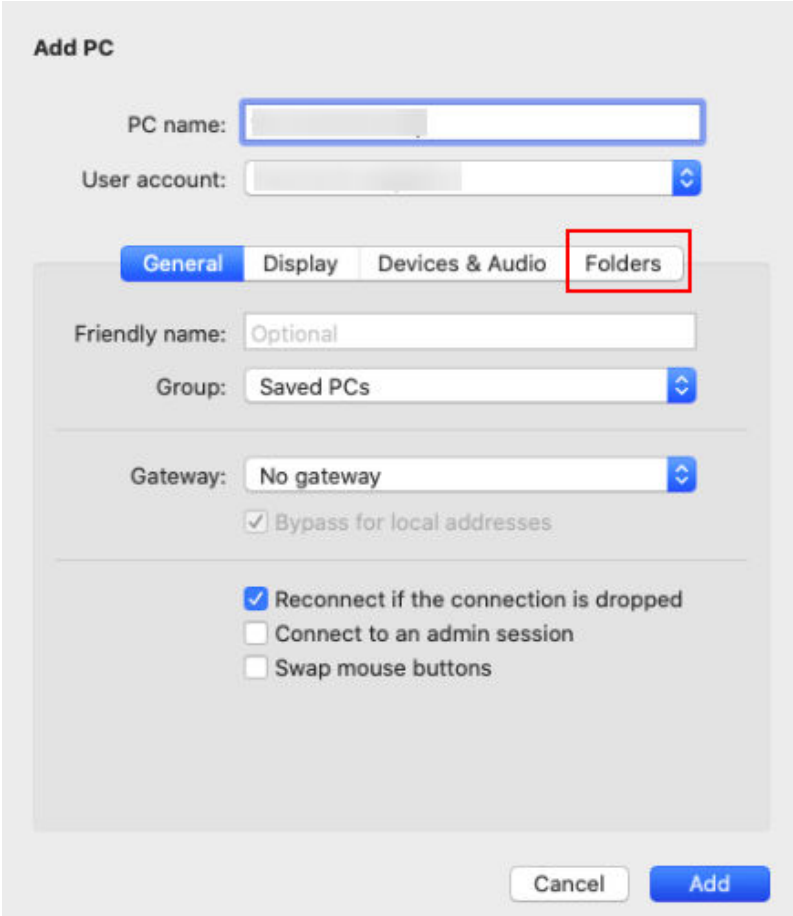
Username:

Password:

Show password

Friendly name:

**Figure 14-83** Add PC



**Add PC**

PC name:

User account:

**General** | Display | Devices & Audio | **Folders**

Friendly name:

Group:

Gateway:

Bypass for local addresses

Reconnect if the connection is dropped

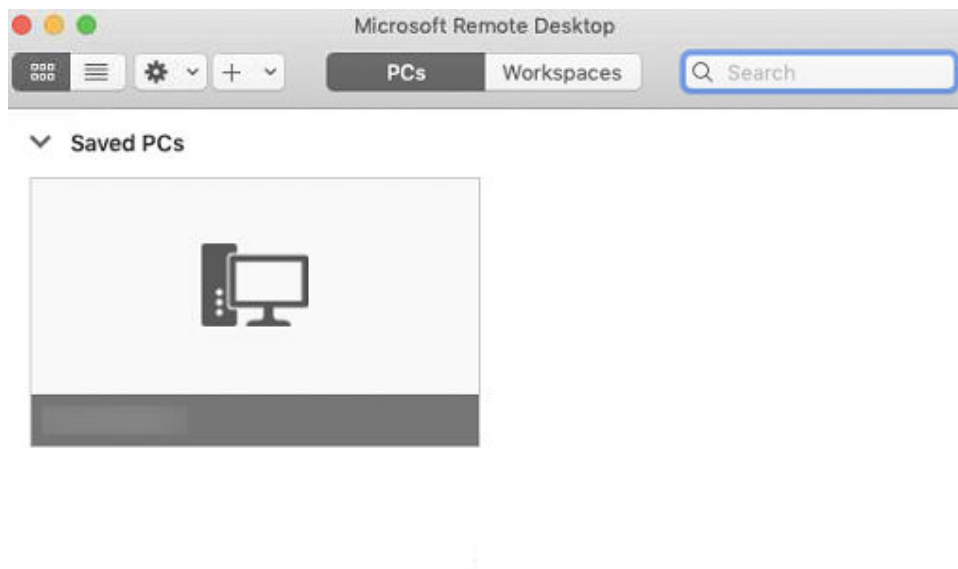
Connect to an admin session

Swap mouse buttons

- 4. Select the folder to be uploaded.
  - a. Click **Folders** and switch to the folder list.
  - b. Click **+** in the lower left corner, select the folder to be uploaded, and click **Add**.

5. On the **Remote Desktop** page, double-click the icon of the target Windows ECS.

**Figure 14-84** Double-click for login



6. Confirm the information and click **Continue**.  
You have connected to the Windows ECS.  
View the shared folder on the ECS.  
Copy the files to be uploaded to the ECS. Alternatively, download the files from the ECS to your local Mac.

## 14.8.5 How Can I Use SCP to Transfer Files Between a Local Linux Computer and a Linux ECS?

### Scenarios

You want to use SCP to transfer files between a local Linux computer and a Linux ECS.

### Procedure

Log in to the management console. On the **Elastic Cloud Server** page, obtain the EIP bound to the target ECS in the **IP Address** column.

- **Uploading files**

Run the following command on the local Linux computer to upload files to the Linux ECS:

```
scp Path in which the files are stored on the local computer
Username@EIP:Path in which the files are to be stored on the Linux ECS
```

For example, to transfer the **/home/test.txt** file on the local computer to the **/home** directory on the ECS whose EIP is 139.x.x.x, run the following command:

```
scp /home/test.txt root@139.x.x.x:/home
```

Enter the login password as prompted.

**Figure 14-85** Setting file uploading

```
[root@ecs-5c83 home]# scp /home/test.txt root@139. :/home
root@139. 's password:
test.txt
```

- **Downloading files**

Run the following command on the local Linux computer to download files from the Linux ECS:

**scp** *Username@EIP:Path in which the files are stored on the Linux ECS Path in which the files are to be stored on the local computer*

For example, to download the **/home/test.txt** file on the ECS whose EIP is 139.x.x.x to the **/home** directory on the local computer, run the following command:

**scp root@139.x.x.x:/home/test.txt /home/**

Enter the login password as prompted.

**Figure 14-86** Setting file downloading

```
[root@ecs-5c83 home]# scp root@139. :/home/test.txt /home
root@139. 's password:
test.txt
[root@ecs-5c83 home]# ls
test.txt
```

## 14.8.6 How Can I Use SFTP to Transfer Files Between a Local Linux Computer and a Linux ECS?

### Scenarios

You want to use SFTP to transfer files between a local Linux computer and a Linux ECS. The following uses CentOS as an example.

### Procedure

1. Log in to the ECS as user **root**.
2. Run the following command to check the OpenSSH version, which is expected to be 4.8p1 or later:

**ssh -V**

Information similar to the following is displayed:

```
OpenSSH_7.4p1, OpenSSL 1.0.2k-fips 26 Jan 2017
```

3. Create a user group and a user (for example, **user1**).

**groupadd sftp**

**useradd -g sftp -s /sbin/nologin user1**

4. Set a password for the user.

**passwd user1**

**Figure 14-87** Setting a password

```
[root@ecs-9a32-0001 ~]# passwd user1
Changing password for user user1.
New password:
BAD PASSWORD: The password is shorter than 8 characters
Retype new password:
passwd: all authentication tokens updated successfully.
[root@ecs-9a32-0001 ~]#
```

5. Assign permissions to directories.

```
chown root:sftp /home/user1
chmod 755 -R /home/user1
mkdir /home/user1/upload
chown -R user1:sftp /home/user1/upload
chmod -R 755 /home/user1/upload
```

6. Run the following command to edit the `sshd_config` configuration file:

```
vim /etc/ssh/sshd_config
```

Comment out the following information:

```
#Subsystem sftp /usr/libexec/openssh/sftp-server
```

Add the following information:

```
Subsystem sftp internal-sftp
Match Group sftp
ChrootDirectory /home/%u
ForceCommand internal-sftp
AllowTcpForwarding no
X11Forwarding no
```

**Figure 14-88** `sshd_config` file with the added information

```
override default of no subsystems
#Subsystem sftp /usr/libexec/openssh/sftp-server

Example of overriding settings on a per-user basis
#Match User anoncvs
X11Forwarding no
AllowTcpForwarding no
PermitTTY no
ForceCommand cvs server
PermitRootLogin yes
PasswordAuthentication yes
UseDNS no
Subsystem sftp internal-sftp
Match Group sftp
ChrootDirectory /home/%u
ForceCommand internal-sftp
AllowTcpForwarding no
X11Forwarding no
```

7. Run the following command to restart the ECS:

```
service sshd restart
```

Alternatively, run the following command to restart `sshd`:

```
systemctl restart sshd
```

8. Run the following command on the local computer to set up the connection:  
**sftp root@IP address**
9. Run the **sftp** command to check the connection.

```
root@ [redacted] 's password:
Connected to [redacted].
sftp> ls
ceshi print_all_tty.sh
s3fs_1.80_centos6.5_x86_64.rpm speedtest.py
uploads
sftp> pwd
Remote working directory: /root
sftp> lpwd
Local working directory: /root
sftp>
```

10. Transfer files or folders.  
To upload files or folders, run the **put -r** command.

```
sftp> put -r ceshi/
Uploading ceshi/ to /root/ceshi
Entering ceshi/
ceshi/mysql57-community-release-el 100% 9224 9.0KB/s 00:00
ceshi/haha 100% 28 0.0KB/s 00:00
sftp>
```

To download files or folders, run the **get -r** command.

```
sftp> get -r s3fs_1.80_centos6.5_x86_64.rpm
Fetching /root/s3fs_1.80_centos6.5_x86_64.rpm to s3fs_1.80_centos6.5_x86_64.rpm
/root/s3fs_1.80_centos6.5_x86_64.r 100% 3250KB 3.2MB/s 00:00
sftp>
```

## 14.8.7 How Can I Use FTP to Transfer Files from a Local Windows Computer to a Windows or Linux ECS?

### Scenarios

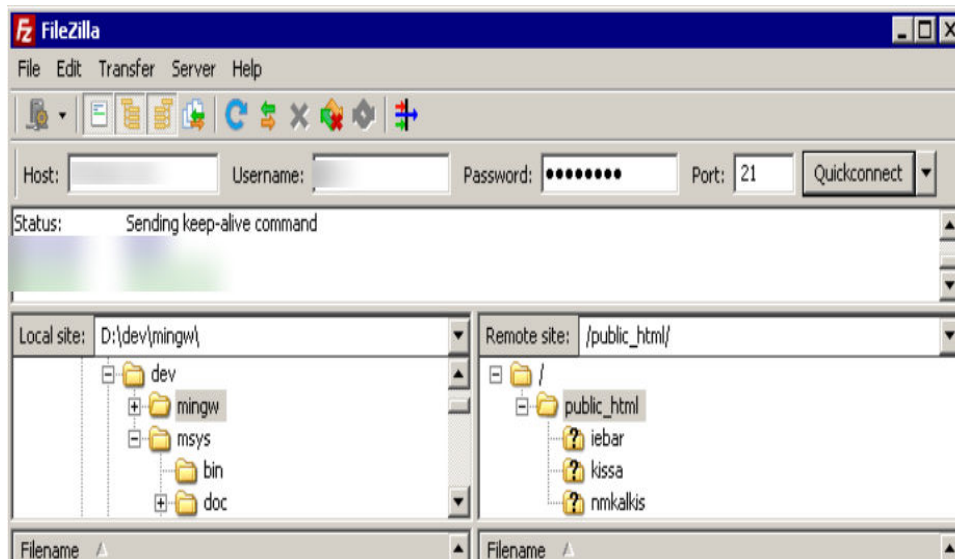
You want to use FTP to transfer files from a local Windows computer to an ECS.

### Prerequisites

- An EIP has been bound to the ECS and access to port 21 is allowed in the inbound direction of the security group to which the ECS belongs.
- You have enabled FTP on the target ECS. If you have not enabled FTP, check the following links to know how to set up an FTP site:

### Procedure

1. [Download FileZilla](#) and install it on the local Windows computer.
2. On the local Windows computer, open FileZilla, enter the information about the target ECS, and click **Quickconnect**.
  - **Host:** EIP bound to an ECS
  - **Username:** username set when the FTP site was set up
  - **Password:** password of the username
  - **Port:** FTP access port, which is port 21 by default

**Figure 14-89** Setting connection parameters

3. Drag files from the local computer on the left to the target ECS on the right to transfer them.

## 14.8.8 How Can I Use FTP to Transfer Files Between a Local Linux Computer and a Linux ECS?

### Scenarios

You want to use FTP on a local Linux computer to transfer files between the computer and a Linux ECS.

### Prerequisites

You have enabled FTP on the target ECS. If you have not enabled FTP, check the following links to know how to set up an FTP site:

- An EIP has been bound to the ECS and access to port 21 is allowed in the inbound direction of the security group to which the ECS belongs.
- You have enabled FTP on the target ECS. If you have not enabled FTP, check the following links to know how to set up an FTP site:

### Procedure

1. Install FTP on the local Linux computer.  
Take CentOS 7.6 as an example. Run the following command to install FTP:  
**yum -y install ftp**
2. Run the following command to access the ECS:  
**ftp** *EIP bound to the ECS*  
Enter the username and password as prompted for login.
  - **Uploading files**  
Run the following command to upload local files to the ECS:  
**put** *Path in which files are stored on the local computer*

For example, to upload the `/home/test.txt` file on the local Linux computer to the ECS, run the following command:

```
put /home/test.txt
```

– **Downloading files**

Run the following command to download files on the ECS to the local computer:

```
get Path in which the files are stored on the ECS Path in which the files are to be stored on the local computer
```

For example, to download the `test.txt` file on the ECS to the local Linux computer, run the following command:

```
get /home/test.txt
```

## 14.8.9 What Should I Do If the Connection Between the Client and the Server Times Out When I Upload a File Using FTP?

### Symptom

When I attempted to access the server from the client to upload a file using FTP, the connection timed out.

### Constraints

The operations described in this section apply to FTP on local Windows only.

### Possible Causes

Data is intercepted by the firewall or security group on the server.

### Solution

1. Check the firewall settings on the server.
2. Disable the firewall or add desired rules to the security group.

## 14.8.10 What Should I Do If Writing Data Failed When I Upload a File Using FTP?

### Symptom

When I attempted to upload a file using FTP, writing data failed. As a result, the file transfer failed.

### Constraints

The operations described in this section apply to FTP on Windows ECSs only.

### Possible Causes

When NAT is enabled on the FTP server, the FTP client must connect to the FTP server in passive mode. In such a case, the public IP address (EIP) of the server




cannot be accessed from the router. You need to add the EIP to the public IP address list on the server. Additionally, set the port range to limit the number of ports with data forwarded by the router.

## Solution

The EIP must be associated with the private IP address using NAT, so the server must be configured accordingly.

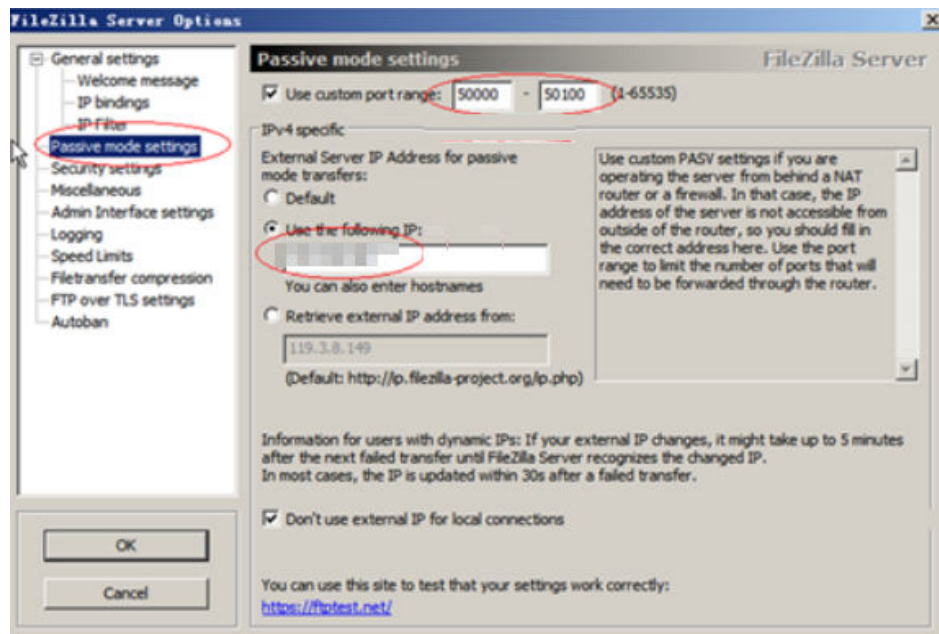
1. Configure the public IP address of the server.  
Choose **Edit > Settings**.

**Figure 14-90** Setting the public IP address



2. Choose **Passive mode settings**, set the port range (for example, 50000-50100) for transmitting data, and enter the target EIP.

**Figure 14-91** Setting the range of ports for data transmission




3. Click **OK**.
4. Allow traffic on TCP ports 50000-50100 and 21 in the security group in the inbound direction.
5. Test the connection on the client.

## 14.8.11 Why Am I Seeing an FTP Folder Error When I Open a Folder on an FTP Server?

### Symptom

An error occurs when you open a folder on an FTP server. The system displays a message asking you to check permissions.

**Figure 14-92** FTP Folder Error



### Possible Causes


The FTP firewall configured for the browser does not allow you to open the folder.

### Solution

The following uses Internet Explorer as an example.

1. Open the Internet Explorer and choose **Tools > Internet options**.
2. Click the **Advanced** tab.
3. Deselect **Use Passive FTP (for firewall and DSL modem compatibility)**.

Figure 14-93 Internet Options



4. Click **OK**, restart Internet Explorer, and open the folder on the FTP server again.

## 14.8.12 Why Do I Fail to Connect to a Linux ECS Using WinSCP?

### Symptom

Connecting to a Linux ECS using WinSCP fails, while using SSH tools like Xshell succeeds.

**Figure 14-94** Connection error using WinSCP

## Root Cause

If you can connect to a Linux ECS using SSH tools, the SSH tools run properly. Check the SFTP configuration file because WinSCP allows you to connect your Linux ECS via SFTP protocol.

Run the following command to view the `/etc/ssh/sshd_config` file:

```
vi /etc/ssh/sshd_config
```

Check the SFTP configuration and the configuration file is `/usr/libexec/openssh/sftp-server`.

**Figure 14-95** SFTP configuration file

```
override default of no subsystems
Subsystem sftp /usr/libexec/openssh/sftp-server
```

If the SFTP configuration file does not exist or the file permission is not 755, connecting to a Linux ECS using WinSCP will fail.

## Solution

- If the SFTP configuration file does not exist, you can transfer the file from an ECS that runs properly to your Linux ECS using SCP or other file transfer tools.
- If the file permission is not 755, you can run the following command to change the file permission to 755:

```
chmod 755 -R /usr/libexec/openssh/sftp-server
```

## 14.9 ECS Migration

## 14.9.1 Can I Migrate an ECS to Another Region, AZ, or Account?

After an ECS is created, it cannot be directly migrated to another region, AZ, or account.

## 14.10 Disk Management

### 14.10.1 Why Can't I Find My Newly Purchased Data Disk After I Log In to My Windows ECS?

#### Symptom

After logging in to my Windows ECS, I cannot find the attached data disk.

---

 **CAUTION**

Formatting a disk will cause data loss. Before formatting a disk, create a backup for it.

---

#### Possible Causes

- A newly added data disk has not been partitioned or initialized.
- The disk becomes offline after the ECS OS is changed or the ECS specifications are modified.

#### Newly Added Data Disk Has Not Been Partitioned or Initialized


A new data disk does not have partitions and file systems by default. That is why it is unavailable in **My Computer**. To resolve this issue, manually initialize the disk.

#### Disk Becomes Offline After the ECS OS Is Changed or the ECS Specifications Are Modified

After the ECS OS is changed, data disks may become unavailable due to file system inconsistency. After the specifications of a Windows ECS are modified, data disks may be offline.

1. Log in to the ECS, open the **cmd** window, and enter **diskmgmt.msc** to switch to the **Disk Management** page.  
Check whether the affected disk is offline.
2. Set the affected disk to be online.  
In the disk list, right-click the affected disk and choose **Online** from the shortcut menu to make it online.

Figure 14-96 Setting disk online



3. In **My Computer**, check whether the data disk is displayed properly. If the fault persists, initialize and partition the disk again. Before initializing the disk, create a backup for it.

## 14.10.2 How Can I Adjust System Disk Partitions?

### Scenarios

If the capacity of system disk partitions is inconsistent with the actual system disk capacity after an ECS is created, you can manually adjust the partitions to expand the system disk.

There are two ways to expand a system disk:

- Consider the empty partition as a new partition and attach this partition to a directory in the root partition after formatting it. For details, see this section.
- Add the empty partition to the root partition to be expanded. For detailed operations, see the following:
  - [How Can I Add the Empty Partition of an Expanded System Disk to the End Root Partition Online?](#)
  - [How Can I Add the Empty Partition of an Expanded System Disk to the Non-end Root Partition Online?](#)

### Procedure

This section uses an ECS running CentOS 7.3 64bit as an example. A 60 GB system disk was created with the ECS. However, the capacity of the system disk partition is displayed as only 40 GB.

To use the 20 GB capacity, performing the following operations:

#### Step 1 View disk partitions.

1. Log in to the Linux ECS.
2. Run the following command to switch to user **root**:  
**sudo su -**
3. Run the following command to view details about the ECS disk:  
**fdisk -l**

In the following command output, **/dev/xvda** or **/dev/vda** indicates the system disk.

