

Elastic Cloud Server (ECS) 8.1.0

User Guide

Issue 01
Date 2021-10-30



Copyright © Huawei Technologies Co., Ltd. 2021. 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 Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are trademarks 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 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 Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: <https://e.huawei.com>

Contents

1 Introduction.....	1
1.1 What Is Elastic Cloud Server?.....	1
1.2 ECS Advantages.....	2
1.3 Application Scenarios.....	4
1.4 Related Services.....	5
1.5 Access Mode and Constraints.....	6
1.6 Implementation Principle.....	8
1.7 Feature List.....	11
2 ECS Type and Flavors.....	12
2.1 General-purpose ECSs.....	12
2.2 GPU-accelerated ECSs.....	13
2.2.1 Overview.....	13
2.2.2 GPU-accelerated ECSs.....	15
2.2.3 Computing-accelerated ECSs.....	17
2.3 vGPU-accelerated ECSs.....	21
2.4 Ultra-high I/O ECSs.....	22
2.5 General Computing-plus ECSs.....	24
2.6 USB Passthrough ECSs.....	26
2.7 Dedicated general-purpose ECSs.....	28
2.8 Memory-optimized ECSs.....	28
2.9 Disk-intensive ECSs.....	29
2.10 Large-memory ECSs.....	31
2.11 AI-accelerated ECSs.....	32
3 Related Concepts.....	34
3.1 Regions and AZs.....	34
3.2 Cloud-Init.....	34
3.3 Local Disk, EVS Disk and Local Pass-through Disk.....	35
3.4 Same Storage.....	38
3.5 Feature Differences Between Arm and x86 Servers.....	39
3.6 ECS Quota.....	43
4 Operation Process.....	45
5 Creating an ECS.....	50

5.1 Overview.....	50
5.2 Applying for an ECS.....	50
5.3 Viewing ECS Creation Status.....	66
5.4 Installing the One-Click Password Reset Plugin.....	67
5.4.1 Overview.....	67
5.4.2 Installing the One-Click Password Reset Plugin for a Windows ECS.....	68
5.4.3 Installing the One-Click Password Reset Plugin for a Linux ECS.....	71
5.5 Initializing EVS Data Disks.....	78
6 Logging In to an ECS.....	79
6.1 Login Mode Overview.....	79
6.2 Logging In to a Linux ECS.....	81
6.2.1 Remotely Logging In to a Linux ECS Using a Key Pair (SSH).....	81
6.2.2 Remotely Logging In to a Linux ECS Using a Password (SSH).....	84
6.2.3 Logging In to a Linux ECS Using VNC (Through the Console).....	86
6.3 Logging In to a Windows ECS.....	88
6.3.1 Obtaining the Password for Logging In to a Windows ECS.....	88
6.3.2 Logging In to a Windows ECS Using VNC (Through the Console).....	89
6.3.3 Logging In to a Windows ECS Using a Password (MSTSC).....	90
7 Managing an ECS.....	93
7.1 Basic Operations.....	93
7.1.1 Viewing ECS Details.....	93
7.1.2 Changing the ECS Name.....	93
7.1.3 Adding and Managing ECS Tags.....	94
7.1.4 Querying ECSs by Filters.....	95
7.1.5 Exporting ECS Details.....	96
7.1.6 Changing the Time Zone for an ECS.....	96
7.2 Life Cycle.....	99
7.2.1 Managing the Life Cycle of an ECS.....	99
7.2.2 Deleting an ECS.....	102
7.2.3 Changing the Validity Period of an ECS.....	103
7.3 Creating a Private Image Using an Existing ECS.....	104
7.4 Modifying the DR or Backup Function of an ECS.....	104
7.5 Cloning an ECS.....	105
7.6 Managing the Watchdog Status of an ECS.....	112
7.7 Managing the HA Status of an ECS.....	113
7.8 ECS Snapshot.....	113
7.9 Creating a CD-ROM Drive and Attaching ISO/UVP VMTools.....	116
8 Passwords and Key Pairs.....	122
8.1 Overview.....	122
8.2 Deleting the Initial Password for Logging In to a Windows ECS.....	123
8.3 Resetting the ECS Password with One Click (Windows and Linux).....	124

8.4 Manually Resetting the Password for Logging In to a Windows ECS.....	125
8.5 Manually Resetting the Password for Logging In to a Linux ECS.....	127
8.6 Creating a Key Pair.....	129
9 ECS Flavors.....	134
9.1 Changing the Flavor of an ECS.....	134
10 EVS Disk.....	140
10.1 Applying for a Data Disk.....	140
10.2 Attaching an EVS Disk.....	144
10.3 Initializing a Data Disk.....	148
10.3.1 Initialization Overview.....	148
10.3.2 Initializing a Windows Data Disk.....	149
10.3.3 Initializing a Linux Data Disk (fdisk).....	154
10.3.4 Initializing a Linux Data Disk (parted).....	159
10.4 Expanding EVS Disk Capacity.....	163
10.4.1 Overview.....	163
10.4.2 Expanding Disk Capacity Online.....	165
10.4.3 Expanding Disk Capacity Offline.....	168
10.4.4 Operations After Expanding Disk Capacity in Windows.....	169
10.4.5 Operations After Expanding Disk Capacity in Linux (Adding Partitions Using fdisk).....	172
10.4.6 Operations After Expanding Disk Capacity in Linux (Adding Partitions Using parted).....	176
10.4.7 Operations After Expanding Disk Capacity in Linux (Replacing Original Partitions Using fdisk)....	180
10.4.8 Operations After Expanding Disk Capacity in Linux (Replacing Original Partitions Using parted).184	
10.4.9 Adding a Data Disk.....	189
10.5 Releasing an EVS Disk.....	189
10.5.1 Detaching an EVS Disk.....	189
10.5.2 Deleting an EVS Disk.....	191
11 ECS Group.....	194
11.1 Creating an ECS Group.....	194
11.2 Adding an ECS to or Removing an ECS from an ECS Group.....	195
12 Network and Security.....	199
12.1 Configuring Intra-VPC Communication and Security Policy for ECSs.....	199
12.2 Enabling Communication Between an ECS and the Internet and Configuring Security Policies.....	199
12.3 Modifying NIC Configurations.....	200
12.4 Configuring Security Group Rules.....	205
12.5 Changing the EIP.....	208
12.6 Binding a Floating Private IP Address (Virtual IP Address).....	209
13 Operating Systems.....	211
13.1 Reinstalling an ECS OS.....	211
13.2 Changing the ECS OS.....	213
14 Monitoring Metrics.....	217

14.1 ECS Monitoring Metrics.....	217
14.2 Viewing ECS Running Status.....	219
14.3 Viewing Information Through Metadata.....	220
15 Load Balancing.....	229
16 Best Practices.....	230
16.1 Creating an Application Allowing Access from External Networks.....	230
16.1.1 Overview.....	230
16.1.2 Implementation Plan.....	230
16.1.3 Requesting and Configuring Services.....	234
16.1.3.1 Applying for a VPC.....	234
16.1.3.2 Creating a Security Group and Configuring Security Group Rules.....	237
16.1.3.3 Creating an ECS.....	240
16.1.3.4 Logging In to an ECS.....	245
16.1.3.5 Initializing a Windows Data Disk.....	246
16.1.3.6 Initializing a Linux Data Disk (fdisk).....	251
16.1.3.7 Initializing a Linux Data Disk (parted).....	255
16.1.4 Deploying the Application and Database.....	260
16.2 Synchronizing the Clock of the Windows ECS.....	260
16.2.1 Overview.....	261
16.2.2 Enabling the NTP Service.....	261
16.2.3 Changing the NTP Server Address.....	263
16.3 Manually Viewing the Disk Mount Point.....	265
16.4 Using Direct Connect to Establish VPN Channels to Implement Cross-Region DR.....	272
16.5 Distributing Traffic on APP ECSs.....	276
16.6 Creating and Attaching a Data Disk to a Database Host.....	311
17 Website Construction Tutorial.....	319
17.1 Building a Discuz Website.....	319
17.1.1 Overview.....	319
17.1.2 Implementation Plan.....	319
17.1.3 Applying for and Configuring Services.....	322
17.1.3.1 Applying for a VPC.....	322
17.1.3.2 Creating a Security Group and Configuring Security Group Rules.....	325
17.1.3.3 Creating an ECS.....	328
17.1.3.4 Logging In to an ECS.....	333
17.1.3.5 Initializing a Linux Data Disk (fdisk).....	334
17.1.3.6 Initializing a Linux Data Disk (parted).....	339
17.1.4 Building a Discuz Website.....	343
17.1.4.1 Installing the Database.....	344
17.1.4.2 Configuring the Web Environment.....	345
17.1.4.3 Deploying the Website Code.....	346
17.1.5 Checking Whether the Website Is Built.....	347

18 FAQs	348
18.1 General FAQs	348
18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?	348
18.1.2 How Do I Handle Error Messages Displayed on ManageOne?	349
18.1.3 Common Methods to Troubleshoot Local Client Plugin Used for Attaching ISO	354
18.1.4 What Is Quota?	355
18.2 Image FAQs	355
18.2.1 What Should I Do If an Image Failed to Be Updated?	355
18.2.2 What Is a Static Injection Image?	356
18.3 ECS FAQs	356
18.3.1 What Is the cloudbase-init Account in Windows ECSs?	357
18.3.2 What Should I Do When an ECS Remains in the Restarting or Stopping State for a Long Time?	357
18.3.3 Can a Deleted ECS Be Provisioned Again?	357
18.3.4 How Can I Change the Static Host Name of a Linux ECS?	357
18.3.5 What Restrictions Are Involved with Using ECSs?	359
18.3.6 What Can I Do with ECSs?	359
18.3.7 How Long Does It Take to Provision an ECS?	359
18.3.8 What Functions Does the Delete Button Provide?	359
18.3.9 What Is a Deleted ECS?	359
18.3.10 Why Does the Task Status Area Show an ECS Creation Failure But the ECS List Displays the Created ECS?	360
18.4 EIP FAQs	360
18.4.1 Can Multiple EIPs Be Bound to an ECS?	360
18.4.2 Will a NIC Added to an ECS Be Identified Automatically?	360
18.5 Login FAQs	360
18.5.1 What Should I Do After I Log In to an ECS Using VNC and Perform an Operation But the Page Does not Respond for a Long Time?	360
18.5.2 What Should I Do If I Cannot View Data After Logging In to an ECS Using VNC?	361
18.5.3 Why Does a Blank Screen Appear While the System Displays a Message Indicating Successful Authentication After I Attempted to Log In to an ECS Using VNC?	361
18.5.4 Why Was My Login to a Linux ECS with a Key File Unsuccessful?	361
18.5.5 Why Does the System Display a Message Indicating that the Password for Logging In to a Windows ECS Cannot Be Queried?	362
18.5.6 What Should I Do If I Cannot Use MSTSC to Log In to an ECS Running Windows Server 2012?	363
18.6 Network and Security FAQs	364
18.6.1 Configuring a Static IP Address for an ECS	364
18.6.2 Why Can I Remotely Connect to an ECS But Cannot Ping It?	365
18.6.3 What Should I Do If a Public Key Cannot Be Imported?	366
18.6.4 What Should I Do If a Public Key Fails to Be Imported to ManageOne After a Key Pair Is Created Using PuTTYgen?	366
18.6.5 How Can I Change the MTU of a Linux ECS NIC?	368
18.6.6 How Can I Change the MTU of a Windows ECS NIC?	370
18.6.7 Accessing the Internet Using an ECS Without a Public IP Address	373
18.6.8 What Do I Do If the Virtual IP Address Cannot Be Pinged After Being Bound to the ECS NIC?	376

18.7 Disk FAQs.....	378
18.7.1 Logging In to the EVS Console as a VDC Administrator or VDC Operator.....	378
18.7.2 How Can I Attach a New EVS Disk to an ECS?.....	379
18.7.3 Can Multiple EVS Disks Be Attached to a Single ECS?.....	379
18.7.4 Can I Attach an EVS Disk to Multiple Instances?.....	380
18.7.5 How Many States Does an EVS Disk Have?.....	380
18.7.6 Does an EVS Disk or Snapshot Generate Metering Information in All States?.....	383
18.7.7 Can I Change the EVS Disk Capacity?.....	383
18.7.8 Will Data in an EVS Disk Be Lost When the EVS Disk Is Detached?.....	384
18.7.9 Device Type.....	384
18.7.10 Shared Disk.....	385
18.7.11 Applying for a Snapshot.....	389
18.7.12 Creating a Backup.....	391
18.8 OS FAQs.....	392
18.8.1 Can I Install or Upgrade the OS by Myself?.....	392
18.8.2 Can the OS of an ECS Be Changed?.....	392
18.8.3 Can I Select Other OSs During ECS OS Reinstallation?.....	393
18.8.4 How Can I Obtain Data Disk Information If Tools Is Deleted?.....	393
18.8.5 What Should I Do If the One-Click Password Reset Plugin Fails to Start?.....	394
18.8.6 What Can I Do If the OS Reinstallation or Change Fails?.....	403
A Supported vGPU Types.....	405
B Installing a GRID Driver on a vGPU-accelerated ECS.....	426
C Supported Driver Versions and OSs.....	430

1 Introduction

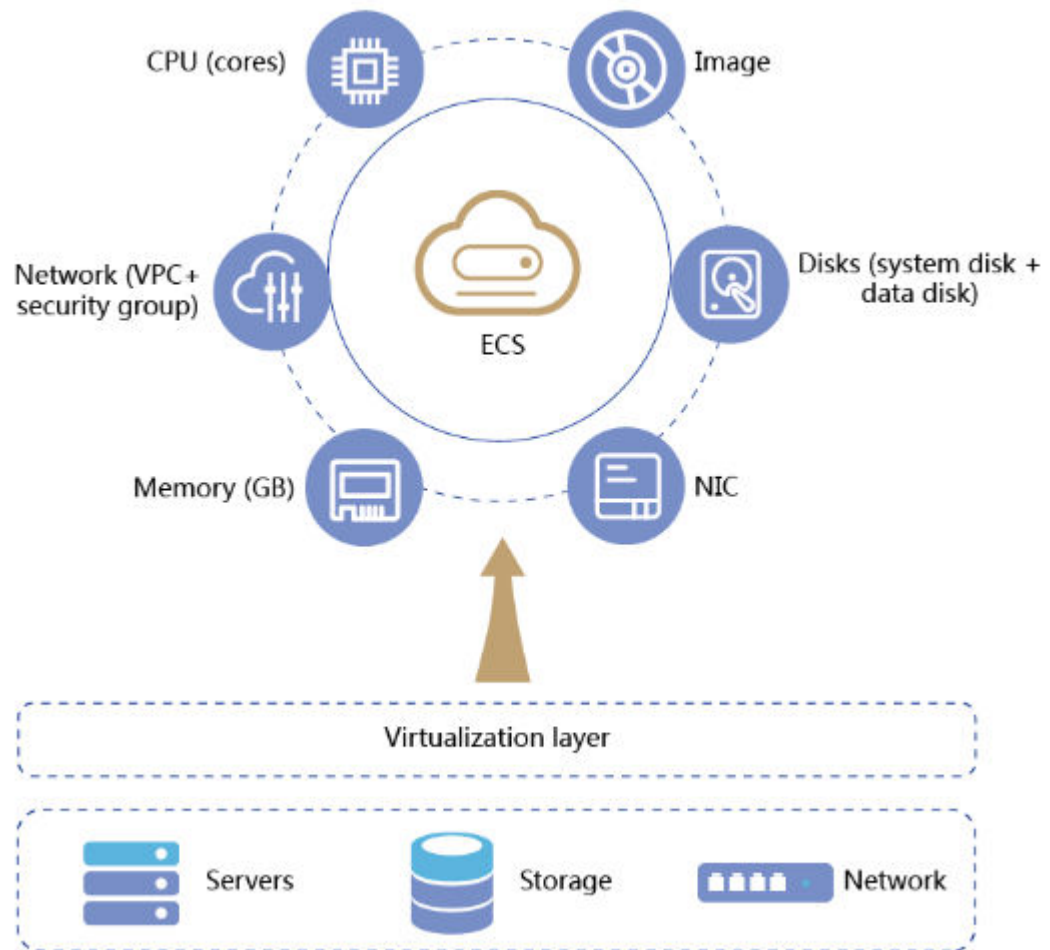
1.1 What Is Elastic Cloud Server?

