

Elastic Cloud Server

User Guide

Issue 01
Date 2024-07-23



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.

Contents

1 Service Overview.....	1
1.1 What Is ECS?.....	1
1.2 ECS Advantages.....	3
1.3 ECS Application Scenarios.....	5
1.4 Notes and Constraints on Using ECSs.....	6
1.5 ECS and Other Services.....	7
1.6 Instances.....	10
1.6.1 ECS Overview.....	10
1.6.2 ECS Lifecycle.....	10
1.6.3 ECS Types.....	11
1.7 x86 ECS Specifications and Types.....	13
1.7.1 A Summary List of x86 ECS Specifications	13
1.7.2 General-Purpose ECSs.....	28
1.7.3 Dedicated General-Purpose ECSs.....	29
1.7.4 Memory-optimized ECSs.....	37
1.7.5 Large-Memory ECSs.....	41
1.7.6 Disk-intensive ECSs.....	43
1.7.7 Ultra-high I/O ECSs.....	45
1.7.8 GPU-accelerated ECSs.....	52
1.7.9 AI-accelerated ECSs.....	59
1.8 Images.....	63
1.8.1 Image Types.....	64
1.8.2 Cloud-Init.....	65
1.9 EVS Disks.....	66
1.10 Network.....	67
1.11 Security.....	69
1.11.1 Data Protection.....	69
1.11.1.1 User Encryption.....	69
1.11.2 Fault Recovery.....	71
1.12 Billing.....	72
1.13 Permissions Management.....	75
1.14 Region and AZ.....	80
2 Getting Started.....	82

2.1 Creating an ECS.....	82
2.1.1 Overview.....	82
2.1.2 Step 1: Configure Basic Settings.....	82
2.1.3 Step 2: Configure Network.....	86
2.1.4 Step 3: Configure Advanced Settings.....	89
2.1.5 Step 4: Confirm.....	92
2.2 Logging In to an ECS.....	92
2.3 Initializing EVS Data Disks.....	94
2.3.1 Scenarios and Disk Partitions.....	94
2.3.2 Initializing a Windows Data Disk (Windows Server 2008).....	95
2.3.3 Initializing a Windows Data Disk (Windows Server 2019).....	102
2.3.4 Initializing a Linux Data Disk (fdisk).....	110
2.3.5 Initializing a Linux Data Disk (parted).....	116
2.3.6 Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2008).....	122
2.3.7 Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2012).....	130
2.3.8 Initializing a Linux Data Disk Larger Than 2 TiB (parted).....	138
3 Instances.....	145
3.1 Selecting an ECS Billing Mode.....	145
3.1.1 Pay-per-Use Billing.....	145
3.1.2 Reserved Instances.....	146
3.1.2.1 Reserved Instance Overview.....	146
3.1.2.2 Enabling and Purchasing a Reserved Instance.....	147
3.1.2.3 Modifying RI Attributes.....	150
3.2 Creating an ECS.....	151
3.2.1 Creating the Same ECS.....	152
3.3 Logging In to a Windows ECS.....	153
3.3.1 Login Overview (Windows).....	153
3.3.2 Logging In to a Windows ECS Using VNC.....	154
3.3.3 Logging In to a Windows ECS Using MSTSC.....	155
3.3.4 Logging In to a Windows ECS from a Linux Computer.....	162
3.3.5 Logging In to a Windows ECS from a macOS Server.....	164
3.3.6 Logging In to a Windows ECS from a Mobile Terminal.....	167
3.4 Logging In to a Linux ECS.....	172
3.4.1 Login Overview (Linux).....	172
3.4.2 Logging In to a Linux ECS Using VNC.....	173
3.4.3 Logging In to a Linux ECS Using an SSH Key Pair.....	175
3.4.4 Logging In to a Linux ECS Using an SSH Password.....	181
3.4.5 Logging In to a Linux ECS from a macOS Server.....	183
3.4.6 Logging In to a Linux ECS from a Mobile Terminal.....	184
3.5 Managing GPU Drivers of GPU-accelerated ECSs.....	196
3.5.1 GPU Driver.....	196
3.5.2 Obtaining a Tesla Driver and CUDA Toolkit.....	197

3.5.3 Installing a GRID Driver on a GPU-accelerated ECS.....	198
3.5.4 Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS.....	208
3.5.5 Uninstalling a GPU Driver from a GPU-accelerated ECS.....	222
3.6 Managing ECSs.....	227
3.6.1 Changing the Time Zone for an ECS.....	227
3.6.2 Obtaining Metadata.....	229
3.6.3 Passing User Data to ECSs.....	238
3.6.4 Changing ECS Names.....	246
3.6.5 Migrating an ECS.....	247
3.6.6 Managing ECS Groups.....	248
3.6.7 Automatically Recovering ECSs.....	251
3.6.8 (Optional) Configuring Mapping Between Hostnames and IP Addresses.....	252
3.6.9 Starting and Stopping ECSs.....	253
3.7 Modifying ECS Specifications.....	254
3.7.1 General Operations.....	254
3.8 Reinstalling or Changing the OS.....	256
3.8.1 Reinstalling the OS.....	256
3.8.2 Changing the OS.....	258
3.9 Viewing ECS Information.....	261
3.9.1 Viewing ECS Creation Statuses.....	262
3.9.2 Viewing Failed Tasks.....	262
3.9.3 Viewing ECS Details (List View).....	263
3.9.4 Exporting ECS Information.....	264
4 Images.....	265
4.1 Overview.....	265
4.2 Creating an Image.....	266
5 EVS Disks.....	269
5.1 Overview.....	269
5.2 Adding a Disk to an ECS.....	269
5.3 Attaching an EVS Disk to an ECS.....	270
5.4 Detaching an EVS Disk from a Running ECS.....	272
5.5 Expanding the Capacity of an EVS Disk.....	274
5.6 Expanding the Local Disks of a Disk-intensive ECS.....	275
5.7 Enabling Advanced Disk.....	276
6 Elastic Network Interface.....	277
6.1 Overview.....	277
6.2 Attaching a Network Interface.....	278
6.3 Detaching a Network Interface.....	279
6.4 Changing a VPC.....	280
6.5 Modifying a Private IP Address.....	282
6.6 Managing Virtual IP Addresses.....	282

6.7 Enabling NIC Multi-Queue.....	288
6.8 Dynamically Assigning IPv6 Addresses.....	291
7 EIP.....	310
7.1 Overview.....	310
7.2 Binding an EIP.....	311
7.3 Unbinding an EIP.....	312
7.4 Changing an EIP.....	313
7.5 Changing an EIP Bandwidth.....	314
7.6 Enabling Internet Connectivity for an ECS Without an EIP.....	314
8 Security.....	318
8.1 Methods for Improving ECS Security.....	318
8.2 Security Groups.....	322
8.2.1 Overview.....	322
8.2.2 Default Security Group and Rules.....	322
8.2.3 Security Group Configuration Examples.....	324
8.2.4 Configuring Security Group Rules.....	329
8.2.5 Changing a Security Group.....	331
8.3 HSS.....	332
8.4 Project and Enterprise Project.....	333
8.5 Protection for Mission-Critical Operations.....	334
9 CBR.....	339
9.1 Overview.....	339
9.2 Backing Up an ECS.....	347
10 Passwords and Key Pairs.....	350
10.1 Passwords.....	350
10.1.1 Application Scenarios for Using Passwords.....	350
10.1.2 Changing the Login Password on an ECS.....	351
10.1.3 Resetting the Password for Logging In to a Windows ECS.....	353
10.1.4 Resetting the Password for Logging In to a Linux ECS.....	355
10.2 Key Pairs.....	358
10.2.1 Application Scenarios for Using Key Pairs.....	358
10.2.2 (Recommended) Creating a Key Pair on the Management Console.....	359
10.2.3 Creating a Key Pair Using PuTTYgen.....	360
10.2.4 Importing a Key Pair.....	364
11 Permissions Management.....	366
11.1 Creating a User and Granting ECS Permissions.....	366
11.2 ECS Custom Policies.....	367
12 Resources and Tags.....	370
12.1 Tag Management.....	370
12.1.1 Overview.....	370

12.1.2 Adding Tags.....	372
12.1.3 Searching for Resources by Tag.....	374
12.1.4 Deleting a Tag.....	375
12.2 Quota Adjustment.....	376
13 Monitoring.....	377
13.1 Monitoring ECSs.....	377
13.2 Basic ECS Metrics.....	378
13.3 OS Monitoring Metrics Supported by ECSs with the Agent Installed.....	384
13.4 Process Monitoring Metrics Supported by ECSs with the Agent Installed.....	426
13.5 Setting Alarm Rules.....	431
13.6 Viewing ECS Metrics.....	432
14 CTS.....	434
14.1 Key Operations Supported by CTS.....	434
14.2 Viewing Traces.....	435
15 FAQs.....	438
15.1 Common Topics.....	438
15.2 ECS Overview.....	438
15.2.1 What Are the Precautions for Using ECSs?.....	439
15.2.2 What Can I Do with ECSs?.....	439
15.2.3 Can ECSs Automatically Recover After the Physical Host Accommodating the ECSs Becomes Faulty?.....	439
15.3 Regions and AZs.....	440
15.3.1 What Is an AZ?.....	440
15.4 Billing.....	440
15.4.1 How Can I Stop an ECS from Being Billed?.....	440
15.5 Creation and Deletion.....	441
15.5.1 What Should I Do If the ECS Resources to Be Purchased Are Sold Out?.....	441
15.5.2 What Is the Creation Time and Startup Time of an ECS?.....	441
15.5.3 Why Does the Failures Area Show an ECS Creation Failure But the ECS List Displays the Created ECS?.....	441
15.5.4 When Does an ECS Become Provisioned?.....	441
15.5.5 Why Cannot I View the ECSs Being Created Immediately After I Pay for Them?.....	442
15.5.6 Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image?.....	442
15.5.7 What Do I Do If I Selected an Incorrect Image for My ECS?.....	443
15.5.8 Should I Choose Windows OS or Linux OS for My ECS?.....	443
15.5.9 How Quickly Can I Obtain an ECS?.....	444
15.5.10 How Can I Manage ECSs by Group?.....	444
15.5.11 Why Did I Fail to Configure an Anti-Affinity ECS Group?.....	445
15.5.12 What Happens After I Click the Delete Button?.....	445
15.5.13 Can a Deleted ECS Be Provisioned Again?.....	445
15.5.14 Can a Deleted ECS Be Restored?.....	445
15.5.15 How Do I Delete or Restart an ECS?.....	446

15.5.16 Can I Forcibly Restart or Stop an ECS?.....	446
15.6 Login and Connection.....	446
15.6.1 What Are the Username and Password for Remote Logins?.....	446
15.6.2 Why Cannot I Use the Username and Password Configured During the Creation of a GPU-accelerated ECS to Log In to the ECS Through SSH?.....	447
15.6.3 Why Can't I Log In to My Windows ECS?.....	448
15.6.4 Why Can't I Log In to My Linux ECS?.....	456
15.6.5 What Should I Do If I Cannot Use MSTSC to Log In to an ECS Running the Windows Server 2012 OS?.....	462
15.6.6 How Can I Change a Remote Login Port?.....	463
15.6.7 Why Cannot I Use a Non-Default SSH Port to Log In to My Linux ECS?.....	466
15.6.8 Why Can't I Obtain the Password for Logging In to My Windows ECS Authenticated Using a Key Pair?.....	468
15.6.9 What Browser Version Is Required to Remotely Log In to an ECS?.....	470
15.6.10 Why Are Garbled Characters Displayed When I Log In to My ECS Using VNC?.....	470
15.6.11 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?.....	471
15.6.12 What Should I Do If I Cannot View Data After Logging In to an ECS Using VNC?.....	471
15.6.13 Why Does a Blank Screen Appear After I Attempted to Log In to an ECS Using VNC?.....	471
15.6.14 What Should I Do If Error Code 1006 or 1000 Is Displayed When I Log In to an ECS Through the Management Console?.....	472
15.6.15 Why No Audio File Can Be Properly Played on My Windows ECS Logged In Using VNC?.....	472
15.6.16 How Can I Change the Resolution of a Windows ECS?.....	474
15.6.17 Why Does an Authentication Failure Occurs After I Attempt to Remotely Log In to a Windows ECS?.....	476
15.6.18 Why Can't I Use the Local Computer to Connect to My Windows ECS?.....	478
15.6.19 How Can I Obtain the Permission to Remotely Log In to a Windows ECS?.....	483
15.6.20 Why Does the System Display No Remote Desktop License Servers Available to Provide a License When I Log In to a Windows ECS?.....	485
15.6.21 Why Does the System Display Error Code 0x112f When I Log In to a Windows ECS?.....	487
15.6.22 Why Does the System Display Error Code 0x1104 When I Log In to a Windows ECS?.....	488
15.6.23 Why Does the System Display Error Code 122.112... When I Log In to a Windows ECS?.....	492
15.6.24 Why Does the System Display Invalid Certificate or Associated Chain When I Log In to a Windows ECS from a Mac?.....	494
15.6.25 Why Does the System Display a Message Indicating Invalid Credentials When I Attempt to Access a Windows ECS?.....	498
15.6.26 Why Does an Internal Error Occur When I Log In to My Windows ECS?.....	503
15.6.27 Why Is My Remote Session Interrupted by a Protocol Error?.....	504
15.6.28 Why Am I Seeing an Error Message That Says Identity of Remote Computer Cannot be Verified When I Log In to a Windows ECS?.....	506
15.6.29 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?.....	507
15.6.30 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?.....	507
15.6.31 Why Does My Remote Desktop Session End Because Another User Logs In When I Log In to a Windows ECS?.....	511

15.6.32 Why Does an ECS Fail to Be Remotely Connected Using RDP and Internal Error Code 4 Is Displayed?.....	514
15.6.33 Why Am I Seeing the Error Message "Module is unknown" When I Remotely Log In to a Linux ECS?.....	515
15.6.34 What Should I Do If Error Message "Permission denied" Is Displayed When I Remotely Log In to a Linux ECS?.....	517
15.6.35 What Should I Do If Error Message "read: Connection reset by peer" Is Displayed When I Remotely Log In to a Linux ECS?.....	519
15.6.36 Why Am I Seeing the Error Message "Access denied" When I Remotely Log In to a Linux ECS?.....	520
15.6.37 What Should I Do If Error Message "Disconnected: No supported authentication methods available" Is Displayed When I Remotely Log In to a Linux ECS?.....	521
15.7 How Do I Handle Error Messages Displayed on the Management Console?.....	521
15.8 How Can I Recover a Windows ECS with an Abnormal Virtualization Driver?.....	524
15.9 Why Is My Windows ECS Muted?.....	527
15.10 How Do I Change an ECS SID?.....	530
15.11 Why Does a Pay-per-Use ECS Fail to Be Started?.....	531
15.12 ECS Management.....	531
15.12.1 How Can a Changed Static Hostname Take Effect Permanently?.....	532
15.12.2 Is an ECS Hostname with Suffix .novalocal Normal?.....	535
15.12.3 Why Is the Hostname of My ECS Restored to the Original Name After the ECS Is Restarted?.....	535
15.12.4 How Can I Set Sequential ECS Names When Creating Multiple ECSs?.....	536
15.12.5 How Can I Modify ECS Specifications?.....	539
15.12.6 Why Do the Disks of a Windows ECS Go Offline After I Modify the ECS Specifications?.....	539
15.12.7 Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?.....	542
15.13 OS Management.....	543
15.13.1 Does OS Change Incur Fees?.....	543
15.13.2 Can I Install or Upgrade the OS of an ECS?.....	543
15.13.3 Can I Change the OS of an ECS?.....	543
15.13.4 How Long Does It Take to Change an ECS OS?.....	544
15.13.5 Will I Lose My Disk Data If I Reinstall ECS OS, Change the OS, or Change the ECS Specifications?.....	544
15.13.6 Does OS Reinstallation Incur Fees?.....	545
15.13.7 Can I Select Another OS During ECS OS Reinstallation?.....	545
15.13.8 How Long Does It Take to Reinstall an ECS OS?.....	545
15.13.9 Do ECSs Support GUI?.....	546
15.13.10 How Can I Install a GUI on an ECS Running CentOS 6?.....	546
15.13.11 How Can I Install a GUI on an ECS Running CentOS 7?.....	546
15.13.12 How Can I Install a GUI on an ECS Running Ubuntu?.....	547
15.13.13 How Can I Install a GUI on an ECS Running Debian?.....	552
15.13.14 Why Does the OS Fail to Respond When kdump Occurs on a Linux ECS?.....	553
15.13.15 How Can I Upgrade the Kernel of a Linux ECS?.....	555
15.13.16 Why Cannot My ECS OS Start Properly?.....	558
15.13.17 How Can I Enable SELinux on an ECS Running CentOS?.....	558
15.13.18 Why Does a Forcibly-Stopped Linux ECS Fail to Be Restarted?.....	559

15.13.19 How Do I View the GPU Usage of a GPU-accelerated ECS?.....	560
15.14 File Upload/Data Transfer.....	562
15.14.1 How Do I Upload Files to My ECS?.....	562
15.14.2 How Can I Transfer Files from a Local Windows Computer to a Windows ECS?.....	563
15.14.3 How Can I Use WinSCP to Transfer Files from a Local Windows Computer to a Linux ECS?.....	566
15.14.4 How Can I Transfer Files from a Local Mac to a Windows ECS?.....	568
15.14.5 How Can I Use SCP to Transfer Files Between a Local Linux Computer and a Linux ECS?.....	571
15.14.6 How Can I Use SFTP to Transfer Files Between a Local Linux Computer and a Linux ECS?.....	572
15.14.7 How Can I Use FTP to Transfer Files from a Local Windows Computer to a Windows or Linux ECS?.....	574
15.14.8 How Can I Use FTP to Transfer Files Between a Local Linux Computer and a Linux ECS?.....	575
15.14.9 How Can I Transfer Data Between a Local Computer and a Windows ECS?.....	576
15.14.10 What Should I Do If the Connection Between the Client and the Server Times Out When I Upload a File Using FTP?.....	579
15.14.11 What Should I Do If Writing Data Failed When I Upload a File Using FTP?.....	580
15.14.12 Why Does Internet Access to an ECS Deployed with FTP Fail?.....	581
15.14.13 Why Am I Seeing an FTP Folder Error When I Open a Folder on an FTP Server?.....	584
15.14.14 Why Do I Fail to Connect to a Linux ECS Using WinSCP?.....	585
15.15 ECS Migration.....	586
15.15.1 Can I Migrate an ECS to Another Region, AZ, or Account?.....	587
15.16 Disk Management.....	587
15.16.1 Why Can't I Find My Newly Purchased Data Disk After I Log In to My Windows ECS?.....	587
15.16.2 How Can I Adjust System Disk Partitions?.....	588
15.16.3 How Can I Obtain the Mapping Between Disk Partitions and Disk Devices on a Windows ECS?.....	594
15.16.4 How Can I Obtain the Mapping Between Disk Partitions and Disk Devices on a Linux ECS?.....	597
15.16.5 How Can I Enable Virtual Memory on a Windows ECS?.....	599
15.16.6 How Can I Add the Empty Partition of an Expanded System Disk to the End Root Partition Online?.....	601
15.16.7 How Can I Add the Empty Partition of an Expanded System Disk to the Non-end Root Partition Online?.....	602
15.16.8 Can I Attach Multiple Disks to an ECS?.....	605
15.16.9 What Are the Requirements for Attaching an EVS Disk to an ECS?.....	606
15.16.10 Which ECSs Can Be Attached with SCSI EVS Disks?.....	606
15.16.11 How Do I Obtain My Disk Device Name in the ECS OS Using the Device Identifier Provided on the Console?.....	607
15.16.12 What Should I Do If Attaching a Disk to a Windows ECS Failed But There Are Still Available Device Names?.....	612
15.16.13 Why Does a Linux ECS with a SCSI Disk Attached Fails to Be Restarted?.....	612
15.16.14 How Can I Check Whether the ECSs Attached with the Same Shared SCSI Disk Are in the Same ECS Group?.....	613
15.16.15 Can All Users Use the Encryption Feature?.....	614
15.16.16 How Can I Add an ECS with Local Disks Attached to an ECS Group?.....	615
15.16.17 Will My EVS Disk Be Deleted When I Delete Its Server?.....	615
15.16.18 Why Does a Disk Attached to a Windows ECS Go Offline?.....	616

15.16.19 Why Does the Disk Drive Letter Change After the ECS Is Restarted?.....	617
15.16.20 How Can I Obtain Data Disk Information If Tools Are Uninstalled?.....	618
15.16.21 How Can I Rectify the Fault That May Occur on a Linux ECS with an NVMe SSD Disk Attached?	619
15.16.22 Why Is the Device Name of My C6 ECS in the sd* Format?.....	621
15.16.23 Why Are Disk Error Logs Printed After a Disk Attached to an ECS Is Formatted with the ext4 File System?.....	622
15.17 Passwords and Key Pairs.....	623
15.17.1 How Can I Change the Password for Logging In to a Linux ECS?.....	623
15.17.2 What Is the Default Password for Logging In to a Linux ECS?.....	623
15.17.3 How Can I Set the Validity Period of the Image Password?.....	623
15.17.4 Changing the Login Password on an ECS.....	624
15.17.5 What Should I Do If the System Displays a Message Indicating that the Password Is Incorrect When I Remotely Log In to My ECS?.....	626
15.17.6 What Should I Do If I Cannot Log In to My ECS Using the Initial Password After I Use It for a Period of Time?.....	626
15.17.7 Disabling SELinux.....	627
15.17.8 How Can I Obtain the Key Pair Used by My ECS?.....	627
15.17.9 How Can I Use a Key Pair?.....	627
15.17.10 What Should I Do If a Key Pair Cannot Be Imported?.....	629
15.17.11 Why Does the Login to My Linux ECS Using a Key File Fail?.....	629
15.17.12 What Should I Do If I Cannot Download a Key Pair?.....	630
15.17.13 Why Does a Key Pair Created Using puttygen.exe Fail to Be Imported on the Management Console?.....	631
15.17.14 What Is the Cloudbase-Init Account in Windows ECSs Used for?.....	633
15.17.15 What Should I Do If Cloud-Init Does Not Work After Python Is Upgraded?.....	634
15.18 Network Configurations.....	635
15.18.1 Can Multiple EIPs Be Bound to an ECS?.....	635
15.18.2 Can an ECS Without an EIP Bound Access the Internet?.....	636
15.18.3 What Should I Do If an EIP Cannot Be Pinged?.....	637
15.18.4 Why Can I Remotely Access an ECS But Cannot Ping It?.....	643
15.18.5 How Do I Query the Egress Public IP Address of My ECS?.....	643
15.18.6 How Can I Configure the NTP and DNS Servers for an ECS?.....	644
15.18.7 What Should I Do If NIC Flapping Occurs After My ECS Specifications Are Modified?.....	649
15.18.8 Will NICs Added to an ECS Start Automatically?.....	651
15.18.9 How Do I Change the CIDR Block of an ECS Subnet?.....	651
15.18.10 How Can I Check Whether the Network Communication Is Normal Between Two ECSs Equipped with an InfiniBand NIC Driver?.....	652
15.18.11 How Can I Manually Configure an IP Address for an InfiniBand NIC?.....	653
15.18.12 Why Is the NIC Not Working?.....	654
15.18.13 How Can I Handle the Issue that a Windows 7 ECS Equipped with an Intel 82599 NIC Reports an Error in SR-IOV Scenarios?.....	656
15.18.14 How Can I Add a Static Route to a CentOS 6.5 OS?.....	657
15.18.15 Why Can't My Linux ECS Obtain Metadata?.....	657

15.18.16 Why Can't My Windows ECS Access the Internet?.....	661
15.18.17 Why Does My Linux ECS Fail to Access the Internet?.....	667
15.18.18 How Do I Troubleshoot an Unresponsive Website Hosted on My ECS?.....	672
15.18.19 Why Did I See "Invalid argument" or "neighbour table overflow" During an Access to a Linux ECS?.....	679
15.18.20 How Can I Obtain the MAC Address of My ECS?.....	680
15.18.21 How Can I Test the Network Performance of Linux ECSs?.....	682
15.18.22 Why Can't I Use DHCP to Obtain a Private IP Address?.....	691
15.18.23 How Can I View and Modify Kernel Parameters of a Linux ECS?.....	693
15.18.24 How Can I Configure Port Redirection?.....	698
15.18.25 Can the ECSs of Different Accounts Communicate over an Intranet?.....	700
15.18.26 Will ECSs That I Purchased Deployed in the Same Subnet?.....	700
15.19 Security Configurations.....	700
15.19.1 Are ECSs with Simple Passwords Easily Attacked?.....	701
15.19.2 How Is ECS Security Ensured?.....	701
15.19.3 How Can I Disable Operation Protection?.....	701
15.20 Resource Management and Tags.....	702
15.20.1 How Can I Create and Delete Tags and Search for ECSs by Tag?.....	702
15.21 Resource Monitoring.....	704
15.21.1 Why Is My Windows ECS Running Slowly?.....	704
15.21.2 Why Is My Linux ECS Running Slowly?.....	708
15.22 Database Applications.....	712
15.22.1 Can a Database Be Deployed on an ECS?.....	712
15.22.2 Does an ECS Support Oracle Databases?.....	712
15.22.3 What Should I Do If a Msg 823 Error Occurs in Oracle, MySQL, or SQL Server System Logs After a Disk Initialization Script Is Executed?.....	712

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 vCPUs, memory, and specifications of an ECS, see [A Summary List of x86 ECS Specifications](#).
- 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](#).

Why ECS

- Rich specifications: A variety of ECS types with custom specifications are available for different scenarios.
- Various image types: Public, private, and shared images are available for you to choose from.
- A broad range of disk types: High I/O, and ultra-high I/O disks are provided to meet the requirements of different service scenarios.
- Reliable data: High-throughput virtual block storage uses the distributed architecture to ensure high availability and it can be scaled out as needed.
- Security protection: The network is isolated and protected using security group rules. Security services, such as Anti-DDoS, Web Application Firewall (WAF), and Vulnerability Scan Service (VSS) can also be used to further enhance ECS security.
- Auto scaling: Elastic computing resources can be automatically adjusted to suit your needs.
- Efficient O&M: ECSs can be efficiently managed through the management console, remote terminals, or APIs with full rights.

- Cloud monitoring: Cloud Eye samples monitored metrics in real time, generates alarms when detecting problems, and immediately notifies related personnel of the alarms.
- Load balancing: Elastic Load Balance (ELB) evenly distributes incoming traffic across ECSs to prevent overload on an individual ECS. Applications are more tolerant of errors and bursty traffic.


For more details, see [ECS Advantages](#) and [ECS Application Scenarios](#).

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) backs up data for EVS disks and ECSs, and uses snapshots and backups to restore the EVS disks and ECSs.

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.

Log in to the management console with your account and choose **Elastic Cloud Server** on the homepage.

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.

For details, see [General-Purpose ECSs](#) and [Dedicated General-Purpose ECSs](#).

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.

For details, see [Memory-optimized ECSs](#).

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.

For details, see [GPU-accelerated ECSs](#).

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.

For details, see [Disk-intensive ECSs](#).

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) or a restart may occur on the ECS.
- Ensure that there is at least 2 GiB of idle memory. Otherwise, BSOD, freezing, or service failures may occur.
- Do not modify the registry. Otherwise, the system startup may fail. 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 disable virtual memory. Otherwise, system performance may deteriorate, or system exceptions may occur.
- Do not delete the VMTool program, or an exception may occur on the ECS.

Precautions for Using Linux ECSs

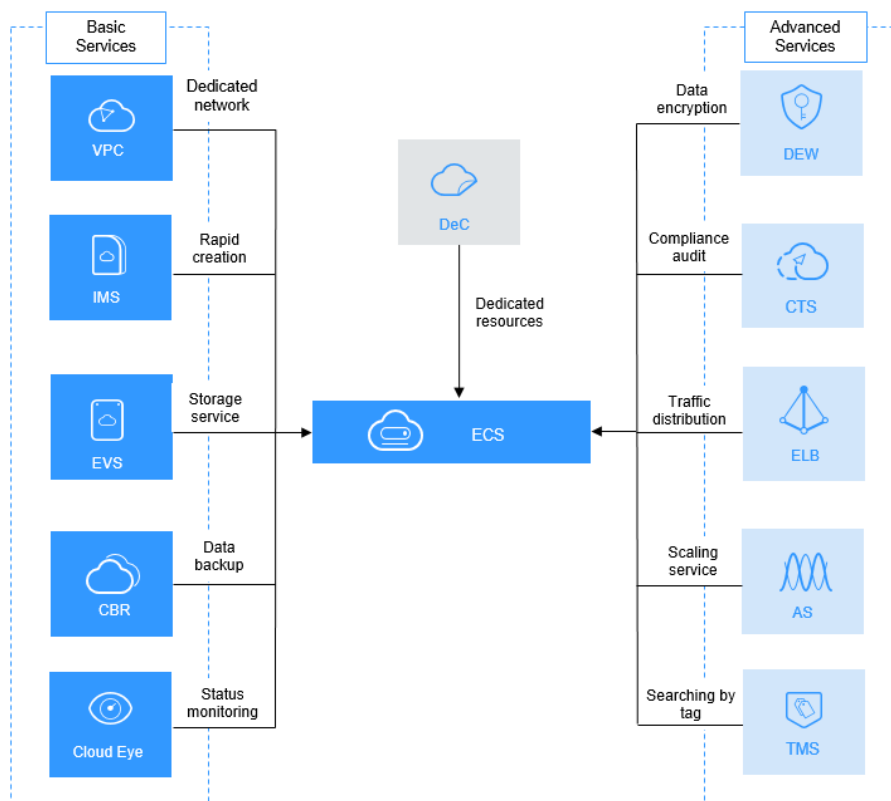
- Do not modify the `/etc/issue` file. Otherwise, the OS distribution will not be identified.
- Do not delete system directories or files. Otherwise, the system may fail to run or start.
- Do not change the permissions for or names of system directories. Otherwise, the system may fail to run or start.
- Do not upgrade the kernel of the Linux unless necessary.

When you have to upgrade the Linux kernel, follow the instructions provided in [How Can I Upgrade the Kernel of a Linux ECS?](#)
- 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 the IP address, subnet mask, or 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

[Figure 1-2](#) shows the relationships between ECS and other services.

Figure 1-2 Relationships between ECS and other services



ECS-related Services

Table 1-1 ECS-related services

Service	Function	Related Operation
Auto Scaling (AS)	Automatically adjusts ECS resources to keep up with changes in demand based on the configured AS policies. This improves resource utilization and reduces costs.	Using an Existing ECS to Create an AS Configuration
Elastic Load Balance (ELB)	Automatically distributes traffic to multiple ECSs. This improves fault tolerance.	Adding Backend Servers
Elastic Volume Service (EVS)	Allows you to attach EVS disks to an ECS and expand their capacity.	<ul style="list-style-type: none"> Attaching a Non-Shared Disk Attaching a Shared Disk

Service	Function	Related Operation
Virtual Private Cloud (VPC)	Allows you to create your own dedicated network on the cloud by customizing security groups, VPNs, IP address ranges, and bandwidths. This simplifies network management. You can also customize the ECS access rules within a security group and between security groups to improve ECS security.	<ul style="list-style-type: none">• Assigning an EIP and Binding It to an ECS• Adding a Security Group Rule
Image Management Service (IMS)	Allows you to create ECSs using images. This improves the efficiency of ECS creation. You can also use an existing ECS to create a private image and export the data of the ECS system disk or data disks.	<ul style="list-style-type: none">• Creating a Data Disk Image from an ECS• Creating a Full-ECS Image from an ECS
Dedicated Computing Cluster (DCC)	To physically isolate your ECS, apply for a DCC before creating the ECS. After you obtain the DCC and set a region for it, your ECS is automatically allocated to the DCC.	<ul style="list-style-type: none">• Enabling a DeC
Cloud Eye	Allows you to check the status of monitored ECS resources. This can be done without requiring additional plug-ins.	<ul style="list-style-type: none">• Basic ECS Metrics• OS Monitoring Metrics Supported by ECSs with the Agent Installed• Process Monitoring Metrics Supported by ECSs with the Agent Installed
Data Encryption Workshop (DEW)	The encryption feature relies on DEW. You can use an encrypted image or EVS disks when creating an ECS. In such a case, you need to use the key provided by DEW to improve data security.	<ul style="list-style-type: none">• EVS Encryption• Encrypting Images
Cloud Trace Service (CTS)	Records ECS-related operations for later query, auditing, and backtracking.	Supported CTS Operations
Cloud Backup and Recovery (CBR)	Backs up EVS disks and ECSs for fault recovery.	Creating a Server Backup Vault

Service	Function	Related Operation
Tag Management Service (TMS)	Identifies ECSs to help classification and search.	<ul style="list-style-type: none">• Adding Tags• Searching for Resources by Tag

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](#).
- For details about ECS specifications, see [A Summary List of x86 ECS Specifications](#).

1.6.2 ECS Lifecycle

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

Table 1-2 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.

Status	Status Attribute	Description
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.
Frozen	Stable	The ECS has been stopped by the administrator because the order has expired or is overdue. An ECS in this state cannot provide services. The system retains it for a period of time. If it is not renewed after the time expires, the system will automatically delete the ECS.

1.6.3 ECS Types

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

- [General-Purpose ECSs](#)
- [Dedicated General-Purpose ECSs](#)
- [Memory-optimized ECSs](#)
- [Large-Memory ECSs](#)
- [Disk-intensive ECSs](#)
- [Ultra-high I/O ECSs](#)
- [GPU-accelerated ECSs](#)
- [AI-accelerated ECSs](#)

ECS Flavor Naming Rules

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

Example: s6.medium.4

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.
- **B** specifies the type ID. For example, **6** in **s6** indicates the sixth-generation general-purpose ECS.
- **C** specifies the flavor size, such as medium, large, xlarge, 2xlarge, 4xlarge, or 8xlarge.
- **D** specifies the ratio of memory to vCPUs expressed in a digit. For example, value **4** indicates that the ratio of memory to vCPUs is 4.

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

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 A Summary List of x86 ECS Specifications

General-Purpose

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.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.medium.4	1	4	0.8/0.1	10	1	1	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
s6.2xlarge.4	8	32	3/0.75	50	2	2	KVM

Dedicated General-Purpose

Table 1-4 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
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

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
c6s.12xlarge.4	48	192	22/22	650	16	8	KVM
c6s.16xlarge.4	64	256	30/30	850	32	8	KVM

Table 1-5 C6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	EVS Basic Bandwidth (Gbit/s)	Virtualization
c6.large.2	2	4	4/1.2	40	2	2	1.5	KVM
c6.xlarge.2	4	8	8/2.4	80	2	3	2	KVM
c6.2xlarge.2	8	16	15/4.5	150	4	4	2.5	KVM
c6.3xlarge.2	12	24	17/7	200	4	6	4	KVM
c6.4xlarge.2	16	32	20/9	280	8	8	5	KVM
c6.6xlarge.2	24	48	25/14	400	8	8	8	KVM
c6.8xlarge.2	32	64	30/18	550	16	8	10	KVM
c6.12xlarge.2	48	96	35/27	750	16	8	15	KVM
c6.16xlarge.2	64	128	40/36	1000	32	8	20	KVM
c6.large.4	2	8	4/1.2	40	2	2	1.5	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	EVS Basic Bandwidth (Gbit/s)	Virtualization
c6.xlarge.4	4	16	8/2.4	80	2	3	2	KVM
c6.2xlarge.4	8	32	15/4.5	150	4	4	2.5	KVM
c6.3xlarge.4	12	48	17/7	200	4	6	4	KVM
c6.4xlarge.4	16	64	20/9	280	8	8	5	KVM
c6.6xlarge.4	24	96	25/14	400	8	8	8	KVM
c6.8xlarge.4	32	128	30/18	550	16	8	10	KVM
c6.12xlarge.4	48	192	35/27	750	16	8	15	KVM
c6.16xlarge.4	64	256	40/36	1000	32	8	20	KVM

Table 1-6 C6nl ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
c6nl.large.1	2	2	4/1	32	2	KVM
c6nl.large.2	2	4	4/1	32	2	KVM
c6nl.large.4	2	8	4/1	32	2	KVM
c6nl.xlarge.1	4	4	8/2	64	2	KVM
c6nl.xlarge.2	4	8	8/2	64	2	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
c6nl.xlarge.4	4	16	8/2	64	2	KVM
c6nl.2xlarge.1	8	8	15/4	120	4	KVM
c6nl.2xlarge.2	8	16	15/4	120	4	KVM
c6nl.2xlarge.4	8	32	15/4	120	4	KVM
c6nl.3xlarge.1	12	12	17/6	160	4	KVM
c6nl.3xlarge.2	12	24	17/6	160	4	KVM
c6nl.3xlarge.4	12	48	17/6	160	4	KVM
c6nl.4xlarge.1	16	16	20/8	224	8	KVM
c6nl.4xlarge.2	16	32	20/8	224	8	KVM
c6nl.4xlarge.4	16	64	20/8	224	8	KVM
c6nl.6xlarge.2	24	48	25/12	320	8	KVM
c6nl.6xlarge.4	24	96	25/12	320	8	KVM
c6nl.8xlarge.2	32	64	30/16	440	16	KVM
c6nl.8xlarge.4	32	128	30/16	440	16	KVM
c6nl.12xlarge.2	48	96	35/27	750	16	KVM
c6nl.12xlarge.4	48	192	35/27	750	16	KVM
c6nl.16xlarge.2	64	128	40/32	800	32	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
c6nl.16xlarge.4	64	256	40/32	800	32	KVM

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

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	Virtualization
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
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

Memory-optimized

Table 1-8 M6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	EVS Basic Bandwidth (Gbit/s)	Virtualization
m6.large.8	2	16	4/1.2	40	2	1.5	KVM

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	EVS Basic Bandwidth (Gbit/s)	Virtualization
m6.xlarge.8	4	32	8/2.4	80	2	2	KVM
m6.2xlarge.8	8	64	15/4.5	150	4	2.5	KVM
m6.3xlarge.8	12	96	17/7	200	4	4	KVM
m6.4xlarge.8	16	128	20/9	280	8	5	KVM
m6.6xlarge.8	24	192	25/14	400	8	8	KVM
m6.8xlarge.8	32	256	30/18	550	16	10	KVM
m6.16xlarge.8	64	512	40/36	1000	32	20	KVM

Table 1-9 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

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
m6nl.8xlarge.8	32	256	30/16	440	16	KVM
m6nl.16xlarge.8	64	512	40/32	800	32	KVM

Table 1-10 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

Large-Memory

Table 1-11 E7 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	Virtualization
e7.12xlarge.20	48	960	30/20	550	16	8	1024	KVM
e7.24xlarge.20	96	1920	44/40	1000	32	8	2048	KVM

Disk-intensive

Table 1-12 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
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

Ultra-high I/O

Table 1-13 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
i3.16xlarge.8	64	512	25/25	500	16	8	8 × 1,600 GiB NVMe	KVM

Table 1-14 I3nl ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Local Disks	Virtualization
i3nl.2xlarge.8	8	64	15/4	120	4	4	1 × 1.6 TiB NVMe	KVM
i3nl.4xlarge.8	16	128	20/8	224	8	8	2 × 1.6 TiB NVMe	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Local Disks	Virtualization
i3n.8xlarge.8	32	256	30/16	440	16	8	4 x 1.6 TiB NVMe	KVM
i3n.12xlarge.8	48	384	35/27	750	16	8	6 x 1.6 TiB NVMe	KVM
i3n.16xlarge.8	64	512	40/32	800	32	8	8 x 1.6 TiB NVMe	KVM

Table 1-15 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 Disk	Virtualization
i7n.2xlarge.4	8	32	10/3.4	120	4	4	64	1 x 1.6T NVMe	KVM
i7n.8xlarge.4	32	128	25/13.5	400	8	8	192	4 x 1.6T NVMe	KVM
i7n.12xlarge.4	48	192	30/20	500	16	8	256	6 x 1.6T NVMe	KVM
i7n.16xlarge.4	64	256	35/27	600	16	8	256	8 x 1.6T NVMe	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	Local Disk	Virtualization
i7n.24xlarge.4	96	420	44/40	800	32	8	256	12 × 1.6T NVMe	KVM

GPU-accelerated

Table 1-16 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 Gen3	1 × 32 GiB	KVM
p2s.4xlarge.8	16	128	15/8	100	8	8	2 × V100	PCIe Gen3	2 × 32 GiB	KVM
p2s.8xlarge.8	32	256	25/15	200	16	8	4 × V100	PCIe Gen3	4 × 32 GiB	KVM
p2s.16xlarge.8	64	512	30/30	400	32	8	8 × V100	PCIe Gen3	8 × 32 GiB	KVM

Table 1-17 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
pi2.2xlarge.4	8	32	10/4	50	4	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
pi2.16xlarge.4	64	256	30/30	400	32	8 × T4	8 × 16	N/A	KVM

Table 1-18 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.16xlarge.4	64	256	30/30	400	32	8 × T4	8 × 16	N/A	KVM

AI-accelerated**Table 1-19** 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
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

Table 1-20 Ai1 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
ai1.large.4	2	8	4/1.3	20	1	8	2	2	KVM
ai1.xlarge.4	4	16	6/2	35	2	16	2	3	KVM
ai1.2xlarge.4	8	32	10/4	50	4	32	4	4	KVM

Flavor	vCPUs	Memory (GiB)	Max. / Assured Bandwidth	Max. PPS (10,000)	Ascend 310 Processors	Ascend RAM (GiB)	Max. NIC Queues	Max. NICs	Virtualization
ai1.4xlarge.4	16	64	15/8	100	8	64	8	8	KVM
ai1.8xlarge.4	32	128	25/15	200	16	128	8	8	KVM

1.7.2 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.

S6 ECSs are suitable for applications that require moderate performance generally but occasionally burstable high performance, such as light-workload web servers, enterprise R&D and testing environments, and low- and medium-performance databases. S6 ECS performance is neither restricted by vCPU credits nor billed for additional credits. You can determine the CPU usage and vCPU credits in monitoring details.

Table 1-21 General-purpose ECS features

Series	Compute	Disk Type	Network
S6	<ul style="list-style-type: none"> vCPU to memory ratio: 1:2 or 1:4 Number of vCPUs: 1 to 8 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.5 GHz 	<ul style="list-style-type: none"> Ultra-high I/O General Purpose SSD High I/O Extreme SSD 	<ul style="list-style-type: none"> An ECS with higher specifications has better network performance. Maximum packets per second (PPS): 500,000 Maximum intranet bandwidth: 3 Gbit/s

Specifications

Table 1-22 S6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
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.medium.4	1	4	0.8/0.1	10	1	1	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
s6.2xlarge.4	8	32	3/0.75	50	2	2	KVM

Scenarios

- Web servers, light-workload applications, and R&D and testing environments
- Small- and medium-sized databases, cache servers, and search clusters

1.7.3 Dedicated General-Purpose ECSs

Overview

Compared with general-purpose ECSs, dedicated general-purpose ECSs provide the combinations of vCPUs and memory with larger specifications. In addition, the ECSs use latest-generation network acceleration engines and Data Plane Development Kit (DPDK) to provide higher network performance.

C6s ECSs use second-generation Intel® Xeon® Scalable processors that feature low latency and high performance, stability, and cost-effectiveness. They are suitable for Internet, gaming, and rendering scenarios, especially those that require high computing and networking performance.

C6 and C6nl ECSs use second-generation Intel® Xeon® Scalable processors to provide powerful and stable computing performance. By using 25GE high-speed

intelligent NICs, they offer ultra-high network bandwidth and packets per second (PPS).

C7n ECSs use third-generation Intel Xeon scalable processors with enhanced performance, security, and stability. They can have a maximum number of 96 vCPUs and a memory speed of 3,200 MHz, and provide a secure, trusted cloud environment.

Table 1-23 Dedicated general-purpose ECS features

Series	Compute	Disk Type	Network
C7n	<ul style="list-style-type: none"> vCPU to memory ratio: 1:2 or 1:4 Number of vCPUs: 2 to 96 3rd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.4 GHz 	<ul style="list-style-type: none"> Ultra-high I/O General Purpose SSD High I/O Extreme SSD 	<ul style="list-style-type: none"> Support for IPv6 Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 8,500,000 Maximum intranet bandwidth: 40 Gbit/s
C6s	<ul style="list-style-type: none"> vCPU to memory ratio: 1:2 or 1:4 Number of vCPUs: 2 to 64 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.5 GHz 	<ul style="list-style-type: none"> Ultra-high I/O General Purpose SSD High I/O Extreme SSD 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 8,500,000 Maximum intranet bandwidth: 30 Gbit/s
C6nl	<ul style="list-style-type: none"> vCPU to memory ratio: 1:1, 1:2, or 1:4 Number of vCPUs: 2 to 64 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.5 GHz 	<ul style="list-style-type: none"> Ultra-high I/O General Purpose SSD High I/O Extreme SSD 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 8,000,000 Maximum intranet bandwidth: 40 Gbit/s

Series	Compute	Disk Type	Network
C6	<ul style="list-style-type: none"> vCPU to memory ratio: 1:2 or 1:4 Number of vCPUs: 2 to 64 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 3.0 GHz/3.4 GHz 	<ul style="list-style-type: none"> Ultra-high I/O General Purpose SSD High I/O Extreme SSD 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 10,000,000 Maximum intranet bandwidth: 40 Gbit/s

Specifications

Table 1-24 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
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

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Virtualization
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 1-25 C6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	EVS Basic Bandwidth (Gbit/s)	Virtualization
c6.large.2	2	4	4/1.2	40	2	2	1.5	KVM
c6.xlarge.2	4	8	8/2.4	80	2	3	2	KVM
c6.2xlarge.2	8	16	15/4.5	150	4	4	2.5	KVM
c6.3xlarge.2	12	24	17/7	200	4	6	4	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	EVS Basic Bandwidth (Gbit/s)	Virtualization
c6.4xlarge.2	16	32	20/9	280	8	8	5	KVM
c6.6xlarge.2	24	48	25/14	400	8	8	8	KVM
c6.8xlarge.2	32	64	30/18	550	16	8	10	KVM
c6.12xlarge.2	48	96	35/27	750	16	8	15	KVM
c6.16xlarge.2	64	128	40/36	1000	32	8	20	KVM
c6.large.4	2	8	4/1.2	40	2	2	1.5	KVM
c6.xlarge.4	4	16	8/2.4	80	2	3	2	KVM
c6.2xlarge.4	8	32	15/4.5	150	4	4	2.5	KVM
c6.3xlarge.4	12	48	17/7	200	4	6	4	KVM
c6.4xlarge.4	16	64	20/9	280	8	8	5	KVM
c6.6xlarge.4	24	96	25/14	400	8	8	8	KVM
c6.8xlarge.4	32	128	30/18	550	16	8	10	KVM
c6.12xlarge.4	48	192	35/27	750	16	8	15	KVM
c6.16xlarge.4	64	256	40/36	1000	32	8	20	KVM

Table 1-26 C6nl ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
c6nl.large.1	2	2	4/1	32	2	KVM
c6nl.large.2	2	4	4/1	32	2	KVM
c6nl.large.4	2	8	4/1	32	2	KVM
c6nl.xlarge.1	4	4	8/2	64	2	KVM
c6nl.xlarge.2	4	8	8/2	64	2	KVM
c6nl.xlarge.4	4	16	8/2	64	2	KVM
c6nl.2xlarge.1	8	8	15/4	120	4	KVM
c6nl.2xlarge.2	8	16	15/4	120	4	KVM
c6nl.2xlarge.4	8	32	15/4	120	4	KVM
c6nl.3xlarge.1	12	12	17/6	160	4	KVM
c6nl.3xlarge.2	12	24	17/6	160	4	KVM
c6nl.3xlarge.4	12	48	17/6	160	4	KVM
c6nl.4xlarge.1	16	16	20/8	224	8	KVM
c6nl.4xlarge.2	16	32	20/8	224	8	KVM
c6nl.4xlarge.4	16	64	20/8	224	8	KVM
c6nl.6xlarge.2	24	48	25/12	320	8	KVM
c6nl.6xlarge.4	24	96	25/12	320	8	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Virtualization
c6nl.8xlarge.2	32	64	30/16	440	16	KVM
c6nl.8xlarge.4	32	128	30/16	440	16	KVM
c6nl.12xlarge.2	48	96	35/27	750	16	KVM
c6nl.12xlarge.4	48	192	35/27	750	16	KVM
c6nl.16xlarge.2	64	128	40/32	800	32	KVM
c6nl.16xlarge.4	64	256	40/32	800	32	KVM

Table 1-27 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

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	Virtualization
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
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

Scenarios

Websites and web applications, generalized databases and cache servers, and medium- and heavy-workload enterprise applications with strict requirements on computing and network performance

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

M7n ECSs use third-generation Intel® Xeon® Scalable processors to provide enhanced computing, security, and stability. They 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.

M6 and M6nl ECSs use second-generation Intel® Xeon® Scalable processors with technologies optimized to offer powerful and stable computing performance. Using 25GE high-speed intelligent NICs, M6 ECSs provide a maximum memory size of 512 GiB based on DDR4 for large-memory applications with high requirements on network bandwidth and packets per second (PPS).

Table 1-28 Memory-optimized ECS features

Series	Compute	Disk Type	Network
M7n	<ul style="list-style-type: none">vCPU to memory ratio: 1:8Number of vCPUs: 2 to 963rd Generation Intel® Xeon® Scalable ProcessorBasic/Turbo frequency: 2.6 GHz/3.4 GHz	<ul style="list-style-type: none">Ultra-high I/OGeneral Purpose SSDHigh I/OExtreme SSD	<ul style="list-style-type: none">Support for IPv6Ultra-high PPS throughputAn ECS with higher specifications has better network performance.Maximum PPS: 8,500,000Maximum intranet bandwidth: 40 Gbit/s

Series	Compute	Disk Type	Network
M6	<ul style="list-style-type: none"> vCPU to memory ratio: 1:8 Number of vCPUs: 2 to 64 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 3.0 GHz/3.4 GHz 		<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 10,000,000 Maximum intranet bandwidth: 40 Gbit/s
M6nl	<ul style="list-style-type: none"> vCPU to memory ratio: 1:8 Number of vCPUs: 2 to 64 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 3.0 GHz/3.5 GHz 		<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 8,000,000 Maximum intranet bandwidth: 40 Gbit/s

Specifications

Table 1-29 M6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	EVS Basic Bandwidth (Gbit/s)	Virtualization
m6.large.8	2	16	4/1.2	40	2	1.5	KVM
m6.xlarge.8	4	32	8/2.4	80	2	2	KVM
m6.2xlarge.8	8	64	15/4.5	150	4	2.5	KVM
m6.3xlarge.8	12	96	17/7	200	4	4	KVM
m6.4xlarge.8	16	128	20/9	280	8	5	KVM

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	EVS Basic Bandwidth (Gbit/s)	Virtualization
m6.6xlarge.8	24	192	25/14	400	8	8	KVM
m6.8xlarge.8	32	256	30/18	550	16	10	KVM
m6.16xlarge.8	64	512	40/36	1000	32	20	KVM

Table 1-30 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 1-31 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

Scenarios

- Applications
Memory-optimized ECSs are suitable for applications that process large volumes of data and require a large amount of memory, rapid data switching and processing, and low-latency storage resources.
- Application scenarios
Big data analysis for precision marketing, e-commerce, and IoV, relational databases, NoSQL databases, and memory data analysis

1.7.5 Large-Memory ECSs

Overview

Large-memory ECSs provide high compute, storage, and network performance. They are suitable for applications that require a large amount of memory, rapid data switching and processing, and low latency. Large-memory ECSs provide large memory and high computing, storage, and network performance.

Available now: E7

Table 1-32 Large-memory ECS features

Series	Compute	Disk Type	Network
E7	<ul style="list-style-type: none">• Memory: 960 GiB to 1,920 GiB• Number of vCPUs: 48 to 96• 3rd Generation Intel® Xeon® Scalable Processor• vCPU to memory ratio: 1:20• Basic/Turbo frequency: 2.6 GHz/3.4 GHz	<ul style="list-style-type: none">• Ultra-high I/O• General Purpose SSD• High I/O• Extreme SSD	<ul style="list-style-type: none">• Support for IPv6• Ultra-high PPS throughput• An ECS with higher specifications has better network performance.• Maximum PPS: 10,000,000• Maximum intranet bandwidth: 44 Gbit/s• Single root I/O virtualization (SR-IOV) and Open vSwitch (OVS) work together to provide up to 10 Gbit/s bandwidth.

Large-Memory E7

Overview

Large-memory E7 ECSs use 3rd generation Intel® Xeon® Scalable processors to provide powerful and stable computing performance. By using 25GE high-speed intelligent NICs, E7 ECSs offer ultra-high network bandwidth and PPS.

Notes

For notes about using E7 ECSs, see [Notes](#).

Scenarios

- Online transaction processing (OLTP) and online analytical processing (OLAP) scenarios

- SAP HANA in-memory databases, such as SAP Business Suite on HANA (SoH), SAP S/4HANA (S4H), SAP Business Warehouse on HANA (BWoH), and SAP BW/4HANA (B4H)
- High-performance databases
- Distributed cache
- Big data processing engines and data mining applications

Specifications

Table 1-33 E7 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	Virtualization
e7.12xlarge.20	48	960	30/20	550	16	8	1024	KVM
e7.24xlarge.20	96	1920	44/40	1000	32	8	2048	KVM

Notes

- Large-memory ECSs do not support NIC hot swapping.
- Affected by the memory loading speed, large-memory ECSs may take longer to start.
- The primary and extension NICs of a large-memory ECS can only be used in the scenarios listed in [Table 1-34](#).

Table 1-34 Application scenarios of the NICs of a large-memory ECS

NIC Type	Application Scenario	Remarks
Primary NIC	Vertical layer 3 communication	N/A
Extension NIC	Horizontal layer 2 communication	To improve network performance, you can set the MTU of an extension NIC to 8888 .

1.7.6 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.

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.

Table 1-35 Disk-intensive ECS features

Series	Compute	Disk Type	Network
D3	<ul style="list-style-type: none"> • vCPU to memory ratio: 1:8 or 1:10 • Number of vCPUs: 4 to 56 • Intel® Xeon® Scalable Processor • Basic/Turbo frequency: 3.0 GHz/3.4 GHz 	<ul style="list-style-type: none"> • Ultra-high I/O • General Purpose SSD • High I/O • Extreme SSD 	<ul style="list-style-type: none"> • Ultra-high PPS throughput • An ECS with higher specifications has better network performance. • Maximum PPS: 5,000,000 • Maximum intranet bandwidth: 40 Gbit/s

Specifications

Table 1-36 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

Notes on Using D3 ECSs

- If the host where a D3 ECS resides becomes 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.
- D3 ECSs do not support specifications modification.
- D3 ECSs do not support local disk snapshots or backups.
- D3 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 D3 ECS.
 - Both EVS disks and local disks can be used as data disks of a D3 ECS.
 - A maximum of 60 disks (including VBD, SCSI, and local disks) can be attached to a D3 ECS. Among the 60 disks, the maximum number of SCSI disks is 30, and the VBD disks (including the system disk) is 24. For details, see [Can I Attach Multiple Disks to an ECS?](#)

 NOTE

The maximum number of disks attached to an existing D3 ECS remains unchanged.

- You can modify the **fstab** file to set automatic disk mounting at ECS start.
- The local disk data of a D3 ECS may be lost 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 D3 ECS is deleted, its local disk data will also be automatically deleted, which can take some time. As a result, a D3 ECS takes a longer time than other ECSs to be deleted. Back up the data before deleting such an ECS.
- Do not store service data in local disks for a long time. Instead, store it in EVS disks. To improve data security, use a high availability architecture and back up data in a timely manner.
- Local disks can only be purchased during ECS creation. The quantity and capacity of your local disks are determined according to the specifications of your ECS.

Application 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

1.7.7 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.

Table 1-37 Ultra-high I/O ECS features

Series	Compute	Disk Type	Network
I3nl	<ul style="list-style-type: none"> vCPU to memory ratio: 1:8 Number of vCPUs: 8 to 64 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 3.0 GHz/3.5 GHz 	<ul style="list-style-type: none"> Ultra-high I/O General Purpose SSD High I/O Extreme SSD 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 8,000,000 Maximum intranet bandwidth: 40 Gbit/s
I3	<ul style="list-style-type: none"> vCPU to memory ratio: 1:8 Number of vCPUs: 8 to 64 Intel® Xeon® Scalable Processor Basic/Turbo frequency: 3.0 GHz/3.4 GHz 	<ul style="list-style-type: none"> Ultra-high I/O General Purpose SSD High I/O Extreme SSD 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 5,000,000 Maximum intranet bandwidth: 25 Gbit/s
I7n	<ul style="list-style-type: none"> vCPU to memory ratio: 1:4 Number of vCPUs: 8 to 96 3rd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.5 GHz 	<ul style="list-style-type: none"> Ultra-high I/O General Purpose SSD High I/O Extreme SSD 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 8,000,000 Maximum intranet bandwidth: 40 Gbit/s

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 1-38 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 Disk	Virtu alization
i7n.2xlarge.4	8	32	10/3.4	120	4	4	64	1 × 1.6T NVMe	KVM
i7n.8xlarge.4	32	128	25/13.5	400	8	8	192	4 × 1.6T NVMe	KVM
i7n.12xlarge.4	48	192	30/20	500	16	8	256	6 × 1.6T NVMe	KVM
i7n.16xlarge.4	64	256	35/27	600	16	8	256	8 × 1.6T NVMe	KVM
i7n.24xlarge.4	96	420	44/40	800	32	8	256	12 × 1.6T NVMe	KVM

Table 1-39 I3 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Local Disks	Virtu alization
i3.2xlarge.8	8	64	2.5/2.5	100	4	4	1 × 1,600 GiB NVMe	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Local Disks	Virtualization
i3.xlarge.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
i3.16xlarge.8	64	512	25/25	500	16	8	8 × 1,600 GiB NVMe	KVM

Table 1-40 I3nl ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Local Disks	Virtualization
i3nl.2xlarge.8	8	64	15/4	120	4	4	1 × 1.6 TiB NVMe	KVM
i3nl.4xlarge.8	16	128	20/8	224	8	8	2 × 1.6 TiB NVMe	KVM
i3nl.8xlarge.8	32	256	30/16	440	16	8	4 × 1.6 TiB NVMe	KVM
i3nl.12xlarge.8	48	384	35/27	750	16	8	6 × 1.6 TiB NVMe	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	Local Disks	Virtualization
i3nl.16xlarge.8	64	512	40/32	800	32	8	8 x 1.6 TiB NVMe	KVM

Local Disk Performance

[Table 1-41](#) and [Table 1-42](#) list the IOPS performance of local disks and specifications of a single local disk attached to an I7n ECS.

Table 1-41 IOPS performance of local disks used by I7n ECSs

Flavor	Maximum IOPS for Random 4 KB Read
i7n.2xlarge.4	900,000
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 1-42 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 1-43](#) and [Table 1-44](#) list the IOPS performance of local disks and specifications of a single local disk attached to an I3 ECS.

Table 1-43 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 1-44 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

Table 1-45 and **Table 1-46** list the IOPS performance of local disks and specifications of a single local disk attached to an I3nl ECS.

Table 1-45 IOPS performance of local disks used by I3nl ECSs

Flavor	Maximum IOPS for Random 4 KB Read
i3nl.2xlarge.8	750,000
i3nl.4xlarge.8	1,500,000
i3nl.8xlarge.8	3,000,000
i3nl.12xlarge.8	4,500,000
i3nl.16xlarge.8	6,000,000

Table 1-46 Specifications of a single local disk attached to an I3nl ECS

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

- 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](#).
- 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`.

1.7.8 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 ECS Types

- P series
 - [Computing-accelerated P2s](#)
 - [Inference-accelerated Pi2](#)
 - [Inference-accelerated Pi2nl](#)

Helpful links:

- [Installing a GRID Driver on a GPU-accelerated ECS](#)
- [Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS](#)

Table 1-47 GPU-accelerated ECSs

Type	Series	GPU	CUDA Cores per GPU	Single-GPU Performance	Application
Computing-accelerated	P2s	NVIDIA V100	5,120	<ul style="list-style-type: none">• 14 TFLOPS of single-precision floating-point computing• 7 TFLOPS of double-precision floating-point computing• 112 TFLOPS tensor core for deep learning acceleration	AI deep learning training, scientific computing, computational fluid dynamics, computational finance, seismic analysis, molecular modeling, and genomics.

Type	Series	GPU	CUDA Cores per GPU	Single-GPU Performance	Application
Inference-accelerated	Pi2	NVIDIA T4 (GPU passthrough)	2,560	<ul style="list-style-type: none"> 8.1 TFLOPS of single-precision floating-point computing 130 INT8 TOPS 260 INT4 TOPS 	Machine learning, deep learning, inference training, scientific computing, seismic analysis, computing finance, rendering, multimedia encoding and decoding
Inference-accelerated	Pi2nl	NVIDIA P4 (GPU passthrough)	2,560	<ul style="list-style-type: none"> 8.1 TFLOPS of single-precision floating-point computing 130 INT8 TOPS 260 INT4 TOPS 	Machine learning, deep learning, inference training, scientific computing, seismic analysis, computing finance, rendering, multimedia encoding and decoding

Images Supported by GPU-accelerated ECSs

Table 1-48 Images supported by GPU-accelerated ECSs

Type	Series	Supported Image
Computing-accelerated	P2s	<ul style="list-style-type: none"> Windows Server 2016 Standard 64bit
Inference-accelerated	Pi2	<ul style="list-style-type: none"> CentOS 7.5 64bit Windows Server 2019 Standard 64bit Windows Server 2016 Standard 64bit

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

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 1-49 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.

- By default, P2s ECSs created using a public image have the Tesla driver installed.
- 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](#).

- 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 1-50 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
pi2.2xlarge.e4	8	32	10/4	50	4	1 × T4	1 × 16	N/A	KVM
pi2.4xlarge.e4	16	64	15/8	100	8	2 × T4	2 × 16	N/A	KVM
pi2.8xlarge.e4	32	128	25/15	200	16	4 × T4	4 × 16	N/A	KVM
pi2.16xlarge.e4	64	256	30/30	400	32	8 × T4	8 × 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](#).
- 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 1-51 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.16xlarge.4	64	256	30/30	400	32	8 × T4	8 × 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](#).
- 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.

1.7.9 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

- [Enhanced AI Inference-accelerated Ai1s \(Type I\)](#)
- [AI Inference-accelerated Ai1 \(Type I\)](#)

Table 1-52 AI-accelerated ECS features

Series	Compute	Disk Type	Network
Ai1s	<ul style="list-style-type: none">• vCPU to memory ratio: 1:4• Number of vCPUs: 2 to 32• 2nd Generation Intel® Xeon® Scalable Processor• Basic/Turbo frequency: 2.6 GHz/3.5 GHz	<ul style="list-style-type: none">• Ultra-high I/O• General Purpose SSD• High I/O• Extreme SSD	<ul style="list-style-type: none">• Ultra-high PPS throughput• An ECS with higher specifications has better network performance.• Maximum PPS: 2,000,000• Maximum intranet bandwidth: 25 Gbit/s

Series	Compute	Disk Type	Network
Ai1	<ul style="list-style-type: none">vCPU to memory ratio: 1:4Number of vCPUs: 2 to 322nd Generation Intel® Xeon® Scalable ProcessorBasic/Turbo frequency: 2.6 GHz/3.5 GHz		

Public Images Supported by AI-accelerated ECSs

Table 1-53 Public images

Type	Series	Public Images
Enhanced AI inference-accelerated (type I)	Ai1s	Ubuntu Server 18.04 64bit CentOS 7.6 64bit
AI inference-accelerated (type I)	Ai1	Ubuntu Server 16.04 64bit CentOS 7.4 64bit

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 1-54 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
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

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.

- Ai1s ECSs support automatic recovery when the hosts accommodating such ECSs become faulty.

AI Inference-accelerated Ai1 (Type I)

Overview

Ai1 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. Ai1 ECSs deliver the computing acceleration capabilities of the Ascend 310 processors on the cloud platform.

Ai1 ECSs are based on Atlas 300I accelerator cards. For details, go to [Ascend Community](#).

Ai1 ECSs are ideal for computer vision, speech recognition, and natural language processing to support smart retail, smart campus, robot cloud brain, and safe city scenarios.

Specifications

Table 1-55 Ai1 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
ai1.large.4	2	8	4/1.3	20	1	8	2	2	KVM
ai1.xlarge.4	4	16	6/2	35	2	16	2	3	KVM
ai1.2xlarge.4	8	32	10/4	50	4	32	4	4	KVM
ai1.4xlarge.4	16	64	15/8	100	8	64	8	8	KVM
ai1.8xlarge.4	32	128	25/15	200	16	128	8	8	KVM

Features

Ai1 ECSs have the following features:

- 1:4 ratio of vCPUs to memory
- 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
- 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
- 16-channel HD video decoder (H.264/H.265) based on built-in hardware video codec engine

Notes

1. Ai1 ECSs support the following OSs:
 - Ubuntu Server 16.04 64bit
 - CentOS 7.4 64bit
2. Ai1 ECSs do not support modification of specifications.
3. Ai1 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 **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 Images

1.8.1 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. 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	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.
Private	<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	<p>A shared image is a private image another user has shared with you.</p> <p>For more information, see "Sharing Images" in <i>Image Management Service User Guide</i>.</p>

1.8.2 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, key pair, 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 that are 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 [Installing Cloud-Init](#).

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, 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 which the ECS belongs to.
- 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](#).

1.9 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 [Do I Need to Install a Driver for SCSI EVS Disks?](#)

Helpful Links

- [Which ECSs Can Be Attached with SCSI EVS Disks?](#)

1.10 Network

VPC

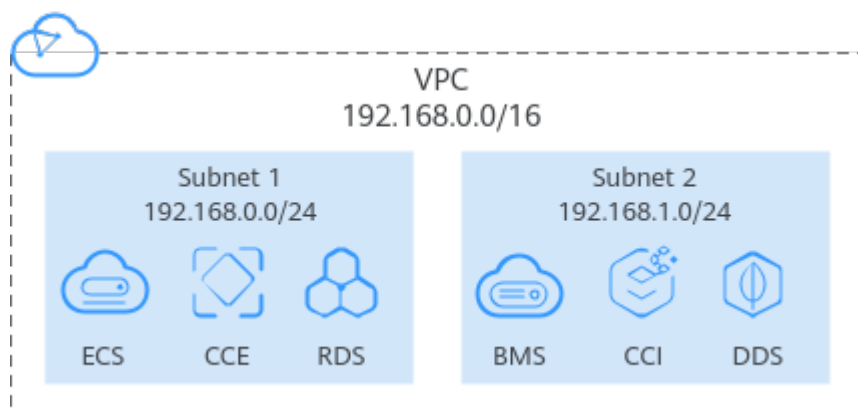
Virtual Private Cloud (VPC) allows you to create customized virtual networks in your logically isolated AZ. Such networks are dedicated zones that are logically isolated, providing secure network environments for your ECSs. You can define security groups, virtual private networks (VPNs), IP address segments, 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.

Figure 1-3 Subnets



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 1-4 Default security group

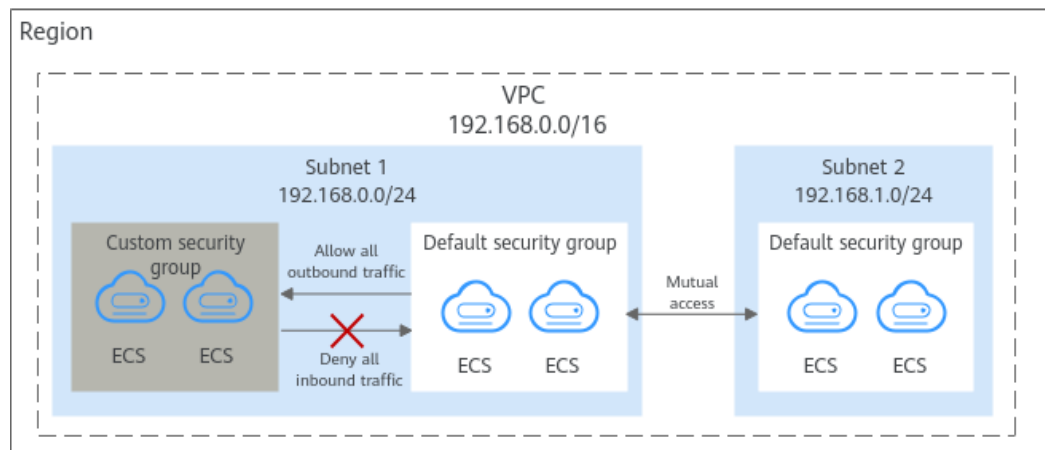


Table 1-56 describes the rules in the default security group.

Table 1-56 Default security group rules

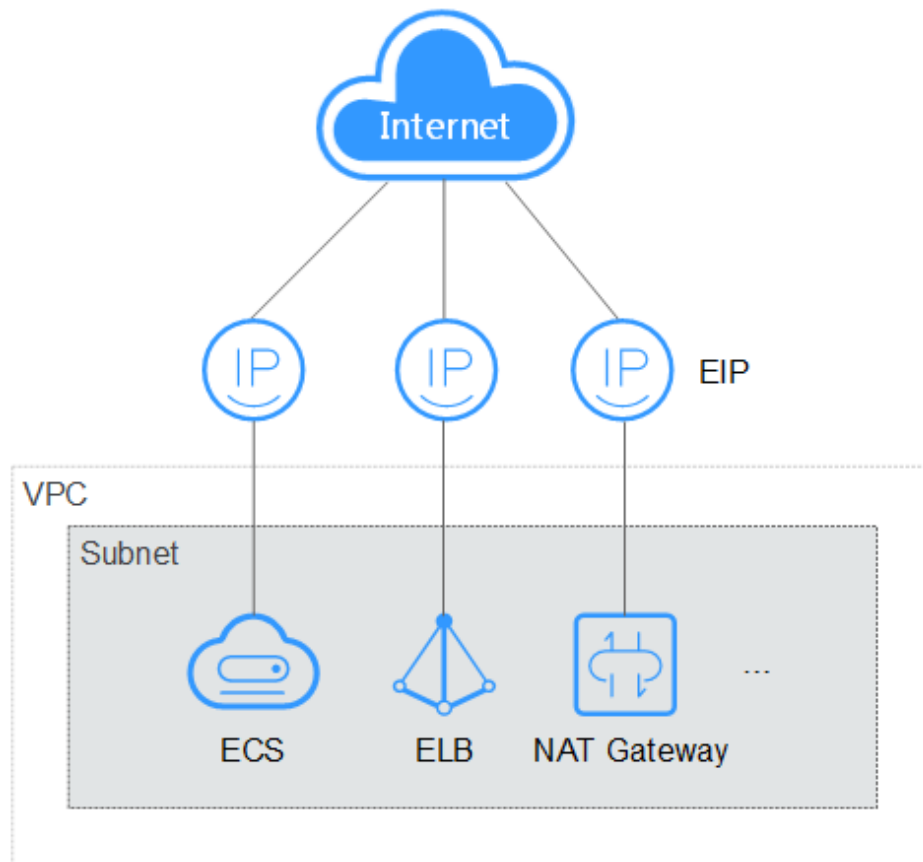
Direction	Protocol	Port/Range	Source/Destination	Description
Outbound	All	All	Destination: 0.0.0.0/0	Allows all outbound traffic.
Inbound	All	All	Source: the current security group (for example, sg-xxxxx)	Allows communications among ECSs within the security group and denies all inbound traffic (incoming data packets).
Inbound	TCP	22	Source: 0.0.0.0/0	Allows all IP addresses to access Linux ECSs over SSH.
Inbound	TCP	3389	Source: 0.0.0.0/0	Allows all IP addresses to access Windows ECSs over RDP.

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, NAT gateways or load balancers.

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

Figure 1-5 Accessing the Internet using an EIP



1.11 Security

1.11.1 Data Protection

1.11.1.1 User Encryption

User encryption allows you to use the encryption feature provided on the cloud platform to encrypt ECS resources, improving data security. User encryption includes image encryption and EVS disk encryption.

Image Encryption

Image encryption supports encrypting private images. When creating an ECS, if you select an encrypted image, the system disk of the created ECS is automatically encrypted, improving data security.

Use either of the following methods to create an encrypted image:

- Use an external image file.
- Use an existing encrypted ECS.

For more information about image encryption, see [Encrypting Images](#).

EVS Disk Encryption

EVS disk encryption supports system disk encryption and data disk encryption.

- When creating an ECS, if you select an encrypted image, the system disk of the created ECS automatically has encryption enabled, and the encryption mode complies with the image encryption mode.
- When creating an ECS, you can encrypt its system disk.
- When creating an ECS, you can encrypt added data disks.

For more information about EVS disk encryption, see [EVS Encryption](#).

Impact on AS

If you use an encrypted ECS to create an Auto Scaling (AS) configuration, the encryption mode of the created AS configuration complies with the ECS encryption mode.

About Keys

The key required for encryption relies on Data Encryption Workshop (DEW). DEW uses a data encryption key (DEK) to encrypt data and uses a customer master key (CMK) to encrypt the DEK.

Figure 1-6 Data encryption process

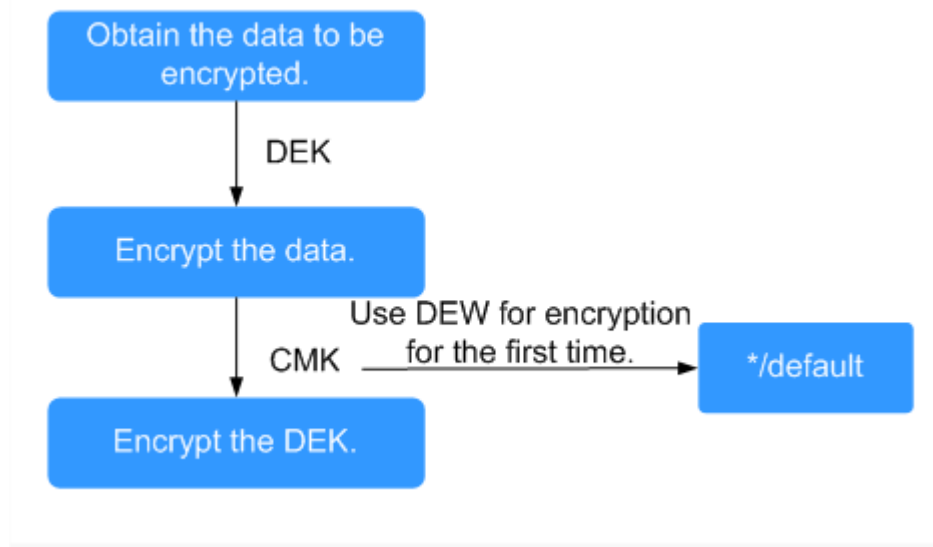


Table 1-57 describes the keys involved in the data encryption process.

Table 1-57 Keys

Name	Description	Function
DEK	An encryption key that is used for encrypting data.	Encrypts specific data.

Name	Description	Function
Custom key	An encryption key created using DEW for encrypting DEKs. A custom key can encrypt multiple DEKs.	Supports CMK disabling and scheduled deletion.
Default key	A master key automatically generated by the system when you use DEW for encryption for the first time. The name extension of a default CMK is /default , for example, evs/default .	<ul style="list-style-type: none">• Supports query of the default key on the DEW console.• Does not support CMK disabling or scheduled deletion.

NOTE

After disabling a CMK or scheduling the deletion of a CMK takes effect, the EVS disk encrypted using this CMK can still be used until the disk is detached from and then attached to an ECS again. During this process, the disk fails to be attached to the ECS because the CMK cannot be obtained, so the EVS disk becomes unavailable.

For details about DEW, see [Data Encryption Workshop User Guide](#).

1.11.2 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 1-58 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.

1.12 Billing

Billing Items

ECSs are billed based on ECS specifications and usage duration.

Table 1-59 Billing items

Billing Item	Pricing Details
ECS	Pricing is based on the ECS type, flavor (including vCPUs and memory), service duration, and the number of purchased ECSs. For details, see the Elastic Cloud Server pricing details.
Image	Public images of the community edition, such as Linux, are free of charge. Other commercial images, such as Windows, generate costs.
EVS disk	EVS disks are mandatory. A system disk is 40 GiB by default. EVS disks can be billed on a pay-per-use basis.
EIP	A public IP address is required only if the ECS accesses the Internet. For details about pricing, see the EIP billing.
Bandwidth	An EIP can be billed by bandwidth or traffic. For further details about pricing, see bandwidth billing.

Billing Modes

An ECS can be billed on a pay-per-use or reserved instance (RI) basis.

- Pay-per-use: The ECS will be billed based on usage frequency and duration. This mode is ideal when you want more flexibility and control on ECS usage.
- RI billing: When the attributes of your pay-per-use ECSs match those of an RI, you can benefit from the RI billing discount.

Table 1-60 describes the differences between the billing modes.

Table 1-60 Billing modes

Billing Mode	Pay-per-Use	RI
Payment Method	Postpaid Billed by service duration.	No upfront payment
Billing Period	Billed by the second and settled by the hour.	Billed by the purchased duration specified in the order.
Billing for Stopped ECS	<ul style="list-style-type: none">• After a common ECS (without local disks or FPGAs attached) is stopped, basic resources (vCPUs, memory, and image) are not billed. Other resources (EVS disks, EIPs, and bandwidth) associated with the ECS will continue to be billed. The vCPU and memory resources of the stopped ECS are released. When the ECS is restarted, the vCPU and memory resources must be requested again. However, if the resources are insufficient, the startup may fail. In such a case, wait several minutes before trying to restart or modify specifications of the ECS.• An ECS with local disks attached, an FPGA-based ECS, or a bare metal ECS will continue to generate costs even after being stopped.	Billed by the purchased duration specified in the order, regardless of whether the ECS is stopped or not.

Billing Mode Change	Can be changed to RIs. You can purchase RIs with the same specifications as the pay-per-use ECSs to change the billing mode without affecting the created ECSs.	Can be changed to pay-per-use. The pay-per-use billing mode takes effect only after the reservation has expired. If early termination is required, early termination fee shall be paid.
Specification Modification	Supported	Not supported
Application Scenarios	This mode is ideal when you want more flexibility and control on ECS usage.	This cost-effective mode is ideal when the duration of ECS usage is predictable. This billing mode is recommended for long-term users.

Billing Examples

In the pay-per-use billing mode, ECSs are billed by the second. The price per second of each type of ECS can be obtained by dividing their hourly price by 3600. Obtain the hourly price on the **Product Pricing Details** page.

For example, if the price of a pay-per-use ECS is \$0.4194 USD/hour, and you purchase such an ECS, the ECS will be billed based on the service duration by second.

- If you use the ECS for 30 minutes, you need to pay for \$0.2097 USD ($0.4194/3600 \times 30 \times 60$).
- If you use the ECS for 1 hour and 30 minutes, you need to pay for \$0.6291 USD ($0.4194/3600 \times 90 \times 60$).

Configuration Changes

- Changing the billing mode
 - Changing from pay-per-use to RI: After you purchase an RI with the same specifications as the pay-per-use ECSs, the billing mode changes to RI without affecting the created ECSs. The new billing mode takes effect in the next month.
 - Changing from RI to pay-per-use: After the billing mode of an ECS is changed from RI to pay-per-use, the new billing mode takes effect only after the reservation has expired. If early termination is required, early termination fee shall be paid.
- Modifying ECS specifications

The specifications of a pay-per-use ECS can be modified, for example, its vCPU and memory specifications can be modified. However, an RI does not support specifications modification.

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 [IAM 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 also need to 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](#).

[Table 1-61](#) lists all the system policies supported by ECS.