**Figure 14-97** Viewing details about the disk

```
[root@ecs-8d6c ~]# df -h
Filesystem Size Used Avail Use% Mounted on
/dev/xvda1 38G 1.2G 35G 4% /
devtmpfs 899M 0 899M 0% /dev
tmpfs 908M 0 908M 0% /dev/shm
tmpfs 908M 8.4M 900M 1% /run
tmpfs 908M 0 908M 0% /sys/fs/cgroup
tmpfs 182M 0 182M 0% /run/user/0
[root@ecs-8d6c ~]# fdisk -l

Disk /dev/xvda: 64.4 GB, 64424509440 bytes, 125829120 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x0004d5e5

 Device Boot Start End Blocks Id System
/dev/xvda1 * 2048 79980543 39989248 83 Linux
/dev/xvda2 79980544 83886079 1952768 82 Linux swap / Solaris
[root@ecs-8d6c ~]# _
```

4. Run the following command to view disk partitions:  
**parted -l /dev/xvda**

**Figure 14-98** Viewing disk partitions

```
[root@ecs-8d6c ~]# parted -l /dev/xvda
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvda: 64.4GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number Start End Size Type File system Flags
 1 1049kB 41.0GB 40.9GB primary ext4 boot
 2 41.0GB 42.9GB 2000MB primary linux-swap(v1)
```

**Step 2** Create a partition for the expanded system disk capacity.

1. Run the following command to switch to the fdisk mode (taking **/dev/xvda** as an example):

**fdisk /dev/xvda**

Information similar to the following is displayed:

```
[root@ecs-8d6c ~]# fdisk /dev/xvda
Welcome to fdisk (util-linux 2.23.2).
```

```
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.
```

```
Command (m for help):
```

2. Enter **n** and press **Enter** to create a new partition.

Because the system disk has two existing partitions, the system automatically creates the third one.

Information similar to the following is displayed.

**Figure 14-99** Creating a new partition

```
[root@ecs-8d6c ~]# fdisk /dev/xvda
Welcome to fdisk (util-linux 2.23.2).

Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Command (m for help): n
Partition type:
 p primary (2 primary, 0 extended, 2 free)
 e extended
Select (default p):
Using default response p
Partition number (3,4, default 3):
First sector (83886080-125829119, default 83886080):
Using default value 83886080
Last sector, +sectors or +size{K,M,G} (83886080-125829119, default 125829119):
Using default value 125829119
Partition 3 of type Linux and of size 20 GiB is set

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 16: Device or resource busy.
The kernel still uses the old table. The new table will be used at
the next reboot or after you run partprobe(8) or kpartx(8)
Syncing disks.
[root@ecs-8d6c ~]#
```

3. Enter the new partition's start cylinder number and press **Enter**.  
The start cylinder number must be greater than the end cylinder numbers of existing partitions. In this example, use the default value for the new partition's start cylinder number and press **Enter**. Information similar to the following is displayed.

**Figure 14-100** Specifying the new partition's start cylinder number

```
First sector (83886080-125829119, default 83886080):
Using default value 83886080
Last sector, +sectors or +size{K,M,G} (83886080-125829119, default 125829119):
```

4. Enter the new partition's end cylinder number and press **Enter**.  
In this example, use the default value for the new partition's end cylinder number and press **Enter**. Information similar to the following is displayed.

**Figure 14-101** Specifying the new partition's end cylinder number

```
Last sector, +sectors or +size{K,M,G} (83886080-125829119, default 125829119):
Using default value 125829119
Partition 3 of type Linux and of size 20 GiB is set
```

5. Enter **p** and press **Enter** to view the created partition.  
Information similar to the following is displayed.



**Figure 14-102** Viewing the created partition

```
Command (m for help): p
Disk /dev/xvda: 64.4 GB, 64424509440 bytes, 125829120 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x0004d5e5

 Device Boot Start End Blocks Id System
/dev/xvda1 * 2048 79980543 39989248 83 Linux
/dev/xvda2 79980544 83886079 1952768 82 Linux swap / Solaris
/dev/xvda3 83886080 125829119 20971520 83 Linux
```

6. Enter **w** and press **Enter**. The system saves and exits the partition. The system automatically writes the partition result into the partition list. Then, the partition is created. Information similar to the following is displayed.

**Figure 14-103** Completing the partition creation

```
Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 16: Device or resource busy.
The kernel still uses the old table. The new table will be used at
the next reboot or after you run partprobe(8) or kpartx(8)
Syncing disks.
```

7. Run the following command to view disk partitions:  
**parted -l /dev/xvda**

**Figure 14-104** Viewing disk partitions

```
Disk Flags:
Number Start End Size Type File system Flags
 1 1049kB 41.0GB 40.9GB primary ext4 boot
 2 41.0GB 42.9GB 2000MB primary linux-swap(v1)
 3 42.9GB 64.4GB 21.5GB primary ext4
```

- Step 3** Run the following command to synchronize the modifications in the partition list with the OS:

**partprobe**

- Step 4** Configure the type of the new partition file system.

1. Run the following command to view the type of the file system:  
**df -TH**

**Figure 14-105** Viewing the file system type

```
[root@ecs-8d6c ~]# df -TH
Filesystem Type Size Used Avail Use% Mounted on
/dev/xvda1 ext4 41G 1.3G 37G 4% /
devtmpfs devtmpfs 943M 0 943M 0% /dev
tmpfs tmpfs 952M 0 952M 0% /dev/shm
tmpfs tmpfs 952M 8.8M 944M 1% /run
tmpfs tmpfs 952M 0 952M 0% /sys/fs/cgroup
tmpfs tmpfs 191M 0 191M 0% /run/user/0
[root@ecs-8d6c ~]#
```

2. Run the following command to format the partition (taking the **ext4** type as an example):

```
mkfs -t ext4 /dev/xvda3
```

#### NOTE

Formatting the partition requires a period of time. During this time, observe the system running status and do not exit the system.

Information similar to the following is displayed:

```
[root@ecs-86dc]# mkfs -t ext4 /dev/xvda3
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
1790544 inodes, 7156992 blocks
357849 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2155872256
219 block groups
32768 blocks per group, 32768 fragments per group
8176 inodes per group
Superblock backups stored on blocks:
 32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
 4096000

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

#### **Step 5** Mount the new partition to the target directory.

If you mount the new partition to a directory that is not empty, the subdirectories and files in the directory will be hidden. It is a good practice to mount the new partition to an empty directory or a newly created directory. If you want to mount the new partition to a directory that is not empty, temporarily move the subdirectories and files in the directory to another directory. After the partition is mounted, move the subdirectories and files back.

Take the newly created directory **/root/new** as an example.

1. Run the following command to create the **/root/new** directory:

```
mkdir /root/new
```

2. Run the following command to mount the new partition to the **/root/new** directory:

```
mount /dev/xvda3 /root/new
```

Information similar to the following is displayed:

```
[root@ecs-86dc]# mount /dev/xvda3 /root/new
[root@ecs-86dc]#
```

3. Run the following command to view the mounted file systems:

**df -TH**

Information similar to the following is displayed:

**Figure 14-106** Viewing the mounted file systems

```
[root@ecs-8d6c ~]# df -TH
Filesystem Type Size Used Avail Use% Mounted on
/dev/xvda1 ext4 41G 1.3G 37G 4% /
devtmpfs devtmpfs 943M 0 943M 0% /dev
tmpfs tmpfs 952M 0 952M 0% /dev/shm
tmpfs tmpfs 952M 8.8M 944M 1% /run
tmpfs tmpfs 952M 0 952M 0% /sys/fs/cgroup
/dev/xvda3 ext4 22G 47M 20G 1% /root/new
tmpfs tmpfs 191M 0 191M 0% /run/user/0
[root@ecs-8d6c ~]# b1
```

- Step 6** Determine whether to set automatic mounting upon system startup for the new disk.

If you do not set automatic mounting upon system startup, you must mount the new partition to the specified directory again after the ECS is restarted.

- If automatic mounting is required, go to [Step 7](#).
- If automatic mounting is not required, no further action is required.

- Step 7** Set automatic mounting upon system startup for the new disk.

**NOTE**

Do not set automatic mounting upon system startup for unformatted disks because this will cause ECS startup failures.

1. Run the following command to obtain the file system type and UUID:

**blkid**

**Figure 14-107** Viewing the file system type

```
[root@ecs-8d6c ~]# blkid
/dev/xvda1: UUID="7c4fce5d-f8f7-4ed6-8463-f2bd22d0ddea" TYPE="ext4"
/dev/xvda2: UUID="5de3cf2c-30c6-4fb2-9e63-830439d4e674" TYPE="swap"
/dev/xvda3: UUID="96e5e028-b0fb-4547-a82a-35ace1086c4f" TYPE="ext4"
[root@ecs-8d6c ~]#
```

According to the preceding figure, the UUID of the new partition is 96e5e028-b0fb-4547-a82a-35ace1086c4f.

2. Run the following command to open the **fstab** file using the vi editor:

**vi /etc/fstab**

3. Press **i** to enter editing mode.
4. Move the cursor to the end of the file and press **Enter**. Then, add the following information:

```
UUID=96e5e028-b0fb-4547-a82a-35ace1086c4f /root/new ext4 defaults 0
0
```

5. Press **Esc**, run the following command, and press **Enter**. The system saves the configurations and exits the vi editor.

```
:wq
```

#### NOTE

If you want to detach a new disk for which automatic mounting upon system startup has been set, you must delete the automatic mounting configuration before you detach the disk. Otherwise, the ECS cannot be started after you detach the disk. To delete the automatic mounting configuration, perform the following operations:

1. Run the following command to open the **fstab** file using the vi editor:

```
vi /etc/fstab
```

2. Press **i** to enter editing mode.
3. Delete the following statement:

```
UUID=96e5e028-b0fb-4547-a82a-35ace1086c4f /root/new ext4 defaults 0 0
```

4. Press **Esc**, run the following command, and press **Enter**. The system saves the configurations and exits the vi editor.

```
:wq
```


----End

### 14.10.3 How Can I Obtain the Mapping Between Disk Partitions and Disk Devices on a Windows ECS?

This section uses an ECS running Windows Server 2008 R2 64bit as an example to describe how to obtain the mapping between disk partitions and disk devices.


1. Log in to the Windows ECS.
2. Click **Start** in the lower left corner of the desktop.
3. Choose **Control Panel > Administrative Tools > Computer Management**.
4. In the navigation pane on the left, choose **Storage > Disk Management**.

Figure 14-108 Disk Management



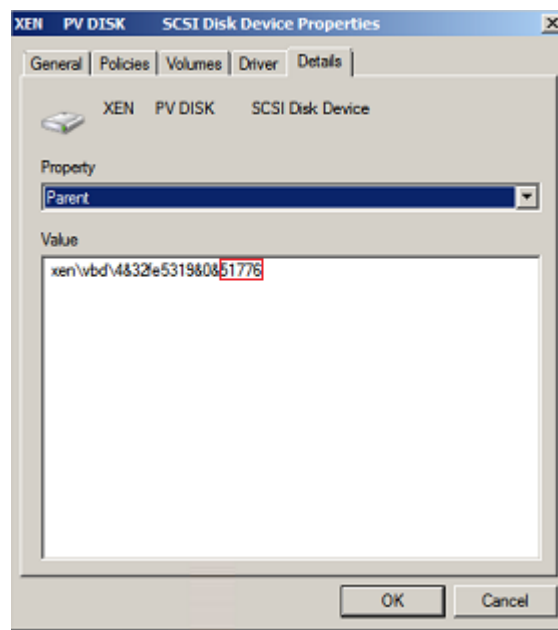
5. Taking disk 1 marked in [Figure 14-108](#) as an example, view the disk device for disk 1.
  - a. Right-click the gray area where disk 1 is located, as shown in the red box in [Figure 14-108](#).
  - b. Click **Properties**.  
The **SCSI Disk Device Properties** dialog box is displayed, as shown in [Figure 14-109](#).

**Figure 14-109** Disk properties



- c. Click the **Details** tab and set **Property** to **Parent**.

**Figure 14-110** Disk device details



- d. Record the digits following **&** in the parameter value, for example, **51776**, which is the master and slave device number corresponding to the disk partition.
- e. Obtain the disk device according to the information listed in [Table 14-4](#). The disk device corresponding to **51776** is **xvde**. The disk device used by disk 1 is xvde.

**Table 14-4** Mapping between disk partitions and disk devices

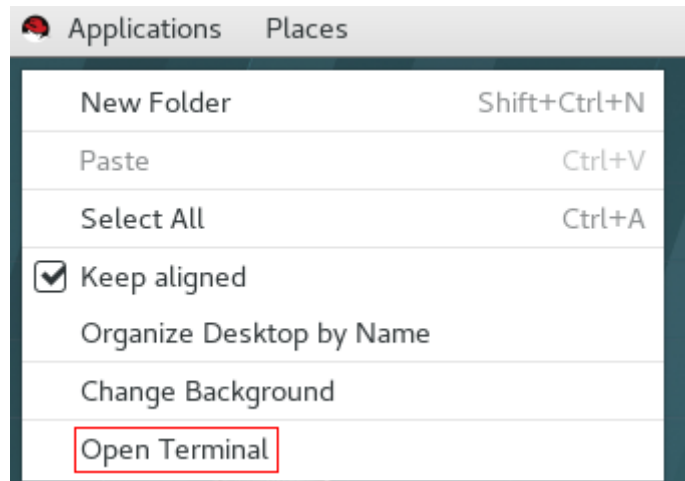
| Master and Slave Device Number for a Disk Partition | Disk Device |
|-----------------------------------------------------|-------------|
| 51712                                               | xvda        |
| 51728                                               | xvdb        |
| 51744                                               | xvdc        |
| 51760                                               | xvdd        |
| 51776                                               | xvde        |
| 51792                                               | xvdf        |
| 51808                                               | xvdg        |
| 51824                                               | xvdh        |
| 51840                                               | xvdi        |
| 51856                                               | xvdj        |
| 51872                                               | xvdk        |
| 51888                                               | xvdl        |
| 51904                                               | xvdm        |
| 51920                                               | xvdn        |
| 51936                                               | xvdo        |
| 51952                                               | xvdp        |
| 268439552                                           | xvdq        |
| 268439808                                           | xvdr        |
| 268440064                                           | xvds        |
| 268440320                                           | xvdt        |
| 268440576                                           | xvdu        |
| 268440832                                           | xvdv        |
| 268441088                                           | xvdw        |
| 268441344                                           | xvdx        |

## 14.10.4 How Can I Obtain the Mapping Between Disk Partitions and Disk Devices on a Linux ECS?

For a Linux ECS, its disk partitions correspond to disk devices. This section uses a Linux ECS running Red Hat Enterprise Linux 7 as an example to describe how to obtain the mapping between disk partitions and disk devices.


1. Log in to the Linux ECS as user **root**.
2. Right-click in the blank area of the desktop and choose **Open Terminal** from the shortcut menu.

**Figure 14-111** open terminal



3. Run the following command to view disk partitions and disk devices:  
**fdisk -l**

**Figure 14-112** Viewing disk partitions and disk devices



**Table 14-5** lists the mapping between disk partitions and disk devices.

**Table 14-5** Mapping between disk partitions and disk devices

| Disk Partition | Disk Device |
|----------------|-------------|
| xvda           | xvda        |
| xvdb           | xvdb        |



| Disk Partition | Disk Device |
|----------------|-------------|
| xvdc           | xvdc        |
| xvdd           | xvdd        |
| xvde           | xvde        |
| xvdf           | xvdf        |
| xvdg           | xvdg        |
| xvdh           | xvdh        |
| xvdi           | xvdi        |
| xvdj           | xvdj        |
| xvdk           | xvdk        |
| xvdl           | xvdl        |
| xvdm           | xvdm        |
| xvdn           | xvdn        |
| xvdo           | xvdo        |
| xvdp           | xvdp        |
| xvdq           | xvdq        |
| xvdr           | xvdr        |
| xvds           | xvds        |
| xvdt           | xvdt        |
| xvdu           | xvdu        |
| xvdv           | xvdv        |
| xvdw           | xvdw        |
| xvdx           | xvdx        |

### 14.10.5 How Can I Enable Virtual Memory on a Windows ECS?

Enabling ECS virtual memory will deteriorate I/O performance. If the memory is insufficient, you are advised to expand the memory by referring to [Modifying ECS Specifications](#). If you really need to enable virtual memory, see the operations described below.


**NOTE**

If the memory usage is excessively high and the I/O performance is not as good as expected, you are not advised to enable virtual memory. The reason is as follows: The excessively high memory usage limits the system performance improvement. Furthermore, frequent memory switching requires massive additional I/O operations, which will further deteriorate the I/O performance and the overall system performance.

The operations described in this section are provided for the ECSs running Windows Server 2008 or later.


1. Right-click **Computer** and choose **Properties** from the shortcut menu.
2. In the left navigation pane, choose **Advanced system settings**.  
The **System Properties** dialog box is displayed.
3. Click the **Advanced** tab and then **Settings** in the **Performance** pane.  
The **Performance Options** dialog box is displayed.

**Figure 14-113** Performance Options



4. Click the **Advanced** tab and then **Background Services** in the **Processor scheduling** pane.
5. Click **Change** in the **Virtual memory** pane.  
The **Virtual Memory** dialog box is displayed.
6. Configure virtual memory based on service requirements.
  - **Automatically manage paging file size for all drives:** Deselect the check box.
  - **Drive:** Select the drive where the virtual memory file is stored.  
You are advised not to select the system disk to store the virtual memory.
  - **Custom size:** Select **Custom size** and set **Initial size** and **Maximum size**.  
Considering **Memory.dmp** caused by blue screen of death (BSOD), you are advised to set **Initial size** to **16** and **Maximum size** to **4,096**.

Figure 14-114 Virtual Memory



7. Click **Set** and then **OK** to complete the configuration.
8. Restart the ECS for the configuration to take effect.

## 14.10.6 How Can I Add the Empty Partition of an Expanded System Disk to the End Root Partition Online?

### Scenarios

If the capacity of system disk partitions is inconsistent with the actual system disk capacity after an ECS is created, you can add the empty partition to the root partition of the system disk.

This section describes how to add the empty partition to the end root partition online.

### Procedure

In the following operations, the ECS that runs CentOS 6.5 64bit and has a 50 GB system disk is used as an example. The system disk has two partitions, **/dev/xvda1: swap** and **/dev/xvda2: root**, and the root partition is the end partition.

1. Run the following command to view disk partitions:

```
parted -l /dev/xvda
```

```
[root@sluo-ecs-5e7d ~]# parted -l /dev/xvda
Disk /dev/xvda: 53.7GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
```

```
Number Start End Size Type File system Flags
1 1049kB 4296MB 4295MB primary linux-swap(v1)
2 4296MB 42.9GB 38.7GB primary ext4 boot
```

2. Run the following command to obtain the file system type and UUID:

**blkid**

```
/dev/xvda1: UUID="25ec3bdb-ba24-4561-bcdc-802edf42b85f" TYPE="swap"
/dev/xvda2: UUID="1a1ce4de-e56a-4e1f-864d-31b7d9dfb547" TYPE="ext4"
```

3. Run the following command to install the growpart tool:

This tool may be integrated in the **cloud-utils-growpart/cloud-utils/cloud-initramfs-tools/cloud-init** package. Run the **yum install cloud-\*** command to ensure it is available.

**yum install cloud-utils-growpart**

4. Run the following command to expand the root partition (the second partition) using growpart:

**growpart /dev/xvda 2**

```
[root@sluo-ecs-5e7d ~]# growpart /dev/xvda 2
CHANGED: partition=2 start=8390656 old: size=75495424 end=83886080 new:
size=96465599,end=104856255
```

5. Run the following command to verify that online capacity expansion is successful:

**parted -l /dev/xvda**

```
[root@sluo-ecs-5e7d ~]# parted -l /dev/xvda
Disk /dev/xvda: 53.7GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
```

```
Number Start End Size Type File system Flags
1 1049kB 4296MB 4295MB primary linux-swap(v1)
2 4296MB 53.7GB 49.4GB primary ext4 boot
```

6. Run the following command to expand the capacity of the file system:

**resize2fs -f \$Partition name**

Suppose the partition name is **/dev/xvda2**, run the following command:

```
[root@sluo-ecs-a611 ~]# resize2fs -f /dev/xvda2
resize2fs 1.42.9 (28-Dec-2013)
Filesystem at /dev/xvda2 is mounted on /; on-line resizing required
old_desc_blocks = 3, new_desc_blocks = 3
....
[root@sluo-ecs-a611 ~] # df -hT //Check file system capacity expansion
```

## 14.10.7 How Can I Add the Empty Partition of an Expanded System Disk to the Non-end Root Partition Online?

### Scenarios

If the capacity of system disk partitions is inconsistent with the actual system disk capacity after an ECS is created, you can add the empty partition to the root partition of the system disk.

This section describes how to add the empty partition to the non-end root partition online.

## Procedure

In the following operations, the ECS that runs CentOS 6.5 64bit and has a 100 GB system disk is used as an example. The system disk has two partitions, **/dev/xvda1: root** and **/dev/xvda2: swap**, and the root partition is not the end partition.

1. Run the following command to view disk partitions:

### **parted -l /dev/xvda**

```
[root@sluo-ecs-a611 ~]# parted -l /dev/xvda
Disk /dev/xvda: 107GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:
```

| Number | Start  | End    | Size   | Type    | File system    | Flags |
|--------|--------|--------|--------|---------|----------------|-------|
| 1      | 1049kB | 41.0GB | 40.9GB | primary | ext4           | boot  |
| 2      | 41.0GB | 42.9GB | 2000MB | primary | linux-swap(v1) |       |

The first is the root partition, and the second is the swap partition.

2. View and edit the fstab partition table to delete the swap partition attachment information.
  - a. Run the following command to view the fstab partition table:

### **tail -n 3 /etc/fstab**

```
[root@sluo-ecs-a611 ~]# tail -n 3 /etc/fstab
#
UUID=7c4fce5d-f8f7-4ed6-8463-f2bd22d0ddea / ext4 defaults 1 1
UUID=5de3cf2c-30c6-4fb2-9e63-830439d4e674 swap swap defaults 0 0
```

- b. Run the following command to edit the fstab partition table and delete the swap partition attachment information.

### **vi /etc/fstab**

#### **tail -n 3 /etc/fstab**

```
[root@sluo-ecs-a611 ~]# vi /etc/fstab
[root@sluo-ecs-a611 ~]# tail -n 3 /etc/fstab
#
UUID=7c4fce5d-f8f7-4ed6-8463-f2bd22d0ddea / ext4 defaults 1 1
```

3. Run the following command to disable the swap partition:

### **swapoff -a**

4. Delete the swap partition.

- a. Run the following command to view the partition:

### **parted /dev/xvda**

```
[root@sluo-ecs-a611 ~]# parted /dev/xvda
GNU Parted 3.1
Using /dev/xvda
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) help
align-check TYPE N check partition N for TYPE(min|opt) alignment
help [COMMAND] print general help, or help on COMMAND
mklabel,mktable LABEL-TYPE create a new disklabel (partition table)
mkpart PART-TYPE [FS-TYPE] START END make a partition
name NUMBER NAME name partition NUMBER as NAME
print [devices]free[list,all][NUMBER] display the partition table, available devices, free space,
all found partitions, or a
particular partition
quit exit program
rescue START END rescue a lost partition near START and END
rm NUMBER delete partition NUMBER
select DEVICE choose the device to edit
disk_set FLAG STATE change the FLAG on selected device
disk_toggle [FLAG] toggle the state of FLAG on selected device
```

```
set NUMBER FLAG STATE change the FLAG on partition NUMBER
toggle [NUMBER [FLAG]] toggle the state of FLAG on partition NUMBER
unit UNIT set the default unit to UNIT
version display the version number and copyright information of GNU
Parted
(parted)
```

b. **Press p.**

```
Disk /dev/xvda: 107GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:
```

| Number | Start  | End    | Size   | Type    | File system    | Flags |
|--------|--------|--------|--------|---------|----------------|-------|
| 1      | 1049kB | 41.0GB | 40.9GB | primary | ext4           | boot  |
| 2      | 41.0GB | 42.9GB | 2000MB | primary | linux-swap(v1) |       |

c. Run the following command to delete the partition:

```
rm 2
```

```
(parted) rm2
```

d. **Press p.**

```
(parted) p
Disk /dev/xvda: 107GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:
```

| Number | Start  | End    | Size   | Type    | File system | Flags |
|--------|--------|--------|--------|---------|-------------|-------|
| 1      | 1049kB | 41.0GB | 40.9GB | primary | ext4        | boot  |

e. Run the following command to edit the fstab partition table:

```
quit
```

```
(parted) quit
Information: You may need to update /etc/fstab.
```

5. Run the following command to view partition after the swap partition is deleted:

```
parted -l /dev/xvda
```

```
[root@sluo-ecs-a611 ~]# parted -l /dev/xvda
Disk /dev/xvda: 107GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:
```

| Number | Start  | End    | Size   | Type    | File system | Flags |
|--------|--------|--------|--------|---------|-------------|-------|
| 1      | 1049kB | 41.0GB | 40.9GB | primary | ext4        | boot  |

6. Run the following command to install the growpart tool:

This tool may be integrated in the **cloud-utils-growpart/cloud-utils/cloud-initramfs-tools/cloud-init** package. Run the **yum install cloud-\*** command to ensure it is available.