Definition

An Elastic Cloud Server (ECS) is an on-demand and scalable compute server consisting of vCPUs, memory, and Elastic Volume Service (EVS) disks. The ECS service works with the Virtual Private Cloud (VPC) and Cloud Server Backup Service (CSBS) services to give you an efficient and reliable computing environment for your data and applications. The resources used by ECSs, including vCPUs and memory, are hardware resources that are consolidated using the virtualization technology.

When creating an ECS, you can customize the number of vCPUs, memory size, image type, and login authentication mode. After an ECS is created, you can use it like using your local computer or physical server. They provide you with relatively inexpensive compute and storage resources on demand. A unified management platform simplifies management and maintenance, enabling you to focus on services.

Figure 1-1 Elastic cloud server



Function

The ECS provides the following functions. For details about the application process and supported functions, see [Table 4-1](#).

- Customize the flavor, image, network, disks, authentication mode, and number of ECSs when creating ECSs.
- Manage the life cycle of an ECS, including starting, stopping, restarting, and deleting an ECS. Clone an ECS, create an ECS snapshot, and manage the watchdog status and HA status. Modify vCPUs and memory of an ECS.
- Expand the capacity of EVS disks attached to an ECS, attach EVS disks to an ECS, detach EVS disks from an ECS, and use shared EVS disks.
- Change and reinstall the ECS OS, and create a private image using an existing ECS.
- Bind an elastic IP address (EIP) to and unbind an EIP from an ECS.

1.2 ECS Advantages

Compared with traditional servers, ECSs are easy to provision and use, and have high reliability, security, and scalability.

Table 1-1 Comparison of ECSs with traditional servers

Item	ECS	Traditional Server
Reliability	<p>The ECS service can work with other cloud services, such as storage services and disaster recovery & backup, to allow flavor modification, data backup, recovery using a backup, and rapid recovery from a fault.</p>	<ul style="list-style-type: none"> • Traditional servers, subject to hardware reliability issues, have a higher likelihood of failure. You need to manually back up their data. • You need to manually restore their data, which may be complex and time-consuming.
Security	<p>The security service ensures that ECSs work in a secure environment. This service protects your data, hosts, and web pages, monitors program execution, and checks whether ECSs are under brute force attacks and whether remote logins are performed. This aims to enhance your system security and mitigate the risks of ECS intrusion by hackers.</p>	<ul style="list-style-type: none"> • You need to purchase and deploy security measures additionally. • It is difficult to perform access control on multiple users to multiple servers.
Scalability	<ul style="list-style-type: none"> • You can modify an ECS flavor, including the number of CPUs and memory size. • You can expand the capacity of the system disk and data disk. • Auto Scaling (AS) is used, which enables you to configure AS policies so that ECSs are automatically added and removed during traffic peaks and lulls, respectively. This ensures that your service requirements are met and maximizes resource utilization. 	<ul style="list-style-type: none"> • Configurations are fixed and are difficult to meet changing needs. • Hardware upgrade is required for modifying configuration, which takes a long time and the service interruption time is uncontrollable. Service scalability and continuity are low.
Easy to use	<ul style="list-style-type: none"> • A simple and easy-to-use unified management console streamlines operations and maintenance. • A wide range of products are provided, including network, storage, security, and big data devices, which can be provisioned and deployed at the one-stop manner. 	<ul style="list-style-type: none"> • Without software support, users must repeat all steps when adding each new server. • It is difficult for you to obtain all required services from one service provider.

Item	ECS	Traditional Server
Easy to provision	After deploying an entire private cloud and finishing necessary configurations, you can customize the number of CPUs, memory size, images, and networks to create ECSs at any time.	When using traditional servers, you must buy and assemble the components and install the operating systems (OSs).

1.3 Application Scenarios

ECSs are virtual machines that can be rapidly provisioned and scaled to suit your changing demands. They provide you with relatively inexpensive compute and storage resources on demand. A unified management platform simplifies management and maintenance, enabling you to focus on services.

HUAWEI CLOUD Stack provides multiple types of ECSs to meet requirements of various scenarios. ECSs are used in a wide range of scenarios, including:

- **Simple applications or small-traffic websites**

Simple applications or small-traffic websites, such as blogs and enterprise websites, have relatively low requirements on the computing and storage performance of the server. A general-purpose ECS will meet the requirements. If you have higher requirements on CPUs, memory, data disks, or the system disk of an ECS, you can modify the ECS flavor or expand disk capacity. You can also create new ECSs at any time.

- **Multimedia making, video making, and image processing**

In multimedia making, video making, or image processing scenarios, ECSs must provide good image processing capabilities. For these scenarios, you can choose ECSs with high CPU and GPU computing performance, such as GPU graphics-accelerated or GPU-computing-accelerated ECSs, to meet your service requirements.

- **Databases and other applications that require fast data exchange and processing**

For high-performance relational databases, NoSQL databases, and other applications that require high I/O performance on servers, you can choose ultra-high I/O ECSs and use high-performance local NVMe SSDs as data disks to provide better read and write performance and lower latency, improving the file read and write rate.

- **Applications with noticeable load peaks and troughs**

For applications that have noticeable load peaks and troughs, such as video websites, school course selection systems, and game companies, the number of visits may increase significantly within a short time. To improve resource utilization and ensure that your applications run properly, you can use AS to work with ECSs. You can configure AS policies so that ECSs are automatically added and removed during traffic peaks and lulls, respectively. This helps maximize resource utilization and also meet service requirements, thereby reducing costs.

- AI inference, machine learning, and deep learning**
 AI-accelerated ECSs use Huawei's Ascend 310 chips. They are suitable for scenarios that require real-time, highly concurrent massive computing, such as, AI inference, machine learning, as well as video encoding and decoding.

For details about the application scenarios, constraints, and recommended flavors for different types of ECSs, see the corresponding sections in [2 ECS Type and Flavors](#).

1.4 Related Services

The ECS service can work with other cloud services to provide you with a stable, secure, highly-available, and easy-to-manage network experience. The following figure shows services that may be used together with ECS. For details, see [Table 1-2](#).

Figure 1-2 ECS-related services

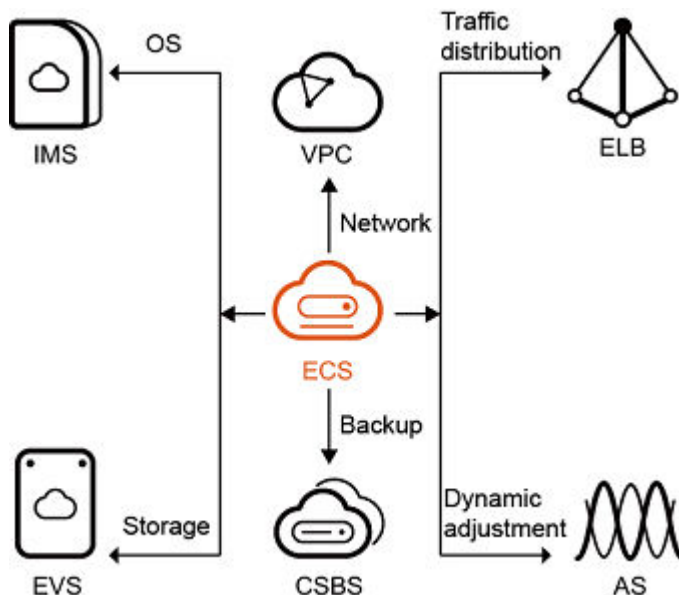


Table 1-2 Relationship between ECS and other cloud services

Service Name	Description
Elastic Volume Service (EVS)	EVS provides storage for ECSs. You can attach EVS disks to an ECS, detach EVS disks from an ECS, and expand the capacity of EVS disks of an ECS.
Image Management Service (IMS)	You can create an ECS using a public image, private image, or shared image. You can create a private image using an ECS.

Service Name	Description
Cloud Server Backup Service (CSBS)	CSBS provides users with on-demand backup service. Users can apply for backup for certain ECSs based on their service requirements so that the ECSs can be automatically and rapidly restored in the event of data loss or damage.
Auto Scaling (AS)	After AS is used and AS policies are configured, the system automatically adds ECSs during traffic peaks and releases ECSs during traffic lulls, meeting your service requirements and maximizing resource utilization.
Elastic Load Balance (ELB)	ELB distributes service loads to multiple ECSs, improving the system's service processing capability. ELB performs health checks on ECSs to automatically remove abnormal ECSs and distribute service loads to healthy ones, ensuring service continuity.
Virtual Private Cloud (VPC)	VPC provides networks for ECSs. You can use the rich functions of VPC to flexibly configure a secure running environment for ECSs.

1.5 Access Mode and Constraints

Access Mode

Two methods are available:

- Web UI
Log in to ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios) as a tenant user and select the cloud service on the **Console** menu.
- API
Use this method if you need to integrate the cloud service into a third-party system for secondary development. For details, see API reference of the service on **Operation Help Center**.

Constraints

ECS constraints

- Virtualization software cannot be installed on ECSs for secondary virtualization.
- The audio card is not supported.

Constraints on Windows

This section describes only the common constraints on Windows. For details about all constraints, visit the official website.

- Do not stop the shutdownmon.exe process of the Windows OS. Otherwise, the ECS may fail to be stopped or restarted.

- Do not rename, delete, or disable the administrator account in Windows. Otherwise, the ECS cannot be used properly.
- The swap partition is not recommended.

Constraints on Linux

- Do not change the permissions of the directories in the partition where the root directory is located, especially the permissions of the **/etc**, **/sbin**, **/bin**, **/boot**, **/dev**, **/usr**, and **/lib** directories. Improper permission modification may cause system exceptions.
- Do not rename, delete, or disable the root account in Linux.
- Do not compile the kernel of the Linux OS or perform any other operations on the kernel.
- The swap partition is not recommended.

Constraints on system capacity specifications

- For details about the number of KVM hosts supported in a single region or a single AZ, see [Table 1-3](#).
- For details about the number of VMs supported in a single region or a single AZ, see [Table 1-3](#).
- The number of VMs in the environment is limited if the number of hosts exceeds 500 to prevent system instability caused by excessive usage.

Table 1-3 KVM hosts and VMs in different deployment scales

Deployment Scale	≤ 50 PMs	≤ 100 PMs	≤ 200 PMs	≤ 500 PMs	≤ 1000 PMs	≤ 4000 PMs (2000 KVMs + 2000 BMSs)
Maximum number of KVM hosts in a region	50	100	200	500	1000	2000 (supporting 2000 centralized BMSs at the same time)
Maximum number of VMs in a region	500	1000	2000	5000	10,000	40000
Maximum number of KVM hosts in an AZ	50	100	200	500	1000	2000
Maximum number of VMs in an AZ	500	1000	2000	5000	10,000	40000