Table 1-61 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
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	Full permissions for ECS. This role must be used together with the Tenant Guest role in the same project. 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. 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.	System role	Server Administrator Policy Content

Table 1-62 lists the common operations supported by each system-defined policy of ECS. Select the policies as required.

Table 1-62 Common operations supported by each system-defined policy

Operation	ECS FullAccess	ECS CommonOperations	ECS ReadOnlyAccess
Creating an ECS	Supported	Not supported	Not supported

Operation	ECS FullAccess	ECS CommonOperations	ECS ReadOnlyAccess
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
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
Adding tags to an ECS	Supported	Supported	Not supported
Modifying ECS specifications	Supported	Not supported	Not supported
Querying the ECS flavor list	Supported	Supported	Supported
Querying ECS groups	Supported	Supported	Supported

Helpful Links

- [IAM Service Overview](#)
- [Creating a User and Granting ECS Permissions](#)

- Permissions Policies and Supported Actions in *Elastic Cloud Server API Reference*

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 1-7 shows the relationship between regions and AZs.

Figure 1-7 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

During the creation process, you can select a random AZ. 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, P2s ECSs are released only in AZ1; P3 ECSs are available in AZ2 and have been sold out in AZ1. If you set **AZ** to **Random**, you can view both P2s and P3 ECSs. If you create a P2s ECS, the system automatically allocates it to AZ1. If you create a P3 ECS, the system allocates it to AZ2.

3. Set **Instance Selection**.

- **By Type:** Select a CPU architecture, ECS type, flavor, vCPUs, and memory to suit your service needs.
- **By Scenario:** Select a category and sub-category of your services and then choose appropriate specifications from the displayed specifications list.

4. (Optional) Select a CPU architecture.

This configuration is displayed when you select **By Type** for **Instance Selection**.

- x86

The x86 CPU architecture uses Complex Instruction Set Computing (CISC). Each instruction can be used to execute low-level hardware operations. There are a large number of instructions and the length of each instruction is different. Therefore, executing such an instruction is complex and time-consuming.

5. 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.

Latest generation shows the types and flavors of newly released ECSs, and **All generations** show the types and flavors of all ECSs provided by the cloud platform.

 **NOTE**

- Before selecting an ECS type, learn about the introduction and notes on each type of ECSs. For details, see [ECS Types](#).
- When purchasing an ECS, sold-out vCPU and memory resources cannot be selected. You can select **Hide sold-out specifications** to hide specifications that have been sold out.
- **Local Disk**: specifies the local storage of the physical server where the ECS is deployed. Only Hard Disk Driver (HDD) disks are supported. If the ECS of the selected type (such as **Disk-intensive**) uses local disks, the system automatically attaches the local disks to the ECS and displays the information of the local disks.
For example, if **Local Disk** is **3x1800 GiB (HDD)**, three HDDs are attached to the ECS and the capacity of each HDD is 1800 GiB.

6. 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.

For instructions about how to create a private image, see [Creating a Private Image](#).

You can also select an encrypted image. For more information about encrypted images, see *Image Management Service User Guide*.

 NOTE

- If you use a full-ECS image to create an ECS, the EVS disks associated with the full-ECS image do not support the function of creating disks using a data disk image.
- 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, and the encryption attributes of the system and data disks of the created ECSs are the same as those of the system and data disks specified in the full-ECS image. Additionally, the SCSI, data encryption, 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, and the encryption attributes of the system and data disks of the created ECSs are the same as those of the system and data disks specified in the full-ECS image. Additionally, the SCSI, data encryption, and sharing attribute settings of data disks can be modified during ECS creation.
- An ISO image created from an ISO file cannot be used to create ECSs. You need to install an OS and drivers on the ECS and use the ECS to create a system disk image first.
- ECSs created from ISO images cannot be directly used because some functions like disk attachment are restricted.
- 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 Enable NIC Multi-Queue for an Image?](#)

- Shared image

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

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

Disks are classified as EVS disks and DSS disks based on whether the storage resources used by the disks are dedicated. DSS disks allow you to use dedicated storage resources.

- If you have requested for a storage pool on the DSS page, click the **DSS** tab and create disks in the obtained storage pool.
- If you have not requested for a dedicated storage pool, click the **Disks** tab and create EVS disks that use public storage resources.


 NOTE

- When you use DSS resources to create a disk, the disk type must be the same as that of the requested storage pool. For example, both are of high I/O type.
- For more information about DSS, see *Dedicated Distributed Storage Service*.

- System disk

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

- If the image based on which an ECS is created is not encrypted, the system disk of the ECS is not encrypted. If the image based on which an ECS is created is encrypted, the system disk of the ECS is automatically encrypted. For details, see [\(Optional\) Encryption-related parameters](#).

- **Encryption:** indicates that the system disk is encrypted if you select this option. For details, see [\(Optional\) Encryption-related parameters](#).
- Data disk
You can create multiple data disks for an ECS and enable required functions for each data disk. During the creation process, you can add a maximum of 23 data disks for each ECS and customize the disk size as needed.
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 [EVS Disks](#).
 - **Share:** indicates that the EVS disk is sharable if you select this option. Such an EVS disk can be attached to multiple ECSs.
 - **Encryption:** indicates that the data disk is encrypted if you select this option. For details, see [\(Optional\) Encryption-related parameters](#).
- (Optional) Encryption-related parameters
To enable encryption, click **Create Agency** to assign KMS access permissions to EVS. If you have rights granting permission, assign the KMS access permissions to EVS. If you do not have the permission, contact the user having the security administrator rights to assign the KMS access permissions. For details, see [Can All Users Use the Encryption Feature?](#)
 - **Encryption:** indicates that the EVS disk has been encrypted.
 - **Create Agency:** assigns KMS access permissions to EVS to obtain KMS keys. After the permissions are assigned, follow-up operations do not require assigning permissions again.
 - **Agency Name:** set to **EVSAccessKMS**, which means that permissions have been assigned to EVS to obtain KMS keys for encrypting or decrypting EVS disks.
 - **KMS Key Name:** specifies the name of the key used by the encrypted EVS disk. You can select an existing key, or click **Create KMS Key** and create a new one on the KMS console. The default value is **evs/default**.
 - **KMS Key ID:** specifies the ID of the key used by the encrypted data disk.

8. 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 which the ECS belongs to.
 - 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: 80
- Source: 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: 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 **Billed By.**

This parameter is mandatory when **EIP** is set to **Auto assign**. If you select **By bandwidth** or **By traffic**, the system will allocate a dedicated bandwidth for you, and the bandwidth is dedicated for one EIP.

- **By bandwidth:** The billing will be based on the duration for which the bandwidth is used.
- **By traffic:** The billing will be based on the total traffic irrespective of the duration for which the bandwidth is used.
- **Shared bandwidth:** The bandwidth can be used by multiple EIPs.

 **NOTE**

- A bandwidth can be shared among a limited number of EIPs. If the number of EIPs cannot meet service requirements, switch to a higher shared bandwidth or apply for expanding the EIP quota of the existing bandwidth.

6. Set **Bandwidth Size.**

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

7. Set **Release Option.**

If you select **Release with ECS**, when you delete the ECS, the assigned EIP will also be deleted together with the ECS.

8. Click **Next: Configure Advanced Settings.**

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 a 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 numbering of the new ECS names continues from the highest existing value. For example, if the ECS with the highest 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 over again 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 name of a Windows ECS can contain a maximum of 15 characters and must be unique. Otherwise, some Windows applications cannot be used.

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

NOTE

The naming rules of hostnames comply with [RFC 952](#) and [RFC 1123](#).

When you set the ECS name and hostname, you are advised to use letters (a-z), digits (0-9), and hyphens (-) to prevent unknown issues. In the ECS:

- Underscores (`_`) will be converted to hyphens (`-`).
- A combination of a hyphen and underscore (`-_`) will be converted to a hyphen (`-`).
- Periods (`.`), hyphens (`-`), underscores (`_`), or Chinese characters at the beginning of the name will be ignored.
- If periods (`.`) and Chinese characters are not at the beginning, they and the content following them will be ignored.

2. Set **Login Mode**.

Key pair authentication is more secure than password authentication. If you select **Password**, ensure that the password meets complexity requirements listed in [Table 2-1](#) to prevent malicious attacks.

– Key pair

A key pair is used for ECS login authentication. You can select an existing key pair, or click **Create Key Pair** and create a desired one.

NOTE

If you use an existing key pair, make sure that you have saved the key file locally. Otherwise, logging in to the ECS will fail.

– 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.• 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.

3. 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.

4. 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

5. 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.

- **As text:** allows you to enter the user data in the text box.
- **As file:** enables the text to automatically inject a script file or other files into a specified directory on an ECS when you create the ECS.

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](#).

– Tag

This configuration is optional. You can tag an ECS to facilitate identification and management. You can add up to 10 tags to an ECS.

 NOTE

Tags added during ECS creation will also be added to the created EIP and EVS disks (including the system disk and data disks) of the ECS. If the ECS uses an existing EIP, the tags will not be added to the EIP.

After creating the ECS, you can view the tags on the pages providing details about the ECS, EIP, and EVS disks.

For details, see [Overview](#).

– 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*.

If you have created an agency in IAM, you can select the agency from the drop-down list and obtain specified operation permissions. For more information about agencies, see [Account Delegation](#).

6. 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 **Enterprise Project**.
This function is provided for enterprise users.
An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is **default**.
Select an enterprise project from the drop-down list. For more details, see [Enterprise Management User Guide](#).
3. Set the number of ECSs to be created.
4. Confirm the configuration and click **Apply Now**.

Follow-up Procedure

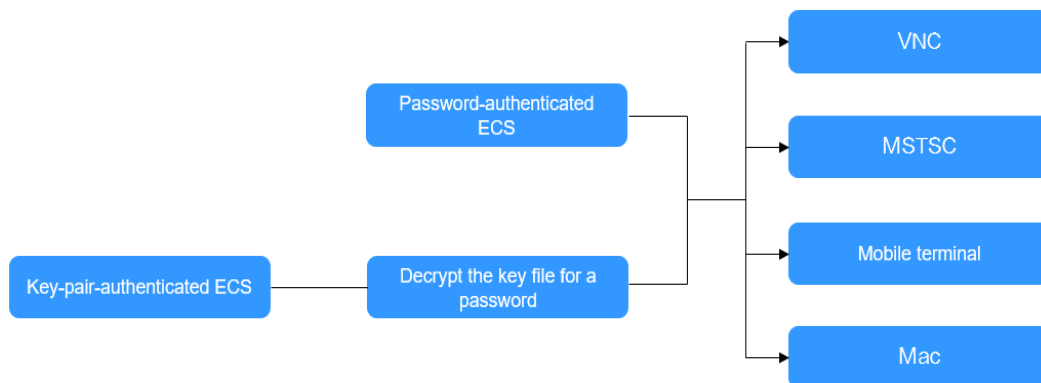
- After an ECS with automatic recovery enabled is created, you can check whether this function has been enabled by viewing **Failures**. For details, see .

2.2 Logging In to an ECS

Logging In to a Windows ECS

[Figure 2-1](#) shows the modes in which a Windows ECS can be logged in.

Figure 2-1 Windows ECS login modes



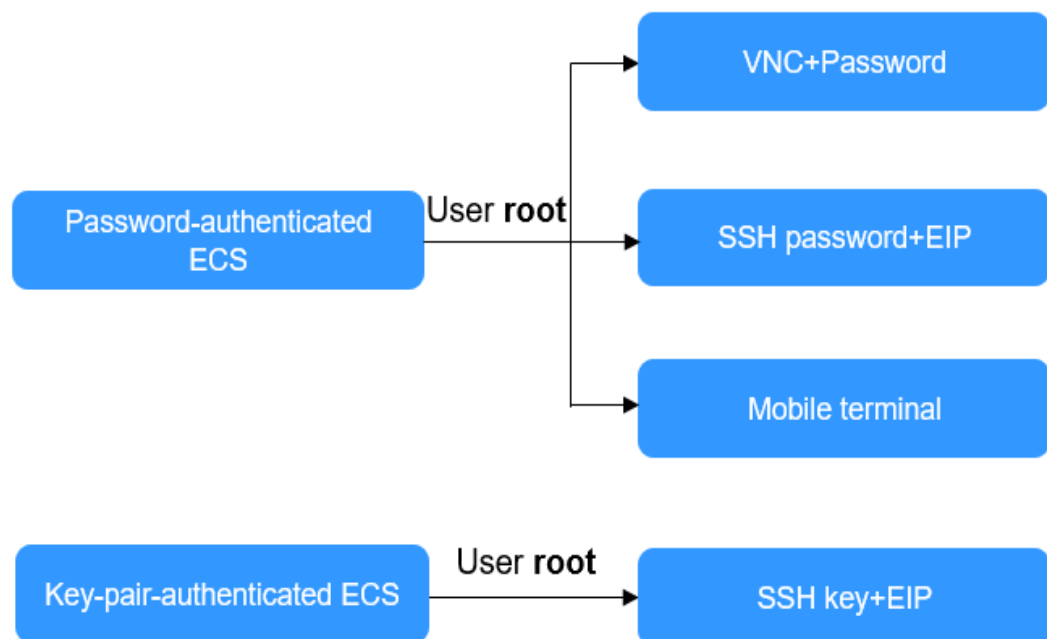
1. (Optional) Retrieve your password from the key file.
To log in to a key-pair-authenticated ECS, use the password obtaining function provided by the management console to decrypt the key file used during ECS creation to obtain a password.
2. Select a login method and log in to the ECS.
 - Using VNC on the management console
The login username is **Administrator**.
For details, see [Logging In to a Windows ECS Using VNC](#).

- Using MSTSC
The login username is **Administrator**, and the ECS must have an EIP bound.
For details, see [Logging In to a Windows ECS Using MSTSC](#).
- From a mobile terminal
The login username is **Administrator**, and the ECS must have an EIP bound.
For details, see [Logging In to a Windows ECS from a Mobile Terminal](#).
- From a Mac
The login username is **Administrator**, and the ECS must have an EIP bound.
For details, see [Logging In to a Windows ECS from a macOS Server](#).

Logging In to a Linux ECS

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

Figure 2-2 Linux ECS login modes



- To log in to a password-authenticated ECS for the first time, use one of the following methods:
 - Using VNC on the management console with login username **root**
For details, see [Logging In to a Linux ECS Using VNC](#).
 - Using an SSH password
The login username is **root**, and the ECS must have an EIP bound.
For details, see [Logging In to a Linux ECS Using an SSH Password](#).
 - From a mobile terminal

The login username is **root**, and the ECS must have an EIP bound.

For details, see [Logging In to a Linux ECS from a Mobile Terminal](#).

- To log in to a key-pair-authenticated ECS for the first time, use a tool, such as PuTTY or XShell, and the desired SSH key as user **root**. Ensure that the ECS has an EIP bound.

For instructions about how to log in to a Linux ECS using an SSH key, see [Logging In to a Linux ECS Using an SSH Key Pair](#).

NOTE

If you want to log in to an ECS using VNC provided on the management console, log in to the ECS using an SSH key, configure the login password, and use the password for login.

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*.
 - For how to log in to a BMS, see the *Bare Metal 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 2-3](#) is displayed, the new disk is offline. Go to [Step 3](#).
- If [Figure 2-6](#) is displayed, the **Initialize Disk** window is prompted. Go to [Step 5](#).

Figure 2-3 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 2-4 Online the disk




NOTE

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

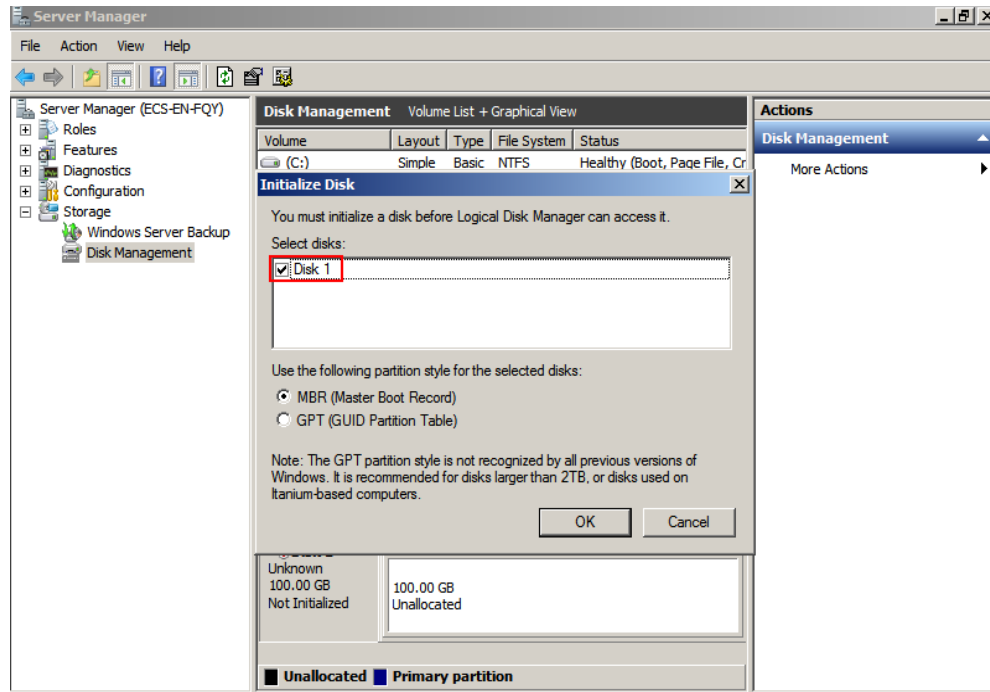
- Step 4** After making the disk online, the status of **Disk 1** changes from **Offline** to **Not Initialized**. Right-click the disk status and choose **Initialize Disk** from the shortcut menu.

Figure 2-5 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 2-6 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 greater 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 2-7 New Simple Volume



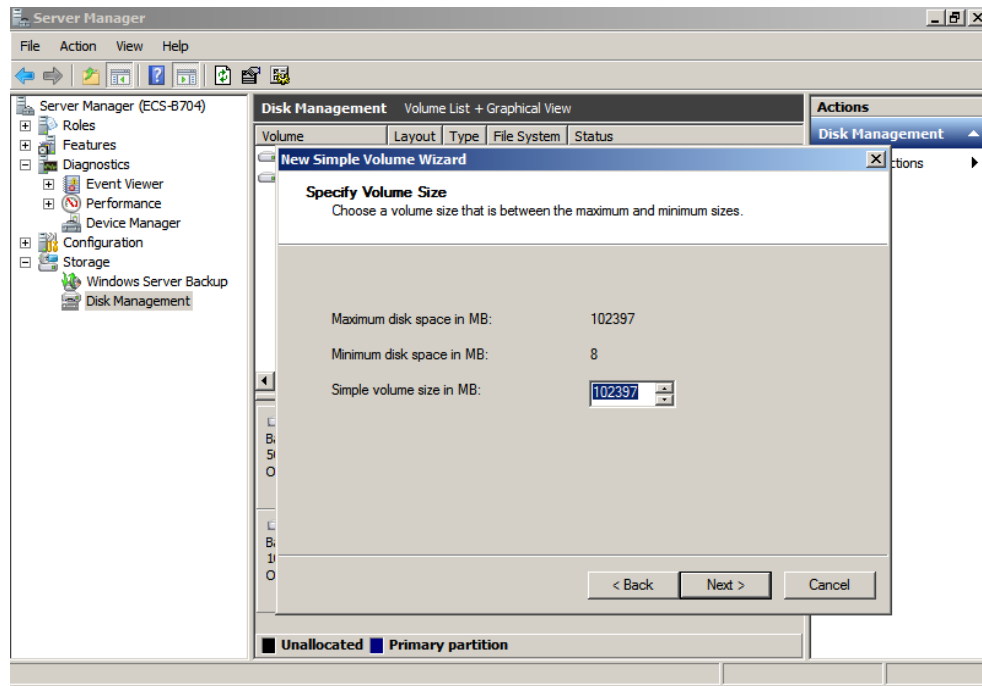
Step 7 On the displayed **New Simple Volume Wizard** window, click **Next**.

Figure 2-8 New Simple Volume Wizard



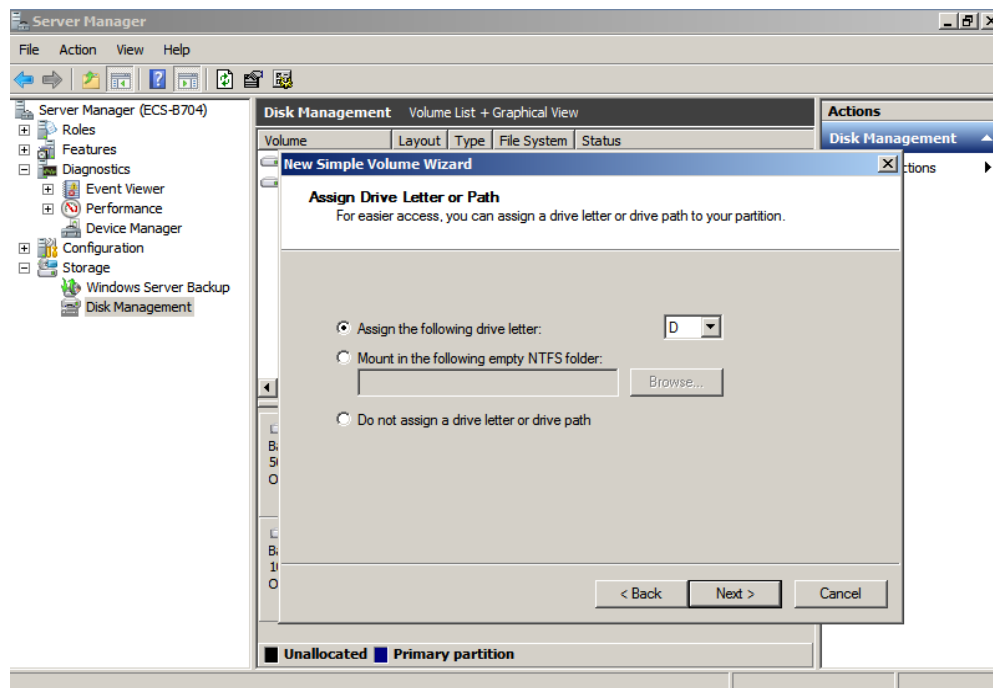
Step 8 Specify the volume size and click **Next**. The default value is the maximum size.

Figure 2-9 Specify Volume Size



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

Figure 2-10 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 2-11 Format Partition





Figure 2-12 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 2-13 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*.
 - For how to log in to a BMS, see the *Bare Metal 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 2-14 Server Manager



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

The **Computer Management** window is displayed.


Figure 2-15 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 2-16 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 2-17 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 greater 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 2-18 New Simple Volume Wizard



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

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


Figure 2-19 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.

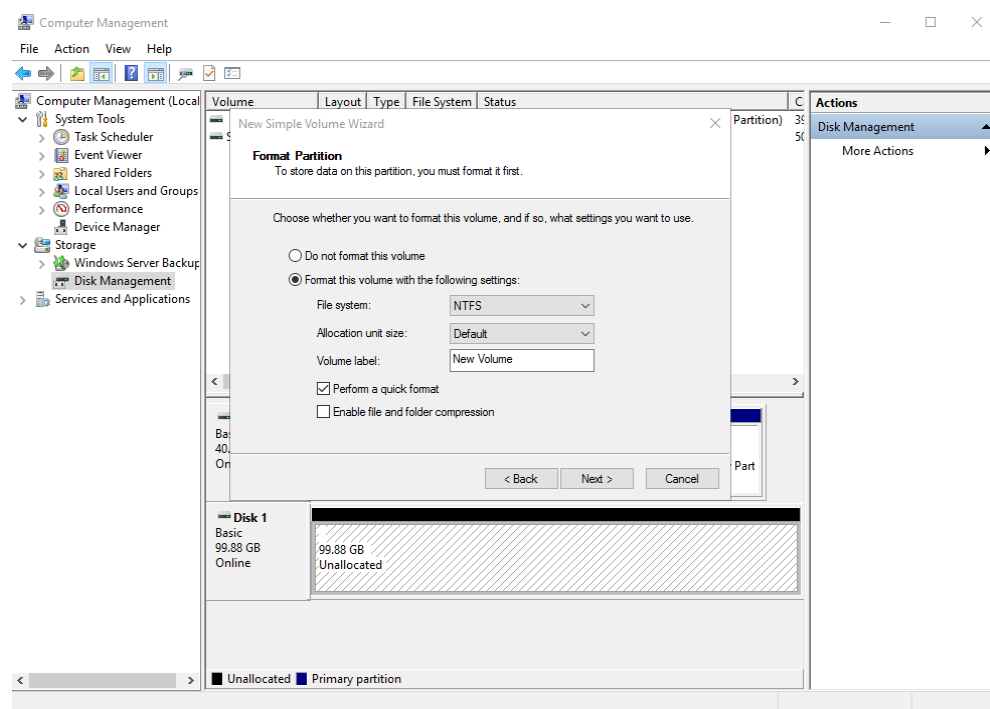
Figure 2-20 Assign Drive Letter or Path



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 2-21 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 2-22 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 2-23 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 2-24 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*.
 - For how to log in to a BMS, see the *Bare Metal 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 `ext4` 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*.
 - For how to log in to a BMS, see the *Bare Metal 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 greater 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:

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 server startup. You can configure the `fstab` file of a server 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 server stop or start. This can even prevent your 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

2.3.6 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*.
 - For how to log in to a BMS, see the *Bare Metal 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 2-25 Server Manager (Windows Server 2008)

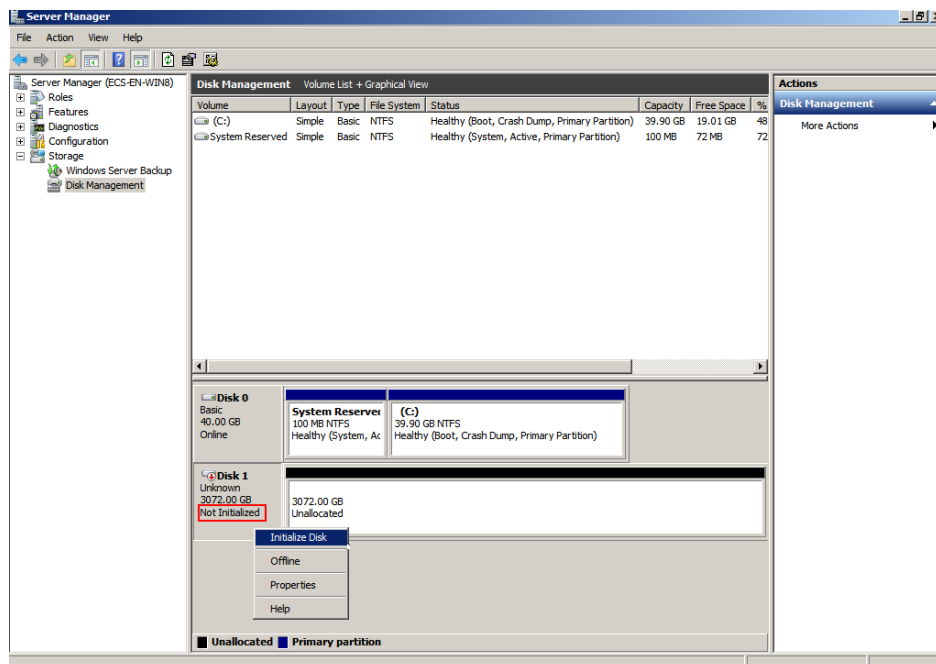


Step 3 Disks are listed in the right pane. If the new disk is offline, bring it online before initializing it.

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

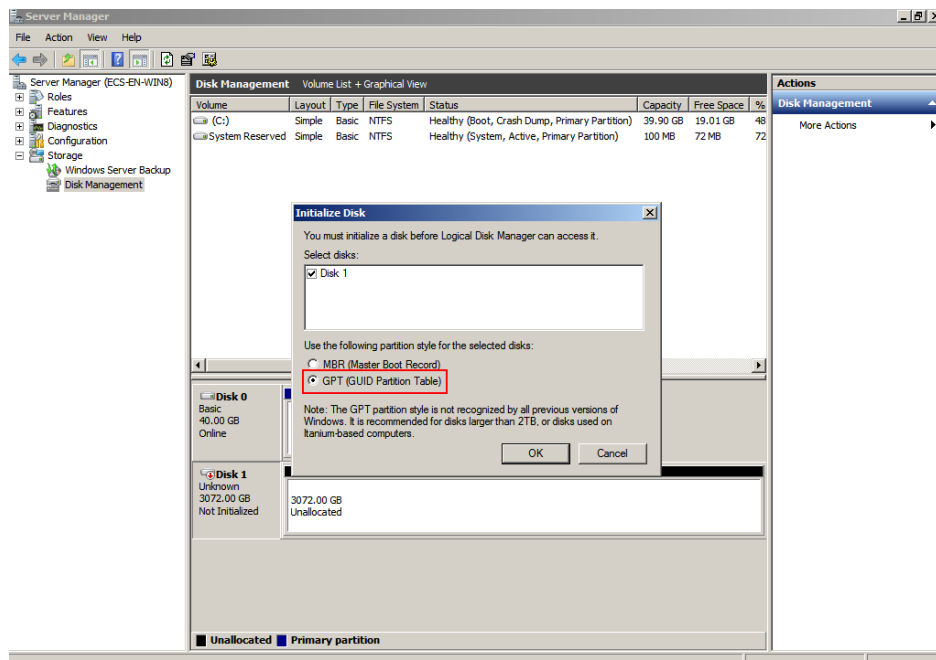
When the status of Disk 1 changes from **Offline** to **Not Initialized**, the disk has been brought online.

Figure 2-26 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 2-27 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 2-28 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 greater 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 2-29 New Simple Volume Wizard (Windows Server 2008)



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


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

Figure 2-30 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 2-31 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 2-32 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 2-33 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 2-34 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 2-35 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*.
 - For how to log in to a BMS, see the *Bare Metal Server User Guide*.

Procedure




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

Figure 2-36 Server Manager (Windows Server 2012)




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

Figure 2-37 Computer Management window (Windows Server 2012)



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

Figure 2-38 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 2-39 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 2-40 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 2-41 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 greater 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 2-42 New Simple Volume Wizard (Windows Server 2012)



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

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


Figure 2-43 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 2-44 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 2-45 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 2-46 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 2-47 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 2-48 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*.
- For how to log in to a BMS, see the *Bare Metal 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 greater 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 server startup. You can configure the **fstab** file of a server 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 server stop or start. This can even prevent your 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

3 Instances

3.1 Selecting an ECS Billing Mode

3.1.1 Pay-per-Use Billing

Concept

Pay-per-use billing is a postpaid billing mode in which an ECS will be billed based on usage frequency and duration. ECSs are billed by second. The system generates a bill every hour based on the usage duration and deducts the billed amount from the account balance. A pay-per-use ECS can be provisioned and deleted at any time.

NOTE

For a stopped pay-per-use ECS, the startup may fail due to insufficient resources. Please wait for several minutes before attempting another restart or changing the ECS specifications.

Billing Examples

In the pay-per-use billing mode, ECSs are billed by the second. The price per second of each type of ECS can be obtained by dividing their hourly price by 3600. Obtain the hourly price on the **Product Pricing Details** page.

For example, if you purchase a pay-per-use ECS priced \$0.68 USD/hour, the ECS will be billed based on the usage duration by the second.

- If you use the ECS for 30 minutes, you need to pay for \$0.34 USD ($0.68/3600 \times 30 \times 60$).
- If you use the ECS for 1 hour and 30 minutes, you need to pay for \$1.02 USD ($0.68/3600 \times 90 \times 60$).

Resources Supporting Pay-per-Use Billing

Resources billed on a pay-per-use basis include:

- ECSs (vCPUs and memory)
- EVS disks purchased together with a yearly/monthly ECS
- Bandwidth purchased with a pay-per-use ECS

3.1.2 Reserved Instances

3.1.2.1 Reserved Instance Overview

Concept

A reserved instance (RI) is not an actual instance, but a billing discount that can be applied to the use of pay-per-use ECSs in your account. When the attributes of your pay-per-use ECSs **match** those of an RI, the RI billing benefit automatically applies to your ECSs. The combination of RIs and pay-per-use billing fully utilizes the flexibility of pay-per-use resources at lower costs.

NOTE

- A purchased RI is billed, regardless of whether it is used or not.

Table 3-1 ECS billing modes

Billing Mode	What It Is	How to Use
RI	A billing discount applied to pay-per-use ECSs.	When the attributes of your pay-per-use ECSs match those of an RI, the RI billing benefit automatically applies to your ECSs.
Pay-per-use	ECS billed based on usage frequency and duration. Such an ECS can be created or deleted at any time.	A basic computing unit that consists of vCPUs, memory, OS, and EVS disks. After purchasing such an ECS, you can use it on the cloud.

- For instructions about how to purchase an RI, see [Enabling and Purchasing a Reserved Instance](#).
- For instructions about how to modify an RI, see [Modifying RI Attributes](#).

What Is Attribute Mapping Between an RI and a Pay-per-Use ECS?

A regional RI is purchased within a region and without an AZ specified.

- Attribute mapping of a regional RI: indicates whether the region, OS type, ECS series, and vCPU/memory ratio of a pay-per-use ECS are the same as those specified in a regional RI.

Application Scenarios

If your ECSs are to be used in a short term, it is a good practice to use pay-per-use rates. If you plan to use ECSs for one or three years, it is a good practice to use RIs. RIs offer discounts for pay-per-use ECSs with matched attributes.

Working Rules

For example, you have a running pay-per-use ECS in your account. After you purchase an RI that matches the attributes of this ECS, the billing benefit of the RI is automatically applied to your ECS when the RI takes effect. A purchased RI takes effect at the next hour.

Table 3-2 lists RI attributes. You can purchase your desired RIs based on these attributes.

Table 3-2 RI attributes

Parameter	Description
Region	<ul style="list-style-type: none">Regional RI: indicates an RI purchased in a region, without an AZ specified. Capacity reservations are not supported for regional RIs.
Flavor	<ul style="list-style-type: none">When purchasing a regional RI, ensure that the ECS series and vCPU/memory ratio specified in the RI are the same as those of the target pay-per-use ECS. <p>NOTE After an RI is purchased, its flavor cannot be directly changed, but you can split or combine it. For details, see Modifying RI Attributes.</p>
OS	The OS of the ECS to be bought, which must match the OS specified in your RI. For example, if you want to use a Linux RI, select a Linux public or private image when purchasing an ECS.
Term	The service duration of an RI. A year is defined as 31,536,000 seconds (365 days).
Offering Class	Standard: Certain attributes, such as the ECS size can be modified during the term. However, the ECS type cannot be changed.
Payment Option	No upfront

3.1.2.2 Enabling and Purchasing a Reserved Instance

A reserved instance (RI) is not an actual instance, but a billing discount that can be applied to pay-per-use ECSs in your account. When the attributes of your pay-per-use ECSs match those of an RI, the RI's discount rate automatically applies to your ECSs.

RIs are suitable for scenarios where the resource usage duration can be predicted. Billing automatically applies your RI's discounted rate when attributes of your ECS usage match attributes of an RI.

- For more information about RIs, see [Reserved Instance Overview](#).
- For instructions about how to modify an RI, see [Modifying RI Attributes](#).

Notes and Constraints

- The quota for the number of RIs that you can purchase in the current region is displayed in the upper left area of the **Reserved Instance** page. The quota for the number of RIs that can be purchased by a user in each region is 20.
- The quota for the number of RIs is automatically reset every month.
- The remaining quota for the number of RIs (Remaining quota = Total quota – Used quota) is reduced only after more RIs are purchased. It will not be changed if RIs are modified, split, combined, or unsubscribed.

Enabling RIs

Before purchasing an RI, contact customer service to apply for the required permissions.

Purchasing an RI

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. In the navigation pane on the left, choose **Reserved Instance**.
4. Click **Buy RI**.
The **Buy RI** page is displayed.
5. Confirm the region.
If the RIs in the selected region do not meet your requirements, select another region.
6. Select an RI type.
The cloud platform provides various RI types for you to choose from based on your application scenarios.
7. Filter for RI specifications.
Set flavor, OS, term, offering class, and payment option to search for the target RI specifications.

[Table 3-3](#) shows specifications parameters.

Table 3-3 RI attributes

Parameter	Description
Region	<ul style="list-style-type: none">• Regional RI: indicates an RI purchased in a region, without an AZ specified. Capacity reservations are not supported for regional RIs.

Parameter	Description
Flavor	<ul style="list-style-type: none">When purchasing a regional RI, ensure that the ECS series and vCPU/memory ratio specified in the RI are the same as those of the target pay-per-use ECS. <p>NOTE After an RI is purchased, its flavor cannot be directly changed, but you can split or combine it. For details, see Modifying RI Attributes.</p>
OS	The OS of the ECS to be bought, which must match the OS specified in your RI. For example, if you want to use a Linux RI, select a Linux public or private image when purchasing an ECS.
Term	The service duration of an RI. A year is defined as 31,536,000 seconds (365 days).
Offering Class	Standard: Certain attributes, such as the ECS size can be modified during the term. However, the ECS type cannot be changed.
Payment Option	No upfront

8. Select specifications.

The cloud platform provides various RI types for you to choose from based on your application scenarios. On the **Buy RI** page, view released RI types and specifications.

Effective Rate: amortized hourly costs of the RI, which is equivalent to the total cost (including any upfront payment) of the RI over the entire term divided by the total number of hours over the entire term. (Effective rate = Total cost of the RI/Entire term of the RI)

Upfront Price: fee that needs to be paid before you purchase an RI.

Hourly Rate: amortized hourly costs of the RI, which is equivalent to the difference between the total cost of the RI and the upfront payment divided by the total number of hours over the entire term (Hourly rate = Total cost of the RI – Upfront payment/Entire term of the RI)

9. Specify an RI name.

The name can be customized. It can contain 1 to 128 characters, which can only be letters, digits, underscores (_), and hyphens (-).

10. Set the number of RIs to be purchased.

- **Quantity:** The system displays the number of RIs that you can purchase.
- **Total Normalization Factors:** measures the ECS size flexibility. The value is determined based on the specifications of the RI to be purchased.

11. Click **Next**.

To learn more about the price, click **Pricing details**.

12. On the page for you to confirm RI specifications, view details and submit the request.

After verifying the configurations and price, click **Submit** and pay for the order as prompted.

- Return to the RI list as prompted and view the purchased RI.

Follow-up Operations

- Purchase a pay-per-use ECS that matches an RI.**

Locate the target RI and click **Buy ECS** in the **Operation** column. The system automatically switches to the page for purchasing ECSs, and the specifications of the ECSs selected by default are the same as those specified in the RI.


NOTE

- If the OS of the target ECS does not match the OS specified in the RI, or the target ECS is not billed on a pay-per-use basis, the RI cannot be used. When the attributes of the ECS match those of the RI, including the ECS series and vCPU/memory ratio, the ECS automatically benefits from the billing discount offered by the RI.
- Check the usage of RIs.**

On the **Reserved Instance** page, click the name of the target RI. In the lower part of the RI details page, view the usage of the RI.

As shown in **Figure 3-1**, the horizontal coordinate indicates the number of in-service days, and the vertical coordinate indicates the usage of RI on the current day. Click the chart to view the RI usage in the selected time period on the current day.

Figure 3-1 Viewing RI usage



3.1.2.3 Modifying RI Attributes

Scenarios

If an RI type cannot meet your computing requirements, you can modify the RI attributes and then apply it to your pay-per-use ECSs.

You can modify the scope and ECS size of a standard RI.

- For more information about RIs, see [Reserved Instance Overview](#).
- For instructions about how to purchase an RI, see [Enabling and Purchasing a Reserved Instance](#).

Constraints

- RIs can be combined only when their attributes, including the OS, payment option, offering class, term, expiration time, region, ECS series, vCPU/memory ratio, and discount are the same.
- The total normalization factors must be the same before and after the modification.
- A maximum of five RIs can be modified in a batch.
- One RI can be split to multiple ones, and multiple RIs can only be combined into one.

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**. On the displayed console, choose **Reserved Instance** from the left navigation pane.
3. On the **Reserved Instance** page, select the target RI and click **Modify RI** in the upper left corner of the list.
4. Modify the RI attributes as required.

Table 3-4 Common operations for modifying an RI

Allowed Operation	Description
Splitting an RI or combining RIs	For example, there are six s3.xlarge.2 RIs in an account, and an s3.xlarge.2 RI has a normalization factor of 4. Then, the six s3.xlarge.2 RIs are equivalent to 24 normalization factors. Then, these RIs can be combined into three s3.2xlarge.2 RIs or split to 24 s3.medium.2 RIs. Just ensure that the splitting or combination matches to the total normalization factor.

NOTICE

Total normalization factors are the number of RIs multiplied by the normalization factor of such an RI. The total normalization factors must be the same before and after the modification.

For example, there are six s3.large.4 RIs with the total normalization factors of 12 (6 x 2) before the modification. These RIs can be split to two s3.xlarge.4 RIs and four s3.medium.4 RIs. After the modification, the total normalization factors are still 12 (2 x 4 + 4 x 1).

5. Verify the modified RI attributes and click **Submit**.

3.2 Creating an ECS

3.2.1 Creating the Same ECS

Scenarios

If you have created an ECS and want to create more with the same configurations, it is a good practice to click **Create Same ECS** on the ECS console.

Notes

Large-memory ECSs and ECSs created using full-ECS images do not support "Create Same ECS".

Procedure









1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Computing**, choose **Elastic Cloud Server**.
4. Select the target ECS and choose **More > Create Same ECS** in the **Operation** column.

Figure 3-2 Creating the same ECS

IP Address	Enterprise Project	Tag	Operation
 (Private IP)	default	--	Remote Login More ▾
 (Private IP)	default	--	<div style="border: 1px solid #ccc; padding: 5px;">Create Same ECS Start Stop Restart Reset Password Modify Specifications Delete Manage Image/Disk/Backup ▶ Manage Network ▶ Migrate ECS</div>
 (Private IP)	default	--	
 (Private IP)	default	--	
 (Private IP)	default	--	
 (Private IP)	default	--	
 (Private IP)	default	--	

The system switches to the ECS creation page and automatically copies the parameter settings of the selected ECS.

 **NOTE**

For security purposes, you must manually configure some of the settings for the new ECSs, including:

- Manually add data disks if the quantity of data disks needed exceeds 10.
 - Manually add NICs if the quantity of NICs needed exceeds 5.
 - Manually add security groups if the quantity of security groups needed exceeds 5.
 - Select a new data disk image if the disks of the source ECS are created using a data disk image.
 - Add disks as needed if the source ECS is created from a full-ECS image and only has the disks included in this image.
 - Select **Encryption** if the disks of the source ECS have been encrypted.
 - Configure the functions in **Advanced Settings**.
 - Configure **EIP** if required (because it is not required by default).
5. Adjust the parameter settings of the created ECSs as needed, confirm the settings, and click **Submit**.

3.3 Logging In to a Windows ECS

3.3.1 Login Overview (Windows)

Constraints

- Only a running ECS can be logged in.
- The username for logging in to a Windows ECS is **Administrator**.
- If the login password is forgotten, reset the password by referring to [Resetting the Password for Logging In to a Windows ECS](#). Resetting the password will interrupt the applications running on the ECS. Exercise caution when performing this operation.
- If an ECS uses key pair authentication, use the password obtaining function available on the management console to decrypt the private key used during ECS creation to obtain a password.
- 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.
- If you log in to a GPU-accelerated ECS using MSTSC, GPU acceleration will fail. This is because MSTSC replaces the WDDM GPU driver with a non-accelerated remote desktop display driver. In such a case, you must log in to the ECS using other methods, such as VNC. If the remote login function available on the management console fails to meet your service requirements, you must install a suitable remote login tool, such as TightVNC, on the ECS.
To download TightVNC, log in at <https://www.tightvnc.com/download.php>.

Login Modes

You can choose from a variety of login modes based on your local OS type.

Table 3-5 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 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 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 Logging In to a Windows ECS from a macOS Server .	
	Mobile terminal	Install a remote connection tool, for example, Microsoft Remote Desktop. For details, see Logging In to a Windows ECS from a Mobile Terminal .	
	Windows	Through the management console. For details, see Logging In to a Windows ECS Using VNC .	No EIP is required.

Helpful Links

- [Login Password Resetting](#)
- [Remote Logins](#)

3.3.2 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.

Prerequisites

If an ECS uses key pair authentication, make sure that the key file has been used to resolve the login password before logging in to the ECS.

Logging In to a Windows ECS



1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Computing**, click **Elastic Cloud Server**.
4. 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.
5. In the **Operation** column of the target ECS, click **Remote Login**.
6. (Optional) When the system displays "Press CTRL+ALT+DELETE to log on", click **Ctrl+Alt+Del** in the upper part of the remote login page to log in to the ECS.

Figure 3-3 Ctrl+Alt+Del



7. Enter the ECS password as prompted.

Helpful Links

- [Login Password Resetting](#)
- [Remote Logins](#)

3.3.3 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](#).
An EIP is not required if you log in to an ECS through an intranet using MSTSC, for example, through VPN or Direct Connect.
- Access to port 3389 is allowed in the inbound direction of the security group which the ECS belongs to. For details, see [Configuring Security Group Rules](#).

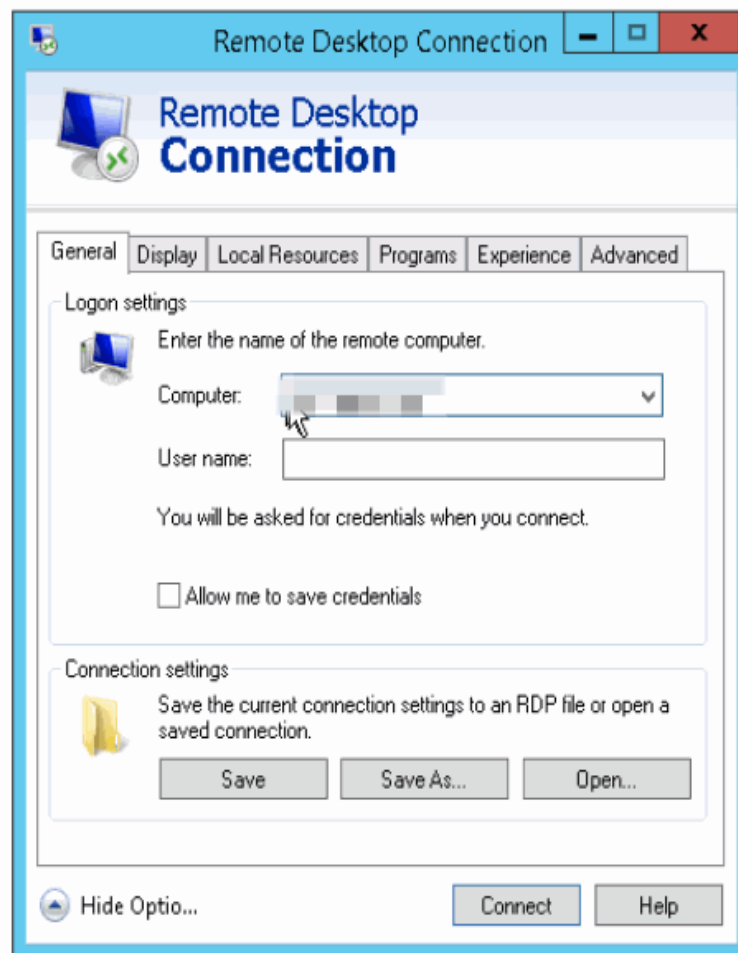
- The network connection between the login tool and the target ECS is normal. For example, the default port 3389 is not blocked by the firewall.
- Remote Desktop Protocol (RDP) needs to be enabled on the target ECS. For ECSs created using public images, RDP has been enabled by default. For instructions 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.

The following uses Windows Server 2012 ECS as an example.


Figure 3-4 Logging in to an ECS using MSTSC



For details, see the following procedure:

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 3-5 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 3-6 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 3-7 Clipboard




To copy files from the local server to your ECS, click **More** and select your desired disks.

Figure 3-8 Drives



6. (Optional) Click the **Display** tab and then adjust the size of the remote desktop.

Figure 3-9 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 [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 3-10 System Properties

3. Click the **Remote** tab and select **Allow remote connections to this computer**.
4. Click **OK**.

Helpful Links

- [Login Password Resetting](#)
- [Remote Logins](#)

3.3.4 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.
An EIP is not required if you log in to an ECS through an intranet using MSTSC, for example, through VPN or Direct Connect.

- Access to port 3389 is allowed in the inbound direction of the security group which the ECS belongs to.
- 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 3-6 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.

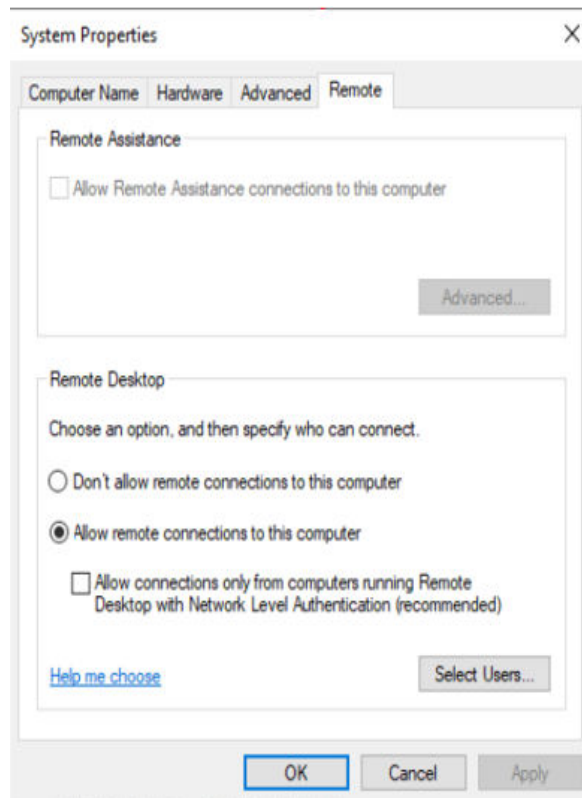
Enabling RDP

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 [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 3-11 System Properties

3. Click the **Remote** tab and select **Allow remote connections to this computer**.
4. Click **OK**.

3.3.5 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. If you have forgotten the password, reset the password by referring to [Resetting the Password for Logging In to a Windows ECS](#).

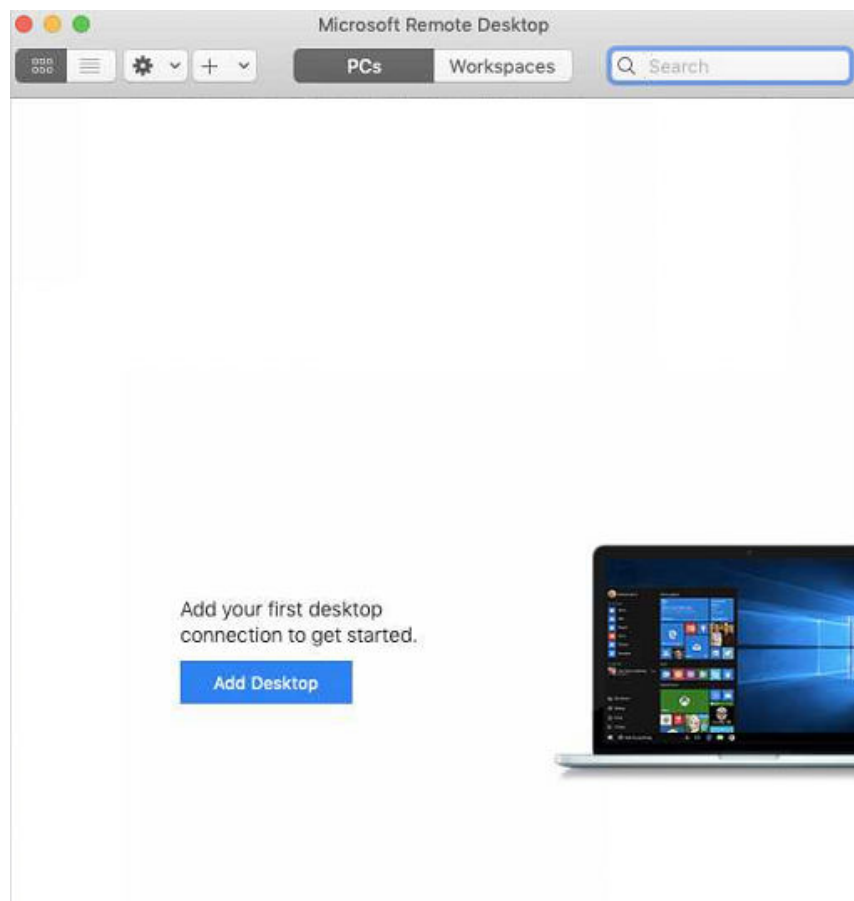
- You have bound an EIP to the ECS. For details, see [Binding an EIP](#).
- Access to port 3389 is allowed in the inbound direction of the security group which the ECS belongs to. For details, see [Configuring Security Group Rules](#).
- The remote access tool supported by Mac, such as Microsoft Remote Desktop for Mac has been installed.

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 3-12 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 3-13 Add user account

Add a User Account

Username:

Password:

Show password

Friendly name:

Figure 3-14 Add PC

Add PC

PC name:

User account:

General | Display | Devices & Audio | Folders

Friendly name:

Group:

Gateway:

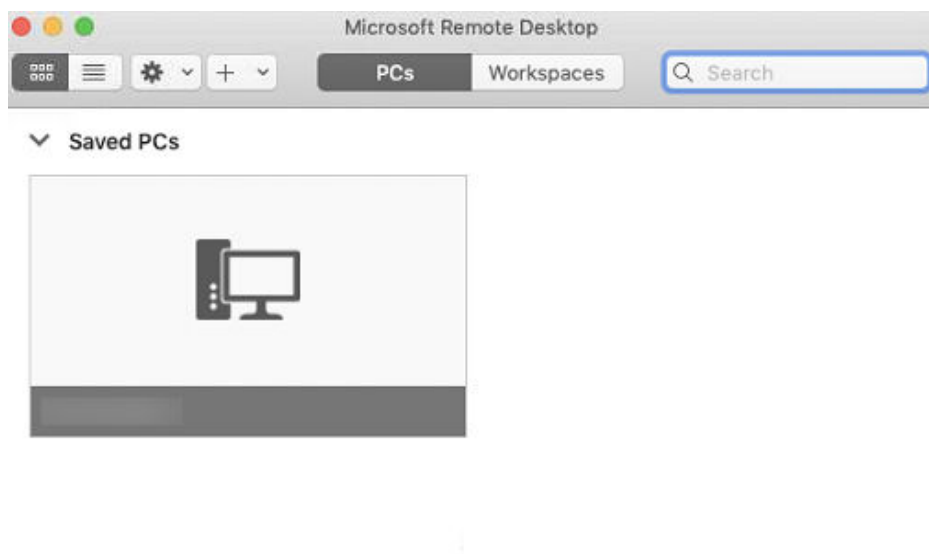
Bypass for local addresses

Reconnect if the connection is dropped

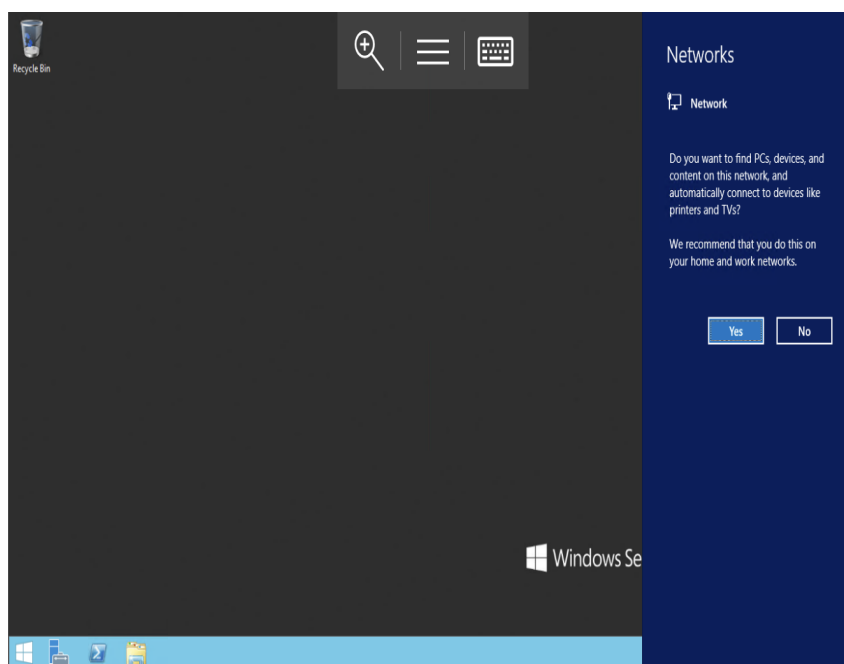
Connect to an admin session

Swap mouse buttons

4. On the **Remote Desktop** page, double-click the icon of the target Windows ECS.

Figure 3-15 Double-click for login

5. Confirm the information and click **Continue**.
You have logged in to the Windows ECS.

Figure 3-16 Successful login

3.3.6 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. If you have forgotten the password, reset the password by referring to [Resetting the Password for Logging In to a Windows 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 which the ECS belongs to. For details, see [Configuring Security Group Rules](#).
- Microsoft Remote Desktop has been installed on the mobile terminal.

Procedure


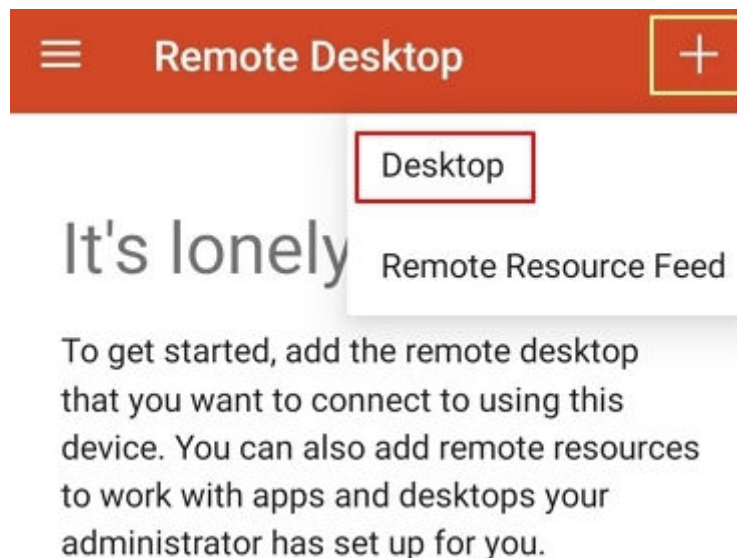
1. Start the Microsoft Remote Desktop client.
2. In the upper right corner of the **Remote Desktop** page, tap  and select **Desktop**.

Figure 3-17 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 3-18 Setting the login information

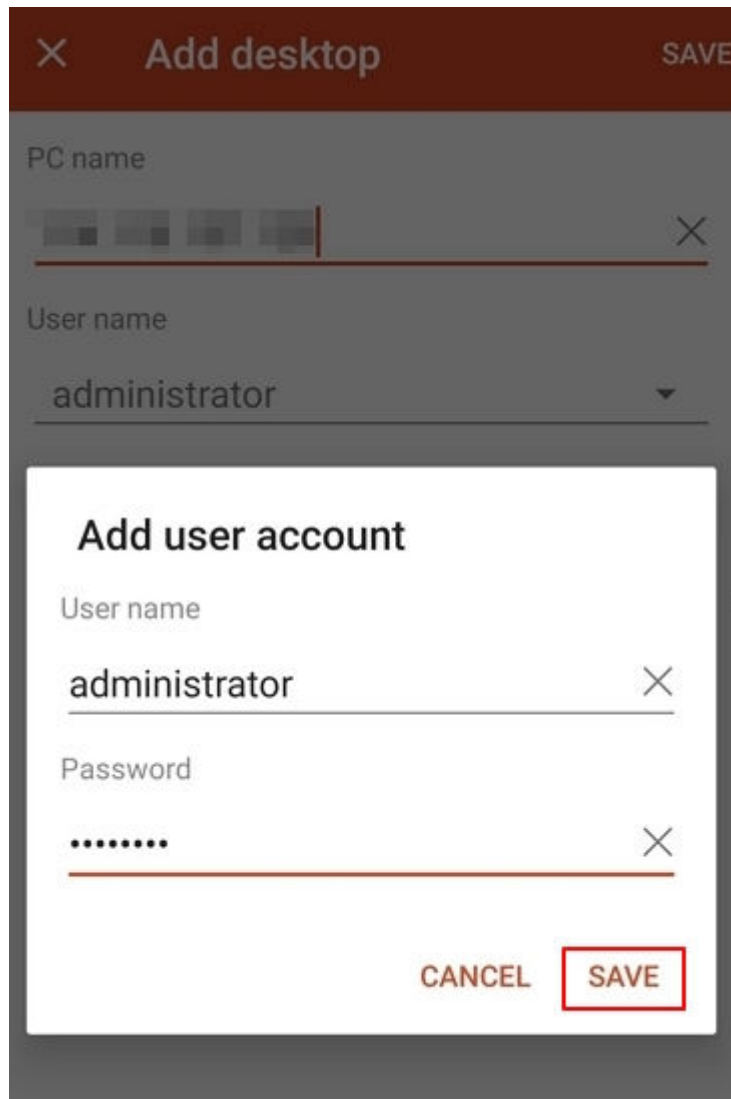

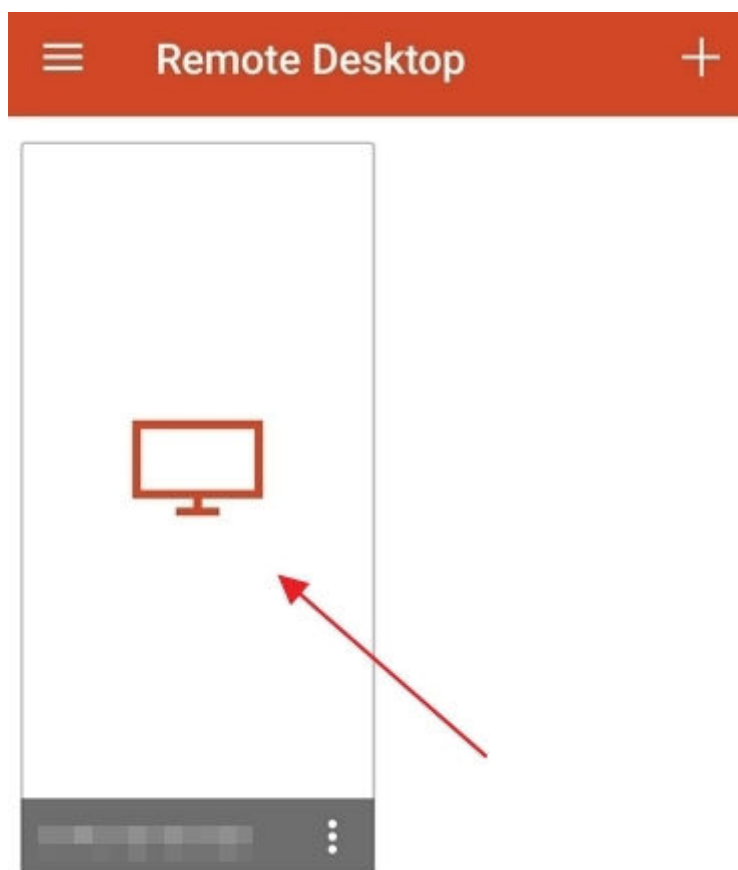


Figure 3-19 Saving the settings



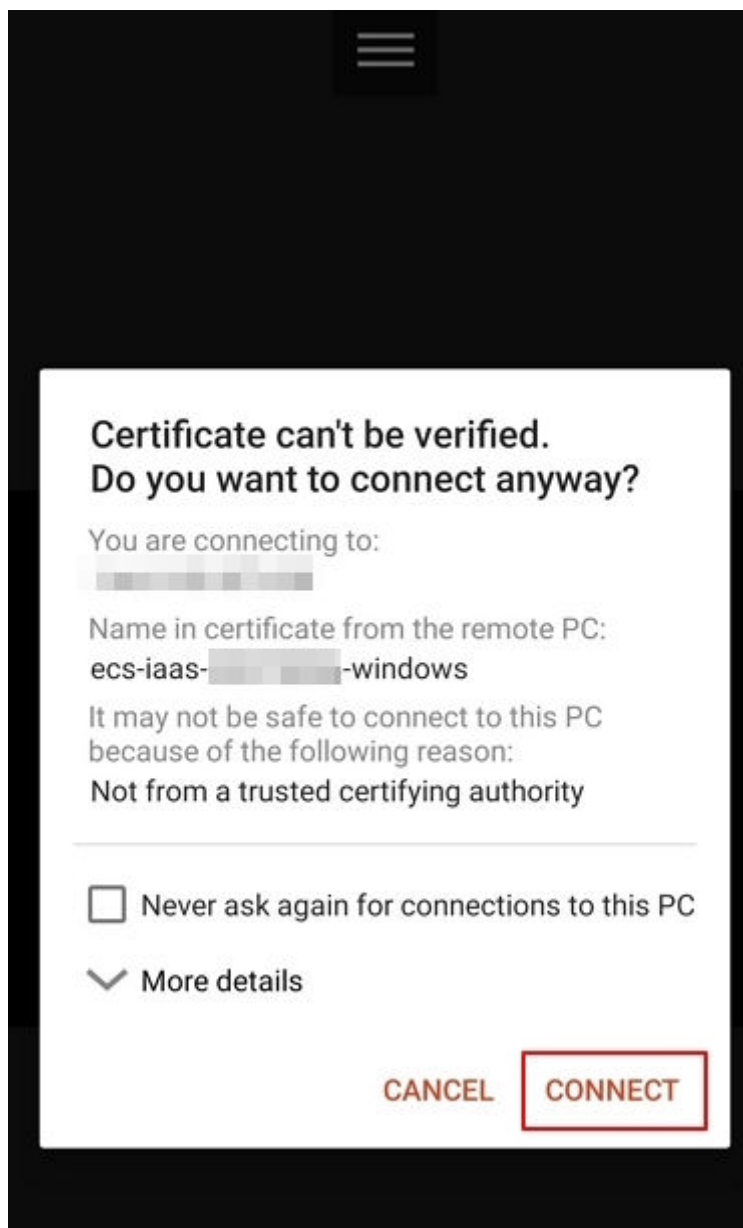
4. On the **Remote Desktop** page, tap the icon of the target Windows ECS.

Figure 3-20 Logging in to the Windows ECS



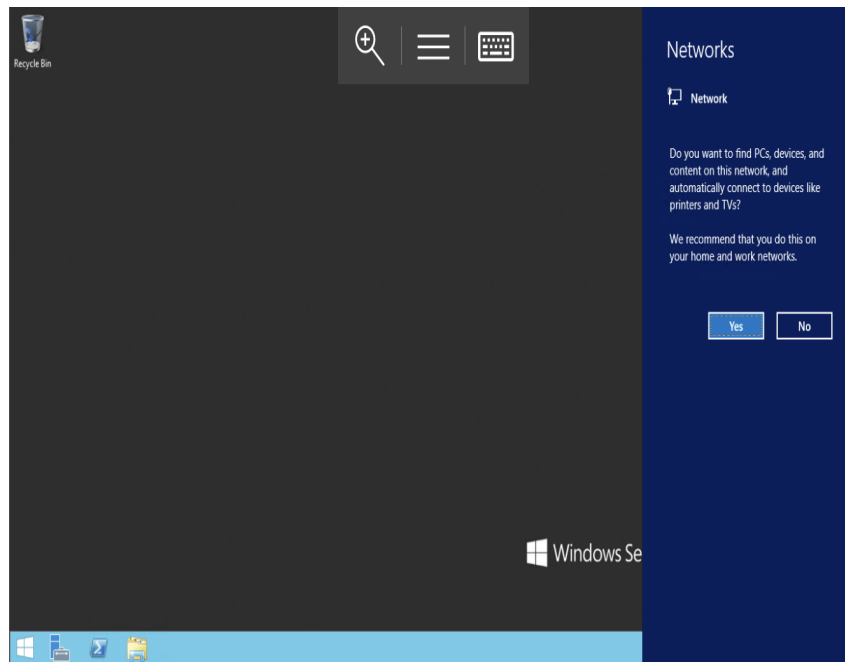
5. Confirm the information and tap **CONNECT**.

Figure 3-21 CONNECT



You have logged in to the Windows ECS.

Figure 3-22 Successful login



3.4 Logging In to a Linux ECS

3.4.1 Login Overview (Linux)

Constraints

- Only a running ECS can be logged in.
- The username for logging in to a Linux ECS is **root**.
- If the login password is forgotten, reset the password by referring to [Resetting the Password for Logging In to a Linux ECS](#). Resetting the password will interrupt the applications running on the ECS. Exercise caution when performing this operation.

Login Modes

You can choose from a variety of login modes based on your local OS type.

Table 3-7 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• Key-pair-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• Key-pair-authenticated: Logging In to a Linux ECS from a Local Linux Server	
	Mobile terminal	Use an SSH client tool, such as Termius or JuiceSSH, to log in to the ECS. Logging In to a Linux ECS from a Mobile Terminal	
	macOS	Use the terminal included in the macOS. Logging In to a Linux ECS from a macOS Server	
	Windows	Use the remote login function available on the management console. For details, see Logging In to a Linux ECS Using VNC .	No EIP is required.

Helpful Links

- [Application Scenarios for Using Passwords](#)
- [Why Can't I Log In to My Linux ECS?](#)

3.4.2 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](#).


NOTE

Before using remote login (VNC) provided on the management console to log in to a Linux ECS authenticated using a key pair, log in to the ECS [using an SSH key](#) and set a login password.

Prerequisites

You have used an SSH key to log in to the Linux ECS authenticated using a key pair and set a login password.

Procedure


1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Computing**, click **Elastic Cloud Server**.
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 **Ctrl+Alt+Del** 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.

6. Enter the ECS password as prompted.


Figure 3-23 Username (root as an example) and password



Follow-up Procedure

Local commands can be copied to an ECS. To do so, perform the following operations:

1. Log in to the ECS using VNC.
2. Click **Paste & Send** in the top area of the page.

Figure 3-24 Paste & Send

3. Press **Ctrl+C** to copy data from the local computer.
4. Press **Ctrl+V** to paste the local data to the **Paste & Send** window.
5. Click **Send**.
Send the copied data to the CLI.

NOTE

There is a low probability that data is lost when you use Paste & Send 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 Paste & Send function.

3.4.3 Logging In to a Linux ECS Using an SSH Key Pair

Scenarios

This section describes how to use an SSH key pair to remotely log in to a Linux ECS from a Windows and a Linux server, respectively.

Prerequisites

- You have obtained the private key file used for creating the ECS. For details about how to create a key pair, see [\(Recommended\) Creating a Key Pair on the Management Console](#).
- You have bound an EIP to the ECS. For details, see [Viewing ECS Details \(List View\)](#).
- You have configured the inbound rules of the security group. 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

You have two methods to log in to a Linux ECS from a local Windows server.

Method 1: Use PuTTY to log in to the ECS.

The following operations use PuTTY as an example. Before using PuTTY to log in, make sure that the private key file has been converted to .ppk format.

1. Check whether the private key file has been converted to .ppk format.
 - If yes, go to step 7.
 - If no, go to step 2.

2. Visit the following website and download PuTTY and PuTTYgen:

<https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html>

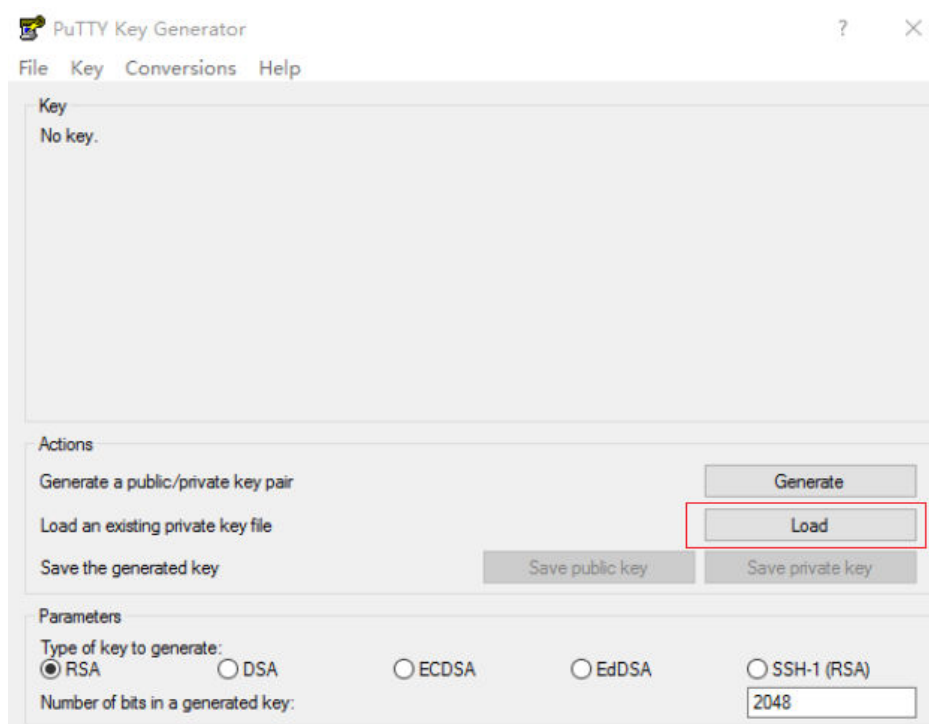
NOTE

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.

3. Run PuTTYgen.
4. In the **Actions** pane, click **Load** and import the private key file that you stored during ECS creation.

Ensure that the format of **All files (*.*)** is selected.

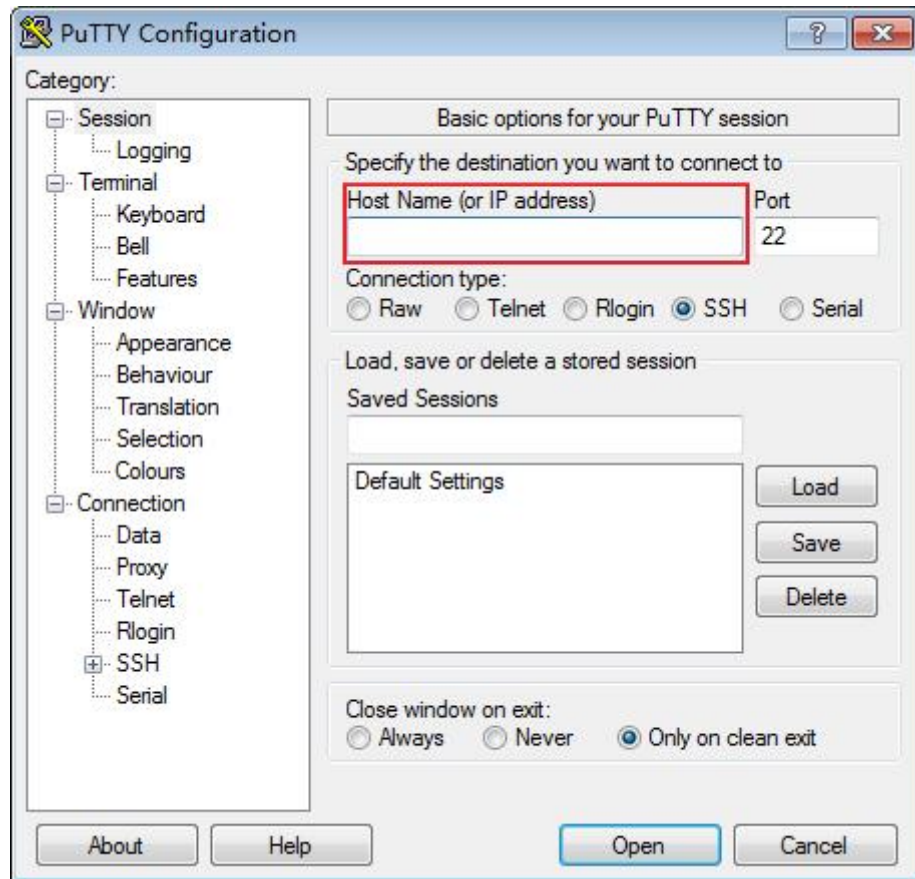
Figure 3-25 Importing the private key file



5. In the **Actions** area, click **Save private key**.
6. Save the converted private key, for example, **kp-123.ppk**, to the local computer.
7. Double-click **PuTTY.EXE**. The **PuTTY Configuration** page is displayed.


8. Choose **Session** and enter the EIP of the ECS under **Host Name (or IP address)**.

Figure 3-26 Configuring the EIP



9. Choose **Connection > Data**. Enter the image username in **Auto-login username**.

Figure 3-27 Entering the username



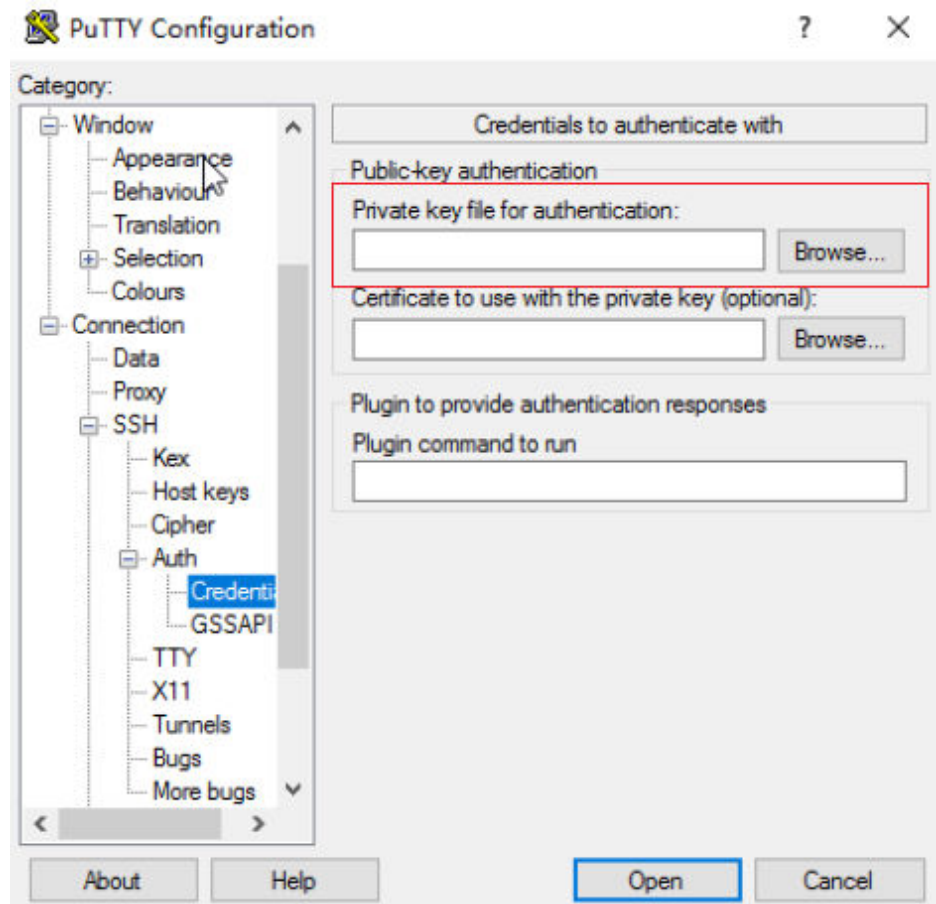
NOTE

When you log in to an ECS using an SSH key:

- The image username is **core** for a CoreOS public image.
- The image username is **root** for a non-CoreOS public image.