```
yum install cloud-utils-growpart
```

7. Run the following command to expand the root partition (the first partition) using growpart:

```
growpart /dev/xvda 1
```

```
[root@sluo-ecs-a611 ~]# growpart /dev/xvda 1
CHANGED: partition=1 start=2048 old: size=79978496 end=79980544 new:
size=209710462,end=209712510
```

8. Run the following command to verify that online capacity expansion is successful:

```
[root@sluo-ecs-a611 ~]# parted -l /dev/xvda
Disk /dev/xvda: 107GB
```

```
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:
```

```
Number Start End Size Type File system Flags
1 1049kB 107GB 107GB primary ext4 boot
```

9. Run the following command to expand the capacity of the file system:

**resize2fs -f \$Partition name**

Suppose the partition name is **/dev/xvda1**, run the following command:

```
[root@sluo-ecs-a611 ~]# resize2fs -f /dev/xvda1
resize2fs 1.42.9 (28-Dec-2013)
Filesystem at /dev/xvda1 is mounted on /; on-line resizing required
old_desc_blocks = 3, new_desc_blocks = 3
....
[root@sluo-ecs-a611 ~] # df -hT //Check file system capacity expansion
```

## 14.10.8 Can I Attach Multiple Disks to an ECS?

Yes. The ECSs created after the disk function upgrade can have up to 60 attached disks.

- When you create an ECS, you can attach 24 disks to it.
- After you create an ECS, you can attach up to 60 disks to it.

**Table 14-6** Numbers of disks that can be attached to a newly created ECS

| ECS Type                | Maximum VBD Disks | Maximum SCSI Disks | Constraint                                                                                                                                       |
|-------------------------|-------------------|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| Xen                     | 60                | 59                 | VBD disks + SCSI disks ≤ 60 (This constraint does not apply to local disks.)<br>The number of local disks is determined based on the ECS flavor. |
| KVM (excluding D3 ECSs) | 24                | 59                 | VBD disks + SCSI disks ≤ 60 (This constraint does not apply to local disks.)<br>The number of local disks is determined based on the ECS flavor. |
| D3                      | 24                | 30                 | VBD disks + SCSI disks ≤ 54 (This constraint does not apply to local disks.)<br>The number of local disks is determined based on the ECS flavor. |

 **NOTE**

- The system disk of an ECS is of VBD type. The maximum number of SCSI disks is 59.
- For a D-series KVM ECS, its local disks use two SCSI controllers, indicating that 30 SCSI drive letters are used. A maximum of 30 SCSI disks can be attached to such an ECS.

The maximum number of disks that you can attach to an ECS that was created before the disk function upgrade remains unchanged, as shown in [Table 14-7](#).

**Table 14-7** Numbers of disks that can be attached to an existing ECS

| ECS Type | Maximum VBD Disks | Maximum SCSI Disks | Maximum Local Disks | Constraint                                |
|----------|-------------------|--------------------|---------------------|-------------------------------------------|
| Xen      | 60                | 59                 | 59                  | VBD disks + SCSI disks + Local disks ≤ 60 |
| KVM      | 24                | 23                 | 59                  | VBD disks + SCSI disks ≤ 24               |

To attach 60 disks, enable advanced disk. For details, see [Enabling Advanced Disk](#).

## How Can I Check Whether an ECS Is Created Before or After the Disk Function Upgrade?

1. Log in to management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Click the name of the target ECS. The page providing details about the ECS is displayed.
4. Click the **Disks** tab.
5. Check the number of disks that can be attached to the ECS to determine the total number of disks.
  - If the total number of disks that can be attached is 24 (including the system disk), the ECS is created before the disk function upgrade. In such a case, you can enable advanced disk as prompted so that up to 60 disks can be attached to the ECS. For details, see [Enabling Advanced Disk](#).
  - If the total number of disks that can be attached is 60 (including the system disk), the ECS is created after the disk function upgrade.

## 14.10.9 What Are the Requirements for Attaching an EVS Disk to an ECS?

- The EVS disk and the target ECS must be located in the same AZ.
- For a non-shared disk, the EVS disk must be in **Available** state.  
For a shared disk, the target EVS disk must be in **In-use** or **Available** state.



- The target ECS must be in **Running** or **Stopped** state.
- The EVS disk must not be frozen.
- Certain ECSs support SCSI EVS disk attachment. For details, see [Which ECSs Can Be Attached with SCSI EVS Disks?](#)

## 14.10.10 Which ECSs Can Be Attached with SCSI EVS Disks?


All types of ECSs can be attached with SCSI EVS disks.

## 14.10.11 How Do I Obtain My Disk Device Name in the ECS OS Using the Device Identifier Provided on the Console?

### Scenarios

You find that the device name displayed in the ECS OS is different from that displayed on the management console and you cannot determine which disk name is correct. This section describes how to obtain the disk name used in an ECS OS according to the device identifier on the console.

### Obtaining the Disk ID of an ECS on the Console

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. Click the target ECS name in the ECS list.  
The ECS details page is displayed.
4. Click the **Disks** tab and then click  to expand the disk information.
5. Check the device type and ID of the disk.

#### NOTE

If **Device Identifier** is not displayed on the page, stop the ECS and restart it.

#### – KVM ECS

- If **Device Type** is **VBD**, use a serial number or BDF to obtain the disk device name.

If you use a serial number (recommended) to obtain the disk name, see [Using a Serial Number to Obtain the Disk Name \(Windows\)](#) and [Using a Serial Number to Obtain a Disk Device Name \(Linux\)](#).

If you use a BDF to obtain the disk device name, see [Using a BDF to Obtain a Disk Device Name \(Linux\)](#). (BDF cannot be used to obtain the disk name of Windows ECSs.)

- If **Device Type** is **SCSI**, use a WWN to obtain the disk name. For details, see [Using a WWN to Obtain the Disk Name \(Windows\)](#) and [Using a WWN to Obtain a Disk Device Name \(Linux\)](#).

## Using a Serial Number to Obtain the Disk Name (Windows)

If a serial number is displayed on the console, use either of the following methods to obtain the disk name.

### cmd

1. Start **cmd** in a Windows OS as an administrator and run either of the following commands:

```
wmic diskdrive get serialnumber
```

```
wmic path win32_physicalmedia get SerialNumber
```

```
wmic path Win32_DiskDrive get SerialNumber
```

#### NOTE

A serial number is the first 20 digits of a disk UUID.

For example, if the serial number of a VBD disk on the console is 97c876c0-54b3-460a-b, run either of the following commands to obtain the serial number of the disk on the ECS OS:

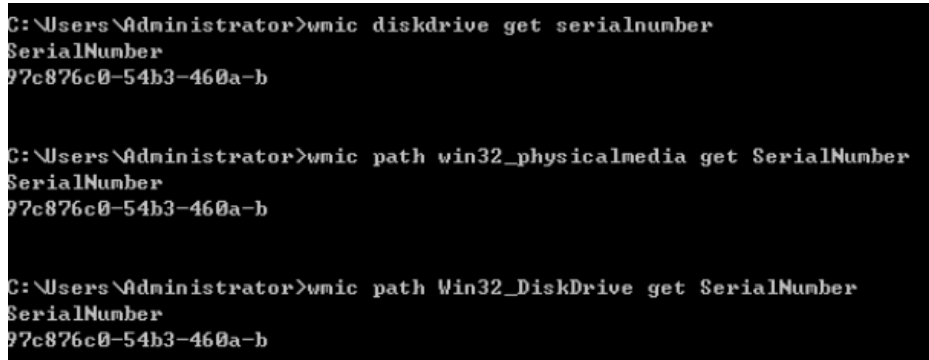
```
wmic diskdrive get serialnumber
```

```
wmic path win32_physicalmedia get SerialNumber
```

```
wmic path Win32_DiskDrive get SerialNumber
```

Information similar to the following is displayed:

**Figure 14-115** Obtaining the disk serial number



```
C:\Users\Administrator>wmic diskdrive get serialnumber
SerialNumber
97c876c0-54b3-460a-b

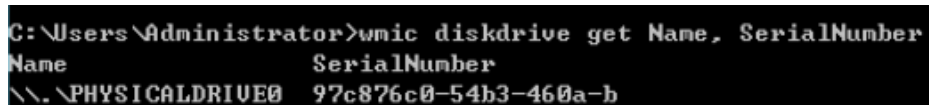
C:\Users\Administrator>wmic path win32_physicalmedia get SerialNumber
SerialNumber
97c876c0-54b3-460a-b

C:\Users\Administrator>wmic path Win32_DiskDrive get SerialNumber
SerialNumber
97c876c0-54b3-460a-b
```

2. Run the following command to check the disk corresponding to the serial number:

```
wmic diskdrive get Name, SerialNumber
```

**Figure 14-116** Checking the disk corresponding to the serial number



```
C:\Users\Administrator>wmic diskdrive get Name, SerialNumber
Name SerialNumber
\\.\PHYSICALDRIVE0 97c876c0-54b3-460a-b
```

### PowerShell

1. Start PowerShell as an administrator in a Windows OS.
2. Run the following command to check the disk on which the logical disk is created:

- Windows Server 2012 or later
  - i. Run the following command to check the disk on which the logical disk is created:  
**Get-CimInstance -ClassName Win32\_LogicalDiskToPartition | select Antecedent, Dependent | fl**  
As shown in [Figure 14-117](#), the disk is **Disk 0**.
  - ii. Run the following command to view the mapping between the serial number and the disk:  
**Get-Disk |select Number, SerialNumber**  
As shown in [Figure 14-117](#), the disk is **Disk 0**.

**Figure 14-117** Viewing the disk on which the logical disk is created

```
PS C:\Users\Administrator> Get-CimInstance -ClassName Win32_LogicalDiskToPartition |select Antecedent, Dependent | fl
Antecedent : Win32_DiskPartition (DeviceID = "Disk #0, Partition #1")
Dependent : Win32_LogicalDisk (DeviceID = "C:")

PS C:\Users\Administrator> Get-Disk |select Number, SerialNumber
Number SerialNumber

0 97c876c0-54b3-460a-b1dswfa16520d39517815206127
```

- Versions earlier than Windows 2012
  - i. Run the following command to check the disk on which the logical disk is created:  
**Get-WmiObject -Class Win32\_PhysicalMedia |select Tag, Serialnumber**
  - ii. Run the following command to view the mapping between the serial number and the disk:  
**Get-WmiObject -Class Win32\_LogicalDiskToPartition |select Antecedent, Dependent |fl**

## Using a Serial Number to Obtain a Disk Device Name (Linux)

If a serial number is displayed on the console, run either of the following commands to obtain the device name.

```
udevadm info --query=all --name=/dev/xxx | grep ID_SERIAL
```

```
ll /dev/disk/by-id/*
```

### NOTE

A serial number is the first 20 digits of a disk UUID.

For example, if the serial number of the VBD disk is 62f0d06b-808d-480d-8, run either of the following commands:

```
udevadm info --query=all --name=/dev/vdb | grep ID_SERIAL
```

```
ll /dev/disk/by-id/*
```

The following information is displayed:

```
[root@ecs-ab63 ~]# udevadm info --query=all --name=/dev/vdb | grep ID_SERIAL
E: ID_SERIAL=62f0d06b-808d-480d-8
```

```
[root@ecs-ab63 ~]# ll /dev/disk/by-id/*
lrwxrwxrwx 1 root root 9 Dec 30 15:56 /dev/disk/by-id/virtio-128d5bfd-f215-487f-9 -> ../vda
lrwxrwxrwx 1 root root 10 Dec 30 15:56 /dev/disk/by-id/virtio-128d5bfd-f215-487f-9-part1 -> ../vda1
lrwxrwxrwx 1 root root 9 Dec 30 15:56 /dev/disk/by-id/virtio-62f0d06b-808d-480d-8 -> ../vdb
```

**/dev/vdb** is the disk device name.

## Using a BDF to Obtain a Disk Device Name (Linux)

1. Run the following command to use a BDF to obtain the device name:

```
ll /sys/bus/pci/devices/BDF disk ID/virtio*/block
```

For example, if the BDF disk ID of the VBD disk is 0000:02:02.0, run the following command to obtain the device name:

```
ll /sys/bus/pci/devices/0000:02:02.0/virtio*/block
```

The following information is displayed:

```
[root@ecs-ab63 ~]# ll /sys/bus/pci/devices/0000:02:02.0/virtio*/block
total 0
drwxr-xr-x 8 root root 0 Dec 30 15:56 vdb
```

**/dev/vdb** is the disk device name.

## Using a WWN to Obtain the Disk Name (Windows)

1. Obtain the device identifier on the console by referring to [Obtaining the Disk ID of an ECS on the Console](#).

2. Manually convert the WWN.

For example, the obtained WWN (device identifier) is 68886030000**3252ffa**16520d39517815.

- a. Obtain the 21st to 17th digits that are counted backwards (**3252f**).
- b. Convert a hexadecimal (**3252f**) to a decimal (**206127**).

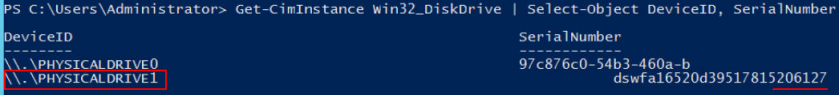
3. Start PowerShell as an administrator in a Windows OS.

4. Run the following command:

```
Get-CimInstance Win32_DiskDrive | Select-Object DeviceID, SerialNumber
```

5. In the command output, the disk whose serial number ends with **206127** is the disk corresponding to the WWN.

**Figure 14-118** Disk with the serial number ending with **206127**



```
PS C:\Users\Administrator> Get-CimInstance Win32_DiskDrive | Select-Object DeviceID, SerialNumber
DeviceID SerialNumber

\\.\PHYSICALDRIVE0 97c876c0-54b3-460a-b
\\.\PHYSICALDRIVE1 dswfa16520d39517815206127
```

## Using a WWN to Obtain a Disk Device Name (Linux)

1. Log in to the ECS as user **root**.
2. Run the following command to view the disk device name:

```
ll /dev/disk/by-id |grep WWN|grep scsi-3
```

For example, if the WWN obtained on the console is 6888603000008b32fa16688d09368506, run the following command:

```
ll /dev/disk/by-id |grep 6888603000008b32fa16688d09368506|grep scsi-3
```

The following information is displayed:

```
[root@host-192-168-133-148 block]# ll /dev/disk/by-id/ |grep 6888603000008b32fa16688d09368506 |
grep scsi-3
lrwxrwxrwx 1 root root 9 May 21 20:22 scsi-36888603000008b32fa16688d09368506 -> ../../sda
```

## 14.10.12 Why Does a Linux ECS with a SCSI Disk Attached Fail to Be Restarted?

### Symptom

For a Linux ECS with a SCSI disk attached, if you have enabled automatic SCSI disk attachment upon ECS startup in **/etc/fstab** and the disk drive letter (for example, **/dev/sdb**) is used, the ECS fails to restart.

### Possible Causes

SCSI disk allocation is determined based on the ID of the slot accommodating the disk as well as the available drive letter in the ECS. Each time you attach a disk to the ECS, an idle drive letter is automatically allocated in sequence. When the ECS starts, the disks are loaded in slot sequence. A slot ID corresponds to a drive letter.

After the SCSI disk is detached from the running ECS, the slot sequence for disks may change, leading to the disk drive letter being changed after the ECS is restarted. As a result, the slot IDs do not correspond to the drive letters, and the ECS fails to restart.

### Solution

1. Log in to the Linux ECS.
2. Run the following command to switch to user **root**:  
**sudo su -**
3. Run the following command to obtain the SCSI ID according to the drive letter of the SCSI disk:

```
ll /dev/disk/by-id/|grep Disk drive letter
```

For example, if the drive letter of the SCSI disk is **/dev/sdb**, run the following command:

```
ll /dev/disk/by-id/|grep sdb
```

```
CNA64_22:/opt/galax/eucalyptus/ecs_scripts # ll /dev/disk/by-id/|grep sdb
lrwxrwxrwx 1 root root 9 Dec 6 11:26 scsi-3688860300001436b005014f890338280 -> ../../sdb
lrwxrwxrwx 1 root root 9 Dec 6 11:26 wwn-0x688860300001436b005014f890338280 -> ../../sdb
```

4. Change the drive letter (for example, **/dev/sdb**) of the SCSI disk to the corresponding SCSI ID in the **/etc/fstab** file.

```
/dev/disk/by-id/SCSI ID
```

For example, if the SCSI ID obtained in step **3** is **scsi-3688860300001436b005014f890338280**, use the following data to replace **/dev/sdb**:


```
/dev/disk/by-id/scsi-3688860300001436b005014f890338280
```

## 14.10.13 Why Does a Disk Attached to a Windows ECS Go Offline?

### Symptom

A disk attached to a Windows ECS goes offline, and the system displays the message "The disk is offline because of policy set by an administrator.", as shown in [Figure 14-119](#).

Figure 14-119 Offline disk



### Possible Causes

Windows has three types of SAN policies: **OnlineAll**, **OfflineShared**, and **OfflineInternal**.

Table 14-8 SAN policies

| SAN Policy      | Description                                                                                                                                             |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| OnlineAll       | Indicates that all newly detected disks are automatically brought online.                                                                               |
| OfflineShared   | Indicates that all newly detected disks on sharable buses, such as FC or iSCSI, are offline by default, whereas disks on non-sharable buses are online. |
| OfflineInternal | Indicates that all newly detected disks are offline.                                                                                                    |

The SAN policy of certain Windows OSs, such as Windows Server 2008/2012 Enterprise Edition and Data Center Edition, is **OfflineShared** by default.

### Solution

Use the disk partition management tool DiskPart to obtain and set the SAN policy on the ECS to **OnlineAll**.

1. Log in to the Windows ECS.
2. Press **Win+R** to run **cmd.exe**.
3. Run the following command to access DiskPart:  
**diskpart**

4. Run the following command to view the SAN policy on the ECS:  
**san**
  - If the SAN policy is **OnlineAll**, run the **exit** command to exit DiskPart.
  - If the SAN policy is not **OnlineAll**, go to step 5.
5. Run the following command to change the SAN policy to **OnlineAll**:  
**san policy=onlineall**
6. (Optional) Use the ECS with the SAN policy changed to create a private image so that the configuration takes effect permanently. After an ECS is created using this private image, the disks attached to the ECS are online by default. You only need to initialize them.

## 14.10.14 Why Does the Disk Drive Letter Change After the ECS Is Restarted?

### Symptom

For a Linux ECS, the drive letter may change after an EVS disk is detached and then attached again, or after an EVS disk is detached and then the ECS is restarted.

### Root Cause

When a Linux ECS has multiple disks attached, it allocates drive letters in the attachment sequence and names the disks as **/dev/vda1**, **/dev/vdb1**, and **/dev/vdc1**, etc.

After a disk is detached and then attached again, or after a disk is detached and the ECS is restarted, the drive letter may change.

For example, an ECS has three disks attached: **/dev/vda1**, **/dev/vdb1**, and **/dev/vdc1**. The mounting parameters in **/etc/fstab** are as follows:

#### cat /etc/fstab

```
UUID=b9a07b7b-9322-4e05-ab9b-14b8050bdc8a / ext4 defaults 0 1
/dev/vdb1 /data1 ext4 defaults 0 0
/dev/vdc1 /data2 ext4 defaults 0 0
```

After **/dev/vdb1** is detached and the ECS is restarted, **/dev/vdc1** becomes **/dev/vdb1** and is mounted to **/data**. In such a case, no disk is mounted to **/data2**.

The change of drive letters can affect the running of applications. To solve this problem, you are advised to use the universally unique identifiers (UUIDs) to replace **/dev/vdx** because a UUID uniquely identifies a disk partition in the Linux OS.

### Solution

1. Log in to the ECS.
2. Run the following command to obtain the partition UUID:

**blkid** *Disk partition*

In this example, run the following command to obtain the UUID of the **/dev/vdb1** partition:

**blkid /dev/vdb1**

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# blkid /dev/vdb1
/dev/vdb1: UUID="b9a07b7b-9322-4e05-ab9b-14b8050cd8cc" TYPE="ext4"
```

The UUID of the **/dev/vdb1** partition is displayed.