1.6 Implementation Principle

Architecture

Figure 1-3 ECS logical architecture

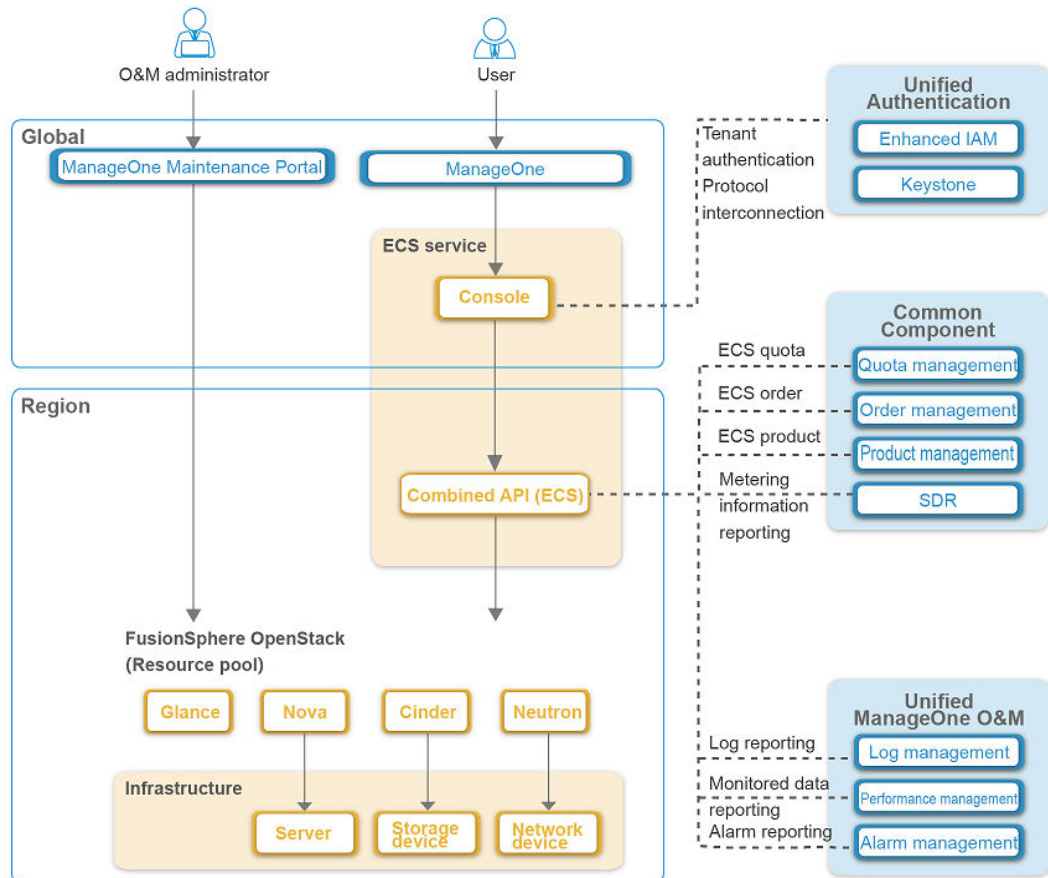


Table 1-4 Component details

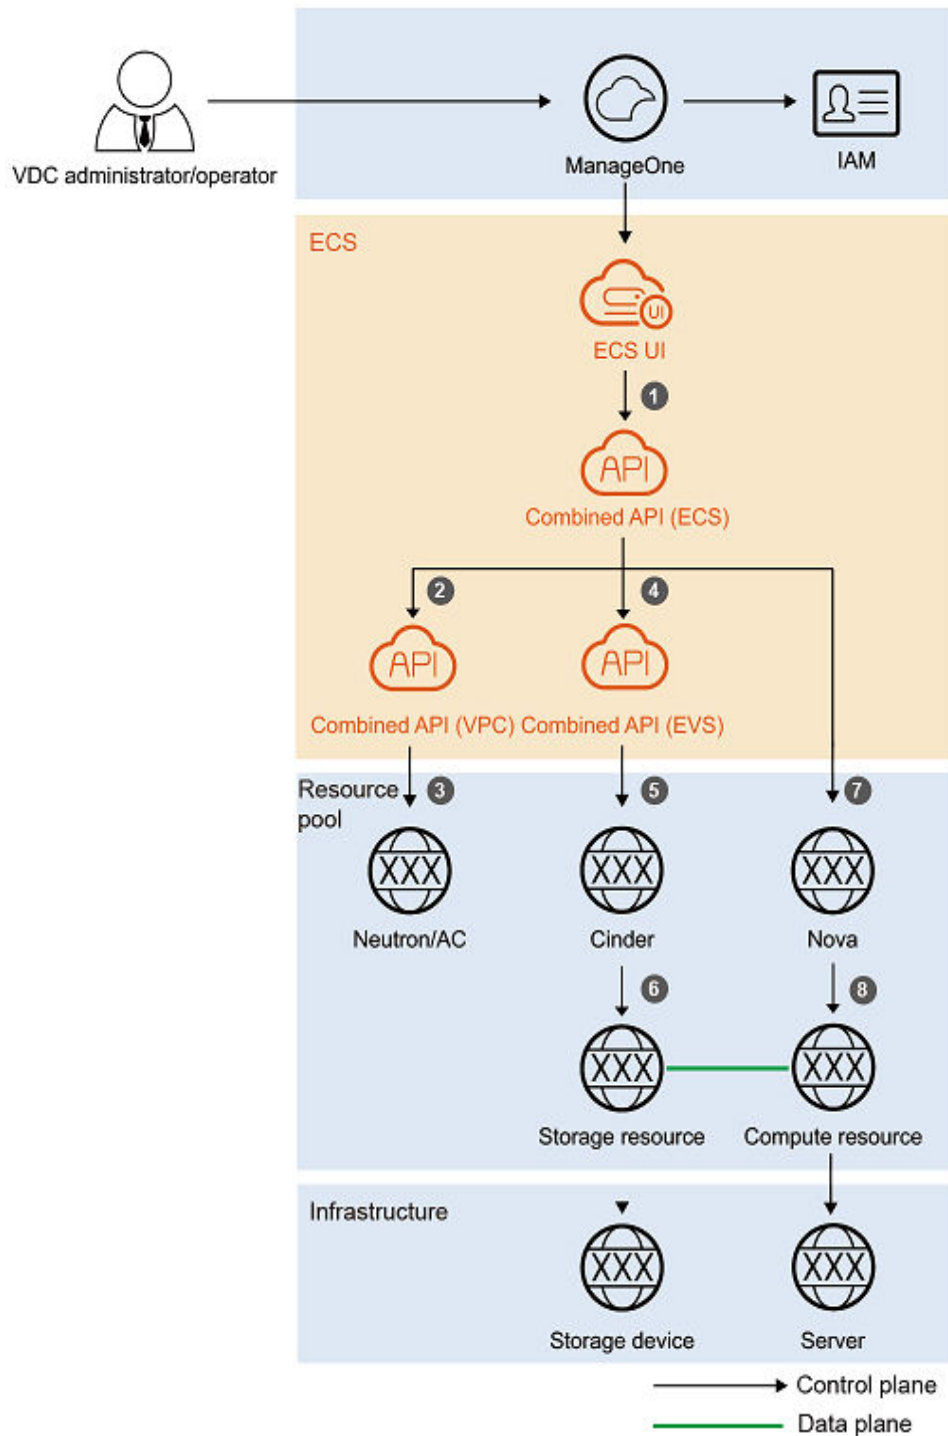
Type	Description
Console	ECS_UI is a console centered on the Elastic Cloud Server (ECS) service and manages relevant resources.
Combined API (ECS)	Provides a backend service for ECSs. It can be seen as the server end of ECS_UI, and can call FusionSphere OpenStack components. Requests sent by an ECS from the console are forwarded by ECS_UI to Combined API and are returned to ECS_UI after being processed by Combined API.

Type	Description
Resource pool	<ul style="list-style-type: none"> ● Glance: Provides image management service. ● Nova: Manages the life cycle of compute instances in the FusionSphere OpenStack environment, for example, creating instances in batches, and scheduling or stopping instances on demand. ● Cinder: Provides persistent block storage for running instances. Its pluggable drives facilitate block storage creation and management. ● Neutron: Provides APIs for network connectivity and addressing.
Unified Authentication	Provides Identity and Access Management (IAM) during login.
Common Component	Combined API reports ECS quota, order, product information, and metering and charging information to the ManageOne operation module.
Unified O&M	Combined API reports ECS log, monitoring, and alarm information to the ManageOne O&M module.

Workflow

Figure 1-4 shows the workflow for creating an ECS.

Figure 1-4 Workflow for creating an ECS



The steps in the figure above are as follows:

1. Submit the application on the ECS page, corresponding to step 1 in the preceding figure.
2. Create network resources, corresponding to step 2 to step 3 in the preceding figure.
 - a. The ECS API of Combined API calls the VPC API of Combined API.

- b. The VPC API calls Neutron to create an EIP or a port.
3. Create storage resources, corresponding to step 4 to step 6 in the preceding figure.
 - a. The ECS API of Combined API calls the EVS API of Combined API.
 - b. The EVS API calls Cinder.
 - c. Cinder creates volumes in the storage pool according to storage resource application policies.
4. Create compute resources, corresponding to step 7 to step 8 in the preceding figure.
 - a. The ECS API sends the request to Nova.
 - b. Nova creates an ECS in the compute resource pool.

1.7 Feature List

Context

The CPUs supported by the service are provided by Intel, Hygon (AMD), Kunpeng, and Phytium. Intel and Hygon (AMD) CPUs use the x86 architecture, and Kunpeng and Phytium CPUs use the Arm architecture. For details about the features and functions supported by servers using different CPUs, see [HUAWEI CLOUD Stack 8.1.0 Basic Service Feature List \(Compute&Network&Basic Management\)](#).

2 ECS Type and Flavors

2.1 General-purpose ECSs

Application Scenarios

General-purpose ECSs provide basic vCPU performance and a balance of computing, memory, and network resources. The performance can be improved based on the working load requirements, providing higher performance within a short period of time. These ECSs are suitable for many applications, such as web servers, enterprise R&D, and small-scale databases.

Recommended Flavors

[Table 2-1](#) list the flavors of general-purpose ECSs.

NOTE

- The naming pattern for flavors of ECSs whose virtualization type is KVM is as follows:
A.B.C
c.8xlarge.8 is an example.
In **A.B.C**:
A specifies the ECS type. For example, **s** indicates a general-purpose ECS, **c** a computing ECS, and **m** a memory-optimized ECS.
B specifies the size in the current series, and can be **medium**, **large**, or **xlarge**.
C specifies the ratio of memory to vCPUs expressed in a digit. For example, value **4** indicates that the ratio of memory to vCPUs is 4.

Table 2-1 Recommended flavors for KVM-based general-purpose ECSs

Type	vCPU	Memory (GB)	Flavor Name
General-purpose	1	1	s3.small.1
	1	2	s3.medium.2

Type	vCPU	Memory (GB)	Flavor Name
	2	4	s3.large.2
	4	8	s3.xlarge.2
	8	16	s3.2xlarge.2
	16	32	s3.4xlarge.2
	1	4	s3.medium.4
	2	8	s3.large.4
	4	16	s3.xlarge.4
	8	32	s3.2xlarge.4
	16	64	s3.4xlarge.4

2.2 GPU-accelerated ECSs

2.2.1 Overview

GPU ECSs provide outstanding floating-point computing capabilities. They are suitable for scenarios that require real-time, highly concurrent massive computing. GPU ECSs are classified as graphics-accelerated and computing-accelerated ECSs. In the preceding information:

- G series ECSs are suitable for 3D animation rendering and CAD.
G series ECSs are divided into G2 and G3 ECSs according to the GPU models.
- P series ECSs are designed for deep learning, scientific computing, and CAE.

 NOTE

- To use GPU ECSs, plan the corresponding host group when you use HUAWEI CLOUD Stack Deploy to install the ECSs. Otherwise, you need to complete a series of configurations. For details about how to configure G3 or GPU ECSs (P series), visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > Configuring GPU-accelerated ECSs**.
- The Arm architecture does not support GPU-accelerated ECSs.
- GPU passthrough Windows ECSs: In the **Device Manager** screen of a Windows ECS, if the emulated VGA device has a yellow exclamation mark on its icon or is disabled, or a GPU has been configured as the primary graphics card on the ECS, the only way to access the ECS is using the ECS IP address via a remote desktop or VNC client (like RealVNC, which is not a console interface) after installing a VNC server on the ECS. This is because the display information of the passthrough GPU cannot be transmitted to the QEMU VNC server. In this case, only the passthrough GPU can be used as the graphics card.
- GPU passthrough Linux ECSs: If a GPU has been configured as the primary graphics card on a Linux ECS, the only way to access the GPU passthrough ECS is using the ECS IP address via a VNC client (like RealVNC, which is not a console interface) after installing a VNC server on the ECS. This is because the display information of the passthrough GPU cannot be transmitted to the QEMU VNC server. In this case, only the passthrough GPU can be used as the graphics card.

Application Scenarios

- Applications
Deep learning, scientific computing, CAE, 3D animation rendering, and CAD
- Scenario characteristics
Real-time massive concurrent computing.
- Application scenarios
 - P series ECSs are suitable for artificial intelligence: Each GPU contains thousands of computing units, providing outstanding parallel computing capabilities. P series ECSs have been optimized for deep learning, supporting massive computing within a short period of time.
 - P series ECSs are suitable for scientific computing: Scientific computing has strict requirements on double-precision computing. During computing emulation, a large number of compute resources are used, and large volumes of data are generated. Therefore, scientific computing also has strict requirements on storage bandwidth and latency. P series ECSs meet these requirements.
 - G series ECSs are suitable for graphic workstation: G series ECSs provide outstanding computing capabilities for professional CAD, video rendering, and graphics processing.

GPU Models

The GPU vendor is NVIDIA whose vendor_id is 0x10de. [Table 2-2](#) describes the GPU models that support GPU ECSs.

Table 2-2 GPU models

GPU ECS	Virtualization Type	GPU Model	GPU Technology	Product ID	API Type	Alias Name	Description
G3 ECSs (GPU-accelerated)	KVM	NVIDIA Tesla T4	GPU passthrough	0x1eb8	PCIe 3.0	nvidia-t4	16 GB video RAM
PI1 ECSs (computing-accelerated)	KVM	NVIDIA Tesla P4	GPU passthrough	0x1bb3	PCIe 3.0	nvidia-p4	8 GB video RAM
PI2 ECSs (computing-accelerated)		NVIDIA Tesla P40		0x1b38	PCIe 3.0 Dual Slot	nvidia-p40	24 GB video RAM, two slots
P1 ECSs (computing-accelerated)		Tesla P100 PCIe 12 GB		0x15f7	PCIe 3.0	nvidia-p100-pcie-12	12 GB video RAM, two slots
		Tesla P100 PCIe 16 GB		0x15f8	PCIe 3.0	nvidia-p100-pcie-16	16 GB video RAM, two slots
P2 ECSs (computing-accelerated)		Tesla V100 PCIe 16 GB		0x1db4	PCIe 3.0	nvidia-v100-pcie-16	16 GB video RAM, two slots
		Tesla V100 PCIe 32 GB		0x1db6	PCIe 3.0	nvidia-v100-pcie-32	32 GB video RAM, two slots

2.2.2 GPU-accelerated ECSs

GPU-accelerated G3 ECSs are supported.

G3 ECSs are based on NVIDIA Tesla T4 hardware passthrough and provide graphics acceleration with 16 GB GPU memory and 2,560 CUDA cores. They support all AI frameworks and network types, and provide powerful performance and high efficiency, maximizing the effectiveness of large-scale deployment.

Supported Range

G3 ECSs have the following features:

- NVIDIA Tesla T4 GPUs
- Graphics acceleration applications
- GPU SR-IOV
- 16 GB GPU memory for graph processing
- 8.1 TFLOPS of single-precision computing, 65 TFLOPS of hybrid-precision computing
- Up to 130 TOPS for INT8 computing, and up to 260 TOPS for INT4 computing

Notes

G3 ECSs

- G3 ECSs support offline but not online flavor changes.
- G3 ECSs support the following OSs:
 - Carrier users: Click [here](#), search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)**, and obtain the latest document. In the document, search for **GPU pass-through** to view the support status.
 - Enterprise users: Click [here](#), search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)**, and obtain the latest document. In the document, search for **GPU pass-through** to view the support status.
- The NVIDIA driver must be installed on the image used to create G3 ECSs. If it is not installed in the images, install the driver after the G3 ECSs are created.

Recommended Flavors

Table 2-3 Recommended flavors for KVM-based GPU ECSs

Type	vCPU	Memory (GB)	Flavor Name	Number of GPUs	GPU Technology
GPU-accelerated (G3)	8	64	G3.2xlarge.8	1×T4	GPU passthrough
	16	128	G3.4xlarge.8	2×T4	
	32	256	G3.8xlarge.8	4×T4	

2.2.3 Computing-accelerated ECSs

Computing-accelerated ECSs are divided into PI1, PI2, P1, and P2 ECSs, whose details are as follows:

- PI1 ECSs use NVIDIA Tesla P4 GPUs dedicated for AI inference. Working with INT8 calculators of NVIDIA Tesla P4 GPUs, PI1 ECSs have shortened the inference latency by 15 times.
- PI2 ECSs use NVIDIA Tesla P40 GPUs dedicated for ultra-high AI inference throughput, providing outstanding real-time inference. The single-precision floating-point computing performance and INT8 integer computing performance of NVIDIA Tesla P40 GPUs are twice times faster than those of NVIDIA Tesla P4 GPUs. When processing deep learning workloads, a server configured with eight NVIDIA Tesla P40 GPUs equals over 100 servers configured only with CPUs.
- P1 ECSs use NVIDIA Tesla P100 GPUs and provide flexibility, high performance, and cost-effectiveness. These ECSs support GPU Direct for direct communication between GPUs, improving data transmission efficiency. P1 ECSs provide outstanding universal computing capabilities and have strengths in deep learning, graphic databases, high-performance databases, Computational Fluid Dynamics (CFD), computing finance, seismic analysis, molecular modeling, and genomics. They are designed for scientific computing.
- Compared with P1 ECSs, P2 ECSs use NVIDIA Tesla V100 GPUs, which have improved both single- and double-precision computing capabilities by 50% and offer 112 TFLOPS of deep learning.

Common Functions Supported by ECSs

Table 2-4 Common ECS functions supported

Function	Supported
Offline flavor change	Yes
Online flavor change	No
Cold migration	Yes
Live migration	No
HA	Yes
ECS snapshot	Yes
Cloning	Yes
CPU QoS	Yes
Memory overcommitment	Yes

Function	Supported
Supported OSs	<p>Supported range:</p> <ul style="list-style-type: none"> Carrier users: Click here, search for FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud), and obtain the latest document. In the document, search for GPU pass-through to view the support status. Enterprise users: Click here, search for FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud), and obtain the latest document. In the document, search for GPU pass-through to view the support status.

PI1 ECSs

Features

PI1 ECSs have the following features:

- NVIDIA Tesla P4 GPUs
- Up to 5.5 TFLOPS by a single GPU for single precision
- Up to 22 TOPS INT8 by a single GPU
- 8 GB ECC GPU memory with a bandwidth of 192 Gbit/s by a single GPU
- Hardware video encoding and decoding engines embedded in GPUs for concurrent real-time 35-channel HD video transcoding and inference

Notes

Ensure that the NVIDIA driver has been installed on the image used to create a PI1 ECS. If not, install the driver after the PI1 ECS is created for computing acceleration.

