Elastic Cloud Server

User Guide

Issue 45

Date 2023-10-30





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

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

Trademarks and Permissions

HUAWEI and other Huawei trademarks are the property of Huawei Technologies Co., Ltd. All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

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

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

Huawei Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road

Qianzhong Avenue Gui'an New District Gui Zhou 550029

People's Republic of China

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

i

Contents

| 1 Service Overview | |
|---|----------|
| 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 | c |
| 1.6.1 ECS Overview | <u>c</u> |
| 1.6.2 ECS Lifecycle | <u>c</u> |
| 1.6.3 ECS Types | 10 |
| 1.7 ECS Specifications and Types | 14 |
| 1.7.1 A Summary List of ECS Specifications | 14 |
| 1.7.2 General-Purpose ECSs | 26 |
| 1.7.3 Dedicated General-Purpose ECSs | 29 |
| 1.7.4 Memory-optimized ECSs | 32 |
| 1.7.5 Disk-intensive ECSs | 34 |
| 1.7.6 Ultra-high I/O ECSs | 44 |
| 1.7.7 GPU-accelerated ECSs | 50 |
| 1.8 Discontinued ECS Specifications | 62 |
| 1.9 Images | 65 |
| 1.9.1 Image Types | 65 |
| 1.9.2 Cloud-Init | 66 |
| 1.10 EVS Disks | 68 |
| 1.11 Network | 69 |
| 1.12 Security | 71 |
| 1.12.1 Identity Authentication and Access Control | 71 |
| 1.12.1.1 Access Control for ECS | 71 |
| 1.12.2 Data Protection | 72 |
| 1.12.2.1 User Encryption | 72 |
| 1.12.3 Fault Recovery | 74 |
| 1.12.4 License Types | 75 |
| 1.13 Billing | 76 |
| 1.14 Permissions Management | 78 |

| 1.15 User Permissions | 83 |
|--|-----|
| 1.16 Region and AZ | 83 |
| 2 Getting Started | 85 |
| 2.1 Creating an ECS | 85 |
| 2.1.1 Overview | 85 |
| 2.1.2 Step 1: Configure Basic Settings | 85 |
| 2.1.3 Step 2: Configure Network | 89 |
| 2.1.4 Step 3: Configure Advanced Settings | 90 |
| 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) | 121 |
| 2.3.7 Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2012) | 129 |
| 2.3.8 Initializing a Linux Data Disk Larger Than 2 TiB (parted) | 137 |
| 3 Instances | 143 |
| 3.1 Viewing ECS Information | |
| 3.1.1 Viewing ECS Creation Statuses | |
| 3.1.2 Viewing Failed Tasks | 144 |
| 3.1.3 Viewing ECS Details (List View) | 145 |
| 3.1.4 Exporting ECS Information | 145 |
| 3.2 Logging In to a Windows ECS | 146 |
| 3.2.1 Login Overview | 146 |
| 3.2.2 Remotely Logging In to a Windows ECS (Using VNC) | 147 |
| 3.2.3 Remotely Logging In to a Windows ECS (Using MSTSC) | 154 |
| 3.2.4 Remotely Logging In to a Windows ECS (from a Linux Computer) | 161 |
| 3.3 Logging In to a Linux ECS | |
| 3.3.1 Login Overview | 163 |
| 3.3.2 Remotely Logging In to a Linux ECS (Using VNC) | 164 |
| 3.3.3 Remotely Logging In to a Linux ECS (Using an SSH Key Pair) | 172 |
| 3.4 Managing ECSs | 178 |
| 3.4.1 Changing ECS Names in a Batch | 178 |
| 3.4.2 Reinstalling the OS | 179 |
| 3.4.3 Changing the OS | 181 |
| 3.4.4 Managing ECS Groups | 185 |
| 3.4.5 Changing the Time Zone for an ECS | 187 |
| 3.4.6 Starting and Stopping ECSs | 189 |
| 3.5 Modifying ECS Specifications | 191 |

| 3.5.1 General Operations | 191 |
|---|-----|
| 3.5.2 Changing a Xen ECS to a KVM ECS (Windows) | 193 |
| 3.5.3 Automatically Changing a Xen ECS to a KVM ECS (Linux) | 198 |
| 3.5.4 Manually Changing a Xen ECS to a KVM ECS (Linux) | 203 |
| 3.6 Migrating an ECS | 207 |
| 3.7 Obtaining Metadata and Passing User Data | 208 |
| 3.7.1 Obtaining Metadata | 208 |
| 3.7.2 Passing User Data to ECSs | 217 |
| 3.8 (Optional) Configuring Mapping Between Hostnames and IP Addresses | 225 |
| 3.9 (Optional) Installing a Driver and Toolkit | 227 |
| 3.9.1 GPU Driver | 227 |
| 3.9.2 Installing a GRID Driver on a GPU-accelerated ECS | 228 |
| 3.9.3 Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS | 237 |
| 3.9.4 Obtaining a Tesla Driver and CUDA Toolkit | 251 |
| 3.9.5 Uninstalling a GPU Driver from a GPU-accelerated ECS | 253 |
| 4 Images | 259 |
| 4.1 Overview | 259 |
| 4.2 Creating an Image | 260 |
| 5 EVS Disks | 263 |
| 5.1 Overview | 263 |
| 5.2 Attaching an EVS Disk to an ECS | 263 |
| 5.3 Detaching an EVS Disk from a Running ECS | 265 |
| 5.4 Expanding the Capacity of an EVS Disk | 267 |
| 5.5 Expanding the Local Disks of a Disk-intensive ECS | 268 |
| 6 Backup Using CBR | 270 |
| 6.1 Overview | |
| 6.2 Backing Up an ECS Data | |
| 7 NICs | 279 |
| 7.1 Overview | |
| 7.2 Attaching a Network Interface | |
| 7.3 Deleting a Network Interface | |
| 7.4 Changing a VPC | |
| 7.5 Managing Virtual IP Addresses | |
| 7.6 Enabling NIC Multi-Queue | 289 |
| 7.7 Dynamically Assigning IPv6 Addresses | |
| 8 EIPs | 305 |
| 8.1 Overview | |
| 8.2 Binding an EIP | |
| 8.3 Unbinding an EIP | |
| 8.4 Changing an EIP | |

| 8.5 Changing an EIP Bandwidth | 309 |
|--|-----|
| 8.6 Enabling Internet Connectivity for an ECS Without an EIP | 309 |
| 9 Security | 313 |
| 9.1 Methods for Improving ECS Security | |
| 9.2 Security Groups | |
| 9.2.1 Overview | |
| 9.2.2 Default Security Group and Rules | 317 |
| 9.2.3 Security Group Configuration Examples | 318 |
| 9.2.4 Configuring Security Group Rules | 323 |
| 9.2.5 Changing a Security Group | 325 |
| 9.3 HSS | 326 |
| 10 Passwords and Key Pairs | 328 |
| 10.1 Passwords | 328 |
| 10.1.1 Application Scenarios for Using Passwords | 328 |
| 10.1.2 Changing the Login Password on an ECS | 329 |
| 10.1.3 Resetting the Password for Logging In to a Windows ECS | 331 |
| 10.1.4 Resetting the Password for Logging In to a Linux ECSECS | 333 |
| 10.2 Key Pairs | 336 |
| 10.2.1 Application Scenarios for Using Key Pairs | 336 |
| 10.2.2 (Recommended) Creating a Key Pair on the Management Console | 337 |
| 10.2.3 Creating a Key Pair Using PuTTYgen | 338 |
| 10.2.4 Importing a Key Pair | 342 |
| 10.2.5 Obtaining and Deleting the Password of a Windows ECS | 343 |
| 10.2.5.1 Obtaining the Password for Logging In to a Windows ECS | 343 |
| 10.2.5.2 Deleting the Initial Password for Logging In to a Windows ECS | 344 |
| 11 Permissions Management | 345 |
| 11.1 Creating a User and Granting ECS Permissions | 345 |
| 11.2 ECS Custom Policies | 346 |
| 12 Resources and Tags | 348 |
| 12.1 Tag Management | 348 |
| 12.1.1 Overview | 348 |
| 12.1.2 Adding Tags | 350 |
| 12.1.3 Searching for Resources by Tag | 352 |
| 12.1.4 Deleting a Tag | 352 |
| 12.2 Quota Adjustment | 354 |
| 13 Monitoring | 355 |
| 13.1 Monitoring ECSs | 355 |
| 13.2 Basic ECS Metrics | 356 |
| 13.3 OS Monitoring Metrics Supported by ECSs with the Agent Installed | 365 |
| 13.4 Process Monitoring Metrics Supported by ECSs with the Agent Installed | 407 |

| 13.5 Setting Alarm Rules | 412 |
|--|-----|
| 13.6 Viewing ECS Metrics | 413 |
| 14 CTS | 415 |
| 14.1 Key Operations Supported by CTS | 415 |
| 14.2 Viewing Audit Logs | 416 |
| 15 Troubleshooting | 418 |
| 15.1 What Can I Do If Switching from a Non-root User to User root Times Out? | 418 |
| 15.2 What Should I Do If Error "command 'gcc' failed with exit status 1" Occurs During PIP-based Software Installation | |
| 15.3 What Should I Do If Packages Are Downloaded Using PIP or wget at a Low Rate? | 420 |
| 15.4 How Can I Handle Slow ECS Startup? | 420 |
| 16 FAQs | 422 |
| 16.1 Common Topics | |
| 16.2 ECS Overview | 422 |
| 16.2.1 What Are the Precautions for Using ECSs? | 423 |
| 16.2.2 What Can I Do with ECSs? | 423 |
| 16.3 Regions and AZs | 423 |
| 16.3.1 What Is an AZ? | 423 |
| 16.3.2 What Is a Region? | 423 |
| 16.3.3 Are Products Different in Different Regions? | 423 |
| 16.3.4 Is Data Transmission Between AZs Billed? | 424 |
| 16.3.5 Can a Load Balancer Distribute Traffic to ECSs in Different Regions? | 424 |
| 16.3.6 Is Application Disaster Recovery Available in Different Regions? | 424 |
| 16.3.7 Are There Any Services Provided for Application Disaster Recovery? | 424 |
| 16.3.8 Can Components Contained in an Application Be Distributed to Different Regions? | 424 |
| 16.4 Billing | 424 |
| 16.4.1 What Are ECS Billing Modes? | 424 |
| 16.4.2 Will Am I Continue to Be Billed After ECSs Are Stopped? | 425 |
| 16.4.3 How Can I Stop an ECS from Being Billed? | |
| 16.5 Creation and Deletion | 426 |
| 16.5.1 What Should I Do If the ECS Resources to Be Purchased Are Sold Out? | 426 |
| 16.5.2 What Is the Creation Time and Startup Time of an ECS? | |
| 16.5.3 Why Does the Failures Area Show an ECS Creation Failure But the ECS List Displays the Creat ECS? | |
| 16.5.4 When Does an ECS Become Provisioned? | |
| 16.5.5 Why Cannot I View the ECSs Being Created Immediately After I Pay for Them? | |
| 16.5.6 Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? | |
| 16.5.7 What Do I Do If I Selected an Incorrect Image for My ECS? | |
| 16.5.8 How Quickly Can I Obtain an ECS? | |
| 16.5.9 How Can I Manage ECSs by Group? | |
| 16.5.10 Why Did I Fail to Configure an Anti-Affinity ECS Group? | |
| 16.5.11 What Happens After I Click the Delete Button? | |

| 16.5.12 Can a Deleted ECS Be Provisioned Again? | 429 |
|---|-------|
| 16.5.13 Can a Deleted ECS Be Restored? | 429 |
| 16.5.14 How Do I Delete or Restart an ECS? | 429 |
| 16.5.15 Can I Forcibly Restart or Stop an ECS? | 429 |
| 16.6 Login and Connection | 430 |
| 16.6.1 What Are the Login Requirements for ECSs? | .430 |
| 16.6.2 What Are the Username and Password for Remote Logins? | . 432 |
| 16.6.3 What Should I Do If Starting an ECS Remains in "Waiting for cloudResetPwdAgent" State? | .432 |
| 16.6.4 Why Can't I Log In to My Windows ECS? | .433 |
| 16.6.5 Why Can't I Log In to My Linux ECS? | 440 |
| 16.6.6 How Can I Change a Remote Login Port? | |
| 16.6.7 Why Cannot I Use a Non-Default SSH Port to Log In to My Linux ECS?ECS | |
| 16.6.8 Why Can't I Obtain the Password for Logging In to My Windows ECS Authenticated Using a Key Pair? | - |
| 16.6.9 What Browser Version Is Required to Remotely Log In to an ECS? | .453 |
| 16.6.10 How Can I Log In to an ECS After Its System Disk Is Exchanged with That Attached to Another ECS Running the Same OS? | |
| 16.6.11 Why Does the System Display a Message Indicating that the Password for Logging In to a Windows ECS Cannot Be Obtained? | . 455 |
| 16.6.12 Why Are Garbled Characters Displayed When I Log In to My ECS Using VNC? | .455 |
| 16.6.13 Why Are Characters Entered Through VNC Still Incorrect After the Keyboard Language Is Switched? | .456 |
| 16.6.14 Why Cannot I Use the French Keyboard to Enter Characters When I Log In to an ECS Using VN | |
| 16.6.15 What Should I Do If the Page Does not Respond After I Log In to an ECS Using VNC and Do N Perform Any Operation for a Long Period of Time? | lot |
| 16.6.16 What Should I Do If I Cannot View Data After Logging In to an ECS Using VNC? | . 457 |
| 16.6.17 Why Does a Blank Screen Appear After I Attempted to Log In to an ECS Using VNC? | .457 |
| 16.6.18 What Should I Do If Error Code 1006 or 1000 Is Displayed When I Log In to an ECS Through the Management Console? | |
| 16.6.19 Why No Audio File Can Be Properly Played on My Windows ECS Logged In Using VNC? | .458 |
| 16.6.20 How Can I Change the Resolution of a Windows ECS? | .460 |
| 16.6.21 Why Does an Authentication Failure Occurs After I Attempt to Remotely Log In to a Windows ECS? | |
| 16.6.22 Why Can't I Use the Local Computer to Connect to My Windows ECS? | .464 |
| 16.6.23 How Can I Obtain the Permission to Remotely Log In to a Windows ECS? | .469 |
| 16.6.24 Why Does the System Display No Remote Desktop License Servers Available to Provide a Licen When I Log In to a Windows ECS? | |
| 16.6.25 Why Does the System Display Error Code 0x112f When I Log In to a Windows ECS? | 473 |
| 16.6.26 Why Does the System Display Error Code 0x1104 When I Log In to a Windows ECS? | .474 |
| 16.6.27 Why Does the System Display Error Code 122.112 When I Log In to a Windows ECS? | . 478 |
| 16.6.28 Why Does the System Display Invalid Certificate or Associated Chain When I Log In to a Wind ECS from a Mac? | |
| 16.6.29 Why Does the System Display a Message Indicating Invalid Credentials When I Attempt to Acc | cess |
| a Windows ECS? | . 484 |

| 16.6.30 Why Does an Internal Error Occur When I Log In to My Windows ECS? | 489 |
|---|------------|
| 16.6.31 Why Is My Remote Session Interrupted by a Protocol Error? | 490 |
| 16.6.32 Why Am I Seeing an Error Message That Says Identity of Remote Computer Cannot be Verified When I Log In to a Windows ECS? | |
| 16.6.33 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? | |
| 16.6.34 Why Am I Seeing an Error Message That Says User Account is not Authorized for Remote Logi When I Log In to a Windows ECS? | |
| 16.6.35 Why Does My Remote Desktop Session End Because Another User Logs In When I Log In to a Windows ECS? | 497 |
| 16.6.36 Why Does an ECS Fail to Be Remotely Connected Using RDP and Internal Error Code 4 Is Displayed? | 500 |
| 16.6.37 Why Am I Seeing the Error Message "Module is unknown" When I Remotely Log In to a Linux ECS? | |
| 16.6.38 What Should I Do If Error Message "Permission denied" Is Displayed When I Remotely Log In t Linux ECS? | to a |
| 16.6.39 What Should I Do If Error Message "read: Connection reset by peer" Is Displayed When I Remotely Log In to a Linux ECS? | |
| 16.6.40 Why Am I Seeing the Error Message "Access denied" When I Remotely Log In to a Linux ECS?. | |
| 16.6.41 What Should I Do If Error Message "Disconnected: No supported authentication methods available" Is Displayed When I Remotely Log In to a Linux ECS? | |
| 16.7 How Do I Handle Error Messages Displayed on the Management Console? | |
| 16.8 How Can I Recover a Windows ECS with an Abnormal Virtualization Driver? | |
| 16.9 Why Is My Windows ECS Muted? | |
| 16.10 How Do I Change an ECS SID? | 516 |
| 16.11 Why Does a Pay-per-Use ECS Fail to Be Started? | |
| 16.12 Why Is the Memory of an ECS Obtained by Running the free Command Inconsistent with the | |
| Actual Memory? | |
| 16.13 ECS Management | |
| 16.13.1 How Can a Changed Static Hostname Take Effect Permanently? | |
| 16.13.2 Is an ECS Hostname with Suffix .novalocal Normal? | |
| 16.13.3 Why Is the Hostname of My ECS Restored to the Original Name After the ECS Is Restarted? | |
| 16.13.4 How Can I Set Sequential ECS Names When Creating Multiple ECSs? | |
| 16.13.5 How Can I Modify ECS Specifications? | |
| 16.13.6 What Should I Do If Executing a Driver Installation Script Failed on an ECS Running CentOS 5? | 526 |
| 16.13.7 What Should I Do If Executing a Driver Installation Script Failed When I Attempted to Modify t Specifications of a Linux ECS? | the 527 |
| 16.13.8 Why Do the Disks of a Windows ECS Go Offline After I Modify the ECS Specifications? | 528 |
| 16.13.9 Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications? | 531 |
| 16.14 OS Management | 533 |
| 16.14.1 Does OS Change Incur Fees? | 533 |
| 16.14.2 Can I Install or Upgrade the OS of an ECS? | 533 |
| 16.14.3 Can I Change the OS of an ECS? | 533 |
| 16.14.4 How Long Does It Take to Change an ECS OS? | 533 |

| 16.14.5 Will I Lose My Disk Data If I Reinstall ECS OS, Change the OS, or Change the ECS Specification | |
|---|---------|
| 16.14.6 Does OS Reinstallation Incur Fees? | |
| 16.14.7 Can I Select Another OS During ECS OS Reinstallation? | 535 |
| 16.14.8 How Long Does It Take to Reinstall an ECS OS? | 535 |
| 16.14.9 Do ECSs Support GUI? | 536 |
| 16.14.10 How Can I Install a GUI on an ECS Running CentOS 6? | 536 |
| 16.14.11 How Can I Install a GUI on an ECS Running CentOS 7? | |
| 16.14.12 How Can I Install a GUI on an ECS Running Ubuntu? | 537 |
| 16.14.13 How Can I Install a GUI on an ECS Running Debian? | 543 |
| 16.14.14 Why Does the OS Fail to Respond When kdump Occurs on a Linux ECS? | |
| 16.14.15 How Can I Upgrade the Kernel of a Linux ECS? | |
| 16.14.16 Why Cannot My ECS OS Start Properly? | 548 |
| 16.14.17 How Can I Enable SELinux on an ECS Running CentOS? | 548 |
| 16.14.18 Why Does a Forcibly-Stopped Linux ECS Fail to Be Restarted? | 549 |
| 16.14.19 How Do I View the GPU Usage of a GPU-accelerated ECS? | |
| 16.15 File Upload/Data Transfer | |
| 16.15.1 How Do I Upload Files to My ECS? | 552 |
| 16.15.2 How Can I Transfer Files from a Local Windows Computer to a Windows ECS? | 553 |
| 16.15.3 How Can I Use WinSCP to Transfer Files from a Local Windows Computer to a Linux ECS? | 556 |
| 16.15.4 How Can I Transfer Files from a Local Mac to a Windows ECS? | 558 |
| 16.15.5 How Can I Use SCP to Transfer Files Between a Local Linux Computer and a Linux ECS? | |
| 16.15.6 How Can I Use SFTP to Transfer Files Between a Local Linux Computer and a Linux ECS? | 562 |
| 16.15.7 How Can I Use FTP to Transfer Files from a Local Windows Computer to a Windows or Linu ECS? | |
| 16.15.8 How Can I Use FTP to Transfer Files Between a Local Linux Computer and a Linux ECS? | 565 |
| 16.15.9 How Can I Transfer Data Between a Local Computer and a Windows ECS? | 566 |
| 16.15.10 What Should I Do If the Connection Between the Client and the Server Times Out When I Upload a File Using FTP? | |
| 16.15.11 What Should I Do If Writing Data Failed When I Upload a File Using FTP? | |
| 16.15.12 Why Does Internet Access to an ECS Deployed with FTP Fail? | |
| 16.15.13 Why Am I Seeing an FTP Folder Error When I Open a Folder on an FTP Server? | |
| 16.15.14 Why Do I Fail to Connect to a Linux ECS Using WinSCP? | |
| 16.16 ECS Migration | |
| 16.16.1 Can I Migrate an ECS to Another Region or Account? | 577 |
| 16.17 Disk Management | 577 |
| 16.17.1 Why Can't I Find My Newly Purchased Data Disk After I Log In to My Windows ECS? | 577 |
| 16.17.2 How Can I Adjust System Disk Partitions? | 578 |
| 16.17.3 How Can I Obtain the Mapping Between Disk Partitions and Disk Devices on a Windows Ed | CS? 584 |
| 16.17.4 How Can I Obtain the Mapping Between Disk Partitions and Disk Devices on a Linux ECS? | |
| 16.17.5 How Can I Enable Virtual Memory on a Windows ECS? | 589 |
| 16.17.6 How Can I Add the Empty Partition of an Expanded System Disk to the End Root Partition Online? | 591 |

| 16.17.7 How Can I Add the Empty Partition of an Expanded System Disk to the Non-end Root Partition Online? | |
|---|-------|
| 16.17.8 Can I Attach Multiple Disks to an ECS? | |
| 16.17.9 What Are the Requirements for Attaching an EVS Disk to an ECS? | . 597 |
| 16.17.10 Which ECSs Can Be Attached with SCSI EVS Disks? | |
| 16.17.11 How Do I Obtain My Disk Device Name in the ECS OS Using the Device Identifier Provided o | |
| 16.17.12 Why Does a Linux ECS with a SCSI Disk Attached Fails to Be Restarted? | |
| 16.17.13 How Can I Check Whether the ECSs Attached with the Same Shared SCSI Disk Are in the Sar | ne |
| 16.17.14 Can All Users Use the Encryption Feature? | |
| 16.17.15 How Can I Add an ECS with Local Disks Attached to an ECS Group? | |
| 16.17.16 Why Does a Disk Attached to a Windows ECS Go Offline? | |
| 16.17.17 Why Does the Disk Drive Letter Change After the ECS Is Restarted? | |
| 16.17.18 How Can I Obtain Data Disk Information If Tools Are Uninstalled? | |
| 16.17.19 How Can I Rectify the Fault That May Occur on a Linux ECS with an NVMe SSD Disk Attache | ed? |
| 16.17.20 Why Is the Device Name of My C6 ECS in the sd* Format? | |
| 16.17.21 Why Are Disk Error Logs Printed After a Disk Attached to an ECS Is Formatted with the ext4 System? | File |
| 16.18 Passwords and Key Pairs | |
| 16.18.1 How Can I Change the Password for Logging In to a Linux ECS? ECS | |
| 16.18.2 How Can I Set the Validity Period of the Image Password? | |
| 16.18.3 Changing the Login Password on an ECS | |
| 16.18.4 Resetting the Password for Logging In to a Windows ECS | |
| 16.18.5 Resetting the Password for Logging In to a Linux ECS | |
| 16.18.6 What Should I Do If the System Displays a Message Indicating that the Password Is Incorrect When I Remotely Log In to My ECS? | |
| 16.18.7 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? | |
| 16.18.8 Disabling SELinux | |
| 16.18.9 How Can I Obtain the Key Pair Used by My ECS? | .621 |
| 16.18.10 How Can I Use a Key Pair? | . 621 |
| 16.18.11 What Should I Do If a Key Pair Cannot Be Imported? | 623 |
| 16.18.12 Why Does the Login to My Linux ECS Using a Key File Fail? | .623 |
| 16.18.13 What Should I Do If I Cannot Download a Key Pair? | . 624 |
| 16.18.14 Why Does a Key Pair Created Using puttygen.exe Fail to Be Imported on the Management Console? | . 625 |
| 16.18.15 What Is the cloudbase-init Account in Windows ECSs Used for? | |
| 16.18.16 What Should I Do If Cloud-Init Does Not Work After Python Is Upgraded? | |
| 16.19 Network Configurations | |
| 16.19.1 Can Multiple EIPs Be Bound to an ECS? | |
| 16.19.2 Can an ECS Without an EIP Bound Access the Internet? | |
| 16.19.3 Why Can't an EIP Be Pinged? | |

| 16.19.4 Why Can I Remotely Access an ECS But Cannot Ping It? | 636 |
|--|-------|
| 16.19.5 How Do I Query the Egress Public IP Address of My ECS? | 636 |
| 16.19.6 How Can I Configure the NTP and DNS Servers for an ECS? | 637 |
| 16.19.7 Configuring DNS | 642 |
| 16.19.8 What Should I Do If NIC Flapping Occurs After My ECS Specifications Are Modified? | 645 |
| 16.19.9 Will NICs Added to an ECS Start Automatically? | 647 |
| 16.19.10 Why Is the NIC Not Working? | 647 |
| 16.19.11 How Can I Handle the Issue that a Windows 7 ECS Equipped with an Intel 82599 NIC Referror in SR-IOV Scenarios? | |
| 16.19.12 How Can I Add a Static Route to a CentOS 6.5 OS? | |
| 16.19.13 Why Can't My Linux ECS Obtain Metadata? | 651 |
| 16.19.14 Why Can't My Windows ECS Access the Internet? | |
| 16.19.15 Why Does My Linux ECS Fail to Access the Internet? | |
| 16.19.16 How Do I Troubleshoot an Unresponsive Website Hosted on My ECS? | |
| 16.19.17 Why Did I See "Invalid argument" or "neighbour table overflow" During an Access to a L | .inux |
| 16.19.18 How Can I Obtain the MAC Address of My ECS? | |
| 16.19.19 How Can I Test Network Performance? | |
| 16.19.20 Why Can't I Use DHCP to Obtain a Private IP Address? | 683 |
| 16.19.21 How Can I View and Modify Kernel Parameters of a Linux ECS? | |
| 16.19.22 How Can I Configure Port Redirection? | |
| 16.19.23 Can the ECSs of Different Accounts Communicate over an Intranet? | |
| 16.19.24 Will ECSs That I Purchased Deployed in the Same Subnet? | 692 |
| 16.20 Security Configurations | 692 |
| 16.20.1 How Is ECS Security Ensured? | 693 |
| 16.21 Resource Management and Tags | 693 |
| 16.21.1 How Can I Create and Delete Tags and Search for ECSs by Tag? | 693 |
| 16.22 Resource Monitoring | 694 |
| 16.22.1 Why Is My Windows ECS Running Slowly? | 694 |
| 16.22.2 Why Is My Linux ECS Running Slowly? | 698 |
| 16.23 Database Applications | 702 |
| 16.23.1 Can a Database Be Deployed on an ECS? | 703 |
| 16.23.2 Does an ECS Support Oracle Databases? | 703 |
| 16.23.3 What Should I Do If a Msg 823 Error Occurs in Oracle, MySQL, or SQL Server System Logs Disk Initialization Script Is Executed? | |
| A Change History | 706 |

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 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, common 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.
- Volume Backup Service (VBS) allows you to create data backups for EVS disks and use the backups to restore the EVS disks. This maximizes user data correctness and security.
- Backup protection: You can back up all EVS disks (including the system disk and data disks) attached to an ECS and use the backup to restore the ECS data.

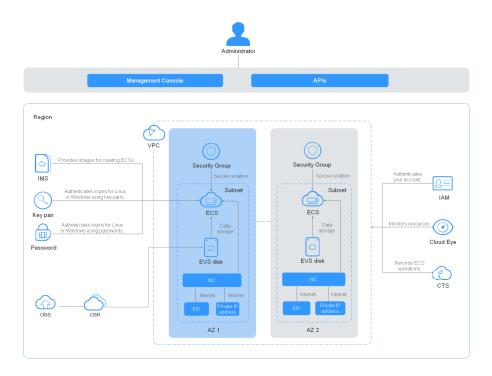


Figure 1-1 System architecture

Access Methods

The cloud platform provides a web-based management console. Use the management console to access ECSs. If you have registered on the cloud platform, log in to the management console 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**.

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 ecommerce websites.
- Do not use ECSs to initiate network attacks, such as DDoS attacks, CC attacks, web attacks, brute force cracking, or to spread viruses and Trojan horses.
- Do not use ECSs for traffic transit.
- Do not use ECSs for web crawling.
- Do not use ECSs to detect other systems like scanning or penetration unless otherwise being authorized.
- Do not deploy any illegal websites or applications on ECSs.
- Do not use ECSs to send spams.

Restricted Operations on ECSs

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

Precautions for Using Windows ECSs

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

Precautions for Using Linux ECSs

- Do not modify the /etc/issue file. Otherwise, the OS distribution will not be identified.
- Do not delete system directories or files. Otherwise, the system may fail to start or run properly.
- Do not change the permissions for or names of system directories. Otherwise, the system may fail to start or run properly.
- Do not upgrade the kernel of the Linux unless necessary.
 When upgrading a 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 IP addresses, subnet mask, and gateway address of an ECS. Otherwise, network exceptions may occur.
- Manually specified IP addresses for Linux ECSs are generally static IP addresses. To avoid network exceptions caused by conflicts between NetworkManager and internal network services, do not enable NetworkManager when not required, such as when installing Kubernetes.

1.5 ECS and Other Services

Figure 1-2 shows the relationships between ECS and other services.

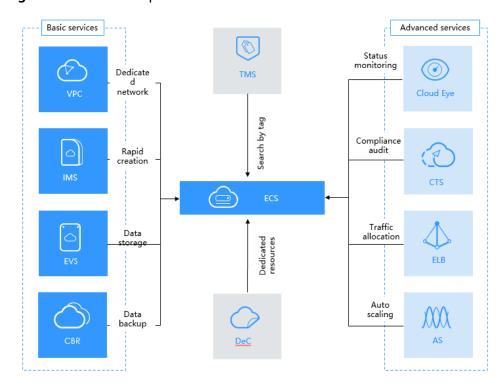


Figure 1-2 Relationships between ECS and other services

ECS-related Services

Auto Scaling (AS)

Automatically adjusts ECS resources based on the configured AS policies. This improves resource usage and reduces resource costs.

- Elastic Load Balance (ELB)
 - Automatically distributes traffic to multiple ECSs. This enhances system service and fault tolerance capabilities.
- Elastic Volume Service (EVS)
 Enables you to attach EVS disks to an ECS and expand their capacity.
- Virtual Private Cloud (VPC)
 - Enables you to configure internal networks and change network configurations by customizing security groups, VPNs, IP address ranges, and bandwidth. This simplifies network management. You can also customize the ECS access rules within a security group and between security groups to improve ECS security.
- Image Management Service (IMS)
 Enables you to create ECSs using images. This improves the efficiency of ECS creation.
- Cloud Eye
 - Allows you to check the status of monitored service objects after you have obtained an ECS. This can be done without requiring additional plug-ins be installed. For details about the metrics supported by ECS, see **Basic ECS**Metrics.
- Key Management Service (KMS)

The encryption feature relies on KMS. 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 KMS to improve data security.

- Cloud Trace Service (CTS)
 - Records ECS-related operations for later query, audit, and backtrack.
- Cloud Server Backup Service (CSBS)
 - Protects ECS backups. CSBS backs up all EVS disks of an ECS, including the system disk and data disks, and uses the backup to restore the ECS.
- Cloud Backup and Recovery (CBR)
 - Backs up EVS disks and ECSs for restoration. You can back up all EVS disks (including the system disk and data disks) attached to an ECS and use the backup to restore the ECS data.
- Tag Management Service (TMS)
 - Tags can be used to identify ECSs for easy management.

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 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 1-1 ECS statuses

| Status | Status Attribute | Description |
|----------|---------------------|---------------------------|
| Creating | Intermediate | The ECS is being created. |

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

1.6.3 ECS Types

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

- General-Purpose ECSs
- Dedicated General-Purpose ECSs
- Memory-optimized ECSs
- Disk-intensive ECSs
- Ultra-high I/O ECSs
- GPU-accelerated ECSs

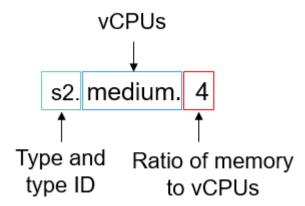
For details about the regions in which a flavor is available, see the information displayed on the management console.

ECS Flavor Naming Rules

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

Example: s2.medium.4

Figure 1-3 ECS flavor naming rules



The format is defined as follows:

- AB indicates the ECS type and type ID.
 - A specifies the ECS type. For example, s indicates a general-computing ECS, c indicates a general computing-plus ECS, and m indicates a memory-optimized ECS.
 - B specifies the type ID. For example, 1 in s1 indicates the first-generation general-computing ECS, and 2 in s2 indicates the second-generation general-computing ECS. Generally, a larger number indicates a newer generation, which is more cost-effective. For example, compared with s1 and s2, s6 is more cost-effective.
- **C** specifies the flavor size (the number of vCPUs), such as small, medium, large, xlarge, 2xlarge, 4xlarge, and 8xlarge.
- **D** specifies the ratio of memory to vCPUs and is expressed in a digit. For example, value **4** indicates that the ratio of memory to vCPUs is 4.

| Flavor Size | vCPUs |
|-------------|--|
| small | 1 |
| medium | 1 |
| large | 2 |
| xlarge | 4 |
| Mxlarge | N x 4. A larger value of N indicates more vCPUs. |

Table 1-2 Mapping between flavor and the number of vCPUs

How Do I Know My ECS Flavor?

When creating an ECS, you can view the flavors in the flavor list.

Figure 1-4 ECS flavors

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 QoS

Network QoS uses basic technologies to improve the quality of network communication. A network with QoS enabled offers predictable network performance and effectively allocates network bandwidth to use network resources.

To obtain the QoS data of an ECS flavor, including the maximum bandwidth, assured bandwidth, maximum intranet PPS, NIC multi-queues, and maximum NICs, see A Summary List of ECS Specifications.

The intranet bandwidth and 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.
 - PPS: packets per second, indicates the number of packets sent per second. It is usually used to measure the network performance.
- NIC multi-queues: allocates NIC interrupt requests to multiple vCPUs for higher PPS performance and bandwidth
- Maximum NICs: indicates the maximum number of NICs that can be attached to an ECS.
- Maximum supplementary NICs: indicates the maximum number of supplementary NICs that can be attached to an ECS.

◯ NOTE

- For instructions about how to test packet transmit and receive, see **How Can I Test Network Performance?**
- For instructions about how to enable NIC multi-queue, see Enabling NIC Multi-Queue.
- 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.

Dedicated and Shared ECSs

Table 1-3 Differences between dedicated and shared ECSs

| Dimensio n | Dedicated ECS | Shared ECS |
|-----------------------------|--|---|
| CPU Allocation | CPUs are exclusively used and there is no CPU contention. | CPUs are shared and CPU contention may occur. |
| Feature | High performance Dedicated and stable computing, storage, and network resources High costs | Unstable performance when loads are high Shared computing, storage, and network resources Low costs |
| Applicatio n Scenario | For enterprises that have high requirements on service stability | For small- and medium-sized websites or individuals that have requirements on cost-effectiveness |
| ECS Specificati ons | Specifications except general- purpose | x86 computing: • General-Purpose ECSs |

1.7 ECS Specifications and Types

1.7.1 A Summary List of ECS Specifications

General-Purpose

Table 1-4 T2 ECS specifications

| Flavor | vCPUs | Memory (GiB) | Max./Assured Bandwidth | Max. PPS | Virtualiz ation |
|------------------|-------|-----------------|---------------------------|----------|---|
| t2.micro | 1 | 1 | Low | Low | KVM |
| t2.small | 1 | 2 | Low | Low | (Applicabl e only to Netherlan ds) |
| t2.micro | 1 | 1 | Low | Low | Xen |
| t2.small | 1 | 2 | Low | Low | (Applicabl e only to |
| t2.large.2 | 2 | 4 | Low | Low | France) |
| t2.xlarge.2 | 4 | 8 | Low | Low | |
| t2.2xlarge. 2 | 8 | 16 | Low | Low | |

Table 1-5 S6 ECS specifications

| Flavor | vCP Us | Memor y (GiB) | Max./ Assured Bandwidth (Gbit/s) | Max. PPS (10,000) | Max. NIC Queue s | Max. NICs | Virtu alizat ion |
|------------------|-----------|---------------------|---|----------------------|---------------------------|--------------|------------------------|
| s6.small. | 1 | 1 | 0.8/0.15 | 10 | 1 | 2 | KVM |
| s6.mediu m.2 | 1 | 2 | 0.8/0.15 | 10 | 1 | 2 | KVM |
| s6.large. 2 | 2 | 4 | 1.5/0.3 | 15 | 1 | 2 | KVM |
| s6.xlarge .2 | 4 | 8 | 2/0.5 | 25 | 1 | 2 | KVM |
| s6.2xlarg e.2 | 8 | 16 | 3/1 | 50 | 2 | 2 | KVM |

| Flavor | vCP Us | Memor y (GiB) | Max./ Assured Bandwidth (Gbit/s) | Max. PPS (10,000) | Max. NIC Queue s | Max. NICs | Virtu alizat ion |
|------------------|-----------|---------------------|---|----------------------|---------------------------|--------------|------------------------|
| s6.4xlarg e.2 | 16 | 32 | 6/1.5 | 100 | 4 | 2 | KVM |
| s6.mediu m.4 | 1 | 4 | 0.8/0.15 | 10 | 1 | 2 | KVM |
| s6.large. 4 | 2 | 8 | 1.5/0.3 | 15 | 1 | 2 | KVM |
| s6.xlarge .4 | 4 | 16 | 2/0.5 | 25 | 1 | 2 | KVM |
| s6.2xlarg e.4 | 8 | 32 | 3/1 | 50 | 2 | 2 | KVM |
| s6.4xlarg e.4 | 16 | 64 | 6/1.5 | 100 | 4 | 2 | KVM |

Table 1-6 S3 ECS specifications

| Flavor | vCP Us | Memor y (GiB) | Max./ Assured Bandwidth (Gbit/s) | Max. PPS (10,000) | Max. NIC Queue s | Max. NICs | Virtu alizat ion |
|------------------|-----------|---------------------|---|----------------------|---------------------------|--------------|------------------------|
| s3.small. 1 | 1 | 1 | 0.5/0.1 | 5 | 1 | 1 | KVM |
| s3.mediu m.2 | 1 | 2 | 0.5/0.1 | 5 | 1 | 1 | KVM |
| s3.large. 2 | 2 | 4 | 0.8/0.2 | 10 | 1 | 1 | KVM |
| s3.xlarge .2 | 4 | 8 | 1.5/0.4 | 15 | 1 | 2 | KVM |
| s3.2xlarg e.2 | 8 | 16 | 3/0.8 | 20 | 2 | 2 | KVM |
| s3.4xlarg e.2 | 16 | 32 | 4/1.5 | 30 | 4 | 2 | KVM |
| s3.mediu m.4 | 1 | 4 | 0.5/0.1 | 5 | 1 | 1 | KVM |
| s3.large. 4 | 2 | 8 | 0.8/0.2 | 10 | 1 | 1 | KVM |

| Flavor | vCP Us | Memor y (GiB) | Max./ Assured Bandwidth (Gbit/s) | Max. PPS (10,000) | Max. NIC Queue s | Max. NICs | Virtu alizat ion |
|------------------|-----------|---------------------|---|----------------------|---------------------------|--------------|------------------------|
| s3.xlarge .4 | 4 | 16 | 1.5/0.4 | 15 | 1 | 2 | KVM |
| s3.2xlarg e.4 | 8 | 32 | 3/0.8 | 20 | 2 | 2 | KVM |
| s3.4xlarg e.4 | 16 | 64 | 4/1.5 | 30 | 4 | 2 | KVM |

Dedicated General-Purpose

Table 1-7 C6 ECS specifications

| Flavor | vCPU s | Memory (GiB) | Max./ Assured Bandwidt h (Gbit/s) | Max. PPS (10,000) | Max. NIC Queu es | Max. NICs | Virtu alizat ion |
|-------------------|-----------|-----------------|---|----------------------|---------------------------|--------------|------------------------|
| c6.larg e.2 | 2 | 4 | 4/1.2 | 40 | 2 | 2 | KVM |
| c6.xlar ge.2 | 4 | 8 | 8/2.4 | 80 | 2 | 3 | KVM |
| c6.2xla rge.2 | 8 | 16 | 15/4.5 | 150 | 4 | 4 | KVM |
| c6.3xla rge.2 | 12 | 24 | 17/7 | 200 | 4 | 6 | KVM |
| c6.4xla rge.2 | 16 | 32 | 20/9 | 280 | 8 | 8 | KVM |
| c6.6xla rge.2 | 24 | 48 | 25/14 | 400 | 8 | 8 | KVM |
| c6.8xla rge.2 | 32 | 64 | 30/18 | 550 | 16 | 8 | KVM |
| c6.16xl arge.2 | 64 | 128 | 40/36 | 1000 | 32 | 8 | KVM |
| c6.larg e.4 | 2 | 8 | 4/1.2 | 40 | 2 | 2 | KVM |
| c6.xlar ge.4 | 4 | 16 | 8/2.4 | 80 | 2 | 3 | KVM |

| Flavor | vCPU s | Memory (GiB) | Max./ Assured Bandwidt h (Gbit/s) | Max. PPS (10,000) | Max. NIC Queu es | Max. NICs | Virtu alizat ion |
|-------------------|-----------|-----------------|---|----------------------|---------------------------|--------------|------------------------|
| c6.2xla rge.4 | 8 | 32 | 15/4.5 | 150 | 4 | 4 | KVM |
| c6.3xla rge.4 | 12 | 48 | 17/7 | 200 | 4 | 6 | KVM |
| c6.4xla rge.4 | 16 | 64 | 20/9 | 280 | 8 | 8 | KVM |
| c6.6xla rge.4 | 24 | 96 | 25/14 | 400 | 8 | 8 | KVM |
| c6.8xla rge.4 | 32 | 128 | 30/18 | 550 | 16 | 8 | KVM |
| c6.16xl arge.4 | 64 | 256 | 40/36 | 1000 | 32 | 8 | KVM |

Table 1-8 C3 ECS specifications

| Flavor | vCPUs | Memory (GiB) | Max./ Assured Bandwidth (Gbit/s) | Max. PPS (10,000) | Max. NIC Queue s | Virtualiz ation |
|-------------------|-------|-----------------|---|-------------------------|---------------------------|--------------------|
| c3.large. 2 | 2 | 4 | 1.5/0.6 | 30 | 2 | KVM |
| c3.xlarge | 4 | 8 | 3/1 | 50 | 2 | KVM |
| c3.2xlarg e.2 | 8 | 16 | 5/2 | 90 | 4 | KVM |
| c3.4xlarg e.2 | 16 | 32 | 10/4 | 130 | 4 | KVM |
| c3.8xlarg e.2 | 32 | 64 | 15/8 | 260 | 8 | KVM |
| c3.15xlar ge.2 | 60 | 128 | 17/16 | 500 | 16 | KVM |

Table 1-9 Cc3 ECS specifications

| Flavor | vCP Us | Memor y (GiB) | Max./ Assured Bandwidth (Gbit/s) | Max. PPS (10,000) | Max. NIC Queue s | Max. NICs | Virtu alizat ion |
|--------------------|-----------|---------------------|---|----------------------|---------------------------|--------------|------------------------|
| cc3.large. | 2 | 8 | 1.5/0.45 | 25 | 2 | 2 | KVM |
| cc3.xlarg e.4 | 4 | 16 | 3/0.9 | 40 | 2 | 3 | KVM |
| cc3.2xlar ge.4 | 8 | 32 | 5/1.8 | 65 | 4 | 4 | KVM |
| cc3.4xlar ge.4 | 16 | 64 | 8/3.6 | 120 | 4 | 8 | KVM |
| cc3.8xlar ge.4 | 32 | 128 | 15/7.2 | 210 | 8 | 8 | KVM |
| cc3.19xla rge.4 | 76 | 304 | 17/17 | 500 | 16 | 8 | KVM |

Memory-optimized

Table 1-10 M6 ECS specifications

| Flavor | vCPUs | Memory (GiB) | Max./ Assured Bandwidt h (Gbit/s) | Max. PPS (10,000) | Max. NIC Queu es | Max. NICs | Virtu alizat ion |
|------------------|-------|-----------------|---|----------------------|---------------------------|--------------|------------------------|
| m6.larg e.8 | 2 | 16 | 4/1.2 | 40 | 2 | 2 | KVM |
| m6.xlar ge.8 | 4 | 32 | 8/2.4 | 80 | 2 | 3 | KVM |
| m6.2xla rge.8 | 8 | 64 | 15/4.5 | 150 | 4 | 4 | KVM |
| m6.3xla rge.8 | 12 | 96 | 17/7 | 200 | 4 | 6 | KVM |
| m6.4xla rge.8 | 16 | 128 | 20/9 | 280 | 8 | 8 | KVM |
| m6.6xla rge.8 | 24 | 192 | 25/14 | 400 | 8 | 8 | KVM |

| Flavor | vCPUs | Memory (GiB) | Max./ Assured Bandwidt h (Gbit/s) | Max. PPS (10,000) | Max. NIC Queu es | Max. NICs | Virtu alizat ion |
|-------------------|-------|-----------------|---|----------------------|---------------------------|--------------|------------------------|
| m6.8xla rge.8 | 32 | 256 | 30/18 | 550 | 16 | 8 | KVM |
| m6.16xl arge.8 | 64 | 512 | 40/36 | 1,000 | 32 | 8 | KVM |
| m6.larg e.16 | 2 | 32 | 4/1.2 | 40 | 2 | 2 | KVM |
| m6.xlar ge.16 | 4 | 64 | 8/2.4 | 80 | 2 | 3 | KVM |
| m6.2xla rge.16 | 8 | 96 | 15/4.5 | 150 | 4 | 4 | KVM |
| m6.3xla rge.16 | 12 | 128 | 17/7 | 200 | 4 | 6 | KVM |
| m6.4xla rge.16 | 16 | 192 | 20/9 | 280 | 8 | 8 | KVM |
| m6.6xla rge.16 | 24 | 256 | 25/14 | 400 | 8 | 8 | KVM |
| m6.8xla rge.16 | 32 | 512 | 30/18 | 550 | 16 | 8 | KVM |

Table 1-11 M2 ECS specifications

| Flavor | vCPUs | Memory (GiB) | Max./ Assured Bandwidth (Gbit/s) | Max. PPS (10,000) | Max. NIC Queue s | Virtual ization |
|------------------|-------|-----------------|---|----------------------|---------------------------|--------------------|
| m2.large .8 | 2 | 16 | 1.5/0.5 | 10 | 1 | KVM |
| m2.xlarg e.8 | 4 | 32 | 3/1 | 15 | 1 | KVM |
| m2.2xlar ge.8 | 8 | 64 | 5/2 | 30 | 2 | KVM |
| m2.4xlar ge.8 | 16 | 128 | 8/4 | 40 | 4 | KVM |
| m2.8xlar ge.8 | 32 | 256 | 13/8 | 60 | 8 | KVM |

Disk-intensive

Table 1-12 D6 ECS specifications

| Flav or | vCPUs | Memor y (GiB) | Max./ Assured Bandwi dth (Gbit/s) | Max. PPS (10,000) | Max. NIC Que ues | Max. NICs | Local Disk (TB) | Virtua lizatio n |
|-----------------------|-------|---------------------|---|-------------------------|---------------------------|--------------|-----------------------|------------------------|
| d6.xl arge. 4 | 4 | 16 | 5/2 | 60 | 2 | 3 | 2 × 7.4 | KVM |
| d6.2 xlarg e.4 | 8 | 32 | 10/4 | 120 | 4 | 4 | 4 × 7.4 | KVM |
| d6.4 xlarg e.4 | 16 | 64 | 20/7.5 | 240 | 8 | 8 | 8 × 7.4 | KVM |
| d6.6 xlarg e.4 | 24 | 96 | 25/11 | 350 | 8 | 8 | 12 × 7.4 | KVM |
| d6.8 xlarg e.4 | 32 | 128 | 30/15 | 450 | 16 | 8 | 16 × 7.4 | KVM |
| d6.1 2xlar ge.4 | 48 | 192 | 40/22 | 650 | 16 | 8 | 24 × 7.4 | KVM |
| d6.1 6xlar ge.4 | 64 | 256 | 42/30 | 850 | 32 | 8 | 32 × 7.4 | KVM |
| d6.1 8xlar ge.4 | 72 | 288 | 44/34 | 900 | 32 | 8 | 36 × 7.4 | KVM |

Table 1-13 D3 ECS specifications

| Flav or | vCPU s | Memor y (GiB) | Max./ Assured Bandwi dth (Gbit/s) | Max. PPS (10,000 | Max. NIC Queu es | Max. NICs | Local Disks (GiB) | Virtua lizatio n |
|----------------------------|-----------|---------------------|---|------------------------|---------------------------|--------------|-------------------------|------------------------|
| d3.xl arge. 8 | 4 | 32 | 2.5/2.5 | 50 | 2 | 3 | 2 × 1,675 | KVM |
| d3.2 xlarg e.8 | 8 | 64 | 5/5 | 100 | 2 | 4 | 4 × 1,675 | KVM |
| d3.4 xlarg e.8 | 16 | 128 | 10/10 | 120 | 4 | 8 | 8 × 1,675 | KVM |
| d3.6 xlarg e.8 | 24 | 192 | 15/15 | 160 | 6 | 8 | 12 × 1,675 | KVM |
| d3.8 xlarg e.8 | 32 | 256 | 20/20 | 200 | 8 | 8 | 16 × 1675 | KVM |
| d3.1 2xlar ge.8 | 48 | 384 | 32/32 | 220 | 16 | 8 | 24 × 1675 | KVM |
| d3.1 4xlar ge.1 0 | 56 | 560 | 40/40 | 500 | 16 | 8 | 28 × 1,675 | KVM |

Table 1-14 D2 ECS specifications

| Flavo r | vCPU s | Memor y (GiB) | Max./ Assured Bandwid th (Gbit/s) | Max. PPS (10,000) | Local Disks (GiB) | Max. NIC Queu es | Virtual ization |
|----------------------|-----------|---------------------|---|----------------------|-------------------------|---------------------------|--------------------|
| d2.xla rge.8 | 4 | 32 | 1/1 | 15 | 2 × 1,675 | 2 | KVM |
| d2.2xl arge. 8 | 8 | 64 | 2/2 | 30 | 4 × 1,675 | 2 | KVM |

| Flavo r | vCPU s | Memor y (GiB) | Max./ Assured Bandwid th (Gbit/s) | Max. PPS (10,000) | Local Disks (GiB) | Max. NIC Queu es | Virtual ization |
|-----------------------|-----------|---------------------|---|----------------------|-------------------------|---------------------------|--------------------|
| d2.4xl arge. 8 | 16 | 128 | 4/4 | 40 | 8 × 1,675 | 4 | KVM |
| d2.6xl arge. 8 | 24 | 192 | 6/6 | 50 | 12 × 1,675 | 6 | KVM |
| d2.8xl arge. 8 | 32 | 256 | 8/8 | 60 | 16 × 1,675 | 8 | KVM |
| d2.12 xlarg e.8 | 48 | 384 | 12/12 | 90 | 24 × 1,675 | 8 | KVM |

Ultra-high I/O

Table 1-15 I3 ECS specifications

| Flavo r | vCPU s | Memor y (GiB) | Max./ Assured Bandwi dth (Gbit/s) | Max. PPS (10,000 | Max. NIC Queu es | Max NIC s | Local Disks | Virtua lizatio n |
|-----------------------|-----------|---------------------|---|------------------------|---------------------------|-----------------|-----------------------------|------------------------|
| i3.2xl arge.8 | 8 | 64 | 2.5/2.5 | 100 | 4 | 4 | 1 × 1,600 GiB NVMe | KVM |
| i3.4xl arge.8 | 16 | 128 | 5/5 | 150 | 4 | 8 | 2 × 1,600 GiB NVMe | KVM |
| i3.8xl arge.8 | 32 | 256 | 10/10 | 200 | 8 | 8 | 4 × 1,600 GiB NVMe | KVM |
| i3.12x large. 8 | 48 | 384 | 15/15 | 240 | 8 | 8 | 6 × 1,600 GiB NVMe | KVM |

| Flavo r | vCPU s | Memor y (GiB) | Max./ Assured Bandwi dth (Gbit/s) | Max. PPS (10,000 | Max. NIC Queu es | Max NIC s | Local Disks | Virtua lizatio n |
|-----------------------|-----------|---------------------|---|------------------------|---------------------------|-----------------|-----------------------------|------------------------|
| i3.15x large. 8 | 60 | 512 | 25/25 | 500 | 16 | 8 | 7 × 1,600 GiB NVMe | KVM |

GPU-accelerated

Table 1-16 G6 ECS specifications

| Flav or | vCP Us | Mem ory (GiB) | Max./ Assur ed Band width (Gbit/ s) | Max. PPS (10,00 0) | Ma x. NIC Que ues | Max. NICs | GPUs | GPU Memo ry (GiB) | Virtu alizat ion |
|---------------------------|-----------|---------------------|---|-----------------------------|-------------------------------|--------------|--------|----------------------------|------------------------|
| g6.4 xlar ge.4 | 16 | 64 | 15/8 | 200 | 8 | 8 | 1 × T4 | 16 | KVM |
| g6.6 xlar ge.4 | 24 | 96 | 25/15 | 200 | 8 | 8 | 1 × T4 | 16 | KVM |
| g6.9 xlar ge.7 | 36 | 252 | 25/15 | 200 | 16 | 8 | 1 × T4 | 16 | KVM |
| g6.1 8xla rge. 7 | 72 | 504 | 30/30 | 400 | 32 | 16 | 2 × T4 | 32 | KVM |

Table 1-17 G5 ECS specifications

| Flav or | vCPUs | Memor y (GiB) | Max./ Assured Bandwi dth (Gbit/s) | Max. PPS (10,000) | Max. NIC Que ues | GPU s | GPU Memor y (GiB) | Virtua lizatio n |
|----------------------|-------|---------------------|---|-------------------------|---------------------------|-------------|----------------------------|------------------------|
| g5.xl arge. 2 | 4 | 8 | 4/1.3 | 20 | 2 | V100 -1Q | 1 | KVM |
| g5.2 xlarg e.2 | 8 | 16 | 6/2 | 35 | 4 | V100 -2Q | 2 | KVM |
| g5.4 xlarg e.4 | 16 | 64 | 10/4 | 50 | 8 | V100 -8Q | 8 | KVM |

Table 1-18 G3 ECS specifications

| Flav or | vCPUs | Memor y (GiB) | Max./ Assured Bandwi dth (Gbit/s) | Max. PPS (10,000) | Max. NIC Que ues | GPU s | GPU Memor y (GiB) | Virtua lizatio n |
|----------------------|-------|---------------------|---|-------------------------|---------------------------|------------|----------------------------|------------------------|
| g3.4 xlarg e.4 | 16 | 64 | 8/2.5 | 50 | 2 | 1 x M60 | 1 x 8 | KVM |
| g3.8 xlarg e.4 | 32 | 128 | 10/5 | 100 | 4 | 2 x M60 | 2 x 8 | KVM |

Table 1-19 P2s ECS specifications

| Flavo r | vCP Us | Me mor y (GiB | Max./ Assure d Bandw idth (Gbit/s) | Max. PPS (10,0 00) | Max NIC Que ues | Ma x. NIC s | GP Us | GPU Con nect ion | GPU Mem ory (GiB) | Virtu alizat ion |
|-----------------------|-----------|------------------------|---|-----------------------------|--------------------------|----------------------|-----------------|---------------------------|--------------------------------|------------------------|
| p2s.2 xlarg e.8 | 8 | 64 | 10/4 | 50 | 4 | 4 | 1 × V10 0 | PCle Gen 3 | 1 × 32 GiB | KVM |

| Flavo r | vCP Us | Me mor y (GiB | Max./ Assure d Bandw idth (Gbit/s) | Max. PPS (10,0 00) | Max NIC Que ues | Ma x. NIC s | GP Us | GPU Con nect ion | GPU Mem ory (GiB) | Virtu alizat ion |
|------------------------|-----------|------------------------|---|-----------------------------|--------------------------|----------------------|-----------------|---------------------------|--------------------------------|------------------------|
| p2s.4 xlarg e.8 | 16 | 128 | 15/8 | 100 | 8 | 8 | 2 × V10 0 | PCle Gen 3 | 2 × 32 GiB | KVM |
| p2s.8 xlarg e.8 | 32 | 256 | 25/15 | 200 | 16 | 8 | 4 × V10 0 | PCle Gen 3 | 4 × 32 GiB | KVM |
| p2s.1 6xlar ge.8 | 64 | 512 | 30/30 | 400 | 32 | 8 | 8 × V10 0 | PCle Gen 3 | 8 × 32 GiB | KVM |

Table 1-20 P2 ECS specifications

| Flav or | vC P Us | Me mor y (Gi B) | Max./ Assure d Bandw idth (Gbit/s) | Max. PPS (10,00 0) | Ma x. NIC Qu eue s | Max NICs | GPU s | GPU Memor y (GiB) | Loc al Disk s (GiB | Virtu aliza tion |
|----------------------|---------------|-----------------------------|---|-----------------------------|-----------------------------------|-------------|-----------------|----------------------------|--------------------------------|------------------------|
| p2.2 xlar ge.8 | 8 | 64 | 10/4 | 50 | 2 | 12 | 1 × V10 0 | 1 × 16 | 1 × 800 | KVM |
| p2.4 xlar ge.8 | 16 | 128 | 15/8 | 100 | 4 | 12 | 2 × V10 0 | 2 × 16 | 2 × 800 | KVM |
| p2.8 xlar ge.8 | 32 | 256 | 25/15 | 200 | 8 | 12 | 4 × V10 0 | 4 × 16 | 4 × 800 | KVM |

| Flav | vC P Us | Me mo ry (Gi B) | Max./ Assure d Band width (Gbit/ s) | Max. PPS (10,0 00) | Max. NIC Queu es | Max. NICs | GPU s | GP U Me mo ry (Gi B) | Loc al Dis ks | Virtua lizatio n |
|---------------------------|---------------|-----------------------------|---|-----------------------------|---------------------------|--------------|-----------|--|------------------------|------------------------|
| pi2. 2xla rge. 4 | 8 | 32 | 10/4 | 50 | 4 | 4 | 1 × T4 | 1 × 16 | N/A | KVM |
| pi2. 4xla rge. 4 | 16 | 64 | 15/8 | 100 | 8 | 8 | 2 × T4 | 2 × 16 | N/A | KVM |
| pi2. 8xla rge. 4 | 32 | 128 | 25/15 | 200 | 16 | 8 | 4 × T4 | 4 × 16 | N/A | KVM |

Table 1-21 Pi2 ECS specifications

1.7.2 General-Purpose ECSs

Overview

General-purpose ECSs provide a baseline level of vCPU performance and a balance of compute, memory, and networking resources. They use a CPU-unbound scheduling scheme. vCPUs are randomly allocated to idle CPU hyper threads based on the system loads. If traffic loads are light, the computing performance is high. However, if traffic loads are heavy, vCPUs of different ECSs compete for physical CPU resources, resulting in unstable computing performance.

Compared with dedicated general-purpose ECSs, general-purpose ECSs focus on resource sharing and cannot ensure stability of compute performance, but they offer lower costs. General-purpose ECSs are suitable for cost-sensitive applications that can tolerate performance jitter, such as web servers, general workloads, development environments, and small-scale databases.

S6 and S3 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.

Scenarios

Applications

General-purpose ECSs are suitable for applications that have no special requirements on CPU performance, memory, disk capacity, or bandwidth, but

have high requirements on security and reliability. Example applications include web servers, R&D and testing environments, and small-scale databases. They feature low initial investment and maintenance costs.

• Application scenarios

Enterprise website deployment, enterprise office environment setup, and enterprise R&D and testing activities

Table 1-22 T2 ECS specifications

| Flavor | vCPUs | Memory (GiB) | Max./Assured Bandwidth | Max. PPS | Virtualiz ation |
|------------------|-------|-----------------|---------------------------|----------|---|
| t2.micro | 1 | 1 | Low | Low | KVM |
| t2.small | 1 | 2 | Low | Low | (Applicabl e only to Netherlan ds) |
| t2.micro | 1 | 1 | Low | Low | Xen |
| t2.small | 1 | 2 | Low | Low | (Applicabl e only to |
| t2.large.2 | 2 | 4 | Low | Low | France) |
| t2.xlarge.2 | 4 | 8 | Low | Low | |
| t2.2xlarge. 2 | 8 | 16 | Low | Low | |

Table 1-23 S6 ECS specifications

| Flavor | vCP Us | Memor y (GiB) | Max./ Assured Bandwidth (Gbit/s) | Max. PPS (10,000) | Max. NIC Queue s | Max. NICs | Virtu alizat ion |
|------------------|-----------|---------------------|---|----------------------|---------------------------|--------------|------------------------|
| s6.small. 1 | 1 | 1 | 0.8/0.15 | 10 | 1 | 2 | KVM |
| s6.mediu m.2 | 1 | 2 | 0.8/0.15 | 10 | 1 | 2 | KVM |
| s6.large. 2 | 2 | 4 | 1.5/0.3 | 15 | 1 | 2 | KVM |
| s6.xlarge .2 | 4 | 8 | 2/0.5 | 25 | 1 | 2 | KVM |
| s6.2xlarg e.2 | 8 | 16 | 3/1 | 50 | 2 | 2 | KVM |

| Flavor | vCP Us | Memor y (GiB) | Max./ Assured Bandwidth (Gbit/s) | Max. PPS (10,000) | Max. NIC Queue s | Max. NICs | Virtu alizat ion |
|------------------|-----------|---------------------|---|----------------------|---------------------------|--------------|------------------------|
| s6.4xlarg e.2 | 16 | 32 | 6/1.5 | 100 | 4 | 2 | KVM |
| s6.mediu m.4 | 1 | 4 | 0.8/0.15 | 10 | 1 | 2 | KVM |
| s6.large. 4 | 2 | 8 | 1.5/0.3 | 15 | 1 | 2 | KVM |
| s6.xlarge .4 | 4 | 16 | 2/0.5 | 25 | 1 | 2 | KVM |
| s6.2xlarg e.4 | 8 | 32 | 3/1 | 50 | 2 | 2 | KVM |
| s6.4xlarg e.4 | 16 | 64 | 6/1.5 | 100 | 4 | 2 | KVM |

Table 1-24 S3 ECS specifications

| Flavor | vCP Us | Memor y (GiB) | Max./ Assured Bandwidth (Gbit/s) | Max. PPS (10,000) | Max. NIC Queue s | Max. NICs | Virtu alizat ion |
|------------------|-----------|---------------------|---|----------------------|---------------------------|--------------|------------------------|
| s3.small. 1 | 1 | 1 | 0.5/0.1 | 5 | 1 | 1 | KVM |
| s3.mediu m.2 | 1 | 2 | 0.5/0.1 | 5 | 1 | 1 | KVM |
| s3.large. 2 | 2 | 4 | 0.8/0.2 | 10 | 1 | 1 | KVM |
| s3.xlarge .2 | 4 | 8 | 1.5/0.4 | 15 | 1 | 2 | KVM |
| s3.2xlarg e.2 | 8 | 16 | 3/0.8 | 20 | 2 | 2 | KVM |
| s3.4xlarg e.2 | 16 | 32 | 4/1.5 | 30 | 4 | 2 | KVM |
| s3.mediu m.4 | 1 | 4 | 0.5/0.1 | 5 | 1 | 1 | KVM |
| s3.large. 4 | 2 | 8 | 0.8/0.2 | 10 | 1 | 1 | KVM |

| Flavor | vCP Us | Memor y (GiB) | Max./ Assured Bandwidth (Gbit/s) | Max. PPS (10,000) | Max. NIC Queue s | Max. NICs | Virtu alizat ion |
|------------------|-----------|---------------------|---|----------------------|---------------------------|--------------|------------------------|
| s3.xlarge .4 | 4 | 16 | 1.5/0.4 | 15 | 1 | 2 | KVM |
| s3.2xlarg e.4 | 8 | 32 | 3/0.8 | 20 | 2 | 2 | KVM |
| s3.4xlarg e.4 | 16 | 64 | 4/1.5 | 30 | 4 | 2 | KVM |

1.7.3 Dedicated General-Purpose ECSs

Overview

Dedicated general-purpose ECSs provide dedicated vCPUs, featuring powerful performance. In addition, the ECSs use latest-generation network acceleration engines and Data Plane Development Kit (DPDK) to provide high network performance, meeting requirements in different scenarios.

- C6 ECSs use second-generation Intel® Xeon® Scalable processors to provide powerful and stable computing performance. By using 25GE high-speed intelligent NICs, C6 ECSs offer ultra-high network bandwidth and packets per second (PPS).
- C3 ECSs deliver powerful and stable computing performance. They use Intel® Xeon® Scalable processors and high-performance NICs to provide high performance and stability for enterprise-grade application requirements.
- CC3 ECSs provide balanced computing, memory, and network performance.
 They are suited for distributed computing clusters and applications with
 moderate loads, such as websites and web applications, generalized databases
 and cache servers, enterprise applications, and data analysis.

Scenarios

- C6 ECSs: Websites and web applications, generalized databases and cache servers, and medium- and heavy-workload enterprise applications with strict requirements on computing and network performance
- C3 ECSs: Small- and medium-scale databases, cache servers, and search clusters with high requirements on stability; enterprise-grade applications
- CC3 ECSs: Small- and medium-scale e-commerce websites, SQL databases, computing clusters, and backend servers for enterprise applications

Table 1-25 C6 ECS specifications

| Flavor | vCPU s | Memory (GiB) | Max./ Assured Bandwidt h (Gbit/s) | Max. PPS (10,000) | Max. NIC Queu es | Max. NICs | Virtu alizat ion |
|-------------------|-----------|-----------------|---|----------------------|---------------------------|--------------|------------------------|
| c6.larg e.2 | 2 | 4 | 4/1.2 | 40 | 2 | 2 | KVM |
| c6.xlar ge.2 | 4 | 8 | 8/2.4 | 80 | 2 | 3 | KVM |
| c6.2xla rge.2 | 8 | 16 | 15/4.5 | 150 | 4 | 4 | KVM |
| c6.3xla rge.2 | 12 | 24 | 17/7 | 200 | 4 | 6 | KVM |
| c6.4xla rge.2 | 16 | 32 | 20/9 | 280 | 8 | 8 | KVM |
| c6.6xla rge.2 | 24 | 48 | 25/14 | 400 | 8 | 8 | KVM |
| c6.8xla rge.2 | 32 | 64 | 30/18 | 550 | 16 | 8 | KVM |
| c6.16xl arge.2 | 64 | 128 | 40/36 | 1000 | 32 | 8 | KVM |
| c6.larg e.4 | 2 | 8 | 4/1.2 | 40 | 2 | 2 | KVM |
| c6.xlar ge.4 | 4 | 16 | 8/2.4 | 80 | 2 | 3 | KVM |
| c6.2xla rge.4 | 8 | 32 | 15/4.5 | 150 | 4 | 4 | KVM |
| c6.3xla rge.4 | 12 | 48 | 17/7 | 200 | 4 | 6 | KVM |
| c6.4xla rge.4 | 16 | 64 | 20/9 | 280 | 8 | 8 | KVM |
| c6.6xla rge.4 | 24 | 96 | 25/14 | 400 | 8 | 8 | KVM |
| c6.8xla rge.4 | 32 | 128 | 30/18 | 550 | 16 | 8 | KVM |
| c6.16xl arge.4 | 64 | 256 | 40/36 | 1000 | 32 | 8 | KVM |

Table 1-26 C3 ECS specifications

| Flavor | vCPUs | Memory (GiB) | Max./ Assured Bandwidth (Gbit/s) | Max. PPS (10,000) | Max. NIC Queue s | Virtualiz ation |
|-------------------|-------|-----------------|---|-------------------------|---------------------------|--------------------|
| c3.large. 2 | 2 | 4 | 1.5/0.6 | 30 | 2 | KVM |
| c3.xlarge | 4 | 8 | 3/1 | 50 | 2 | KVM |
| c3.2xlarg e.2 | 8 | 16 | 5/2 | 90 | 4 | KVM |
| c3.4xlarg e.2 | 16 | 32 | 10/4 | 130 | 4 | KVM |
| c3.8xlarg e.2 | 32 | 64 | 15/8 | 260 | 8 | KVM |
| c3.15xlar ge.2 | 60 | 128 | 17/16 | 500 | 16 | KVM |

Table 1-27 Cc3 ECS specifications

| Flavor | vCP Us | Memor y (GiB) | Max./ Assured Bandwidth (Gbit/s) | Max. PPS (10,000) | Max. NIC Queue s | Max. NICs | Virtu alizat ion |
|--------------------|-----------|---------------------|---|----------------------|---------------------------|--------------|------------------------|
| cc3.large. | 2 | 8 | 1.5/0.45 | 25 | 2 | 2 | KVM |
| cc3.xlarg e.4 | 4 | 16 | 3/0.9 | 40 | 2 | 3 | KVM |
| cc3.2xlar ge.4 | 8 | 32 | 5/1.8 | 65 | 4 | 4 | KVM |
| cc3.4xlar ge.4 | 16 | 64 | 8/3.6 | 120 | 4 | 8 | KVM |
| cc3.8xlar ge.4 | 32 | 128 | 15/7.2 | 210 | 8 | 8 | KVM |
| cc3.19xla rge.4 | 76 | 304 | 17/17 | 500 | 16 | 8 | KVM |

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

- M6 ECSs use the second-generation Intel® Xeon® Scalable processors with technologies optimized to offer powerful and stable computing performance. Using 25GE high-speed intelligent NICs, M6 ECSs provide a maximum memory size of 512 GiB based on DDR4 for memory-intensive applications with high requirements on network bandwidth and Packets Per Second (PPS).
- M2 ECSs use Intel Xeon E5-2690 v4 CPUs and are designed for memoryoptimized applications.

Scenarios

Applications

Memory-optimized ECSs are suitable for applications that require a large amount of memory such as relational databases, NoSQL databases, and memory data analysis

• Application scenarios

Big data analysis for precision marketing, e-commerce, and IoV, relational databases, NoSQL databases, and memory data analysis

Table 1-28 M6 ECS specifications

| Flavor | vCPUs | Memory (GiB) | Max./ Assured Bandwidt h (Gbit/s) | Max. PPS (10,000) | Max. NIC Queu es | Max. NICs | Virtu alizat ion |
|------------------|-------|-----------------|---|----------------------|---------------------------|--------------|------------------------|
| m6.larg e.8 | 2 | 16 | 4/1.2 | 40 | 2 | 2 | KVM |
| m6.xlar ge.8 | 4 | 32 | 8/2.4 | 80 | 2 | 3 | KVM |
| m6.2xla rge.8 | 8 | 64 | 15/4.5 | 150 | 4 | 4 | KVM |
| m6.3xla rge.8 | 12 | 96 | 17/7 | 200 | 4 | 6 | KVM |
| m6.4xla rge.8 | 16 | 128 | 20/9 | 280 | 8 | 8 | KVM |

| Flavor | vCPUs | Memory (GiB) | Max./ Assured Bandwidt h (Gbit/s) | Max. PPS (10,000) | Max. NIC Queu es | Max. NICs | Virtu alizat ion |
|-------------------|-------|-----------------|---|----------------------|---------------------------|--------------|------------------------|
| m6.6xla rge.8 | 24 | 192 | 25/14 | 400 | 8 | 8 | KVM |
| m6.8xla rge.8 | 32 | 256 | 30/18 | 550 | 16 | 8 | KVM |
| m6.16xl arge.8 | 64 | 512 | 40/36 | 1,000 | 32 | 8 | KVM |
| m6.larg e.16 | 2 | 32 | 4/1.2 | 40 | 2 | 2 | KVM |
| m6.xlar ge.16 | 4 | 64 | 8/2.4 | 80 | 2 | 3 | KVM |
| m6.2xla rge.16 | 8 | 96 | 15/4.5 | 150 | 4 | 4 | KVM |
| m6.3xla rge.16 | 12 | 128 | 17/7 | 200 | 4 | 6 | KVM |
| m6.4xla rge.16 | 16 | 192 | 20/9 | 280 | 8 | 8 | KVM |
| m6.6xla rge.16 | 24 | 256 | 25/14 | 400 | 8 | 8 | KVM |
| m6.8xla rge.16 | 32 | 512 | 30/18 | 550 | 16 | 8 | KVM |

Table 1-29 M2 ECS specifications

| Flavor | vCPUs | Memory (GiB) | Max./ Assured Bandwidth (Gbit/s) | Max. PPS (10,000) | Max. NIC Queue s | Virtual ization |
|------------------|-------|-----------------|---|----------------------|---------------------------|--------------------|
| m2.large .8 | 2 | 16 | 1.5/0.5 | 10 | 1 | KVM |
| m2.xlarg e.8 | 4 | 32 | 3/1 | 15 | 1 | KVM |
| m2.2xlar ge.8 | 8 | 64 | 5/2 | 30 | 2 | KVM |
| m2.4xlar ge.8 | 16 | 128 | 8/4 | 40 | 4 | KVM |

| Flavor | vCPUs | Memory (GiB) | Max./ Assured Bandwidth (Gbit/s) | Max. PPS (10,000) | Max. NIC Queue s | Virtual ization |
|------------------|-------|-----------------|---|----------------------|---------------------------|--------------------|
| m2.8xlar ge.8 | 32 | 256 | 13/8 | 60 | 8 | KVM |

1.7.5 Disk-intensive ECSs

Overview

Disk-intensive ECSs are delivered with local disks for high storage bandwidth and IOPS. 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.

D6 ECSs, with a vCPU/memory ratio of 1:4, use 2nd Generation Intel® Xeon® Scalable processors to offer powerful and stable computing performance. Equipped with proprietary 25GE high-speed intelligent NICs and local SATA disks, D6 ECSs offer ultra-high network bandwidth, PPS, and local storage. The capacity of a single SATA disk is up to 7.4 TB, and an ECS can have up to 36 such disks attached.

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.

D2 ECSs are KVM-based. They use local storage for high storage performance and intranet bandwidth.