10. Choose **Connection > SSH > Auth > Credentials**. In the configuration item **Private key file for authentication**, click **Browse** and select the private key converted in step 6.

Figure 3-28 Importing the private key file



11. Click **Open** to log in to the ECS.

Method 2: Use Xshell to log in to the ECS.


1. Start the Xshell tool.
2. Run the following command using the EIP to remotely log in to the ECS through SSH:

```
ssh Username@EIP
```

NOTE

When you log in to an ECS using an SSH key:

- The image username is **core** for a CoreOS public image.
 - The image username is **root** for a non-CoreOS public image.
3. (Optional) If the system displays the **SSH Security Warning** dialog box, click **Accept & Save**.

Figure 3-29 SSH Security Warning

4. Select **Public Key** and click **Browse** beside the user key text box.
5. In the user key dialog box, click **Import**.
6. Select the locally stored key file and click **Open**.
7. Click **OK** to log in to the ECS.

Logging In to a Linux ECS from a Local Linux Server

To log in to the Linux ECS from local Linux, perform the operations described in this section. The following operations use private key file **kp-123.pem** as an example to log in to the ECS. The name of your private key file may differ.

1. On the Linux CLI, run the following command to change operation permissions:

```
chmod 400 /path/kp-123.pem
```

NOTE

In the preceding command, replace *path* with the actual path where the key file is saved.

2. Run the following command to log in to the ECS:

```
ssh -i /path/kp-123.pem Default username@EIP
```

For example, if the default username is **root** and the EIP is **123.123.123.123**, run the following command:

```
ssh -i /path/kp-123.pem root@123.123.123.123
```

NOTE

In the preceding command:

- *path* refers to the path under which the key file is stored.
- *EIP* is the EIP bound to the ECS.

Follow-up Procedure

- After logging in to the ECS using the SSH key, you can set a password (by using the **passwd** command) to log in to the ECS using VNC.

Helpful Links

- [Application Scenarios for Using Passwords](#)
- [Why Can't I Log In to My Linux ECS?](#)

3.4.4 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.

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 which the ECS belongs to. 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 3-30 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.

```
root@xx.xx.xx.xx's password:
```

```
Welcome to Elastic Cloud Service
```

Helpful Links

- [Application Scenarios for Using Passwords](#)
- [Why Can't I Log In to My Linux ECS?](#)

3.4.5 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. If the password is forgotten, reset the password by referring to [Resetting the Password for Logging In to a Linux 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 which the ECS belongs to. 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
```

- Using an SSH key
 - a. Open the terminal of the macOS and run the following command to change permissions. The following operations use private key file **kp-123.pem** as an example. Replace it with your actual private key file.

```
chmod 400 /path/kp-123.pem
```

NOTE

In the preceding command, *path* refers to the path where the key file is saved.

- b. Run the following command to log in to the ECS:

```
ssh -i /path/kp-123.pem Username@EIP
```

NOTE

- The username is **core** for a CoreOS public image.
- The username is **root** for a non-CoreOS public image.

Follow-up Procedure

- After logging in to the ECS using the SSH key, you can set a password (by using the **passwd** command) to log in to the ECS using VNC.

3.4.6 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 Termius, 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. If the password is forgotten, reset the password by referring to [Resetting the Password for Logging In to a Linux 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 which the ECS belongs to. 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 3-31 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 3-32 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 3-33 Login information



If the following page is displayed, you have connected to the Linux ECS.

Figure 3-34 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 3-35 Starting JuiceSSH




2. On the **Connections** page, tap .

Figure 3-36 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.


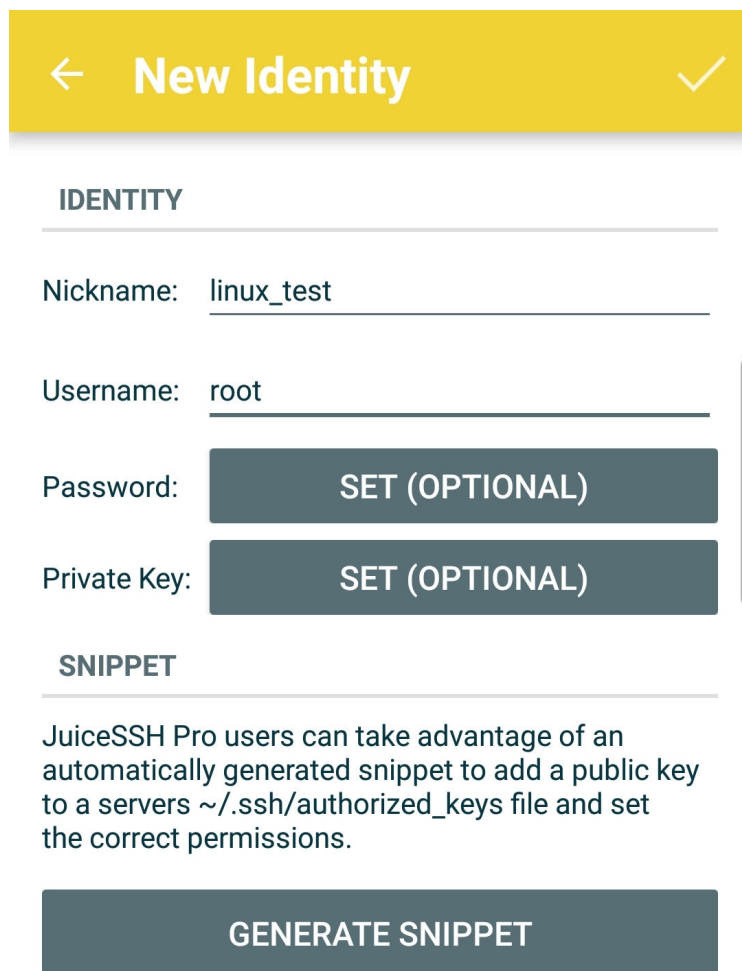
- ii. On the **New Identity** page, set the following parameters and tap .
 - **Nickname:** Set an identity name as required to facilitate subsequent management. This parameter is optional. In this example, set it to **linux_test**.
 - **Username:** Enter **root**.
 - **Password:** Tap **SET (OPTIONAL)**, enter the login password, and tap **OK**.

Figure 3-37 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 3-38 Port

← **New Connection** ✓

BASIC SETTINGS

Nickname:

Type:

Address:

Identity:

ADVANCED SETTINGS

Port:

Connect Via:

Run Snippet:

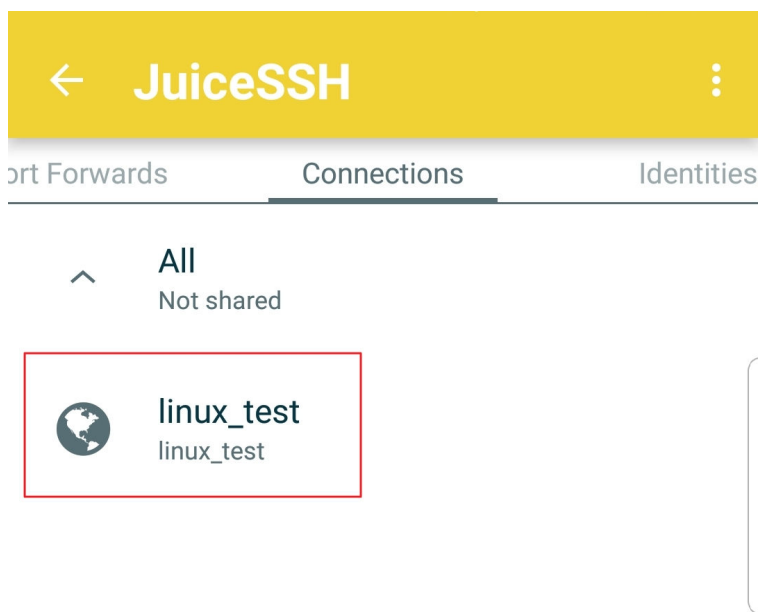
Backspace:

GROUPS

ADD TO GROUP


4. On the **Connections** page, tap the created connection.

Figure 3-39 Connections




5. Confirm the information that is displayed and tap **ACCEPT**.

Figure 3-40 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 3-41 Tutorial



You have logged in to the Linux ECS.

Figure 3-42 Successful login



3.5 Managing GPU Drivers of GPU-accelerated ECSs

3.5.1 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 3-8 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.

Driver	License	CUDA	OpenGL	DirectX	Vulkan	Application Scenario	Description
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.

3.5.2 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 3-9 Mapping between Tesla drivers and ECS types

ECS Type	Driver	Product Series	Product
P2s	Tesla	V	V100
Pi2	Tesla	T	T4
Pi2nl	Tesla	T	T4

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 3-43 Selecting the CUDA Toolkit version

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 3-44 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](#)

3.5.3 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 3-45 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 3-10](#).

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 3-10 GRID driver versions supported by GPU-accelerated ECSs

ECS Type	GPU Attachment	OS	Driver Version	CPU Architecture
P2s	GPU passthrough	<ul style="list-style-type: none"> Windows Server 2016 Standard 64bit 	Select a version as needed.	x86_64

ECS Type	GPU Attachment	OS	Driver Version	CPU Architecture
PI2	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](#).
- Refer to [Figure 3-46](#) to obtain the Product Activation Key (PAK) from the email indicating successful registration with NVIDIA.

Figure 3-46 PAK




5. Enter the PAK obtained in step 4 on the **Redeem Product Activation Keys** page and click **Redeem**.

Figure 3-47 Redeem Product Activation Keys



6. Specify **Username** and **Password** and click **LOGIN**.

Figure 3-48 Logging in to the official NVIDIA website

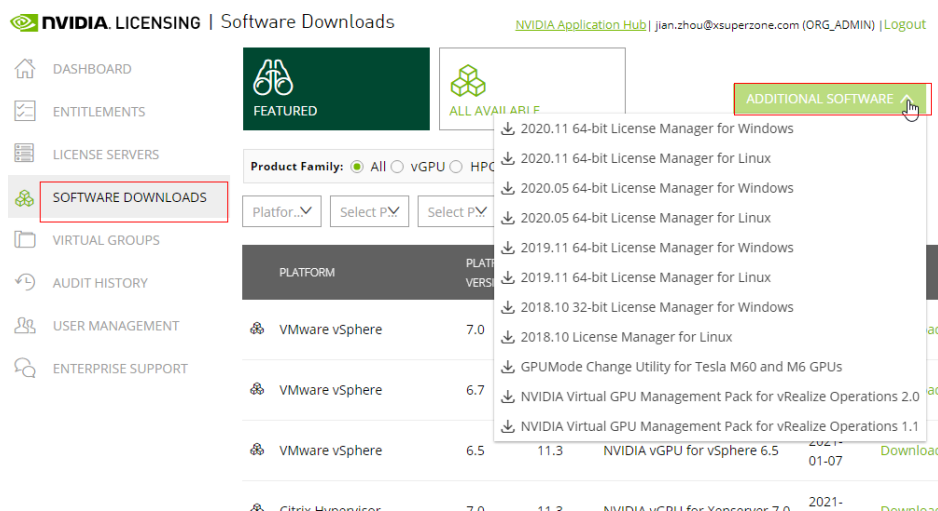


7. Log in at the official NVIDIA website as prompted and select **SOFTWARE DOWNLOADS**.

Figure 3-49 SOFTWARE DOWNLOADS page

PLATFORM	PLATFORM VERSION	PRODUCT VERSION	DESCRIPTION	RELEASE DATE
VMware vSphere	7.0	11.3	NVIDIA vGPU for vSphere 7.0	2021-01-07
VMware vSphere	6.7	11.3	NVIDIA vGPU for vSphere 6.7	2021-01-07
VMware vSphere	6.5	11.3	NVIDIA vGPU for vSphere 6.5	2021-01-07
Citrix Hypervisor	7.0	11.3	NVIDIA vGPU for Xenserver 7.0	2021-01-07

8. Download the GRID driver of the required version. For details, see [Table 3-10](#).
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 3-50 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, ensure that the license server and the GPU-accelerated ECS with the GRID driver installed are in the same VPC subnet.
 - If the license server is to be accessed using a public IP address, configure the security group which license server belongs to 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/flexnetls/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
-----

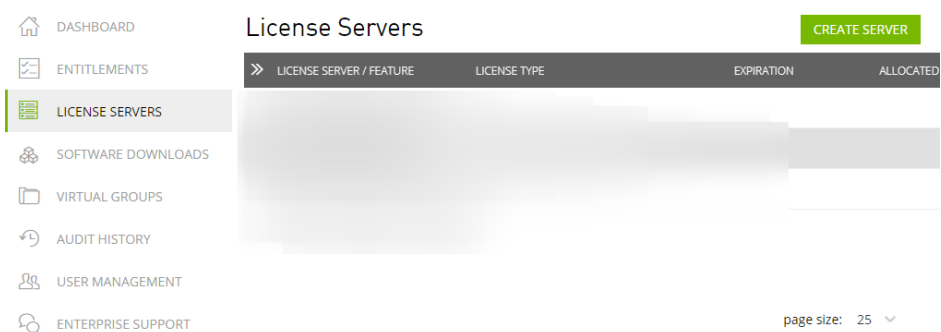
License Server has been successfully installed to:

  /opt/flexnetls/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 3-51 LICENSE SERVERS



- b. Click **CREATE SERVER**.
 - c. On the displayed **Create License Server** page, configure parameters.


Figure 3-52 Create License Server

Table 3-11 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. You can log in to the ECS and run ipconfig -a to query the MAC address.
Feature	Select a feature, enter the number of required licenses in the Licenses text box, and click ADD . In active/standby deployment, enter the name of the standby server in Failover License Server and enter the MAC address in Failover MAC Address .

- d. Click **CREATE LICENSE SERVER**.
- e. Download the license file.

Figure 3-53 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 3-54 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 3-55 License server management console

3.5.4 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++`

- Download the NVIDIA driver package.

Select a driver version at [NVIDIA Driver Downloads](#) based on the ECS type. Click **SEARCH**.

Figure 3-56 Selecting a NVIDIA driver version

NVIDIA Driver Downloads

Advanced Driver Search

Product Type: <input type="text" value="Tesla"/>	Operating System: <input type="text" value="Linux 64-bit"/>
Product Series: <input type="text"/>	CUDA Toolkit: <input type="text" value="10.1"/>
Product: <input type="text"/>	Language: <input type="text" value="English (US)"/>
	Recommended/Beta: <input type="text" value="All"/> ?

- Select a driver version as required. The following uses Tesla 418.67 as an example.

Figure 3-57 Selecting a driver version

NVIDIA Driver Downloads

Advanced Driver Search

Product Type: <input type="text" value="Tesla"/>	Operating System: <input type="text" value="Linux 64-bit"/>
Product Series: <input type="text"/>	CUDA Toolkit: <input type="text" value="10.1"/>
Product: <input type="text"/>	Language: <input type="text" value="English (US)"/>
	Recommended/Beta: <input type="text" value="All"/> ?

Name	Version	Release Date	CUDA Toolkit
Tesla Driver for Linux x64	418.126.02	February 28, 2020	10.1
Tesla Driver for Linux x64	418.116.00	December 9, 2019	10.1
Tesla Driver for Linux x64	418.87.01	October 3, 2019	10.1
Tesla Driver for Linux x64	418.87.00	August 14, 2019	10.1
Tesla Driver for Linux x64	418.67	May 7, 2019	10.1
Tesla Driver for Linux x64	418.40.04	March 25, 2019	10.1
Tesla Driver for Linux x64	418.40.04	March 25, 2019	10.1

- Click the driver to be downloaded. On the **TESLA DRIVER FOR LINUX X64** page that is displayed, click **DOWNLOAD**.
- Copy the download link.

Figure 3-58 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 3-59 Obtaining the installation package

```
root@ecs-474b:~# wget http://us.download.nvidia.com/tesla/418.67/NVIDIA-Linux-x86_64-418.67.run
--2020-03-26 17:59:31-- http://us.download.nvidia.com/tesla/418.67/NVIDIA-Linux-x86_64-418.67.run
Resolving us.download.nvidia.com (us.download.nvidia.com)... 129.227.66.140, 129.227.66.139
Connecting to us.download.nvidia.com (us.download.nvidia.com)|129.227.66.140:80... connected.
HTTP request sent, awaiting response... 301 Moved Permanently
Location: https://us.download.nvidia.cn/tesla/418.67/NVIDIA-Linux-x86_64-418.67.run [following]
--2020-03-26 17:59:34-- https://us.download.nvidia.cn/tesla/418.67/NVIDIA-Linux-x86_64-418.67.run
Resolving us.download.nvidia.cn (us.download.nvidia.cn)... 60.222.11.61, 60.222.11.11, 123.134.184.166, ...
Connecting to us.download.nvidia.cn (us.download.nvidia.cn)|60.222.11.61:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 107232512 (102M) [application/octet-stream]
Saving to: 'NVIDIA-Linux-x86_64-418.67.run'

NVIDIA-Linux-x86_64-418.67.run 100%[=====] 102.26M 1.07MB/s in 5m 2s
2020-03-26 18:04:40 (346 KB/s) - 'NVIDIA-Linux-x86_64-418.67.run' saved [107232512/107232512]
```

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 3-60 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
```

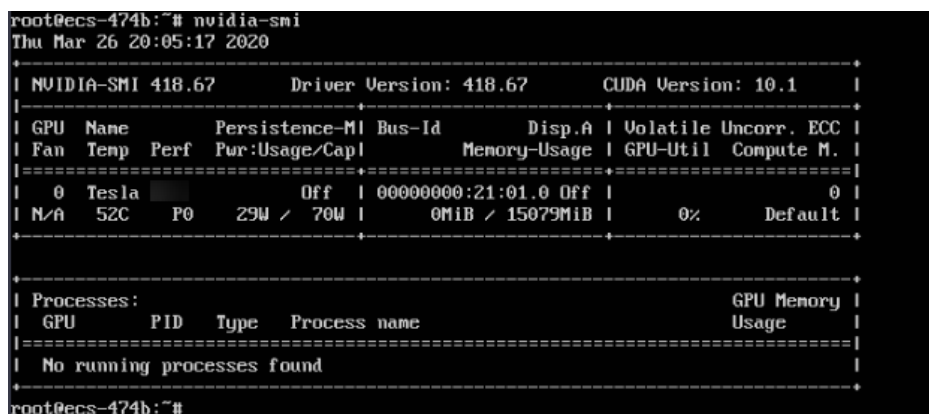
- c. Run the following command to back up and create an initramfs application:
 - Ubuntu
sudo update-initramfs -u
 - CentOS:
mv /boot/initramfs-\$(uname -r).img /boot/initramfs-\$(uname -r).img.bak
dracut -v /boot/initramfs-\$(uname -r).img \$(uname -r)
 - d. Restart the ECS:
reboot
10. Select **OK** for three consecutive times as prompted to complete the driver installation.

Figure 3-61 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 3-62 Viewing the NVIDIA driver version



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 3-63 Selecting a driver type (Windows)

Advanced Driver Search

Product Type: <input type="text" value="Tesla"/>	Operating System: <input type="text" value="Windows Server 2016"/>
Product Series: <input type="text"/>	CUDA Toolkit: <input type="text" value="10.1"/>
Product: <input type="text"/>	Language: <input type="text" value="English (US)"/>
	Recommended/Beta: <input type="text" value="All"/> ?

3. Select a driver version as required. The following uses Tesla 425.25 as an example.

Figure 3-64 Selecting a driver version (Windows)

Advanced Driver Search

Product Type: <input type="text" value="Tesla"/>	Operating System: <input type="text" value="Windows Server 2016"/>
Product Series: <input type="text"/>	CUDA Toolkit: <input type="text" value="10.1"/>
Product: <input type="text"/>	Language: <input type="text" value="English (US)"/>
	Recommended/Beta: <input type="text" value="All"/> ?

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 3-65 Downloading the driver installation package

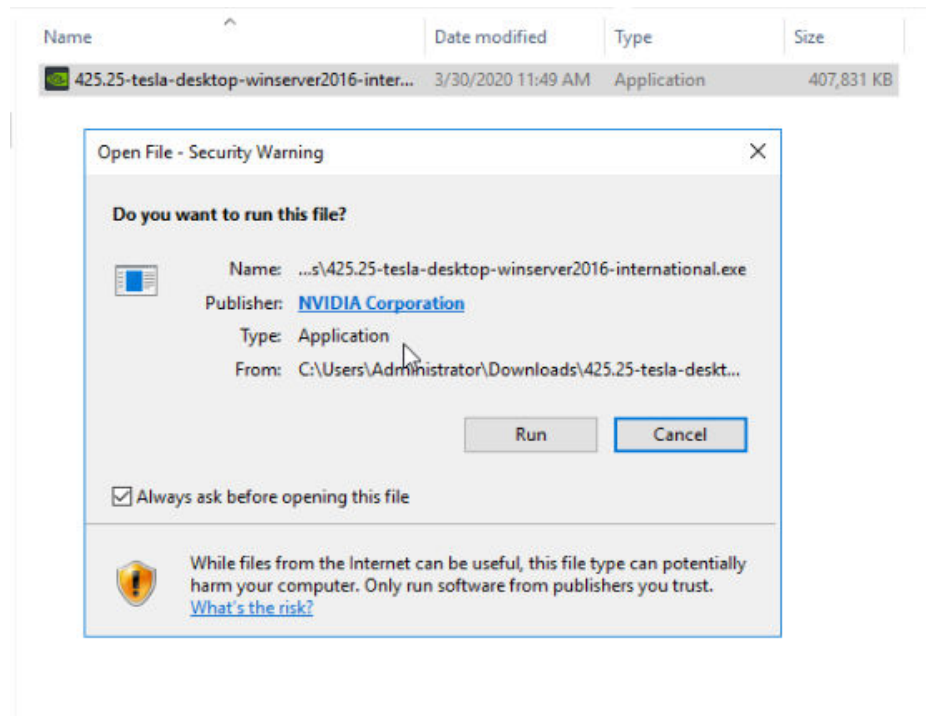
Download

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 3-66 Running the NVIDIA driver installation program




7. Select an installation path and click **OK**.

Figure 3-67 Selecting an installation path




8. Install the NVIDIA program as prompted.

Figure 3-68 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 3-69 Display adapters

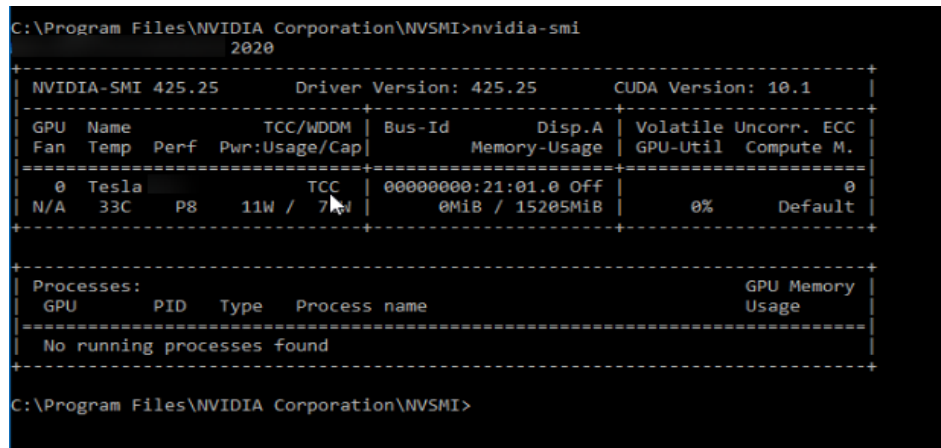


- b. Open the **cmd** window on the ECS and run the following commands:
cd C:\Program Files\NVIDIA Corporation\NVSMI

nvidia-smi

If the command output contains the installed driver version, the driver has been installed.

Figure 3-70 Viewing the NVIDIA driver version

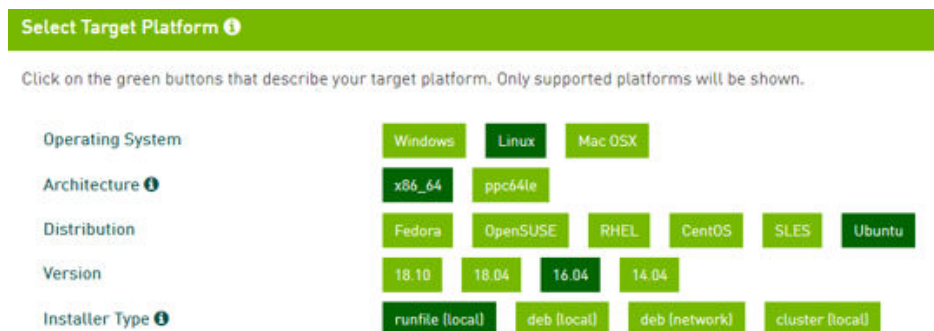


Installing the CUDA Toolkit on a Linux ECS

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.


- Log in to the ECS.
- 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++**
- On the CUDA download page, set parameters according to the information shown in [Obtaining a Tesla Driver and CUDA Toolkit](#).

Figure 3-71 Selecting a CUDA version



- Find the link for downloading CUDA 10.1 and copy the link.

Figure 3-72 Copying the link for downloading CUDA

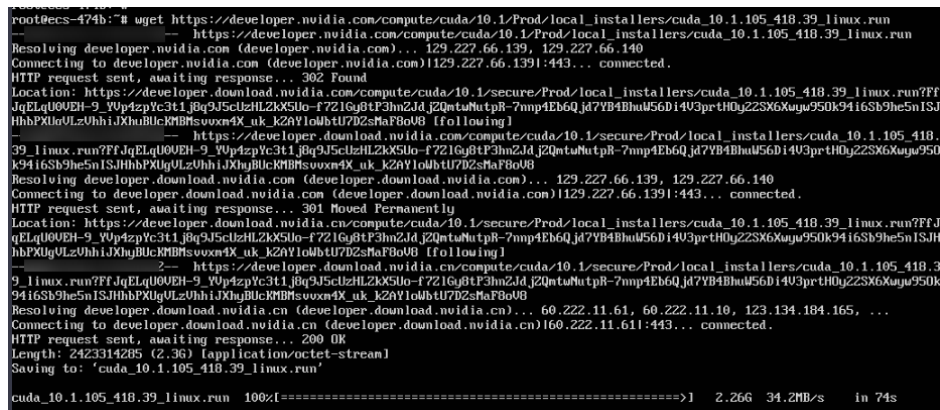


- Run the following command on the ECS to download CUDA:

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**

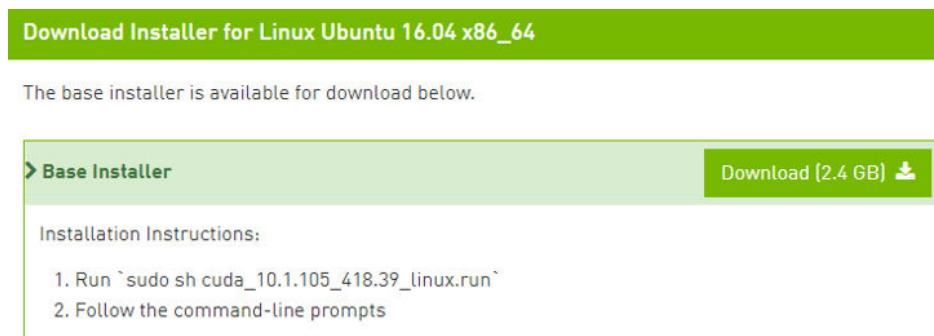
Figure 3-73 Downloading CUDA



- Install CUDA.

Follow the instructions provided on the official NVIDIA website.

Figure 3-74 Installing CUDA



- 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 3-75 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 3-76 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 3-77 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_Uilities/deviceQuery`:
cd /usr/local/cuda-10.1/samples/1_Uilities/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 3-78 deviceQuery common output

```
root@ecs-474b:/usr/local/cuda-10.1/samples/1_Uilities/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_Uilities/deviceQuery#
```

13. Check the CUDA version.
/usr/local/cuda/bin/nvcc -V

Figure 3-79 Checking the CUDA version

```
root@ecs-474b:/usr/local/cuda-10.1/samples/1_Uilities/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 3-80 Selecting a CUDA version




3. Find the link for downloading CUDA 10.1.

Figure 3-81 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 3-82 Installing CUDA




6. On the **CUDA Setup Package** page, select an installation path and click **OK**.

Figure 3-83 Selecting an installation path




7. Install the CUDA toolkit as prompted.

Figure 3-84 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 3-85 Successful installation



3.5.5 Uninstalling a GPU Driver from a GPU-accelerated ECS

Scenarios

You can manually uninstall the GPU driver from a GPU-accelerated ECS.
This section describes how to uninstall a GPU driver from a Windows ECS and a Linux ECS.

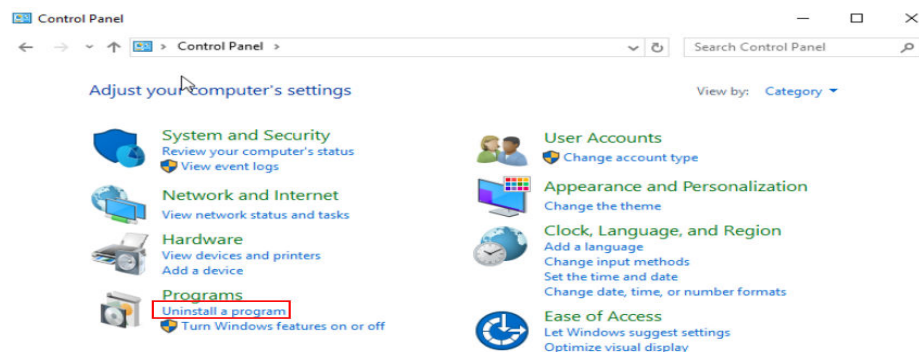
- [Uninstalling a GPU Driver from a Windows ECS](#)
- [Uninstalling a GPU Driver from a Linux ECS](#)

Uninstalling a GPU Driver from a Windows ECS

This section uses Windows Server 2016 Datacenter Edition 64-bit as an example to describe how to uninstall the NVIDIA driver (driver version: 462.31) from a GPU-accelerated ECS.


1. Log in to the ECS.
2. Click **Start** in the task bar and choose **Control Panel**.
3. In Control Panel, click **Uninstall a program** under **Programs**.

Figure 3-86 Uninstalling a program.




4. Right-click the NVIDIA driver to be uninstalled and choose **Uninstall/Change** from the shortcut menu.

Figure 3-87 Uninstalling a NVIDIA driver




5. In the displayed **NVIDIA Uninstaller** window, click **UNINSTALL**.

Figure 3-88 Confirming the uninstallation



6. After the uninstallation is complete, click **RESTART LATER**.
7. Check whether the NVIDIA driver has been uninstalled.
 - a. In Control Panel, click **Device Manager**.
If no NVIDIA graphics cards are not displayed under **Display adapters**, the driver is uninstalled successfully.

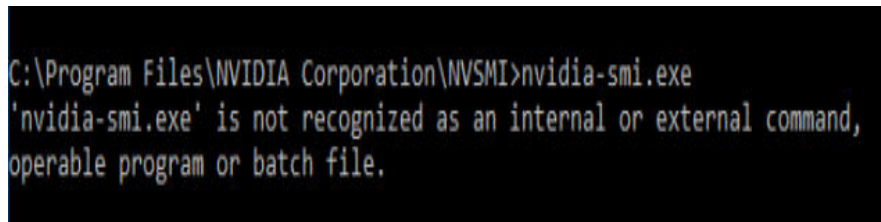
Figure 3-89 Viewing Display adapters



- b. Open the cmd window of the ECS and run the following commands:

```
cd C:\Program Files\NVIDIA Corporation\NVSMI  
nvidia-smi.exe
```

Figure 3-90 Command output



```
C:\Program Files\NVIDIA Corporation\NVSMI>nvidia-smi.exe  
'nvidia-smi.exe' is not recognized as an internal or external command,  
operable program or batch file.
```

If the command output indicates that the file does not exist, the driver is uninstalled successfully.

After the NVIDIA driver is uninstalled, you can install a new NVIDIA driver without restarting the ECS.

Uninstalling a GPU Driver from a Linux ECS

For NVIDIA Tesla drivers installed using .run Packages, you are advised to perform the following steps to uninstall it.

NOTE

If you use .run Packages to install the NVIDIA Grid driver, you only need to perform [step 1](#) to uninstall the NVIDIA driver.

The following uses 64-bit Ubuntu Server 20.04 as an example to describe how to uninstall Tesla 460.73.01 and CUDA 11.2.

1. Uninstall the NVIDIA driver.
 - a. Query the path where **nvidia-uninstall** is stored.
whereis nvidia-uninstall
Generally, **nvidia-uninstall** is stored in the **/usr/bin/** directory.

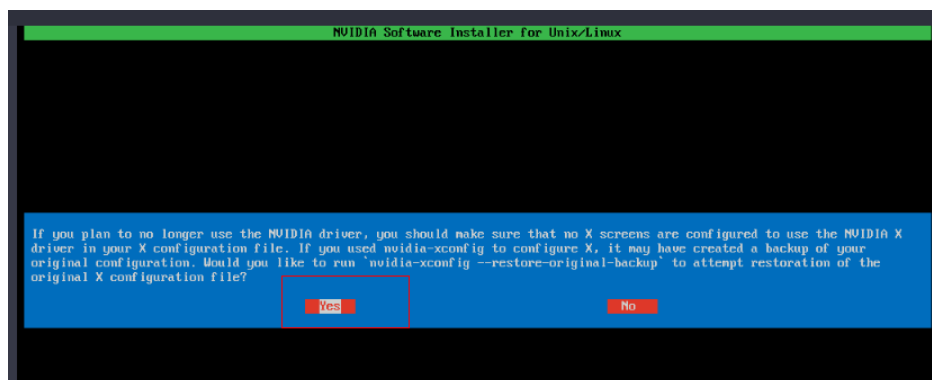
Figure 3-91 Querying the nvidia-uninstall path



```
root@ecs-... # whereis nvidia-uninstall  
nvidia-uninstall: /usr/bin/nvidia-uninstall
```


- b. Uninstall the driver from the path where **nvidia-uninstall** is stored.
/usr/bin/nvidia-uninstall
 - c. Select **Yes** and press **Enter**.

Figure 3-92 NVIDIA driver uninstallation (1)




- d. Select **OK** and press **Enter**.

Figure 3-93 NVIDIA driver uninstallation (2)



- e. After the driver is uninstalled, press **Enter**.

Figure 3-94 NVIDIA driver uninstallation (3)



2. Uninstall the CUDA and CUDA Deep Neural Network (cuDNN) libraries.
To upgrade the CUDA driver version, uninstall the corresponding CUDA library and then install a new one with the target version.

- a. Uninstall the CUDA library.

/usr/local/cuda/bin/cuda-uninstaller


Generally, **cuda-uninstaller** is stored in the **/usr/local/cuda/bin** directory.

NOTE

The uninstallation command varies depending on CUDA versions. If the **cuda-uninstaller** file is not found, check whether a file starting with **uninstall_cuda** exists in the **/usr/local/cuda/bin/** directory.

If such a file exists, replace **cuda-uninstaller** in the preceding command with the file name.

- b. On the uninstallation page, select all options, move the cursor to **Done**, and press **Enter**.

Figure 3-95 Uninstalling a CUDA driver

If the CUDA library is uninstalled, the message "Successfully uninstalled" is displayed.

- c. Remove the CUDA and cuDNN libraries.
`rm -rf /usr/local/cuda-11.2`

3.6 Managing ECSs

3.6.1 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 `/user/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 `/user/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
9. Run the following command to restart the ECS so that all services and applications running on the ECS use the new time zone:
reboot
10. 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

1. Log in to the ECS.
2. Click the time display on the far right side of the task bar located at the bottom of your screen. In the displayed dialog box, click **Date and time settings**.

The **Date & time** window is displayed.

Figure 3-96 Date and time

Date & time

Date and time

9:52 AM, Tuesday, March 28, 2023

Set time automatically



Set time zone automatically



Change date and time

Change

Time zone

(UTC+08:00) Taipei

Adjust for daylight saving time automatically



Show additional calendars in the taskbar

Simplified Chinese (Lunar)

3. Select the target time zone from the **Time zone** drop-down list.

NOTE

You can also enable **Set time zone automatically**. The system automatically updates your time zone based on the physical location of your ECSs.

3.6.2 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 3-12](#). 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]).Value
```

```
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 3-12 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 3-12 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 3-13 .

Metadata Type	Metadata Item	Description
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, authorize agency permissions to the op_svc_ecs account and ECSs in IAM. NOTE You can determine what permissions are granted to the agency based on the principal of least privilege (PoLP). ECSs will not use agencies to perform operations on resources.
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.
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.

Metadata Type	Metadata Item	Description
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 3-13 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?
enterprise_project_id	String	Specifies the ID of the enterprise project accommodating an ECS.

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+dNwlydGlxJHUp+00poeUsAdBFkbPbYQTmpNwpoEb43k9z+96TyrekNKS+iLYDdRny4kKGoNPEVBCC05Hg1TcDbIAPfJwgJS1okqEtlcofUhKml3K0fto+5KXEDU3GNuGwyZXjdVb9HQWU+E1jztAJjqsahnU+g/tawABTVySLBKlAT8fMGax1mTGgArucn/WzDcy19DGioKPE7F8lLtSQ4Ww3VCLK5VYB/h0x+4r7IVHrPmYX/bi1Yhm3Dc4rRYNaTjdOV5gUOsbo3oAeQkmKwQ/NO0N8qw5Ya4l8ZUW4tMav4mOsRySOOB35v0bvaJc6p+50DTbWNeX5A2MLiEhTP3vsPrmvk4LRF7CLz2J2TGIM14OoVBw7LARwmv9cz532zHki/c8tlhRzLmOTXh/wL36zFW10DeuReUGmxth7lGNmRMQKV6+mil78jm/KMPpgAdK3vwYF/GcelOFJD2HghMUUCeMbwYnvijLTejuBpwhJMNIHA/NvIEsxDxqBCoss/Jfe+yCmUFyxovj+L8oNkTzkmCNzw3Ra0hiKchGhqK3BleToV/kVx5DdF081xrEA+qyoM6CVyfJtEoz1zRRyoo9bJ65Eg6Jd8dj1UCVsDqRY1pljgzE/Mzsw6AaaCVhaMJL7u7YMVdyKzA6z65Xtvujz0Vo=",
  "uuid": "ca9e8b7c-f2be-4b6d-a639-f10b4d994d04",
  "availability_zone": "lt-test-1c",
  "enterprise_project_id": "0",
  "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"
}
```

User Data (OpenStack Metadata 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.

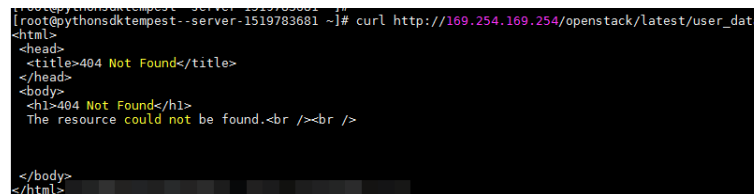
- URI
/169.254.169.254/openstack/latest/user_data
- Usage method
Supports GET requests.
- Example
Linux:
curl http://169.254.169.254/openstack/latest/user_data
Windows:

Invoke-RestMethod http://169.254.169.254/openstack/latest/user_data

```
ICAgICAgDQoiQSBjbG91ZCBkb2VzIG5vdCBrbm93IHdoeSBpdCBtb3ZlcyBpbiBqdXN0IHN1Y2ggYSBkaXJlY  
3Rpb24gYW5kIGF0IHN1Y2ggYSBzcGVlZC4uLkl0IGZlZWxzIGFuIGltcHVsc2lvbi4uLnRoXMGaXMgdGhllH  
BsYWNlIHhrvIGdviG5vdy4gQnV0IHRoZSBza3kga25vd3MgdGhllHJlYXNvbnMgYW5kIHRoZSBwYXR0ZXJ  
cyBiZWphbmQgYWxslGNsb3VkcycwYmV5b3VkaW5kIGVudSB3aWxslGtub3csIHVbywgd2hbiB5b3UgbGlm  
dCB5b3VkaW5kIGVudSB3aWxslGtub3csIHVbywgd2hbiB5b3UgbGlm dCB5b3VkaW5kIGVudSB3aWxslG  
tub3csIHVbywgd2hbiB5b3UgbGlm dCB5b3VkaW5kIGVudSB3aWxslGtub3csIHVbywgd2hbiB5b3UgbGlm  
dCB5b3VkaW5kIGVudSB3aWxslGtub3csIHVbywgd2hbiB5b3UgbGlm dCB5b3VkaW5kIGVudSB3aWxslG
```

NOTE

If user data was not passed to the ECS during ECS creation, the query result is 404.

Figure 3-97 404 Not Found

```
[root@python3ktempest--server-1519783681 ~]# curl http://169.254.169.254/openstack/latest/user_data
<html>
<head>
<title>404 Not Found</title>
</head>
<body>
<h1>404 Not Found</h1>
The resource could not be found.<br /><br />
</body>
</html>
```

Network Data (OpenStack Metadata API)

This API is used to query information about all NICs attached to an ECS, including their DNS server addresses, network bandwidth, IDs, private IP addresses, EIPs, and MAC addresses.

- URI
/openstack/latest/network_data.json
- Usage method
Supports GET requests.
- Example

NOTE

instance_max_bandwidth and **instance_min_bandwidth** are in the unit of Mbit/s. If the value is -1, the bandwidth is not limited.

Linux:

curl http://169.254.169.254/openstack/latest/network_data.json

Windows:

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
```

```
ICAgICAgDQoiQSBjbG91ZCBkb2VzIG5vdCBRbm93IHdoeSBpdCBtb3ZlcyBpbjBqdXN0IHN1Y2ggYSBkaXJlY3Rpb24gYW5kIGF0IHN1Y2ggYSBzcGVlZC4uLkl0IGZlZWxzIGFuIGltcHVsc2lubi4uLnRoaXMgaXMgdGhllHBsYWVWNIHRvIGdvlG5vdy4gQnV0IHRoZSBza3kga25vd3MgdGhllHJlYXNvbnMgYW5kIHRoZSBwYXR0ZXJuYyBiZWphbmQgYWxslGNsb3VkcycgYW5kIHlvdSB3aWxslGtub3csIHRvbywgd2h1b3UgYmVudCB5b3Vyc2VsZiBoaWdoIGVub3VnaCB0byBzZWUgYmV5b25kIGhvcml6b25zLiINCg0KLVJpY2hhcmQgQmFjaA=
```

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?](#)

3.6.3 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.

Figure 3-98 User data injection



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 3-14 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 and SSH key, 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
```

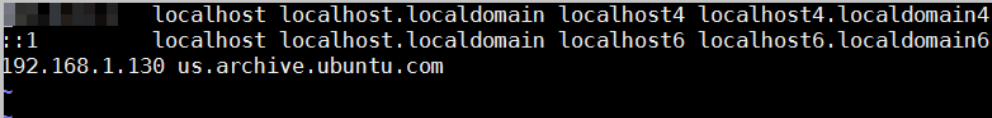
Example 2: Inject a Cloud-Config data script.

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 3-99 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 3-15 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.

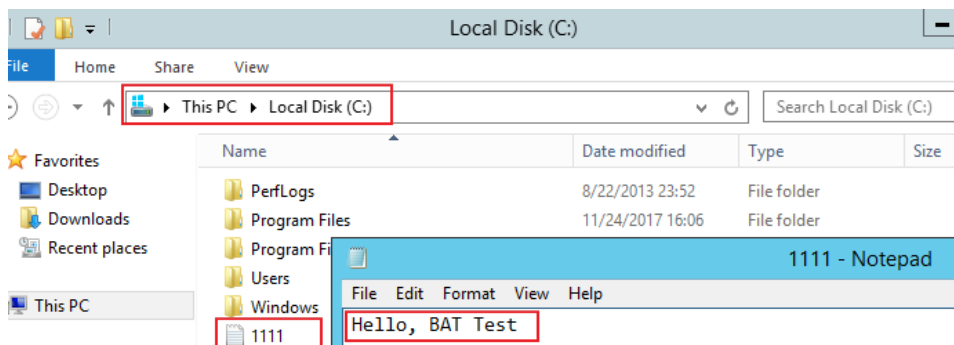
- How can I view the customized user data passed into a Windows ECS?
 - a. Log in to the ECS.
 - b. Access the following URL in the address box of the browser and view the user data:
http://169.254.169.254/openstack/latest/user_data
- Script usage examples
This section describes how to inject scripts 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 3-100 Creating text file (Batch)

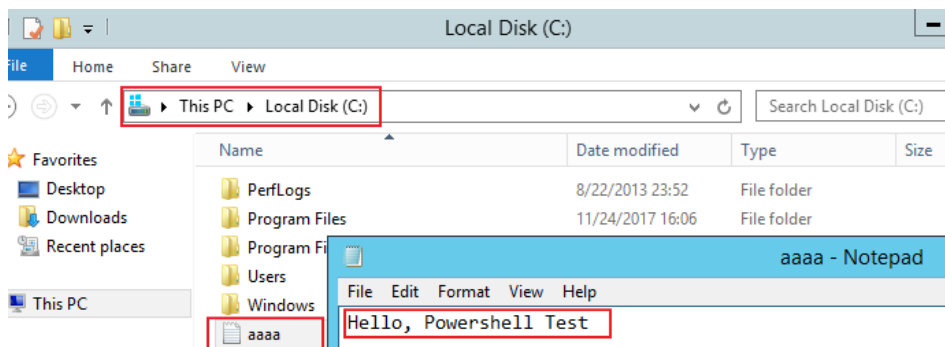
To view the user data passed to the Windows ECS, log in at http://169.254.169.254/openstack/latest/user_data.

Figure 3-101 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 3-102 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 3-103 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 3-16](#).

Table 3-16 Password complexity requirements

Parameter	Requirement
Password	<ul style="list-style-type: none">• Consists of 8 to 26 characters.• 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 3-17](#).

Table 3-17 Password complexity requirements

Parameter	Requirement
Password	<ul style="list-style-type: none">• Consists of 8 to 26 characters.• 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 3-17](#).

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
```

Case 6

This case illustrates how to pass the user data to assign user **root** permissions for remotely logging in to a Linux ECS. After passing the file to an ECS, you can log in to the ECS as user **root** using SSH key pair authentication.

User data example:

```
#cloud-config
disable_root: false
runcmd:
- sed -i 's/^PermitRootLogin.*$/PermitRootLogin without-password/' /etc/ssh/sshd_config
- sed -i '/^KexAlgorithms.*$/d' /etc/ssh/sshd_config
- service sshd restart
```

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/>



3.6.4 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. Click  in the upper left corner and select your region and project.
3. Under **Computing**, choose **Elastic Cloud Server**.
4. Click the name of the target ECS.
5. 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.

6. Click  next to the ECS name.
7. Click **OK**.

Changing the Names of Multiple ECSs in a Batch

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Computing**, click **Elastic Cloud Server**.
4. Select the target ECSs.
5. Click **More** above the ECS list and select **Change ECS Name** from the drop-down list.
6. Enter a new name.
7. Click **OK**.

If you change ECS names in a batch, the new ECS names are the same, for example, all are **ecs-test**.

3.6.5 Migrating an ECS

Scenarios

ECSs can be migrated:

- Between Dedicated Hosts (DeHs)
- From a DeH to a public resource pool
- From a public resource pool to a DeH

This section describes how to migrate ECSs from a public resource pool to a DeH.

NOTE

- Before migrating an ECS, ensure that there are available DeH resources.
- For details about migrating ECSs from a DeH to another DeH or to a public resource pool, see "Migrating ECSs" in the *Dedicated Host User Guide*.

Constraints

- Only stopped ECSs can be migrated.
- To ensure that the migration is successful, there must be an available DeH.
- CBR and CSBS backups are not affected by migrations.
- ECS IDs remain unchanged after a migration.

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. Locate the row that contains the target ECS and choose **More > Migrate ECS** in the **Operation** column.

4. In the displayed dialog box, select the target DeH.

 **NOTE**

If no DeHs are available, create a DeH first. For details, see *Purchasing DeHs on a Pay-per-Use Basis* in the *Dedicated Host User Guide*.

5. Click **OK**.

3.6.6 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.

An ECS group supports the following policies:

- **Anti-affinity policy:** ECSs in the same ECS group run on different hosts for improved reliability, high availability, and disaster recovery.


You can associate the anti-affinity policy with an ECS group so that the ECSs in this group can run on different hosts for improved service HA and underlying DR.

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**


Constraints

- ECS groups support the anti-affinity policy only.
- ECSs are deployed on different physical hosts.
- If the maximum number of ECS groups is reached, you need to contact customer service to increase the quota.
- The maximum number of ECSs that can be added to an ECS group varies depending on the region. You can view the quota on the **ECS Group** page, as shown in **Figure 3-104**.

Figure 3-104 Maximum number of ECSs that can be added to 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. Click  in the upper left corner and select your region and project.
3. Under **Computing**, choose **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **ECS Group**.
5. On the **ECS Group** page, click **Create ECS Group**.
6. Enter the name of an ECS group.
7. Select a policy for the ECS group.
8. 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. Click  in the upper left corner and select your region and project.
 3. Under **Computing**, choose **Elastic Cloud Server**.


4. In the navigation pane on the left, choose **ECS Group**.
5. Locate the row that contains the target ECS group and click **Add ECS** in the **Operation** column.
6. On the **Add ECS** page, select an ECS to be added.
7. (Optional) Stop the ECS.

Certain types of ECSs need to be stopped before being added to an ECS group. If **Stop** is displayed in the row containing the selected ECS, you need to stop the ECS first.

- a. Click **Stop** in the **Operation** column.
 - b. Select a stop option.
 - **Stop**: The ECS will be stopped normally.
 - **Forcibly stop the preceding ECSs**: This operation will cause loss of unsaved data. Exercise caution when performing this operation.
 - c. Click **Yes**.
8. 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. Click  in the upper left corner and select your region and project.
3. Under **Computing**, choose **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **ECS Group**.
5. Expand the ECS group information and view the ECSs in the ECS group.
6. Locate the ECS to be removed and click **Remove** in the **Operation** column.
7. 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. Click  in the upper left corner and select your region and project.
3. Under **Computing**, choose **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **ECS Group**.
5. Locate the ECS group to be deleted and click **Delete** in the **Operation** column.
6. In the displayed dialog box, click **Yes**.

3.6.7 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.

3.6.8 (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
```

3.6.9 Starting and Stopping ECSs

You can start, stop, restart, or delete the ECS.

- To prevent a sudden load increase, you are advised to start or stop a small number of ECSs at a time.
- If an ECS remains in the **Restarting** or **Stopping** state for a long time, you can forcibly restart or stop it. In such a case, any unsaved data on the ECS will be lost. Therefore, exercise caution when forcibly restarting or stopping an ECS.

Starting ECSs

1. Log in to the management console.
2. Under **Computing**, select **Elastic Cloud Server**.
3. In the ECS list, select the target ECSs.
4. Click **Start** in the upper left corner of the list.
5. In the displayed window, click **Yes**.

NOTE

Contact the administrator if the ECS has been in the **Starting** state for more than 30 minutes.

Stopping ECSs

1. Log in to the management console.
2. Under **Computing**, select **Elastic Cloud Server**.
3. In the ECS list, select the target ECSs.
4. Click **Stop** in the upper left corner of the list.
5. In the displayed dialog box, select a stop option based on your service requirements.

NOTICE

After an ECS is forcibly stopped, unsaved data on the ECS will be lost.

6. Click **Yes** to stop the ECSs.

NOTE

Contact the administrator if the ECS has been in the **Stopping** state for more than 30 minutes.

Restarting ECSs

1. Log in to the management console.
2. Under **Computing**, select **Elastic Cloud Server**.
3. In the ECS list, select the target ECSs.
4. Click **Restart** in the upper left corner of the list.

NOTICE

After an ECS is forcibly restarted, unsaved data on the ECS will be lost.

5. Click **Yes** to stop the ECSs.

NOTE

Contact the administrator if the ECS has been in the **Restarting** state for more than 30 minutes.

Deleting an ECS

1. Log in to the management console.
2. Under **Computing**, select **Elastic Cloud Server**.
3. In the ECS list, select the target ECSs.
4. Click **Delete** in the upper left corner of the list.

NOTE

Contact the administrator if the ECS has been in the **Deleting** state for more than 30 minutes.

3.7 Modifying ECS Specifications

3.7.1 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

- When modifying the specifications of an ECS, sold-out vCPU and memory resources are unavailable for selection.
- 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.
- Before modifying specifications, make sure that the ECS has been stopped.

Step 1: Modify Specifications

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Computing**, click **Elastic Cloud Server**.
4. On the **Elastic Cloud Server** page, view the status of the target ECS.
If the ECS is not in **Stopped** state, click **More** in the **Operation** column and select **Stop**.
5. Click **More** in the **Operation** column and select **Modify Specifications**.
The **Modify ECS Specifications** page is displayed.
6. Select the new ECS type, vCPUs, and memory as prompted.
7. Click **Next**.
8. Confirm the settings, read and select the disclaimer, and then click **Submit Application**.
9. Check whether the specifications have been modified.
After modifying the specifications, you can check whether the specifications have been modified in **Failures**.
 - a. Check whether **Failures** is displayed on the management console. For details, see [Viewing Failed Tasks](#).
 - If yes, go to step **9.b**.
 - If no, the specifications have been modified.
 - b. Click **Failures**. Then, in the **Failures** dialog box, click **Operation Failures** and check whether the task is contained in the list by **Name/ID**, **Operated At**, or **Task**.
 - If yes, the specifications modification failed. See [Follow-up Procedure](#) for failure causes.
 - If no, the specifications have been modified.

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?](#)

Follow-up Procedure

Perform the following operations in the event of a specifications modification failure:

1. Log in to the management console.
2. Under **Management & Deployment**, choose **Cloud Trace Service**.
3. In the navigation pane on the left, choose **Trace List**.
4. In the **Trace Name** column, locate the **resizeServer** event by resource ID.
The resource ID is the ID of the ECS on which the specifications modification failed.
5. Click **View Trace** in the **Operation** column to view the failure cause.
If the fault cannot be rectified based on logs, contact customer service.

3.8 Reinstalling or Changing the OS

3.8.1 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 or key. Otherwise, the injection may fail, and the ECS cannot be logged in to.

- You can choose to encrypt the system disk of an ECS during OS reinstallation.


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.
- If an ECS OS is to be reinstalled using a full-ECS image, the ECS system disk can be encrypted.

Prerequisites

- The target ECS is stopped.
- The target ECS has a system disk attached.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Computing**, choose **Elastic Cloud Server**.
4. Locate the row containing the target ECS and choose **More > Manage Image/Disk/Backup > Reinstall OS** in the **Operation** column.

Only stopped ECSs support OS reinstallation. If the ECS is not stopped, stop it before proceeding with reinstallation.

5. (Optional) Select the **Encrypted** option to encrypt the system disk during OS reinstallation.

To enable encryption, click **Create Agency** to assign KMS access permissions to EVS. If you have rights granting permission, assign the KMS access permissions to EVS. If you do not have the permission, contact the user having the security administrator rights to assign the KMS access permissions. For details, see [Can All Users Use the Encryption Feature?](#)

Encryption parameters are as follows:

- **Encryption**: indicates that the EVS disk has been encrypted.
- **Create Agency**: assigns KMS access permissions to EVS to obtain KMS keys. After the permissions are assigned, follow-up operations do not require assigning permissions again.
- **Agency Name**: set to **EVSAccessKMS**, which means that permissions have been assigned to EVS to obtain KMS keys for encrypting or decrypting EVS disks.
- **KMS Key Name**: specifies the name of the key used by the encrypted EVS disk. You can select an existing key, or click **Create KMS Key** and create a new one on the KMS console. The default value is **evs/default**.
- **KMS Key ID**: specifies the ID of the key used by the encrypted data disk.

6. Select the login mode.

If the target ECS uses key pair authentication, you can replace the original key pair.

Figure 3-105 Reinstall OS

×

Reinstall OS

Note the following points before you reinstall the OS:

1. This operation will reinstall the original OS image on the ECS and delete data on all system partitions. [Back up](#) the data before you continue.
2. There are no impacts to data disks when you reinstall an OS.
3. The ECS will automatically restart after the OS has been reinstalled.
4. Custom settings such as DNS and hostname need to be configured again after the OS has been reinstalled.

Image	CentOS 7.5 64bit
	System Disk Capacity: 40 GiB OS Architecture: 64-bit
	<input type="checkbox"/> Encrypted ?
Login Mode	Password Key pair
Password	Keep your password secure. The system cannot detect your password.
	<input style="width: 100%;" type="password" value="....."/>
Confirm Password	<input style="width: 100%;" type="password" value="....."/>

OK Cancel

7. Click **OK**.

8. 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 **3** to **8** again to retry the OS installation.

If the second reinstallation attempt still fails, contact customer service for manual recovery at the backend.

3.8.2 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 or key. Otherwise, the login will fail.
- The ECS for which you want to change the OS must be in any of the following states: **Stopped**, **Reinstallation failed**, or **Failed to change the OS**.
- 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.
- The system disk can be encrypted.
- Windows and Linux cannot be changed to each other.
- For details about the change between different OSs, see [Notes on Change Between Windows and Linux](#).
- If an ECS OS is to be changed using a full-ECS image, the ECS system disk can be encrypted.
- The boot mode (BIOS or UEFI) cannot be changed.

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. For details, see [Backing Up an ECS](#).
- 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.
- You can choose to encrypt the system disk of an ECS during OS change.

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.
For details, see [Backing Up an ECS](#).
- If you want to change the login authentication mode from password to key pair during the OS change, create a key file in advance.
For details, see [\(Recommended\) Creating a Key Pair on the Management Console](#).
- 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* [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. Click  in the upper left corner and select your region and project.
3. Under **Computing**, choose **Elastic Cloud Server**.
4. Locate the row containing the target ECS and choose **More > Manage Image/Disk/Backup > Change OS** in the **Operation** column.
Only stopped ECSs support OS change. If the ECS is not stopped, stop it before proceeding with changing.
5. Select the target image.
6. (Optional) Select the **Encryption** option to encrypt the system disk during OS change.

To enable encryption, click **Create Agency** to assign KMS access permissions to EVS. If you have rights granting permission, assign the KMS access permissions to EVS. If you do not have the permission, contact the user having the security administrator rights to assign the KMS access permissions. For details, see [Can All Users Use the Encryption Feature?](#)

Encryption parameters are as follows:

- **Encryption**: indicates that the EVS disk has been encrypted.
- **Create Agency**: assigns KMS access permissions to EVS to obtain KMS keys. After the permissions are assigned, follow-up operations do not require assigning permissions again.

- **Agency Name:** set to **EVSAccessKMS**, which means that permissions have been assigned to EVS to obtain KMS keys for encrypting or decrypting EVS disks.
 - **KMS Key Name:** specifies the name of the key used by the encrypted EVS disk. You can select an existing key, or click **Create KMS Key** and create a new one on the KMS console. The default value is **evs/default**.
 - **KMS Key ID:** specifies the ID of the key used by the encrypted data disk.
7. Configure the login mode.
If the target ECS uses key pair authentication, you can replace the original key pair.
 8. Click **OK**.
 9. In the **Change OS** dialog box, confirm the specifications, and click **OK**.
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 **3** to **9** again to retry the OS change.
- If the second OS change attempt is unsuccessful, contact customer service for manual recovery at the backend.

3.9 Viewing ECS Information

3.9.1 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. Click  in the upper left corner and select your region and project.
3. Under **Computing**, click **Elastic Cloud Server**.
4. After creating an ECS, view the creation status above the ECS list beside the common operations (**Start**, **Stop**, **Restart**, and **More**).

Figure 3-106 ECS creation status



5. 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?](#)

3.9.2 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 3-18](#) lists the types of failures that can be recorded in the **Failures** area.

Table 3-18 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 .
Operation failures	<ul style="list-style-type: none">Modifying ECS specifications If an ECS specifications modification failed, this operation is recorded in Failures .

Procedure



- Log in to the management console.
- Click  in the upper left corner and select your region and project.
- Under **Computing**, choose **Elastic Cloud Server**.
- View **Failures** on the right side of common operations.

Figure 3-107 Failures

- 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.


3.9.3 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

- Log in to the management console.
- Click  in the upper left corner and select your region and project.



3. 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.
4. 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**.
5. Click the name of the target ECS.
The page providing details about the ECS is displayed.
6. View the ECS details.
You can click the tabs and perform operations. For details, see [Changing a Security Group](#), [Attaching a Network Interface](#), [Adding Tags](#), and [Binding an EIP](#).

3.9.4 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. Click  in the upper left corner and select your region and project.
3. Under **Computing**, choose **Elastic Cloud Server**.
4. 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.

5. In the default download path, view the exported ECS information.

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.
- Encrypted images cannot be shared.
- Only the full-ECS images created using CBR can be shared.

Helpful Links

- [Creating a Private Image](#)

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. Click  in the upper left corner and select your region and project.
3. Under **Computing**, click **Elastic Cloud Server**.
4. In the ECS list, choose **More** > **Manage Image** > **Create Image** in the **Operation** column.
5. 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 System disk image .
Source	Click the ECS tab and select an ECS with required configurations.

Table 4-3 Image information

Parameter	Description
Encryption	This parameter specifies whether the image will be encrypted. The value is provided by the system and cannot be changed. <ul style="list-style-type: none">• Only an unencrypted private image can be created from an unencrypted ECS.• Only an encrypted private image can be created from an encrypted ECS.
Name	Set a name for the image.

Parameter	Description
Enterprise Project	Select an enterprise project from the drop-down list. This parameter is available only if you have enabled enterprise projects or your account is an enterprise account. To enable this function, contact your customer manager. An enterprise project provides central management of cloud resources on a project.
Tag	(Optional) Set a tag key and a tag value for the image to make identification and management of your images easier.
Description	(Optional) Enter a description of the image.

6. Click **Apply Now**.

5 EVS Disks

5.1 Overview

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](#).

Helpful Links

- [Attaching an EVS Disk to an ECS](#)
- [Introduction to Data Disk Initialization Scenarios and Partition Styles](#)
- [Why Can't I Find My Newly Purchased Data Disk After I Log In to My Windows ECS?](#)
- [How Can I Adjust System Disk Partitions?](#)
- [Can I Attach Multiple Disks to an ECS?](#)
- [What Are the Requirements for Attaching an EVS Disk to an ECS?](#)

5.2 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 **Storage**, click **Elastic Volume Service**.
3. On the **Disks** page, click **Create Disk**.
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*.

For instructions about how to set EVS disk parameters, see [Create an EVS 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](#).

5.3 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.

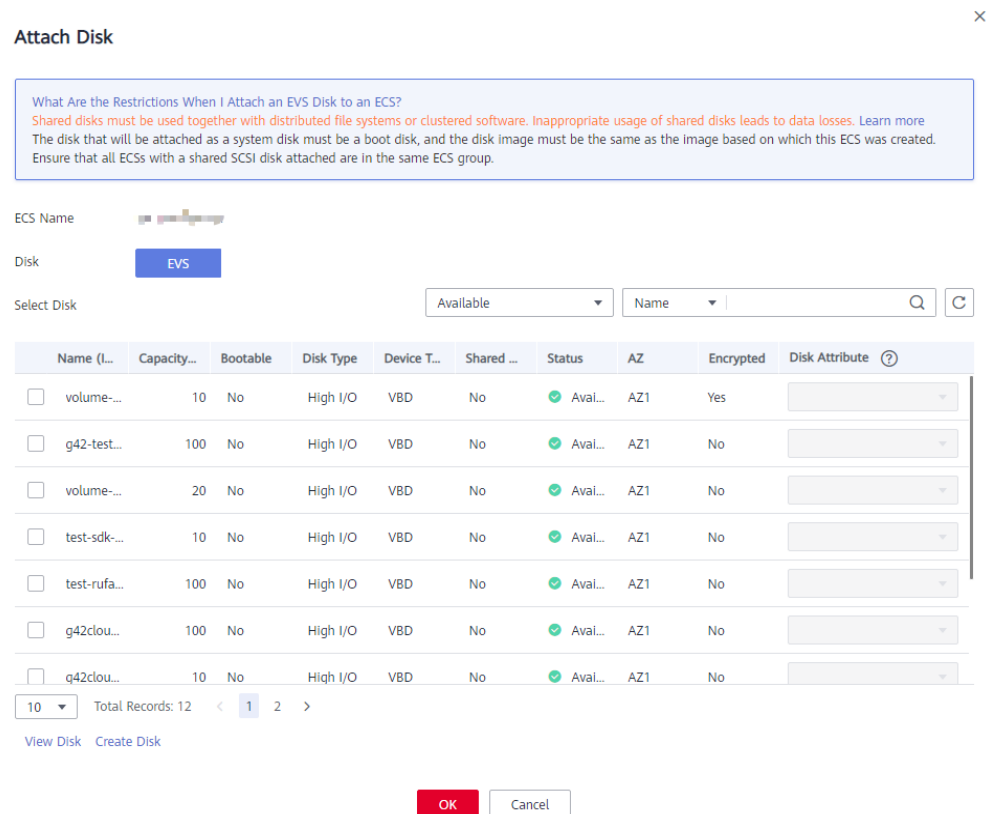
Prerequisites

- EVS disks are available.
For details about how to create an EVS disk, see [Creating an EVS Disk](#).

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Computing**, choose **Elastic Cloud Server**.
4. In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID for search.

- Click the name of the target ECS.
The page providing details about the ECS is displayed.
- Click the **Disks** tab. Then, click **Attach Disk**.
The **Attach Disk** dialog box is displayed.

Figure 5-1 Attach Disk (Xen)

- 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?](#)
- 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](#).

5.4 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 [External Image File Formats and Supported OSs](#).
- [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
	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

OS	Version
	14.04 64bit
	14.04.4 64bit
Windows	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**.

5.5 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:

- Apply for 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 Capacity for an In-use EVS Disk](#).

- 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 Capacity for an Available EVS Disk](#).

For details, see [Expansion Overview](#).

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.

Related Operations

For a Windows ECS, if you want to expand the disk capacity by clearing disk files, you can reduce the size of the WinSxS folder using tools built into Windows. For details, see [Clean Up the WinSxS Folder](#).

5.6 Expanding the Local Disks of a Disk-intensive ECS

Scenarios

Disk-intensive ECSs can use both local disks and EVS disks to store data. Local disks are generally used to store service data and feature higher throughput than EVS disks.

Disk-intensive ECSs do not support specifications modification. When the capacity of local disks is insufficient, you can create a new disk-intensive ECS with higher specifications for capacity expansion. The data stored in the original ECS can be migrated to the new ECS through EVS.

Procedure

1. Create an EVS disk according to the volume of data to be migrated.
2. Attach the EVS disk to the disk-intensive ECS for which you want to expand the capacity.
3. Back up the data stored in the local disks to the EVS disk that is newly attached to the disk-intensive ECS.
4. Detach the EVS disk from the ECS.
 - a. On the **Elastic Cloud Server** page, select this disk-intensive ECS and ensure that it has been stopped.
If the ECS is running, choose **More > Stop** to stop the ECS.
 - b. Click the name of the disk-intensive ECS. The page providing details about the ECS is displayed.
 - c. Click the **Disks** tab. Locate the row containing the EVS data disk and click **Detach** to detach the disk from the ECS.
5. Ensure that a new disk-intensive ECS with higher specifications than the original one is available.
The local disk capacity is sufficient enough to meet your requirements.
6. Attach the EVS disk to the new disk-intensive ECS.

- On the **Elastic Cloud Server** page, click the name of the ECS described in step [5](#) to view details.
- Click the **Disks** tab. Then, click **Attach Disk**.
In the displayed dialog box, select the EVS disk detached in step [4](#) and the device name.
 - Migrate the data from the EVS disk to the local disks of the new disk-intensive ECS.


5.7 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

- Log in to the management console.
- Click  in the upper left corner and select your region and project.
- Under **Computing**, click **Elastic Cloud Server**.
- Click the name of the target ECS. The page providing details about the ECS is displayed.
- Click the **Disks** tab.
- 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.
- Click **OK**.
- Stop and then start the target ECS.
This operation allows advanced disk to take effect.
- 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, try again later or contact customer service.

6 Elastic Network Interface

6.1 Overview

An elastic network interface (referred to as a network interface in this documentation) is a virtual network card. You can create and configure network interfaces and attach them to your instances (ECSs and BMSs) to obtain flexible and highly available network configurations.

Network Interface Types

- A primary network interface is created together with an instance by default, and cannot be detached from the instance.
- An extended network interface is created on the **Network Interfaces** console, and can be attached to or detached from an instance.

Application Scenarios

- Flexible migration
You can detach a network interface from an instance and then attach it to another instance. The network interface retains its private IP address, EIP, and security group rules. In this way, service traffic on the faulty instance can be quickly migrated to the standby instance, implementing quick service recovery.
- Traffic management
You can attach multiple network interfaces that belong to different subnets in a VPC to the same instance, and configure the network interfaces to carry the private network traffic, public network traffic, and management network traffic of the instance. You can configure access control policies and routing policies for each subnet, and configure security group rules for each network interface to isolate networks and service traffic.

Helpful Links

- [Will NICs Added to an ECS Start Automatically?](#)
- [How Do I Change the CIDR Block of an ECS Subnet?](#)


- [How Can I Check Whether the Network Communication Is Normal Between Two ECSs Equipped with an InfiniBand NIC Driver?](#)
- [How Can I Manually Configure an IP Address for an InfiniBand NIC?](#)
- [Why Is the NIC Not Working?](#)

6.2 Attaching a Network Interface

Scenarios

If your ECS requires multiple network interfaces (NICs), you can attach them to your ECS.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Computing**, click **Elastic Cloud Server**.
4. Click the name of the target ECS.
The page providing details about the ECS is displayed.
5. On the **Network Interfaces** tab, click **Attach Network Interface**.
6. 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.
7. 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 customer service.

6.3 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.

6.4 Changing a VPC

Scenarios

This section describes how to change a VPC.

Constraints

- Only running or stopped ECSs support VPC change.
- A VPC can be changed for an ECS only if the ECS has one NIC.
- If you have reinstalled or changed the OS of an ECS before changing the VPC, log in to the ECS and check whether the password or key pair configured during the reinstallation or change is successfully injected.
 - If the login is successful, the password or key pair is injected. Perform operations as required.
 - Otherwise, the system is injecting the password or key pair. During this period, do not perform any operations on the ECS.
- During the VPC switchover, do not bind, unbind, or replace the EIP. Otherwise, a message indicating insufficient permissions will be displayed, but you do not need to take any action.
- If an ECS NIC has an IPv6 address, the VPC of the ECS cannot be changed.

Notes

- A VPC can be changed on a running ECS, but the ECS network connection will be interrupted during the change process.

NOTE

If you intend to change the VPC for a running ECS, the VPC change may fail when traffic is routed to the ECS NIC. In this case, you are advised to try again later or stop the ECS first and then try to change the VPC.

- After the VPC is changed, the subnet, private IP address, MAC address, and OS NIC name of the ECS will change.
- After the VPC is changed, the source/destination check and virtual IP address must be configured again.
- After the VPC is changed, you are required to reconfigure network-related application software and services, such as ELB, VPN, NAT, and DNS.

Prerequisites

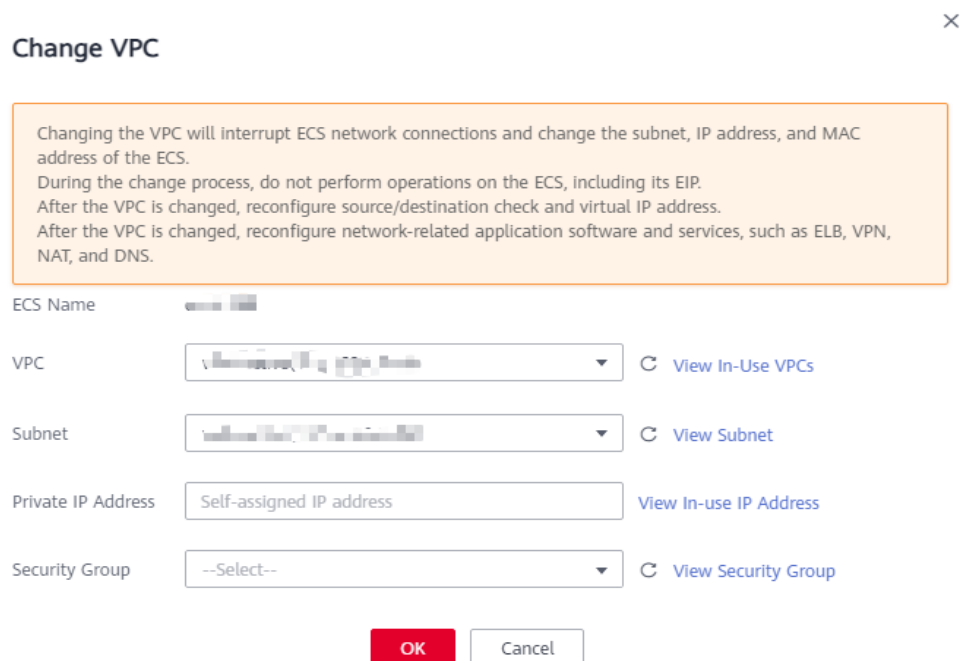
The target VPC, subnet, private IP address, and security group are available.

Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. In the ECS list, locate the row that contains the target ECS. Click **More** in the **Operation** column and select **Manage Network > Change VPC**.

The **Change VPC** dialog box is displayed.

Figure 6-1 Change VPC



Change VPC ×

Changing the VPC will interrupt ECS network connections and change the subnet, IP address, and MAC address of the ECS.
During the change process, do not perform operations on the ECS, including its EIP.
After the VPC is changed, reconfigure source/destination check and virtual IP address.
After the VPC is changed, reconfigure network-related application software and services, such as ELB, VPN, NAT, and DNS.

ECS Name:

VPC: [View In-Use VPCs](#)

Subnet: [View Subnet](#)

Private IP Address: [View In-use IP Address](#)

Security Group: [View Security Group](#)

4. Select an available VPC and subnet from the drop-down lists, and set the private IP address and security group as prompted.

You can select multiple security groups. In such a case, the access rules of all the selected security groups will apply to the ECS.

 NOTE

Using multiple security groups may deteriorate ECS network performance. You are suggested to select no more than five security groups.

5. Click **OK**.

6.5 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.
- To change the private IP address for a backend server of a load balancer, remove the backend server from the backend server group first.

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.

6.6 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**.

Logging In to the ECS to Configure Virtual IP Addresses

After you bind one or more virtual IP addresses to an ECS on the console, you must log in to the ECS to manually configure these virtual IP address.

The following OSs are used as examples here. For other OSs, see the help documents on their official websites.

- Linux: CentOS 7.2 64bit and Ubuntu 22.04 server 64bit
- Windows: Windows Server

Linux (CentOS 7.2 64bit is used as an example.)

1. Obtain the NIC that the virtual IP address is to be bound and the connection of the NIC:

nmcli connection

Information similar to the following is displayed:

```
[172.16.0.247 ~]# nmcli connection
NAME                UUID                                  TYPE      DEVICE
Wired connection 1  5e72ec5a-6165-3bd6-a34b-ce43981acb27  ethernet  eth0
docker0             cd351a91-c5eb-4b69-83eb-df092a2ccf6b  bridge    docker0
```

The command output in this example is described as follows:

- **eth0** in the **DEVICE** column indicates the NIC that the virtual IP address is to be bound.
 - **Wired connection 1** in the **NAME** column indicates the connection of the NIC.
2. Add the virtual IP address for the connection:

nmcli connection modify "Connection name of the NIC" +ipv4.addresses
Virtual IP address

Configure the parameters as follows:

- *Connection name of the NIC*: The connection name of the NIC obtained in **1**. In this example, the connection name is **Wired connection 1**.
- *Virtual IP address*: Enter the virtual IP address to be added. If you add multiple virtual IP addresses at a time, separate every two with a comma (,).

Example commands:

- Adding a single virtual IP address: **nmcli connection modify "Wired connection 1" +ipv4.addresses 172.16.0.125**
- Adding multiple virtual IP addresses: **nmcli connection modify "Wired connection 1" +ipv4.addresses 172.16.0.125,172.16.0.126**

3. Make the configuration in 2 take effect:

nmcli connection up "Connection name of the NIC"

In this example, run the following command:

nmcli connection up "Wired connection 1"

Information similar to the following is displayed:

```
root@ecs-X-ubuntu:~# nmcli connection up "Wired connection 1"
Connection successfully activated (D-Bus active path: /org/freedesktop/NetworkManager/ActiveConnection/6)
```

4. Check whether the virtual IP address has been bound:

ip a

Information similar to the following is displayed. In the command output, the virtual IP address 172.16.0.125 is bound to NIC eth0.

```
1: 172.16.0.247_subnet0-ecs-pod6-gaea-dpdk-ipv6 ~]# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default qlen 1000
    link/ether fa:16:3e:e5:d5:cd brd ff:ff:ff:ff:ff:ff
    inet 172.16.0.247/24 brd 172.16.0.255 scope global noprefixroute dynamic eth0
        valid_lft 86398sec preferred_lft 86398sec
    inet 172.16.0.125/32 brd 172.16.0.125 scope global noprefixroute eth0
        valid_lft forever preferred_lft forever
    inet6 2001:db8:a583:62c:7dd3:a19a:4031:d6fb/128 scope global tentative noprefixroute dynamic
        valid_lft 86400sec preferred_lft 86400sec
    inet6 fe80::5371:9bf9:b652:e35b/64 scope link noprefixroute
        valid_lft forever preferred_lft forever
```

NOTE

To delete an added virtual IP address, perform the following steps:

1. Delete the virtual IP address from the connection of the NIC:

nmcli connection modify "Connection name of the NIC" -ipv4.addresses *Virtual IP address*

To delete multiple virtual IP addresses at a time, separate every two with a comma (,). Example commands are as follows:

- Deleting a single virtual IP address: **nmcli connection modify "Wired connection 1" -ipv4.addresses 172.16.0.125**
- Deleting multiple virtual IP addresses: **nmcli connection modify "Wired connection 1" -ipv4.addresses 172.16.0.125,172.16.0.126**

2. Make the deletion take effect by referring to 3.

Linux (Ubuntu 22.04 server 64bit is used as an example.)

If an ECS runs Ubuntu 22 or Ubuntu 20, perform the following operations:

1. Obtain the NIC that the virtual IP address is to be bound:

ifconfig

Information similar to the following is displayed. In this example, the NIC bound to the virtual IP address is **eth0**.

```
root@ecs-X-ubuntu:~# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 172.16.0.210 netmask 255.255.255.0 broadcast 172.16.0.255
    inet6 fe80::f816:3eff:fe01:f1c3 prefixlen 64 scopeid 0x20<link>
```



```
ether fa:16:3e:01:f1:c3 txqueuelen 1000 (Ethernet)
RX packets 43915 bytes 63606486 (63.6 MB)
RX errors 0 dropped 0 overruns 0 frame 0
TX packets 3364 bytes 455617 (455.6 KB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
...
```

2. Switch to the **/etc/netplan** directory:

```
cd /etc/netplan
```

3. Add a virtual IP address to the NIC.

- a. Open the configuration file **01-netcfg.yaml**:

```
vim 01-netcfg.yaml
```

- b. Press **i** to enter the editing mode.

- c. In the NIC configuration area, add a virtual IP address.

In this example, add a virtual IP address for **eth0**:

```
addresses:
```

```
- 172.16.0.26/32
```

The file content is as follows:

```
network:
  version: 2
  renderer: NetworkManager
  ethernets:
    eth0:
      dhcp4: true
      addresses:
        - 172.16.0.26/32
    eth1:
      dhcp4: true
    eth2:
      dhcp4: true
    eth3:
      dhcp4: true
    eth4:
      dhcp4: true
```

- d. Press **Esc**, enter **:wq!**, save the configuration, and exit.