Recommended flavors

Table 2-5 Flavors of PI1 computing-accelerated ECSs

Type	vCPU	Memory (GB)	Flavor Name	GPU
PI1 computing-accelerated ECSs	8	64	PI1.2xlarge.8	1×P4
	16	128	PI1.4xlarge.8	2×P4
	32	256	PI1.8xlarge.8	4×P4

PI2 ECSs

Notes

Ensure that the NVIDIA driver has been installed on the image used to create a P12 ECS. If not, install the driver after the P12 ECS is created for computing acceleration.

Recommended flavors

Table 2-6 Recommended flavors for P12 computing-accelerated ECSs

Type	vCPU	Memory (GB)	Flavor Name	GPU
P12 computing-accelerated	8	64	P12.2xlarge.8	1×P40
	16	128	P12.4xlarge.8	2×P40
	32	256	P12.8xlarge.8	4×P40

P1 ECSs

Features

P1 ECSs have the following features:

- NVIDIA Tesla P100 GPUs
- 9.3 TFLOPS for single precision and 4.7 TFLOPS for double precision
- Comprehensive basic capabilities
Networks are user-defined, subnets can be divided, and network access policies can be configured as needed. Mass storage is used, and elastic capacity expansion as well as backup and restoration is supported to make data more secure. Auto Scaling allows you to add or reduce the number of ECSs quickly.
- 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 P1 ECSs.

Notes

Ensure that the NVIDIA driver has been installed on the image used to create a P1 ECS. If not, install the driver after the P1 ECS is created for computing acceleration.

Recommended flavors

Table 2-7 Recommended flavors for P1 computing-accelerated ECSs

Type	vCPU	Memory (GB)	Flavor Name	GPU
P1 computing-accelerated	8	64	P1.2xlarge.8	1×P100
	16	128	P1.4xlarge.8	2×P100

Type	vCPU	Memory (GB)	Flavor Name	GPU
	32	256	P1.8xlarge.8	4×P100

P2 ECSs

Features

P2 ECSs have the following features:

- NVIDIA Tesla V100 GPUs
- 14 TFLOPS of single-precision computing, 7 TFLOPS of double-precision computing, and 112 TFLOPS of deep learning
- Comprehensive basic capabilities

Networks are user-defined, subnets can be divided, and network access policies can be configured as needed. Mass storage is used, and elastic capacity expansion as well as backup and restoration is supported to make data more secure. Auto Scaling allows you to add or reduce the number of ECSs quickly.

- 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.

Notes

Ensure that the NVIDIA driver has been installed on the image used to create a P2 ECS. If not, install the driver after the P2 ECS is created for computing acceleration.

Recommended flavors

Table 2-8 Recommended flavors for P2 computing-accelerated ECSs

Type	vCPU	Memory (GB)	Flavor Name	GPU
P2 computing-accelerated	8	64	P2.2xlarge.8	1×V100
	16	128	P2.4xlarge.8	2×V100
	32	256	P2.8xlarge.8	4×V100

2.3 vGPU-accelerated ECSs

Application Scenarios

GPU virtualization indicates that a physical GPU can be virtualized into multiple virtual GPUs (vGPUs) with the support of hardware so that multiple vGPUs can be used by multiple VMs. GPU virtualization accelerates 2D graphics processing and 3D graphics rendering for multiple VMs at the same time. In GPU virtualization scenarios, each VM can directly access some hardware resources of the physical GPU through the bound vGPU device. (All vGPU devices can access and share the 3D graphics engine and video encoding and decoding engine of the physical GPU at different times and have independent video RAM.) VMs using vGPUs are like using passthrough physical GPUs, providing good GPU performance.

- **Designer (computing- and rendering-intensive):** suitable for 3D graphics final assembly designers and professional design personnel in the computer-aided design (CAD), computer-aided engineering (CAE), and computer-aided manufacturing (CAM) fields. These fields require a large number of computing and rendering resources, high display resolution, and 3D API compatibility.
- **Power User (computing and rendering medium-load):** suitable for 3D graphics component designers to process services, such as component-level editing or drawing viewing. In addition to basic 2D functions, 3D hardware acceleration, computing and rendering performance, and high cost-effectiveness are required.

NOTE

To use vGPU-accelerated ECSs, plan the corresponding host group when you use HUAWEI CLOUD Stack Deploy to install the ECSs. Otherwise, you need to complete a series of configurations. For details, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > Configuring vGPU-accelerated ECSs**.

Constraints

- Only one type of GPU can be installed on a host.
- For details about the servers that support GPU virtualization, see [A Supported vGPU Types](#).
- VM live migration is not supported.
- Flavors of running VMs cannot be changed.
- A vGPU device can be attached to only one vGPU VM.
- Each physical GPU can be virtualized into only the same type of vGPU devices.
- A physical GPU that is being used in passthrough mode cannot be virtualized into a vGPU device.
- A physical GPU that has been virtualized into a vGPU device cannot be used for passthrough.
- The smart memory overcommitment function can be enabled for vGPU VMs. Before starting a VM, you need to enable the memory overcommitment function on the node where the VM is deployed. In this case, the memory reservation value of the vGPU VM must be 100%.

- Only the x86 architecture is supported.

2.4 Ultra-high I/O ECSs

Application Scenarios

Ultra-high I/O ECSs use high-performance local NVMe SSDs as data disks to provide high storage input/output operations per second (IOPS) and low read/write latency. The ratio of memory to vCPU is 8:1, excepting the ECSs with 60 vCPUs. You can create such ECSs with high-performance NVMe SSDs attached on the management console.

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

NOTE

- When KVM virtualization and ultra-high I/O ECSs are used, you are advised to plan the corresponding host group during installation and deployment using HUAWEI CLOUD Stack Deploy. Otherwise, you need to complete a series of configurations. For details, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > Configuring Ultra-high I/O ECSs**.
- When the Arm architecture is used, ultra-high I/O ECSs must be used with general computing-plus ECSs for better performance. It is recommended that ultra-high I/O ECSs be used with general computing-plus ECSs (Arm-DPDK).

Constraints

Table 2-9 Constraints

Category	Description
ECSs	<ul style="list-style-type: none">• Ultra-high I/O ECSs support KVM virtualization only. They do not support flavor change, cold migration, live migration, HA, ECS snapshot, clone, or resource release upon ECS shutdown.• Ultra-high I/O ECSs cannot be used in a converged deployment scenario. A converged deployment means that Huawei Distributed Block Storage is deployed on compute nodes.• After an ultra-high I/O ECS is deleted, the data on the local NVMe SSD is automatically deleted. Back up the data before deleting it.

Category	Description
OSs	<ul style="list-style-type: none"> ● Ultra-high I/O ECSs support the following OSs: <ul style="list-style-type: none"> – Carrier users: Click here, search for FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud), and obtain the latest document. In the document, search for SSD card to view the support status. – Enterprise users: Click here, search for FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud), and obtain the latest document. In the document, search for SSD card to view the support status. ● When Huawei ES3000 series NVMe SSDs are used, to create an ultra-high I/O ECS using a private image, make sure that the NVMe driver has been installed during the private image creation. If not, install the driver after the ultra-high I/O ECS is created.
Disks	<ul style="list-style-type: none"> ● Local NVMe SSDs can only be used as data disks. Their reliability depends on the reliability of physical servers and hard disks, which are SPOF-prone. Therefore, you are advised to configure data redundancy at the application layer to ensure data availability. Use EVS disks to store data that needs to be retained for long periods of time. ● Local NVMe SSDs are not hot-swappable. To replace them, power off the physical server first. ● SSDs of the same type and different capacities can be installed on the same server, for example, Huawei ES3000 series 1.6 TB and 3.2 TB SSDs. However, disks from different vendors cannot be installed on the same server, for example, Samsung and Huawei disks. ● Disks and cards on a single server can be used together. ● When using Huawei ES3000 series NVMe SSDs to create ultra-high I/O ECSs, ensure that the NVMe SSD firmware version of the physical server is 3.10 or later. Otherwise, the creation will fail. For details about how to query and upgrade the firmware version of the NVMe SSD, see "Product Management" > "Resource Pools" > "FusionSphere OpenStack" > "Compute" > "Upgrading the Firmware Version of the NVMe SSD" in HUAWEI CLOUD Stack 8.1.0 O&M Guide.

Recommended Flavors

Table 2-10 Recommended flavors for ultra-high I/O ECSs

Type	vCPU (U)	Memory (GB)	Flavor Name	Local SSD Specifications (for Reference)	Disk Type
Ultra-high I/O ECSs	8	64	i3.2xlarge.8	1 x 1600 GB NVMe SSD	NVMe_SSD
	16	128	i3.4xlarge.8	2 x 1600 GB NVMe SSDs	NVMe_SSD
	32	256	i3.8xlarge.8	4 x 1600 GB NVMe SSDs	NVMe_SSD
	48	384	i3.12xlarge.8	6 x 1600 GB NVMe SSDs	NVMe_SSD
	60	512	i3.15xlarge.8	7 x 1600 GB NVMe SSDs	NVMe_SSD

NOTE

- Configure the capacity and quantity of local SSDs based on server hardware configuration.
- The I/O performance of SSDs varies depending on the vendor and model. The disk hardware specifications shall prevail. The random read/write performance of Samsung NVMe SSDs (PM1725b) is about 30% lower than that of Huawei ES3600 series SSDs.
- When the quantity of vCPUs is greater than or equal to 48, Kunpeng 920 CPU models 5220, 3210, 5221K, and 3211K are not supported.

2.5 General Computing-plus ECSs

Application Scenarios

Compared with general-purpose ECSs, general computing-plus ECSs provide vCPUs and memory of larger specifications, and higher network performance.

- Where x86 servers are used, general computing-plus ECSs can use Data Plane Development Kit (DPDK) to accelerate packet processing and provide higher network performance.
- When the Arm architecture is used, general computing-plus ECSs support both DPDK and user-mode OVS with hardware acceleration to provide high network performance. General computing-plus ECSs (Arm-DPDK) support bonding across NICs for physical link-level HA. General computing-plus ECSs (Arm-hardware acceleration) only support bonding within a single NIC. It does not provide physical link-level HA, but delivers higher network performance.

When the Arm architecture is used, you are advised to use general computing-plus ECSs (Arm-DPDK) for better performance and user experience.

NOTE

- Currently, general computing-plus ECSs support only KVM as the virtualization platform. To use general computing-plus ECSs, plan the corresponding host group when you use HUAWEI CLOUD Stack Deploy to install the ECSs. Otherwise, you need to complete a series of configurations. For details, see visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > Configuring General Computing-Plus ECSs**.
- The memory of VMs that support live migration cannot exceed 512 GB.

Constraints

Table 2-11 Common functions supported by general computing-plus ECSs

Function	Scenario	Supported
Offline flavor change	x86	Yes. An x86-based ECS flavor can only be changed to another x86-based one.
	Arm	Yes. An Arm-based ECS flavor can only be changed to another Arm-based one.
Online flavor change	x86	No
	Arm	No
Cold migration	x86/Arm	Yes
Live migration	x86/Arm	Yes
HA	x86/Arm	Yes
ECS snapshot	x86/Arm	Yes
Cloning	x86/Arm	Yes
CPU QoS	x86/Arm	Yes
Memory overcommitment	x86	No
	Arm	No

Recommended Flavors

Table 2-12 Recommended flavors for general computing-plus ECSs

Type	vCPU	Memory (GB)	Flavor Name
General computing-plus ECSs	2	8	c3.large.4
	4	16	c3.xlarge.4
	8	32	c3.2xlarge.4
	16	64	c3.4xlarge.4
	32	128	c3.8xlarge.4
	60	256	c3.16xlarge.4

NOTICE

When the quantity of vCPUs is greater than or equal to 48, Kunpeng 920 CPU models 5220, 3210, 5221K, and 3211K are not supported.

2.6 USB Passthrough ECSs

Application Scenarios

When selecting the USB-passthrough type, you can create an ECS to which a USB device on a physical server is attached. In addition, the administrator can detach the USB device and attach it to another ECS on Service OM. Some applications will run properly only when they work with a license device, such as a USB dongle. You can use USB passthrough ECSs to deploy applications of this kind.

NOTE

- In the scenario where the x86 or Arm architecture and KVM virtualization is used, to use USB passthrough ECSs, plan the corresponding host group when you use HUAWEI CLOUD Stack Deploy to install the ECSs. Otherwise, you need to complete a series of configurations. For details, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > Configuring USB Passthrough ECSs**.
- When the Arm architecture is used, USB passthrough ECSs must be used with general computing-plus ECSs for better performance. It is recommended that ECSs of this type be used with general computing-plus ECSs (Arm-DPDK).

Constraints

- Supported USB Models
USB 1.0, USB 2.0, and USB 3.0 when the x86 architecture is used
USB 2.0 and USB 3.0 when the Arm architecture is used

- OSs supported by USB-passthrough ECSs in KVM scenarios
 - Carrier users: Click [here](#), search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)**, and obtain the latest document. In the document, search for **USB passthrough** to view OSs that can be used by USB-passthrough ECSs. (FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Enterprise Virtualization) / USB device pass-through)
 - Enterprise users: Click [here](#), search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)**, and obtain the latest document. In the document, search for **USB passthrough** to view OSs that can be used by USB-passthrough ECSs.
- Common functions supported by USB passthrough ECSs

Table 2-13 Common functions supported by USB passthrough ECSs

Function	Supported
Offline flavor change	Yes, but a USB-passthrough ECS can only be changed to another type of ECS through offline flavor change.
Online flavor change	Yes
Cold migration	No
Live migration	No
HA	No
ECS snapshot	Yes, but an ECS snapshot does not contain data stored on the USB device.
Cloning	No
CPU QoS	Yes
Memory overcommitment	Yes

Recommended Flavors

Type	vCPU	Memory (GB)	Flavor Name	USB Device	Virtualization Type
USB passthrough ECSs	1	4	u.medium.4	1 x USB	KVM
	2	8	u.large.4	1 x USB	KVM
	4	16	u.xlarge.4	1 x USB	KVM
	8	32	u.2xlarge.4	1 x USB	KVM
	16	64	u.4xlarge.4	1 x USB	KVM

2.7 Dedicated general-purpose ECSs

Dedicated general-purpose ECSs provide stable and dedicated CPU and memory resources, and are suitable for web services and small-scale database applications with high demands on CPU and memory performance.

Application Scenarios

Dedicated general-purpose ECSs are suitable for applications, such as large games, that require dedicated CPU and memory resources.

Recommended Flavors

Table 2-14 Recommended flavors for KVM-based dedicated general-purpose ECSs

Type	vCPU	Memory (GB)	Flavor Name	Virtualization Type
Dedicated general-purpose ECSs	2	8	cc3.large.4	KVM
	4	16	cc3.xlarge.4	KVM
	8	32	cc3.2xlarge.4	KVM
	16	64	cc3.4xlarge.4	KVM
	32	128	cc3.8xlarge.4	KVM
	76	304	cc3.19xlarge.4	KVM

NOTICE

When the quantity of vCPUs is greater than or equal to 48, Kunpeng 920 CPU models 5220, 3210, 5221K, and 3211K are not supported.

2.8 Memory-optimized ECSs

Application Scenarios

Memory-optimized ECSs are developed based on the KVM virtualization platform and designed for processing large-scale data sets in the memory. They provide a maximum memory size of 512 GB based on DDR4 for high-memory computing applications. Memory-optimized ECSs are suitable for applications that require a large amount of memory, process large volumes of data, and demand rapid data switching and processing. The scenarios include precision advertising, e-commerce big data analysis, and IoT big data analysis.