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

Table 1-30 D6 ECS specifications

| Flav or | vCPUs | Memor y (GiB) | Max./ Assured Bandwi dth (Gbit/s) | Max. PPS (10,000) | Max. NIC Que ues | Max. NICs | Local Disk (TB) | Virtua lizatio n |
|-----------------------|-------|---------------------|---|-------------------------|---------------------------|--------------|-----------------------|------------------------|
| d6.xl arge. 4 | 4 | 16 | 5/2 | 60 | 2 | 3 | 2 × 7.4 | KVM |
| d6.2 xlarg e.4 | 8 | 32 | 10/4 | 120 | 4 | 4 | 4 × 7.4 | KVM |
| d6.4 xlarg e.4 | 16 | 64 | 20/7.5 | 240 | 8 | 8 | 8 × 7.4 | KVM |
| d6.6 xlarg e.4 | 24 | 96 | 25/11 | 350 | 8 | 8 | 12 × 7.4 | KVM |
| d6.8 xlarg e.4 | 32 | 128 | 30/15 | 450 | 16 | 8 | 16 × 7.4 | KVM |
| d6.1 2xlar ge.4 | 48 | 192 | 40/22 | 650 | 16 | 8 | 24 × 7.4 | KVM |
| d6.1 6xlar ge.4 | 64 | 256 | 42/30 | 850 | 32 | 8 | 32 × 7.4 | KVM |
| d6.1 8xlar ge.4 | 72 | 288 | 44/34 | 900 | 32 | 8 | 36 × 7.4 | KVM |

Table 1-31 D3 ECS specifications

| Flav or | vCPU s | Memor y (GiB) | Max./ Assured Bandwi dth (Gbit/s) | Max. PPS (10,000 | Max. NIC Queu es | Max. NICs | Local Disks (GiB) | Virtua lizatio n |
|----------------------------|-----------|---------------------|---|------------------------|---------------------------|--------------|-------------------------|------------------------|
| d3.xl arge. 8 | 4 | 32 | 2.5/2.5 | 50 | 2 | 3 | 2 × 1,675 | KVM |
| d3.2 xlarg e.8 | 8 | 64 | 5/5 | 100 | 2 | 4 | 4 × 1,675 | KVM |
| d3.4 xlarg e.8 | 16 | 128 | 10/10 | 120 | 4 | 8 | 8 × 1,675 | KVM |
| d3.6 xlarg e.8 | 24 | 192 | 15/15 | 160 | 6 | 8 | 12 × 1,675 | KVM |
| d3.8 xlarg e.8 | 32 | 256 | 20/20 | 200 | 8 | 8 | 16 × 1675 | KVM |
| d3.1 2xlar ge.8 | 48 | 384 | 32/32 | 220 | 16 | 8 | 24 × 1675 | KVM |
| d3.1 4xlar ge.1 0 | 56 | 560 | 40/40 | 500 | 16 | 8 | 28 × 1,675 | KVM |

Table 1-32 D2 ECS specifications

| Flavo r | vCPU s | Memor y (GiB) | Max./ Assured Bandwid th (Gbit/s) | Max. PPS (10,000) | Local Disks (GiB) | Max. NIC Queu es | Virtual ization |
|----------------------|-----------|---------------------|---|----------------------|-------------------------|---------------------------|--------------------|
| d2.xla rge.8 | 4 | 32 | 1/1 | 15 | 2 × 1,675 | 2 | KVM |
| d2.2xl arge. 8 | 8 | 64 | 2/2 | 30 | 4 × 1,675 | 2 | KVM |

| Flavo r | vCPU s | Memor y (GiB) | Max./ Assured Bandwid th (Gbit/s) | Max. PPS (10,000) | Local Disks (GiB) | Max. NIC Queu es | Virtual ization |
|-----------------------|-----------|---------------------|---|----------------------|-------------------------|---------------------------|--------------------|
| d2.4xl arge. 8 | 16 | 128 | 4/4 | 40 | 8 × 1,675 | 4 | KVM |
| d2.6xl arge. 8 | 24 | 192 | 6/6 | 50 | 12 × 1,675 | 6 | KVM |
| d2.8xl arge. 8 | 32 | 256 | 8/8 | 60 | 16 × 1,675 | 8 | KVM |
| d2.12 xlarg e.8 | 48 | 384 | 12/12 | 90 | 24 × 1,675 | 8 | KVM |

Performance of a Single SATA HDD Disk Attached to a D6 ECS

Table 1-33 Performance of a single SATA HDD disk attached to a D6 ECS

| Metric | Performance |
|--------------------|---------------------|
| Disk capacity | 3,600 GiB |
| Maximum throughput | 198 MBps |
| Access latency | Within milliseconds |

Specifications of a Single SAS HDD Disk Attached to a D3 ECS

Table 1-34 Specifications of a single SAS HDD disk attached to a D3 ECS

| Metric | Performance |
|--------------------|---------------------|
| Disk capacity | 1,675 GiB |
| Maximum throughput | 247 Mbit/s |
| Access latency | Within milliseconds |

Specifications of a Single SAS HDD Disk Attached to a D2 ECS

Table 1-35 Specifications of a single SAS HDD disk attached to a D2 ECS

| Metric | Performance |
|--------------------|---------------------|
| Disk capacity | 1,675 GiB |
| Maximum throughput | 230 Mbit/s |
| Access latency | Within milliseconds |

Notes on Using D6 ECSs

- Currently, the following operating systems are supported (subject to the information displayed on the console):
 - CentOS 6.8/6.9/7.2/7.3/7.4/7.5/7.6/7.7/7.8/7.9 64bit
 - SUSE Enterprise Linux Server 11 SP3/SP4 64bit
 - SUSE Enterprise Linux Server 12 SP1/SP2/SP3/SP4 64bit
 - Red Hat Enterprise Linux
 6.4/6.5/6.6/6.7/6.8/6.9/6.10/7.0/7.1/7.2/7.3/7.4/7.5/7.6/8.0 64bit
 - Windows Server 2008 R2 Enterprise 64bit
 - Windows Server 2012 R2 Standard 64bit
 - Windows Server 2016 Standard 64bit
 - Debian 8.1.0/8.2.0/8.4.0/8.5.0/8.6.0/8.7.0/8.8.0/9.0.0 64bit
 - EulerOS 2.2/2.3/2.5/2.9 64bit
 - Fedora 22/23/24/25/26/27/28 64bit
 - OpenSUSE 13.2/15.0/15.1/42.2/42.3 64bit
 - Ubuntu 20.04 64bit
- If the host where a D6 ECS is deployed is faulty, the ECS cannot be restored through live migration.
 - If the host is faulty or subhealthy and needs to be repaired, you need to stop the ECS.
 - 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.
- D6 ECSs do not support specifications modification.
- D6 ECSs do not support local disk snapshots or backups.
- D6 ECSs can use both local disks and EVS disks to store data. 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 a D6 ECS.
 - Both EVS disks and local disks can be used as data disks of a D6 ECS.
 - A maximum of 60 disks (including VBD, SCSI, and local disks) can be attached to a D6 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?**

The maximum number of disks attached to an existing D6 ECS remains unchanged.

- You can modify the **fstab** file to set automatic disk mounting at ECS start.
- The local disk data of a D6 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 D6 ECS is deleted, its local disk data will also be automatically deleted, which can take some time. As a result, a D6 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. They cannot be separately purchased after the ECS has been created. The quantity and capacity of your local disks are determined according to the specifications of your ECS.

Notes on Using D3 ECSs

- Currently, the following operating systems are supported (subject to the information displayed on the console):
 - CentOS 6.8/6.9/7.2/7.3/7.4/7.5/7.6/7.7/7.8/7.9 64bit
 - Red Hat Enterprise Linux 6.9/7.4/7.5/7.7/7.8/8.2 64bit
 - Windows Server 2008 R2 Enterprise 64bit
 - Windows Server 2012 R2 Standard 64bit
 - Windows Server 2016 Standard 64bit
 - Windows 10 Enterprise 64bit
 - Windows Server 1709 64bit
 - Windows Server 1909/2004 SAC 64bit
 - Windows Server 2012 R2 Datacenter 64bit
 - Windows Server 2012 R2 Standard 64bit
 - Windows Server 2016/2019 Standard 64bit
 - SUSE Enterprise Linux Server 12/15 64bit
 - Debian 8.6/8.10/9.0/10 64bit
 - EulerOS 2.5 64bit
 - Fedora 27/28 64bit
 - CoreOS 1298.6 64bit
 - Oracle Linux 6.9/7.3/7.4/7.5 64bit
 - Ubuntu 14.04/16.04/18.04/20.04 64bit
- 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.

Notes on Using D2 ECSs

- Currently, the following operating systems are supported (subject to the information displayed on the console):
 - CentOS 6.8/6.9/7.2/7.3/7.4/7.5/7.6/7.7/7.8/7.9 64bit
 - Red Hat Enterprise Linux 6.9/7.4/7.5/7.7/7.8/8.2 64bit
 - Windows 10 Enterprise 64bit
 - Windows Server 1709 64bit
 - Windows Server 1909/2004 SAC 64bit
 - Windows Server 2012 R2 Datacenter 64bit
 - Windows Server 2012 R2 Standard 64bit

- Windows Server 2016/2019 Standard 64bit
- SUSE Enterprise Linux Server 12/15 64bit
- Debian 8.6/8.10/9.0/10 64bit
- Fedora 27/28 64bit
- CoreOS 1298.6 64bit
- Oracle Linux 6.9/7.3/7.4/7.5 64bit
- Ubuntu 14.04/16.04/18.04/20.04 64bit
- If the host where a D2 ECS is deployed becomes faulty, the ECS cannot be migrated.
- To improve network performance, you can set the NIC MTU of a D2 ECS to 8888.
- D2 ECSs do not support specifications modification.
- D2 ECSs do not support local disk snapshots or backups.
- D2 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. 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 a D1 ECS.
 - Both EVS disks and local disks can be used as data disks of a D1 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?

| \square | NOTE |
|-----------|------|
| | |

The maximum number of disks attached to an existing D2 ECS remains unchanged.

 You are advised to use World Wide Names (WWNs), but not drive letters, in applications to perform operations on local disks to prevent drive letter drift (low probability) on Linux. Take local disk attachment as an example:

If the local disk WWN is wwn-0x50014ee2b14249f6, run the mount /dev/disk/by-id/wwn-0x50014ee2b14249f6 command.

∩ NOTE

How can I view the local disk WWN?

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

ll /dev/disk/by-id

- The local disk data of a D2 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 D2 ECS is deleted, its local disk data will also be automatically deleted, which can take some time. As a result, a D2 ECS takes a longer time than other ECSs to be deleted. Back up the data before deleting such an ECS.

- Do not store long-term service data in local disks. Instead, back up data in a timely manner and use a high availability data architecture. Use EVS disks to store service data that needs to be stored for a long time.
- 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.
- The basic resources, including vCPUs, memory, and image of a stopped D2
 ECS will continue to be billed. To stop the ECS from being billed, delete it and
 its associated resources. For details, see Will Am I Continue to Be Billed
 After ECSs Are Stopped?

Handling Damaged Local Disks Attached to an ECS of D Series

If a local disk attached to an ECS is damaged, perform the following operations to handle this issue:

For a Linux ECS:

- 1. Detach the faulty local disk.
 - a. Run the following command to query the mount point of the faulty disk:df -Th

Figure 1-5 Querying the mount point

```
l# df
Filesystem
                Type
                          Size
                                 Used Avail Usez Mounted on
devtmpfs
                devtmpfs
                          4.0M
                                    0
                                      4.0M
                                               8% /dev
tmpfs
                tmpfs
                           16G
                                    0
                                        16G
                                               0% /dev/shm
tmpfs
                tmpfs
                           16G
                                 8.6M
                                        16G
                                               1% /run
                                               0% /sys/fs/cgroup
tmpfs
                tmpfs
                          4.0M
                                    Ø
                                       4.0M
                           40G
/dev/vda1
                                 2.1G
                                        36G
                ext4
                                        16G
                                               1% /tmp
tmpfs
                tmpfs
                            16G
                                  32K
/dev/sda1
                            .7T
                                        1.6T
                                               1% /mnt/sda1
```

b. Run the following command to detach the faulty local disk:

umount *Mount point*

In the example shown in **Figure 1-5**, the mount point of **/dev/sda1** is **/mnt/sda1**. Run the following command:

umount /mnt/sda1

- Check whether the mount point of the faulty disk is configured in /etc/fstab
 of the ECS. If yes, comment out the mount point to prevent the ECS from
 entering the maintenance mode upon ECS startup after the faulty disk is
 replaced.
 - a. 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/sda1 partition:

blkid /dev/sda1

Information similar to the following is displayed:

/dev/sda1: UUID="b9a07b7b-9322-4e05-ab9b-14b8050cd8cc" TYPE="ext4"

b. Run the following command to check whether **/etc/fstab** contains the automatic mounting information about the disk partition:

cat /etc/fstab

Information similar to the following is displayed:

UUID=b9a07b7b-9322-4e05-ab9b-14b8050cd8cc /mnt ext4 defaults 0 0

- If the mounting information exists, perform the following steps to delete it.
 - i. Run the following command to edit /etc/fstab:

vi /etc/fstab

Use the UUID obtained in **2.a** to check whether the mounting information of the local disk is contained in **/etc/fstab**. If yes, comment out the information. This prevents the ECS from entering the maintenance mode upon ECS startup after the local disk is replaced.

- ii. Press i to enter editing mode.
- Delete or comment out the automatic mounting information of the disk partition.
 - For example, add a pound sign (#) at the beginning of the following command line to comment out the automatic mounting information: # UUID=b9a07b7b-9322-4e05-ab9b-14b8050cd8cc /mnt ext4 defaults 0 0
- iv. Press **Esc** to exit editing mode. Enter :**wq** and press **Enter** to save the settings and exit.
- Run the following command to obtain the WWN of the local disk:
 For example, if the sdc disk is faulty, obtain the WWN of the sdc disk.
 Il /dev/disk/by-id/ | grep wwn-

Figure 1-6 Querying the WWN of the faulty local disk

4. Stop the ECS and provide the WWN of the faulty disk to technical support personnel to replace the local disk.

After the local disk is replaced, restart the ECS to synchronize the new local disk information to the virtualization layer.

For a Windows ECS:

- Open Computer Management, choose Computer Management (Local) > Storage > Disk Management, and view the disk ID, for example, Disk 1.
- Open Windows PowerShell as an administrator and obtain the serial number of the faulty disk according to the mapping between the disk ID and serial number.

Get-Disk | select Number, SerialNumber

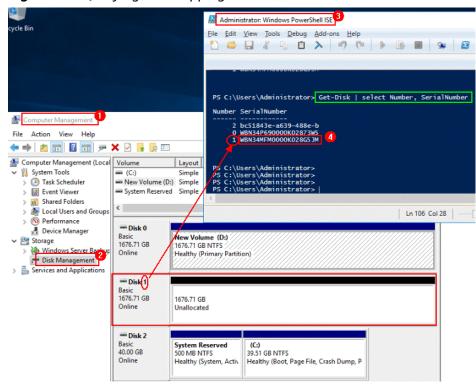


Figure 1-7 Querying the mapping between the disk ID and serial number

3. Stop the ECS and provide the serial number of the faulty disk to technical support personnel to replace the local disk.

After the local disk is replaced, restart the ECS to synchronize the new local disk information to the virtualization layer.

1.7.6 Ultra-high I/O ECSs

Overview

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

Scenarios

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

Specifications

Table 1-36 I3 ECS specifications

| Flavo r | vCPU s | Memor y (GiB) | Max./ Assured Bandwi dth (Gbit/s) | Max. PPS (10,000 | Max. NIC Queu es | Max NIC s | Local Disks | Virtua lizatio n |
|-----------------------|-----------|---------------------|---|------------------------|---------------------------|-----------------|-----------------------------|------------------------|
| i3.2xl arge.8 | 8 | 64 | 2.5/2.5 | 100 | 4 | 4 | 1 × 1,600 GiB NVMe | KVM |
| i3.4xl arge.8 | 16 | 128 | 5/5 | 150 | 4 | 8 | 2 × 1,600 GiB NVMe | KVM |
| i3.8xl arge.8 | 32 | 256 | 10/10 | 200 | 8 | 8 | 4 × 1,600 GiB NVMe | KVM |
| i3.12x large. 8 | 48 | 384 | 15/15 | 240 | 8 | 8 | 6 × 1,600 GiB NVMe | KVM |
| i3.15x large. 8 | 60 | 512 | 25/25 | 500 | 16 | 8 | 7 × 1,600 GiB NVMe | KVM |

Features

Table 1-37 and **Table 1-38** list the IOPS performance of local disks and specifications of a single local disk attached to an I3 ECS.

Table 1-37 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 |

| Flavor | Maximum IOPS for Random 4 KB Read |
|---------------|-----------------------------------|
| i3.15xlarge.8 | 5,250,000 |
| i3.16xlarge.8 | 6,000,000 |

Table 1-38 Specifications of a single I3 local disk

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

Notes

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

∩ NOTE

EulerOS 2.2 and Ubuntu Server 16.04 are recommended.

- If the host where an ultra-high I/O ECS is deployed is faulty, the ECS cannot be restored through live migration.
 - If the host is faulty or subhealthy, you need to stop the ECS for hardware repair.
 - In case of system maintenance or hardware faults, the ECS will be redeployed (to ensure HA) and cold migrated to another host. The local disk data of the ECS will not be retained.
- Ultra-high I/O ECSs do not support specifications change.
- Ultra-high I/O ECSs do not support local disk snapshots or backups.
- Ultra-high I/O ECSs can use local disks, and can also have EVS disks attached to provide a larger storage size. Note the following when using the two types of storage media:

- Only an EVS disk, not a local disk, can be used as the system disk of an ultra-high I/O ECS.
- Both EVS disks and local disks can be used as data disks of an ultra-high I/O ECS.
- An ultra-high I/O ECS can have a maximum of 60 attached disks (including VBD, SCSI, and local disks). An ECS can have a maximum of 60 attached disks, including the system disk. For details about constraints, see Can I Attach Multiple Disks to an ECS?
- 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.
- Local disks attached to Ir3 ECSs can be split for multiple ECSs to use. If a local disk is damaged, the ECSs that use this disk will be affected.
 - You are advised to add Ir3 ECSs to an ECS group during the creation process to prevent such failures. For details, see **Managing ECS Groups**.
- The basic resources, including vCPUs, memory, and image of an ultra-high I/O ECS will continue to be billed after the ECS is stopped. To stop the ECS from being billed, delete it and its associated resources. For more information, see Will Am I Continue to Be Billed After ECSs Are Stopped?

Handling Damaged Local Disks Attached to an ECS of I Series

If a local disk attached to an ECS is damaged, perform the following operations to handle this issue:

For a Linux ECS:

- 1. Detach the faulty local disk.
 - Run the following command to query the mount point of the faulty disk:
 df -Th

Figure 1-8 Querying the mount point

```
1# df -Th
                          Size
Filesystem
                                Used Avail Use: Mounted on
               Type
devtmpfs
                          4.0M
                                      4.0M
               devtmpfs
                                    0
                                               8% /dev
tmpfs
               tmpfs
                           16G
                                    A
                                        16G
                                               8% /dev/shm
tmpfs
               tmpfs
                           16G
                                8.6M
                                        16G
                                               1% /run
tmpfs
               tmpfs
                          4.0M
                                   A
                                       4.0M
                                               0% /sys/fs/cgroup
/dev/vda1
                           60G
                                 2.4G
                                        54G
                                               5% /
               ext4
tmpfs
                tmpfs
                           16G
                                  32K
                                        16G
                                               1% /tmp
dev/nume@n1
                          1.5T
                                 28K
                                       1.4T
                                               1% /mnt/nume0
```

b. Run the following command to detach the faulty local disk:

umount *Mount point*

In the example shown in **Figure 1-8**, the mount point of **/dev/nvme0n1** is **/mnt/nvme0**. Run the following command:

umount /mnt/nvme0

- Check whether the mount point of the faulty disk is configured in /etc/fstab
 of the ECS. If yes, comment out the mount point to prevent the ECS from
 entering the maintenance mode upon ECS startup after the faulty disk is
 replaced.
 - a. 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/nvme0n1 partition:

blkid /dev/nvme0n1

Information similar to the following is displayed:

/dev/nvme0n1: UUID="b9a07b7b-9322-4e05-ab9b-14b8050cd8cc" TYPE="ext4"

b. Run the following command to check whether **/etc/fstab** contains the automatic mounting information about the disk partition:

cat /etc/fstab

Information similar to the following is displayed:

```
UUID=b9a07b7b-9322-4e05-ab9b-14b8050cd8cc /mnt ext4 defaults 0 0
```

- c. If the mounting information exists, perform the following steps to delete it.
 - i. Run the following command to edit /etc/fstab:

vi /etc/fstab

Use the UUID obtained in **2.a** to check whether the mounting information of the local disk is contained in **/etc/fstab**. If yes, comment out the information. This prevents the ECS from entering the maintenance mode upon ECS startup after the local disk is replaced.

- ii. Press i to enter editing mode.
- iii. Delete or comment out the automatic mounting information of the disk partition.

For example, add a pound sign (#) at the beginning of the following command line to comment out the automatic mounting information: # UUID=b9a07b7b-9322-4e05-ab9b-14b8050cd8cc /mnt ext4 defaults 0 0

iv. Press **Esc** to exit editing mode. Enter :wq and press **Enter** to save the settings and exit.

3. Run the following command to obtain the SN of the local disk:

For example, if the nvme0n1 disk is faulty, obtain the serial number of the nvme0n1 disk.

ll /dev/disk/by-id/

Figure 1-9 Querying the serial number of the faulty local disk

4. Stop the ECS and provide the serial number of the faulty disk to technical support personnel to replace the local disk.

After the local disk is replaced, restart the ECS to synchronize the new local disk information to the virtualization layer.

For a Windows ECS:

- Open Computer Management, choose Computer Management (Local) > Storage > Disk Management, and view the disk ID, for example, Disk 1.
- 2. Open Windows PowerShell as an administrator and run the following command to query the disk on which the logical disk is created:

Get-CimInstance -ClassName Win32_LogicalDiskToPartition |select | Antecedent, Dependent | fl

Figure 1-10 Querying 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 #1, Partition #1")
Dependent: Win32_LogicalDisk (DeviceID = "C:")
```

3. Run the following command to obtain the serial number of the faulty disk according to the mapping between the disk ID and serial number:

Get-Disk | select Number, SerialNumber

Figure 1-11 Querying the mapping between the disk ID and serial number

```
PS C:\Vsers\Administrator> Get-Disk | select Number, SerialNumber

Number SerialNumber

0 0100 0000 0000 0000 0022 A100 30A4 0D5A.
1 2e38cae8-85b9-436b-b
```

Ⅲ NOTE

If the serial number cannot be obtained by running the preceding command, see Using a Serial Number to Obtain the Disk Name (Windows).

4. Stop the ECS and provide the serial number of the faulty disk to technical support personnel to replace the local disk.

After the local disk is replaced, restart the ECS to synchronize the new local disk information to the virtualization layer.

1.7.7 GPU-accelerated ECSs

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

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

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

GPU-accelerated ECS Types

Recommended: Computing-accelerated P2s and Inference-accelerated Pi2

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

- G series
 - Graphics-accelerated Enhancement G6
 - Graphics-accelerated Enhancement G5
 - Graphics-accelerated Enhancement G3
- P series
 - Computing-accelerated P2s (recommended)
 - Computing-accelerated P2
 - Inference-accelerated Pi2 (recommended)

Helpful links:

- Installing a GRID Driver on a GPU-accelerated ECS
- Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS

Images Supported by GPU-accelerated ECSs

Table 1-39 Images supported by GPU-accelerated ECSs

| Category | ECS Type | Supported Image | |
|-------------|----------|------------------------------------|----------------------|
| Graphics- | G6 | CentOS 8.2 64bit | |
| accelerated | | CentOS 7.6 64bit | |
| | | | • Ubuntu 20.04 64bit |
| | | | • Ubuntu 18.04 64bit |
| | | Windows Server 2019 Standard 64bit | |
| | | Windows Server 2016 Standard 64bit | |

| Category | ECS Type | Supported Image |
|---------------------------|----------|--|
| Graphics- accelerated | G5 | CentOS 8.2 64bit CentOS 7.6 64bit CentOS 7.5 64bit Ubuntu 20.04 64bit Ubuntu 18.04 64bit Windows Server 2019 Standard 64bit Windows Server 2016 Standard 64bit Windows Server 2019 Datacenter 64bit Windows Server 2016 Datacenter 64bit |
| Graphics- accelerated | G3 | Windows Server 2019 Standard 64bitWindows Server 2019 Standard 64bit |
| Computing- accelerated | P2s | Windows Server 2016 Standard 64bit |
| Computing- accelerated | P2 | Ubuntu 16.04 Server 64bit |
| Inference- accelerated | Pi2 | CentOS 7.5 64bit Windows Server 2019 Standard 64bit Windows Server 2016 Standard 64bit |

Graphics-accelerated Enhancement G6

Overview

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

Select your desired GPU-accelerated ECS type and specifications.

| Flav or | vCP Us | Mem ory (GiB) | Max./ Assur ed Band width (Gbit/ s) | Max. PPS (10,00 0) | Ma x. NIC Que ues | Max. NICs | GPUs | GPU Memo ry (GiB) | Virtu alizat ion |
|---------------------------|-----------|---------------------|---|-----------------------------|-------------------------------|--------------|--------|----------------------------|------------------------|
| g6.4 xlar ge.4 | 16 | 64 | 15/8 | 200 | 8 | 8 | 1 × T4 | 16 | KVM |
| g6.6 xlar ge.4 | 24 | 96 | 25/15 | 200 | 8 | 8 | 1 × T4 | 16 | KVM |
| g6.9 xlar ge.7 | 36 | 252 | 25/15 | 200 | 16 | 8 | 1 × T4 | 16 | KVM |
| g6.1 8xla rge. 7 | 72 | 504 | 30/30 | 400 | 32 | 16 | 2 × T4 | 32 | KVM |

Table 1-40 G6 ECS specifications

G6 ECS Features

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

Supported Common Software

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

AutoCAD

- 3DS MAX
- MAYA
- Agisoft PhotoScan
- ContextCapture

Notes

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

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

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

For details about how to configure a GRID license, see **Installing a GRID Driver on a GPU-accelerated ECS**.

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

For details, see Installing a GRID Driver on a GPU-accelerated ECS.

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

Graphics-accelerated Enhancement G5

Overview

G5 ECSs use NVIDIA GRID vGPUs and provide comprehensive, professional graphics acceleration. They use NVIDIA Tesla V100 GPUs and support DirectX, OpenGL, and Vulkan. These ECSs provide 1 GiB of GPU memory, meeting requirements from entry-level through professional graphics processing.

Select your desired GPU-accelerated ECS type and specifications.

| Flav or | vCPUs | Memor y (GiB) | Max./ Assured Bandwi dth (Gbit/s) | Max. PPS (10,000) | Max. NIC Que ues | GPU s | GPU Memor y (GiB) | Virtua lizatio n | |
|----------------------|-------|---------------------|---|-------------------------|---------------------------|-------------|----------------------------|------------------------|--|
| g5.xl arge. 2 | 4 | 8 | 4/1.3 | 20 | 2 | V100 -1Q | 1 | KVM | |
| g5.2 xlarg e.2 | 8 | 16 | 6/2 | 35 | 4 | V100 -2Q | 2 | KVM | |
| g5.4 xlarg e.4 | 16 | 64 | 10/4 | 50 | 8 | V100 -8Q | 8 | KVM | |

Table 1-41 G5 ECS specifications

□ NOTE

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

G5 ECS Features

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

Supported Common Software

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

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

Notes

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

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

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

For details about how to configure a GRID license, see **Installing a GRID Driver on a GPU-accelerated ECS**.

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

For details, see Installing a GRID Driver on a GPU-accelerated ECS.

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

Graphics-accelerated Enhancement G3

Overview

G3 ECSs are based on PCI passthrough and exclusively use GPUs for professional graphics acceleration. In addition, G3 ECSs use NVIDIA Tesla M60 GPUs and support DirectX and OpenGL with up to 16 GiB of GPU memory and 4096 x 2160 resolution. They are ideal for professional graphics workstations.

| Flav or | vCPUs | Memor y (GiB) | Max./ Assured Bandwi dth (Gbit/s) | Max. PPS (10,000) | Max. NIC Que ues | GPU s | GPU Memor y (GiB) | Virtua lizatio n |
|----------------------|-------|---------------------|---|-------------------------|---------------------------|------------|----------------------------|------------------------|
| g3.4 xlarg e.4 | 16 | 64 | 8/2.5 | 50 | 2 | 1 x M60 | 1 x 8 | KVM |
| g3.8 xlarg e.4 | 32 | 128 | 10/5 | 100 | 4 | 2 x M60 | 2 x 8 | KVM |

Table 1-42 G3 ECS specifications

G3 ECS Features

- CPU: Intel® Xeon® E5-2697 v4 processors (2.3 GHz of base frequency and 3.5 GHz of turbo frequency)
- Provide professional graphics acceleration APIs
- NVIDIA M60 GPUs
- Graphics applications accelerated
- GPU passthrough
- Automatic scheduling of G3 ECSs to AZs where NVIDIA M60 GPUs are used
- A maximum specification of 16 GiB of GPU memory and 4096 x 2160 resolution for processing graphics and videos

Notes

 After a G3 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.

Resources will be released after a G3 ECS is stopped. If resources are insufficient at the next start, the start may fail. If you want to use such an ECS for a long period of time, do not stop the ECS.

- When the Windows OS running on a G3 ECS is started, the GRID driver is loaded by default, and vGPUs are used for video output by default. In such a case, the remote login function provided on the management console is not supported. To access such an ECS, use RDP, such as Windows MSTSC. Then, install a third-party VDI tool on the ECS for remote login, such as VNC.
- By default, G3 ECSs created using a public image have had the GRID driver of a specific version installed.
- If a G3 ECS is created using a private image, make sure that the GRID driver
 was installed during the private image creation. If not, install the driver for
 graphics acceleration after the ECS is created. For details, see Installing a
 GRID Driver on a GPU-accelerated ECS.

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

Computing-accelerated P2s

Overview

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

Specifications

Table 1-43 P2s ECS specifications

| Flavo r | vCP Us | Me mor y (GiB | Max./ Assure d Bandw idth (Gbit/s) | Max. PPS (10,0 00) | Max NIC Que ues | Ma x. NIC s | GP Us | GPU Con nect ion | GPU Mem ory (GiB) | Virtu alizat ion |
|------------------------|-----------|------------------------|---|-----------------------------|--------------------------|----------------------|-----------------|---------------------------|--------------------------------|------------------------|
| p2s.2 xlarg e.8 | 8 | 64 | 10/4 | 50 | 4 | 4 | 1 × V10 0 | PCle Gen 3 | 1 × 32 GiB | KVM |
| p2s.4 xlarg e.8 | 16 | 128 | 15/8 | 100 | 8 | 8 | 2 × V10 0 | PCle Gen 3 | 2 × 32 GiB | KVM |
| p2s.8 xlarg e.8 | 32 | 256 | 25/15 | 200 | 16 | 8 | 4 × V10 0 | PCle Gen 3 | 4 × 32 GiB | KVM |
| p2s.1 6xlar ge.8 | 64 | 512 | 30/30 | 400 | 32 | 8 | 8 × V10 0 | PCle 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.

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.

Computing-accelerated P2

Overview

P2 ECSs use NVIDIA Tesla V100 GPUs, which have improved both single- and double-precision computing capabilities by 50% than the performance of the previous-generation GPUs and offer 112 TFLOPS of deep learning.

Specifications

Table 1-44 P2 ECS specifications

| Flav or | vC P Us | Me mor y (Gi B) | Max./ Assure d Bandw idth (Gbit/s) | Max. PPS (10,00 0) | Ma x. NIC Qu eue s | Max NICs | GPU s | GPU Memor y (GiB) | Loc al Disk s (GiB | Virtu aliza tion |
|----------------------|---------------|-----------------------------|---|-----------------------------|-----------------------------------|-------------|-----------------|----------------------------|--------------------------------|------------------------|
| p2.2 xlar ge.8 | 8 | 64 | 10/4 | 50 | 2 | 12 | 1 × V10 0 | 1 × 16 | 1 × 800 | KVM |
| p2.4 xlar ge.8 | 16 | 128 | 15/8 | 100 | 4 | 12 | 2 × V10 0 | 2 × 16 | 2 × 800 | KVM |
| p2.8 xlar ge.8 | 32 | 256 | 25/15 | 200 | 8 | 12 | 4 × V10 0 | 4 × 16 | 4 × 800 | KVM |

P2 ECS Features

- CPU: Intel® Xeon® Processor E5-2690 v4 (2.6 GHz)
- NVIDIA Tesla V100 GPUs
- GPU hardware passthrough
- 14 TFLOPS of single-precision computing, 7 TFLOPS of double-precision computing, and 112 TFLOPS of deep learning
- Maximum network bandwidth of 12 Gbit/s
- 16 GiB of HBM2 GPU memory with a bandwidth of 900 Gbit/s
- 800 GiB NVMe SSDs for temporary local storage
- 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, P2 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 P2 ECSs.

Supported Common Software

P2 ECSs are used in computing acceleration scenarios, such as deep learning training, inference, scientific computing, molecular modeling, and seismic analysis. If the software requires GPU CUDA parallel computing, use P2 ECSs. P2 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

- P2 ECSs have local NVMe SSDs attached, which will continue to be billed after the ECSs are stopped. To stop the ECS from being billed, delete it and its associated resources. For more information, see Will Am I Continue to Be Billed After ECSs Are Stopped?
- The local NVMe SSDs attached to P2 ECSs are dedicated for services with strict requirements on storage I/O performance, such as deep learning training and HPC. Local disks are attached to the ECSs of specified flavors and cannot be separately bought. In addition, you are not allowed to detach a local disk and then attach it to another ECS.

□ NOTE

Data may be lost on the local NVMe SSDs attached to P2 ECSs due to a fault, for example, due to a disk or host fault. Therefore, you are suggested to store only temporary data in local NVMe SSDs. If you store important data in such a disk, securely back up the data.

- P2 ECSs do not support specifications modification.
- P2 ECSs do not support automatic recovery.
- After you delete a P2 ECS, the data stored in local NVMe SSDs is automatically cleared.
- By default, P2 ECSs created using a public image have the Tesla driver installed.
- If a P2 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 1-45 Pi2 ECS specifications

| Flav | vC P Us | Me mo ry (Gi B) | Max./ Assure d Band width (Gbit/ s) | Max. PPS (10,0 00) | Max. NIC Queu es | Max. NICs | GPU s | GP U Me mo ry (Gi B) | Loc al Dis ks | Virtua lizatio n |
|---------------------------|---------------|-----------------------------|---|-----------------------------|---------------------------|--------------|-----------|--|------------------------|------------------------|
| pi2. 2xla rge. 4 | 8 | 32 | 10/4 | 50 | 4 | 4 | 1 × T4 | 1 × 16 | N/A | KVM |
| pi2. 4xla rge. 4 | 16 | 64 | 15/8 | 100 | 8 | 8 | 2 × T4 | 2 × 16 | N/A | KVM |
| pi2. 8xla rge. 4 | 32 | 128 | 25/15 | 200 | 16 | 8 | 4 × T4 | 4 × 16 | N/A | KVM |

Pi2 ECS Features

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

Supported Common Software

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

Pi2 ECSs support the following commonly used software:

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

Notes

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

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.

1.8 Discontinued ECS Specifications

Disk-intensive D1

Table 1-46 D1 ECS specifications

| Flavor | vCPUs | Memory (GiB) | Max./ Assured Bandwidt h | Max. PPS | Local Disks (GiB) | Virtuali zation |
|----------------|-------|-----------------|-----------------------------------|----------|-------------------------|--------------------|
| d1.xlarg e | 6 | 55 | Medium | Medium | 3 × 1,675 | Xen |
| d1.2xlar ge | 12 | 110 | Medium | Medium | 6 × 1,675 | Xen |
| d1.4xlar ge | 24 | 220 | Medium | Medium | 12 × 1,675 | Xen |
| d1.8xlar ge | 48 | 440 | Medium | Medium | 24 × 1,675 | Xen |

High-Performance Computing H1

Table 1-47 H1 ECS specifications

| Flavor | vCPU s | Memory (GiB) | Max./ Assured Bandwidt h (Gbit/s) | Max. PPS (10,000) | Max. NIC Queu es | Max. NICs | Virtu alizat ion |
|------------------|-----------|-----------------|---|----------------------|---------------------------|--------------|------------------------|
| h1.large. 2 | 2 | 4 | 1.5/0.5 | 10 | 1 | 3 | KVM |
| h1.xlarge .2 | 4 | 8 | 3/1 | 15 | 1 | 3 | KVM |
| h1.2xlarg e.2 | 8 | 16 | 5/2 | 30 | 2 | 4 | KVM |
| h1.4xlarg e.2 | 16 | 32 | 8/4 | 40 | 4 | 8 | KVM |
| h1.8xlarg e.2 | 32 | 64 | 13/8 | 60 | 8 | 8 | KVM |
| h1.large. 4 | 2 | 8 | 1.5/0.5 | 10 | 1 | 3 | KVM |
| h1.xlarge .4 | 4 | 16 | 3/1 | 15 | 1 | 3 | KVM |
| h1.2xlarg e.4 | 8 | 32 | 5/2 | 30 | 2 | 4 | KVM |
| h1.4xlarg e.4 | 16 | 64 | 8/4 | 40 | 4 | 8 | KVM |
| h1.8xlarg e.4 | 32 | 128 | 13/8 | 60 | 8 | 8 | KVM |

Graphics-accelerated G2

Table 1-48 G2 ECS specifications

| Flavo r | vCPUs | Memor y (GiB) | Max./ Assured Bandwid th | Max. PPS | GPU | GPU Memor y (GiB) | Virtual ization |
|----------------|-------|------------------|-----------------------------------|----------|------------|-------------------------|--------------------|
| g2.2xl arge | 8 | 64 | Medium | Medium | 1 × M60 | 8 | Xen |

Graphics-accelerated G1

Table 1-49 G1 ECS specifications

| Flavo r | vCPUs | Memor y (GiB) | Max./ Assured Bandwid th | Max. PPS | GPU | GPU Memor y (GiB) | Virtual ization |
|----------------|-------|---------------------|-----------------------------------|----------|-------------------|----------------------------|--------------------|
| g1.xla rge | 4 | 8 | Medium | Medium | 1 × M60-1 Q | 1 | Xen |
| g1.2xl arge | 8 | 16 | Medium | Medium | 1 × M60-1 Q | 1 | Xen |
| g1.4xl arge | 16 | 32 | Medium | Medium | 1 × M60-4 Q | 4 | Xen |

Memory-optimized M1

Table 1-50 M1 ECS specifications

| Flavor | vCPUs | Memory (GiB) | Max./Assured Bandwidth | Max. PPS | Virtualiz ation |
|------------|-------|-----------------|---------------------------|----------|--------------------|
| m1.large | 2 | 16 | Low | Low | Xen |
| m1.xlarge | 4 | 32 | Medium | Medium | Xen |
| m1.2xlarge | 8 | 64 | Medium | Medium | Xen |
| m1.4xlarge | 16 | 128 | Medium | Medium | Xen |

General-Purpose S1

Table 1-51 S1 ECS specifications

| Flavor | vCPUs | Memory (GiB) | Max./Assured Bandwidth | Max. PPS | Virtualiz ation |
|------------|-------|-----------------|---------------------------|----------|--------------------|
| s1.medium | 1 | 4 | Low | Low | Xen |
| s1.large | 2 | 8 | Low | Low | Xen |
| s1.xlarge | 4 | 16 | Medium | Medium | Xen |
| s1.2xlarge | 8 | 32 | Medium | Medium | Xen |

| Flavor | vCPUs | Memory (GiB) | Max./Assured Bandwidth | Max. PPS | Virtualiz ation |
|------------|-------|-----------------|---------------------------|----------|--------------------|
| s1.4xlarge | 16 | 64 | Medium | Medium | Xen |
| s1.8xlarge | 32 | 128 | Medium | Medium | Xen |

General-Purpose C1

Table 1-52 C1 ECS specifications

| Flavor | vCPUs | Memory (GiB) | Max./Assured Bandwidth | Max. PPS | Virtualiz ation |
|------------|-------|-----------------|---------------------------|----------|--------------------|
| c1.large | 2 | 2 | Low | Low | Xen |
| c1.xlarge | 4 | 4 | Medium | Medium | Xen |
| c1.2xlarge | 8 | 8 | Medium | Medium | Xen |
| c1.4xlarge | 16 | 16 | Medium | Medium | Xen |

Computing II

Table 1-53 Computing II (C2) ECS specifications

| Flavor | vCPUs | Memory (GiB) | Max./ Assured Bandwidth | Max. PPS | Virtualizati on |
|------------|-------|-----------------|-------------------------------|----------|--------------------|
| c2.large | 2 | 4 | Low | Low | Xen |
| c2.xlarge | 4 | 8 | Medium | Medium | Xen |
| c2.2xlarge | 8 | 16 | Medium | Medium | Xen |

1.9 Images

1.9.1 Image Types

What Is Image?

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

Images can be public or private. Public images are provided by the system by default, and private images are manually created. You can use any type of image

to create an ECS. You can also create a private image using an existing ECS or external image. This provides you with a simple and fast way to create ECSs tailored to your needs. For example, if you use web services, your image can contain web server configurations, static configurations, and dynamic page code. After you use this image to create an ECS, the web server will run on the created ECS.

Image Types

| Image Type | Description |
|------------------|---|
| Public image | A public image is a standard, widely used image. It contains an OS and preinstalled public applications and is available to all users. Public images are very stable and their OS and any included software have been officially authorized for use. If a public image does not contain the environments or software you need, you can use a public image to create an ECS and then deploy the required environments or software on it. |
| Private image | 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. |
| | A private image can be a system disk image, data disk image, ISO image, or full-ECS image. |
| | 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 of the same size. |
| Shared image | A shared image is a private image another user has shared with you. |
| | For more information about shared images, see "Sharing Images" in <i>Image Management Service User Guide</i> . |

1.9.2 Cloud-Init

Cloud-Init is an open-source cloud initialization program, which initializes some of the customized configurations of a newly created ECS, such as the hostname, 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 created using a private image support customized 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. For details, see Passing User Data to ECSs.
- If Cloud-Init is supported, ECSs do not support password authentication anymore. All newly created ECSs use key pair authentication. This change will influence your ECS logins. For details, see the following sections:
 - Login Overview
 - What Is the cloudbase-init Account in Windows ECSs Used for?
 - Why Does the Login to My Linux ECS Using a Key File Fail?
 - Why Does the System Display a Message Indicating that the Password for Logging In to a Windows ECS Cannot Be Obtained?
- If Cloud-Init is supported, you can view and use metadata to configure and manage running ECSs. For details, see **Obtaining Metadata**.

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?**

Notes

- When using Cloud-Init, enable DHCP in the VPC to which the ECS belongs.
- When using Cloud-Init, ensure that security group rules for the outbound direction meet the following requirements:

Protocol: TCP

Port: 80

- Destination: 169.254.0.0/16

□ NOTE

If you use the default security group rules for the outbound direction, the metadata can be accessed because the default rules meet the preceding requirements. For details about the default security group rules for the outbound direction, see **Default Security Group and Rules**.

1.10 EVS Disks

What Is Elastic Volume Service?

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

Disk Types

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

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

Device Types

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

VBD

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

SCSI

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

■ NOTE

For more information about how to use SCSI EVS disks, for example, how to install a driver for SCSI EVS disks, see **Device Types and Usage Instructions**.

Helpful Links

Which ECSs Can Be Attached with SCSI EVS Disks?

1.11 Network

VPC

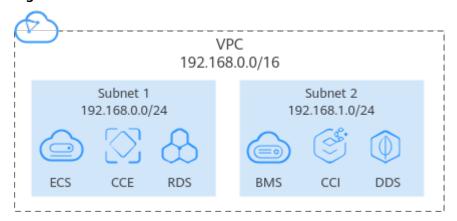
Virtual Private Cloud (VPC) allows you to create customized virtual networks in your logically isolated AZ. Such networks are dedicated zones that are logically isolated, providing secure network environments for your ECSs. You can define security groups, virtual private networks (VPNs), IP address segments,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

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

Subnet

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

Figure 1-12 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.

Region VPC 192.168.0.0/16 Subnet 1 Subnet 2 192.168.0.0/24 192.168.1.0/24 Custom security Default security group Default security group Allow all group outbound traffic Mutual ECS ECS ECS inbound traffic

Figure 1-13 Default security group

Table 1-54 describes default security group rules.

Table 1-54 Default security group rules

| Directi on | Protoc ol | Port/ Range | Source/ Destination | Description |
|---------------|--------------|----------------|--|--|
| Outbo und | All | All | Destination: 0.0.0.0/0 | Allows all outbound traffic. |
| Inboun d | All | All | Source: the current security group (for example, sg-xxxxx) | Allows communications among ECSs within the security group and denies all inbound traffic (incoming data packets). |

EIP

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

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

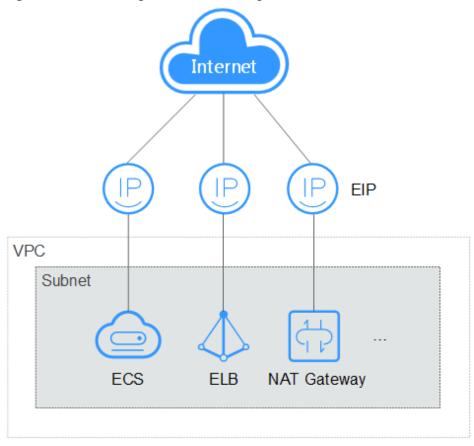


Figure 1-14 Accessing the Internet using an EIP

1.12 Security

1.12.1 Identity Authentication and Access Control

1.12.1.1 Access Control for ECS

IAM Identity Authentication

IAM provides fine-grained permissions management, user identity authentication, and resource access control.

You can use your account to create IAM users, and assign permissions to the IAM users to control their access to specific resources. IAM permissions define which actions on your cloud resources are allowed or denied.

- For details about permissions management, see Permissions Management.
- For details about how to grant ECS permissions, see Creating a User and Granting ECS Permissions.
- For details about custom policies, see ECS Custom Policies.
- For details about policies and supported actions, see Permissions Policies and Supported Actions.

Access Control

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.

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.

For details about how to configure a VPC and security group, see **Step 2**: **Configure Network**.

1.12.2 Data Protection

1.12.2.1 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 *Image Management Service User Guide*.

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 added data disks.

For more information about EVS disk encryption, see *Elastic Volume Service User Guide*.

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 used for encryption relies on the Key Management Service (KMS). KMS uses a data encryption key (DEK) to encrypt data and uses a customer master key (CMK) to encrypt the DEK.

Figure 1-15 Data encryption process

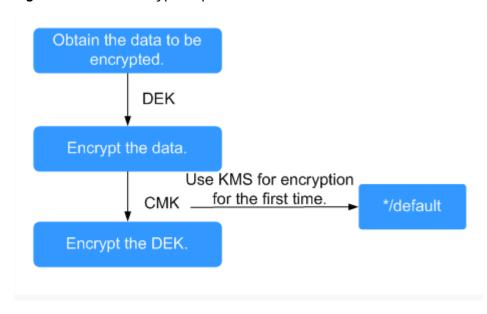


Table 1-55 describes the keys involved in the data encryption process.

Table 1-55 Keys

| Name | Description | Function |
|-------------|--|---|
| DEK | An encryption key that is used for encrypting data. | Encrypts specific data. |
| Custom key | An encryption key created using KMS 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 KMS for encryption for the first time. | Supports query of the default key on the KMS console. |
| | The name extension of a default CMK is /default, for example, evs/default. | Does not support CMK disabling or scheduled deletion. |

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. Therefore, the EVS disk becomes unavailable.

For details about KMS, see Key Management Service User Guide.

1.12.3 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 1-56 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) |
| Recommend ed 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. |

Helpful Links

- Creating a Server Backup Vault
- Creating a Disk Backup Vault
- Restoring Data Using a Cloud Server Backup
- Using a Backup to Create an Image
- Restoring Data Using a Cloud Disk Backup

1.12.4 License Types

Using License from the System

You can use OS licenses provided by the cloud platform. After creating an ECS with a license authorized, you can use the authorized OS. The platform manages license compliance for you.

BYOL

What Is BYOL?

Bring your own license (BYOL) allows you to use your existing OS license. In such a case, you do not need to apply for a license again. In BYOL license type, you do not need to pay for the license fee when creating an ECS.

How to Use BYOL?

If you select the BYOL license type, you are required to manage licenses by yourself. The cloud platform provides functions for you to maintain license compliance during the license lifecycle. If you have obtained an OS license, you do not need to apply for a license.

The OSs supporting BYOL are as follows:

Application Scenarios

The system does not support dynamic license type changing. ECSs support BYOL in the following scenarios:

Creating an ECS

After creating an ECS, you cannot change its license type. If the license type must be changed, reinstall or change the ECS OS.

Reinstalling an ECS OS

When reinstalling an ECS OS, you can set the license type for the ECS.

Changing an ECS OS

When changing an ECS OS, you can set the license type for the ECS.

• Attaching a system disk

The license type of a system disk is determined by the ECS license type after the ECS is created, the ECS OS is reinstalled, or the ECS OS is changed. If the system disk is detached and then attached to a new ECS or the original ECS, ensure that the ECS license type is the same as the system disk license type.

1.13 Billing

Billing Items

ECSs are billed based on ECS specifications and service duration.

Table 1-57 Billing details

| Billing Item | Description |
|----------------------|---|
| ECS | ECSs are billed by instance type and flavor (vCPUs and memory), service duration, and purchased ECS quantity. |
| | For pricing details, see Elastic Cloud Server Price Calculator . |
| Image | Public images of the community edition, such as Linux, are free of charge. Other commercial images, such as Windows images, are billed. |
| EVS disk (mandatory) | EVS disks are billed by disk capacity on a pay-per-use basis. For pricing details, see Elastic Volume Service Price Calculator. You are advised to select the same duration of EVS disks as ECSs. |
| EIP (optional) | A public IP address is required for public accessibility. For pricing details, see Elastic IP Price Calculator . |
| Bandwidth (optional) | An EIP can be billed by bandwidth or traffic. For pricing details, see Bandwidth Price Calculator . |

Billing Modes

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

- Pay-per-use: a flexible mode with the billing accurately down to the second. An ECS is billed from the time when it is provisioned to the time when it is deleted.
 - Common ECSs refer to ECSs without local disks or FPGAs attached. After a common ECS is stopped, it is billed as follows:
 - Basic resources (vCPUs, memory, and image) are not billed. Associated resources such as its EVS disks, EIPs, and bandwidth are billed.
 - When you try to start the ECS the next time, the system will allocate vCPUs and memory again, but if resources are insufficient, the startup may fail. In this case, you can try again later or resize the ECS specifications before attempting to start it.

Special pay-per-use ECSs will continue to be billed after being stopped. After a special ECS is stopped, its resources such as vCPUs and memory will be retained.

Special ECSs include:

- BMSs
- ECSs attached with local disks, such as disk-intensive ECSs and ultra-high I/O ECSs
- FPGA-based ECSs

To stop billing for special ECSs, delete them and their associated resources.

RI: ECSs will be billed for one year, two years, three years, five years, without
or with all upfront payment. For pricing details, see Elastic Cloud Server
Price Calculator and obtain it from the O&M personnel.

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 3,600. Obtain the hourly price on the **Product Pricing Details** page.

For example, if you purchase a pay-per-use ECS priced €0.0160 EUR/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 €0.008 EUR (0.0160/3,600 x 30 x 60).
- If you use the ECS for 1 hour and 30 minutes, you need to pay for €0.024 EUR (0.0160/3,600 x 90 x 60).

- If the usage duration is less than 650 hours, the ECS will be billed by the actual usage duration.
- If the usage duration exceeds 650 hours, the ECS will still be billed by 650 hours. The exceeded duration is not billed.

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 current or 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

ECSs billed on a pay-per-use basis support specification modification, while ECSs purchased using RIs do not support specification modification.

Helpful Links

- Will ECSs Continue to Be Billed After They are Stopped?
- How Can I Stop an ECS from Being Billed?

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

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

Table 1-58 lists all the system policies supported by ECS.

Table 1-58 System-defined policies supported by ECS

| Policy Name | Description | Туре | 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 CommonOp erations | Common user permissions for ECS. Users granted these permissions can start, stop, restart, and query ECSs. | System- defined policy | ECS CommonOp erations Policy Content |
| ECS ReadOnlyAc cess | Read-only permissions for ECS. Users granted these permissions can only view ECS data. | System- defined policy | ECS ReadOnlyAc cess Policy Content |
| Server Administrat or | 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. | System role | Server Administrat or Policy Content |
| | 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. | | |

Table 1-59 lists the common operations supported by each system-defined policy of ECS. Select the policies as required.

Table 1-59 Common operations supported by each system-defined policy

| Operation | ECS FullAccess | ECS CommonOper ations | ECS ReadOnlyAccess |
|-----------------|----------------|-----------------------------|-----------------------|
| Creating an ECS | Supported | Not supported | Not supported |

| Operation | ECS FullAccess | ECS CommonOper ations | 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:cloudServerFpgaImags: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:createlmage",
                  "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:floatinglps:create",
"vpc:floatinglps: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

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:privatelps:*"
         "vpc:securityGroups:*",
        "vpc:securityGroupRules:*",
         "vpc:floatinglps:*",
         "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"
   }
]
```

1.15 User Permissions

Two types of permissions are provided by default: user management and resource management.

- User management refers to the management of users, user groups, and user group rights.
- Resource management refers to the control operations that can be performed by users on cloud service resources.

For further details, see **System Permissions**.

1.16 Region and AZ

Concept

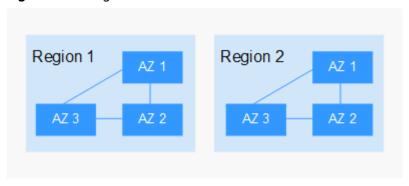
A region and availability zone (AZ) identify the location of a data center. You can create resources in a specific region and AZ.

• 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 1-16 shows the relationship between regions and AZs.

Figure 1-16 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 ondemand. 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

Notes

Dedicated physical resources

To make an ECS run on isolated physical hardware, apply for a dedicated host (DeH) before creating the ECS.

For details about DeH, see **Dedicated Host User Guide**.

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 \bigcirc 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.

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

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

For example, S3 ECSs are available only in AZ 1; S1 ECSs are available for sale in AZ 2 and AZ 3 and have been sold out in AZ 1. If you set **AZ** to **Random**, you can view both S3 and S1 ECSs. If you create an S3 ECS, the system automatically allocates it to AZ 1. If you create an S1 ECS, the system randomly allocates it to AZ 2 or AZ 3.

Set **DeH**.

This configuration is optional. This parameter is available only when you click **Create ECS** on the **Dedicated Host** page. It is unavailable when you click **Create ECS** on the **Elastic Cloud Server** page.

DeH refers to physical server resources dedicated for a specified user. You can deploy ECSs on DeHs for better isolation, security, and performance of your ECSs. You can continue using your existing server software licenses of ECSs on DeHs to reduce costs. For more details, see *Dedicated Host User Guide*.

4. Set **Specifications**.

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

∩ NOTE

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

5. Select an image.

- Public image

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

- Private image

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

□ NOTE

- If 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, sharing attribute, and data encryption 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, sharing attribute, and data encryption settings of data disks can be modified during ECS creation.
- For more details about how to use CSBS backups to create images, see "Using Backups to Create Images" in *Cloud Server Backup Service User Guide* and "Creating a Full-ECS Image Using a CSBS Backup" in *Image Management* Service User Guide.
- To ensure that NIC multi-queue is enabled on an ECS created using a private image, configure NIC multi-queue when creating such a private image. NIC multi-queue routes NIC interrupt requests among multiple vCPUs for higher network packets per second (PPS) and bandwidth.

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

Shared image

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

6. (Optional) Set License Type.

Specifies a license type for using an OS or software. This parameter is displayed only when the selected image is billed.

- Using License from the System
 - Allows you to use the license provided by the cloud service platform. Obtaining the authorization of such a license is billed.
- Bring your own license (BYOL)
 - Allows you to use your existing OS license. In such a case, you do not need to apply for a license again.

For more information about license types, see License Types.

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

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) Encryptionrelated 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. When creating an ECS, you can add up to 23 data disks with customized sizes to it. After the ECS is created, you can add up to 23 VBD disks or 59 SCSI disks to it.

Click **Show** \vee 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 Xrole** 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.

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

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

□ NOTE

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

- This IP address serves as the start IP address.
- Ensure that the IP addresses required by the NICs are within the subnet, consecutive, and available.
- The subnet with the specified IP address cannot overlap with other subnets.
- 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: TCPPort Range: 80
- Remote End: 169.254.0.0/16

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

- Protocol: ANYPort Range: ANYRemote End: 0.0.0.0/16
- 4. Set EIP.

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

The following options are provided:

Auto assign

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

Specify

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

Do not use

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

5. Set Bandwidth Size.

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

2.1.4 Step 3: Configure Advanced Settings

Advanced Settings

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.

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

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

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

2. Set **Login Mode**.

Key pair: allows you to use a key pair for 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.

3. Set Cloud Backup and Recovery.

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

The following options are provided:

- Create new
 - i. Set the name of the cloud backup vault, which consists of 1 to 64 characters, containing only letters, digits, underscores (_), and

- hyphens (-). For example, **vault-f61e**. The default naming rule is **vault** *xxxx*.
- ii. Enter the vault capacity, which is required for backing up the ECS. The vault capacity cannot be smaller than that of the ECS to be backed up. Its value ranges from the total capacity of the ECS to 10,485,760 in the unit of GB.
- iii. Select a backup policy from the drop-down list, or log in to the CBR console and configure a desired one.
- Use existing
 - i. Select an existing cloud backup vault from the drop-down list.
 - ii. Select a backup policy from the drop-down list, or log in to the CBR console and configure a desired one.
- Do not use

Skip this configuration if CBR is not required. If you need to enable CBR after creating an ECS, log in to the CBR console, locate the target vault, and bind the ECS to the vault.

4. Set ECS Group (Optional).

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

∩ NOTE

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

- 5. 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. 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*.

6. Click Next: Confirm.

2.1.5 Step 4: Confirm

Confirming the Order

- 1. On the **Confirm** page, view details about the ECS configuration.
 - To learn more about the price, click **Pricing details**.
- 2. Set the number of ECSs to be created.
 - After the configuration, click **Price Calculator** to view the ECS configuration fee.
- 3. Confirm the configuration and click **Apply Now**.

Follow-up Procedure

• After an ECS with an EIP bound is created, the system automatically generates a reverse domain name in the format of "ecs-xx-xx-xx-compute.xxx.com" for each EIP by default. In the format, "xx-xx-xx" indicates the EIP, and "xxx" indicates the domain name of the cloud service provider. You can use the reverse domain name to access the ECS.

Use one of the following commands to obtain the default reverse domain name of an EIP:

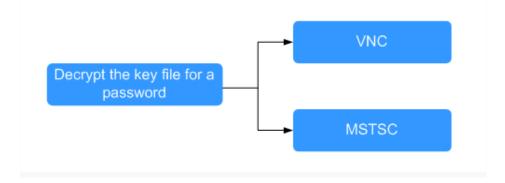
- ping -a EIP
- nslookup [-qt=ptr] EIP
- dig -x EIP

2.2 Logging In to an ECS

Logging In to a Windows ECS

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

Figure 2-1 Windows ECS login modes



1. Obtain the password.

Use the password obtaining function provided by the management console to decrypt the key file to obtain a password.

For details, see Obtaining the Password for Logging In to a Windows ECS.

- 2. Select a login method and log in to the ECS.
 - Management console (VNC)
 For details, see Remotely Logging In to a Windows ECS (Using VNC).
 - Remote desktop connection (MSTSC)
 For details, see Remotely Logging In to a Windows ECS (Using MSTSC).

Logging In to a Linux ECS

When you log in to a newly created Linux ECS for the first time, use an SSH key as user **cloud** (without default password) and make sure that the ECS has had an EIP bound. **Figure 2-2** shows the login mode.

Figure 2-2 Linux ECS login modes



Use the key file obtained during ECS creation to log in to the ECS through SSH.

For details, see Remotely Logging In to a Linux ECS (Using an SSH Key Pair).

□ NOTE

To log in to the ECS in remote VNC login mode provided on the management console, log in to the ECS using an SSH key and set the password of the login user, such as user **root**. Then, use the username and password to log in to the ECS.

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 as well as their functions and usage, see "Notes" in ECS Types.

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 2-1 lists the common disk partition styles. In Linux, different disk partition styles require different partitioning tools.

Table 2-1 Disk partition styles

| Disk Partition Style | Maximum Disk Capacity Supported | Maximum Number of Partitions Supported | Linux Partitioning Tool |
|----------------------------------|--|--|-------------------------------|
| Master Boot Record (MBR) | 2 TiB | 4 primary partitions 3 primary partitions and 1 extended partition 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: 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 | • fdisk • parted |
| GUID Partition Table (GPT) | 18 EiB 1 EiB = 1048576 TiB | into 5 logical partitions 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.

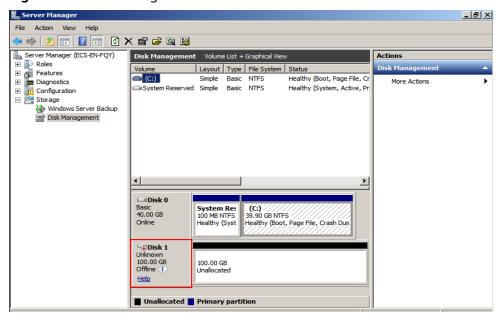
The **Server Manager** window is displayed.

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.

_ B × Action View Help Server Manager (ECS-EN-FQY)

Roles Disk Management Actions Volume List + Graphical View Volume Layout Type File System Status Healthy (Boot, Page File, Cr More Actions System Reserved Simple Basic NTFS Healthy (System, Active, Pr 🐌 Windows Server Backup Disk Management □Disk 0 (C:) 39.90 GB NTFS Healthy (Boot, Page File, Crash Dun Basic 40.00 GB System Res 100 MB NTFS Healthy (Syst Disk 1 Offline 🕕 Help

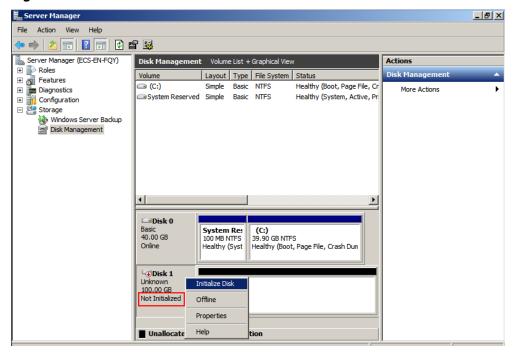
Figure 2-4 Online the disk

Ⅲ NOTE

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

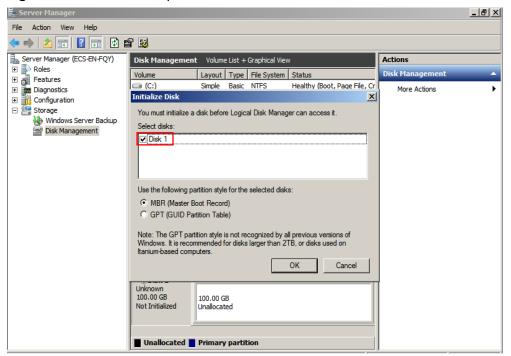
Step 4 After making the disk online, the disk status changes from **Offline** to **Not Initialized**. Right-click the disk status and choose **Initialize Disk** from the shortcut menu.

Figure 2-5 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 larger than 2 TiB.

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

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

Server Manager _ B × File Action View Help ⊨ 🤿 | 🚈 📊 | 🛭 ந | 🗗 📽 🐯 Lack Server Manager (ECS-B704) Actions Volume List + Graphical Vie Roles

Features

Diagnostics

Full Event Viewer

Performance Layout Type File System Status Volume Simple Basic NTFS Healthy (Boot, Page File, Crash Dump, More Actions System Reserved Simple Basic NTFS Healthy (System, Active, Primary Parti Device Manager

Configuration
Storage Windows Server Backup Disk Management Disk 0 System Reser 100 MB NTFS (C:) 49.90 GB NTFS 50.00 GB Online Healthy (Boot, Page File, Cra Healthy (System □Disk 1 Properties Basic 100.00 GB 100.00 GB Unallocated Primary partition

Figure 2-7 New Simple Volume

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

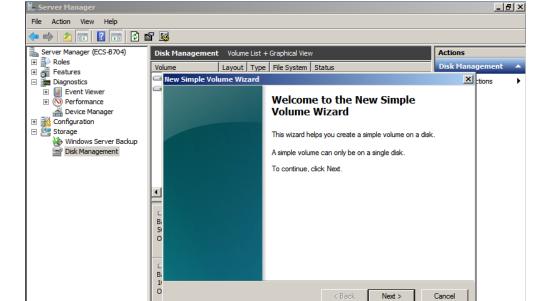


Figure 2-8 New Simple Volume Wizard

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

■ Unallocated ■ Primary partition

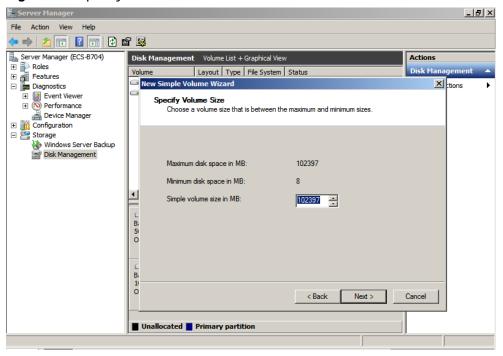
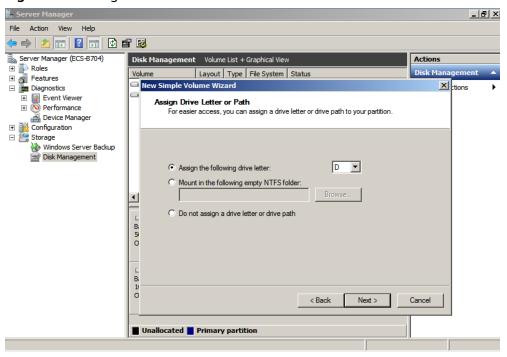


Figure 2-9 Specify Volume Size

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





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

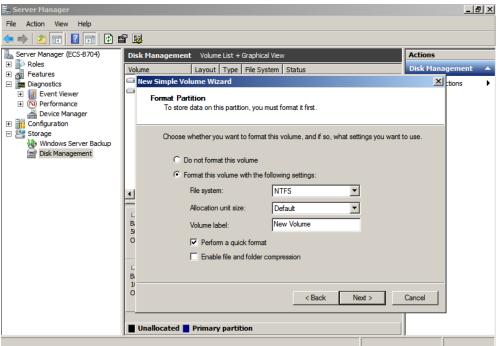
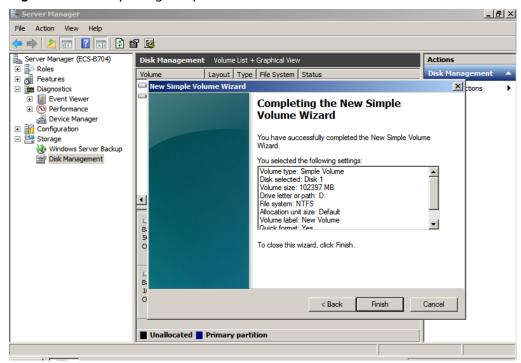


Figure 2-12 Completing the partition creation

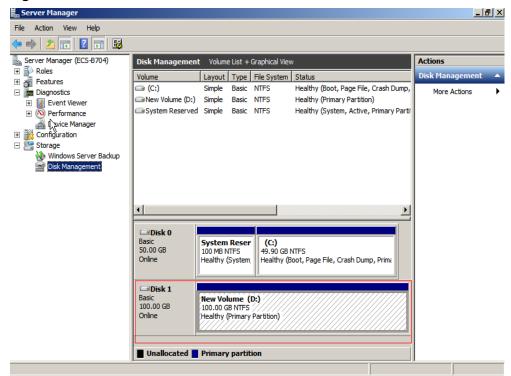


NOTICE

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

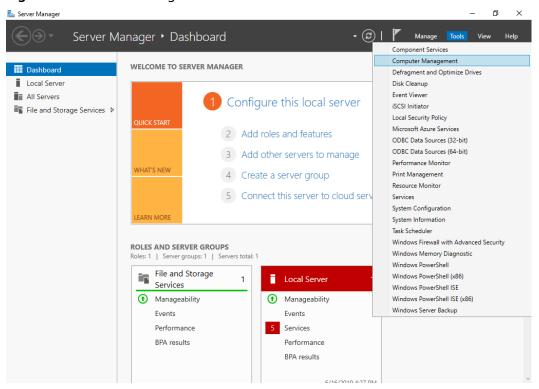
- You have attached a data disk to a server but not initialized the disk.
- You have logged in to the server.
 - For how to log in to an ECS, see the Elastic Cloud Server User Guide.
 - For how to log in to a BMS, see the Bare Metal Server User Guide.

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.

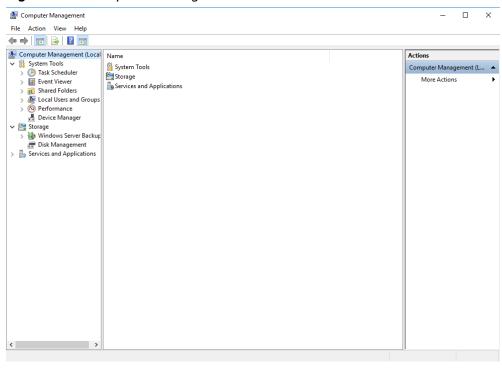


Figure 2-15 Computer Management

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.

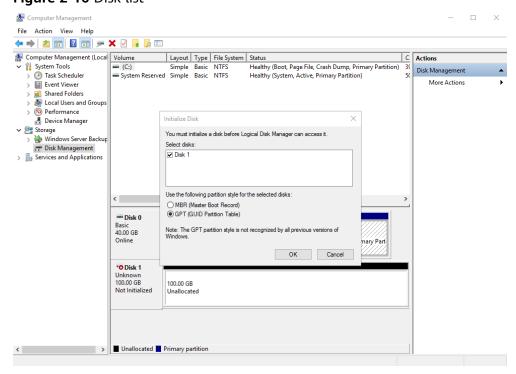
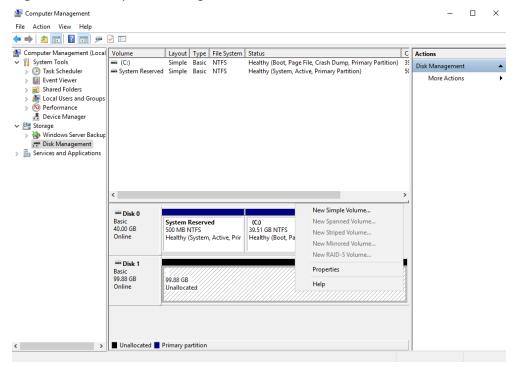


Figure 2-16 Disk list

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 larger than 2 TiB.

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

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

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

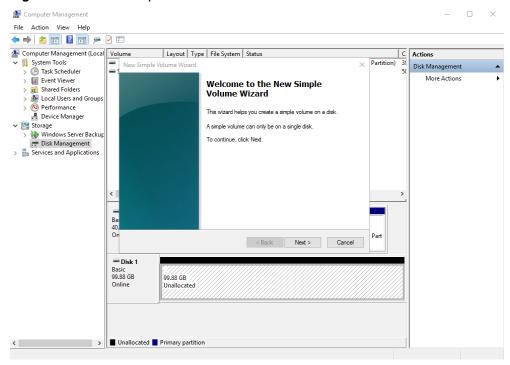
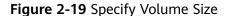
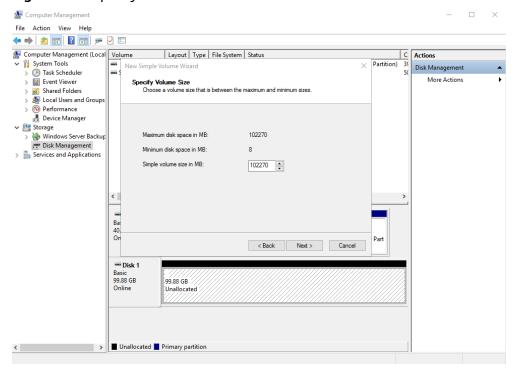


Figure 2-18 New Simple Volume Wizard

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

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

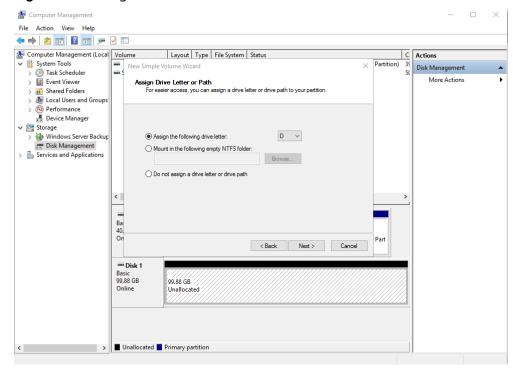




Step 8 Specify the volume size and click **Next**. The system selects the maximum volume size by default. You can specify the volume size as required. In this example, the default setting is used.

The Assign Drive Letter or Path page is displayed.

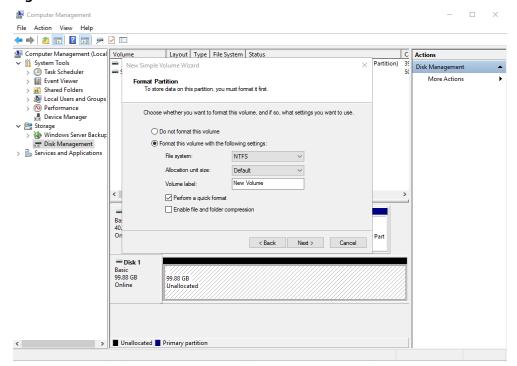
Figure 2-20 Assign Drive Letter or Path



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.

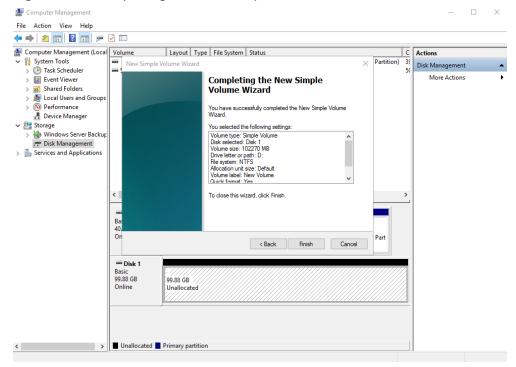
Figure 2-21 Format Partition



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.

Computer Management × File Action View Help 🐤 🔿 🙎 📆 🛭 🗗 🤛 🖸 🜆 Computer Management (Local Volume Layout Type File System Status C Actions C:) Simple Basic NTFS
 New Volume (D:) Simple Basic NTFS
 System Reserved Simple Basic NTFS Healthy (Boot, Page File, Crash Dump, Primary Partition) 35
Healthy (Primary Partition) 95
Healthy (System, Active, Primary Partition) 56 System Tools Disk Management (Task Scheduler Event Viewer More Actions Services and Applications Basic 40.00 GB Online System Reserved 500 MB NTFS Healthy (System, Active, Prir (C:) 39.51 GB NTFS Healthy (Boot, Page File, Crash Dump, Primary Part Disk 1 Basic 99.88 GB Online New Volume (D:) 99.87 GB NTFS Healthy (Primary Partition)

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.

│ 📝 🧻 🖛 │ This PC File Computer View ← → · ↑ 🖳 › This PC ∨ Folders (6) ♣ Quick access Desktop Downloads Documents Pictures This PC ∨ Devices and drives (2) Desktop New Volume (D:) Documents 18.7 GB free of 39.5 GB 99.7 GB free of 99.8 GE Downloads Music Pictures Wideos Network 8 items

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

- You have attached a data disk to a server but not initialized the disk.
- 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.

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
                 Size Used Avail Use% Mounted on
Filesystem
          Type
/dev/vda1
           ext4
                 43G 1.9G 39G 5%/
           devtmpfs 2.0G 0 2.0G 0% /dev
devtmpfs
          tmpfs 2.0G 0 2.0G 0% /dev/shm
tmpfs
tmpfs
          tmpfs
                 2.0G 9.1M 2.0G 1% /run
         tmpfs
                 2.0G 0 2.0G 0% /sys/fs/cgroup
tmpfs
                 398M 0 398M 0% /run/user/0
tmpfs
          tmpfs
/dev/vdb1 ext4 106G 63M 101G 1% /mnt/sdc
```

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

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.

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
```

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

- You have attached a data disk to a server but not initialized the disk.
- 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.