3. Run the following command to open the **fstab** file using the vi editor:

```
vi /etc/fstab
```

4. Press **i** to enter the editing mode.
5. Move the cursor to the end of the file and press **Enter**. Then, add the following information:

```
UUID=b9a07b7b-9322-4e05-ab9b-14b8050cd8cc /data1 ext4 defaults 0 0
```

The parameters are defined as follows:

- **UUID=b9a07b7b-9322-4e05-ab9b-14b8050cd8cc**: UUID of a disk partition.
- **/data1**: directory on which the partition is mounted. You can run **df -TH** to query the directory.
- **ext4**: File system format of the partition. You can run **df -TH** to query the format.
- **defaults**: partition mount option. Normally, this parameter is set to **defaults**.
- **0** (the first one): whether to use Linux dump backup.
  - **0**: Linux dump backup is not used. Normally, dump backup is not used, and you can set this parameter to **0**.
  - **1**: Linux dump backup is used.
- **0** (the second one): fsck option, that is, whether to use fsck to check disks during startup.
  - **0**: fsck is not used.
  - If the mount point is the root partition (**/**), this parameter must be set to **1**.

When this parameter is set to **1** for the root partition, this parameter for other partitions must start with **2** so that the system checks the partitions in the ascending order of the values.

6. Repeat steps **2** to **5** to replace the UUID of **/dev/vdc1**.
7. Run the following command again to check the disk mounting parameters:

```
cat /etc/fstab
```

The following information is displayed:

```
UUID=b9a07b7b-9322-4e05-ab9b-14b8050bdc8a / ext4 defaults 0 1
UUID=b9a07b7b-9322-4e05-ab9b-14b8050cd8cc /data1 ext4 defaults 0 0
UUID=b9a07b7b-9322-4e05-ab9b-14b8050ab6bb /data2 ext4 defaults 0 0
```

## 14.10.15 How Can I Obtain Data Disk Information If Tools Are Uninstalled?

If you uninstall Tools from a Linux ECS in a non-PVOPS system, data disks cannot be identified. In such a case, you can create a new ECS and attach the data disks



of the original ECS to the new ECS and view information about the data disks. The procedure is as follows:

1. Log in to the management console and create a new ECS.

 **NOTE**

Ensure that the new ECS is located in the same AZ and has the same parameter settings as the original ECS.

2. (Optional) On the **Elastic Cloud Server** page, locate the row containing the original ECS, click **More** in the **Operation** column, and select **Stop**. On the **Stop ECS** page, select **Forcibly stop the preceding ECSs** and click **Yes** to forcibly stop the original ECS.

Manually refresh the **Elastic Cloud Server** page. The original ECS is stopped once the **Status** changes to **Stopped**.

 **NOTE**

The ECSs running certain OSs support online data disk detaching. If your OS supports this feature, you can detach data disks from the running ECS.

3. View information about the data disks attached to the original ECS.

 **NOTE**

If the original ECS has multiple data disks attached, repeat steps 4 to 6 to attach each data disk to the new ECS.

4. Click a data disk. The **Elastic Volume Service** page is displayed.
5. Select the data disk to be detached and click **Detach** in the **Operation** column. On the **Detach Disk** page, select the original ECS and click **OK** to detach the data disk from the original ECS.

Manually refresh the **Elastic Volume Service** page. The data disk is detached from the original ECS once the **Status** changes to **Available**.

6. Select the detached data disk and click **Attach** in the **Operation** column. On the **Attach Disk** page, click the new ECS, select a device name, and click **OK** to attach the data disk to the new ECS.

Manually refresh the EVS list. The data disk is attached to the new ECS once the **Status** value changes to **In-use**. You can then log in to the management console and view information about the data disk of the new ECS.

## 14.11 Passwords and Key Pairs

### 14.11.1 How Can I Set the Validity Period of the Image Password?

If an ECS cannot be logged in because of expired image password, you can contact the administrator for handling.

If the ECS can still be logged in, you can perform the following operations to set the password validity period.

## Procedure

The following operations use EulerOS 2.2 as an example.

1. Log in to the ECS.
2. Run the following command to check the password validity period:  
**vi /etc/login.defs**  
The value of parameter **PASS\_MAX\_DAYS** is the password validity period.
3. Run the following command to change the value of parameter **PASS\_MAX\_DAYS**:  
**chage -M 99999 user\_name**  
*99999* is the password validity period, and *user\_name* is the system user, for example, user **root**.

### NOTE

You are advised to configure the password validity period as needed and change it at a regular basis.

4. Run command **vi /etc/login.defs** to verify that the configuration has taken effect.

**Figure 14-120** Configuration verification

```
Password aging controls:
#
PASS_MAX_DAYS Maximum number of days a password may be used.
PASS_MIN_DAYS Minimum number of days allowed between password changes.
PASS_MIN_LEN Minimum acceptable password length.
PASS_WARN_AGE Number of days warning given before a password expires.
#
PASS_MAX_DAYS 99999
PASS_MIN_DAYS 0
PASS_MIN_LEN 5
PASS_WARN_AGE 7
```

## 14.11.2 Why Does Login to My ECS Using the Reset Password Fail?

### Scenarios

You cannot use the new password that you set to log in to the ECS.

### NOTE

Ensure that the one-click password reset plug-in is not blocked by security software. Otherwise, the one-click password reset function is unavailable.

After the password is reset, you must restart the ECS for the new password to take effect.

### Windows

Perform the following operations to locate the fault:

- Step 1** Check whether port 80 in the outbound direction of the security group is permitted.
1. Log in to management console.

2. Select the target ECS to switch to the page that provides details about the ECS.
3. On the **Security Groups** tab, check whether the outbound rule allows access from port 80.

In the default security group rule, all ports are allowed in the outbound direction.

**Step 2** Check whether DHCP is enabled in the VPC of the ECS.

1. On the ECS details page, click the VPC name to navigate to the VPC console.
2. In the VPC list, click the VPC name.
3. In the **Network Components** area, click the number in the **Subnets** row to go to the **Subnets** page.
4. In the subnet list, click the subnet name to view its details.
5. In the **Gateway and DNS Information** area, check whether DHCP is enabled.

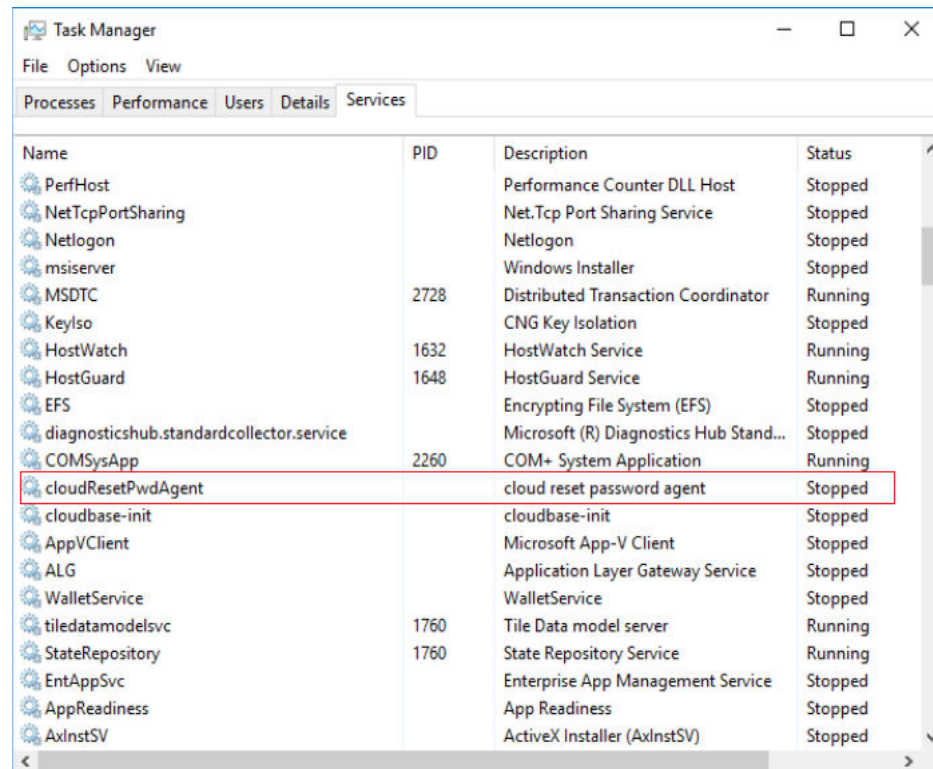
**Step 3** If both the security group and DHCP are properly configured but one-click password reset fails to take effect, use the original password to log in to the ECS.

- If the original password is invalid, reset the password. For details, see [Resetting the Password for Logging In to a Windows ECS](#).
- If the original password is valid, use it to log in to the ECS and reset the password. For details, see [Changing the Login Password on an ECS](#).

**Step 4** Check whether the password reset plug-in CloudResetPwdAgent has been installed on the ECS. To do so, perform the following operations:

Start the **Task Manager** and check whether **cloudResetPwdAgent** is displayed on the **Services** tab. As shown in the [Figure 14-121](#), the password reset plug-in has been installed on the ECS. If no, the one-click password reset plug-in has not been installed on the ECS.

To install the plug-in, see [Installing the One-Click Password Reset Plug-in on an ECS](#).

**Figure 14-121** Successful plug-in installation

----End

## Linux

Perform the following operations to locate the fault:

- Step 1** Check whether port 80 in the outbound direction of the security group is permitted.
1. Log in to management console.
  2. Select the target ECS to switch to the page that provides details about the ECS.
  3. On the **Security Groups** tab, check whether the outbound rule allows access from port 80.  
In the default security group rule, all ports are allowed in the outbound direction.
- Step 2** Check whether DHCP is enabled in the VPC of the ECS.
1. On the ECS details page, click the VPC name to navigate to the VPC console.
  2. In the VPC list, click the VPC name.
  3. In the **Network Components** area, click the number in the **Subnets** row to go to the **Subnets** page.
  4. In the subnet list, click the subnet name to view its details.
  5. In the **Gateway and DNS Information** area, check whether DHCP is enabled.
- Step 3** If both the security group and DHCP are properly configured but one-click password reset fails to take effect, use the original password to log in to the ECS.

- If the original password is invalid, enter the single-user mode and reset the password.
- If the original password can be used, perform the following operations for further check:
  - a. Use the original password to log in to the ECS.
  - b. Run the `curl http://169.254.169.254/openstack/latest/resetpwd_flag` command to check whether the one-click password reset function is available.
    - If the returned value is **true**, the password can be reset with a few clicks.
    - If any other value is returned, the password cannot be reset.

```
root@ecs-f7e2 ~]# service cloudResetPwdAgent status
cloudResetPwdAgent is not running.
root@ecs-f7e2 ~]# curl http://169.254.169.254/openstack/latest/reset_pwd_flag
{"message":"API not found","request_id":"c3b0eb06-156d-44c7-a044-891926965403"}
root@ecs-f7e2 ~]# curl http://169.254.169.254/openstack/latest/resetpwd_flag
{"resetpwd_flag": "True"}root@ecs-f7e2 ~]#
```

#### Step 4 Check whether **CloudResetPwdAgent** has been installed.

1. Check whether the **CloudResetPwdAgent** directory is available in the root directory on the ECS.
  - If the directory is available, go to [Step 4.2](#).
  - If the directory is not available, the one-click password reset plug-in has not been installed on the ECS.

To install the plug-in, see [Installing the One-Click Password Reset Plug-in on an ECS](#).

2. Run the following command to check the CloudResetPwdAgent status:

##### **service cloudResetPwdAgent status**

If the command output is "unrecognized service", the one-click password reset plug-in has not been installed on the ECS.

To install the plug-in, see [Installing the One-Click Password Reset Plug-in on an ECS](#).

----End

### 14.11.3 Why Am I Seeing the Message Indicating That the Port Is Used by a One-Click Password Reset Plug-in?

#### Symptom

When you attempt to run an application on an ECS, the system displays a message indicating that the required port is used by a one-click password reset plug-in.

#### Possible Causes

If an ECS works in AUTO mode, when its one-click password reset plug-in starts, the plug-in randomly use a port, which may be a service port.

**NOTE**

The one-click password reset plug-in has been upgraded to work in PIPE mode by default.

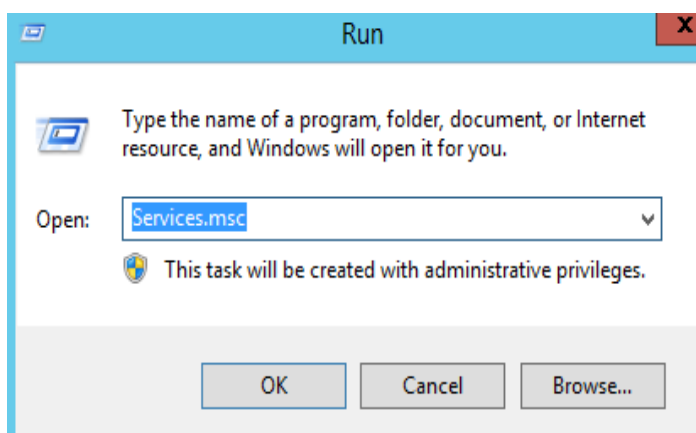
- Newly created ECSs run in PIPE mode by default, and their one-click password reset plug-ins will not use service ports.
- Existing ECSs still work in AUTO mode, in which the plug-ins randomly select idle ports with the smallest port numbers ranging from 31000 to 32999.

## Method 1 (Recommended): Modifying the wrapper Files of the One-Click Password Reset Plug-in for the PIPE Mode

In the wrapper files, change **AUTO (SOCKET)** to **PIPE**. After the change, the plug-in will not use service ports.


1. Open CloudResetPwdAgent configuration files.
  - Linux  
**/CloudResetPwdAgent/conf/wrapper.conf** and **/CloudResetPwdUpdateAgent/conf/wrapper.conf**
  - Windows  
**C:\CloudResetPwdUpdateAgent\conf\wrapper.conf** and **C:\CloudResetPwdUpdateAgent\conf\wrapper.conf**
2. Add the following data to the end of the files:  
**wrapper.backend.type=PIPE**
3. Restart CloudResetPwdUpdateAgent.
  - Linux  
**/CloudResetPwdUpdateAgent/bin/cloudResetPwdUpdateAgent.script restart**
  - Windows
    - i. Press **Win+R** to start the **Run** text box.
    - ii. Enter **services.msc** and click **OK**.

**Figure 14-122** Run



- iii. Right-click **cloud reset password update agent** and choose **Restart** from the shortcut menu.

Figure 14-123 Services (Local)



## Method 2: Modifying the Configuration to Change the Port Range

Modify the CloudResetPwdAgent configuration to change the default port range for the password reset plug-in so that the service port is out of the port range.

For example, to change the port range for the password reset plug-in to 40000-42000, perform the following operations:

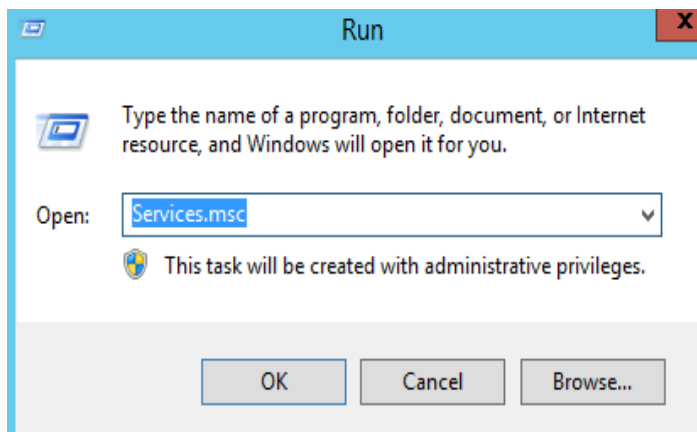
1. Open CloudResetPwdAgent configuration files.
  - Linux  
`/CloudResetPwdAgent/conf/wrapper.conf` and `/CloudResetPwdUpdateAgent/conf/wrapper.conf`
  - Windows  
`C:\CloudResetPwdUpdateAgent\conf\wrapper.conf` and `C:\CloudResetPwdUpdateAgent\conf\wrapper.conf`
2. Add the following data to the configuration files:  
`wrapper.port.min=40000`  
`wrapper.port.max=41000`  
`wrapper.jvm.port.min=41001`  
`wrapper.jvm.port.max=42000`

Figure 14-124 Modifying the configuration files

```
[root@ ~]# tail -n 4 /CloudResetPwdUpdateAgent/conf/wrapper.conf
wrapper.port.min=40000
wrapper.port.max=41000
wrapper.jvm.port.min=41001
wrapper.jvm.port.max=42000
[root@sluo-ecs-9545 ~]# tail -n 4 /CloudResetPwdAgent/conf/wrapper.conf
wrapper.port.min=40000
wrapper.port.max=41000
wrapper.jvm.port.min=41001
wrapper.jvm.port.max=42000
```


3. Restart CloudResetPwdUpdateAgent.
  - Linux  
`/CloudResetPwdUpdateAgent/bin/cloudResetPwdUpdateAgent.script`  
`restart`
  - Windows
    - i. Press **Win+R** to start the **Run** text box.
    - ii. Enter `services.msc` and click **OK**.

Figure 14-125 Run



- iii. Right-click **cloud reset password update agent** and choose **Restart** from the shortcut menu.

Figure 14-126 Services (Local)



## 14.11.4 Why Does the One-Click Password Reset Plug-in Use Too Much VIRT and SHR?

### Symptom

The one-click password reset plug-in uses too much VIRT and SHR.

Figure 14-127 Viewing the virtual memory usage

```
top - 14:56:06 up 4 days, 3:22, 1 user, load average: 0.00, 0.02, 0.05
Tasks: 1 total, 0 running, 1 sleeping, 0 stopped, 0 zombie
%Cpu(s): 0.5 us, 0.5 sy, 0.0 ni, 99.0 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 3864028 total, 137704 free, 183268 used, 3543056 buff/cache
KiB Swap: 0 total, 0 free, 0 used, 3358852 avail Mem

 PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND
18240 root 20 0 2513828 61876 13036 S 0.0 1.6 0:00.75 java
```

### Root Causes

Memory used by applications is from the glibc memory pool. In versions earlier than glibc 2.10, there is only one memory pool named main arena. In glibc 2.10 and later versions, there is a memory pool named thread arena. Therefore, applications can use memory from two memory pools, which results in high usage of VIRT and SHR. You can limit the memory that can be used by the one-click password reset plug-in.



## Procedure

1. Modify the parameters of the one-click password reset plug-in.
  - a. Run the following command to view the configuration file.  
**vim /CloudResetPwdUpdateAgent/conf/wrapper.conf**

Figure 14-128 wrapper.conf before modification

```
Initial Java Heap Size (in MB)
#wrapper.java.initmemory=16

Maximum Java Heap Size (in MB)
#wrapper.java.maxmemory=64
```

- b. Delete the comment tag (#) at the beginning of **wrapper.java.initmemory=16** and **wrapper.java.maxmemory=64**.

Figure 14-129 wrapper.conf after modification

```
Initial Java Heap Size (in MB)
wrapper.java.initmemory=16

Maximum Java Heap Size (in MB)
wrapper.java.maxmemory=64
#
```

2. Run the following commands to restart the one-click password reset plug-in.  
**cd /CloudResetPwdUpdateAgent/bin/**  
**./cloudResetPwdUpdateAgent.script restart**

## 14.11.5 How Can I Obtain the Key Pair Used by My ECS?

### Symptom

You have created multiple key pairs, and you are trying to find the key pair to log in to the target ECS.


### Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. On the **Elastic Cloud Server** page, select the target ECS.
4. Click the name of the target ECS.  
The page providing details about the ECS is displayed.

5. Obtain the **Key Pair** value.  
The value is the key pair used by the ECS.

### 14.11.6 What Should I Do If a Key Pair Cannot Be Imported?

If you use Internet Explorer 9 to access the management console, the key pair may fail to import. In this case, perform the following steps to modify browser settings and then try again:

1. Click  in the upper right corner of the browser.
2. Select **Internet Options**.
3. Click the **Security** tab in the displayed dialog box.
4. Click **Internet**.
5. If the security level indicates **Custom**, click **Default Level** to restore to the default settings.
6. Move the scroll bar to set the security level to **Medium** and click **Apply**.
7. Click **Custom Level**.
8. Set **Initialize and script ActiveX controls not marked as safe for scripting** to **Prompt**.
9. Click **Yes**.

### 14.11.7 Why Does the Login to My Linux ECS Using a Key File Fail?

#### Symptom

When you use the key file created during your Linux ECS creation to log in to the ECS, the login fails.

#### Possible Causes

Possible causes vary depending on the image used to create the Linux ECS.

- Cause 1: The image that you used to create the Linux ECS is a private image, on which Cloud-Init is not installed.
- Cause 2: Cloud-Init is installed on the image, but you did not obtain the key pair when you created the ECS.

#### Solution

- If the issue is a result of cause 1, proceed as follows:  
If you created a private image without installing Cloud-Init, you cannot customize the ECS configuration. As a result, you can log in to the ECS only using the original image password or key pair.  
The original image password or key pair is the OS password or key pair you configured when you created the private image.
- If the issue is a result of cause 2, proceed as follows:
  - a. Locate the row containing the target ECS, click **More** in the **Operation** column, and select **Restart**.

- b. Use the key file to log in to the ECS again and check whether the login is successful.
  - If the login is successful, no further action is required.
  - If the login fails, contact the administrator for technical support.

## 14.11.8 Why Does a Key Pair Created Using `puttygen.exe` Fail to Be Imported on the Management Console?

### Symptom

When you try to import a key pair that you created using **puttygen.exe** on the management console, the system displays a message indicating that the import failed.

### Possible Causes

The format of the public key content does not meet system requirements.