Recommended Flavors

Table 2-15 Recommended flavors for memory-optimized ECSs

Type	Flavor Name	vCPU	Memory (GB)	Virtualization Type
Memory-optimized ECSs	m3.large.8	2	16	KVM
	m3.xlarge.8	4	32	KVM
	m3.2xlarge.8	8	64	KVM
	m3.3xlarge.8	12	96	KVM
	m3.4xlarge.8	16	128	KVM
	m3.6xlarge.8	24	192	KVM
	m3.8xlarge.8	32	256	KVM
	m3.15xlarge.8	60	512	KVM

NOTICE

When the quantity of vCPUs is greater than or equal to 48, Kunpeng 920 CPU models 5220, 3210, 5221K, and 3211K are not supported.

2.9 Disk-intensive ECSs

Application Scenarios

Disk-intensive ECSs use local storage, and provide better sequential read and write performance and lower latency by using pass-through HDDs as data disks. They use a vCPU to memory ratio of 1:8 and provide powerful and stable computing performance, ensuring efficient data processing. They provide high intranet performance, including high intranet bandwidth and packets per second (pps), meeting the requirements for data exchange between ECSs during peak hours.

Disk-intensive ECSs are suitable for scenarios that require high I/O performance and rapid data switching and processing to handle massive data sets. Such scenarios include MapReduce computing, distributed Hadoop computing, large data warehouse, distributed file system, data processing, and log processing.

 NOTE

- To use disk-intensive ECSs, plan the corresponding host group when you use HUAWEI CLOUD Stack Deploy to install the ECSs. Otherwise, you need to complete a series of configurations. For details about how to manually configure disk-intensive ECSs when KVM virtualization is used, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > Configuring Disk-intensive ECSs**.
- When the Arm architecture is used, disk-intensive ECSs must be used with general computing-plus ECSs for better performance. It is recommended that ECSs of this type be used with general computing-plus ECSs (Arm-DPDK).

Notes

- Disk-intensive ECSs do not support live migration, cold migration, VM HA, cloning, or ECS snapshot, but they support memory overcommitment and user-defined CPU QoS settings. A disk-intensive ECS cannot be changed into another type of ECS by changing its flavor. Disk-intensive ECSs cannot be used in a converged deployment scenario. A converged deployment means that Huawei Distributed Block Storage is deployed on compute nodes.
- Disk-intensive ECSs do not support hot-swapping, snapshot, or backup for pass-through HDDs.
- Local and EVS disks can both be used to store data, but pass-through HDDs can only be used as data disks. For details about the disks that can be used by disk-intensive ECSs and the number of disks that can be attached, see [Table 3-1](#).
- The local disk data of a disk-intensive ECS may be lost due to some events, such as host breakdown or local disk damage. Do not use local disks for persistent storage. Instead, back up data in a timely manner and use a high availability data architecture. If your application cannot provide the desired data reliability, you are advised to use EVS disks to build your ECS and store data persistently on EVS disks.
- When you delete a disk-intensive ECS, the data on the local disk is automatically deleted.
- When modifying the flavor of a disk-intensive ECS, you can add more local disks to expand local disk capacity. You cannot expand or reduce the capacity of existing local disks.
- Resources will not be released for disk-intensive ECSs when they are shut down.

Flavors

Table 2-16 Recommended flavors for disk-intensive ECSs

Type	vCPU	Memory (GB)	Flavor Name	Virtualization Type	Local Disks	Capacity of One Local Disk
Disk-intensive ECSs	4	32	d2.xlarge.8	KVM	2	1800 GB
	8	64	d2.2xlarge.8	KVM	4	1800 GB

Type	vCPU	Memory (GB)	Flavor Name	Virtualization Type	Local Disks	Capacity of One Local Disk
	16	128	d2.4xlarge.8	KVM	8	1800 GB
	24	192	d2.6xlarge.8	KVM	12	1800 GB
	32	256	d2.8xlarge.8	KVM	16	1800 GB
	60	540	d2.15xlarge.9	KVM	24	1800 GB

NOTICE

When the quantity of vCPUs is greater than or equal to 48, Kunpeng 920 CPU models 5220, 3210, 5221K, and 3211K are not supported.

2.10 Large-memory ECSs

Application Scenarios

Large-memory ECSs are used for applications that require a large amount of memory, rapid data switching, and low latency, and process large volumes of data. These ECSs are suitable for OLAP scenarios, such as in-memory databases, big data processing engines, and data mining.

Currently, only 2288H V5 and 2488H V5 servers can be used to create large-memory ECSs.

 **NOTE**

ARM servers do not support large-memory ECSs.

Recommended Flavors

Table 2-17 Recommended flavors for large-memory ECSs

Type	vCPU	Memory (GB)	Flavor Name
Large-memory ECSs	12	174	e3.3xlarge.14
	24	348	e3.6xlarge.14
	48	696	e3.12xlarge.14
	12	353	e3.3xlarge.28

Type	vCPU	Memory (GB)	Flavor Name
	24	706	e3.6xlarge.28
	48	1412	e3.12xlarge.28
	96	2824	e3.24xlarge.28

NOTICE

When the quantity of vCPUs is greater than or equal to 48, Kunpeng 920 CPU models 5220, 3210, 5221K, and 3211K are not supported.

2.11 AI-accelerated ECSs

Application Scenarios

Huawei's Ascend 310 chips provide extraordinary computational power. They provide a built-in hardware video codec engine that supports a 16-channel HD video decoder. AI-accelerated ECSs use Huawei Ascend 310 chips. Leveraging the remarkable features of Ascend 310 chips, such as low power consumption and high computing capacity, AI-accelerated ECSs significantly improve energy efficiency. They are suitable for scenarios that require real-time massive concurrent computing, such as, AI inference, machine learning, and video encoding and decoding. AI-accelerated ECSs can be used to run workloads related to machine vision, voice recognition, and natural language processing and support a range of scenarios, including smart retail, smart campus, robot brains on cloud, and safe city.

NOTE

- To use AI-accelerated ECSs, plan the corresponding host group when you use HUAWEI CLOUD Stack Deploy to install the ECSs. Otherwise, you need to complete a series of configurations. For details about how to manually configure AI-accelerated ECSs, visit **Operation Help Center** and choose **Operation** > **Compute Services** > **Elastic Cloud Server (ECS)** > **Configuring AI-accelerated ECSs**.
- When Arm servers are used, AI-accelerated ECSs must be used with general computing-plus ECSs for better performance. It is recommended that AI-accelerated ECSs be used with general computing-plus ECSs (Arm-DPDK).
- Ascend chip development guide: <https://ascend.huawei.com/home>

Chip Information

- Each accelerator card contains four Ascend 310 chips.
- Single-chip half-precision (FP16): 8TeraFLOPS. Single-chip integer precision (INT8): 16TeraOPS.
- A single chip provides 8 GB of GPU memory with a memory bandwidth of 50 GB/s.
- The built-in hardware video codec engine supports a 16-channel HD video decoder (H.264/H.265).

Constraints

- Both Arm and x86 servers are supported, but only KVM virtualization is supported.
- Supported physical servers and OSs: Click [here](#) to start a query.
- Common functions supported by ECSs

Table 2-18 Common functions supported by ECSs

Function	Supported
Offline flavor change	Yes
Online flavor change	No
Cold migration	Yes
Live migration	No
HA	Yes
ECS snapshot	Yes
Cloning	Yes
CPU QoS	Yes
Memory overcommitment	No

Recommended Flavors

Table 2-19 Recommended flavors for AI-accelerated ECSs

Flavor	vCPUs	Memory (GB)	Quantity of Ascend 310 Chips	Ascend RAM (GB)	Virtualization Type
Ai1.large.4	2	8	1	8	KVM
Ai1.xlarge.4	4	16	2	16	KVM
Ai1.2xlarge.4	8	32	4	32	KVM
Ai1.4xlarge.4	16	64	8	64	KVM

3 Related Concepts

3.1 Regions and AZs

A region is a geographic area where resources used by your ECSs are located. ECSs can be created in different regions so that applications can be designed to meet specific user requirements, reduce network latency, or comply with local laws or regulations.

An availability zone (AZ) is a physical region where resources use independent power supply and networks. AZs in the same region can communicate with each other through the internal network and provide cost-effective and low-latency network connections. AZs are physically isolated from each other. An AZ is not adversely affected by another faulty AZ because each AZ uses independent power supply and networks. Therefore, you can create your ECSs in multiple AZs to ensure that your applications in one AZ will not be adversely affected by a fault in another AZ.

3.2 Cloud-Init

Cloud-Init is an open-source cloud initialization program that can initialize custom configurations, such as the host name, key, and user data, for an ECS.

To use Cloud-Init, the following requirements must be met:

- For an ECS that runs Windows, install Cloudbase-Init.
- For an ECS that runs Linux, install Cloud-Init.

After Cloud-Init or Cloudbase-Init is installed in an image, you can configure the initial attributes of an ECS when creating the ECS.

Note the following when using Cloud-Init:

- When creating an ECS, if the selected image supports Cloud-Init, you can use user data injection to inject customized initial configurations into the ECS, such as the ECS login password, to initialize the configurations of the ECS. For details, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > Configuration Before ECS Creation > Creating a User Data Injection Script**.

- If a running ECS supports Cloud-Init, you can view and use the ECS metadata to configure and manage the ECS. For more information, see [14.3 Viewing Information Through Metadata](#).

3.3 Local Disk, EVS Disk and Local Pass-through Disk

Disks used by an ECS include local disks, EVS disks, and local passthrough disks.

- A local disk refers to a disk residing on the local host of the ECS to provide non-persistent storage. This type of storage provides high I/O throughput and low latency but cannot be used for persistent data storage.

An ECS loses its local disk data after it is rebuilt on another host. Local disks cannot be live migrated, but can be cold migrated. Cold migration does not cause data loss, but takes a long time because disk files need to be copied from the source to the destination node during the migration.

- Elastic Volume Service (EVS) is a virtual block storage service that provisions block storage to Elastic Cloud Servers (ECSs) and Bare Metal Servers (BMSs) from storage backends. Users can create EVS disks online and attach them to ECSs. Users can use EVS disks the same way they use traditional hard disks on servers. EVS disks deliver higher data reliability and I/O throughput and are easy to use. They can be used for file systems, databases, or other system software and applications that require block storage resources. EVS disks provide persistent storage, meaning that the data stored on them does not get lost when ECS instances are started, stopped, or migrated.

EVS disks are categorized as Virtual Block Device (VBD) or Small Computer System Interface (SCSI), depending on whether advanced SCSI commands are supported.

- A local pass-through disk allows an ECS to have direct access to the disk space on its host. Local pass-through disks offer high read/write speeds and low latency and are suitable for scenarios that require high I/O performance and rapid data switching and processing to handle massive data sets.

Currently, ECSs that use local pass-through disks are as follows:

- Ultra-high I/O ECSs: They use high-performance local NVMe SSDs as data disks. They do not support flavor changes (online or offline), cold migration, live migration, HA, ECS snapshot, cloning, memory overcommitment, or user-defined CPU QoS settings. For more information about ultra-high I/O ECSs, see [2.4 Ultra-high I/O ECSs](#).
- Disk-intensive ECSs: They use pass-through HDDs as data disks. They do not support live migration, cold migration, HA, cloning, or ECS snapshot, but they support memory overcommitment and user-defined CPU QoS settings. A disk-intensive ECS cannot be changed into another type of ECS by changing its flavor. For more information about disk-intensive ECSs, see [2.9 Disk-intensive ECSs](#).

The total number of system and data disks cannot exceed 60. [Table 3-1](#) lists the categories of local disks, EVS disks and local pass-through disks and the quantity of disks that can be attached.

Table 3-1 Disk type description

Category	Disk Type	Purpose	Description
Local disk	HDD	Used as both the system disk and data disks. When used as data disks, only one local HDD can be attached to an ECS.	<p>The performance of local disks varies with the load on the physical host and single points of failure (SPOFs) may exist. Local disks are suitable for systems that run only for a short period of time and have relatively low stability and reliability standards.</p> <p>You are advised to configure data redundancy at the application layer and synchronize or back up important data from local disks to other ECSs or EVS disks in a timely manner, ensuring data availability.</p>
EVS disk	SCSI disk	<p>Used only as data disks.</p> <p>A maximum of 59 such disks can be attached.</p>	<p>EVS disks of this type support transparent SCSI command transmission and allow the ECS OS to directly access the underlying storage media. SCSI EVS disks support advanced SCSI commands (such as SCSI-3 persistent pre-lock) in addition to basic SCSI read and write commands. They can be used in cluster scenarios where data security is enhanced by using the SCSI lock mechanism, such as the Windows MSCS cluster.</p>
	VBD disk	<p>Used as both the system disk and data disks.</p> <p>The number of disks that can be attached to an ECS depends on Disk Device Type set for the ECS image when the image was registered on Service OM. For details, see Table 3-2.</p>	<p>EVS disks of this type support only basic SCSI read/write commands. They are mostly used in common scenarios like OA and testing, or common Linux clusters such as RHCS.</p>

Category	Disk Type	Purpose	Description
Local pass-through disk	NVMe SSD	Used only as data disks. A maximum of eight NVMe SSDs can be used to create an ultra-high I/O ECS.	Ultra-high I/O ECSs use high-performance local NVMe SSDs as data disks to provide high storage IOPS and low read/write latency. Disks of this type can be used for high-performance relational databases, NoSQL databases (such as Cassandra and MongoDB), and ElasticSearch.
	Pass-through HDD	Used only as data disks. A maximum of 59 pass-through HDDs can be used to create a disk-intensive ECS.	Disk-intensive ECSs use pass-through HDDs as data disks to provide a higher sequential reading performance and a lower latency, improving file read and write performance. Disks of this type are suitable for scenarios that require high I/O performance and rapid data switching and processing to handle massive data sets. The scenarios include MapReduce computing, Hadoop distributed computing, large data warehouse, distributed file system, data processing, and log processing.

 **NOTE**

- If you create an ECS earlier than FusionSphere Service 6.3.1, a maximum of 12 disks can be attached to your ECS.
- If the number of disks that can be attached to an ECS is less than the number that you specify, some drive letters have been pre-occupied by the system.

Table 3-2 Total number of VBD disks that can be attached

Disk Device Type	Total VBD Disks
ide	4 (x86)
	0 (ARM) NOTE When ARM servers are used, Disk Device Type cannot be set to ide during image registration.

Disk Device Type	Total VBD Disks
virtio	24 NOTE <ul style="list-style-type: none">When ARM servers are used, if the ECS bus type is Virtio, the total number of NICs cannot exceed 16, the total number of VBD disks cannot exceed 24, and the total number of NICs, disks (EVS disks and pass-through disks), and NPU cards cannot exceed 24.When x86 servers are used, if Boot Mode of the image is set to UEFI during ECS creation: During online disk attachment, the mount point must be between vda and vdp. A maximum of 16 disks (including system disks) can be attached. If the mount point exceeds vdp, for example, vdq, you must shut down the ECS, attach the target disk, and then start the ECS. Such a process is an offline disk attachment process.
scsi	60

3.4 Same Storage

Same storage means that the system and data disks of an ECS reside on the same storage backend. If the backend storage configures the storage tag, the ECS supports DR and backup.

- When you create an ECS, if **Boot Device** of the selected ECS flavor is set to **Cloud Disk** and the selected ECS supports the customization of **Same Storage**, select **Yes**, indicating that the system and data disks of the ECS are located in the same storage. If you select **No**, there is no such restriction.
- When creating an ECS, you can customize **Same Storage** for the selected ECS. However, if **Same Storage** is set to **No** for an ECS, you can set whether to support DR or backup after the ECS is created. For details, see [7.4 Modifying the DR or Backup Function of an ECS](#).