Welcome to GNU Parted! Type 'help' to view a list of commands. (parted)

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

Information similar to the following is displayed:

(parted) p

Error: /dev/vdb: unrecognised disk label Model: Virtio Block Device (virtblk)

Disk /dev/vdb: 107GiB

Sector size (logical/physical): 512B/512B

Partition Table: unknown

Disk Flags: (parted)

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

Step 4 Set the disk partition style.

mklabel Disk partition style

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

mklabel gpt

NOTICE

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

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

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

Information similar to the following is displayed:

(parted) mklabel gpt (parted) p

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

Number Start End Size File system Name Flags

(parted)
```

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

- **Step 6** Enter **unit s** and press **Enter** to set the measurement unit of the disk to sector.
- **Step 7** Create a new partition.

mkpart Partition name Start sector End sector

In this example, run the following command:

mkpart test 2048s 100%

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

```
Information similar to the following is displayed:
```

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

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

Information similar to the following is displayed:

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

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

(parted)

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

Information similar to the following is displayed:

(parted) q

Information: You may need to update /etc/fstab.

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

Step 10 View the disk partition information.

lsblk

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
__vda1 253:1 0 40G 0 part /
```

```
vdb 253:16 0 100G 0 disk

—vdb1 253:17 0 100G 0 part
```

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

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

mkfs -t File system format /dev/vdb1

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

mkfs -t ext4 /dev/vdb1

Information similar to the following is displayed:

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

Allocating group tables: done Writing inode tables: done Creating journal (32768 blocks): done

Writing superblocks and filesystem accounting information: done

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

NOTICE

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

Step 12 Create a mount point.

mkdir Mount point

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

mkdir /mnt/sdc

□ NOTE

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

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

mount Disk partition Mount point

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

mount /dev/vdb1 /mnt/sdc

Step 14 Check the mount result.

df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH
                Size Used Avail Use% Mounted on
Filesystem
         Type
                43G 1.9G 39G 5%/
/dev/vda1
          ext4
devtmpfs
          devtmpfs 2.0G 0 2.0G 0% /dev
         tmpfs 2.0G 0 2.0G 0% /dev/shm
tmpfs
         tmpfs
tmpfs
               2.0G 9.0M 2.0G 1% /run
               tmpfs
         tmpfs
tmpfs
         tmpfs
/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 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.

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

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

- You have attached a data disk to a server but not initialized the disk.
- 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.

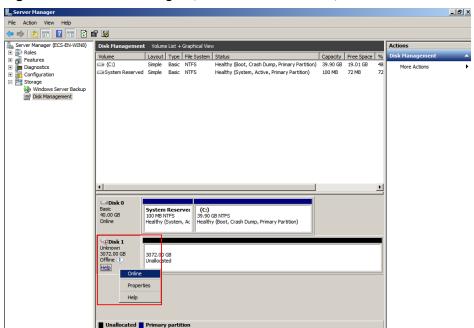


Figure 2-25 Server Manager (Windows Server 2008)

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.

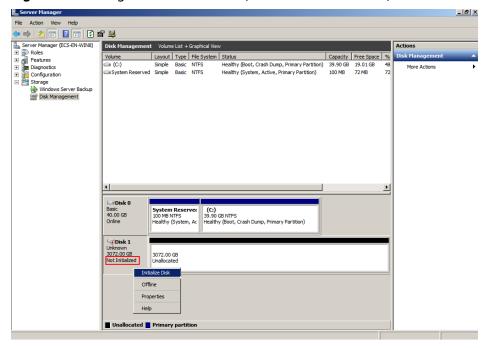


Figure 2-26 Bring online succeeded (Windows Server 2008)

Step 4 In the Disk 1 area, right-click and choose Initialize Disk from the shortcut menu.

The Initialize Disk dialog box is displayed.

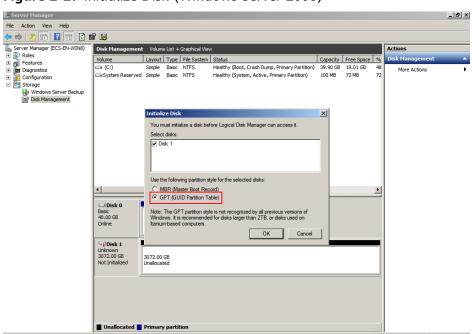


Figure 2-27 Initialize Disk (Windows Server 2008)

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.

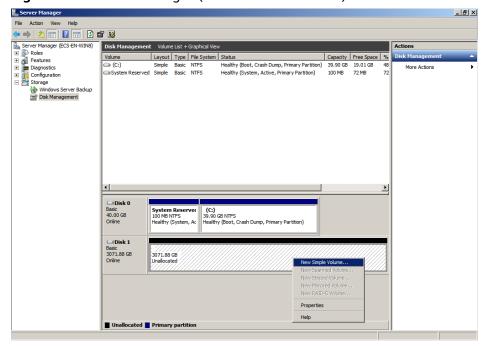


Figure 2-28 Server Manager (Windows Server 2008)

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 larger than 2 TiB.

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

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

The New Simple Volume Wizard window is displayed.

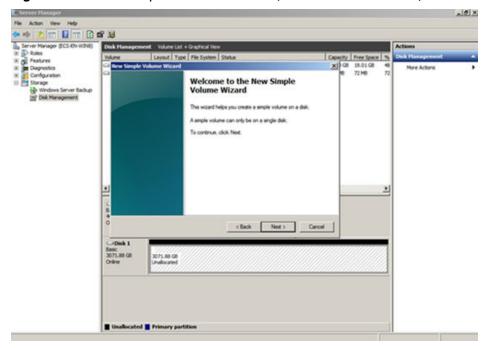


Figure 2-29 New Simple Volume Wizard (Windows Server 2008)

Step 7 Follow the prompts and click Next.

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

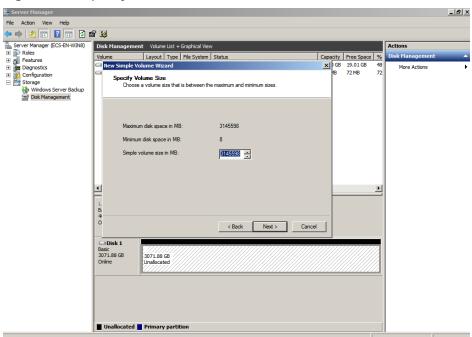


Figure 2-30 Specify Volume Size (Windows Server 2008)

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.

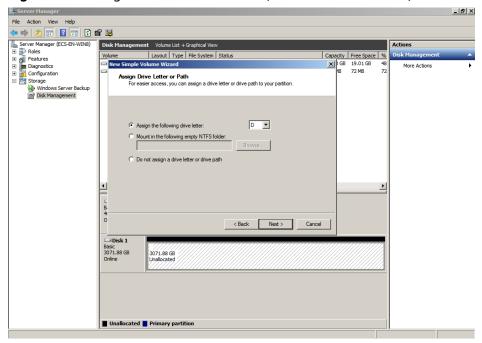
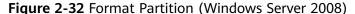
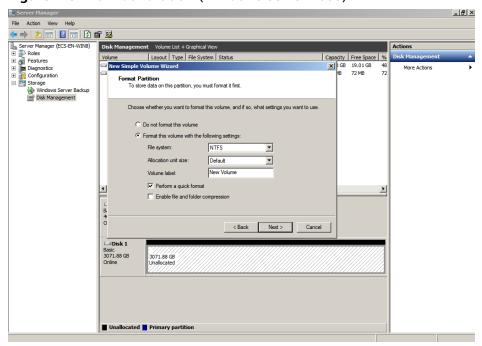


Figure 2-31 Assign Drive Letter or Path (Windows Server 2008)

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.





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.

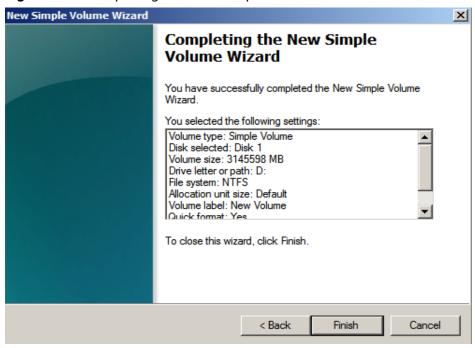


Figure 2-33 Completing the New Simple Volume Wizard

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.

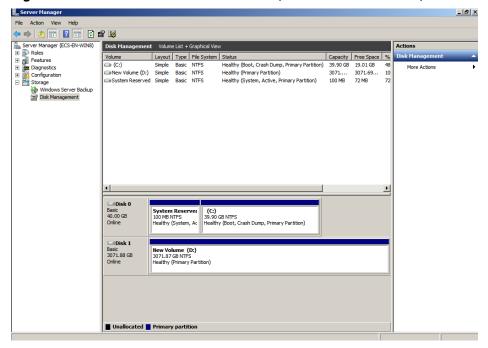


Figure 2-34 Disk initialization succeeded (Windows Server 2008)

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.

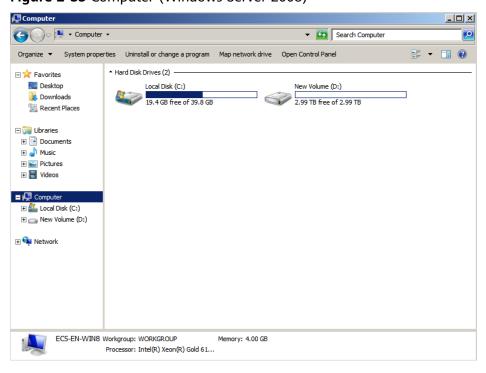


Figure 2-35 Computer (Windows Server 2008)

----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

- You have attached a data disk to a server but not initialized the disk.
- 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.

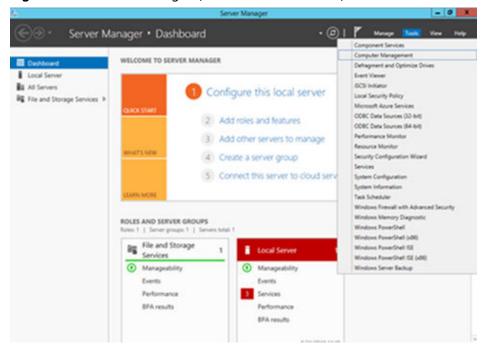
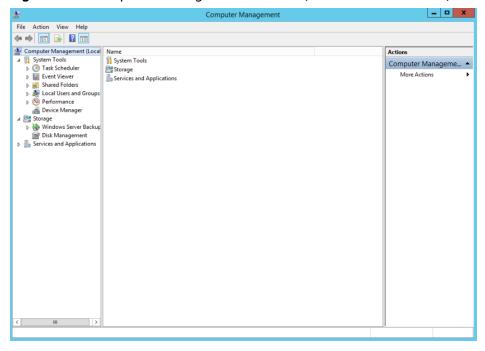


Figure 2-36 Server Manager (Windows Server 2012)

Step 2 In the upper right corner, choose **Tools** > **Computer Management**.

The **Computer Management** window is displayed.





Step 3 Choose Storage > Disk Management.

Disks are displayed in the right pane.

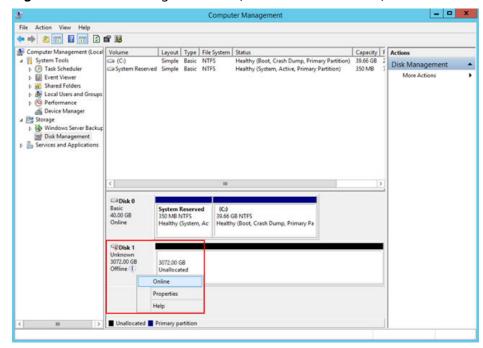


Figure 2-38 Disk Management list (Windows Server 2012)

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.

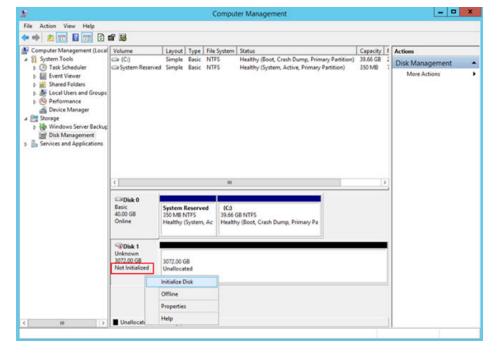
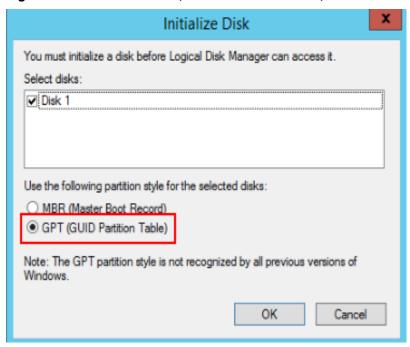


Figure 2-39 Bring online succeeded (Windows Server 2012)

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

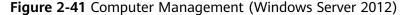
The Initialize Disk dialog box is displayed.

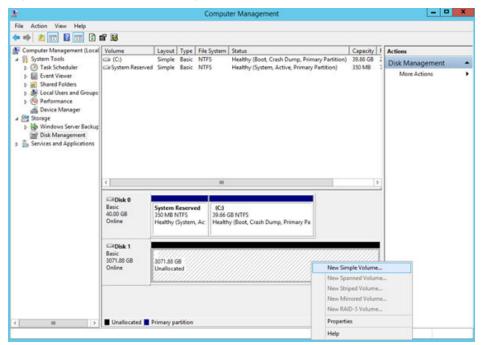
Figure 2-40 Initialize Disk (Windows Server 2012)



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.





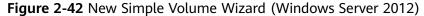
NOTICE

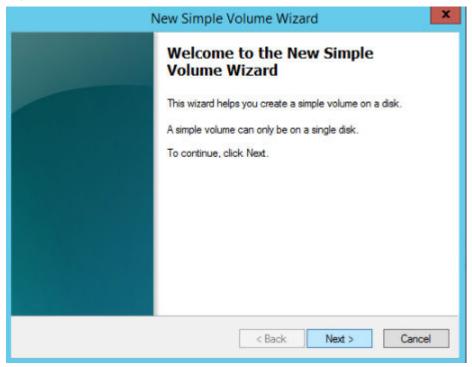
The maximum disk size supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Because an EVS data disk currently supports up to 32 TiB, use GPT if your disk size is larger than 2 TiB.

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

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

The New Simple Volume Wizard window is displayed.





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

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

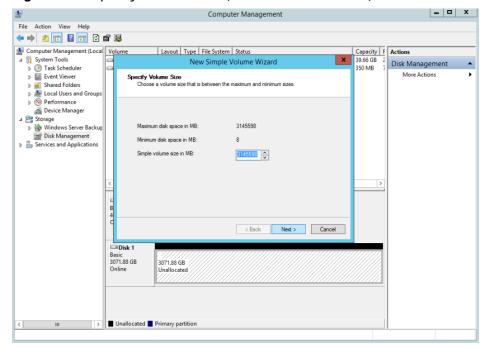
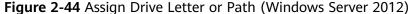
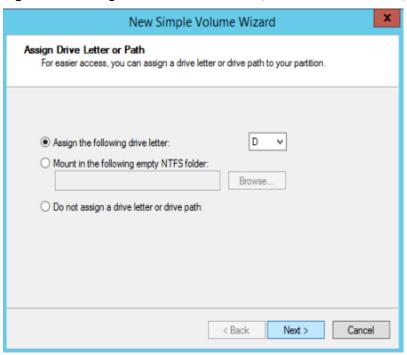


Figure 2-43 Specify Volume Size (Windows Server 2012)

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.

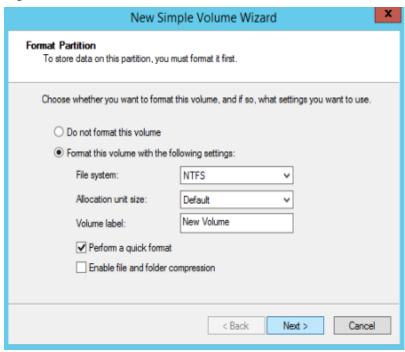




Step 10 Assign a drive letter or path to your partition and click **Next**. The system assigns drive letter D by default. In this example, the default setting is used.

The Format Partition page is displayed.

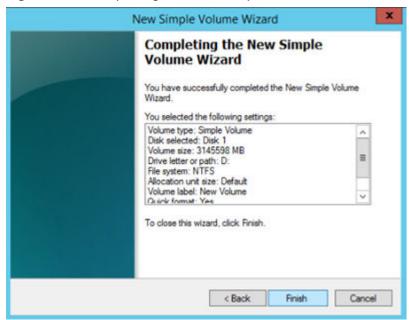
Figure 2-45 Format Partition (Windows Server 2012)



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.





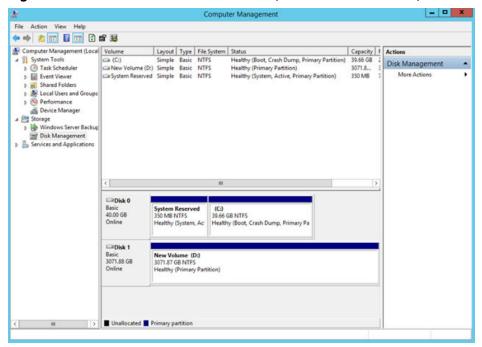
NOTICE

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.

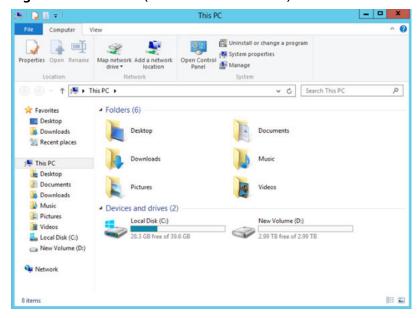


Figure 2-48 This PC (Windows Server 2012)

----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

- You have attached a data disk to a server but not initialized the disk.
- 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 larger than 2 TiB.

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

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

Information similar to the following is displayed:

(parted) mklabel gpt (parted) p

Model: Virtio Block Device (virtblk)

Disk /dev/vdb: 3299GiB

Sector size (logical/physical): 512B/512B

Partition Table: gpt

Disk Flags:

Number Start End Size File system Name Flags

(parted)

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

Step 7 Create a new partition.

mkpart Partition name Start sector End sector

In this example, run the following command:

mkpart opt 2048s 100%

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

Information similar to the following is displayed:

(parted) mkpart opt 2048s 100%

Warning: The resulting partition is not properly aligned for best performance.

Ignore/Cancel? Ignore

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

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

Information similar to the following is displayed:

(parted) p

Model: Virtio Block Device (virtblk) Disk /dev/vdb: 6442450944s

```
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start End Size File system Name Flags
1 2048s 6442448895s 6442446848s opt
```

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

- **Step 9** Enter **q** and press **Enter** to exit parted.
- **Step 10** View the disk partition information.

lsblk

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
—vda1 253:1 0 1G 0 part /boot
—vda2 253:2 0 39G 0 part /
vdb 253:16 0 3T 0 disk
—vdb1 253:17 0 3T 0 part
```

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

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

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

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

mkfs -t ext4 /dev/vdb1

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# mkfs -t ext4 /dev/vdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
201326592 inodes, 805305856 blocks
40265292 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2952790016
24576 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
     32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
     4096000, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
     102400000, 214990848, 512000000, 550731776, 644972544
Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

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

NOTICE

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

Step 12 Create a mount point.

mkdir Mount point

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

mkdir /mnt/sdc

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
          Type Size Used Avail Use% Mounted on
Filesystem
/dev/vda2
           ext4
                  42G 1.5G 38G 4% /
          devtmpfs 2.0G 0 2.0G 0% /dev
devtmpfs
          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
tmpfs
/dev/vda1
          ext4
                 1.1G 153M 801M 17% /boot
          tmpfs
                 398M 0 398M 0% /run/user/0
tmpfs
/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
```

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

The system saves the configurations and exits the vi editor.

- **Step 6** Verify that the disk is auto-mounted at startup.
 - 1. Unmount the partition.

umount Disk partition

In this example, run the following command:

umount /dev/vdb1

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

mount -a

3. Query the file system mounting information.

mount | grep Mount point

In this example, run the following command:

mount | grep /mnt/sdc

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

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

----End

3 Instances

3.1 Viewing ECS Information

3.1.1 Viewing ECS Creation Statuses

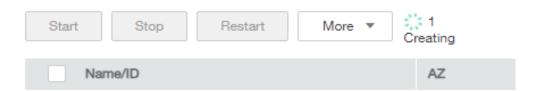
Scenarios

After submitting the request for creating an ECS, you can view the creation status. This section describes how to view the creation status of an ECS.

Procedure

- 1. Log in to the management console.
- 2. Under Computing, click Elastic Cloud Server.
- 3. After creating an ECS, view the creation status above the ECS list beside the common operations (**Start**, **Stop**, **Restart**, and **More**).

Figure 3-1 ECS creation status



4. Click the number displayed above **Creating** and view task details.

◯ NOTE

- An ECS that is being created is in one of the following states:
 - **Creating**: The ECS is being created.
 - Faulty: Creating the ECS failed. In such a case, the system automatically rolls back the task and displays an error code on the GUI, for example, Ecs.0013 Insufficient EIP quota.
 - Running: The request of creating the ECS has been processed, and the ECS is running properly. An ECS in this state can provide services for you.
 - See **How Do I Handle Error Messages Displayed on the Management Console?** for troubleshooting.
- If you find that the task status area shows an ECS creation failure but the ECS has been created successfully and displayed in the ECS list, see Why Does the Failures Area Show an ECS Creation Failure But the ECS List Displays the Created ECS?

3.1.2 Viewing Failed Tasks

Scenarios

You can view the details of failed task (if any) in the **Failures** area, including the task names and statuses. This section describes how to view failures.

Failure Types

Table 3-1 lists the types of failures that can be recorded in the **Failures** area.

Table 3-1 Failure types

| Failure Type | Description |
|----------------------|--|
| Creation failures | A task failed. For a failed task, the system rolls back the task and displays an error code, for example, Ecs.0013 Insufficient EIP quota. For details about how to handle the failure, see How Do I Handle Error Messages Displayed on the Management Console? |
| Operation failures | Modifying ECS specifications If an ECS specifications modification failed, this operation is recorded in Failures . |
| | Automatic recovery enabled during ECS creation Automatic recovery is enabled during ECS creation. After the ECS is created, if the system fails to enable automatic recovery, this operation is recorded in Failures. |

Procedure

- 1. Log in to the management console.
- 2. Under Computing, choose Elastic Cloud Server.
- 3. View **Failures** on the right side of common operations.

□ NOTE

If **Failures** is not displayed on the management console, the following tasks have been successfully executed:

- The ECS specifications are modified.
- Automatic recovery is enabled during ECS creation.
- 4. Click the number displayed in the **Failures** area to view task details.
 - Creation Failures: show the failed ECS creation tasks.
 - Operation Failures: show the tasks with failed operations and error codes, which help you troubleshoot the faults.

3.1.3 Viewing ECS Details (List View)

Scenarios

After obtaining ECSs, you can view and manage them on the management console. This section describes how to view detailed ECS configurations, including its name, image, system disk, data disks, VPC, NIC, security group, and EIP.

To view the private IP address of an ECS, view it on the **Elastic Cloud Server** page.

Procedure

- 1. Log in to the management console.
- 2. Under Computing, choose Elastic Cloud Server.

The **Elastic Cloud Server** page is displayed. On this page, you can view your ECSs and the basic information about the ECSs, such as their specifications, images, and IP addresses.

- 3. In the search box above the ECS list, select a filter (such as ECS name, ID, or private IP address), enter the corresponding information, and press **Enter**.
- 4. Click the name of the target ECS.

The page providing details about the ECS is displayed.

5. View the ECS details.

You can click the tabs and perform operations. For details, see **Changing a Security Group**, **Attaching a Network Interface**, **Adding Tags**, and **Binding an EIP**.

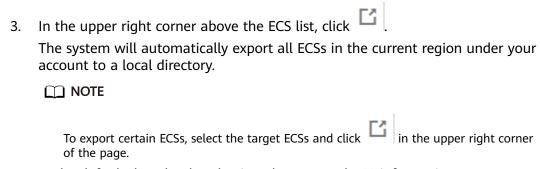
3.1.4 Exporting ECS Information

Scenarios

The information of all ECSs under your account can be exported in a CSV file to a local directory. The file includes the IDs, private IP addresses, and EIPs of your ECSs.

Procedure

- Log in to the management console.
- 2. Under **Computing**, choose **Elastic Cloud Server**.



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

3.2 Logging In to a Windows ECS

3.2.1 Login Overview

Constraints

- Only a running ECS can be logged in.
- 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 3-2 Windows login modes

| ECS OS | Local OS | Connection Method | Requirement |
|---------|----------|---|--|
| Windows | Windows | Use MSTSC. Click Start on the local computer. In the Search programs and files text box, enter mstsc to open the Remote Desktop Connection dialog box. For details, see Remotely Logging In to a Windows ECS (Using MSTSC). | The target ECS has had an EIP bound. |
| | Linux | Install a remote connection tool, for example, rdesktop. For details, see Remotely Logging In to a Windows ECS (from a Linux Computer). | |
| | Windows | Through the management console. For details, see Remotely Logging In to a Windows ECS (Using VNC). | No EIP is required. |

Helpful Links

- Login Password Resetting
- Remote Logins

3.2.2 Remotely Logging In to a Windows ECS (Using VNC)

Scenarios

This section describes how to use VNC provided on the management console to log in to an ECS. This function applies to emergency O&M. In other scenarios, you are advised to log in to ECSs using SSH or MSTSC.

Constraints

- 1. The remote login function is implemented using custom ports. Therefore, before attempting to log in remotely, ensure that the port to be used is not blocked by the firewall. For example, if the remote login link is xxx:8002, ensure that port 8002 is not blocked by the firewall.
- 2. If the client OS uses a local proxy and the firewall port cannot be configured on the local proxy, disable the proxy mode and then try logging in remotely.
- 3. The French keyboard supports the following browsers: Google Chrome 55, Mozilla Firefox 50, and Internet Explorer 11.
 - If your browser version is too early to support the French keyboard, use both the soft keyboard and physical keyboard for data input.

4. The password for logging in to a Linux ECS has been set.

If you have not set the password, log in to the ECS using an SSH key and then set the login password.

Login Notes

1. When you log in to the ECS using VNC, four types of keyboards will be used. These are described in **Table 3-3**.

Table 3-3 Keyboard types

| Keyboard Type | Description | Keyboard Language |
|---|---|--|
| Physical keyboard | Used by the terminal and allows terminal data input. | Selected by users locally. |
| Input method keyboard on the terminal | Used for logging in to the management console from a terminal, such as a computer. The keyboard input method of the terminal must comply with the physical keyboard language type. In this way, the entered data can be correctly transferred from the physical keyboard to the VNC client. | Selected by users locally. |
| VNC keyboard | Used for VNC logins. The VNC keyboard input method must comply with the physical keyboard language type. In this way, the entered data can be correctly transferred from the VNC client to the ECS OS. NOTE The English keyboard is used by default. The system also supports other keyboard languages. | Can be configured using the management console. For instructions about how to select a VNC keyboard language, see Logging In to an ECS Using an English Keyboard and Logging In to an ECS Using a Non-English Keyboard. |

| Keyboard Type | Description | Keyboard Language |
|--------------------|--|--|
| ECS OS keyboard | Input method keyboard configured in the ECS OS. Ensure that this input method complies with the physical keyboard language type for correct response to the entered data transferred from the VNC client. NOTE The default OS keyboard language of an ECS created using a public image is English. The OS keyboard language of an ECS created using a private image is customized. | Configured by users locally. For instructions about how to change an ECS OS keyboard language, see Changing the OS Keyboard Language. |

2. When you log in to the ECS using VNC, ensure that your configured keyboard language is correct.

The entered data is as expected only if the input method keyboard on the terminal, the VNC keyboard, and the ECS OS keyboard languages are the same as the physical keyboard language. For details about language configuration in the four types of keyboards, see **Table 3-4**.

Table 3-4 Language configuration in the four types of keyboards

| Physical Keyboard | Input Method Keyboard on the Terminal | VNC Keyboard | ECS OS Keyboard | Permission |
|----------------------|---|-----------------|--------------------|------------|
| English | English | English | English | Yes |
| | | | French | No |
| | | French | English | No |
| | | | French | No |
| | French | English | English | No |
| | | | French | No |
| | | French | English | No |
| | | | French | No |
| French | English | English | English | No |
| | | | French | No |
| | | French | English | No |

| Physical Keyboard | Input Method Keyboard on the Terminal | VNC Keyboard | ECS OS Keyboard | Permission |
|----------------------|---|-----------------|--------------------|------------|
| | | | French | No |
| | French | English | English | No |
| | | | French | No |
| | | French | English | No |
| | | | French | Yes |

3. If the password used when you create the ECS is entered using the English keyboard, you must use the English keyboard to enter the password when logging in to the ECS later.

Logging In to an ECS Using an English Keyboard

- Log in to the management console.
- 2. Under Computing, click Elastic Cloud Server.
- Obtain the password for logging in to the ECS.
 Before logging in to the ECS using VNC, you must have the login password.
 - For instructions about how to obtain the password for logging in to a Windows ECS, see Obtaining the Password for Logging In to a Windows ECS.
 - Do as follows to obtain the password for logging in to a Linux ECS created using a private image:
 - If the image that is used to create the ECS meets the following conditions, the default image password is the login password:
 - When Cloud-Init is installed, **lock_passwd** is set to **false**. This indicates that password authentication is enabled.
 - o **passwd** is set to the default image password.
 - If password authentication is disabled or no default image password is set when Cloud-Init is installed, you must log in to the ECS by following the instructions provided in Remotely Logging In to a Linux ECS (Using an SSH Key Pair). Then, you can set the ECS login password.
 - If a Linux ECS is created using a public image, you must log in to the ECS by following the instructions provided in Remotely Logging In to a Linux ECS (Using an SSH Key Pair). Then, you can set the ECS login password.
- 4. In the search box above the upper right corner of the ECS list, enter the ECS name and click for search.
- 5. Locate the row containing the ECS and click **Remote Login** in the **Operation** column.

- 6. In the **Configure Keyboard Layout for Remote Login** dialog box, select English keyboard.
- 7. Click **Remote Login**.
- 8. (Optional) If you have changed the system language, in the dialog box that is displayed, click **Start Remote Login**.
- (Optional) When the system displays "Press CTRL+ALT+DELETE to log on", click **Send CtrlAltDel** in the upper part of the remote login page to log in to the ECS.

Figure 3-2 Send CtrlAltDel



10. Enter the password obtained in 3.

Logging In to an ECS Using a Non-English Keyboard

- 1. Log in to the management console.
- 2. Under Computing, click Elastic Cloud Server.
- Obtain the password for logging in to the ECS.
 Before logging in to the ECS using VNC, you must have the login password.
 - For instructions about how to obtain the password for logging in to a Windows ECS, see Obtaining the Password for Logging In to a Windows ECS.
 - Do as follows to obtain the password for logging in to a Linux ECS created using a private image:
 - If the image that is used to create the ECS meets the following conditions, the default image password is the login password:
 - When Cloud-Init is installed, lock_passwd is set to false. This
 indicates that password authentication is enabled.
 - o **passwd** is set to the default image password.
 - If password authentication is disabled or no default image password is set when Cloud-Init is installed, you must log in to the ECS by following the instructions provided in Remotely Logging In to a Linux ECS (Using an SSH Key Pair). Then, you can set the ECS login password.
 - If a Linux ECS is created using a public image, you must log in to the ECS by following the instructions provided in Remotely Logging In to a Linux ECS (Using an SSH Key Pair). Then, you can set the ECS login password.
- 4. In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID, and click for search.
- 5. Locate the row containing the ECS and click **Remote Login** in the **Operation** column.

- 6. In the **Configure Keyboard Layout for Remote Login** dialog box, select your desired keyboard language.
 - When logging in to the ECS using VNC for the first time, select the default English keyboard. The ECS OS uses the English keyboard by default.
 - If you have changed the keyboard language of the ECS OS, select the keyboard language to which you have changed.
- 7. Click Remote Login.
- 8. (Optional) If you have changed the system language, in the dialog box that is displayed, click **Start Remote Login**.
- 9. (Optional) When the system displays "Press CTRL+ALT+DELETE to log on", click **Send CtrlAltDel** in the upper part of the remote login page to log in to the ECS.

Figure 3-3 Send CtrlAltDel



- 10. Enter the password obtained in 3.
 - When logging in to the ECS using VNC for the first time, use the English keyboard to enter the password. After you have logged in to the ECS, see Changing the OS Keyboard Language to change the keyboard language of the ECS OS. You can then select the keyboard language and enter the password the next time you log in.
 - If you have changed the keyboard language of the ECS OS, ensure that the keyboard language in use, the keyboard language selected in step 6, and the changed OS keyboard language are all the same.

Changing the OS Keyboard Language

• If the ECS is running Linux, run the following command:

loadkeys keymapfile

The *keymapfile* parameter indicates the name of the file containing the mappings between the keys and displayed characters.

For example, if the name of a French keyboard mapping file is **fr**, run the **loadkeys fr** command.

If the ECS is running Windows, perform the following operations:
 Switch the input method or open the soft keyboard before entering characters. To do so, click the function menu icon and select soft keyboard and keyboard layout.

Configuration Example

Scenarios

If you attempt to log in to an ECS created using a public image for the first time, the languages of the four types of keyboards before the configuration are as follows (**Before configuration** row in **Table 3-5**):

• Physical keyboard: French

• Input method keyboard on the terminal: English

VNC keyboard: EnglishECS OS keyboard: English

In this case, you must change the languages of the other three types of keyboards to the same language as the physical keyboard for expected data entering. For details, see the **Solution 1** row in **Table 3-5**.

Table 3-5 Languages in the four types of keyboards

| - | Physical Keyboard | Input Method Keyboard on the Terminal | VNC Keyboard | ECS OS Keyboard |
|-----------------------------|----------------------|---|-----------------|-----------------|
| Before configu ration | French | English | English | English |
| Solutio n 1 | French | French | French | French |
| Solutio n 2 | English | English | English | English |

Procedure

- 1. Locally configure the language, for example, French, in the input method keyboard on the terminal.
- 2. Set the VNC keyboard language to English.
 - □ NOTE

When you log in to the ECS using VNC for the first time, the default ECS OS keyboard language is English. Therefore, you must set the VNC keyboard language to English.

- Log in to the ECS and change the ECS OS language to French.For details, see Changing the OS Keyboard Language.
- Change the VNC keyboard language to French.
 For details, see Logging In to an ECS Using a Non-English Keyboard.

To set the languages on the four types of keyboards to all be the same, repeat steps 1 to 4.

During the configuration, if English characters cannot be entered using the current physical keyboard, use the English soft keyboard to modify the configuration described in the **Solution 2** row of **Table 3-5**.

- To enable the Windows English soft keyboard, choose **Start** > **Run**, enter **osk**, and press **Enter**.
- The method of enabling the Linux English soft keyboard varies depending on the OS version and is not described in this document.

Related Links

For FAQs about VNC-based ECS logins, see the following links:

- What Browser Version Is Required to Remotely Log In to an ECS?
- Why Cannot I Use the French Keyboard to Enter Characters When I Log In to an ECS Using VNC?
- 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?
- What Should I Do If I Cannot View Data After Logging In to an ECS Using VNC?
- Why Are Characters Entered Through VNC Still Incorrect After the Keyboard Language Is Switched?
- Why Does a Blank Screen Appear After I Attempted to Log In to an ECS Using VNC?

3.2.3 Remotely Logging In to a Windows ECS (Using MSTSC)

Scenarios

This section describes how to use the remote login tool MSTSC to log in to a Windows ECS from a local computer.

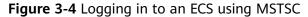
Prerequisites

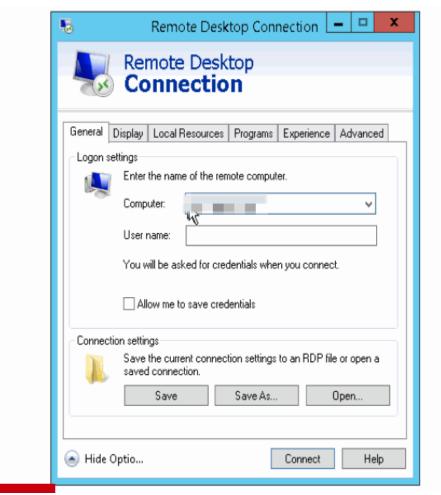
- The target ECS is running.
- You have obtained the password for logging in to the Windows ECS. For details, see Obtaining the Password for Logging In to a Windows ECS.
- You have bound an EIP to the ECS. For details, see Binding an EIP.
- Access to port 3389 is allowed in the inbound direction of the security group to which the ECS belongs. For details, see Configuring Security Group Rules.
- The network connection between the login tool and the target ECS is normal. For example, the default port 3389 is not blocked by the firewall.
- Remote Desktop Protocol (RDP) needs to be enabled on the target ECS. For ECSs created using public images, RDP has been enabled by default. For instructions about how to enable RDP, see Enabling RDP.

Logging In to a Windows ECS Using MSTSC

If your local server runs Windows, you can use the remote desktop connection tool MSTSC delivered with the Windows OS to log in to a Windows ECS.

The following uses Windows Server 2012 ECS as an example.

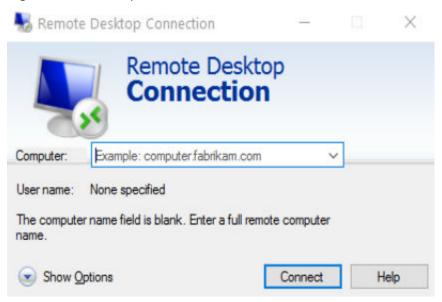




For details, see the following procedure:

- 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**.

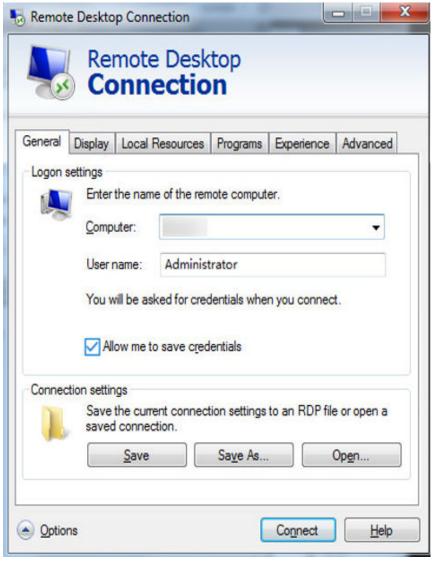
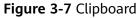
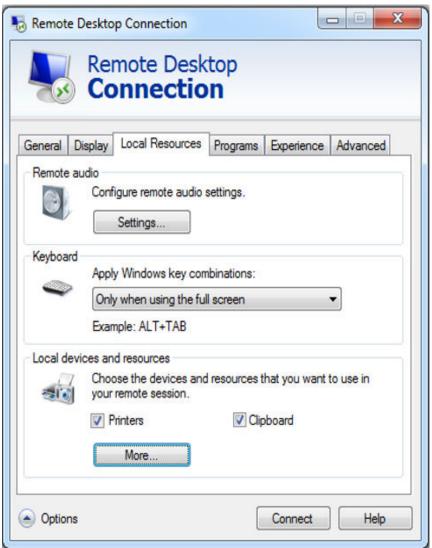


Figure 3-6 Remote Desktop Connection

5. (Optional) To use local server resources in a remote session, configure parameters on the **Local Resources** tab.

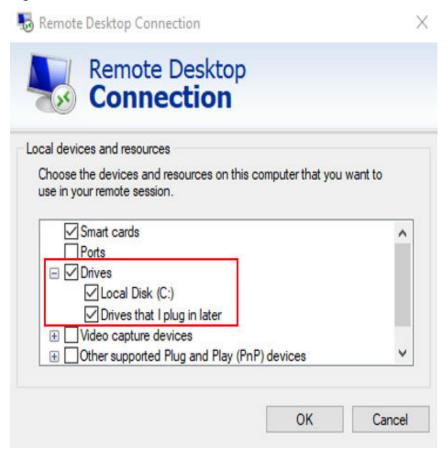
To copy data from the local server to your ECS, select **Clipboard**.





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.

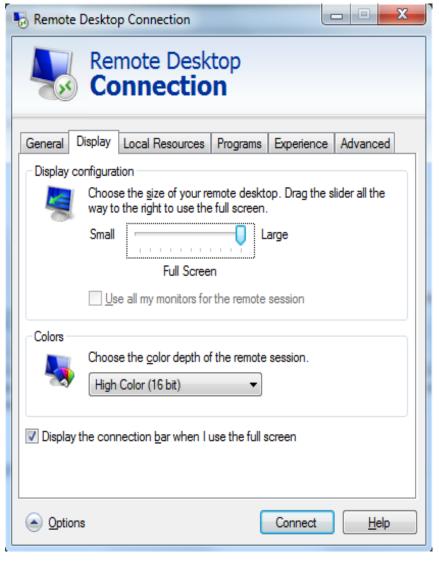


Figure 3-9 Adjusting the size of the desktop

- 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

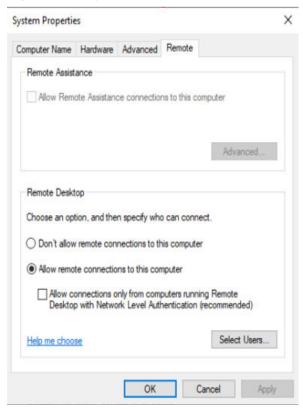
By default, RDP has been enabled on the ECSs created using a public image.

Log in to the Windows ECS using VNC.
 For details, see Remotely Logging In to a Windows ECS (Using VNC).

2. Click **Start** in the task bar and choose **Control Panel** > **System and Security** > **System** > **Remote settings**.

The **System Properties** dialog box is displayed.

Figure 3-10 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

3.2.4 Remotely Logging In to a Windows ECS (from a Linux Computer)

Scenarios

This section describes how to log in to a Windows ECS from a Linux computer.

Prerequisites

- The target ECS is running.
- You have bound an EIP to the ECS.
- Access to port 3389 is allowed in the inbound direction of the security group to which the ECS belongs.

- Data can be exchanged between the login tool and the target ECS. For example, the default port 3389 is not blocked by the firewall.
- RDP has been enabled on the target ECS. By default, RDP has been enabled on the ECSs created using a public image. For instructions about how to enable RDP, see Enabling RDP.

Procedure

To log in to a Windows ECS from a local Linux computer, use a remote access tool, such as rdesktop.

1. Run the following command to check whether rdesktop has been installed on the ECS:

rdesktop

If the message "command not found" is displayed, rdesktop is not installed. In such a case, obtain the rdesktop installation package at the **official rdesktop website**.

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

rdesktop -u *Username* -p *Password* -g *Resolution EIP*For example, run rdesktop -u administrator -p password -g 1024*720
121.xx.xx.xx.

Table 3-6 Parameters in the remote login command

| Parameter | Description |
|-----------|---|
| -u | Username, which defaults to Administrator for Windows ECSs |
| -р | 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

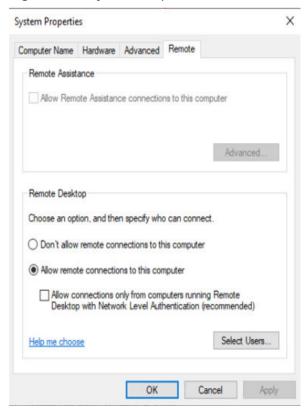
By default, RDP has been enabled on the ECSs created using a public image.

Log in to the Windows ECS using VNC.
 For details, see Remotely Logging In to a Windows ECS (Using VNC).

2. Click **Start** in the task bar and choose **Control Panel** > **System and Security** > **System** > **Remote settings**.

The **System Properties** dialog box is displayed.

Figure 3-11 System Properties



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

3.3 Logging In to a Linux ECS

3.3.1 Login Overview

Constraints

- Only a running ECS can be logged in.
- The username for logging in to a Linux ECS created using a public image is **cloud**.

Login Modes

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

| ECS OS | Local OS | Connection Method | Requirement |
|--------|----------|--|----------------------------------|
| Linux | Windows | Use a remote login tool, such as PuTTY or Xshell. For details, see Logging In to a Linux ECS from a Local Windows Server. | The target ECS has an EIP bound. |
| | Linux | Run commands. For details, see Logging In to a Linux ECS from a Local Linux Server. | |
| | Windows | Use the remote login function available on the management console. For details, see Remotely Logging In to a Linux ECS (Using VNC). | No EIP is required. |

Table 3-7 Linux ECS login modes

Helpful Links

- Application Scenarios for Using Passwords
- Why Can't I Log In to My Linux ECS?

3.3.2 Remotely Logging In to a Linux ECS (Using VNC)

Scenarios

This section describes how to use VNC provided on the management console to log in to an ECS. This function applies to emergency O&M. In other scenarios, you are advised to log in to ECSs using SSH.

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.

Constraints

- 1. The remote login function is implemented using custom ports. Therefore, before attempting to log in remotely, ensure that the port to be used is not blocked by the firewall. For example, if the remote login link is xxx:8002, ensure that port 8002 is not blocked by the firewall.
- 2. If the client OS uses a local proxy and the firewall port cannot be configured on the local proxy, disable the proxy mode and then try logging in remotely.
- 3. The French keyboard supports the following browsers: Google Chrome 55, Mozilla Firefox 50, and Internet Explorer 11.
 - If your browser version is too early to support the French keyboard, use both the soft keyboard and physical keyboard for data input.

4. The password of user **root** or a common user for logging in to the Linux ECS has been set.

If you have not set the password, log in to the ECS using an SSH key and then set the login password.

Login Notes

1. When you log in to the ECS using VNC, four types of keyboards will be used, as described in **Table 3-8**.

Table 3-8 Keyboard types

| Keyboard Type | Description | Keyboard Language |
|---|---|--|
| Physical keyboard | Used by the terminal and allows terminal data input. | Selected by users locally. |
| Input method keyboard on the terminal | Used for logging in to the management console from a terminal, such as a computer. The keyboard input method of the terminal must comply with the physical keyboard language type. In this way, the entered data can be correctly transferred from the physical keyboard to the VNC client. | Selected by users locally. |
| VNC keyboard | Used for VNC logins. The VNC keyboard input method must comply with the physical keyboard language type. In this way, the entered data can be correctly transferred from the VNC client to the ECS OS. NOTE The English keyboard is used by default. The system also supports other keyboard languages. | Can be configured using the management console. For instructions about how to select a VNC keyboard language, see Logging In to an ECS Using an English Keyboard and Logging In to an ECS Using a Non-English Keyboard. |

| Keyboard Type | Description | Keyboard Language |
|--------------------|--|--|
| ECS OS keyboard | Input method keyboard configured in the ECS OS. Ensure that this input method complies with the physical keyboard language type for correct response to the entered data transferred from the VNC client. NOTE The default OS keyboard language of an ECS created using a public image is English. The OS keyboard language of an ECS created using a private image is customized. | Configured by users locally. For instructions about how to change an ECS OS keyboard language, see Changing the OS Keyboard Language. |

2. When you log in to the ECS using VNC, ensure that your configured keyboard language is correct.

The entered data is as expected only if the input method keyboard on the terminal, the VNC keyboard, and the ECS OS keyboard languages are the same as the physical keyboard language. For details about language configuration in the four types of keyboards, see **Table 3-9**.

Table 3-9 Language configuration in the four types of keyboards

| Physical Keyboard | Input Method Keyboard on the Terminal | VNC Keyboard | ECS OS Keyboard | Supported or Not |
|----------------------|---|-----------------|--------------------|---------------------|
| English | English | English | English | Yes |
| | | | French | No |
| | | French | English | No |
| | | | French | No |
| | French | English | English | No |
| | | | French | No |
| | | French | English | No |
| | | | French | No |
| French | English | English | English | No |
| | | | French | No |
| | | French | English | No |

| Physical Keyboard | Input Method Keyboard on the Terminal | VNC Keyboard | ECS OS Keyboard | Supported or Not |
|----------------------|---|-----------------|--------------------|---------------------|
| | | | French | No |
| | French | English | English | No |
| | | | French | No |
| | | French | English | No |
| | | | French | Yes |

3. If the password used when you create the ECS is entered using the English keyboard, you must use the English keyboard to enter the password when logging in to the ECS later.

Logging In to an ECS Using an English Keyboard

- 1. Log in to the management console.
- 2. Under Computing, click Elastic Cloud Server.
- 3. Obtain the password for logging in to the ECS.

 Before logging in to the ECS, you must have the login password.
 - Do as follows to obtain the password for logging in to a Linux ECS created using a private image:
 - If the image used when creating the ECS meets the following conditions, the default image password is the login password: When Cloud-Init is installed, lock_passwd is set to false (this indicates that password authentication is enabled). passwd is set to the default image password.
 - If password authentication is disabled or no default image password is set when Cloud-Init is installed, you must log in to the ECS by following the instructions provided in Remotely Logging In to a Linux ECS (Using an SSH Key Pair). Then, you can set the ECS login password.
 - If a Linux ECS is created using a public image, you must log in to the ECS by following the instructions provided in Remotely Logging In to a Linux ECS (Using an SSH Key Pair). Then, you can set the ECS login password.
- 4. In the search box above the upper right corner of the ECS list, enter the ECS name and click .
- 5. Locate the row containing the ECS and click **Remote Login** in the **Operation** column.
- 6. In the **Configure Keyboard Layout for Remote Login** dialog box, select English keyboard.

- 7. Click **Remote Login**.
- 8. (Optional) If you have changed the system language, in the dialog box that is displayed, click **Start Remote Login**.
- 9. (Optional) When the system displays "Press CTRL+ALT+DELETE to log on", click **Send CtrlAltDel** in the upper part of the remote login page to log in to the ECS.

Figure 3-12 Send CtrlAltDel



10. Enter the password obtained in 3.

Logging In to an ECS Using a Non-English Keyboard

- 1. Log in to the management console.
- 2. Under **Computing**, click **Elastic Cloud Server**.
- 3. Obtain the password for logging in to the ECS.

 Before logging in to the ECS, you must have the login password.
 - Do as follows to obtain the password for logging in to a Linux ECS created using a private image:
 - If the image used when creating the ECS meets the following conditions, the default image password is the login password: When Cloud-Init is installed, lock_passwd is set to false (this indicates that password authentication is enabled). passwd is set to the default image password.
 - If password authentication is disabled or no default image password is set when Cloud-Init is installed, you must log in to the ECS by following the instructions provided in Remotely Logging In to a Linux ECS (Using an SSH Key Pair). Then, you can set the ECS login password.
 - If a Linux ECS is created using a public image, you must log in to the ECS by following the instructions provided in Remotely Logging In to a Linux ECS (Using an SSH Key Pair). Then, you can set the ECS login password.
- 4. In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID, and click
- 5. Locate the row containing the ECS and click **Remote Login** in the **Operation** column.
- 6. In the **Configure Keyboard Layout for Remote Login** dialog box, select your desired keyboard language.
 - When logging in to the ECS using VNC for the first time, select the default English keyboard. The ECS OS uses the English keyboard by default.

- If you have changed the keyboard language of the ECS OS, select the keyboard language to which you have changed.
- 7. Click **Remote Login**.
- 8. (Optional) If you have changed the system language, in the dialog box that is displayed, click **Start Remote Login**.
- 9. (Optional) When the system displays "Press CTRL+ALT+DELETE to log on", click **Send CtrlAltDel** in the upper part of the remote login page to log in to the ECS.

Figure 3-13 Send CtrlAltDel



- 10. Enter the password obtained in 3.
 - When logging in to the ECS using VNC for the first time, use the English keyboard to enter the password. After you have logged in to the ECS, see Changing the OS Keyboard Language to change the keyboard language of the ECS OS. You can then select the keyboard language and enter the password the next time you log in.
 - If you have changed the keyboard language of the ECS OS, ensure that the keyboard language in use, the keyboard language selected in step 6, and the changed OS keyboard language are all the same.

Changing the OS Keyboard Language

If the ECS is running Linux, run the following command:

loadkeys keymapfile

The *keymapfile* parameter indicates the name of the file containing the mappings between the keys and displayed characters.

For example, if the name of a French keyboard mapping file is **fr**, run the **loadkeys fr** command.

Configuration Example

Scenarios

If you attempt to log in to an ECS created using a public image for the first time, the languages of the four types of keyboards before the configuration are as follows (**Before configuration** row in **Table 3-10**):

- Physical keyboard: French
- Input method keyboard on the terminal: English
- VNC keyboard: English
- ECS OS keyboard: English

In this case, you must change the languages of the other three types of keyboards to the same language as the physical keyboard for expected data entering. For details, see the **Solution 1** row in **Table 3-10**.

| - | Physical Keyboard | Input Method Keyboard on the Terminal | VNC Keyboard | ECS OS Keyboard |
|-----------------------------|----------------------|---|-----------------|-----------------|
| Before configu ration | French | English | English | English |
| Solutio n 1 | French | French | French | French |
| Solutio n 2 | English | English | English | English |

Table 3-10 Languages in the four types of keyboards

Procedure

- 1. Locally configure the language, for example, French, in the input method keyboard on the terminal.
- 2. Set the VNC keyboard language to English.
 - □ NOTE

When you log in to the ECS using VNC for the first time, the default ECS OS keyboard language is English. Therefore, you must set the VNC keyboard language to English.

- 3. Log in to the ECS and change the ECS OS language to French.
 - For details, see Changing the OS Keyboard Language.
- Change the VNC keyboard language to French.
 For details, see Logging In to an ECS Using a Non-English Keyboard.

To set the languages on the four types of keyboards to all be the same, perform 1 to 4.

□ NOTE

During the configuration, if English characters cannot be entered using the current physical keyboard, use the English soft keyboard to modify the configuration described in the **Solution 2** row of **Table 3-10**. In such a case, you only need to use the English soft keyboard to enter characters.

- To enable the Windows English soft keyboard, choose **Start** > **Run**, enter **osk**, and press **Enter**.
- The method of enabling the Linux English soft keyboard varies depending on the OS version and is not described in this document.

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 **Input Commands** in the upper right corner of the page.

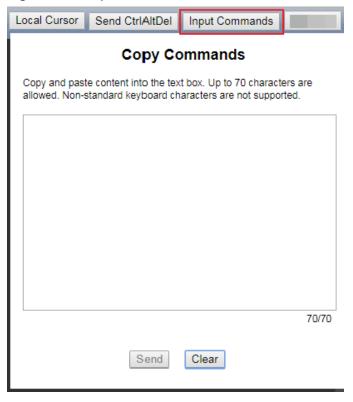


Figure 3-14 Input Commands

- 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.
- Click **Send**.Send the copied data to the CLI.

□ NOTE

There is a low probability that data is lost when you use Input Commands on the VNC page of a GUI-based Linux ECS. This is because the number of ECS vCPUs fails to meet GUI requirements. In such a case, it is a good practice to send a maximum of 5 characters at a time or switch from GUI to CLI (also called text interface), and then use the command input function.

Related Links

For FAQs about VNC-based ECS logins, see the following links:

- What Browser Version Is Required to Remotely Log In to an ECS?
- Why Cannot I Use the French Keyboard to Enter Characters When I Log In to an ECS Using VNC?
- 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?
- What Should I Do If I Cannot View Data After Logging In to an ECS Using VNC?
- Why Are Characters Entered Through VNC Still Incorrect After the Keyboard Language Is Switched?
- Why Does a Blank Screen Appear After I Attempted to Log In to an ECS Using VNC?

3.3.3 Remotely 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.
- You have obtained the image username. If a public image is used, the image username is **cloud**.

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.

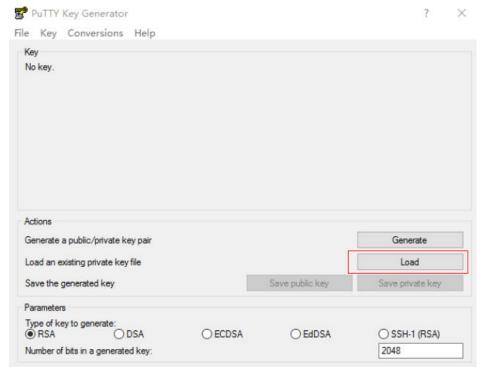


Figure 3-15 Importing the private key file

- 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)**.

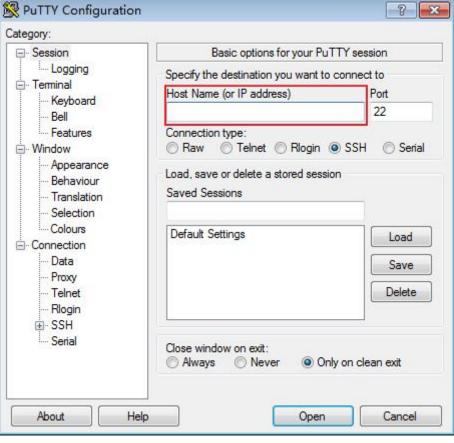


Figure 3-16 Configuring the EIP

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

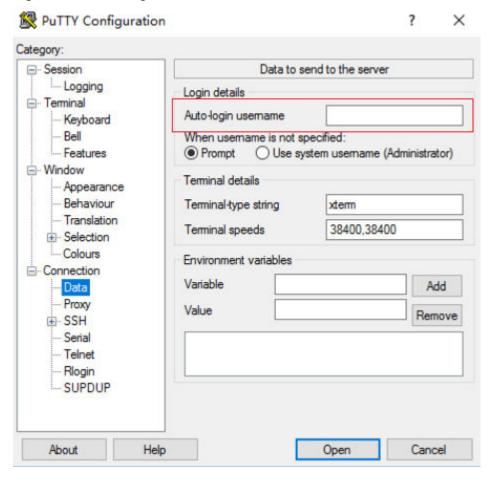


Figure 3-17 Entering the username

□ NOTE

If a public image is used, the image username is **cloud**.

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**.

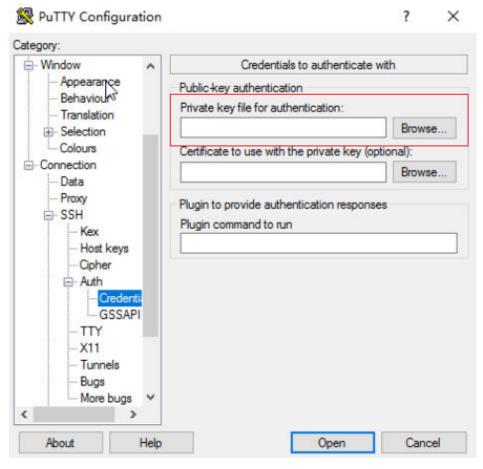


Figure 3-18 Importing the private key file

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:

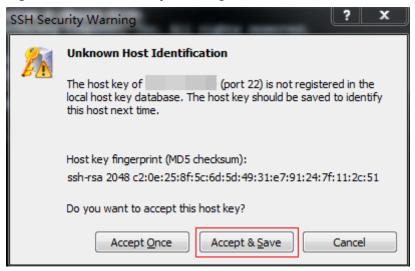
ssh *Username*@*EIP*

Ⅲ NOTE

If a public image is used, the image username is **cloud**.

3. (Optional) If the system displays the **SSH Security Warning** dialog box, click **Accept & Save**.

Figure 3-19 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

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 **cloud** and the EIP is **123.123.123.123**, run the following command:

ssh -i /path/kp-123.pem cloud@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. • After logging in to an ECS as a common user, you can run the following command to switch to user **root** without entering a password:

sudo su -

• After switching to user **root**, you can run the following command to change the password of user **root**:

sudo passwd root

Helpful Links

- Application Scenarios for Using Passwords
- Why Can't I Log In to My Linux ECS?

3.4 Managing ECSs

3.4.1 Changing ECS Names in a Batch

Scenarios

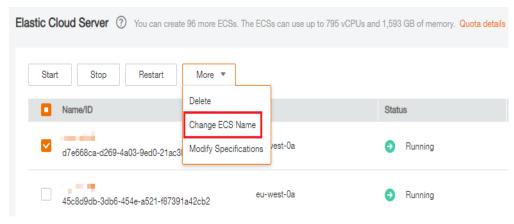
The name of a created ECS can be changed as needed.

Multiple ECS names can be changed in a batch. After the change, the ECS names are the same.

Procedure

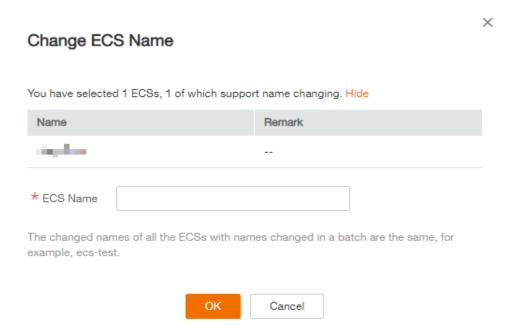
- 1. Log in to the management console.
- 2. Under Computing, click Elastic Cloud Server.
- 3. Select the ECS whose name you want to change.
- 4. Click **More** above the ECS list and select **Change ECS Name** from the drop-down list.

Figure 3-20 Changing the ECS name



5. Enter a new name.

Figure 3-21 Entering a new name



6. Click OK.

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

3.4.2 Reinstalling the OS

Scenarios

If the OS of an ECS fails to start or requires optimization, reinstall the OS.

Notes

- After the OS is reinstalled, the IP and MAC addresses of the ECS remain unchanged.
- Reinstalling the OS clears the data in all partitions of the EVS system disk, including the system partition. Therefore, back up data before reinstalling the OS
- Reinstalling the OS does not affect data disks.
- Do not perform any operations on the ECS immediately after its OS is reinstalled. Wait for several minutes until the system successfully injects the password or key. Otherwise, the injection may fail, and the ECS cannot be logged in to.
- After the OS is reinstalled, the password for logging in to the ECS will be reset. To retrieve the password, perform the following operations:
 - For a Linux ECS, log in to it using the key and set a new password. For instructions about how to log in to an ECS using a key pair, see Remotely Logging In to a Linux ECS (Using an SSH Key Pair).
 - For a Windows ECS, retrieve the password by following the instructions provided in Obtaining the Password for Logging In to a Windows ECS.

• 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.

Prerequisites

• The target ECS has a system disk attached.

Procedure

- 1. Log in to the management console.
- 2. Under Computing, choose Elastic Cloud Server.
- 3. Locate the row containing the target ECS and choose More > Manage Image/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.
- 4. (Optional) Select the **Encrypted** option to encrypt the system disk during OS reinstallation.

To enable encryption, click **Create Xrole** 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.

Encryption parameters are as follows:

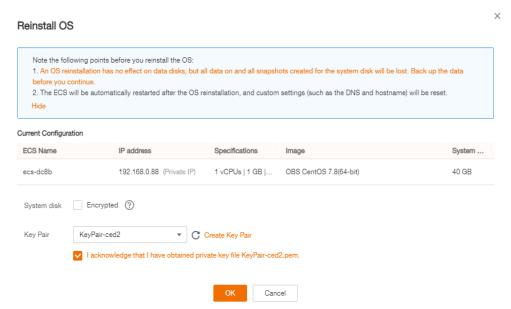
- Encryption: indicates that the EVS disk has been encrypted.
- Create Xrole: assigns KMS access permissions to EVS to obtain KMS keys.
 After the permissions are assigned, follow-up operations do not require assigning permissions again.
- Xrole Name: EVSAccessKMS: specifies that permissions have been 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.
- 5. (Optional) Select a **License Type** (**Use license from the system** or **Bring your own license (BYOL)**) if the reinstalled OS running on your ECS is billed. For more details, see **License Types**.

The following OSs are billed:

6. Select the login mode.

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

Figure 3-22 Reinstall OS



- 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.

Ⅲ NOTE

A temporary ECS is created during the reinstallation process. After reinstallation, this ECS will be automatically deleted. Do not perform any operation on the temporary ECS during the reinstallation process.

Follow-up Procedure

If the reinstallation fails, perform steps 2 to 8 again to retry the OS installation.

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

3.4.3 Changing the OS

Scenarios

Changing an ECS OS will change the system disk attached to the ECS. After the change, the system disk ID of the ECS will be changed, and the original system disk will be deleted.

If the OS running on an ECS cannot meet service requirements, change the ECS OS.

The cloud platform supports changing between image types (public images, private images, and shared images) and between OSs. You can change your OS by changing your ECS image.

Constraints

- The OS change takes about 10 to 20 minutes During this process, the ECS status is **Changing OS**.
- Do not perform any operations on the ECS before the system injects the password or key pair. 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.
- For details about the change between different OSs, see **Notes on Change Between Windows and Linux**.
- An ISO image created from an ISO file cannot be used to change the OS of an ECS. You need to install an OS and drivers on the ECS and use the ECS to create a system disk image first.
- 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 Data.
- 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.
- Do not perform any operations on the ECS immediately after its OS is changed. 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.
- After the OS is changed, the password for logging in to the ECS is reset. To retrieve the password, perform the following operations:
 - For a Linux ECS, log in to it using the key and set a new password. For instructions about how to log in to an ECS using a key pair, see Remotely Logging In to a Linux ECS (Using an SSH Key Pair).
 - For a Windows ECS, retrieve the password by following the instructions provided in Obtaining the Password for Logging In to a Windows ECS.

- The system disk capacity of an ECS with OS changed may change because the system disk capacity specified by the image of the changed OS may be changed.
- 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 Data.
- 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.

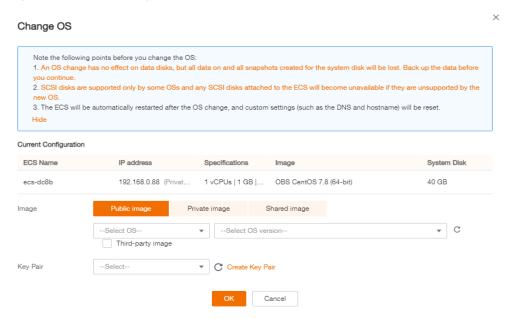
- If you plan to use a private image to change the OS, ensure that a private image is available. For details about how to create a private image, see Image Management Service User Guide.
 - If the image of a specified ECS is required, make sure that a private image has been created using this ECS.
 - If a local image file is required, make sure that the image file has been imported to the cloud platform and registered as a private image.
 - If a private image from another region is required, make sure that the image has been copied.
 - If a private image from another user account is required, make sure that the image has been shared with you.

Procedure

- 1. Log in to the management console.
- 2. Under **Computing**, choose **Elastic Cloud Server**.
- 3. Locate the row containing the target ECS and choose More > Manage Image/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.
- 4. Select the target image.

For more details, see Creating an ECS.

Figure 3-23 OS Change



5. (Optional) Select the **Encryption** option to encrypt the system disk during OS change.

To enable encryption, click **Create Xrole** 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.

Encryption parameters are as follows:

- Encryption: indicates that the EVS disk has been encrypted.
- Create Xrole: assigns KMS access permissions to EVS to obtain KMS keys.
 After the permissions are assigned, follow-up operations do not require assigning permissions again.
- Xrole Name: EVSAccessKMS: specifies 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. (Optional) Select a License Type (Use license from the system or Bring your own license (BYOL)) if the changed OS running on your ECS is billed. For more details, see License Types.

The following OSs are billed:

7. Configure the login mode.

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

8. Click OK.

9. On the **Change OS** page, confirm the specifications, and click **Submit**.

After the application is submitted, the status **Changing OS** is displayed. When this status disappears, the OS change is complete.

□ NOTE

A temporary ECS is created during the OS change process. After the process is complete, this ECS will be automatically deleted.

Follow-up Procedure

- If the OSs before and after the OS change are both Linux, and automatic mounting upon system startup has been enabled for data disks, the data disk partition mounting information will be lost after the OS is changed. In such a case, you need to update the /etc/fstab configuration.
 - a. Write the new partition information into /etc/fstab.

It is a good practice to back up the **/etc/fstab** file before writing data into it.

To enable automatic partition mounting upon system startup, see **Initializing a Linux Data Disk (fdisk)**.

- b. Mount the partition so that you can use the data disk.
 - mount Disk partition Device name
- c. Check the mount result.

df -TH

- If the OS change is unsuccessful, perform steps 2 to 9 again to retry the OS change.
- If the second OS change attempt is unsuccessful, contact customer service for manual recovery at the backend.

3.4.4 Managing ECS Groups

Scenarios

An ECS group logically classifies ECSs. ECSs in an ECS group comply with the same policy. Only the anti-affinity policy is available now, which enables the ECSs in the same ECS group to be deployed on different hosts to improve service reliability.

You can use an ECS group to deploy target ECSs on different hosts to ensure high service availability and underlying DR capabilities.

An ECS group supports the following functions:

- Creating an ECS Group
- Adding an ECS to an ECS Group
- Removing an ECS from an ECS Group
- Deleting an ECS Group

NOTICE

You are not allowed to add an existing ECS to an ECS group. Instead, add it only when creating the ECS.

Creating an ECS Group

Create an ECS group to apply the same policy to all group members. ECS groups are independent from each other.

- 1. Log in to the management console.
- 2. Under **Computing**, click **Elastic Cloud Server**.
- 3. In the navigation pane on the left, choose **ECS Group**.
- 4. On the ECS Group page, click Create ECS Group.
- 5. Enter an ECS group name.
- 6. Select a policy for the ECS group.
- 7. Click **OK**.

Adding an ECS to an ECS Group

Add an ECS to an ECS group only when creating the ECS.

NOTICE

- When you add an ECS to an ECS group, it will be deployed on a host different from other ECSs in this group for improved service reliability. 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.
- Only stopped ECSs can be added to an ECS group.
- Do not try to start the ECS when it is being added to an ECS group.
- 1. Log in to the management console.
- 2. Under Computing, choose Elastic Cloud Server.
- 3. In the navigation pane on the left, choose **ECS Group**.
- 4. Locate the row that contains the target ECS group and click **Add ECS** in the **Operation** column.
- 5. On the **Add ECS** page, select an ECS to be added.
- 6. Click **OK**. The ECS is added to the ECS group.

Removing an ECS from an ECS Group

After an ECS is removed from an ECS group, the ECS does not comply with the anti-affinity policy anymore.

- 1. Log in to the management console.
- 2. Under Computing, click Elastic Cloud Server.

- 3. In the navigation pane on the left, choose **ECS Group**.
- 4. Expand the ECS group information and view the ECSs in the ECS group.
- 5. Click **Remove** in the **Operation** column of the target ECS.
- 6. In the displayed dialog box, click **Yes**.

 The ECS is removed from the ECS group.

Deleting an ECS Group

After an ECS group is deleted, the policy does not apply to the ECSs in the ECS group anymore.

- 1. Log in to the management console.
- 2. Under Computing, click Elastic Cloud Server.
- 3. In the navigation pane on the left, choose **ECS Group**.
- 4. Click **Delete** in the **Operation** column of the target ECS group.
- 5. In the displayed dialog box, click Yes.

3.4.5 Changing the Time Zone for an ECS

Scenarios

The default time zone for an ECS is the one you selected when creating the image that was used to create the ECS. This section describes how to change the time zone for an ECS to the local one or to another time zone in your network.

After you log in to your ECS, if you find that the time on the ECS is different from the local time, you can change the time zone for the ECS so that the time on the ECS is the same as the local time.

For Linux ECSs

The process of changing the time zone for a Linux ECS depends on the OS. In this section, the CentOS 6.x 64bit OS is used to demonstrate how to change the time zone for a Linux ECS.

- 1. Log in to the ECS.
- 2. Run the following command to switch to user **root**:

su - root

3. Run the following command to obtain the time zones supported by the ECS:

ls /usr/share/zoneinfo/

In the terminal display, the /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.

For example:

- If you are to use the time zone for Paris, France, run the ls /usr/share/ zoneinfo/Europe command to obtain the directory /usr/share/zoneinfo/ Europe/Paris.
- 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.

For example:

If the target time zone is for Paris, France, change the ZONE entry value as follows:

ZONE="Europe/Paris"

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:

In -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

- 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 dialog box that is displayed, click **Change date and time settings**.

The **Date and Time** page is displayed.



Figure 3-24 Date and Time

- 3. Click **Change time zone**.
 - The **Time Zone Settings** page is displayed.
- 4. In the **Set the time zone** pane, choose the target time zone from the **Time zone** drop-down list.
- 5. Click OK.

3.4.6 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

- 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.

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

3.5 Modifying ECS Specifications

3.5.1 General Operations

Scenarios

If ECS specifications do not meet service requirements, you can modify the ECS specifications, including vCPUs and memory. Certain ECSs allow you to change their types when you modify their specifications.

- Before changing a Xen ECS to a KVM ECS, you need to manually install the
 required drivers on the ECS first, or the ECS will be unavailable after the
 modification is complete. For example, starting the OS will fail. The following
 section describes how to change a Xen ECS to a KVM ECS. For Linux,
 automatically changing a Xen ECS to a KVM ECS is recommended.
 - Changing a Xen ECS to a KVM ECS (Windows)
 - Automatically Changing a Xen ECS to a KVM ECS (Linux)
 - Manually Changing a Xen ECS to a KVM ECS (Linux)
- For specifications modifications between KVM ECSs or between Xen ECSs, see this section.

NOTE

- ECSs can be classified as the following based on the virtualization types:
 - Xen ECSs: S1, C1, C2, T2, and M1 ECSs.
 - KVM ECSs: See the Virtualization column in ECS Specifications and Types.
- Before changing a Xen ECS to a KVM ECS, install the required drivers on the ECS first, or the ECS will be unavailable after the modification is complete. For example, starting the OS will fail.
- When changing Xen to KVM for a Linux ECS, follow the instructions provided in Automatically Changing a Xen ECS to a KVM ECS (Linux).

Notes

- Downgrading ECS specifications (vCPU or memory) will reduce performance.
- Certain ECS types do not support specifications modification currently. For details about available ECS types as well as their functions and usage, see "NOTE" in ECS Types.
- When the disk status is **Expanding**, you are not allowed to modify the specifications of the ECS where the disk is attached.
- Before modifying the specifications of a Windows ECS, modify the SAN policy by following the instructions provided in Why Does a Disk Attached to a

Windows ECS Go Offline? to prevent disks from going offline after the specifications are modified.

Step 1: Modify Specifications

The following section provides general operations for modifying specifications. For instructions about how to change a Xen ECS to a KVM ECS, see the following operations.

- Changing a Xen ECS to a KVM ECS (Windows)
- Automatically Changing a Xen ECS to a KVM ECS (Linux)
- Manually Changing a Xen ECS to a KVM ECS (Linux)
- 1. Log in to the management console.
- 2. Under **Computing**, click **Elastic Cloud Server**.
- Click More in the Operation column and select Modify Specifications.
 The Modify ECS Specifications page is displayed.
- 4. Select the new ECS type, vCPUs, and memory as prompted.

 Before modifying the specifications, stop the ECS or select **Automatically stop the ECSs and then modify specifications**.
- 5. (Optional) Set **DeH**.
 - If the ECS is created on a DeH, the system allows you to change the DeH. To do so, select the target DeH from the drop-down list. If no DeH is available in the drop-down list, remaining DeH resources are insufficient and cannot be used to create the ECS with specifications modified.
- 6. (Optional) Select the check box to confirm the ECS configuration.
 If a Xen ECS is changed to a KVM ECS, configure the ECS by following the instructions provided in Changing a Xen ECS to a KVM ECS (Windows), Automatically Changing a Xen ECS to a KVM ECS (Linux), or Manually Changing a Xen ECS to a KVM ECS (Linux).
- 7. Click **Next**.
- 8. Confirm the settings and click **Submit Application**.

Step 2: Check Disk Attachment

After specifications are modified, disk attachment may fail. Therefore, check disk attachment after specifications modification. If disks are properly attached, the specifications modification is successful.

- Windows ECS
 - For details, see Why Do the Disks of a Windows ECS Go Offline After I Modify the ECS Specifications?
- Linux ECS

For details, see Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?

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.

3.5.2 Changing a Xen ECS to a KVM ECS (Windows)

Scenarios

Before changing a Xen ECS that runs Windows to a KVM ECS, make sure that PV driver and UVP VMTools have been installed on the ECS.

This section describes how to install the PV driver and UVP VMTools and change Xen to KVM.

- ECSs can be classified as the following based on the virtualization types:
 - Xen ECSs: S1, C1, C2, T2, and M1 ECSs.
 - KVM ECSs: See the Virtualization column in ECS Specifications and Types.
- Before changing a Xen ECS to a KVM ECS, install the required drivers on the ECS first, or the ECS will be unavailable after the modification is complete. For example, starting the OS will fail.
- When changing Xen to KVM for a Linux ECS, follow the instructions provided in Automatically Changing a Xen ECS to a KVM ECS (Linux).

Constraints

- If a Windows ECS is attached with a cross-region disk, the ECS specifications cannot be modified. Otherwise, ECS data may be lost.
- A Xen ECS with more than 24 VBD disks attached cannot be changed to a KVM ECS.
- A Xen ECS can be changed to a KVM ECS, but a KVM ECS cannot be changed to a Xen ECS.

Procedure

Figure 3-25 shows the flowchart for changing a Xen ECS to a KVM ECS.