If you store a public key by clicking **Save public key** on PuTTY Key Generator, the format of the public key content will change. You cannot import the key on the management console.

### Solution


Use the locally stored private key and **PuTTY Key Generator** to restore the format of the public key content. Then, import the public key to the management console.

1. Double-click **puttygen.exe** to open **PuTTY Key Generator**.

**Figure 14-130** PuTTY Key Generator



2. Click **Load** and select the private key.  
The system automatically loads the private key and restores the format of the public key content in **PuTTY Key Generator**. The content in the red box in **Figure 14-131** is the public key whose format meets system requirements.

**Figure 14-131** Restoring the format of the public key content

3. Copy the public key content to a .txt file and save the file in a local directory.
4. Import the public key to the management console.
  - a. Log in to the management console.
  - b. Under **Computing**, click **Elastic Cloud Server**.
  - c. In the navigation pane on the left, choose **Key Pair**.
  - d. On the key pair page, click **Import Key Pair**.
  - e. Copy the public key content in the .txt file to **Public Key Content** and click **OK**.

## 14.11.9 What Is the Cloudbase-Init Account in Windows ECSs Used for?

### Description

In Windows ECSs, **cloudbase-init** is the default account of the Cloudbase-Init agent program. It is used to obtain the metadata and execute configurations when an ECS starts.

#### NOTE

This account is unavailable on Linux ECSs.

Do not modify or delete this account or uninstall the Cloudbase-Init agent program. Otherwise, you will be unable to insert data to initialize an ECS created using a Windows private image.

## Security Hardening for Randomized `cloudbase-init` Passwords

In Cloudbase-Init 0.9.10, the security of randomized `cloudbase-init` passwords has been hardened to ensure that the hash values (LM-HASH and NTLM-HASH) of the passwords are different.

In Windows, the hash passwords are in the format of "Username:RID:LM-HASH value:NT-HASH value".

For example, in

```
"Administrator:500:C8825DB10F2590EAAAD3B435B51404EE:683020925C5D8569C23AA724774CE9CC:::",
```

- Username: **Administrator**
- RID: **500**
- LM-HASH value: **C8825DB10F2590EAAAD3B435B51404EE**
- NT-HASH value: **683020925C5D8569C23AA724774CE9CC**

Use an image to create two ECSs, `ecs01` and `ecs02`. Then, verify that the hash values of the `cloudbase-init` account for the two ECSs are different.

- LM-HASH and NTLM-HASH values of the `cloudbase-init` account for `ecs01`

Figure 14-132 `ecs01`

```
----- BEGIN DUMP -----
c\cloudbase-init:1003:AAD3B435B51404EEAAD3B435B51404EE:CCA3BDDEB517A0E2342AEB34C0473C39:::
Guest:501:AAD3B435B51404EEAAD3B435B51404EE:31D6CFE0D16AE931B73C59D7E0C089C0:::
Administrator:500:AAD3B435B51404EEAAD3B435B51404EE:27CF57575EB83D9A6D7D27831157A947:::
----- END DUMP -----
3 dumped accounts
```

- LM-HASH and NTLM-HASH values of the `cloudbase-init` account for `ecs02`

Figure 14-133 `ecs02`

```
----- BEGIN DUMP -----
c\cloudbase-init:1003:AAD3B435B51404EEAAD3B435B51404EE:5B635D5F5306E26E0EE66915D7C1CA98:::
Guest:501:AAD3B435B51404EEAAD3B435B51404EE:31D6CFE0D16AE931B73C59D7E0C089C0:::
Administrator:500:AAD3B435B51404EEAAD3B435B51404EE:0501525C0083243750D23927A82070B6:::
----- END DUMP -----
3 dumped accounts
```

## 14.11.10 What Should I Do If Cloud-Init Does Not Work After Python Is Upgraded?

### Symptom

Take an ECS running CentOS 6.8 as an example. After Python was upgraded from 2.6 to 2.7, Cloud-Init did not work. Data, such as the login password, key, and hostname could not be imported to the ECS using Cloud-Init.

After the `cloud-init -v` command was executed to view the Cloud-Init version, the system displayed errors, as shown in [Figure 14-134](#).

Figure 14-134 Improper running of Cloud-Init

```
[root@ecs-8560 ~]# cloud-init -v
Traceback (most recent call last):
 File "/usr/bin/cloud-init", line 39, in <module>
 from cloudinit import patcher
ImportError: No module named cloudinit
[root@ecs-8560 ~]# cloud-init init --local
Traceback (most recent call last):
 File "/usr/bin/cloud-init", line 39, in <module>
 from cloudinit import patcher
ImportError: No module named cloudinit
[root@ecs-8560 ~]#
```

### Possible Causes

The Python version used by Cloud-Init was incorrect.

### Solution

Change the Python version used by Cloud-Init to the source version. To do so, change the environment variable value of `/usr/bin/cloud-init` from the default value `#!/usr/bin/python` to `#!/usr/bin/python2.6`.

Figure 14-135 Changing the Python version

```
[root@ecs-8560 ~]# head -n 1 /usr/bin/cloud-init
#!/usr/bin/python2.6
[root@ecs-8560 ~]# ls /usr/bin/python* -lh
lrwxrwxrwx 1 root root 24 Jul 19 10:55 /usr/bin/python -> /usr/local/bin/python2.7
lrwxrwxrwx 1 root root 6 Jun 9 2017 /usr/bin/python2 -> python
-rwxr-xr-x 1 root root 8.9K Aug 18 2016 /usr/bin/python2.6
```

## 14.12 Network Configurations

### 14.12.1 Can Multiple EIPs Be Bound to an ECS?

#### Scenarios

Multiple EIPs can be bound to an ECS, but this operation is not recommended.

If an ECS has multiple NICs attached and you want to bind multiple EIPs to this ECS, you need to configure policy-based routes for these NICs so that these extension NICs can communicate with external works. For details, see [Configuration Example](#).

#### Configuration Example

[Table 14-9](#) lists ECS configurations.

Table 14-9 ECS configurations

| Parameter | Configuration |
|-----------|---------------|
| Name      | ecs_test      |

| Parameter     | Configuration    |
|---------------|------------------|
| Image         | CentOS 6.5 64bit |
| EIP           | 2                |
| Primary NIC   | eth0             |
| Secondary NIC | eth1             |

### Example 1:

If you intend to access public network 11.11.11.0/24 through standby NIC **eth1**, perform the following operations to configure a route:

1. Log in to the ECS.
2. Run the following command to configure a route:

```
ip route add 11.11.11.0/24 dev eth1 via 192.168.2.1
```

In the preceding command, **192.168.2.1** is the gateway IP address of standby NIC **eth1**.

### Example 2:

Based on example 1, if you intend to enable routing for default public network traffic through standby NIC **eth1**, perform the following operations to configure a route:

1. Log in to the ECS.
2. Run the following command to delete the default route:

```
ip route delete default
```

---

#### NOTICE

Exercise caution when deleting the default route because this operation will interrupt the network and result in SSH login failures.

3. Run the following command to configure a new default route:

```
ip route add 0.0.0.0/0 dev eth1 via 192.168.2.1
```

In the preceding command, **192.168.2.1** is the gateway IP address of standby NIC **eth1**.

## 14.12.2 Can an ECS Without an EIP Bound Access the Internet?

Yes.

You can configure the SNAT server so that the ECS without an EIP bound can access the Internet.

For details, see "Configuring an SNAT Server" in *Virtual Private Cloud User Guide*.



## 14.12.3 What Should I Do If an EIP Cannot Be Pinged?

### Symptom

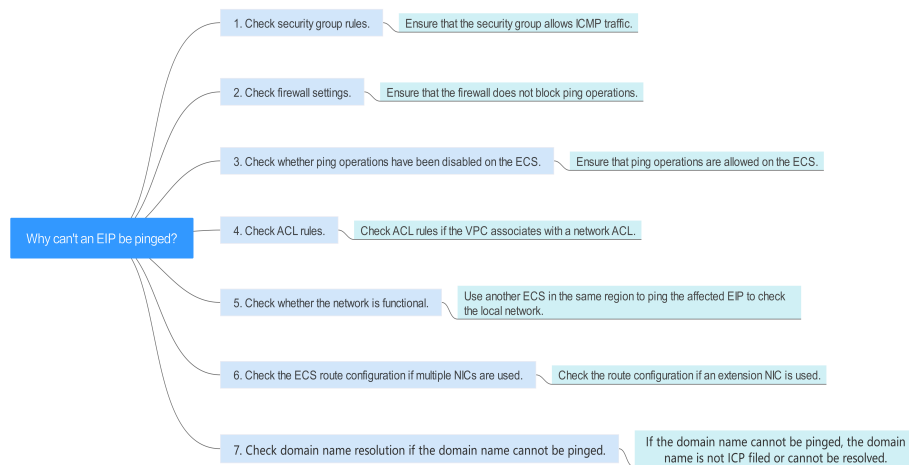
After you purchase an EIP and bind it to an ECS, the local host or other cloud servers cannot ping the EIP of the ECS.

### Fault Locating

The following fault causes are sequenced based on their occurrence probability.

If the fault persists after you have ruled out a cause, check other causes.

**Figure 14-136** Method of locating the failure to ping an EIP



**Table 14-10** Method of locating the failure to ping an EIP

| Possible Cause                                         | Solution                                                                                                                                                           |
|--------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ICMP access rules are not added to the security group. | Add ICMP access rules to the security group. For details, see <a href="#">Checking Security Group Rules</a> .                                                      |
| Ping operations are prohibited on the firewall.        | Allow ping operations on the firewall. For details, see <a href="#">Checking Firewall Settings</a> .                                                               |
| Ping operations are prohibited on the ECS.             | Allow ping operations on the ECS. For details, see <a href="#">Checking Whether Ping Operations Have Been Disabled on the ECS</a> .                                |
| Network ACL is associated.                             | If the VPC is associated with a network ACL, check the network ACL rules. For details, see <a href="#">Checking ACL Rules</a> .                                    |
| A network exception occurred.                          | Use another ECS in the same region to check whether the local network is functional. For details, see <a href="#">Checking Whether the Network Is Functional</a> . |

| Possible Cause                                               | Solution                                                                                                                                                                                                                             |
|--------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Routes are incorrectly configured if multiple NICs are used. | If the network is inaccessible due to an extension NIC, the fault is generally caused by incorrect route configurations. To resolve this issue, see <a href="#">Checking the ECS Route Configuration If Multiple NICs Are Used</a> . |

## Checking Security Group Rules

ICMP is used for the ping command. Check whether the security group accommodating the ECS allows ICMP traffic.

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. On the **Elastic Cloud Server** page, click the name of the target ECS.  
The page providing details about the ECS is displayed.
4. Click the **Security Groups** tab, expand the information of the security group, and view security group rules.
5. Click the security group ID.  
The system automatically switches to the **Security Group** page.
6. On the **Outbound Rules** page, click **Add Rule**. In the displayed dialog box, set required parameters to add an outbound rule.

**Table 14-11** Security group rules

| Transfer Direction | Type | Protocol/Port Range | Source                                             |
|--------------------|------|---------------------|----------------------------------------------------|
| Outbound           | IPv4 | ICMP/Any            | 0.0.0.0/0<br>0.0.0.0/0 indicates all IP addresses. |

7. On the **Inbound Rules** tab, click **Add Rule**. In the displayed dialog box, set required parameters to add an inbound rule.

**Table 14-12** Security group rules

| Transfer Direction | Type | Protocol/Port Range | Source                                             |
|--------------------|------|---------------------|----------------------------------------------------|
| Inbound            | IPv4 | ICMP/Any            | 0.0.0.0/0<br>0.0.0.0/0 indicates all IP addresses. |

8. Click **OK** to complete the security rule configuration.

## Checking Firewall Settings

If a firewall is enabled on the ECS, check whether the firewall blocks the ping operations.

### Linux

1. Consider CentOS 7 as an example. Run the following command to check the firewall status:

#### **firewall-cmd --state**

If **running** is displayed in the command output, the firewall has been enabled.

2. Check whether there is any ICMP rule blocking the ping operations.

#### **iptables -L**

If the command output shown in [Figure 14-137](#) is displayed, there is no ICMP rule blocking the ping operations.

**Figure 14-137** Checking firewall rules

```
[root@ecs-3c4e ~]# iptables -L
Chain INPUT (policy ACCEPT)
target prot opt source destination icmp echo-request
ACCEPT icmp -- anywhere anywhere

Chain FORWARD (policy ACCEPT)
target prot opt source destination

Chain OUTPUT (policy ACCEPT)
target prot opt source destination icmp echo-reply
ACCEPT icmp -- anywhere anywhere

[root@ecs-3c4e ~]#
```

If the ping operations are blocked by an ICMP rule, run the following commands to modify the rule for unblocking:

```
iptables -A INPUT -p icmp --icmp-type echo-request -j ACCEPT
```

```
iptables -A OUTPUT -p icmp --icmp-type echo-reply -j ACCEPT
```

### Windows

1. Log in to the Windows ECS, click the Windows icon in the lower left corner of the desktop, and choose **Control Panel > Windows Firewall**.
2. Click **Turn Windows Firewall on or off**.  
View and set the firewall status.
3. If the firewall is **On**, go to [4](#).
4. Check the ICMP rule statuses in the firewall.

- a. In the navigation pane on the **Windows Firewall** page, click **Advanced settings**.

- b. Enable the following rules:

**Inbound Rules: File and Printer Sharing (Echo Request - ICMPv4-In)**

**Outbound Rules: File and Printer Sharing (Echo Request - ICMPv4-Out)**

If IPv6 is enabled, enable the following rules:

## Inbound Rules: File and Printer Sharing (Echo Request - ICMPv6-In) Outbound Rules: File and Printer Sharing (Echo Request - ICMPv6-Out)

Figure 14-138 Inbound Rules

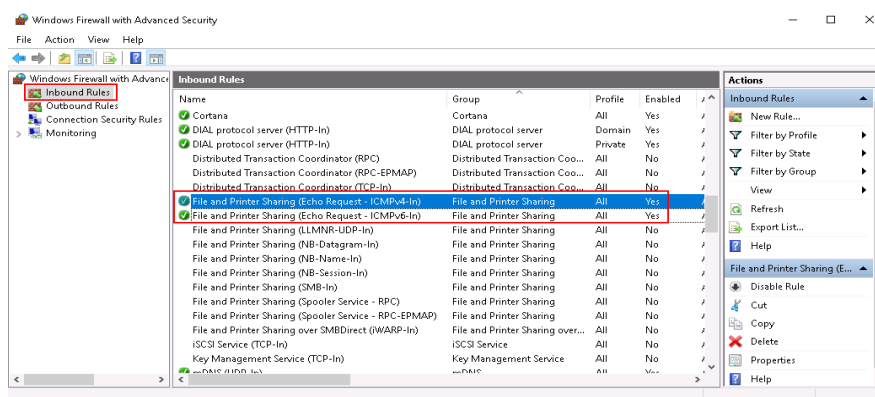
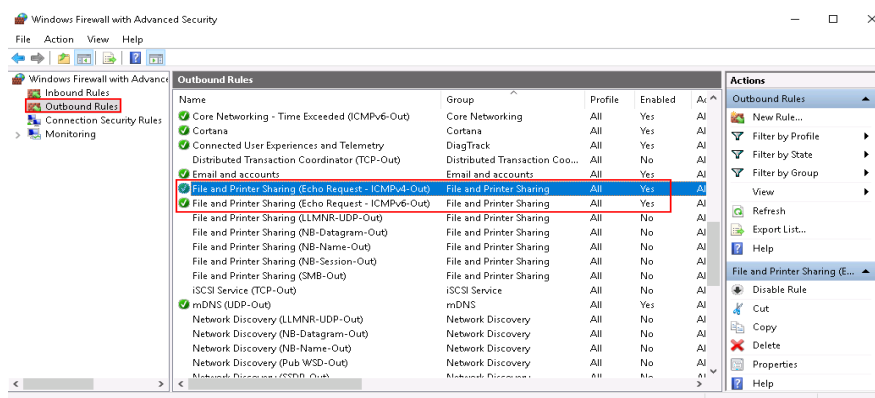


Figure 14-139 Outbound Rules



## Checking Whether Ping Operations Have Been Disabled on the ECS

### Windows

Enable ping operations using the CLI.

1. Start the **Run** dialog box. Enter **cmd** and press **Enter**.
2. Run the following command to enable ping operations:  
**netsh firewall set icmpsetting 8**

### Linux

Check the ECS kernel parameters.

1. Check the **net.ipv4.icmp\_echo\_ignore\_all** value in the **/etc/sysctl.conf** file. Value **0** indicates that ping operations are allowed, and value **1** indicates that ping operations are prohibited.
2. Allow ping operations.
  - Run the following command to temporarily allow the ping operations:  
**#echo 0 >/proc/sys/net/ipv4/icmp\_echo\_ignore\_all**

- Run the following command to permanently allow the ping operations:  
`net.ipv4.icmp_echo_ignore_all=0`

## Checking ACL Rules

By default, no ACL is configured for a VPC. If a network ACL is associated with a VPC, check the ACL rules.

1. Check whether the subnet of the ECS has been associated with a network ACL.  
If an ACL name is displayed, the network ACL has been associated with the ECS.
2. Click the ACL name to view its status.
3. If the network ACL is enabled, add an ICMP rule to allow traffic.

### NOTE

The default network ACL rule denies all incoming and outgoing packets. If a network ACL is disabled, the default rule is still effective.

## Checking Whether the Network Is Functional

1. Use another ECS in the same region to check whether the local network is functional.  
Use another ECS in the same region to ping the affected EIP. If the EIP can be pinged, the VPC is functional. In such a case, rectify the local network fault and ping the affected EIP again.
2. Check whether the link is accessible.  
A ping failure is caused by packet loss or long delay, which may be caused by link congestion, link node faults, or heavy load on the ECS.

## Checking the ECS Route Configuration If Multiple NICs Are Used

Generally, the default route of an OS will preferentially select the primary NIC. If an extension NIC is selected in a route and the network malfunctions, this issue is typically caused by incorrect route configuration.

- If the ECS has multiple NICs, check whether the default route is available.
  - a. Log in to the ECS and run the following command to check whether the default route is available:

**ip route**

**Figure 14-140** Default route

```
[root@do-not-del-scy ~]# ip route
default via 192.168.2.1 dev eth0
169.254.0.0/16 dev eth0 scope link metric 1002
169.254.169.254 via 192.168.2.1 dev eth0 proto static
192.168.2.0/24 dev eth0 proto kernel scope link src 192.168.2.112
```

- b. If the route is unavailable, run the following command to add it:  
**ip route add default via XXXX dev eth0**

 NOTE

In the preceding command, *XXXX* specifies a gateway IP address.

- If the ECS has multiple NICs and the EIP is bound to an extension NIC, configure policy routing on the ECS for network communication with the extension NIC.

## 14.12.4 Why Can I Remotely Access an ECS But Cannot Ping It?

### Symptom

You can remotely access an ECS but when you ping the EIP bound to the ECS, the ping operation fails.

### Possible Causes

A desired inbound rule is not added for the security group, and ICMP is not enabled.

### Solution

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. On the **Elastic Cloud Server** page, click the name of the target ECS.  
The page providing details about the ECS is displayed.
4. Click the **Security Groups** tab, expand the information of the security group, and click the security group ID.
5. On the **Inbound Rules** tab of the **Security Group** page, click **Add Rule**.
6. Add an inbound rule for the security group and enable ICMP.
  - **Protocol:** ICMP
  - **Source:** IP address 0.0.0.0/0

## 14.12.5 Will NICs Added to an ECS Start Automatically?

Based on test results, if the ECS runs CentOS 7.0, NICs added to the ECS cannot start automatically. You must start the NICs manually.

## 14.12.6 How Can I Check Whether the Network Communication Is Normal Between Two ECSs Equipped with an InfiniBand NIC Driver?

For high-performance H2 ECSs equipped with an InfiniBand NIC driver (InfiniBand ECSs for short), perform the following operations to check whether the driver installation is successful and whether the network communication between the ECSs is normal.

**NOTE**

During the check, if your ECS has no command tool installed, such as `ibstat`, obtain the tool from the installation package for the InfiniBand NIC driver and install the tool.

**Step 1** Check whether the NICs of the InfiniBand ECSs are functional.

1. Log in to the ECS.
2. Run the following command to check whether the NIC is functional:

**ibstat**

- If it is functional, go to [Step 2](#).
- If it is not functional, contact the administrator for technical support.

**Step 2** Check whether the network communication between two InfiniBand ECSs is normal.

1. Log in to one InfiniBand ECS and run the following command:

```
ib_write_bw -x 0 --pkey_index 0
```

2. Log in to the other InfiniBand ECS and run the following command:

```
ib_write_bw -x 0 --pkey_index 0 ip_addr
```

In the preceding command, *ip\_addr* is the NIC IP address of the first InfiniBand ECS.

3. Check whether the terminal display is correct.

**Figure 14-141** Normal network communication

```
root@host-11-11-11-111 MLNX_OFED_LINUX-3.4-1.0.0.0-rhel7.2-x86_64]# ib_write_bw -x 0 --pkey_index 0 4.29.43.20

RDMA Write BW Test
Dual-port : OFF Device : mlx5_0
Number of qps : 1 Transport type : IB
Connection type : RC Using SRQ : OFF
TX depth : 128
CQ Moderation : 100
Mtu : 4096[B]
Link type : IB
GID index : 0
Max inline data : 0[B]
rdma_cm QPs : OFF
Data ex. method : Ethernet

Local address: LID 0x05 QPN 0x0067 PSN 0xaaccfb RKey 0x001c0c VAddr 0x007fb3cd1b0000
GID: 254:128:00:00:00:00:00:00:03:00:135:40:178
remote address: LID 0x05 QPN 0x006a PSN 0xebbf6d RKey 0x001c10 VAddr 0x007fdad5990000
GID: 254:128:00:00:00:00:00:00:01:03:00:135:40:178

#bytes #iterations BW peak[MB/sec] BW average[MB/sec] MsgRate[Mpps]
65536 5000 12132.78 11900.18 0.190403
```

- If the terminal display is shown in [Figure 14-141](#), the network communication between the two InfiniBand ECSs is normal.
- If the InfiniBand network is inaccessible, contact the administrator for technical support.