4. Make the configuration in **3** take effect:

```
netplan apply
```

5. Check whether the virtual IP address has been bound:

```
ip a
```

Information similar to the following is displayed. In the command output, the virtual IP address 172.16.0.26 is bound to NIC eth0.

```
root@ecs-X-ubuntu:/etc/netplan# ip a
...
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default
qlen 1000
    link/ether fa:16:3e:01:f1:c3 brd ff:ff:ff:ff:ff:ff
    altname enp0s3
    altname ens3
    inet 172.16.0.26/32 scope global noprefixroute eth0
        valid_lft forever preferred_lft forever
    inet 172.16.0.210/24 brd 172.16.0.255 scope global dynamic noprefixroute eth0
        valid_lft 107999971sec preferred_lft 107999971sec
    inet6 fe80::f816:3eff:fe01:f1c3/64 scope link
        valid_lft forever preferred_lft forever
```

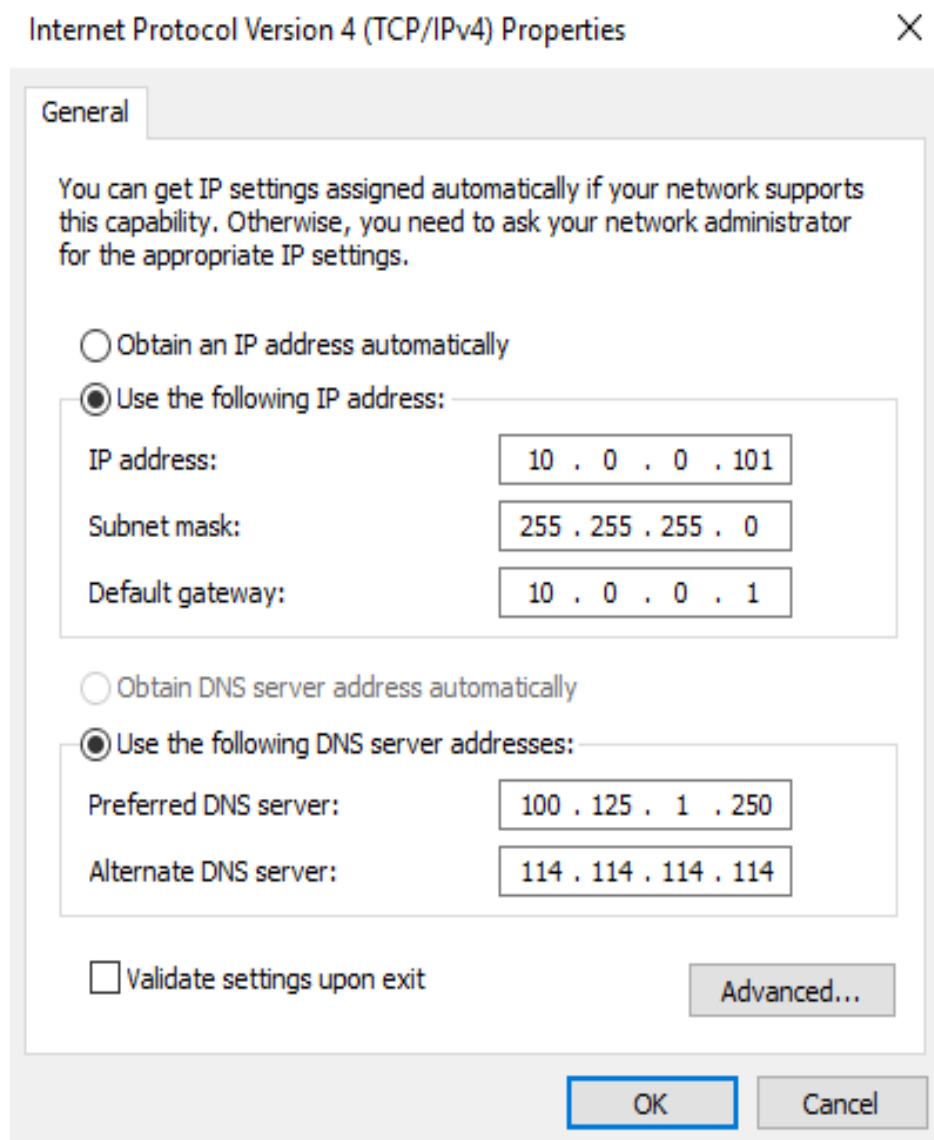
NOTE

To delete an added virtual IP address, perform the following steps:

1. Open the configuration file **01-netcfg.yaml** and delete the virtual IP address of the corresponding NIC by referring to **3**.
2. Make the deletion take effect by referring to **4**.

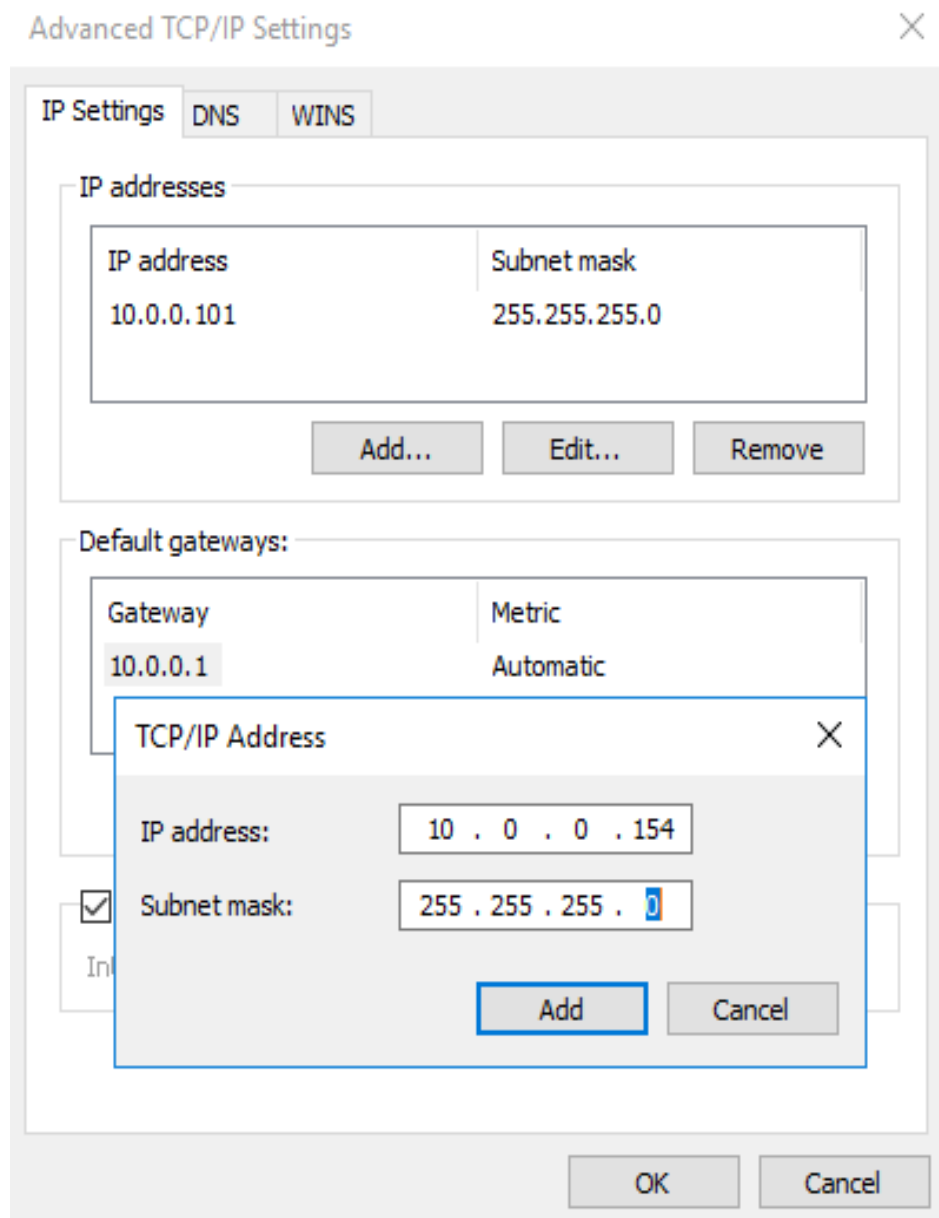
Windows OS (Windows Server is used as an example here.)

1. In **Control Panel**, click **Network and Sharing Center**, and click the corresponding local connection.
2. On the displayed page, click **Properties**.
3. On the **Network** tab page, select **Internet Protocol Version 4 (TCP/IPv4)**.
4. Click **Properties**.
5. Select **Use the following IP address** and set **IP address** to the private IP address of the ECS, for example, 10.0.0.101.

Figure 6-2 Configuring private IP address

6. Click **Advanced**.
7. On the **IP Settings** tab, click **Add** in the **IP addresses** area. Add the virtual IP address, for example, 10.0.0.154.

Figure 6-3 Configuring virtual IP address



8. Click **OK**.
9. In the **Start** menu, open the Windows command line window and run the following command to check whether the virtual IP address has been configured:

ipconfig /all

In the command output, **IPv4 Address** is the virtual IP address 10.0.0.154, indicating that the virtual IP address of the ECS NIC has been correctly configured.

6.7 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](#)

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 6-2](#) support NIC multi-queue.

NOTE

- The PV driver of a Windows ECS dynamically adjusts the number of NIC queues based on the number of vCPUs of the ECS, and you do not need to set the number of Windows NIC multi-queues.
- Public images that contain Windows Server 2008 are no longer available. However, you can still use private images that contain Windows Server 2008.
- 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 customer service to upgrade the kernel.

Table 6-1 Support of NIC multi-queue for Windows ECSs

Image	Support of NIC Multi-Queue	NIC Multi-Queue Enabled by Default
Windows Server 2008 R2 Standard/Enterprise/DataCenter 64bit	Yes	Yes
Windows Server 2008 Enterprise SP2 64bit	Yes	Yes
Windows Server 2008 Web R2 64-bit	Yes	Yes
Windows Server 2008 R2 Enterprise 64bit_WithGPUdriver	Yes	Yes
Windows Server 2012 R2 Standard 64bit_WithGPUdriver	Yes	Yes
Windows Server 2012 R2 Standard/DataCenter 64 bit	Yes	Yes
Windows Server 2016 Standard/DataCenter 64 bit	Yes	Yes
Windows Server 2019 DataCenter 64 bit	Yes	Yes

Table 6-2 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
Fedora 24/25/30 64bit	Yes	Yes
EulerOS 2.2/2.3/2.5 64bit	Yes	Yes

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*. After the image file is imported, view the value of **NIC Multi-Queue** on the page providing details about the image.

- If the value is **Supported**, go to [Creating an ECS Using a Private Image](#).
- If the value is **Not supported**, go to [Setting NIC Multi-Queue for the Image](#).

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:

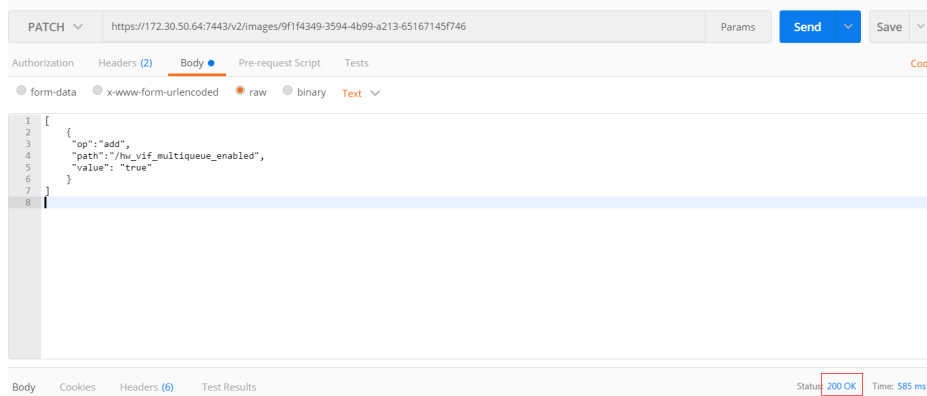
```
PATCH /v2/images/{image_id}
```

The request body is as follows:

```
[
  {
    "op": "add",
    "path": "/hw_vif_multiqueue_enabled",
    "value": "true"
  }
]
```

Figure 6-4 shows an example request body for modifying the NIC multi-queue attribute.

Figure 6-4 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.

Viewing the Number of Queues of the NIC

NIC multi-queue has been enabled.

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
```

Example:

```
[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 four queues have been enabled.
```

6.8 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:

- 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 6-3 Enabling dynamic assignment of IPv6 addresses for different OSs

OS	Automatically/ Manually Enabling	Reference
Windows Server 2012	Automatically	Windows Server 2012
Windows Server 2008	Automatically	Windows Server 2008
Linux	Automatically (recommended)	Linux (Automatically Enabling Dynamic Assignment of IPv6 Addresses)
Linux	Manually	Linux (Manually Enabling Dynamic Assignment of IPv6 Addresses)

Enabling IPv6 for an ECS

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**.

Windows Server 2012


Step 1 Check whether IPv6 is enabled for the ECS.

Run the following command in the CMD window to check it:

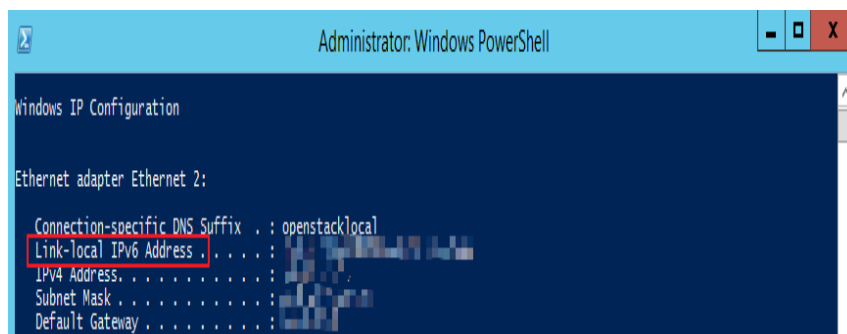
ipconfig

- If an IPv6 address and a link-local IPv6 address are displayed, IPv6 is enabled and dynamic IPv6 assignment is also enabled.

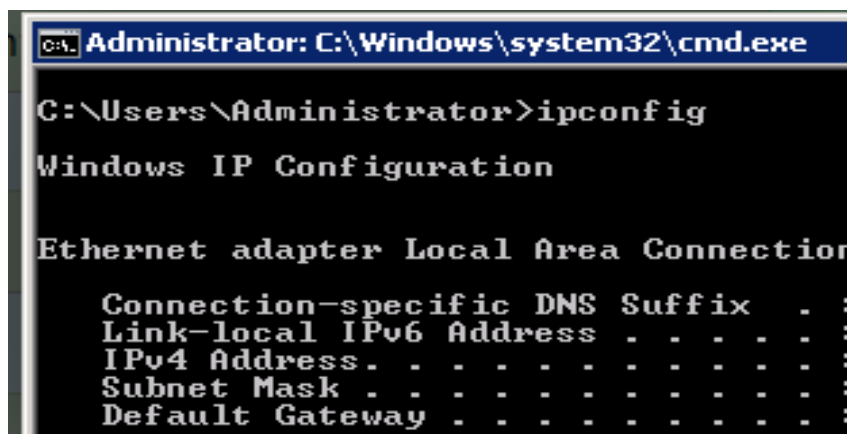
Figure 6-5 Querying the IPv6 address



- If only a link-local IPv6 address is displayed, IPv6 is enabled but dynamic IPv6 assignment is not enabled. Go to [Step 2](#).

Figure 6-6 Link-local IPv6 address


- If neither an IPv6 address nor link-local IPv6 address is displayed, IPv6 is disabled. Go to [Step 3](#).

Figure 6-7 IPv6 disabled**NOTE**

By default, dynamic IPv6 address assignment is enabled for Windows public images, as shown in [Figure 6-5](#). 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 6-8 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 6-9 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 6-10 Configuring an IPv6 address and a DNS server address

- (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
- Perform [Step 1](#) to check whether dynamic IPv6 address assignment is enabled.

----End

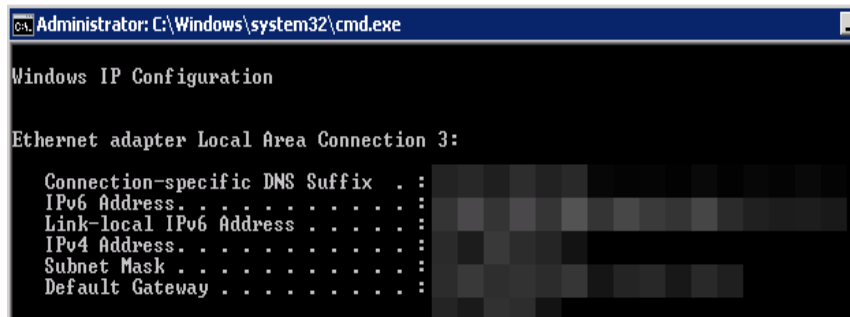
Windows Server 2008

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.

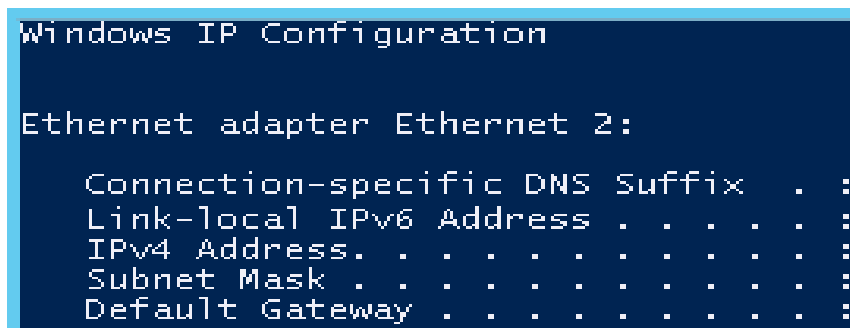
Figure 6-11 Querying the IPv6 address

```
Administrator: C:\Windows\system32\cmd.exe
Windows IP Configuration

Ethernet adapter Local Area Connection 3:

    Connection-specific DNS Suffix  . : 
    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](#).

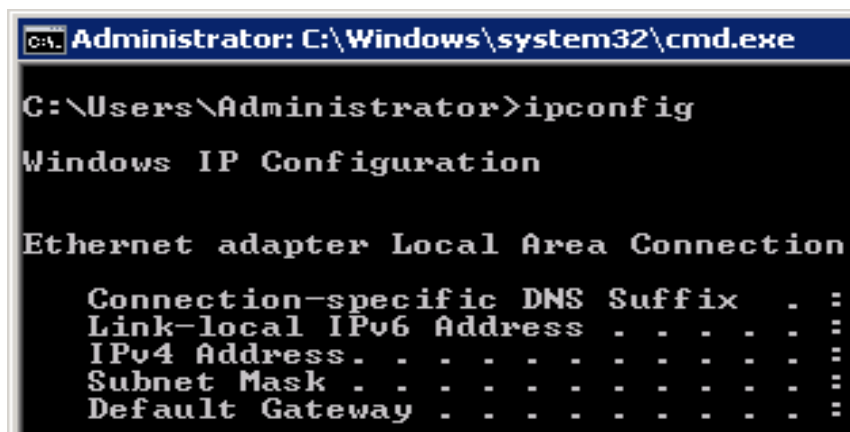
Figure 6-12 Link-local IPv6 address

```
Windows IP Configuration

Ethernet adapter Ethernet 2:

    Connection-specific DNS Suffix  . : 
    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 6-13 IPv6 disabled

```
Administrator: C:\Windows\system32\cmd.exe
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

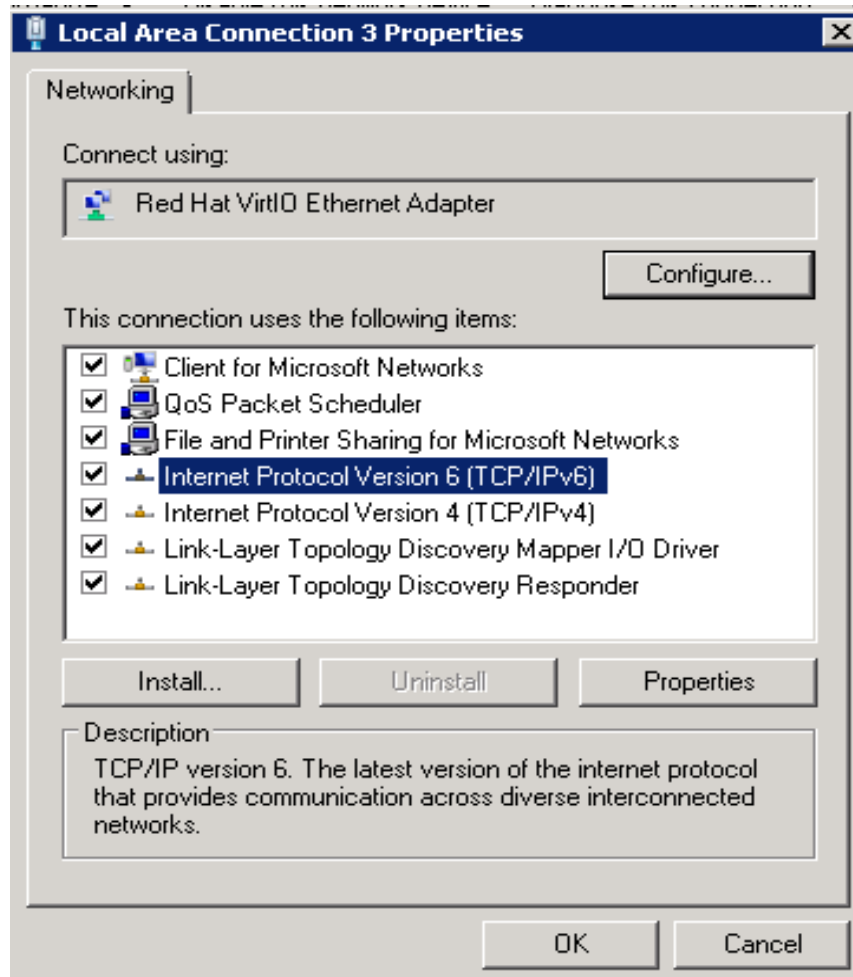
By default, dynamic IPv6 address assignment is enabled for Windows public images, as shown in [Figure 6-11](#). No additional configuration is required.

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 6-14 Configuring dynamic IPv6 address assignment




6. Perform **Step 1** to check whether dynamic IPv6 address assignment is enabled.

Step 3 Enable and configure IPv6.

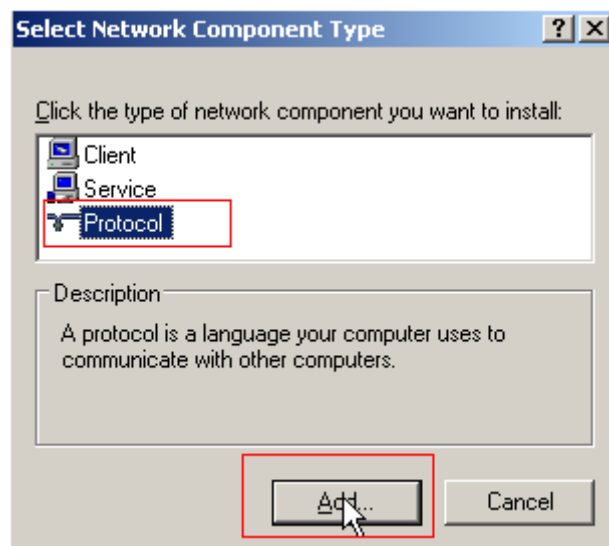
1. Choose **Start > Control Panel > Network Connection > Local Connection**.
2. Select **Properties**, select the following options, and click **Install**.

Figure 6-15 Enabling and configuring IPv6

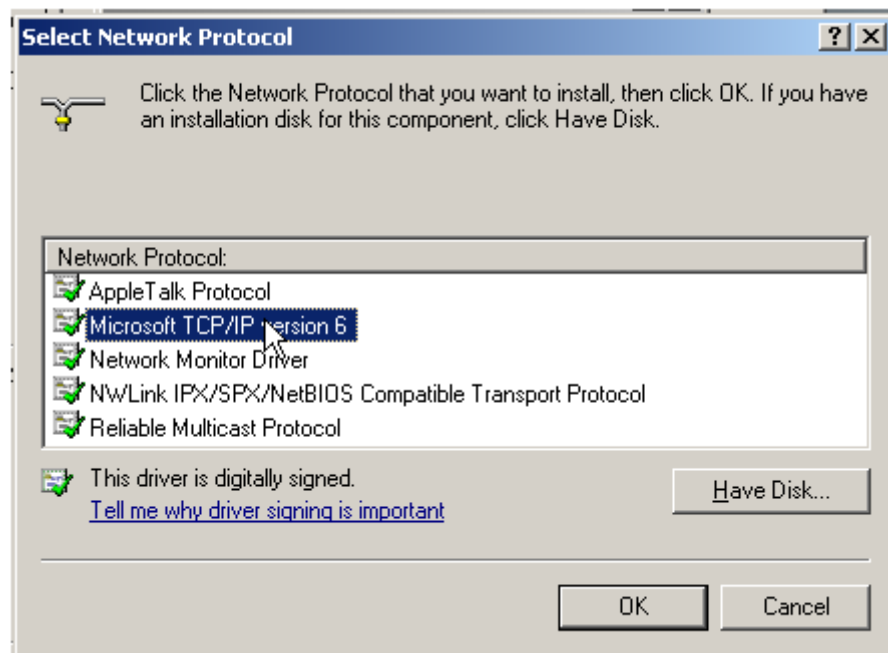


3. Select **Protocol** and click **Add**.

Figure 6-16 Adding the protocol



4. Select **Microsoft TCP/IP Version 6** and click **OK**.

Figure 6-17 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 6-18 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 6-19 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 6-20 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 6-19](#).

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.
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 6-21 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 6-22 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 6-23 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 6-22](#).

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 6-24 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 6-25 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 6-26 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 6-27 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 6-28 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 6-29 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 6-30 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 6-31 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     log_usage_msg
153     ;;
154 esac
```

Figure 6-32 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         log_action_end_msg $?
162     else
163         log_action_end_msg $?
164     fi
165     ;;
166
167 force-reload|restart)
168     if init_is_upstart; then
169         exit 1
170     fi
171     process_options
172
173     log_warning_msg "Running $0 $1 is deprecated because it may not re-enable some interfaces"
174     log_action_begin_msg "Reconfiguring network interfaces"
175     if /usr/bin/timeout 30 ifdown -a --exclude=lo $verbose || true
176     then
177         set -f
178         exclusions=$(process_exclusions)
179         if /usr/bin/timeout 30 ifup -a --exclude=lo $exclusions $verbose && ifup_hotplug $exclusions $verbose
180     then
181         log_action_end_msg $?
182     else
183         log_action_end_msg $?
184     fi
185     fi
186     ;;
187 *)
188     log_usage_msg
189     ;;
190 esac
```

- 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**.

7 EIP


7.1 Overview

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, NAT gateways or load balancers.

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

Figure 7-1 Accessing the Internet using an EIP



Helpful Links

- [Binding an EIP](#)
- [Changing an EIP](#)
- [Changing an EIP Bandwidth](#)

7.2 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. Click  in the upper left corner and select your region and project.
3. Under **Computing**, click **Elastic Cloud Server**.
4. On the **Elastic Cloud Server** page, click the name of the target ECS.
The page providing details about the ECS is displayed.
5. On the **EIPs** tab, click **Bind EIP**.

Figure 7-2 EIPs

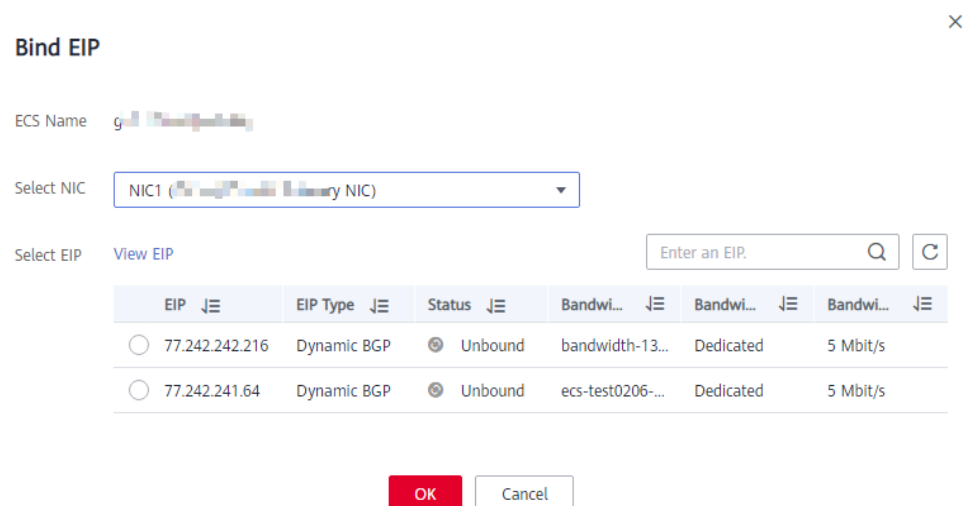


6. Select an EIP and click **OK**.

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.

Figure 7-3 Binding an EIP




View the bound EIP on the **EIPs** tab.

7.3 Unbinding an EIP

Scenarios

This section describes how to unbind an EIP from an ECS.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Computing**, click **Elastic Cloud Server**.
4. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network > Unbind EIP**.
5. 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.

7.4 Changing an EIP

Scenarios

You can change the EIP bound to your ECS as needed.

 **NOTE**

Currently, the EIP bound to the ECS cannot be directly replaced. You need to unbind the EIP first and then bind a new one to the ECS.


If there are no available EIPs, purchase one first.

Restrictions


To avoid unintended actions, the system caches the EIP that you released for 24 hours. If you change the EIP within this period, the system preferentially assigns this EIP.

If you want to purchase a new EIP and bind it to your ECS, you are advised to purchase one first before unbinding the original EIP.

Unbinding an EIP

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network > Unbind EIP**.
4. Confirm the displayed information and click **Yes**.

Binding a New EIP

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network > Bind EIP**.
4. Select the desired EIP and click **OK**.

 **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.

Figure 7-4 Binding a new EIP

Bind EIP ×

ECS Name: g-...

Select NIC: NIC1 (Primary NIC)

Select EIP: [View EIP](#)


EIP	EIP Type	Status	Bandwi...	Bandwi...	Bandwi...
<input type="radio"/> 77.242.242.216	Dynamic BGP	<input checked="" type="radio"/> Unbound	bandwidth-13...	Dedicated	5 Mbit/s
<input type="radio"/> 77.242.241.64	Dynamic BGP	<input checked="" type="radio"/> Unbound	ecs-test0206...	Dedicated	5 Mbit/s

7.5 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. Click  in the upper left corner and select your region and project.
3. Under **Computing**, click **Elastic Cloud Server**.
4. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network > Modify Bandwidth**.
5. Change the bandwidth name, billing mode, or bandwidth size as prompted.

7.6 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. Click  in the upper left corner and select your region and project.
3. Under **Computing**, click **Elastic Cloud Server**.
4. In the search box above the upper right corner of the ECS list, enter the proxy ECS name for search.
5. Click the name of the proxy ECS. The page providing details about the ECS is displayed.
6. On the **Network Interfaces** tab, click . Then, disable **Source/Destination Check**.

Figure 7-5 Disabling 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.

7. Log in to the proxy ECS.
For more details, see [Login Overview \(Linux\)](#).
8. 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 7-6 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
```

9. 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 **10**.
- If **1** (enabled) is displayed, go to **15**.

10. Run the following command to open the IP forwarding configuration file in the vi editor:

```
vi /etc/sysctl.conf
```

11. Press **i** to enter editing mode.
12. 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
```

13. Press **Esc**, type **:wq**, and press **Enter**.
The system saves the configurations and exits the vi editor.
14. Run the following command to make the modification take effect:

```
sysctl -p /etc/sysctl.conf
```

15. 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.

16. 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 16 into `rc.local`, press `Esc` to exit Insert mode, and enter `:wq` to save the settings and exit.

17. Run the following commands to save the iptables configuration and make it start up automatically upon ECS startup:

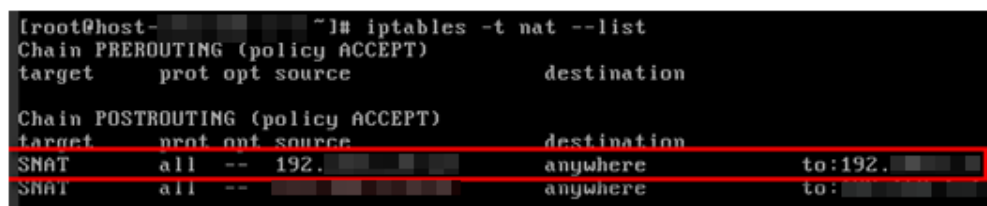
```
service iptables save  
chkconfig iptables on
```

18. Run the following command to check whether SNAT has been configured:


```
iptables -t nat --list
```

SNAT has been configured if information similar to [Figure 7-7](#) is displayed.

Figure 7-7 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 -- anywhere to:192.168.125.4
```

19. Add a route.
 - a. Log in to the management console.
 - b. Click  in the upper left corner and select your region and project.
 - c. Under **Network**, click **Virtual Private Cloud**.
 - d. 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**.

- e. 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.

20. 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
```

8 Security

8.1 Methods for Improving ECS Security

Scenarios

If ECSs are not protected, they may be attacked by viruses, resulting in data leakage or data loss.

You can use the methods introduced below to protect your ECSs from viruses or attacks.

Protection Types

ECS can be protected externally and internally.

Table 8-1 Methods for improving ECS security

Type	Description	Protection Method
External security	Trojan horses or other viruses are common external security issues. To address these issues, you can choose services such as Host Security Service (HSS) based on your service requirements:	<ul style="list-style-type: none">• Enabling HSS• Monitoring ECSs• Backing Up Data Periodically
Internal security	Weak passwords and incorrect ports opening may cause internal security issues. Improving the internal security is the key to improving the ECS security. If the internal security is not improved, external security solutions cannot effectively intercept and block various external attacks.	<ul style="list-style-type: none">• Enhancing the Login Password Strength• Improving the Port Security• Periodically Upgrading the Operating System

Enabling HSS

Host Security Service (HSS) is designed to improve the overall security for ECSs. It helps you identify and manage the assets on your servers, eliminate risks, and defend against intrusions and web page tampering. There are also advanced protection and security operations functions available to help you easily detect and handle threats.

Before using the HSS service, install the HSS agent on your ECSs first and you will be able to check the ECS security status and risks in a region on the HSS console.

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.

Server monitoring includes basic monitoring, OS monitoring, and process monitoring for servers.

- Basic monitoring

Basic monitoring does not require the agent to be installed and automatically reports ECS metrics to Cloud Eye. Basic monitoring for KVM ECSs is performed every 5 minutes.

- OS monitoring

By installing the Agent on an ECS, OS monitoring provides system-wide, active, and fine-grained monitoring. OS monitoring for KVM ECSs is performed every minute.

To enable OS monitoring for a created ECS:

You need to manually install the agent.


For instructions about how to install and configure the Agent, see [Agent Installation and Configuration](#).

- Process monitoring

Process monitoring provides monitoring of active processes on ECSs and it requires the Agent to be installed on the ECSs to be monitored. Processes are monitored at an interval of 1 minute (for KVM ECSs).

After server monitoring is enabled, you can set ECS alarm rules to customize the monitored objects and notification policies and learn about the ECS running status at any time.

On the ECS console, click  to view monitoring metrics.

Figure 8-1 Viewing ECS metrics

Backing Up Data Periodically

Data backup is a process of storing all or part of data in different ways to prevent data loss. The following uses Cloud Backup and Recovery (CBR) as an example. For more backup methods, see [Overview](#).

CBR enables you to back up ECSs 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.

You can use the cloud server backup and cloud disk backup to [back up your ECS data](#).

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

Enhancing the Login Password Strength

Key pair authentication is recommended because it is more secure than password-based authentication. If you select the password-based authentication, ensure that the password meets the strength requirements listed in [Table 8-2](#) to prevent malicious attacks.

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

The password must conform to the following rules:

- The password must consist of at least 10 characters.
- Do not use easily guessed passwords (for example, passwords in common rainbow tables or passwords with adjacent keyboard characters). The password must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters.
- Do not include accounts in passwords, such as administrator, test, root, oracle, and mysql.
- Change the password at least every 90 days.
- Do not reuse the latest five passwords.

- Set different passwords for different applications. Do not use the same password for multiple applications.

Table 8-2 Password complexity requirements

Parameter	Requirement
Password	<ul style="list-style-type: none">• Consists of 8 to 26 characters.• 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.

Improving the Port Security

You can use security groups to protect the network security of your ECSs. A security group controls inbound and outbound traffic for your ECSs. Inbound traffic originates from the outside to the ECS, while outbound traffic originates from the ECS to the outside.

You can configure security group rules to grant access to or from specific ports. You are advised to disable high-risk ports and only enable necessary ports.

Table 8-3 lists common high-risk ports. You are advised to change these ports to non-high-risk ports.

Table 8-3 Common high-risk ports

Protocol	Port
TCP	42, 135, 137, 138, 139, 444, 445, 593, 1025, 1068, 1434, 3127, 3128, 3129, 3130, 4444, 4789, 5554, 5800, 5900, and 9996
UDP	135 to 139, 1026, 1027, 1028, 1068, 1433, 1434, 4789, 5554, and 9996

Periodically Upgrading the Operating System

After ECSs are created, you need to maintain and periodically upgrade the operating system.

8.2 Security Groups

8.2.1 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. For details, see [VPC Connectivity](#).

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](#).

Notes and 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.
- A security group can have no more than 6,000 instances associated, or its performance will deteriorate.
- For inbound security group rules, the sum of the rules with **Source** set to **Security group**, the rules with **Source** set to **IP address group**, and the rules with inconsecutive ports, cannot exceed 128. Outbound rules also have this restriction.
 - When **Source** is set to **Security group**, you can select the current security group or a different security group.
 - An example of inconsecutive ports is 22,25,27.

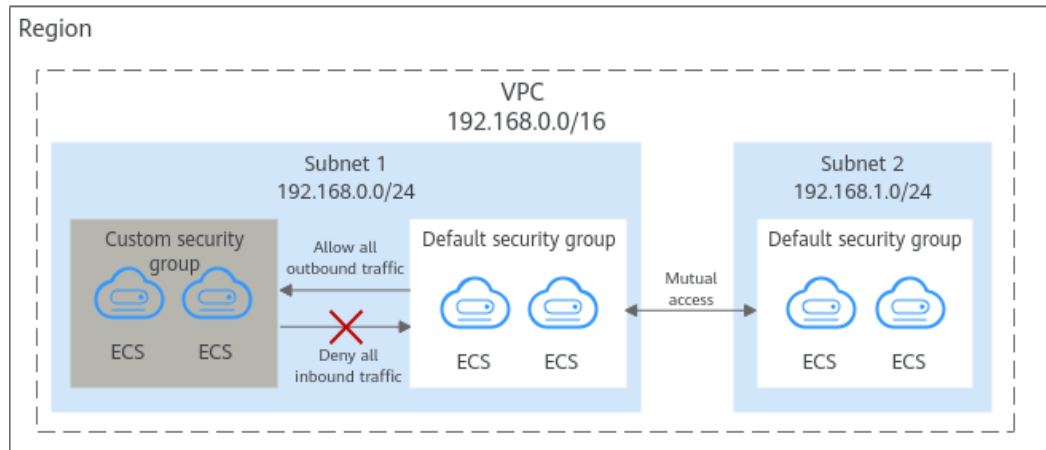
8.2.2 Default Security Group and Rules

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 8-2 shows the default security group.

Figure 8-2 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 8-4 describes the rules in the default security group.

Table 8-4 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).
Inbound	TCP	22	Source: 0.0.0.0/0	Allows all IP addresses to access Linux ECSs over SSH.

Direction	Protocol	Port/Range	Source/Destination	Description
Inbound	TCP	3389	Source: 0.0.0.0/0	Allows all IP addresses to access Windows ECSs over RDP.

8.2.3 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 8-5](#).

Table 8-5 Default outbound rules in a security group

Direction	Type	Protocol & Port	Destination	Description
Outbound	IPv4	All	0.0.0.0/0	Allows the instances in the security group to access any IPv4 address over any port.

Direction	Type	Protocol & Port	Destination	Description
Outbound	IPv6	All	::/0	Allows the instances in the security group to access any IPv6 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 8-6](#).
- To remotely log in to a Windows ECS using RDP, enable port 3389. For details, see [Table 8-7](#).

Table 8-6 Remotely logging in to a Linux ECS using SSH

Direction	Type	Protocol & Port	Source
Inbound	IPv4	TCP: 22	IP address: 0.0.0.0/0

Table 8-7 Remotely logging in to a Windows ECS using RDP

Direction	Type	Protocol & Port	Source
Inbound	IPv4	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 8-8](#).

Table 8-8 Remotely logging in to an ECS using a specified IP address

ECS Type	Direction	Type	Protocol & Port	Source
Linux ECS	Inbound	IPv4	TCP: 22	IP address: 192.168.0.0/24
Windows ECS	Inbound	IPv4	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 8-9 Remotely connecting to an ECS from a local server to upload or download files

Direction	Type	Protocol & Port	Source
Inbound	IPv4	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 8-10 Setting up a website on an ECS to provide services externally

Direction	Type	Protocol & Port	Source
Inbound	IPv4	TCP: 80	IP address: 0.0.0.0/0
Inbound	IPv4	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 8-11 Using **ping** command to verify network connectivity

Direction	Type	Protocol & Port	Source
Inbound	IPv4	ICMP: All	IP address: 0.0.0.0/0
Inbound	IPv6	ICMP: All	IP address: ::/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 8-12 Enabling communications between instances in different security groups

Direction	Type	Protocol & Port	Source
Inbound	IPv4	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 8-13 Allowing external instances to access the database deployed on an ECS

Direction	Type	Protocol & Port	Source	Description
Inbound	IPv4	TCP: 3306	Security group: sg-A	Allows the ECSs in security group sg-A to access the MySQL database service.
Inbound	IPv4	TCP: 1521	Security group: sg-B	Allows the ECSs in security group sg-B to access the Oracle database service.
Inbound	IPv4	TCP: 1433	IP address: 172.16.3.21/32	Allows the ECS whose private IP address is 172.16.3.21 to access the MS SQL database service.
Inbound	IPv4	TCP: 5432	IP address: 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.
Inbound	IPv4	TCP: 6379	IP address group: ipGroup-A	Allows ECSs whose private IP addresses are in IP address group ipGroup-A to access the Redis 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 8-15](#) 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 8-14 Allowing ECSs to access specific external websites

Direction	Type	Protocol & Port	Destination	Description
Outbound	IPv4	TCP: 80	IP address: 132.15.XX.XX	Allows ECSs in the security group to access the external website at http://132.15.XX.XX:80 .

Direction	Type	Protocol & Port	Destination	Description
Outbound	IPv4	TCP: 443	IP address: 145.117.XX.XX	Allows ECSs in the security group to access the external website at https://145.117.XX.XX:443.

2. Delete the original outbound rules that allow all traffic.

Table 8-15 Default outbound rules in a security group

Direction	Type	Protocol & Port	Destination	Description
Outbound	IPv4	All	0.0.0.0/0	Allows the instances in the security group to access any IPv4 address over any port.
Outbound	IPv6	All	::/0	Allows the instances in the security group to access any IPv6 address over any port.

8.2.4 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](#).

For details about the default security group rules, see [Default Security Groups and Security Group Rules](#). 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 8-16 Inbound rule parameter description

Parameter	Description	Example Value
Type	Source IP address version. You can select: <ul style="list-style-type: none"> • IPv4 • IPv6 	IPv4
Protocol & Port	The network protocol used to match traffic in a security group rule. The value can be All , TCP , UDP , GRE , and ICMP .	TCP
	Destination port used to match traffic in a security group rule. The value can be from 1 to 65535. Inbound rules control incoming traffic over specific ports to instances in the security group.	22, or 22-30
Source	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. If you select IP address for Source , you can enter multiple IP addresses in the same IP address box. Each IP address represents a different security group rule. <ul style="list-style-type: none"> • IP address: <ul style="list-style-type: none"> - Single IP address: 192.168.10.10/32 - All IP addresses: 0.0.0.0/0 - IP address range: 192.168.1.0/24 If the source is a security group, this rule will apply to all instances associated with the selected security group.	0.0.0.0/0
Description	Supplementary information about the security group rule. This parameter is optional. The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	N/A

7. Configure required parameters.
You can click + to add more outbound rules.

Table 8-17 Outbound rule parameter description

Parameter	Description	Example Value
Type	Destination IP address version. You can select: <ul style="list-style-type: none">• IPv4• IPv6	IPv4
Protocol & Port	The network protocol used to match traffic in a security group rule. The value can be All , TCP , UDP , GRE , and ICMP .	TCP
	Destination port used to match traffic in a security group rule. The value can be from 1 to 65535. Outbound rules control outgoing traffic over specific ports from instances in the security group.	22, or 22-30
Destination	Destination of the security group rule. The value can be an IP address or a security group to allow access to IP addresses or instances in the security group. For example: If you select IP address for Destination , you can enter multiple IP addresses in the same IP address box. Each IP address represents a different security group rule. <ul style="list-style-type: none">• IP address:<ul style="list-style-type: none">- Single IP address: 192.168.10.10/32- All IP addresses: 0.0.0.0/0- IP address range: 192.168.1.0/24	0.0.0.0/0
Description	Supplementary information about the security group rule. This parameter is optional. The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	N/A

8. Click **OK** to complete the security rule configuration.

8.2.5 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

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. 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.
4. Select the target NIC and security groups.

You can select multiple security groups. In such a case, the rules of all the selected security groups will apply to the ECS.

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.

5. Click **OK**.

8.3 HSS

What Is HSS?

Host Security Service (HSS) helps you identify and manage the assets on your servers, eliminate risks, and defend against intrusions and web page tampering. There are also advanced protection and security operations functions available to help you easily detect and handle threats.

After installing the HSS agent on your ECSs, you will be able to check the ECS security status and risks in a region on the HSS console.

How Do I Check Host Security Statuses?

On the **Server** tab, you can view the ECS security statuses in the current region.


1. Log in to the management console.
2. Click  and choose **Security > Host Security Service**.
3. Go to the **Servers** page to view the protection status of the target servers.

Figure 8-3 ECS list





Table 8-18 Statuses

Parameter	Description
Agent Status	<ul style="list-style-type: none">● Not installed: The agent has not been started or even has not been installed.● Online: The agent is running properly.● Offline: The agent fails to communicate with the HSS server. Therefore, HSS cannot protect your ECS. Click Offline. Then, the ECSs with agent being offline and the offline reasons are displayed.
Protection Status	<ul style="list-style-type: none">● Enabled: The ECS is properly protected using HSS.● Disabled: HSS has been disabled on the ECS. If an ECS does not need protection, disable HSS on it to reduce its resource consumption.
Detection Result	<ul style="list-style-type: none">● Risky: The ECS is risky.● Safe: No risks are detected.● Pending risk detection: HSS is not enabled for the ECS.

8.4 Project and Enterprise Project

Creating a Project and Assigning Permissions

- **Creating a project**
 - a. Log in to the management console.
 - b. Click . Under **Management & Deployment**, choose **Identity and Access Management**.
 - c. On the IAM console, choose **Projects** from the navigation pane, and click **Create Project**.
 - d. Select a region in which you want to create a subproject.
 - e. Enter a project name.
 - f. Click **OK**.
- **Assigning permissions**

You can assign permissions (of resources and operations) to user groups to associate projects with user groups. You can add users to a user group to control projects that users can access and the resources on which users can perform operations. To do so, perform the following operations:

 - a. On the **User Groups** page of the IAM console, locate the target user group and click **Authorize** in the **Operation** column.
 - b. Select policies or roles from the list.
 - c. Click **Next** and select **Region-specific projects**.
 - d. In the displayed regional project list, select one or more projects and click **OK**.

- e. On the **Users** page, locate the target user and click **Authorize** in the **Operation** column.
- f. Select **Inherit permissions from user groups** and select the user group authorized in **a**.
- g. Click **OK**.

Creating an Enterprise Project and Assigning Permissions

- **Creating an enterprise project**

On the management console, click the username in the upper right corner and choose **Enterprise Management** from the drop-down list. On the **Enterprise Project Management** console, click **Create Enterprise Project**.

 **NOTE**

Enterprise is available on the management console only if you have enabled the enterprise project, or your account is the primary account. To enable this function, contact customer service.

- **Assigning permissions**

You can add a user group to an enterprise project and configure a policy to associate the enterprise project with the user group. You can add users to a user group to control projects users can access and the resources on which users can perform operations. To do so, perform the following operations:

- a. On the **Enterprise Management** console, click the name of an enterprise project to go to the enterprise project details page.
- b. On the **Permissions** page, click **Assign Permissions** on the **User Groups** tab, select user groups to be associated with the enterprise project, and then attach policies to the user groups.

- **Associating ECSs with enterprise projects**

You can use enterprise projects to manage cloud resources.

- Select enterprise projects when purchasing ECSs.

On the page for buying an ECS, select an enterprise project from the **Enterprise Project** drop-down list.

- Add ECSs to an enterprise project.

On the **Enterprise Project Management** page, you can add existing ECSs to an enterprise project.

Value **default** indicates the default enterprise project. Resources that are not allocated to any enterprise projects under your account are displayed in the default enterprise project.

For more details, see [Enterprise Management User Guide](#).

8.5 Protection for Mission-Critical Operations

Scenarios

ECS protects against mission-critical operations. If you want to perform a mission-critical operation on the management console, you must enter a credential for identity verification. You can perform the operation only after your identity is

verified. For account security, it is a good practice to enable operation protection. The setting will take effect for both the account and users under the account.

The following operations can be protected: Stop, restart, or delete an ECS; detach a disk from an ECS; unbind an EIP from an ECS.

Enabling Operation Protection

Operation protection is disabled by default. Perform the following operations to enable it:


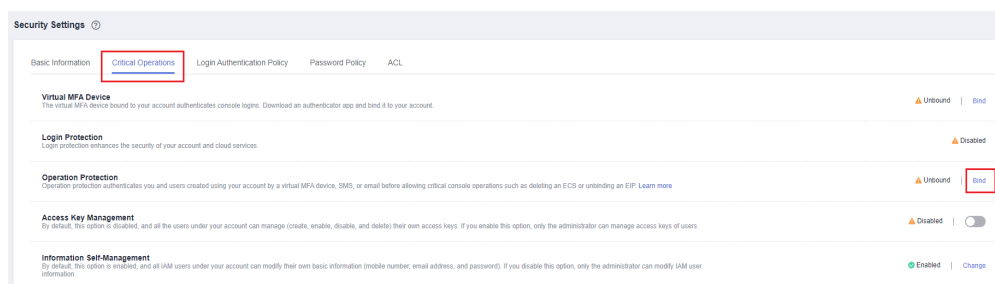
1. Log in to the management console.
2. Click . Under **Management & Deployment**, choose **Identity and Access Management**.
3. In the left navigation pane of the IAM console, choose **Security Settings**.
4. On the **Security Settings** page, choose **Critical Operations > Operation Protection > Enable**.


Figure 8-4 Critical Operations



5. On the **Operation Protection** page, select **Enable** to enable operation protection.

When you or the IAM users under your account perform critical operations, for example, deleting ECS resources, you are required to enter a verification code based on the selected verification method.

Figure 8-5 Operation Protection



- Operation Protection Enable
You and users created using your account will need to perform identity verification by using the method you specify here.
- Self-verification Verification by another person
- Disable
Identity verification will not be required for performing a critical operation.

NOTE

- When performing a critical operation, you will be asked to choose a verification method from email, SMS, and virtual MFA device.
 - If you have bound only a mobile number, only SMS verification is available.
 - If you have bound only an email address, only email verification is available.
 - If you have not bound an email address, mobile number, or virtual MFA device, you are required to bind one to continue with the critical operation.
- You can change the mobile number, email address, and virtual MFA device on the [Account Settings](#) page.

Verifying an Identity

After operation protection is enabled, when you perform a mission-critical operation, the system will verify your identity.

- If you have bound an email address, enter the email verification code.
- If you have bound a mobile number, enter the SMS verification code.
- If you have bound a virtual MFA device, enter a 6-digit dynamic verification code of the MFA device.

When you attempt to stop an ECS, select a verification method.

Figure 8-6 Identity Verification


The screenshot shows a dialog box titled "Identity Verification" with a close button (X) in the top right corner. Inside the dialog, there is an information message box with a blue border and an 'i' icon. The message reads: "You have enabled operation protection. If you do not require operation protection for critical operations, go to Security Settings > Critical Operations > Operation Protection to disable it. [Disable Identity Verification](#)". Below the message, the "Verification Method" is set to "SMS" (indicated by a selected radio button). Other options are "Email" and "Virtual MFA device", both with unselected radio buttons and a help icon (?). The "Mobile Number" field is partially obscured by a grey box, with a "Change" link to its right. The "Verification Code" section contains a text input field with "6-digit code" and a "Send Code" button. At the bottom, there are "OK" and "Cancel" buttons.

Figure 8-7 Identity Verification

The screenshot shows a dialog box titled "Identity Verification" with a close button (X) in the top right corner. Inside the dialog, there is an information message box with a blue border, an 'i' icon, and a close button (X) in its top right corner. The message reads: "The administrator has enabled operation protection. For account security purposes, verify your identity to proceed with the operation. If you do not require operation protection for critical operations, request the administrator to disable it." Below the message, the "Verification Method" is set to "Email" (indicated by a selected radio button). Other options are "SMS" and "Virtual MFA device", both with unselected radio buttons and a help icon (?). The "Email" field is partially obscured by a grey box, with a "Change" link to its right. The "Verification Code" section contains a text input field with "8-digit code" and a "Send Code" button. At the bottom, there are "OK" and "Cancel" buttons.


Disabling Operation Protection

Perform the following operations to disable operation protection.

1. Log in to the management console.
2. Click . Under **Management & Deployment**, choose **Identity and Access Management**.
3. In the left navigation pane of the IAM console, choose **Security Settings**.

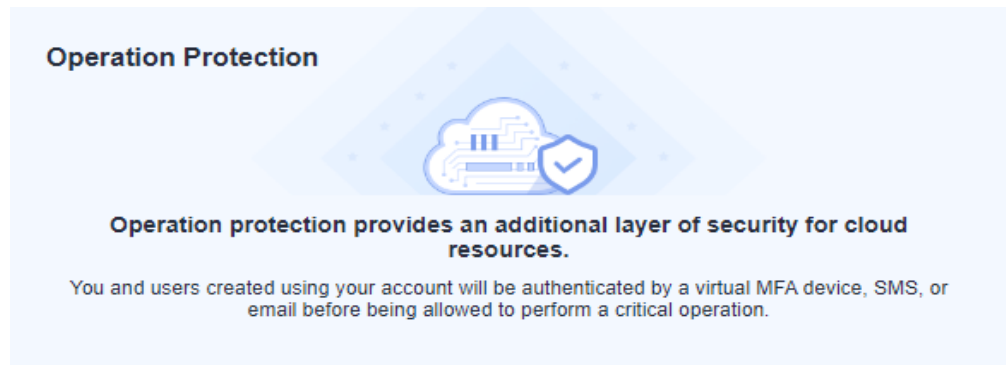
4. On the **Security Settings** page, choose **Critical Operations > Operation Protection > Change**.

Figure 8-8 Modifying operation protection settings



5. On the **Operation Protection** page, select **Disable** and click **OK**.

Figure 8-9 Disabling Operation Protection



- Operation Protection
- Enable**
You and users created using your account will need to perform identity verification by using the method you specify here.
 - Disable**
Identity verification will not be required for performing a critical operation.

Helpful Links

- [How Do I Bind a Virtual MFA Device?](#)
- [How Do I Obtain MFA Verification Codes?](#)

9 CBR

9.1 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, Snapshot, 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	Backup Method	Restoration Method
Cloud server backup	All disks (system and data disks) on an ECS	<ul style="list-style-type: none"> ● Hacker attacks and viruses 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. ● Accidental data deletion You can use cloud server backup to restore data to the backup point prior to the accidental deletion. ● Application update errors You can use cloud server backup to restore data to the backup point prior to the application update. ● System breakdown You can use cloud server backup to restore an ECS to the backup point in time prior to system breakdown. 	<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>	<p>Creating a Cloud Server Backup</p>	<ul style="list-style-type: none"> ● Restoring Data Using a Cloud Server Backup

Backup Type	Backup Object	Application Scenario	Differences and Advantages	Backup Method	Restoration Method
Cloud disk backup	One or more specified disks (system or data disks)	<ul style="list-style-type: none"> ● Only data disks need to be backed up, because the system disk does not contain users' application data. 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. ● Use backups as baseline data. 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. 	Backup data is stored in OBS, instead of disks. This ensures data restoration upon disk data loss or corruption. Backup cost is reduced without compromising data security.	Creating a Cloud Disk Backup	<ul style="list-style-type: none"> ● Restoring Data Using a Cloud Disk Backup ● Using a Backup to Create a Disk

Backup Type	Backup Object	Application Scenario	Differences and Advantages	Backup Method	Restoration Method
Snapshot	One or more specified disks (system or data disks)	<ul style="list-style-type: none"> ● Routine data backup You can create snapshots for disks on a timely basis and use snapshots to recover your data in case that data is lost or inconsistent due to unintended actions, viruses, or attacks. ● Rapid data restoration You can create a snapshot or multiple snapshots before an application software upgrade or a service data migration. If an exception occurs during the upgrade or migration, service data can be rapidly restored to the time point when the snapshot was created. For example, if ECS A cannot be started due to a fault occurred in system disk A, you can create disk B using an existing snapshot of system disk A and attach disk B to a properly running ECS, for example ECS B. In this case, ECS B can read the data of system disk A from the disk B. ● Rapid deployment of multiple services You can use a snapshot to create multiple EVS disks containing the same 	<ul style="list-style-type: none"> ● The snapshot data is stored with the disk data to facilitate rapid data backup and restoration. ● You can create snapshots to rapidly save disk data as it was at specified points in time. You can also use snapshots to create new disks so that the created disks will contain the snapshot data in the beginning 	Creating a Snapshot	Rolling Back Data from a Snapshot

Backup Type	Backup Object	Application Scenario	Differences and Advantages	Backup Method	Restoration Method
		<p>initial data, and these disks can be used as data resources for various services, for example data mining, report query, and development and testing.</p> <p>This method protects the initial data and creates disks rapidly, meeting the diversified service data requirements.</p> <p>NOTE</p> <ul style="list-style-type: none">• A snapshot can be rolled back only to its source disk. Rollback to another disk is not possible.• If you have reinstalled or changed the ECS OS, snapshots of the system disk are automatically deleted. Snapshots of the data disks can be used as usual.			

Backup Type	Backup Object	Application Scenario	Differences and Advantages	Backup Method	Restoration Method
System disk image	System disk	<ul style="list-style-type: none"> Rapid system recovery 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. Rapid deployment of multiple services You can use a system disk image to quickly create multiple ECSs with the same OS, thereby quickly deploying services these ECSs. 	A system disk image can help an ECS with OS damaged to quickly change its OS.	Creating a System Disk Image	<ul style="list-style-type: none"> Changing the OS of a Faulty ECS Using a System Disk Image Creating an ECS from a System Disk Image
Data disk image	Specific data disk	Rapid data replication 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.	Creating a Data Disk Image	Creating a Data Disk from a Data Disk Image

Backup Type	Backup Object	Application Scenario	Differences and Advantages	Backup Method	Restoration Method
Full-ECS image	All disks (system and data disks) on an ECS	<ul style="list-style-type: none">• Rapid system recovery 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.• Rapid deployment of multiple services You can use a full-ECS image to quickly create multiple ECSs with the same OS and data, thereby quickly deploying services these ECSs.	A full-ECS image facilitates service migration.	Creating a Full-ECS Image	Creating an ECS from a Full-ECS Image

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 9-1 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

Item	One-Off Backup	Periodic Backup
Backup name	User-defined backup name, which is manualbk_XXXX by default	System-assigned backup name, which is autobk_XXXX 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.

9.2 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. Click  in the upper left corner and select your region and project.
3. Under **Computing**, choose **Elastic Cloud Server**.
4. In the ECS list, locate the target ECS and choose **More > Manage Image/Backup > Create Server 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 [Creating a Server Backup Vault](#).
5. 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.
After the backup is complete, you can restore server data or create images on the **Backups** tab page. For details, see [Restoring Data Using a Cloud Server Backup](#) and [Using a Backup to Create an Image](#).

EVS Disk Backup Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Computing**, choose **Elastic Cloud Server**.
4. In the ECS list, locate the target ECS and choose **More > Manage Image/Backup > Create Disk 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 [Creating a Disk Backup Vault](#).
5. 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.

After the backup is complete, you can restore disk data on the **Backups** tab page. For details, see [Restoring Data Using a Cloud Disk Backup](#).

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
Changing the Login Password on an ECS	N/A NOTE The reference is for Windows or Linux ECSs.
Resetting the Password for Logging In to a Windows ECS	The password reset plug-in has not been installed.
Resetting the Password for Logging In to a Linux ECS	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">• Consists of 8 to 26 characters.• 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.

10.1.2 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.

Prerequisites

The ECS can be logged in.

Background

[Table 10-3](#) shows the ECS password complexity requirements.

Table 10-3 Password complexity requirements

Parameter	Requirement
Password	<ul style="list-style-type: none">• Consists of 8 to 26 characters.• 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.

Windows

1. Log in to the ECS.
For details, see [Login Overview \(Windows\)](#).
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 10-3](#)):

```
net user Administrator New password
```

Linux

1. Use the existing key file to log in to the ECS as user **root** through SSH.
For details, see [Logging In to a Linux ECS Using an SSH Key Pair](#).
2. Run the following command to reset the password of user **root**:
passwd
To reset the password of another user, replace **passwd** with **passwd username**.
3. Enter the new password as prompted. Ensure that the new password meets the requirements described in [Table 10-3](#).

```
New password:  
Retype new password:
```

```
If the following information is displayed, the password has been changed:  
passwd: all authentication tokens updates successfully
```

10.1.3 Resetting the Password for Logging In to a Windows ECS

Scenarios

You can reset your ECS password if:

- The password is forgotten.
- The password has expired.

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.

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. Click  in the upper left corner and select your region and project.
 - c. Under **Computing**, click **Elastic Cloud Server**.
 - d. 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.

- e. 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.e 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 10-4):
net user Administrator New password

Table 10-4 Password complexity requirements

Parameter	Requirement
Password	<ul style="list-style-type: none">• Consists of 8 to 26 characters.• Contains at least three of the following character types:<ul style="list-style-type: none">- Uppercase letters- Lowercase letters- Digits- Special characters for Windows: <code>!@%-_=[]:./,?</code>- Special characters for Linux: <code>!@%-_=[]:./^,{}?</code>• 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.

10.1.4 Resetting the Password for Logging In to a Linux ECS

Scenarios

Keep your password secure. Reset the password if:


- The password is forgotten.
- The password has expired.

This section describes how to reset the password of user **root**. After resetting the password, you can log in to the ECS, and change the private key or reset the password of a non-**root** user.

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. 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. Click  in the upper left corner and select your region and project.
 - c. Under **Computing**, click **Elastic Cloud Server**.
 - d. 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.

- e. 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.e and attach it to the temporary ECS.
 3. 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 3.b as prompted.

If the following information is displayed, the password has been changed:
set password success.

4. (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.

5. 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 1.
 - 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 5.b.
6. Restart the original Linux ECS.

10.2 Key Pairs

10.2.1 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*.

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.

10.2.2 (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. Click  in the upper left corner and select your region and project.
3. Under **Computing**, click **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **Key Pair**.
5. On the displayed page, click **Create Key Pair**.

6. Enter a key pair name.
A key pair name consists of two parts: KeyPair and four random digits (KeyPair-xxxx).
7. Click **OK**.
8. 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.

10.2.3 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 10-2 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 10-3](#) is the public key.

Figure 10-3 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 10-4 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 10-5 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](#).

10.2.4 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. Click in the upper left corner and select your region and project.

3. Under **Computing**, click **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **Key Pair**.
5. On the **Key Pair Service** page, click **Import Key Pair**.
6. 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](#)).
 -  **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.
- 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**.

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?](#)

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 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 the user to the group created in 1.
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](#). The following provides examples of common ECS custom policies.

Example Custom Policies

- Example 1: Only allowing users to start, stop, and restart ECSs in batches

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ecs:cloudServerFlavors:get",
        "ecs:cloudServers:reboot",
        "ecs:cloudServers:start",
        "ecs:cloudServers:get",
        "ecs:cloudServers:list",
        "ecs:cloudServers:stop"
      ]
    }
  ]
}
```

- Example 2: Only allowing users to stop and delete ECSs in batches

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ecs:cloudServers:get",
        "ecs:cloudServers:delete",
        "ecs:cloudServers:list",
        "ecs:cloudServers:stop"
      ]
    }
  ]
}
```