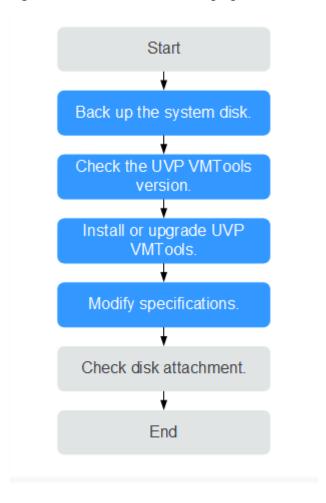


Figure 3-25 Flowchart for changing a Xen ECS to a KVM ECS

Table 3-11 describes the operations for changing a Xen ECS to a KVM ECS.

Table 3-11 Procedure for changing a Xen ECS to a KVM ECS

| Step | Operation |
|------|--|
| 1 | (Optional) Step 1: Back Up the System Disk |
| 2 | Step 2: Check the UVP VMTools Version |
| 3 | Step 3: Install or Upgrade UVP VMTools |
| 4 | Step 4: Modify Specifications |
| 5 | (Optional) Step 5: Check Disk Attachment |

(Optional) Step 1: Back Up the System Disk

If you modify the specifications of an ECS without installing the driver, the ECS may become unavailable and the data on the system disk may be lost. Therefore, back up the system disk first.

- 1. Before you create a system disk backup, check the ECS.
- 2. For instructions about how to back up the system disk, see**Creating a VBS Backup**.

Step 2: Check the UVP VMTools Version

Before modifying specifications, check the UVP VMTools version.

- 1. Log in to the ECS.
- 2. Download the driver check script.

Execute the script as the administrator and wait for the check result.

URL for downloading the script: https://latin-server-resize.obs.na-mexico-1.myhuaweicloud.com/windows/server_resize/check kvm drivers.vbs

After checking that the required driver has been installed, the system automatically tags the ECS. The specifications of only the tagged ECSs can be modified.

- If the check result is "Check version success!", the driver version meets service requirements and the ECS is tagged. Then, go to Step 4: Modify Specifications.
- If the check result is "Check version success but set metadata failed!
 Please run this script again later.", the driver version meets service requirements but tagging the ECS failed. In such a case, try again later.
- If the check result is "Check version failed! Please install drivers at first.", the driver version does not meet service requirements. In such a case, install or upgrade UVP VMTools by following the instructions provided in Step 3: Install or Upgrade UVP VMTools.

Step 3: Install or Upgrade UVP VMTools

When you install or upgrade UVP VMTools, if the PV driver has been installed on the ECS, the system will check the PV driver version. Ensure that the PV driver version meets service requirements. Otherwise, installing UVP VMTools will fail on the ECS. This section describes how to check the installation of the PV driver and UVP VMTools.

CAUTION

Before installing the PV driver or upgrading UVP VMTools, ensure that the ECS meets the following requirements:

- The available system disk size of the ECS is greater than 2 GB.
- Third-party virtualization platform tools, such as Citrix Xen Tools and VMware Tools, have been uninstalled to prevent driver installation failures. For instructions about how to uninstall the tools, see the official documents of the tools
- Antivirus software or intrusion detection software has been disabled. You can enable them after the driver is installed.

1. Check whether the PV driver version meets the UVP VMTools dependency requirements.

Switch to the C:\Program Files (x86)\Xen PV Drivers\bin directory, open the version.ini file, and view the PV driver version.

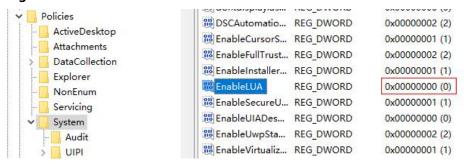
pvdriverVersion=5.0.104.010

- If the directory is available and the driver version is 5.0 or later, the PV driver meeting service requirements has been installed. In such a case, go to step 6 to install UVP VMTools.
- If the directory is unavailable or the driver version is earlier than 5.0, the PV driver has not been properly installed or the version does not meet service requirements. Then, see the following steps to uninstall the PV driver and install a new one.
- 2. Record the User Account Control (UAC) configuration of the ECS.

If the PV driver version is earlier than 5.0, DisableLUA is added to the registry during PV driver installation to prevent too many pop-up windows during driver upgrade, and EnableLUA is added to the registry during PV driver uninstallation (this has been resolved in PV driver 5.0 and later versions). To prevent adverse impact on your services, you need to record the UAC configuration before uninstalling the PV driver. Then check and restore the EnableLUA configuration in the registry after installing the new version. For details about UAC configurations, see official Microsoft documents.

- a. In the **Run** dialog box, enter **regedit** and click **OK** to open the registry editor.
- b. Record the HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows \CurrentVersion\Policies\System\EnableLUA value.

Figure 3-26 EnableLUA



- 3. Uninstall the PV driver of the old version.
 - a. On the ECS OS, choose **Start** > **Control Panel**.
 - b. Click **Uninstall a program**.
 - c. Uninstall **GPL PV Drivers for Windows** *x.x.x.xx* as prompted.
 - d. Restart the ECS on the management console.
- 4. Install the PV driver of the new version.
 - Download the PV driver installation package.
 Download the PV driver at https://ecs-instance-driver.obs.cn-north-1.myhuaweicloud.com/pvdriver-windows.zip.
 - b. Decompress the PV driver software package.

- c. Double-click pvdriver-windows.iso.
- d. Run Setup.exe and install the PV driver as prompted.
 Wait until the driver installation is complete. Do not click Setup.exe during the installation.
- e. Restart the ECS as prompted for the PV driver to take effect.
- 5. Check and restore the UAC configuration.
 - a. In the **Run** dialog box, enter **regedit** and click **OK** to open the registry editor.
 - b. Check the HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows \CurrentVersion\Policies\System\EnableLUA value and compare it with the value you recorded. If they are different, change the value to the one recorded in step 2.
- 6. Install or upgrade UVP VMTools.
 - Download the UVP VMTools installation package.
 Download UVP VMTools at https://ecs-instance-driver.obs.cn-north-1.myhuaweicloud.com/vmtools-windows.zip.
 - b. Decompress the UVP VMTools installation package.
 - c. Double-click vmtools-windows.iso.
 - d. Run **Setup.exe** and install UVP VMTools as prompted.

The installation program will automatically adapt to the OS version and identify whether UVP VMTools is newly installed or upgraded.

Wait until the installation is complete. Do not click **Setup.exe** during the installation.

- e. Restart the ECS as prompted for UVP VMTools to take effect.
- f. Check whether UVP VMTools has been installed. For details, see Step 2: Check the UVP VMTools Version.

Step 4: Modify Specifications

- 1. Log in to management console.
- 2. Under Computing, click Elastic Cloud Server.
- 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.
- 4. Click **More** in the **Operation** column and select **Modify Specifications**. The **Modify ECS Specifications** page is displayed.
- 5. Select the new ECS type, vCPUs, and memory as prompted.
- 6. (Optional) Set **DeH**.

If the ECS is created on a DeH, you can change the DeH where the ECS resides.

To do so, select the target DeH from the drop-down list. If no DeH is available in the drop-down list, it indicates that DeH resources are insufficient and cannot be used to create the ECS with specifications modified.

7. Select the check box to confirm that step ECS has been performed.

8. Click OK.

(Optional) Step 5: Check Disk Attachment

After a Xen ECS is changed to a KVM ECS, 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?

Follow-up Procedure

If the ECS specifications have been modified but the OS cannot be started after remote login, reinstall the ECS OS to resolve this issue. For details, see **Reinstalling the OS**.

3.5.3 Automatically Changing a Xen ECS to a KVM ECS (Linux)

Scenarios

Before changing a Xen ECS that runs Linux to a KVM ECS, make sure that the required drivers have been installed and configured on the ECS.

This section describes how to use a script to automatically install drivers on the ECS, configure the device name, and change Xen to KVM.

- Xen ECSs include S1, C1, C2, T2, and M1 ECSs.
- To obtain KVM ECSs, see the Virtualization column in ECS Specifications and Types.
- To support both Xen and KVM, Linux ECSs require the xen-pv and virtio drivers. Before
 changing a Xen ECS to a KVM ECS, make sure that the Linux ECS has been configured,
 including driver installation and automatic disk attachment.

Constraints

- To prevent data loss, the specifications of Linux ECSs that use LVM or RAID arrays cannot be modified.
- A Xen ECS with more than 24 VBD disks attached cannot be changed to a KVM ECS.
- A Xen ECS can be changed to a KVM ECS, but a KVM ECS cannot be changed to a Xen ECS.

Procedure

Figure 3-27 shows the flowchart for automatically changing a Xen ECS to a KVM ECS.

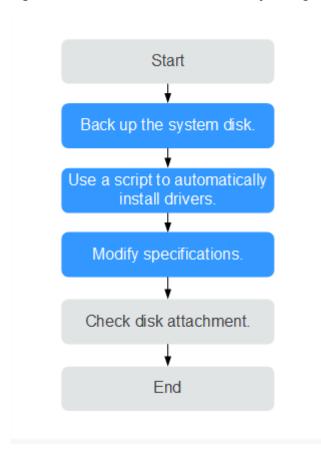


Figure 3-27 Flowchart for automatically changing a Xen ECS to a KVM ECS

Table 3-12 OSs that support automatic configuration using a script

| OS | Version |
|----------|----------------|
| CentOS | CentOS 7 |
| | CentOS 6 |
| Debian | Debian 9 |
| | Debian 8 |
| EulerOS | • EulerOS 2 |
| OpenSUSE | OpenSUSE 42 |
| | OpenSUSE 15 |
| SUSE | • SUSE 15 |
| | • SUSE 12 |
| | • SUSE 11 |
| SUSE-SAP | SUSE-SAP 12 |
| Ubuntu | • Ubuntu 18.04 |
| | • Ubuntu 16.04 |
| | • Ubuntu 14.04 |

| OS | Version |
|--------------|----------------|
| Red Hat | • Red Hat 7 |
| | Red Hat 6 |
| Fedora | • Fedora 29 |
| | Fedora 28 |
| | Fedora 27 |
| | Fedora 26 |
| Oracle Linux | Oracle Linux 7 |

Table 3-13 describes the operations for automatically changing a Xen ECS to a KVM ECS using a script.

Table 3-13 Procedure for automatically changing a Xen ECS to a KVM ECS using a script

| Step | Operation |
|------|---|
| 1 | (Optional) Step 1: Back Up System Disk Data |
| 2 | Step 2: Using a Script to Automatically Install Drivers |
| 3 | Step 3: Modify Specifications |
| 4 | (Optional) Step 4: Check Disk Attachment |

(Optional) Step 1: Back Up System Disk Data

If you modify the specifications of an ECS without installing the driver, the ECS may become unavailable and the data on the system disk may be lost. Therefore, back up the system disk first.

- Before you create a system disk backup, check the ECS.
 Stop and then start the ECS to ensure that services can run properly after the ECS is started. Back up the system disk.
- 2. For instructions about how to back up the system disk, see**Creating a VBS Backup**.

Step 2: Using a Script to Automatically Install Drivers

Use a script to install drivers on an ECS. If your ECS does not support configuration using a script, manually configure it by referring to **Manually Changing a Xen ECS to a KVM ECS (Linux)**.

□ NOTE

For details about the ECS OSs that support automatic configuration using a script, see Table 3-12.

- Log in to the ECS.
- 2. Run the following command to download the driver installation script to the **root** directory:

curl URL > ~/resize_ecs_modify_linux.sh

In the preceding command, *URL* is the address for downloading the specifications modification script.

Select an address for downloading the optimization script based on the region where the ECS is located:

URL for downloading the script: https://latin-server-resize.obs.namexico-1.myhuaweicloud.com/linux/server_resize/ resize_ecs_modify_linux.sh

Run the following command to execute the script which automatically checks and installs the native Xen PV driver and virtio driver:

bash resize_ecs_modify_linux.sh

Figure 3-28 Executing the script

4. Wait until the script is executed.

After checking that the required driver has been installed, the system automatically tags the ECS. The specifications of only the tagged ECSs can be modified.

If the check result is "{Image name} already contain xen and virtio driver", the driver has been installed.

- If the check result is "Success to set kvm meta!" or "this server already has kvm meta.", the ECS has been tagged. Then, go to Step 3: Modify Specifications.
- If the check result is "Failed to set metadata, please try again.", tagging the ECS failed. In such a case, try again later.

If the installation failed, manually configure the ECS by following the instructions provided in Manually Changing a Xen ECS to a KVM ECS (Linux) or contact customer service.

Figure 3-29 Successful script execution

```
3762] device-mapper: uevent: version 1.8.3
753] device-mapper: ioctl: 4.37.1-ioctl (2018-04-03) initialised: dm-devel@redhat.com
ymuk configuration file ...
x image: /boot/vmlinuz-3.18.9-1862.12.1.e17.x86.64
rd image: /boot/vmlinuz-3.18.9-1862.12.1.e17.x86.64.img
x image: /boot/vmlinuz-3.18.9-957.e17.x86.64.img
x image: /boot/vmlinuz-3.18.9-957.e17.x86.64.img
x image: /boot/vmlinuz-9-rescue-0875b01876e04499976c7ec5ad9839ef
rd image: /boot/vmlinuz-9-rescue-0875b01876e04499976c7ec5ad9839ef
rd image: /boot/vmlitramfs-9-rescue-0875b01876e04499976c7ec5ad9839ef
rd image: /boot/vmlitramfs-9-rescue-0875b01876e054499976c7ec5ad8839ef.img
361] SGI XS with RCLs, security attributes, no debug enabled
514] xor: automatically using best checksumming function:
252448.800 MB2 yttributes.
2661 raid6: sse2x1 gen() 7382 MB/s
2661 raid6: sse2x2 gen() 8589 MB/s
381 raid6: sse2x4 gen() 18273 MB/s
391 raid6: avv2x2 gen() 16589 MB/s
661 raid6: avv2x2 gen() 16589 MB/s
662 raid6: swig algorithm avx2x2 gen() (18976 MB/s)
393 raid6: using algorithm avx2x2 gen() (18976 MB/s)
394 raid6: using algorithm avx2x2 gen() (18976 MB/s)
395 raid6: using algorithm avx2x2 gen() (18976 MB/s)
395 raid6: using algorithm avx2x2 gen() (18976 MB/s)
396 raid6: using algorithm avx2x2 gen() (18976 MB/s)
397 raid6: using algorithm avx2x2 gen() (18976 MB/s)
398 raid6: using algorithm avx2x2 gen() (18976 MB/s)
399 raid6: using algorithm avx2x2 gen() (18976 MB/s)
    09-24 15:12:13 Info:check xen∕ide driver is already exist in /boot/initramfs-0-rescue-0f5b010f6eb344909f6efec5ad0039ef.in
        .
19-24 15:12:23 Info:xen driver:yes
19-24 15:12:23 Info:ide driver:no
19-24 15:12:23 Info:check virtio driver is already exist in /boot/initramfs-0-rescue-0f5b010f6eb344909f6cfec5ad0039ef.img
89-24 15:12:41 Info:check xen/ide driver: yes
89-24 15:12:41 Info:check xen/ide driver is already exist in /boot/initramfs-3.18.8-1862.12.1.e17.x86_64.img or not
89-24 15:12:45 Info:check xen/ide driver: yes
89-24 15:12:45 Info:check virtio driver: yes
89-24 15:12:45 Info:check virtio driver: yes
89-24 15:12:52 Info:check ven/ide driver is already exist in /boot/initramfs-3.18.8-1862.12.1.e17.x86_64.img or not
89-24 15:12:52 Info:check xen/ide driver is already exist in /boot/initramfs-3.18.8-957.e17.x86_64.img or not
89-24 15:12:55 Info:xen driver: yes
89-24 15:12:55 Info:check ven/ide driver is already exist in /boot/initramfs-3.18.8-957.e17.x86_64.img or not
89-24 15:12:55 Info:check virtio driver: yes
89-24 15:12:55 Info:check virtio driver: yes
89-24 15:12:38 Info:check virtio driver: yes
89-24 15:13:38 Info:check virtio driver: yes
```

◯ NOTE

- Make sure that the ECS has been configured successfully, or the ECS may become unavailable after the specifications are modified. If the operation failed, follow the instructions provided in Manually Changing a Xen ECS to a KVM ECS (Linux) for manual operations.
- FAQs related to a script installation failure:
 - What Should I Do If Executing a Driver Installation Script Failed on an **ECS Running CentOS 5?**
 - What Should I Do If Executing a Driver Installation Script Failed When I Attempted to Modify the Specifications of a Linux ECS?

Step 3: Modify Specifications

- 1. Log in to management console.
- Under Computing, click Elastic Cloud Server.
- 3. 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**.
- 4. Click **More** in the **Operation** column and select **Modify Specifications**. The **Modify ECS Specifications** page is displayed.
- Select the new ECS type, vCPUs, and memory as prompted.
- (Optional) Set **DeH**.

If the ECS is created on a DeH, you can change the DeH where the ECS

To do so, select the target DeH from the drop-down list. If no DeH is available in the drop-down list, it indicates that DeH resources are insufficient and cannot be used to create the ECS with specifications modified.

- 7. Select the check box to confirm that the configuration is complete.
- 8. Click OK.

(Optional) Step 4: Check Disk Attachment

After a Xen ECS is changed to a KVM ECS, disk attachment may fail. Therefore, check disk attachment after specifications modification. If disks are properly attached, the specifications modification is successful.

Linux ECS

For details, see Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?

Follow-up Procedure

If the ECS specifications have been modified but the OS cannot be started after remote login, reinstall the ECS OS to resolve this issue. For details, see **Reinstalling the OS**.

3.5.4 Manually Changing a Xen ECS to a KVM ECS (Linux)

Scenarios

Before changing a Xen ECS that runs Linux to a KVM ECS, install and configure required drivers.

This section describes how to manually install drivers on a Linux ECS, configure automatic disk attachment, and change Xen to KVM.

For instructions about how to use a script to automatically install drivers, see **Automatically Changing a Xen ECS to a KVM ECS (Linux)**.

- Xen ECSs include S1, C1, C2, T2, and M1 ECSs.
- To obtain KVM ECSs, see the Virtualization column in ECS Specifications and Types.
- To support both Xen and KVM, Linux ECSs require the xen-pv and virtio drivers. Before changing a Xen ECS to a KVM ECS, make sure that the Linux ECS has been configured, including driver installation and automatic disk attachment.

Constraints

- To prevent data loss, the specifications of Linux ECSs that use LVM or RAID arrays cannot be modified.
- A Xen ECS with more than 24 VBD disks attached cannot be changed to a KVM ECS.
- A Xen ECS can be changed to a KVM ECS, but a KVM ECS cannot be changed to a Xen ECS.

Procedure

Figure 3-30 shows the flowchart for manually changing a Xen ECS to a KVM ECS.

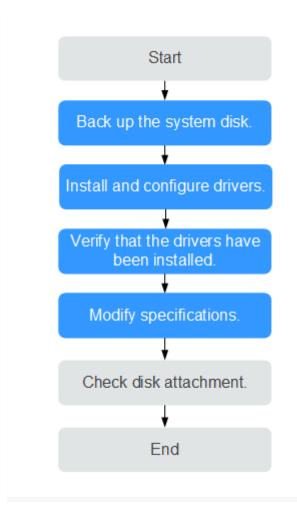


Figure 3-30 Flowchart for manually changing a Xen ECS to a KVM ECS

Table 3-14 Procedure for manually changing a Xen ECS to a KVM ECS

| Step | Task |
|------|---|
| 1 | (Optional) Step 1: Back Up System Disk Data |
| 2 | Step 2: Install Drivers |
| 3 | Step 3: Check Whether the ECS Is Configured Correctly |
| 4 | Step 4: Modify Specifications |
| 5 | (Optional) Step 5: Check Disk Attachment |

(Optional) Step 1: Back Up System Disk Data

If you modify the specifications of an ECS without installing the driver, the ECS may become unavailable and the data on the system disk may be lost. Therefore, back up the system disk first.

- Before you create a system disk backup, check the ECS.
 Stop and then start the ECS to ensure that services can run properly after the ECS is started. Back up the system disk.
- 2. For instructions about how to back up the system disk, see**Creating a VBS Backup**.

Step 2: Install Drivers

Perform the following operations to manually install drivers on an ECS.

- 1. Log in to the ECS.
- 2. Uninstall tools from the ECS.

For details, see "Optimizing a Linux Private Image" in *Image Management Service User Guide*.

Change the GRUB disk ID to UUID.

For details, see "Optimizing a Linux Private Image" in *Image Management Service User Guide*.

4. Change the fstab disk ID to UUID.

For details, see "Optimizing a Linux Private Image" in *Image Management Service User Guide*.

5. Install native Xen and KVM drivers.

For details, see "Optimizing a Linux Private Image" in *Image Management Service User Guide*.

Step 3: Check Whether the ECS Is Configured Correctly

Perform the following operations to check whether the drivers have been installed and the configuration files have been modified.

◯ NOTE

Before manually modifying specifications, make sure that the ECS has been configured correctly.

- 1. Log in to the ECS.
- 2. Run the following command to check whether the root partition is in UUID format:

cat /boot/grub/grub.cfg

- If yes, the disk ID in the GRUB configuration file has been changed to UUID.
- If no, the modification failed. In such a case, change the GRUB disk ID to UUID again by referring to Step 2: Install Drivers.

...menuentry 'Ubuntu Linux, with Linux 3.13.0-24-generic' --class ubuntu --class gnu-linux --class gnu --class os --unrestricted \$menuentry_id_option 'gnulinux-3.13.0-24-generic-advanced-ec51d860-34bf-4374-ad46-a0c3e337fd34' {
recordfail
load_video
gfxmode \$linux_gfx_mode
insmod gzio
insmod part_msdos
insmod ext2
if [x\$feature_platform_search_hint = xy]; then

```
search --no-floppy --fs-uuid --set=root ec51d860-34bf-4374-ad46-a0c3e337fd34 else search --no-floppy --fs-uuid --set=root ec51d860-34bf-4374-ad46-a0c3e337fd34 fi echo 'Loading Linux 3.13.0-24-generic ...' linux /boot/vmlinuz-3.13.0-24-generic root=UUID=ec51d860-34bf-4374-ad46-a0c3e337fd34 ro echo 'Loading initial ramdisk ...' initrd /boot/initrd.img-3.13.0-24-generic }
```

∩ NOTE

The path in which the GRUB configuration file is stored varies depending on the OS. For example, the path can be /boot/grub/menu.lst, /boot/grub/grub.cfg, /boot/grub2/grub.cfg, or /boot/grub/grub.conf.

3. Run the following command to check whether the disk ID in the fstab configuration file is UUID:

cat /etc/fstab

- If yes, the disk ID has been changed to UUID.
- If no, the modification failed. In such a case, change the fstab disk ID to UUID again by referring to Step 2: Install Drivers.

```
[root@****** ~]# cat /etc/fstab

UUID=4eb40294-4c6f-4384-bbb6-b8795bbb1130 / xfs defaults 0 0

UUID=2de37c6b-2648-43b4-a4f5-40162154e135 swap swap defaults 0 0
```

- 4. Check whether the native Xen and KVM drivers have been installed.
 - If the boot virtual file system is initramfs, run the following commands: lsinitrd /boot/initramfs-`uname -r`.img | grep `uname -r` | grep xen lsinitrd /boot/initramfs-`uname -r`.img | grep `uname -r` | grep virtio
 - If the boot virtual file system is initrd, run the following commands: lsinitrd /boot/initrd-`uname -r` | grep `uname -r` | grep xen lsinitrd /boot/initrd-`uname -r` | grep `uname -r` | grep virtio

If the names of the native Xen and KVM drivers are displayed in the command output, the drivers have been installed.

```
[root@CTU10000xxxxx home]# lsinitrd /boot/initramfs-`uname -r`.img | grep `uname -r`| grep xen
                             54888 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
-rwxr--r-- 1 root root
block/xen-blkfront.ko
                             45664 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/
-rwxr--r-- 1 root root
drivers/net/xen-netfront.ko
[root@CTU10000xxxxx home]# lsinitrd /boot/initramfs-`uname -r`.img | grep `uname -r`| grep virtio
                             23448 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
-rwxr--r-- 1 root root
block/virtio_blk.ko
-rwxr--r-- 1 root root
                             50704 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/
drivers/net/virtio_net.ko
-rwxr--r-- 1 root
                   root
                             28424 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
scsi/virtio_scsi.ko
                                 0 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
drwxr-xr-x 2 root
                   root
virtio
-rwxr--r-- 1 root
                             14544 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
                    root
virtio/virtio.ko
-rwxr--r-- 1 root
                             21040 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
                    root
virtio/virtio_pci.ko
-rwxr--r-- 1 root
                             18016 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
                   root
virtio/virtio_ring.ko
```

Ⅲ NOTE

Make sure that the ECS has been configured successfully, or the ECS may become unavailable after the specifications are modified.

Step 4: Modify Specifications

- 1. Log in to management console.
- 2. Under Computing, click Elastic Cloud Server.
- 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.
- 4. Click **More** in the **Operation** column and select **Modify Specifications**. The **Modify ECS Specifications** page is displayed.
- 5. Select the new ECS type, vCPUs, and memory as prompted.
- 6. (Optional) Set **DeH**.

If the ECS is created on a DeH, you can change the DeH where the ECS resides.

To do so, select the target DeH from the drop-down list. If no DeH is available in the drop-down list, it indicates that DeH resources are insufficient and cannot be used to create the ECS with specifications modified.

- 7. Select the check box to confirm that the configuration is complete.
- 8. Click OK.

(Optional) Step 5: Check Disk Attachment

After a Xen ECS is changed to a KVM ECS, disk attachment may fail. Therefore, check disk attachment after specifications modification. If disks are properly attached, the specifications modification is successful.

Linux ECS

For details, see Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?

Follow-up Procedure

If the ECS specifications have been modified but the OS cannot be started after remote login, reinstall the ECS OS to resolve this issue. For details, see **Reinstalling the OS**.

3.6 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.

If no DeHs are available, create a DeH first. For details, see Allocating DeHs in the *Dedicated Host User Guide*.

5. Click OK.

3.7 Obtaining Metadata and Passing User Data

3.7.1 Obtaining Metadata

Scenarios

ECS metadata includes basic information of an ECS on the cloud platform, such as the ECS ID, hostname, and network information. ECS metadata can be obtained using either OpenStack or EC2 compatible APIs, as shown in **Table 3-15**. 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 3-15 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 3-15 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 3-16. |
| 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. |

| Metadata Type | Metadata Item | Description | | | |
|----------------|--|---|--|--|--|
| 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.jso n | Displays ECS network information. | | | |
| OpenStack | /securitykey | Obtains temporary AKs and SKs. | | | |
| | | Before enabling an ECS to obtain a temporary AK and SK, make sure that the op_svc_ecs account has been authorized on IAM and that the desired ECS resources have been authorized for management. | | | |
| EC2-compatible | /meta-data/ hostname | Displays the name of the host accommodating an ECS. | | | |
| | | To remove the suffix .novalocal from an ECS, see: | | | |
| | | Is an ECS Hostname with Suffix .novalocal Normal? | | | |
| EC2-compatible | /meta-data/ local-hostname | The meaning of this field is the same as that of hostname. | | | |
| EC2-compatible | /meta-data/ public-hostname | The meaning of this field is the same as that of hostname. | | | |
| EC2-compatible | /meta-data/ instance-type | Displays an ECS flavor. | | | |
| EC2-compatible | /meta-data/ | Displays the fixed IP address of an ECS. | | | |
| | local-ipv4 | If there are multiple NICs, only the IP address of the primary NIC is displayed. | | | |
| EC2-compatible | /meta-data/ placement/ availability-zone | Displays the AZ accommodating an ECS. | | | |
| EC2-compatible | /meta-data/ | Displays the EIP bound to the ECS. | | | |
| | public-ipv4 | If there are multiple NICs, only the EIP of the primary NIC is displayed. | | | |

| Metadata Type | Metadata Item | Description |
|----------------|--|--|
| 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 3-16 Metadata key fields

| Parameter | Туре | Description |
|-----------------------|--------|--|
| uuid | String | Specifies an ECS ID. |
| availability_zon e | String | Specifies the AZ where an ECS locates. |
| meta | Dict | Specifies the metadata information, including the image name, image ID, and VPC ID. |
| hostname | String | Specifies the name of the host accommodating an ECS. To remove the suffix .novalocal from an ECS, see: Is an ECS Hostname with Suffix .novalocal Normal? |

Prerequisites

- The target ECS has been logged in.
- Security group rules in the outbound direction meet the following requirements:

Protocol: TCP

- Port: 80

- Destination: 169.254.0.0/16

Ⅲ 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+dNwlYdGIxJHUp+00poeUsAdBFkbPbYQTmpNwpoEb43k9z+96TyrekNKS
+iLYDdRNy4kKGoNPEVBCc05Hg1TcDblAPfJwgJS1okgEtlcofUhKmL3K0fto
+5KXEDU3GNuGwyZXjdVb9HQWU+E1jztAJjjqsahnU+g/tawABTVySLBKlAT8fMGax1mTGgArucn/
WzDcy19DGioKPE7F8ILtSQ4Ww3VClK5VYB/h0x+4r7IVHrPmYX/
bi1Yhm3Dc4rRYNaTjdOV5qUOsbO3oAeQkmKwQ/
NO0N8qw5Ya4l8ZUW4tMav4mOsRySOOB35v0bvaJc6p
+50DTbWNeX5A2MLiEhTP3vsPrmvk4LRF7CLz2J2TGIM14OoVBw7LARwmv9cz532zHki/c8tlhRzLmOTXh/
wL36zFW10DeuReUGmxth7IGNmRMQKV6+mil78jm/KMPpgAdK3vwYF/
GcelOFJD2HghMUUCeMbwYnvijLTejuBpwhJMNiHA/NvlEsxJDxqBCoss/Jfe+yCmUFyxovJ
+L8oNkTzkmtCNzw3Ra0hiKchGhqK3BleToV/kVx5DdF081xrEA
+qyoM6CVyfJtEoz1zlRRyoo9bJ65Eg6JJd8dj1UCVsDqRY1pIjgzE/
Mzsw6AaaCVhaMJL7u7YMVdyKzA6z65Xtvujz0Vo="
  "uuid": "ca9e8b7c-f2be-4b6d-a639-f10b4d994d04",
  "availability zone": "lt-test-1c",
  "hostname": "ecs-ddd4.novalocal",
  "launch_index": 0,
  "meta": {
     "metering.image_id": "3a64bd37-955e-40cd-ab9e-129db56bc05d",
    "metering.imagetype": "gold",
"metering.resourcespeccode": "s3.medium.2.linux",
    "image_name": "CentOS 7.6 64bit",
     "metering.resourcetype": "1"
     "vpc_id": "3b6c201f-aeb3-4bce-b841-64756e66cb49",
     "os_bit": "64",
     "cascaded.instance_extrainfo": "pcibridge:1",
    "os_type": "Linux",
"charging_mode": "0"
  "project_id": "6e8b0c94265645f39c5abbe63c4113c6",
  "name": "ecs-ddd4"
```

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 3Rpb24gYW5kIGF0IHN1Y2ggYSBcGVlZC4uLkl0IGZlZWxzIGFuIGltcHVsc2lvbi4uLnRoaXMgaXMgdGhlIH BsYWNlIHRvIGdvIG5vdy4gQnV0IHRoZSBza3kga25vd3MgdGhlIHJlYXNvbnMgYW5kIHRoZSBwYXR0ZXJu cyBiZWhpbmQgYWxsIGNsb3VkcywgYW5kIHlvdSB3aWxsIGtub3csIHRvbywgd2hlbiB5b3UgbGlmdCB5b3 Vyc2VsZiBoaWdoIGVub3VnaCB0byBzZWUgYmV5b25kIGhvcml6b25zLiINCg0KLVJpY2hhcmQgQmFjaA=

◯ NOTE

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

Figure 3-31 404 Not Found

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.

Ⅲ 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

ICAgICAgDQoiQSBjbG91ZCBkb2VzIG5vdCBrbm93IHdoeSBpdCBtb3ZlcyBpbiBqdXN0IHN1Y2ggYSBkaXJlY 3Rpb24gYW5kIGF0IHN1Y2ggYSBcGVIZC4uLkl0IGZIZWxzIGFuIGltcHVsc2lvbi4uLnRoaXMgaXMgdGhlIH BsYWNllHRvIGdvIG5vdy4gQnV0IHRoZSBza3kga25vd3MgdGhlIHJlYXNvbnMgYW5kIHRoZSBwYXR0ZXJu cyBiZWhpbmQgYWxsIGNsb3VkcywgYW5kIHlvdSB3aWxsIGtub3csIHRvbywgd2hlbiB5b3UgbGlmdCB5b3 Vyc2VsZiBoaWdoIGVub3VnaCB0byBzZWUgYmV5b25kIGhvcml6b25zLiINCg0KLVJpY2hhcmQgQmFjaA=

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.

LIRI

/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.

a IID

/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 AAAAB3NzaC1yc2EAAAADAQABAAABAQDI5Fw5k8Fgzajn1zJwLoV3+wMP+6CyvsSilc/hioggSnYu/AD0Yqm8vVO0kWlun1rFbdO+QUZKyVr/OPUjQSw4SRh4qsTKf/+eFoWTjplFvd1WCBZzS/WRenxlwR00KkczHSJro763+wYcwKieb4eKRxaQoQvoFgVjLBULXAjH4eKoKTVNtMXAvPP9aMy2SLgsJNtMb9ArfziAiblQynq7UIfLnN3VclzPeiWrqtzjyOp6CPUXnL0lVPTvbLe8sUteBsJZwlL6K4i+Y0lf3ryqnmQgC21yW4Dzu+kwk8FVT2MgWkCwiZd8gQ/+uJzrJFyMfUOBIklOBfuUENIJUhABGenerated-by-Nova

Helpful Links

Why Can't My Linux ECS Obtain Metadata?

3.7.2 Passing User Data to ECSs

Scenarios

Specify **User Data** to pass user data to ECSs to:

- Simplify ECS configuration.
- Initialize the ECS OS configuration.
- Upload your scripts to ECSs during ECS creation.
- Perform other tasks using scripts.

Use Restrictions

- Linux
 - The image that is used to create ECSs must have Cloud-Init installed.
 - The user data to be specified must be less than or equal to 32 KB.
 - If user data is uploaded as text, the data can contain only ASCII characters. If user data is uploaded using a file, the file can contain any characters and the file size cannot exceed 32 KB.
 - The image that is used to create ECSs must be a public image, a private image created based on a public image, or a private image with Cloud-Init installed.
 - The format of the customized scripts must be supported by Linux ECSs.
 - DHCP must be enabled on the VPC network, and port 80 must be enabled for the security group in the outbound direction.

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 3-32 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.

By default, the scripts are executed as user **root**.

 Script type: Both user-data scripts and Cloud-Config data scripts are supported.

| Table 3-17 Linux ECS script type: |
|--|
|--|

| - | 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. | The first line must be #cloud-config, and no space is allowed in front of it. |
| | 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. | |
| 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

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 3-33 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.

| | Batch-Processing Prog | | | |
|------------------------------------|-----------------------|--|--|--|
| Table 3-18 Windows ECS script type | | | | |

| - | 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.

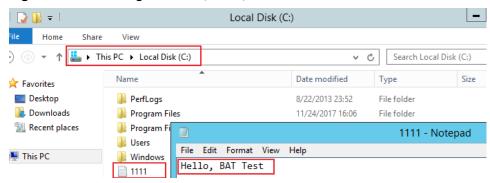
Example 1: Inject a batch-processing program script.

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 3-34 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 3-35 Viewing user data (Batch)



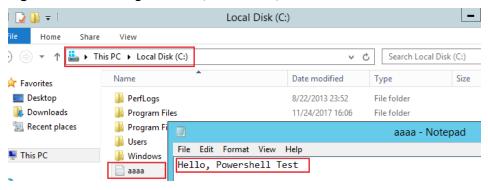
Example 2: Inject a PowerShell script.

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 3-36 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 3-37 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 3-19.

| Parameter | Requirement | Example Value | | |
|-----------|---|---|--|--|
| Password | Consists of 8 to 26 characters. Contains at least three of the following character types: Uppercase letters Lowercase letters Digits Special characters: \$!@%=+[]:./^,{}? 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 | YNbUwp! dUc9MClnv NOTE The example password is generated randomly. Do not use it. | | |

Table 3-19 Password complexity requirements

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 3-20.

Table 3-20 Password complexity requirements

| Parameter | Requirement | Example Value |
|-----------|---|--|
| Password | Consists of 8 to 26 characters. Contains at least three of the following character types: Uppercase letters Lowercase letters Digits Special characters: \$!@%=+[]:./^,{}? 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. | YNbUwp! dUc9MClnv NOTE The example password is generated randomly. Do not use it. |

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 3-20.

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/

3.8 (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

3.9 (Optional) Installing a Driver and Toolkit

3.9.1 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 purchased and 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 GPUaccelerated ECS.

Table 3-21 Acceleration supported by GPU drivers

| Dri ver | Lice nse | CUDA | Open GL | Direct X | Vulka n | Applicati on Scenario | Description |
|------------|--------------|---------------|---------------|---------------|---------------|--|--|
| GRI D | Requ ired | Suppo rted | Suppo rted | Suppo rted | Suppo rted | 3D rendering, graphics workstati on, and game accelerati on | The GRID driver must be paid and requires a license to accelerate graphics and image applications. |

| Dri ver | Lice nse | CUDA | Open GL | Direct X | Vulka n | Applicati on Scenario | Description |
|------------|---------------------|---------------|----------------------|----------------------|----------------------|--|--|
| Tes la | Not requi red | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted | Scientific computin g, 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. |

3.9.2 Installing a GRID Driver on a GPU-accelerated ECS

Scenarios

To use graphics acceleration, such as OpenGL, DirectX, or Vulkan, install a GRID driver and separately purchase and configure a GRID license. The GRID driver with a vDWS license also supports CUDA for both computing and graphics acceleration.

- A graphics-accelerated (G series) ECS created using a public image has had a GRID driver of a specified version installed by default, but the GRID license must be purchased and configured separately.
- If a GPU-accelerated ECS is created using a private image, install a GRID driver and separately purchase and configure a GRID license.

This section describes how to install a GRID driver, purchase or apply for a GRID license, and configure the license server.

Process of installing a GRID driver:

- 1. Purchasing 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

- 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.

Purchasing a GRID License

Purchase a license.

To obtain an official license, contact NVIDIA or their NVIDIA agent in your local country or region.

Apply for a trial license.

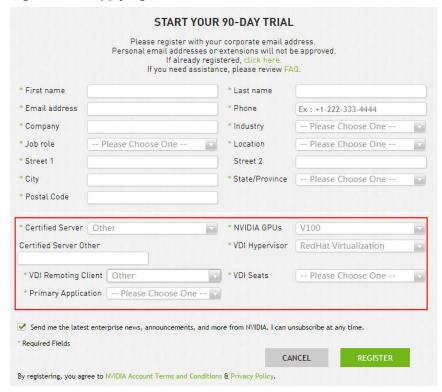
Log in at the official NVIDIA website and enter desired information.

For details about how to register for an account and apply for a trial license, see official NVIDIA help page.

Ⅲ 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. Purchase an official license then.

Figure 3-38 Applying for a trial license



Downloading GRID Driver and Software License Packages

 Obtain the driver installation package required for an OS. For details, see Table 3-22.

For more information about the GRID driver, see **NVIDIA vGPU Software Documentation**.

MOTE

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.

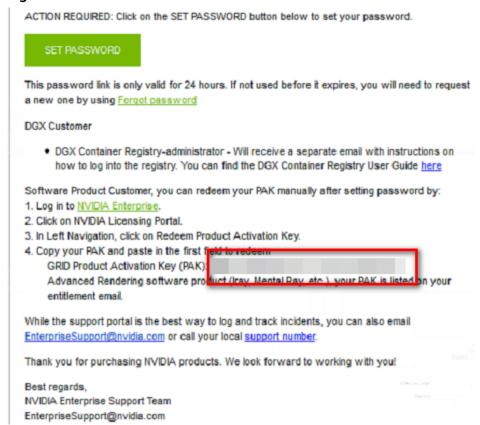
Table 3-22 GRID driver versions supported by GPU-accelerated ECSs

| ECS Type | GPU Attachme nt | OS | Driver Version | CPU Architect ure |
|------------------|------------------------|--|-----------------------------|-------------------------|
| G5.8xl arge.4 | GPU passthrou gh | CentOS 8.2 64bit CentOS 7.6 64bit CentOS 7.5 64bit Ubuntu 20.04 64bit Ubuntu 18.04 64bit Windows Server 2019 Standard 64bit Windows Server 2016 Standard 64bit Windows Server 2019 Datacenter 64bit Windows Server 2019 Datacenter 64bit | Select a version as needed. | x86_64 |
| G3 | GPU passthrou gh | Windows Server 2019 Standard 64bit Windows Server 2019 Standard 64bit | Select a version as needed. | x86_64 |
| P2s | GPU passthrou gh | Windows Server 2016 Standard 64bit | Select a version as needed. | x86_64 |
| P2 | GPU passthrou gh | Ubuntu 16.04 Server 64bit | Select a version as needed. | x86_64 |
| PI2 | GPU passthrou gh | CentOS 7.5 64bit Windows Server 2019 Standard 64bit Windows Server 2016 Standard 64bit | Select a version as needed. | x86_64 |

- 2. After the registration, log in at the **official NVIDIA website** and enter the account.
- 3. Check whether NVIDIA is used for the first time.

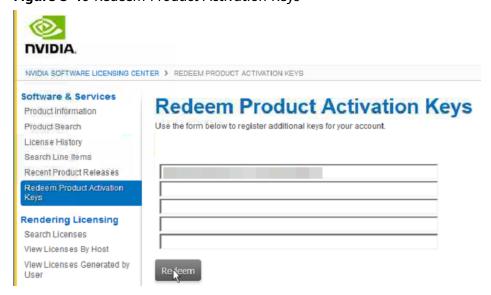
- a. If yes, go to step 4.
- b. If no, go to step 6.
- 4. Obtain the Product Activation Key (PAK) from the email indicating successful registration with NVIDIA.

Figure 3-39 PAK



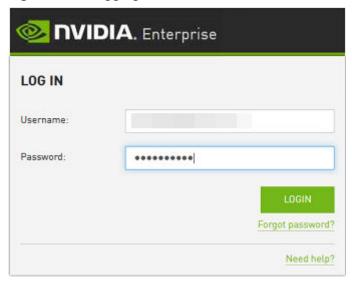
5. Enter the PAK obtained in step 4 on the **Redeem Product Activation Keys** page and click **Redeem**.

Figure 3-40 Redeem Product Activation Keys



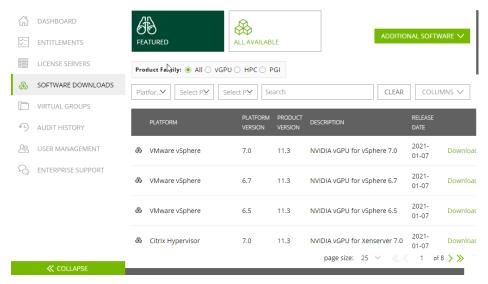
6. Specify **Username** and **Password** and click **LOGIN**.

Figure 3-41 Logging in to the official NVIDIA website



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

Figure 3-42 SOFTWARE DOWNLOADS page



- 8. Download the GRID driver of the required version. For details, see **Table 3-22**.
- 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.

NVIDIA, LICENSING | Software Downloads NVIDIA Application Hub | jian.zhou@xsuperzone.com (ORG_ADMIN) | Logout ENTITLEMENTS ALL AVAILABLE FEATURED 业 2020.11 64-bit License Manager for Windows Product Family:
All VGPU HPC 2020.11 64-bit License Manager for Linux ₹ 2020.05 64-bit License Manager for Windows SOFTWARE DOWNLOADS Platfor... Select P.M. Select P.M. & 2020.05 64-bit License Manager for Linux VIRTUAL GROUPS على 2019.11 64-bit License Manager for Windows ✓ 2019.11 64-bit License Manager for Linux AUDIT HISTORY **SQ. USER MANAGEMENT** ENTERPRISE SUPPORT VMware vSphere 6.7 Leading NVIDIA Virtual GPU Management Pack for vRealize Operations 2.0 11.3 NVIDIA vGPU for vSphere 6.5 01-07 & VMware vSphere 11.2 NIVIDIA vCDI I for Vancaniar 7.0 A Citriu Humanicar

Figure 3-43 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.

- 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 to which license server belongs and add inbound rules for TCP 7070 and TCP 8080.
- 2. Install the license server.
 - a. Run the following command to decompress the installation package. The **Installer.zip** in the command indicates the name of the software package obtained in **10**.

unzip Installer.zip

b. Run the following command to assign execution permissions to the installer:

chmod +x setup.bin

c. Run the installer as user root:

sudo ./setup.bin -i console

d. In the Introduction section, press **Enter** to continue.

```
Introduction
InstallAnywhere will guide you through the installation of License Server.

It is strongly recommended that you quit all programs before continuing with this installation.

Respond to each prompt to proceed to the next step in the installation. If you want to change something on a previous step, type 'back'.

You may cancel this installation at any time by typing 'quit'.

PRESS <ENTER> TO CONTINUE:
```

e. In the License Agreement section, press **Enter** to turn to last pages and accept the license agreement.

Enter Y and press Enter.

DO YOU ACCEPT THE TERMS OF THIS LICENSE AGREEMENT? (Y/N): Y

- f. In the Choose Install Folder section, press **Enter** to retain the default path for installing the License Server software.
- g. In the Choose Local Tomcat Server Path section, enter the Tomcat's local path in the "/var/lib/*Tomcat version*" format, for example, /var/lib/tomcat8.
- h. In the Choose Firewall Options section, confirm the port to be enabled in the firewall and press **Enter**.

```
Choose Firewall Options
-----
The license server listens on port 7070. This port must be opened in the firewall for other machines to obtain licenses from this server.

The license server's management interface listens on port 8080. Leave this port closed to prevent unauthorized access to the management interface.

->1- License server (port 7070)
2- Management interface (port 8080)

ENTER A COMMA-SEPARATED LIST OF NUMBERS REPRESENTING THE DESIRED CHOICES, OR PRESS <ENTER> TO ACCEPT THE DEFAULT:
```

i. In the Pre-Installation Summary section, confirm the information and press **Enter** to start the installation.

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:
```

- Obtain the license file.
 - a. Log in to the **NVIDIA website** on a new tab and select **LICENSE SERVERS**.

Figure 3-44 LICENSE SERVERS



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

Create License Server X Server Name Feature Licenses Name this license server Select a feature Added Features Description No features have been added yet MAC Address MAC Address (XX:XX:XX:XX:XX or XX-XX-XX-XX-XX) (i) Failover server configuration is optional.

If configuring, you must provide a name AND MAC address Failover License Server Failover License Server Failover MAC Address Failover MAC Address CANCEL RESET CREATE LICENSE SERVER

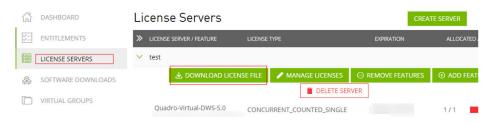
Figure 3-45 Create License Server

Table 3-23 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 3-46 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 3-47 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.

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 3-48 License server management console



3.9.3 Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS

Scenarios

Before using a GPU-accelerated ECS, make sure that the desired Tesla driver and CUDA toolkit have been installed on the ECS for computing acceleration.

- A computing-accelerated (P series) ECS created using a public image has had a Tesla driver of a specified version installed by default.
- After a GPU-accelerated ECS is created using a private image, it must have a Tesla driver installed. Otherwise, computing acceleration will not take effect.

This section describes how to install a Tesla driver and CUDA toolkit on a GPU-accelerated ECS.

Notes

- The ECS must have an EIP bound.
- Check whether the CUDA toolkit and Tesla driver have been installed on the ECS.

Ⅲ 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 <code>/var/log/nvidia-installer.log</code>. If the log shows that the failure was caused by a driver compilation error, for example, the <code>get_user_pages</code> 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.
- 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 qcc-c++

3. Download the NVIDIA driver package.

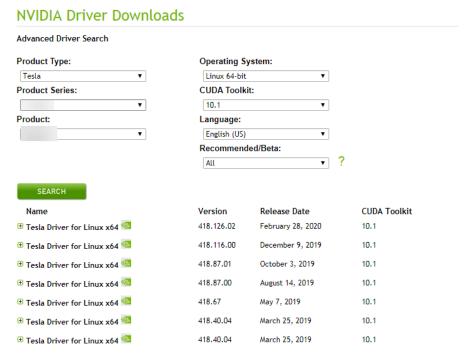
Select a driver version at **NVIDIA Driver Downloads** based on the ECS type. Click **SEARCH**.

Figure 3-49 Selecting a NVIDIA driver version



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

Figure 3-50 Selecting a driver version



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

Figure 3-51 Copying the download link



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 3-52 Obtaining the installation package

8. Run the following command to install the driver:

sh NVIDIA-Linux-x86_64-418.67.run

(Optional) If the following information is displayed after the command for installing the driver is executed, disable the Nouveau driver.

Figure 3-53 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

- Run the following command to back up and create an initramfs application:
 - Ubuntusudo update-initramfs -u
 - CentOS:

mv /boot/initramfs-\$(uname -r).img /boot/initramfs-\$(uname -r).img.bak

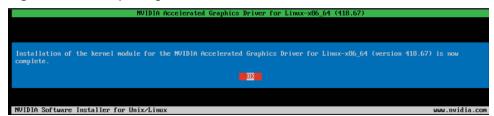
dracut -v /boot/initramfs-\$(uname -r).img \$(uname -r)

d. Restart the ECS:

reboot

10. Select **OK** for three consecutive times as prompted to complete the driver installation.

Figure 3-54 Completing the NVIDIA driver installation

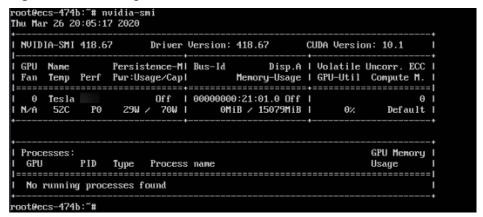


11. Run the following command to set systemd:

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 3-55 Viewing the NVIDIA driver version

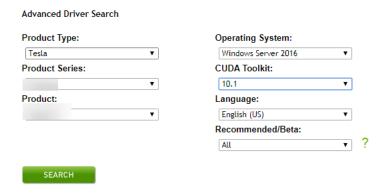


Installing a Tesla Driver on a Windows ECS

The following uses Windows Server 2016 Standard 64bit as an example to describe how to install a Tesla driver on a GPU-accelerated ECS.

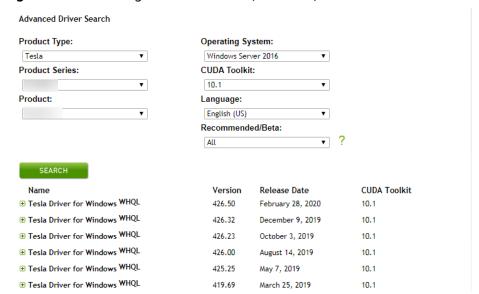
- 1. Log in to the ECS.
- Download the NVIDIA driver package.
 Select a driver version at NVIDIA Driver Downloads based on the ECS type.

Figure 3-56 Selecting a driver type (Windows)



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

Figure 3-57 Selecting a driver version (Windows)



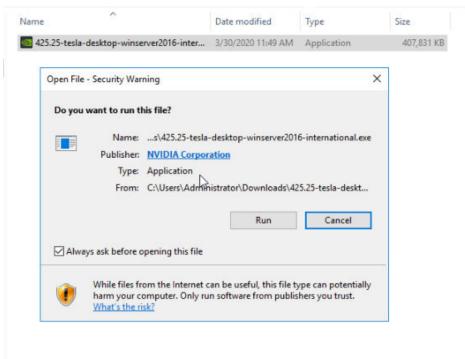
- 4. Click the driver to be downloaded. On the **TESLA DRIVER FOR WINDOWS** page that is displayed, click **DOWNLOAD**.
- 5. Click **AGREE & DOWNLOAD** to download the installation package.

Figure 3-58 Downloading the driver installation package



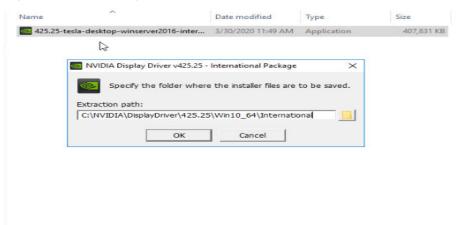
Double-click the driver and click Run.

Figure 3-59 Running the NVIDIA driver installation program



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

Figure 3-60 Selecting an installation path



8. Install the NVIDIA program as prompted.

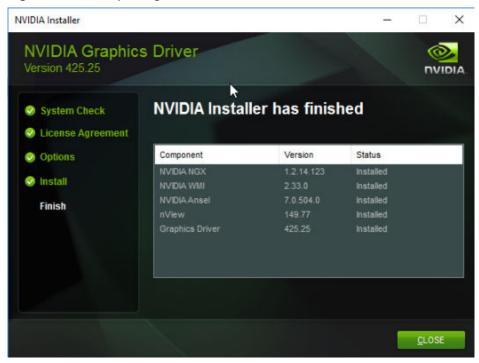


Figure 3-61 Completing the driver installation

- 9. Restart the ECS.
- 10. Check whether the NVIDIA driver has been installed.
 - a. Switch to **Device Manager** and click **Display adapters**.

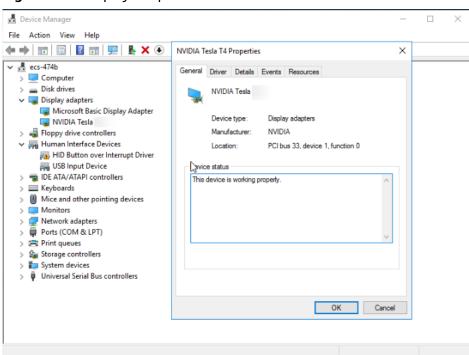


Figure 3-62 Display adapters

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 3-63 Viewing the NVIDIA driver version

```
:\Program Files\NVIDIA Corporation\NVSMI>nvidia-smi
                    2020
 NVIDIA-SMI 425.25
                          Driver Version: 425.25
                                                          CUDA Version: 10.1
                       TCC/WDDM
                                                  Disp.A
                                                          | Volatile Uncorr. ECC
| GPU-Util Compute M.
      Temp Perf Pwr:Usage/Cap
                                           Memory-Usage
                            TCC
7 N
                                   00000000:21:01.0 Off
                                                                 0%
               Р8
                     11W /
                                         0MiB / 15205MiB
                                                                          Default
                                                                       GPU Memory
Processes:
                  Type
                          Process name
  No running processes found
 Program Files\NVIDIA Corporation\NVSMI>
```

Installing the CUDA Toolkit on a Linux ECS

The following uses Ubuntu 16.04 64bit as an example to describe how to install the CUDA 10.1 toolkit on a GPU-accelerated ECS.

- 1. Log in to the ECS.
- 2. Update the system software based on the OS.
 - Ubuntu
 Update the software installation source: apt-get -y update
 Install necessary programs: apt-get install gcc g++ make
 - CentOS
 Update the software installation source: yum -y update exclude=kernel* --exclude=centos-release* --exclude=initscripts*
 Install the desired program: yum install -y kernel-devel-`uname -r` gcc gcc-c++
- 3. On the CUDA download page, set parameters according to the information shown in **Obtaining a Tesla Driver and CUDA Toolkit**.

Figure 3-64 Selecting a CUDA version



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

Figure 3-65 Copying the link for downloading CUDA



5. Run the following command on the ECS to download CUDA:

wget Copied link

For example, wget https://developer.nvidia.com/compute/cuda/10.1/Prod/local_installers/cuda_10.1.105_418.39_linux.run

Figure 3-66 Downloading CUDA

6. Install CUDA.

Follow the instructions provided on the official NVIDIA website.

Figure 3-67 Installing CUDA



7. Run the following command to install CUDA:

sh cuda_10.1.243_418.87.00_linux.run

8. Select **accept** on the installation page and press **Enter**.

Figure 3-68 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 3-69 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 3-70 Completing the installation

Run the following command to switch to /usr/local/cuda-10.1/samples/
 Utilities/deviceQuery:

cd /usr/local/cuda-10.1/samples/1_Utilities/deviceQuery

- 11. Run the **make** command to automatically compile the deviceQuery program.
- 12. Run the following command to check whether CUDA has been installed: ./deviceQuery

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

Figure 3-71 deviceQuery common output

13. Check the CUDA version.

/usr/local/cuda/bin/nvcc -V

Figure 3-72 Checking the CUDA version

```
[root@ecs-474b deviceQuery]# /usr/local/cuda/bin/nvcc -U
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
[root@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 3-73 Selecting a CUDA version



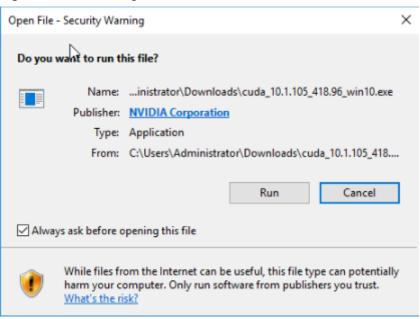
3. Find the link for downloading CUDA 10.1.

Figure 3-74 Finding the link for downloading CUDA



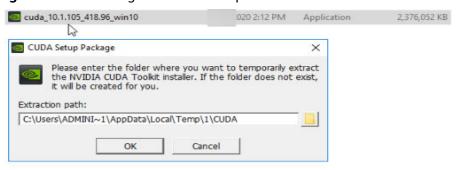
- 4. Click **Download** to download the CUDA toolkit.
- 5. Double-click the installation file and click **Run** to install the CUDA toolkit.

Figure 3-75 Installing CUDA



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

Figure 3-76 Selecting an installation path



7. Install the CUDA toolkit as prompted.

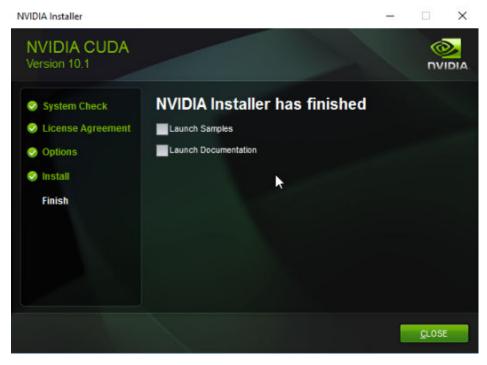


Figure 3-77 Completing the installation

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 3-78 Successful installation

```
C:\Users\Administrator>nvcc -V
nvcc: NVIDIA (R) Cuda compiler driver
Copyright (c) 2005-2019 NVIDIA Corporation
Built on Fri_Feb__8_19:08:26_Pacific_Standard_Time_2019
Cuda compilation tools, release 10.1, V10.1.105
C:\Users\Administrator>_
```

3.9.4 Obtaining a Tesla Driver and CUDA Toolkit

Scenarios

Before using a GPU-accelerated ECS, make sure that the desired Tesla driver and CUDA toolkit have been installed on the ECS. Otherwise, computing acceleration will not take effect. This section describes how to obtain a Tesla driver and CUDA toolkit. Select a driver version based on your ECS type.

For instructions about how to install the Tesla driver and CUDA toolkit, see **Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS**.

Downloading a Tesla Driver

Download a driver based on your ECS type.

Table 3-24 Mapping between Tesla drivers and ECS types

| ECS Type | Driver | Product Series | Product |
|----------|--------|----------------|---------|
| P2s | Tesla | V | V100 |
| P2 | Tesla | V | V100 |
| Pi2 | Tesla | Т | T4 |
| G5 | Tesla | V | V100 |

Downloading a CUDA Toolkit

Download the **CUDA software package** and select the corresponding CUDA Toolkit software package based on the instance type and driver version.

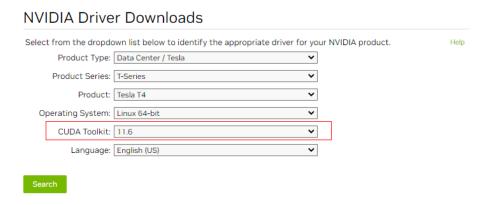
Ⅲ NOTE

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 3-79 Selecting the CUDA Toolkit version



2. Select a CUDA Toolkit 11.6 package to download.

Figure 3-80 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 11.4.1 (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

3.9.5 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.

© Control Panel

Adjust you computer's settings

View by: Category ▼

System and Security
Review your computer's status
View event logs

Network and Internet
View network status and tasks

Hardware
View devices and printers
Add a device

Programs

Control Panel

View by: Category ▼

User Accounts
Change account type

Appearance and Personalization
Change the theme

Clock, Language, and Region
Add a language
Change input methods
Set the time and date
Change date, time, or number formats

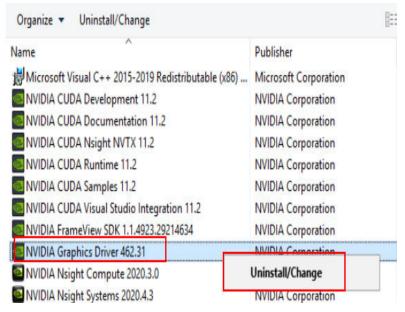
Ease of Access Let Windows suggest settings Optimize visual display

Figure 3-81 Uninstalling a program.

Programs
Uninstall a program
Turn Windows features on or off

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

Figure 3-82 Uninstalling a NVIDIA driver



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

Graphics Driver
Version 462.31

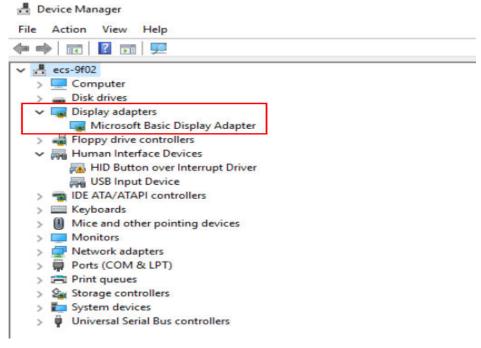
Do you really want to remove this software?

Figure 3-83 Confirming the uninstallation

- 6. After the uninstallation is complete, click **RESTART LATER**.
- 7. Check whether the NVIDIA driver has been uninstalled.
 - In Control Panel, click **Device Manager**.

 If no NVIDIA graphics cards are not displayed under **Display adapters**, the driver is uninstalled successfully.

Figure 3-84 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 3-85 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 NVDIA 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 3-86 Querying the nvidia-uninstall path



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

Figure 3-87 NVIDIA driver uninstallation (1)



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

Figure 3-88 NVIDIA driver uninstallation (2)

```
ERROR: Failed to run '/wsr/bin/nvidia-xconfig --restore-original-backup':

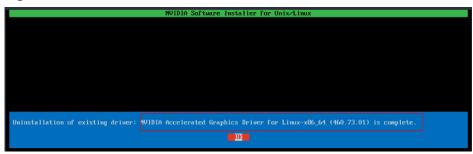
WARNING: Unable to locate/open X configuration file.

WARNING: Unable to parse X.Org version string.

ERROR: Unable to restore from original backup file '/etc/X11/xorg.conf.nvidia-xconfig-original' (Mo such file or directory)
```

e. After the driver is uninstalled, press Enter.

Figure 3-89 NVIDIA driver uninstallation (3)



- 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.

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 3-90 Uninstalling a CUDA driver

```
CUDA Uninstaller

[X1 CUDA_Demo_Suite_11.2

[X1 CUDA_Toolkit_11.2

[X1 CUDA_Documentation_11.2

[X1 CUDA_Samples_11.2

Done

Up/Down: Move | 'Enter': Select
```

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

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.

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.

Procedure

- 1. Log in to the management console.
- 2. Under **Computing**, click **Elastic Cloud Server**.
- 3. In the ECS list, choose **More** > **Manage Image/Backup** > **Create Image** in the **Operation** column.
- 4. Configure the following information:

Table 4-2 and Table 4-3 list the parameters in the Image Type and Source and Image Information areas, respectively.

Table 4-2 Image type and source

| Parameter | Description |
|------------|--|
| Туре | Select Create Image. |
| 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. |
| | 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 |
|-------------|---|
| 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. |

5. Click **Next** and submit the request.

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 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

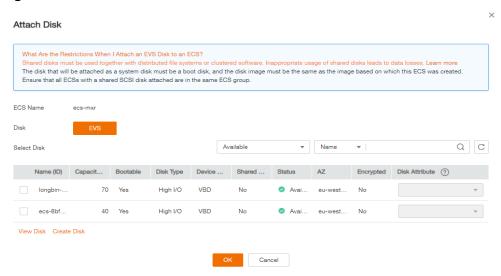
- 1. Log in to the management console.
- 2. Under Computing, choose Elastic Cloud Server.
- 3. In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID for search.
- 4. Click the name of the target ECS.

The page providing details about the ECS is displayed.

5. Click the **Disks** tab. Then, click **Attach Disk**.

The Attach Disk dialog box is displayed.

Figure 5-1 Attach Disk



- 6. Select the target disk and specify the disk as the system disk or data disk
 - For KVM ECSs, you can specify a disk as a system disk or data disk but cannot specify a device name for the disk.
 - For Xen ECSs, you can specify the device name of a disk, such as /dev/vdb.

- If no EVS disks are available, click Create Disk in the lower part of the list.
- For details about restrictions on attaching a disk, see What Are the Requirements for Attaching an EVS Disk to an ECS?

7. Click OK.

After the disk is attached, you can view the information about it on the **Disks** tab.

Follow-up Procedure

If the attached disk is newly created, the disk can be used only after it is initialized.

For details about how to initialize a data disk, see Scenarios and Disk Partitions.

5.3 Detaching an EVS Disk from a Running ECS

Scenarios

You can detach EVS disks from an ECS.

- System disks (mounted to /dev/sda or /dev/vda) can only be detached offline. They must be stopped before being detached.
- Data disks (mounted to points other than dev/sda) can be detached online if the attached ECS is running certain OSs. You can detach these data disks without stopping the ECS.

This section describes how to detach a disk from a running ECS.

Constraints

- The EVS disk to be detached must be mounted to a point other than /dev/sda or /dev/vda.
 - EVS disks mounted to /dev/sda or /dev/vda are system disks and cannot be detached from running ECSs.
- Before detaching an EVS disk from a running Windows ECS, make sure that UVP VMTools have been installed on the ECS and that the tools are running properly.
- Before detaching an EVS disk from a running Windows ECS, ensure that no programs are reading data from or writing data to the disk. Otherwise, data will be lost.
- SCSI EVS disks cannot be detached from running Windows ECSs.
- Before detaching an EVS disk from a running Linux ECS, you must log in to the ECS and run the **umount** command to cancel the association between the disk and the file system. In addition, ensure that no programs are reading data from or writing data to the disk. Otherwise, detaching the disk will fail.

Notes

 On a Windows ECS, if the disk is in non-offline state, the system forcibly detaches the EVS disk. If this occurs, the system may generate a xenvbd alarm. You can ignore this alarm.

□ NOTE

To view the status of an EVS disk, perform the following operations:

- 1. Click **Start** in the task bar. In the displayed **Start** menu, right-click **Computer** and choose **Manage** from the shortcut menu.
 - The Server Manager page is displayed.
- 2. In the navigation pane on the left, choose **Storage** > **Disk Management**. The EVS disk list is displayed in the right pane.
- 3. View the status of each EVS disk.
- Do not detach an EVS disk from an ECS that is being started, stopped, or restarted.
- Do not detach an EVS disk from a running ECS whose OS does not support this feature. OSs supporting EVS disk detachment from a running ECS are listed in OSs Supporting EVS Disk Detachment from a Running ECS.
- For a running Linux ECS, the drive letter may be changed after an EVS disk is detached from it and then attached to it again. This is a normal case due to the drive letter allocation mechanism of the Linux system.
- For a running Linux ECS, the drive letter may be changed after an EVS disk is detached from it and the ECS is restarted. This is a normal case due to the drive letter allocation mechanism of the Linux system.

OSs Supporting EVS Disk Detachment from a Running ECS

OSs supporting EVS disk detachment from a running ECS include two parts:

- For the first part, see 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 | |

| os | Version | |
|---|---|--|
| OpenSUSE | 42.2 64bit | |
| | 42.1 64bit | |
| Oracle Linux Server | 7.3 64bit | |
| release | 7.2 64bit | |
| | 6.8 64bit | |
| | 6.7 64bit | |
| Ubuntu Server | 16.04 64bit | |
| | 14.04 64bit | |
| | 14.04.4 64bit | |
| Windows (SCSI EVS disks | Windows Server 2008 R2 Enterprise 64bit | |
| cannot be detached from a running ECS.) | 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

- On the Elastic Cloud Server page, click the name of the ECS from which the EVS disk is to be detached. The page providing details about the ECS is displayed.
- 2. Click the **Disks** tab. Locate the row containing the EVS disk to be detached and click **Detach**.

5.4 Expanding the Capacity of an EVS Disk

Scenarios

You can expand the disk capacity if the disk space is insufficient. The capacities of both system disks and data disks can be expanded.

Procedure

The capacity of an EVS disk can be expanded in either of the following ways:

- Create an EVS disk and attach it to an ECS.
- Expand the capacity of an existing EVS disk. The capacities of both system disks and data disks can be expanded.

You can expand the disk capacities when the EVS disks are in the **In-use** or **Available** state.

- Expanding an In-use EVS disk means expanding the capacity of an EVS disk that has been attached to an ECS. Only certain OSs support the expansion of In-use EVS disks. For details, see Expanding 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**.

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**.

5.5 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.
- 8. Migrate the data from the EVS disk to the local disks of the new diskintensive ECS.

6 Backup Using CBR

6.1 Overview

What Is CBR?

Cloud Backup and Recovery (CBR) enables you to back up cloud servers and disks with ease. In case of a virus attack, accidental deletion, or software or hardware fault, you can restore data to any point in the past when the data was backed up.

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

What Are the Differences Between Backup, 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.

| Back up Type | Backup Object | Application Scenario | Differences and Advantages | Back up Meth od | Restor ation Meth od |
|---------------------------------------|---|---|---|----------------------------------|--|
| Clou d serve r back up | All disks (system and data disks) on an ECS | 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 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. | All disks on an ECS are backed up at the same time, ensuring data consistency. In addition, you can configure backup policies for automatic backup. | Creating a Cloud Serve r Back up | • Restoring Data Using a Cloud Server Backup |

| Back up Type | Backup Object | Application Scenario | Differences and Advantages | Back up Meth od | Restor ation Meth od |
|---------------------------------|--|--|--|---|---|
| Clou d disk back up | One or more specified disks (system or data disks) | 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 compromisin g data security. | Creat ing a Cloud Disk Back up | Res tori ng Dat a Usi ng a Clo ud Dis k Bac ku p Usi ng a Bac ku p to Cre ate a Dis k |

| Back up Type | Backup Object | Application Scenario | Differences and Advantages | Back up Meth od | Restor ation Meth od |
|---------------------------------|------------------|--|---|--------------------------------|--|
| Syste m disk imag e | System disk | 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 Syste m Disk Imag e | Ch an gin g the OS of a Fau lty ECS Usi ng a Sys te m Dis k Im age The OS of a Fau lty ECS of a Fau |

| Back up Type | Backup Object | Application Scenario | Differences and Advantages | Back up Meth od | Restor ation Meth od |
|---------------------------|---|---|--|---|--|
| Data disk imag e | 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. | Creat ing a Data Disk Imag e | Creati ng a Data Disk from a Data Disk Image |
| Full- ECS imag e | All disks (system and data disks) on an ECS | 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. | Creat ing a Full- ECS Imag e | Creati ng 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 6-1 One-off backup and periodic backup

| Item | One-Off Backup | Periodic Backup |
|------------------------------|--|---|
| Backup policy | Not required | Required |
| Number of backup tasks | One manual backup task | Periodic tasks driven by a backup policy |
| Backup name | User-defined backup name, which is 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 |
| Applicatio n scenario | Executed before patching or upgrading the OS or upgrading an application on a resource. A one-off backup can be used to restore the resource to the original state if the patching or upgrading fails. | Executed for routine maintenance of a resource. The latest backup can be used for restoration if an unexpected failure or data loss occurs. |

6.2 Backing Up an ECS Data

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.

□ NOTE

For Windows ECSs, you can install the Windows Server Backup tool provided by the Windows OS and use CBR to back up full ECS data.

ECS Backup Procedure

- 1. Log in to the management console.
- 2. Under Computing, choose Elastic Cloud Server.
- In the ECS list, locate the target ECS and choose More > Manage Image/ Backup > Create Server Backup in the Operation column.
 - 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.

4. Click **OK**. The system automatically creates a backup for the ECS.

On the **Backups** tab page, if the status of the backup is **Available**, the backup task is successful.

The ECS can be restarted if the backup progress of an ECS exceeds 10%. However, to ensure data integrity, restart it after the backup is complete.

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. Under **Computing**, choose **Elastic Cloud Server**.
- In the ECS list, locate the target ECS and choose More > Manage Image/ Backup > Create Disk Backup in the Operation column.
 - 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.
- 4. Click **OK**. The system automatically creates a backup for the disk.

On the **Backups** tab page, 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**.

7 NICs

7.1 Overview

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.

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.

Helpful Links

- Will NICs Added to an ECS Start Automatically?
- Why Is the NIC Not Working?

7.2 Attaching a Network Interface

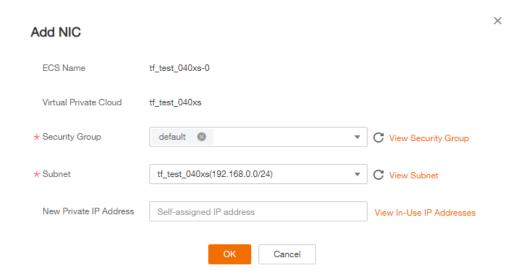
Scenarios

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

Procedure

- 1. Log in to the management console.
- 2. Under Computing, click Elastic Cloud Server.
- Click the name of the target ECS.
 The page providing details about the ECS is displayed.
- 4. On the NICs tab, click Add NIC.
- 5. Set the subnet and security group for the network interface to be added.

Figure 7-1 Configuring the subnet and security group



- Security Group: You can select multiple security groups. In such a case, the rules of these security groups are effectively aggregated to create one set of rules to apply to the ECS.
- Subnet: specifies the subnet to which the network interface to be added belongs.
- New Private IP Address: If you want to add a network interface with a specified IP address, enter an IP address into the New Private IP Address field. I
- 6. Click OK.

Follow-up Procedure

Some OSs cannot identify newly added network interfaces. In this case, you must manually activate the network interfaces. Ubuntu is used as an example in the following network interface activation procedure. Required operations may vary among systems. For additional information, see the documentation for your OS.

- 1. Locate the row containing the target ECS and click **Remote Login** in the **Operation** column.
 - Log in to the ECS.
- 2. Run the following command to view the network interface name: **ifconfig -a**

In this example, the network interface name is eth2.

3. Run the following command to switch to the target directory:

cd /etc/network

4. Run the following command to open the **interfaces** file:

vi interfaces

5. Add the following information to the **interfaces** file:

auto eth2

iface eth2 inet dhcp

6. Run the following command to save and exit the **interfaces** file:

:wq

- 7. Run either the **ifup eth2** command or the **/etc/init.d/networking restart** command to make the newly added network interface take effect.
 - *X* in the preceding command indicates the network interface name and SN, for example, **ifup eth2**.
- 8. Run the following command to check whether the network interface name obtained in step 2 is displayed in the command output:

ifconfig

For example, check whether **eth2** is displayed in the command output.

- If yes, the newly added network interface has been activated, and no further action is required.
- If no, the newly added network interface failed to be activated. Go to step 9.
- 9. Log in to the management console. Locate the row containing the target ECS, click **More** in the **Operation** column, and select **Restart**.
- 10. Run the following command to check whether the network interface name obtained in step 2 is displayed in the command output:
 - If yes, no further action is required.
 - If no, contact customer service.

7.3 Deleting 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 delete 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. Click the **NICs** tab. Then, click **Delete** in the row of the target network interface.

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.

Certain ECSs do not support network interface deletion when they are running. For details, see the GUI display. To delete a network interface from such an ECS, stop the ECS first.

7.4 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.

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

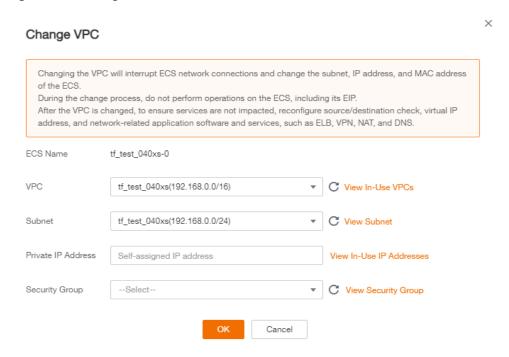
The target VPC, subnet, private IP address, and security group are available.

Procedure

- 1. Log in to the management console.
- 2. Under Computing, click Elastic Cloud Server.
- 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 7-2 Change VPC



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 apply on 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**.

7.5 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.
- Click OK.

Configuring a Virtual IP Address for an ECS

Manually configure the virtual IP address bound to an ECS.

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

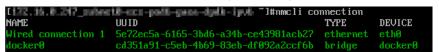
- Linux: CentOS 7.2 64bit and Ubuntu 22.04 server 64bit
- Windows: Windows Server

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:



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:

I#mmcli commection up "Wired commection 1" Commection successfully activated (D-Bus active path: /org/freedesktop/NetworkManager/ActiveCommection/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.

□ 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-ubantu:~# 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.
- 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-ubantu:/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
    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.)

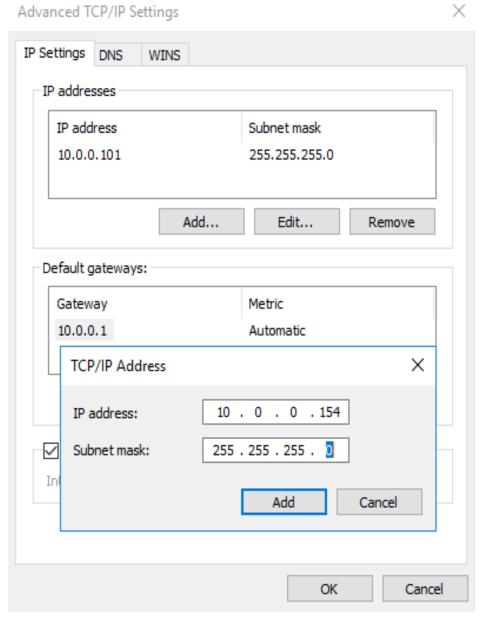
- 1. In **Control Panel**, click **Network and Sharing Center**, and click the corresponding local connection.
- 2. On the displayed page, click **Properties**.
- 3. On the Network tab page, select Internet Protocol Version 4 (TCP/IPv4).
- 4. Click Properties.
- 5. Select **Use the following IP address**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

Figure 7-3 Configuring private IP address

Internet Protocol Version 4 (TCP/IPv4) Properties Х General You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings. Obtain an IP address automatically Use the following IP address: IP address: 10 . 0 . 0 . 101 Subnet mask: 255 . 255 . 255 . 0 Default gateway: 10 . 0 . 0 . 1 Obtain DNS server address automatically Use the following DNS server addresses: Preferred DNS server: 100 . 125 . 1 . 250 Alternate DNS server: 114 . 114 . 114 . 114 Validate settings upon exit Advanced... OK Cancel

- 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 7-4 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.

7.6 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 multiqueue enables multiple CPUs to process ECS NIC interruptions, thereby improving packets per second (PPS) and I/O performance.

The ECS described in this section is assumed to comply with the requirements on specifications and virtualization type.

- If the ECS was created using a public image listed in Support of NIC Multi-Queue, NIC multi-queue has been enabled on the ECS by default. Therefore, you do not need to perform the operations described in this section.
- If the ECS was created using a private image and the OS of the external image file is listed in Support of NIC Multi-Queue, perform the following operations to enable NIC multi-queue:
 - a. Importing the External Image File to the IMS Console
 - b. Setting NIC Multi-Queue for the Image
 - c. Creating an ECS Using a Private Image
 - d. Enabling NIC Multi-Queue

Support of NIC Multi-Queue

NIC multi-queue can be enabled on an ECS only when the ECS specifications, virtualization type, and image OS meet the requirements described in this section.

 For details about the ECS specifications that support NIC multi-queue, see ECS Specifications and Types.

If the number of NIC queues is greater than 1, NIC multi-queue is supported.

- The virtualization type must be KVM.
- The Linux public images listed in **Table 7-2** support NIC multi-queue.

∩ NOTE

- Windows public images have not supported NIC multi-queue. If you enable NIC multi-queue in a Windows public image, starting an ECS created using such an image may be slow.
- It is a good practice to upgrade the kernel version of the Linux ECS to 2.6.35 or later. Otherwise, NIC multi-queue is not supported.

Run the **uname -r** command to obtain the kernel version. If the kernel version is earlier than 2.6.35, contact customer service to upgrade the kernel.

Table 7-1 Support of NIC multi-queue for Linux ECSs

| Image | Support of NIC Multi- Queue | NIC Multi- Queue Enabled by Default |
|---|-----------------------------------|---|
| Ubuntu 14.04/16.04/18.04/20.04 server 64bit | Yes | Yes |
| OpenSUSE 42.2/15.* 64bit | Yes | Yes |
| SUSE Enterprise 12 SP1/SP2 64bit | Yes | Yes |
| CentOS 6.8/6.9/7.*/8.* 64bit | Yes | Yes |
| Debian 8.0.0/8.8.0/8.9.0/9.0.0/10.0.0/10.2.0 64bit | Yes | Yes |
| Fedora 24/25/30 64bit | Yes | Yes |
| EulerOS 2.2/2.3/2.5 64bit | Yes | Yes |

Table 7-2 Support of NIC multi-queue for KVM ECSs

| OS | Image | Status |
|---------|---|--------------------------------|
| Windows | Windows Server 2008 Web R2 64-bit | Supported using private images |
| | Windows Server 2008 R2 Standard/ DataCenter/Enterprise 64bit | Supported using private images |
| | Windows Server 2012 R2 Standard/DataCenter 64bit | Supported using private images |
| | Windows Server 2016 Standard/DataCenter 64bit | Supported using private images |
| Linux | Ubuntu 14.04/16.04 server 64bit | Supported |
| | OpenSUSE 42.2 64bit | Supported |
| | SUSE Enterprise 12 SP1/SP2 64bit | Supported |
| | CentOS 6.8/6.9/7.0/7.1/7.2/7.3/7.4/7.5/7.6 64bit | Supported |
| | Debian 8.0.0/8.8.0/8.9.0/9.0.0 64bit | Supported |
| | Fedora 24/25 64bit | Supported |
| | EulerOS 2.2 64bit | Supported |

Importing the External Image File to the IMS Console

For details, see "Registering an Image File as a Private Image" in *Image Management Service User Guide*.

Setting NIC Multi-Queue for the Image

Windows OSs have not commercially supported NIC multi-queue. If you enable NIC multi-queue in a Windows image, starting an ECS created using such an image may be slow.

Use one of the following methods to set the NIC multi-queue attribute:

Method 1:

- 1. Log in to the management console.
- 2. Under Computing, click Image Management Service.
- 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:

- 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.

- 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.

Add Content-Type to the request header.

The value of **Content-Type** is **application/openstack-images-v2.1-json-patch**.

The request URI is in the following format:

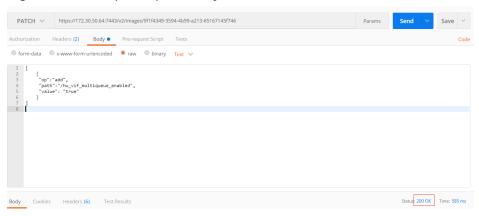
PATCH /v2/images/{image_id}

The request body is as follows:

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

Figure 7-5 shows an example request body for modifying the NIC multiqueue attribute.

Figure 7-5 Example request body



Creating an ECS Using a Private Image

Create an ECS using a registered private image. For details, see **Creating an ECS**. Note the following when setting the parameters:

- **Region**: Select the region where the private image is located.
- **Image**: Select **Private image** and then the desired image from the drop-down list.

Enabling NIC Multi-Queue

KVM Windows ECSs use private images to support NIC multi-queue. For details, see "How Do I Set NIC Multi-queue Feature of an Image?" in *Image Management Service User Guide*.

This section uses a Linux ECS running CentOS 7.4 as an example to describe how to enable NIC multi-queue.

Step 1 Enable NIC multi-queue.

- 1. Log in to the ECS.
- 2. Run the following command to obtain the number of queues supported by the NIC and the number of queues with NIC multi-queue enabled:

ethtool -l N/C

3. Run the following command to configure the number of queues used by the NIC:

ethtool -L N/C combined Number of queues

An example is provided as follows:

```
[root@localhost ~]# ethtool -l eth0 #View the number of queues used by NIC eth0.
Channel parameters for eth0:
Pre-set maximums:
RX:
            0
TX:
            0
Other:
Combined: 4 #Indicates that a maximum of four queues can be enabled for the NIC.
Current hardware settings:
RX:
            0
TX:
                 n
Other:
Combined: 1 #Indicates that one queue has been enabled.
```

[root@localhost ~]# ethtool -L eth0 combined 4 #Enable four queues on NIC eth0.

Step 2 (Optional) Enable irqbalance so that the system automatically allocates NIC interrupts on multiple vCPUs.

1. Run the following command to enable irgbalance:

service irqbalance start

2. Run the following command to view the irqbalance status:

service irqbalance status

If the **Active** value in the command output contains **active** (running), irqbalance has been enabled.

Figure 7-6 Enabled irqbalance

Step 3 (Optional) Enable interrupt binding.

Enabling irqbalance allows the system to automatically allocate NIC interrupts, improving network performance. If the improved network performance still fails to meet your requirements, manually configure interrupt affinity on the ECS.

To do so, perform the following operations:

Configure the following script so that one ECS vCPU serves the interrupt requests initialized by one queue. One queue corresponds to one interrupt, and one interrupt binds to one vCPU.