----End

## 14.12.7 How Can I Manually Configure an IP Address for an InfiniBand NIC?

IP over InfiniBand (IPoIB) allows IP data transmission over InfiniBand. For SUSE high-performance H2 and HL1 ECSs, if IPoIB is required, you must manually configure an IP address for the InfiniBand NIC after installing the InfiniBand NIC driver.

## Prerequisites

The InfiniBand NIC driver has been installed on the high-performance H2 or HL1 ECSs.

## Background

To prevent IP address conflict of the InfiniBand NICs configured for the ECSs of a tenant, determine the IP address to be configured for an InfiniBand NIC according to the IP addresses available in the VPC. The method is as follows:

For example, if the first two eight-bits of the IP address (specified by **IPADDR**) to be configured for the InfiniBand NIC are consistently **169.254**, the latter two eight-bits must be the same as those of the **eth0** IP address, and the subnet mask must be the same as that of the **eth0** NIC.

An example is provided as follows:

If the IP address of the **eth0** NIC is 192.168.0.100/24, the IP address to be configured for the InfiniBand NIC is 169.254.0.100/24.

## Procedure

1. Log in to the ECS.
2. Run the following command to switch to user **root**:  
**sudo su -**
3. Run the following command to edit the **/etc/sysconfig/network/ifcfg-ib0** file:  
**vi /etc/sysconfig/network/ifcfg-ib0**
4. Enter the following information:  
**DEVICE=ib0**  
**BOOTPROTO=static**  
**IPADDR=IP address to be configured for the InfiniBand NIC**  
**NETMASK=Subnet mask**  
**STARTMODE=auto**

### NOTE

For instructions about how to obtain the IP address and subnet mask for an InfiniBand NIC, see [Background](#).

5. Run the following command to restart the network for the configuration to take effect:  
**service network restart**

## 14.12.8 How Can I Handle the Issue that a Windows 7 ECS Equipped with an Intel 82599 NIC Reports an Error in SR-IOV Scenarios?

### Symptom

When the 20.4.1 driver package downloaded at Intel website <https://downloadcenter.intel.com/search?keyword=Intel++Ethernet+Connections+CD>



was installed in a Windows 7 64bit ECS with SR-IOV passthrough enabled, the system displayed the message "No Intel adapter found".

## Cause Analysis

The OS identifies an Intel 82599 passthrough NIC without a driver installed as an Ethernet controller. When the 20.4.1 driver package was installed, the OS did not identify the Intel NIC, leading to the error.

## Solution

Run **Autorun.exe** in the folder where the 20.4.1 driver package is stored. Install a driver on the NIC before installing the driver package so that the NIC can be identified as an Intel 82599 virtual function (VF) device by the OS. Use either of the following methods to install the driver:

- Method 1: Update the version.
  - a. Download the 18.6 driver package at the Intel website.
  - b. Run **Autorun.exe**.
  - c. Run **Autorun.exe** in the folder where the 20.4.1 driver package is stored to update the driver.
- Method 2: Use the device manager.
  - a. Start the Windows resource manager. Right-click **Computer** and choose **Manage** from the shortcut menu. In the **Device Manager** window, locate the NIC. When the NIC has no driver installed, the NIC locates in **Other devices** and is named **Ethernet Controller**.
  - b. Right-click **Ethernet Controller** and choose **Update Driver Software**.
  - c. Click **Browse**, select the path where the driver package is stored, and click **Next**.
  - d. Locate the NIC in **Network Adapter** of **Device Manager**.
  - e. Run **Autorun.exe** to install the 20.4.1 driver package.

## 14.12.9 Why Can't My Linux ECS Obtain Metadata?

### Symptom

The security group of the Linux ECS has been configured based on the prerequisites in [Obtaining Metadata](#) in the outbound direction, but the ECS still cannot obtain the metadata through the route with the destination of 169.254.169.254.

### Root Cause

Run the following command on the Linux ECS configured with a static IP address:

```
ip route| grep 169.254
```

The route with the destination of 169.254.169.254 does not exist, but the route with the destination of 169.254.0.0/16 exists.

**Figure 14-142** Route information

```
1# ip route | grep 169.254
169.254.0.0/16 dev eth0 scope link
1#
```

After the network is restarted, the original route with the destination of 169.254.169.254 is changed to the route with the destination of 169.254.0.0/16 without a next hop, as shown in [Figure 14-142](#). As a result, the Linux ECS cannot obtain metadata.

## Solution

1. Add the route with the destination of 169.254.169.254, and specify the next hop (gateway) and the output device (primary NIC of the Linux ECS). The following is an example:  
**# ip route add 169.254.169.254 via 192.168.1.1 dev eth0**  
192.168.1.1 is the gateway address of the subnet that the primary NIC resides, and eth0 is the primary NIC.
2. Run the following command to verify that the metadata can be obtained:  
**# curl http://169.254.169.254**

**Figure 14-143** Obtaining metadata

```
ecs-test [~] # ip route add 169.254.169.254 via 192.168.1.1 dev eth0
ecs-test [~] # curl http://169.254.169.254
1.0
2007-01-19
2007-03-01
2007-08-29
2007-10-10
2007-12-15
2008-02-01
2008-09-01
2009-04-04
latestecs-test [~] #
```

3. Run the following command to create or modify the `/etc/sysconfig/network-scripts/route-eth0` file to prevent the static route from being changed after network restart:  
**# vi /etc/sysconfig/network-scripts/route-eth0**  
Add the following content to the file:  
In this example, the primary NIC is eth0 and gateway address is 192.168.1.1. Replace them based on site requirements.  
**# 169.254.169.254 via 192.168.1.1**

## 14.12.10 How Can I Obtain the MAC Address of My ECS?

This section describes how to obtain the MAC address of an ECS.

### NOTE

The MAC address of an ECS cannot be changed.

## Linux (CentOS 6)

1. Log in to the Linux ECS.
2. Run the following command to view the MAC address of the ECS:

**ifconfig**

**Figure 14-144** Obtaining the MAC address

```
[root@CentOS68-XEN ~]# ifconfig
eth0 Link encap:Ethernet HWaddr FA:16:3E:2A:36:DE
 inet addr:192.168.22.227 Bcast:192.168.22.255 Mask:255.255.255.0
 inet6 addr: fe80::f816:3eff:fe2a:36de/64 Scope:Link
 UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
 RX packets:4699 errors:0 dropped:0 overruns:0 frame:0
 TX packets:2213 errors:0 dropped:0 overruns:0 carrier:0
 collisions:0 txqueuelen:1000
 RX bytes:472826 (461.7 KiB) TX bytes:438396 (428.1 KiB)

lo Link encap:Local Loopback
 inet addr:127.0.0.1 Mask:255.0.0.0
 inet6 addr: ::1/128 Scope:Host
 UP LOOPBACK RUNNING MTU:65536 Metric:1
 RX packets:1 errors:0 dropped:0 overruns:0 frame:0
 TX packets:1 errors:0 dropped:0 overruns:0 carrier:0
 collisions:0 txqueuelen:0
 RX bytes:28 (28.0 b) TX bytes:28 (28.0 b)
```

## Linux (CentOS 7)

1. Log in to the Linux ECS.
2. Run the following command to view the MAC address of the ECS:

**ifconfig**

**Figure 14-145** Obtaining the NIC information

```
[root@ecs-683a ~]# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
 inet 192.168.0.65 netmask 255.255.255.0 broadcast 192.168.0.255
 inet6 fe80::f816:3eff:fec3:46fc prefixlen 64 scopeid 0x20<link>
 ether fa:16:3e:c3:46:fc txqueuelen 1000 (Ethernet)
 RX packets 14457 bytes 20617950 (19.6 MiB)
 RX errors 0 dropped 0 overruns 0 frame 0
 TX packets 1867 bytes 245185 (239.4 KiB)
 TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
 inet netmask 255.0.0.0
 inet6 ::1 prefixlen 128 scopeid 0x10<host>
 loop txqueuelen 1000 (Local Loopback)
 RX packets 0 bytes 0 (0.0 B)
 RX errors 0 dropped 0 overruns 0 frame 0
 TX packets 0 bytes 0 (0.0 B)
 TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

3. Run the following command to view the MAC address of NIC **eth0**:  
**ifconfig eth0 |grep "ether"**

**Figure 14-146** Obtaining the MAC address of eth0

```
[root@ecs-683a ~]# ifconfig eth0 |egrep "ether"
ether fa:16:3e:c3:46:fc txqueuelen 1000 (Ethernet)
[root@ecs-683a ~]#
```

4. Obtain the returned MAC address.

```
ifconfig eth0 |egrep "ether" |awk '{print $2}'
```

**Figure 14-147** Obtaining the MAC address of eth0

```
[root@ecs-683a ~]# ifconfig eth0 |egrep "ether" |awk '{print $2}'
fa:16:3e:c3:46:fc
[root@ecs-683a ~]#
```

## Windows

1. Press **Win+R** to start the **Run** text box.
2. Enter **cmd** and click **OK**.
3. Run the following command to view the MAC address of the ECS:

```
ipconfig /all
```

```
Ethernet adapter Ethernet 2:
Connection-specific DNS Suffix . :
Description :
Physical Address. :
DHCP Enabled. :
Autoconfiguration Enabled :
Link-local IPv6 Address :
IPv4 Address. :
Subnet Mask :
Lease Obtained. :
Lease Expires :
Default Gateway :
DHCP Server :
DHCPv6 IAID :
DHCPv6 Client DUID. :
DNS Servers :
NetBIOS over Tcpi. :
```

## 14.12.11 How Can I Test Network Performance?

Use netperf and iperf3 to test network performance between ECSs. The test operations include preparations, TCP bandwidth test, UDP PPS test, and latency test.

### Background

- Tested ECS: an ECS that is tested for network performance. Such an ECS functions as the client (TX end) or server (RX end) in netperf tests.
- Auxiliary ECS: an ECS that is used to exchange test data with the tested ECS. The auxiliary ECS functions as the client (TX end) or server (RX end) in netperf tests.
- [Table 14-13](#) and [Table 14-14](#) list the common netperf and iperf3 parameters.

**Table 14-13** Common netperf parameters

| Parameter | Description                                                                                      |
|-----------|--------------------------------------------------------------------------------------------------|
| -p        | Port number                                                                                      |
| -H        | IP address of the RX end                                                                         |
| -t        | Protocol used in packet transmitting, the value of which is <b>TCP_STREAM</b> in bandwidth tests |
| -l        | Test duration                                                                                    |
| -m        | Data packet size, which is suggested to be <b>1440</b> in bandwidth tests                        |

**Table 14-14** Common iperf3 parameters

| Parameter | Description                                                                                                                                                                         |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -p        | Port number                                                                                                                                                                         |
| -c        | IP address of the RX end                                                                                                                                                            |
| -u        | UDP packets                                                                                                                                                                         |
| -b        | TX bandwidth                                                                                                                                                                        |
| -t        | Test duration                                                                                                                                                                       |
| -l        | Data packet size, which is suggested to be <b>16</b> in PPS tests                                                                                                                   |
| -A        | ID of the vCPU used by iperf3<br>In this section, the maximum number of 16 vCPUs is used as an example for each ECS. If an ECS has 8 vCPUs, the <b>-A</b> value ranges from 0 to 7. |

## Test Preparations

### Step 1 Prepare ECSs.

Ensure that both type and specifications of the tested ECS and auxiliary ECSs are the same. In addition, ensure that these ECSs are deployed in the same ECS group with anti-affinity enabled.

**Table 14-15** Preparations

| Category   | Quantity | Image                          | Specifications       | IP Address   |
|------------|----------|--------------------------------|----------------------|--------------|
| Tested ECS | 1        | CentOS 7.4 64bit (recommended) | At least eight vCPUs | 192.168.2.10 |

| Category      | Quantity | Image                          | Specifications   | IP Address                    |
|---------------|----------|--------------------------------|------------------|-------------------------------|
| Auxiliary ECS | 8        | CentOS 7.4 64bit (recommended) | At least 8 vCPUs | 192.168.2.11-19<br>2.168.2.18 |

**Step 2** Install the netperf, iperf3, and sar test tools on both the tested ECS and auxiliary ECSs.

**Table 14-16** lists the procedures for installing these tools.

**Table 14-16** Installing test tools

| Tool    | Procedure                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| netperf | <ol style="list-style-type: none"><li>1. Run the following command to install gcc:<br/><b>yum -y install unzip gcc gcc-c++</b></li><li>2. Run the following command to download the netperf installation package:<br/><b>wget https://github.com/HewlettPackard/netperf/archive/refs/tags/netperf-2.7.0.zip</b></li><li>3. Run the following commands to decompress the installation package and install netperf:<br/><b>unzip netperf-2.7.0.zip</b><br/><b>cd netperf-netperf-2.7.0/</b><br/><b>./configure &amp;&amp; make &amp;&amp; make install</b></li></ol> |
| iperf3  | <ol style="list-style-type: none"><li>1. Run the following command to download the iperf3 installation package:<br/><b>wget --no-check-certificate https://codeload.github.com/esnet/iperf/zip/master -O iperf3.zip</b></li><li>2. Run the following commands to decompress the installation package and install iperf3:<br/><b>unzip iperf3.zip</b><br/><b>cd iperf-master/</b><br/><b>./configure &amp;&amp; make &amp;&amp; make install</b></li></ol>                                                                                                          |
| sar     | Run the following command to install sar:<br><b>yum -y install sysstat</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

**Step 3** Enable NIC multi-queue.

Perform the following operations on both tested ECS and auxiliary ECSs.

1. Run the following command to check the number of queues supported by the ECSs:  
**ethtool -l eth0 | grep -i Pre -A 5 | grep Combined**
2. Run the following command to enable NIC multi-queue:  
**ethtool -L eth0 combined X**

In the preceding command, *X* is the number of queues obtained in [Step 3.1](#).

----End

## TCP Bandwidth Test (Using netperf)

Perform the test on multiple flows. This section considers 16 flows that are evenly distributed to eight ECSs, as an example.

### NOTE

The TCP bandwidth test uses the multi-flow model.

- When testing the TCP transmission (TX) bandwidth, use the one-to-many model to ensure that the capability of the receiver is sufficient.
- When testing the TCP receiver (RX) bandwidth, use the many-to-one model to ensure that the capability of the sender is sufficient.

### Step 1 Test the TCP TX bandwidth.

1. Run the following commands on all auxiliary ECSs to start the netserver process:

```
netserver -p 12001
```

```
netserver -p 12002
```

In the preceding commands, **-p** specifies the listening port.

2. Start the netperf process on the tested ECS and specify a netserver port for each auxiliary ECS. For details about common netperf parameters, see [Table 14-13](#).

##The IP address is for the first auxiliary ECS.

```
netperf -H 192.168.2.11 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
```

```
netperf -H 192.168.2.11 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

##The IP address is for the second auxiliary ECS.

```
netperf -H 192.168.2.12 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
```

```
netperf -H 192.168.2.12 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

##The IP address is for the third auxiliary ECS.

```
netperf -H 192.168.2.13 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
```

```
netperf -H 192.168.2.13 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

##The IP address is for the fourth auxiliary ECS.

```
netperf -H 192.168.2.14 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
```

```
netperf -H 192.168.2.14 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

##The IP address is for the fifth auxiliary ECS.

```
netperf -H 192.168.2.15 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
```

```
netperf -H 192.168.2.15 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

##The IP address is for the sixth auxiliary ECS.

```
netperf -H 192.168.2.16 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
```

```
netperf -H 192.168.2.16 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

##The IP address is for the seventh auxiliary ECS.

```
netperf -H 192.168.2.17 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
netperf -H 192.168.2.17 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

##The IP address is for the eighth auxiliary ECS.

```
netperf -H 192.168.2.18 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
netperf -H 192.168.2.18 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

**Step 2** Test the TCP RX bandwidth.

1. Start the netserver process on the tested ECS.

##The port number is for the first auxiliary ECS.

```
netserver -p 12001
```

```
netserver -p 12002
```

##The port number is for the second auxiliary ECS.

```
netserver -p 12003
```

```
netserver -p 12004
```

##The port number is for the third auxiliary ECS.

```
netserver -p 12005
```

```
netserver -p 12006
```

##The port number is for the fourth auxiliary ECS.

```
netserver -p 12007
```

```
netserver -p 12008
```

##The port number is for the fifth auxiliary ECS.

```
netserver -p 12009
```

```
netserver -p 12010
```

##The port number is for the sixth auxiliary ECS.

```
netserver -p 12011
```

```
netserver -p 12012
```

##The port number is for the seventh auxiliary ECS.

```
netserver -p 12013
```

```
netserver -p 12014
```

##The port number is for the eighth auxiliary ECS.

```
netserver -p 12015
```

```
netserver -p 12016
```

2. Start the netperf process on all auxiliary ECSs.

Log in to auxiliary ECS 1.

```
netperf -H 192.168.2.10 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
netperf -H 192.168.2.10 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

Log in to auxiliary ECS 2.

```
netperf -H 192.168.2.10 -p 12003 -t TCP_STREAM -l 300 -- -m 1440 &
netperf -H 192.168.2.10 -p 12004 -t TCP_STREAM -l 300 -- -m 1440 &
```



Log in to auxiliary ECS 3.

```
netperf -H 192.168.2.10 -p 12005 -t TCP_STREAM -l 300 -- -m 1440 &
netperf -H 192.168.2.10 -p 12006 -t TCP_STREAM -l 300 -- -m 1440 &
```

Log in to auxiliary ECS 4.

```
netperf -H 192.168.2.10 -p 12007 -t TCP_STREAM -l 300 -- -m 1440 &
netperf -H 192.168.2.10 -p 12008 -t TCP_STREAM -l 300 -- -m 1440 &
```

Log in to auxiliary ECS 5.

```
netperf -H 192.168.2.10 -p 12009 -t TCP_STREAM -l 300 -- -m 1440 &
netperf -H 192.168.2.10 -p 12010 -t TCP_STREAM -l 300 -- -m 1440 &
```

Log in to auxiliary ECS 6.

```
netperf -H 192.168.2.10 -p 12011 -t TCP_STREAM -l 300 -- -m 1440 &
netperf -H 192.168.2.10 -p 12012 -t TCP_STREAM -l 300 -- -m 1440 &
```

Log in to auxiliary ECS 7.

```
netperf -H 192.168.2.10 -p 12013 -t TCP_STREAM -l 300 -- -m 1440 &
netperf -H 192.168.2.10 -p 12014 -t TCP_STREAM -l 300 -- -m 1440 &
```

Log in to auxiliary ECS 8.

```
netperf -H 192.168.2.10 -p 12015 -t TCP_STREAM -l 300 -- -m 1440 &
netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 -- -m 1440 &
```

### Step 3 Analyze the test result.

After the test is complete, the output of the netperf process on one TX end is shown in [Figure 14-148](#). The final result is the sum of the test results of the netperf processes on all TX ends.

**Figure 14-148** Output of the netperf process on one TX end

```
Recv Send Send
Socket Socket Message Elapsed
Size Size Size Time Throughput
bytes bytes bytes secs. 10^6bits/sec
TX buffer Test duration Throughput
87380 16384 1440 120.02 956.30
RX buffer Data packet size
```

#### NOTE

There are a large number of netperf processes. To facilitate statistics collection, it is a good practice to run the following command to view test data on the tested ECS using sar:

```
sar -n DEV 1 60
```

----End

## UDP PPS Test (Using iperf3)

### Step 1 Test the UDP TX PPS.

1. Log in to an auxiliary ECS.
2. Run the following commands on all auxiliary ECSs to start the server process:

```
iperf3 -s -p 12001 &
iperf3 -s -p 12002 &
```

In the preceding commands, **-p** specifies the listening port.