- Example 3: Only allowing VNC login

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ecs:cloudServerFlavors:get",
        "ecs:cloudServers:vnc",
        "ecs:cloudServers:get",
        "ecs:cloudServers:list"
      ]
    }
  ]
}
```

- Example 4: 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 which the user belongs to. 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 and Tags

12.1 Tag Management

12.1.1 Overview

Scenarios

A tag identifies an ECS. Adding tags to an ECS facilitates ECS identification and management.

You can add a tag to an ECS during the ECS creation or after the ECS is created. You can add a maximum of 10 tags to each ECS.

NOTE

Tags added during the ECS creation will also be added to the EIP and EVS disks (including the system disk and data disks) of the ECS. If the ECS uses an existing EIP, the tags will not be added to the EIP.

After creating the ECS, you can view the tags on the pages providing details about the ECS, EIP, and EVS disks.

Basics of Tags

Tags are used to identify cloud resources. When you have many cloud resources of the same type, you can use tags to classify cloud resources by dimension (for example, use, owner, or environment).

Figure 12-1 Example tags




Figure 12-1 shows how tags work. In this example, you assign two tags to each cloud resource. Each tag contains a key and a value that you define. The key of one tag is **Owner**, and the key of another tag is **Usage**. Each tag has a value.

You can quickly search for and filter specific cloud resources based on the tags added to them. For example, you can define a set of tags for cloud resources in an account to track the owner and usage of each cloud resource, making resource management easier.

Tag Naming Rules

- Each tag consists of a key-value pair.
- A maximum of 10 tags can be added to an ECS.
- For each resource, a tag key must be unique and can have only one tag value.
- A tag consists of a tag key and a tag value. **Table 12-1** lists the tag key and value requirements.

Table 12-1 Tag key and value requirements

Parameter	Requirement	Example Value
Key	<ul style="list-style-type: none"> • Cannot be left blank. • The key value must be unique for an ECS. • Can contain a maximum of 36 characters. • Can only consist of digits, letters, hyphens (-), and underscores (_). • Cannot contain the following special characters: =*<>_ / • Cannot be started or ended with a space. 	Organization

Parameter	Requirement	Example Value
Value	<ul style="list-style-type: none">• Can contain a maximum of 43 characters.• Can only consist of digits, letters, hyphens (-), underscores (_), and periods (.).• Cannot contain the following special characters: =*<>\, /• Cannot be started or ended with a space.	Apache

12.1.2 Adding Tags


Tags are used to identify cloud resources, such as ECSs, images, and disks. If you have multiple types of cloud resources which are associated with each other, you can add tags to the resources to classify and manage them easily. For more details, see [Overview](#).

You can add tags to an ECS in any of the following ways:

- [Adding Tags During ECS Creation](#)
- [Adding Tags on the ECS Details Page](#)
- [Adding Tags on the TMS Console](#)

For details about how to use predefined tags, see [Using Predefined Tags](#).

Adding Tags During ECS Creation

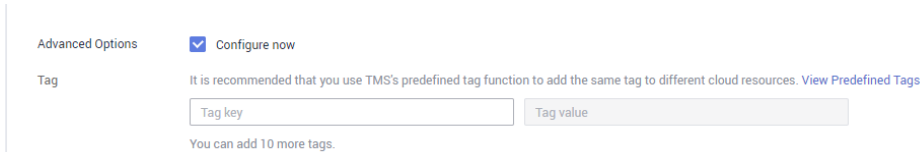
1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Computing**, choose **Elastic Cloud Server**.
4. Click **Create ECS**.
5. Configure parameters for the ECS.

Select **Configure now** for **Advanced Options**. Then, add a tag key and tag value. For the tag key and tag value requirements, see [Table 12-1](#).

NOTE

For details about other parameters, see [Creating an ECS](#).

Figure 12-2 Adding tags



Advanced Options Configure now

Tag

It is recommended that you use TMS's predefined tag function to add the same tag to different cloud resources. [View Predefined Tags](#)

Tag key Tag value


You can add 10 more tags.

Adding Tags on the ECS Details Page

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. In the ECS list, click the name of the target ECS.
The ECS details page is displayed.
4. Click the **Tags** tab and then **Add Tag**. In the displayed dialog box, enter the tag key and tag value. For the tag key and tag value requirements, see [Table 12-1](#).

You can change the tag value after the tag is added.

Figure 12-3 Adding tags on the Tags tab



Adding Tags on the TMS Console


NOTE

This method is suitable for adding tags with the same tag key to multiple resources.

1. Log in to the management console.
2. Under **Management & Deployment**, choose **Tag Management Service**.
3. On the displayed **Resource Tags** page, select the region where the resource is located, select **ECS-ECS** for **Resource Type**, and click **Search**.




All ECSs matching the search criteria are displayed.


4. In the **Search Result** area, click **Create Key**. In the displayed dialog box, enter a key (for example **project**) and click **OK**.

After the tag is created, the tag key is added to the resource list. If the key is not displayed in the resource list, click  and select the created key from the drop-down list.

By default, the value of the tag key is **Not tagged**. You need to set a value for the tag of each resource to associate the tag with the resource.

Figure 12-4 Resource list

<input type="checkbox"/>	Resource Type	Resource Name	Region	Total Tags	project
<input checked="" type="checkbox"/>	ECS-ECS		UAE-Abu Dhabi (ae-ad-1)	0	Not tagged
<input checked="" type="checkbox"/>	ECS-ECS		UAE-Abu Dhabi (ae-ad-1)	0	Not tagged
<input checked="" type="checkbox"/>	ECS-ECS		UAE-Abu Dhabi (ae-ad-1)	0	Not tagged
<input checked="" type="checkbox"/>	ECS-ECS		UAE-Abu Dhabi (ae-ad-1)	0	Not tagged
<input checked="" type="checkbox"/>	ECS-ECS		UAE-Abu Dhabi (ae-ad-1)	0	Not tagged

5. Click **Edit** to make the resource list editable.
6. Locate the row containing the target ECS, click , and enter a value (for example **A**).

After a value is set for a tag key, the number of tags is incremented by 1. Repeat the preceding steps to add tag values for other ECSs.

Figure 12-5 Setting a tag value

Resource Type	Resource Name	Region	Total Tags	project
ECSECS		UAE-Abu Dhabi (ae-ae-1)	0	A
ECSECS		UAE-Abu Dhabi (ae-ae-1)	0	Not tagged
ECSECS		UAE-Abu Dhabi (ae-ae-1)	0	Not tagged
ECSECS		UAE-Abu Dhabi (ae-ae-1)	0	Not tagged

Using Predefined Tags

If you want to add the same tag to multiple ECSs or other resources, you can create a predefined tag on the TMS console and then select the tag for the ECSs or resources. This frees you from having to repeatedly enter tag keys and values. To do so, perform the following operations:

1. Log in to the management console.
2. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.
3. Choose **Predefined Tags** in the left navigation pane and click **Create Tag**. In the displayed dialog box, enter a key (for example, **project**) and a value (for example, **A**).
4. Choose **Service List > Computing > Elastic Cloud Server**, and select the predefined tag by following the procedure for adding a tag.

12.1.3 Searching for Resources by Tag

After tags are added to resources, you can search for resources by tag using either of the following methods.

Searching for ECSs by Tag

On the **Elastic Cloud Server** page, search for ECSs by tag key or value.

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. Click **Search by Tag** above the upper right corner of the ECS list to expand the search area.
4. Enter the tag of the ECS to be queried.

Neither the tag key nor value can be empty. When the tag key or value is matched, the system automatically shows the target ECSs.

5. Add tags.

The system supports multiple tags and uses the intersection set of all tags to search for ECSs.

6. Click **Search**.

The system searches for ECSs based on tag keys and values.

Filtering Resources on the TMS Console

1. Log in to the management console.
2. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.

3. On the **Resource Tags** page, set the search criteria, including **Region**, **Resource Type**, and **Resource Tag**.
4. Click **Search**.
All the resources that meet the search criteria will be displayed in the **Search Result** area.

12.1.4 Deleting a Tag

If you no longer need a tag, delete it in any of the following ways:


- [Deleting a Tag on the ECS Details Page](#)
- [Deleting a Tag on the TMS Console](#)
- [Batch Deleting Tags on the TMS Console](#)



Deleting a Tag on the ECS Details Page

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. In the ECS list, click the name of the target ECS.
The ECS details page is displayed.
4. Click the **Tags** tab. Locate the row containing the tag to be deleted and click **Delete** in the **Operation** column. In the **Delete Tag** dialog box, click **Yes**.

Deleting a Tag on the TMS Console

1. Log in to the management console.
2. Under **Management & Deployment**, choose **Tag Management Service**.
3. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.
4. On the **Resource Tags** page, set the search criteria for ECSs and click **Search**.
5. In the **Search Result** area, click **Edit** to make the resource tag list editable.


If the key of a tag you want to delete is not contained in the list, click  and select the tag key from the drop-down list. It is a good practice to select at most 10 keys to display.

6. Locate the row containing the target ECS and click .
7. (Optional) Click  in the upper right of the **Search Result** area.
The resource list is refreshed and the refresh time is updated.

Batch Deleting Tags on the TMS Console

NOTICE

Exercise caution when deleting tags in a batch. After you delete the tags, they will be removed from all the associated ECSs and cannot be recovered.

1. Log in to the management console.
2. Under **Management & Deployment**, choose **Tag Management Service**.
3. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.
4. On the **Resource Tags** page, set the search criteria for ECSs and click **Search**.
5. Select the target ECSs.
6. Click **Manage Tag** in the upper left corner of the list.
7. In the displayed **Manage Tag** dialog box, click **Delete** in the **Operation** column. Click **OK**.
8. (Optional) Click  in the upper right of the **Search Result** area.
The resource list is refreshed and the refresh time is updated.



12.2 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. Click  in the upper left corner and select the desired region and project.
3. In the upper right corner of the page, click  .
The **Service Quota** page is displayed.
4. 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?

1. Log in to the management console.
2. In the upper right corner of the page, choose **Resources > My Quotas**.
The **Service Quota** page is displayed.
3. Click **Increase Quota** in the upper right corner of the page.
4. On the **Create Service Ticket** page, configure parameters as required.
In the **Problem Description** area, fill in the content and reason for adjustment.
5. After all necessary parameters are configured, select **I have read and agree to the Ticket Service Protocol and Privacy Statement** and click **Submit**.

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.

Server monitoring includes basic monitoring, OS monitoring, and process monitoring for servers.

- Basic monitoring includes metrics automatically reported by ECSs, such as CPU usage.
- OS monitoring provides proactive, fine-grained OS monitoring for ECSs, and it requires the Agent (a plug-in) to be installed on all ECSs to be monitored. OS monitoring supports metrics such as CPU usage and memory usage (Linux).
- Process monitoring provides monitoring of active processes on ECSs and it requires the Agent to be installed on the ECSs to be monitored. By default, Cloud Eye collects CPU usage, memory usage, and number of opened files of active processes.

This section covers the following content:

- Viewing basic ECS metrics
- Viewing OS metrics (Agent installed on ECS)
- Viewing process monitoring metrics (Agent installed on ECS)
- Customizing ECS alarm rules
- Viewing ECS running statuses for routine monitoring

Helpful Links

- [Why Is My Windows ECS Running Slowly?](#)
- [Why Is My Linux ECS Running Slowly?](#)

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](#).

NOTE

- Certain ECS metrics require the installation of UVP VMTools on the image from which the ECS is created. For details about how to install UVP VMTools, see <https://github.com/UVP-Tools/UVP-Tools/>.
- Certain ECS metrics require the installation of the Agent on the ECS. After the Agent is installed, log in to the management console and choose **Cloud Eye** under **Management & Deployment**. On the Cloud Eye console, choose **Server Monitoring** > **Elastic Cloud Server** from the left navigation pane to view ECS metrics, such as **AGT. User Space CPU Usage**. For details, see [OS Monitoring Metrics Supported by ECSs with the Agent Installed](#).
 - For details about how to install the Agent on a Windows ECS, see "Installing and Configuring the Agent (Windows)" in *Cloud Eye User Guide*.
 - For details about how to install the Agent on a Linux ECS, see "Installing and Configuring the Agent (Linux)" in *Cloud Eye User Guide*.

Table 13-1 Basic ECS metrics

Metric	Windows		Linux	
	Xen	KVM	Xen	KVM
None	Xen	KVM	Xen	KVM
CPU Usage	Supported	Supported	Supported	Supported
Memory Usage	Supported	Supported	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Not supported

Metric	Windows		Linux	
Disk Usage	Supported	Supported	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Not supported
Disk Read Bandwidth	Supported	Supported	Supported	Supported
Disk Write Bandwidth	Supported	Supported	Supported	Supported
Disk Read IOPS	Supported	Supported	Supported	Supported
Disk Write IOPS	Supported	Supported	Supported	Supported
Inband Incoming Rate	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Supported	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Not supported
Inband Outgoing Rate	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Supported	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Not supported
Outband Incoming Rate	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported

Metric	Windows		Linux	
Outband Outgoing Rate	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	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 Unit: Percent Formula: CPU usage of an ECS/Number of vCPUs in the ECS	≥ 0	ECS	5 minutes
mem_util	Memory Usage	Memory usage of an ECS This metric is unavailable if the image has no UVP VMTools installed. Unit: Percent Formula: Used memory of an ECS/ Total memory of the ECS NOTE The memory usage of QingTian ECSs cannot be monitored.	≥ 0	ECS	5 minutes

Metric ID	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Interval (Raw Metrics and KVM Only)
disk_util_inband	Disk Usage	Disk usage of an ECS This metric is unavailable if the image has no UVP VMTools installed. Unit: Percent Formula: Used capacity of an ECS-attached disk/Total capacity of the ECS-attached disk	≥ 0	ECS	5 minutes
disk_read_bytes_rate	Disk Read Bandwidth	Number of bytes read from an ECS-attached disk per second Unit: byte/s Formula: Total number of bytes read from an ECS-attached disk/ Monitoring interval $\text{byte_out} = (\text{rd_bytes} - \text{last_rd_bytes})/\text{Time difference}$	≥ 0	ECS	5 minutes
disk_write_bytes_rate	Disk Write Bandwidth	Number of bytes written to an ECS-attached disk per second Unit: byte/s Formula: Total number of bytes written to an ECS-attached disk/ Monitoring interval	≥ 0	ECS	5 minutes

Metric ID	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Interval (Raw Metrics and KVM Only)
disk_read_requests_rate	Disk Read IOPS	Number of read requests sent to an ECS-attached disk per second Unit: request/s Formula: Total number of read requests sent to an ECS-attached disk/ Monitoring interval $req_out = (rd_req - last_rd_req) / \text{Time difference}$	≥ 0	ECS	5 minutes
disk_write_requests_rate	Disk Write IOPS	Number of write requests sent to an ECS-attached disk per second Unit: request/s Formula: Total number of write requests sent to an ECS-attached disk/ Monitoring interval $req_in = (wr_req - last_wr_req) / \text{Time difference}$	≥ 0	ECS	5 minutes
network_incoming_bytes_rate_inband	Inband Incoming Rate	Number of incoming bytes on an ECS per second Unit: byte/s Formula: Total number of inband incoming bytes on an ECS/Monitoring interval	≥ 0	ECS	5 minutes

Metric ID	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Interval (Raw Metrics and KVM Only)
network_outgoing_bytes_rate_inband	Inband Outgoing Rate	Number of outgoing bytes on an ECS per second Unit: byte/s Formula: Total number of inband outgoing bytes on an ECS/Monitoring interval	≥ 0	ECS	5 minutes
network_incoming_bytes_aggregate_rate	Outband Incoming Rate	Number of incoming bytes on an ECS per second on the hypervisor Unit: byte/s Formula: Total number of outband incoming bytes on an ECS/Monitoring interval This metric is unavailable if SR-IOV is enabled.	≥ 0	ECS	5 minutes
network_outgoing_bytes_aggregate_rate	Outband Outgoing Rate	Number of outgoing bytes on an ECS per second on the hypervisor Unit: byte/s Formula: Total number of outband outgoing bytes on an ECS/Monitoring interval This metric is unavailable if SR-IOV is enabled.	≥ 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"> Linux: Check metric value changes in file / proc/stat in a collection period. Run the top command to check the %Cpu(s) value. Windows: Obtain the metric value using the Windows API GetSystemTimes. 	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"> Linux: Check metric value changes in file / proc/stat in a collection period. Windows: Obtain the metric value using the Windows API GetSystemTimes. 	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 Unit: percent <ul style="list-style-type: none">Linux: Check metric value changes in file / proc/stat in a collection period. Run the top command to check the %Cpu(s) us value.Windows: Obtain the metric value using the Windows API GetSystemTimes.	0-100	ECS	1 minute
cpu_usage_system	(Agent) Kernel Space CPU Usage	Percentage of time that the CPU is used by kernel space Unit: percent <ul style="list-style-type: none">Linux: Check metric value changes in file / proc/stat in a collection period. Run the top command to check the %Cpu(s) sy value.Windows: Obtain the metric value using the Windows API GetSystemTimes.	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"> Linux: Other Process CPU Usage = 1 - Idle CPU Usage - Kernel Space CPU Usage - User Space CPU Usage Windows: Other Process CPU Usage = 1 - Idle CPU Usage - Kernel Space CPU Usage - User Space CPU Usage 	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"> Linux: Check metric value changes in file / proc/stat in a collection period. Run the top command to check the %Cpu(s) ni value. Windows is not supported currently. 	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 Unit: percent <ul style="list-style-type: none">Linux: Check metric value changes in file / proc/stat in a collection period. Run the top command to check the %Cpu(s) wa value.Windows is not supported currently.	0-100	ECS	1 minute
cpu_usage_irq	(Agent) CPU Interrupt Time	Percentage of time that the CPU is servicing interrupts Unit: percent <ul style="list-style-type: none">Linux: Check metric value changes in file / proc/stat in a collection period. Run the top command to check the %Cpu(s) hi value.Windows is not supported currently.	0-100	ECS	1 minute
cpu_usage_softirq	(Agent) CPU Software Interrupt Time	Percentage of time that the CPU is servicing software interrupts Unit: percent <ul style="list-style-type: none">Linux: Check metric value changes in file / proc/stat in a collection period. Run the top command to check the %Cpu(s) si value.Windows is not supported currently.	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 Linux: Obtain the metric value from the number of logic CPUs in load1/ in file /proc/loadavg . Run the top command to check the load1 value.	≥ 0	ECS	1 minute
load_averge5	(Agent) 5-Minute Load Average	CPU load averaged from the last 5 minutes Linux: Obtain the metric value from the number of logic CPUs in load5/ in file /proc/loadavg . Run the top command to check the load5 value.	≥ 0	ECS	1 minute
load_averge15	(Agent) 15-Minute Load Average	CPU load averaged from the last 15 minutes Linux: Obtain the metric value from the number of logic CPUs in load15/ in file /proc/loadavg . Run the top command to check the load15 value.	≥ 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">Linux: Obtain the metric value from /proc/meminfo.<ul style="list-style-type: none">If MemAvailable is displayed in /proc/meminfo, obtain the value.If MemAvailable is not displayed in /proc/meminfo, MemAvailable = MemFree + Buffers+CachedWindows: The metric value is calculated by available memory minuses used memory. The value is obtained by calling the Windows API GlobalMemoryStatusEx.	≥ 0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
mem_usedPercent	(Agent) Memory Usage	Memory usage of the monitored object Unit: percent <ul style="list-style-type: none">Linux: Obtain the metric value from the /proc/meminfo file: (MemTotal - MemAvailable)/ MemTotal<ul style="list-style-type: none">If MemAvailable is displayed in /proc/meminfo, MemUsedPercent = (MemTotal - MemAvailable)/ MemTotalIf MemAvailable is not displayed in /proc/meminfo, MemUsedPercent = (MemTotal - MemFree - Buffers - Cached)/ MemTotalWindows: The calculation formula is as follows: Used memory size/Total memory size*100%.	0-100	ECS	1 minute
mem_free	(Agent) Idle Memory	Amount of memory that is not being used Unit: GB <ul style="list-style-type: none">Linux: Obtain the metric value from /proc/meminfo.Windows is not supported currently.	≥ 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 Unit: GB <ul style="list-style-type: none">Linux: Obtain the metric value from /proc/meminfo. Run the top command to check the KiB Mem:buffers value.Windows is not supported currently.	≥ 0	ECS	1 minute
mem_cached	(Agent) Cache	Amount of memory that is being used for file caches Unit: GB <ul style="list-style-type: none">Linux: Obtain the metric value from /proc/meminfo. Run the top command to check the KiB Swap:cached Mem value.Windows is not supported currently.	≥ 0	ECS	1 minute
total_open_files	(Agent) Total File Handles	Total handles used by all processes Unit: count <ul style="list-style-type: none">Linux: Use the /proc/{pid}/fd file to summarize the handles used by all processes.Windows is not supported currently.	≥ 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 Unit: GB <ul style="list-style-type: none">• Linux: Run the df -h command to check the value in the Avail 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 (~).• 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 (~).	≥ 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 Unit: GB</p> <ul style="list-style-type: none">Linux: Run the df -h command to check the value in the Size 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 (~).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 (~).	≥ 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 Unit: GB <ul style="list-style-type: none"> Linux: Run the df -h command to check the value in the Used 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 (~). 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 (~). 	≥ 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: Disk Usage = Used Disk Space/Disk Storage Capacity</p> <p>Unit: percent</p> <ul style="list-style-type: none"> 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 (~). 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 (~). 	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">Linux: The disk read rate is calculated based on the data changes in the sixth column of the corresponding device in file /proc/diskstats 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 (~).Windows:<ul style="list-style-type: none">Use Win32_PerformanceData_PerfDisk_LogicalDisk object in the WMI to obtain disk I/O 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 (~).	≥ 0 bytes/s	<ul style="list-style-type: none">EC S - DiskEC S - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
		<ul style="list-style-type: none"> - When the CPU usage is high, monitoring data obtaining timeout 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_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">Linux: The disk read requests are calculated based on the data changes in the fourth column of the corresponding device in file /proc/diskstats 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 (~).Windows:<ul style="list-style-type: none">Use Win32_PerformanceData_PerfDisk_LogicalDisk object in the WMI to obtain disk I/O 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 (~).When the CPU usage is high, monitoring data obtaining timeout	≥ 0 requests/s	<ul style="list-style-type: none">EC S - DiskEC S - Mount point	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">Linux: The disk write rate is calculated based on the data changes in the tenth column of the corresponding device in file /proc/diskstats in a collection period. <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">Windows:<ul style="list-style-type: none">Use Win32_PerformanceData_PerfDisk_LogicalDisk object in the WMI to obtain disk I/O 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 (~).When the CPU usage is high, monitoring data obtaining timeout	≥ 0 bytes/s	<ul style="list-style-type: none">EC S - DiskEC S - Mount point	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">Linux: The disk write requests are calculated based on the data changes in the eighth column of the corresponding device in file /proc/diskstats in a collection period. <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">Windows:<ul style="list-style-type: none">Use Win32_PerformanceData_PerfDisk_LogicalDisk object in the WMI to obtain disk I/O 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 (~).When the CPU usage is high, monitoring data obtaining timeout	≥ 0 requests/s	<ul style="list-style-type: none">EC S - DiskEC S - Mount point	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 Unit: ms/count</p> <ul style="list-style-type: none"> Linux: The average read request time is calculated based on the data changes in the seventh column of the corresponding device in file /proc/diskstats 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 (~). Windows is not supported currently. 	≥ 0 ms/Count	<ul style="list-style-type: none"> EC S - Disk EC S - Mount point 	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 Unit: ms/count</p> <ul style="list-style-type: none"> Linux: The average write request time is calculated based on the data changes in the eleventh column of the corresponding device in file /proc/diskstats 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 (~). Windows is not supported currently. 	≥ 0 ms/Count	<ul style="list-style-type: none"> EC S - Disk EC S - Mount point 	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"> Linux: The disk I/O usage is calculated based on the data changes in the thirteenth column of the corresponding device in file /proc/diskstats 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 (~). Windows is not supported currently. 	0-100	<ul style="list-style-type: none"> EC S - Disk EC S - Mount point 	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"> Linux: The average disk queue length is calculated based on the data changes in the fourteenth column of the corresponding device in file /proc/diskstats 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 (~). Windows is not supported currently. 	≥ 0	<ul style="list-style-type: none"> EC S - Disk EC S - Mount point 	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"> 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 /proc/diskstats in a collection period. <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"> Windows is not supported currently. 	≥ 0 bytes/op	<ul style="list-style-type: none"> EC S - Disk EC S - Mount point 	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_read_bytes_per_operation	(Agent) 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"> 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 /proc/diskstats in a collection period. <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"> Windows is not supported currently. 	≥ 0 bytes/op	<ul style="list-style-type: none"> EC S - Disk EC S - Mount point 	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"> 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 /proc/diskstats 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 (~). Windows is not supported currently. 	≥ 0	<ul style="list-style-type: none"> EC S - Disk EC S - Mount point 	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. Calculation formula: Used storage space of all mounted disk partitions/ Total disk storage space</p> <ul style="list-style-type: none"> 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. Windows ECSs do not support this metric. 	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 Possible values are 0 (read and write) and 1 (read only).</p> <p>Linux: Check file system information in the fourth column in file /proc/mounts.</p>	<ul style="list-style-type: none"> 0: readable and writable 1: read-only 	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 Linux: Run the df -i command to check the value in the Inodes 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	ECS - Mount point	1 minute
disk_inodesUsed	(Agent) Total inode Used	Number of used index nodes on the disk Linux: Run the df -i command to check the value in the IUsed 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	ECS - Mount point	1 minute
disk_inodesUsedPercent	(Agent) Percentage of Total inode Used	Number of used index nodes on the disk Unit: percent Linux: Run the df -i command to check the value in the IUse% 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 Unit: bit/s <ul style="list-style-type: none">Linux: Check metric value changes in file / proc/net/dev in a collection period.Windows: Use the MibIfRow object in the WMI to obtain network metric data.	≥ 0 bit/s	ECS	1 minute
net_bitSent	(Agent) Inbound Bandwidth	Number of bits sent by this NIC per second Unit: bit/s <ul style="list-style-type: none">Linux: Check metric value changes in file / proc/net/dev in a collection period.Windows: Use the MibIfRow object in the WMI to obtain network metric data.	≥ 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 Unit: count/s <ul style="list-style-type: none">Linux: Check metric value changes in file / proc/net/dev in a collection period.Windows: Use the MibIfRow object in the WMI to obtain network metric data.	≥ 0 Counts/s	ECS	1 minute
net_packetSent	(Agent) NIC Packet Send Rate	Number of packets sent by this NIC per second Unit: count/s <ul style="list-style-type: none">Linux: Check metric value changes in file / proc/net/dev in a collection period.Windows: Use the MibIfRow object in the WMI to obtain network metric data.	≥ 0 Counts/s	ECS	1 minute
net_errin	(Agent) Receive Error Rate	Percentage of receive errors detected by this NIC per second Unit: percent <ul style="list-style-type: none">Linux: Check metric value changes in file / proc/net/dev in a collection period.Windows is not supported currently.	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 Unit: percent <ul style="list-style-type: none"> Linux: Check metric value changes in file / proc/net/dev in a collection period. Windows is not supported currently. 	0-100	ECS	1 minute
net_dropin	(Agent) Received Packet Drop Rate	Percentage of packets received by this NIC which were dropped per second Unit: percent <ul style="list-style-type: none"> Linux: Check metric value changes in file / proc/net/dev in a collection period. Windows is not supported currently. 	0-100	ECS	1 minute
net_dropout	(Agent) Transmitted Packet Drop Rate	Percentage of packets transmitted by this NIC which were dropped per second Unit: percent <ul style="list-style-type: none"> Linux: Check metric value changes in file / proc/net/dev in a collection period. Windows is not supported currently. 	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 Unit: ms Collection method for Linux ECSs: Run chronyc sources -v to obtain the offset.	≥ 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 Unit: count <ul style="list-style-type: none">Linux: Obtain TCP connections in all states from the /proc/net/tcp file, and then collect the number of connections in each state.Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 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 Unit: count <ul style="list-style-type: none">Linux: Obtain TCP connections in all states from the /proc/net/tcp file, and then collect the number of connections in each state.Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute
net_tcp_sys_sent	(Agent) TCP SYS_SENT	Number of TCP connections that are being requested by the client Unit: count <ul style="list-style-type: none">Linux: Obtain TCP connections in all states from the /proc/net/tcp file, and then collect the number of connections in each state.Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute
net_tcp_sys_rcv	(Agent) TCP SYS_RECEIVED	Number of pending TCP connections received by the server Unit: count <ul style="list-style-type: none">Linux: Obtain TCP connections in all states from the /proc/net/tcp file, and then collect the number of connections in each state.Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_fin_wait1	(Agent) TCP FIN_WAIT1	Number of TCP connections waiting for ACK packets when the connections are being actively closed by the client Unit: count <ul style="list-style-type: none">Linux: Obtain TCP connections in all states from the /proc/net/tcp file, and then collect the number of connections in each state.Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute
net_tcp_fin_wait2	(Agent) TCP FIN_WAIT2	Number of TCP connections in the FIN_WAIT2 state Unit: count <ul style="list-style-type: none">Linux: Obtain TCP connections in all states from the /proc/net/tcp file, and then collect the number of connections in each state.Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute
net_tcp_time_wait	(Agent) TCP TIME_WAIT	Number of TCP connections in TIME_WAIT state Unit: count <ul style="list-style-type: none">Linux: Obtain TCP connections in all states from the /proc/net/tcp file, and then collect the number of connections in each state.Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 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 Unit: count <ul style="list-style-type: none">Linux: Obtain TCP connections in all states from the /proc/net/tcp file, and then collect the number of connections in each state.Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute
net_tcp_close_wait	(Agent) TCP CLOSE_WAIT	Number of TCP connections in CLOSE_WAIT TCP state Unit: count <ul style="list-style-type: none">Linux: Obtain TCP connections in all states from the /proc/net/tcp file, and then collect the number of connections in each state.Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 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 Unit: count <ul style="list-style-type: none">Linux: Obtain TCP connections in all states from the /proc/net/tcp file, and then collect the number of connections in each state.Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 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 Unit: count <ul style="list-style-type: none">Linux: Obtain TCP connections in all states from the /proc/net/tcp file, and then collect the number of connections in each state.Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 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 Unit: count <ul style="list-style-type: none">Linux: Obtain TCP connections in all states from the /proc/net/tcp file, and then collect the number of connections in each state.Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute
net_tcp_retrans	(Agent) TCP Retransmission Rate	Percentage of packets that are resent Unit: percent <ul style="list-style-type: none">Linux: Obtain the metric value from the /proc/net/snmp file. The value is the ratio of the number of sent packets to the number of retransmitted packages in a collection period.Windows: Obtain the metric value using WindowsAPI GetTcpStatistics.	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 Unit: none <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card.	<ul style="list-style-type: none">0: The GPU is healthy.1: The GPU is subhealthy.2: The GPU is faulty.	<ul style="list-style-type: none">ECSECS - GPU	1 minute
gpu_usage_encoder	Encoding Usage	Encoding capability usage on the GPU Unit: percent <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card.	0-100%	<ul style="list-style-type: none">ECSECS - GPU	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">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card.	0-100%	<ul style="list-style-type: none">ECSECS - GPU	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 0 each time the GPU is reset.</p> <p>Unit: count</p> <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card.	≥ 0	<ul style="list-style-type: none">ECSECS - GPU	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_volatile_uncorrectable	Volatile Uncorrectable ECC Errors	<p>Number of uncorrectable ECC errors since the GPU is reset. The value is reset to 0 each time the GPU is reset.</p> <p>Unit: count</p> <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card.	≥ 0	<ul style="list-style-type: none">ECSECS - GPU	1 minute
gpu_aggregate_correctable	Aggregate Correctable ECC Errors	<p>Aggregate correctable ECC errors on the GPU</p> <p>Unit: count</p> <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card.	≥ 0	<ul style="list-style-type: none">ECSECS - GPU	1 minute
gpu_aggregate_uncorrectable	Aggregate Uncorrectable ECC Errors	<p>Aggregate uncorrectable ECC Errors on the GPU</p> <p>Unit: count</p> <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card.	≥ 0	<ul style="list-style-type: none">ECSECS - GPU	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">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card.	≥ 0	<ul style="list-style-type: none">ECSECS - GPU	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">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card.	≥ 0	<ul style="list-style-type: none">ECSECS - GPU	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_performance_state	(Agent) Performance Status	GPU performance of the monitored object Unit: none <ul style="list-style-type: none"> Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card. Windows: Obtain the metric value using the nvml.dll library of the graphics card. 	P0-P15, P32 <ul style="list-style-type: none"> P0: indicates the maximum performance status. P15: indicates the minimum performance status. P32: indicates the unknown performance status. 	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 Unit: percent <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card.	0-100	ECS - GPU	1 minute
gpu_usage_gpu	(Agent) GPU Usage	GPU usage of the monitored object Unit: percent <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card.	0-100	ECS - GPU	1 minute

Dimensions

Dimension	Key	Value
ECS	instance_id	Specifies the ECS ID.

13.4 Process Monitoring Metrics Supported by ECSs with the Agent Installed

Description

Process monitoring provides monitoring of active processes on ECSs and it requires the Agent to be installed on the ECSs to be monitored. By default, Cloud Eye

collects CPU usage, memory usage, and number of opened files of active processes.

This section describes process monitoring metrics reported to Cloud Eye.

Namespace

AGT.ECS

Process Metrics

After the agent is installed, you can view the default process metrics listed in the following table, regardless of ECS types and OSs.

Table 13-13 Process metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
proc_pHashId_cpu	CPU Usage	CPU consumed by a process. pHashId (process name and process ID) is the value of md5 . Unit: percent <ul style="list-style-type: none">Linux: Check metric value changes in file / proc/pid/stat.Windows: Call the Windows API <code>GetProcessTimes</code> to obtain the CPU usage of the process.	0-100%	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
proc_pHashId_memory	Memory Usage	<p>Memory consumed by a process. pHashId (process name and process ID) is the value of md5.</p> <p>Unit: percent</p> <ul style="list-style-type: none">Linux: RSS*PAGESIZE/ MemTotal <p>Obtain the RSS value by checking the second column of file /proc/pid/statm.</p> <p>Obtain the PAGESIZE value by running the getconf PAGESIZE command.</p> <p>Obtain the MemTotal value by checking file /proc/meminfo.</p> <ul style="list-style-type: none">Windows: Call the Windows API <code>procGlobalMemoryStatusEx</code> to obtain the total memory size. Call <code>GetProcessMemoryInfo</code> to obtain the used memory size. Use the used memory size to divide the total memory size to get the memory usage.	0-100%	ECS	1 minute
proc_pHashId_files	Open Files	<p>Number of files opened by a process. pHashId (process name and process ID) is the value of md5.</p> <ul style="list-style-type: none">Linux: Run the ls -l /proc/pid/fd command to view the number of opened files.Windows is not supported currently.	≥0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
proc_running_count	(Agent) Running Processes	Number of processes that are running <ul style="list-style-type: none">Linux: You can obtain the state of each process by checking the Status value in the /proc/pid/status file, and then collect the total number of processes in each state.Windows is not supported currently.	≥0	ECS	1 minute
proc_idle_count	(Agent) Idle Processes	Number of processes that are idle <ul style="list-style-type: none">Linux: You can obtain the state of each process by checking the Status value in the /proc/pid/status file, and then collect the total number of processes in each state.Windows is not supported currently.	≥0	ECS	1 minute
proc_zombie_count	(Agent) Zombie Processes	Number of zombie processes <ul style="list-style-type: none">Linux: You can obtain the state of each process by checking the Status value in the /proc/pid/status file, and then collect the total number of processes in each state.Windows is not supported currently.	≥0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
proc_blocked_count	(Agent) Blocked Processes	Number of processes that are blocked <ul style="list-style-type: none">Linux: You can obtain the state of each process by checking the Status value in the /proc/pid/status file, and then collect the total number of processes in each state.Windows is not supported currently.	≥0	ECS	1 minute
proc_sleeping_count	(Agent) Sleeping Processes	Number of processes that are sleeping <ul style="list-style-type: none">Linux: You can obtain the state of each process by checking the Status value in the /proc/pid/status file, and then collect the total number of processes in each state.Windows is not supported currently.	≥0	ECS	1 minute
proc_total_count	(Agent) Total Processes	Total number of processes on the monitored object <ul style="list-style-type: none">Linux: You can obtain the state of each process by checking the Status value in the /proc/pid/status file, and then collect the total number of processes in each state.Windows: Obtain the total number of processes by using the system process status support module psapi.dll.	≥0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
proc_specified_count	(Agent) Specified Processes	Number of specified processes <ul style="list-style-type: none">Linux: You can obtain the state of each process by checking the Status value in the <code>/proc/pid/status</code> file, and then collect the total number of processes in each state.Windows: Obtain the total number of processes by using the system process status support module psapi.dll.	≥0	<ul style="list-style-type: none">ECSECS - Process	1 minute

Dimensions

Dimension	Key	Value
ECS	instance_id	Specifies the ECS ID.

13.5 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. On the **Modify Alarm Rule** page, set parameters as prompted.
- 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*.

13.6 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 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 CTS

14.1 Key Operations Supported by CTS

Scenarios

Cloud Trace Service (CTS) records user operations performed on ECSs and related resources for further query, auditing, and backtracking.

Prerequisites

CTS has been provisioned.

Key ECS Operations Recorded by CTS

Table 14-1 ECS operations recorded by CTS

Operation	Resource Type	Event Name
Creating an ECS	ecs	createServer createServerV2 createServerV21
Deleting an ECS	ecs	deleteServer deleteServerV2 deleteServerV21
Starting an ECS	ecs	startServer
Restarting an ECS	ecs	rebootServer
Stopping an ECS	ecs	stopServer
Adding an ECS NIC	ecs	addNic
Deleting an ECS NIC	ecs	deleteNic delNic

Operation	Resource Type	Event Name
Attaching a disk (on the EVS console)	ecs	attachVolume2
Reinstalling an OS	ecs	reinstallOs
Changing an OS	ecs	changeOs
Modifying specifications	ecs	resizeServer
Enabling automatic recovery on an ECS	ecs	addAutoRecovery
Disabling automatic recovery on an ECS	ecs	deleteAutoRecovery

14.2 Viewing Traces


Scenarios




After you enable CTS and the management tracker is created, CTS starts recording operations on cloud resources. After a data tracker is created, the system starts recording operations on data in OBS buckets. CTS stores operation records generated in the last seven days.

This section describes how to query and export operation records of the last seven days on the CTS console.


- [Viewing Real-Time Traces in the Trace List of the New Edition](#)
- [Viewing Real-Time Traces in the Trace List of the Old Edition](#)



Viewing Real-Time Traces in the Trace List of the New Edition

1. Log in to the management console.
2. Click  in the upper left corner and choose Management & Deployment > **Cloud Trace Service**. The CTS console is displayed.
3. Choose **Trace List** in the navigation pane on the left.
4. On the **Trace List** page, use advanced search to query traces. You can combine one or more filters.
 - **Trace Name:** Enter a trace name.
 - **Trace ID:** Enter a trace ID.
 - **Resource Name:** Enter a resource name. If the cloud resource involved in the trace does not have a resource name or the corresponding API operation does not involve the resource name parameter, leave this field empty.
 - **Resource ID:** Enter a resource ID. Leave this field empty if the resource has no resource ID or if resource creation failed.
 - **Trace Source:** Select a cloud service name from the drop-down list.

- **Resource Type:** Select a resource type from the drop-down list.
 - **Operator:** Select one or more operators from the drop-down list.
 - **Trace Status:** Select **normal**, **warning**, or **incident**.
 - **normal:** The operation succeeded.
 - **warning:** The operation failed.
 - **incident:** The operation caused a fault that is more serious than the operation failure, for example, causing other faults.
 - Time range: Select **Last 1 hour**, **Last 1 day**, or **Last 1 week**, or specify a custom time range.
5. On the **Trace List** page, you can also export and refresh the trace list, and customize the list display settings.
 - Enter any keyword in the search box and press Enter to filter desired traces.
 - Click **Export** to export all traces in the query result as an .xlsx file. The file can contain up to 5000 records.
 - Click  to view the latest information about traces.
 - Click  to customize the information to be displayed in the trace list. If **Auto wrapping** is enabled (), excess text will move down to the next line; otherwise, the text will be truncated. By default, this function is disabled.
 6. For details about key fields in the trace structure, see section "Trace References" > "Trace Structure" and section "Trace References" > "Example Traces".
 7. (Optional) On the **Trace List** page of the new edition, click **Go to Old Edition** in the upper right corner to switch to the **Trace List** page of the old edition.

Viewing Real-Time Traces in the Trace List of the Old Edition

1. Log in to the management console.
2. Click  in the upper left corner and choose **Management & Deployment** > **Cloud Trace Service**. The CTS console is displayed.
3. Choose **Trace List** in the navigation pane on the left.
4. Each time you log in to the CTS console, the new edition is displayed by default. Click **Go to Old Edition** in the upper right corner to switch to the trace list of the old edition.
5. Set filters to search for your desired traces. The following filters are available:
 - **Trace Type**, **Trace Source**, **Resource Type**, and **Search By:** Select a filter from the drop-down list.
 - If you select **Resource ID** for **Search By**, specify a resource ID.
 - If you select **Trace name** for **Search By**, specify a trace name.
 - If you select **Resource name** for **Search By**, specify a resource name.

- **Operator:** Select a user.
 - **Trace Status:** Select **All trace statuses**, **Normal**, **Warning**, or **Incident**.
 - Time range: You can query traces generated during any time range in the last seven days.
 - Click **Export** to export all traces in the query result as a CSV file. The file can contain up to 5000 records.
6. Click **Query**.
 7. On the **Trace List** page, you can also export and refresh the trace list.
 - Click **Export** to export all traces in the query result as a CSV file. The file can contain up to 5000 records.
 - Click  to view the latest information about traces.
 8. Click  on the left of a trace to expand its details.
 9. Click **View Trace** in the **Operation** column. The trace details are displayed.

View Trace ×

```
{
  "request": "",
  "trace_id": "XXXXXXXXXXXXXXXXXXXX",
  "code": "200",
  "trace_name": "createDockerConfig",
  "resource_type": "dockerlogincmd",
  "trace_rating": "normal",
  "api_version": "",
  "message": "createDockerConfig, Method: POST Url=/v2/manage/utills/secret, Reason:",
  "source_ip": "XXXXXXXXXX",
  "domain_id": "XXXXXXXXXXXXXXXXXXXX",
  "trace_type": "ApiCall",
  "service_type": "SWR",
  "event_type": "system",
  "project_id": "XXXXXXXXXXXXXXXXXXXX",
  "response": "",
  "resource_id": "",
  "tracker_name": "system",
  "time": "Nov 16, 2023 10:54:04 GMT+08:00",
  "resource_name": "dockerlogincmd",
  "user": {
    "domain": {
      "name": "XXXXXXXXXX",
      "id": "XXXXXXXXXXXXXXXXXXXX"
    }
  }
}
```

10. For details about key fields in the trace structure, see section "Trace References" > "Trace Structure" and section "Trace References" > "Example Traces" in the *CTS User Guide*.
11. (Optional) On the **Trace List** page of the old edition, click **New Edition** in the upper right corner to switch to the **Trace List** page of the new edition.

15 FAQs

15.1 Common Topics

Remote Logins

- [Why Can't I Log In to My Windows ECS?](#)
- [Why Can't I Log In to My Linux ECS?](#)
- [What Are the Username and Password for Remote Logins?](#)

ECS Failures or Slow ECS Responses

- [Why Is My Windows ECS Running Slowly?](#)
- [Why Is My Linux ECS Running Slowly?](#)

Internet Access Failures

- [Why Can't My Windows ECS Access the Internet?](#)
- [Why Does My Linux ECS Fail to Access the Internet?](#)
- [Can an ECS Without an EIP Bound Access the Internet?](#)

Passwords and Key Pairs

- [What Are the Username and Password for Remote Logins?](#)
- [Resetting the Password for Logging In to a Windows ECS](#)
- [Resetting the Password for Logging In to a Linux ECS](#)

Ping Failures

- [What Should I Do If an EIP Cannot Be Pinged?](#)
- [Why Can I Remotely Access an ECS But Cannot Ping It?](#)

15.2 ECS Overview

15.2.1 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.

15.2.2 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.

15.2.3 Can ECSs Automatically Recover After the Physical Host Accommodating the ECSs Becomes Faulty?

Yes, ECS can automatically recover if the physical host becomes faulty.

ECSs run on physical hosts. Although there are multiple mechanisms to ensure system reliability, fault tolerance, and high availability, host hardware might be damaged or power failures might occur. If physical hosts cannot be powered on or restarted due to damage, CPU and memory data will be lost and live migration cannot be used to recovery ECSs.

The cloud platform provides automatic recovery by default to restart ECSs through cold migration, ensuring high availability and dynamic ECS migration. If a physical host accommodating ECSs breaks down, the ECSs will automatically be migrated to a functional physical host to minimize the impact on your services. During the process, the ECSs will restart.

NOTE

- 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

15.3 Regions and AZs

15.3.1 What Is an AZ?

AZ

An availability zone (AZ) is a physical region where resources use independent power supplies and networks. AZs are physically isolated but interconnected through an internal network.

There are multiple AZs in each region. If one AZ becomes faulty, other AZs in the same region continue to provide services.

AZs in the same region can communicate with each other through an internal network.

Selecting an AZ

You can select an AZ when you are purchasing an ECS. After the ECS is created, the AZ cannot be changed. If there is only one AZ displayed in a region, it means the region only provides one AZ.

15.4 Billing

15.4.1 How Can I Stop an ECS from Being Billed?

This section uses a pay-per-use ECS as an example to describe how you are billed after the ECS is deleted. [Table 15-1](#) lists the resources associated with the ECS.

Table 15-1 Billing example of a pay-per-use ECS

Resources	Description	Billing Mode
ECS basic resources	vCPUs, memory, and image	Pay-per-use
EVS disks	System disk	Pay-per-use
	Data disk	Pay-per-use
EIP	N/A	Pay-per-use

After the ECS is deleted, it is billed as follows:

- ECS basic resources: no longer billed
- EVS disks
 - System disk: no longer billed

- Data disks: no longer billed if you have selected **Delete the data disks attached to the following ECSs** when you were deleting the ECS. Otherwise, the data disks will continue to be billed.
- EIP: no longer billed if you have selected **Release the EIPs bound to the following ECSs** when you were deleting the ECS. Otherwise, the EIP will continue to be billed.

15.5 Creation and Deletion

15.5.1 What Should I Do If the ECS Resources to Be Purchased Are Sold Out?

Each region has two or three AZs. If resources in an AZ are sold out, you can change the AZ and purchase resources in another AZ.

15.5.2 What Is the Creation Time and Startup Time of an ECS?

Creation time: time when the ECS is created on the cloud platform.

Startup time: time when the ECS is started for the first time.

15.5.3 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.

15.5.4 When Does an ECS Become Provisioned?

Pay-per-use ECS: The ECS is automatically provisioned after it is created.

15.5.5 Why Cannot I View the ECSs Being Created Immediately After I Pay for Them?

You can view the ECSs being created only after the system disks attached to the ECSs are created. This requires a period of time.

15.5.6 Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image?

Symptom

When you use a full-ECS image that was created using a CSBS backup to create ECSs, the process is time-consuming or the system displays a message indicating that the image cannot be used to rapidly create ECSs.

Cause Analysis

If your full-ECS image is in the old backup format provided by CSBS, this issue occurs.

NOTE

- CSBS has a new backup format. You can rapidly create ECSs if the full-ECS image is in this format.
- This issue does not occur if a full-ECS image is created using a CBR backup.

Solution (Using CBR)

If you want to use a full-ECS image to rapidly create ECSs, ensure that the full-ECS image is created using a CSBS backup in the new format. The procedure is as follows:

- Scenario 1: The ECS based on which the target CSBS backup is created is available.
Use an ECS to create a CBR backup and use the backup to create a full-ECS image. You can use this full-ECS image to rapidly create ECSs.
 - For instructions about how to back up an ECS, see *Cloud Backup and Recovery User Guide*.
 - For instructions about how to create a full-ECS image, see *Image Management Service User Guide*.
- Scenario 2: The ECS based on which the target CSBS backup is created is unavailable.
 - a. Use the full-ECS image to create an ECS.
 - b. Use the ECS to create a CBR backup.
For details, see *Cloud Backup and Recovery User Guide*.
 - c. Use the CBR backup to create a full-ECS image.
For details, see *Image Management Service User Guide*.
You can use the full-ECS image to rapidly create ECSs.

Solution (Using CSBS)

If you want to use a full-ECS image to rapidly create ECSs, ensure that the full-ECS image is created using a CSBS backup in the new format. The procedure is as follows:

- Scenario 1: The ECS based on which the target CSBS backup is created is available.

Back up the original ECS on the **Cloud Server Backup Service** page and use the new format to create a full-ECS image. You can use the full-ECS image to rapidly create ECSs.

- For instructions about how to back up an ECS, see *Cloud Server Backup Service User Guide*.
- For instructions about how to create a full-ECS image, see *Image Management Service User Guide*.

- Scenario 2: The ECS based on which the target CSBS backup is created is unavailable.

- a. Use the full-ECS image to create a new ECS.
- b. Back up the newly created ECS.
For details, see *Cloud Server Backup Service User Guide*.
- c. Use the CSBS backup to create a full-ECS image.
For details, see *Image Management Service User Guide*.
You can use the full-ECS image to rapidly create ECSs.

15.5.7 What Do I Do If I Selected an Incorrect Image for My ECS?

You can change the image for your ECS on the ECS console.

1. Select the target ECS and click **Stop** in the upper left corner of the ECS list.
2. Locate the row that contains the target ECS, choose **More > Manage Image/Disk/Backup > Change OS** in the **Operation** column.
The **Change OS** dialog box is displayed.
3. Select the target image type and image.
4. Set the login mode. You can select **Password** or **Key pair**.
5. Set the other parameters and click **OK**.

After the application is submitted, the ECS status changes to **Changing OS**. The OS changing has been successfully completed when the ECS status changes to **Running**.

For details, see [Changing the OS](#).

15.5.8 Should I Choose Windows OS or Linux OS for My ECS?

Difference Between Windows OS and Linux OS

The following table shows the difference between Windows OS and Linux OS. Select an OS based on your service requirements.

Table 15-2 Differences

OS	Developer Language	Database
Windows	ASP.NET, MFC and C#	ACCESS and SQL Server
Linux	Shell	MySQL and SQLite
Both Window and Linux support developer languages HTML, C, Java, and PHP.		

Windows OS

Windows Server 2012, Windows Server 2016, and Windows Server 2019 have some advanced features, such as network performance and system compatibility optimization. For the best performance, Windows Server 2019 is recommended.

Linux OS

There are a variety of Linux versions. You can select an appropriate version based on your service requirements.

OS Change

If you want to change the OS of your ECS, perform the following operations:

1. Select the target ECS and click **Stop** in the upper left corner of the ECS list.
2. Locate the row that contains the target ECS, choose **More > Manage Image/Disk/Backup > Change OS** in the **Operation** column.

The **Change OS** dialog box is displayed.

3. Select the target image type and image.
4. Set the login mode. You can select **Password** or **Key pair**.
5. Set the other parameters and click **OK**.

After the application is submitted, the ECS status changes to **Changing OS**. The OS changing has been successfully completed when the ECS status changes to **Running**.

For details, see [Changing the OS](#).

15.5.9 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 customer service.

15.5.10 How Can I Manage ECSs by Group?

You cannot manage ECSs by folders or groups, but you can use tags to organize your ECSs

Tags help you group your ECSs by things by whatever categories are useful to you.

For more information, see [Overview](#).

15.5.11 Why Did I Fail to Configure an Anti-Affinity ECS Group?

When you configure an anti-affinity ECS group during ECS purchase, an error occurred. This may be caused by insufficient resources.

In this case, you can try the following measures:

- Wait for a while and try again.
- Purchase ECSs in small batches.
- Select another AZ with sufficient resources to purchase ECSs.


15.5.12 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 **More > Delete**.

Figure 15-1 Deleting selected ECSs



15.5.13 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.

15.5.14 Can a Deleted ECS Be Restored?

No. The data of a deleted ECS cannot be restored. Therefore, before deleting an ECS, back up or migrate its data.

15.5.15 How Do I Delete or Restart an ECS?

Deleting an ECS


1. Log in to the management console.
2. Select the region where the ECS is located.
3. Under **Computing**, choose **Elastic Cloud Server**.
4. Locate the row containing the target ECS and choose **More > Delete** in the **Operation** column.

Restarting an ECS

1. Log in to the management console.
2. Select the region where the ECS is located.
3. Under **Computing**, choose **Elastic Cloud Server**.
4. Locate the row containing the target ECS and choose **More > Restart** in the **Operation** column.

15.5.16 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. Click  in the upper left corner and select your region and project.
3. Under **Computing**, click **Elastic Cloud Server**.
4. 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.
5. Select **Forcibly restart the preceding ECSs** or **Forcibly stop the preceding ECSs**.
6. Click **OK**.

15.6 Login and Connection

15.6.1 What Are the Username and Password for Remote Logins?

Username for logging in to an ECS:

- For Windows: **Administrator**
- For Linux: **root**

If you forgot the login password or did not set a password when creating the ECS, you can reset the password.

15.6.2 Why Cannot I Use the Username and Password Configured During the Creation of a GPU-accelerated ECS to Log In to the ECS Through SSH?

Solution

Log in to the ECS using VNC, modify the configuration file, and log in to the ECS through SSH.

1. On the ECS console, locate the ECS and click **Remote Login**.
2. On the login page, enter user **root** and its password.

NOTE

The password is the one you set during the ECS creation.

```
Connected (encrypted) to: QEMU (i-000FA82E) Before you exit,ensure that computer is locked.
ec2: #####
ec2: ----BEGIN SSH HOST KEY FINGERPRINTS----
ec2: 256 a4:9c:e9:d9:35:68:26:27:c1:0c:43:77:ce:db:17:35 (ECDSA)
ec2: 2048 67:e0:3d:0e:1a:0b:7a:ee:46:5a:1c:4e:44:c3:6f:b7 (RSA)
ec2: ----END SSH HOST KEY FINGERPRINTS----
ec2: #####
----BEGIN SSH HOST KEY KEYS----
ecdsa-sha2-nistp256 AAAAE2UjZHNhLXNoYTItbmlzdHAyNTYAAAAIbmlzdHAyNTYAAABBBGgDOEd
5y0ug132daqNO11YL3VBR1ZFx91ywQT8mBGUxh7X72y1opMBhQxP2E7t0o5JXt5i831P1+YPLRi9X0w=

ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQCB8xDnU4ZXP8+4pqD810A7fUz_jhhwR487z8uHa+eEvG
H1dWAVDtY4XrSZE73y_jhSvXyaGY/1GLpeczo6MgdQfW7p8/rnu+TnJ+CHUZ/<\/x0cCDSpInZpYe2cWTrsg
P8GpvZK6ZgqxFCWmkJMMZEYR_j51BtUARU8HCeh7A8bbGJaOUzCuLuUwH0edpdMU1u1BD4bGP/5zsPDGo
y.jexLlavWvsRReaWZAWQ6nTxJ55qx2fs54Gb53SUItleIE2u3aH4DtwCeSoxI+/7jc3tSmcc/PHvWnb5
56ZU0sI1c6p+9xmcI8Rm8KncKr8NMUv3xR/BbGIKcY4dniZZC81Q51B7yAs7
----END SSH HOST KEY KEYS----
cloud-init[3732]: Cloud-init v. 0.7.5 finished at Wed, 17 Jan 2018 06:39:54 +000
0. Datasource DataSourceEc2. Up 36.21 seconds

CentOS Linux 7 (Core)
Kernel 3.10.0-123.el7.x86_64 on an x86_64

Login with linux/cloud.1234, sudo for root.
ecs-dec7 login:
```

3. In the `/etc/ssh/` directory, modify the three configuration items in the `ssh_config` file, as shown in the following figure.

```
SyslogFacility AUTH
PermitRootLogin yes
# Do not enable sshd passwd auth without ensuring really strong passwords
PasswordAuthentication yes
ChallengeResponseAuthentication no
GSSAPIAuthentication no
GSSAPICleanupCredentials yes
UsePAM yes
AcceptEnv LANG LC_CTYPE LC_NUMERIC LC_TIME LC_COLLATE LC_MONETARY LC_MEASUREMENT
AcceptEnv LC_PAPER LC_NAME LC_ADDRESS LC_TELEPHONE LC_MEASUREMENT
AcceptEnv LC_IDENTIFICATION LC_ALL LANGUAGE
AcceptEnv XMODIFIERS
X11Forwarding yes
Subsystem sftp /usr/libexec/openssh/sftp-server
#UseDNS no
Ciphers aes256-ctr,aes192-ctr,aes128-ctr
KexAlgorithms diffie-hellman-group-exchange-sha256
AllowTcpForwarding no
GatewayPorts no
X11UseLocalhost yes
AllowAgentForwarding yes
PermitTunnel no
LogLevel VERBOSE
RSAAuthentication yes
PubkeyAuthentication yes
PermitEmptyPasswords no
RhostsRSAAuthentication no
HostbasedAuthentication no
IgnoreRhosts yes
AllowUsers root
~
~
~
~
"sshd_config" 31L, 938C written
bash-4.1#
```

4. Save the modification and exit. Then, run the following command to restart SSH:
service sshd restart
5. After the restart, use the SSH password to log in again.
6. If the fault persists, contact customer service.

15.6.3 Why Can't I Log In to My Windows ECS?

Symptom

A Windows ECS cannot be logged in to due to some reasons. For example, the network is abnormal, the firewall does not allow access to the local port for accessing the remote desktop, or the ECS vCPUs are overloaded.

This section describes how to troubleshoot login failures on a Windows ECS.

If you cannot log in to your Windows ECS, follow the instructions provided in [Checking the VNC Login](#). Then, locate the login fault by referring to [Fault Locating](#).

Checking the VNC Login

Check whether you can log in to the ECS using VNC on the management console.

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. In the **Operation** column of the target ECS, click **Remote Login**.

If the VNC login still fails, record the resource details and fault occurred time for further fault locating and analysis.

Fault Locating

If you can log in to the ECS using VNC but cannot log in to the ECS using a remote desktop connection, locate the fault as follows.

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.

Table 15-3 Possible causes and solutions


Possible Cause	Solution
The ECS is frozen or stopped.	Make sure that the ECS is in the Running state. For details, see Checking the ECS Status .
The entered username or password is incorrect.	The default username for Windows ECSs is Administrator . If the password is incorrect, reset the password on the management console. For details, see Checking the Login Mode .
The ECS is overloaded.	If the bandwidth or CPU usage of the ECS is excessively high, login failures may occur. For details, see Checking Whether the ECS Is Overloaded .
The ECS has no EIP bound.	To log in to an ECS using RDP or MSTSC, ensure that the ECS has an EIP bound. For details, see Checking Whether an ECS Has an EIP Bound .
The access is blocked by the Internet service provider (ISP).	Check whether you can access the ECS using another hotspot or network. For details, see Checking Whether the Network Is Normal .
The access is blocked by the firewall.	Disable the firewall and try again. For details, see Checking Whether the Firewall Is Correctly Configured .
The remote login port has been disabled in the security group or on the ECS.	Check whether the security group and the ECS allow traffic on the remote login port. For details, see Checking Whether the Remote Access Port Is Correctly Configured .
An IP address whitelist for SSH logins has been configured.	Check whether an IP address whitelist for SSH logins has been configured after HSS is enabled. For details, see Checking the IP Address Whitelist for SSH Logins (with HSS Enabled) .

Possible Cause	Solution
The remote desktop protocol has been disabled on the ECS.	Make sure that the remote desktop protocol has been enabled on the ECS (only required for RDP and MSTSC logins). For details, see Checking the Remote Desktop Protocol on the ECS .
The access is blocked by third-party antivirus software.	Disable or uninstall the third-party antivirus software and try again. For details, see Checking Whether the Access Is Blocked by Antivirus Software .
The cause is displayed in the error message.	If an error message is displayed during remote login, check the operation guide based on the error information. For details, see Checking Whether an Error Occurred During a Remote Login .

Checking the ECS Status

Check whether the ECS is in the **Running** state on the management console. If the ECS is stopped, start it and try to log in to the ECS again.

Figure 15-2 Checking the ECS status



Checking the Login Mode

Check the login mode you set when you created the ECS.

Figure 15-3 Login Mode

Login Mode: Password Key pair

Before using SSH to log in to a Linux ECS in password login mode, you must click Remote Login to log in to the ECS and enable SSH login. [Learn how to enable SSH login.](#)
Before changing the password of a Windows ECS in password login mode, you must click Remote Login to log in to the ECS.

Username: root

Password: Keep your password secure. The system cannot detect your password.

Confirm Password:

- **Password:** Check whether the login password is correct. If you forgot your password, reset the password. After you reset the password, restart the ECS for the new password to take effect.

- **Key pair:** If your ECS is authenticated using a key pair, parse the private key file to obtain a password.

Checking Whether the ECS Is Overloaded

If the bandwidth or CPU usage of the ECS is excessively high, login failures may occur.

If you have created an alarm rule in Cloud Eye, the system automatically sends an alarm notification to you when the bandwidth or CPU usage reaches the threshold specified in the rule.

To resolve this issue, perform the operations described in [Why Is My Windows ECS Running Slowly?](#)

- If the login failure is caused by high CPU usage, perform the following operations to reduce the CPU usage:
 - Stop certain processes that are not used temporarily and try again.
 - Verify that the Windows Update process is not running on the backend.
 - Restart the ECS.
 - Reinstall the ECS OS. Back up important data before the reinstallation.
 - If the ECS OS cannot be reinstalled due to important data, replace the disk attached to the ECS. To do so, back up data on the original disk, detach the disk from the ECS, attach the new disk to the ECS, and copy data to the new disk.

You can also upgrade the vCPUs and memory by [modifying the ECS specifications](#).

- If the login fails because the bandwidth exceeds the limit, perform the following operations:

For instructions about how to increase the bandwidth, see [Changing an EIP Bandwidth](#).

After you perform the preceding operations, try to remotely log in to the ECS again.

Checking Whether an ECS Has an EIP Bound

An ECS can access the Internet only after it has an EIP bound.

Before logging in to an ECS using MSTSC, make sure that an EIP has been bound to the ECS. For details, see [Assigning an EIP and Binding It to an ECS](#).

NOTE

If you log in to an ECS over an intranet, for example, using VPN or Direct Connect, you do not need to bind an EIP to the ECS.

Checking Whether the Network Is Normal

Use a local PC in another network or use another hotspot to access the ECS. Check whether the fault occurs on the local network. If so, contact the carrier to resolve this issue.


After you perform the preceding operations, try to remotely log in to the ECS again.

Checking Whether the Firewall Is Correctly Configured


Check whether the firewall is enabled.

1. Log in to the Windows ECS.
2. Click the Windows icon in the lower left corner of the desktop and choose **Control Panel > System and Security > Windows Firewall**.

Figure 15-4 Windows Firewall



3. Choose **Check firewall status > Turn Windows Firewall on or off**.
View and set the firewall status.

Figure 15-5 Turn off Windows Firewall

Ensure that the remote access port on the local end is allowed on the firewall. The default port is TCP 3389.

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

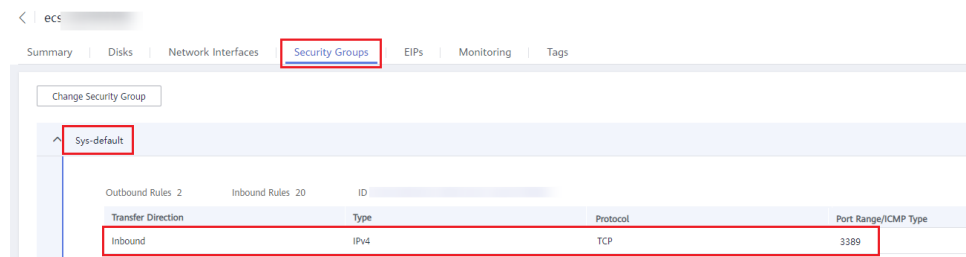
The default port is 3389. If you use another port, add that port in the inbound rule of the firewall.

After you perform the preceding operations, try to remotely log in to the ECS again.

Checking Whether the Remote Access Port Is Correctly Configured


1. 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 ECS details page, click the **Security Groups** tab and check port 3389 in the inbound rule of the security group.

Figure 15-6 Checking remote access ports

If you need to modify security group rules, see [Modifying a Security Group Rule](#).


2. Check whether the remote connection port is changed.
 - a. Choose **Start > Run**, enter **cmd**, and press **Enter**. In the CLI, enter **regedit** to open **Registry Editor**.
 - b. In **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\TerminalServer\WinStations\RDP-Tcp\PortNumber**, check whether the port is the default port 3389. If not, change the port to port 3389.



3. Check whether the number of connections is limited.

Check the internal remote desktop configuration of the ECS.

 - a. Choose **Start > Run**, enter **cmd**, and press **Enter**. In the CLI, enter **gpedit.msc** to open **Local Group Policy Editor**.
 - b. Choose **Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Connections**. Then, in the **Limit number of connections** dialog box, check whether the number of connections is limited.



NOTE

If **Limit number of connections** is set to **Enabled**, a remote connection to the Windows ECS may fail when the number of connections exceeds the limit. In such a case, disable **Limit number of connections** or set a larger limit for connections.

After you perform the preceding operations, try to remotely log in to the ECS again.

Checking the IP Address Whitelist for SSH Logins (with HSS Enabled)

After HSS is enabled, you can configure an IP address whitelist for SSH logins as required. The IP address whitelist controls SSH access to ECSs, effectively preventing account cracking.

After you configure the allowlist, SSH logins will be allowed only from IP addresses in the allowlist.

1. On the **Events** page, check whether a local host IP address is intercepted due to brute force cracking.
2. Check whether the IP address whitelist for SSH logins has been enabled. If it has been enabled, ensure that the IP address of the local host has been added to the IP address whitelist.

CAUTION

- Before enabling this function, ensure that all IP addresses that need to initiate SSH logins are added to the allowlist. Otherwise, you cannot remotely log in to your ECS through SSH.
- Exercise caution when adding a local IP address to the allowlist. This will make HSS no longer restrict access from this IP address to your ECSs.

Checking the Remote Desktop Protocol on the ECS

Make sure that the remote desktop protocol has been enabled on the ECS (only required for MSTSC logins).

Log in to the ECS using VNC and enable the remote desktop protocol.

For details, see [Enabling RDP](#).

Checking Whether the Access Is Blocked by Antivirus Software

Third-party antivirus software may lead to a failure in accessing the ECS.

If third-party antivirus software is running, check whether the remote connection is blocked by the software. If the remote connection is blocked, add the EIP bound to the ECS to the whitelist of the antivirus software and try to access the ECS again.

You can also disable or uninstall the third-party antivirus software and try to remotely log in to the ECS again.

Checking Whether an Error Occurred During a Remote Login

If an error message is displayed during remote login, check the operation guide based on the error information.

If the fault persists, record the resource details and fault occurred time, and contact technical support for assistance

If the fault persists after the preceding operations are performed, record the resource details and fault occurred time, and contact customer service for technical support.

15.6.4 Why Can't I Log In to My Linux ECS?

Symptom

A Linux ECS cannot be logged in to due to some reasons. For example, the network is abnormal, the firewall does not allow access to the local port for accessing the remote desktop, or the ECS vCPUs are overloaded.

This section describes how to troubleshoot login failures on a Linux ECS.

If you cannot log in to your Linux ECS, follow the instructions provided in [Checking the VNC Login](#). Then, locate the login fault by referring to [Fault Locating](#).

Checking the VNC Login

Check whether you can log in to the ECS using VNC on the management console.

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. In the **Operation** column of the target ECS, click **Remote Login**.

- (Optional) When the system displays "Press CTRL+ALT+DELETE to log on", click **Ctrl+Alt+Del** 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.

If the VNC login still fails, record the resource details and fault occurred time for further fault locating and analysis.

Fault Locating

If you can log in to the ECS using VNC but cannot log in to the ECS using a remote desktop connection, locate the fault as follows.

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.

Table 15-4 Possible causes and solutions


Possible Cause	Solution
The ECS is frozen or stopped.	Make sure that the ECS is in the Running state. For details, see Checking the ECS Status .
The entered username or password is incorrect.	The default username for Linux ECSs is root . If the password is incorrect, reset the password on the management console. For details, see Checking the Login Mode .
The ECS is overloaded.	If the bandwidth or CPU usage of the ECS is excessively high, login failures may occur. For details, see Checking Whether the ECS Is Overloaded .
The ECS has no EIP bound.	To log in to an ECS using RDP or MSTSC, ensure that the ECS has an EIP bound. For details, see Checking Whether an ECS Has an EIP Bound .
The access is blocked by the ISP.	Check whether you can access the ECS using another hotspot or network. For details, see Checking Whether the Network Is Normal .
The security group of the ECS denies inbound traffic on the remote login port.	Check whether the security group allows inbound traffic on the remote login port. For details, see Checking Whether the Security Group Is Correctly Configured .
The remote access port is incorrectly configured.	Check whether the remote access port is correctly configured on the local computer and the ECS. For details, see Checking Whether the Remote Access Port Is Correctly Configured .

Possible Cause	Solution
An IP address whitelist for SSH logins has been configured.	Check whether an IP address whitelist for SSH logins has been configured after HSS is enabled. For details, see Checking the IP Address Whitelist for SSH Logins (with HSS Enabled) .
An OS fault has occurred.	The file system is damaged. For details, see Checking Whether an OS Fault Has Occurred .
The access is blocked by third-party antivirus software.	Disable or uninstall the third-party antivirus software and try again. For details, see Checking Whether the Access Is Blocked by Antivirus Software .
The cause is displayed in the error message.	If an error message is displayed during remote login, check the operation guide based on the error information. For details, see Checking Whether an Error Occurred During a Remote Login .

Checking the ECS Status

Check whether the ECS is in the **Running** state on the management console. If the ECS is stopped, start it and try to log in to the ECS again.


Figure 15-7 Checking the ECS status



Checking the Login Mode

Check the login mode you set when you created the ECS.

Figure 15-8 Login Mode



- Password:** Check whether the login password is correct. If you forgot your password, reset the password. After you reset the password, restart the ECS for the new password to take effect.

- **Key pair**
 - For the first login, use an SSH key. For details, see [Logging In to a Linux ECS Using an SSH Key Pair](#).
 - For a non-first login, if you want to use the remote login function (VNC) provided by the management console, log in to the ECS using the SSH key and set the password.

Checking Whether the ECS Is Overloaded

If the bandwidth or CPU usage of the ECS is excessively high, login failures may occur.

If you have created an alarm rule in Cloud Eye, the system automatically sends an alarm notification to you when the bandwidth or CPU usage reaches the threshold specified in the rule.

To resolve this issue, perform the operations described in [Why Is My Linux ECS Running Slowly?](#)

- If the login failure is caused by high CPU usage, perform the following operations to reduce the CPU usage:
 - Stop certain processes that are not used temporarily and try again.
 - Restart the ECS.
 - Reinstall the ECS OS. Back up important data before the reinstallation.
 - If the ECS OS cannot be reinstalled due to important data, replace the disk attached to the ECS. To do so, back up data on the original disk, detach the disk from the ECS, attach the new disk to the ECS, and copy data to the new disk.

You can also upgrade the vCPUs and memory by [modifying the ECS specifications](#).

- If the login fails because the bandwidth exceeds the limit, perform the following operations:
For instructions about how to increase the bandwidth, see [Changing an EIP Bandwidth](#).

After you perform the preceding operations, try to remotely log in to the ECS again.

Checking Whether an ECS Has an EIP Bound

If you need to use a remote login tool (such as PuTTY or Xshell) to access the ECS, bind an EIP to the ECS.

For details, see [Assigning an EIP and Binding It to an ECS](#).

Checking Whether the Network Is Normal

Use a local PC in another network or use another hotspot to access the ECS. Check whether the fault occurs on the local network. If so, contact the carrier to resolve this issue.

After you perform the preceding operations, try to remotely log in to the ECS again.

Checking Whether the Security Group Is Correctly Configured

Check whether the local host can access port 22 on the ECS.

Run the following command to check whether port 22 is accessible:

```
telnet ECS private IP address
```

If port 22 is inaccessible, check whether port 22 is opened in the security group rule.

On the ECS details page, click the **Security Groups** tab and check that port 22 is configured in the inbound rule of the security group.

For details about how to modify a security group rule, see [Modifying a Security Group Rule](#).

After you perform the preceding operations, try to remotely log in to the ECS again.

Checking Whether the Remote Access Port Is Correctly Configured

Check ECS settings.

1. Check whether the sshd process is running.
2. Check whether your local PC is denied by the ECS.
 - a. Log in to the ECS and run the following command:

```
vi /etc/hosts.deny
```
 - b. If the IP address of the local PC is in the **hosts.deny** file, the ECS denies connection attempts from the local PC. In such a case, delete the IP address from the file.
3. Open the **/etc/ssh/ssh_config** file in the local PC and view the default login port. Then, open the **/etc/ssh/sshd_config** file in the ECS and check whether the SSH port is the default port 22.

```
# semanage port -a -t ssh_port_t -p tcp #PORTNUMBER
#
#Port 22
#AddressFamily any
```

After you perform the preceding operations, try to remotely log in to the ECS again.

Checking the IP Address Whitelist for SSH Logins (with HSS Enabled)

After HSS is enabled, you can configure an IP address whitelist for SSH logins as required. The IP address whitelist controls SSH access to ECSs, effectively preventing account cracking.

After you configure the allowlist, SSH logins will be allowed only from IP addresses in the allowlist.

1. On the **Events** page, check whether a local host IP address is intercepted due to brute force cracking.

2. Check whether the IP address whitelist for SSH logins has been enabled. If it has been enabled, ensure that the IP address of the local host has been added to the IP address whitelist.

 **CAUTION**

- Before enabling this function, ensure that all IP addresses that need to initiate SSH logins are added to the allowlist. Otherwise, you cannot remotely log in to your ECS through SSH.
 - Exercise caution when adding a local IP address to the allowlist. This will make HSS no longer restrict access from this IP address to your ECSs.
-

Checking Whether an OS Fault Has Occurred

- Password injection failure
The password failed to be injected using Cloud-Init.
- File system damaged after a forcible stop
There is a low probability that the file system is damaged after a forcible stop, which causes the ECS fails to be restarted. For details, see [Why Does a Forcibly-Stopped Linux ECS Fail to Be Restarted?](#)

After you perform the preceding operations, try to remotely log in to the ECS again.

Checking Whether the Access Is Blocked by Antivirus Software

Third-party antivirus software may lead to a failure in accessing the ECS.

If third-party antivirus software is running, check whether the remote connection is blocked by the software. If the remote connection is blocked, add the EIP bound to the ECS to the whitelist of the antivirus software and try to access the ECS again.

You can also disable or uninstall the third-party antivirus software and try to remotely log in to the ECS again.

Checking Whether an Error Occurred During a Remote Login

If an error message is displayed during remote login, check the operation guide based on the error information.

If the fault persists, record the resource details and fault occurred time, and contact technical support for assistance.

If the fault persists after the preceding operations are performed, record the resource details and fault occurred time, and contact customer service for technical support.

15.6.5 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.

15.6.6 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. Click  in the upper left corner and select your region and project.
 - c. Under **Computing**, click **Elastic Cloud Server**.
 - d. On the ECS list, click the name of an ECS for which you want to modify the security group rule.
 - e. On the ECS details page, click the security group in the **Security Groups** area to go to the security group details page.
 - f. 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](#).
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 15-9 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 15-10 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 15-11 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. Click  in the upper left corner and select your region and project.
 - c. Under **Computing**, click **Elastic Cloud Server**.
 - d. On the ECS list, click the name of an ECS for which you want to modify the security group rule.
 - e. On the ECS details page, click the security group in the **Security Groups** area to go to the security group details page.
 - f. 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](#).
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 15-12 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**

15.6.7 Why Cannot I Use a Non-Default SSH Port to Log In to My Linux ECS?

Symptom

After changing the default SSH port, you could not use the new port to log in to the ECS.

Possible Causes

- The access to the new port is not allowed in the security group.
- The new port is not enabled on the firewall.
- The new port is not added to the SSH configuration file.
- The hosts configuration file is incorrectly configured.

Checking Security Group Rules

Check whether the security group is correctly configured.

For example, if the new SSH port number is 2020, ensure that there is a security group rule without restriction in the outbound direction and allowing access to this port in the inbound direction.

Checking Firewall Rules

Run the **iptables** command to check whether the new SSH port, for example, port 2020 is enabled on the firewall.

1. Log in to the Linux ECS.
2. Take CentOS 7.5 as an example. Run the following command to edit the iptables file:

```
vi /etc/sysconfig/iptables
```

3. Add a rule for port 2020.

```
-A INPUT -m state --state NEW -m tcp -p tcp -dport 2020 -j ACCEPT
```

4. Restart iptables.

```
systemctl restart iptables
```

Checking the SSH Configuration File

Log in to the ECS and check the SSH configuration file.

1. Run the following command to check whether port 2020 has been configured:

```
vi /etc/ssh/sshd_config
```

2. If the port has not been configured, replace **#Port 22** with **Port 2020**.
3. Run the following command to restart SSH:

```
service sshd restart
```

Checking the hosts Configuration File

The **/etc/hosts.allow** and **/etc/hosts.deny** files of a Linux ECS are used to permit or deny an IP address or an IP address segment, respectively, to remotely access the ECS using SSH.

1. Add the following statement to **/etc/hosts.allow** to allow the IP address 192.168.1.3 to access the ECS using SSH:

```
sshd: 192.168.1.3
```

2. Check **/etc/hosts.deny**. If **sshd:all:deny** is contained, comment it out.

NOTE

If a rule is set in both **hosts.allow** and **hosts.deny**, the rule in **hosts.allow** takes precedence. For example, if "sshd: 192.168.1.3" is set in **hosts.allow** and "sshd:all:deny" is set in **hosts.deny**, the ECS allows only the SSH login from IP address 192.168.1.3.

15.6.8 Why Can't I Obtain the Password for Logging In to My Windows ECS Authenticated Using a Key Pair?

Symptom

A private key cannot be used to obtain the password for logging in to a Windows ECS that is authenticated using a key pair.

Possible Causes

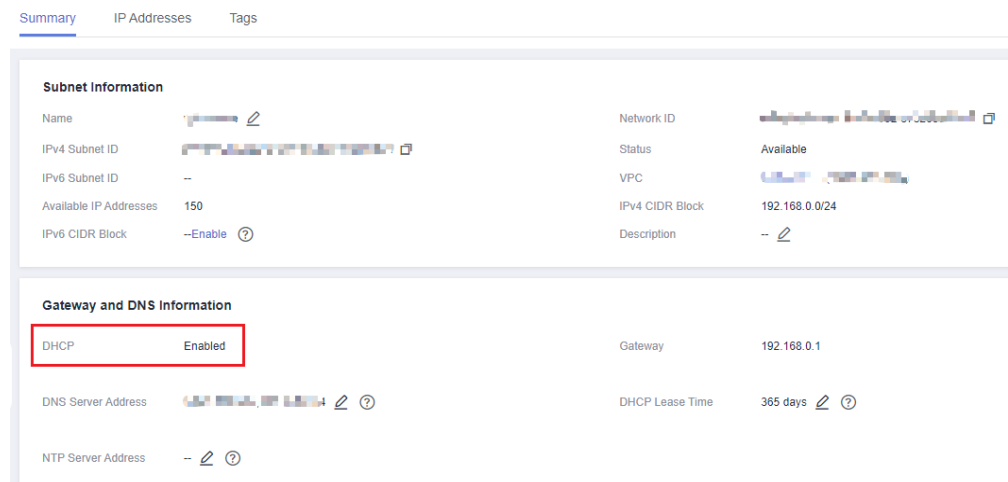
The password fails to inject using Cloudbase-Init due to:

- A network fault, leading to the failure of the connection from the ECS to the Cloudbase-Init server.
- No configuration on the image for Cloudbase-Init to obtain the password.
- Other reasons.


Solution

If logging in to an ECS with Cloudbase-Init enabled failed, perform the following operations to locate the fault:


1. Ensure that Cloudbase-Init has been correctly configured on the image that was used to create the ECS.
 - If Cloudbase-Init has not been configured, your ECS will not allow customized configurations, and you can log in to it only by using the original image password.
 - The ECSs created using a public image have Cloudbase-Init installed by default. You do not need to install and configure Cloudbase-Init anymore.
 - If you created your ECS by using an external image file, install and configure Cloudbase-Init.
For details, see [Installing and Configuring Cloudbase-Init](#).
2. Ensure that the key pair for logging in to the ECS is correct.
The key used for obtaining the password must be the key used during the ECS creation.
3. Ensure that DHCP is enabled in the VPC to which the ECS belongs.
On the management console, check whether DHCP has been enabled in the target subnet.

Figure 15-13 DHCP

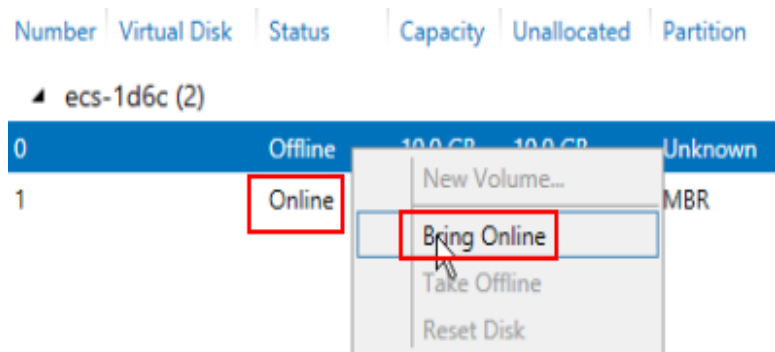
4. Ensure that the ECS has an EIP bound.
5. Ensure that traffic to and from port 80 is allowed in security group rules.

Figure 15-14 Security group rules for port 80


6. Check Cloudbase-Init logs to identify the cause.
 - a. Stop the affected ECS and detach the system disk from it.

Figure 15-15 Detaching the system disk

- b. Use a public image to create a temporary Windows ECS and attach the system disk detached in 6.a to the ECS.
- c. Log in to the temporary ECS, open the **Server Manager** page, choose **File and Storage Services > Volumes > Disks**, right-click the offline disk, and choose **Online** from the shortcut menu.

Figure 15-16 Setting disk online

- d. Switch to the **cloudbase-init** file in **/Program Files/Cloudbase Solution/Cloudbase-Init/log** of this disk to view the log for fault locating.

Figure 15-17 cloudbase-init

15.6.9 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 15-5](#).

Table 15-5 Browser version requirements

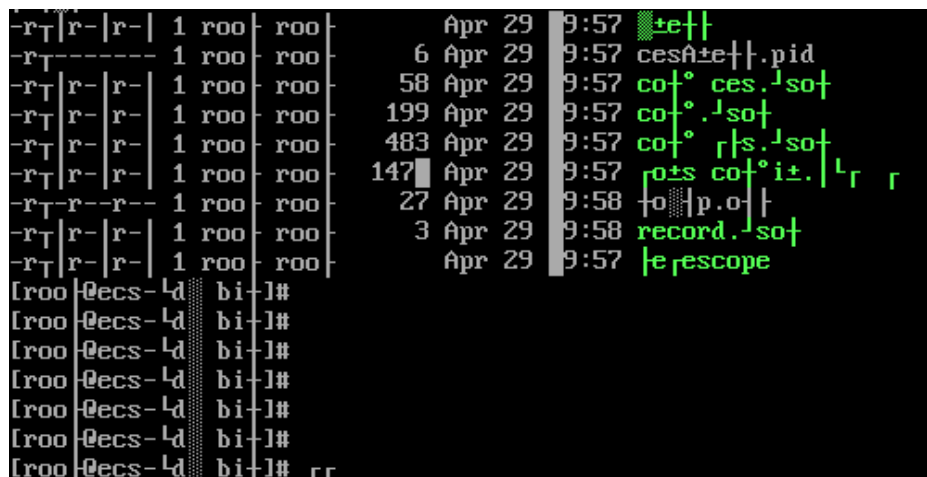
Browser	Version
Google Chrome	31.0-75.0
Mozilla Firefox	27.0-62.0
Internet Explorer	10.0-11.0

15.6.10 Why Are Garbled Characters Displayed When I Log In to My ECS Using VNC?

Symptom

After I attempt to log in to my Linux ECS using VNC, garbled characters are displayed, as shown in [Figure 15-18](#).

Figure 15-18 Garbled characters on the VNC-based login page



```
-rT | r- | r- | 1 roo | roo | Apr 29 | 9:57 | te++
-rT | r- | r- | 1 roo | roo | 6 Apr 29 | 9:57 | cesAte++ | .pid
-rT | r- | r- | 1 roo | roo | 58 Apr 29 | 9:57 | cof° | ces. | so+
-rT | r- | r- | 1 roo | roo | 199 Apr 29 | 9:57 | cof° | . | so+
-rT | r- | r- | 1 roo | roo | 483 Apr 29 | 9:57 | cof° | r | s. | so+
-rT | r- | r- | 1 roo | roo | 147 Apr 29 | 9:57 | ros | cof° | i±. | Lr | r
-rT | r- | r- | 1 roo | roo | 27 Apr 29 | 9:58 | to | p. | o | t
-rT | r- | r- | 1 roo | roo | 3 Apr 29 | 9:58 | record. | so+
-rT | r- | r- | 1 roo | roo | Apr 29 | 9:57 | e | rescope
[roo | @ecs- | Ld | bi- | ]#
[roo | @ecs- | Ld | bi- | ]#
[roo | @ecs- | Ld | bi- | ]#
[roo | @ecs- | Ld | bi- | ]#
[roo | @ecs- | Ld | bi- | ]#
[roo | @ecs- | Ld | bi- | ]#
[roo | @ecs- | Ld | bi- | ]#
```

Possible Causes

The **cat** command was executed to display a large binary file, leading to garbled characters.

Solution

Log in to the ECS as user **root** and run the following command for recovery:

```
reset
```

NOTE

The **reset** command is used to re-initialize the ECS and refresh the terminal display. After the **reset** command is executed, the garbled characters are cleared and the fault is rectified.

15.6.11 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.

15.6.12 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.

15.6.13 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.

15.6.14 What Should I Do If Error Code 1006 or 1000 Is Displayed When I Log In to an ECS Through the Management Console?

Symptom

When I attempted to remotely log in to an ECS using VNC, the system displayed error code 1006, as shown in [Figure 15-19](#).

Figure 15-19 Error message displayed in a VNC-based remote login



Possible Causes

- The ECS is abnormal.
- Another user has logged in to the ECS.
- No operations are performed on the ECS and it is automatically disconnected.

Troubleshooting

1. Log in to the ECS again using VNC.
 - If the login is successful, no further action is required.
 - If the fault persists, go to [2](#).
2. Check whether the ECS is normal.

Error code 1006 is displayed if the ECS is stopped, deleted, being migrated or restarted, or encounters a connection timeout.
3. Check whether another user has logged in to the ECS.

If yes, you can log in to the ECS only after that user logs out.

15.6.15 Why No Audio File Can Be Properly Played on My Windows ECS Logged In Using VNC?

Symptom

When I logged in to my Windows ECS using MSTSC, audio files can be properly played. However, when I logged in to that ECS using VNC, audio files failed to be played.

Possible Causes

VNC does not support audio playing.

Solution

Use your local PC (running Windows 7, for example) to play the audio files.

1. Start your local PC.


NOTE

Start your local PC, instead of logging in to your Windows ECS.

2. Press **Win+R** to start the **Run** text box.
3. Enter **mstsc** and click **OK**.


The **Remote Desktop Connection** window is displayed.

Figure 15-20 Remote Desktop Connection




4. Click **Options** in the lower left corner and click the **Local Resources** tab.

Figure 15-21 Local Resources



5. In the **Remote audio** pane, click **Settings**.

Figure 15-22 Setting remote audio playback

6. In the **Remote audio playback** pane, select **Play on this computer**.

15.6.16 How Can I Change the Resolution of a Windows ECS?


Scenarios

You can change the resolution of Windows ECSs.

Solution 1: Using VNC

The operations of changing an ECS resolution vary according to the Windows OS. This section uses the Windows Server 2016 Standard 64-bit edition as an example to describe how to change the resolution of a Windows ECS.

1. Log in to the ECS using VNC.
2. Right-click the desktop and choose **Display settings** from the shortcut menu.


Figure 15-23 Display settings

3. On the **Settings** page, click the **Display** tab and then **Advanced display settings**.

 **NOTE**


If the remote desktop is not fully displayed, set **Change the size of text, apps, and other items** to **100%**.

Figure 15-24 Settings



4. In the **Resolution** drop-down list, select the desired resolution.

Figure 15-25 Setting a resolution



5. Click **Apply**.


Solution 2: Using MSTSC

Before remotely logging in to your ECS using MSTSC, change the resolution of the Windows ECS.

1. On your local computer (client), click **Start**.

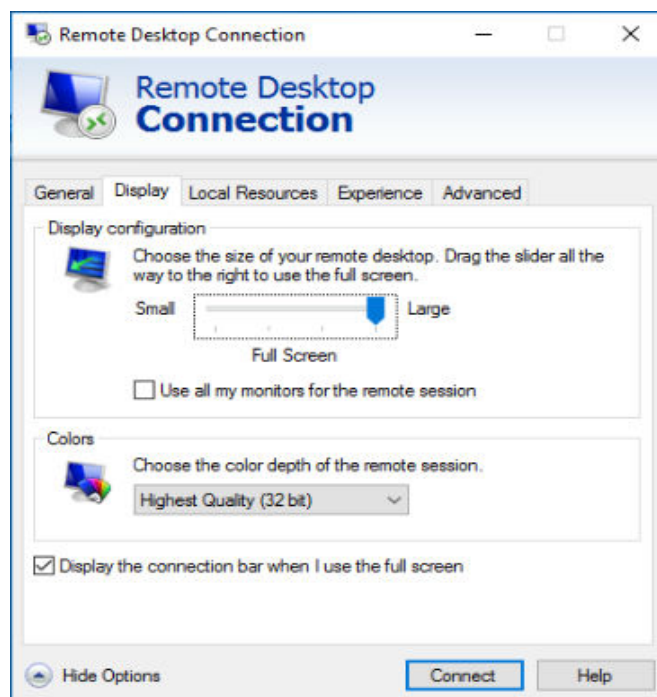
2. In the **Search programs and files** text box, enter **mstsc**.
3. In the **Remote Desktop Connection** window, click **Show Options** in the lower left corner.

Figure 15-26 Remote Desktop Connection



4. Click the **Display** tab. Then, in the **Display configuration** pane, set the resolution.

Figure 15-27 Display



5. Use MSTSC to log in to the ECS.

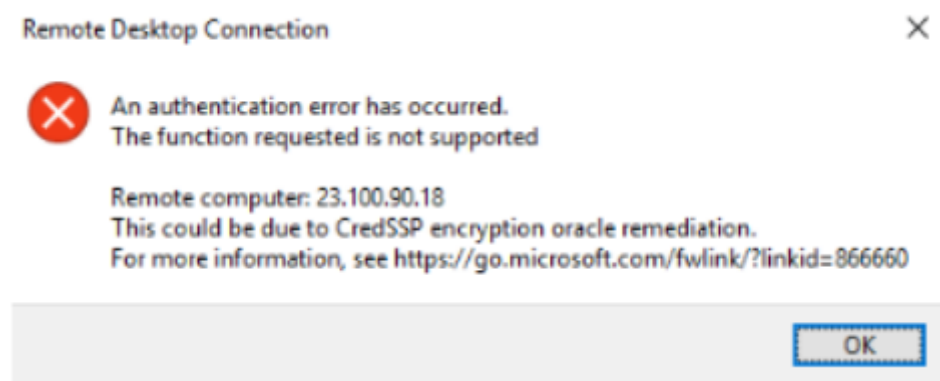
15.6.17 Why Does an Authentication Failure Occurs After I Attempt to Remotely Log In to a Windows ECS?

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 15-28](#), 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 15-28 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 15-29 Remote settings



5. Click **OK**.

15.6.18 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 15-30 Cannot connect to the remote computer



Possible Cause

- 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.

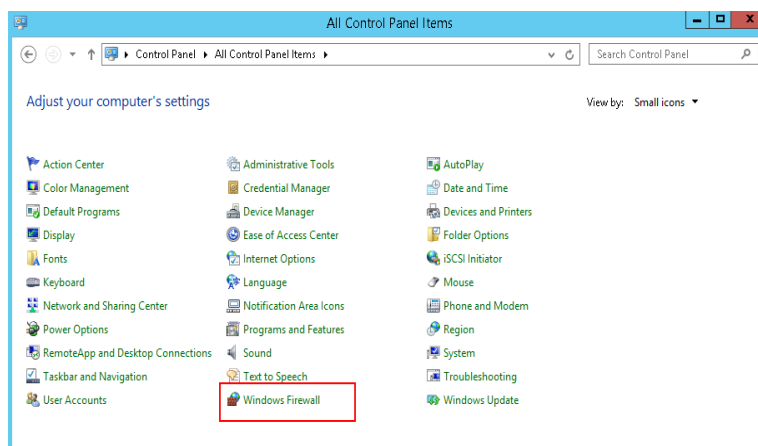
On the ECS details page, click the **Security Groups** tab and check port 3389 in the inbound rule of the security group.

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 15-31 Windows Firewall



3. Click **Turn Windows Firewall on or off**.
View and set the firewall status.


Figure 15-32 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 15-33 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 of the Windows ECS: Select **Allow remote connections to this computer**. The procedure is as follows:

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 remote connections to this computer**.

Figure 15-34 Remote settings

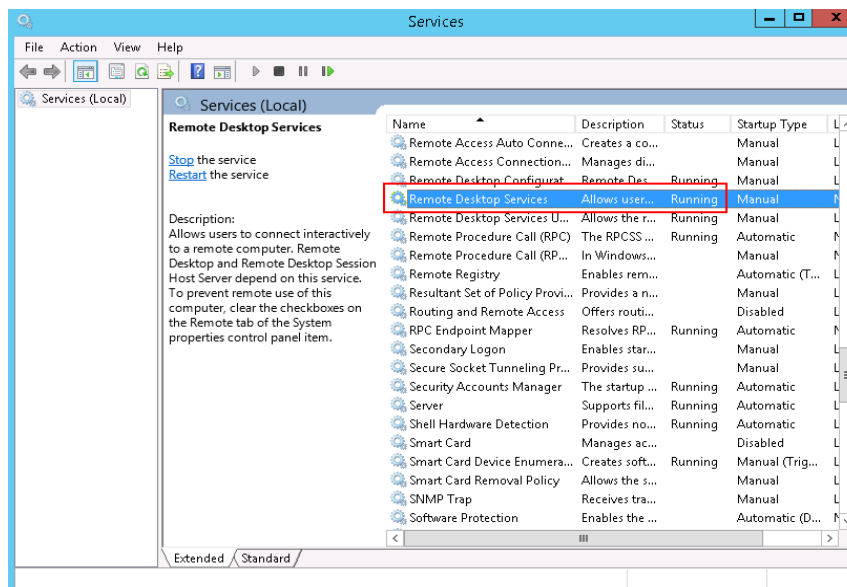


5. Click **OK**.

Checking Remote Desktop Services

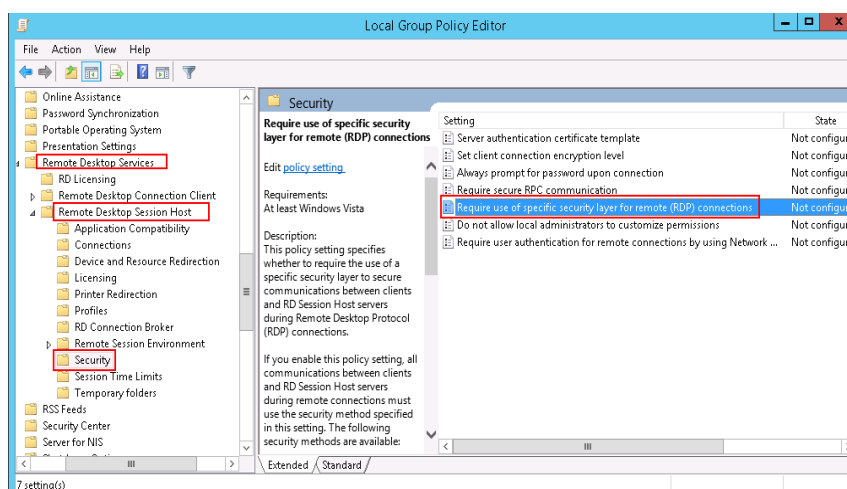
1. Open the Windows search box, enter **services**, and select **Services**.

2. In the **Services** window, restart **Remote Desktop Services**. Ensure that **Remote Desktop Services** is in the **Running** status.


Figure 15-35 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 15-36 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 15-37 Setting security layer to RDP

15.6.19 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.

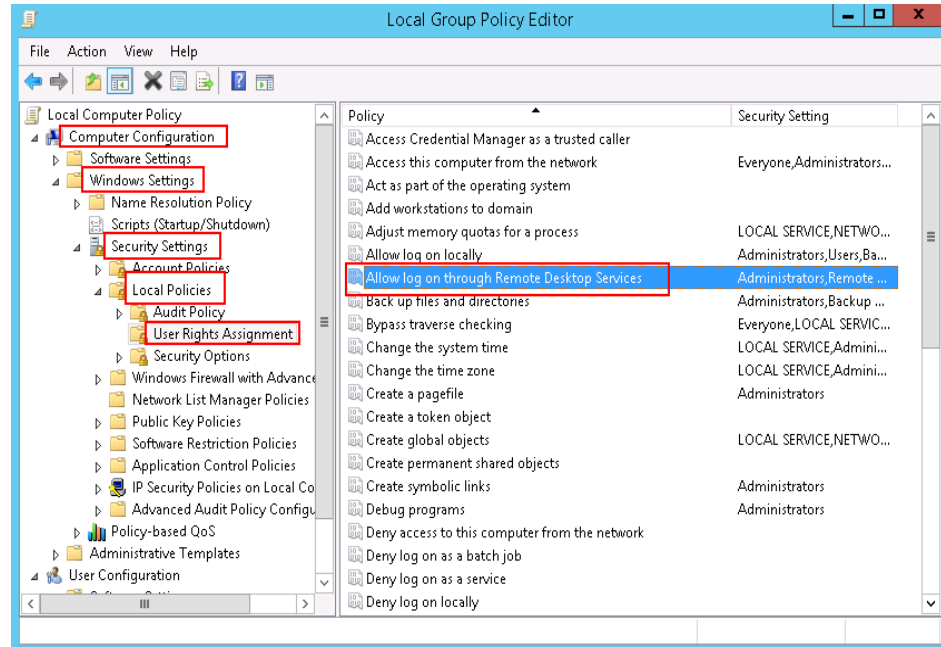
Figure 15-38 Remote login right missing.

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**.

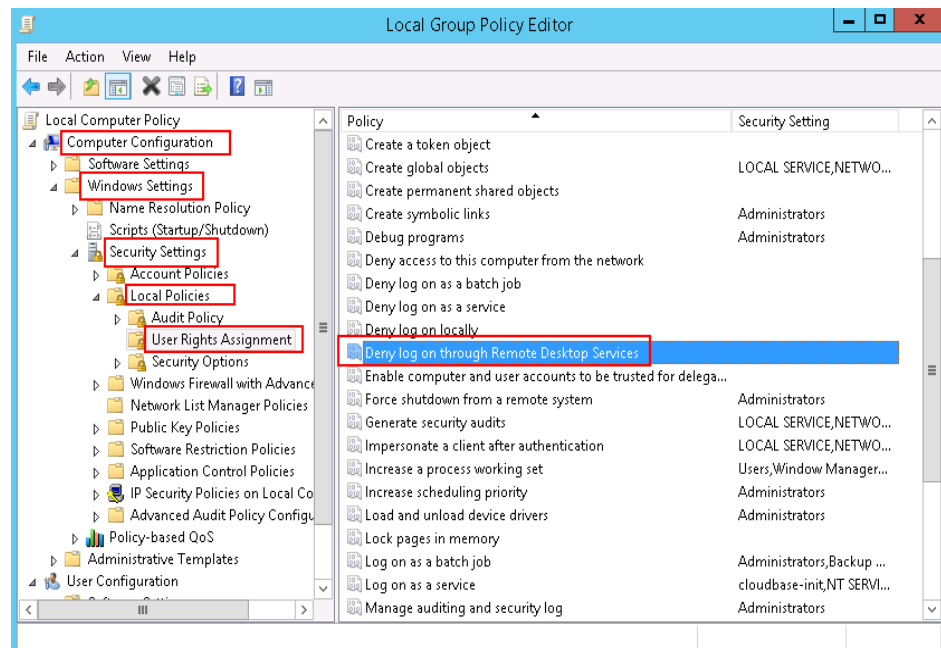
- a. Locate and double-click **Allow log on through Remote Desktop Services**. Ensure that **Administrators** and **Remote Desktop Users** have been added.

Figure 15-39 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 15-40 Deny log on through Remote Desktop Services properties




15.6.20 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 15-41 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 15-42 Deleting roles



3. In the displayed dialog box, deselect **Remote Desktop Session Host** and keep clicking **Next** till you finish the operation.

Figure 15-43 Deselecting Remote Desktop Session Host

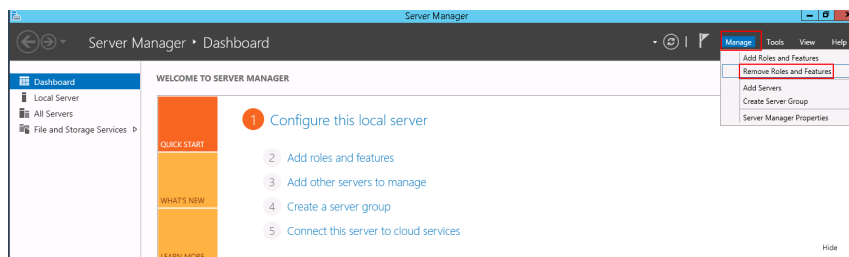


4. Click **Delete**.
5. Restart the ECS.

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 15-44 Deleting roles and features



3. Select the destination server and click **Next**.
4. Deselect **Remote Desktop Services**.

Figure 15-45 Deselecting Remote Desktop Services




5. Click **Delete**.
6. Restart the ECS.

15.6.21 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 15-46](#).

Figure 15-46 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.

15.6.22 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 15-47 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.

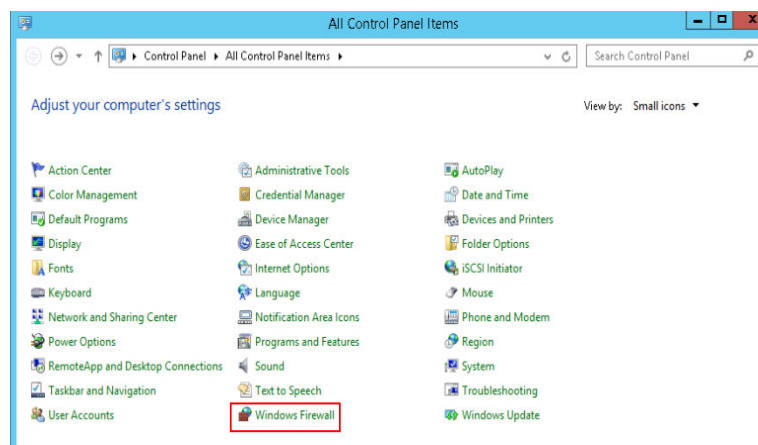
Solution

Step 1 Check security group settings.

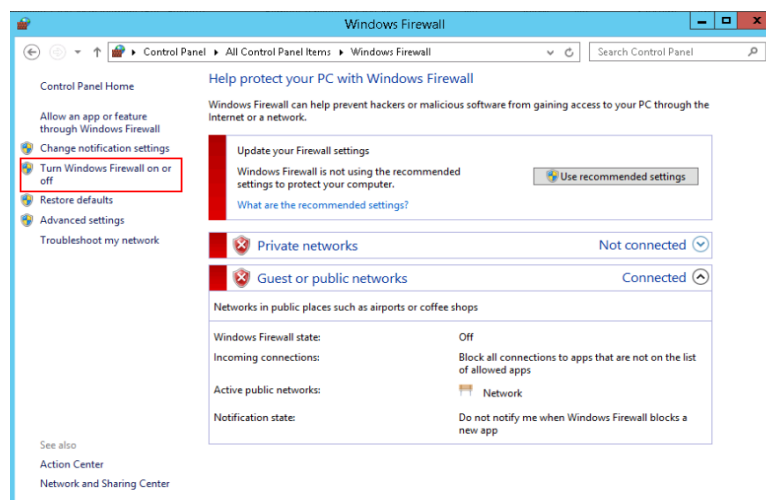
Check whether port 3389 is allowed in inbound direction. If it is allowed, go to [Step 2](#).

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 15-48 Checking port 3389




As shown in **Figure 15-48**, 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 15-49 Checking System process



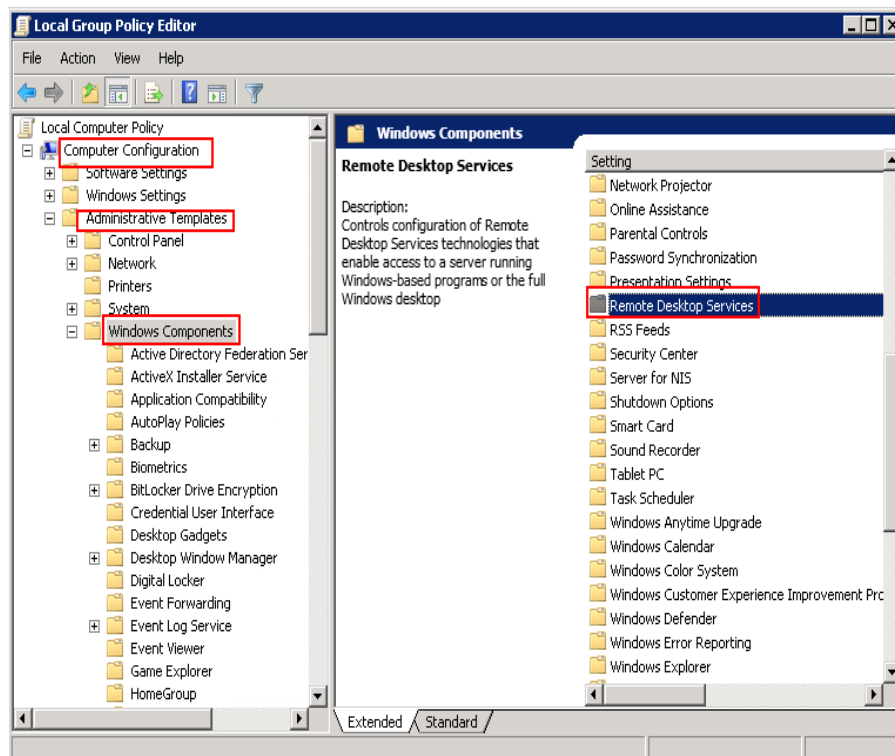
4. If port 3389 is used by HTTP protocols, it indicates that the port is used by IIS.
5. Enter **http://127.0.0.1:3389** in the address box of the browser and press **Enter**. Check whether the website can be visited normally.
6. Change the port used by IIS and restart IIS.

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.

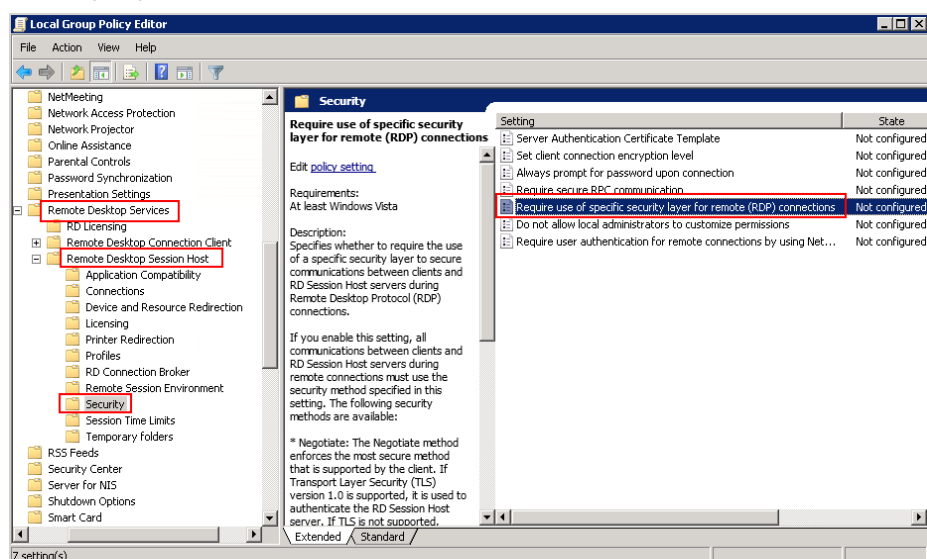
Step 5 Check the remote desktop session host configuration.

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**.

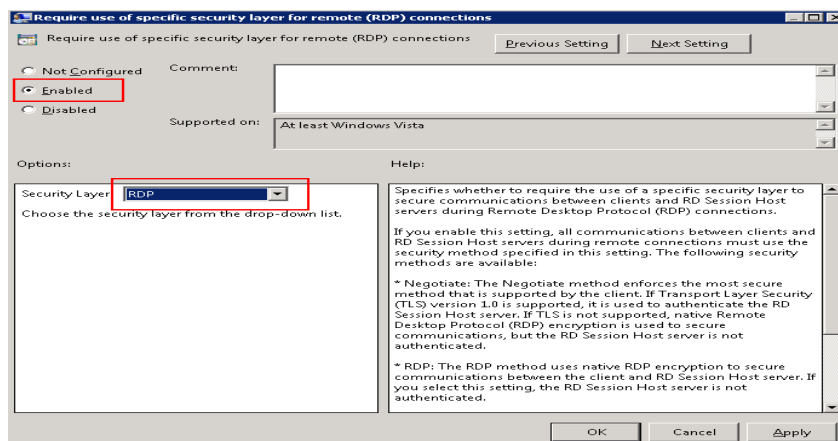
Figure 15-50 Remote Desktop Services

5. **Remote Desktop Session Host > Security.**

Figure 15-51 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 15-52 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 15-53 Updating the group policy



----End

15.6.23 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 15-54 Event viewer



3. In the **Event Level** pane, select event levels.

Figure 15-55 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


15.6.24 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 15-56 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 15-57 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 15-58 Preferences setting

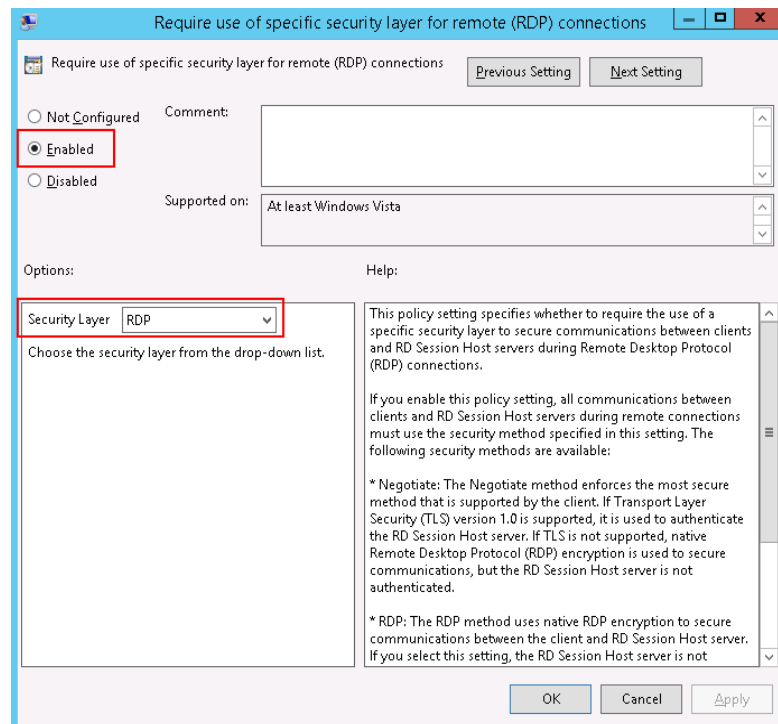
2. Select **Security** and modify the parameter settings according the following figure.

Figure 15-59 Security setting

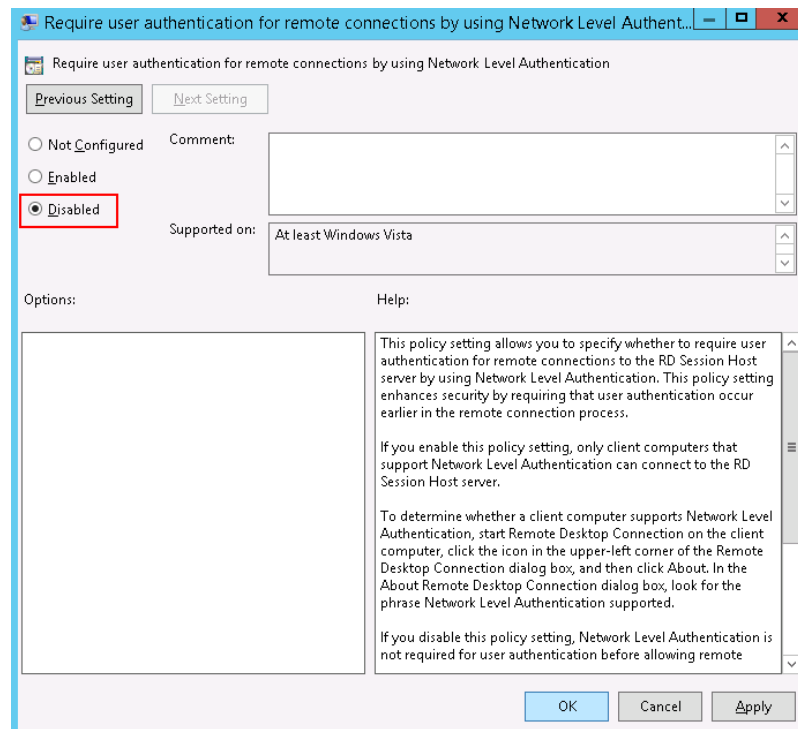
3. Remotely connect to the Windows ECS again. If the error message **Invalid certificate or associated chain** is still displayed, go to 4.
4. Log in to the Windows ECS using VNC.
5. Press **Win+R** to start the **Open** text box.
6. Enter **gpedit.msc** to access the Local Group Policy Editor.
7. In the left navigation pane, choose **Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Security**.

Figure 15-60 Remote Desktop Session Host

8. Modify the following parameters as prompted:
 - Enable **Require use of specific security layer for remote (RDP) connections**.

Figure 15-61 Require use of specific security layer for remote (RDP) connections

- Disable **Require user authentication for remote connections by using Network Level Authentication**.

Figure 15-62 Remote connection authentication

9. Close the group policy editor and restart the ECS.

15.6.25 Why Does the System Display a Message Indicating Invalid Credentials When I Attempt to Access a Windows ECS?

Symptom

When you use a local PC running Windows to access a Windows ECS using RDP (for example, MSTSC), the system displays a message indicating that the credentials are invalid.

Solution

Perform the following steps to rectify the fault. After completing each step, try to access the ECS to check whether the fault is rectified. If the fault persists, go to the next step.

Step 1: Change Network Access Policy

Step 2: Modify Credentials Delegation

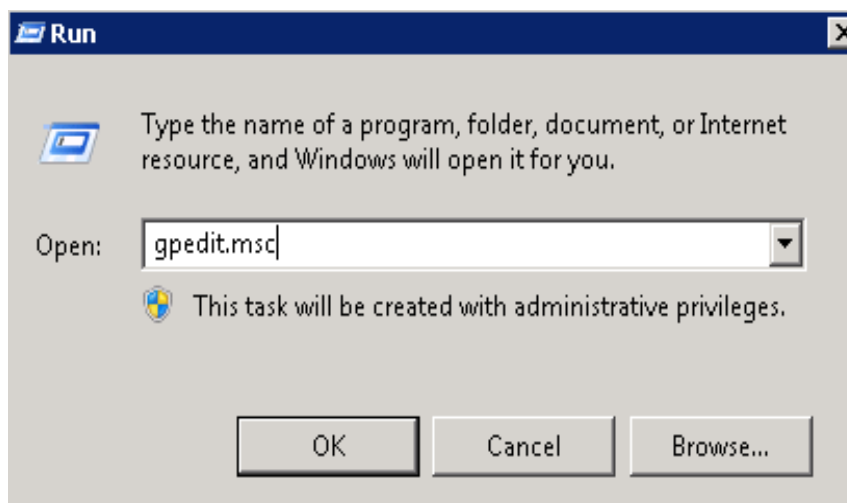
Step 3: Set the Credentials of the Local Server

Step 4: Disable Password Protected Sharing

Step 1: Change Network Access Policy

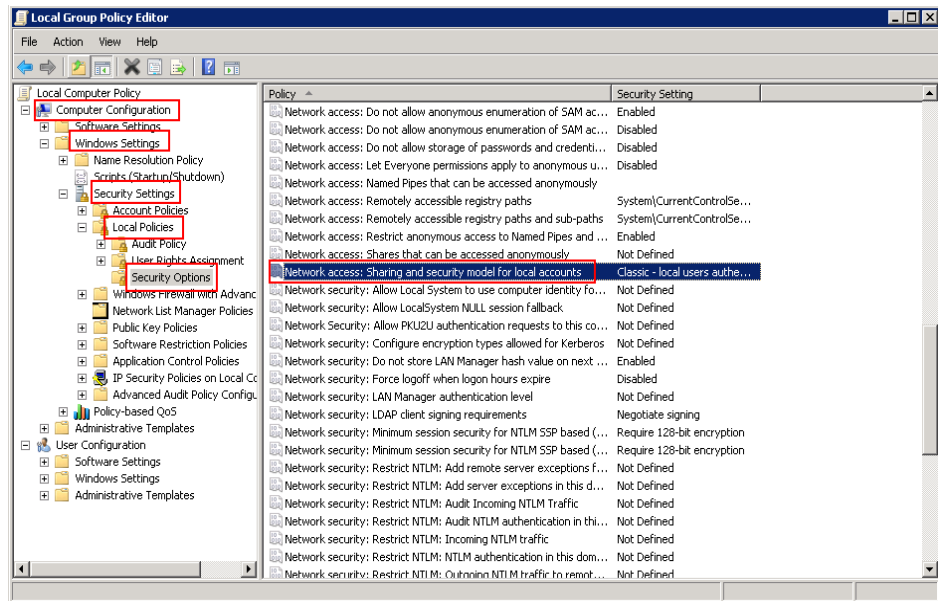
1. Log in to the ECS using VNC on the management console.
2. Choose **Start > Run**. In the **Run** dialog box, enter **gpedit.msc** and click **OK** to start **Local Group Policy Editor**.

Figure 15-63 gpedit.msc



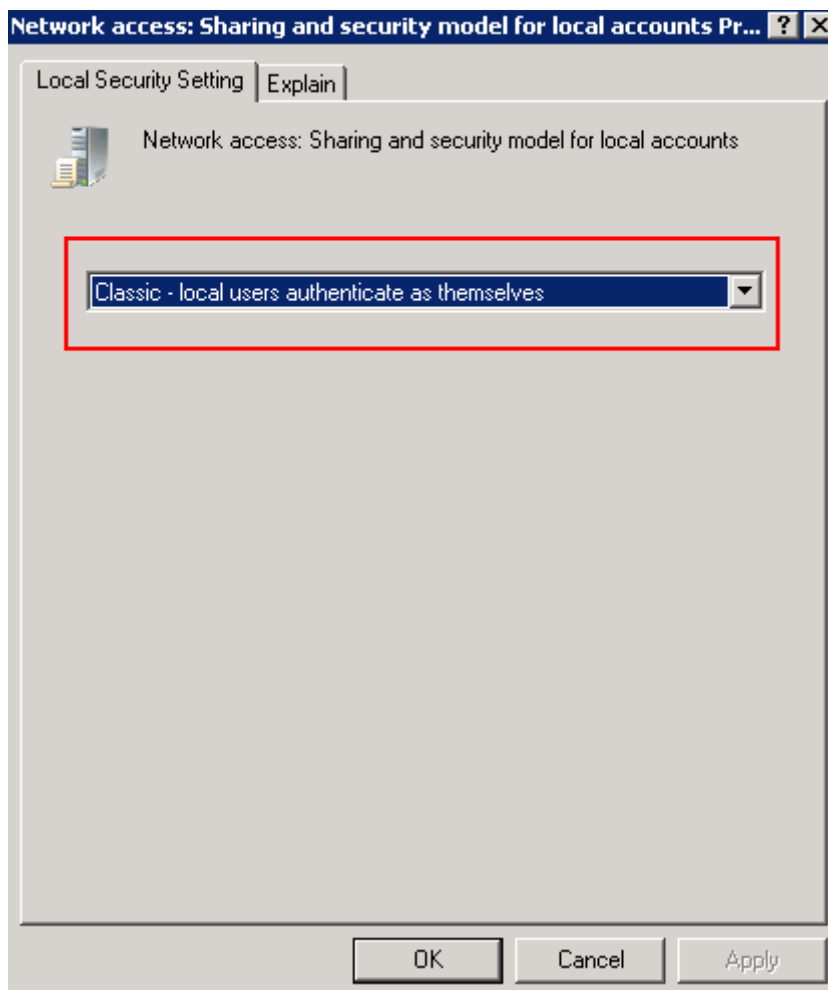
3. Choose **Computer Configuration > Windows Settings > Security Settings > Local Policies > Security Options** and click **Network access: Sharing and security model for local accounts**.

Figure 15-64 Locating the network access policy



4. Select **Classic - local users authenticate as themselves** and click **OK**.

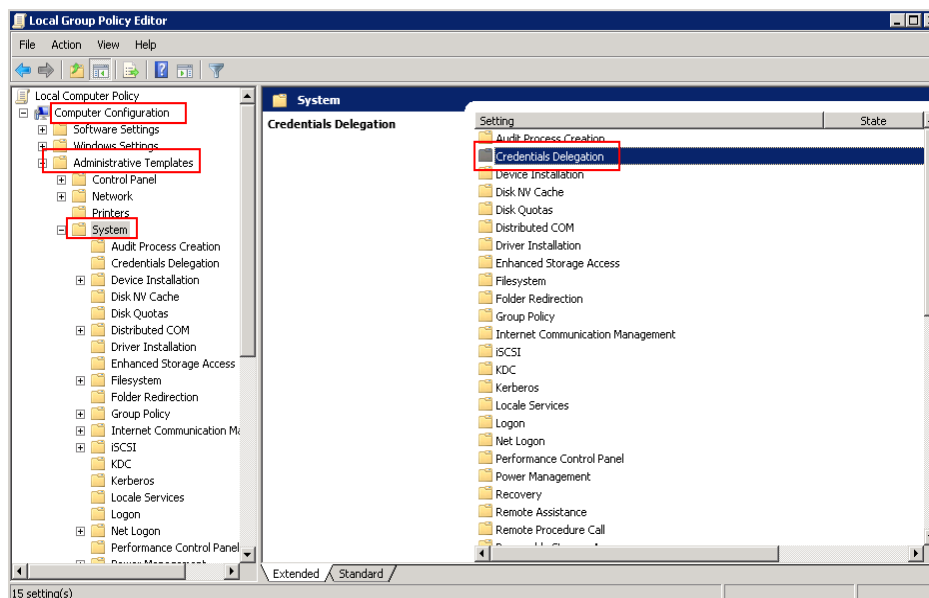
Figure 15-65 Changing the network access policy



Step 2: Modify Credentials Delegation


1. Log in to the ECS using VNC on the management console.
2. Choose **Start > Run**. In the **Run** dialog box, enter **gpedit.msc** and click **OK** to start **Local Group Policy Editor**.
3. Choose **Computer Configuration > Administrative Templates > System** and locate **Credentials Delegation**.

Figure 15-66 Locating the network access policy




4. Double-click **Allow Delegating Saved Credentials with NTLM-only Server Authentication** and click **OK**.

Figure 15-67 Allow Delegating Saved Credentials with NTLM-only Server Authentication




5. Select **Enabled** and enter **TERMSRV/*** in the **Show Contents** text box. **TERMSRV/*** indicates the terminal server running on all computers.

Figure 15-68 Enabled



6. Refresh the group policy for the settings to take effect.
7. Choose **Start > Run**. In the **Run** dialog box, enter **gpupdate /force** and press **OK** to update the group policy.


Figure 15-69 Updating the group policy



Step 3: Set the Credentials of the Local Server


1. Open the control panel on the local server and choose **Credential Manager > Windows Credentials**.

Figure 15-70 Credential Manager



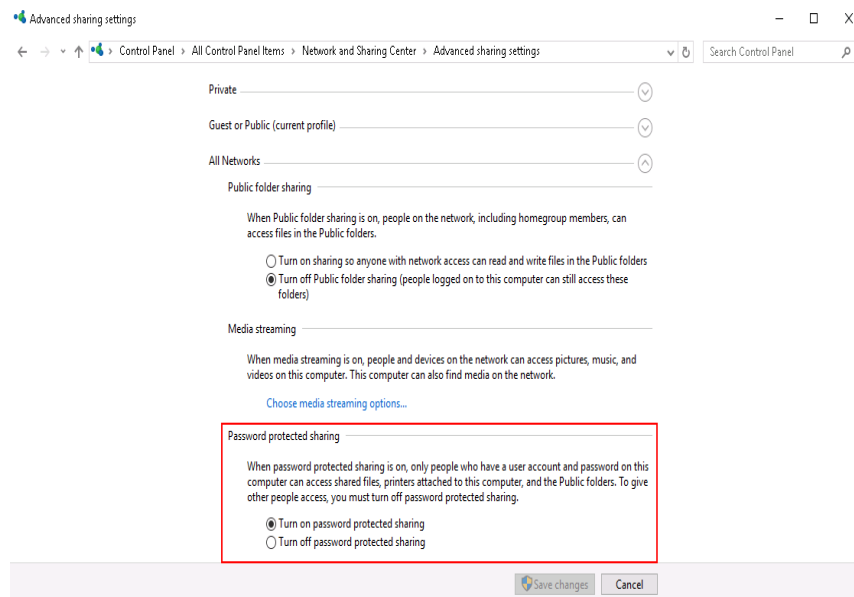
2. Check whether the credential of the target ECS is contained in the Windows credentials. If there is no credential, add one.
 - Internet or network address: IP address of the ECS
 - **User name:** Username for logging in to the ECS
 - **Password:** Password for logging in to the ECS

Figure 15-71 Add a Windows Credential



Step 4: Disable Password Protected Sharing

1. Log in to the ECS.
2. Choose **Start > Control Panel > All Control Panel Items > Network and Sharing Center > Change advanced sharing settings.**
3. In the **Password protected sharing** pane, select **Turn off password protected sharing.**

Figure 15-72 Turn off password protected sharing

4. Click **Save changes**.

15.6.26 Why Does an Internal Error Occur When I Log In to My Windows ECS?

Symptom

When you attempt to log in to your Windows ECS using MSTSC, the system displays an error message indicating an internal error.

Solution

1. On the local server, run **cmd** as an administrator.
2. Run the **netsh winsock reset** command.

```
Administrator: C:\Windows\system32\cmd.exe
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>netsh winsock reset

Sucessfully reset the Winsock Catalog.
You must restart the computer in order to complete the reset.
```

3. Restart the local server.
4. Log in to the ECS again.

If you still cannot log in to the ECS, check your local network. Change the network (for example, use your phone's mobile data) and check whether you can log in to the ECS remotely.

If you can remotely log in to the ECS using your phone's mobile data, your local network is abnormal. Restart your local network (for example, restart the router).


If the fault still persists, record the resource details and fault occurred time, and contact customer service for technical support.

15.6.27 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 15-73 Protocol error



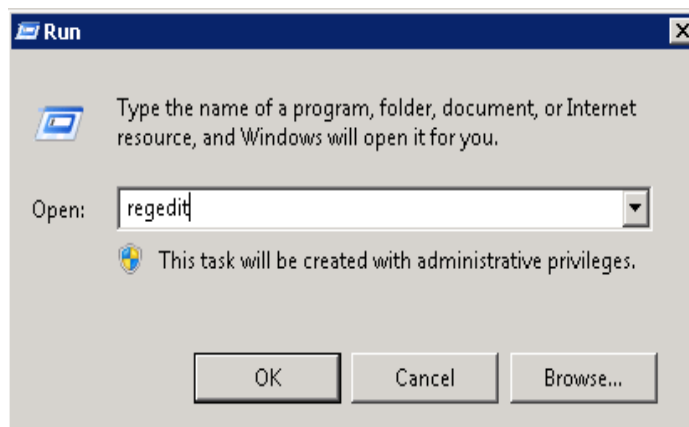
Possible Causes

The registry subkey Certificate is damaged.

Solution

1. In the **Run** dialog box, enter **regedit** and click **OK** to open the registry editor.

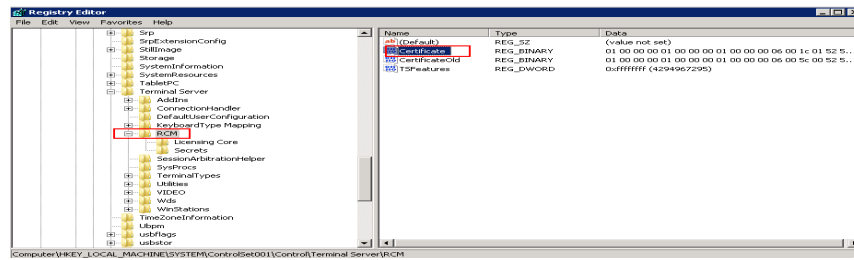
Figure 15-74 Opening the registry editor



2. Choose **HKEY_LOCAL_MACHINE > SYSTEM > ControlSet001 > Control > Terminal Server > RCM**.


3. Delete Certificate.

Figure 15-75 Deleting Certificate



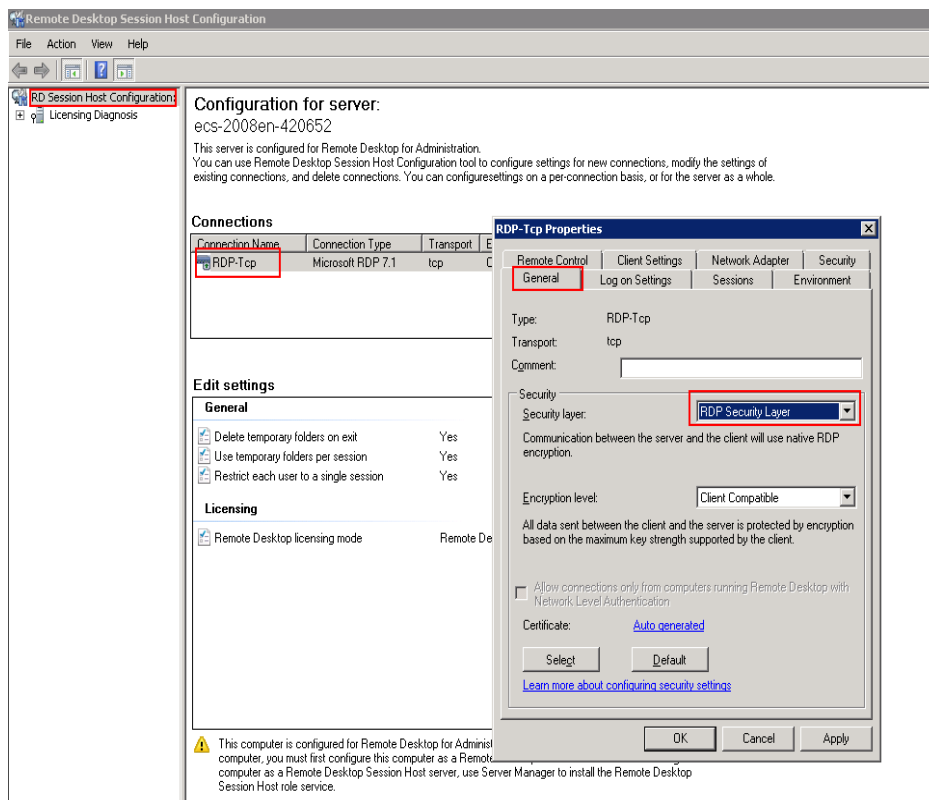
4. Restart the ECS.
5. Choose **Start > Administrative Tools > Remote Desktop Services > Remote Desktop Session Host Configuration**.

Figure 15-76 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 15-77 RDP-Tcp properties




15.6.28 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 15-78 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.

15.6.29 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 15-79 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.

15.6.30 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 15-80 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 15-81 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 15-82 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 15-83 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 15-84 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

15.6.31 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.

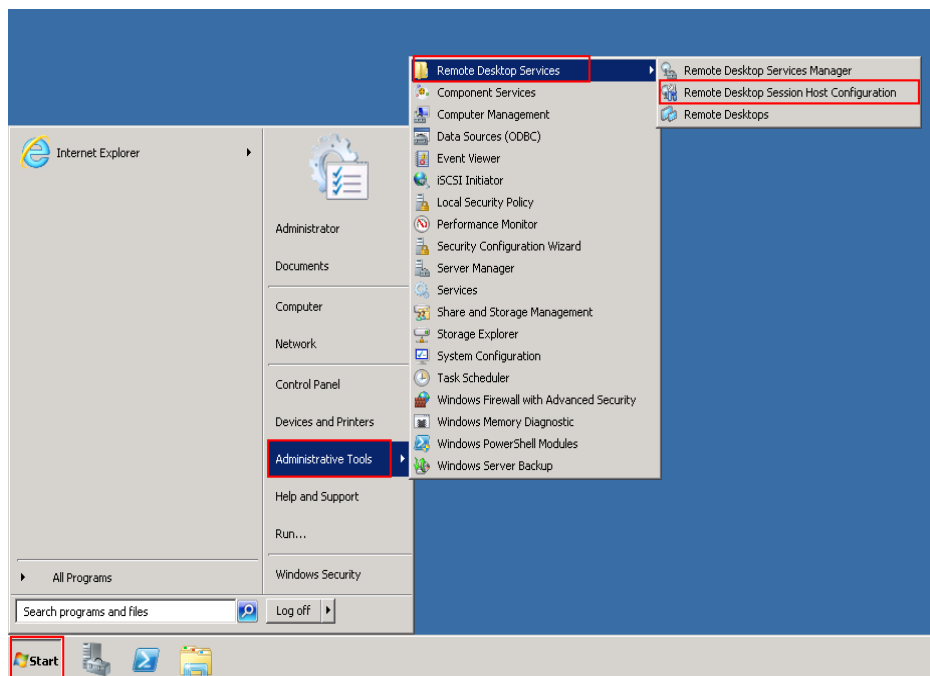
Figure 15-85 Ended remote desktop session



Windows Server 2008

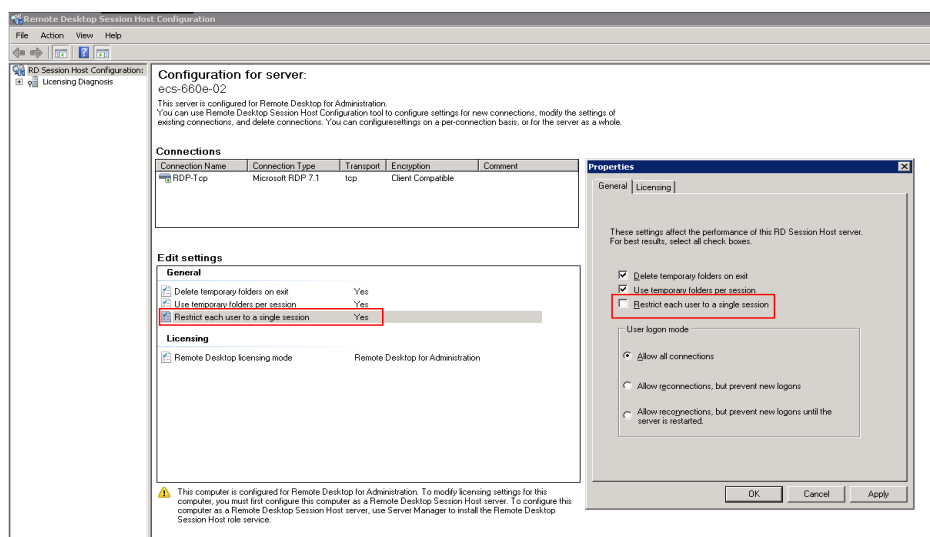
1. Choose **Start > Administrative Tools > Remote Desktop Services > Remote Desktop Session Host Configuration**.

Figure 15-86 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 15-87 Modifying the configuration

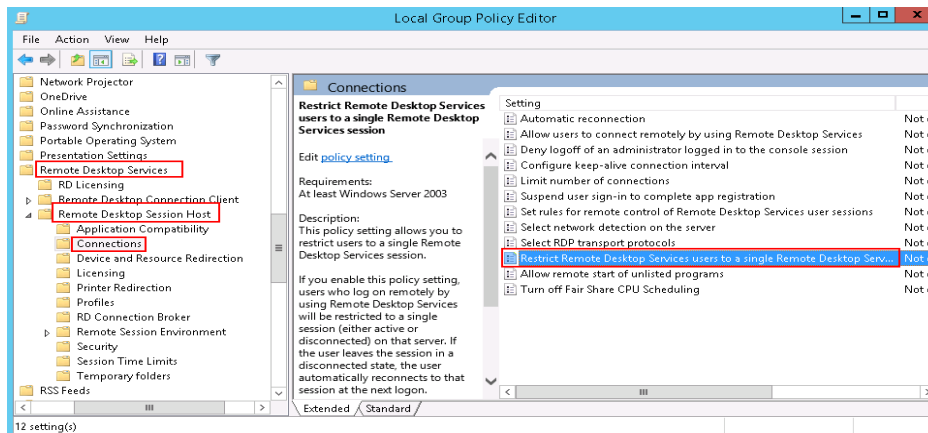


Windows Server 2012

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 15-88 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 15-89 Modifying the configuration



4. Run **gpupdate/force** to update the group policy.

15.6.32 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 15-90 Remote Desktop Services



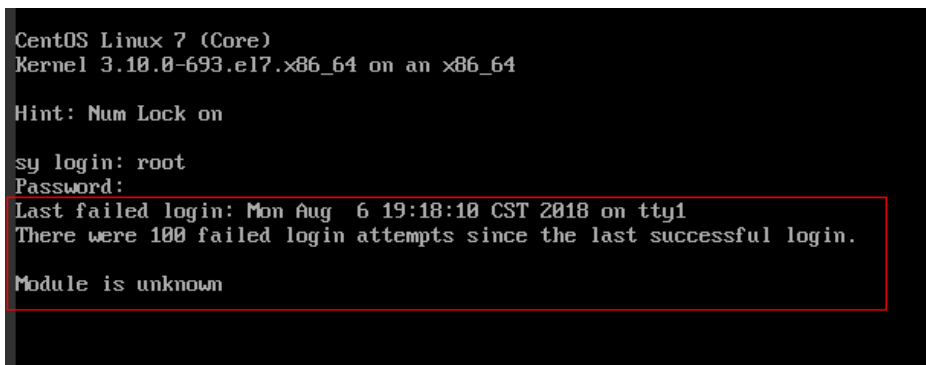
4. Remotely connect to the ECS again.
If the connection still fails, run the `cmd` command on the local server as the administrator, run the `netsh winsock reset` command to restore the default network connection configurations, and then retry the remote connection.

15.6.33 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 15-91 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.

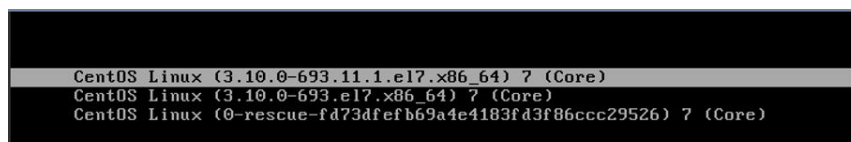
Root Cause

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 15-92 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 customer service 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 15-93 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 15-94 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 15-95 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 15-96 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
```

15.6.34 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 15-97 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 15-98 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 customer service 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 15-99 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 15-100 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
```

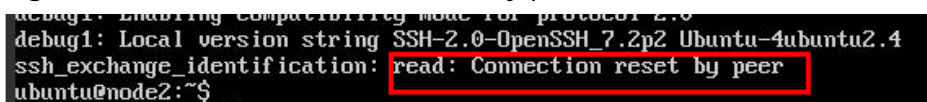
- g. Run the following command to go to the `/sysroot` directory:
`# chroot /sysroot`
2. Run the following command to view the `fs.nr_open` value:
`sysctl fs.nr_open`
3. Change the `nofile` value in `/etc/security/limits.conf` so that the value is smaller than the `fs.nr_open` value obtained in 2.
`vi /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.

15.6.35 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 15-101 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.

15.6.36 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

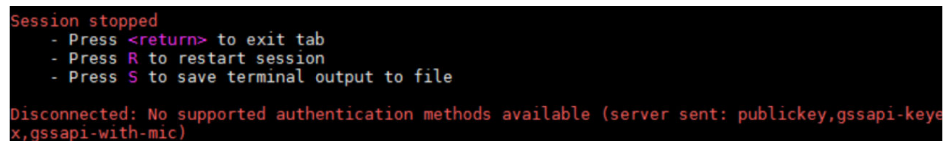
```
systemctl restart sshd
```

15.6.37 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 15-102 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`
2. Modify the following settings:
Change **PasswordAuthentication no** to **PasswordAuthentication yes**.
Alternatively, delete the comment tag (`#`) before **PasswordAuthentication yes**.
3. Restart SSH.
 - CentOS 6
`service sshd restart`
 - CentOS 7
`systemctl restart sshd`

15.7 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 15-6](#).

Table 15-6 Error codes and solution suggestions

Error Code	Message Displayed on the Management Console	Solution Suggestion
Ecs.0000	Request error. Try again later or contact customer service.	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 customer service and request an ECS quota increase.	Contact customer service and request an ECS quota increase. NOTE 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 or your balance is insufficient.	Contact customer service to check your account information.
Ecs.0005	System error. Try again later or contact customer service.	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.
Ecs.0011	Invalid password. 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.

Error Code	Message Displayed on the Management Console	Solution Suggestion
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 customer service and request an EIP quota increase.	Contact customer service 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. 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 customer service 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 customer service for troubleshooting.

15.8 How Can I Recover a Windows ECS with an Abnormal Virtualization Driver?

Background

An error occurs in the virtualization driver on a Windows ECS because of improper running of Tools. To ensure proper ECS running, handle this issue by following the instructions provided in this section.

Issue Description

The virtualization driver of an ECS became abnormal and this affected the data security, availability, and performance of the ECS.

The impact of this issue is as follows:

1. The file system may be damaged.
When you stop or restart such an ECS on the management console, the ECS will be forcibly stopped or restarted due to the lack of the virtualization driver Tools.
2. Services on the ECS may become unavailable
The affected ECS cannot be hot migrated between physical servers. If the host accommodating such an ECS becomes faulty or the hardware of the host is maintained, the ECS cannot be migrated to another host, affecting service high availability.
3. The network and storage performance of the ECS deteriorates.
The virtualization driver can improve the ECS network and storage performance. When the virtualization driver becomes abnormal, the network and storage performance will deteriorate.

Scenarios


- Scenario 1
Tools is not installed on the Windows ECS.
- Scenario 2
Tools has been uninstalled from the ECS.

Procedure

To install Tools on the Windows ECS, do as follows:


1. Obtain the Tools installation package **pvdriever-windows.zip**.
Contact the administrator to obtain the installation package.
2. Decompress the software package and double-click **setup.exe** to start the installation.

Figure 15-103 Installing the virtualization driver




3. Click **Next** and select "I accept the terms in the License Agreement".
4. Click **Install** to start the installation.

Figure 15-104 Installation progress




5. Click **Finish** to complete the installation.

Figure 15-105 Installation completed




6. Restart the ECS, as shown in [Figure 15-106](#).

Figure 15-106 Determining whether to restart the ECS



7. View the virtualization driver status in the bottom right corner of the ECS desktop. The yellow icon indicates that the virtualization driver is running properly.

Figure 15-107 Proper running status of the virtualization driver



15.9 Why Is My Windows ECS Muted?

Symptom

You cannot play audio files on a Windows ECS that is remotely accessed using MSTSC.

Constraints

This section applies only to ECSs running Windows Server 2008 R2 or Windows Server 2016.

Possible Causes


The audio function is disabled on Windows ECSs by default. As a result, audio files cannot be played on them. To enable the audio function, perform the operations described in this section.

Step 1: Enable Windows Audio

Enable Windows audio and set it to run automatically.

1. Start the **Run** dialog box.
2. Enter **services.msc** to access the service management console.
3. Find **Windows Audio** and set it as follows:
 - **Startup type: Automatic**
 - **Service status: Start**

The following figure uses Windows Server 2012 as an example.



4. Disable the remote connection.

Step 2: Enable Audio and Video Playback


The method of enabling audio and video playback varies depending on the ECS OS.

Windows Server 2008

Step 1 Enable RDP-TCP Audio and video playback and Audio recording.

1. Log in to the **Remote Desktop Session Host Configuration** management console.
 - a. Choose **Start > Control Panel**.
 - b. In the upper right corner of the page, choose **Category** for **View by**.
 - c. Choose **System and Security > Administrative Tools > Remote Desktop Services > Remote Desktop Session Host Configuration**.
2. Deselect **Audio and video playback** and **Audio recording**.
 In the **Connections** pane, double-click **RDP-Tcp**. In the **RDP-Tcp Properties** dialog box, click the **Client Settings** tab and deselect **Audio and video playback** and **Audio recording**.

Figure 15-108 Remote Desktop Session Host Configuration

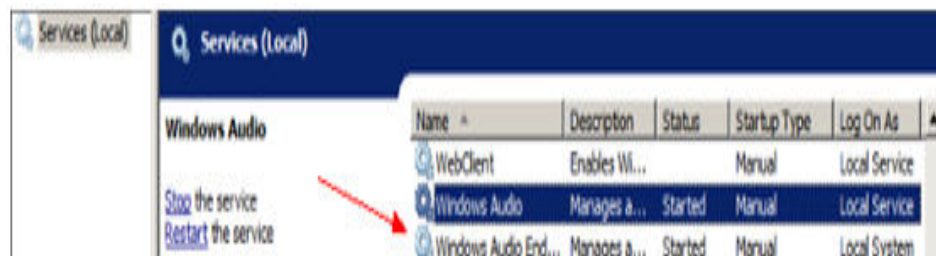


3. Click **OK** to enable the audio function.

Step 2 Click **Send CtrlAltDel** to restart the ECS and log in to it.

Step 3 Enable the audio service.

Figure 15-109 Enabling the audio service



Step 4 Play an audio file to verify the service.

----End

Windows Server 2012

Step 1 Start the **Run** dialog box.

Step 2 Run the `gpedit.msc` command to start **Local Group Policy Editor**.

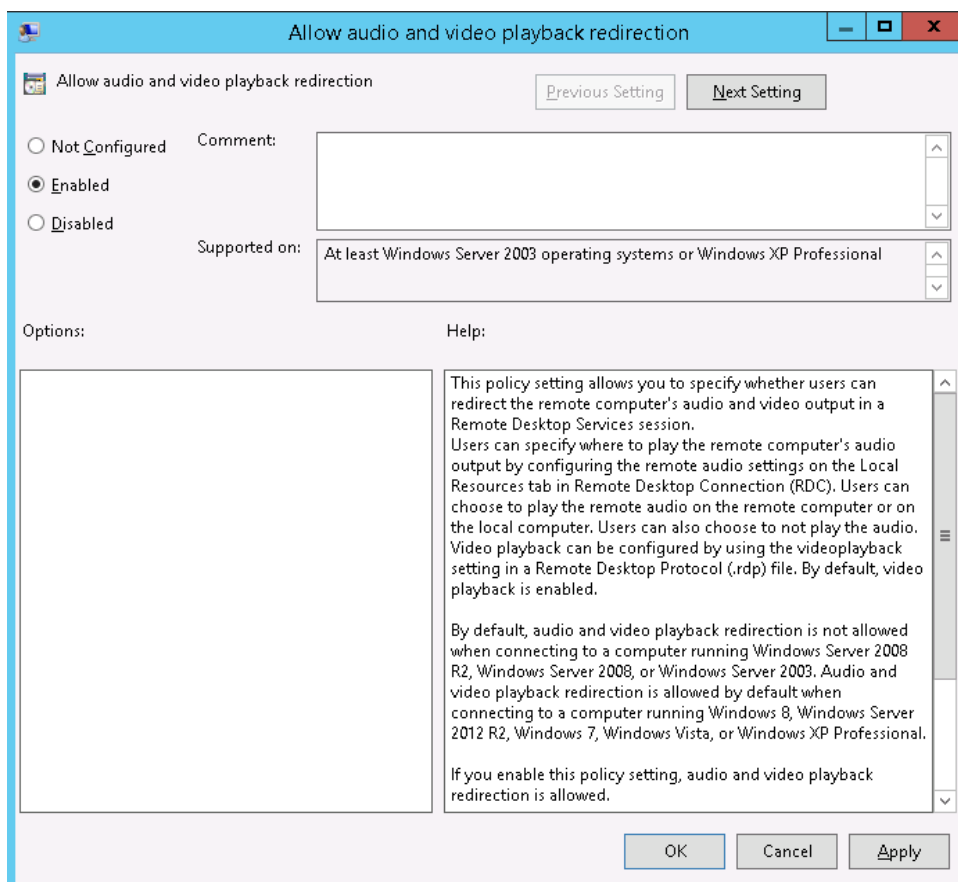
```
Microsoft Windows [6.3.9600]
(c) 2013 Microsoft Corporation.

C:\Users\Administrator>gpedit.msc

C:\Users\Administrator>
```

Step 3 Choose **Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Device and Resource Redirection**. Then, enable **Allow audio and video playback redirection**.

Step 4 Select **Enabled** and click **Apply**.



Retain the default settings of MSTSC.


Step 5 Run the following command to update the group policy:

```
gpupdate
```

```
----End
```

Step 3: Configure Remote Audio Settings

Start the local remote desktop software MSTSC, choose **Options > Local Resources**, and click **Settings** in **Remote audio**. Then, select **Play on this computer** in **Remote audio playback** and click **OK**.



Log in to the ECS using MSTSC and check whether audio files can be played properly.

15.10 How Do I Change an ECS SID?

Microsoft identifies computers and users by security identifier (SID). The ECSs created using an image have the same SID. If such ECSs are required to join in a Windows domain, they must use different SIDs.

This section describes how to use SIDCHG to change an ECS SID.

To change SIDs in a batch, use a private image and follow the operations provided in [Running Sysprep](#).

NOTE

Changing an ECS SID may lead to data loss or system damage, so back up ECS data before changing the SID.

Procedure

1. Click [SIDCHG](#) to download it.

NOTE

For the server edition, download the 64-bit version.

Figure 15-110 Downloading SIDCHG

SIDCHG 2.0o

[SIDCHG](#) and [SIDCHG64 \(64-bit Windows\)](#)

These are directly executables of SIDCHG SID Change Utility. There is no installation program.

It is important to not interrupt SID change in process. Additionally, on Windows 10, **Do not Log in into the computer during SID change!** Logging in will affect Start Menu and modern Windows interfaces and apps.

2. Run the following command to change the ECS SID:

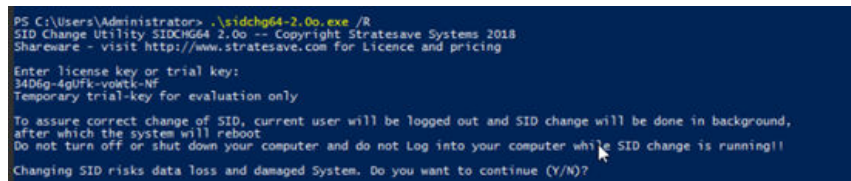
sidchg64-2.0n.exe /R

NOTE

In the preceding command, **/R** indicates that the ECS will automatically restart after its SID is changed, and **/S** indicates that the ECS will not automatically restart.

3. Enter the trial key or license and press **Enter**.
[Obtain the latest trail key and learn how to use SIDCHG.](#)
4. When the system displays a message asking you whether to continue, press **y**.

Figure 15-111 Risk prompt




```
PS C:\Users\Administrator> .\sidchg64-2.0n.exe /R
SID Change Utility SIDCHG64 2.0n -- Copyright Stratesave Systems 2018
Shareware - visit http://www.stratesave.com for Licence and pricing

Enter license key or trial key:
34D6g-4gUfk-v0etk-NF
Temporary trial-key for evaluation only

To assure correct change of SID, current user will be logged out and SID change will be done in background,
after which the system will reboot
Do not turn off or shut down your computer and do not Log into your computer while SID change is running!!
Changing SID risks data loss and damaged System. Do you want to continue (Y/N)?
```

5. Log in to the ECS again.

Figure 15-112 Re-login



6. After the ECS is restarted, run the **cmd** command to open the CLI and run **whoami /user** to verify that the SID has been changed.

15.11 Why Does a Pay-per-Use ECS Fail to Be Started?

After a pay-per-use ECS is stopped, its resources such as vCPUs and memory are released. When it is being restarted, the startup may fail due to insufficient resources.

In this case, you can try to start it again or modify the ECS specifications by referring to [General Operations](#).

15.12 ECS Management

15.12.1 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
```

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
```

- 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. 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.

15.12.2 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 15-7](#) lists the hostnames (obtained by running the **hostname** command) of ECSs created using different images and those displayed after the ECSs are restarted.

Table 15-7 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?](#)

15.12.3 Why Is the Hostname of My ECS Restored to the Original Name After the ECS Is Restarted?

The following uses an ECS running CentOS 7 as an example:

1. Log in to the Linux ECS and view the Cloud-Init configuration file.
2. In the `/etc/cloud/cloud.cfg` file, comment out or delete `update_hostname`.

NOTE

- `update_hostname` indicates that the hostname is changed in Cloud-Init each time the ECS is restarted.
- For an ECS created from a public image, Cloud-Init has been installed on it by default. You do not need to manually install Cloud-Init for it. For details about how to modify a private image, see [Installing Cloud-Init](#).

15.12.4 How Can I Set Sequential ECS Names When Creating Multiple ECSs?

Scenarios

When creating multiple ECSs at the same time, you can use either of the following methods to sequentially name the ECSs:

- Automatic naming: The system automatically adds a hyphen followed by a four-digit incremental number to the end of each ECS name.
- Customizable naming: You can customize a naming rule in the format "name_prefix[begin_number,bits]name_suffix". The system will automatically name the ECSs according to naming rule you specify.

This section describes how to use the two methods to name ECSs.

Automatic Naming

You can customize the name according to the following naming rules: The name must contain 1 to 64 characters that can be only letters, digits, underscores (`_`), and hyphens (`-`).

When you create multiple ECSs at the same time, the system automatically adds a hyphen followed by a four-digit incremental number to the end of each ECS name. In this case, the customized name is 1 to 59 characters long. For example, if you are creating multiple ECSs and enter `ecs` for the ECS name, the created ECSs will be named `ecs-0001`, `ecs-0002`, and so on. 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.

- Example 1: If there is no existing ECS and you enter `ecs-f526`, the ECSs will be named `ecs-f526-0001`, `ecs-f526-0002`, `ecs-f526-0003`,
- Example 2: If there is an ECS named `ecs-f526-0010` and you enter `ecs-f526`, the ECSs will be named `ecs-f526-0011`, `ecs-f526-0012`, `ecs-f526-0013`,
- Example 3: If there is an ECS named `ecs-0010` and you select **Allow duplicate ECS name**, all the ECSs will be named `ecs-0010`.

Customizable Naming

You can customize a naming rule in the format "name_prefix[begin_number,bits]name_suffix". The system will automatically name the ECSs according to naming rule you specify.

Field Description for a Customizable Naming Rule

Figure 15-113 shows the format of a customizable naming rule.

Figure 15-113 Format of a customizable naming rule




Table 15-8 describes these parameters.

Table 15-8 Parameters in a customizable naming rule

Field	Mandatory	Description	Example
name_prefix	Yes	ECS name prefix The name prefix can contain only letters, digits, underscores (_), and hyphens (-).	ecs
[begin_number,bits]	Yes	Sequence numbers that increase in ascending order to differentiate multiple ECSs.	[0,4]
name_suffix	No	ECS name suffix The name suffix can contain only letters, digits, underscores (_), and hyphens (-).	f526

Table 15-9 [begin_number,bits] parameters

Field	Mandatory	Description	Example
begin_number	No	Begin number of ECS names. The begin number ranges from 0 to 9999. The default value is 0.	0

Field	Mandatory	Description	Example
bits	No	Number of bits for the sequential numbers in ECS names. The value ranges from 1 to 4. The default value is 4.	4

Notes on Using Customizable Naming

- Customized names cannot be duplicate.
- No space is allowed in [begin_number,bits].
- If the bits of "Begin number + Number of ECSs to be created - 1" is greater than the specified bits, the bits of "Begin number + Number of ECSs to be created - 1" will be used.

For example, if [begin_number,bits] is set to [8,1] and the number of ECSs to be created is 2, the bits of "Begin number + Number of ECSs to be created - 1" is the same as the specified bits (1). Then, the ECSs will be named *name_prefix8name_suffix* and *name_prefix9name_suffix*.

If [begin_number,bits] is set to [8,1] and the number of ECSs to be created is 3, the specified bits is 1, the bits of "Begin number + Number of ECSs to be created - 1" (value 10, bits 2) is different from the specified bits (1). Therefore, the bits of "Begin number + Number of ECSs to be created - 1" will be used, which is 2.

The ECSs will be named *name_prefix08name_suffix*, *name_prefix09name_suffix*, and *name_prefix10name_suffix*.

- If the value of "Begin number + Number of ECSs to be created" is greater than the maximum value **9999**, the sequential numbers that exceed **9999** will consistently to be **9999**.
- If [begin_number,bits] is set to [] or [,], the begin number starts from **0**, and the number of bits is **4** by default.
- If [begin_number,bits] is set to [99] or [99,], the begin number starts from **99**, and the number of bits is **4** by default.

Customizable Naming Examples

- Example 1: If you select customizable naming and enter *name_prefix[,]name_suffix*,
The ECSs will be named *name_prefix0000name_suffix*, *name_prefix0001name_suffix*, *name_prefix0002name_suffix*,
- Example 2: If you select customizable naming and enter *name_prefix[]name_suffix*,
The ECSs will be named *name_prefix0000name_suffix*, *name_prefix0001name_suffix*, *name_prefix0002name_suffix*,
- Example 3: If you select customizable naming and enter *name_prefix[9,]name_suffix*,
The ECSs will be named *name_prefix0009name_suffix*, *name_prefix0010name_suffix*, *name_prefix0011name_suffix*,

- Example 4: If you select customizable naming and enter *name_prefix[3]name_suffix*,
The ECSs will be named *name_prefix000name_suffix*,
name_prefix001name_suffix, *name_prefix002name_suffix*,
- Example 5: If you select customizable naming and enter *name_prefix[8]name_suffix*,
The ECSs will be named *name_prefix0008name_suffix*,
name_prefix0009name_suffix, *name_prefix0010name_suffix*,
- Example 6: If you select customizable naming and enter *name_prefix[9999]name_suffix*,
All the ECSs will be named *name_prefix9999name_suffix*.
- Example 7: If you select customizable naming and enter *name_prefix[8]*,
The ECSs will be named *name_prefix0008*, *name_prefix0009*,
name_prefix0010,

15.12.5 How Can I Modify ECS Specifications?

If the specifications of an existing ECS cannot meet service requirements, modify the ECS specifications as needed, for example, upgrading the vCPUs and memory.

To do so, switch to the list view on the **Elastic Cloud Server** page, locate the row containing the target ECS and choose **More > Modify Specifications** in the **Operation** column.

Specification modifications include specification upgrade and downgrade.

- For pay-per-use ECSs, the specifications upgrade and downgrade take effect immediately. You are billed based on the new specifications.

15.12.6 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 15-114 Disks before modifying ECS specifications

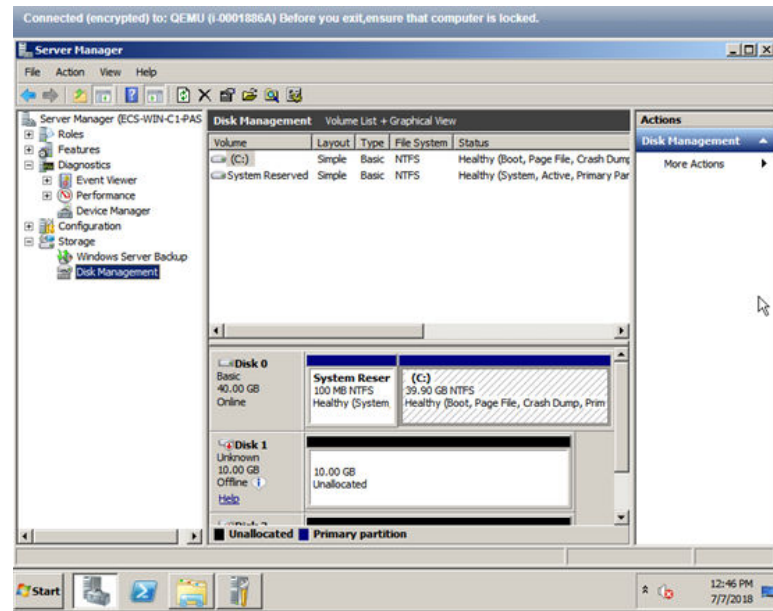
After the specifications are modified, check the number of disks.

Figure 15-115 Disks after modifying ECS specifications

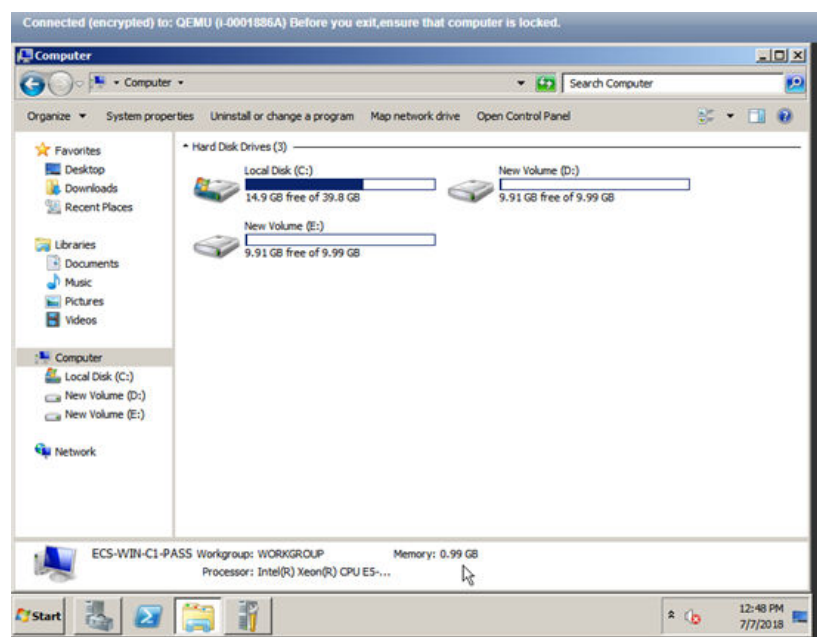
Only one system disk is displayed. The data disks are offline after you modify the specifications.

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 15-116 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 customer service.

Figure 15-117 Disks after you bring the disks online

15.12.7 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 15-118 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 15-118](#), 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 15-119 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 15-119](#), 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.

6. 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 customer service.

Figure 15-120 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 ~]#
[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 15-120](#), the numbers of disks before and after specifications modifications are the same. The disks are `/dev/vda`, `/dev/vdb`, and `/dev/vdc`.

15.13 OS Management

15.13.1 Does OS Change Incur Fees?

After the OS is changed, different images are used and system disk capacity may increase. You will be billed based on the new configurations.

For details about OS change, see [Changing the OS](#).

15.13.2 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. For details about image selection, see [Should I Choose Windows OS or Linux OS for My ECS?](#)
- You can change your ECS OS through the management console, for example, you can upgrade CentOS 7.2 to CentOS 7.3.

15.13.3 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](#).

15.13.4 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/Backup > Change OS** in the **Operation** column.

During this process, the ECS is in **Changing OS** state.

15.13.5 Will I Lose My Disk Data If I Reinstall ECS OS, Change the OS, or Change the ECS Specifications?

Table 15-10 Impact

Item	OS Reinstallation	OS Change	Specifications Modification
Application scenario	Initialize an ECS. The ECS OS remains unchanged after OS change.	Change the OS of an ECS by changing its image. For details about OS change constraints, see Changing the OS .	Change ECS specifications, such as increasing the number of vCPUs or adding memory, to meet your service requirements.
Billing	OS reinstallation is free of charge. The ECS price remains unchanged.	OS change is free of charge. However, you will be billed based on your new image type after OS change.	Modifying ECS specifications is free of charge. However, you will be billed based on the new specifications after modification.
IP address	The private IP address, EIP, and MAC address remain unchanged.	The private IP address, EIP, and MAC address remain unchanged.	The private IP address, EIP, and MAC address remain unchanged.
System disk	Reinstalling OS will clear the data in all partitions of the ECS system disk. Back up data before reinstalling the OS.	Changing OS will clear the data in all partitions of the ECS system disk. Back up data before changing the OS.	No impact on system disk.
Data disk	No impact on data disk.	No impact on data disk	No impact on data disk.

Item	OS Reinstallation	OS Change	Specifications Modification
Backup	Back up data before reinstalling the OS to prevent data loss.	Back up data before changing the OS to prevent data loss.	Create a system disk snapshot before modifying ECS specifications to prevent data loss.

15.13.6 Does OS Reinstallation Incur Fees?

Reinstalling an OS for an ECS allows you to use the original image to reinstall the ECS and does not incur fees.

15.13.7 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](#).

15.13.8 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.

Figure 15-121 Reinstall OS

Reinstall OS
×

Note the following points before you reinstall the OS:

1. This operation will reinstall the original OS image on the ECS and delete data on all system partitions. [Back up](#) the data before you continue.
2. There are no impacts to data disks when you reinstall an OS.
3. The ECS will automatically restart after the OS has been reinstalled.
4. Custom settings such as DNS and hostname need to be configured again after the OS has been reinstalled.

Image CentOS 7.5 64bit

System Disk Capacity: 40 GiB OS Architecture: 64-bit

Encrypted ?

Login Mode
Password
Key pair

Password Keep your password secure. The system cannot detect your password.

Confirm Password Keep your password secure. The system cannot detect your password.

OK
Cancel

15.13.9 Do ECSs Support GUI?

Windows ECSs are managed through a GUI but Linux ECSs are managed through the CLI. You can configure a GUI if required.

Before installing a GUI on an ECS, ensure that the idle memory is greater than or equal to 2 GiB. Otherwise, the GUI installation may fail or the ECS cannot be started after the installation.

15.13.10 How Can I Install a GUI on an ECS Running CentOS 6?

Scenarios

To provide a pure system, the ECSs running CentOS 6 do not have a GUI installed by default. You can install a GUI on such ECSs as needed.

Constraints

- Before installing a GUI on an ECS, ensure that the idle memory is greater than or equal to 2 GB. Otherwise, the GUI installation may fail or the ECS cannot be started after the installation.

Procedure

1. Run the following command to obtain the installation component provided by the OS:
yum groupinstall "Desktop"
2. Run the following command to set the default startup level to 5 (GUI):
sed -i 's/id:3:initdefault:/id:5:initdefault:/' /etc/inittab
3. Run the following command:
startx

15.13.11 How Can I Install a GUI on an ECS Running CentOS 7?

Scenarios

You want to install a GUI on an ECS running CentOS 7 series.

Constraints

- Before installing a GUI on an ECS, ensure that the idle memory is greater than or equal to 2 GB. Otherwise, the GUI installation may fail or the ECS cannot be started after the installation.

Procedure

1. Run the following command to install the GUI desktop component:
yum groupinstall "Server with GUI"

 NOTE

If the following message is displayed after the installation is complete:

```
Failed : python -urllib3.noarch 0:1.10.2-7.e17
```

Run the following command:

```
mv /usr/lib/python2.7/site-packages/urllib3/packages/  
ssl_match_hostname /usr/lib/python2.7/site-packages/urllib3/packages/  
ssl_match_hostname.bak
```

```
yum install python-urllib3 -y
```

2. After the installation is complete, run the following command to set the default startup level to **graphical.target**:
systemctl set-default graphical.target
3. Run the following command to start **graphical.target**:
systemctl start graphical.target
4. Restart the ECS.
5. Log in to the ECS using VNC provided on the management console. Set the language, time zone, username, and password as prompted.

15.13.12 How Can I Install a GUI on an ECS Running Ubuntu?

Scenarios

To provide a pure system, the ECSs running Ubuntu do not have a GUI installed by default. You can install a GUI on such ECSs as needed.

For GPU-accelerated ECSs, after installing a GUI, you need to configure X Server, x11vnc, and lightdm to make sure that:

- The graphics system and VNC server are automatically started upon the ECS startup.
- Applications can invoke GPUs properly after a remote login using VNC.

You can perform the following steps to install a GUI on an Ubuntu ECS:

- **Installing a GUI**
- **(Optional) Configuring X Server, x11vnc, and lightdm**: required only for GPU-accelerated ECSs.
- **(Optional) Verifying Drivers on GPU-accelerated ECSs**: required only for GPU-accelerated ECSs.

Constraints

- This document applies to ECSs running Ubuntu 16.04, 18.04, and 20.04.
- The Ubuntu ECS must have an EIP bound or have an intranet image source configured.
- Before installing a GUI on an ECS, ensure that the idle memory is greater than or equal to 2 GB. Otherwise, the GUI installation may fail or the ECS cannot be started after the installation.
- GPU-accelerated ECSs must have a correct GPU driver installed. For details, see [GPU Driver](#).

Installing a GUI

1. Log in to the ECS and install a GUI desktop environment.
 - a. Run the following command to update the software library:
apt-get update
 - b. Run the following command to install the Ubuntu GUI desktop component:
 - For Ubuntu 16.04, run the following command:
apt-get install -y scite xorg xubuntu-desktop
 - For Ubuntu 18.04 and 20.04, run the following command:
apt-get install -y ubuntu-desktop

2. Run the following command to edit the **root/.profile** file:

```
vim /root/.profile
```

Press **i** to enter the editing mode and change **mesg n || true** in the last line to **tty -s && mesg n || true**. After the modification, the file content is as follows:

```
# ~/.profile: executed by Bourne-compatible login shells.

if [ "$BASH" ]; then
  if [ -f ~/.bashrc ]; then
    . ~/.bashrc
  fi
fi
tty -s && mesg n || true
```

3. Press **Esc** to exit editing mode.
4. Run the following command to save and exit the configuration file:
:wq
5. (Mandatory for Ubuntu 20.04) Add a member account.

After GUI desktop component is installed on the ECS, you cannot log in to the Ubuntu 20.04 OS as user **root** **user**. You need to add a member account for logging in to the GUI desktop.

Run the following command to add user **user01**:

```
adduser user01
```

Set a password for **user01** as prompted.

```
Adding user `user01' ...
Adding new group `user01' (1001) ...
Adding new user `user01' (1001) with group `user01' ...
Creating home directory `/home/user01' ...
Copying files from `/etc/skel' ...
New password:
Retype new password:
passwd: password updated successfully
```

Set information about **user01**. You can press **Enter** to skip the setting. Then the system prompts you to check whether the entered information is correct.

Enter **Y**.

```
Changing the user information for user01
Enter the new value, or press ENTER for the default
Full Name []:
Room Number []:
Work Phone []:
Home Phone []:
```



```
Other []:  
Is the information correct? [Y/n] Y
```

6. Run the reboot command to restart the ECS.
7. Log in to the ECS using VNC provided on the management console and log in to the GUI desktop using the member account created in [5](#) or the **root** account.
 - For Ubuntu 20.04 OS, you need to use the member account to log in to the GUI desktop.
 - For GPU-accelerated ECSs, you also need to [configure X Server, x11vnc, and lighthdm](#).

(Optional) Configuring X Server, x11vnc, and lighthdm

For GPU-accelerated ECSs, you need to configure X Server, x11vnc, and lighthdm when installing a GUI.

1. Remotely log in to the ECS.
2. Query the BusID of the GPU.