```
#!/bin/bash
service irqbalance stop
eth_dirs=$(ls -d /sys/class/net/eth*)
if [ $? -ne 0 ];then
  echo "Failed to find eth* , sleep 30" >> $ecs_network_log
  sleep 30
  eth_dirs=$(ls -d /sys/class/net/eth*)
for eth in $eth_dirs
do
  cur_eth=$(basename $eth)
  cpu_count=`cat /proc/cpuinfo| grep "processor"| wc -l`
  virtio_name=$(ls -l /sys/class/net/"$cur_eth"/device/driver/ | grep pci |awk {'print $9'})
  affinity_cpu=0
  virtio_input="$virtio_name""-input"
  irqs_in=$(grep "$virtio_input" /proc/interrupts | awk -F ":" '{print $1}')
  for irq in ${irqs_in[*]}
     echo $((affinity_cpu%cpu_count)) > /proc/irq/"$irq"/smp_affinity_list
     affinity_cpu=$[affinity_cpu+2]
  done
  affinity cpu=1
  virtio_output="$virtio_name""-output"
  irqs_out=$(grep "$virtio_output" /proc/interrupts | awk -F ":" '{print $1}')
  for irq in ${irqs_out[*]}
```

```
echo $((affinity_cpu%cpu_count)) > /proc/irq/"$irq"/smp_affinity_list
affinity_cpu=$[affinity_cpu+2]
done
done
```

Step 4 (Optional) Enable XPS and RPS.

XPS allows the system with NIC multi-queue enabled to select a queue by vCPU when sending a data packet.

```
#!/bin/bash
# enable XPS feature
cpu_count=$(grep -c processor /proc/cpuinfo)
dec2hex(){
 echo $(printf "%x" $1)
eth_dirs=$(ls -d /sys/class/net/eth*)
if [ $? -ne 0 ];then
  echo "Failed to find eth* , sleep 30" >> $ecs_network_log
  eth_dirs=$(ls -d /sys/class/net/eth*)
for eth in $eth_dirs
do
  cpu_id=1
  cur_eth=$(basename $eth)
  cur_q_num=$(ethtool -l $cur_eth | grep -iA5 current | grep -i combined | awk {'print $2'})
  for((i=0;i < cur\_q\_num;i++))
     if [ $i -eq $cpu_count ];then
        cpu_id=1
     xps_file="/sys/class/net/${cur_eth}/queues/tx-$i/xps_cpus"
     rps_file="/sys/class/net/${cur_eth}/queues/rx-$i/rps_cpus"
     cpuset=$(dec2hex "$cpu_id")
     echo $cpuset > $xps_file
     echo $cpuset > $rps_file
     let cpu_id=cpu_id*2
  done
done
```

----End

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 N/C
```

Example:

```
[root@localhost ~]# ethtool -l eth0 #View the number of queues used by NIC eth0.
Channel parameters for eth0:
Pre-set maximums:
RX:
            0
            0
TX:
Other:
Combined: 4 #Indicates that a maximum of four queues can be enabled for the NIC.
Current hardware settings:
RX:
            0
TX:
            0
Other:
Combined: 1 #Indicates that four queues have been enabled.
```

7.7 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.
 For details about how to enable IPv6 on a subnet, see Enabling IPv6 on the Subnet Where the ECS Works.
- 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

• Linux: Dynamic assignment of IPv6 addresses can be enabled automatically (recommended) or manually.

If a private image created from a CentOS 6.x or Debian ECS with automatic IPv6 address assignment enabled is used to create an ECS in an environment that does not support IPv6, the ECS may start slow because of IPv6 address assignment timeout. You can set the timeout duration for assigning IPv6 addresses by referring to Setting the Timeout Duration for IPv6 Address Assignment.

Table 7-3 Enabling dynamic assignment of IPv6 addresses for different OSs

| OS | Automatically/ Manually Enabling | Reference |
|-------|-------------------------------------|--|
| Linux | Automatically (recommended) | Linux (Automatically Enabling Dynamic Assignment of IPv6 Addresses) |

| OS | Automatically/ Manually Enabling | Reference |
|-------|-------------------------------------|--|
| Linux | Manually | Linux (Manually Enabling Dynamic Assignment of IPv6 Addresses) |

Enabling IPv6 on the Subnet Where the ECS Works

- 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.

Linux (Automatically Enabling Dynamic Assignment of IPv6 Addresses)

The **ipv6-setup-**xxx tool can be used to enable Linux OSs to automatically acquire IPv6 addresses. xxx indicates a tool, which can be rhel or debian.

You can also enable dynamic IPv6 address assignment by following the instructions in Linux (Manually Enabling Dynamic Assignment of IPv6 Addresses).

⚠ CAUTION

- When you run **ipv6-setup-***xxx*, the network service will be automatically restarted. As a result, the network is temporarily disconnected.
- If a private image created from a CentOS 6.x or Debian ECS with automatic IPv6 address assignment enabled is used to create an ECS in an environment that does not support IPv6, the ECS may start slow because of IPv6 address assignment timeout. Set the timeout duration for assigning IPv6 addresses to 30s by referring to Setting the Timeout Duration for IPv6 Address Assignment and try to create a new private image again.

Step 1 Run the following command to check whether IPv6 is enabled for the ECS:

ip addr

• If only an IPv4 address is displayed, IPv6 is disabled. Enable it by referring to **Setting the Timeout Duration for IPv6 Address Assignment**.

Figure 7-7 IPv6 disabled

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000 link/ether fa:16:3e: brd ff:ff:ff:ff:ff: ff inet brd scope global noprefixroute dynamic eth0 valid_lft 1193sec preferred_lft 1193sec
```

• If a link-local address (starting with fe80) is displayed, IPv6 is enabled but dynamic assignment of IPv6 addresses is not enabled.

Figure 7-8 IPv6 enabled

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000 link/ether fa:16:3e: brd ff:ff:ff:ff:ff:ff:inet brd scope global noprefixroute dynamic eth0 valid_lft 76391sec preferred_lft 76391sec inet6 fe80::f816: /64 scope link valid_lft forever preferred_lft forever
```

• If the following address is displayed, IPv6 is enabled and an IPv6 address has been assigned:

Figure 7-9 IPv6 enabled and an IPv6 address assigned

□ NOTE

IPv6 is enabled for Linux public images by default, as shown in Figure 7-8.

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
```

- 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)



If a private image created from a CentOS 6.x or Debian ECS with automatic IPv6 address assignment enabled is used to create an ECS in an environment that does not support IPv6, the ECS may start slow because of IPv6 address assignment timeout. Set the timeout duration for assigning IPv6 addresses to 30s by referring to **Setting the Timeout Duration for IPv6 Address Assignment** and try to create a new private image again.

Step 1 Run the following command to check whether IPv6 is enabled for the ECS:

ip addr

 If only an IPv4 address is displayed, IPv6 is disabled. Enable it by referring to Step 2.

Figure 7-10 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
inet brd scope global noprefixroute dynamic eth0
valid_lft 1193sec preferred_lft 1193sec
```

• If a link-local address (starting with fe80) is displayed, IPv6 is enabled but dynamic assignment of IPv6 addresses is not enabled.

Figure 7-11 IPv6 enabled

• If the following address is displayed, IPv6 is enabled and an IPv6 address has been assigned:

Figure 7-12 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: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
```

Ⅲ NOTE

IPv6 is enabled for Linux public images by default, as shown in Figure 7-11.

Step 2 Enable IPv6 for the ECS.

Run the following command to check whether IPv6 is enabled for the kernel:

sysctl -a | grep ipv6

- If a command output is displayed, IPv6 is enabled.
- If no information is displayed, IPv6 is disabled. Go to Step 2.2 to load the IPv6 module.
- 2. Run the following command to load the IPv6 module:

modprobe ipv6

3. Add the following content to the /etc/sysctl.conf file:

net.ipv6.conf.all.disable_ipv6=0

4. Save the configuration and exit. Then, run the following command to load the configuration:

sysctl-p

Step 3 Enable dynamic IPv6 address assignment for the ECS.

- Ubuntu 18.04/20.04
 - a. Run the following command to access /etc/netplan/:

cd /etc/netplan

b. Run the following command to list the configuration file:

ls

Figure 7-13 Configuration file name

c. Run the following command to edit the configuration file:

vi 01-network-manager-all.yaml

d. Append the following content to the configuration file (pay attention to the yaml syntax and text indentation):

```
ethernets:
eth0:
dhcp6: true
```

Figure 7-14 Edited configuration file

```
# Let NetworkManager manage all devices on thin system
network:
    version: 2
    renderer: NetworkManager
    ethernets:
       eth0:
            dhcp6: true
```

Save the changes and exit.

e. Run the following command to make the changes take effect: sudo netplan apply

- Ubuntu 22.04
 - a. Run the following command to access /etc/netplan/:

cd /etc/netplan

b. Run the following command to list the configuration file:ls

Figure 7-15 Configuration file name

```
root@ecs-485b:/etc/netplan# ls
01-netcfg.yaml
```

c. Run the following command to edit the configuration file:

vi 01-netcfg.yaml

d. Append the following content to the configuration file **01-netcfg.yaml** (pay attention to the yaml syntax and text indentation):

```
ethernets:
eth0:
dhcp6: true
```

Figure 7-16 Edited configuration file

```
network:

version: 2

renderer: NetworkManager
ethernets:

eth0:

dhcp4: true
dhcp6: true
eth1:
dhcp4: true
eth2:
dhcp4: true
eth3:
dhcp4: true
eth4:
dhcp4: true
```

Save the changes and exit.

e. Run the following command to make the changes take effect:

sudo netplan apply

f. Run the following command to edit /etc/NetworkManager/ NetworkManager.conf:

vi /etc/NetworkManager/NetworkManager.conf

g. Append the following content to the configuration file NetworkManager.conf (pay attention to the file format and indentation):

```
[main] plugins=ifupdown,keyfile
```

dhcp=dhclient [ifupdown] managed=true [device] wifi.scan-rand-mac-address=no

Figure 7-17 Modification result

```
[main]
plugins=ifupdown,keyfile
dhcp=dhclient

[ifupdown]
managed=true

[device]
wifi.scan-rand-mac-address=no
```

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
 - 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

- Edit the /etc/sysconfig/network file to add or modify the following line: NETWORKING IPV6=yes
- c. For an ECS running CentOS 6, you need to edit the configuration files of its extension NICs. For example, if the extension NIC is eth1, you need to edit /etc/sysconfig/network-scripts/ifcfg-eth1.

Add the following configuration items to the file:

IPV6INIT=yes DHCPV6C=yes

In CentOS 6.3, dhcpv6-client requests are filtered by **ip6tables** by default. So, you also need to add a rule allowing the dhcpv6-client request to the **ip6tables** file.

i. Run the following command to add the rule to **ip6tables**:

ip6tables -A INPUT -m state --state NEW -m udp -p udp --dport 546 -d fe80::/64 -j ACCEPT

ii. Run the following command to save the rule in **ip6tables**:

service ip6tables save

Figure 7-18 Example command

```
[root@ecs-cd02 log]# ip6tables -A IMPUT -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/ip6table[ 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 guery the NIC information:

nmcli con

Figure 7-19 Querying NIC information

| [root@ecs-166 | b ~]# | nmcli con | | |
|---------------|-------|--------------------------------------|----------|--------|
| NAME | | UUID | TYPE | DEVICE |
| System eth0 | | 5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03 | ethernet | eth0 |
| Wired connect | ion 1 | 9c92fad9-6ecb-3e6c-eb4d-8a47c6f50c04 | ethernet | eth1 |
| Wired connect | ion 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 gueried NIC information.

iii. Run the following commands to bring eth1 down and up:

ifdown eth1

ifup eth1

- e. Restart the network service.
 - i. For CentOS 6, run the following command to restart the network service:

service network restart

ii. For CentOS 7/EulerOS/Fedora, run the following command to restart the network service:

systemctl restart NetworkManager

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

SUSE, openSUSE, or CoreOS

SUSE 11 SP4 does not support dynamic IPv6 address assignment.

No additional configuration is required for SUSE 12 SP1 or SUSE 12 SP2.

No additional configuration is required for openSUSE 13.2 or openSUSE 42.2.

No additional configuration is required for CoreOS 10.10.5.

----End

Setting the Timeout Duration for IPv6 Address Assignment

After automatic IPv6 address assignment is configured on an ECS running CentOS 6.x or Debian, the ECS will be created as a private image. When this image is used to create an ECS in an environment that IPv6 is unavailable, the ECS may start slow because acquiring an IPv6 address times out. Before creating the private image, you can set the timeout duration for acquiring IPv6 addresses to 30s as follows:

- CentOS 6.x.
 - a. Run the following command to edit the **dhclient.conf** file:
 - vi /etc/dhcp/dhclient.conf
 - b. Press i to enter editing mode and add the timeout attribute to the file. timeout 30;
 - c. Enter :wq to save the settings and exit.
- Debian 7.5:
 - a. Run the following command to edit the **networking** file:
 - vi /etc/init.d/networking
 - b. Press i to enter editing mode and add the timeout attribute.

Figure 7-20 Modification 1

Figure 7-21 Modification 2

- 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
 - Press i to enter editing mode and add the timeout attribute to the file.
 [Service]
 TimeoutStartSec=30
- Debian 9.0
 - Run the following command to edit the networking.service file:
 vi /etc/system/system/network-online.target.wants/ networking.service
 - b. Press i to enter editing mode and change **TimeoutStartSec=5min** to **TimeoutStartSec=30**.

 $\mathbf{8}_{\scriptscriptstyle \mathsf{EIPs}}$

8.1 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.

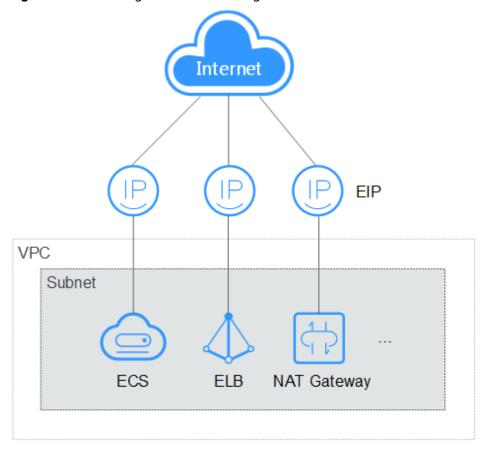


Figure 8-1 Accessing the Internet using an EIP

Helpful Links

- Binding an EIP
- Changing an EIP
- Changing an EIP Bandwidth

8.2 Binding an EIP

Scenarios

You can assign an EIP and bind it to an ECS to enable the ECS to access the Internet.

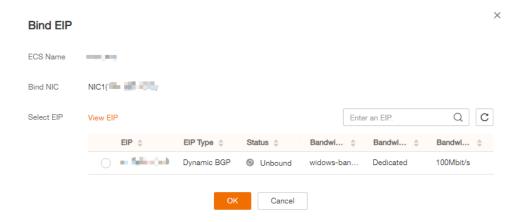
Procedure

- 1. Log in to the management console.
- 2. Under **Computing**, click **Elastic Cloud Server**.
- 3. In the ECS list, select the ECS to which an EIP is to be bound, and choose More > Manage Network > Bind EIP in the Operation column.
- 4. In the displayed dialog box, select an EIP

◯ NOTE

If no EIP is available in the current region, the EIP list is empty. In such a case, allocate an EIP and then bind it.

Figure 8-2 Binding an EIP



Click OK.

After the EIP is bound, view it in the ECS list on the **Elastic Cloud Server** page.

8.3 Unbinding an EIP

Scenarios

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

Procedure

- 1. Log in to the management console.
- 2. Under **Computing**, click **Elastic Cloud Server**.
- 3. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network** > **Unbind EIP**.

Figure 8-3 Unbinding an EIP



4. Confirm the EIP to be unbound and click Yes.

Unreleased EIPs will continue to be billed. To stop the EIPs from being billed, release them.

8.4 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. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network** > **Unbind EIP**.
- 3. Confirm the displayed information and click Yes.

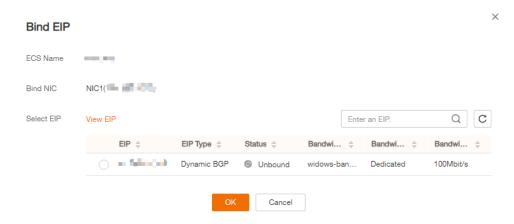
Binding a New EIP

- Log in to the management console.
- 2. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network** > **Bind EIP**.
- 3. 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 8-4 Binding a new EIP



8.5 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. Under Computing, click Elastic Cloud Server.
- 3. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network** > **Modify Bandwidth**.
- 4. Change the bandwidth name, billing mode, or bandwidth size as prompted.

8.6 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.

NAT Gateway is recommended, which provides both the SNAT and DNAT functions for your ECSs in a VPC and allows the ECSs to access or provide services accessible from the Internet. For details, see *NAT Gateway User Guide*.

Prerequisites

- A proxy ECS with an EIP bound is available.
- The IP address of the proxy ECS is in the same network and same security group as the ECSs that need to access the Internet.

Linux Proxy ECS

In this example, the proxy ECS runs CentOS 6.5.

- 1. Log in to the management console.
- 2. Under Computing, click Elastic Cloud Server.
- 3. In the search box above the upper right corner of the ECS list, enter the proxy ECS name for search.
- 4. Click the name of the proxy ECS. The page providing details about the ECS is displayed.
- 5. On the **Network Interfaces** tab, click . Then, disable **Source/Destination Check**.

Figure 8-5 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.

- 6. Log in to the proxy ECS.
 - For more details, see Login Overview.
- 7. Run the following command to check whether the proxy ECS can access the Internet:

ping www.google.com

The proxy ECS can access the Internet if information similar to the following is displayed:

```
64 bytes from 220.181.111.148: icmp_seq=1 ttl=51 time=9.34 ms 64 bytes from 220.181.111.148: icmp_seq=2 ttl=51 time=9.11 ms 64 bytes from 220.181.111.148: icmp_seq=3 ttl=51 time=8.99 ms
```

8. Run the following command to check whether IP forwarding is enabled on the proxy ECS:

cat /proc/sys/net/ipv4/ip_forward

- If 0 (disabled) is displayed, go to 9.
- If 1 (enabled) is displayed, go to 14.
- 9. Run the following command to open the IP forwarding configuration file in the vi editor:

vi /etc/sysctl.conf

- 10. Press i to enter editing mode.
- 11. Set the net.ipv4.ip_forward value to 1.

Set the **net.ipv4.ip_forward** value to **1**.

◯ NOTE

If the **sysctl.conf** file does not contain the **net.ipv4.ip_forward** parameter, run the following command to add it:

echo net.ipv4.ip_forward=1 >> /etc/sysctl.conf

12. Press **Esc**, type :wq, and press **Enter**.

The system saves the configurations and exits the vi editor.

13. Run the following command to make the modification take effect:

sysctl -p /etc/sysctl.conf

14. Run the following commands to configure default **iptables** rules:

iptables -P INPUT ACCEPT

iptables -P OUTPUT ACCEPT

iptables -P FORWARD ACCEPT



Running **iptables -P INPUT ACCEPT** will set default INPUT policy to ACCEPT, which poses security risks. You are advised to set security group rules to restrict inbound access.

15. Run the following command to configure source network address translation (SNAT) to enable ECSs in the same network segment to access the Internet through the proxy ECS:

iptables -t nat -A POSTROUTING -o eth0 -s *subnet/netmask-bits* **-j SNAT -- to** *nat-instance-ip*

For example, if the proxy ECS is in network 192.168.125.0, the subnet mask has 24 bits, and the private IP address is 192.168.125.4, run the following command:

iptables -t nat -A POSTROUTING -o eth0 -s 192.168.125.0/24 -j SNAT --to 192.168.125.4

◯ NOTE

To retain the preceding configuration even after the ECS is restarted, run the **vi /etc/rc.local** command to edit the **rc.local** file. Specifically, copy the rule described in step **15** into **rc.local**, press **Esc** to exit Insert mode, and enter :**wq** to save the settings and exit.

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

service iptables save

chkconfig iptables on

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

iptables -t nat --list

SNAT has been configured if information similar to **Figure 8-6** is displayed.

Figure 8-6 Successful SNAT configuration



- 18. Add a route.
 - a. Log in to the management console.
 - b. Under **Network**, click **Virtual Private Cloud**.
 - c. Choose **Route Tables** in the left navigation pane. In the route table list, click a target route table. On the displayed page, click **Add Route**.
 - d. Set route information on the displayed page.
 - Destination: indicates the destination network segment. The default value is 0.0.0.0/0.
 - Next Hop: indicates the private IP address of the proxy ECS.
 You can obtain the private IP address of the ECS on the Elastic Cloud Server page.
- 19. To delete the added iptables rules, run the following command:

iptables -t nat -D POSTROUTING -o eth0 -s *subnet/netmask-bits* **-j SNAT -- to** *nat-instance-ip*

For example, if the proxy ECS is in network segment 192.168.125.0, the subnet mask has 24 bits, and the private IP address is 192.168.125.4, run the following command:

iptables -t nat -D POSTROUTING -o eth0 -s 192.168.125.0/24 -j SNAT --to 192.168.125.4

9 Security

9.1 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 9-1 Methods for improving ECS security

| Туре | 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: | Enabling HSSMonitoring ECSsBacking 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. | Improving the Port Security Periodically Upgrading the Operating System |

Enabling HSS

HSS is designed to improve the overall security for ECSs. It helps you identify and manage the information on your ECSs, eliminate risks, and defend against intrusions and web page tampering.

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.

For details, see Installing an Agent on the Linux OS and Enabling the Basic/Enterprise/Premium Edition.

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.

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.

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.

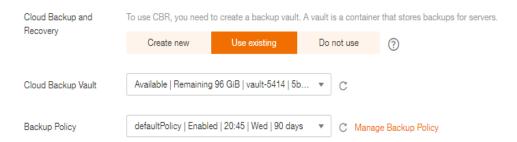
To enable CBR when purchasing an ECS:

Set CBR when purchasing an ECS. The system will associate the ECS with a cloud backup vault and the selected backup policy to periodically back up the ECS.

Auto assign

- a. Set the name of the cloud backup vault, which is a character string consisting of 1 to 64 characters, including letters, digits, underscores (_), and hyphens (-). For example, **vault-f61e**. The default naming rule is **vault** xxxx.
- b. Enter the vault capacity, which is required for backing up the ECS. The vault capacity cannot be smaller than that of the ECS to be backed up. Its value ranges from the total capacity of the ECS to 10,485,760 in the unit of GB.
- c. Select a backup policy from the drop-down list, or log in to the CBR console and configure a desired one.
- Use existing
 - a. Select an existing cloud backup vault from the drop-down list.
 - b. Select a backup policy from the drop-down list, or log in to the CBR console and configure a desired one.
- Do not use: Skip the CBR setting. If you require this function after purchasing the ECS, log in to the CBR console and bind the desired cloud backup vault to your ECS.

Figure 9-1 Setting CBR



To back up data for a created ECS:

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.

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 9-2 lists common high-risk ports. You are advised to change these ports to non-high-risk ports.

Table 9-2 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.

9.2 Security Groups

9.2.1 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**.

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**.

Security Group Constraints

- By default, you can add up to 50 security group rules to a security group.
- By default, you can add an ECS or extension NIC to up to five security groups. In such a case, the rules of all the selected security groups are aggregated to take effect.

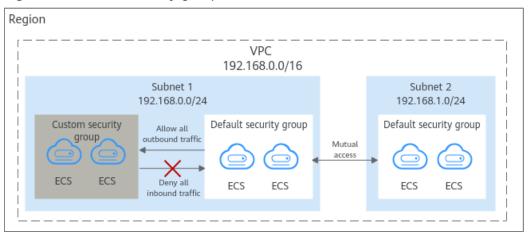
9.2.2 Default Security Group and Rules

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 9-2 shows the default security group.

Figure 9-2 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 9-3 describes the rules in the default security group.

| Directi on | Protoc ol | Port/ Range | Source/ Destination | Description |
|---------------|--------------|----------------|--|--|
| Outbo und | All | All | Destination: 0.0.0.0/0 | Allows all outbound traffic. |
| Inboun d | 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). |

Table 9-3 Default security group rules

9.2.3 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 9-4.

Table 9-4 Default outbound rules in a security group

| Directi on | Protoc ol & Port | Destination | Description |
|---------------|------------------------|-------------|--|
| Outbou nd | All | 0.0.0.0/0 | This rule allows the instances in the security group to access any IPv4 address over any port. |

Remotely Logging In to an ECS from a Local Server

A security group denies all external requests by default. To remotely log in to an ECS in a security group from a local server, add an inbound rule based on the OS running on the ECS.

- To remotely log in to a Linux ECS using SSH, enable port 22. For details, see Table 9-5.
- To remotely log in to a Windows ECS using RDP, enable port 3389. For details, see **Table 9-6**.

Table 9-5 Remotely logging in to a Linux ECS using SSH

| Direction | Protocol & Port | Source |
|-----------|-----------------|-----------------------|
| Inbound | TCP: 22 | IP address: 0.0.0.0/0 |

Table 9-6 Remotely logging in to a Windows ECS using RDP

| Direction | Protocol & Port | Source |
|-----------|-----------------|-----------------------|
| Inbound | TCP: 3389 | IP address: 0.0.0.0/0 |

NOTICE

If the source is set to 0.0.0.0/0, any IP address can be used to remotely log in to the ECS. To ensure security, set the source to a specific IP address based on service requirements. For details about the configuration example, see **Table 9-7**.

Table 9-7 Remotely logging in to an ECS using a specified IP address

| ECS Type | Direction | Protocol & Port | Source |
|-----------|-----------|--------------------|----------------------------|
| Linux ECS | Inbound | TCP: 22 | IP address: 192.168.0.0/24 |

| ECS Type | Direction | Protocol & Port | Source |
|-------------|-----------|--------------------|--------------------------|
| Windows ECS | Inbound | TCP: 3389 | IP address: 10.10.0.0/24 |

Remotely Connecting to an ECS from a Local Server to Upload or Download FTP Files

By default, a security group denies all external requests. If you need to remotely connect to an ECS from a local server to upload or download files, you need to enable FTP ports 20 and 21.

Table 9-8 Remotely connecting to an ECS from a local server to upload or download files

| Direction | Protocol & Port | Source |
|-----------|-----------------|-----------------------|
| Inbound | TCP: 20-21 | IP address: 0.0.0.0/0 |

NOTICE

You must first install the FTP server program on the ECSs and check whether ports 20 and 21 are working properly.

Setting Up a Website on an ECS to Provide Services Externally

A security group denies all external requests by default. If you have set up a website on an ECS that can be accessed externally, you need to add an inbound rule to the ECS security group to allow access over specific ports, such as HTTP (80) and HTTPS (443).

Table 9-9 Setting up a website on an ECS to provide services externally

| Direction | Protocol & Port | Source |
|-----------|-----------------|-----------------------|
| Inbound | TCP: 80 | IP address: 0.0.0.0/0 |
| Inbound | TCP: 443 | IP address: 0.0.0.0/0 |

Using ping Command to Verify Network Connectivity

Ping works by sending an Internet Control Message Protocol (ICMP) Echo Request. To ping an ECS from your PC to verify the network connectivity, you need to add an inbound rule to the security group of the ECS to allow ICMP traffic.

Table 9-10 Using ping command to verify network connectivity

| Direction | Protocol & Port | Source |
|-----------|-----------------|-----------------------|
| Inbound | ICMP: All | IP address: 0.0.0.0/0 |

Enabling Communications Between Instances in Different Security Groups

Instances in the same VPC but associated with different security groups cannot communicate with each other. If you want ECSs in security group **sg-A** to access MySQL databases in security group **sg-B**, you need to add an inbound rule to security group **sg-B** to allow access from ECSs in security group **sg-A**.

Table 9-11 Enabling communications between instances in different security groups

| Direction | Protocol & Port | Source |
|-----------|-----------------|----------------------|
| Inbound | TCP: 3306 | Security group: sg-A |

Allowing External Instances to Access the Database Deployed on an ECS

A security group denies all external requests by default. If you have deployed a database on an ECS and want the database to be accessed from external instances on a private network, you need to add an inbound rule to the security group of the ECS to allow access over corresponding ports. Here are some common ports for databases:

MySQL: port 3306
Oracle: port 1521
MS SQL: port 1433
PostgreSQL: port 5432

Redis: port 6379

Table 9-12 Allowing external instances to access the database deployed on an ECS

| Direction | Protocol & Port | Source | Description |
|-----------|--------------------|-------------------------|--|
| Inbound | TCP: 3306 | Security group: sg-A | This rule allows the ECSs in security group sg-A to access the MySQL database service. |
| Inbound | TCP: 1521 | Security group: sg-B | This rule allows the ECSs in security group sg-B to access the Oracle database service. |

| Direction | Protocol & Port | Source | Description |
|-----------|--------------------|-------------------------------|---|
| Inbound | TCP: 1433 | IP address: 172.16.3.21/32 | This rule allows the ECS whose private IP address is 172.16.3.21 to access the MS SQL database service. |
| Inbound | 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. |

NOTICE

In this example, the source is for reference only. Set the source address based on your requirements.

Allowing ECSs to Access Specific External Websites

By default, a security group allows all outbound traffic. **Table 9-14** lists the default rules. If you want to allow ECSs to access specific websites, configure the security group as follows:

1. Add outbound rules to allow traffic over specific ports to specific IP addresses.

Table 9-13 Allowing ECSs to access specific external websites

| Direct ion | Protocol & Port | Destination | Description |
|---------------|--------------------|------------------------------|--|
| Outbo und | TCP: 80 | IP address: 132.15.XX.XX | This rule allows ECSs in the security group to access the external website at http://132.15.XX.XX:80. |
| Outbo und | TCP: 443 | IP address: 145.117.XX.XX | This rule 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.

| | | | · · · · · · · · · · · · · · · · · · · |
|---------------|------------------------|-------------|--|
| Directi on | Protoc ol & Port | Destination | Description |
| Outbou nd | All | 0.0.0.0/0 | This rule allows the instances in the security group to access any IPv4 address over any port. |

Table 9-14 Default outbound rules in a security group

9.2.4 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 **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 9-15 Inbound rule parameter description

| Param eter | Description | Example Value |
|---------------|--|------------------|
| Туре | Source IP address version. You can select: | IPv4 |
| | • IPv4 | |
| | • IPv6 | |

| Param eter | Description | Example Value |
|------------------------|--|------------------|
| Protoc ol & Port | The network protocol used to match traffic in a security group rule. The value can be All , TCP , UDP , GRE , and ICMP . | ТСР |
| | 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. IP address: Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6) IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) 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 |
| Descrip tion | 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 9-16 Outbound rule parameter description

| Param eter | Description | Example Value |
|------------------------|---|------------------|
| Туре | Destination IP address version. You can select:IPv4IPv6 | IPv4 |
| Protoc ol & Port | The network protocol used to match traffic in a security group rule. The value can be All , TCP , UDP , GRE , and ICMP . | ТСР |

| Param eter | Description | Example Value |
|-----------------|---|------------------|
| | Destination port used to match traffic in a security group rule. The value can be from 1 to 65535. | 22, or 22-30 |
| | Outbound rules control outgoing traffic over specific ports from instances in the security group. | |
| Destina tion | 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: | 0.0.0.0/0 |
| | IP address: | |
| | Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) | |
| | All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6) | |
| | IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) | |
| Descrip tion | Supplementary information about the security group rule. This parameter is optional. | N/A |
| | The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >). | |

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

9.2.5 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.
- 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.

Figure 9-3 Change Security Group



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 be aggregated to apply on the ECS.

To create a security group, click Create Security Group.

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

5. Click OK.

9.3 HSS

What Is HSS?

Host Security Service (HSS) is designed to improve the overall security for ECSs. It helps you identify and manage the information on your ECSs, eliminate risks, and defend against intrusions and web page tampering.

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 Use HSS?

Before using the HSS service, install the HSS agent on your ECS. For details, see **Installing an Agent on the Linux OS** and **Enabling the Enterprise/Premium Edition**.

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 & Compliance** > **Host Security Service**.
- 3. On the **Server** tab, check the ECS security statuses.

Figure 9-4 ECS list

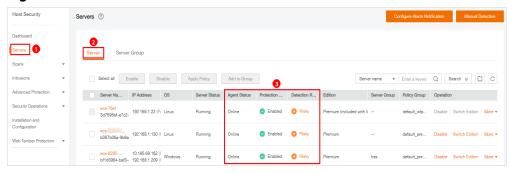


Table 9-17 Statuses

| Parameter | Description |
|-----------------|---|
| Agent Status | 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 | Enabled: The ECS is properly protected using HSS. |
| Status | 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 | Risky: The ECS is risky. |
| Result | Safe: No risks are detected. |
| | Pending risk detection: HSS is not enabled for the ECS. |

For more details, see **HSS**.

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.

Parameter Requirement **Example Value** Password Consists of 8 to 26 characters. YNbUwp! dUc9MClnv Contains at least three of the following NOTE character types: The example Uppercase letters password is generated Lowercase letters randomly. Do not Digits use it. - Special characters: \$!@%-_=+[]:./^,{}? • 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

• Cannot start with a slash (/) for Windows

applies only to Windows ECSs.)

Table 10-2 Password complexity requirements

10.1.2 Changing the Login Password on an ECS

ECSs.

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.

| Parameter | Requirement | Example Value |
|-----------|---|--|
| Password | Consists of 8 to 26 characters. Contains at least three of the following character types: Uppercase letters Lowercase letters Digits Special characters: \$!@%=+[]:./^,{}? 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. | YNbUwp! dUc9MClnv NOTE The example password is generated randomly. Do not use it. |

Table 10-3 Password complexity requirements

Windows

- 1. Log in to the ECS.
 - For details, see **Login Overview**.
- 2. Press Win+R to start the Run dialog box.
- 3. Enter **cmd** to open the command-line interface (CLI) window.
- 4. Run the following command to change the password (the new password must meet the requirements described in **Table 10-3**):

net user Administrator New password

Linux

- Use the existing key file to log in to the ECS as user root through SSH.
 For details, see Remotely 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. Under **Computing**, click **Elastic Cloud Server**.
 - c. Stop the original Windows ECS, switch to the page providing details about the ECS, and click the **Disks** tab.

□ NOTE

Do not forcibly stop the Windows ECS. Otherwise, password reset may fail.

- d. Locate the row containing the system disk to be detached and click **Detach** to detach the system disk from the ECS.
- 2. Attach the system disk to the temporary ECS.
 - On the page providing details about the temporary ECS, click the **Disks** tab.
 - b. Click **Attach Disk**. In the displayed dialog box, select the system disk detached in step **1.d** and attach it to the temporary ECS.
 - c. Remotely log in to the temporary ECS.

d. Run the following command to view the directory of the system disk detached from the original Windows ECS now attached to the temporary ECS:

fdisk -l

e. Run the following command to mount the file system of the detached system disk to the temporary ECS:

mount -t ntfs-3g /dev/Result obtained in step 2.d /mnt/

For example, if the result obtained in step **2.d** is **xvde2**, run the following command:

mount -t ntfs-3g /dev/xvde2 /mnt/

If the following error information is displayed after the preceding command is executed, the NTFS file systems may be inconsistent. In such a case, rectify the file system inconsistency.

The disk contains an unclean file system (0, 0).

Metadata kept in Windows cache, refused to mount.

Failed to mount '/dev/xvde2': Operation not permitted

The NTFS partition is in an unsafe state. Please resume and shutdown

Windows fully (no hibernation or fast restarting), or mount the volume read-only with the 'ro' mount option.

Back up the disk data, run the following command to rectify the NTFS file system inconsistency, and attach the system disk:

ntfsfix /dev/Result obtained in step 2.d

For example, if the result obtained in step **2.d** is **xvde2**, run the following command:

ntfsfix /dev/xvde2

- 3. Change the password of the specified user and clear the original password.
 - a. Run the following command to back up the SAM file:

cp /mnt/Windows/System32/config/SAM /mnt/Windows/System32/config/SAM.bak

b. Run the following command to change the password of the specified user:

chntpw -u Administrator /mnt/Windows/System32/config/SAM

c. Enter **1**, **q**, and **y** as prompted, and press **Enter**.

The password has been reset if the following information is displayed:

Select: [q] > 1 Password cleared! Select: [q] > q Hives that have changed: #Name 0<SAM> Write hive files? (y/n) [n] : y 0<SAM> - OK

- 4. Stop the temporary ECS, detach the system disk, and attach the system disk to the original Windows ECS.
 - a. Stop the temporary ECS, switch to the page providing details about the ECS, and click the **Disks** tab.
 - b. Click **Detach** to detach the data disk temporarily attached in step **2.b**.
 - c. On the page providing details about the original Windows ECS, click the **Disks** tab.

- d. Click **Attach Disk**. In the displayed dialog box, select the data disk detached in step **4.b** and device name **/dev/sda**.
- 5. Start the original Windows ECS and set a new login password.
 - a. Click **Start** to start the original Windows ECS. After the status becomes **Running**, click **Remote Login** in the **Operation** column.
 - b. Click **Start**. Enter **CMD** in the search box and press **Enter**.
 - c. Run the following command to change the password (the new password must meet the requirements described in **Table 10-4**):

net user Administrator New password

Table 10-4 Password complexity requirements

| Parameter | Requirement | Example Value |
|-----------|--|--|
| Password | Consists of 8 to 26 characters. Contains at least three of the following character types: Uppercase letters Lowercase letters Digits Special characters: \$!@%=+[]:./^, {}? 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. | YNbUwp! dUc9MClnv NOTE The example password is generated randomly. Do not use it. |

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. Download the script for resetting the password and upload the script to the temporary ECS.

Contact customer service to obtain the password reset script. Use a connection tool, such as WinSCP, to upload the obtained **changepasswd.sh** script to the temporary ECS.

To download WinSCP, log in at https://winscp.net/.

- 2. Stop the original Linux ECS, detach the system disk from it, and attach the system disk to the temporary ECS.
 - a. Log in to the management console.
 - b. Under Computing, click Elastic Cloud Server.
 - c. Stop the original ECS, switch to the page providing details about the ECS, and click the **Disks** tab.
 - ∩ NOTE

Do not forcibly stop the original ECS. Otherwise, password reset may fail.

- d. Locate the row containing the system disk to be detached and click **Detach** to detach the system disk from the ECS.
- 3. Attach the system disk to the temporary ECS.
 - a. On the page providing details about the temporary ECS, click the **Disks** tab.
 - b. Click **Attach Disk**. In the displayed dialog box, select the system disk detached in step **2.d** and attach it to the temporary ECS.
- 4. Log in to the temporary ECS remotely and reset the password.
 - a. Locate the row containing the temporary ECS and click **Remote Login** in the **Operation** column.
 - 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
                 2048 83884031 83881984 40G 83 Linux
/dev/vda1 *
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
                2048 83886079 83884032 40G 83 Linux
/dev/vdb1 *
```

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.

If the original ECS and the temporary ECS both run CentOS 7, a mount failure may occur during script execution. To resolve this issue, replace **mount \$dev \$mountPath** with **mount -o nouuid \$dev \$mountPath** in the script.

d. Enter the new password and the directory obtained in step **4.b** as prompted.

If the following information is displayed, the password has been changed: set password success.

5. (Optional) Enable remote root login for non-root users.

vi /etc/ssh/sshd_config

Modify the following settings:

- Change PasswordAuthentication no to PasswordAuthentication yes.
 Alternatively, uncomment PasswordAuthentication yes.
- Change PermitRootLogin no to PermitRootLogin yes.
 Alternatively, uncomment PermitRootLogin yes.
- Change the value of **AllowUsers** to **root**.

Search for **AllowUsers** in the file. If **AllowUsers** is missing, add it at the end of the file.

- 6. Stop the temporary ECS, detach the system disk, attach the system disk to the original Linux ECS, and restart the original Linux ECS.
 - a. Stop the temporary ECS, switch to the page providing details about the ECS, and click the **Disks** tab.
 - b. Click **Detach** to detach the data disk attached in step 2.
 - c. On the page providing details about the original Linux ECS, click the **Disks** tab.
 - d. Click **Attach Disk**. In the displayed dialog box, select the data disk detached in **6.b**.
- 7. Restart the original Linux ECS.

10.2 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. Therefore, you do not need to worry about password interception, cracking, or leakage.

Scenarios

When purchasing an ECS, you are advised to select the key pair login mode. For Windows ECSs, key pairs are required to decrypt the passwords so that you can use the decrypted password to log in.

- Logging in to a Linux ECS
 - You can directly use a key pair to log in.
 - When you are creating the ECS, select the key pair login mode. For details, see "Set Login Mode" in Step 3: Configure Advanced Settings.
 - After the ECS is created, bind a key pair to the ECS by referring to "Binding a Key Pair" in the *Data Encryption Workshop User Guide*.
- Logging in to a Windows ECS

You can use the key pair to obtain a password for login. The password is randomly generated and therefore is more secure.

For details, see Obtaining the Password for Logging In to a Windows ECS.

Creating a Key Pair

You can create a key pair or use an existing one for remote login authentication.

Creating a key pair

You can create a key pair using either of the following method:

- Follow the instructions in (Recommended) Creating a Key Pair on the Management Console. The public key is automatically stored in the system, and the private key is stored locally.
- Follow the instructions in **Creating a Key Pair Using PuTTYgen**. Both the public and private keys are stored locally.
 - After the key pair is created, import the key pair following the instructions provided in **Importing a Key Pair** so that you can use it.
- Using an existing key pair

If an existing key pair (created using PuTTYgen, for example) is available, you can import the public key by referring to **Importing a Key Pair** on the management console to let the system maintain the public key for you.

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. Under **Computing**, click **Elastic Cloud Server**.
- 3. In the navigation pane on the left, choose **Key Pair**.

- 4. On the displayed page, click Create Key Pair.
- 5. Enter a key pair name.

A key pair name consists of two parts: KeyPair and four random digits (KeyPair-xxxx).

- Click OK.
- 7. Manually or automatically download a .pem private key file with the name that you specify as the key name. Store it in a secure place and click **OK**.

◯ NOTE

This is the only chance for you to save the private key file. Keep it secure. You'll need to provide the key pair name when you create an ECS, and the corresponding private key each time you connect to the ECS through SSH.

10.2.3 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**.

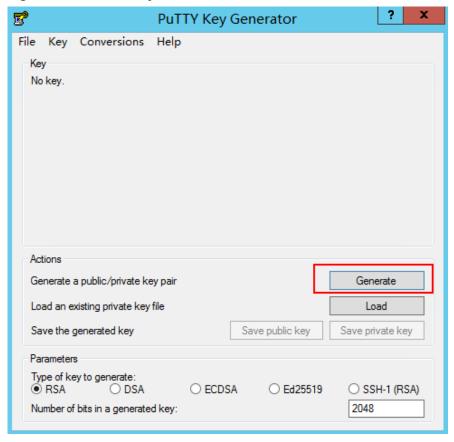


Figure 10-2 PuTTY Key Generator

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.

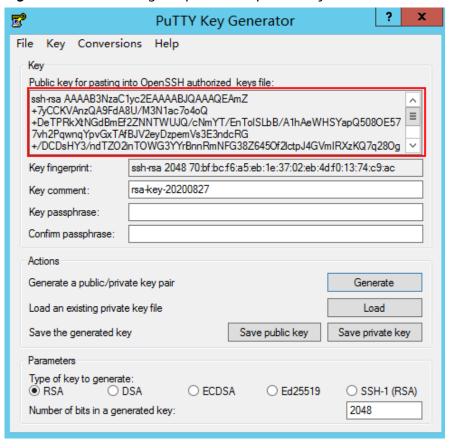


Figure 10-3 Generating the public and private keys

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**.

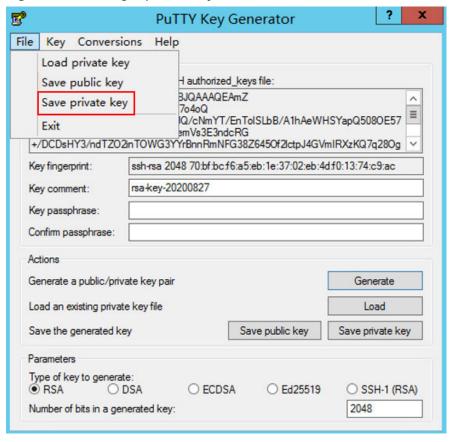


Figure 10-4 Saving a private key

- 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.

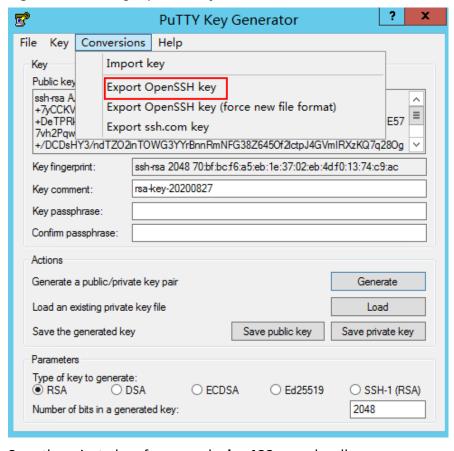


Figure 10-5 Saving a private key

- 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. Under Computing, click Elastic Cloud Server.

- 3. In the navigation pane on the left, choose **Key Pair**.
- 4. On the **Key Pair Service** page, click **Import Key Pair**.
- 5. Use either of the following methods to import the key pair:
 - Selecting a file
 - i. In the Import Key Pair dialog box of the management console, click Select File and select the locally stored public key file (for example, the .txt file saved in 3 in Creating a Key Pair Using PuTTYgen).

□ 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?

10.2.5 Obtaining and Deleting the Password of a Windows ECS

10.2.5.1 Obtaining the Password for Logging In to a Windows ECS

Scenarios

Password authentication is required to log in to a Windows ECS. Therefore, you must use the key file used when you created the ECS to obtain the administrator password generated during ECS creation. The administrator user is **Administrator** or the user configured using Cloudbase-Init. This password is randomly generated, offering high security.

You can obtain the initial password for logging in to a Windows ECS through the management console. For details, see this section.

Obtaining the Password Through the Management Console

- 1. Obtain the private key file (.pem file) used when you created the ECS.
- 2. Log in to the management console.
- 3. Under Computing, click Elastic Cloud Server.
- 4. On the **Elastic Cloud Server** page, select the target ECS.
- 5. In the **Operation** column, click **More** and select **Get Password**.

Figure 10-6 Get Password

- Use either of the following methods to obtain the password through the key file:
 - Click Select File and upload the key file from a local directory.
 - Copy the key file content to the text field.
- 7. Click **Get Password** to obtain a random password.

10.2.5.2 Deleting the Initial Password for Logging In to a Windows ECS

Scenarios

After you obtain the initial password, it is a good practice to delete it to ensure system security.

Deleting the initial password does not affect ECS operation or login. Once deleted, the password cannot be retrieved. Before you delete a password, it is a good practice to record it.

Procedure

- 1. Log in to the management console.
- 2. Under Computing, click Elastic Cloud Server.
- 3. On the **Elastic Cloud Server** page, select the target ECS.
- 4. In the **Operation** column, click **More** and select **Delete Password**. The system displays a message, asking you whether you want to delete the password.
- 5. Click **OK** to delete the password.

1 1 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 employees 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**).

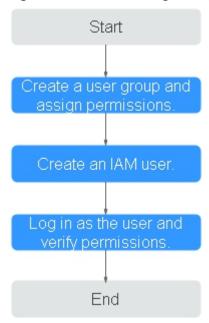
Prerequisites

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



- Create a user group and assign permissions to it.
 Create a user group on the IAM console, and attach the ECSReadOnlyAccess policy to the group.
- 2. Create an IAM user.
 - Create a user on the IAM console and add the user to the group created in 1.
- 3. Log in and verify permissions.
 - Log in to the ECS console by using the created user, and verify that the user only has read permissions for ECS.
 - 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 section contains examples of common ECS custom policies.

Example Custom Policies

• Example 1: Allowing users to stop and delete multiple ECSs at a time

• Example 2: Denying ECS deletion

A policy with only "Deny" permissions must be used in conjunction with other policies to take effect. If the permissions assigned to a user contain both "Allow" and "Deny", the "Deny" permissions take precedence over the "Allow" permissions.

The following method can be used if you need to assign permissions of the **ECSFullAccess** policy to a user but you want to prevent the user from deleting ECSs. Create a custom policy for denying ECS deletion, and attach both policies to the group to which the user belongs. Then, the user can perform all operations on ECSs except deleting ECSs. The following is an example of a deny policy:

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.

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

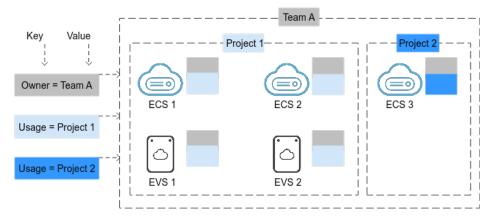


Figure 12-1 shows how tags work. In this example, you assign two tags to each cloud resource. Each tag contains a key and a value that you define. The key of one tag is **Owner**, and the key of another tag is **Usage**. Each tag has a value.

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 | Cannot be left blank. | Organization |
| | 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 (_). | |
| Value | Can contain a maximum of 43 characters. | Apache |
| | Can only consist of digits, letters, hyphens (-), underscores (_), and periods (.). | |

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

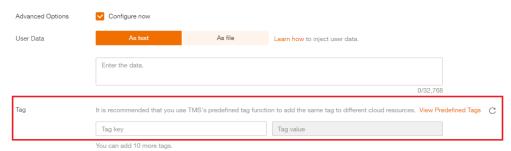
- 1. Log in to the management console.
- 2. Under Computing, choose Elastic Cloud Server.
- 3. Click Create ECS.
- 4. Configure parameters for the ECS.

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

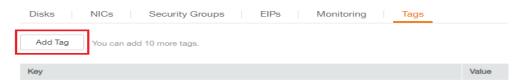


Adding Tags on the ECS Details Page

- 1. Log in to the management console.
- Under Computing, choose Elastic Cloud Server.
- 3. In the ECS list, click the name of the target ECS.
 - The ECS details page is displayed.
- 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.

- 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



- 5. Click **Edit** to make the resource list editable.
- 6. Locate the row containing the target ECS, click $^{\bigodot}$, 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



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:

- Log in to the management console.
- 2. Under Management & Deployment, choose Tag Management Service.
- 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 gueried.
 - 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. Under Management & Deployment, choose Tag Management Service.
- 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.
- 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 $oldsymbol{oldsymbol{arphi}}$.
- 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. On the **Resource Tags** page, set the search criteria for ECSs and click **Search**.
- 4. Select the target ECSs.
- 5. Click Manage Tag in the upper left corner of the list.
- 6. In the displayed **Manage Tag** dialog box, click **Delete** in the **Operation** column. Click **OK**.
- 7. (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. In the upper right corner of the page, click The **Service Quota** page is displayed.
- View the used and total quota of each type of resources on the displayed page.

If a quota cannot meet service requirements, apply for a higher quota.

How Do I Apply for a Higher Quota?

The system does not support online quota adjustment.

If you need to adjust a quota, contact the operations administrator.

$13_{\text{Monitoring}}$

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)
- 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. √ indicates that the metric is supported, and x indicates that the metric is not supported.

■ 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/.

Table 13-1 Basic ECS metrics

| Metric | Windows | | Linux | | |
|----------------------------|-----------|-----------|---|---------------|--|
| 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 | |
| Disk Usage | Supported | Supported | Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.) | Not supported | |
| Disk Read Bandwidt h | Supported | Supported | Supported | Supported | |

| Metric | Windows | | Linux | | |
|--------------------------------|--|-----------|--|---------------|--|
| Disk Write Bandwidt h | 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 | |
| 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 | |
| Network Connectio ns | Not supported | Supported | Not supported | Supported | |

| Metric | Windows | | Linux | | |
|-------------------------------------|---------------|---------------|---------------|---------------|--|
| Inbound Bandwidt h | Not supported | Supported | Not supported | Supported | |
| Outbound Bandwidt h | Not supported | Supported | Not supported | Supported | |
| Inbound PPS | Not supported | Supported | Not supported | Supported | |
| Outbound PPS | Not supported | Supported | Not supported | Supported | |
| New Connectio ns | Not supported | Supported | Not supported | Supported | |
| System Status Check Failed | Supported | Not supported | Supported | Not 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 minutesKVM ECSs: 5 minutes

Table 13-2 Basic metric description

| Metric ID | Parame ter | Description | Value Range | Monitore d Object & Dimensio n | 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 |

| Metric ID | Parame ter | Description | Value Range | Monitore d Object & Dimensio n | Monitoring Interval (Raw Metrics and KVM Only) |
|------------------------------|-------------------------------|---|----------------|--|---|
| 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 |
| disk_util_i nband | 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_rat e | Disk Read Bandwi dth | 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 byte_out = (rd_bytes - last_rd_bytes)/Time difference | ≥ 0 | ECS | 5 minutes |

| Metric ID | Parame ter | Description | Value Range | Monitore d Object & Dimensio n | Monitoring Interval (Raw Metrics and KVM Only) |
|----------------------------------|--------------------------------|--|----------------|--|---|
| disk_write _bytes_rat e | Disk Write Bandwi dth | 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 |
| 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)/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)/Time difference | ≥ 0 | ECS | 5 minutes |

| Metric ID | Parame ter | Description | Value Range | Monitore d Object & Dimensio n | Monitoring Interval (Raw Metrics and KVM Only) |
|---|----------------------------------|---|----------------|--|---|
| network_i ncoming_ bytes_rate _inband | Inband Incomin g 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 |
| network_ outgoing_ bytes_rate _inband | Inband Outgoin g 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_i ncoming_ bytes_agg regate_rat e | Outban d Incomin g 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 |

| Metric ID | Parame ter | Description | Value Range | Monitore d Object & Dimensio n | Monitoring Interval (Raw Metrics and KVM Only) |
|---|----------------------------------|---|----------------|--|---|
| network_ outgoing_ bytes_agg regate_rat e | Outban d Outgoin g 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 |
| network_ vm_conne ctions | Networ k Connect ions | Total number of TCP and UDP connections to an ECS Unit: count NOTE This metric is collected out-of-band and its value may be greater than the number of network connections queried in the OS. | ≥ 0 | ECS | 5 minutes |
| network_ vm_band width_in | Inbound Bandwi dth | Number of public and private bits received by the ECS per second Unit: byte/s | ≥ 0 | ECS | 5 minutes |
| network_ vm_band width_out | Outbou nd Bandwi dth | Number of public and private bits sent by the ECS per second Unit: byte/s | ≥ 0 | ECS | 5 minutes |
| network_ vm_pps_i n | Inbound PPS | Number of public and private packets received by the ECS per second Unit: packet/s | ≥ 0 | ECS | 5 minutes |

| Metric ID | Parame ter | Description | Value Range | Monitore d Object & Dimensio n | Monitoring Interval (Raw Metrics and KVM Only) |
|---------------------------------------|------------------------|--|----------------|--|---|
| network_ vm_pps_o ut | Outbou nd PPS | Number of public and private packets sent by the ECS per second Unit: packet/s | ≥ 0 | ECS | 5 minutes |
| network_ vm_newc onnection s | New Connect ions | Number of new connections (including TCP, UDP, and ICMP) created on the ECS Unit: count | ≥ 0 | ECS | 5 minutes |

| Metric ID | Parame ter | Description | Value Range | Monitore d Object & Dimensio n | Monitoring Interval (Raw Metrics and KVM Only) |
|-------------------------------|-------------------------------------|--|----------------------|--|---|
| inst_sys_s tatus_erro r | System Status Check Failed | Status of the cloud platform for running ECSs, which can be 0 or 1 • 0: The system is running properly. All check items are normal. • 1: The system is not running properly. At least one check item failed. When the power source of the physical server fails or the hardware/ software becomes faulty, the check result is 1. The system periodically checks the status and returns check results using value 0 or 1. • If the detected fault does not affect ECS functions, certain management operations performed on the ECS, such as starting, stopping, or specifications modifications, may be affected. • If the detected fault affects ECS functions, such as a server power supply failure, the system will | 0 or 1 | ECS | 5 minutes |

| Metric ID | Parame ter | Description | Value Range | Monitore d Object & Dimensio n | Monitoring Interval (Raw Metrics and KVM Only) |
|-----------|---------------|--------------------------------------|----------------|--|---|
| | | recover the ECS as soon as possible. | | | |

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 | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|--------------------|------------------------------|--|----------------|--|---|
| cpu_usa ge | (Agent) CPU Usage | CPU usage of the monitored object Unit: percent 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_usa ge_idle | (Agent) Idle CPU Usage | Percentage of time that CPU is idle Unit: percent 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 | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|--------------------------|--|---|----------------|--|---|
| cpu_usa ge_user | (Agent) User Space CPU Usage | Percentage of time that the CPU is used by user space Unit: percent • 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_usa ge_syst em | (Agent) Kernel Space CPU Usage | Percentage of time that the CPU is used by kernel space Unit: percent • 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 | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|-------------------------|---|---|----------------|--|---|
| cpu_usa ge_othe r | (Agent) Other Process CPU Usage | Percentage of time that the CPU is used by other processes Unit: percent 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_usa ge_nice | (Agent) Nice Process CPU Usage | Percentage of time that the CPU is in user mode with low-priority processes which can easily be interrupted by higher-priority processes Unit: percent • 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 | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|---------------------------|---|--|----------------|--|---|
| cpu_usa ge_iowa it | (Agent) iowait Process CPU Usage | Percentage of time that the CPU is waiting for I/O operations to complete Unit: percent • 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_usa ge_irq | (Agent) CPU Interrupt Time | Percentage of time that the CPU is servicing interrupts Unit: percent 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_usa ge_softi rq | (Agent) CPU Software Interrupt Time | Percentage of time that the CPU is servicing software interrupts Unit: percent 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 | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|--------------------|---|--|----------------|--|---|
| load_av erage1 | (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_av erage5 | (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_av erage15 | (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 | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|-------------------|--------------------------------|---|----------------|--|---|
| mem_av ailable | (Agent) Available Memory | Amount of memory that is available and can be given instantly to processes Unit: GB Linux: Obtain the metric value from / proc/meminfo. If MemAvailable is displayed in /proc/meminfo, obtain the value. If MemAvailable is not displayed in / proc/meminfo, MemAvailable = MemFree + Buffers+Cached Windows: 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 | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|-------------------------|----------------------------|---|----------------|--|---|
| mem_us edPerce nt | (Agent) Memory Usage | Memory usage of the monitored object Unit: percent Linux: Obtain the metric value from the /proc/meminfo file: (MemTotal - MemAvailable)/ MemTotal If MemAvailable is displayed in /proc/meminfo, MemUsedPercent = (MemTotal-MemAvailable)/ MemTotal If MemAvailable is not displayed in /proc/meminfo, MemUsedPercent = (MemTotal - MemFree - Buffers - Cached)/ MemTotal Windows: The calculation formula is as follows: Used memory size/Total memory size*100%. | 0-100 | ECS | 1 minute |
| mem_fr ee | (Agent) Idle Memory | Amount of memory that is not being used Unit: GB Linux: Obtain the metric value from / proc/meminfo. Windows is not supported currently. | ≥ 0 | ECS | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|----------------------|----------------------------------|--|----------------|--|---|
| mem_buf fers | (Agent) Buffer | Amount of memory that is being used for buffers Unit: GB 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_ca ched | (Agent) Cache | Amount of memory that is being used for file caches Unit: GB 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_op en_files | (Agent) Total File Handles | Total handles used by all processes Unit: count • 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 | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|---------------|---------------------------------------|---|----------------|--|---|
| disk_fre e | (Agent) Available Disk Space | Free space on the disks Unit: GB 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 - Moun t point | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|----------------|-------------------------------|--|----------------|--|---|
| disk_tot al | (Agent) Disk Storage Capacity | Total space on the disks, including used and free Unit: GB 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 - Moun t point | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|------------|----------------------------------|--|----------------|--|---|
| disk_use d | (Agent) Used Disk Space | Used space on the disks Unit: GB 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 - Moun t point | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|--------------------------|--------------------------|--|----------------|--|---|
| disk_use dPercen t | (Agent) Disk Usage | Percentage of total disk space that is used, which is calculated as follows: Disk Usage = Used Disk Space/Disk Storage Capacity Unit: percent 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 - Moun t point | 1 minute |

OS Metric: Disk I/O

Table 13-7 Disk I/O metrics

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|--------------------------------------|----------------------------------|---|----------------|--|---|
| disk_agt _read_b ytes_rat e | (Agent) Disks Read Rate | Number of bytes read from the monitored disk per second Unit: byte/s • 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: - Use Win32_PerfFormatt edData_PerfDisk_Lo gicalDisk 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 | • EC S - Dis k • EC S - Mo unt poi nt | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|--------|---------------|---|----------------|--|---|
| | | When the CPU usage is high, monitoring data obtaining timeout may occur and result in the failure of obtaining monitoring data. | | | |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|---|-----------------------------|--|-----------------|--|---|
| disk_agt _read_re quests_r ate | (Agent) Disks Read Requests | Number of read requests sent to the monitored disk per second Unit: request/s 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: Use Win32_PerfFormatt edData_PerfDisk_Lo gicalDisk 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 | EC S - Dis k EC S - Mo unt poi nt | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|--------|---------------|---|----------------|--|---|
| | | may occur and result in the failure of obtaining monitoring data. | | | |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|---------------------------------------|--------------------------|--|----------------|--|---|
| disk_agt _write_b ytes_rat e | (Agent) Disks Write Rate | Number of bytes written to the monitored disk per second Unit: byte/s • 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. 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 Win32_PerfFormatt edData_PerfDisk_Lo gicalDisk 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 | • EC S - Dis k • EC S - Mo unt point | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|--------|---------------|---|----------------|--|---|
| | | may occur and result in the failure of obtaining monitoring data. | | | |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|--|------------------------------|---|-----------------|--|---|
| disk_agt _write_r equests_ rate | (Agent) Disks Write Requests | Number of write requests sent to the monitored disk per second Unit: request/s • 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. 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 Win32_PerfFormatt edData_PerfDisk_Lo gicalDisk 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 | EC S - Dis k EC S - Mo unt poi nt | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|-------------------|---|--|------------------|--|---|
| | | may occur and result in the failure of obtaining monitoring data. | | | |
| disk_rea dTime | (Agent) Average Read Request Time | Average amount of time that read requests have waited on the disks Unit: ms/count 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 | • EC S - Dis k • EC S - Mo unt poi nt | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|--------------------|--|--|------------------|---|---|
| disk_wri teTime | (Agent) Average Write Request Time | Average amount of time that write requests have waited on the disks Unit: ms/count 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 | EC S - Dis k EC S - Mo unt poi nt | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|------------------|------------------------------|--|----------------|---|---|
| disk_ioU tils | (Agent) Disk I/O Usage | Percentage of the time that the disk has had I/O requests queued to the total disk operation time Unit: percent 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 | EC S - Dis k EC S - Mo unt poi nt | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|---------------------------|---------------------------|---|----------------|--|---|
| disk_que ue_lengt h | (Agent) Disk Queue Length | Average number of read or write requests queued up for completion for the monitored disk in the monitoring period Unit: count 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 | EC S - Dis k EC S - Mo unt poi nt | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|--|---|---|-----------------|--|---|
| disk_wri te_bytes _per_op eration | (Agent) Average Disk Write Size | Average number of bytes in an I/O write for the monitored disk in the monitoring period Unit: byte/op 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. 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 bytes/op | • EC S - Dis k • EC S - Mo unt point | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|---|---|---|-----------------|--|---|
| disk_rea d_bytes_ per_ope ration | (Agent) Average Disk Read Size | Average number of bytes in an I/O read for the monitored disk in the monitoring period Unit: byte/op 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. 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 bytes/op | EC S - Dis k EC S - Mo unt poi nt | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|-------------------|--|---|----------------|--|---|
| disk_io_ svctm | (Agent) Disk I/O Service Time | Average time in an I/O read or write for the monitored disk in the monitoring period Unit: ms/op • 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 | • EC S - Dis k • EC S - Mo unt point | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|----------------------------------|--------------------------|---|----------------|--|---|
| disk_dev ice_used _percent | Block Device Usage | Percentage of the physical disk usage of the monitored object. Calculation formula: Used storage space of all mounted disk partitions/ Total disk storage space 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 | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|---------------------|---|---|--|--|---|
| disk_fs_r wstate | (Agent) File System Read/ Write Status | Read and write status of the mounted file system of the monitored object Possible values are 0 (read and write) and 1 (read only). Linux: Check file system information in the fourth column in file /proc/mounts. | • 0: reada ble and writa ble • 1: read- only | ECS - Moun t point | 1 |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|--------------------------------|---|---|----------------|--|---|
| disk_ino desTotal | (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 - Moun t point | 1 minute |
| disk_ino desUsed | (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 - Moun t point | 1 minute |
| disk_ino desUsed Percent | (Agent) Percentag e 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 - Moun t point | 1 minute |

□ NOTE

The Windows OS does not support the file system metrics.

OS Metric: NIC

Table 13-9 NIC metrics

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|-----------------|--|--|----------------|--|---|
| net_bitR ecv | (Agent) Outboun d Bandwidt h | Number of bits received by this NIC per second Unit: bit/s • 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_bitS ent | (Agent) Inbound Bandwidt h | Number of bits sent by this NIC per second Unit: bit/s • 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 | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|--------------------|---|---|-----------------|--|---|
| net_pac ketRecv | (Agent) NIC Packet Receive Rate | Number of packets received by this NIC per second Unit: count/s • 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_pac ketSent | (Agent) NIC Packet Send Rate | Number of packets sent by this NIC per second Unit: count/s • 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_erri n | (Agent) Receive Error Rate | Percentage of receive errors detected by this NIC per second Unit: percent 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 | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitori ng Period (Raw Data) |
|-----------------|---|--|----------------|--|---|
| net_erro ut | (Agent) Transmit Error Rate | Percentage of transmit errors detected by this NIC per second Unit: percent 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_dro pin | (Agent) Received Packet Drop Rate | Percentage of packets received by this NIC which were dropped per second Unit: percent 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_dro pout | (Agent) Transmitt ed Packet Drop Rate | Percentage of packets transmitted by this NIC which were dropped per second Unit: percent 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 | Para meter | Description | Value Range | Moni tored Obje ct & Dime nsion | Monito ring Period (Raw Data) |
|----------------|-----------------|---|----------------|--|---|
| ntp_offs et | (Agen t) NTP | NTP offset of the monitored object | ≥ 0 ms | ECS | 1 minute |
| | Offset | Unit: ms | | | |
| | | Collection method for Linux ECSs: Run chronyc sources -v to obtain the offset. | | | |

OS Metric: TCP

Table 13-11 TCP metrics

| Metric | Para meter | Description | Value Range | Moni tored Obje ct & Dime nsion | Monito ring Period (Raw Data) |
|-------------------|--------------------------|--|----------------|--|---|
| net_tcp_ total | (Agen t) TCP TOTAL | Total number of TCP connections in all states Unit: count • 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 | Para meter | Description | Value Range | Moni tored Obje ct & Dime nsion | Monito ring Period (Raw Data) |
|-----------------------------|--|--|----------------|--|---|
| net_tcp_ establis hed | (Agen t) TCP ESTAB LISHE D | Number of TCP connections in ESTABLISHED state Unit: count 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 | (Agen t) TCP SYS_S ENT | Number of TCP connections that are being requested by the client Unit: count 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_recv | (Agen t) TCP SYS_R ECV | Number of pending TCP connections received by the server Unit: count • 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 | Para meter | Description | Value Range | Moni tored Obje ct & Dime nsion | Monito ring Period (Raw Data) |
|---------------------------|----------------------------------|---|----------------|--|---|
| net_tcp_f in_wait1 | (Agen t) TCP FIN_W AIT1 | Number of TCP connections waiting for ACK packets when the connections are being actively closed by the client Unit: count 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_f in_wait2 | (Agen t) TCP FIN_W AIT2 | Number of TCP connections in the FIN_WAIT2 state Unit: count • 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_w ait | (Agen t) TCP TIME_ WAIT | Number of TCP connections in TIME_WAIT state Unit: count 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 | Para meter | Description | Value Range | Moni tored Obje ct & Dime nsion | Monito ring Period (Raw Data) |
|----------------------------|-----------------------------------|--|----------------|--|---|
| net_tcp_ close | (Agen t) TCP CLOSE | Number of closed TCP connections Unit: count 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_w ait | (Agen t) TCP CLOSE _WAIT | Number of TCP connections in CLOSE_WAIT TCP state Unit: count • 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 | (Agen t) TCP LAST_ ACK | Number of TCP connections waiting for ACK packets when the connections are being passively closed by the client Unit: count • 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 | Para meter | Description | Value Range | Moni tored Obje ct & Dime nsion | Monito ring Period (Raw Data) |
|---------------------|---|---|----------------|--|---|
| net_tcp_ listen | (Agen t) TCP LISTE N | Number of TCP connections in the LISTEN state Unit: count 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 | (Agen t) TCP CLOSI NG | Number of TCP connections to be automatically closed by the server and the client at the same time Unit: count • 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 | (Agen t) TCP Retran smissi on Rate | Percentage of packets that are resent Unit: percent 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 | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitor ing Period (Raw Data) |
|---------------------------|-------------------------|--|--|--|---|
| gpu_sta tus | GPU Health Status | Overall measurement of the GPU health Unit: 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: The GPU is heal thy. 1: The GPU is subh ealt hy. 2: The GPU is fault y. | • ECS • ECS - GP U | 1 minute |
| gpu_usa ge_enco der | Encoding Usage | Encoding capability usage on the GPU Unit: percent • 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 • ECS - GP U | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitor ing Period (Raw Data) |
|--------------------------------------|---|---|----------------|--|---|
| gpu_usa ge_deco der | Decoding Usage | Decoding capability usage on the GPU Unit: percent 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 • ECS - GP U | 1 minute |
| gpu_vol atile_co rrectabl e | Volatile Correctab le ECC Errors | Number of correctable ECC errors since the GPU is reset. The value is reset to 0 each time the GPU is reset. Unit: count 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 | • ECS • ECS - GP U | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitor ing Period (Raw Data) |
|---|--|---|----------------|--|---|
| gpu_vol atile_un correcta ble | Volatile Uncorrect able ECC Errors | Number of uncorrectable ECC errors since the GPU is reset. The value is reset to 0 each time the GPU is reset. Unit: count 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 | • ECS • ECS - GP U | 1 minute |
| gpu_ag gregate _correct able | Aggregat e Correctab le ECC Errors | Aggregate correctable ECC errors on the GPU Unit: count 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 | • ECS • ECS - GP U | 1 minute |
| gpu_ag gregate _uncorr ectable | Aggregat e Uncorrect able ECC Errors | Aggregate uncorrectable ECC Errors on the GPU Unit: count 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 | • ECS • ECS - GP U | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitor ing Period (Raw Data) |
|---|---|---|----------------|--|---|
| gpu_reti red_pag e_single _bit | Retired Page Single Bit Errors | Number of retired page single bit errors, which indicates the number of single-bit pages blocked by the graphics card Unit: count • Linux: Obtain the metric value using the librvidia-ml.so.1 library file of the graphics card. • Windows: Obtain the metric value using the nvml.dll library of the graphics card. | ≥ 0 | • ECS • ECS - GP U | 1 minute |
| gpu_reti red_pag e_doubl e_bit | Retired Page Double Bit Errors | Number of retired page double bit errors, which indicates the number of double-bit pages blocked by the graphics card Unit: count 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 | • ECS • ECS - GP U | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitor ing Period (Raw Data) |
|-------------------------------|-----------------------------|--|---|--|---|
| gpu_per formanc e_state | (Agent) Performa nce Status | GPU performance of the monitored object Unit: 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 P0: indic ates the maximu m perform ance stat us. P15: indic ates the minimu m perform ance stat us. P32: indic ates the unk now n perform ance stat us. | ECS - GPU | 1 minute |

| Metric | Paramet er | Description | Value Range | Monit ored Objec t & Dime nsion | Monitor ing Period (Raw Data) |
|-----------------------|-----------------------------------|--|----------------|--|---|
| gpu_usa ge_me m | (Agent) GPU Memory Usage | GPU memory usage of the monitored object Unit: percent • 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_usa ge_gpu | (Agent) GPU Usage | GPU usage of the monitored object Unit: percent 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 | Param eter | Description | Value Range | Moni tored Objec t & Dime nsion | Monito ring Period (Raw Data) |
|--------------------------|---------------|--|----------------|--|---|
| proc_pH ashId_cp u | CPU Usage | CPU consumed by a process. pHashId (process name and process ID) is the value of md5 . | 0-100% | ECS | 1 minute |
| | | Unit: percent Linux: Check metric value changes in file / proc/pid/stat. Windows: Call the Windows API | | | |
| | | GetProcessTimes to obtain the CPU usage of the process. | | | |

| Metric | Param eter | Description | Value Range | Moni tored Objec t & Dime nsion | Monito ring Period (Raw Data) |
|---------------------------|---------------------|---|----------------|--|---|
| proc_pH ashId_m em | Memo ry Usage | Memory consumed by a process. pHashId (process name and process ID) is the value of md5 . Unit: percent | 0-100% | ECS | 1 minute |
| | | Linux: RSS*PAGESIZE/ MemTotal | | | |
| | | Obtain the RSS value by checking the second column of file / proc/pid/statm. | | | |
| | | Obtain the PAGESIZE value by running the getconf PAGESIZE command. | | | |
| | | Obtain the MemTotal value by checking file / proc/meminfo . | | | |
| | | Windows: Call the Windows API procGlobalMemoryStatusEx to obtain the total memory size. Call GetProcessMemoryInfo to obtain the used memory size. Use the used memory size to divide the total memory size to get the memory usage. | | | |
| proc_pH ashId_fil e | Open Files | Number of files opened by a process. pHashId (process name and process ID) is the value of md5 . | ≥0 | ECS | 1 minute |
| | | Linux: Run the ls -l / proc/pid/fd command to view the number of opened files. | | | |
| | | Windows is not supported currently. | | | |

| Metric | Param eter | Description | Value Range | Moni tored Objec t & Dime nsion | Monito ring Period (Raw Data) |
|----------------------------|---|--|----------------|--|---|
| proc_run ning_cou nt | (Agent) Runnin g Proces ses | Number of processes that are running • 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 Proces ses | Number of processes that are idle 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_zo mbie_co unt | (Agent) Zombi e Proces ses | Number of zombie processes 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 | Param eter | Description | Value Range | Moni tored Objec t & Dime nsion | Monito ring Period (Raw Data) |
|-----------------------------|--|--|----------------|--|---|
| proc_blo cked_cou nt | (Agent) Blocke d Proces ses | Number of processes that are blocked 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_sle eping_co unt | (Agent) Sleepi ng Proces ses | Number of processes that are sleeping 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_tot al_count | (Agent) Total Proces ses | Total number of processes on the monitored object 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 | Param eter | Description | Value Range | Moni tored Objec t & Dime nsion | Monito ring Period (Raw Data) |
|------------------------------|---|--|----------------|--|---|
| proc_spe cified_co unt | (Agent) Specifi ed Proces ses | Number of specified processes 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 | • EC S - Pr oc ess | 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.

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.

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 | ecs | attachVolume attachVolumeV2 |
| Attaching a disk (on the EVS console) | ecs | attachVolume2 |
| Detaching a disk | ecs | detachVolume |
| 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 Audit Logs

Scenarios

CTS starts to record ECS operations after it is provisioned. You can view the operation records of the last seven days on the management console.

This section describes how to view the operation records.

Procedure

- 1. Log in to the management console.
- Click Service List. Under Management & Deployment, choose Cloud Trace Service.
- 3. In the navigation pane on the left, choose **Trace List**.
- 4. Specify filter criteria as needed. The following filter criteria are available:
 - Trace Type, Trace Source, Resource Type, and Search By:
 Select a filter criterion from the drop-down list.
 If you select Resource ID for Search By, you need to select or enter a specific resource ID.
 - Operator: Select a specific operator (which is a user rather than the tenant).
 - Trace Status: Available options include All trace statuses, Normal,
 Warning, and Incident. You can only select one of them.
 - Time Range: In the upper-right corner, you can select any time range of the last seven days to view traces generated during that period.
- 5. Expand the trace for details.

6. Click **View Trace**. A dialog box is displayed, in which the trace structure details are displayed.

For more information about CTS, see *Cloud Trace Service User Guide*.

15 Troubleshooting

15.1 What Can I Do If Switching from a Non-root User to User root Times Out?

Symptom

When you run the **sudo** command to switch to user **root** on an Ubuntu or Debian ECS, the system prompts connection timeout.

Figure 15-1 Connection timeout