3. Start the client process on the tested ECS. For details about common iperf3 parameters, see [Table 14-14](#).

```
##Auxiliary ECS 1
```

```
iperf3 -c 192.168.2.11 -p 12001 -u -b 100M -t 300 -l 16 -A 0 &
iperf3 -c 192.168.2.11 -p 12002 -u -b 100M -t 300 -l 16 -A 1 &
```

```
##Auxiliary ECS 2
```

```
iperf3 -c 192.168.2.12 -p 12001 -u -b 100M -t 300 -l 16 -A 2 &
iperf3 -c 192.168.2.12 -p 12002 -u -b 100M -t 300 -l 16 -A 3 &
```

```
##Auxiliary ECS 3
```

```
iperf3 -c 192.168.2.13 -p 12001 -u -b 100M -t 300 -l 16 -A 4 &
iperf3 -c 192.168.2.13 -p 12002 -u -b 100M -t 300 -l 16 -A 5 &
```

```
##Auxiliary ECS 4
```

```
iperf3 -c 192.168.2.14 -p 12001 -u -b 100M -t 300 -l 16 -A 6 &
iperf3 -c 192.168.2.14 -p 12002 -u -b 100M -t 300 -l 16 -A 7 &
```

```
##Auxiliary ECS 5
```

```
iperf3 -c 192.168.2.15 -p 12001 -u -b 100M -t 300 -l 16 -A 8 &
iperf3 -c 192.168.2.15 -p 12002 -u -b 100M -t 300 -l 16 -A 9 &
```

```
##Auxiliary ECS 6
```

```
iperf3 -c 192.168.2.16 -p 12001 -u -b 100M -t 300 -l 16 -A 10 &
iperf3 -c 192.168.2.16 -p 12002 -u -b 100M -t 300 -l 16 -A 11 &
```

```
##Auxiliary ECS 7
```

```
iperf3 -c 192.168.2.17 -p 12001 -u -b 100M -t 300 -l 16 -A 12 &
iperf3 -c 192.168.2.17 -p 12002 -u -b 100M -t 300 -l 16 -A 13 &
```

```
##Auxiliary ECS 8
```

```
iperf3 -c 192.168.2.18 -p 12001 -u -b 100M -t 300 -l 16 -A 14 &
iperf3 -c 192.168.2.18 -p 12002 -u -b 100M -t 300 -l 16 -A 15 &
```

**Step 2** Test the UDP RX PPS.

1. Start the server process on the tested ECS. For details about common iperf3 parameters, see [Table 14-14](#).

##The port number is for the first auxiliary ECS.

```
iperf3 -s -p 12001 -A 0 -i 60 &
```

```
iperf3 -s -p 12002 -A 1 -i 60 &
```

##The port number is for the second auxiliary ECS.

```
iperf3 -s -p 12003 -A 2 -i 60 &
```

```
iperf3 -s -p 12004 -A 3 -i 60 &
```

##The port number is for the third auxiliary ECS.

```
iperf3 -s -p 12005 -A 4 -i 60 &
```

```
iperf3 -s -p 12006 -A 5 -i 60 &
```

##The port number is for the fourth auxiliary ECS.

```
iperf3 -s -p 12007 -A 6 -i 60 &
```

```
iperf3 -s -p 12008 -A 7 -i 60 &
```

##The port number is for the fifth auxiliary ECS.

```
iperf3 -s -p 12009 -A 8 -i 60 &
```

```
iperf3 -s -p 12010 -A 9 -i 60 &
```

##The port number is for the sixth auxiliary ECS.

```
iperf3 -s -p 12011 -A 10 -i 60 &
```

```
iperf3 -s -p 12012 -A 11 -i 60 &
```

##The port number is for the seventh auxiliary ECS.

```
iperf3 -s -p 12013 -A 12 -i 60 &
```

```
iperf3 -s -p 12014 -A 13 -i 60 &
```

##The port number is for the eighth auxiliary ECS.

```
iperf3 -s -p 12015 -A 14 -i 60 &
```

```
iperf3 -s -p 12016 -A 15 -i 60 &
```

2. Start the client process on all auxiliary ECSs. For details about common iperf3 parameters, see [Table 14-14](#).

Log in to auxiliary ECS 1.

```
iperf3 -c 192.168.2.10 -p 12001 -u -b 100M -t 300 -l 16 -A 0 &
```

```
iperf3 -c 192.168.2.10 -p 12002 -u -b 100M -t 300 -l 16 -A 1 &
```

Log in to auxiliary ECS 2.

```
iperf3 -c 192.168.2.10 -p 12003 -u -b 100M -t 300 -l 16 -A 0 &
```

```
iperf3 -c 192.168.2.10 -p 12004 -u -b 100M -t 300 -l 16 -A 1 &
```

Log in to auxiliary ECS 3.

```
iperf3 -c 192.168.2.10 -p 12005 -u -b 100M -t 300 -l 16 -A 0 &
```

```
iperf3 -c 192.168.2.10 -p 12006 -u -b 100M -t 300 -l 16 -A 1 &
```

Log in to auxiliary ECS 4.

```


iperf3 -c 192.168.2.10 -p 12007 -u -b 100M -t 300 -l 16 -A 0 &
iperf3 -c 192.168.2.10 -p 12008 -u -b 100M -t 300 -l 16 -A 1 &
Log in to auxiliary ECS 5.
iperf3 -c 192.168.2.10 -p 12009 -u -b 100M -t 300 -l 16 -A 0 &
iperf3 -c 192.168.2.10 -p 12010 -u -b 100M -t 300 -l 16 -A 1 &
Log in to auxiliary ECS 6.
iperf3 -c 192.168.2.10 -p 12011 -u -b 100M -t 300 -l 16 -A 0 &
iperf3 -c 192.168.2.10 -p 12012 -u -b 100M -t 300 -l 16 -A 1 &
Log in to auxiliary ECS 7.
iperf3 -c 192.168.2.10 -p 12013 -u -b 100M -t 300 -l 16 -A 0 &
iperf3 -c 192.168.2.10 -p 12014 -u -b 100M -t 300 -l 16 -A 1 &
Log in to auxiliary ECS 8.
iperf3 -c 192.168.2.10 -p 12015 -u -b 100M -t 300 -l 16 -A 0 &
iperf3 -c 192.168.2.10 -p 12016 -u -b 100M -t 300 -l 16 -A 1 &

```

**Step 3** Analyze the test result.

**Figure 14-149** shows an example of the UDP PPS test result.

**Figure 14-149** UDP PPS test result



#### NOTE

There are a large number of iperf3 processes. To facilitate statistics collection, it is a good practice to run the following command to view test data on the tested ECS using sar:

```
sar -n DEV 1 60
```

----End

## Latency Test

**Step 1** Run the following command to start the qperf process on the tested ECS:

```
qperf &
```

**Step 2** Log in to auxiliary ECS 1 and run the following command to perform a latency test:

```
qperf 192.168.2.10 -m 64 -t 60 -vu udp_lat
```

After the test is complete, the **lat** value in the command output is the latency between ECSs.

----End

## 14.12.12 Why Can't I Use DHCP to Obtain a Private IP Address?

### Symptom

You attempt to use DHCP to obtain a private IP address, but you cannot obtain the IP address.

- For Linux, a private IP address cannot be assigned.
- For Windows, a private IP address is changed to an IP address in the 169.254 network segment, which is different from the private IP address displayed on the ECS console.

#### NOTE

You are advised to use a public image to create an ECS. All public images support DHCP continuous discovery mode.

### Solution (Linux)

The following uses CentOS 7.2 as an example. For solutions about other OSs, see the corresponding help documentation.


1. Log in to the ECS and run the following command:  
**ps -ef | grep dhclient**
2. If the dhclient process does not exist, restart the NIC or run any of the following commands to initiate a DHCP request:  
**dhclient eth0, ifdown eth0 + ifup eth0, or dhcpcd eth0**
3. If the DHCP client does not send any requests for a long time, for example, the issue recurs after the NIC is restarted, do the following:
  - a. Run the following command to configure a static IP:  
**vi /etc/sysconfig/network-scripts/ifcfg-eth0**

```
BOOTPROTO=static
IPADDR=192.168.1.100 #IP address (modified)
NETMASK=255.255.255.0 #Mask (modified)
GATEWAY=192.168.1.1 #Gateway IP address (modified)
```
  - b. Restart the ECS to make the network settings take effect.
  - c. Select an image in which DHCP runs stably.
4. If the fault persists, obtain the messages in **/var/log/messages** on the affected ECS, use the MAC address of the affected NIC to filter the desired log, and check whether there is any process that prevents DHCP from obtaining an IP address.
5. If the fault persists, contact technical support.

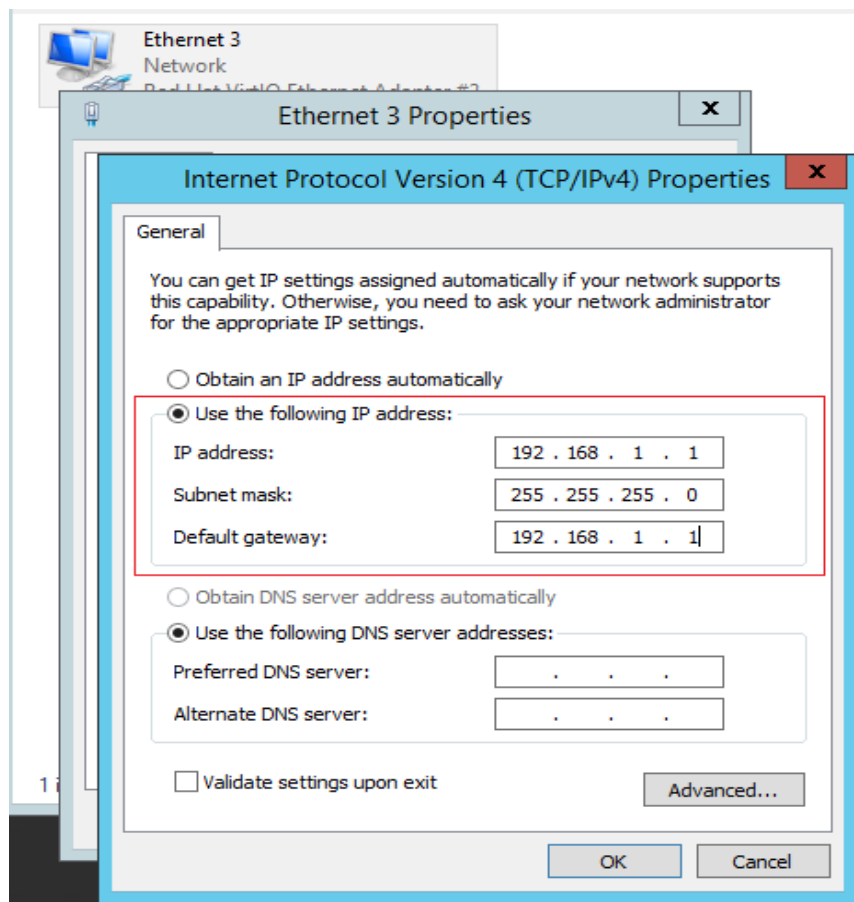
### Solution (Windows)

The following uses Windows 2012 as an example. For solutions about other OSs, see the corresponding help documentation.

1. Right-click a local area connection and choose **Disable** from the shortcut menu. Then, choose **Enable**.



2. If the DHCP client does not send any requests for a long time, for example, the issue recurs after the NIC is restarted, do the following:
  - a. Right-click **Local Area Connection** and choose **Properties** from the shortcut menu.
  - b. In the displayed dialog box, select **Internet Protocol Version 4 (TCP/IPv4)**, click **Properties**, and modify parameter settings.



- c. Restart the ECS to make the network settings take effect.

3. If the fault persists, contact technical support.

## 14.12.13 How Can I View and Modify Kernel Parameters of a Linux ECS?

Modify the kernel parameters only if the parameter settings affect your services. Kernel parameters vary depending on OS versions. If the parameter settings must be modified,

- Ensure that the target parameter settings meet service requirements.
- Modify the correct kernel parameters. For details about common kernel parameters, see [Table 14-17](#).
- Back up key ECS data before modifying kernel parameter settings.

### Background

**Table 14-17** Common Linux kernel parameters

| Parameter                   | Description                                                                                                                                                                                        |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| net.core.rmem_default       | Specifies the default size (in bytes) of the window for receiving TCP data.                                                                                                                        |
| net.core.rmem_max           | Specifies the maximum size (in bytes) of the window for receiving TCP data.                                                                                                                        |
| net.core.wmem_default       | Specifies the default size (in bytes) of the window for transmitting TCP data.                                                                                                                     |
| net.core.wmem_max           | Specifies the maximum size (in bytes) of the window for transmitting TCP data.                                                                                                                     |
| net.core.netdev_max_backlog | Specifies the maximum number of packets that can be sent to a queue when the rate at which each network port receives packets is faster than the rate at which the kernel processes these packets. |
| net.core.somaxconn          | Defines the maximum length of the listening queue for each port in the system. This parameter applies globally.                                                                                    |
| net.core.optmem_max         | Specifies the maximum size of the buffer allowed by each socket.                                                                                                                                   |

| Parameter                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| net.ipv4.tcp_mem              | <p>Uses the TCP stack to show memory usage in memory pages (4 KB generally).</p> <p>The first value is the lower limit of memory usage.</p> <p>The second value is the upper limit of the load added to the buffer when the memory is overloaded.</p> <p>The third value is the upper limit of memory usage. When this value is reached, packets can be discarded to reduce memory usage. For a large BDP, increase the parameter value as needed. The unit of this parameter is memory page but not byte.</p>         |
| net.ipv4.tcp_rmem             | <p>Specifies the memory used by sockets for automatic optimization.</p> <p>The first value is the minimum number of bytes allocated to the socket buffer for receiving data.</p> <p>The second value is the default value, which is overwritten by <b>rmem_default</b>. The buffer size can increase to this value when the system load is not heavy.</p> <p>The third value is the maximum number of bytes allocated to the socket buffer for receiving data. This value is overwritten by <b>rmem_max</b>.</p>       |
| net.ipv4.tcp_wmem             | <p>Specifies the memory used by sockets for automatic optimization.</p> <p>The first value is the minimum number of bytes allocated to the socket buffer for transmitting data.</p> <p>The second value is the default value, which is overwritten by <b>wmem_default</b>. The buffer size can increase to this value when the system load is not heavy.</p> <p>The third value is the maximum number of bytes allocated to the socket buffer for transmitting data. This value is overwritten by <b>wmem_max</b>.</p> |
| net.ipv4.tcp_keepalive_time   | <p>Specifies the interval at which keepalive detection messages are sent in seconds for checking TCP connections.</p>                                                                                                                                                                                                                                                                                                                                                                                                  |
| net.ipv4.tcp_keepalive_intvl  | <p>Specifies the interval at which keepalive detection messages are resent in seconds when no response is received.</p>                                                                                                                                                                                                                                                                                                                                                                                                |
| net.ipv4.tcp_keepalive_probes | <p>Specifies the maximum number of keepalive detection messages that are sent to determine a TCP connection failure.</p>                                                                                                                                                                                                                                                                                                                                                                                               |



| Parameter                   | Description                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| net.ipv4.tcp_sack           | Enables selective acknowledgment (value <b>1</b> indicates enabled). This configuration allows the transmitter to resend only lost packets, thereby improving system performance. However, this configuration will increase the CPU usage. You are suggested to enable selective acknowledgment for WAN communication.                                         |
| net.ipv4.tcp_fack           | Enables forwarding acknowledgment for selective acknowledgment (SACK), thereby reducing congestion. You are suggested to enable forwarding acknowledgment.                                                                                                                                                                                                     |
| net.ipv4.tcp_timestamps     | Specifies a TCP timestamp, which will add 12 bytes in the TCP packet header. This configuration calculates RTT using RFC1323, a more precise retransmission method upon timeout than retransmission. You are suggested to enable this parameter for higher system performance.                                                                                 |
| net.ipv4.tcp_window_scaling | Enables RFC1323-based window scaling by setting the parameter value to <b>1</b> if the TCP window is larger than 64 KB. The maximum TCP window is 1 GB. This parameter takes effect only when window scaling is enabled on both ends of the TCP connection.                                                                                                    |
| net.ipv4.tcp_syncookies     | Specifies whether to enable TCP synchronization ( <b>syncookie</b> ). This configuration prevents socket overloading when a large number of connections are attempted to set up. <b>CONFIG_SYN_COOKIES</b> must be enabled in the kernel for compilation. The default value is <b>0</b> , indicating that TCP synchronization is disabled.                     |
| net.ipv4.tcp_tw_reuse       | Specifies whether a <b>TIME-WAIT</b> socket ( <b>TIME-WAIT</b> port) can be used for new TCP connections.<br><b>NOTE</b><br>This parameter is valid only for clients and takes effect only when <b>net.ipv4.tcp_timestamps</b> is enabled.<br>This parameter cannot be set to <b>1</b> if NAT is enabled. Otherwise, an error will occur in remote ECS logins. |
| net.ipv4.tcp_tw_recycle     | Allows fast recycle of <b>TIME-WAIT</b> sockets.<br><b>NOTE</b><br>This parameter is valid only when <b>net.ipv4.tcp_timestamps</b> is enabled.<br>Do not set this parameter to <b>1</b> if NAT is enabled. Otherwise, an error will occur during remote ECS logins.                                                                                           |

| Parameter                      | Description                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| net.ipv4.tcp_fin_timeout       | Specifies the time (in seconds) during which a socket TCP connection that is disconnected from the local end remains in the <b>FIN-WAIT-2</b> state. Process suspension may be caused by the disconnection from the peer end, continuous connection from the peer end, or other reasons.                                                                                                           |
| net.ipv4.ip_local_port_range   | Specifies local port numbers allowed by TCP/UDP.                                                                                                                                                                                                                                                                                                                                                   |
| net.ipv4.tcp_max_syn_backlog   | Specifies the maximum number of connection requests that are not acknowledged by the peer end and that can be stored in the queue. The default value is <b>1024</b> . If the server is frequently overloaded, try to increase the value.                                                                                                                                                           |
| net.ipv4.tcp_low_latency       | This option should be disabled if the TCP/IP stack is used for high throughput, low latency.                                                                                                                                                                                                                                                                                                       |
| net.ipv4.tcp_westwood          | Enables the congestion control algorithm on the transmitter end to evaluate throughput and improve the overall bandwidth utilization. You are suggested to enable the congestion control algorithm for WAN communication.                                                                                                                                                                          |
| net.ipv4.tcp_bic               | Enables binary increase congestion for fast long-distance networks so that the connections with operations being performed at a rate of Gbit/s can be functional. You are suggested to enable binary increase congestion for WAN communication.                                                                                                                                                    |
| net.ipv4.tcp_max_tw_buckets    | Specifies the number of TIME_WAIT buckets, which defaults to <b>180000</b> . If the number of buckets exceeds the default value, extra ones will be cleared.                                                                                                                                                                                                                                       |
| net.ipv4.tcp_synack_retries    | Specifies the number of times that SYN+ACK packets are retransmitted in <b>SYN_RECV</b> state.                                                                                                                                                                                                                                                                                                     |
| net.ipv4.tcp_abort_on_overflow | When this parameter is set to <b>1</b> , if the system receives a large number of requests within a short period of time but fails to process them, the system will send reset packets to terminate the connections. It is recommended that you improve system processing capabilities by optimizing the application efficiency instead of performing reset operations.<br>Default value: <b>0</b> |
| net.ipv4.route.max_size        | Specifies the maximum number of routes allowed by the kernel.                                                                                                                                                                                                                                                                                                                                      |
| net.ipv4.ip_forward            | Forward packets between interfaces.                                                                                                                                                                                                                                                                                                                                                                |

| Parameter                                          | Description                                                                  |
|----------------------------------------------------|------------------------------------------------------------------------------|
| net.ipv4.ip_default_ttl                            | Specifies the maximum number of hops that a packet can pass through.         |
| net.netfilter.nf_conntrack_tcp_timeout_established | Clears iptables connections that are inactive for a specific period of time. |
| net.netfilter.nf_conntrack_max                     | Specifies the maximum value of hash entries.                                 |

## Viewing Kernel Parameters

- Method 1: Run the cat command in **/proc/sys** to view file content.  
**/proc/sys/** is a pseudo directory generated after the Linux kernel is started. The **net** folder in this directory stores all kernel parameters that have taken effect in the system. The directory tree structure is determined based on complete parameter names. For example, **net.ipv4.tcp\_tw\_recycle** corresponds to the **/proc/sys/net/ipv4/tcp\_tw\_recycle** file, and the content of the file is the parameter value.

Example:

To view the **net.ipv4.tcp\_tw\_recycle** value, run the following command:

```
cat /proc/sys/net/ipv4/tcp_tw_recycle
```

- Method 2: Use the **/etc/sysctl.conf** file.  
Run the following command to view all parameters that have taken effect in the system:

```
/usr/sbin/sysctl -a
```

```
net.ipv4.tcp_syncookies = 1
net.ipv4.tcp_max_tw_buckets = 4096
net.ipv4.tcp_tw_reuse = 1
net.ipv4.tcp_tw_recycle = 1
net.ipv4.tcp_keepalive_time = 1800
net.ipv4.tcp_fin_timeout = 30
.....
net.ipv4.tcp_keepalive_time = 1200
net.ipv4.ip_local_port_range = 1024 65000
net.ipv4.tcp_max_syn_backlog = 8192
net.ipv4.tcp_rmem = 16384 174760 349520
net.ipv4.tcp_wmem = 16384 131072 262144
net.ipv4.tcp_mem = 262144 524288 1048576
.....
```

## Modifying Kernel Parameter Settings

- Method 1: Run the echo command in **/proc/sys** to modify the file for the target kernel parameters.  
The parameter values changed using this method take effect only during the current running and will be reset after the system is restarted. To make the modification take effect permanently, see method 2.

**/proc/sys/** is a pseudo directory generated after the Linux kernel is started. The **net** folder in this directory stores all kernel parameters that have taken effect in the system. The directory tree structure is determined based on complete parameter names. For example, **net.ipv4.tcp\_tw\_recycle**

corresponds to the `/proc/sys/net/ipv4/tcp_tw_recycle` file, and the content of the file is the parameter value.

Example:

To change the `net.ipv4.tcp_tw_recycle` value to `0`, run the following command:

```
echo "0" > /proc/sys/net/ipv4/tcp_tw_recycle
```

- Method 2: Use the `/etc/sysctl.conf` file.

The parameter values changed using this method take effect permanently.

- a. Run the following command to change the value of a specified parameter:

```
/sbin/sysctl -w kernel.domainname="example.com"
```

Example:

```
sysctl -w net.ipv4.tcp_tw_recycle="0"
```

- b. Run the following command to change the parameter value in the `/etc/sysctl.conf` file:

```
vi /etc/sysctl.conf
```

- c. Run the following command for the configuration to take effect:

```
/sbin/sysctl -p
```

## 14.12.14 How Can I Configure Port Redirection?

### Symptom

It is expected that the EIP and port on ECS 1 accessed from the Internet can be automatically redirected to the EIP and port on ECS 2.

### Windows

For example, to redirect port 8080 on ECS 1 bound with EIP 192.168.10.43 to port 18080 on ECS 2 bound with EIP 192.168.10.222, perform the following operations on ECS 1.

#### NOTE

Ensure that the desired ports have been enabled on the ECS security group and firewall.

1. Open the `cmd` window on the ECS and run the following command: The ECS running Windows Server 2012 is used as an example.

```
netsh interface portproxy add v4tov4 listenaddress=192.168.10.43
listenport=8080 connectaddress=192.168.10.222 connectport=18080
```

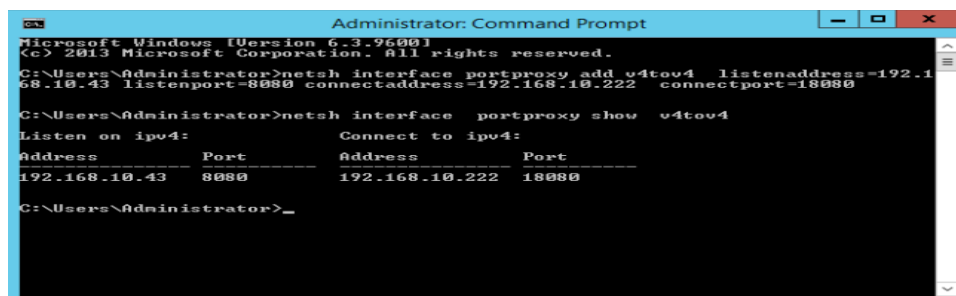
To cancel port redirection, run the following command:

```
netsh interface portproxy delete v4tov4 listenaddress=192.168.10.43
listenport=8080
```

2. Run the following command to view all port redirections configured on the ECS:

```
netsh interface portproxy show v4tov4
```

Figure 14-150 Port redirections on Windows



## Linux

For example, to redirect port 1080 on ECS 1 to port 22 on ECS 2 with the following configurations:

Private IP address and EIP of ECS 1: 192.168.72.10 and 123.xxx.xxx.456

Private IP address of ECS 2: 192.168.72.20

### NOTE

- Ensure that the desired ports have been enabled on the ECS security group and firewall.
- Ensure that the source/destination check function is disabled.