To support DR or backup, the following conditions must be met:

- Boot Device** of the ECS flavor must be **Cloud Disk**.
- If the ECS only has the system disk, make sure that the storage backend where the system disk resides has the storage tag configured. For details, visit [Operation Help Center > Operation > Compute Services > Elastic Cloud Server \(ECS\) > FAQs > Disk FAQs > \(Optional\) Creating a Disk Type](#).
- If the ECS has both system and data disks, these disks must reside in the same storage backend, that is, they must have the same storage tag. For details, visit [Operation Help Center > Operation > Compute Services > Elastic Cloud Server \(ECS\) > FAQs > Disk FAQs > \(Optional\) Creating a Disk Type](#).

3.5 Feature Differences Between Arm and x86 Servers

Advanced RISC Machine (Arm) uses the reduced instruction set computer (RISC), and the Intel processor uses the complex instruction set computer (CISC).

The following table lists the ECS features supported by the Arm server, and the features similar to what the x86 server provides are not presented:

Table 3-3 Feature differences between Arm and x86 servers

Category	Feature	Arm Server (TaiShan 200 Server)	x86 Server
Basic features	ECS Type	<ul style="list-style-type: none">The following types of ECSs are not supported: GPU and large-memory ECSs.You are advised to use general computing-plus ECSs with ultra-high I/O, disk-intensive, USB-passthrough, or AI-accelerated ECSs for better performance.	All available types of ECSs are supported.
	Availability Zone	In an AZ, only Arm servers or x86 servers are deployed.	
	Virtualization Type	Only KVM is supported.	Only KVM is supported.
ECS flavors	NUMA Affinity	It is strongly recommended that you enable NUMA affinity in Arm scenarios. If you do not enable NUMA affinity, ECSs may not provide excellent performance.	Configure the parameter based on the ECS type. For details, see visit Operation Help Center and choose Operation > Compute Services > Elastic Cloud Server (ECS) > Configuration Before ECS Creation > Creating a Flavor .

Category	Feature	Arm Server (TaiShan 200 Server)	x86 Server
	NUMA Nodes	<p>For the Arm server, you are advised to deploy the following quantity of NUMA nodes depending on the CPU vendor:</p> <ul style="list-style-type: none"> Phytium: Enter an integer between 1 to 8 that can be exactly divided by vCPU and Memory. Recommended setting: Number of vCPUs/ Number of NUMA nodes ≤ 8 For example, if the number of vCPUs is less than or equal to 8, set this parameter to 1. If the number of vCPUs is greater than 8 and less than or equal to 16, set this parameter to 2. Hisilicon: Enter 1, 2, or 4 that can be exactly divided by vCPU and Memory. Recommended setting: When the number of vCPUs is less than or equal to 16, set this parameter to 1. If the number of vCPUs is greater than 16 and less than or equal to 32, set this parameter to 2. If the number of vCPUs is greater than 32, set this parameter to 4. 	<p>Enter an integer between 1 to 8 that can be exactly divided by vCPU and Memory.</p>

Cat ego ry	Feature	Arm Server (TaiShan 200 Server)	x86 Server
		<p>NOTICE</p> <ul style="list-style-type: none"> When the quantity of vCPUs is less than or equal to 48 and the Kunpeng 920 CPU model is 5220, 3210, 5221K, or 3211K, NUMA Nodes can be set to 1 or 2. When the quantity of vCPUs is greater than or equal to 48, Kunpeng 920 CPU models 5220, 3210, 5221K, and 3211K are not supported. 	
	CPU QoS	Only KVM is supported.	Only the KVM scenario is supported.
	Hyperthreade d Core Sharing Mode	Not supported	Supported
Host Gro up	Memory Overcommit ment Ratio	Not supported	Supported NOTE If the flavor of the hugepage memory is configured, memory overcommitment is not supported.
	Host Type	A host group can accommodate only Arm or x86 hosts.	
Spe cific atio ns	vCPU	The quantity of vCPUs must not exceed the quantity of vCPUs available on physical servers. For details about how to view the quantity of vCPUs available on a physical host, visit Operation Help Center and choose Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > How Do I Check the Number of vCPUs Available on a Physical Host?	255 or fewer vCPUs are supported and the quantity of vCPUs must not exceed the quantity of vCPUs available on physical servers. For details about how to view the quantity of vCPUs available on a physical host, visit Operation Help Center and choose Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > How Do I Check the Number of vCPUs Available on a Physical Host?
	Memory	256 GB or smaller	4096 GB or smaller

Cat ego ry	Feature	Arm Server (TaiShan 200 Server)	x86 Server
	Online flavor change	Not supported	Supported
	Offline flavor change	An existing ECS flavor can be replaced with only an Arm ECS flavor.	An existing ECS flavor can be replaced with only an x86 ECS flavor.
Sup por ted OSs	Supported OS	<ul style="list-style-type: none"> Carrier users: Click here and search for FusionSphere SIA Huawei Guest OS Compatibility Guide (ARM). Enterprise users: Click here and search for FusionSphere SIA Huawei Guest OS Compatibility Guide (ARM). 	For details about supported OSs, see the information provided on the server website.
	Changing an ECS OS	An existing ECS OS can be replaced with only an Arm OS.	An existing ECS OS can be replaced with only an x86 OS.
	OSs that support the one-click password reset plugin	All Windows OSs do not support it. For details about Linux OSs that support it, see Table 5-4 .	For Windows OSs, see Table 5-3 . For Linux OSs, see Table 5-4 .
	Image Boot Mode	Only the UEFI mode is supported.	Both the BIOS and UEFI modes are supported.
Dis ks	Quantity of Attached Disks	For details, see Table 3-1 and Table 3-2 .	
Net wor ks	Number of NICs that can be attached to an ECS	16 If the ECS bus type is Virtio, the total number of NICs, disks (EVS disks and pass-through disks), and NPU cards cannot exceed 24.	16
	Number of NICs	When Arm servers are used, the total number of NICs for all ECSs on a physical host cannot exceed 158.	The number of NICs is related to the model of the physical server in use.

Category	Feature	Arm Server (TaiShan 200 Server)	x86 Server
Others	Watchdog	The 6300ESB watchdog is used.	The IPMI watchdog is used.
	ECS Bus Type	The IDE bus is not supported, but the Virtio bus and the SCSI bus are supported.	The IDE bus, the Virtio bus, and the SCSI bus are supported.

Table 3-4 Feature differences among Intel, Kunpeng, Hygon, and Phytium servers

Category	Feature	Intel (x86)	Kunpeng (Arm)	Hygon (x86)	Phytium (Arm)
Resource management	ECS resource QoS (CPU QoS)	Supported	Supported	Not supported	Not supported
	Cross-generation CPU live migration	Supported	Supported	Not supported	Not supported
VM performance tuning	High-precision VM	Supported	Supported	Not supported	Not supported
	6300ESB Watchdog	Supported	Supported	Not supported	Supported
	VM installation in PXE mode	Supported	Supported	Not supported	Not supported
	PCI passthrough	Supported	Supported	Supported	Not supported
I/O-optimization	NVMe SSD passthrough	Supported	Supported	Supported	Not supported
	Da Vinci card passthrough (Arm +x86)	Supported	Supported	Not supported	Not supported
	PV channel (fault notification channel/short-circuit notification)	Supported	Supported	Not supported	Not supported

3.6 ECS Quota

A quota is a resource management and control technology that allocates and manages the maximum number of resources (including resource capacity and quantity) available to a single VDC, preventing resources from being overused by

users in a single VDC and affecting other VDCs. The platform allows users to set ECS quotas for VDCs at all levels.

ECS quotas include:

- vCPUs
- Memory (GB)
- GPUs
- NPUs
- ECS snapshots
- ECS instances

If the number of resources in a VDC reaches the quota value, the resources cannot be requested. Delete idle resources or contact the administrator to modify the quota. For details about how to change quotas, visit **Operation Help Center** and choose **Operation > Operation Management > Managing Organizations > Managing Quotas**.

4 Operation Process

Figure 4-1 and **Table 4-1** show the operation process and main functions for creating an ECS.

Figure 4-1 Procedure for creating and managing an ECS

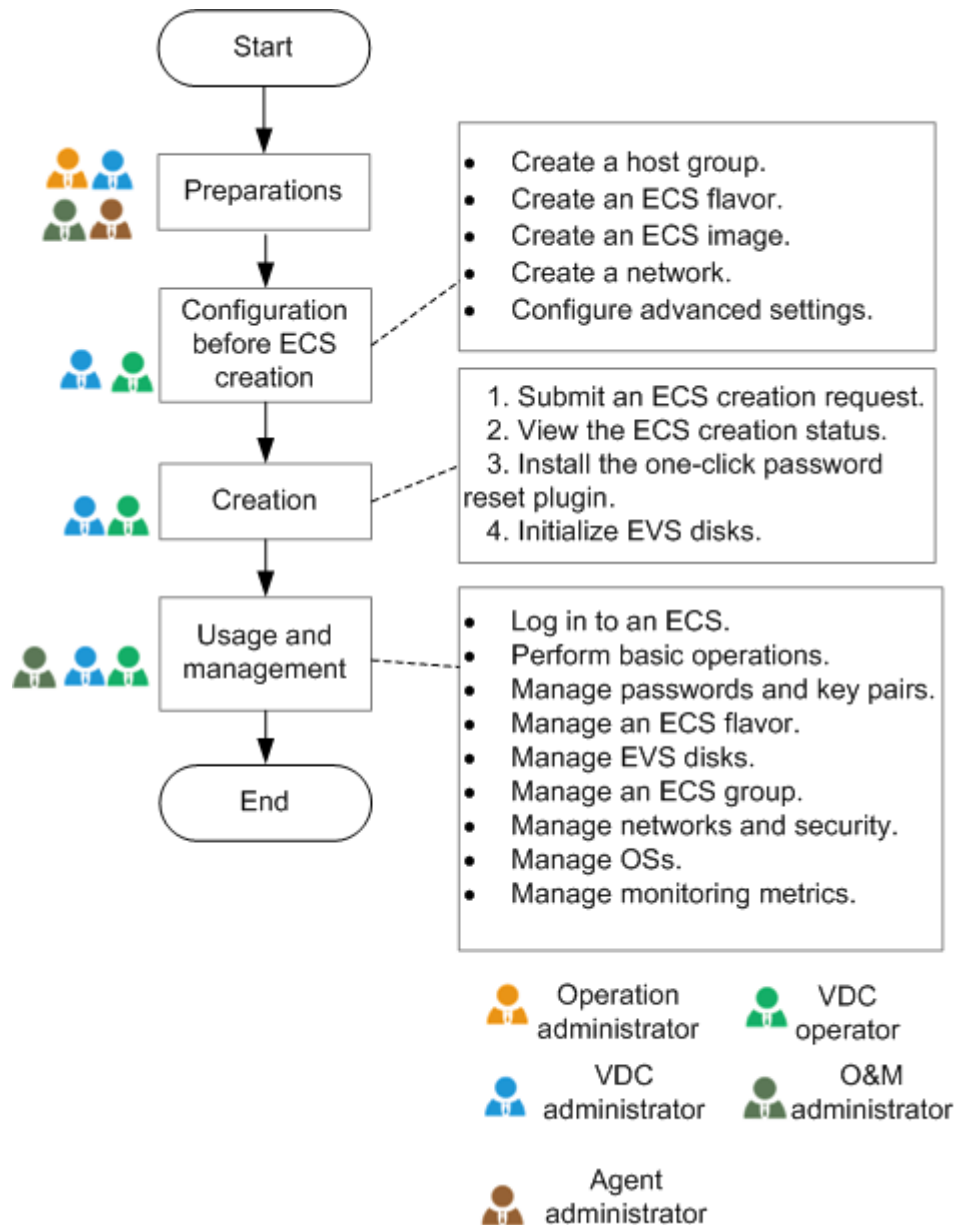


Table 4-1 Process description

Operation	Description and Reference	Operator and UI Portal
Preparation	<p>Dividing cloud resources:</p> <p>If no accounts are available, contact an operation administrator to create a VDC and VDC administrator, and then use the VDC administrator account to create a VDC operator.</p> <p>For details, visit Operation Help Center and choose Operation > VDC Tenant Modeling.</p>	<p>ManageOne Operation Portal (ManageOne Operation Management Portal in B2B scenarios): operation administrator</p> <p>ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios):</p> <ul style="list-style-type: none"> • VDC administrator • Agent administrator
Configuration before ECS creation	<p>For details, visit Operation Help Center and choose Operation > Compute Services > Elastic Cloud Server (ECS) > Configuration Before ECS Creation.</p>	<ul style="list-style-type: none"> • Service OM: <ul style="list-style-type: none"> - O&M administrator - Agent administrator • ManageOne Operation Portal (ManageOne Operation Management Portal in B2B scenarios): operation administrator • ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios): <ul style="list-style-type: none"> - VDC administrator - Agent administrator
Creation	<p>When creating an ECS, configure the parameters for creating an ECS and submit the application.</p> <p>5.2 Applying for an ECS</p>	<p>ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios):</p>

Operation	Description and Reference	Operator and UI Portal
	<p>After submitting an ECS creation request, you can query the task status in the Task Status area.</p> <p>5.3 Viewing ECS Creation Status</p> <p>For an ECS whose virtualization type is KVM, you can install the one-click password reset plugin to reset the password with one click when the password is lost or expires. Without this plugin, you need to manually reset the password of an ECS when the password is lost or expires.</p> <p>5.4 Installing the One-Click Password Reset Plugin</p> <p>If you add data disks when creating an ECS, you need to initialize the data disks after creating the ECS.</p> <p>5.5 Initializing EVS Data Disks</p>	<ul style="list-style-type: none"> ● VDC administrator ● VDC operator
Usage and Management	<p>After logging in to an ECS, you can view information about the ECS, modify configuration of the ECS, and manage the ECS.</p> <p>6 Logging In to an ECS</p> <p>Basic operations on an ECS include viewing the ECS details, changing the ECS name, managing tags, and managing the life cycle.</p> <p>7.1 Basic Operations to 7.2 Life Cycle</p> <p>Operations on an ECS include creating an image, setting DR and backup capabilities, changing the watchdog status and VM HA status, cloning an ECS, creating an ECS snapshot, and creating a CD-ROM drive and attaching an ISO File to an ECS.</p> <p>7.3 Creating a Private Image Using an Existing ECS to 7.9 Creating a CD-ROM Drive and Attaching ISO/UVP VMTools</p> <p>You can obtain, clear, and modify the ECS password, and you can create a key pair.</p> <p>8 Passwords and Key Pairs</p>	<p>ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios):</p> <ul style="list-style-type: none"> ● VDC administrator ● VDC operator

Operation	Description and Reference	Operator and UI Portal
	You can change the number of vCPUs or memory size of an ECS. 9 ECS Flavors	
	You can request data disks, attach data disks to an ECS, expand the capacity of data disks of an ECS, and detach data disks from an ECS. 10 EVS Disk	
	You can configure the mode through which ECSs communicate with each other and with the Internet, and manage the NIC, security group, EIP, and private IP address. 12 Network and Security	
	You can reinstall or change an ECS OS. 13 Operating Systems	
	You can view the monitoring metrics, running status, and metadata of an ECS. 14 Monitoring Metrics	

5 Creating an ECS

5.1 Overview

After configuration selection, you can start creating ECSs. A complete procedure for creating ECSs is as follows:

- Step 1** Create ECSs.
- Step 2** View the ECS creation status.
- Step 3** Initialize the data disks attached to the ECS after the ECS is created.
- Step 4** For an ECS whose virtualization type is KVM, you can install the one-click password reset plugin to reset the password with one click when the password is lost or expires. Without this plugin, you need to manually reset the password of an ECS when the password is lost or expires.