```
linux@ubantu-test-1:/etc$
linux@ubantu-test-1:/etc$ sudo su
sudo: unable to resolve host ubantu-test-1: Connection timed out
root@ubantu-test-1:/etc#
```

Solution

- 1. Log in to the ECS.
- 2. Run the following command to edit the hosts configuration file:

vi /etc/hosts

- 3. Press i to enter editing mode.
- 4. Add the IP address and hostname to the last line of the hosts configuration file.

Private IP address hostname

An example is provided as follows:

If the ECS hostname is **hostname** and the private IP address of the ECS is 192.168.0.1, add the following statement:

192.168.0.1 hostname

- 5. Press **Esc** to exit editing mode.
- 6. Run the following command to save the configuration and exit:

:wq

To update the hostname of an Ubuntu or Debian ECS, set the value of parameter manage_etc_hosts in the /etc/cloud/cloud.cfg file to false and update the new hostname in the /etc/hosts file. When editing the /etc/hosts file, do not delete the statement in the line where 127.0.0.1 is located. Otherwise, switching from a non-root user to user root will time out.

15.2 What Should I Do If Error "command 'gcc' failed with exit status 1" Occurs During PIP-based Software Installation

Symptom

When installing the Python library software, you need to configure the PIP source. Take the official image source as an example:

```
[root@test home]# cat /root/.pip/pip.conf
[global]
index-url = https://pypi.python.org/simple
trusted-host = pypi.python.org
```

During the installation, the system displays the message "command 'gcc' failed with exit status 1". However, GCC has been installed by running the yum command before the Python library software is installed using the PIP.

Figure 15-2 Installation error

```
creating Guild/Kemp.linux-xoo_64-2.7/psutil
gcc-pthread-fno-strict-aliasing -02 -g -pipe -Wall -Wp,-D_FORTIFY_SOURCE-2 -fexceptions -fstack-protector-strong --params-sp-biffer-size-4 -grecord-gcc-switches -m64 -mtune-generic -D_GNU_SOURCE -FFIC -fwrapv -DNDEBUG -02 -g -pipe -Wall -Wp,-D_FORTIFY_SOURCE-2 -fexceptions -fstack-protector-strong --params-sp-buffer-size-4 -grecord-gcc-switches -m64 -mtune-generic -D_GNU_SOURCE -FFIC -fwrapv -DNDEBUG -02 -g -pipe -Wall -Wp,-D_FORTIFY_SOURCE-2 -fexceptions -fstack-protector-strong --params-sp-buffer-size-4 -grecord-gcc-switches -m64 -mtune-generic -D_GNU_SOURCE-2FFIC -fwrapv -PFIC -DPSUTIL_POSIX=1 -DPSUTIL_UERSION-543 -DPSUTIL_LINUX=1 -1/usr/include/python2.7 -c psutil/_psutil_common.c -o build/temp.linux-x66_64-2.7/psutil/_psutil_common.o psutil/_psutil_common.c9:28: fatal error: Python.h: No such file or directory

#include <Python.h>

compilation terminated.
error: command 'gcc' failed with exit status 1

Command 'wsr/bin/python2 -u -c 'import setuptools, tokenize; file_='/tmp/pip-build-uhwaPs/psutil/setup.py':exec(compile(getat tr(tokenize, 'open', open') -gen') - gen') - gen')
```

Possible Causes

openssl-devel is not supported.

Solution

The following operations use psutil as an example:

- 1. Run the following command to install openssl-devel: yum install gcc libffi-devel python-devel openssl-devel -y
- 2. Use PIP to install the Python library software again. The error message is cleared.

Figure 15-3 Successful installation

```
| TrootRecs-9x80-test1.pip1# pip install psutil
| Collecting psuti|
| Using cached https://mirrors.aliyun.com/pypi/packages/e2/e1/600326635f97fee09bf0426fef14c5c29f4049c79f60fd79f433d0c1bd96/psutil-5.4.3.tar.gz
| Installing collected packages: psutil | Running setup.py install for psutil ... done
| Successfully installed psutil-5.4.3|
```

15.3 What Should I Do If Packages Are Downloaded Using PIP or wget at a Low Rate?

Symptom

When a user runs the wget command to download software packages, the download rate is far less than the bandwidth.

Figure 15-4 wget-based package downloading

Possible Causes

The official PIP website is accessed using HTTPS. Each time PIP is used to install a third-party Python module, the source code package must be downloaded at the official PIP website. Therefore, openssl packages are required.

Solution

Run the following command to install openssl packages:

yum install openssl openssl-devel

15.4 How Can I Handle Slow ECS Startup?

If an ECS requires a long period of time to start, you can change the default timeout to speed up the startup.

- 1. Log in to the ECS.
- 2. Run the following command to switch to user **root**:

sudo su

3. Run the following command to obtain the grub version:

rpm -qa | grep grub

Figure 15-5 Viewing the grub version

```
[root@ ]# rpm -qa | grep grub
grub2-2.02-0.44.el7.centos.x86_64
```

- 4. Change the timeout in the grub file to 0s.
 - If the grub version is earlier than 2:
 Open the /boot/grub/grub.cfg or /boot/grub/menu.lst file and change the timeout value to 0.
 - If the grub version is 2:
 Open the /boot/grub2/grub.cfg file and change the timeout value to 0.

Figure 15-6 Changing timeout duration

```
#boot=/dev/sda
default=8
timeout=8
splashimage=(hd0,1)/boot/grub/splash.xpm.gz
hiddenmenu
title CentOS (2.6.32-696.16.1.el6.x86_64)
root (hd0,1)
kernel /boot/vmlinuz-2.6.32-696.16.1.el6.x86_64 ro root=UUID=2bc0f5fd-e0
19-4ba5-8ce0-0fe12b6efc24 rd_NO_LUKS rd_NO_LUM LANG=en_US.UTF-8 rd_NO_DM SYSFONT
= latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us rd_NO_DM rhgb q
```

16 FAQS

16.1 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

- Why Can't an EIP Be Pinged?
- Why Can I Remotely Access an ECS But Cannot Ping It?

16.2 ECS Overview

16.2.1 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 FCSs
- Do not change NIC MAC addresses. Otherwise, the network connection will fail.

16.2.2 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.

16.3 Regions and AZs

16.3.1 What Is an AZ?

ΑZ

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.

16.3.2 What Is a Region?

When you create an ECS, select the nearest region for low network latency and quick resource access.

16.3.3 Are Products Different in Different Regions?

Yes. Currently, each region contains different products. Certain products are available for trial release in certain regions only.

16.3.4 Is Data Transmission Between AZs Billed?

Data transmission between AZs in the same region is free of charge. However, data transmission between AZs in different regions will be billed.

16.3.5 Can a Load Balancer Distribute Traffic to ECSs in Different Regions?

Only dedicated load balancers support this. Backend servers can be from VPCs in different regions.

16.3.6 Is Application Disaster Recovery Available in Different Regions?

Yes.

You can deploy active and standby application nodes in different regions. If the active application node is faulty, the standby application node continues to provide services.

16.3.7 Are There Any Services Provided for Application Disaster Recovery?

No. Currently, the standard application disaster recovery service is unavailable now. If you have such a requirement, please contact us. We will customize an application disaster recovery solution based on your application scenarios.

16.3.8 Can Components Contained in an Application Be Distributed to Different Regions?

Yes. However, such a deployment mode is not recommended.

You are advised to deploy the components contained in an application in the same region. In this manner, these components can communicate with each other over an internal network, reducing bandwidth costs of using public networks and ensuring communication quality between the components.

16.4 Billing

16.4.1 What Are ECS Billing Modes?

ECSs can be billed in pay-per-use payments or applied with reserved instances (RIs). Choose a billing mode best suit you.

- In pay-per-use billing mode, fees are billed by hour. The service duration less than one hour is billed as one hour.
- For RI billing details, visit https://cloud.orange-business.com/offres/ infrastructure-iaas/flexible-engine/elastic-cloud-server/

16.4.2 Will Am I Continue to Be Billed After ECSs Are Stopped?

For pay-per-use ECSs:

- After a common instance (an ECS without local disks or FPGAs attached) is stopped, basic resources (vCPUs, memory, and image) are not billed, but system disks are still billed based on the capacity you use. The associated resources such as EVS disks, EIPs, and bandwidth, are separately billed.
- Special instances (ECSs with local disks attached, FPGA-based ECSs, or BMSs) will continue to be billed after being stopped. To stop the ECS from being billed, delete it and its associated resources.

□ 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.

For details, see How Can I Stop an ECS from Being Billed?

16.4.3 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 16-1** lists the resources associated with the ECS.

| Table | 16-1 | Rilling | example | of a | pay-per-use | FCS |
|-------|------|---------|-----------|------|-------------|-----|
| IUDIC | 101 | DILLIIA | CAUITIDIC | OI G | Day Dei ase | - |

| 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.

16.5 Creation and Deletion

16.5.1 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.

16.5.2 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.

16.5.3 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.

16.5.4 When Does an ECS Become Provisioned?

Pay-per-use ECS: The ECS is automatically provisioned after it is created.

16.5.5 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.

16.5.6 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.

 CSBS has a new backup format. You can rapidly create ECSs if the full-ECS image is in this format

Solution

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 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*.
 - 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.

16.5.7 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/ 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.

16.5.8 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.

If it takes a long time to obtain your ECS, contact customer service.

16.5.9 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.

16.5.10 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.

16.5.11 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 16-1 Deleting selected ECSs



16.5.12 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.

16.5.13 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.

16.5.14 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.

16.5.15 Can I Forcibly Restart or Stop an ECS?

Yes. If an ECS remains in the **Restarting** or **Stopping** state for over 30 minutes after it is restarted, you can forcibly restart or stop the ECS as follows:

- 1. Log in to the management console.
- 2. Under Computing, click Elastic Cloud Server.

- 3. Select the target ECS and click **Restart** or **Stop**.

 A dialog box is displayed to confirm whether you want to restart or stop the
- 4. Select Forcibly restart the preceding ECSs or Forcibly stop the preceding ECSs.
- 5. Click **OK**.

16.6 Login and Connection

16.6.1 What Are the Login Requirements for ECSs?

Windows

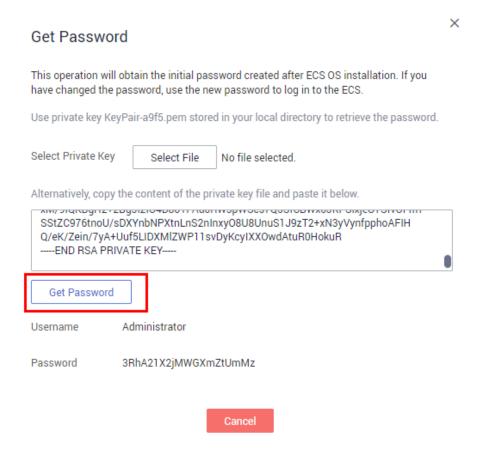
- Ensure that the ECS has an EIP bound (only required for MSTSC logins).
 For details, see Assigning an EIP and Binding It to an ECS.
- Make sure that the remote desktop protocol has been enabled on the ECS (only required for MSTSC logins).
 - If MSTSC-based remote desktop connection is used, log in to the ECS using VNC and enable the remote desktop protocol.

For details, see Remotely Logging In to a Windows ECS (Using MSTSC).

More information:

- If you created your ECS by using an external image file and the ECS does not have the password reset plug-in installed, reset the ECS login password by attaching a disk to the ECS. For details, see Resetting the Password for Logging In to a Windows ECS.
- If your ECS is authenticated using a key pair, parse the private key file to a password before you log in to the ECS.
 - a. Locate the target ECS.
 - b. In the **Operation** column, click **More** and select **Get Password**.
 - Copy the content of the private key file and paste it into the text box.
 Click **Get Password** to obtain a random password.

Figure 16-2 Get Password



Linux

- Ensure that the ECS has an EIP bound (only required for SSH logins).

 SSH logins are available for Linux ECSs only. You can use a remote login tool to log in to your ECS. In such a case, the ECS must have an EIP bound.
 - Verify that the ECS has an EIP bound.
 For details, see Assigning an EIP and Binding It to an ECS.
 - Check whether the EIP bound to the ECS can be pinged.
 - If you use a public IP address, see Why Can't an EIP Be Pinged? for troubleshooting.

More information:

- If your ECS was created using an external image file and the ECS has not had password reset plug-in installed, reset the ECS login password by attaching a disk to the ECS. For details, see Resetting the Password for Logging In to a Windows ECS.
- For a Linux ECS authenticated using a key pair:
 - For the first login, use an SSH key. For details, see Login Using an SSH Key.

- 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.
- For a key-pair-authenticated ECS, using a private key file to obtain its login password will fail because Cloud-Init may fail to inject the password.

16.6.2 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, do as follows:

- Resetting the Password for Logging In to a Windows ECS
- Resetting the Password for Logging In to a Linux ECS

16.6.3 What Should I Do If Starting an ECS Remains in "Waiting for cloudResetPwdAgent" State?

Symptom

During ECS startup, it remains in "Waiting for cloudResetPwdAgent" state for 20-30 seconds.

Figure 16-3 Starting cloudResetPwdAgent

```
Starting rpcbind: [ OK ]
Starting MFS statd: [ OK ]
Starting cloudResetPwdAgent...
Waiting for cloudResetPwdAgent....
```

Possible Causes

This issue is caused by the intranet DNS and user-defined DNS configurations.

Solution

- Log in to the ECS as user root.
- Run the following command to modify the /etc/cloud/cloud.cfg configuration file:
 - vi /etc/cloud/cloud.cfg
- 3. Add the following statement to the configuration file:

manage_etc_hosts: true

Figure 16-4 Editing the configuration file

```
users:
- name: root
lock_passwd: False

disable_root: 0
ssh_pwauth: 1

datasource_list: ['OpenStack']
manage_etc_hosts: true

datasource:
OpenStack:
# timeout: the timeout value for a request at metadata service
timeout: 50
# The length in seconds to wait before giving up on the metadata
# service. The actual total wait could be up to
# len(reso]vable_metadata_urls)*timeout
max_wait: 120
```

16.6.4 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 **Login Using VNC on the Management Console**. Then, locate the login fault by referring to **Fault Locating**.

Login Using VNC on the Management Console

Check whether you can log in to the ECS using VNC on the management console.

- 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**.
- 4. (Optional) When the system displays "Press CTRL+ALT+DELETE to log on", click **Send CtrlAltDel** in the upper part of the remote login page to log in to the ECS.

Figure 16-5 Send CtrlAltDel



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 16-2 Possible causes and solutions

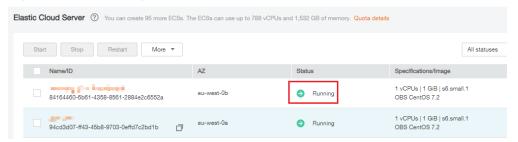
| Possible Cause | Solution |
|--|--|
| The ECS is frozen or stopped. | Make sure that the ECS is in the Running state. For details, see Checking the ECS Status . |
| The entered username or password is incorrect. | The default username for Windows ECSs is Administrator . For details, see Checking the 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. |
| No EIP is bound to the ECS. | To log in to an ECS using RDP or MSTSC, ensure that the ECS has an EIP bound. For details, see Verifying that the 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 Functional. |
| 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 SSH login IP address whitelist is configured in HSS. For details, see Checking the IP Address Allowlist for SSH Logins (with HSS Enabled). |
| 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. |

| Possible Cause | Solution |
|--|--|
| 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 16-6 Checking the ECS status



Checking the Login Mode

Check the login mode you set when you created the ECS.

Figure 16-7 Login Mode



- **Key pair**: If your ECS is authenticated using a key pair, parse the private key file to obtain a password.
 - a. Locate the target ECS.
 - b. In the **Operation** column, click **More** and select **Get Password**.
 - c. Copy the content of the private key file and paste it into the text box. Click **Get Password** to obtain a new random 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 using 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.

Verifying that the ECS Has an EIP Bound

An ECS can access the Internet only after it is bound with an EIP.

Before logging in to an ECS using RDP or MSTSC, make sure that an EIP has been bound to the ECS. For details, see **Assigning an EIP and Binding It to an ECS**.

Ⅲ NOTE

If you are logging in to an ECS through 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 Functional

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.

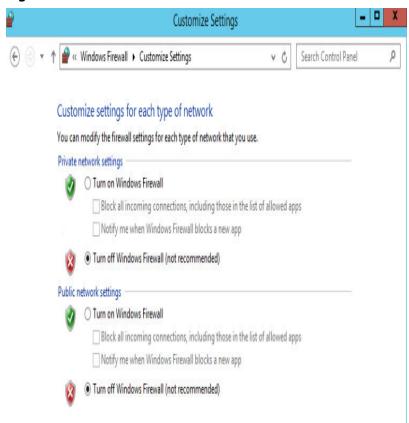
- 1. Log in to the Windows ECS.
- 2. Click the Windows icon in the lower left corner of the desktop and choose Control Panel > System and Security > Windows Firewall.

Figure 16-8 Windows Firewall



Choose Check firewall status > Turn Windows Firewall on or off.View and set the firewall status.

Figure 16-9 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.

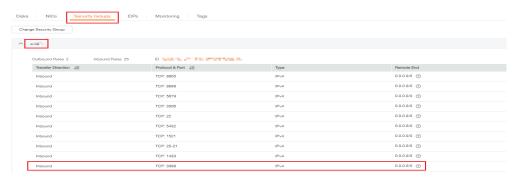
The default port is 3389. If you use another port, add that port in the inbound rule of the firewall.

After you perform the preceding operations, try to remotely log in to the ECS again.

Checking Whether the Remote Access Port Is Correctly Configured

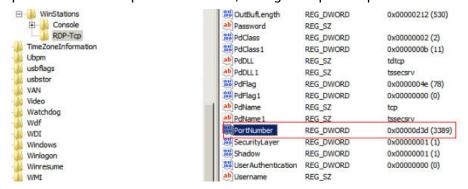
Check whether port 3389 (used by default) on the ECS is accessible.
 Ensure that port 3389 has been added in the inbound rule.
 On the page providing details about the ECS, click the Security Groups tab and view port 3389 in the inbound rule of the security group.

Figure 16-10 Checking remote access ports



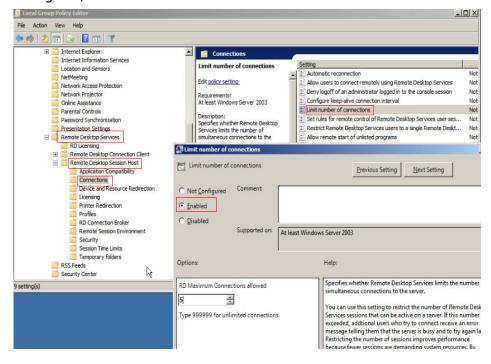
If you need to modify security group rules, see **Modifying a Security Group Rule**.

- 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 Allowlist for SSH Logins (with HSS Enabled)

After HSS is enabled, you can configure an IP address allowlist for SSH logins as required. The IP address allowlist 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 allowlist 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 allowlist.

! 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 Remotely Logging In to a Windows ECS (Using MSTSC).

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 of the ECS to the allowlist 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.

16.6.5 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 **Login Using VNC on the Management Console**. Then, locate the login fault by referring to **Fault Locating**.

Login Using VNC on the Management Console

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**.
- 4. (Optional) When the system displays "Press CTRL+ALT+DELETE to log on", click **Send CtrlAltDel** in the upper part of the remote login page to log in to the ECS.

Do not press **CTRL+ALT+DELETE** on the physical keyboard because this operation does not take effect.

If the VNC login still fails, record resource information and failure 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 16-3 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 Linux ECSs is root . 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. |
| No EIP is bound to the ECS. | To log in to an ECS using RDP or MSTSC, ensure that the ECS has an EIP bound. For details, see Verifying that the 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 Functional. |

| Possible Cause | Solution |
|--|--|
| 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 Remote Access Port 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. |
| An IP address whitelist for SSH logins has been configured. | Check whether an SSH login IP address whitelist is configured in HSS. For details, see Checking the IP Address Allowlist 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 16-11 Checking the ECS status



Checking the Login Mode

Check the login mode you set when you created the ECS.

Figure 16-12 Login Mode



• Key pair

- For the first login, use an SSH key. For details, see Remotely 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 using 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.

Verifying that the 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 Functional

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.

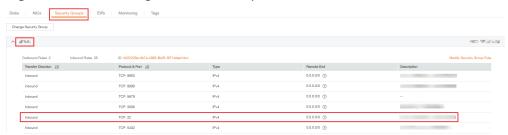
Run the following command to check whether port 22 is accessible:

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.

Figure 16-13 Checking remote access ports



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.

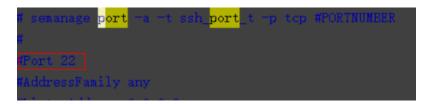
Checking Whether the Remote Access Port Is Correctly Configured

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.



After you perform the preceding operations, try to remotely log in to the ECS again.

Checking the IP Address Allowlist for SSH Logins (with HSS Enabled)

After HSS is enabled, you can configure an IP address allowlist for SSH logins as required. The IP address allowlist 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 allowlist 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 allowlist.



- 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 of the ECS to the allowlist 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.

16.6.6 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. Under **Computing**, click **Elastic Cloud Server**.
 - c. On the ECS list, click the name of an ECS for which you want to modify the security group rule.
 - d. On the ECS details page, click the security group in the **Security Groups** area to go to the security group details page.
 - e. On the **Inbound Rules** tab, click **Add Rule**. In the displayed dialog box, set **Protocol & Port** as follows:
 - Protocols: TCP (Custom ports)
 - Port: 2020

For details, see **Adding a Security Group Rule**.

- 2. Log in to the ECS.
- 3. In the **Run** dialog box, enter **regedit** to access the registry editor.
- 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.

Value name:

PortNumber

Value data:

Base

Hexadecimal

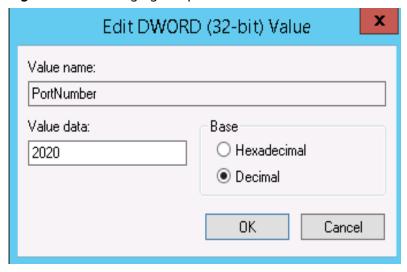
OK

Cancel

Figure 16-14 Changing the port number to 2020

- 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 16-15 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 16-16 Selecting Services



- 8. In the **Services** window, restart **Remote Desktop Services** or the ECS.
- 9. Use "IP address:Port" to remotely access the ECS.

Linux

The following procedure uses an ECS running CentOS 7.3 as an example. The default login port of a Linux ECS is 22. To change it to port 2020, for example, do as follows:

- 1. Modify the security group rule.
 - a. Log in to the management console.
 - b. Under **Computing**, click **Elastic Cloud Server**.
 - c. On the ECS list, click the name of an ECS for which you want to modify the security group rule.
 - d. On the ECS details page, click the security group in the **Security Groups** area to go to the security group details page.
 - e. On the **Inbound Rules** tab, click **Add Rule**. In the displayed dialog box, set **Protocol & Port** as follows:
 - Protocols: TCP (Custom ports)
 - Port: 2020

For 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 16-17 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 2020

16.6.7 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.

- 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.

| \sim | NIOTE |
|--------|-------|
| | |
| | |

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.

16.6.8 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:

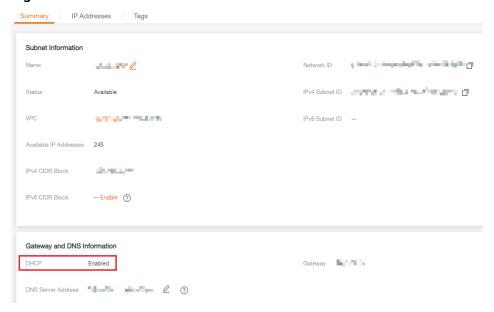
- 1. Ensure that Cloudbase-Init has been correctly configured on the image that was used to create the ECS.
 - If Cloudbase-Init has not been configured, your ECS will not allow customized configurations, and you can log in to it only by using the original image password.
 - The ECSs created using a public image have Cloudbase-Init installed by default. Therefore, 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.

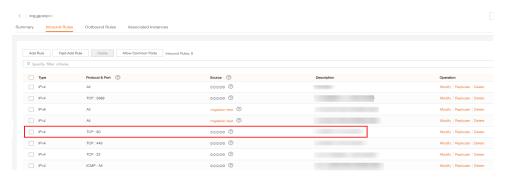
 The key used for obtaining the password must be the key used during the ECS creation.
- Ensure that DHCP is enabled in the VPC to which the ECS belongs.On the management console, check whether DHCP has been enabled in the target subnet.

Figure 16-18 DHCP



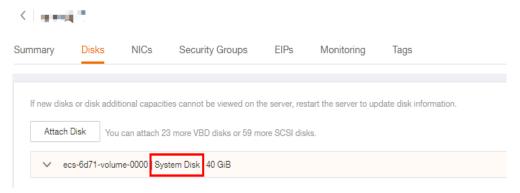
- 4. Ensure that the ECS has an EIP bound.
- 5. Ensure that traffic to and from port 80 is allowed in security group rules.

Figure 16-19 Security group rules for port 80



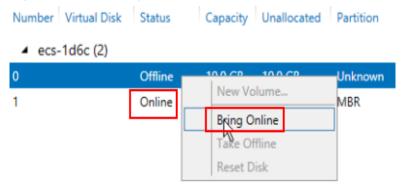
- 6. Check Cloudbase-Init logs to identify the cause.
 - a. Stop the affected ECS and detach the system disk from it.

Figure 16-20 Detaching the system disk



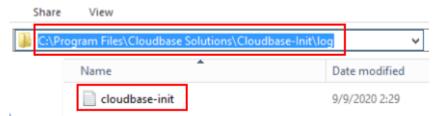
- b. Use a public image to create a temporary Windows ECS and attach the system disk detached in **6.a** to the ECS.
- c. Log in to the temporary ECS, open the **Server Manager** page, choose **File and Storage Services** > **Volumes** > **Disks**, right-click the offline disk, and choose **Online** from the shortcut menu.

Figure 16-21 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 16-22 cloudbase-init



16.6.9 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 16-4**.

Table 16-4 Browser version requirements

| Browser | Version |
|-------------------|-----------|
| Google Chrome | 31.0-75.0 |
| Mozilla Firefox | 27.0-62.0 |
| Internet Explorer | 10.0-11.0 |

16.6.10 How Can I Log In to an ECS After Its System Disk Is Exchanged with That Attached to Another ECS Running the Same OS?

Symptom

Two ECSs run the same OS, for example, both run Windows or Linux. The system disks attached to the two ECSs are exchanged offline. After the exchanging, the login keys of the ECSs may change. In such a case, how can I log in to the ECSs?

□ NOTE

Before stopping an ECS for disk detachment, release the IP address assigned to the ECS using DHCP so that ECS can correctly obtain an IP address later. To do so, perform the following operations:

- 1. Log in to the Windows ECS.
- 2. Run the following command to release the IP address:

ipconfig /release

This operation will interrupt network connections and affect the use of the ECS. After the ECS is restarted, network connections will automatically recover.

Windows

For example, there are two Windows ECSs with parameters configured in **Table 16-5**.

Table 16-5 Parameter configurations

| ECS | System Disk | Key Pair |
|--------|-------------|------------|
| ecs_01 | vol_01 | Keypair_01 |
| ecs_02 | vol_02 | Keypair_02 |

System disk vol_01 is detached from ecs_01 offline and then attached to ecs_02 as the system disk. How can I log in to ecs_02?

The random password for logging in to ecs_02 must be resolved again. The procedure is as follows:

1. Delete the initial password for logging in to ecs_02.

Locate the row containing ecs_02, click **More** in the **Operation** column, and select **Delete Password** from the drop-down list. Then, click **Delete**.

ecs_02 must be in **Stopped** state.

2. Start ecs 02.

Locate the row containing ecs_02, click **More** in the **Operation** column, and select **Start** from the drop-down list. Then, in the **Start ECS** dialog box, click **OK**.

- 3. Obtain the password for logging in to ecs_02.
 - a. Locate the row containing ecs_02, click **More** in the **Operation** column, and select **Get Password** from the drop-down list.
 - b. Click **Select File** and upload private key file **Keypair_02** of ecs_02.
 - c. Click **Get Password** to obtain a new random password.
- 4. Use the random password obtained in step 3 to log in to ecs_02 with the system disk replaced.

Linux

For example, there are two Linux ECSs with parameters configured in Table 16-6.

Table 16-6 Parameter configurations

| ECS | System Disk | Key Pair |
|--------|-------------|------------|
| ecs_01 | vol_01 | Keypair_01 |
| ecs_02 | vol_02 | Keypair_02 |

System disk vol_01 is detached from ecs_01 offline and then attached to ecs_02 as the system disk. How can I log in to ecs_02?

Use either of the following methods to log in to ecs_02:

- Use private key file **Keypair_01** of ecs_01.
- Use private key file **Keypair_02** of ecs_02.

16.6.11 Why Does the System Display a Message Indicating that the Password for Logging In to a Windows ECS Cannot Be Obtained?

Symptom

Password authentication is required to log in to a Windows ECS. Therefore, you require a key file to obtain the initial password for logging in to the ECS. However, after you click **Get Password** (see **Obtaining the Password for Logging In to a Windows ECS**), the system displays a message indicating that the password could not be obtained, resulting in an ECS login failure.

Possible Causes

Possible causes vary depending on the image used to create the Windows ECS.

- Cause 1: The image used to create the Windows ECS is a private image, on which Cloudbase-Init has not been installed.
- Cause 2: Cloudbase-Init has been installed on the image, but the key pair has not been obtained when the Windows ECS was created.

Solution

- If the issue is a result of cause 1, proceed as follows:
 - If a private image is created without Cloudbase-Init installed, the ECS configuration cannot be customized. As a result, you can log in to the ECS only using the original image password.
 - The original image password is the OS password configured when the private image was created.
- 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. Click **More** in the **Operation** column and select **Get Password** to check whether the password can be obtained.
 - If you can obtain the password, no further action is required.
 - If you cannot obtain the password, contact customer service.

16.6.12 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 16-23**.

roo 9:57 co ces. Iso-58 Apr 29 roo 199 Apr 29 9:57 corooroo483 Apr 29 coroo rooApr 29 roo9:58 Apr 29 roo9:58 record. Apr 29 roo roo9:57 e rescope roo roo Ьi 1#]#]# Eroo [roo]# 1# [roo -1# [roo

Figure 16-23 Garbled characters on the VNC-based login page

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.

16.6.13 Why Are Characters Entered Through VNC Still Incorrect After the Keyboard Language Is Switched?

During ECS login using VNC, changing the remote login keyboard language ensures only that characters entered in the VNC window for an ECS are correctly mapped. It does not change the output language of the ECS OS. If characters are not entered correctly, you must log in to the ECS and configure its keyboard output language.

• For Linux ECSs, perform the following:

Run the following command to load the keyboard mapping file:

loadkeys keymapfile

The *keymapfile* parameter indicates the name of the file containing the mappings between the keys and displayed characters.

For example, if the name of a French keyboard mapping file is **fr**, run the **loadkeys fr** command.

• For Windows ECSs, switch the input method or open the soft keyboard to enter characters. To open the soft keyboard, perform the following:

- 1. Click the **Option Menu** icon.
- 2. Select Soft Keyboard.
- 3. Select a keyboard layout.

16.6.14 Why Cannot I Use the French Keyboard to Enter Characters When I Log In to an ECS Using VNC?

 For ECSs running Windows, the following characters cannot be input normally: ~ `

To input these characters, use either of the following methods:

Method 1: Click **AltGr** on the VNC page and then press the suitable key on the keyboard twice.

Method 2: Press AltGr+~(`)+Space.

• For ECSs running Windows and Linux, the following characters cannot be input normally: #, {, [, |, \, ^, @,], }, ¤, €

To input these characters, do as follows:

Click **AltGr** on the VNC page and then press the suitable key on the keyboard.

• No L'accent circonflexe or Le tréma can be input on characters A, E, I, O, and U using the physical keyboard.

To input these characters, do as follows:

Click ^: on the VNC page and then press the suitable key on the keyboard.

No character can be input using the French keypad.

To input these characters, do as follows:

Use the French keyboard.

16.6.15 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.

16.6.16 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.

16.6.17 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.

16.6.18 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 16-24**.

Figure 16-24 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.

16.6.19 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.

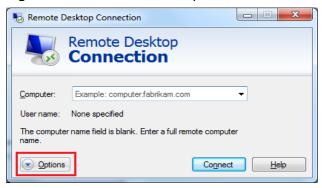
1. Start your local PC.

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.

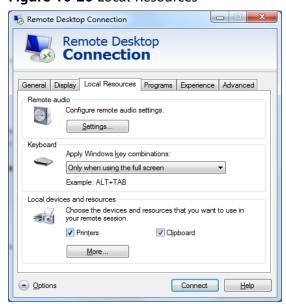
The **Remote Desktop Connection** window is displayed.

Figure 16-25 Remote Desktop Connection



4. Click **Options** in the lower left corner and click the **Local Resources** tab.

Figure 16-26 Local Resources



5. In the **Remote audio** pane, click **Settings**.

Remote Desktop
Connection

Remote Desktop
Connection

Remote audio playback

Play on this computer

Do not play
Play on remote computer

Remote audio recording

Remote audio recording

Record from this computer

Do not record

Figure 16-27 Setting remote audio playback

6. In the Remote audio playback pane, select Play on this computer.

16.6.20 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.

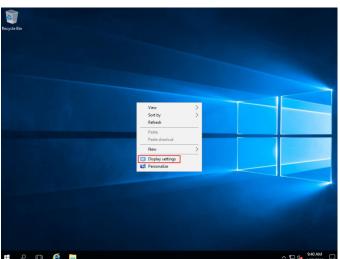


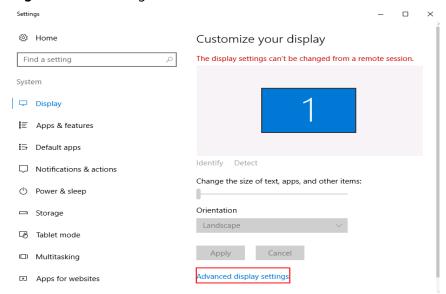
Figure 16-28 Display settings

On the Settings page, click the Display tab and then Advanced display settings.

◯ NOTE

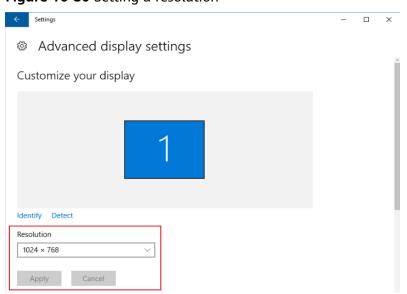
If the remote desktop is not fully displayed, set Change the size of text, apps, and other items to 100%.

Figure 16-29 Settings



4. In the **Resolution** drop-down list, select the desired resolution.

Figure 16-30 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**.

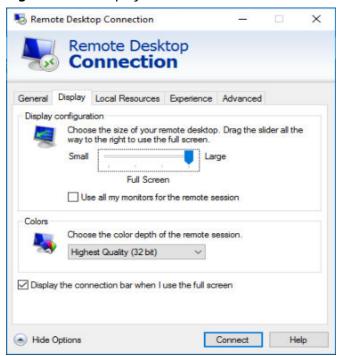
- 2. In the **Search programs and files** text box, enter **mstsc**.
- In the Remote Desktop Connection window, click Show Options in the lower left corner.

Figure 16-31 Remote Desktop Connection



4. Click the **Display** tab. Then, in the **Display configuration** pane, set the resolution.

Figure 16-32 Display



5. Use MSTSC to log in to the ECS.

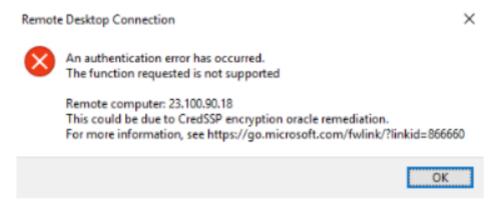
16.6.21 Why Does an Authentication Failure Occurs After I Attempt to Remotely Log In to a Windows ECS?

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 16-33, 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 16-33 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 navigation pane on the left, 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).

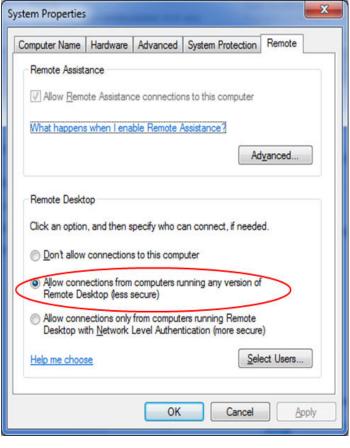


Figure 16-34 Remote settings

5. Click OK.

16.6.22 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 16-35 Cannot connect to the remote computer



Possible Causes

- Port 3389 of the security group on the ECS is disabled. For details, see
 Checking Port Configuration on the ECS.
- The firewall on the ECS is disabled. For details, see **Checking Whether the Firewall Is Correctly Configured**.
- The remote desktop connection is not correctly configured. For details, see
 Checking Remote Desktop Connection Settings.
- Remote Desktop Services are not started. For solution, see Checking Remote Desktop Services.
- Remote Desktop Session Host is not correctly configured. For details, see
 Checking Remote Desktop Session Host Configuration.

Checking Port Configuration on the ECS

Check whether port 3389 (used by default) on the ECS is accessible.

Ensure that port 3389 has been added in the inbound rule.

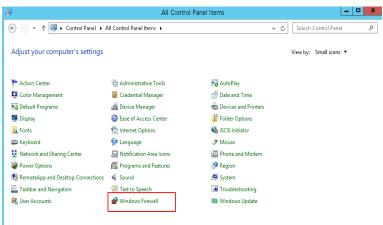
On the page providing details about the ECS, click the **Security Groups** tab and view port 3389 in the inbound rule of the security group.

Checking Whether the Firewall Is Correctly Configured

Check whether the firewall is enabled on the ECS.

- 1. Log in to the ECS using VNC available on the management console.
- 2. Click the Windows icon in the lower left corner of the desktop and choose **Control Panel** > **Windows Firewall**.

Figure 16-36 Windows Firewall



3. Click Turn Windows Firewall on or off.

View and set the firewall status.

(←) → ↑ 💣 ➤ Control Panel ➤ All Control Panel Items ➤ Windows Firewall ✓ ♂ Search Control Panel Help protect your PC with Windows Firewall Control Panel Home $Windows \ Firewall \ can help \ prevent \ hackers \ or \ malicious \ software from \ gaining \ access \ to \ your \ PC \ through \ the \ Internet \ or \ a \ network.$ Allow an app or feature through Windows Firewall Private networks Change notification settings Not connected (v) Turn Windows Firewall on or off Guest or public networks Connected (A) Restore defaults Networks in public places such as airports or coffee shops Advanced settings Troubleshoot my network Incoming connections: Block all connections to apps that are not on the list of allowed apps Active public networks: Metwork Notification state: Do not notify me when Windows Firewall blocks a Action Center Network and Sharing Cente

Figure 16-37 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

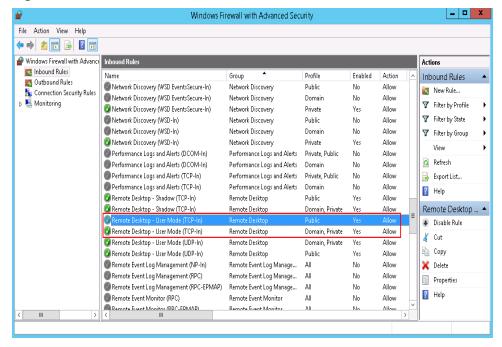


Figure 16-38 Inbound Rules

If the port configured in the inbound rule of the firewall is different from that configured on the remote server, the remote login will fail. If this occurs, add the port configured on the remote server in the inbound rule of the firewall.

□ NOTE

The default port is 3389. If you use another port, add that port in the inbound rule of the firewall.

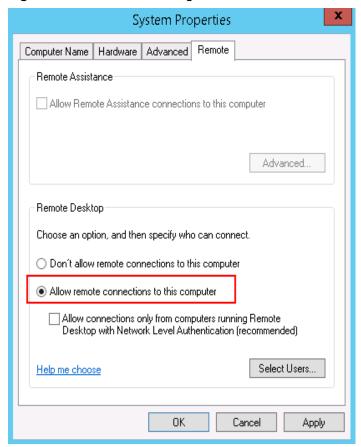
After you perform the preceding operations, try to remotely log in to the ECS again.

Checking Remote Desktop Connection Settings

Modify the remote desktop connection settings on the Windows ECS:

- 1. Log in to the ECS.
- 2. Click **Start** in the lower left corner, right-click **Computer**, and choose **Properties** from the shortcut menu.
- 3. In the left navigation pane, choose **Remote settings**.
- Click the Remote tab. In the Remote Desktop pane, select Allow connections from computers running any version of Remote Desktop (less secure).

Figure 16-39 Remote settings



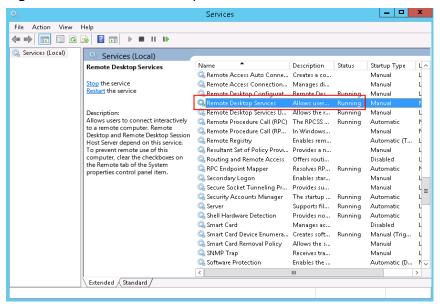
5. Click **OK**.

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.

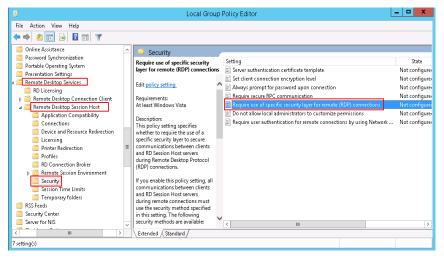
Figure 16-40 Remote Desktop Services



Checking Remote Desktop Session Host Configuration

- 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 16-41 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.

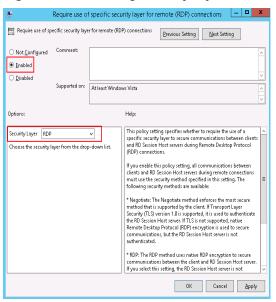


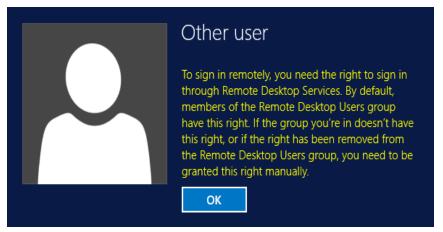
Figure 16-42 Setting security layer to RDP

16.6.23 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 16-43 Remote login right missing.

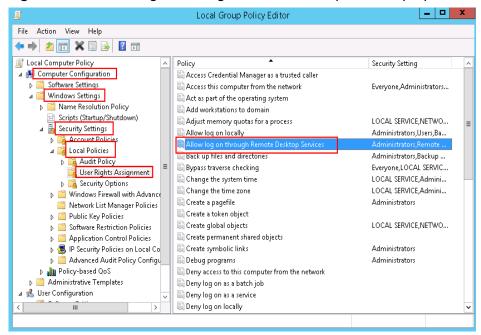


Solution

- 1. Open the **cmd** window and enter **gpedit.msc**.
- 2. Click **OK** to start Local Group Policy Editor.
- Choose Computer Configuration > Windows Settings > Security Settings > Local Policies > User Rights Assignment.

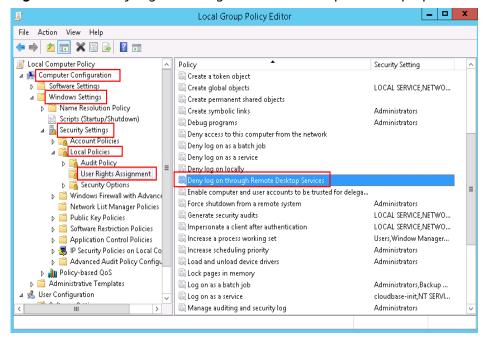
Locate and double-click Allow log on through Remote Desktop
 Services. Ensure that Administrators and Remote Desktop Users have been added.





b. Locate and double-click **Deny log on through Remote Desktop Services**. If the administrator account exists, delete it.

Figure 16-45 Deny log on through Remote Desktop Services properties



16.6.24 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 16-46 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.

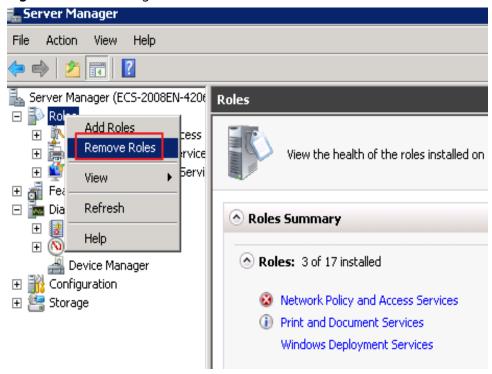
Precautions

- 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

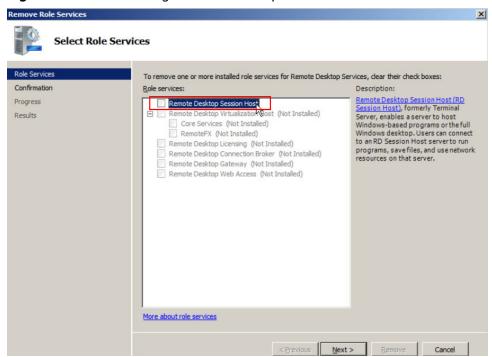
- 1. Log in to the Windows ECS using VNC available on the management console.
- 2. Open **Server Manager**, right-click **Remote Desktop Services** under **Roles**, and choose **Remove Roles** from the shortcut menu.

Figure 16-47 Deleting roles



3. In the displayed dialog box, deselect **Remote Desktop Session Host** and keep clicking **Next** till you finish the operation.

Figure 16-48 Deselecting Remote Desktop Session Host

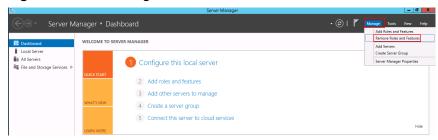


- 4. Click Delete.
- 5. Restart the ECS.

Windows Server 2012

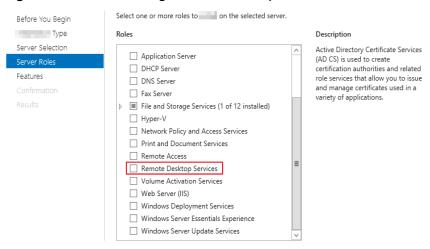
- 1. Log in to the Windows ECS using VNC available on the management console.
- Open Server Manager, choose Manage > Remove Roles and Features, and click Next.

Figure 16-49 Deleting roles and features



- 3. Select the destination server and click **Next**.
- 4. Deselect **Remote Desktop Services**.

Figure 16-50 Deselecting Remote Desktop Services



- 5. Click **Delete**.
- 6. Restart the ECS.

16.6.25 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.

Figure 16-51 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?**

MOTE

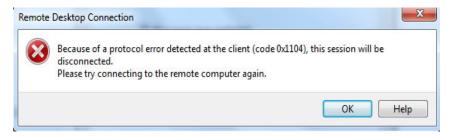
This method will deteriorate the disk I/O performance. Therefore, use this method only when necessary.

16.6.26 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 16-52 Protocol error (code: 0x1104)



Possible Causes

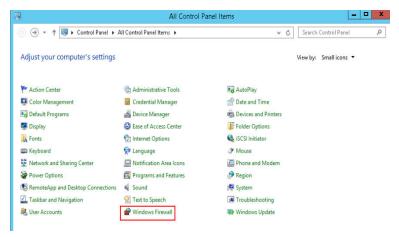
- Port 3389 of the security group on the ECS is disabled.
- The firewall on the ECS is disabled.
- Port 3389 on the ECS is used by other processes.
- The Remote Desktop Session Host is incorrectly configured.

Solution

Step 1 Check security group settings.

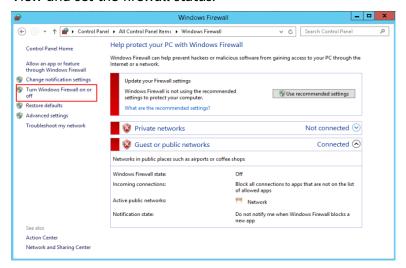
Check whether port 3389 is allowed in inbound direction. If it is allowed, go to **Step 2**.

- 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 16-53 Checking port 3389

As shown in Figure 16-53, 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 16-54 Checking System process

```
Administrator: Command Prompt

Snapshot of HTTP service state (Server Session View):

Server session ID: FF00000020000001

Version: 1.0

State: Active
Properties:

Max bandwidth: 4294967295

Timeouts:

Entity body timeout (secs): 120

Drain entity body timeout (secs): 120

Request queue timeout (secs): 120

Header wait timeout (secs): 120

Hinimum send rate (bytes/sec): 150

URL groups:

URL group ID: FE00000040000001

State: Active
Request queue name: Request queue is unnamed.
Properties:

Max bandwidth: inherited
Max connections: inherited
Timeouts:

Timeout values inherited
Number of registered URLs:

HTTPS://+:3389/

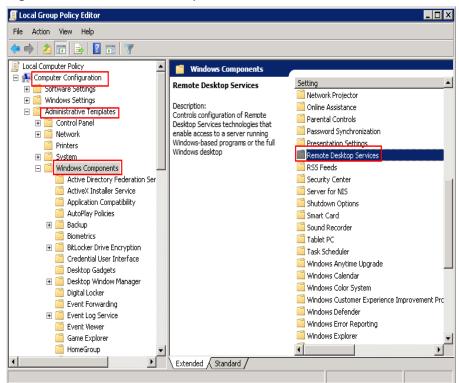
HTTPS://+:3389/

Request queues:
```

- 4. If port 3389 is used by HTTP protocols, it indicates that the port is used by IIS.
- 5. Enter http://127.0.0.1:3389 in the address box of the browser and press Enter. Check whether the website can be visited normally.
- 6. Change the port used by IIS and restart IIS.
- **Step 4** If no error occurs during the preceding steps, go to step **Step 5** to check whether error 0x1104 is caused by the configuration of Remote Desktop Session Host.
- **Step 5** Check the remote desktop session host configuration.
 - 1. Log in to the ECS using VNC.

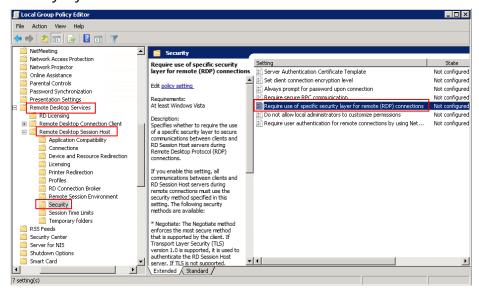
- 2. Open the **cmd** window and enter **gpedit.msc**.
- 3. Click **OK** to start Local Group Policy Editor.
- 4. Choose Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services.

Figure 16-55 Remote Desktop Services



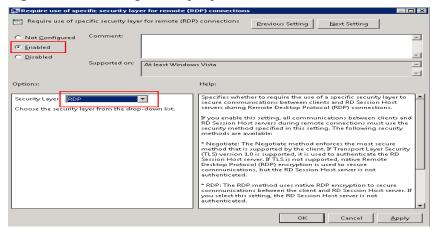
5. Remote Desktop Session Host > Security.

Figure 16-56 Remote (RDP) Connection requires the use of the specified security layer



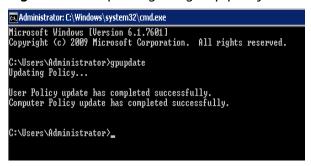
6. Set Require use of specific security layer for remote (RDP) connections to Enabled and Security layer to RDP.

Figure 16-57 Setting security layer



- 7. Click OK.
- 8. After the configuration is complete, open the **cmd** window.
- Run the following command to update the group policy: gpupdate

Figure 16-58 Updating the group policy



----End

16.6.27 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

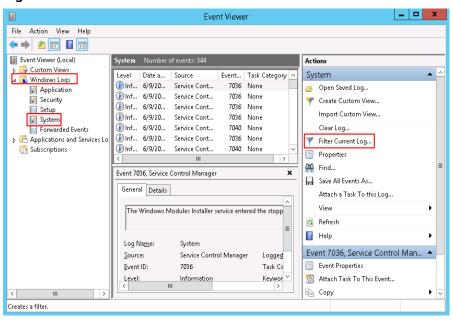
- 1. System resources are insufficient or unavailable.
- 2. The services cannot be started.

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 16-59 Event viewer



3. In the **Event Level** pane, select event levels.

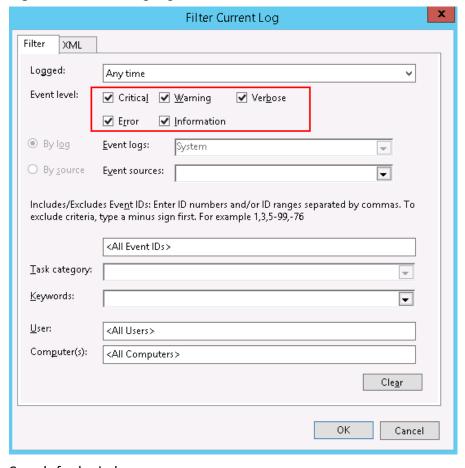


Figure 16-60 Filtering logs

- 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.

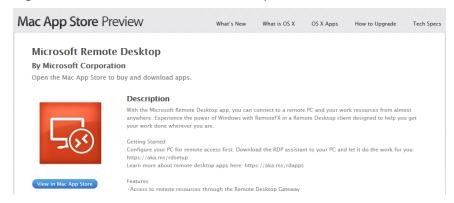
----End

16.6.28 Why Does the System Display Invalid Certificate or Associated Chain When I Log In to a Windows ECS from a Mac?

Symptom

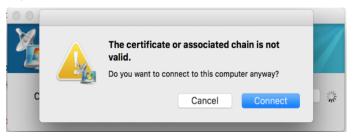
When you use Microsoft Remote Desktop for Mac to remotely access a Windows ECS, the system displays invalid certificate or associated chain.

Figure 16-61 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 16-62 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 16-63 Preferences setting



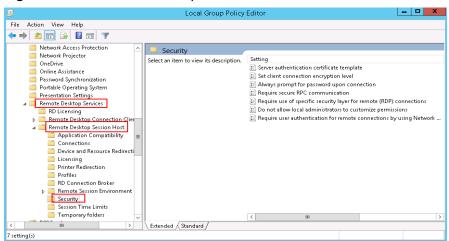
2. Select **Security** and modify the parameter settings according the following figure.

Figure 16-64 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.
- In the left navigation pane, choose Computer Configuration >
 Administrative Templates > Windows Components > Remote Desktop
 Services > Remote Desktop Session Host > Security.

Figure 16-65 Remote Desktop Session Host



- 8. Modify the following parameters as prompted:
 - Enable Require use of specific security layer for remote (RDP) connections.

Require use of specific security layer for remote (RDP) connections Require use of specific security layer for remote (RDP) connections Previous Setting O Not <u>C</u>onfigured Comment: • Enabled O <u>D</u>isabled Supported on: At least Windows Vista Help: Options: This policy setting specifies whether to require the use of a Security Layer RDP specific security layer to secure communications between clients and RD Session Host servers during Remote Desktop Protocol (RDP) connections. Choose the security layer from the drop-down list. If you enable this policy setting, all communications between clients and RD Session Host servers during remote connection must use the security method specified in this setting. The following security methods are available: * Negotiate: The Negotiate method enforces the most secure method that is supported by the client. If Transport Layer Security (TLS) version 1.0 is supported, it is used to authenticate the RD Session Host server. If TLS is not supported, native Remote Desktop Protocol (RDP) encryption is used to secure communications, but the RD Session Host server is not authenticated. RDP: The RDP method uses native RDP encryption to secure communications between the client and RD Session Host server. If you select this setting, the RD Session Host server is not OK Cancel

Figure 16-66 Require use of specific security layer for remote (RDP) connections

 Disable Require user authentication for remote connections by using Network Level Authentication.

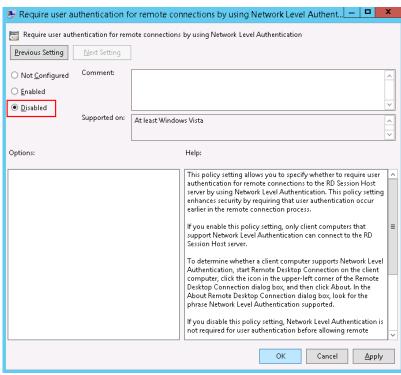


Figure 16-67 Remote connection authentication

9. Close the group policy editor and restart the ECS.

16.6.29 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

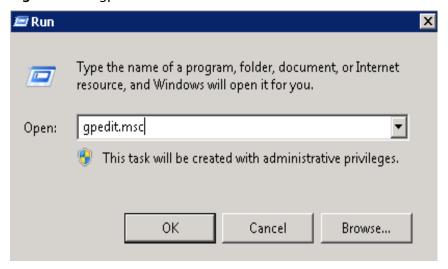
Step 3: Set the Credentials of the Local Server

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 16-68 gpedit.msc



Choose Computer Configuration > Windows Settings > Security Settings >
 Local Policies > Security Options and locate Network access: Sharing and
 security model for local accounts.

Local Group Policy Editor Local Computer Policy

Computer Configuration

Software Settings

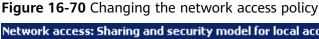
Windows Settings

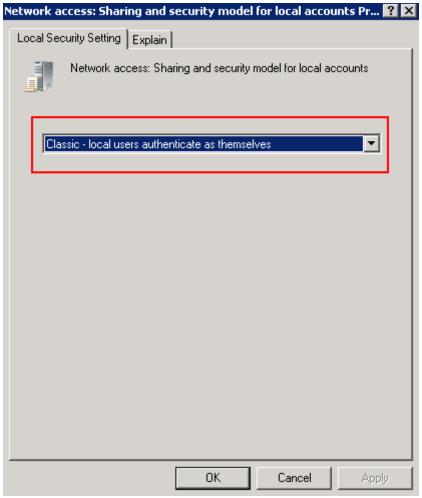
Name Resolution Policy Security Setting Network access: Do not allow anonymous enumeration of SAM ac... Enabled Network access: Do not allow anonymous enumeration of SAM ac... Disabled Network access: Do not allow storage of passwords and credenti. Network access: Let Everyone permissions apply to anonymous u... Disabled Scripts (Startup/Shutdown)

Security Settings Network access: Named Pipes that can be accessed anonymously Network access: Remotely accessible registry paths Account Policies Network access: Remotely accessible registry paths and sub-paths System\CurrentControlSe... Local Policies Network access: Restrict anonymous access to Named Pipes and ... Enabled Audit Policy Not Defined Security Options Network security: Allow LocalSystem NULL session fallback Network List Manager Policies Not Defined Public Key Policies
Software Restriction Policies Network Security: Allow PKU2U authentication requests to this co... Not Defined Network security: Configure encryption types allowed for Kerberos Application Control Policies Network security: Do not store LAN Manager hash value on next ... Enabled Network security: Force logoff when logon hours expire Network security: LAN Manager authentication level Disabled Not Defined Network security: LDAP client signing requirements Negotiate signing Network security: Minimum session security for NTLM SSP based (... Require 128-bit encryption with the security: Minimum session security for NTLM SSP based (... Require 128-bit encryption with the security: Minimum session security for NTLM SSP based (... Require 128-bit encryption Network security: Restrict NTLM: Add remote server exceptions f... Not Defined
Network security: Restrict NTLM: Add server exceptions in this d... Not Defined Administrative Templates Network security: Restrict NTLM: Audit Incoming NTLM Traffic Not Defined Network security: Restrict NTLM: Audit NTLM authentication in thi... Not Defined
Notwork security: Restrict NTLM: Incoming NTLM traffic Not Defined Network security: Restrict NTLM: NTLM authentication in this dom... Not Defined Network security: Restrict NTLM: Outdoing NTLM traffic to remot.... Not Defined

Figure 16-69 Locating the network access policy

4. Select Classic - local users authenticate as themselves and click OK.

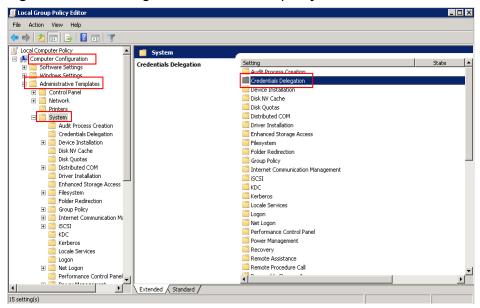




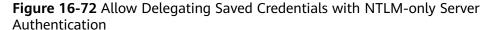
Step 2: Modify Credentials Delegation

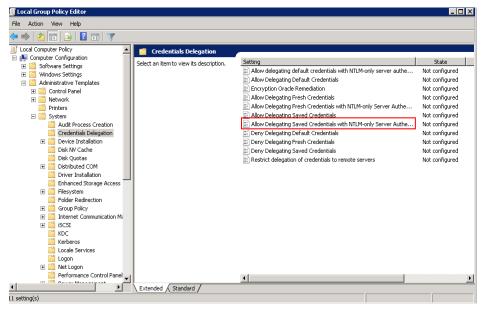
- 1. Log in to the ECS using VNC on the management console.
- 2. Choose **Start** > **Run**. In the **Run** dialog box, enter **gpedit.msc** and click **OK** to start **Local Group Policy Editor**.
- 3. Choose Computer Configuration > Administrative Templates > System and locate Credentials Delegation.

Figure 16-71 Locating the network access policy



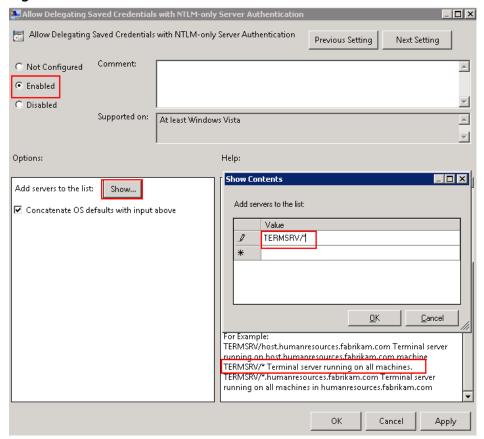
 Double-click Allow Delegating Saved Credentials with NTLM-only Server Authentication and click OK.





5. Select **Enabled** and enter **TERMSRV/*** in the **Show Contents** text box. **TERMSRV/*** indicates the terminal server running on all computers.

Figure 16-73 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.

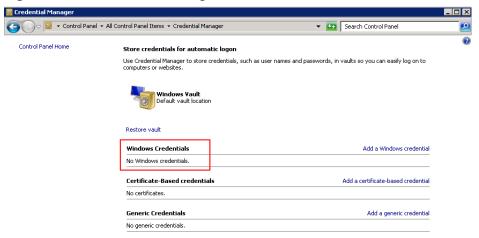
Figure 16-74 Updating the group policy

```
■ C:\Windows\system32\gpupdate.exe
Updating Policy...
User Policy update has completed successfully.
-
```

Step 3: Set the Credentials of the Local Server

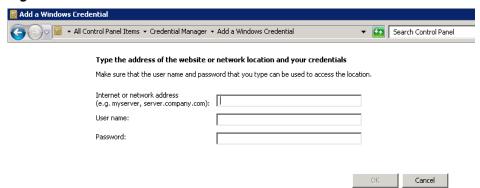
 Open the control panel on the local server and choose Credential Manager > Windows Credentials.

Figure 16-75 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 16-76 Add a Windows Credential



Step 4: Disable Password Protected Sharing

- 1. Log in to the ECS.
- 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**.

Advanced sharing settings

Private

Guest or Public (current profile)

All Networks

Public folder sharing

When Public folder sharing is on, people on the network, including homegroup members, can access files in the Public folder sharing (people logged on to this computer can still access these folders)

Media streaming

When media streaming is on, people and devices on the network can access pictures, music, and vides on this computer. This computer can also find media on the network.

Choose media streaming options...

Password protected sharing

When password protected sharing

When password protected sharing

Only people who have a user account and password on this computer can access shared files, printers attached to this computer, and the Public folders. To give other people access, you must turn off password protected sharing.

O Turn on password protected sharing

Turn off password protected sharing

Turn off password protected sharing

Figure 16-77 Turn off password protected sharing

4. Click Save changes.

16.6.30 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 ECS using your phone's mobile data, your local network is abnormal. Restart your local network (for example, restart the router).

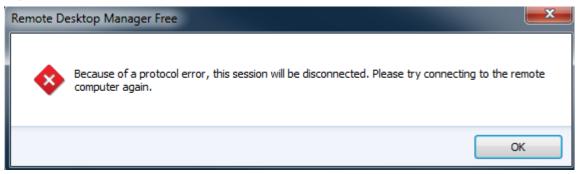
If the fault still persists, record the resource details and fault occurred time, and contact customer service for technical support.

16.6.31 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 16-78 Protocol error



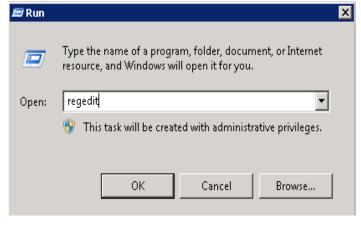
Possible Causes

The registry subkey Certificate is damaged.

Solution

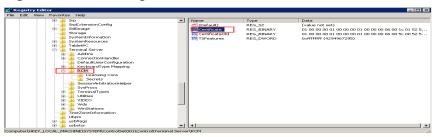
1. In the **Run** dialog box, enter **regedit** and click **OK** to open the registry editor.

Figure 16-79 Opening the registry editor



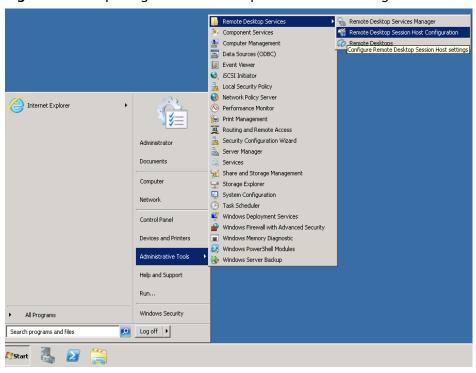
 Choose HKEY_LOCAL_MACHINE > SYSTEM > ControlSet001 > Control > Terminal Server > RCM. 3. Delete Certificate.

Figure 16-80 Deleting Certificate



- 4. Restart the ECS.
- 5. Choose Start > Administrative Tools > Remote Desktop Services > Remote Desktop Session Host Configuration.

Figure 16-81 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**.

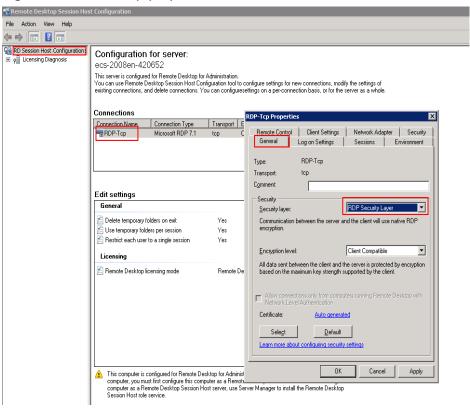


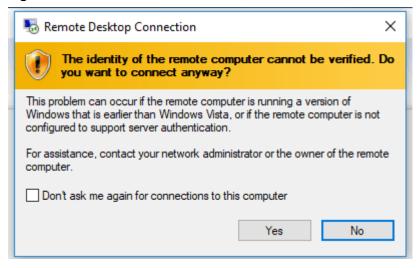
Figure 16-82 RDP-Tcp properties

16.6.32 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 16-83 Protocol error



Possible Causes

Security software installed on the ECS prevents logins from unknown IP addresses.

Solution

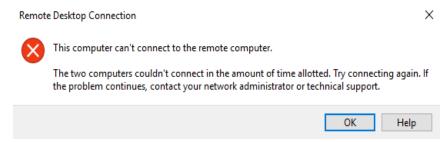
- Uninstall the security software.
- Open the security software and enable the default login mode.

16.6.33 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 16-84 Error message



Solution

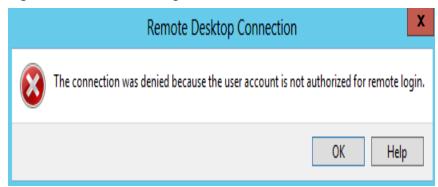
- 1. On the local computer, click on the **Start** icon, type **cmd** into the box, and run the command as an administrator.
- 2. Run the **netsh winsock reset** command.
- 3. Restart the local computer as prompted and reconnect to the ECS.

16.6.34 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 16-85 Error message



Possible Causes

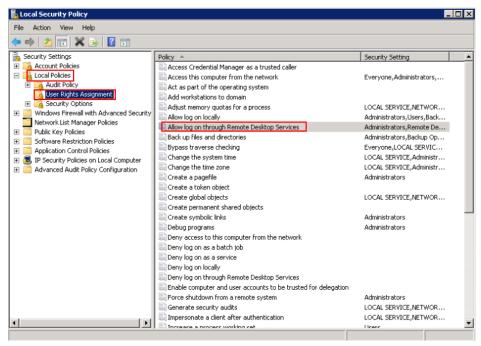
The remote desktop connection permissions have been incorrectly configured.

Solution

Step 1 Check remote desktop permissions on the ECS.

- 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 16-86 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 16-87 Allow log on through Remote Desktop Services properties

Step 2 Check the target user group.

- 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**.

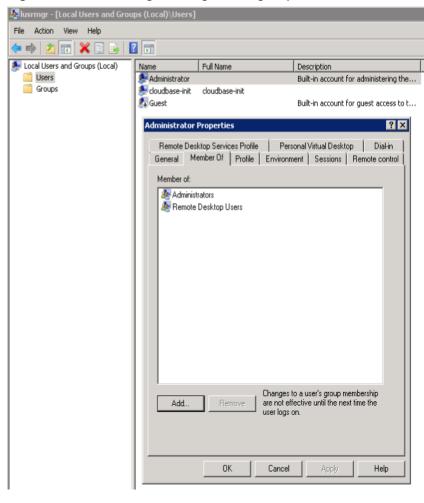


Figure 16-88 Checking the target user group

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.

##Remote Desktop Session Host Configuration
File Action View Help
| Configuration for server:
ecs_2008en-420652
This server is configured for Remote Desktop for Administration.
You can use Remote Desktop Session Host Configure settings for new connections, modify the settings of existing connections, and delete connections. You can use Remote Desktop Session Host Configure settings for new connections, modify the settings of existing connections, and delete connections. You can configure settings for new connections, modify the settings of existing connections, and delete connections. You can configure settings for new connections, modify the settings of existing connections, and delete connections. You can configure settings for new connections, modify the settings of existing connections.

Connections

Figure 16-89 Security

- 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

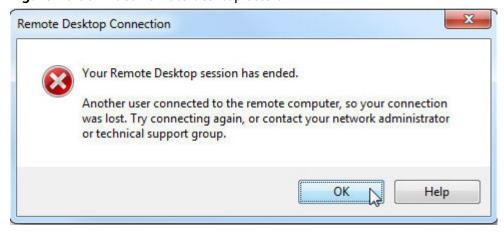
----End

16.6.35 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.

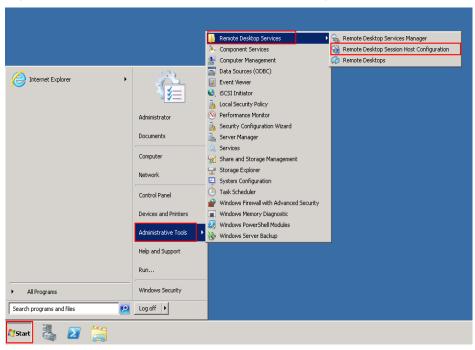
Figure 16-90 Ended remote desktop session



Windows Server 2008

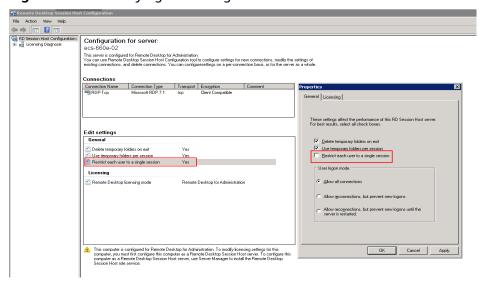
1. Choose Start > Administrative Tools > Remote Desktop Services > Remote Desktop Session Host Configuration.

Figure 16-91 Remote Desktop Session Host Configuration



 Double-click Restrict each user to a single session and deselect Restrict each user to a single session, and click OK.

Figure 16-92 Modifying the configuration

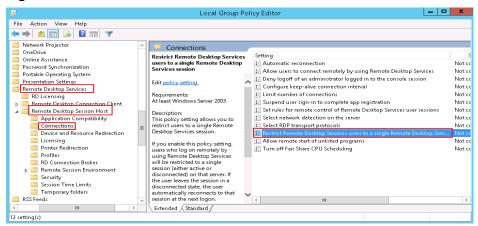


Windows Server 2012

1. Choose **Start** > **Run**. In the **Run** dialog box, enter **gpedit.msc** and click **OK** to start Local Group Policy Editor.

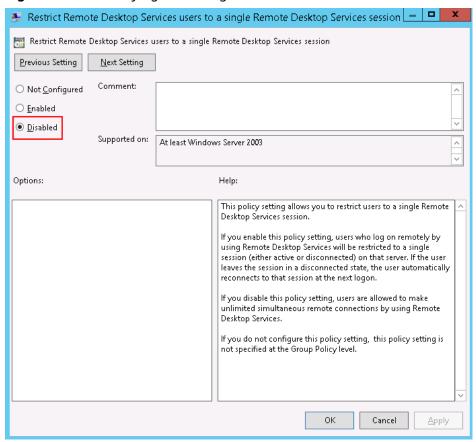
 Choose Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Connections.

Figure 16-93 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 16-94 Modifying the configuration



4. Run **gpupdate/force** to update the group policy.

16.6.36 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.

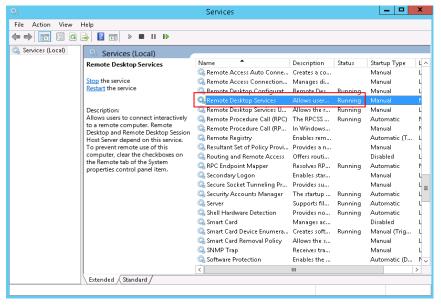


Figure 16-95 Remote Desktop Services

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.

16.6.37 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 16-96 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:

- 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 16-97 Entering the kernel editing mode

```
CentOS Linux (3.10.0-693.11.1.el7.x86_64) 7 (Core)
CentOS Linux (3.10.0-693.el7.x86_64) 7 (Core)
CentOS Linux (0-rescue-fd73dfefb69a4e4183fd3f86ccc29526) 7 (Core)
```


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 16-98 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
Initrd1b /boot/initramfs=3.10.0-327.b2.59.83.h1b2.x8b_b4.1mg
```

Figure 16-99 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 16-100 System log

```
Aug 6 18:88:89 sy login: pam_succeed_if(login:auth): requirement "uid >= 1888" not met by user "root"
Aug 6 18:88:11 sy login: PAILED_LUGIN 1 PRUM tty1 FUN root, Authentication Failure
Aug 6 18:88:15 sy login: pam_unix(login:session): session opened for user root by LOGIN(uid=8)
Aug 6 18:88:15 sy login: PAMU Le is unknown
Aug 6 18:18:41 sy login: PAMU unable to dlopen(/lib/security/pam_limits.so): /lib/security/pam_limits.so: cannot open shared object file: No such file or directory
Aug 6 18:18:41 sy login: PAM adding faulty module: /lib/security/pam_limits.so
Aug 6 18:18:44 sy login: pam_unix(login:session): session opened for user root by LOGIN(uid=8)
Aug 6 18:18:44 sy login: pam_unix(login:session): session opened for user root by LOGIN(uid=8)
Aug 6 18:18:44 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 16-101 Modifying the error information

```
reguired
                        pam selinux.so open
                        pam_namespace.so
session
           required
session
           optional
                        pam_keyinit.so force revoke
           include
session
                        system-auth
session
           include
                        postlogin
-session
          optional
                        pam_ck_connector.so
# session required /lib/security/pam_limits.so
```

4. 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

16.6.38 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 16-102 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

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 16-103 Entering the kernel editing mode

```
CentOS Linux (3.10.0-693.11.1.el7.x86_64) 7 (Core)
CentOS Linux (3.10.0-693.el7.x86_64) 7 (Core)
CentOS Linux (0-rescue-fd73dfefb69a4e4183fd3f86ccc29526) 7 (Core)
```


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.

- 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 16-104 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
Initralb /boot/initramfs=3.10.0-327.b2.59.83.h1b2.x8b_b4.1mg
```

Figure 16-105 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

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.