```
lspci | grep -i nvidia
```

Figure 15-122 GPU's BusID

```
00:0d.0 3D controller: NVIDIA Corporation GV100GL [Tesla V100 PCIe 32GB] (rev a1)
```

3. Generate the X Server configuration.
nvidia-xconfig --enable-all-gpus --separate-x-screens
4. Configure the GPU's BusID in "Section Device" in the generated `/etc/X11/xorg.conf`.
 - a. Edit `/etc/X11/xorg.conf`.
vi /etc/X11/xorg.conf
 - b. Press **i** to enter editing mode.
 - c. Add the GPU's BusID in "Section "Device"".

Figure 15-123 Adding the GPU's BusID

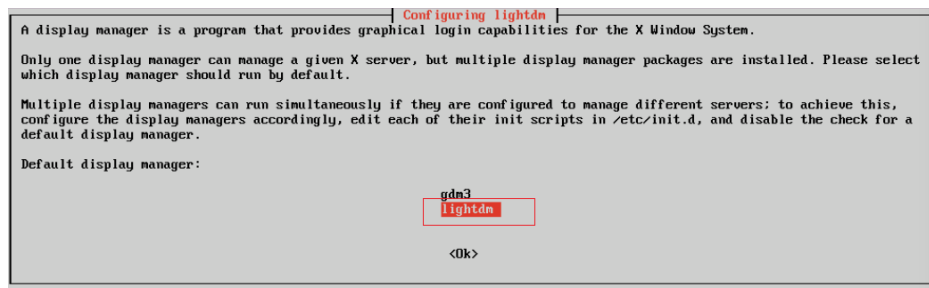
```
Section "Device"  
    Identifier      "Device0"  
    Driver          "nvidia"  
    VendorName     "NVIDIA Corporation"  
    BoardName      "Tesla V100-PCIe-32GB"  
    BusID          "PCI:00:13:0"  
EndSection
```

NOTE

The BusID queried in step [2](#) is a hexadecimal number. You need to convert it to a decimal number before adding it to "Section Device" in `/etc/X11/xorg.conf`.

1. For example, the queried BusID is `00.0d.0` (a hexadecimal number) and needs to be converted to `PCI:00:13:0` (a decimal number).
- d. Press **Esc** to exit editing mode.

- e. Run the following command to save and exit the configuration file:
:wq
5. Install x11vnc.
apt-get -y install x11vnc
6. Install lighthdm.
apt-get -y install lightdm
7. Select **lightdm** as the default display manager.

Figure 15-124 Selecting a display manager

8. Configure the GUI desktop environment to automatically start upon ECS startup.
systemctl set-default graphical.target
9. (Optional) Configure the x11vnc to automatically start upon ECS startup.
 - a. Add the **/lib/systemd/system/myservice.service** file.
vi /lib/systemd/system/myservice.service
 - b. Press **i** to enter editing mode.
 - c. Add the following content to the file:

```
[Unit]
Description=My Service
After=network.target lightdm.service

[Service]
Type=oneshot
ExecStart=/usr/bin/x11vnc -forever -loop -noxdamage -repeat -rfbport 5902 -shared -bg -auth guess -o /var/log/vnc.log

[Install]
WantedBy=multi-user.target
Alias=myservice.service
```
 - d. Press **Esc** to exit editing mode.
 - e. Run the following command to save and exit the configuration file:
:wq
10. Load configuration files.
systemctl daemon-reload
systemctl enable myservice.service
11. Run the reboot command to restart the ECS.


(Optional) Verifying Drivers on GPU-accelerated ECSs

After installing a GUI on a GPU-accelerated ECS, perform the following operations to check whether the driver is working properly:

1. Log in to the management console.
2. Configure a security group for the ECS.
 - a. On the ECS list, click the name of an ECS for which you want to configure the security group rule. On the ECS details page, click **Security Groups**.
 - b. Expand the security group and in the upper right corner of the security group rule list, click **Modify Security Group Rule**.
 - c. On the **Inbound Rules** page, click **Add Rule**.
 - d. In the **Add Inbound Rule** dialog box, follow the prompts to add the following security group rule:

Allow inbound access through TCP port *5902*. The port number is determined by the **rfbport** parameter in step [9.c](#).
3. Log in to the ECS using VNC.

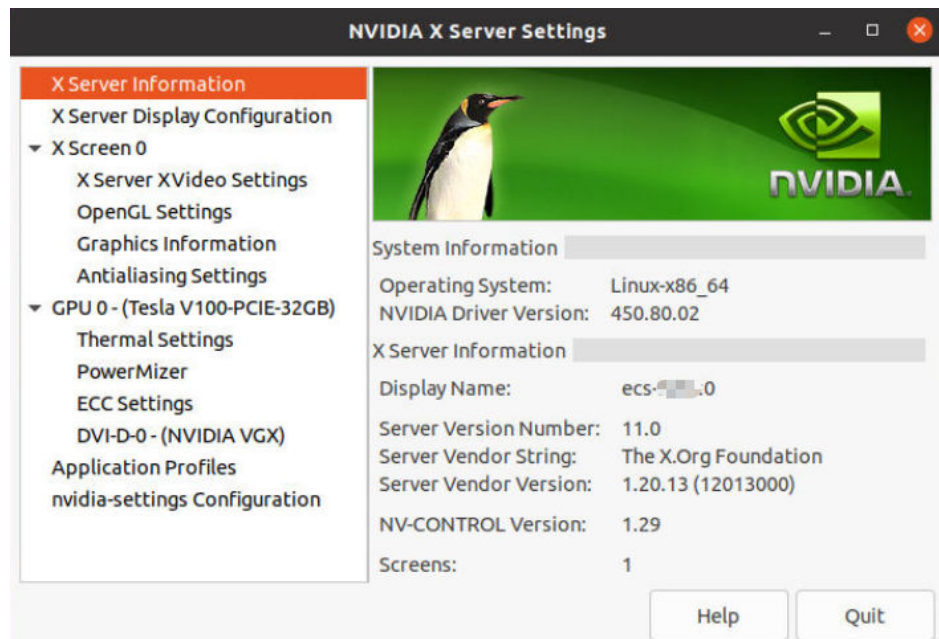
The following uses TightVNC as an example.

Figure 15-125 TightVNC client

4. Right-click on the blank area and choose **Open in Terminal** from the shortcut menu.
5. Run the following command on the terminal. If the graphics card information is displayed as follows, the driver is working properly.

nvidia-settings

Figure 15-126 Graphics card information

**NOTE**

If a GPU-accelerated ECS has a GRID driver installed, you need to configure a license to use the GPU rendering capability. For details, see [Installing a GRID Driver on a GPU-accelerated ECS](#).

15.13.13 How Can I Install a GUI on an ECS Running Debian?

Scenarios

To provide a pure system, the ECSs running Debian do not have a GUI installed by default. You can install a GUI on such ECSs as needed.

Constraints

- The operations described in this section apply to ECSs running Debian 8, Debian 9, or Debian 10 only.
- Before installing a GUI on an ECS, ensure that the memory is no less than 2 GB to prevent GUI installation or ECS startup failures.

Procedure

1. Log in to the ECS and run the following command to update the software library:
apt update
2. Run the following command to upgrade the software library:
apt upgrade
3. Run the following command to install taskset:
apt install taskset

4. Run the following command to use tasksel to install the GNOME GUI:
tasksel install desktop gnome-desktop
The installation takes a long time. Please wait.
5. Run the following command to set the GUI as the default startup target:
systemctl set-default graphical.target
6. Create a member account.
After GUI desktop component is installed on the ECS, you cannot log in to the Debian OS as user root **user**. Therefore, you need to add a member account for logging in to the GUI desktop.
Run the following command to add user **user01**:
adduser user01
Set a password for **user01** as prompted.

```
Adding user `user01' ...
Adding new group `user01' (1001) ...
Adding new user `user01' (1001) with group `user01' ...
Creating home directory `/home/user01' ...
Copying files from `/etc/skel' ...
New password:
Retype new password:
passwd: password updated successfully
```

Set information about **user01**. You can press **Enter** to skip the setting. Then the system prompts you to check whether the entered information is correct.
Enter **Y**.

```
Changing the user information for user01
Enter the new value, or press ENTER for the default
  Full Name []:
  Room Number []:
  Work Phone []:
  Home Phone []:
  Other []:
Is the information correct? [Y/n] Y
```
7. Run the reboot command to restart the ECS.
8. Log in to the ECS using VNC provided on the management console and log in to the GUI desktop using the member account added in [6](#).

15.13.14 Why Does the OS Fail to Respond When kdump Occurs on a Linux ECS?

Symptom

When kdump occurs on a Xen Linux ECS, the OS fails to respond and cannot be automatically recovered. For example, if you run the **echo c>/proc/sysrq-trigger** command to trigger kdump, this fault occurs.

Figure 15-127 Triggering kdump

```
root@ecs-xen01 linux1# systemctl status kdump
■ kdump.service - Crash recovery kernel arming
   Loaded: loaded (/usr/lib/systemd/system/kdump.service; enabled; vendor preset: enabled)
   Active: active (exited) since Wed 2018-01-17 06:15:35 UTC; 6min ago
   Process: 1397 ExecStart=/usr/bin/kdumpctl start (code=exited, status=0/SUCCESS)
   Main PID: 1397 (code=exited, status=0/SUCCESS)
   CGroup: /system.slice/kdump.service

Jan 17 06:15:05 ecs-xen01.novalocal systemd[1]: Starting Crash recovery kernel arming...
Jan 17 06:15:35 ecs-xen01.novalocal kdumpctl[1397]: kexec: loaded kdump kernel
Jan 17 06:15:35 ecs-xen01.novalocal kdumpctl[1397]: Starting kdump: [OK]
Jan 17 06:15:35 ecs-xen01.novalocal systemd[1]: Started Crash recovery kernel arming.
root@ecs-xen01 linux1# echo c > /proc/sysrq-trigger
```

NOTE

Generally, kdump is disabled for public images. This issue does not occur on the ECSs created using public images.

Possible Causes

- Certain Linux kernel versions are incompatible with Xen virtualization.
- If kdump is enabled in the ECS with the kernel not supporting `soft_rest`, the ECS stops responding during dump.

Solution

Method 1: Disable kdump.

CentOS 7.5 is used as an example in the following.

1. Forcibly restart the ECS.
 - a. Log in to management console.
 - b. Under **Computing**, choose **Elastic Cloud Server**.
 - c. In the ECS list, select the target ECS and click **Restart**.
 - d. Select **Forcibly restart the preceding ECSs** or **Forcibly stop the preceding ECSs**.
 - e. Click **OK**.
2. Disable kdump.
 - a. Log in to the forcibly restarted ECS as user **root**.
 - b. Run the following command to disable kdump:
service kdump stop

Method 2:

If the target ECS supports the `crash_kexec_post_notifiers` function, add the function to the ECS startup configuration file (`menu.lst` or `grub.cfg`). To do so, perform the following operations:

1. Run the following command to check whether the ECS supports the `crash_kexec_post_notifiers` function:
cat /proc/kallsyms |grep crash_kexec_post_notifiers

Figure 15-128 Support for the `crash_kexec_post_notifiers` function

```
^Clinux-EVdrQm:~ # cat /proc/kallsyms |grep crash_kexec_post_notifiers
ffffffff816c3a20 r __param_str_crash_kexec_post_notifiers
ffffffff819c3da8 r __param_crash_kexec_post_notifiers
ffffffff81d58ec4 B crash_kexec_post_notifiers
```

- If yes, go to step 2.
 - If no, use method 1.
2. Add the `crash_kexec_post_notifiers` function to the startup configuration file (`menu.lst` or `grub.cfg`).

Take `menu.lst` as an example.

- a. Run the following command to open the `menu.lst` file:
`vi /boot/grub/menu.lst`
- b. Add the `crash_kexec_post_notifiers` function to the startup item.

Figure 15-129 Editing the `menu.lst` file

```
Modified by YaST2. Last modification on Thu Feb 22 10:51:10 UTC 2018
default 2
timeout 5
password --encrypted $6$XxIhQxs0E6KxQF8$hb7SVqVz3DFxV6q7LSUmzp0Fw4RTX16Ce3Y.FpbIdOfsitbSC0v7F.L.m$waroAFLeAanR1OtsqHluYQM/dh7/

##Don't change this comment - YaST2 identifier: Original name: linux##
title UVP Linux Enterprise Server V200R003C00 - 3.0.93-0.8
root (hd0,0)
kernel /vmlinuz-3.0.93-0.8-default root=/dev/disk/by-id/scsi-35000c5001ce8b6a7-part5 resume=/dev/sda1 splash=silent showopts
initrd /initrd-3.0.93-0.8-default

##Don't change this comment - YaST2 identifier: Original name: failafe##
title Failafe -- UVP Linux Enterprise Server V200R003C00 - 3.0.93-0.8
root (hd0,0)
kernel /vmlinuz-3.0.93-0.8-default root=/dev/disk/by-id/scsi-35000c5001ce8b6a7-part5
initrd /initrd-3.0.93-0.8-default

title UVP Linux Enterprise Server V200R003C00
root (hd0,0)
kernel /boot/xen.gz dom0_mem=8192M mem_for_icache=4096M balloon_zone=32768M dom0_max_vcpus=4 dom0_reserve_vcpus=4 numa=on console=
ed guest=0 x2apic=1 crashkernel=192M@16M Watchdog=1 shm dev num=0 Evm client2server size=128 shm server2client size=64 extra_guest_in
s_ic_enable=0 gnttab_max_nr_frames=3072 ple_gap=128 ple_window=4096 ecRed credit_default_yield=0 apic=1 crash_kexec_post_notifiers
module /boot/vmlinuz-3.0.93-0.8-xen console=tty0 console=ttyS0,115200 root=/dev/disk/by-id/scsi-35000c5001ce8b6a7-part5 vga=0x317
module /boot/initrd-3.0.93-0.8-xe
```

- c. Run the following command to restart the ECS for the modification to take effect:
`reboot`

15.13.15 How Can I Upgrade the Kernel of a Linux ECS?

Upgrade Notes

If tools have been installed on the Linux ECS, you must uninstall the tools before upgrading the ECS kernel. Otherwise, the following issues may occur after the kernel is upgraded:

- The Linux ECS cannot identify the NIC, leading to network access failure.
- The Linux ECS cannot identify data disks. As a result, starting system mount points fails, and the ECS cannot start.

Background

PVOPS is the Xen driver delivered with Linux distributions.

Procedure

1. Log in to the ECS.
2. Check whether the Tools have been installed on the Linux ECS, taking the SUSE Linux Enterprise Server 11 SP1 as an example.

- a. Run the following command on any directory to view the ECS driver:

```
lsmod | grep xen
```

Figure 15-130 Viewing the ECS driver

```
Linux:~/Desktop # lsmod | grep xen
xen_vbd                23600  3
cdrom                  40567  2 sr_mod,xen_vbd
xen_vmdq               4295   0
xen_vnif               36374   0
xen_balloon            14925   1 xen_vnif
xen_hcall              1867   0
xen_platform_pci      94554   5 xen_vbd,xen_vmdq,xen_vnif,xen_balloon,xen_hcall,[permanent]
```

- b. Run the following command to view the driver path, taking a disk driver as an example:

```
modinfo xen_vbd
```

Figure 15-131 Viewing the driver path

```
Linux:~/Desktop # modinfo xen_vbd
filename:      /lib/modules/2.6.32.12-0.7-default/updates/pvdriver/xen-vbd/xen-vbd.ko
license:      Dual BSD/GPL
alias:         xen:vbd
srcversion:   5D8B666F0EA3F1E31B58F0C
depends:       xen-platform-pci,cdrom
vermagic:     2.6.32.12-0.7-default SMP mod_unload modversions
```

- c. Check whether **pvdriver** is contained in the driver path.
 - If so, the tools have been installed in the ECS. Then, go to step **3**.
 - If no, go to step **4**.

3. Uninstall the tools.

- a. Run the following command to switch to user **root**:

```
su root
```

- b. Run the following command to uninstall Tools in the root directory:

```
/etc/uvp-monitor/uninstall
```

NOTE

After Tools is uninstalled, ECS monitoring metrics may be lost and monitoring data cannot be collected. To resolve this issue, you can compile and install the UVP Tools. For details, see <https://github.com/UVP-Tools/UVP-Tools/>.

4. Upgrade the kernel using the method determined by yourself.
5. Check whether the Linux ECS driver supports PVOPS. Use any one of the following methods:
 - Method 1:
Determine based on the ECS OS.
 - All Linux distribution OSs are delivered with a Xen open-source driver, which supports PVOPS.

- The SUSE Linux Enterprise Server 11 SP3 provided by the OS competence center is not delivered with any Xen open-source driver and does not support PVOPS.
- Method 2:
Check whether the ECS driver has a Xen driver module. If so, the ECS driver supports PVOPS. To obtain the data, run the following command in any directory:

lsmod | grep xen

Figure 15-132 Viewing the ECS driver

```
[root@localhost ~]# lsmod | grep xen
xen_vnif          59585  0 [permanent]
xen_vbd           50857  0
xen_balloon      45641  1 xen_vnif,[permanent]
xen_platform_pci 118125  3 xen_vnif,xen_vbd,xen_balloon,[permanent]
```

NOTE

The name of a Xen driver module varies depending on the Linux distribution OS. You only need to check whether the driver has a driver module with the **XEN** field.

- Method 3:
Run the **cat /boot/config* | grep -i xen** command in any directory and check whether the **XEN** field is contained in the command output. If so, the ECS driver supports PVOPS.

Figure 15-133 Viewing the **XEN** field

```
root@ubuntu:/home# cat /boot/config* | grep -i xen
CONFIG_XEN=y
CONFIG_XEN_DOM0=y
CONFIG_XEN_PVHVM=y
CONFIG_XEN_MAX_DOMAIN_MEMORY=500
CONFIG_XEN_SAVE_RESTORE=y
# CONFIG_XEN_DEBUG_FS is not set
CONFIG_XEN_PVH=y
CONFIG_PCI_XEN=y
```

6. Upgrade the kernel based on the result obtained in step 5.
 - If the Linux ECS driver supports PVOPS, go to step 8.
 - If the Linux ECS driver does not support PVOPS, go to step 7.
7. Install the open-source component xen-kmp so that the ECS driver supports PVOPS. For instructions about how to use PVOPS, see "Optimizing a Linux Private Image" in *Image Management Service User Guide*.
8. (Optional) Configure required parameters based on the defect list for certain Linux distribution OSs.


To obtain the defect list, go to following URL:

<https://github.com/UVP-Tools/UVP-Tools/tree/master/docs>

15.13.16 Why Cannot My ECS OS Start Properly?

1. Check the image based on which the ECS was created. If the image is a public one, this issue is not caused by private image sources.

Figure 15-134 Image type



2. Click **Apply for Server** and check whether the same ECS can be created. If not, this image may have been canceled.
3. Change the ECS OS to one that is available on the management console.

15.13.17 How Can I Enable SELinux on an ECS Running CentOS?

Symptom

SELinux is disabled on ECSs running CentOS 7.5 by default. After I enable SELinux by running `/etc/selinux/config` and enter the login password, the login failed.

This section describes how to resolve this issue based on enabled SELinux.

Solution

The operations described in this section are performed on ECSs running CentOS 7.5.

1. Run the following command to change **SELINUX=disabled** in the SELinux configuration file to **SELINUX=enforcing**:

vim /etc/selinux/config

```
# This file controls the state of SELinux on the system.
# SELINUX= can take one of these three values:
#   enforcing - SELinux security policy is enforced.
#   permissive - SELinux prints warnings instead of enforcing.
#   disabled - No SELinux policy is loaded.
SELINUX=enforcing
# SELINUXTYPE= can take one of three two values:
#   targeted - Targeted processes are protected,
#   minimum - Modification of targeted policy. Only selected processes are protected.
#   mls - Multi Level Security protection.
SELINUXTYPE=targeted
```

2. Run the following command to automatically enable SELINUX on the file system upon ECS restarting:
- touch /.autorelabel**
3. Run the following command to restart the ECS for the configuration to take effect:

reboot

 NOTE

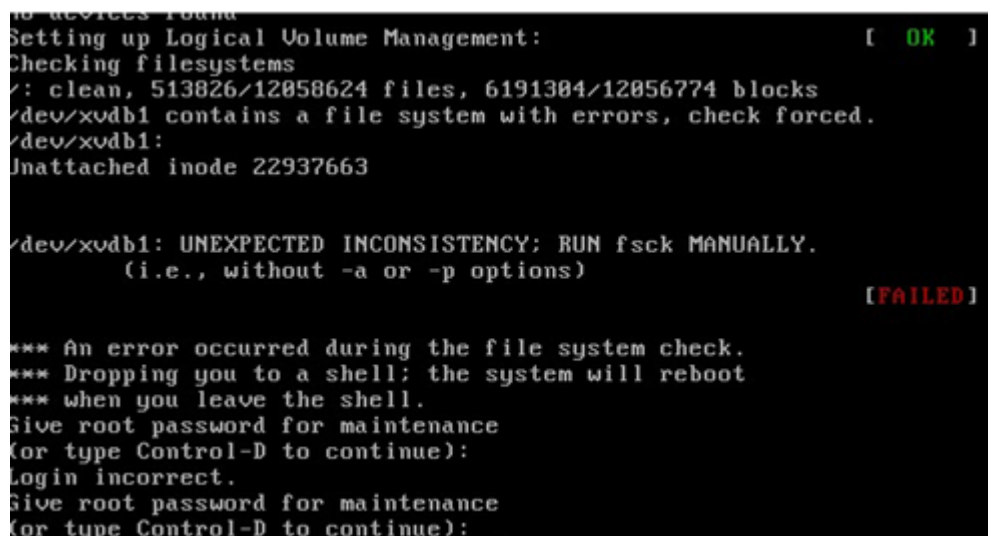
After the preceding command is executed, the system automatically restarts twice.

15.13.18 Why Does a Forcibly-Stopped Linux ECS Fail to Be Restarted?

Symptom

When you try to restart a forcibly-stopped Linux ECS, the ECS failed to be restarted, as shown in [Figure 15-135](#).

Figure 15-135 Restart failure



```
Setting up Logical Volume Management: [ OK ]
Checking filesystems
/: clean, 513826/12858624 files, 6191384/12856774 blocks
/dev/xvdb1 contains a file system with errors, check forced.
/dev/xvdb1:
Unattached inode 22937663

/dev/xvdb1: UNEXPECTED INCONSISTENCY; RUN fsck MANUALLY.
(i.e., without -a or -p options) [FAILED]

*** An error occurred during the file system check.
*** Dropping you to a shell; the system will reboot
*** when you leave the shell.
Give root password for maintenance
(or type Control-D to continue):
Login incorrect.
Give root password for maintenance
(or type Control-D to continue):
```

Possible Cause

As shown in [Figure 15-135](#), the ECS cannot be restarted because the file system was damaged. Forcibly stopping or restarting an ECS is highly risky because this operation may cause inconsistent metadata in the file system, leading to the file system damage.

Solution

Use the disk repair tool (fsck) delivered with the Linux OS to rectify the fault.

The following procedure considers the affected disk partition as `/dev/xvdb1`, which is the partition shown in [Figure 15-135](#).

1. Enter the password of user **root** as prompted.
2. Run the following command to check whether the affected disk partition has been mounted:

```
mount | grep xvdb1
```

- If yes, go to step [3](#).
- If no, go to step [4](#).

3. Run the following command to unmount the affected disk partition:
umount /dev/xvdb1
4. Run the following command to rectify the file system of the affected disk partition:
fsck -y /dev/xvdb1
5. Run the following command to restart the ECS:
reboot

 **NOTE**

If the fault persists, contact customer service for technical support.

15.13.19 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 15-136 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
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 15-137 GPU usage



15.14 File Upload/Data Transfer

15.14.1 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.

15.14.2 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.

15.14.3 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 an SSH key pair,
 - The username is **core** for a CoreOS public image.
 - The username is **root** for a non-CoreOS public image.
 - 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 or converted using a key.
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.

15.14.4 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 15-138 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 15-139 Add user account





Figure 15-140 Add PC

The screenshot shows the 'Add PC' dialog box. At the top, there are two input fields: 'PC name' and 'User account'. Below these are four tabs: 'General', 'Display', 'Devices & Audio', and 'Folders'. The 'Folders' tab is selected and highlighted with a red box. Under the 'Folders' tab, there are several settings: 'Friendly name' (text box with 'Optional' as a placeholder), 'Group' (dropdown menu with 'Saved PCs'), 'Gateway' (dropdown menu with 'No gateway'), and three checkboxes: 'Bypass for local addresses' (checked), 'Reconnect if the connection is dropped' (checked), 'Connect to an admin session' (unchecked), and 'Swap mouse buttons' (unchecked). At the bottom right, there are 'Cancel' and 'Add' 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 15-141 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.

15.14.5 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 15-142 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 15-143 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
```

15.14.6 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 15-144 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 15-145 `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>
```

15.14.7 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 15-146 Setting connection parameters

3. Drag files from the local computer on the left to the target ECS on the right to transfer them.

15.14.8 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
```

15.14.9 How Can I Transfer Data Between a Local Computer and a Windows ECS?

Method 1: Install a Data Transfer Tool


Install a data transfer tool, such as FileZilla on both the local computer and the Windows ECS to transmit data.

Method 2: Configure Local Disk Mapping

Use MSTSC to transfer data. This method does not support resumable transmission. Do not use this method to transfer large files. If you want to transfer a large file, use FTP.

1. Log in to the local computer.
2. Press **Win+R** to open the **Run** text box.
3. Enter `mstsc` to start the remote desktop connection.

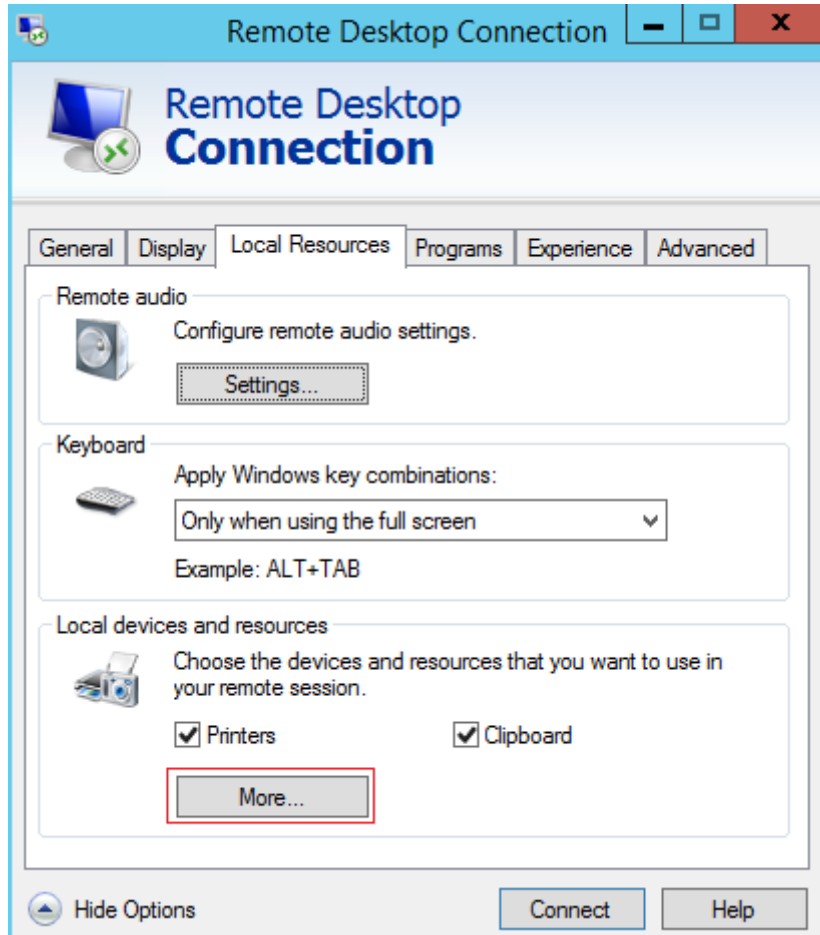
Figure 15-147 Remote Desktop Connection



4. In the **Remote Desktop Connection** window, click  in the lower left corner.

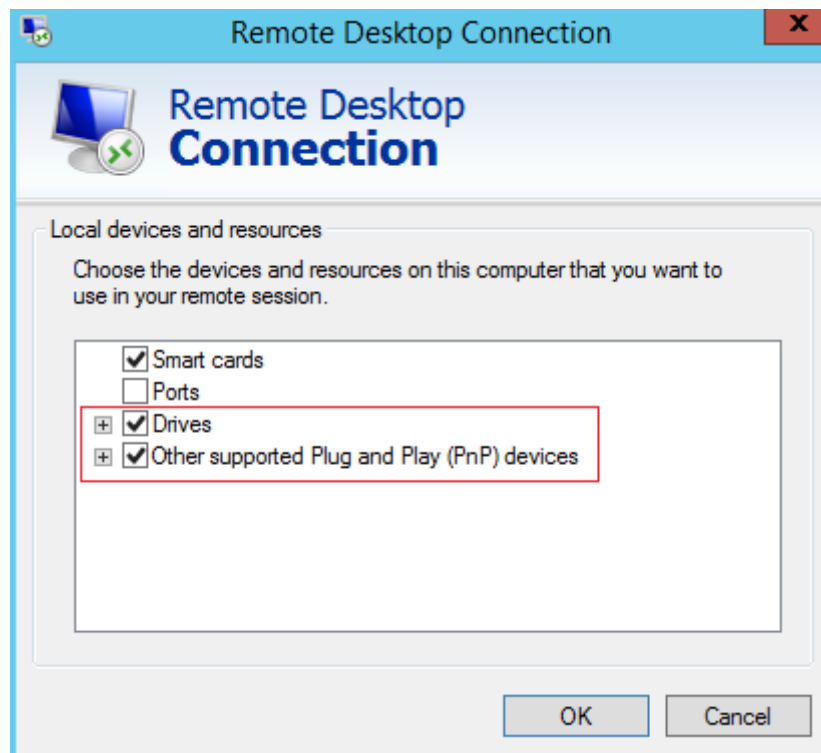
5. Click the **Local Resources** tab and then click **More** in the **Local devices and resources** pane.

Figure 15-148 Local Resources




6. Select **Drives** and **Other supported Plug and Play (PnP) devices** and click **OK** to map all disks on the local computer to the Windows ECS.
If you want to map only certain disks on the local computer to the Windows ECS, expand **Drives** and select the desired ones.

Figure 15-149 Local devices and resources




7. Open the **Remote Desktop Connection** window again and enter the EIP bound to the Windows ECS in the **Computer** text box.

Figure 15-150 Connecting a remote desktop to the Windows ECS



8. Click **Connect**.
Log in to the Windows ECS.
9. Check the disks of the Windows ECS. If the disk information of the local computer is displayed, data can be transmitted between your local computer and the Windows ECS.

Figure 15-151 Viewing disks



Method 3: Set Up an FTP Site

Set up an FTP site and transfer files to the ECS.

15.14.10 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.

15.14.11 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.

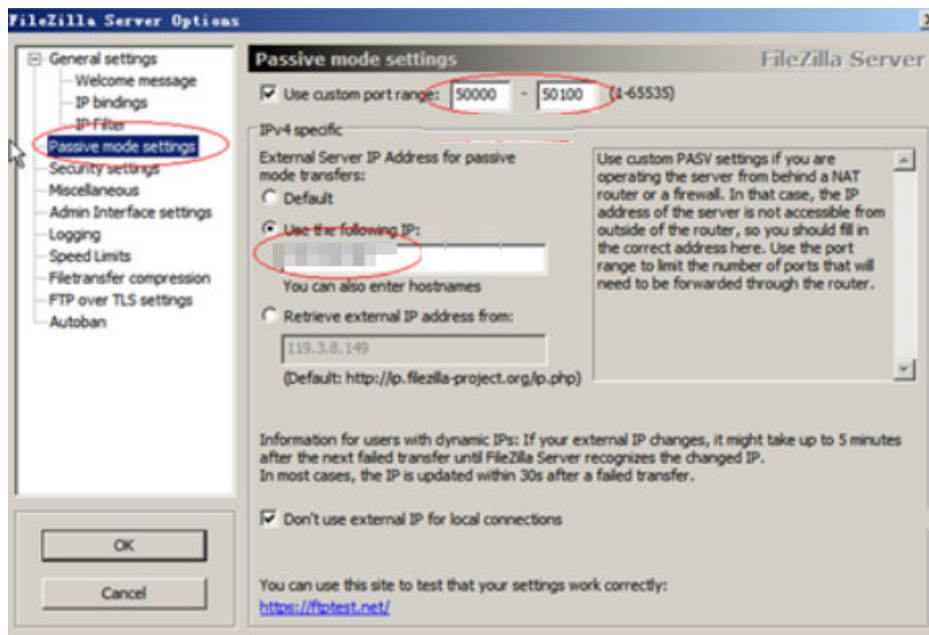
1. Configure the public IP address of the server.
Choose **Edit > Settings**.

Figure 15-152 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 15-153 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.

15.14.12 Why Does Internet Access to an ECS Deployed with FTP Fail?

Symptom

- You cannot access a Windows ECS with FTP deployed by using an EIP.
- The FTP client cannot access the FTP server, and the connection times out.
- It takes a lot of time to upload files.


Possible Causes

- The security group associated with the target ECS denies inbound traffic.
- The firewall of the ECS blocks the FTP process.


Enabling FTP Firewall Support

To allow a server to access an FTP server deployed on an ECS using an EIP, the FTP server must work in passive mode. In this case, enable FTP firewall support.

1. Log in to the management console and then log in to the ECS using .
2. Choose **Start > Server Manager**.
3. In **Server Manager**, choose **Dashboard > Tools > Internet Information Services (IIS) Manager**.




4. Double-click **FTP Firewall Support**.



5. Set parameters and click **Apply**.

- **Data Channel Port Range:** specifies the range of ports used for passive connections. The port range is 1025-65535. Configure this parameter based on site requirements.
- **External IP Address of Firewall:** Enter the public IP address of the ECS.



- Restart the ECS for the firewall configuration to take effect.

Setting the Security Group and Firewall

After deploying FTP, add a rule to the target security group to allow access to the FTP port in the inbound direction.

After [enabling FTP firewall support](#), allow access to the ports used by the FTP site and the data channel ports used by the FTP firewall in the security group.

By default, the firewall allows access to TCP port 21 for FTP. If another port is used, add an inbound rule that allows access to that port on the firewall.

- Log in to the management console.
- Under **Computing**, click **Elastic Cloud Server**.
- On the **Elastic Cloud Server** page, click the name of the target ECS.
The page providing details about the ECS is displayed.
- Click the **Security Groups** tab and view security group rules.
- Click the security group ID.
The system automatically switches to the **Security Group** page.
- On the **Inbound Rules** tab, click **Add Rule** and configure the access rule for the inbound direction.

Set **Source** to the IP address segment containing the IP addresses allowed to access the ECS over the Internet.

The valid port range that can be specified in [Enabling FTP Firewall Support](#) is 1025-65535. For example, the configured data port range is 5000-6000.

NOTE

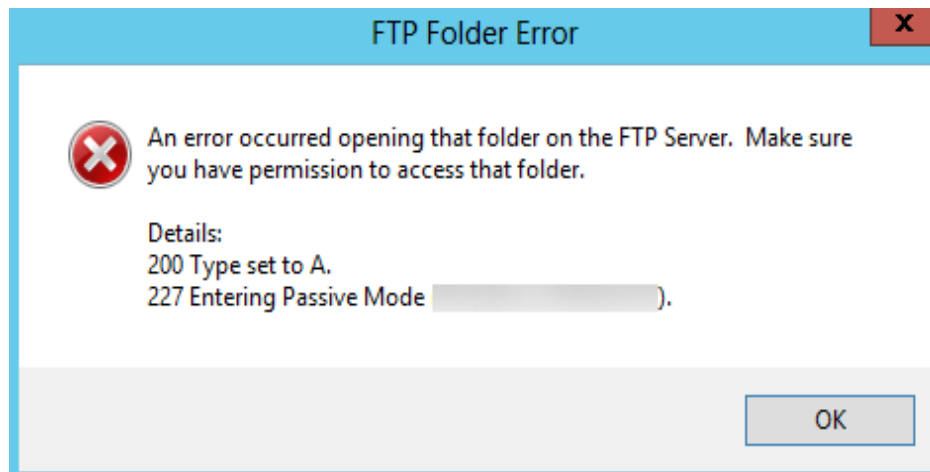
The default source IP address **0.0.0.0/0** indicates that all IP addresses can access ECSs in the security group.

15.14.13 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 15-154 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 15-155 Internet Options



4. Click **OK**, restart Internet Explorer, and open the folder on the FTP server again.

15.14.14 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 15-156 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 15-157 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
```

15.15 ECS Migration

15.15.1 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.

15.16 Disk Management

15.16.1 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.


For details, see [Introduction to Data Disk Initialization Scenarios and Partition Styles](#).

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 15-158 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.

15.16.2 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 ECS as user **root**.
2. 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 15-159 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 ~]# _
```

3. Run the following command to view disk partitions:
parted -l /dev/xvda

Figure 15-160 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 15-161 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 15-162 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 15-163 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 15-164 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 15-165 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 15-166 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 15-167 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 15-168 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 15-169 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

15.16.3 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 15-170 Disk Management




5. Taking disk 1 marked in [Figure 15-170](#) 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 15-170](#).
 - b. Click **Properties**.
The **SCSI Disk Device Properties** dialog box is displayed, as shown in [Figure 15-171](#).

Figure 15-171 Disk properties



- c. Click the **Details** tab and set **Property** to **Parent**.

Figure 15-172 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 15-11](#). The disk device corresponding to **51776** is **xvde**. The disk device used by disk 1 is xvde.

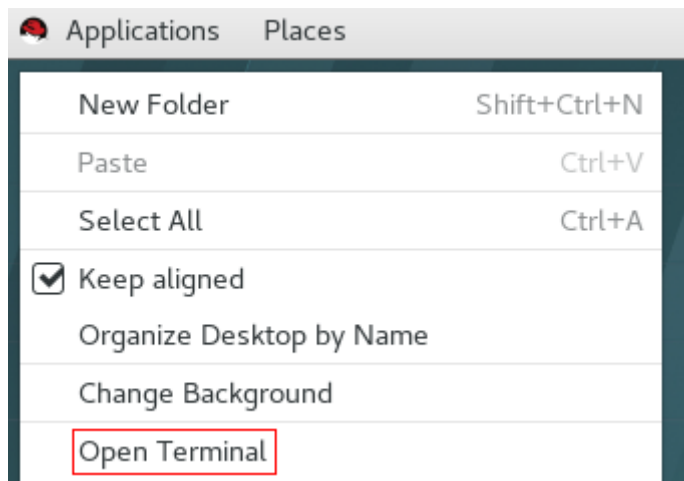
Table 15-11 Mapping between disk partitions and disk devices

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

15.16.4 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 15-173 open terminal

3. Run the following command to view disk partitions and disk devices:
fdisk -l

Figure 15-174 Viewing disk partitions and disk devices

```
root@localhost:~  
File Edit View Search Terminal Help  
[root@localhost ~]# fdisk -l  
  
Disk /dev/xvda: 10.7 GB, 10737418240 bytes, 20971520 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: 0x000ba575  
  
   Device Boot      Start         End      Blocks   Id  System  
/dev/xvda1    *          2048        2099199       1048576   83   Linux  
/dev/xvda2                2099200       16777215       7339008   83   Linux  
/dev/xvda3                16777216       20971519        2097152   82   Linux swap / Solaris  
  
Disk /dev/xvdb: 21.5 GB, 21474836480 bytes, 41943040 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
```

Table 15-12 lists the mapping between disk partitions and disk devices.

Table 15-12 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

15.16.5 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.

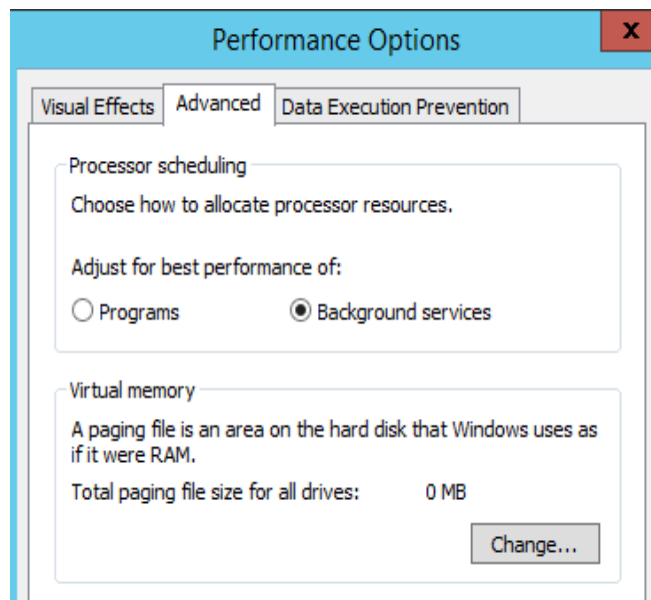
NOTE

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.

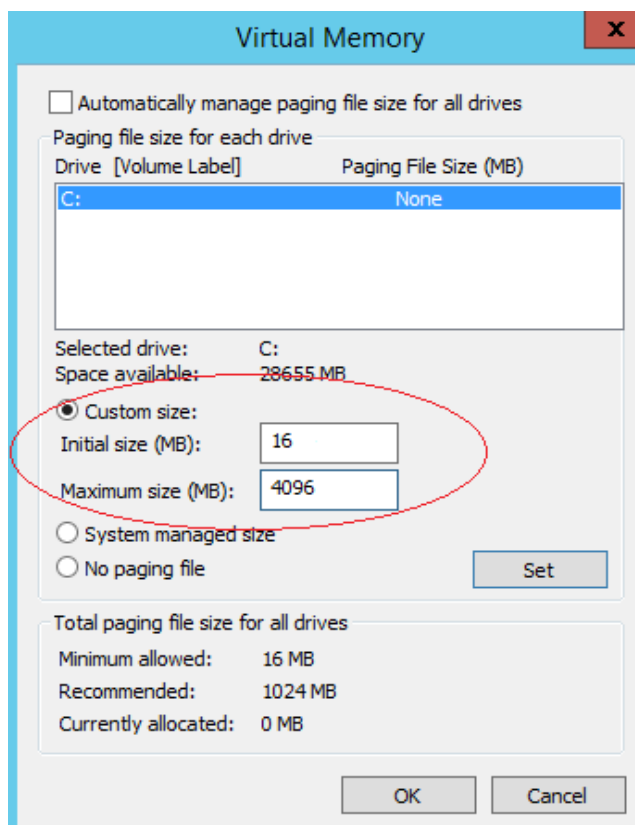
1. Right-click **Computer** and choose **Properties** from the shortcut menu.
2. In the left navigation pane, choose **Advanced system settings**.
The **System Properties** dialog box is displayed.
3. Click the **Advanced** tab and then **Settings** in the **Performance** pane.
The **Performance Options** dialog box is displayed.

Figure 15-175 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 15-176 Virtual Memory



7. Click **Set** and then **OK** to complete the configuration.
8. Restart the ECS for the configuration to take effect.

15.16.6 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
```

15.16.7 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
```

15.16.8 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 15-13 Numbers of disks that can be attached to a newly created ECS

ECS Type	Maximum VBD Disks	Maximum SCSI Disks	Constraint
KVM (excluding D3 ECSs)	24	59	VBD disks + SCSI disks ≤ 60 (This constraint does not apply to local disks.) 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.) 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 15-14](#).

Table 15-14 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

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.
 - 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.

15.16.9 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.
- 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?](#)

15.16.10 Which ECSs Can Be Attached with SCSI EVS Disks?

A Xen ECS running one of the following OSs supports SCSI EVS disks:

- Windows
- SUSE Enterprise Linux Server 11 SP4 64bit
- SUSE Enterprise Linux Server 12 64bit
- SUSE Enterprise Linux Server 12 SP1 64bit
- SUSE Enterprise Linux Server 12 SP2 64bit

All KVM ECSs support SCSI EVS disks.


15.16.11 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.

For details about how to attach disks, see [Attaching an EVS Disk to an ECS](#).

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 device name, see [Using a Serial Number to Obtain the Disk Device 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 device name of Windows ECSs.)
 - If **Device Type** is **SCSI**, use a WWN to obtain the disk device name.
For details, see [Using a WWN to Obtain the Disk Name \(Windows\)](#) and [Using a WWN to Obtain a Disk Device Name \(Linux\)](#).
- QingTian ECS
For QingTian ECSs such as C7 or M7 series, the **Device Type** is **VBD** or **SCSI**, and the device ID is **WWN**. You can use the WWN to obtain the disk device name of the ECS. For details, see [Using a WWN to Obtain the Disk Name \(Windows\)](#) and [Using a WWN to Obtain a Disk Device Name \(Linux\)](#).
- Xen ECS
 - For details, see [Obtaining the Disk Device Name of a Xen Instance](#).

Using a Serial Number to Obtain the Disk Device 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:

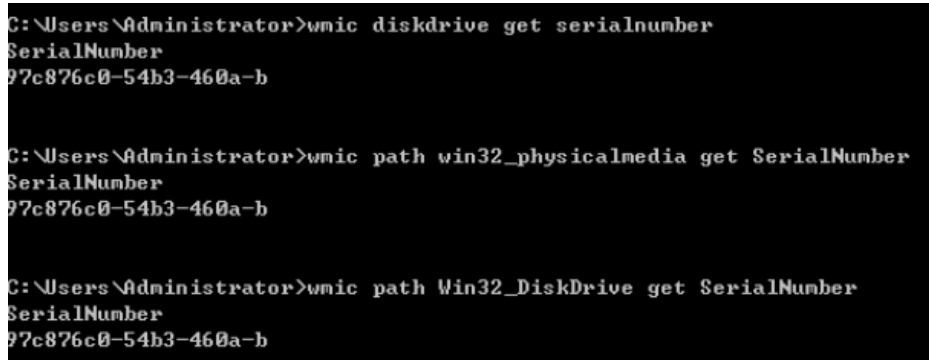
```
wmic diskdrive get serialnumber
```

```
wmic path win32_physicalmedia get SerialNumber
```

```
wmic path Win32_DiskDrive get SerialNumber
```

Information similar to the following is displayed:

Figure 15-177 Obtaining the disk serial number



```
C:\Users\Administrator>wmic diskdrive get serialnumber
SerialNumber
97c876c0-54b3-460a-b

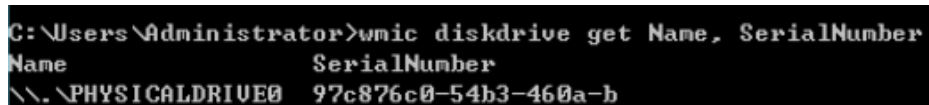
C:\Users\Administrator>wmic path win32_physicalmedia get SerialNumber
SerialNumber
97c876c0-54b3-460a-b

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 15-178 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 15-179](#), 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 15-179](#), the disk is **Disk 0**.

Figure 15-179 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
1
```

- 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.

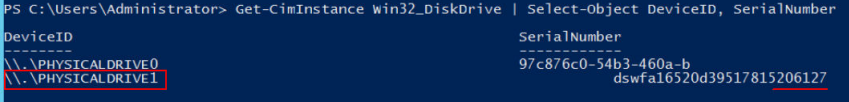
- a. Obtain the 21st to 17th digits that are counted backwards (**3252f**).
- b. Convert a hexadecimal (**3252f**) to a decimal (**206127**).

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 15-180 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
```

Obtaining the Disk Device Name of a Xen Instance

Step 1 Obtain the disk information displayed 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  to expand the disk information.
5. Check the device name, type, and ID of the disk.
 - If the device type is **VBD**, go to [Step 2](#).
 - If the device name is **SCSI**, go to [Step 3](#).

NOTE

If **Device Identifier** is not displayed on the page, stop the ECS and restart it.

Step 2 Check the device name attached to the VBD disk in the ECS.

For a VBD disk, the device name displayed on the management console corresponds to the disk device name in the ECS OS. For details, see [Table 15-15](#).

Table 15-15 Mapping between disk device names displayed on the management console and those obtained on the ECS

Device Name (on Management Console)	Device Name (in ECS)
/dev/sd***	/dev/xvd***
/dev/vd***	/dev/xvd***
/dev/xvd***	/dev/xvd***

An example is provided as follows:

If the device name displayed on the management console is **/dev/sdb**, the device name of the device attached to the ECS is **/dev/xvdb**.

Step 3 Check the device name of the SCSI disk attached to the ECS.

1. Obtain the disk device ID.
The device ID of the SCSI disk is the disk WWN on the ECS.
2. Log in to the ECS as user **root**.
3. Run the following command to view the disk device name:
ll /dev/disk/by-id |grep *WWN*|grep scsi-3

```
[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
```

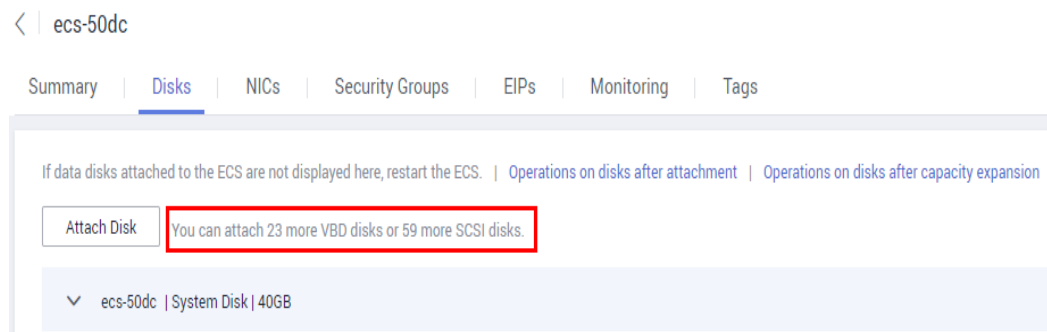
----End

15.16.12 What Should I Do If Attaching a Disk to a Windows ECS Failed But There Are Still Available Device Names?

Symptom

On the Windows ECS details page, the system displays a message indicating that at most n more disks can be attached to the ECS. However, after you clicked **Attach Disk**, the attachment failed.

Figure 15-181 Disk attachment



Possible Causes

If an EVS disk in arrears is not renewed, the system forcibly uninstalls it, which may cause a residual drive letter on the Windows ECS. As a result, the actual number of available device names on the ECS is less than the displayed number.

Solution

Restart the ECS and attach the disk again.

If the attaching still fails, contact customer service for technical support.

15.16.13 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 ECS as user **root**.
2. 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
```

3. 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 2 is **scsi-3688860300001436b005014f890338280**, use the following data to replace **/dev/sdb**:

```
/dev/disk/by-id/scsi-3688860300001436b005014f890338280
```

15.16.14 How Can I Check Whether the ECSs Attached with the Same Shared SCSI Disk Are in the Same ECS Group?

Scenarios

Shared EVS disks of the SCSI type support SCSI locks. To improve data security, the shared EVS disks of the SCSI type must be attached to the ECSs in the same anti-affinity ECS group. This section describes how to check whether the ECSs attached with the same shared SCSI disk are in the same ECS group.

- For details about ECS groups, see [Managing ECS Groups](#).
- For details about using shared EVS disks, see [Shared EVS Disks and Usage Instructions](#).

Procedure

1. Log in to the management console.
2. Under **Storage**, click **Elastic Volume Service**.
3. Click the target shared SCSI disk to view its details.
4. In the **Servers** pane on the right side of the page, the ECSs to which the shared SCSI disk is attached are displayed.

In this example, the ECSs to which the shared SCSI disk **volume-0001** is attached are **ecs-0001** and **ecs-0002**.

5. Click the names of these ECSs, respectively. On the page that provides details about an ECS, you can view the ECS group to which the current ECS belongs.

 **NOTE**

If the ECS group field is left blank, the ECS has not been added to any ECS group.

15.16.15 Can All Users Use the Encryption Feature?


The permissions of users in a user group to use the encryption feature are as follows:

- The user who has security administrator permissions can grant KMS access permissions to EVS for using the encryption feature.
- When a common user who does not have security administrator permissions attempts to use the encryption feature, the condition varies depending on whether the user is the first one in the user group to use this feature.
 - If the common user is the first one in the user group to use the encryption feature, the common user must request a user who has security administrator permissions to grant the common user permissions. Then, the common user can use the encryption feature.
 - If the common user is not the first one in the user group to use the encryption feature, the user directly has the permissions to use the encryption feature.

The following section uses a user group as an example to describe how to grant KMS access permissions to EVS for using the encryption feature.

For example, a user group shown in [Figure 15-182](#) consists of four users, user 1 to user 4. User 1 has security administrator permissions. Users 2, 3, and 4 are common users who do not have security administrator permissions.

Figure 15-182 User group



Scenario 1: User 1 Uses the Encryption Feature

In this user group, if user 1 uses the encryption feature for the first time, the procedure is as follows:

1. User 1 creates Xrole to grant KMS access permissions to EVS.

After user 1 grants permissions, the system automatically creates key **evs/default** for encrypting EVS disks.

NOTE

When user 1 uses the encryption feature for the first time, the user must grant the KMS access permissions to EVS. Then, all the users in the user group can use the encryption feature by default.

2. User 1 selects a key.

One of the following keys can be used:

- Default key **evs/default**
- Custom key, which was created before using the EVS disk encryption feature
- Newly created key (For instructions about how to create a key, see "Creating a Key Pair" in *Data Encryption Workshop User Guide*.)

After user 1 uses the encryption feature, all other users in the user group can use this feature, without requiring to contact user 1 for permissions granting.

Scenario 2: Common User Uses the Encryption Feature

In this user group, when user 3 uses the encryption feature for the first time:

1. The system displays a message indicating that the user has no permissions.
2. User 3 asks user 1 to create Xrole to grant KMS access permissions to EVS.

After user 1 grants the permissions, user 3 and all other users in the user group can use the encryption feature by default.

15.16.16 How Can I Add an ECS with Local Disks Attached to an ECS Group?

An ECS group logically isolates ECSs. The ECSs in an ECS group support anti-affinity and are allocated on different hosts.

An ECS with local disks attached cannot be added to an ECS group after the ECS is created. Such ECSs can be added to an ECS group only during the ECS creation.

15.16.17 Will My EVS Disk Be Deleted When I Delete Its Server?


- For pay-per-use disks:
 - If such a disk is separately purchased and has been attached, the system will prompt you whether to delete the disk when you delete the server, and you can make the decision based on your requirements.

15.16.18 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 15-183](#).

Figure 15-183 Offline disk



Possible Causes

Windows has three types of SAN policies: **OnlineAll**, **OfflineShared**, and **OfflineInternal**.

Table 15-16 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.

15.16.19 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
```

15.16.20 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.

15.16.21 How Can I Rectify the Fault That May Occur on a Linux ECS with an NVMe SSD Disk Attached?


Symptom

When a Linux ECS with an NVMe SSD disk attached, such as a P1 ECS, becomes faulty, you must contact the administrator to remotely rebuild the ECS again.

If automatic NVMe SSD disk attachment upon ECS startup is enabled in `/etc/fstab` on the faulty ECS, the system disk recovers after the ECS is created. However, the attached NVMe SSD disk does not have a file system, and automatic NVMe SSD disk attachment upon ECS startup fails to take effect. As a result, the ECS enters the emergency mode, as shown in [Figure 15-184](#).

6. Verify that the ECS recovers and can be logged in.

Figure 15-187 Logging in to the ECS



15.16.22 Why Is the Device Name of My C6 ECS in the sd* Format?

Symptom

The device name of previously purchased C6 ECSs is in vd* format, for example, vda and vdb, but the device name of newly purchased C6 ECSs is in sd* format.

This section describes the reason why the device name is changed to the sd* format and how to handle the sd* device name in common scenarios.

Root Cause

The device name of the Linux system is automatically generated based on certain rules that are related to the disk protocol and disk sequence number, which brings some uncertainties. When disks are attached to C6 ECSs, either virtio-blk or virtio-scsi is used.

- If virtio-blk is allocated, the device name format is vd*.
- If virtio-scsi is allocated, the device name format is sd*.

Disk Partitioning and Formatting

Problem: Before using an ECS for the first time, you need to partition or format the attached data disks. If the ECS device name is in sd* format, running **/dev/vd*** will fail.

Solution: Dynamically obtain the device name and then perform operations on the disk. You can dynamically obtain device names in either of the following ways:

- Method 1: Run fdisk to query the device name.

Log in to the ECS and run the following command to query the data disk list:

```
fdisk -l
```

Information similar to the following is displayed, indicating the ECS has two disks attached. **/dev/vda** is the system disk, and **/dev/vdb** is the new data disk.

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

```
Disk /dev/vda: 42.9 GB, 42949672960 bytes, 83886080 sectors  
Units = sectors of 1 x 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 GB, 107374182400 bytes, 209715200 sectors
Units = sectors of 1 x 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

This is a convenient method to obtain the device name, but you cannot obtain the mapping between the EVS disks attached to the ECS and the device names in the OS. If you want to know the mapping, obtain the device name by referring to method 2.

- Method 2: Use serial-id or wwn to obtain the device name.

For details, see [How Do I Obtain My Disk Device Name in the ECS OS Using the Device Identifier Provided on the Console?](#).

Automatic Mounting of File Systems

You are advised to use UUIDs to identify disks in the file because they are unique identifiers for disk partitions and do not change with device names. Use the UUID of the file system to configure automatic mounting for a system disk.

- Automatic Mounting for a System Disk
 - If a public image or a private image created from a public image is used, UUIDs are used for automatic disk mounting and no action is required.
 - If a private image created using a non-public image is used, select **Enable automatic configuration** when creating the image. Then, the system automatically uses UUIDs for automatic disk mounting.
 - If **Enable automatic configuration** is not selected when you are creating a private image, see [Changing the Disk Identifier in the fstab File to UUID](#).

15.16.23 Why Are Disk Error Logs Printed After a Disk Attached to an ECS Is Formatted with the ext4 File System?

Symptom

When a VBD disk is attached to an ECS and the partition is in ext4 format, the following log may be displayed on the console:

```
blk_update_request: operation not supported error, dev vdb, sector 826298624 op 0x9:(WRITE_ZEROES) flags 0x800 phys_seg 0 prio class 0
```

Figure 15-188 Printed logs

```
[ 1732.062294] blk_update_request: operation not supported error, dev vdb, sector 8504 op 0x9:(WRITE_ZEROES) flags 0x800 phys_seg 0 prio class 0
[ 1732.366291] blk_update_request: operation not supported error, dev vdb, sector 12592 op 0x9:(WRITE_ZEROES) flags 0x800 phys_seg 0 prio class 0
[ 1732.654268] blk_update_request: operation not supported error, dev vdb, sector 16688 op 0x9:(WRITE_ZEROES) flags 0x800 phys_seg 0 prio class 0
[ 1732.942279] blk_update_request: operation not supported error, dev vdb, sector 20784 op 0x9:(WRITE_ZEROES) flags 0x800 phys_seg 0 prio class 0
[ 1733.230277] blk_update_request: operation not supported error, dev vdb, sector 24880 op 0x9:(WRITE_ZEROES) flags 0x800 phys_seg 0 prio class 0
```

Involved OSs: Ubuntu 20.04, CentOS 8.0, CentOS 8.1, and other ECSs whose kernel versions are 4.18 or later

Root Cause

VBD disks do not support the advanced SCSI command `WRITE_ZEROES`.

If the ECS OS kernel version is 4.18 or later and the disk partition is formatted with the ext4 file system, the `WRITE_ZEROES` command is delivered. The system does not support the command and prints a log, which has no impact on the ECS performance and you can ignore it.

15.17 Passwords and Key Pairs

15.17.1 How Can I Change the Password for Logging In to a Linux ECS?

Solution

1. Use the existing key file to log in to the Linux ECS as user **root**.
2. Run the following command to reset the password of user **root**:

```
passwd
```

To reset the password of another user, replace **passwd** with **passwd *username***.

3. Enter the new password as prompted.

```
New password:
```

```
Retype new password:
```

```
If the following information is displayed, the password has been reset:  
passwd: all authentication tokens updates successfully
```

15.17.2 What Is the Default Password for Logging In to a Linux ECS?

The default username for logging in to an ECS running Linux (such as CentOS or Ubuntu) is **root**, and the password is the one you set during ECS creation.

15.17.3 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 15-189 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
```

15.17.4 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.

Prerequisites

The ECS can be logged in.

Background

[Table 15-17](#) shows the ECS password complexity requirements.

Table 15-17 Password complexity requirements

Parameter	Requirement
Password	<ul style="list-style-type: none">• Consists of 8 to 26 characters.• 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.

Windows

1. Log in to the ECS.
For details, see [Login Overview \(Windows\)](#).
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 15-17](#)):

```
net user Administrator New password
```

Linux

1. Use the existing key file to log in to the ECS as user **root** through SSH.
For details, see [Logging In to a Linux ECS Using an SSH Key Pair](#).
 2. Run the following command to reset the password of user **root**:
- ```
passwd
```
- To reset the password of another user, replace **passwd** with **passwd username**.
3. Enter the new password as prompted. Ensure that the new password meets the requirements described in [Table 15-17](#).

```
New password:
Retype new password:
```

```
If the following information is displayed, the password has been changed:
passwd: all authentication tokens updates successfully
```

## 15.17.5 What Should I Do If the System Displays a Message Indicating that the Password Is Incorrect When I Remotely Log In to My ECS?

### Solution

Check the network configuration of the ECS and determine whether the fault is caused by a **Cloud-Init** failure.

- Verify that port 80 is bypassed in both inbound and outbound directions in the security group to which the target ECS belongs.
- Verify that DHCP is enabled in the subnet to which the target ECS belongs.

#### NOTE

After verifying the preceding configurations, restart the ECS, wait for 3 to 5 minutes, and remotely log in to the ECS using a password or key.

## 15.17.6 What Should I Do If I Cannot Log In to My ECS Using the Initial Password After I Use It for a Period of Time?

### Solution

Check whether the remote login page can be displayed.

- If the login page cannot be displayed, an error may have occurred in the GuestOS process on the ECS. In such a case, contact customer service for troubleshooting.
- If the login page can be displayed, log in to the OS in single-user mode for troubleshooting. The procedure is as follows:
  - Check whether the password can be changed in single-user mode.  
If the password can be changed, change it and contact customer service to check whether the password has been maliciously changed due to an attack.
  - If the password cannot be changed, verify that the values of **hard** and **soft** in **/etc/security/limits.conf** are not greater than 65535.

```
- nice - max nice priority allowed to raise to values: 1-20, 19
- rtprio - max realtime priority
#
#<domain> <type> <item> <value>
#
#* soft core 0
#* hard rss 10000
#@student hard nproc 20
#@faculty soft nproc 20
#@faculty hard nproc 50
#ftp hard nproc 0
#@student - maxlogins 4
End of file
```

Change the password in single-user mode and try to log in to the ECS again.

## 15.17.7 Disabling SELinux

### NOTE

SUSE does not have the SELinux configuration files. You can skip this section.

### Procedure

1. Use the vi editor to open `/etc/selinux/config`.  
**vi /etc/selinux/config**
2. Press **i** to enter insert mode and set the value of **SELINUX** to **disabled**.

```
This file controls the state of SELinux on the system.
SELINUX- can take one of these three values:
enforcing - SELinux security policy is enforced.
permissive - SELinux prints warnings instead of enforcing.
disabled - No SELinux policy is loaded.
SELINUX=disabled
SELINUXTYPE- can take one of three two values:
targeted - Targeted processes are protected.
minimum - Modification of targeted policy. Only selected processes
mls - Multi Level Security protection.
SELINUXTYPE=targeted
```


3. Press **Esc** and enter `:wq` to save and exit the file.

## 15.17.8 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. Click  in the upper left corner and select your region and project.
3. Under **Computing**, click **Elastic Cloud Server**.
4. On the **Elastic Cloud Server** page, select the target ECS.
5. Click the name of the target ECS.  
The page providing details about the ECS is displayed.
6. Obtain the **Key Pair** value.  
The value is the key pair used by the ECS.

## 15.17.9 How Can I Use a Key Pair?

### Symptom


When you purchase an ECS, the system asks you to select a login mode. If you select **Key pair**, you are required to select an existing key pair or create a new pair.

If no key pair is available, create one on the management console.

## Solution


1. In the navigation pane of the ECS console, choose **Key Pair**. Then, click **Create Key Pair**.

**Figure 15-190** Creating a key pair



2. After the key pair is created, download the private key to a local directory.

**Figure 15-191** Downloading a key pair



3. When purchasing an ECS, select the created or existing key pair in **Key pair**.



**Figure 15-192** Selecting a key pair


The screenshot shows the 'Login Mode' section with two tabs: 'Password' and 'Key pair'. The 'Key pair' tab is selected. Below the tabs, there is a text instruction: 'To click Remote Login to log in to a Linux ECS in key pair login mode, you must set a login password after the ECS is created. Learn how to set the password. The private key will be required for logging in to the ECS and for reinstalling or changing the OS. Keep it secure.' Below this, the 'Key Pair' section has a dropdown menu with 'KeyPair-389e' selected, a 'Create Key Pair' button, and a help icon. A checkbox is checked with the text: 'I acknowledge that I have obtained private key file KeyPair-389e.pem and that without this file I will not be able to log in to my ECS.' Below the checkbox, there is a note: 'After a Linux ECS is created, use this key pair to log in to the ECS. After a Windows ECS is created, locate the row that contains the ECS in the ECS list, click Get I password.'

## Helpful Links

- [\(Recommended\) Creating a Key Pair on the Management Console](#)
- [Logging In to a Linux ECS Using an SSH Key Pair](#)

### 15.17.10 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**.

### 15.17.11 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 customer service for technical support.

## 15.17.12 What Should I Do If I Cannot Download a Key Pair?


The private key file of a key pair can be downloaded only once.

If your private key file has been lost, create a key pair and download the private key file again.

## Solution


1. Log in to the management console and choose **Key Pair**.

Figure 15-193 Key Pair

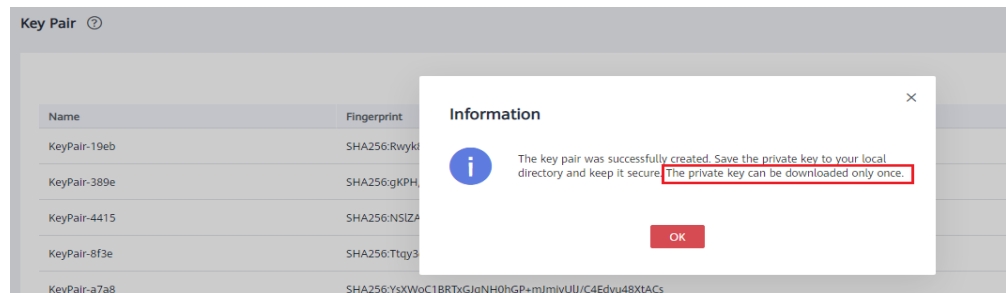


2. Click **Create Key Pair**.

Figure 15-194 Creating a key pair



3. Click **OK** to save the private key to your local directory.

**Figure 15-195** Saving the private key

## 15.17.13 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

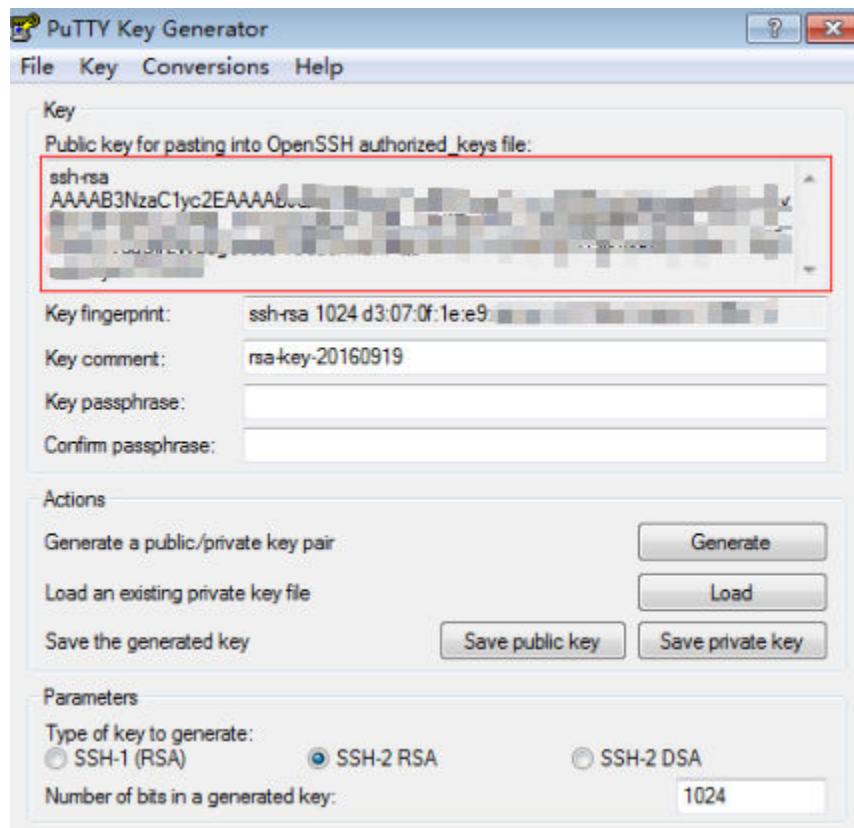
Use the locally stored private key and **PuTTY Key Generator** to restore the format of the public key content. Then, import the public key to the management console.


1. Double-click **puttygen.exe** to open **PuTTY Key Generator**.

**Figure 15-196** 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 15-197](#) is the public key whose format meets system requirements.

**Figure 15-197** 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. Click  in the upper left corner and select your region and project.
  - c. Under **Computing**, click **Elastic Cloud Server**.
  - d. In the navigation pane on the left, choose **Key Pair**.
  - e. On the key pair page, click **Import Key Pair**.
  - f. Copy the public key content in the .txt file to **Public Key Content** and click **OK**.

## 15.17.14 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 15-198 ecs01

```
----- BEGIN DUMP -----
cloudbase-init:1003:AAD3B435B51404EEAAD3B435B51404EE:CCA3B0DEB517A0E2342AEB34C0473C39:::
Guest:501:AAD3B435B51404EEAAD3B435B51404EE:31D6CFE0D16AC931B73C59D7E0C089C0:::
Administrator:500:AAD3B435B51404EEAAD3B435B51404EE:27CF57575EB83D9A6D7D27831157A947:::
----- END DUMP -----
3 dumped accounts
```

- LM-HASH and NTLM-HASH values of the **cloudbase-init** account for ecs02

Figure 15-199 ecs02

```
----- BEGIN DUMP -----
cloudbase-init:1003:AAD3B435B51404EEAAD3B435B51404EE:5B635D5F5306E26E0EE66915D7C1CA9B:::
Guest:501:AAD3B435B51404EEAAD3B435B51404EE:31D6CFE0D16AC931B73C59D7E0C089C0:::
Administrator:500:AAD3B435B51404EEAAD3B435B51404EE:0501525C0083243750D23927A82070B6:::
----- END DUMP -----
3 dumped accounts
```

## 15.17.15 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 15-200](#).

Figure 15-200 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 15-201 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
```

## 15.18 Network Configurations

### 15.18.1 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 15-18](#) lists ECS configurations.

Table 15-18 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**.

## 15.18.2 Can an ECS Without an EIP Bound Access the Internet?

Yes.

You can use the NAT Gateway service to allow ECSs in a VPC to access the Internet using an EIP. The SNAT function provided by the NAT Gateway service allows the ECSs in a VPC to access the Internet without requiring an EIP. Additionally, SNAT supports a large number of concurrent connections for applications that have a large number of requests and connections. For more information about NAT Gateway, see *NAT Gateway Service Overview*.



## 15.18.3 What Should I Do If an EIP Cannot Be Pinged?

### Symptom

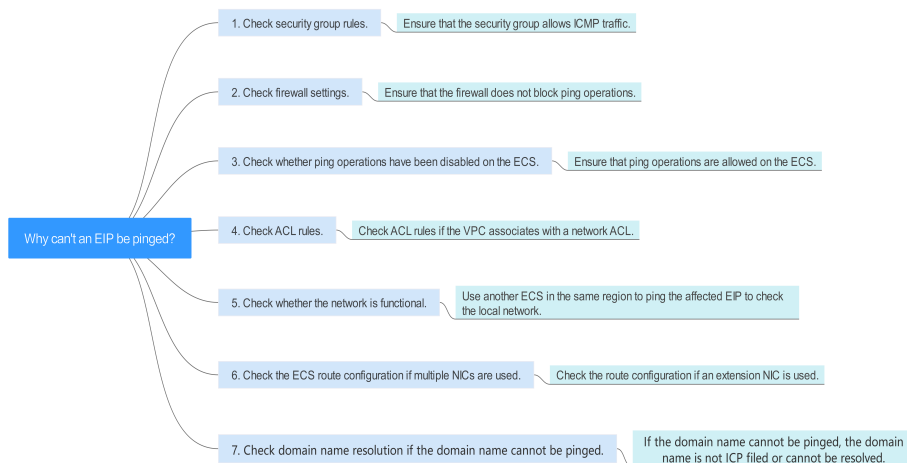
After you purchase an EIP and bind it to an ECS, the local host or other cloud servers cannot ping the EIP of the ECS.

### 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 15-202** Method of locating the failure to ping an EIP



**Table 15-19** 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> . |

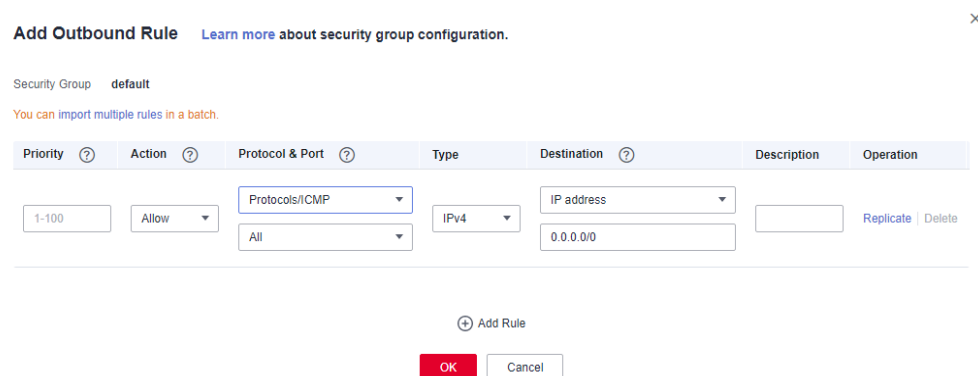
| Possible Cause                                               | Solution                                                                                                                                                                                                                             |
|--------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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> . |
| The domain name is not ICP licensed.                         | If the domain name cannot be pinged or cannot be resolved, see <a href="#">Checking Domain Name Resolution If the Domain Name Cannot Be Pinged</a> to resolve this issue.                                                            |

## Checking Security Group Rules

ICMP is used for the ping command. Check whether the security group accommodating the ECS allows ICMP traffic.

- Under **Computing**, click **Elastic Cloud Server**.
- On the **Elastic Cloud Server** page, click the name of the target ECS.  
The page providing details about the ECS is displayed.
- Click the **Security Groups** tab, expand the information of the security group, and view security group rules.
- Click the security group ID.  
The system automatically switches to the **Security Group** page.
- On the **Outbound Rules** page, click **Add Rule**. In the displayed dialog box, set required parameters to add an outbound rule.

**Figure 15-203** Adding an outbound rule



**Table 15-20** 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. |

6. On the **Inbound Rules** tab, click **Add Rule**. In the displayed dialog box, set required parameters to add an inbound rule.

**Figure 15-204** Adding an inbound rule

Add Inbound Rule [Learn more about security group configuration.](#)

Security Group default

You can import multiple rules in a batch.

| Priority | Action | Protocol & Port | Type | Source     | Description | Operation          |
|----------|--------|-----------------|------|------------|-------------|--------------------|
| 1-100    | Allow  | Protocols/ICMP  | IPv4 | IP address |             | Replicate   Delete |
|          |        | All             |      | 0.0.0.0/0  |             |                    |

+ Add Rule

OK Cancel

**Table 15-21** 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. |

7. 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 15-205](#) is displayed, there is no ICMP rule blocking the ping operations.