On the ECS details page, click **Network Interfaces** and disable **Source/Destination Check**.

By default, the source/destination check function is enabled. When this function is enabled, the system checks whether source IP addresses contained in the packets sent by ECSs are correct. If the IP addresses are incorrect, the system does not allow the ECSs to send the packets. This mechanism prevents packet spoofing, thereby improving system security. However, this mechanism prevents the packet sender from receiving returned packets. You need to disable the source/destination check.

**Step 1** Log in to Linux ECS 1.

1. Run the following command to modify the configuration file:

```
vi /etc/sysctl.conf
```

2. Add **net.ipv4.ip\_forward = 1** to the file.
3. Run the following command to complete the modification:

```
sysctl -p /etc/sysctl.conf
```

**Step 2** Run the following commands to add rules to the **nat** table in **iptables** so that the access to port 1080 on ECS 1 can be redirected to port 22 on ECS 2:

```
iptables -t nat -A PREROUTING -d 192.168.72.10 -p tcp --dport 1080 -j DNAT --to-destination 192.168.72.20:22
```

```
iptables -t nat -A POSTROUTING -d 192.168.72.20 -p tcp --dport 22 -j SNAT --to 192.168.72.10
```

**Step 3** Run the following command to log in to port 1080 on ECS 1 for check:

```
ssh -p 1080 123.xxx.xxx.456
```

**Figure 14-151** Port redirections on Linux

```
[root@linux-centos ~]# ssh -p 1080 [redacted]
root@[redacted]'s password: [redacted]
```

Enter the password to log in to ECS 2 with hostname **ecs-inner**.

**Figure 14-152** Logging in to ECS 2

```
[root@ecs-inner ~]#
```

----End

## 14.12.15 Can the ECSs of Different Accounts Communicate over an Intranet?

No. The ECSs of different accounts cannot communicate with each other over an intranet.

## 14.12.16 Will ECSs That I Purchased Deployed in the Same Subnet?

You can customize your network to deploy the ECSs. Therefore, whether they are in the same subnet is totally up to you.

## 14.13 Resource Monitoring

### 14.13.1 Why Is My Windows ECS Running Slowly?

If your ECS runs slowly or is inaccessible unexpectedly, the bandwidth or vCPU usage of the ECS may be excessively high. If you have created an alarm rule using Cloud Eye, the system automatically sends an alarm to you when the bandwidth or CPU usage reaches the threshold specified in the rule.

To handle this issue, perform the following operations:


1. Fault locating:  
Identify the drivers from unknown sources and processes leading to high bandwidth or CPU usage.  
Windows offer multiple tools to locate faults, including Task Manager, Performance Monitor, Resource Monitor, Process Explorer, Xperf (supported by versions later than Windows Server 2008), and full memory dump.
2. Check whether the processes and drivers are normal and handle the issue accordingly.
  - If the processes are not malicious, optimize their programs or modify ECS specifications.
  - If the processes are malicious, stop these processes manually or use a third-party tool to stop them automatically.

- If the drivers are from official sources, there is no need to deal with system built-in drivers. Determine whether to uninstall the third-party software based on your requirements.
- If the drivers are from unknown sources, you are advised to uninstall them by using commercial antivirus software or third-party security management tools.

## Fault Locating


1. Log in to the ECS using VNC available on the management console.
2. Start the **Run** dialog box, and then enter **perfmon -res**.

Figure 14-153 Starting the Resource Monitor



3. On the **Resource Monitor** page, click the **CPU** or **Network** tab to view the CPU or bandwidth usage.

Figure 14-154 Resource Monitor

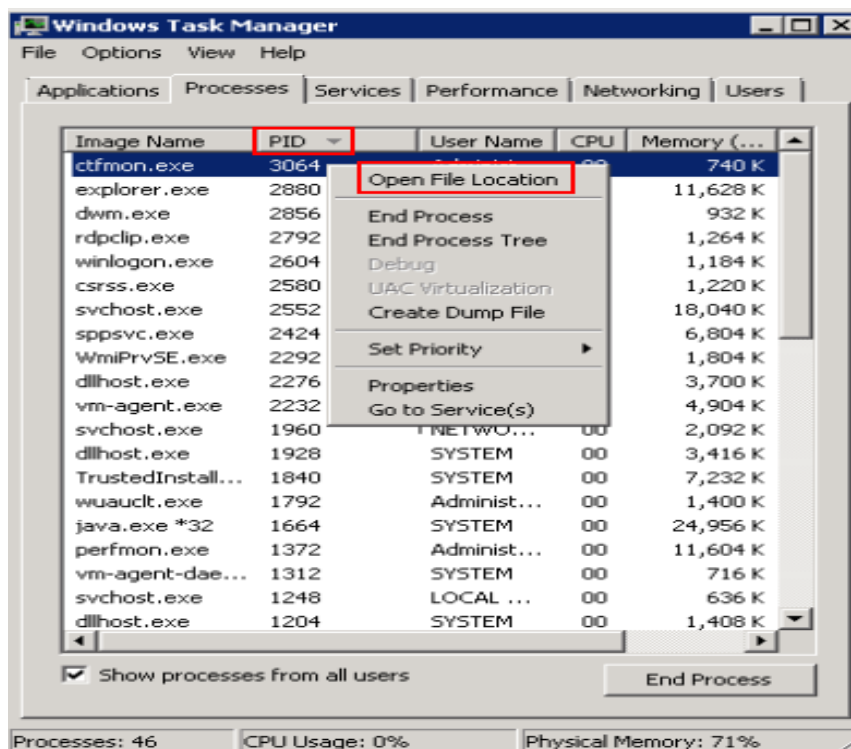


4. Obtain the IDs and names of the processes with high CPU or bandwidth usage.
5. On the remote login page, click **Ctrl+Alt+Del** to start the **Windows Task Manager**.

Alternatively, start the **Run** dialog box and enter **taskmgr** to start the **Windows Task Manager**.

The following describes how to display PIDs in **Windows Task Manager**, locate a process, and check whether it is malicious.

- a. Click the **Details** tab.
- b. Click **PID** to sort the data.
- c. Right-click the process with high CPU or bandwidth usage and choose **Open File Location** from the shortcut menu.
- d. Check whether the process is malicious.

**Figure 14-155** Checking the process

6. Open the **Run** dialog box and enter **fltmc** to view the filter drivers of the system.

The following figure uses Windows 10 as an example. Different OSs have different built-in drivers. For details, see their official websites. If a third-party driver is installed, it is also displayed in this figure.

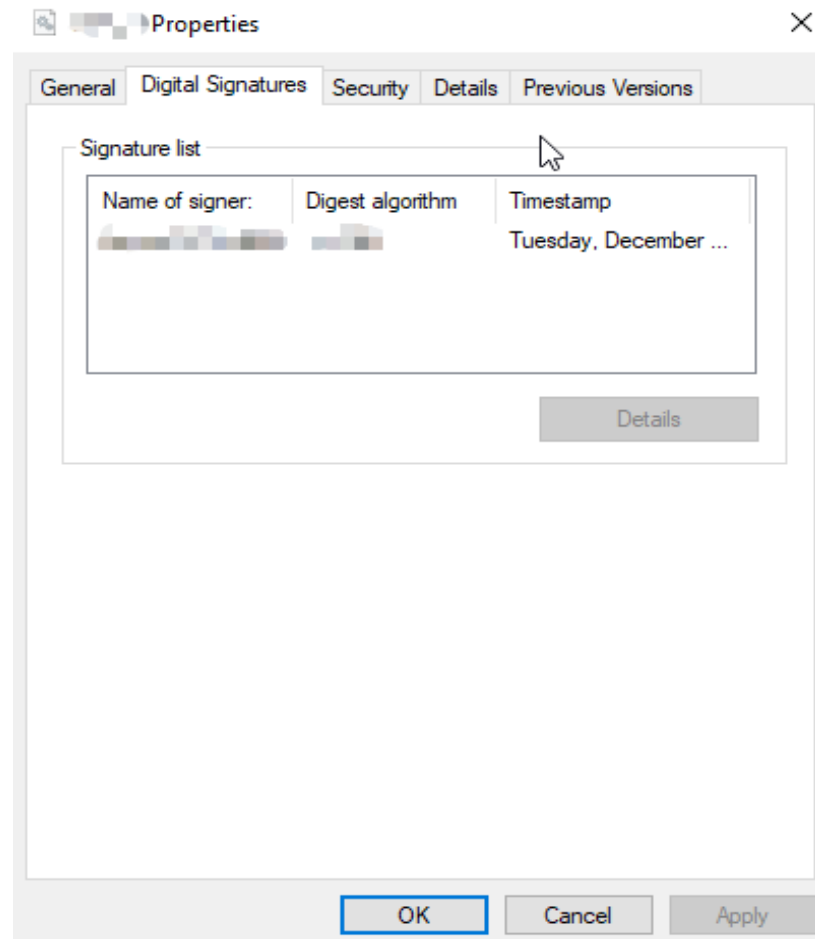
**Figure 14-156** Viewing the system drivers

| Filter Name | Num Instances | Altitude | Frame |
|-------------|---------------|----------|-------|
| WdFilter    | 3             | 328010   | 0     |
| storqosflt  | 0             | 244000   | 0     |
| wcifs       | 0             | 189900   | 0     |
| CldFlt      | 0             | 180451   | 0     |
| FileCrypt   | 0             | 141100   | 0     |
| luaflv      | 1             | 135000   | 0     |
| npsvctrig   | 1             | 46000    | 0     |
| Wof         | 1             | 40700    | 0     |

The following describes how to view a driver source and check whether the source is unknown.



- a. Go to the **C:\Windows\System32\drivers** directory on the local PC.
- b. Click the name of the unknown driver and choose **Properties** to view its details.
- c. Click the **Digital Signatures** tab to view the driver source.

**Figure 14-157** Viewing the driver source

## Troubleshooting

Before the troubleshooting, check whether the processes or drivers leading to the high CPU or bandwidth usage are normal, and handle the issue accordingly.

### Suggestions for non-malicious processes

1. If your ECS runs Windows Server 2008 or 2012, ensure that the available memory is 2 GiB or larger.
2. Check whether Windows Update is running.
3. Check whether the antivirus software is scanning files and programs on the backend.
4. Check whether any applications requiring high CPU or bandwidth resources are running on the ECS. If yes, modify ECS specifications or increase bandwidth.

5. If the ECS configuration meets the application requirements, deploy applications separately. For example, deploy the database and applications separately.

#### Suggestions for malicious processes

If the high CPU or bandwidth usage is caused by viruses or Trojan horses, manually stop the affected processes. You are advised to troubleshoot the issue as follows:

1. Use the commercial-edition antivirus software or install [Microsoft Safety Scanner](#) to scan for viruses in security mode.
2. Install the latest patches for Windows.
3. Run **MSconfig** to disable all drivers that are not delivered with Microsoft and check whether the fault is rectified. For details, see [How to perform a clean boot in Windows](#).

#### Suggestions for drivers from unknown sources

Some viruses and Trojan horses are loaded through the filter drivers of the system. If you find a driver from an unknown source, you are advised to uninstall it. You can also use commercial antivirus software or third-party security management tools to delete it.

If an unknown driver cannot be deleted, or will appear again after being deleted, it is usually a virus or Trojan horse driver. If the driver cannot be completely deleted using commercial antivirus software or third-party security management tools, you are advised to reinstall the OS and back up data before the reinstallation.

### 14.13.2 Why Is My Linux ECS Running Slowly?

If your ECS runs slowly or is inaccessible unexpectedly, the bandwidth or vCPU usage of the ECS may be excessively high. If you have created an alarm rule using Cloud Eye, the system automatically sends an alarm to you when the bandwidth or CPU usage reaches the threshold specified in the rule.

To handle this issue, perform the following operations:

1. Fault locating  
Identify the processes leading to high bandwidth or CPU usage.
2. Check whether the processes are normal and handle the issue accordingly.
  - If the processes are normal, optimize them or modify ECS specifications.
  - If the processes are malicious, use a third-party tool to automatically stop the processes or manually stop them.

#### Common Commands

The following uses the CentOS 7.2 64bit OS as an example to describe common commands. The commands may vary depending on Linux OS editions. For details, see the official documentation for the specific OS edition.

The common commands for checking Linux ECS performance metrics, such as the CPU usage, are as follows:

- `ps -aux`
- `ps -ef`
- `top`

## Locating High CPU Usage

1. Log in to the ECS using VNC.
2. Run the following command to check the OS running status:

### `top`

Information similar to the following is displayed.

```
top - 20:56:02 up 37 days, 9:09, 1 user, load average: 0.00, 0.01, 0.05
Tasks: 80 total, 1 running, 79 sleeping, 0 stopped, 0 zombie
%Cpu(s): 0.2 us, 0.3 sy, 0.0 ni, 99.5 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 3880024 total, 2963304 free, 178384 used, 738336 buff/cache
KiB Swap: 0 total, 0 free, 0 used. 3434808 avail Mem

 PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND
 8115 root 20 0 161896 2216 1564 R 0.3 0.1 0:00.01 top
 1 root 20 0 125480 3884 2604 S 0.0 0.1 0:11.32 systemd
 2 root 20 0 0 0 0 S 0.0 0.0 0:00.00 kthreadd
 3 root 20 0 0 0 0 S 0.0 0.0 0:00.04 ksoftirqd/0
 5 root 0 -20 0 0 0 S 0.0 0.0 0:00.00 kworker/0:0H
 7 root rt 0 0 0 0 S 0.0 0.0 0:00.18 migration/0
 8 root 20 0 0 0 0 S 0.0 0.0 0:00.00 rcu_bh
 9 root 20 0 0 0 0 S 0.0 0.0 7:32.18 rcu_sched
 10 root 0 -20 0 0 0 S 0.0 0.0 0:00.00 lru-add-drain
```

3. View the command output.
  - The first line in the command output is "20:56:02 up 37 days, 1 user, load average: 0.00, 0.01, 0.05", indicating that:  
The current system time is 20:56:02; the ECS has been running for 37 days; there is one login user; the last three values indicate the average CPU load in the last 1 minute, 5 minutes, and 15 minutes, respectively.
  - The third line in the command output shows the overall CPU usage.
  - The fourth line in the command output shows the overall memory usage.
  - The lower part of the command output shows the resource usage of each process.

### NOTE

1. On the `top` page, enter `q` or press `Ctrl+C` to exit.
2. Alternatively, click **Input Command** in the upper right corner of the VNC login page, paste or enter commands in the displayed dialog box, and click **Send**.
3. Common parameters in top commands are as follows:
  - `s`: Change the image update frequency.
  - `l`: Show or hide the first line for the top information.
  - `t`: Show or hide the second line for tasks and the third line for CPUs.
  - `m`: Show or hide the fourth line for Mem and the fifth line for Swap.
  - `N`: Sort processes by PID in ascending or descending order.
  - `P`: Sort processes by CPU usage in ascending or descending order.
  - `M`: Sort processes by memory usage in ascending or descending order.
  - `h`: Show help for commands.
  - `n`: Set the number of processes displayed in the process list.

4. Run the `ll /proc/PID/exe` command to obtain the program file specified by a PID.

```
root@elb-mq01:~# ll /proc/4243/exe
lrwxrwxrwx 1 root root 0 Mar 18 11:46 /proc/4243/exe -> /CloudResetPwdUpdateAgent/depend/jre1.8.0_131/bin/java
```

## Troubleshooting High CPU Usage

If the processes leading to high CPU usage are malicious, run the `top` command to stop them. If the `kswapd0` process leads to high CPU usage, optimize the program for the process or upgrade the ECS specifications for a larger memory capacity.

`kswapd0` is a virtual memory management process. When the physical memory becomes insufficient, `kswapd0` runs to allocate disk swap capacity for caching. This uses a large number of CPU resources.

- For the detected malicious processes

Quickly stop such processes on the `top` page. To do so, perform the following operations:

- a. Press the `k` key during the execution of the `top` command.
- b. Enter the PID of the process to be stopped.

The PID of the process is the value in the first column of the `top` command output. For example, to stop the process with PID 52, enter `52` and press `Enter`.

```
top - 21:07:38 up 37 days, 9:21, 1 user, load average: 0.01, 0.02, 0.05
Tasks: 81 total, 1 running, 79 sleeping, 1 stopped, 0 zombie
%Cpu(s): 0.0 us, 3.2 sy, 0.0 ni, 96.8 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 3880024 total, 2961520 free, 178960 used, 739544 buff/cache
KiB Swap: 0 total, 0 free, 0 used. 3434216 avail Mem
PID to signal/kill [default pid = 1, 52]
 PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND
 1 root 20 0 125480 3884 2604 S 0.0 0.1 0:11.32 systemd
 2 root 20 0 0 0 0 S 0.0 0.0 0:00.00 kthreadd
```

- c. After the operation is successful, information similar to the following is displayed. Press `Enter`.

```
top - 21:07:38 up 37 days, 9:21, 1 user, load average: 0.01, 0.02, 0.05
Tasks: 81 total, 1 running, 79 sleeping, 1 stopped, 0 zombie
%Cpu(s): 0.0 us, 3.2 sy, 0.0 ni, 96.8 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 3880024 total, 2961520 free, 178960 used, 739544 buff/cache
KiB Swap: 0 total, 0 free, 0 used. 3434216 avail Mem
Send pid 52 signal [15/sigterm]
 PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND
 1 root 20 0 125480 3884 2604 S 0.0 0.1 0:11.32 systemd
 2 root 20 0 0 0 0 S 0.0 0.0 0:00.00 kthreadd
```

- For the `kswapd0` process

To check the memory usage of a process, perform the following operations:

- a. Run the `top` command to check the resource usage of the `kswapd0` process.
- b. If the process remains in non-sleeping state for a long period of time, you can preliminarily determine that the system is consistently paging. In such a case, the high CPU usage is caused by insufficient memory.

```
Tasks: 81 total, 1 running, 79 sleeping, 1 stopped, 0 zombie
%Cpu(s): 0.2 us, 52.2 sy, 0.0 ni, 99.7 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 3880024 total, 3014820 free, 179024 used, 686180 buff/cache
KiB Swap: 0 total, 0 free, 0 used. 3433948 avail Mem
 PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND
 36 root 20 0 0 0 0 S 99.0 0.0 964:10.45 kswapd0
 4595 nginx 20 0 125392 3576 1040 S 0.3 0.1 60:04.91 nginx
 1 root 20 0 125480 3884 2604 S 0.0 0.1 0:11.47 systemd
```

- c. Run the **vmstat** command to check the virtual memory usage of the system.  
If the **si** and **so** values are large, the system is frequently paging and the physical memory of the system is insufficient.
  - **si**: Volume of data written from the swap partition to the memory per second, which is transferred from the disk to the memory.
  - **so**: Volume of data written from the memory to the swap partition per second, which is transferred from the memory to the disk.
- d. Further identify the causes of high memory usage. Run commands, such as **free** and **ps** to check the memory usage of the system and processes in the system.
- e. Restart the application or release the memory when traffic is light.  
To handle this issue, expand the ECS memory. If memory expansion is not allowed, optimize the application and enable hugepage memory.

## Handling High Bandwidth Usage

If the high bandwidth usage is caused by normal service access of non-malicious processes, enlarge the bandwidth to handle this issue. If the high bandwidth usage is caused by abnormal service access, for example, malicious access from certain IP addresses, CC attacks on the ECS, or malicious processes, use the traffic monitoring tool **nethogs** to monitor the bandwidth usage of each process in real time and identify faulty processes.

- Using **nethogs** for troubleshooting
  - a. Run the following command to install **nethogs**:  
**yum install nethogs -y**  
After the installation, run the **netgos** command to check bandwidth usage.  
Parameters in the **nethogs** command are as follows:
    - **-d**: Set the update interval in the unit of second. The default value is 1s.
    - **-t**: Enable tracing.
    - **-c**: Set the number of updates.
    - **device**: Set the NIC to be monitored. The default value is **eth0**.The following parameters are involved in command execution:
    - **q**: Exit **nethogs**.
    - **s**: Sort processes in the process list by TX traffic in ascending or descending order.
    - **r**: Sort processes in the process list by RX traffic in ascending or descending order.
    - **m**: Switch the display unit in the sequence of KB/s, KB, B, and MB.

- b. Run the following command to check the bandwidth usage of each process on the specified NIC:

```
nethogs eth1
```

```
NetHogs version 0.8.5
```

| PID   | USER  | PROGRAM                                | DEV  | SENT   | RECEIVED     |
|-------|-------|----------------------------------------|------|--------|--------------|
| 4596  | nginx | nginx: worker process                  | eth1 | 34.368 | 3.267 KB/sec |
| ?     | root  | 192.168.0.92:90-100.125.68.19:17873    |      | 0.179  | 0.246 KB/sec |
| ?     | root  | 192.168.0.92:11211-213.32.10.149:44945 |      | 0.000  | 0.000 KB/sec |
| ?     | root  | 192.168.0.92:20101-185.176.26.66:43408 |      | 0.000  | 0.000 KB/sec |
| ?     | root  | unknown TCP                            |      | 0.000  | 0.000 KB/sec |
| TOTAL |       |                                        |      | 34.540 | 3.512 KB/sec |

The parameters in the command output are as follows:

- **PID:** ID of the process.
  - **USER:** user who runs the process.
  - **PROGRAM:** IP addresses and port numbers of the process and connection, respectively. The former is for the server and the latter is for the client.
  - **DEV:** Network port to which the traffic is destined.
  - **SENT:** Volume of data sent by the process per second.
  - **RECEIVED:** Volume of data received by the process per second.
- c. Stop malicious programs or blacklist malicious IP addresses.  
To stop a malicious process, run the **kill** *PID* command.  
To blacklist a malicious IP address or limit its rate, use iptables.

## 14.14 Database Applications

### 14.14.1 Can a Database Be Deployed on an ECS?

Yes. You can deploy a database of any type on an ECS.

### 14.14.2 Does an ECS Support Oracle Databases?

Yes. You are advised to perform a performance test beforehand to ensure that the Oracle database can meet your requirements.

---

# A Change History

---

| Released On | Description                               |
|-------------|-------------------------------------------|
| 2024-04-15  | This issue is the first official release. |