16.6.39 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 16-106 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 | | |
|----------------|---------------------|--|
| То | Action From | |
| | | |
| 22 | ALLOW Anywhere | |
| 22 (v6) | ALLOW Anywhere (v6) | |

Try to remotely log in to the ECS again.

16.6.40 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

16.6.41 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 16-107 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-keye
```

Possible Causes

A policy that denies password-authenticated logins is enabled on the SSH server.

Solution

- Open the /etc/ssh/sshd_config file and check the following settings: vi /etc/ssh/sshd_config
- Modify the following settings:
 Change PasswordAuthentication no to PasswordAuthentication yes.
 Alternatively, delete the comment tag (#) before PasswordAuthentication
- 2. Restart SSH.

yes.

- CentOS 6
 service sshd restart
- CentOS 7
 systemctl restart sshd

16.7 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.

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 16-7**.

Table 16-7 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 the 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 Elastic Cloud Server API Reference. |
| 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. |

16.8 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:

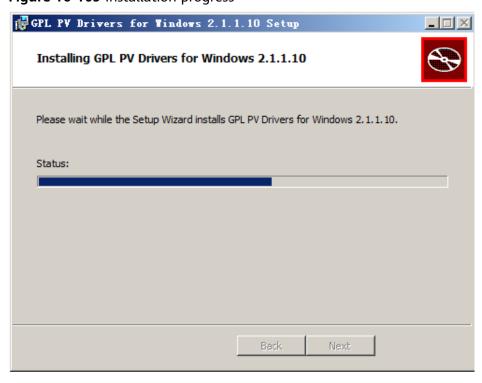
- Obtain the Tools installation package pvdriver-windows.zip.
 Contact the administrator to obtain the installation package.
- 2. Decompress the software package and double-click **setup.exe** to start the installation.



Figure 16-108 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 16-109 Installation progress



5. Click **Finish** to complete the installation.

Completed the GPL PV Drivers for Windows 2.1.1.10 Setup

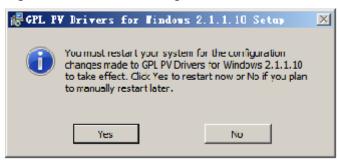
Completed the GPL PV Drivers for Windows 2.1.1.10 Setup Wizard

Click the Finish button to exit the Setup Wizard.

Figure 16-110 Installation completed

6. Restart the ECS, as shown in Figure 16-111.

Figure 16-111 Determining whether to restart the ECS



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 16-112 Proper running status of the virtualization driver



16.9 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

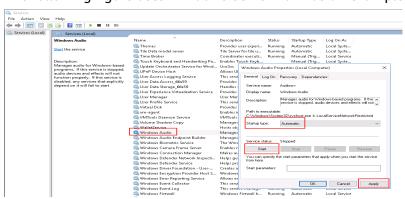
The audio function is disabled on Windows ECSs by default. As a result, audio files cannot be played on them. To enable the audio function, perform the operations described in this section.

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.

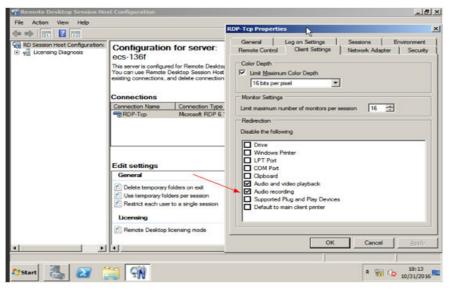
Windows Server 2008

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 16-113 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 16-114 Enabling the audio service



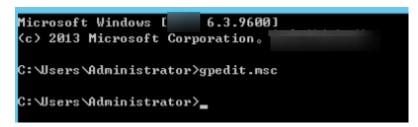
Step 4 Play an audio file to verify the service.

----End

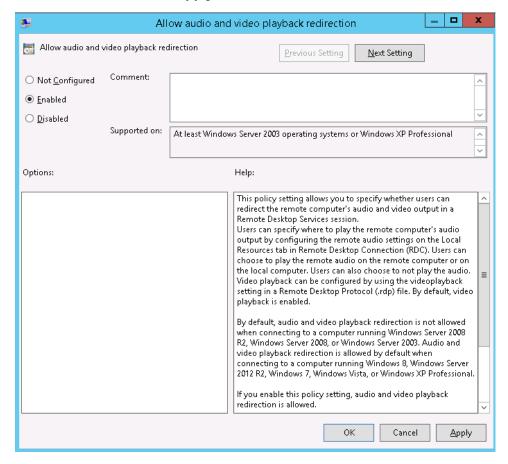
Windows Server 2012

Step 1 Start the **Run** dialog box.

Step 2 Run the **gpedit.msc** command to start **Local Group Policy Editor**.



- 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.

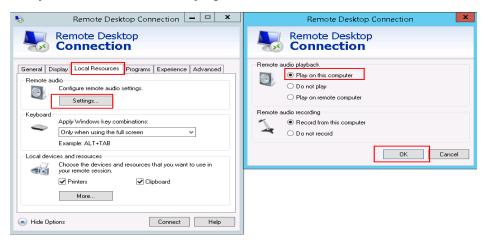
Step 5 Run the following command to update the group policy:

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.

16.10 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**.

Changing an ECS SID may lead to data loss or system damage. Therefore, 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 16-115 Downloading SIDCHG

SIDCHG 2.00 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 16-116 Risk prompt

```
PS C:\Users\Administrator> .\sidchg64-2.0o.exe /R
SID Change Utility SIDCHG64 2.0o -- Copyright Stratesave Systems 2018
Shareware - visit http://www.stratesave.com for Licence and pricing
Enter license key or trial key:
34D0g-4gUfk-vookt-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 16-117 Re-login

```
SID of this computer is being changed after which the computer will reboot

Do not log on to this computer and do not turn off the computer at this time!

At Stopping Service
```

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.

16.11 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**.

16.12 Why Is the Memory of an ECS Obtained by Running the free Command Inconsistent with the Actual Memory?

Symptom

After you create an ECS, you run the **free -m** command to view the ECS memory. The ECS memory is less than the memory configured during ECS creation.

For example:

When you are creating an ECS, the configured memory size is 4194304 KB (4096 MB). After the ECS is created, you run the **free -m** command to view its memory. The command output is as follows:

[root@localhost ~]# free -m total used free shared buff/cache available Mem: 3790 167 3474 8 147 3414 Swap: 1022 0 1022

The memory in the command output is 3790 MB, which is less than the configured 4096 MB.

Run the **dmidecode -t memory** command to check the actual memory configured for the ECS. The command output is as follows:

[root@localhost ~]# dmidecode -t memory # dmidecode 3.0 Getting SMBIOS data from sysfs. SMBIOS 2.8 present.

Handle 0x1000, DMI type 16, 23 bytes Physical Memory Array Location: Other Use: System Memory Error Correction Type: Multi-bit ECC Maximum Capacity: 4 GB Error Information Handle: Not Provided Number Of Devices: 1

Handle 0x1100, DMI type 17, 40 bytes Memory Device

Array Handle: 0x1000

Error Information Handle: Not Provided

Total Width: Unknown Data Width: Unknown Size: 4096 MB Form Factor: DIMM Set: None Locator: DIMM 0

Bank Locator: Not Specified

Type: RAM
Type Detail: Other
Speed: Unknown
Manufacturer: QEMU
Serial Number: Not Specified
Asset Tag: Not Specified
Part Number: Not Specified

Rank: Unknown

Configured Clock Speed: Unknown Minimum Voltage: Unknown Maximum Voltage: Unknown Configured Voltage: Unknown

The memory in the command output is the same as that configured during ECS creation.

Possible Causes

When the OS is started, related devices are initialized, which occupies memory. In addition, when the kernel is started, it also occupies memory. The memory occupied by kdump can be set. Unless otherwise specified, do not change the memory size occupied by kdump.

The command output of **free -m** shows the available memory of the ECS, and that of **dmidecode -t memory** shows the hardware memory.

Therefore, the memory obtained by running the **free** -**m** command is less than the memory configured for the ECS. This is a normal phenomenon.

Ⅲ NOTE

This is a normal phenomenon even for physical servers.

16.13 ECS Management

16.13.1 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:
- 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*

If there is no **HOSTNAME** in the configuration file, manually add this parameter and set it to the changed hostname.

For example:

HOSTNAME=new hostname

c. Run the following command to save and exit the configuration file:

:wq

- 3. Modify the /etc/cloud/cloud.cfg configuration file.
 - a. Run the following command to edit the configuration file:sudo vim /etc/cloud/cloud.cfg
 - b. Use either of the following methods to modify the configuration file:
 - Method 1: Change the preserve_hostname parameter value or add the preserve_hostname parameter to the configuration file.

If preserve_hostname: false is already available in the /etc/cloud/cloud.cfg configuration file, change it to preserve_hostname: true. If preserve_hostname is unavailable in the /etc/cloud/cloud.cfg configuration file, add preserve_hostname: true before cloud init modules.

If you use method 1, the changed hostname still takes effect after the ECS is stopped or restarted. However, if the ECS is used to create a private image and the image is used to create a new ECS, the hostname of the new ECS is the hostname (**new_hostname**) used by the private image, and user-defined hostnames cannot be injected using Cloud-Init.

Method 2 (recommended): Delete or comment out update_hostname.

If you use method 2, the changed hostname still takes effect after the ECS is stopped or restarted. If the ECS is used to create a private image and the image is used to create a new ECS, the changed hostname permanently takes effect, and user-defined hostnames (such as **new_new_hostname**) can be injected using Cloud-Init.

4. Run the following command to restart the ECS:

sudo reboot

5. Run the following command to check whether the hostname has been changed:

sudo hostname

If the changed hostname is displayed in the command output, the hostname has been changed and the new name permanently takes effect.

Modifying the Mapping Between the ECS Hostname and IP Address (Modifying the hosts File)

If you want to use the changed hostname as the preferred localhost and localhost.localdomain, update the mapping between the hostname and IP address after the hostname is changed and then save the configuration to the corresponding Cloud-Init configuration file so that the new hostname takes effect permanently.

The changed hostname is assumed to be **new hostname**.

- 1. Modify the /etc/hostname configuration file.
 - a. Run the following command to edit the configuration file: sudo vim /etc/hostname
 - b. Change the hostname to the new one.

- c. Run the following command to save and exit the configuration file:
- 2. Modify the /etc/sysconfig/network configuration file.
 - 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.
 - 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
- Modify the /etc/cloud/cloud.cfg configuration file.
 - 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.

16.13.2 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 16-8** lists the hostnames (obtained by running the **hostname** command) of ECSs created using different images and those displayed after the ECSs are restarted.

Table 16-8 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?

16.13.3 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.

16.13.4 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**.

- 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 16-118 shows the format of a customizable naming rule.

Figure 16-118 Format of a customizable naming rule

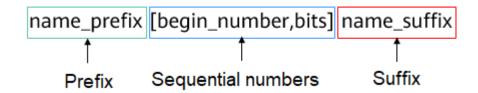


Table 16-9 describes these parameters.

Table 16-9 Parameters in a customizable naming rule

| Field | Mandato ry | Description | Example |
|----------------------|---------------|--|---------|
| name_prefix | Yes | ECS name prefix The name prefix can contain only letters, digits, underscores (_), and hyphens (-). | ecs |
| [begin_number,bit s] | 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 |

| Field | Mandato ry | Description | Example |
|--------------|---------------|---|---------|
| begin_number | No | Begin number of ECS names. The begin number ranges from 0 to 9999. The default value is 0 . | 0 |
| 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 |

Table 16-10 [begin_number,bits] parameters

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_prefix***08***name_suffix*, *name_prefix***10***name_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,

- Example 3: If you select customizable naming and enter name prefix[9,] name suffix,
 - The ECSs will be named name_prefix0009 name_suffix, name_prefix0010 name_suffix, name_prefix0011 name_suffix,
- Example 4: If you select customizable naming and enter name_prefix[,3] name_suffix,
 - The ECSs will be named name_prefix000 name_suffix, name_prefix001 name_suffix, name_prefix002 name_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,

16.13.5 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.

16.13.6 What Should I Do If Executing a Driver Installation Script Failed on an ECS Running CentOS 5?

Scenarios

After executing the script for installing the Virtio driver on an ECS running CentOS 5, users cannot determine whether the driver has been successfully installed. This section describes how to check driver installation.

Procedure

- Log in to the ECS and create a temporary directory check.
 mkdir /check
- Copy the image file to the current directory.
 cp /boot/initrd-2.6.18-308.el5.img /check/
- 3. Run the following commands to convert the file format to .gz:

cd /check

mv initrd-2.6.18-308.el5.img initrd-2.6.18-308.el5.img.gz

4. Decompress the package.

gzip -d initrd-2.6.18-308.el5.img.gz

5. Check whether the driver has been successfully installed.

```
cpio -t -F initrd-2.6.18-308.el5.img | grep virtio
```

The check process is shown in the following figure.

Figure 16-119 Checking driver installation

```
Iroot@sto-saas2pri checkl# cpio -t -F initrd-2.6.18-308.el5.img | grep virtio 14562 blocks
lib/virtio.ko
lib/virtio_pci.ko
lib/virtio_ping.ko
lib/virtio_blk.ko
lib/virtio_blk.ko
lib/virtio_net.ko
lib/virtio_net.ko
```

If the command output contains **virtio**, **virtio_blk**, **virtio_net**, and **virtio_pci**, the driver has been successfully installed.

In the preceding figure, the image is of an early version and has no virtio_scsi driver installed. As a result, SCSI disks cannot be attached to such an ECS.

16.13.7 What Should I Do If Executing a Driver Installation Script Failed When I Attempted to Modify the Specifications of a Linux ECS?

Symptom

During ECS specifications modification, a script was used to automatically install drivers on a Linux ECS.

During the installation, the following information is displayed, indicating that the script has been executed. However, the installation check failed. The possible cause is that certain drivers were not installed on the ECS.

```
...
Info:ECS modify success (mkinitrd)
Info:Check xen and virtio driver again!
...
Error:ECS modify error!
```

Figure 16-120 Successful driver installation with a failed installation check

```
cat: /boot/config-2.6.32: No such file or directory
cat: /boot/config-2.6.32: No such file or directory
cat: /boot/config-2.6.32: No such file or directory
2019-12-18 17:15:38 Inforberation for directory
2019-12-18 17:15:38 Inforberations.
2019-12-18 17:15:38 Inforberations.
2019-12-18 17:15:38 Inforberations.
2019-12-18 17:15:38 Inforberations.
2019-12-18 Inforberations.
2019-12-18 Inforberations.
2019-12-18 Inforberations.
2019-12-18 Inforberations.
2019-12-18 17:16:42 Inforberations.
2019-12-18 17:16:42 Inforberation for directory
2019-12-18 17:16:42 Inforberation for skinding
2019-12-18 17:17:56 Inforberation for skinding
2019-12-18 17:17:56 Inforberation file: /boot/instranfs-2.6.32 ing
2019-12-18 17:17:58 Inforberation file or directory
2019-12-18 17:18:01 Inforberat
```

Solution

Check whether the desired drivers have been successfully installed by following the instructions provided in "Check Whether the ECS Has Been Configured" in Manually Changing a Xen ECS to a KVM ECS (Linux).

If the drivers failed to install, the script may not match the ECS OS. In such a case, manually install the drivers on the Linux ECS by following the instructions provided in Manually Changing a Xen ECS to a KVM ECS (Linux).

Certain Linux ECSs do not have the virtio_scsi driver installed because the kernel version is too early or the kernel has been modified. Such ECSs cannot be attached with SCSI disks. However, this issue will not affect the ECS specifications modification from Xen to KVM. If the ECSs do not use SCSI disks, you can still modify their specifications without the virtio_scsi driver.

16.13.8 Why Do the Disks of a Windows ECS Go Offline After I Modify the ECS Specifications?

Scenarios

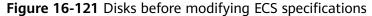
After you modify specifications of a Windows ECS, the disks may go offline. Therefore, you need to check the number of disks after you modify the specifications.

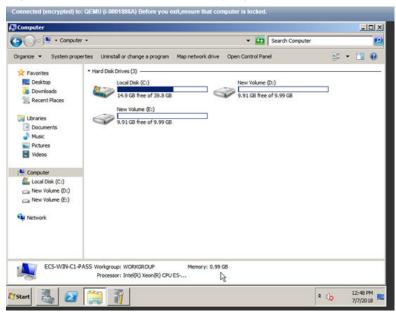
Procedure

- 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.





After the specifications are modified, check the number of disks.

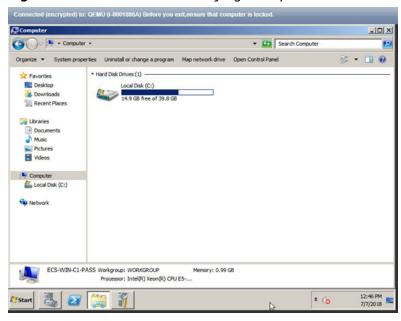


Figure 16-122 Disks after modifying ECS specifications

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.
 - In the navigation pane on the left, 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.

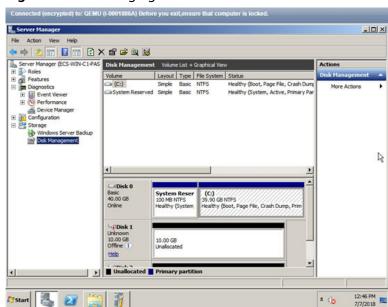
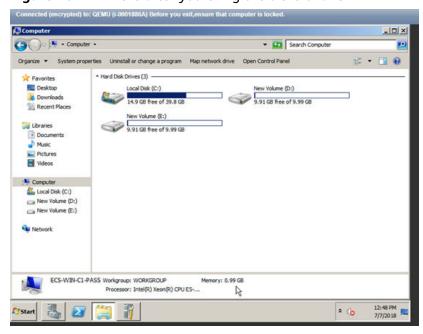


Figure 16-123 Bringing the disk online

- 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 16-124 Disks after you bring the disks online



16.13.9 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. Therefore, you need to check the disk attachment after you modify the specifications.

Procedure

- 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 16-125 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 16-125**, 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 16-126 Viewing disks attached after specifications modification

As shown in Figure 16-126, 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/vbd1 /mnt/vbd1

In the preceding command, /dev/vbd1 is the disk to be attached, and /mnt/vbd1 is the path for disk attachment.

NOTICE

Ensure that /mnt/vbd1 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 16-127 Checking the number of disks attached

As shown in **Figure 16-127**, the numbers of disks before and after specifications modifications are the same. The disks are **/dev/vda**, **/dev/vdb**, and **/dev/vdc**.

16.14 OS Management

16.14.1 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.

16.14.2 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.
- You can change your ECS OS through the management console, for example, you can upgrade CentOS 7.2 to CentOS 7.3.

16.14.3 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.

16.14.4 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/Backup** > **Change OS** in the **Operation** column.

During this process, the ECS is in **Changing OS** state.

Change OS Note the following points before you change the OS: 1. An OS change has no effect on data disks, but all data on and all snapshots created for the system disk will be lost. Back up the data before 2. SCSI disks are supported only by some OSs and any SCSI disks attached to the ECS will become unavailable if they are unsupported by the 3. The ECS will be automatically restarted after the OS change, and custom settings (such as the DNS and hostname) will be reset Hide Current Configuration IP address ECS Name 192.168.0.88 (Privat... 1 vCPUs | 1 GB |... OBS CentOS 7.8 (64-bit) 40 GB Private image Shared image ▼ C ▼ --Select OS version--Third-party image Key Pair --Select--▼ C Create Key Pair OK Cancel

Figure 16-128 OS Change

16.14.5 Will I Lose My Disk Data If I Reinstall ECS OS, Change the OS, or Change the ECS Specifications?

Table 16-11 Impact

| Item | OS Reinstallation | OS Change | Specifications Modification |
|-----------------------------|--|--|---|
| Applicat ion 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. |

| Item | OS Reinstallation | OS Change | Specifications Modification |
|----------------|---|--|---|
| 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. |
| 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. |

16.14.6 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.

16.14.7 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**.

16.14.8 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/Backup** > **Reinstall OS** in the **Operation** column.

During this process, the ECS is in **Reinstalling OS** state.

Note the following points before you reinstall the OS:

1. An OS reinstallation has no effect on data disks, but all data on and all snapshots created for the system disk will be lost. Back up the data before you continue.

2. The ECS will be automatically restarted after the OS reinstallation, and custom settings (such as the DNS and hostname) will be reset.

Hide

Current Configuration

ECS Name | P address | Specifications | Image | System ...

ecs-dc8b | 192.168.0.88 (Private |P) | 1 vCPUs | 1 GB |... OBS CentOS 7.8(64-bit) | 40 GB

System disk | Encrypted ?

Key Pair | KeyPair-ced2 | V C Create Key Pair | Cancel

Figure 16-129 Reinstall OS

16.14.9 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 GB. Otherwise, the GUI installation may fail or the ECS cannot be started after the installation.

16.14.10 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

16.14.11 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"

If the following message is displayed after the installation is complete:

Failed: python-urllibs3.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 **tgraphical.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.

16.14.12 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 and x11vnc 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 ligthdm: 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
```

- press Esc to exit editing mode.
- 4. Run the following command to save and exit the configuration file:
- 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**. 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:
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 ligthdm.

(Optional) Configuring X Server, x11vnc, and ligthdm

For GPU-accelerated ECSs, you need to configure X Server, x11vnc, and ligthdm when installing a GUI.

- 1. Remotely log in to the ECS.
- 2. Query the BusID of the GPU.

lspci | grep -i nvidia

Figure 16-130 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
```

- 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 16-131 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
```


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.

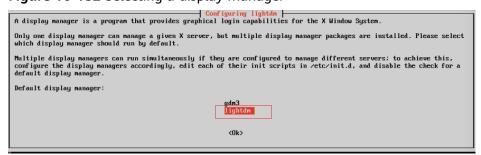
apt-get -y install x11vnc

6. Install ligthdm.

apt-get -y install lightdm

7. Select **ligthdm** as the default display manager.

Figure 16-132 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:

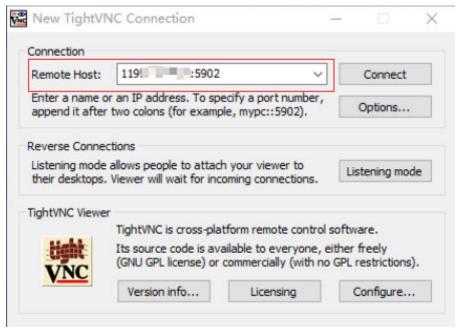
- 1. Log in to the management console.
- 2. Configure a security group for the ECS.
 - a. On the ECS list, click the name of an ECS for which you want to configure the security group rule. On the ECS details page, click **Security Groups**.
 - b. Expand the security group and in the upper right corner 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 **Step 9.c.**

3. Log in to the ECS using VNC.

The following uses TightVNC as an example.

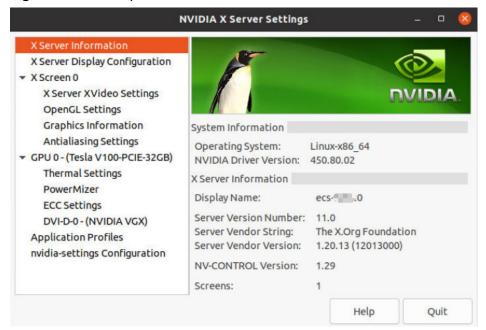
Figure 16-133 TightVNC client



- 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 16-134 Graphics card information



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**.

16.14.13 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 tasksel:

apt install tasksel

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
```

- 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.

16.14.14 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 16-135 Triggering kdump

```
Iroot@ecs-xen01 linuxl# 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.
Iroot@ecs-xen01 linuxl# echo c > /proc/sysrq-trigger
-
```


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.
 - 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:

 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 16-136 Support for the crash_kexec_post_notifiers function

- If yes, go to step 2.
- If no, use method 1.
- Add the crash_kexec_post_notifiers function to the startup configuration file.
 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 16-137 Editing the menu.lst file

c. Run the following command to restart the ECS for the modification to take effect:

reboot

16.14.15 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 16-138 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_bcall
        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 16-139 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: 5D8B666FOEA3F1E31B58FOC
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/.

- 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 16-140 Viewing the ECS driver

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 16-141 Viewing the XEN field

```
root@ubuntu:/home# cat /boot/config* | grep -i xen

CONFIG_XEN=y

CONFIG_XEN_DOMO=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

16.14.16 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 16-142 Image type



- 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.

16.14.17 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.

16.14.18 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 16-143**.

Figure 16-143 Restart failure

Possible Causes

As shown in **Figure 16-143**, 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 16-143.

- 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 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

If the fault persists, contact customer service for technical support.

16.14.19 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

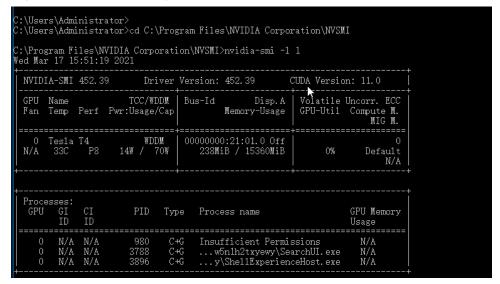
1. Log in to the GPU-accelerated ECS.

- 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 16-144 GPU usage



□ 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.

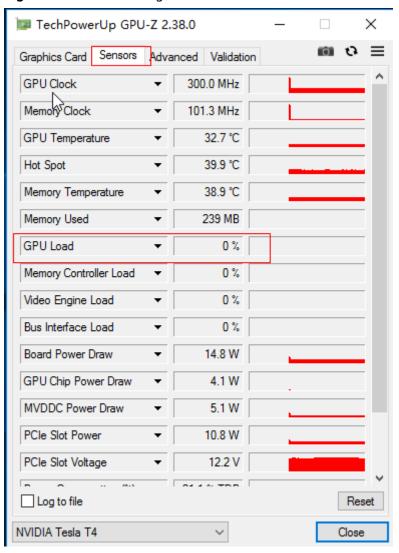


Figure 16-145 GPU usage

16.15 File Upload/Data Transfer

16.15.1 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.
- (Recommended) Local disk mapping
 - Use MSTSC to transfer files. This method does not support resumable transmission. Therefore, 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.

16.15.2 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

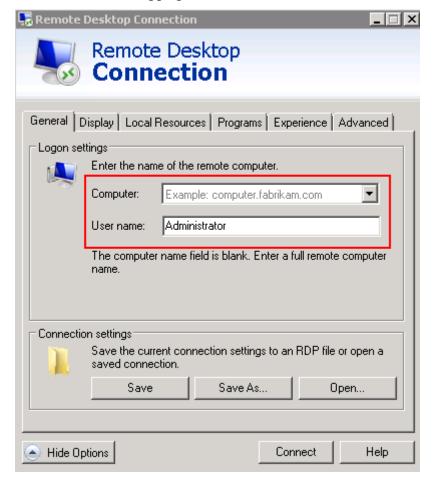
1. On the local Windows computer, click **Start**. In the **Search programs and files** text box, enter **mstsc**.

The **Remote Desktop Connection** window is displayed.

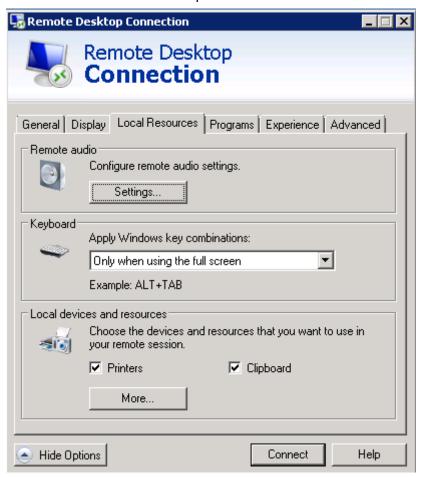
2. Click **Options**.



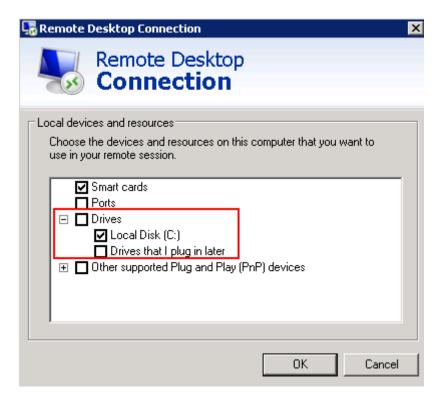
3. On the **General** tab, enter the EIP bound to the ECS and username **Administrator** for logging in to the ECS.



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.

16.15.3 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.

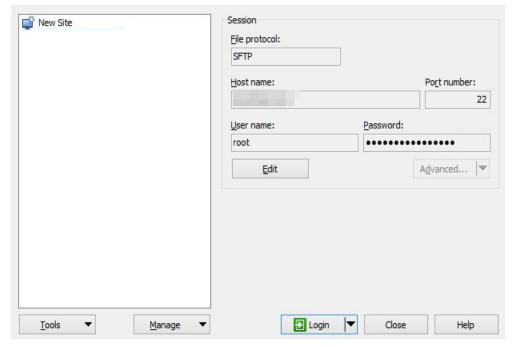
To transfer a file from a local Windows computer to a Linux ECS, WinSCP is commonly used. This section describes how to transfer files from a local Windows computer to a Linux ECS using WinSCP. In this example, the ECS running CentOS 7.2 is used as an example.

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.

16.15.4 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 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**.

Microsoft Remote Desktop

Workspaces

Add your first desktop connection to get started.

Add Desktop

Figure 16-146 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 username **administrator** and password for logging in to the Windows ECS and click **Add**.

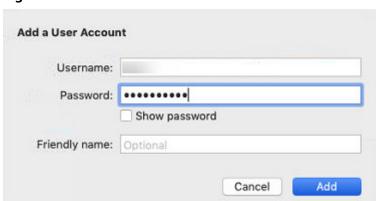


Figure 16-147 Add user account

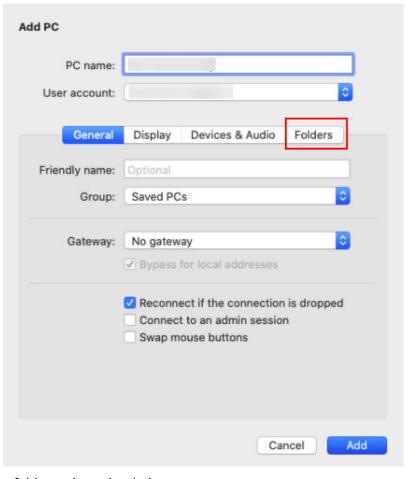
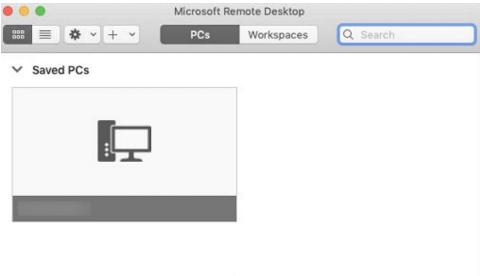


Figure 16-148 Add PC

- 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 16-149 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.

16.15.5 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 16-150 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 16-151 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
```

16.15.6 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

- 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 -q sftp -s /sbin/nologin user1

4. Set a password for the user.

passwd user1

Figure 16-152 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 16-153 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.

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>
```

16.15.7 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 the 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

F FileZilla File Edit Transfer Server Help i 🖺 🖰 🕏 🗙 🔌 🐠 Password: Port: 21 Host: Quickconnect Username: Status: Sending keep-alive command Local site: D:\dev\mingw\ Remote site: /public_html/ • 🖹 🫅 dev **□ □** / 🛨 🗀 mingw 🖃 🗀 public_html ⊟ 🛅 msys 갑 iebar in 🛅 🚹 kissa 🕀 🧀 doc nmkalkis 👚 ▲ Filename Filename

Figure 16-154 Setting connection parameters

Drag files from the local computer on the left to the target ECS on the right to transfer them.

16.15.8 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

16.15.9 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. Therefore, 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.

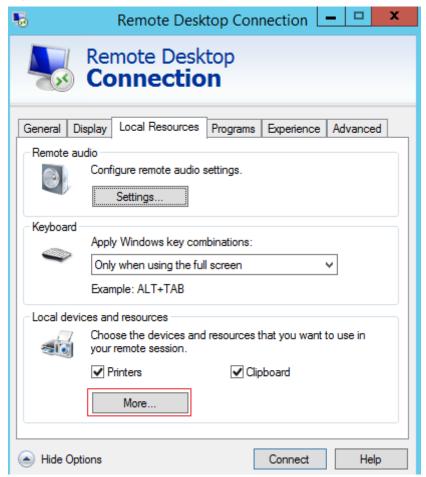
Figure 16-155 Remote Desktop Connection



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 16-156 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.

Remote Desktop
Connection

Local devices and resources
Choose the devices and resources on this computer that you want to use in your remote session.

Smart cards
Ports
Ports
Other supported Plug and Play (PnP) devices

OK Cancel

Figure 16-157 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 16-158 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.

Computer View Uninstall or change a program 9 System properties Properties Open Rename Map network Add a network drive ▼ location Open Control Panel Manage Location Network System ♠ ♠ ↑ № ► This PC ► ∨ ♂ Search This PC * Favorites ◆ Folders (6) Desktop **Downloads** Recent places 1 This PC Network Pictures Devices and drives (2) Local Disk (C:) C on ECS-2677 28.2 GB free of 39.6 GB 8 items

Figure 16-159 Viewing disks

Method 3: Set Up an FTP Site

Set up an FTP site and transfer files to the ECS.

16.15.10 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.

16.15.11 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

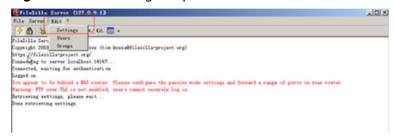
When NAT is enabled on the FTP server, the FTP client must connect to the FTP server in passive mode. In such a case, the public IP address (EIP) of the server cannot be accessed from the router. Therefore, you need to add the EIP to the public IP address list on the server. Additionally, set the port range to limit the number of ports with data forwarded by the router.

Solution

The public IP address must be associated with the private IP address using NAT. Therefore, the server must be configured accordingly.

Configure the public IP address of the server.
 Choose Edit > Settings.

Figure 16-160 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 public IP address.

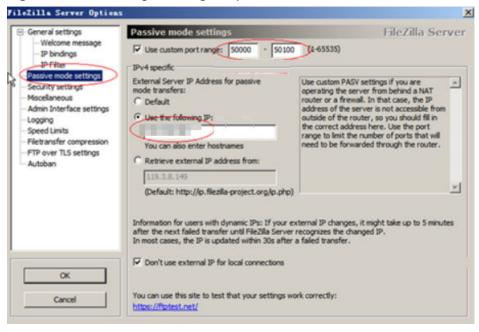


Figure 16-161 Setting the range of ports for data transmission

- 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.

16.15.12 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.

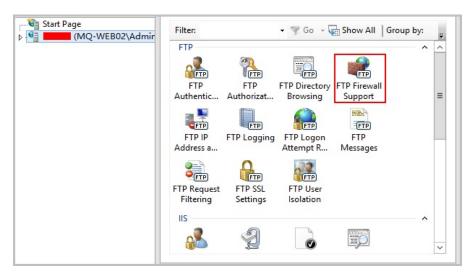
Possible Causes

- The security group associated with the target ECS denies inbound traffic.
- The firewall of the ECS blocks the FTP process.

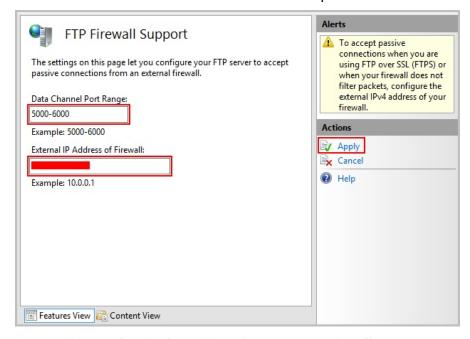
Enabling FTP Firewall Support

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. Double-click FTP Firewall Support.



- 2. 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.



3. 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.

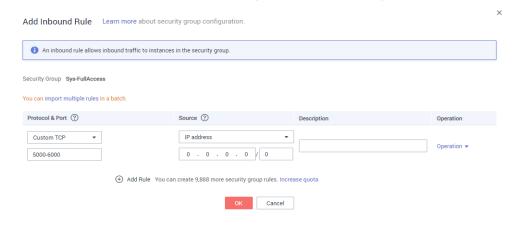
After **enabling FTP firewall support**, allow access to the ports used by the FTP site and the data channel ports used by the FTP firewall in the security group.

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.

- 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 and view security group rules.
- 5. Click the security group ID.
 - The system automatically switches to the **Security Group** page.
- 6. 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

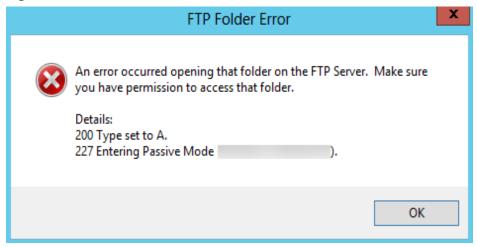
The default source IP address **0.0.0.0/0** indicates that all IP addresses can access ECSs in the security group.

16.15.13 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 16-162 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).

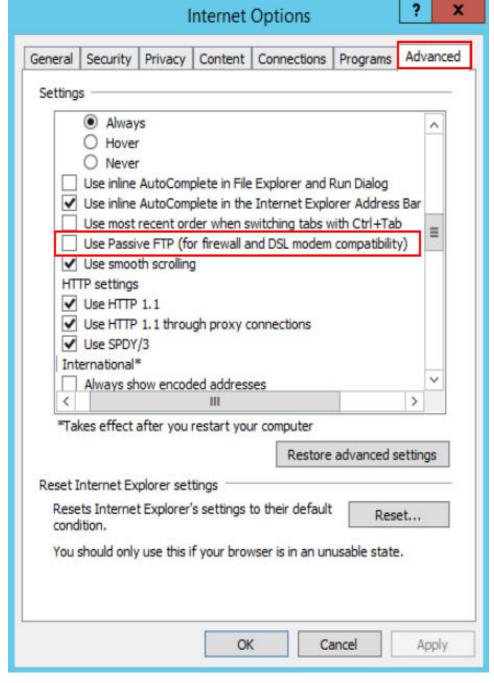


Figure 16-163 Internet Options

4. Click **OK**, restart Internet Explorer, and open the folder on the FTP server again.

16.15.14 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.

Connection has been unexpectedly closed. Server sent command exit status

Cannot initialize SFTP protocol. Is the host running a SFTP server?

OK

Reconnect (1 s)

Help

Figure 16-164 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 16-165 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

16.16 ECS Migration

16.16.1 Can I Migrate an ECS to Another Region or Account?

After an ECS is created, it cannot be directly migrated to another region or account.

16.17 Disk Management

16.17.1 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.



Formatting a disk will cause data loss. Therefore, 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.

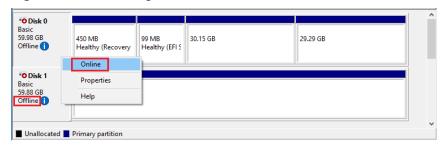
For details, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

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.
- 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 16-166 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.

16.17.2 How Can I Adjust System Disk Partitions?

Scenarios

If the capacity of system disk partitions is inconsistent with the actual system disk capacity after an ECS is created, you can manually adjust the partitions to expand the system disk.

There are two ways to expand a system disk:

- Consider the empty partition as a new partition and attach this partition to a directory in the root partition after formatting it. For details, see this section.
- Add the empty partition to the root partition to be expanded. For detailed operations, see the following:
 - How Can I Add the Empty Partition of an Expanded System Disk to the End Root Partition Online?
 - How Can I Add the Empty Partition of an Expanded System Disk to the Non-end Root Partition Online?

Procedure

This section uses an ECS running CentOS 7.3 64bit as an example. A 60 GB system disk was created with the ECS. However, the capacity of the system disk partition is displayed as only 40 GB.

To use the 20 GB capacity, performing the following operations:

Step 1 View disk partitions.

- 1. Log in to the Linux ECS.
- 2. Run the following command to switch to user **root**:

sudo su -

3. Run the following command to view details about the ECS disk:

fdisk -l

In the following command output, /dev/xvda or /dev/vda indicates the system disk.

Figure 16-167 Viewing details about the disk

```
]# df
root0ecs-8d6c
                        Used Avail Use% Mounted on
Filesystem
                 Size
/dev/xvda1
                  38G
                        1.2G
                                35G
                               899M
devtmpfs
                 899M
                           и
                                       0% /dev
tmpfs
                 908M
                           и
                               908M
                                       0% /dev/shm
tmpfs
                  908M
                        8.4M
                               900M
                                       1% /run
tmpfs
                 908M
                               908M
                                       0% /sys/fs/cgroup
                           0
                  182M
                            0
                               182M
                                       0% /run/user/0
tmpfs
[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
                                                           Ιd
                                                               System
/dev/xvda1
                        2048
                                 79980543
                                               39989248
                                                               Linux
                                                           83
                    79980544
                                 83886079
                                                           82 Linux swap / Solaris
/deu/xuda2
                                                1952768
[root@ecs-8d6c ~1# _
```

4. Run the following command to view disk partitions:

parted -l /dev/xvda

Figure 16-168 Viewing disk partitions

```
[root@ecs- ~]# 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
                         Size
                                          File system
                End
                                 Type
                                                           Flags
        1049kB
                41.0GB
                         40.9GB
                                 primary
                                          ext4
                                                           boot
 1
 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 16-169 Creating a new partition

```
| Control | Fraction |
```

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 16-170 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 16-171 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 16-172 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
                            Start
     Device Boot
                                              End
                                                           Blocks
                                                                       14
                                                                            System
                                        79980543
                                                        39989248
 /dev/xvda1
                             2048
                                                                       83
                                                                            Linux
                        79980544
                                                                       82
 /dev/xvda2
                                        83886079
                                                         1952768
                                                                            Linux swap / Solaris
                                       125829119
 /dev/xvda3
                        83886080
                                                        20971520
```

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 16-173 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 16-174 Viewing disk partitions

```
Disk Flags:
        Start
                         Size
                                            File system
Number
                 End
                                                             Flags
                                  Type
        1049kB
                 41.0GB
                         40.9GB
                                            ext4
                                                             boot
1
                                  primary
2
        41.0GB
                42.9GB
                         2000MB
                                            linux-swap(v1)
                                  primary
                         21.5GB
 3
        42.9GB
                64.4GB
                                  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 16-175 Viewing the file system type

| [root@ecs-8d6c | ~]# df -T | Н | | | | |
|----------------|-----------|--------------|------|--------------|------------|----------------|
| Filesystem | Туре | Size | Used | A∨ail | Usez | Mounted on |
| /dev/xvda1 | ext4 | 41G | 1.3G | 37G | 4% | / |
| devtmpfs | devtmpfs | 943M | 0 | 943M | 0 % | ∕de∨ |
| tmpfs | tmpfs | 952M | 0 | 952M | 0 % | /dev/shm |
| tmpfs | tmpfs | 952 M | 8.8M | 944M | 1% | /run |
| tmpfs | tmpfs | 952 M | 0 | 952 M | 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

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,
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 16-176 Viewing the mounted file systems

| Filesystem | Туре | 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 | 8% | /dev/shm |
| tmpfs | tmpfs | 952M | 8.8M | 944M | 1% | /run |
| tmpfs | tmpfs | 952M | 0 | 952M | 8% | /sys/fs/cgroup |
| /dev/xvda3 | ext4 | 22G | 47M | 20G | 1% | /root/new |
| tmpfs | tmpfs | 191M | 0 | 191M | 0% | /run/user/0 |
| [root@ecs-8d6 | oc ~1# bl | | | | | |

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 16-177 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="96e5e028-60fb-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

16.17.3 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**.

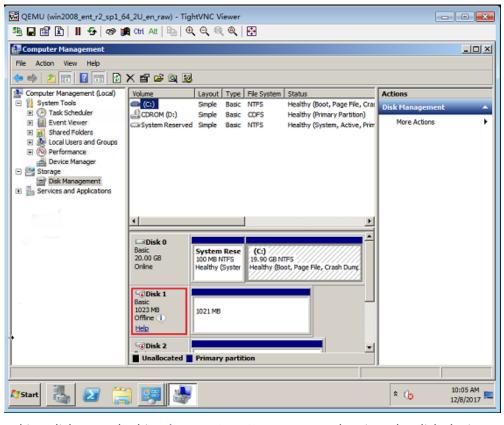


Figure 16-178 Disk Management

- 5. Taking disk 1 marked in **Figure 16-178** 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 16-178.
 - b. Click **Properties**.

The SCSI Disk Device Properties dialog box is displayed, as shown in Figure 16-179.

General Policies Volumes Driver Details

XEN PV DISK SCSI Disk Device

XEN PV DISK SCSI Disk Device

Device type: Disk drives

Manufacturer: (Standard disk drives)

Location: Bus Number 0, Target Id 0, LUN 0

Device status

This device is working properly.

Figure 16-179 Disk properties

c. Click the **Details** tab and set **Property** to **Parent**.

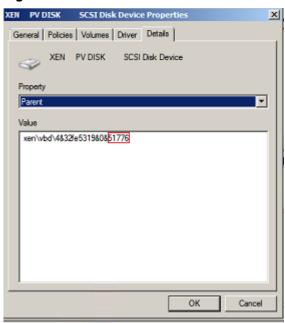


Figure 16-180 Disk device details

- 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 16-12.
 The disk device corresponding to 51776 is xvde. The disk device used by disk 1 is xvde.

Table 16-12 Mapping between disk partitions and disk devices

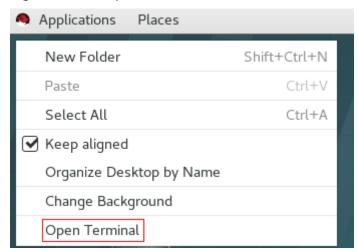
| Master and Slave Device Number for a Disk Partition | Disk Device |
|--|-------------|
| 51712 | xvda |
| 51728 | xvdb |
| 51744 | xvdc |
| 51760 | xvdd |
| 51776 | xvde |
| 51792 | xvdf |
| 51808 | xvdg |
| 51824 | xvdh |
| 51840 | xvdi |
| 51856 | xvdj |
| 51872 | xvdk |
| 51888 | xvdl |
| 51904 | xvdm |
| 51920 | xvdn |
| 51936 | xvdo |
| 51952 | xvdp |
| 268439552 | xvdq |
| 268439808 | xvdr |
| 268440064 | xvds |
| 268440320 | xvdt |
| 268440576 | xvdu |
| 268440832 | xvdv |
| 268441088 | xvdw |
| 268441344 | xvdx |

16.17.4 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 16-181 open terminal



3. Run the following command to view disk partitions and disk devices:

fdisk -l

Figure 16-182 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
                                      End
                                               Blocks
                                                         Id System
                      Start
/dev/xvdal *
                       2048
                                 2099199
                                              1048576
                                                         83
                                                             Linux
                    2099200
                                16777215
                                              7339008
/dev/xvda2
                                                         83
                                                             Linux
                                                         82 Linux swap / Solaris
                   16777216
                                20971519
                                              2097152
/dev/xvda3
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 16-13 lists the mapping between disk partitions and disk devices.

Table 16-13 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 |

16.17.5 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.

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.
- In the navigation pane on the left, choose Advanced system settings.
 The System Properties dialog box is displayed.
- Click the Advanced tab and then Settings in the Performance pane.
 The Performance Options dialog box is displayed.

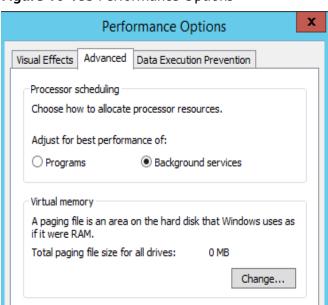


Figure 16-183 Performance Options

- 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 4096.

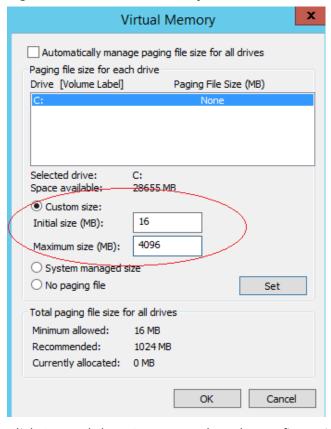


Figure 16-184 Virtual Memory

- Click **Set** and then **OK** to complete the configuration.
- Restart the ECS for the configuration to take effect.

16.17.6 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.

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
```

16.17.7 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 attaching 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 attaching 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
                             exit program
                                  rescue a lost partition near START and END
 rescue START END
 rm NUMBER
                                  delete partition NUMBER
 select DEVICE
                                choose the device to edit
 disk_set FLAG STATE
                                   change the FLAG on selected device
                                  toggle the state of FLAG on selected device
 disk_toggle [FLAG]
```

set NUMBER FLAG STATE toggle [NUMBER [FLAG]] unit UNIT version Parted change the FLAG on partition NUMBER toggle the state of FLAG on partition NUMBER set the default unit to UNIT

display the version number and copyright information of GNU

b. Press **p**.

(parted)

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

16.17.8 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 16-14 Numbers of disks that can be attached to a newly created ECS

| ECS Type | Maximu m VBD Disks | Maximu m SCSI Disks | Constraint |
|---|--------------------------|---------------------------|--|
| Xen (excluding D1 ECSs) | 60 | 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. |
| D1 | 60 | 59 | VBD disks + SCSI disks + Local disks ≤ 60 The number of local disks is determined based on the ECS flavor. |
| KVM (excluding D2 and 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. |

| ECS Type | Maximu m VBD Disks | Maximu m SCSI Disks | Constraint |
|-----------|--------------------------|---------------------------|--|
| D2 and 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. Therefore, 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. Therefore, 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 16-15**.

Table 16-15 Numbers of disks that can be attached to an existing ECS

| ECS Type | Maximu m VBD Disks | Maximu m SCSI Disks | Maximu m Local Disks | Constraint |
|----------|--------------------------|---------------------------|----------------------------|--|
| Xen | 60 | 59 | 59 | VBD disks + SCSI disks + Local disks ≤ 60 |
| KVM | 24 | 23 | 59 | VBD disks + SCSI disks ≤ 24 |

To attach 60 disks to an exiting ECS, create the same ECS, modify its specifications, reinstall/change its OS, or migrate the ECS.

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.

16.17.9 What Are the Requirements for Attaching an EVS Disk to an ECS?

- The EVS disk and the target ECS must be located in the same AZ.
- For a non-shared disk, the EVS disk must be in Available state.
 For a shared disk, the target EVS disk must be in In-use or Available state.
- The target ECS must be in **Running** or **Stopped** state.
- The EVS disk must not be frozen.
- Certain ECSs support SCSI EVS disk attachment. For details, see Which ECSs
 Can Be Attached with SCSI EVS Disks?

16.17.10 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.

16.17.11 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**.

Using a Serial Number to Obtain the Disk Name (Windows)

If a serial number is displayed on the console, use either of the following methods to obtain the disk name.

cmd

1. Start **cmd** in a Windows OS as an administrator and run either of the following commands:

wmic diskdrive get serialnumber wmic path win32_physicalmedia get SerialNumber wmic path Win32_DiskDrive get SerialNumber

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

Figure 16-185 Obtaining the disk serial number

Information similar to the following is displayed:

```
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 Vin32_DiskDrive get SerialNumber
SerialNumber
```

2. Run the following command to check the disk corresponding to the serial number:

wmic diskdrive get Name, SerialNumber

Figure 16-186 Checking the disk corresponding to the serial number

```
C:\Users\Administrator>wmic diskdrive get Name, SerialNumber
Name SerialNumber
\\.\PHYSICALDRIVEØ 97c876c0-54b3-460a-b
```

PowerShell

- 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 16-187, 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 16-187, the disk is Disk 0.

Figure 16-187 Viewing the disk on which the logical disk is created

```
PS C:\Users\Administrator> Get-CimInstance -ClassName Win32_LogicalDiskToPartition |select Anteceder

Antecedent : Win32_DiskPartition (DeviceID = "Disk #0, Partition #1")

PS C:\Users\Administrator> Get-Disk |select Number, SerialNumber

Number SerialNumber

0 97c876c0-54b3-460a-b
1 dswfa16520d39517815206127
```

- Versions earlier than Windows 2012
 - 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

Background

Disk information displayed varies according to the ECS virtualization type. For the sake of convenience, ECSs that use the KVM virtualization type are called KVM instances, and ECSs that use the Xen virtualization type are called Xen instances.

Obtaining the Disk Device Name of a KVM Instance

Step 1 Obtain the disk information displayed on the console.

- 1. Log in to the management console.
- 2. Under Computing, click Elastic Cloud Server.
- Click the target ECS name in the ECS list.
 The page providing details about the ECS is displayed.
- 4. Click the **Disks** tab and then to expand the disk information.
- 5. Check the device 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**.



If **Device Identifier** is not displayed on the web page, stop the ECS and restart it.

Step 2 Check the device name of a VBD disk attached to the ECS.

- Obtain the disk device ID by referring to Step 1.
 The device ID of the VBD disk shows the PCI address of the disk on the ECS.
 The address is in the format of "domain:bus:slot.function".
- 2. Log in to the ECS as user root.
- 3. In /sys/bus/pci/devices/DOMIN:BUS:SLOT.FUNCTION/virtio*/block, view the device name.

For example, if the device ID of the VBD disk is **0000:00:05.0**, the device name is shown as follows:

A90CF6C6-BEC0-0C44-8082-8C8610755B61:/sys/bus/pci/devices/0000:00:05.0/virtiol/block # ll /sys/bus/pci/devices/0000:00:05.0/virtio1/block total 0 drwxr-xr-x 10 root root 0 May 22 11:01 **vda**

The displayed information is the disk device name, /dev/vda in the preceding figure.

Step 3 Check the device name of a SCSI disk attached to the ECS.

- Obtain the disk device ID by referring to Step 1.
 The device ID of the SCSI disk displays the disk WWN on the ECS.
- 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 WWMgrep 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

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.
 - 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 ECS.

Table 16-16 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*** |

| Device Name (on Management Console) | Device Name (in ECS) |
|--|----------------------|
| /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.

- Obtain the disk device ID.
 The device ID of the SCSI disk is the disk WWN on the ECS.
- 2. Log in as user root.
- 3. Run the following command to view the disk device name:

ll /dev/disk/by-id |grep WWMgrep 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

16.17.12 Why Does a Linux ECS with a SCSI Disk Attached Fails 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. Therefore, a slot ID corresponds to a drive letter.

After the SCSI disk is detached from the running ECS, the slot sequence for disks may change, leading to the disk drive letter being changed after the ECS is restarted. As a result, the slot IDs do not correspond to the drive letters, and the ECS fails to restart.

Solution

- 1. Log in to the Linux ECS.
- 2. Run the following command to switch to user **root**:

sudo su -

3. Run the following command to obtain the SCSI ID according to the drive letter of the SCSI disk:

ll /dev/disk/by-id/|grep Disk drive letter

For example, if the drive letter of the SCSI disk is **/dev/sdb**, run the following command:

ll /dev/disk/by-id/|grep sdb

CNA64_22:/opt/galax/eucalyptus/ecs_scripts # ll /dev/disk/by-id/|grep sdb lrwxrwxrwx 1 root root 9 Dec 6 11:26 scsi-3688860300001436b005014f890338280 -> ../../sdb lrwxrwxrwx 1 root root 9 Dec 6 11:26 wwn-0x688860300001436b005014f890338280 -> ../../sdb

4. Change the drive letter (for example, /dev/sdb) of the SCSI disk to the corresponding SCSI ID in the /etc/fstab file.

/dev/disk/by-id/SCS/ ID

For example, if the SCSI ID obtained in step **3** is scsi-3688860300001436b005014f890338280, use the following data to replace **/dev/sdb**:

/dev/disk/by-id/scsi-3688860300001436b005014f890338280

16.17.13 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 antiaffinity 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" in the *Elastic Volume Service User Guide*.

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.

| \square |) not | Е |
|-----------|-------|---|
| | | |

If the ECS group field is left blank, the ECS has not been added to any ECS group.

16.17.14 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 16-188** 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.

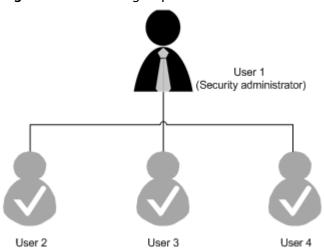


Figure 16-188 User group

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:

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" 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.

16.17.15 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 antiaffinity and are allocated on different hosts.

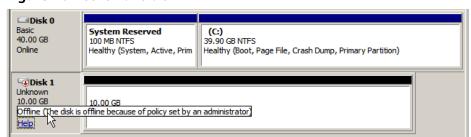
An ECS with local disks attached cannot be added to an ECS group after the ECS is created. Such ECSs can be added to an ECS group only during the ECS creation.

16.17.16 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."

Figure 16-189 Offline disk



Possible Causes

Windows has three types of SAN policies: **OnlineAll**, **OfflineShared**, and **OfflineInternal**.

Table 16-17 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.

16.17.17 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
```

16.17.18 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. | Loa | in t | o the | manac | iement | console | and | create a | new | ECS. |
|-----|-----|------|-------|-------|--------|---------|-----|-----------|-----|------|
| • • | | | ~~ | | , | | 44 | c. care c | | |

□ NOTE

Ensure that the new ECS is located in the same AZ and has the same parameter settings as the original ECS.

 (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**.

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.

Ⅲ 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.

16.17.19 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 create the ECS again for reconstruction.

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 16-190.

Figure 16-190 Emergency mode

```
Connected (encrypted) to: QEMU (instance-00000067) Before you exit, ensure that computer is locked.

5.3478961 intel_rapl: no valid rapl domains found in package 0
5.5063741 EXT4-fs (numeOn1): UFS: Can't find ext4 filesystem
Welcome to emergency mode! After logging in, type "journalctl -xb" to view
system logs, "systemctl reboot" to reboot, "systemctl default" or ^D to
try again to boot into default mode.
Give root password for maintenance
(or type Control-D to continue): _
```

To ensure that the new ECS is functional, you must manually delete the attachment information in /etc/fstab.

Ⅲ NOTE

If the NVMe SSD disk is faulty, data on it will be lost. The operations provided in this section are only used to restore automatic NVMe SSD disk attachment to an ECS, but not restoring the data on the disk.

Solution

- 1. Log in to the ECS.
- 2. Enter the password of user **root** to log in to the ECS.

Figure 16-191 Logging in to the ECS

```
Connected (encrypted) to: QEMU (instance-00000067) Before you exit,ensure that co

[ 5.347896] intel_rapl: no valid rapl domains found in package 0

[ 5.506374] EXT4-fs (nvmeOn1): UFS: Can't find ext4 filesystem

Welcome to emergency mode! After logging in, type "journalctl -xb" to view
system logs, "systemctl reboot" to reboot, "systemctl default" or ^D to

try again to boot into default mode.

Give root password for maintenance

(or type Control-D to continue):

[root0localhost ~1#
```

3. Run the following command to edit the /etc/fstab file:

vi /etc/fstab

4. Delete the attaching information of the NVMe SSD disk and save the file.

Figure 16-192 Deleting the automatic attaching information

```
# /etc/fstab
# created by anaconda on Wed Aug 9 09:22:35 Z017
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
# /dev/mapper/cl-root / xfs defaults 0 0
UUID=17cbcc3f-0b23-4eaa-04f6-6bc68583b521 /boot xfs defaults 0 0
/dev/mapper/cl-swap swap swap defaults 0 0
/dev/nvmc0n1 /for_nvme ext3 defaults 0 0

"/etc/fstab" 12L, 506C
```

5. Run the following command to restart the ECS:

reboot

6. Verify that the ECS recovers and can be logged in.

Figure 16-193 Logging in to the ECS

```
Connected (encrypted) to: QEMU (instance
CentOS Linux 7 (Core)
Kernel 3.10.0-514.el7.x86_64 on an x86_64
localhost login: _
```

16.17.20 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
                 2048 83886079 41942016 83 Linux
/dev/vda1 *
Disk /dev/vdb: 107.4 GB, 107374182400 bytes, 209715200 sectors
Units = sectors of 1 \times 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.

16.17.21 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 16-194 Printed logs



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.

16.18 Passwords and Key Pairs

16.18.1 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

16.18.2 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.

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**.

Ⅲ 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 16-195 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
```

16.18.3 Changing the Login Password on an ECS

Scenarios

This section describes how to change the password for logging in to an ECS when the password is about to expire, the password is forgotten, or you are logging in to the ECS for the first time. It is a good practice to change the initial password upon the first login.

Prerequisites

The ECS can be logged in.

Background

Table 16-18 shows the ECS password complexity requirements.

Table 16-18 Password complexity requirements

| Parameter | Requirement | Example Value |
|-----------|---|--|
| Password | Consists of 8 to 26 characters. Contains at least three of the following character types: Uppercase letters Lowercase letters Digits Special characters: \$!@%=+[]:./^,{}? 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. | YNbUwp! dUc9MClnv NOTE The example password is generated randomly. Do not use it. |

Windows

- 1. Log in to the ECS.
 - For details, see **Login Overview**.
- 2. Press Win+R to start the Run dialog box.
- 3. Enter **cmd** to open the command-line interface (CLI) window.
- 4. Run the following command to change the password (the new password must meet the requirements described in **Table 16-18**):
 - net user Administrator New password

Linux

- Use the existing key file to log in to the ECS as user root through SSH.
 For details, see Remotely 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 16-18**.

New password:

Retype new password:

If the following information is displayed, the password has been changed: passwd: all authentication tokens updates successfully

16.18.4 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. Under Computing, click Elastic Cloud Server.
- c. Stop the original Windows ECS, switch to the page providing details about the ECS, and click the **Disks** tab.

∩ NOTE

Do not forcibly stop the Windows ECS. Otherwise, password reset may fail.

- d. Locate the row containing the system disk to be detached and click **Detach** to detach the system disk from the ECS.
- 2. Attach the system disk to the temporary ECS.
 - a. On the page providing details about the temporary ECS, click the **Disks** tab.
 - b. Click **Attach Disk**. In the displayed dialog box, select the system disk detached in step **1.d** and attach it to the temporary ECS.
 - c. Remotely log in to the temporary ECS.
 - d. Run the following command to view the directory of the system disk detached from the original Windows ECS now attached to the temporary ECS:

fdisk -l

e. Run the following command to mount the file system of the detached system disk to the temporary ECS:

mount -t ntfs-3g /dev/Result obtained in step 2.d /mnt/

For example, if the result obtained in step **2.d** is **xvde2**, run the following command:

mount -t ntfs-3g /dev/xvde2 /mnt/

If the following error information is displayed after the preceding command is executed, the NTFS file systems may be inconsistent. In such a case, rectify the file system inconsistency.

The disk contains an unclean file system (0, 0). Metadata kept in Windows cache, refused to mount. Failed to mount '/dev/xvde2': Operation not permitted The NTFS partition is in an unsafe state. Please resume and shutdown Windows fully (no hibernation or fast restarting), or mount the volume read-only with the 'ro' mount option.

Back up the disk data, run the following command to rectify the NTFS file system inconsistency, and attach the system disk:

ntfsfix /dev/Result obtained in step 2.d

For example, if the result obtained in step **2.d** is **xvde2**, run the following command:

ntfsfix /dev/xvde2

- 3. Change the password of the specified user and clear the original password.
 - a. Run the following command to back up the SAM file:

cp /mnt/Windows/System32/config/SAM /mnt/Windows/System32/ config/SAM.bak

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 16-19**):

net user Administrator New password

Parameter Requirement Example Value Password Consists of 8 to 26 characters. YNbUwp! dUc9MClnv Contains at least three of the NOTE following character types: The example Uppercase letters password is generated Lowercase letters randomly. Do Digits not use it. Special characters: \$!@%-_=+[]:./^, 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.

Table 16-19 Password complexity requirements

16.18.5 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. Download the script for resetting the password and upload the script to the temporary ECS.

Contact customer service to obtain the password reset script. Use a connection tool, such as WinSCP, to upload the obtained **changepasswd.sh** script to the temporary ECS.

To download WinSCP, log in at https://winscp.net/.

- 2. Stop the original Linux ECS, detach the system disk from it, and attach the system disk to the temporary ECS.
 - a. Log in to the management console.
 - b. Under Computing, click Elastic Cloud Server.
 - c. Stop the original ECS, switch to the page providing details about the ECS, and click the **Disks** tab.

Do not forcibly stop the original ECS. Otherwise, password reset may fail.

- d. Locate the row containing the system disk to be detached and click **Detach** to detach the system disk from the ECS.
- 3. Attach the system disk to the temporary ECS.
 - a. On the page providing details about the temporary ECS, click the **Disks** tab.
 - b. Click **Attach Disk**. In the displayed dialog box, select the system disk detached in step **2.d** and attach it to the temporary ECS.
- 4. Log in to the temporary ECS remotely and reset the password.
 - a. Locate the row containing the temporary ECS and click **Remote Login** in the **Operation** column.
 - 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 16-196 Viewing the directory of the system disk

```
root@ecs----:~# fdisk -1
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
                 2048 83884031 83881984 40G 83 Linux
/dev/vda1
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
                2048 83886079 83884032 40G 83 Linux
/dev/vdb1 ×
```

c. Run the following commands in the directory where the **changepasswd.sh** script is stored to run the script for resetting the password:

chmod +x changepasswd.sh

./changepasswd.sh

When you run the password reset script, if the system displays a message indicating that there is no command related to logical volume manager (LVM), such as the message "no lvs command", install an LVM tool on the temporary ECS. The LVM2 tool is recommended, which can be installed by running the **yum install lvm2** command.

∩ NOTE

If the original ECS and the temporary ECS both run CentOS 7, a mount failure may occur during script execution. To resolve this issue, replace **mount \$dev \$mountPath** with **mount -o nouuid \$dev \$mountPath** in the script.

d. Enter the new password and the directory obtained in step **4.b** as prompted.

If the following information is displayed, the password has been changed: set password success.