----End

5.2 Applying for an ECS

This section describes how to configure the parameters for creating an ECS, including the ECS flavor, image, network, and authentication mode.


Context

ECSs are more cost-effective than physical servers. Within minutes, you can obtain flexible and on-demand ECS resources from the cloud platform.

Procedure

- Step 1** Log in to ManageOne as a VDC operator using a browser.
 - URL in non-B2B scenarios: **https://Address for accessing ManageOne Operation Portal**, for example, **https://console.demo.com**
URL in B2B scenarios: **https://Address for accessing ManageOne Tenant Portal**, for example, **https://tenant.demo.com**

- Username and password: username and password of the VDC operator.

Step 2 Click  in the upper left corner, select a region, and choose **Computing > Elastic Cloud Server**.

The **Elastic Cloud Server** page is displayed.

Step 3 Select a resource set from the menu bar on the top of the page.

Step 4 Click **Apply for ECS**.

The **Select Service** page is displayed.

Step 5 Select a service and click **Apply Now**.

The **Apply for ECS** page is displayed.

Step 6 Complete basic configurations for the ECS.

 **NOTE**

- Customizable settings vary depending on the product you select. The ECS you selected in **Step 5** determines whether **AZ**, **ECS Type**, **vCPUs**, **Memory**, **Image Type**, and **Image** can be customized. During the configuration, you can skip the parameters that cannot be customized.
- The screenshot is only an example. If the actual environment is different from the screenshot, use the actual environment.

Table 5-1 Parameter description

Parameter	Description	Example Value
AZ	Specifies a physical region where resources use independent power supplies and networks. AZs are physically isolated but interconnected through an internal network. To enhance application availability, create ECSs in different AZs.	kvm_az

Parameter	Description	Example Value
Creation Method	<p>Select a creation method for an ECS.</p> <ul style="list-style-type: none"> • New: Customize parameters to create an ECS. • Create from Template: Create an ECS using an ECS image or existing ECS backup as a template. To create an ECS using ECS backup, the CSBS service must be deployed on the platform. The ECS image and ECS backup displayed on the page are in the current region and support cross-AZ deployment, but do not support cross-region deployment. A full-ECS image and ECS backup that have expired or been deleted cannot be used to create an ECS. <p>When an ECS is requested, the OS and the boot mode of the OS cannot be changed. You can click Upgrade Flavor to modify the ECS type and flavor. The type and capacity of the system disk or data disk can be changed, but the capacity cannot be less than that of the source ECS disk. Other parameters can be customized as required.</p>	New
ECS Type	<p>The platform provides various ECSs for you to select based on application scenarios.</p> <p>The ECS type is determined by the ECS type tag selected during flavor creation. For details, visit Operation Help Center and choose Operation > Compute Services > Elastic Cloud Server (ECS) > Configuration Before ECS Creation > Creating a Flavor.</p>	General-purpose

Parameter	Description	Example Value
Boot Mode	<p>Specifies the ECS boot mode, which can be BIOS or UEFI.</p> <p>Basic Input/Output System (BIOS) is used to load the basic computer code to initialize hardware, check hardware functions, and boot the OS.</p> <p>Unified Extensible Firmware Interface (UEFI) does not need a long self-check as BIOS does, simplifying hardware initialization and OS boot. In addition, UEFI is easy to use because it supports graphical user interfaces (GUIs), various operation modes, and hardware driver insertion.</p> <p>NOTE</p> <p>This parameter is available only when the following requirements are met. Otherwise, this parameter is not available.</p> <ul style="list-style-type: none"> • The virtualization type of the selected AZ is KVM. • After you select vCPUs and memory (MB), the system filters image files based on the selected memory size. An image will be displayed here only if its Min Memory (MB) specified during image registration is smaller than the memory size of the selected flavor. <p>This parameter is available only if 1) at least one of the displayed image files is configured to use the UEFI boot mode; 2) UEFI boot is selected during image registration on Service OM. Otherwise, this parameter is unavailable, indicating that all image files use BIOS as the default boot mode.</p> <ul style="list-style-type: none"> • In Arm scenarios, the ECS boot mode can only be UEFI and cannot be changed. 	BIOS
Image Type	<ul style="list-style-type: none"> • 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. • Private Image Image available only to the user who created it using an existing ECS or external image file. It contains an OS, pre-installed public applications, and your private applications. Using a private image to create ECSs removes the need to configure multiple ECSs repeatedly. • Shared Image A shared image is a private image shared by another user. 	Public Image

Parameter	Description	Example Value
Image	<ul style="list-style-type: none"> ● Windows Used for development platforms or production workloads that run on Windows. It is recommended that the memory capacity be at least 1 GB. ECSs created using a Windows image support the installation of Internet Information Services (IIS) and SQL servers. ● Linux Used for development platforms or production workloads that run on Linux. <p>NOTE Select a 64-bit OS if the required memory capacity is 4 GB or larger. This is because 32-bit OSs allow addressing only within a 4 GB memory range.</p> <p>During ECS creation, the system filters image files based on the selected flavor.</p> <ul style="list-style-type: none"> ● For a flavor whose Boot Device is set to Cloud Disk, an image is displayed here only when its Min Memory (MB) specified during image registration is less than or equal to the selected memory. ● For a flavor whose Boot Device is set to Local Disk, an image is displayed here only when its Min Memory (MB) specified during image registration and the minimum disk required by the image are less than or equal to the Memory and Root Disk (GB) of the selected flavor, respectively. <p>If you select BIOS or UEFI as the boot mode, the image files that use this boot mode will be displayed. If the boot mode parameter is not displayed, it indicates that all the image files use BIOS as the default boot mode.</p> <p>NOTE If you select vGPU-accelerated for ECS Type and the driver is not installed in the image, install the GRID driver by referring to B Installing a GRID Driver on a vGPU-accelerated ECS after the ECS is created.</p>	CentOS

Parameter	Description	Example Value
Joint Windows Domain	<p>This parameter is available if the virtualization type of the selected AZ is KVM, the ECS is running a Windows OS, and the service selected in Step 5 has been configured with domain information. If the selected image is a static injection image or does not have Cloudbase-Init installed, it cannot be added to a domain.</p> <p>The administrator can perform unified authentication for ECSs added to the same domain. The following functions will be available for ECSs added to a domain: manage compute resources, reduce network management complexity and costs, enhance security, and support account roaming and folder redirection. Resources can be shared among ECSs in the same domain. For more information about domain servers and their functions, click here.</p> <p>Specify whether to add an ECS to a Windows domain. You can select a domain from the drop-down list. Available options are those defined by the administrator during product creation.</p>	-

Parameter	Description	Example Value
Same Storage	<p>This parameter is available only when Boot Device of the ECS flavor is set to Cloud Disk.</p> <ul style="list-style-type: none"> • If this parameter is set to Yes, the created ECS will support backup, DR, and ECS snapshot. • If this parameter is set to No, the created ECS will not support backup or DR. Whether it will support ECS snapshot depends on whether all disks of the ECS reside in the same storage backend. <p>If this parameter is set to Yes, the system selects the same storage backend configured with a storage tag to create the system and data disks of the ECS. The ECS can be provisioned successfully only when a storage backend that meets the requirements above is available in the environment. For details about how to configure the storage tag, visit Operation Help Center > Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > Disk FAQs > (Optional) Creating a Disk Type.</p> <p>After the ECS is provisioned, you can change the Same Storage setting if certain conditions are met. For details, see 7.4 Modifying the DR or Backup Function of an ECS.</p>	Yes
System Disk	<p>Select a disk type and set the disk size. To ensure that the ECS runs properly, the minimum allowed capacity of the system disk is related to the selected image file.</p> <p>If the system disk capacity cannot be changed on the web page, log in to DMK and change the value of is_supported_modify_sys_disk_size to true. For details, visit Operation Help Center and choose Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > Modifying Configuration Items on DMK.</p>	10GB

Parameter	Description	Example Value
Data Disk	<p>This parameter is displayed after you click Add Data Disk.</p> <p>Select a disk type and set the disk size. You can create multiple data disks for an ECS.</p> <p>NOTE</p> <ul style="list-style-type: none">• If Shared Disk is selected, the data disk can be attached to multiple ECSs. A shared disk can be attached to a maximum of 16 ECSs.• If you select Yes for Same Storage, ensure that the system disk and data disks of the ECS reside in the same storage backend, and the storage backend has the storage tag configured. Otherwise, the ECS cannot be provisioned.• If you select SCSI, transparent SCSI command transmission is supported. Therefore, when the VM HA function is used, lock protection is not supported, and the disk may be dual written. If SCSI is not displayed, log in to DMK and set the value of is_supported_volume_device_type in the ECS_UI configuration file to true. For details about how to change the value, visit Operation Help Center and choose Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > Modifying Configuration Items on DMK.	100GB
Quantity	<p>Set the number of ECSs to be created.</p> <p>NOTE</p> <p>By default, a maximum of 30 ECSs can be requested at a time.</p>	1

Step 7 Click **Next**. The **Network Configuration** page is displayed.

Step 8 Complete network configurations for the ECS.

Table 5-2 Parameter description

Parameter	Description	Example Value
Resource Set	<p>Select the current resource set or another resource set from the drop-down list. You can view the current resource set in the navigation bar at the top. Assume that the current resource set is Resource Set A and another resource set available is Resource Set B.</p> <ul style="list-style-type: none"> • When you select the current resource set, VPCs available will be those in Resource Set A. • If you select Resource Set B, VPCs available will be those in Resource Set B. By selecting Resource Set B, you create ECSs in Resource Set A by using the network resources of Resource Set B. With other configurations including security groups, you enable these ECSs to communicate with all those in the VPCs that belong to Resource Set B, allowing ECSs of different projects to share the same VPCs. <p>NOTE</p> <ul style="list-style-type: none"> • This parameter is available when VPC sharing is enabled on Service OM and the shared VPC permission is configured for the resource set on ManageOne. Otherwise, this parameter is not displayed. By default, this function is disabled. For details, visit Operation Help Center and choose Operation > Network Services > Virtual Private Cloud (VPC) > Shared VPC Best Practices. • If you select Resource Set B, security groups available will be those in Resource Set A, but the EIPs will be those that belong to Resource Set B. 	project_02
VPC	<p>Provides network functions for ECSs. The network functions include the subnet and security group.</p> <p>You can select an existing VPC, or click Create VPC to create one.</p>	-

Parameter	Description	Example Value
NIC	<p>Includes primary and extension NICs. You can add a maximum of 15 extension NICs to an ECS.</p> <ul style="list-style-type: none"> If you select VPC Subnet, all subnets in the VPC are available for you to choose from. In this case, the NIC supports layer 3 communication, allowing the ECS to communicate with networks (for example, the public network or other VPCs) beyond the VPC. If you select Intra-Project Subnet, all project-level subnets in the project are available for you to choose from. All NICs configured with the same subnet can communicate with each other at layer 2 on the project level. Layer 2 communication is supported within the same VPC and between different VPCs. <p>The Primary NIC network type must be VPC Subnet. Otherwise, you cannot access the Internet through the allocated EIP, and you cannot access the Object Storage Service (OBS) or a security service. You can select the network type for an Extension NIC as required.</p> <p>You can choose to use the automatically assigned IP address or manually assign one. If you choose to manually assign IP addresses when creating ECSs in batches, you can set the value of Specify IP Address By to IP Address Range or Single IP Address. Then, click OK.</p> <p>NOTE</p> <ul style="list-style-type: none"> If the selected subnet has only an IPv4 address segment, the NIC will only have one IPv4 address. If the selected subnet has both IPv4 and IPv6 address segments, the NIC will have one IPv4 address and one IPv6 address. If the selected subnet does not have DHCP enabled and the selected image does not support static IP address injection, after an ECS is created, you need to manually configure IP addresses for the ECS. Otherwise, the NIC cannot be reached. For details, see 18.6.1 Configuring a Static IP Address for an ECS. 	VPC Subnet subnet-c869(192.168.0.0/24)

Parameter	Description	Example Value
Security Group	<p>A security group implements access control for ECSs within the security group, enhancing security protection on ECSs. This enhances ECS security.</p> <p>When creating an ECS, you can select multiple security groups. Multiple security groups may affect the ECS network performance. You are advised to select a maximum of five security groups. In such a case, the access rules of all the selected security groups apply to the ECS.</p> <p>NOTE</p> <p>If the image has Cloud-Init or Cloudbase-Init installed, you need to initialize the ECS after it is created. Before initializing an ECS, ensure that the security group rule in the outbound direction meets the following requirements:</p> <ul style="list-style-type: none"> • Protocol: TCP • Port Range: 80 • Remote End: 169.254.0.0/16 <p>If you use the default security group rules in the outbound direction, the preceding requirements are already met, and this parameter does not need to be set.</p> <p>If you want to be able to log in to an ECS in SSH mode, you need to configure the inbound rules of the security group to allow your local computer to access the ECS. For details, see 12.4 Configuring Security Group Rules.</p>	-

Parameter	Description	Example Value
EIP	<p>A static public IP address bound to an ECS in a VPC. Using the EIP, the ECS provides services externally.</p> <p>The following options are provided:</p> <ul style="list-style-type: none"> ● Do Not Use Without an EIP, the ECS cannot access the Internet and is used only in the private network or cluster. ● Automatically Assign: The system automatically assigns an EIP for the ECS. In the Region Type I scenario, you also need to select External Network, Subnet, Bandwidth Type, and Bandwidth to which the EIP belongs. <ul style="list-style-type: none"> – If Bandwidth Type is set to Dedicated Bandwidth, you need to configure the bandwidth of the EIP. The EIP occupies bandwidth exclusively. Select this mode when you desire stable and large bandwidth. – If Bandwidth Type is set to Shared Bandwidth, you need to set Bandwidth Name for the shared bandwidth. The EIP shares the bandwidth with other EIPs that are added to the shared bandwidth. If the shared bandwidth contains three EIPs and the peak bandwidth is 10 MB/s, the total traffic of the three EIPs cannot exceed 10 MB/s. Select this mode when your application does not have a high bandwidth requirement and there is a bandwidth cap. <p>NOTE If Automatically Assign is selected, ensure that the EIP quota is sufficient. Otherwise, ECS provisioning will fail.</p> <ul style="list-style-type: none"> ● Specify: An existing EIP is assigned for the ECS. When using an existing EIP, you cannot create ECSs in batches. <p>NOTE If you are not authorized to used EIP, contact the administrator to change your permissions.</p>	Do Not Use

Step 9 Click **Next**. The **Advanced Settings** page is displayed.

1. Set the ECS name.

When you create ECSs in batches, the system automatically adds an incremental number to the end of the custom ECS name, for example,

ecs-0001, ecs-0002, and so on. The value ranges from 0001 to 9999 by default.

To start with a specific number, click **Change Suffix Start Number** to customize the value. For example, if you set the value to 1126, the ECS names will be xxx-1126, xxx-1127, and so on.

 **NOTE**

If you want to create a Windows ECS that needs to be added to a domain, or if you require that the host name of the ECS (that is, the computer name shown in the ECS OS) must be unique, set the ECS name by following the instructions provided in **Operation Help Center > Operation > Compute Services > Elastic Cloud Server (ECS) > ECS Host Name > Rules for Configuring ECS Names (Unique Host Names)**.