**Figure 15-205** Checking firewall rules

```
[root@ecs-3c4e ~]# iptables -L
Chain INPUT (policy ACCEPT)
target prot opt source destination
ACCEPT icmp -- anywhere anywhere icmp echo-request

Chain FORWARD (policy ACCEPT)
target prot opt source destination

Chain OUTPUT (policy ACCEPT)
target prot opt source destination
ACCEPT icmp -- anywhere anywhere icmp echo-reply
[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)**


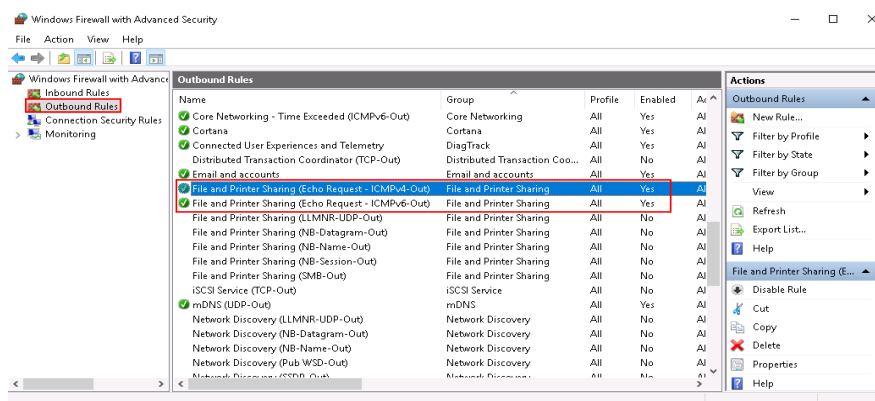
**Figure 15-206** Inbound Rules

Figure 15-207 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 15-208** 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.

## Checking Domain Name Resolution If the Domain Name Cannot Be Pinged

If you can ping the EIP but not the domain name, the possible cause is that an error occurred in domain name resolution.

1. Check the domain name resolution.

If the domain name records are incorrectly configured, the domain name may fail to be resolved.

Switch to the DNS management console to view details about the domain name resolution.

2. Check the DNS server configuration.

If the system shows no server found after you ping a domain name, this issue may be caused by slow response from the DNS server.

## 15.18.4 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**

## 15.18.5 How Do I Query the Egress Public IP Address of My ECS?

### Scenarios

After servers are migrated to the cloud, they usually use EIPs to access the Internet.

You can log in to the management console and view the EIP bound to the ECS in the ECS list. For details, see [Viewing ECS Details \(List View\)](#).

If you want to query the EIP bound to the ECS without logging in to the management console, do as follows.

This section uses an ECS running CentOS 7.5 as an example.

### Procedure

1. [Log in to an ECS](#).
2. Run any of the following commands to query the EIP of the ECS:

- **curl icanhazip.com**
- **curl ifconfig.me**
- **curl ipinfo.io/ip**
- **curl ipecho.net/plain**
- **curl www.trackip.net/i**

## 15.18.6 How Can I Configure the NTP and DNS Servers for an ECS?

### For Linux OSs

Take the NTP and DNS servers running SUSE as an example.

#### Step 1 Configure the NTP server for the ECS.

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 edit the **ntp.conf** configuration file:  
**vim /etc/ntp.conf**
4. Add the following statement to configure the NTP server:  
**server Domain name or IP address of the NTP server**

Example:

If the IP address of the NTP server is 192.168.56.1, add the following statement:

**server 192.168.56.1**

5. Run the following command to start the NTP service upon system restart:  
**service ntp restart**
6. Run the following command to check the status of the NTP server:  
**service ntp status**

#### NOTE

If you want to disable NTP, perform the following steps:

1. Run the **service ntp stop** command to stop NTP.
2. Run the **systemctl disable ntp** command to disable the function of automatically starting NTP upon ECS startup.

#### Step 2 Configure the DNS server for the ECS.

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 edit the **resolv.conf** configuration file:  
**vi /etc/resolv.conf**
4. Add the following statement to configure the DNS server:  
**nameserver = IP addresses of the DNS servers**



Example:

If the IP addresses of the DNS servers are 8.8.8.8 and 4.4.4.4, add the following statements:

```
nameserver = 8.8.8.8
```

```
nameserver = 4.4.4.4
```

 NOTE

The IP addresses of the DNS servers must be the same as those in the VPC subnet. Otherwise, the DNS modification cannot persistently take effect.

5. Run the following command to restart the network:

```
rcnetwork restart
```

```
service network restart
```

```
/etc/init.d/network restart
```

----End

## Windows


Take an ECS running Windows Server 2012 as an example.

**Step 1** Log in to the Windows ECS as user **Administrator**.

**Step 2** Enable the local area connection.

1. In the lower right corner of the taskbar, right-click the network connection icon.
2. Click **Open Network and Sharing Center**.

**Figure 15-209** Open Network and Sharing Center




3. In the navigation pane on the left, click **Change adapter settings**.

**Step 3** Configure the DNS server for the ECS.


1. Double-click network connections.
2. Click **Properties** in the lower left corner.

Figure 15-210 Local area connection




3. Select **Internet Protocol Version 4 (TCP/IPv4)** and click **Properties**.

Figure 15-211 Selecting a protocol type



4. Select **Use the following DNS server addresses** and set the IP addresses of the DNS servers.

**Figure 15-212** Setting the IP addresses of the DNS servers**Step 4** Configure the NTP server for the ECS.

1. Start the **Run** dialog box. Enter **regedit** and click **OK**.
2. Modify the registry entries.
  - In **HKEY\_LOCAL\_MACHINE > SYSTEM > CurrentControlSet > Services > W32Time > TimeProviders > NtpClient**, set the value of **Enabled** to **1**, indicating that the NTP client is used.
  - In **HKEY\_LOCAL\_MACHINE > SYSTEM > CurrentControlSet > Services > W32Time > TimeProviders > NtpServer**, set the value of **Enabled** to **0**, indicating that the NTP server is stopped.
  - Choose **HKEY\_LOCAL\_MACHINE > SYSTEM > CurrentControlSet > Services > W32Time > Parameters** file and set the **NtpServer** data. Set the data of **TYPE** to **NTP**.
  - In **HKEY\_LOCAL\_MACHINE \ SYSTEM \ CurrentControlSet \ Services \ W32Time \ TimeProviders \ NtpClient**, set the value of **SpecialPollInterval** to **60** and that of **Base** to **Decimal**, indicating the clock synchronization cycle is 60s.
  - In **HKEY\_LOCAL\_MACHINE \ SYSTEM \ CurrentControlSet \ Services \ W32Time \ config**, set the values of **MaxPosPhaseCorrection** and **MaxNegPhaseCorrection** to **fffffff** and that of **Base** to **Hexadecimal**.

3. Open the **Run** dialog box, enter **services.msc**, and click **OK**. The **Services** window is displayed.
4. View the service named **Windows Time** and set the **Start Type** to **Automatic** to synchronize time from the NTP server.
5. Open the **Run** dialog box and run the following commands in sequence to restart the Windows Time service:  
**net stop w32time**  
**net start w32time**
6. Manually change the time on the client to make it different from that on the NTP server. One minute later, check whether the time on the client is the same as that on the NTP server. If yes, the time is synchronized.

----End

## 15.18.7 What Should I Do If NIC Flapping Occurs After My ECS Specifications Are Modified?

### Symptom

Take a Linux ECS as an example. After the user modified ECS specifications and ran the **ifconfig** command, the user found that the original eth0 and eth1 NICs were changed to eth2 and eth3 NICs, indicating that NIC flapping occurred.

### Root Cause

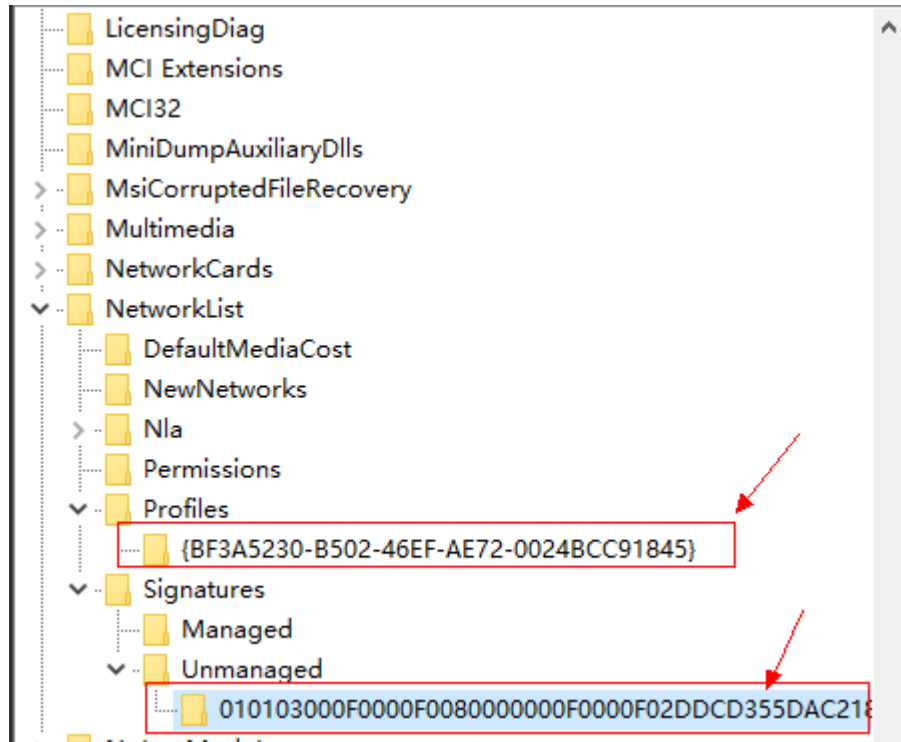
NIC flapping occurs because NIC retaining is enabled in the image from which the ECS is created.

### Solution to Windows

For a Windows ECS, delete the directories in the following registries and restart the ECS to resolve this issue:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion
\NetworkList\Profiles
```

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion
\NetworkList\Signatures\Unmanaged
```



## Solution to Linux

For a Linux ECS, perform the following operations and restart the ECS to resolve this issue:

1. Run the following command to view the files in the network rule directory:  
**ls -l /etc/udev/rules.d**
2. Run the following commands to delete the files with both **persistent** and **net** included in file names from the network rule directory:  
**rm -fr /etc/udev/rules.d/\*net\*persistent\*.rules**  
**rm -fr /etc/udev/rules.d/\*persistent\*net\*.rules**
3. Run the following command to check whether the initrd image file with a name starting with **initrd** and ending with **default** contains both **persistent** and **net** network rules (change the italic data in the following command to the actual OS version):  
**lsinitrd /boot/initrd-2.6.32.12-0.7-default |grep persistent|grep net**
  - If yes, go to steps **4** and **5**.
  - If no, no further action is required.
4. Run the following command to back up the initrd image file (change the italic data in the following command to the actual OS version):  
**cp /boot/initrd-2.6.32.12-0.7-default /boot/initrd-2.6.32.12-0.7-default\_bak**
5. Run the following command to regenerate the initrd image file:  
**mkinitrd**

Perform the following operations when an OS, such as Ubuntu, uses the initramfs image:

1. Run the following command to check whether the initramfs image file with a name starting with **initrd** and ending with **generic** contains both **persistent** and **net** network rules:

```
lsinitramfs /boot/initrd.img-3.19.0-25-generic|grep persistent|grep net
```

- If yes, go to steps [2](#) and [3](#).
- If no, no further action is required.

2. Run the following command to back up the initrd image file:

```
cp /boot/initrd.img-3.19.0-25-generic /boot/initrd.img-3.19.0-25-generic_bak
```

3. Run the following command to regenerate the initramfs image file:

```
update-initramfs -u
```

## 15.18.8 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.

## 15.18.9 How Do I Change the CIDR Block of an ECS Subnet?

### Scenarios


You want to change the CIDR block of an ECS subnet. After you create a subnet, you cannot directly change its CIDR block.

To change a CIDR block, you need to change the subnet.

### Prerequisites

The ECS has been stopped.

### Procedure

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. In the search box above the ECS list, enter the ECS name, IP address, or ID, and click  for search.
4. Click the name of the ECS whose subnet needs to be modified.  
The page providing details about the ECS is displayed.
5. Click the **NICs** tab. Locate the row containing the NIC and click **Modify Private IP**.  
The **Modify Private IP** dialog box is displayed.
6. Change the subnet and private IP address of the primary NIC as required.

#### NOTE

- You can only change to a subnet within the same VPC.
- If you do not specify the target private IP address, the system will automatically assign one to the primary NIC.

For example, the original subnet is **subnet-demo (192.168.0.0/24)** and the new subnet is **subnet-fe21 (192.168.6.0/25)**. Therefore, you change the ECS subnet CIDR block by changing the ECS subnet.

## 15.18.10 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 customer service 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 15-213 Normal network communication

```
iroote@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 4.29.43.28

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
CID index : 0
Max inline data : 8[B]
rdma_cm QPs : OFF
Data ex. method : Ethernet

local address: LID 0x05 QPN 0x0067 PSN 0xaaccfb RKey 0x001c0c VAddr 0x007fb3cd1b0000
CID: 254:128:00:00:00:00:00:00:00:03:00:135:40:178
remote address: LID 0x05 QPN 0x006a PSN 0xebbf6d RKey 0x001c10 VAddr 0x007fdad5990000
CID: 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 12132.70 11900.18 0.190403
```



- If the terminal display is shown in [Figure 15-213](#), the network communication between the two InfiniBand ECSs is normal.
- If the InfiniBand network is inaccessible, contact customer service for technical support.

----End

## 15.18.11 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
```

## 15.18.12 Why Is the NIC Not Working?

### Symptom

The NIC fails to work.

### Possible Causes

The NIC driver has not been correctly installed.

### Solution

D1 and H1 ECSs use passthrough NICs to improve network performance. You must install the passthrough NIC driver on the ECSs or the image that is used for creating the ECSs.

**NOTE**

If you mount the CD/DVD-ROM driver over a VPN, ensure that the VPN bandwidth is greater than 8 Mbit/s.

To install the passthrough NICE driver, do as follows:

- Step 1** Obtain the passthrough NIC driver.

Passthrough NIC driver versions vary depending on the OS. For details, see [Table 15-22](#).

**Table 15-22** NIC driver versions and OSs

| NIC Driver Version | OS               | How to Obtain                                                                                                                                           |
|--------------------|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| ixgbev 2.16.4      | CentOS 7.2 64bit | <a href="https://sourceforge.net/projects/e1000/files/ixgbev%20stable/2.16.4/">https://sourceforge.net/projects/e1000/files/ixgbev%20stable/2.16.4/</a> |

- Step 2** Log in to the ECS.

For more details, see [Login Overview \(Linux\)](#).

- Step 3** Install the passthrough NIC driver on the ECS. In this procedure, CentOS 7.2 64bit is used as an example.

1. Configure the passthrough NIC.

Not all ECS OSs identify passthrough NICs using the standard NIC naming rule of **eth $x$** , where  $x$  is a number. If this is the case, you must configure the ECS so that it can identify the passthrough NIC. The procedure is as follows:

- a. Run the following command to view all NICs on the ECS and identify the passthrough NIC:

```
ifconfig -a
```

- b. Run the following command to switch to the directory where configuration files are stored:

```
cd /etc/sysconfig/network-scripts/
```

- c. Run the following command to create a configuration file for the passthrough NIC:

```
cp ifcfg-eth0 ifcfg-NIC_name
```

In the preceding command, *NIC\_name* specifies the name of the passthrough NIC.

- d. Use the vi editor to edit this configuration file:

```
vi ifcfg-NIC_name
```

- e. Set the **DEVICE** parameter in the configuration file to the name of the passthrough NIC. The following is an example configuration:

```
DEVICE="NIC_name"
BOOTPROTO="dhcp"
ONBOOT="yes"
STARTMODE="onboot"
```

- f. Run the following command to restart the network service and allow the configuration to take effect:

```
service network restart
```

2. Upload the obtained passthrough NIC driver to a directory on the ECS, for example, **/home**.
3. Switch to user **root** on the ECS CLI and open the target directory.  
In this example, the passthrough NIC driver is stored in the **/home** directory. Run the **cd /home** command to switch to the target directory.
4. Run the following command to decompress the software package.  

```
tar -zxvf ixgbevf-2.16.4.tar.gz
```
5. Run the following command to switch to the generated **src** directory:  

```
cd ixgbevf-2.16.4/src
```
6. Run the following commands to install the driver:  

```
make
make install
```
7. Run the following command to restart the ECS to make the drive take effect:  

```
reboot
```
8. Switch to user **root** on the ECS CLI and open the **src** directory, for example, by running the **cd /home/ixgbevf-2.16.4/src** command. Then, run the following commands to check whether the driver has been installed:  

```
rmmmod ixgbevf
insmod ./ixgbevf.ko
ethtool -i NIC_name
```

In the preceding command, *NIC\_name* specifies the passthrough NIC name, for example, **ens5**.

 **NOTE**

- After you run the **rmmod ixgbevf** command, the system may display an error message. This message does not affect the installation of the passthrough NIC driver and can be ignored.
  - *NIC\_name* specifies the passthrough NIC name, for example, **ens5**.
9. Check the driver status based on the displayed information.  
In this example, the driver is installed if **driver** is **ixgbevf** and **version** is **2.16.4**.

----End

## 15.18.13 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.

## 15.18.14 How Can I Add a Static Route to a CentOS 6.5 OS?

### Scenarios

After the system restarts, non-static routes are lost, affecting network availability. To prevent this issue from occurring, you must add static routes to the system.

### Procedure

The following section uses a CentOS 6.5 OS as an example.

1. Log in to the ECS.
2. Create or modify the static route configuration file.

If the **static-routes** configuration file is not in the **/etc/sysconfig/** directory, create this file. If such a file is available, run the following command to add a static route into this file:

```
any net 192.168.2.0 netmask 255.255.255.0 gw 192.168.1.34
```

After the configuration, save and exit the file. The following figure shows the modified file content.

```
[root@lsw-centos65-0001 sysconfig]# cat static-routes
any net 192.168.2.0 netmask 255.255.255.0 gw 192.168.1.34
```

3. Run the following command to restart the network service to make the static route take effect:

```
service network restart
```

```
[root@lsw-centos65-0001 sysconfig]# service network restart
Shutting down interface eth0: [OK]
Shutting down loopback interface: [OK]
Bringing up loopback interface: [OK]
Bringing up interface eth0:
Determining IP information for eth0... done. [OK]
```

4. Run the following command to view routes:

```
route -n
```

```
[root@lsw-centos65-0001 sysconfig]# route -n
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
169.254.169.254 192.168.1.1 255.255.255.255 UGH 0 0 0 eth0
192.168.2.0 192.168.1.34 255.255.255.0 UG 0 0 0 eth0
192.168.1.0 0.0.0.0 255.255.255.0 U 0 0 0 eth0
169.254.0.0 0.0.0.0 255.255.0.0 U 1002 0 0 eth0
0.0.0.0 192.168.1.1 0.0.0.0 UG 0 0 0 eth0
```

## 15.18.15 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.

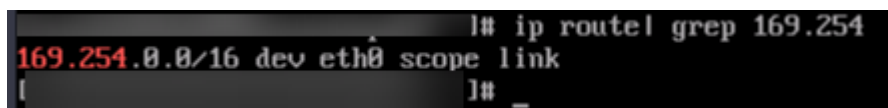
## Root Cause

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 15-214 Route information



```
]|# ip route | grep 169.254
169.254.0.0/16 dev eth0 scope link
]|# _
```

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 15-214](#). 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.

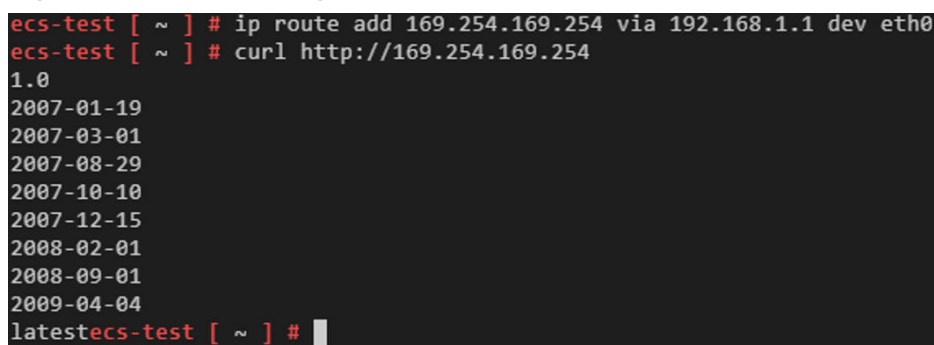
[How Do I View the Primary NIC?](#)

[How Do I View the Gateway Address?](#)

2. Run the following command to verify that the metadata can be obtained:

```
curl http://169.254.169.254
```

Figure 15-215 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

## How Do I View the Primary NIC?


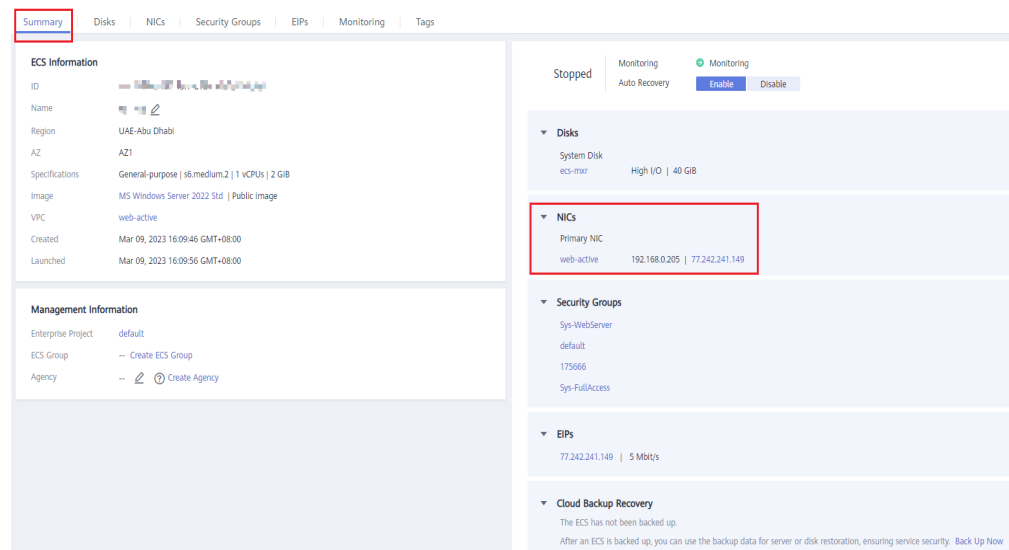
1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Computing**, click **Elastic Cloud Server**.
4. Click the name of the target ECS.  
The page providing details about the ECS is displayed.
5. Click the **Summary** tab to view details about the primary NIC.

Figure 15-216 Primary NIC details



## How Do I View the Gateway Address?


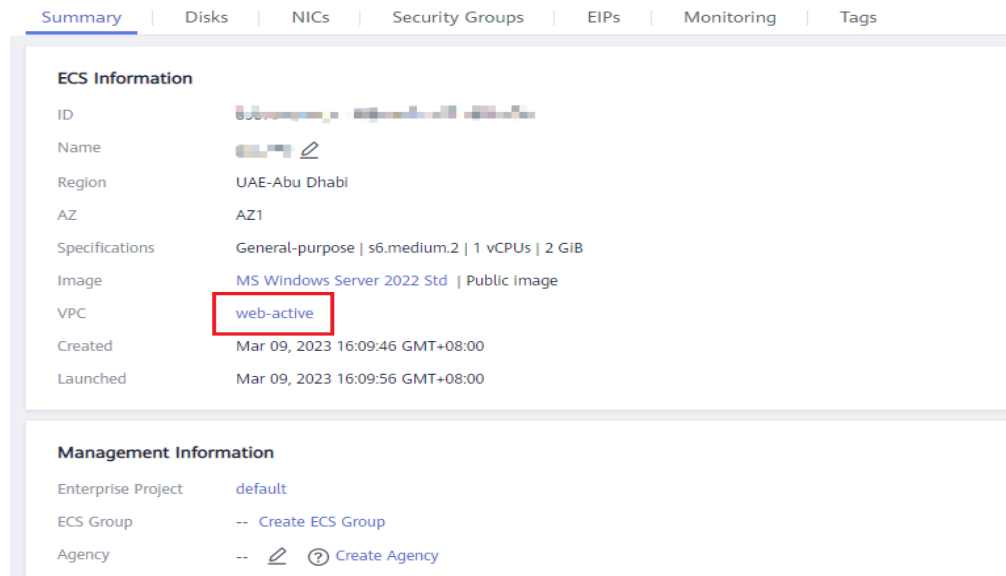

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Computing**, click **Elastic Cloud Server**.
4. Click the name of the target ECS.  
The page providing details about the ECS is displayed.
5. Click the VPC name to go to the VPC list page.

Figure 15-217 VPC name



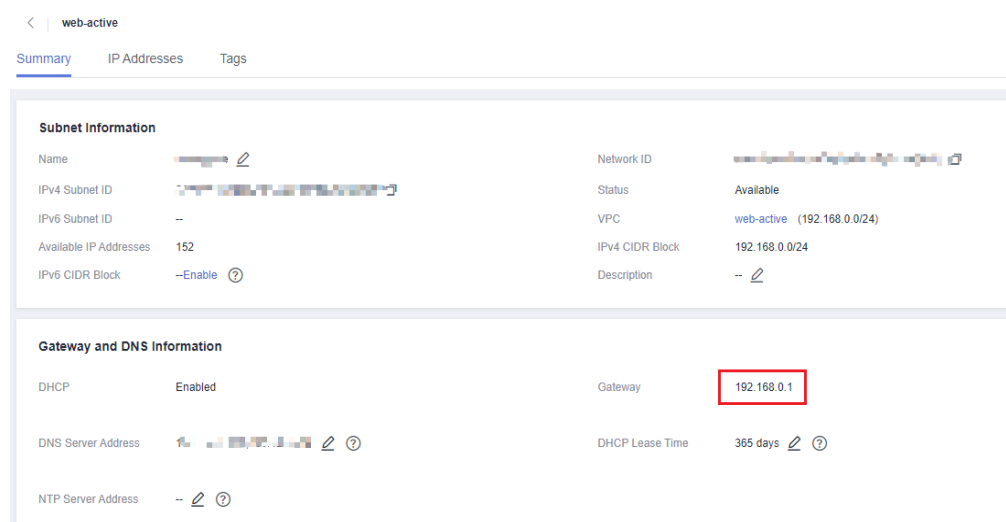
6. Locate the row that contains the target VPC and click the number in the **Subnets** column to go to the subnet list page.

Figure 15-218 Number in the Subnets column



7. Click the target subnet name to go to the subnet details page and view the gateway address.

Figure 15-219 Gateway address





## 15.18.16 Why Can't My Windows ECS Access the Internet?

### Symptom

Your attempt to access the Internet from your Windows ECS failed.

### 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.

**Table 15-23** Possible causes and solutions

| Possible Cause                                     | Solution                                                                                                                                                                                |
|----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The ECS is frozen or stopped, or has no EIP bound. | Check whether the ECS is in <b>Running</b> state and has an EIP bound. For details, see <a href="#">Checking the ECS Status</a> .                                                       |
| The ECS is overloaded.                             | Check whether the bandwidth and vCPU usage of the ECS are too high. For details, see <a href="#">Checking Whether the ECS Is Overloaded</a> .                                           |
| The EIP bandwidth exceeds the limit.               | Increase the bandwidth and try again. For details, see <a href="#">Checking Whether the EIP Bandwidth Exceeded the Limit</a> .                                                          |
| The access is blocked by the ISP.                  | Check whether you can access the ECS using another hotspot or network. For details, see <a href="#">Checking Whether the ISP Network Is Functional</a> .                                |
| The network configuration on the ECS is incorrect. | Check whether the NIC and DNS configurations are correct. For details, see <a href="#">Checking the NIC Configuration</a> .                                                             |
| Routing is incorrectly configured.                 | Check whether the default route of 0.0.0.0 designates to the default gateway. For details, see <a href="#">Checking Whether the Default Route Is Destined for the Default Gateway</a> . |
| The security group is incorrectly configured.      | Check whether the security group allows the network traffic in the outbound direction. For details, see <a href="#">Checking Whether the Security Group Is Correctly Configured</a> .   |
| A network ACL has been associated with the ECS.    | Disassociate the network ACL with the ECS and try again. For details, see <a href="#">Checking ACL Rules</a> .                                                                          |
| The EIP is blocked.                                | If the EIP is blocked, the ECS cannot access the Internet. For details, see <a href="#">Checking Whether the EIP Is Blocked</a> .                                                       |

| Possible Cause                                           | Solution                                                                                                                                                               |
|----------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The access is blocked by the firewall.                   | Disable the firewall and try again. For details, see <a href="#">Checking the Firewall Configuration</a> .                                                             |
| The gateway is inaccessible.                             | Run the <b>ping</b> command to check whether the DNS server is running properly. For details, see <a href="#">Checking Whether the Gateway Is Accessible</a> .         |
| The ECS performance cannot meet service requirements.    | Run the <b>netstat</b> command to check the network connection status. For details, see <a href="#">Checking the ECS Performance</a> .                                 |
| The access is blocked by third-party antivirus software. | Disable or uninstall the third-party antivirus software and try again. For details, see <a href="#">Checking Whether the Access Is Blocked by Antivirus Software</a> . |
| The ECS has been attacked by viruses or Trojan horses.   | Check whether the ECS is affected by viruses or Trojan horses. For details, see <a href="#">Checking the ECS Security Status</a> .                                     |

## Checking the ECS Status

- Check whether the ECS is in the **Running** state on the management console.
- Check whether an ECS has an EIP bound.  
An ECS can access the Internet only if it has an EIP bound.  
For details, see [Binding an EIP](#).

## Checking Whether the ECS Is Overloaded

If the bandwidth and CPU usage of an ECS are too high, the network may be disconnected.

If you have created an alarm rule in Cloud Eye, the system automatically sends an alarm notification to you when the bandwidth or CPU usage reaches the threshold specified in the rule.

To resolve this issue, perform the operations described in [Why Is My Windows ECS Running Slowly?](#)

## Checking Whether the EIP Bandwidth Exceeded the Limit

An ECS with an EIP bound accesses the Internet using the bandwidth configured for the EIP.

If Internet access fails, check whether the EIP bandwidth exceeds the limit.

## Checking Whether the ISP Network Is Functional

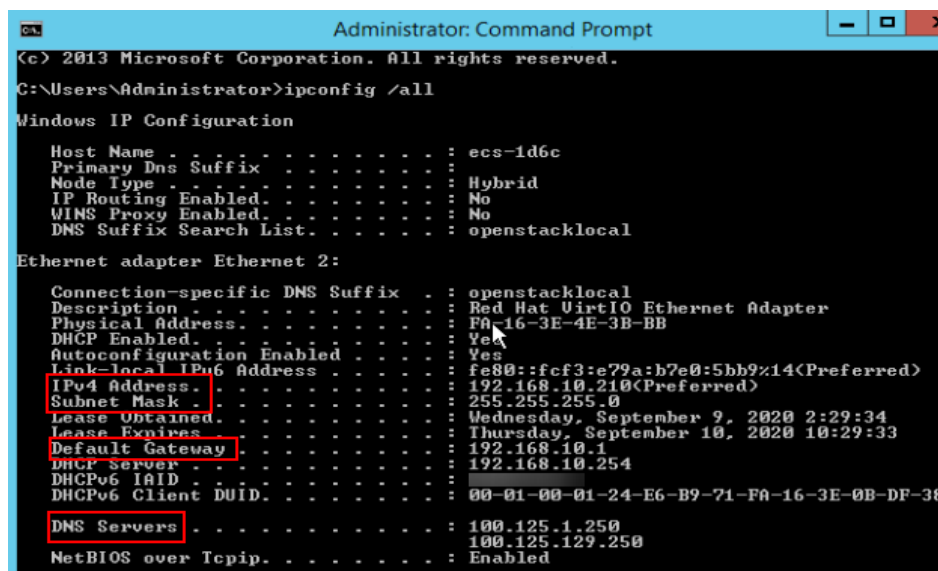
Check whether the fault occurs for a specific IP address. If so, the IP address may be blocked by the ISP.

Try another hotspot for access. If the access is successful, the fault may lie in the local carrier network. Contact the carrier to resolve this issue.

## Checking the NIC Configuration

- Check whether the NIC and DNS configurations on the ECS are consistent with those displayed on the ECS management console.
  - a. On the CLI of the ECS, run the **ipconfig /all** command to check whether the NIC and DNS configurations are correct, as shown in [Figure 15-220](#).

Figure 15-220 NIC and DNS configurations



```
Administrator: Command Prompt
(c) 2013 Microsoft Corporation. All rights reserved.
C:\Users\Administrator>ipconfig /all

Windows IP Configuration


Host Name : ecs-1d6c
Primary Dns Suffix :
Node Type : Hybrid
IP Routing Enabled. : No
WINS Proxy Enabled. : No
DNS Suffix Search List. : openstacklocal

Ethernet adapter Ethernet 2:

Connection-specific DNS Suffix . . . : openstacklocal
Description : Red Hat VirtIO Ethernet Adapter
Physical Address. : 8A-16-3E-4E-3B-BB
Dhcp Enabled. : Yes
Autoconfiguration Enabled : Yes
Link-local IPv6 Address : fe80::fcf3:e79a:b7e0:5bb9%14(Preferred)
IPv4 Address. : 192.168.10.210(Preferred)
Subnet Mask : 255.255.255.0
Lease Obtained. : Wednesday, September 9, 2020 2:29:34
Lease Expires : Thursday, September 10, 2020 10:29:33
Default Gateway : 192.168.10.1
Dhcp server : 192.168.10.254
Dhcpv6 IAID :
Dhcpv6 Client DUID. : 00-01-00-01-24-E6-B9-71-FA-16-3E-0B-DF-38
DNS Servers : 100.125.1.250
 100.125.129.250
NetBIOS over Tcpip. : Enabled
```

- b. Log in to the management console. On the ECS list page, click the name of the target ECS.
  - c. On the page providing details about the ECS, click the VPC name.
  - d. On the VPC list page, click the number displayed in the **Subnets** column.
  - e. On the subnet list page, click the name of the target subnet. The subnet details page is displayed .
- Open the **cmd** window, run the **ncpa.cpl** command to start Network and Sharing Center, and check whether the NIC is functional.

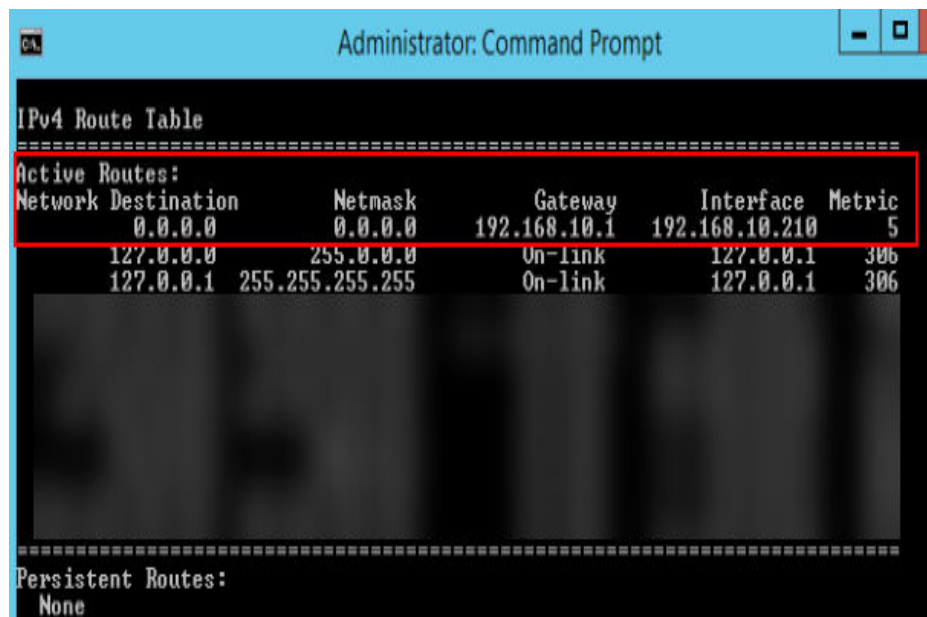
Figure 15-221 NIC status



### Checking Whether the Default Route Is Destined for the Default Gateway

Run the **route print** command to obtain the routing table of the ECS and check whether the default route of 0.0.0.0 is destined for the default gateway.

Figure 15-222 Default route settings



```
Administrator: Command Prompt

IPv4 Route Table
=====
Active Routes:
Network Destination Netmask Gateway Interface Metric
0.0.0.0 0.0.0.0 192.168.10.1 192.168.10.210 5
127.0.0.0 255.0.0.0 On-link 127.0.0.1 306
127.0.0.1 255.255.255.255 On-link 127.0.0.1 306

Persistent Routes:
None
```

## Checking Whether the Security Group Is Correctly Configured

Check whether the security group of the ECS is correctly configured. If an allowlist is configured for the outbound rules of the security group, the network traffic in the outbound direction is permitted.

## Checking ACL Rules

By default, no ACL rules are 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. Disassociate the network ACL from the subnet of the ECS.

On the page providing details about the network ACL, choose **Associated Subnets > Disassociate**.

### NOTE

The default network ACL rule denies all incoming and outgoing packets. If a network ACL is disabled, the default rule is still effective.

4. Try to access the Internet through the ECS again.

## Checking Whether the EIP Is Blocked

IP address blocking indicates that all traffic is destined to a null route. If the EIP is blocked, the ECS cannot access the Internet.

Generally, blocked EIPs will be automatically unblocked after 24 hours if no subsequent attack occurs.

## Checking the Firewall Configuration

Disable firewall rules for the ECS and check whether the Internet connection is restored.

If the connection is restored, check the firewall settings.


1. Log in to the Windows ECS.
2. Click the Windows icon in the lower left corner of the desktop and choose **Control Panel > System and Security > Windows Firewall**.

Figure 15-223 Windows Firewall



3. Choose **Check firewall status > Turn Windows Firewall on or off**. View and set the firewall status.

Figure 15-224 Turn off Windows Firewall



## Checking Whether the Gateway Is Accessible

1. Run the **ping** command to check whether data can be exchanged between the ECS and the gateway.  
Use an IP address in a different network segment to ping the gateway to check network connections.
2. Run the **ping** command to obtain the IP address of the DNS server.  
Compare the time required for pinging the DNS server and the time for pinging a specific IP address, and determine whether the DNS server is running properly.

## Checking the ECS Performance

Run the **netstat** command to check whether SYN-SENT, CLOSE\_WAIT, or FIN\_WAIT is found.

If any of them is found, port resources are used up. This issue is generally caused by a software bug. After the bug is fixed, restart the ECS.

Figure 15-225 Checking network connection



## Checking Whether the Access Is Blocked by Antivirus Software

Disable or uninstall the third-party antivirus software on the ECS, and check whether the fault is rectified.

## Checking the ECS Security Status

Check the ECS security status and determine whether the ECS is affected by viruses or Trojan horses.

## 15.18.17 Why Does My Linux ECS Fail to Access the Internet?

### Symptom

Your attempt to access the Internet from your Linux ECS failed.

### 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.

Table 15-24 Possible causes and solutions

| Possible Cause                                     | Solution                                                                                                                          |
|----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| The ECS is frozen or stopped, or has no EIP bound. | Check whether the ECS is in <b>Running</b> state and has an EIP bound. For details, see <a href="#">Checking the ECS Status</a> . |

| Possible Cause                                              | Solution                                                                                                                                                                                                 |
|-------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The ECS is overloaded.                                      | Check whether the bandwidth and vCPU usage of the ECS are too high. For details, see <a href="#">Checking Whether the ECS Is Overloaded</a> .                                                            |
| The EIP bandwidth exceeds the limit.                        | Increase the bandwidth and try again. For details, see <a href="#">Checking Whether the EIP Bandwidth Exceeded the Limit</a> .                                                                           |
| The DNS configuration is incorrect.                         | Change the DNS server to a private one. For details, see <a href="#">Checking the DNS Configuration</a> .                                                                                                |
| Specified resolution has been configured in the hosts file. | Check whether the mappings in the hosts configuration file are correct. For details, see <a href="#">Checking the hosts Configuration File</a> .                                                         |
| Both Network and NetworkManager are enabled.                | Use either of the two tools to prevent incompatibility issues. For details, see <a href="#">Checking Whether Both Network and NetworkManager Have Been Enabled</a> .                                     |
| The security group is incorrectly configured.               | Check whether the security group allows the network traffic in the outbound direction. For details, see <a href="#">Checking Whether the Security Group Is Correctly Configured</a> .                    |
| A network ACL has been associated with the ECS.             | Disassociate the network ACL with the ECS and try again. For details, see <a href="#">Checking ACL Rules</a> .                                                                                           |
| The EIP is blocked.                                         | If the EIP is blocked, the ECS cannot access the Internet. For details, see <a href="#">Checking Whether the EIP Is Blocked</a> .                                                                        |
| The private IP address is lost.                             | Check whether the <b>dhclient</b> process is running. If it is not running, the private IP address may be lost. For details, see <a href="#">Checking Whether a Private IP Address Can Be Obtained</a> . |
| NICs are incorrectly configured.                            | Check whether the NIC and DNS configurations are correct. For details, see <a href="#">Checking the NIC Configuration</a> .                                                                              |
| Firewall is enabled on the ECS.                             | Disable the firewall and try again. For details, see <a href="#">Checking the Firewall Configuration</a> .                                                                                               |

## Checking the ECS Status

- Check whether the ECS is in the **Running** state on the management console.
- Check whether an ECS has an EIP bound.  
An ECS can access the Internet only if it has an EIP bound.  
For details, see [Binding an EIP](#).



## Checking Whether the ECS Is Overloaded

If the bandwidth and CPU usage of an ECS are too high, the network may be disconnected.

If you have created an alarm rule in Cloud Eye, the system automatically sends an alarm notification to you when the bandwidth or CPU usage reaches the threshold specified in the rule.

To resolve this issue, perform the operations described in [Why Is My Linux ECS Running Slowly?](#)

## Checking Whether the EIP Bandwidth Exceeded the Limit

An ECS with an EIP bound accesses the Internet using the bandwidth configured for the EIP.

If Internet access fails, check whether the EIP bandwidth exceeds the limit.

## Checking the DNS Configuration

Private DNS servers resolve domain names for the ECSs created using a public image by default. The private DNS servers do not affect the domain name resolution for the ECSs to access the Internet. Additionally, you can use the private DNS servers to directly access the internal addresses of other cloud services, such as OBS. Compared with the access through the Internet, this access mode features high performance and low latency.

For Linux ECSs, run the following command to check the DNS configuration:

```
cat /etc/resolv.conf
```

If the command output shown in [Figure 15-226](#) is displayed, the domain name is resolved using the private DNS server.

**Figure 15-226** DNS configuration

```
[root@ecs-bae5 ~]# cat /etc/resolv.conf
; generated by /sbin/dhclient-script
search openstacklocal
options single-request-reopen
nameserver 100.125.135.29
nameserver 100.125.17.29
```

If the domain name of the ECS is resolved using a non-private DNS server and you want to switch to a private DNS server, change the DNS server to a private one.

For details, see [How Can I Configure the NTP and DNS Servers for an ECS?](#)

## Checking the hosts Configuration File

If the DNS configuration is correct but the ECS still cannot access the Internet, check whether the mapping information in the hosts configuration file is correct. In case of any incorrect mapping, comment them out.

For Linux, run the following command to view the hosts configuration:

```
vim /etc/hosts
```

If there is an incorrect domain name mapping, comment it out and save the hosts file.

## Checking Whether Both Network and NetworkManager Have Been Enabled

Network and NetworkManager are two network management tools, and either one of them can be enabled each time. If both of them are enabled, they are incompatible with each other.

Take CentOS 7 as an example. NetworkManager is recommended for CentOS 7.

1. Check the Network or NetworkManager running status.

```
systemctl status network
```

```
systemctl status NetworkManager
```

2. Run the following commands to disable Network:

```
systemctl stop network
```

```
systemctl disable network
```

3. Run the following commands to enable NetworkManager:

```
systemctl start NetworkManager
```

```
systemctl enable NetworkManager
```

## Checking Whether the Security Group Is Correctly Configured

Check whether the security group of the ECS is correctly configured. If an allowlist is configured for the outbound rules of the security group, the network traffic in the outbound direction is permitted.

## Checking ACL Rules

By default, no ACL rules are 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. Disassociate the network ACL from the subnet of the ECS.

On the page providing details about the network ACL, choose **Associated Subnets > Disassociate**.

 NOTE

The default network ACL rule denies all incoming and outgoing packets. If a network ACL is disabled, the default rule is still effective.

4. Try to access the Internet through the ECS again.

## Checking Whether the EIP Is Blocked

IP address blocking indicates that all traffic is destined to a null route. If the EIP is blocked, the ECS cannot access the Internet.

Generally, blocked EIPs will be automatically unblocked after 24 hours if no subsequent attack occurs.

## Checking Whether a Private IP Address Can Be Obtained

Private IP addresses may be lost if the dhclient process is not running or the target NIC is not managed by NetworkManager because NetworkManager automatic startup is not enabled. Perform the following operations to locate the fault:

Consider an ECS running CentOS 7 as an example.

1. Run the following command to check whether dhclient is running:

```
ps -ef |grep dhclient |grep -v grep
```

2. If dhclient is not detected, run the following command to check whether NetworkManager is running:

```
systemctl status NetworkManager
```

- If NetworkManager is in **Active: inactive (dead)** state, NetworkManager is not enabled. Run the following command to check whether NetworkManager is automatically started upon system startup:

```
systemctl is-enabled NetworkManager
```

If the command output is **disabled**, run the following command to enable NetworkManager automatic startup:

```
systemctl enable NetworkManager && systemctl start NetworkManager
```

- If NetworkManager is in **Active: active (running)** state, run the following command to check whether the target NIC is managed by NetworkManager:

```
nmcli device status
```

If the NIC is in **unmanaged** state, run the following command to enable it to be managed by NetworkManager:

```
nmcli device set eth0 managed yes
```

3. Run the following commands to restart NetworkManager:

```
systemctl restart NetworkManager
```

4. Run the following command to check whether the private IP address can be allocated:

```
ip add
```

## Checking the NIC Configuration

1. Run the following command to open the `/etc/sysconfig/network-scripts/ifcfg-eth0` file:

```
vi /etc/sysconfig/network-scripts/ifcfg-eth0
```

2. Modify the following configuration in this file.

Consider an ECS running CentOS 7 as an example.

```
DEVICE="eth0"
BOOTPROTO="dhcp"
ONBOOT="yes"
TYPE="Ethernet"
PERSISTENT_DHCLIENT="yes"
```

3. Run the following command to restart the network:

```
service network restart
```

## Checking the Firewall Configuration

Consider an ECS running CentOS 7 as an example. Check whether the firewall is enabled.

```
firewall-cmd --state
```

The command output is as follows:

```
[root@ecs-centos7 ~]# firewall-cmd --state
running
```

Run the following command to disable the firewall:

```
systemctl stop firewalld.service
```

---

### CAUTION

Enabling a firewall and configuring a security group protect your ECSs. If you disable a firewall, exercise caution when you enable ports in the security group.

---

## 15.18.18 How Do I Troubleshoot an Unresponsive Website Hosted on My ECS?

### Symptom

Websites running on an ECS might become unreachable for multiple reasons. Check whether the configurations of network, port, firewall, or security group of the ECS are correct.

### Fault Locating

If an error is displayed when you access a website, identify possible causes based on the error message.

**NOTE**

If the error message cannot help you locate the fault, record the resource details and fault occurred time, and contact customer service for technical support.

You can also locate the fault based on the following possible causes which are listed in order of their probability.

If the fault persists after you have ruled out one cause, move on to the next one.

**Table 15-25** Possible causes and solutions

| Possible Cause         | Solution                                                                                                                                                 |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port communication     | Check whether the web port used by the target website is properly listened to on the ECS. For details, see <a href="#">Checking Port Communication</a> . |
| Security group rules   | Check whether the access to the port is allowed in the security group of the ECS. For details, see <a href="#">Checking Security Group Rules</a> .       |
| Firewall configuration | Disable the firewall and try again. For details, see <a href="#">Checking the Firewall Configuration</a> .                                               |
| Route configuration    | Check whether the gateway configurations in the ECS route table are correct. For details, see <a href="#">Checking the ECS Route Configuration</a> .     |
| Local network          | Check whether you can use another hotspot or network to access the website. For details, see <a href="#">Checking the Local Network</a> .                |
| CPU usage              | Identify and optimize the processes leading to high vCPU usage. For details, see <a href="#">Checking the CPU usage</a> .                                |

## Checking Port Communication

Ensure that service processes and ports are in **LISTEN** state. [Table 15-26](#) lists the common TCP statuses.

- Linux  
Run the **netstat -antpu** command to check whether the port used by the target website is in **LISTEN** status,  
for example, **netstat -antpu |grep sshd**.

**Figure 15-227** Checking port listening status

```
[root@elb-mq02 ~]# netstat -antpu | grep sshd
tcp 0 0 0.0.0.0:22 0.0.0.0:* LISTEN 7178/sshd
```

- If the port status is **LISTEN**, go to [Checking Security Group Rules](#).
- If the port status is not **LISTEN**, check whether the web service process has been started and correctly configured.

- Windows
  - Perform the following operations to check port communication:
  - a. Run **cmd.exe**.
  - b. Run the **netstat -ano | findstr "Port number"** command to obtain the port number used by the process.  
For example, run **netstat -ano | findstr "80"**.

**Figure 15-228** Checking port listening status

```
C:\Users\Administrator>netstat -ano |findstr "80"
TCP 0.0.0.0:80 0.0.0.0:0 LISTENING 4
TCP 0.0.0.0:49155 0.0.0.0:0 LISTENING 880
TCP [::]:80 [::]:0 LISTENING 4
TCP [::]:49155 [::]:0 LISTENING 880
UDP 0.0.0.0:123 *:* 808
UDP [::]:123 *:* 808
```

- If the port is in **LISTENING** state, go to [Checking Security Group Rules](#).
- If the port is not in **LISTENING** state, check whether the web service process has been started and correctly configured.

**Table 15-26** Common TCP statuses

| TCP Status  | Description                                                                                          | Application Scenario                                                                                                                                                                                                                                                  |
|-------------|------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LISTEN      | Listens for network connection requests from a remote TCP port.                                      | The TCP server is running properly.                                                                                                                                                                                                                                   |
| ESTABLISHED | Indicates that a connection has been set up.                                                         | A TCP connection is properly set up.                                                                                                                                                                                                                                  |
| TIME-WAIT   | Waits until the remote TCP server receives the acknowledgment after sending a disconnection request. | The TCP connection is disconnected, and this state is cleared in 1 minute.                                                                                                                                                                                            |
| CLOSE-WAIT  | Waits for a disconnection request sent by a local user.                                              | An application program fault leads to an open socket. This state is displayed after the network is disconnected, indicating that a process is in an infinite loop or waiting for certain requirements to be met. To resolve this issue, restart the affected process. |
| FIN-WAIT-2  | Waits for the network disconnection request from a remote TCP server.                                | The network has been disconnected and requires 12 minutes to automatically recover.                                                                                                                                                                                   |

| TCP Status | Description                                                                                               | Application Scenario                                                                                                                                                                |
|------------|-----------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SYN-SENT   | Waits for the matched network connection request after a network connection request is sent.              | The TCP connection request failed, which is generally caused by the delayed handling of high CPU usage on the server or by a DDoS attack.                                           |
| FIN-WAIT-1 | Waits for the remote TCP disconnection request, or the acknowledgment for previous disconnection request. | If the network has been disconnected, this state may not automatically recover after 15 minutes. If the port has been used for a long period, restart the OS to resolve this issue. |

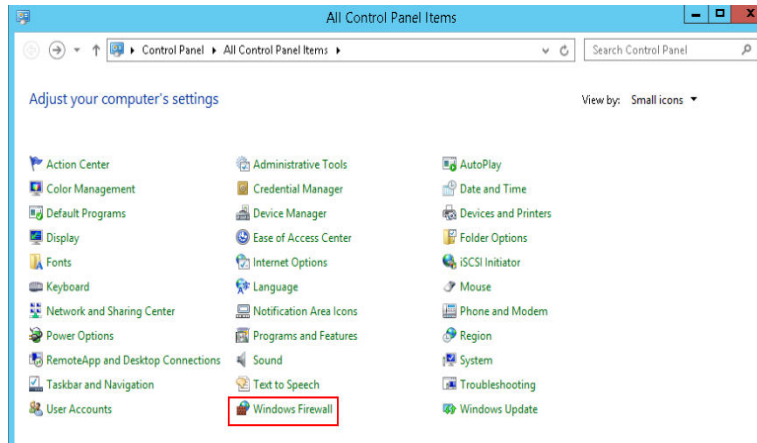
## Checking Security Group Rules

If the port used by the target website is denied in the security group, add a rule to the security group to allow the access of the port.


1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. In the ECS list, click the name of the target ECS.
4. On the **Security Groups** tab, view security group rules.
5. Click **Modify Security Group Rule**.
6. Configure the rule to allow the access of the port used by the website.

## Checking the Firewall Configuration

- Linux ECS  
The following uses port 80 and CentOS 6.8 as an example.
  - a. Run the **iptables -nvL --line-number** command to obtain firewall policies.
  - b. Run the following commands to allow access to port 80:  
**iptables -A INPUT -p tcp --dport 80 -j ACCEPT**  
**iptables -A OUTPUT -p tcp --sport 80 -j ACCEPT**
  - c. Run the **service iptables save** command to save the added rules.
  - d. Run the **service iptables restart** command to restart iptables.
  - e. Run the **iptables -nvL --line-number** command to check whether the added rules have taken effect.
  - f. Disable the firewall and try again.
- Windows ECS
  - a. Log in to the Windows ECS.
  - b. Click the Windows icon in the lower left corner of the desktop and choose **Control Panel > Windows Firewall**.



- c. Click **Turn Windows Firewall on or off**.  
View and set the firewall status.



- d. Disable the firewall and try again.

## Checking the ECS Route Configuration

- Linux ECS
  - a. Run the **route** command to check the routing policy. Ensure that the default route of 0.0.0.0 is destined for the gateway and that the IP address and the gateway are in the same network segment, as shown in the first and third lines in the following figure.

```
[root@ ~]# route
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
default gateway 0.0.0.0 UG 100 0 0 eth0
gateway 255.255.255.255 UGH 100 0 0 eth0
0.0.0.0 255.255.255.0 U 100 0 0 eth0
0.0.0.0 255.255.255.0 U 101 0 0 eth1
0.0.0.0 255.255.255.0 U 102 0 0 eth2
[root@ ~]#
```

- b. Run the **ifconfig** or **ip addr** command to obtain the ECS IP address.



Figure 15-229 ifconfig command output

```
[root@... ~]# ifconfig -a
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
 inet [REDACTED] netmask 255.255.255.0 broadcast 1[REDACTED]
 inet6 fe80::f816:3eff:fe24:1e7f prefixlen 64 scopeid 0x20<link>
 ether fa:16:3e:24:1e:7f txqueuelen 1000 (Ethernet)
 RX packets 227250083 bytes 21176207838 (19.7 GiB)
 RX errors 0 dropped 0 overruns 0 frame 0
 TX packets 149514101 bytes 276209392634 (257.2 GiB)
 TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
 inet 127.0.0.1 netmask 255.0.0.0
 inet6 ::1 prefixlen 128 scopeid 0x10<host>
 loop txqueuelen 1000 (Local Loopback)
 RX packets 14 bytes 1088 (1.0 KiB)
 RX errors 0 dropped 0 overruns 0 frame 0
 TX packets 14 bytes 1088 (1.0 KiB)
 TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

Figure 15-230 ip addr command output

```
[root@... ~]# ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
 link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
 inet 127.0.0.1/8 scope host lo
 valid_lft forever preferred_lft forever
 inet6 ::1/128 scope host
 valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default qlen 1000
 link/ether fa:16:3e:24:1e:7f brd ff:ff:ff:ff:ff:ff
 inet [REDACTED]/24 brd [REDACTED] scope global noprefixroute dynamic eth0
 valid_lft 77109sec preferred_lft 77109sec
 inet6 fe80::f816:3eff:fe24:1e7f/64 scope link
 valid_lft forever preferred_lft forever
```

- c. Run the **route -n** command to obtain the gateway in the routing table. The following is an example just for reference.

Figure 15-231 route -n command output

```
[root@... ~]# route -n
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
0.0.0.0 [REDACTED] 0.0.0.0 UG 100 0 0 eth0
1 [REDACTED] 255.255.255.255 UGH 100 0 0 eth0
1 [REDACTED] 0.0.0.0 255.255.255.0 U 100 0 0 eth0
```

- Windows ECS
  - a. Run **cmd.exe**.
  - b. Run the **ipconfig** command to obtain the ECS IP address.

Figure 15-232 ipconfig command output

```
Administrator: Command Prompt
Microsoft Windows [Version 10.0.14393]
(c) 2016 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>ipconfig

Windows IP Configuration

Ethernet adapter Ethernet 4:

 Connection-specific DNS Suffix . : openstacklocal
 Link-local IPv6 Address : [REDACTED]
 IPv4 Address. : [REDACTED]
 Subnet Mask : [REDACTED]
 Default Gateway : [REDACTED]
```

- c. Run the **route print** command to obtain the gateway in the routing table.

**Figure 15-233 route print command output**

```
Select Administrator: Command Prompt

C:\Users\Administrator>route print
=====
Interface List
10...fa 16 3e 90 4b b3Red Hat VirtIO Ethernet Adapter
1.....Software Loopback Interface 1
2...00 00 00 00 00 00 e0 Microsoft ISATAP Adapter
9...00 00 00 00 00 00 e0 Microsoft Teredo Tunneling Adapter
=====

IPv4 Route Table
=====
Active Routes:
Network Destination Netmask Gateway Interface Metric
=====
```

## Checking the Local Network

Try another hotspot or network for access.

If the access is successful, the fault may occur in the local carrier network. In such a case, rectify the local network fault and try again.

## Checking the CPU usage

If the bandwidth or vCPU usage of an ECS is too high, website access failures may occur. If you have in Cloud Eye, the system automatically sends an alarm to you when the bandwidth or CPU usage reaches the threshold specified in the rule.

1. Identify the processes leading to a high bandwidth or vCPU usage.
  - Windows  
Windows offers 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 analysis.
  - Linux  
Run the **top** command to check the OS running status.
2. Check whether the processes are malicious and handle the issue accordingly.
  - If the processes are normal, optimize them or modify ECS configurations by referring to .
  - If the processes are malicious, stop these processes manually or use a third-party tool to stop them automatically.

## 15.18.19 Why Did I See "Invalid argument" or "neighbour table overflow" During an Access to a Linux ECS?

### Symptom

1. When a Linux ECS sends a request to a server in the same subnet, the server has received the request but does not return a response. When the server pings the client, the message "sendmsg: Invalid argument" is displayed.  

```
64 bytes from 192.168.0.54: icmp_seq=120 ttl=64 time=0.064 ms
64 bytes from 192.168.0.54: icmp_seq=122 ttl=64 time=0.071 ms
ping: sendmsg: Invalid argument
ping: sendmsg: Invalid argument
ping: sendmsg: Invalid argument
```
2. "neighbor table overflow" is displayed in the `/var/log/messages` log file or the `dmesg` command output of a Linux ECS.  

```
[21208.317370] neighbour: ndisc_cache: neighbor table overflow!
[21208.317425] neighbour: ndisc_cache: neighbor table overflow!
[21208.317473] neighbour: ndisc_cache: neighbor table overflow!
[21208.317501] neighbour: ndisc_cache: neighbor table overflow!
```

### Root Cause

The Neighbour table references the ARP cache. When the Neighbour table overflows, the ARP table is full and will reject connections.

You can run the following command to check the maximum size of the ARP cache table:

```
cat /proc/sys/net/ipv4/neigh/default/gc_thresh3
```

Check the following parameters in the ARP cache table:

```
/proc/sys/net/ipv4/neigh/default/gc_thresh1
/proc/sys/net/ipv4/neigh/default/gc_thresh2
/proc/sys/net/ipv4/neigh/default/gc_thresh3
```

- `gc_thresh1`: The minimum number of entries to keep in the ARP cache. The garbage collector will not run if there are fewer than this number of entries in the cache.
- `gc_thresh2`: The soft maximum number of entries to keep in the ARP cache. The garbage collector will allow the number of entries to exceed this for 5 seconds before collection will be performed.
- `gc_thresh3`: The hard maximum number of entries to keep in the ARP cache. The garbage collector will always run if there are more than this number of entries in the cache.

To verify the actual number of IPv4 ARP entries, run the following command:

```
ip -4 neigh show nud all | wc -l
```

### Solution

1. Make sure that the number of servers in a subnet is less than the `default.gc_thresh3` value.
2. Adjust parameters: change `gc_thresh3` to a value much greater than the number of servers in the same VPC network segment, and make sure that the `gc_thresh3` value is greater than the `gc_thresh2` value, and the `gc_thresh2` value is greater than the `gc_thresh1` value.

For example, if a subnet has a 20-bit mask, the network can accommodate a maximum of 4,096 servers. The **default.gc\_thresh3** value of this network segment must be a value much greater than 4,096.

Temporary effective:

```
sysctl -w net.ipv4.neigh.default.gc_thresh1=2048
sysctl -w net.ipv4.neigh.default.gc_thresh2=4096
sysctl -w net.ipv4.neigh.default.gc_thresh3=8192
```

Always effective:

Add the following content to the `/etc/sysctl.conf` file:

```
net.ipv4.neigh.default.gc_thresh1 = 2048
net.ipv4.neigh.default.gc_thresh2 = 4096
net.ipv4.neigh.default.gc_thresh3 = 8192
```

Add IPv6 configuration if required:

```
net.ipv6.neigh.default.gc_thresh1 = 2048
net.ipv6.neigh.default.gc_thresh2 = 4096
net.ipv6.neigh.default.gc_thresh3 = 8192
```

## 15.18.20 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 15-234** 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 15-235** 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 [REDACTED] 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 15-236** Obtaining the MAC address of eth0

```
[root@ecs-683a ~]# ifconfig eth0 |grep "ether"
 ether fa:16:3e:c3:46:fc txqueuelen 1000 (Ethernet)
[root@ecs-683a ~]#
```

4. Obtain the returned MAC address.  
**ifconfig eth0 |grep "ether" |awk '{print \$2}'**

**Figure 15-237** Obtaining the MAC address of eth0

```
[root@ecs-683a ~]# ifconfig eth0 |grep "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.
```

## 15.18.21 How Can I Test the Network Performance of Linux ECSs?

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 15-27](#) and [Table 15-28](#) list the common netperf and iperf3 parameters.

**Table 15-27** 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 15-28** Common iperf3 parameters

| Parameter | Description              |
|-----------|--------------------------|
| -p        | Port number              |
| -c        | IP address of the RX end |
| -u        | UDP packets              |

| Parameter | Description                                                                                                                                                                         |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -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 15-29** Preparations

| Category      | Quantity | Image                          | Specifications       | IP Address                    |
|---------------|----------|--------------------------------|----------------------|-------------------------------|
| Tested ECS    | 1        | CentOS 7.4 64bit (recommended) | At least eight vCPUs | 192.168.2.10                  |
| 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 15-30** lists the procedures for installing these tools.

**Table 15-30** 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 15-27](#).

##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 15-238](#). The final result is the sum of the test results of the netperf processes on all TX ends.

**Figure 15-238** 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 15-28](#).

```
##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 15-28](#).

```
##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 15-28](#).

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 15-239** shows an example of the UDP PPS test result.

**Figure 15-239** 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

## 15.18.22 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.

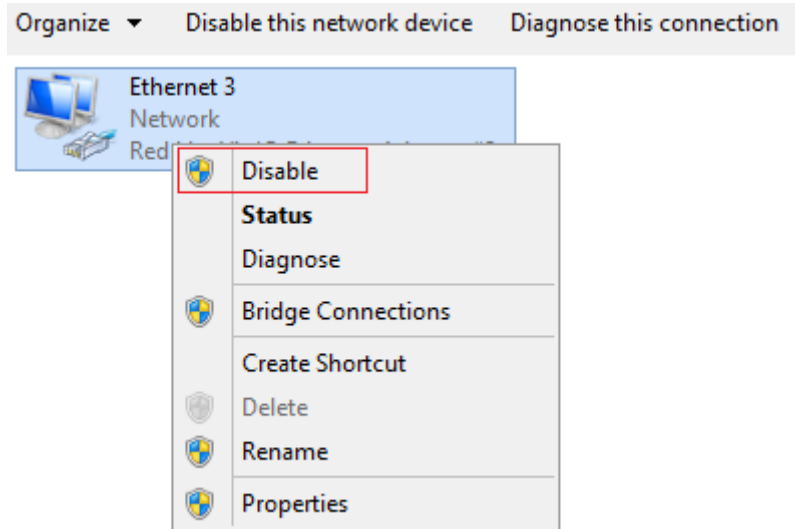
1. Log in to the ECS and run the following command:  
**ps -ef | grep dhclient**
2. If the dhclient process does not exist, restart the NIC or run any of the following commands to initiate a DHCP request:  
**dhclient eth0, ifdown eth0 + ifup eth0, or dhcpcd eth0**
3. If the DHCP client does not send 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.

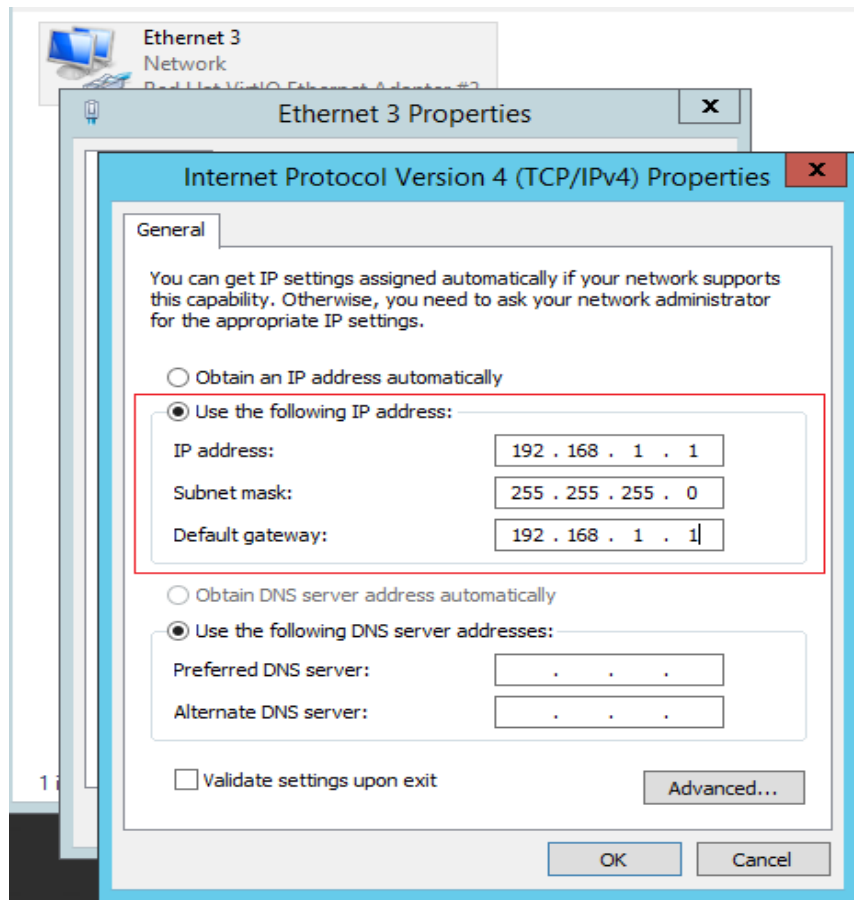
### Solution (Windows)

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.

## 15.18.23 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 15-31](#).
- Back up key ECS data before modifying kernel parameter settings.

### Background

**Table 15-31** 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
```

## 15.18.24 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 15-240 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 15-241** Port redirections on Linux

```
[root@linux-centos ~]# ssh -p 1080
root@: 's password:
```

Enter the password to log in to ECS 2 with hostname **ecs-inner**.

**Figure 15-242** Logging in to ECS 2

```
[root@ecs-inner ~]#
```

----End

## 15.18.25 Can the ECSs of Different Accounts Communicate over an Intranet?

No. The ECSs of different accounts cannot communicate with each other over an intranet.

To enable the communication over an intranet, use the methods provided in the following table.

| Scenario           | Billing        | Method                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| In the same region | Free of charge | Use VPC peering to enable the communication over an intranet. <ul style="list-style-type: none"><li>• <a href="#">VPC Peering Connection Creation Procedure</a></li><li>• <a href="#">Creating a VPC Peering Connection with a VPC in Another Account</a></li></ul>                                                                                                   |
| In the same region | Billed         | Use VPC Endpoint to enable the communication over an intranet. <ul style="list-style-type: none"><li>• <a href="#">What Is VPC Endpoint?</a></li><li>• <a href="#">Configuring a VPC Endpoint for Communications Across VPCs of Different Domains</a></li><li>• <a href="#">What Are the Differences Between VPC Endpoints and VPC Peering Connections?</a></li></ul> |

## 15.18.26 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.

## 15.19 Security Configurations



## 15.19.1 Are ECSs with Simple Passwords Easily Attacked?

It is recommended that your password contain 8 to 26 characters that consists of digits, uppercase and lowercase letters, and special characters.

If your ECS has been intruded, contact customer service for technical support.

**Table 15-32** Password complexity requirements

| Parameter | Requirement                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Password  | <ul style="list-style-type: none"><li>• Consists of 8 to 26 characters.</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> |

## 15.19.2 How Is ECS Security Ensured?

Host Security Service (HSS) helps you identify and manage the assets on your servers, eliminate risks, and defend against intrusions and web page tampering. There are also advanced protection and security operations functions available to help you easily detect and handle threats.

After installing the HSS agent on your ECSs, you will be able to check the ECS security status and risks in a region on the HSS console.


## 15.19.3 How Can I Disable Operation Protection?

### Symptom

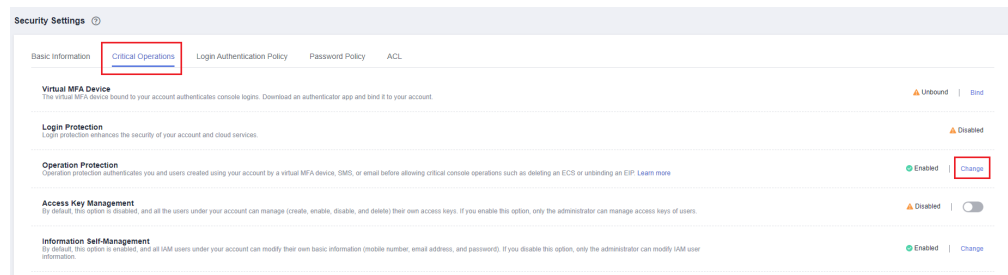
When I perform critical operations on my ECS with operation protection enabled, for example, deleting my ECS or modifying ECS specifications, I have to enter the password and verification code for authentication. To disable operation protection, perform the operations described in this section.

### Procedure

1. Log in to the management console.

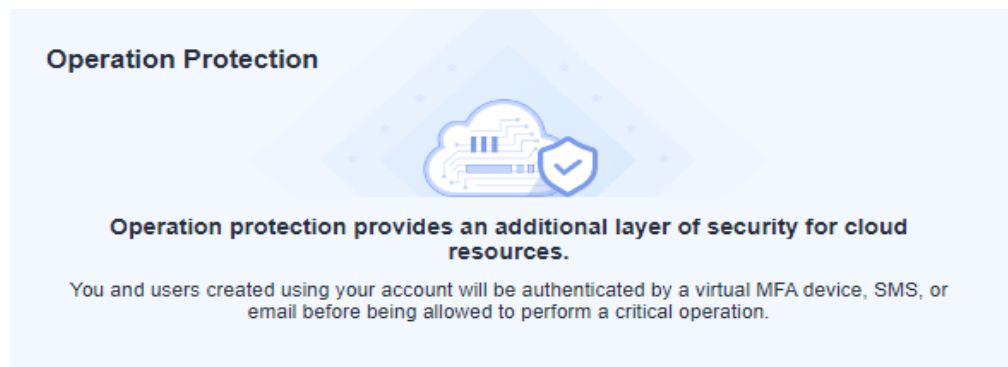
2. Click . Under **Management & Deployment**, choose **Identity and Access Management**.
3. In the left navigation pane of the IAM console, choose **Security Settings**.
4. On the **Security Settings** page, choose **Critical Operations > Operation Protection > Change**.

**Figure 15-243** Modifying operation protection settings



5. On the **Operation Protection** page, select **Disable** and click **OK**.

**Figure 15-244** Disabling Operation Protection



- Operation Protection
- Enable**  
You and users created using your account will need to perform identity verification by using the method you specify here.
  - Disable**  
Identity verification will not be required for performing a critical operation.

## 15.20 Resource Management and Tags

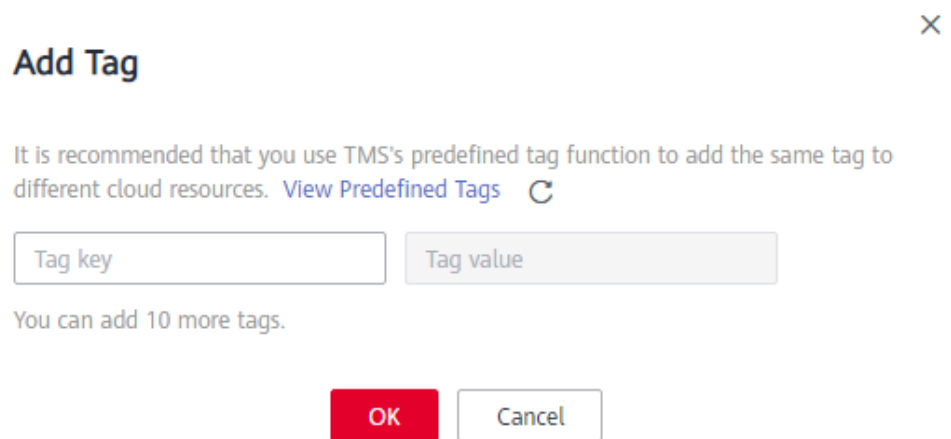
### 15.20.1 How Can I Create and Delete Tags and Search for ECSs by Tag?

#### Creating a Tag

1. Log in to the management console.

2. Select the region where the ECS is located.
3. Under **Computing**, click **Elastic Cloud Server**.
4. Click the name of the target ECS.  
The page providing details about the ECS is displayed.
5. Click **Tags** and then **Add Tag**.
6. Enter the tag key and value, and click **OK**.

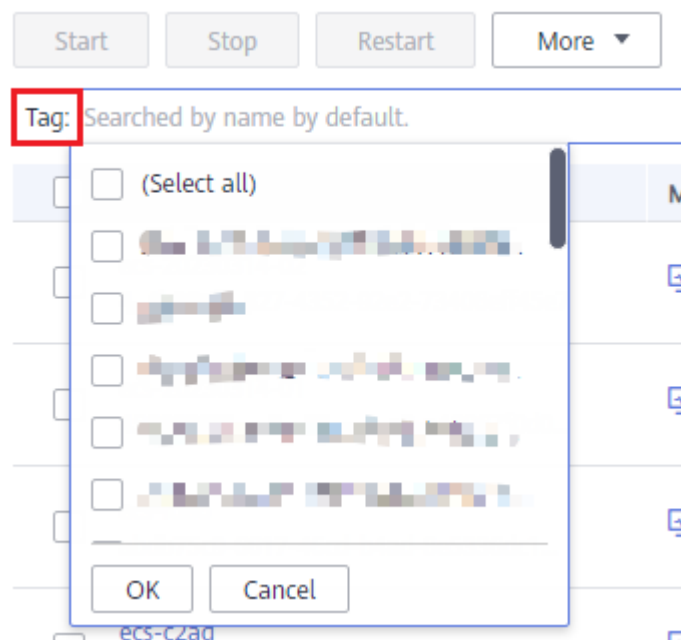
Figure 15-245 Adding tags



## Searching for ECSs by Tag

1. Log in to the management console.
2. Select the region where the ECS is located.
3. On the **Elastic Cloud Server** page, search for ECSs by tag.

Figure 15-246 Searching for ECSs by tag




4. In the search bar, choose **Tag** and then select the tag key and value, and click **OK**.

## Deleting a Tag

1. Log in to the management console.
2. Select the region where the ECS is located.
3. Click **Elastic Cloud Server**.
4. Click the name of the target ECS.
5. On the page providing details about the ECS, click **Tags**, locate the row containing the target tag, and click **Delete** in the **Operation** column.

Figure 15-247 Deleting a tag



## 15.21 Resource Monitoring

### 15.21.1 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 in 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 malicious 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 15-248 Starting the Resource Monitor



3. On the **Resource Monitor** page, click the **CPU** or **Network** tab to view the CPU or bandwidth usage.

Figure 15-249 Resource Monitor



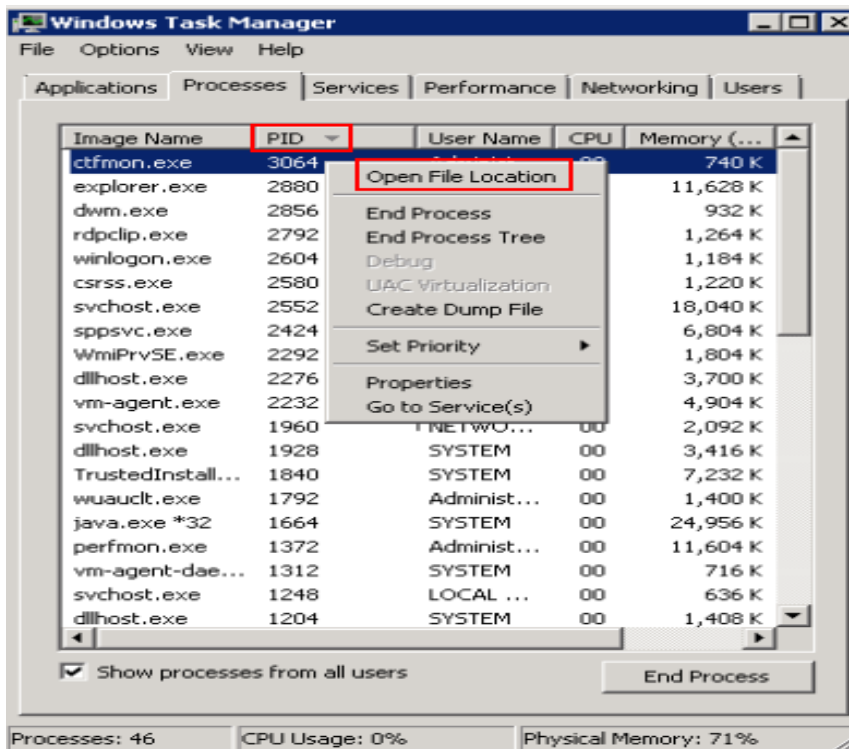
4. Obtain the IDs and names of the processes with high CPU or bandwidth usage.
5. On the remote login page, click **Ctrl+Alt+Del** to start the **Windows Task Manager**.

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 15-250 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 15-251 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     |
| luafv       | 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 15-252** 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.

## 15.21.2 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 in 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 malicious and handle the issue accordingly.
  - If the processes are normal, optimize them or modify ECS specifications.
  - If the processes are malicious, stop these processes manually or use a third-party tool to stop them automatically.

### 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 sysconfig]# 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:

- Press the **k** key during the execution of the top command.
- 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
```

- 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:

- Run the top command to check the resource usage of the **kswapd0** process.
- If the process remains in non-sleeping state for a long period, 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 1840 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
```

- 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 0.0.5
```

| PID   | USER  | PROGRAM                                | DEV  | SENT   | RECEIVED     |
|-------|-------|----------------------------------------|------|--------|--------------|
| 4596  | nginx | nginx: worker process                  | eth1 | 34.360 | 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  | 192.168.0.92:20101-185.176.26.66:43400 |      | 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.

## 15.22 Database Applications

### 15.22.1 Can a Database Be Deployed on an ECS?

Yes. You can deploy a database of any type on an ECS.

### 15.22.2 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.

### 15.22.3 What Should I Do If a Msg 823 Error Occurs in Oracle, MySQL, or SQL Server System Logs After a Disk Initialization Script Is Executed?

#### Symptom

After a disk is added to an ECS and the disk initialization script is automatically executed upon ECS startup, the Msg 823 error occurs in the database system logs of the Oracle, MySQL, and SQL Server databases.

#### Possible Causes

During the execution of the disk initialization script **WinVMDataDiskAutoInitialize.ps1**, diskpart is invoked to enable the virtual disk service. After the execution is complete, diskpart exits and the virtual disk service

is disabled. The automatic startup period of the built-in WinVMDaDataDiskAutoInitialize.ps1 overlaps the automatic startup period of the customer's database services, which may cause I/O operation errors.

The database uses Windows APIs (for example, ReadFile, WriteFile, ReadFileScatter, WriteFileGather) to perform file I/O operations. After performing these I/O operations, the database checks for any error conditions associated with these API calls. If the API calls fail with an operating system error, the database reports error 823. For details, see o obtain Microsoft official instructions, see [MSSQLSERVER error 823](#).

The 823 error message contains the following information:


- Whether the I/O operation is a read or write request
- The offset within the file where the I/O operation was attempted
- The database file against which the I/O operation was performed
- The operating system error code and error description in parentheses

The 823 error message usually indicates that there is a problem with underlying storage system or the hardware or a driver that is in the path of the I/O request. You can encounter this error when there are inconsistencies in the file system or if the database file is damaged.

## Solution

1. Log in to the ECS, open the **Run** dialog box, enter **services.msc**, and press **Enter**.
2. Search for the virtual disk service and ensure that it has been stopped.

**Figure 15-253** Checking the virtual disk status




If the virtual disk service is running, stop it in either of the following ways:

- On the **Services** page of the Windows operating system, right-click **Virtual Disk** and choose **Stop**.
- Open PowerShell and run the following command to stop the virtual disk service:

```
Get-Service -Name "vds" | Where {$_.status -eq 'Running'} | Stop-Service -Force
```

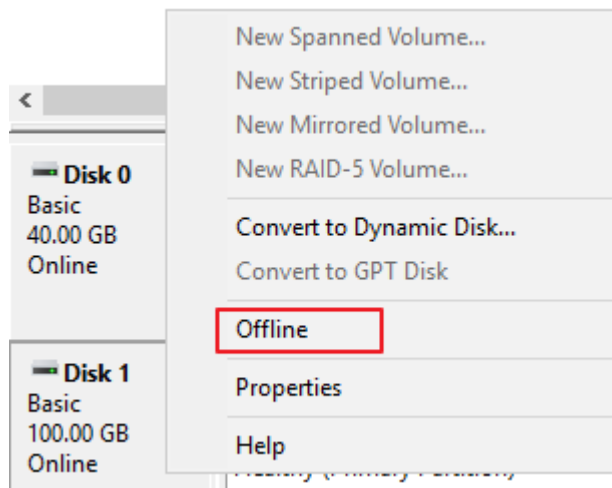
3. Disable the disk initialization script WinVMDaDataDiskAutoInitialize.ps1 from automatically initializing Windows data disks upon ECS startup.
  - a. Open the **Run** dialog box, enter **taskschd.msc**, and press **Enter**. The **Task Scheduler** window is displayed.
  - b. Open **Task Scheduler Library**, right-click **WinVMDaDataDiskInitialize** in the scheduled task list, and choose **End**.

Figure 15-254 Ending WinVMDaDataDiskInitialize



4. Restart the ECS or take the data disk offline and then online.
  - a. Open the **Run** dialog box, enter **diskmgmt.msc**, and press **Enter**. The **Disk Management** window is displayed.
  - b. Right-click the block to which the disk belongs and choose **Offline**.

Figure 15-255 Setting disk offline



- c. Right-click the block to which the disk belongs and choose **Online**.

**Figure 15-256** Setting disk online