5. (Optional) Enable remote root login for non-root users.

vi /etc/ssh/sshd_config

Modify the following settings:

- Change PasswordAuthentication no to PasswordAuthentication yes.
 Alternatively, uncomment PasswordAuthentication yes.
- Change PermitRootLogin no to PermitRootLogin yes.
 Alternatively, uncomment PermitRootLogin yes.
- Change the value of AllowUsers to root.
 Search for AllowUsers in the file. If AllowUsers is missing, add it at the end of the file.
- 6. Stop the temporary ECS, detach the system disk, attach the system disk to the original Linux ECS, and restart the original Linux ECS.
 - a. Stop the temporary ECS, switch to the page providing details about the ECS, and click the **Disks** tab.
 - b. Click **Detach** to detach the data disk attached in step 2.
 - c. On the page providing details about the original Linux ECS, click the **Disks** tab.
 - d. Click **Attach Disk**. In the displayed dialog box, select the data disk detached in **6.b**.
- 7. Restart the original Linux ECS.

16.18.6 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

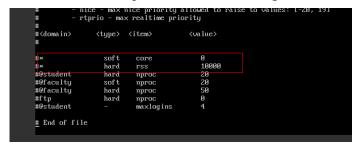
After verifying the preceding configurations, restart the ECS, wait for 3 to 5 minutes, and remotely log in to the ECS using a password or key.

16.18.7 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.



Change the password in single-user mode and try to log in to the ECS again.

16.18.8 Disabling SELinux

Ⅲ 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.

16.18.9 How Can I Obtain the Key Pair Used by My ECS?

Symptom

You have created multiple key pairs, and you are trying to find the key pair to log in to the target ECS.

Procedure

- 1. Log in to the management console.
- 2. Under Computing, click Elastic Cloud Server.
- 3. On the **Elastic Cloud Server** page, select the target ECS.
- Click the name of the target ECS.
 The page providing details about the ECS is displayed.
- Obtain the **Key Pair** value.
 The value is the key pair used by the ECS.

16.18.10 How Can I Use a Key Pair?

Symptom

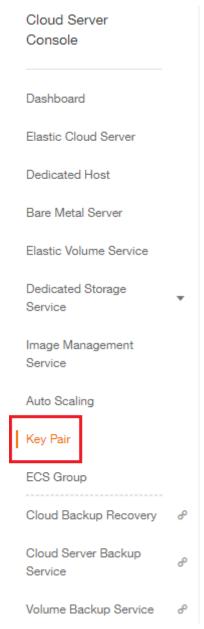
When you purchase an ECS, the system asks you to select a login mode. If you select **Key pair**, you are required to select an existing key pair or create a new pair.

If no key pair is available, create one on the management console.

Solution

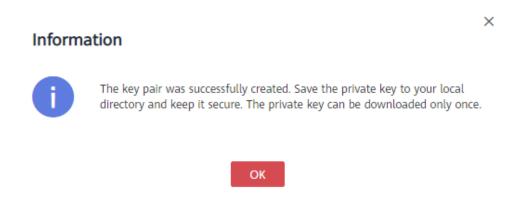
1. In the navigation pane of the ECS console, choose **Key Pair**. Then, click **Create Key Pair**.

Figure 16-197 Creating a key pair



2. After the key pair is created, download the private key to a local directory.

Figure 16-198 Downloading a key pair



When purchasing an ECS, select the created or existing key pair in Key pair.

Figure 16-199 Selecting a key pair



16.18.11 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.
- Move the scroll bar to set the security level to Medium and click Apply.
- Click Custom Level.
- 8. Set Initialize and script ActiveX controls not marked as safe for scripting to Prompt.
- 9. Click Yes.

16.18.12 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.

16.18.13 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 16-200 Key Pair



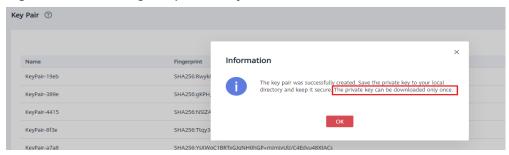
Click Create Key Pair.

Figure 16-201 Creating a key pair



3. Click **OK** to save the private key to your local directory.

Figure 16-202 Saving the private key



16.18.14 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. Therefore, 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**.

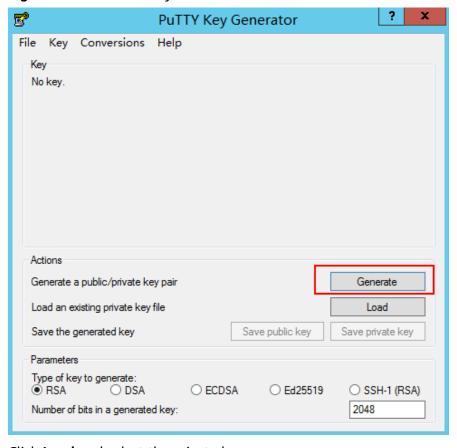


Figure 16-203 PuTTY Key Generator

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 16-204** is the public key whose format meets system requirements.

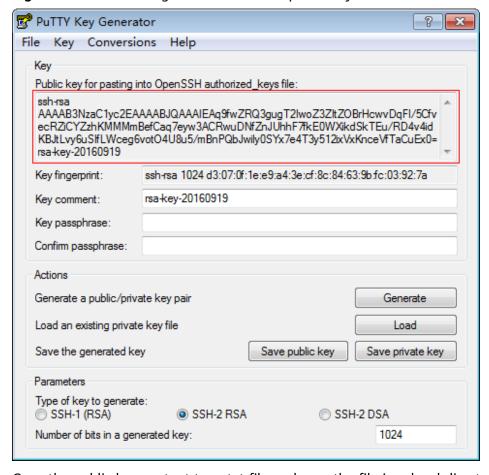


Figure 16-204 Restoring the format of the public key content

- 3. Copy the public key content to a .txt file and save the file in a local directory.
- 4. Import the public key to the management console.
 - a. Log in to the management console.
 - b. Under Computing, click Elastic Cloud Server.
 - c. In the navigation pane on the left, choose **Key Pair**.
 - d. On the key pair page, click Import Key Pair.
 - e. Copy the public key content in the .txt file to **Public Key Content** and click **OK**.

16.18.15 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 the 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:683020925C5D8569C 23AA724774CE9CC:::",

- 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 16-205 ecs01

• LM-HASH and NTLM-HASH values of the **cloudbase-init** account for ecs02

Figure 16-206 ecs02

16.18.16 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 16-207**.

Figure 16-207 Improper running of Cloud-Init

```
Iroot@ecs-8560 ~ 1# cloud init v
Traceback (most recent call last):
    File "/usr/bin/cloud-init", line 39, in <module>
        from eloudinit import patcher
ImportError: No module named cloudinit
Iroot@ecs-8560 ~ 1# 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
Iroot@ecs-8560 ~ 1#
```

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 16-208 Changing the Python version

```
frootGecs-8560 ~ 1# head -n 1 /usr/bin/cloud-init

#/usr/bin/python2.6

frootGecs-8560 ~ 1# Is /usr/bin/python* -lh

lrwxrwxrwx 1 root root 24 Jul 19 18:55 /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
```

16.19 Network Configurations

16.19.1 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 16-20 lists ECS configurations.

Table 16-20 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**.

16.19.2 Can an ECS Without an EIP Bound Access the Internet?

Yes.

You can configure the SNAT server so that the ECS without an EIP bound can access the Internet.

For details, see **Configuring an SNAT Server**.

16.19.3 Why Can't an EIP Be Pinged?

Symptom

After you purchase an EIP and bind it to an ECS, the EIP cannot be pinged on a local server or other cloud servers.

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 16-209 Method of locating the failure to ping an EIP

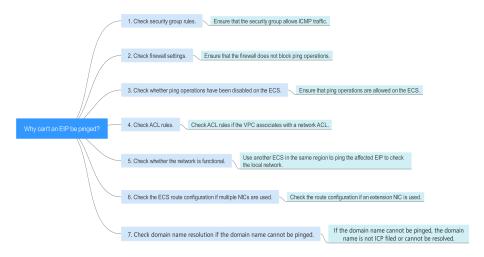


Table 16-21 Method of locating the failure to ping an EIP

| Possible Cause | Solution |
|--|---|
| ICMP access rules are not added to the security group. | Add ICMP access rules to the security group. For details, see Checking Security Group Rules. |
| Ping operations are prohibited on the firewall. | Allow ping operations on the firewall. For details, see Checking Firewall Settings. |
| Ping operations are prohibited on the ECS. | Allow ping operations on the ECS. For details, see Checking Whether Ping Operations Have Been Disabled on the ECS. |
| Network ACL is associated. | If the VPC is associated with a network ACL, check the network ACL rules. For details, see Checking ACL Rules . |
| A network exception occurred. | Use another ECS in the same region to check whether the local network is functional. For details, see Checking Whether the Network Is Functional. |

| 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 Checking the ECS Route Configuration If Multiple NICs Are Used. |

Checking Security Group Rules

ICMP is used for the ping command. Check whether the security group accommodating the ECS allows ICMP traffic.

- 1. Log in to the management console.
- 2. Under Computing, click Elastic Cloud Server.
- 3. On the **Elastic Cloud Server** page, click the name of the target ECS. The page providing details about the ECS is displayed.
- 4. Click the **Security Groups** tab, expand the information of the security group, and view security group rules.
- 5. Click the security group ID.
 - The system automatically switches to the **Security Group** page.
- 6. On the **Outbound Rules** page, click **Add Rule**. In the displayed dialog box, set required parameters to add an outbound rule.

Table 16-22 Security group rules

| Transfer Direction | Туре | Protocol/Port Range | Source |
|-----------------------|------|------------------------|---|
| Outboun d | IPv4 | ICMP/Any | 0.0.0.0/0 0.0.0.0/0 indicates all IP addresses. |

7. On the **Inbound Rules** tab, click **Add Rule**. In the displayed dialog box, set required parameters to add an inbound rule.

Table 16-23 Security group rules

| Transfer Direction | Туре | Protocol/Port Range | Source |
|-----------------------|------|------------------------|---|
| Inbound | IPv4 | ICMP/Any | 0.0.0.0/0 0.0.0.0/0 indicates all IP addresses. |

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

Checking Firewall Settings

If a firewall is enabled on the ECS, check whether the firewall blocks the ping operations.

Linux

1. Consider CentOS 7 as an example. Run the following command to check the firewall status:

firewall-cmd --state

If **running** is displayed in the command output, the firewall has been enabled.

2. Check whether there is any ICMP rule blocking the ping operations.

iptables -L

If the command output shown in **Figure 16-210** is displayed, there is no ICMP rule blocking the ping operations.

Figure 16-210 Checking firewall rules

```
[root@ecs-3c4e ~]# iptables -L
Chain INPUT (policy ACCEPT)
target
          prot opt source
ACCEPT
          icmp -- anywhere
                                        anywhere
                                                              icmp echo-request
Chain FORWARD (policy ACCEPT)
                                        destination
target
Chain OUTPUT (policy ACCEPT)
          prot opt source
                                        destination
target
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.
 - 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 16-211 Inbound Rules

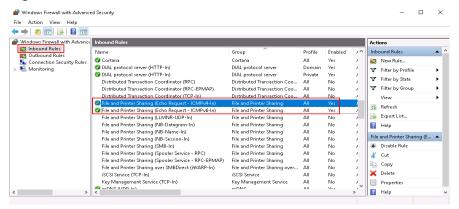
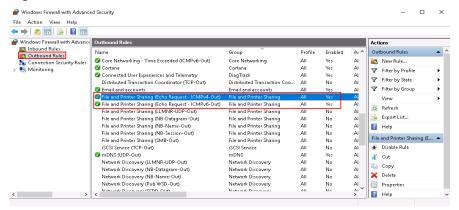


Figure 16-212 Outbound Rules



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**.
- 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.
- 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.

Ⅲ 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 16-213 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
```

If the route is unavailable, run the following command to add it:
 ip route add default via XXXX dev eth0

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.

16.19.4 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

16.19.5 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

16.19.6 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

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

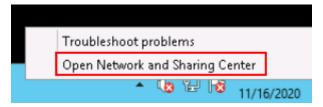
----End

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 16-214 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.

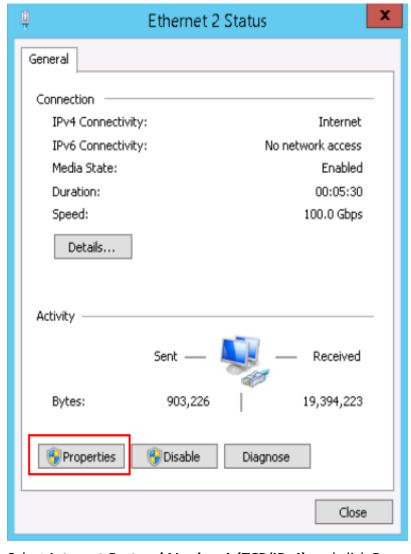


Figure 16-215 Local area connection

3. Select Internet Protocol Version 4 (TCP/IPv4) and click Properties.

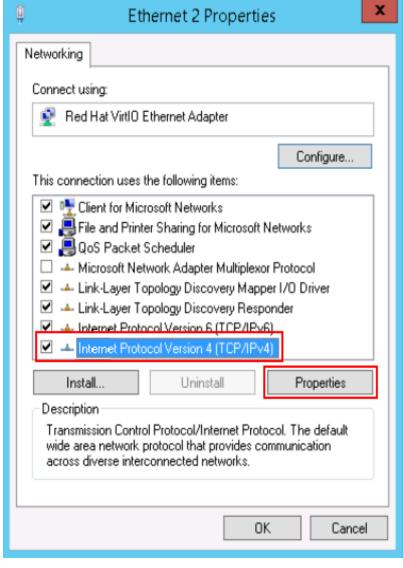


Figure 16-216 Selecting a protocol type

Select Use the following DNS server addresses and set the IP addresses of the DNS servers.

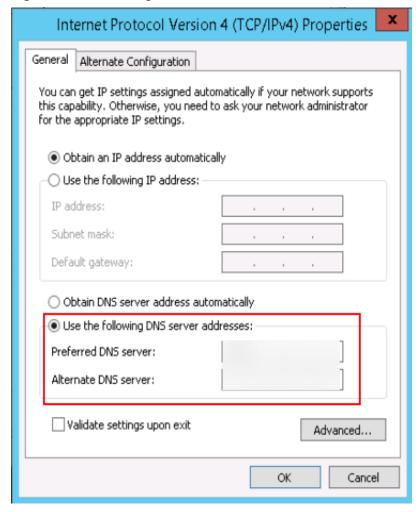


Figure 16-217 Setting the IP addresses of the DNS servers

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 ffffffff 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

16.19.7 Configuring DNS

A DNS server is used to resolve domain names of file systems. For details about DNS server IP addresses, see **What Are Private DNS Servers and What Are Their Addresses?**

Scenarios

By default, the IP address of the DNS server used to resolve domain names of file systems is automatically configured on ECSs when creating ECSs. No manual configuration is needed except when the resolution fails due to a change in the DNS server IP address.

Windows Server 2012 is used as an example in the operation procedures for Windows.

Procedure (Linux)

- **Step 1** Log in to the ECS as user **root**.
- Step 2 Run the vi /etc/resolv.conf command to edit the /etc/resolv.conf file. Add the DNS server IP address above the existing nameserver information. See Figure 16-218.

Figure 16-218 Configuring DNS

```
; generated by /sbin/dhclient-script
search openstacklocal
nameserver nameser
```

The format is as follows:

nameserver 100.125.1.250 nameserver 100.125.17.29

- **Step 3** Press **Esc**, input :wq, and press **Enter** to save the changes and exit the vi editor.
- **Step 4** Run the following command to check whether the IP address is successfully added: cat /etc/resolv.conf

Step 5 Run the following command to check whether an IP address can be resolved from the file system domain name:

nslookup File system domain name

Obtain the file system domain name from the file system mount point.

- **Step 6** (Optional) In a network environment of the DHCP server, edit the /etc/resolv.conf file to prevent the file from being automatically modified upon an ECS startup, and prevent the DNS server IP address added in **Step 2** from being reset.
 - 1. Run the following command to lock the file:

chattr +i /etc/resolv.conf

Run the chattr -i /etc/resolv.conf command to unlock the file if needed.

2. Run the following command to check whether the editing is successful:

lsattr /etc/resolv.conf

If the information shown in Figure 16-219 is displayed, the file is locked.

Figure 16-219 A locked file

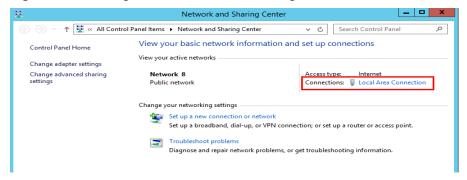
```
[root@dailing list /]# lsattr /etc/resolv.conf
----i----e- /etc/resolv.conf
```

----End

Procedure (Windows)

- **Step 1** Go to the ECS console and log in to the ECS running Windows Server 2012.
- Step 2 Click This PC in the lower left corner.
- **Step 3** On the page that is displayed, right-click **Network** and choose **Properties** from the drop-down list. The **Network and Sharing Center** page is displayed, as shown in **Figure 16-220**. Click **Local Area Connection**.

Figure 16-220 Page for network and sharing center



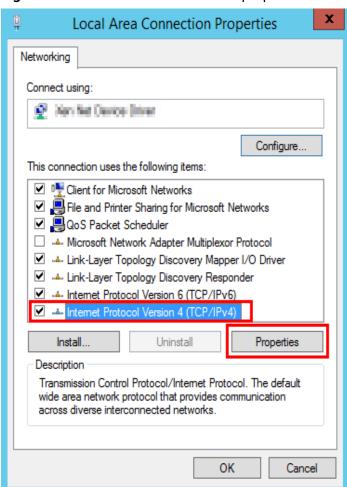
Step 4 In the **Activity** area, select **Properties**. See **Figure 16-221**.

Figure 16-221 Local area connection



Step 5 In the Local Area Connection Properties dialog box that is displayed, select Internet Protocol Version 4 (TCP/IPv4) and click Properties. See Figure 16-222.

Figure 16-222 Local area connection properties



Step 6 In the dialog box that is displayed, select **Use the following DNS server addresses:** and configure DNS, as shown in **Figure 16-223**. The DNS server IP address is 100.125.1.250. After completing the configuration, click **OK**.

Obtain DNS server address automatically

Use the following DNS server addresses:

Preferred DNS server:

Alternate DNS server:

Validate settings upon exit

OK

Cancel

Figure 16-223 Configuring DNS on Windows

----End

16.19.8 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

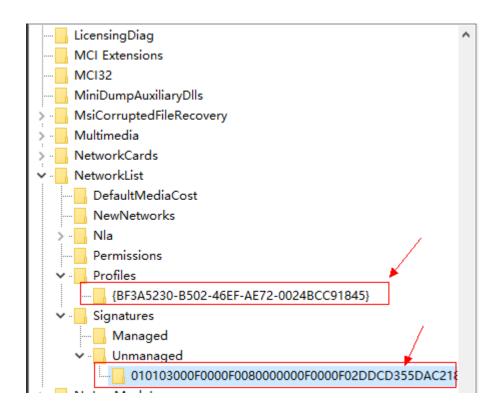
NIC flapping occurs because NIC retaining is enabled in the image from which the ECS is created.

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:

- 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:

 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.
- 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

16.19.9 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.

16.19.10 Why Is the NIC Not Working?

Symptom

The NIC equipped on a disk-intensive ECS does not work.

Possible Causes

The NIC driver has not been correctly installed.

Solution

Disk-intensive 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 16-24**.

| NIC Driver Version | OS | How to Obtain |
|-----------------------|--|---|
| ixgbevf 2.16.4 | CentOS 7.2 64bit | https://sourceforge.net/ |
| | Red Hat Enterprise Linux 7.2 64bit | projects/e1000/files/ ixgbevf%20stable/ 2.16.4/ |
| ixgbevf 2.16.1 | SUSE Linux Enterprise Server 11 SP3 64bit | https://sourceforge.net/ projects/e1000/files/ ixgbevf%20stable/ 2.16.1/ |
| | SUSE Linux Enterprise Server 11 SP4 64bit | |

Table 16-24 NIC driver versions and OSs

Step 2 Log in to the ECS.

For more details, see **Login Overview**.

- **Step 3** Install the passthrough NIC driver on the ECS. In this procedure, Red Hat Enterprise Linux 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 $\mathbf{eth}x$, 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-N/C_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**.

- 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 ixqbevf-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:

rmmod 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.
- 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

16.19.11 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.

16.19.12 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

```
Troot@lsw-centos65-0001 sysconfig]# service network restart
Shutting down interface eth0:
Shutting down loopback interface:
Bringing up loopback interface:
Bringing up interface eth0:
Determining IP information for eth0... done.
```

4. Run the following command to view routes:

route -n

```
root@lsw-centos65-0001 sysconfig]# route
Kernel IP routing table
Destination
169.254.169.254
                                                    Flags Metric Ref
                                                                          Use Iface
                                   255.255.255.255 UGH
                 192.168.1.1
                                                                              eth0
                                   255.255.255.0
                 192.168.1.34
92.168.2.0
                                                    UG
                                                                             0 eth0
92.168.1.0
                 0.0.0.0
                                   255.255.255.0
                                                                             0 eth0
   254.0.0
                 0.0.0.0
                                   255.255.0.0
                                                           1002
                                                                   И
                                                                             0 eth0
                                   Й.Й.
                                       . П. П
```

16.19.13 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 16-224 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 a result, the Linux ECS cannot obtain metadata.

Solution

1. Add the route with the destination of 169.254.169.254, and specify the next hop (gateway) and the output device (primary NIC of the Linux ECS). The following is an example:

ip route add 169.254.169.254 via 192.168.1.1 dev eth0

192.168.1.1 is the gateway address of the subnet that the primary NIC resides, and eth0 is the primary NIC.

2. Run the following command to verify that the metadata can be obtained:

curl http://169.254.169.254

Figure 16-225 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 [ ~ ] # ■
```

 Run the following command to create or modify the /etc/sysconfig/networkscripts/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

16.19.14 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.

| Possible Cause | Solution |
|--|--|
| The ECS is frozen or stopped, or has no EIP bound. | Check whether the ECS is running and has an EIP bound. For details, see Checking the ECS Status . |
| The ECS is overloaded. | Check whether the bandwidth and vCPU usage of the ECS are too high. For details, see Checking Whether the ECS Is Overloaded. |

| Possible Cause | Solution |
|--|--|
| The EIP bandwidth exceeds the bandwidth limit. | Increase the bandwidth and try again. For details, see Checking Whether the EIP Bandwidth Exceeded the Limit. |
| 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 ISP Network Is Functional. |
| The network configuration on the ECS is incorrect. | Check whether the NIC and DNS configurations are correct. For details, see Checking the NIC Configuration. |
| Routing is incorrectly configured. | Check whether the default route of 0.0.0.0 designates to the default gateway. For details, see Checking Whether the Default Route Is Destined for the Default Gateway. |
| The security group is incorrectly configured. | Check whether the security group allows the network traffic in the outbound direction. For details, see Checking Whether the Security Group Is Correctly Configured. |
| A network ACL has been associated with the ECS. | Disassociate the network ACL with the ECS and try again. For details, see Checking ACL Rules . |
| The EIP is blocked. | If the EIP is blocked, the ECS cannot access the Internet. For details, see Checking Whether the EIP Is Blocked. |
| The access is blocked by the firewall. | Disable the firewall and try again. For details, see Checking the Firewall Configuration. |
| The gateway is inaccessible. | Run the ping command to check whether the DNS server is running properly. For details, see Checking Whether the Gateway Is Accessible . |
| The ECS performance cannot meet service requirements. | Run the netstat command to check the network connection status. For details, see Checking the ECS Performance . |
| 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 ECS has been attacked by viruses or Trojan horses. | Check whether the ECS is affected by viruses or Trojan horses. For details, see Checking the ECS Security Status. |

Checking the ECS Status

- Check whether the ECS is in the **Running** state on the management console.
- Check whether an EIP has been bound to the ECS.
 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 using 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

If an EIP is bound to the ECS, the ECS can access the Internet through the bandwidth.

If Internet access fails, check whether the EIP bandwidth exceeds the bandwidth 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.
 - 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 16-226.

Figure 16-226 NIC and DNS configurations

```
_ 0
                                          Administrator: Command Prompt
                             Corporation. All rights reserved.
    .Users\Administrator>ipconfig /all
Windows IP Configuration
     ecs-1d6c
                                                            Hybrid
No
No
openstacklocal
   hernet adapter Ethernet 2:
      onnection-specific DNS Suffix
                                                             openstacklocal
Red Hat VirtIO Ethernet Adapter
FA_16-3E-4E-3B-BB
      onnection spectors
escription
hysical Address
HCP Enabled
utoconfiguration Enabled
ink-local IPu6 Address
                                                               es
e88::fcf3:e79a:b7e0:5bb9x14(Preferred)
92.168.10.210(Preferred)
52.255.255.
ednesday, September 9, 2020 2:29:34
hursday, September 10, 2020 10:29:33
                                                                 -01-00-01-24-E6-B9-71-FA-16-3E-0B-DF-38
                                                             100.125.1.250
100.125.129.250
Enabled
    DNS Servers .
    NetBIOS over Tcpip. . . . .
```

- 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.

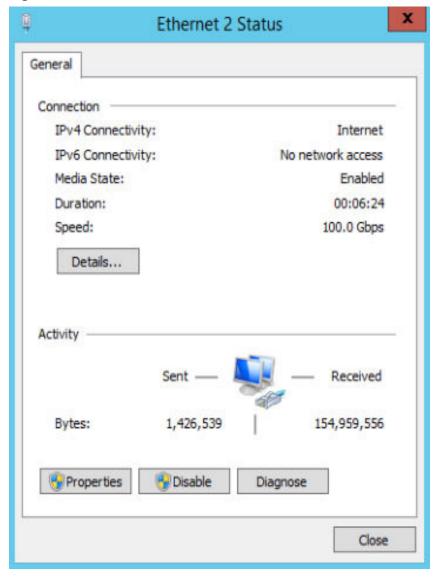


Figure 16-227 NIC status

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.

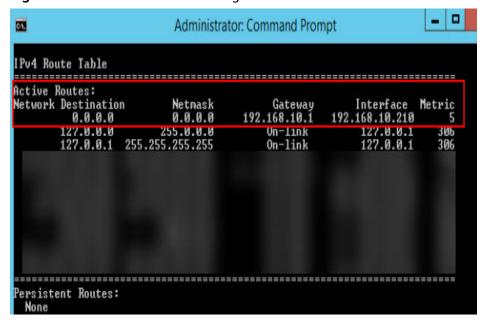


Figure 16-228 Default route settings

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.
- Disassociate the network ACL from the subnet of the ECS.
 On the page providing details about the network ACL, choose Associated Subnets > Disassociate.

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 16-229 Windows Firewall



3. Choose Check firewall status > Turn Windows Firewall on or off. View and set the firewall status.

Figure 16-230 Turn off Windows Firewall



Checking Whether the Gateway Is Accessible

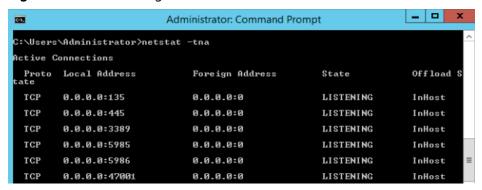
- 1. Run the **ping** command to check whether data can be exchanged between the ECS and the gateway.
 - Use an IP address in a different network segment to ping the gateway to check network connections.
- 2. Run the **ping** command to obtain the IP address of the DNS server. Compare the time required 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 16-231 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.

16.19.15 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 16-26 Possible causes and solutions

| Possible Cause | Solution |
|--|--|
| The ECS is frozen or stopped, or has no EIP bound. | Check whether the ECS is in the Running state and has an EIP bound. For details, see Checking the ECS Status . |

| Possible Cause | Solution |
|---|---|
| The ECS is overloaded. | Check whether the bandwidth and vCPU usage of the ECS are too high. For details, see Checking Whether the ECS Is Overloaded. |
| The EIP bandwidth exceeds the bandwidth limit. | Increase the bandwidth and try again. For details, see Checking Whether the EIP Bandwidth Exceeded the Limit. |
| The DNS configuration is incorrect. | Change the DNS server to a private one. For details, see Checking the DNS Configuration. |
| Specified resolution has been configured in the hosts file. | Check whether the mappings in the hosts configuration file are correct. For details, see Checking the hosts Configuration File. |
| Both Network and NetworkManager are enabled. | Use either of the two tools to prevent incompatibility issues. For details, see Checking Whether Both Network and NetworkManager Have Been Enabled. |
| The security group is incorrectly configured. | Check whether the security group allows the network traffic in the outbound direction. For details, see Checking Whether the Security Group Is Correctly Configured. |
| A network ACL has been associated with the ECS. | Disassociate the network ACL with the ECS and try again. For details, see Checking ACL Rules . |
| The EIP is blocked. | If the EIP is blocked, the ECS cannot access the Internet. For details, see Checking Whether the EIP Is Blocked. |
| The private IP address is lost. | Check whether the dhclient process is running. If it is not running, the private IP address may be lost. For details, see Checking Whether a Private IP Address Can Be Obtained . |
| NICs are incorrectly configured. | Check whether the NIC and DNS configurations are correct. For details, see Checking the NIC Configuration. |
| Firewall is enabled on the ECS. | Disable the firewall and try again. For details, see Checking the Firewall Configuration. |

Checking the ECS Status

- Check whether the ECS is in the **Running** state on the management console.
- Check whether an EIP has been bound to the ECS.
 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 using 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

If an EIP is bound to the ECS, the ECS can access the Internet through the bandwidth.

If Internet access fails, check whether the EIP bandwidth exceeds the bandwidth 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.

Run the following command to view the DNS configuration:

cat /etc/resolv.conf

If the command output shown in **Figure 16-232** is displayed, the domain name is resolved using the private DNS server.

Figure 16-232 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.
- 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

 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



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.

16.19.16 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.

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 16-27 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 Checking Port Communication. |
| Security group rules | Check whether the access to the port is allowed in the security group of the ECS. For details, see Checking Security Group Rules. |
| Firewall configuration | Disable the firewall and try again. For details, see Checking the Firewall Configuration. |
| Route configuration | Check whether the gateway configurations in the ECS route table are correct. For details, see Checking the ECS Route Configuration. |
| Local network | Check whether you can use another hotspot or network to access the website. For details, see Checking the Local Network. |
| CPU usage | Identify and optimize the processes leading to high vCPU usage. For details, see Checking the CPU usage. |

Checking Port Communication

Ensure that service processes and ports are in **LISTEN** state. **Table 16-28** lists the common TCP statuses.

Linux

Run the **netstat -antp** command to check whether the port used by the target website is in **LISTEN** state.

For example, run netstat -ntulp | grep 80.

Figure 16-233 Checking port listening status



- 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 16-234 Checking port listening status

```
C:\Users\Administrator>netstat -ano |findstr "80"
         0.0.0.0:80
0.0.0.0:49155
                                   0.0.0.0:0
0.0.0.0:0
 TCP
                                                              LISTENING
 TCP
                                                              LISTENING
                                                                                880
 TCP
                                    [::]:0
                                                              LISTENING
          [::]:80
 TCP
          [::]:49155
                                                              LISTENING
                                                                                889
                                    T::1:0
         0.0.0.0:123
 UDP
                                    *:*
                                                                                808
 UDP
          [::1:123
                                                                                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 16-28 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 of time, 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.

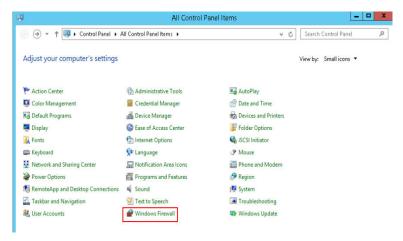
- 1. Log in to the management console.
- 2. Under Computing, choose Elastic Cloud Server.
- 3. On the ECS list, click the ECS for which you want to change the security group rules.
- 4. On the **Security Groups** tab, view security group rules.
- 5. Click Modify Security Group Rule.
- 6. Configure the rule to allow the access of the port used by the website.

Checking the Firewall Configuration

Linux ECS

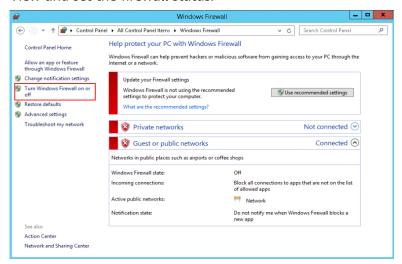
The following uses port 80 and CentOS 6.8 as an example.

- 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.

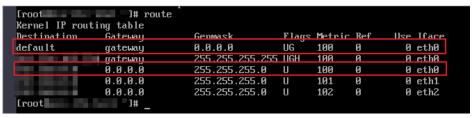
View and set the firewall status.



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.



b. Run the **ifconfig** or **ip addr** command to obtain the ECS IP address.

Figure 16-235 ifconfig command output

```
[root@elb-mq02 ~]# ifconfig -a
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 192.168.0.145    netmask 255.255.255.0    broadcast 192.168.0.255
    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,L00PBACK,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 16-236 ip addr command output

```
[root@elb-mq02 ~]# ip addr
1: lo: <L00PBACK,UP,L0WER_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
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:16:30:24:1e:7f brd ff:ff:ff:ff:ff:
inet 192.168.0.145/24 brd 192.168.0.255 scope global noprefixroute dynamic eth0
    valid_lft 77109sec preferred_lft 77109sec
    inet6 f880::f816:3eff:fe24:le7f/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 16-237 route -n command output

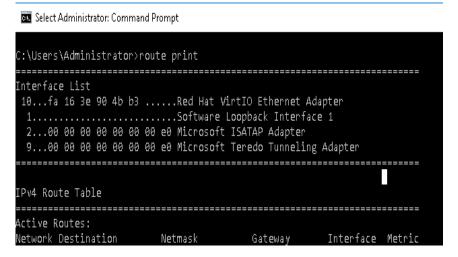
```
[root@elb-mg02 ~]# route
Kernel IP routing table
Destination Gateway
0.0.0.0 192.168.0
                                                                 Flags Metric Ref
                                                                                             Use Iface
                                           Genmask
                                           0.0.0.0
255.255.255.255
                                                                 UG
                                                                          100
                                                                                               0 eth0
                                                                                   Θ
                                                                                   Θ
                     192.168.0.1
                                                                 UGH
                                                                          100
                                                                                                0
                                                                                                  eth0
```

- Windows ECS
 - a. Run cmd.exe.
 - b. Run the **ipconfig** command to obtain the ECS IP address.

Figure 16-238 ipconfig command output

c. Run the **route print** command to obtain the gateway in the routing table.

Figure 16-239 route print command output



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 created an alarm rule using Cloud Eye, the system automatically sends an alarm to you when the bandwidth or CPU usage reaches the threshold specified in the rule.

- 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 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.

16.19.17 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=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
```

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
```

ping: sendmsg: Invalid argument

- 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
```

16.19.18 How Can I Obtain the MAC Address of My ECS?

This section describes how to obtain the MAC address of an ECS.

The MAC address of an ECS cannot be changed.

Linux (CentOS 6)

- 1. Log in to the Linux ECS.
- Run the following command to view the MAC address of the ECS: ifconfig

Figure 16-240 Obtaining the MAC address

```
[root@Cent0S68-XEN ~]# ifconfig
           Link encap:Ethernet HWaddr FA:16:3E:2A:36:DE
eth0
          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.
- Run the following command to view the MAC address of the ECS: ifconfig

Figure 16-241 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
               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
```

Run the following command to view the MAC address of NIC eth0: ifconfig eth0 |egrep "ether"

Figure 16-242 Obtaining the MAC address of eth0

4. Obtain the returned MAC address.

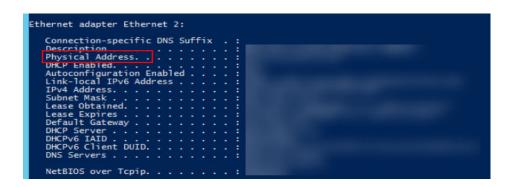
ifconfig eth0 |egrep "ether" |awk '{print \$2}'

Figure 16-243 Obtaining the MAC address of eth0

```
[root@ecs-683a ~]# ifconfig eth0 |egrep "ether" |awk '{print $2}'
fa:16:3e:c3:46:fc
[root@ecs-683a ~]#
```

Windows

- 1. Press Win+R to start the Run text box.
- 2. Enter cmd and click OK.
- Run the following command to view the MAC address of the ECS: ipconfig /all



16.19.19 How Can I Test Network Performance?

Use netperf and iperf3 to test network performance between ECSs. The test operations include preparations, Transmission Control Protocol (TCP) bandwidth test, User Datagram Protocol (UDP) packets per second (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 16-29 and Table 16-30 list the common netperf and iperf3 parameters.

Table 16-29 Common netperf parameters

| Parameter | Description |
|-----------|--|
| -р | Port number |
| -H | IP address of the RX end |
| -t | Protocol used in packet transmitting, the value of which is TCP_STREAM in bandwidth tests |
| -l | Test duration |
| -m | Data packet size, which is suggested to be 1440 in bandwidth tests |

Table 16-30 Common iperf3 parameters

| Parameter | Description |
|-----------|--------------------------|
| -р | 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 16 in PPS tests |
| -A | ID of the vCPU used by iperf3 In this section, the maximum number of 16 vCPUs is used as an example for each ECS. If an ECS has 8 vCPUs, the -A 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 16-31 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 2.168.2.18 |

Step 2 Install the netperf, iperf3, and sar test tools on both the tested ECS and auxiliary ECSs.

Table 16-32 lists the procedures for installing these tools.

Table 16-32 Installing test tools

| Tool | Procedure |
|---------|---|
| netperf | Run the following command to install gcc: yum -y install unzip gcc gcc-c++ |
| | Run the following command to download the netperf installation package: wget https://github.com/HewlettPackard/netperf/archive/ refs/tags/netperf-2.7.0.zip |
| | Run the following commands to decompress the installation package and install netperf: unzip netperf-2.7.0.zip |
| | cd netperf-netperf-2.7.0/ |
| | ./configure && make && make install |
| iperf3 | Run the following command to download the iperf3 installation package: wgetno-check-certificate https://codeload.github.com/esnet/iperf/zip/master -O iperf3.zip |
| | Run the following commands to decompress the installation package and install iperf3: unzip iperf3.zip |
| | cd iperf-master/ |
| | ./configure && make && make install |
| sar | Run the following command to install sar: |
| | yum -y install sysstat |

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.

Ⅲ 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.

 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 16-29.

##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
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 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 --- -m 1440 & netperf -H 192.168.2.10 -p 120
```

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 16-244**. The final result is the sum of the test results of the netperf processes on all TX ends.

Figure 16-244 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.

Run the following commands on all auxiliary ECSs to start the server process: 2. iperf3 -s -p 12001 & iperf3 -s -p 12002 & In the preceding commands, **-p** specifies the listening port. Start the client process on the tested ECS. For details about common iperf3 parameters, see Table 16-30. ##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.

Start the server process on the tested ECS. For details about common iperf3 parameters, see Table 16-30.

##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 &
Start the client process on all auxiliary ECSs. For details about common iperf3
parameters, see Table 16-30.
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 &
```

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 5.

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 16-245 shows an example of the UDP PPS test result.

Figure 16-245 UDP PPS test result

```
297.00-298.00 sec
298.00-299.00 sec
299.00-300.00 sec
                                                                       Mbits/sec
Mbits/sec
                                                                                           Jitter
                                         4.43 MBytes
4.31 MBytes
                                                                                          70622
                                                               36.2 Mbits/sec
[ ID] Interval
[ 5] 0.00-30
[ 5] 0.00-30
                                          Transfer
                                                               Bitrate
                                                                                                          Lost/Total Datagrams
0/19403434 (0%) sender
87/19403434 (0.00045%)
             0.00-300.00 sec 1.16 GBytes
0.00-300.05 sec 1.16 GBytes
                                                               33.1 Mbits/sec
33.1 Mbits/sec
                                                                                          0.000 ms
0.019 ms
iperf Done.
                                                                                                                                    Packet loss rate
       Test duration 300s Total TX traffic
                                                                                                          Lost packets/
                                                                  Average tranmission rate
                                                                                                            Total packets
```

Ⅲ 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:

gperf 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

16.19.20 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.

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

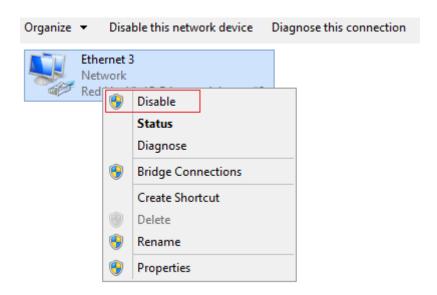
BOOTPROTO=static IPADDR=192.168.1.100 #IP address (modified) NETMASK=255.255.255.0 #Mask (modified) GATEWAY=192.168.1.1 #Gateway IP address (modified)

- 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.

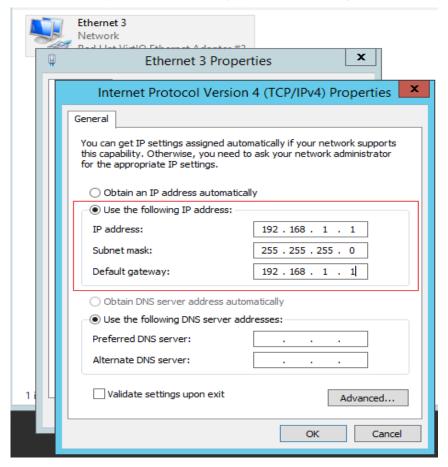
Solution (Windows)

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.

16.19.21 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 16-33**.
- Back up key ECS data before modifying kernel parameter settings.

Background

Table 16-33 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_bac klog | 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 | Uses the TCP stack to show memory usage in memory pages (4 KB generally). |
| | The first value is the lower limit of memory usage. |
| | The second value is the upper limit of the load added to the buffer when the memory is overloaded. |
| | 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. |
| net.ipv4.tcp_rmem | Specifies the memory used by sockets for automatic optimization. |
| | The first value is the minimum number of bytes allocated to the socket buffer for receiving data. |
| | The second value is the default value, which is overwritten by rmem_default . The buffer size can increase to this value when the system load is not heavy. |
| | The third value is the maximum number of bytes allocated to the socket buffer for receiving data. This value is overwritten by rmem_max . |
| net.ipv4.tcp_wmem | Specifies the memory used by sockets for automatic optimization. |
| | The first value is the minimum number of bytes allocated to the socket buffer for transmitting data. |
| | The second value is the default value, which is overwritten by wmem_default . The buffer size can increase to this value when the system load is not heavy. |
| | The third value is the maximum number of bytes allocated to the socket buffer for transmitting data. This value is overwritten by wmem_max . |
| net.ipv4.tcp_keepalive_ti me | Specifies the interval at which keepalive detection messages are sent in seconds for checking TCP connections. |
| net.ipv4.tcp_keepalive_int vl | Specifies the interval at which keepalive detection messages are resent in seconds when no response is received. |
| net.ipv4.tcp_keepalive_pr obes | Specifies the maximum number of keepalive detection messages that are sent to determine a TCP connection failure. |

| Parameter | Description |
|---------------------------------|--|
| net.ipv4.tcp_sack | Enables selective acknowledgment (value 1 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_scali ng | Enables RFC1323-based window scaling by setting the parameter value to 1 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 (syncookie). This configuration prevents socket overloading when a large number of connections are attempted to set up. CONFIG_SYN_COOKIES must be enabled in the kernel for compilation. The default value is 0, indicating that TCP synchronization is disabled. |
| net.ipv4.tcp_tw_reuse | Specifies whether a TIME-WAIT socket (TIME-WAIT port) can be used for new TCP connections. NOTE This parameter is valid only for clients and takes effect only when net.ipv4.tcp_timestamps is enabled. This parameter cannot be set to 1 if NAT is enabled. Otherwise, an error will occur in remote ECS logins. |
| net.ipv4.tcp_tw_recycle | Allows fast recycle of TIME-WAIT sockets. NOTE This parameter is valid only when net.ipv4.tcp_timestamps is enabled. Do not set this parameter to 1 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 FIN-WAIT-2 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_ran ge | Specifies local port numbers allowed by TCP/UDP. |
| net.ipv4.tcp_max_syn_bac klog | 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 1024 . 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_buc kets | Specifies the number of TIME_WAIT buckets, which defaults to 180000 . If the number of buckets exceeds the default value, extra ones will be cleared. |
| net.ipv4.tcp_synack_retrie s | Specifies the number of times that SYN+ACK packets are retransmitted in SYN_RECV state. |
| net.ipv4.tcp_abort_on_ove rflow | When this parameter is set to 1, 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. Default value: 0 |
| 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.tcp_keepalive_time = 1000
net.ipv4.tcp_max_syn_backlog = 8192
net.ipv4.tcp_max_syn_backlog = 8192
net.ipv4.tcp_mem = 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 <code>/proc/sys/net/ipv4/tcp_tw_recycle</code> 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"

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

16.19.22 How Can I Configure Port Redirection?

Requirement

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.

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

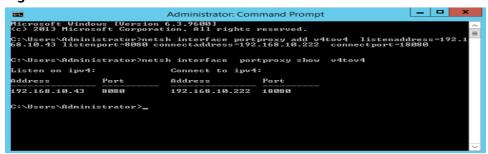
To cancel port redirection, run the following command:

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 16-246 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 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. Therefore, 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.
- 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 16-247 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 16-248 Logging in to ECS 2



----End

16.19.23 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 | Free of charge | Use VPC peering to enable the communication over an intranet. | |
| region | | VPC Peering Connection Creation Procedure | |
| | | Creating a VPC Peering Connection with a VPC in Another Account | |
| In the Billed same | | Use VPC Endpoint to enable the communication over an intranet. | |
| region | | What Is VPC Endpoint? | |
| | | Configuring a VPC Endpoint for Communication Across VPCs of Different Domains | |
| | | What Are the Differences Between VPC Endpoints and VPC Peering Connections? | |
| In different regions | Billed | Use VPN to enable the communication over an intranet. • Virtual Private Network | |
| | | Does a VPN Allow for Communication Between Two VPCs? | |

16.19.24 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.

16.20 Security Configurations

16.20.1 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.

For details, see Installing an Agent on a Linux Server and Enabling the Basic/Enterprise/Premium Edition.

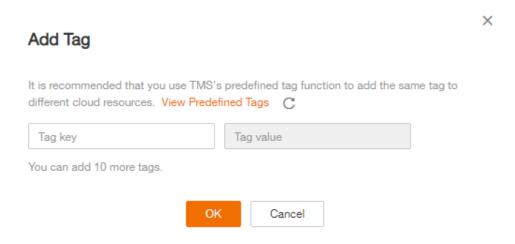
16.21 Resource Management and Tags

16.21.1 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.
- 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 16-249 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 16-250 Searching for ECSs by tag



4. In the search bar, choose **Tag** and then select the tag key and value, and click **Search**.

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 16-251 Deleting a tag



16.22 Resource Monitoring

16.22.1 Why Is My Windows ECS Running Slowly?

If your ECS runs slowly or is inaccessible unexpectedly, the bandwidth or vCPU usage of the ECS may be excessively high. If you have created an alarm rule using Cloud Eye, the system automatically sends an alarm to you when the bandwidth or CPU usage reaches the threshold specified in the rule.

To handle this issue, perform the following operations:

- 1. Fault locating:
 - Identify the drivers from unknown sources and processes leading to high bandwidth or CPU usage.
 - Windows offer multiple tools to locate faults, including Task Manager, Performance Monitor, Resource Monitor, Process Explorer, Xperf (supported by versions later than Windows Server 2008), and full memory dump.
- 2. Check whether the processes and drivers are normal and handle the issue accordingly.
 - If the processes are not malicious, optimize their programs or modify ECS specifications.

- If the processes are malicious, stop these processes manually or use a third-party tool to stop them automatically.
- If the drivers are from official sources, there is no need to deal with system built-in drivers. Determine whether to uninstall the third-party software based on your requirements.
- If the drivers are from unknown sources, you are advised to uninstall them by using commercial antivirus software or third-party security management tools.

Fault Locating

- 1. Log in to the ECS using VNC available on the management console.
- 2. Start the **Run** dialog box, and then enter **perfmon -res**.

Figure 16-252 Starting the Resource Monitor



3. On the **Resource Monitor** page, click the **CPU** or **Network** tab to view the CPU or bandwidth usage.

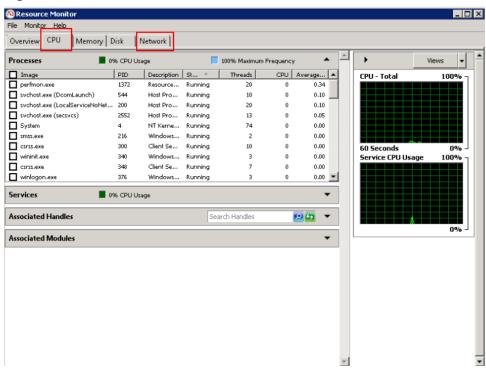


Figure 16-253 Resource Monitor

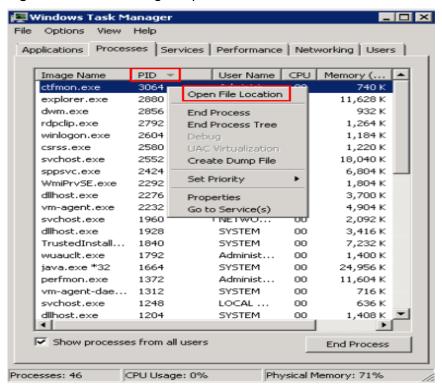
- 4. Obtain the IDs and names of the processes with high CPU or bandwidth usage.
- 5. On the remote login page, click **Ctrl+Alt+Del** to start the **Windows Task Manager**.

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 16-254 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.

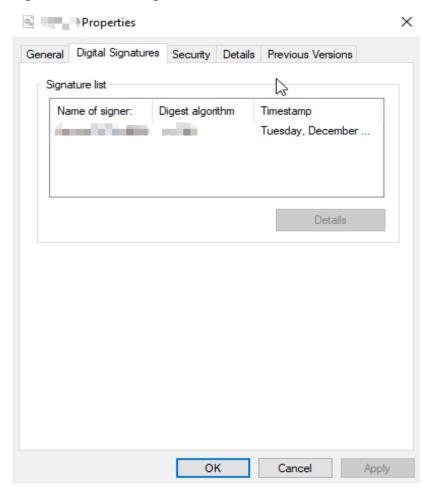
Figure 16-255 Viewing the system drivers

| Filter Name | Num Instances | Altitude | Frame |
|-------------|---------------|----------|-------|
| WdFilter | 3 | 328010 | 0 |
| storgosflt | 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.

- 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 16-256 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.

16.22.2 Why Is My Linux ECS Running Slowly?

If your ECS runs slowly or is inaccessible unexpectedly, the bandwidth or vCPU usage of the ECS may be excessively high. If you have created an alarm rule using Cloud Eye, the system automatically sends an alarm to you when the bandwidth or CPU usage reaches the threshold specified in the rule.

If your ECS runs slowly or is disconnected suddenly, the possible causes are as follows:

Your ECS is a shared ECS.

Multiple ECSs share CPU resources. When resources are insufficient, ECSs may contend for CPU resources, causing slow responses.

If your ECS is a shared ECS, perform the following steps:

- 1. Fault locating: Check the instance type. For details about dedicated and shared ECSs, see **ECS Types**.
- 2. Troubleshooting: If you have high requirements on service stability, you are advised to change a shared ECS to a dedicated ECS by referring to **General Operations for Modifying Specifications**.

To handle this issue, perform the following operations:

- Fault locating
 Identify the processes leading to high bandwidth or CPU usage.
- 2. Check whether the processes are normal and handle the issue accordingly.
 - If the processes are normal, optimize them or modify ECS specifications.
 - If the processes are malicious, use a third-party tool to automatically stop the processes or manually stop them.

Common Commands

The following uses the CentOS 7.2 64bit OS as an example to describe common commands. The commands may vary depending on Linux OS editions. For details, see the official documentation for the specific OS edition.

The common commands for checking Linux ECS performance metrics, such as the CPU usage, are as follows:

- ps -aux
- ps -ef
- top

Locating High CPU Usage

- 1. Log in to the ECS using VNC.
- 2. Run the following command to check the OS running status:

top

Information similar to the following is displayed.

| top - | 20:56 | 5:02 up 3 | 37 da | ys, 9:09 | 9, 1 us | er, | loa | ıd aver | age: | 0.00, 0.0 | l, 0.05 |
|----------|-------|------------------------|-------|----------|----------|-------|-----|-----------------|------|---------------------|---------------|
| | | | | unning, | | | | | | | |
| 2.Cpu (s | s): (| a.2 us, | 0.3 | sy, 0.0 | ni, 99. | 5 id, | E |).0 wa, | 0.0 | 3 hi, 0.0 | si, 0.0 st |
| KiB Me | em : | 3880024 | tota | 1, 2963 | 304 free | . 17 | 183 | 1 84 use | ed, | 738336 but | f/cache |
| KiB Su | ωap: | 0 | tota | ι1, | 0 free | , | | 0 use | d. | 3 434808 avi | ail Mem |
| | | | | | | | | | | | |
| PID | USER | PR | NI | VIRT | RES | SHR | S | иCPU | >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 | ksoftirgd/0 |
| 5 | root | 0 | -20 | 0 | 0 | 0 | S | 0.0 | 0.0 | 0:00.00 | kworker/0:0H |
| 7 | root | $\mathbf{r}\mathbf{t}$ | 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 | |
| 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.

- 1. On the **top** page, enter **q** or press **Ctrl+C** to exit.
- 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.
- Run the ll /proc/PID/exe command to obtain the program file specified by a PID.

```
[root@elb-mq01 sysconfig]# 11 /proc/4243/exe
| lrwxrwxrwx 1 root root 0 Mar 18 11:46 /proc/4243/exe -> /CloudBesetPwdUpdateAgent/depend/jre1.8.0_131/bin/java
```

Troubleshooting High CPU Usage

If the processes leading to high CPU usage are malicious, run the top command to stop them. If the **kswapd0** process leads to high CPU usage, optimize the program for the process or upgrade the ECS specifications for a larger memory capacity.

kswapd0 is a virtual memory management process. When the physical memory becomes insufficient, **kswapd0** runs to allocate disk swap capacity for caching. This uses a large number of CPU resources.

- For the detected malicious processes
 - Quickly stop such processes on the top page. To do so, perform the following operations:
 - a. Press the \mathbf{k} key during the execution of the top command.
 - b. Enter the PID of the process to be stopped.

The PID of the process is the value in the first column of the top command output. For example, to stop the process with PID 52, enter **52** and press **Enter**.

```
sks: 81 total, 1 running, 79 sleeping,
ou(s): 0.0 us, 3.2 sy, 0.0 ni, 96.8 id,
                                                      1 stopped,
                                                                      0 zombie
                                                               0.0 hi, 0.0 si,
                                                     0.0 wa,
           3880024 total,
                              2961520 free,
                                                  178960 used,
                                                                   739544 buff/cache
                     total,
                                      И
                                                        0 used.
                                                                  3434216 avail Mem
PID to signal/kill [default pid
                                     = 11 52
  PID USER
                                                       %CPU %MEM
```

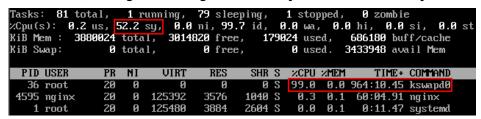
c. After the operation is successful, information similar to the following is displayed. Press **Enter**.



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.
- b. If the process remains in non-sleeping state for a long period of time, you can preliminarily determine that the system is consistently paging. In such a case, the high CPU usage is caused by insufficient memory.



c. Run the **vmstat** command to check the virtual memory usage of the system.

If the **si** and **so** values are large, the system is frequently paging and the physical memory of the system is insufficient.

- **si**: Volume of data written from the swap partition to the memory per second, which is transferred from the disk to the memory.
- **so**: Volume of data written from the memory to the swap partition per second, which is transferred from the memory to the disk.
- d. Further identify the causes of high memory usage. Run commands, such as **free** and **ps** to check the memory usage of the system and processes in the system.
- e. Restart the application or release the memory when traffic is light.

 To handle this issue, expand the ECS memory. If memory expansion is not allowed, optimize the application and enable hugepage memory.

Handling High Bandwidth Usage

If the high bandwidth usage is caused by normal service access of non-malicious processes, enlarge the bandwidth to handle this issue. If the high bandwidth usage is caused by abnormal service access, for example, malicious access from certain IP addresses, CC attacks on the ECS, or malicious processes, use the traffic monitoring tool **nethogs** to monitor the bandwidth usage of each process in real time and identify faulty processes.

- Using nethogs for troubleshooting
 - a. Run the following command to install nethogs:yum install nethogs -y

After the installation, run the **netgos** command to check bandwidth usage.

Parameters in the **nethogs** command are as follows:

- -d: Set the update interval in the unit of second. The default value is
 1.
- **-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



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.

16.23 Database Applications

16.23.1 Can a Database Be Deployed on an ECS?

Yes. You can deploy a database of any type on an ECS.

16.23.2 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.

16.23.3 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

WinVMDataDiskAutoInitialize.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:

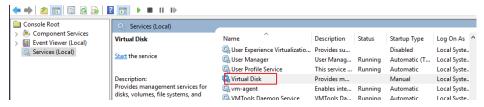
- Whether the I/O operation is a read or write request
- The offset within the file where the I/O operation was attempted
- The database file against which the I/O operation was performed
- The operating system error code and error description in parentheses

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 16-257 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 WinVMDataDiskAutoInitialize.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 **WinVMDataDiskInitialize** in the scheduled task list, and choose **End**.

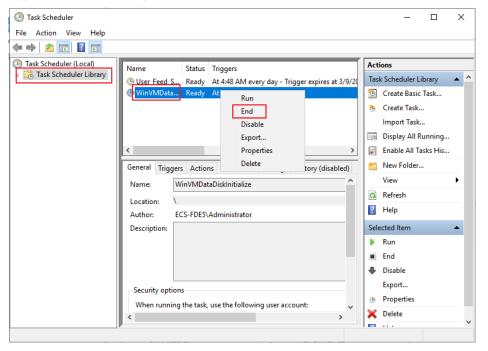
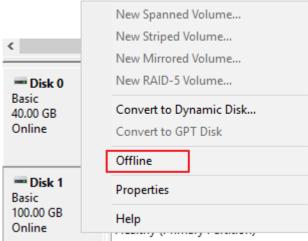


Figure 16-258 Ending WinVMDataDiskInitialize

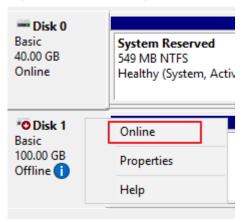
- 4. Restart the ECS or take the data disk offline and then online.
 - Open the Run dialog box, enter diskmgmt.msc, and press Enter. The Disk Management window is displayed.
 - b. Right-click the block to which the disk belongs and choose **Offline**.

Figure 16-259 Setting disk offline



c. Right-click the block to which the disk belongs and choose **Online**.

Figure 16-260 Setting disk online



A Change History

| Released On | Description |
|-------------|---|
| 2023-10-30 | This issue is the forty-fifth official release. Added the following content: Added m6.large.16, m6.xlarge.16, m6.2xlarge.16, m6.3xlarge.16, m6.4xlarge.16, m6.6xlarge.16, and m6.8xlarge.16 in Memory-optimized ECSs. |
| 2023-08-30 | This issue is the forty-fourth official release. Added Billing. |
| 2023-07-30 | This issue is the forty-third official release. Modified the following content: • Added D6 ECS specifications in Disk-intensive ECSs. • Added G6 ECS specifications in GPU-accelerated ECSs. |

| Released On | Description |
|-------------|---|
| 2023-04-30 | This issue is the forty-second official release. |
| | Added the following content: |
| | A Summary List of ECS Specifications |
| | Access Control for ECS |
| | Fault Recovery |
| | Discontinued ECS Specifications |
| | Permissions Management |
| | Viewing Failed Tasks |
| | Starting and Stopping ECSs |
| | Uninstalling a GPU Driver from a GPU-accelerated ECS |
| | Overview |
| | Overview |
| | Dynamically Assigning IPv6 Addresses |
| | Overview |
| | Changing an EIP |
| | Methods for Improving ECS Security |
| | • HSS |
| | Permissions Management |
| | • OS Monitoring Metrics Supported by ECSs with the Agent Installed |
| | Process Monitoring Metrics Supported by ECSs with the Agent Installed |
| | Common Topics |
| | What Is an AZ? |
| | What Is a Region? |
| | Are Products Different in Different Regions? |
| | Is Data Transmission Between AZs Billed? |
| | • Can a Load Balancer Distribute Traffic to ECSs in Different Regions? |
| | • Is Application Disaster Recovery Available in Different Regions? |
| | • Are There Any Services Provided for Application Disaster Recovery? |
| | Can Components Contained in an Application Be Distributed to Different Regions? |
| | • What Should I Do If the ECS Resources to Be Purchased Are Sold Out? |
| | What Is the Creation Time and Startup Time of an ECS? |
| | When Does an ECS Become Provisioned? |
| | Why Cannot I View the ECSs Being Created Immediately After I Pay for Them? |

| Released On | Description |
|-------------|--|
| | What Do I Do If I Selected an Incorrect Image for My ECS? |
| | How Can I Manage ECSs by Group? |
| | • Why Did I Fail to Configure an Anti-Affinity ECS Group? |
| | • Can a Deleted ECS Be Restored? |
| | How Do I Delete or Restart an ECS? |
| | What Are the Login Requirements for ECSs? |
| | • What Are the Username and Password for Remote Logins? |
| | What Should I Do If Starting an ECS Remains in "Waiting for cloudResetPwdAgent" State? |
| | Why Are Garbled Characters Displayed When I Log In to My ECS Using VNC? |
| | What Should I Do If Error Code 1006 or 1000 Is Displayed When I Log In to an ECS Through the Management Console? |
| | Why No Audio File Can Be Properly Played on My Windows ECS Logged In Using VNC? |
| | How Can I Change the Resolution of a Windows ECS? |
| | Why Can't I Log In to My Windows ECS? |
| | Why Can't I Log In to My Linux ECS? |
| | • Why Cannot I Use a Non-Default SSH Port to Log In to My Linux ECS? |
| | Why Does the System Display a Message Indicating Invalid Credentials When I Attempt to Access a Windows ECS? |
| | Why Does an Internal Error Occur When I Log In to My Windows ECS? |
| | Why Is the Hostname of My ECS Restored to the Original Name After the ECS Is Restarted? |
| | How Can I Set Sequential ECS Names When Creating Multiple ECSs? |
| | How Can I Modify ECS Specifications? |
| | How Can I Recover a Windows ECS with an Abnormal Virtualization Driver? |
| | Why Is My Windows ECS Muted? |
| | How Do I Change an ECS SID? |
| | Why Does a Pay-per-Use ECS Fail to Be Started? |
| | Does OS Change Incur Fees? |
| | Will I Lose My Disk Data If I Reinstall ECS OS, Change the OS, or Change the ECS Specifications? |
| | Does OS Reinstallation Incur Fees? |
| | Do ECSs Support GUI? |
| | How Can I Install a GUI on an ECS Running CentOS 6? |

| Released On | Description |
|-------------|---|
| | How Can I Install a GUI on an ECS Running CentOS 7? |
| | How Can I Install a GUI on an ECS Running Ubuntu? |
| | How Can I Install a GUI on an ECS Running Debian? |
| | Why Does the OS Fail to Respond When kdump Occurs on a Linux ECS? |
| | How Can I Upgrade the Kernel of a Linux ECS? |
| | Why Cannot My ECS OS Start Properly? |
| | How Can I Enable SELinux on an ECS Running CentOS? |
| | Why Does a Forcibly-Stopped Linux ECS Fail to Be Restarted? |
| | How Do I View the GPU Usage of a GPU-accelerated ECS? |
| | • How Can I Transfer Data Between a Local Computer and a Windows ECS? |
| | Why Does Internet Access to an ECS Deployed with FTP Fail? |
| | How Can I Check Whether the ECSs Attached with the Same Shared SCSI Disk Are in the Same ECS Group? |
| | How Can I Add an ECS with Local Disks Attached to an ECS Group? |
| | How Can I Rectify the Fault That May Occur on a Linux ECS with an NVMe SSD Disk Attached? |
| | • Why Is the Device Name of My C6 ECS in the sd* Format? |
| | Why Are Disk Error Logs Printed After a Disk Attached to an ECS Is Formatted with the ext4 File System? |
| | • How Can I Change the Password for Logging In to a Linux ECS? |
| | Changing the Login Password on an ECS |
| | Resetting the Password for Logging In to a Windows ECS |
| | Resetting the Password for Logging In to a Linux ECS |
| | What Should I Do If the System Displays a Message Indicating that the Password Is Incorrect When I Remotely Log In to My ECS? |
| | 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? |
| | Disabling SELinux |
| | How Can I Use a Key Pair? |
| | What Should I Do If I Cannot Download a Key Pair? |
| | How Do I Query the Egress Public IP Address of My ECS? |
| | Configuring DNS |
| | What Should I Do If NIC Flapping Occurs After My ECS Specifications Are Modified? |
| | How Can I Add a Static Route to a CentOS 6.5 OS? |

| Released On | Description |
|-------------|--|
| | Why Can't My Windows ECS Access the Internet? Why Does My Linux ECS Fail to Access the Internet? |
| | How Do I Troubleshoot an Unresponsive Website Hosted on My ECS? |
| | Why Did I See "Invalid argument" or "neighbour table overflow" During an Access to a Linux ECS? |
| | How Is ECS Security Ensured? |
| | 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? |
| | Modified the following content: |
| | Put S1 and C1 ECS specifications in General-Purpose ECSs offline. |
| | Put D1 ECS specifications in Disk-intensive ECSs offline. |
| | Put G2 and G1 ECS specifications in GPU-accelerated ECSs offline. |
| | Put high-performance computing and computing II specifications offline. |
| 2023-02-28 | This issue is the thirteenth official release. |
| | Permissions Management |
| | Permissions Management |
| 2022-07-25 | This issue is the twelfth official release. |
| | Modified the following content: |
| | Added s6.4xlarge.2 and s6.4xlarge.4 in General-Purpose ECSs . |
| 2022-05-06 | This issue is the eleventh official release. |
| | Modified the following content: |
| | General-Purpose ECSs |
| | Dedicated General-Purpose ECSs |
| | Memory-optimized ECSs |
| | 1.9.33 High-Performance Computing ECSs |
| | GPU-accelerated ECSs |

| Released On | Description |
|-------------|--|
| 2022-03-23 | This issue is the tenth official release. Added the following content: Changing a Xen ECS to a KVM ECS (Windows) Automatically Changing a Xen ECS to a KVM ECS (Linux) Automatically Changing a Xen ECS to a KVM ECS (Linux) What Should I Do If Executing a Driver Installation Script Failed on an ECS Running CentOS 5? What Should I Do If Executing a Driver Installation Script Failed When I Attempted to Modify the Specifications of a Linux ECS? |
| 2021-07-22 | This issue is the ninth official release. Modified the following content: Deleted section "Modifying a Private IP Address". |
| 2021-06-23 | This issue is the eighth official release. Added: ECS Advantages ECS Application Scenarios Notes and Constraints on Using ECSs |
| 2021-03-04 | This issue is the seventh official release. Added the address for downloading the password reset script in Resetting the Password for Logging In to a Linux ECS. |
| 2020-12-17 | This issue is the sixth official release. Modified the GPU memory of P2s flavors in GPU-accelerated ECSs. |
| 2020-11-30 | This issue is the fifth official release. Terminated automatic ECS recovery. |
| 2020-10-30 | This issue is the fourth official release. Added the following content: • Added Migrating an ECS. • Added Pi2 ECS specifications in GPU-accelerated ECSs. |

| This issue is the third official release. Added the following content: • Added P2s ECS specifications in GPU-accelerated ECSs. • Added S6 ECS specifications in Disk-intensive ECSs. • Added GPU Driver to describe GRID and Tesla drivers. • Added GPU Driver to describe GRID and Tesla drivers. • Added GPU Driver to describe GRID and Tesla drivers. • Added GPU Driver to describe GRID and Tesla drivers. • Added GPU Driver to describe GRID and Tesla drivers. • Added GPU Driver to describe GRID and Tesla drivers. • Added GPU Driver to describe GRID and Tesla drivers. • Added GPU Driver to describe GRID and Tesla drivers. • Added GPU Driver to describe GRID and Tesla drivers. • Added an example for setting a time zone in Changing the Time Zone for an ECS. • Added GECS specifications in Dedicated General-Purpose ECSs. • Added M6 ECS specifications in Memory-optimized ECSs. • Overview • Ultra-high I/O ECSs • Added g5.zxlarge.2 and g5.4xlarge.4 in GPU-accelerated ECSs. • Modifying a Private IP Address 2018-09-30 This issue is the second official release. Added the following content: • Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? • Added newly released C3 and G3 ECSs. • Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: • Adjusted the structure of the document. • Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. • Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. • Updated supported images in Enabling NIC Multi-Queue. • Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | Released On | Description |
|---|-------------|---|
| Added P2s ECS specifications in GPU-accelerated ECSs. Added S6 ECS specifications in General-Purpose ECSs. Added D3 ECS specifications in Disk-intensive ECSs. Added GPU Driver to describe GRID and Tesla drivers. Added flavors t2.large.2, t2.xlarge.2, and t2.2xlarge.2 in General-Purpose ECSs. 19.5 OS Monitoring Metrics Supported by ECSs with the Agent Installed Added an example for setting a time zone in Changing the Time Zone for an ECS. Added C6 ECS specifications in Dedicated General-Purpose ECSs. Added M6 ECS specifications in Memory-optimized ECSs. Overview Ultra-high I/O ECSs Added g5.2xlarge.2 and g5.4xlarge.4 in GPU-accelerated ECSs. Modifying a Private IP Address 2018-09-30 This issue is the second official release. Added the following content: Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | 2020-08-10 | This issue is the third official release. |
| Added S6 ECS specifications in General-Purpose ECSs. Added D3 ECS specifications in Disk-intensive ECSs. Added GPU Driver to describe GRID and Tesla drivers. Added flavors t2.large.2, t2.xlarge.2, and t2.2xlarge.2 in General-Purpose ECSs. 19.5 OS Monitoring Metrics Supported by ECSs with the Agent Installed Added an example for setting a time zone in Changing the Time Zone for an ECS. Added C6 ECS specifications in Dedicated General-Purpose ECSs. Added M6 ECS specifications in Memory-optimized ECSs. Overview Ultra-high I/O ECSs Added g5.2xlarge.2 and g5.4xlarge.4 in GPU-accelerated ECSs. Modifying a Private IP Address 2018-09-30 This issue is the second official release. Added the following content: Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | Added the following content: |
| Added D3 ECS specifications in Disk-intensive ECSs. Added GPU Driver to describe GRID and Tesla drivers. Added flavors t2.large.2, t2.xlarge.2, and t2.2xlarge.2 in General-Purpose ECSs. 19.5 OS Monitoring Metrics Supported by ECSs with the Agent Installed Added an example for setting a time zone in Changing the Time Zone for an ECS. Added C6 ECS specifications in Dedicated General-Purpose ECSs. Added M6 ECS specifications in Memory-optimized ECSs. Overview Ultra-high I/O ECSs Added g5.2xlarge.2 and g5.4xlarge.4 in GPU-accelerated ECSs. Modifying a Private IP Address 2018-09-30 This issue is the second official release. Added the following content: Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | Added P2s ECS specifications in GPU-accelerated ECSs. |
| Added GPU Driver to describe GRID and Tesla drivers. Added flavors t2.large.2, t2.xlarge.2, and t2.2xlarge.2 in General-Purpose ECSs. 19.5 OS Monitoring Metrics Supported by ECSs with the Agent Installed Added an example for setting a time zone in Changing the Time Zone for an ECS. Added C6 ECS specifications in Dedicated General-Purpose ECSs. Added M6 ECS specifications in Memory-optimized ECSs. Overview Ultra-high I/O ECSs Added g5.2xlarge.2 and g5.4xlarge.4 in GPU-accelerated ECSs. Modifying a Private IP Address This issue is the second official release. Added the following content: Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | Added S6 ECS specifications in General-Purpose ECSs. |
| Added flavors t2.large.2, t2.xlarge.2, and t2.2xlarge.2 in General-Purpose ECSs. 19.5 OS Monitoring Metrics Supported by ECSs with the Agent Installed Added an example for setting a time zone in Changing the Time Zone for an ECS. Added C6 ECS specifications in Dedicated General-Purpose ECSs. Added M6 ECS specifications in Memory-optimized ECSs. Overview Ultra-high I/O ECSs Added g5.2xlarge.2 and g5.4xlarge.4 in GPU-accelerated ECSs. Modifying a Private IP Address This issue is the second official release. Added the following content: Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | Added D3 ECS specifications in Disk-intensive ECSs. |
| General-Purpose ECSs. 19.5 OS Monitoring Metrics Supported by ECSs with the Agent Installed Added an example for setting a time zone in Changing the Time Zone for an ECS. Added C6 ECS specifications in Dedicated General-Purpose ECSs. Added M6 ECS specifications in Memory-optimized ECSs. Overview Ultra-high I/O ECSs Added g5.2xlarge.2 and g5.4xlarge.4 in GPU-accelerated ECSs. Modifying a Private IP Address This issue is the second official release. Added the following content: Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | Added GPU Driver to describe GRID and Tesla drivers. |
| Agent Installed Added an example for setting a time zone in Changing the Time Zone for an ECS. Added C6 ECS specifications in Dedicated General-Purpose ECSs. Added M6 ECS specifications in Memory-optimized ECSs. Overview Ultra-high I/O ECSs Added g5.2xlarge.2 and g5.4xlarge.4 in GPU-accelerated ECSs. Modifying a Private IP Address This issue is the second official release. Added the following content: Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | |
| Time Zone for an ECS. Added C6 ECS specifications in Dedicated General-Purpose ECSs. Added M6 ECS specifications in Memory-optimized ECSs. Overview Ultra-high I/O ECSs Added g5.2xlarge.2 and g5.4xlarge.4 in GPU-accelerated ECSs. Modifying a Private IP Address This issue is the second official release. Added the following content: Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | l '' ' |
| ECSs. Added M6 ECS specifications in Memory-optimized ECSs. Overview Ultra-high I/O ECSs Added g5.2xlarge.2 and g5.4xlarge.4 in GPU-accelerated ECSs. Modifying a Private IP Address This issue is the second official release. Added the following content: Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | |
| Overview Ultra-high I/O ECSs Added g5.2xlarge.2 and g5.4xlarge.4 in GPU-accelerated ECSs. Modifying a Private IP Address This issue is the second official release. Added the following content: Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | |
| Ultra-high I/O ECSs Added g5.2xlarge.2 and g5.4xlarge.4 in GPU-accelerated ECSs. Modifying a Private IP Address This issue is the second official release. Added the following content: Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | Added M6 ECS specifications in Memory-optimized ECSs. |
| Added g5.2xlarge.2 and g5.4xlarge.4 in GPU-accelerated ECSs. Modifying a Private IP Address This issue is the second official release. Added the following content: Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: | | Overview |
| ECSs. Modifying a Private IP Address This issue is the second official release. Added the following content: Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | Ultra-high I/O ECSs |
| This issue is the second official release. Added the following content: Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | |
| Added the following content: Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | Modifying a Private IP Address |
| Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | 2018-09-30 | This issue is the second official release. |
| Full-ECS Image? Added newly released C3 and G3 ECSs. Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | Added the following content: |
| Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | |
| accelerated ECS Modified the following content: Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | Added newly released C3 and G3 ECSs. |
| Adjusted the structure of the document. Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | |
| Modified description in Creating an ECS, Logging In to an ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | Modified the following content: |
| ECS, and Backup Using CBR to support full-ECS images. Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | Adjusted the structure of the document. |
| Updated supported images in Enabling NIC Multi-Queue. Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | · · · · · · · · · · · · · · · · · · · |
| Modified description in Managing ECS Groups, allowing you to add a created ECS to an ECS group or remove it from an ECS group. | | Added g1.4xlarge and P2 ECSs in GPU-accelerated ECSs. |
| to add a created ECS to an ECS group or remove it from an ECS group. | | Updated supported images in Enabling NIC Multi-Queue. |
| 2018-08-15 This issue is the first official release. | | to add a created ECS to an ECS group or remove it from an |
| | 2018-08-15 | This issue is the first official release. |