2. Set the host name prefix of the ECS.

The host name prefix of the ECS and a suffix of 5 random characters (0-9 and a-z) form the ECS host name, that is, the computer name shown in the OS. The value is in the format "Host Name Prefix-5 random characters".

- This parameter needs to be configured if this parameter is to be displayed. The system generates the host name prefix according to the ECS name you configured in [Step 9.1](#) and by automatically filtering out unallowed characters based on the naming rules displayed on the page. You can change this prefix. The generated ECS host name is unique in the region or resource set. For details, see visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > ECS Host Name > Enabling ECS Host Name Uniqueness**.
- Skip this parameter if it is not displayed. The system generates host names for ECSs based on the default rules. The host names of different ECSs may be the same. For details, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > ECS Host Name > Default Rules for Generating ECS Host Names**.

3. Set the running status of the ECS. This parameter is available if the virtualization type of the AZ where the ECS resides is KVM.

- **Stopped:** A newly obtained ECS stays in the **Stopped** state.

 **NOTE**

- If **ECS Initial Status** is set to **Stopped** and the host group where the ECS resides is configured with tags for releasing resources upon ECS shutdown, the ECS does not occupy any of the following resources: vCPU, memory, GPU, NPU, USB, and volume connections. For details, see step "Configure custom tags" in **Operation Help Center > Operation > Compute Services > Elastic Cloud Server (ECS) > Configuration Before ECS Creation > Creating a Host Group**.
- When Arm servers are used, NIC connections can be released for general computing-plus ECSs when they are shut down.
- Resources will not be released for disk-intensive or ultra-high I/O ECSs when they are shut down.
- **Running:** A newly obtained ECS stays in the **Running** state.

4. Set a key or new password. Select the image password as the ECS password or reset a new password or key for the ECS.

 **NOTE**

This parameter is displayed only when the following conditions are met:

- The image selected in [Step 6](#) must have the Cloud-Init (Linux OS) plugin installed, and Cloud-Init is selected during image registration.
- The image selected in [Step 6](#) can use the image password as the ECS password.
- The image password cannot be used as the ECS password in the Windows OS.
- **No:** Use the password set during image creation as the ECS password.
- **Yes:** Set a key or new password for the ECS.

5. Select an authentication mode for the ECS.

 **NOTE**

The image selected in [Step 6](#) must have the Cloud-Init (Linux OS) or Cloudbase-Init (Windows OS) plugin installed, and Cloud-Init is selected during image registration.

- Key pair

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

 **NOTE**

- If you use an existing key pair, make sure that you have saved the key file to a local directory. Otherwise, logging in to the ECS will fail.
- Windows ECSs support only the password authentication mode. If the login mode is set to **Key pair**, you must use the key file used during ECS creation to obtain the password of user **Administrator** or Cloudbase-Init account generated during ECS installation for subsequent logins. For details, see [6.3.1 Obtaining the Password for Logging In to a Windows ECS](#).
- For Windows, if the selected image supports static IP address injection, the key pair authentication is dimmed. Only the password mode can be selected.

- Password

A password is used for ECS login authentication. If the ECS runs a Linux OS, you can use username **root** and its password to log in to the ECS. If the ECS runs a Windows OS, you can use username **Administrator** and its initial password to log in to the ECS.

When using password login authentication on an ECS whose virtualization type is KVM, you can select **Customize user** and customize a username and password to create a user.

- On a Windows ECS, this is a common user without administrator rights. To perform operations that require administrator rights, switch to the administrator role first.
 - On a Linux ECS, this is a common user without administrator rights. You can use this user to log in to the ECS over SSH. If you need to run a script or system command after login, enter the password of the **root** user to upgrade rights. After this user is created, logging in to the ECS as the **root** user over SSH is disabled by default. To enable the function, see [Enabling Login Using a Password over SSH](#).
6. If you need to configure advanced settings for the ECS, select **Configure**. Otherwise, go to [Step 10](#).

- **Watchdog:** Enable or disable watchdog for an ECS.
The watchdog function provides a heartbeat mechanism used to monitor the health status of ECSs. When an ECS does not work properly, an alarm is generated, and the system attempts to restart the ECS. If restarting the ECS is successful, the alarm is cleared.

NOTICE

- Before enabling the watchdog function in x86 scenarios, ensure that the watchdog program that complies with the standard IPMI watchdog has been installed on the image selected in [Step 6](#). Otherwise, the ECS may restart repeatedly.
 - Before enabling the watchdog function in Arm scenarios, ensure that the watchdog program that complies with the standard 6300ESB watchdog has been installed on the image selected in [Step 6](#). Otherwise, the ECS may restart repeatedly.
-
- **Watchdog Alarm Policy:** In Arm scenarios, this parameter is available if **Watchdog** is set to **Enable**. If the 6300ESB watchdog does not detect watchdog information in the specified time, an alarm will be generated. The ECS will determine, based on this alarm policy, whether to get restarted.
 - **HA:** Enable or disable the HA function for an ECS.
To support HA, ECSs must meet the following requirements: the global HA function is enabled, the HA function of the host group where the ECS resides is enabled or not configured, and the HA function of the ECS is enabled. When HA is enabled, an ECS is automatically rebuilt on another host whenever the ECS or its host becomes faulty, ensuring service continuity.

NOTE

- For details about how to enable the global HA function, see "Product Management" > "Resource Pool" > "FusionSphere OpenStack" > "Compute" > "Configuring the VM HA Function" in *HUAWEI CLOUD Stack 8.1.0 O&M Guide*.
 - To check whether the HA function of the host group is enabled, log in to Service OM and check it in the **Custom Tag** area on the **Configuration** tab page of the host group details page. If a custom tag whose tag name is **_ha_enabled** and tag value is **False** exists, the HA function of the host group is disabled. If the tag does not exist or its value is **True**, the HA function of the host group is enabled.
You are advised not to enable this function for the management host group. Otherwise, services may be affected.
 - Resources need to be reserved for HA. Otherwise, the ECS HA function may fail. To ensure that the ECS HA functions properly, you need to clear alarms such as host exceptions and insufficient resources in a timely manner.
- **CD-ROM Drive**
For ECSs whose virtualization type is KVM, you can select **Use** or **Not use** for **CD-ROM Drive**.

- Select **Use** if you want to attach the local CD-ROM drive or an ISO file to the ECS.
- If the CD-ROM drive is used and UVP VMTools is installed in the image you have selected, UVP VMTools will be automatically upgraded after the ECS is provisioned. If the CD-ROM drive is not used, even if UVP VMTools is installed in the image you have selected, UVP VMTools cannot be automatically upgraded after ECS provisioning. You will need to manually upgrade UVP VMTools.

NOTE

UVP VMTools collects internal monitoring metrics of ECSs to monitor their running status and supports communication between ECSs and physical hosts. UVP VMTools also provides the following functions:

- Improves disk I/O performance and network I/O performance for Windows ECSs.
- Reports alarms when faults occur on Linux ECSs.

– ECS Group

An ECS group is a logical group with affinity, anti-affinity, weak affinity, or weak anti-affinity rules configured. ECSs added to an ECS group will be scheduled to the same or different hosts according to the affinity rules of the group. For details about how to create an ECS group, see [11.1 Creating an ECS Group](#).

NOTE

- If the policy of the ECS group is affinity or anti-affinity, the ECS will fail to be created when existing hosts or resources are insufficient to fulfill the affinity or anti-affinity rules of the ECS group.
- ECSs whose virtualization type is KVM can be added to an ECS group configured with any affinity rule. ECSs of other virtualization types cannot be added to ECS groups.

– Tag

Specifies the ECS tags. This parameter is optional and helps you identify and manage your ECSs.

Click **Add Tag**, and select an existing key and value from the drop-down list box. The tags come from the Tag Management Service. To add or modify a tag, ask the administrator to do it after choosing **Console > Mgmt & Deployment > Tag Management** from the top menu bar.

NOTE

- When multiple tags are added to an ECS, each tag must have a unique key.
- During ECS creation, the system also automatically generates a built-in tag. The tag key is identical to the VPC ID and is invisible on the UI.

For details about tag management, see [7.1.3 Adding and Managing ECS Tags](#).

– File Injection


This parameter is optional. Enables the ECS to automatically inject a script file or other files into a specified directory when you create the ECS.

For details about file injection, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > Configuration Before ECS Creation > Creating a File Injection Script**.

– User Data Injection

Enables the ECS to automatically inject user data when the ECS starts for the first time. This configuration is optional. After this function is enabled, the ECS automatically injects the user data upon its first startup. For details about user data injection, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > Configuration Before ECS Creation > Creating a User Data Injection Script**.

Step 10 Click **Next** to go to the **Confirmation** page.

1. Check whether all settings are correct. If you need to modify a configuration item, click the  icon next it.
2. Confirm **Required Duration**.

 **NOTE**

Specifies the required duration for an ECS. It begins from the time when the ECS was created. You can use it within the use duration. When this duration expires, the ECS status becomes **Expired**.

If the ECS is running properly before the expiration, the ECS will still run properly and the system will not be shut down. In this case, you can only **Extend** and **Delete** the ECS.

Step 11 Click **Add to Cart** or **Apply Now**.

- **Add to Cart:** Add the configured ECS to the shopping cart, and submit the order after you confirm all the resources you need, including network and storage resources.
- **Apply Now:** Submit the task.

 **NOTE**

- If the ECS you requested needs administrator approval, it will be provisioned after your request is approved. Otherwise, the ECS will be provisioned immediately.
- If you create an ECS with additional data disks, initialize the data disks after the ECS is created. For details about how to initialize the data disks, see [10.3 Initializing a Data Disk](#).
- If your ECS is assigned both IPv4 and IPv6 addresses and runs CentOS 7.5 or Ubuntu Server 18.04.1, or the network communication is abnormal after the application is successful, visit **Operation Help Center** and choose **Operation > Compute Services > Image Management Service (IMS) > FAQs > Configuring a VM to Dynamically Obtain IPv6 Addresses**.

----End

5.3 Viewing ECS Creation Status

After obtaining an ECS, you can view the creation status. The task involves several subtasks, such as creating an ECS, binding an EIP, and attaching an EVS disk.

Context

The task status can be **Processing** or **Failed**:

- **Processing**: The system is processing the task.
- **Failed**: The system failed to process the task. For a failed task, the system automatically rolls back the task and displays an easy-to-understand error code, for example, **Ecs.0013Insufficient EIP quota**. For details, see [18.1.2 How Do I Handle Error Messages Displayed on ManageOne?](#)

This section describes how to view ECS creation progress and the information displayed in the **Task Status** area.

Procedure

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 2** **Task Status** is displayed next to **Operation** above the ECS list. After you perform the task for creating an ECS, the status of the task is displayed in the **Task Status** column.
- Step 3** Click the number displayed in the **Task Status** area and view details about the processing tasks and failed tasks for ECS creation statuses.
 - Locate the row that contains the ECS in the **Processing** state, and click **Cancel** in the **Operation** column to cancel the subtask.
 - Click **Reapply** in the **Operation** column of the ECS in the **Failed** state to reapply for all subtasks that fail to be executed or have been canceled. You can also use the name of the ECS that failed to be applied for.

NOTE

If the **Task Status** area shows an ECS creation failure but the ECS list displays the created ECS, see section [18.3.10 Why Does the Task Status Area Show an ECS Creation Failure But the ECS List Displays the Created ECS?](#)

----End

5.4 Installing the One-Click Password Reset Plugin

5.4.1 Overview

You can use the one-click password reset plugin to conveniently reset the password of a KVM ECS on the ECS console. Without this plugin, you need to manually reset the password of an ECS when the password is lost or expires. For details about how to manually reset the password of an ECS, see [8.4 Manually Resetting the Password for Logging In to a Windows ECS](#) and [8.5 Manually Resetting the Password for Logging In to a Linux ECS](#).

It is up to you to determine whether to install the one-click password reset plugin.

 NOTE

- If the one-click password reset plugin has been installed in the image used for creating the ECS, you do not need to reinstall this plugin. For details about how to verify whether this plugin has been installed on a Windows ECS, go to [Step 1](#). For details about how to verify whether this plugin has been installed on a Linux ECS, go to [Step 1](#).
- If the plugin cannot automatically start upon ECS startup, rectify the fault by referring to [18.8.5 What Should I Do If the One-Click Password Reset Plugin Fails to Start?](#).

5.4.2 Installing the One-Click Password Reset Plugin for a Windows ECS

Before You Start

- After the one-click password reset plug-in is installed on the ECS, you need to restart the ECS to use the one-click password reset function.
- If you perform operations that have impact on the original system disk, including reinstalling an ECS OS, changing an ECS OS, or replacing the system disk, you must reinstall the one-click password reset plugin.
- It is up to you to determine whether to install this plugin.
- When the password of an ECS is lost or expires, you can only manually reset the password. In this case, you cannot install and use this plugin to reset the password. For details about how to reset the password, see [8.4 Manually Resetting the Password for Logging In to a Windows ECS](#).
- This plugin can be automatically updated only when the ECS is bound with an EIP. Alternatively, you can manually download the upgrade package to install or update the plugin.
- If the one-click password reset plugin has been installed in the image used for creating the ECS, you do not need to reinstall this plugin. For details about how to verify whether this plugin has been installed, go to [Step 1](#).
- [Table 5-3](#) lists the Windows OSs that support this plugin. You cannot use this plugin in any Windows OS that is not listed here.

Table 5-3 Windows OSs that support the one-click password reset plugin

OS	OS Version
Windows	Windows Server 2008 R2 Enterprise 64-bit (English) Windows Server 2008 R2 Enterprise 64-bit (Chinese) Windows Server 2008 R2 Standard 64-bit (English) Windows Server 2008 SP2 Enterprise 64-bit Windows Server 2008 R2 Datacenter 64-bit Windows Server 2008 R2 Standard 64-bit Windows Server 2012 R2 Standard 64-bit (English) Windows Server 2012 R2 Standard 64-bit (Chinese) Windows Server 2012 R2 Datacenter 64-bit (English) Windows Server 2012 R2 Datacenter 64-bit (Chinese) Windows Web Server 2008 R2 64-bit Windows 2008 Enterprise R2 64-bit (English) Windows 2012 R2 Standard Windows 2012 R2 Datacenter Windows 2012 R2 Datacenter (English) Windows 2012 R2 Standard (English) Windows 2016 Datacenter 64-bit (English) Windows 2016 Datacenter 64-bit (Chinese) Windows Server 2016 for Hygon 64-bit

Prerequisites

- The ECS must have more than 300 MB remaining space, and data can be written to its C drive.
- The ECS network connectivity is normal and DHCP is enabled.
- The outbound rules of the security group to which the ECS belongs must meet the following configuration requirements:
 - **Protocol:** TCP
 - **Port Range:** 80
 - **Remote End:** 169.254.0.0/16

If the default security group rule is used, you do not need to modify any configuration because the default security group rule meets all the preceding configuration requirements. In the default security group rule, **Protocol**, **Port Range**, and **Remote End** are set to **ANY**, **ANY**, and **0.0.0.0/16**, respectively.

Procedure

Step 1 Check whether the one-click password reset plugin has been installed on the ECS.

1. Log in to the ECS. For details, see [6.3.2 Logging In to a Windows ECS Using VNC \(Through the Console\)](#).

2. Press **Ctrl+Shift+Esc**. In the task manager, check whether there are **cloudResetPwdAgent** and **cloudResetPwdUpdateAgent**.
 - If there are **cloudResetPwdAgent** and **cloudResetPwdUpdateAgent**, as shown in [Figure 5-1](#), the one-click password reset plugin has been installed on the ECS. In this case, you do not need to install this plugin again.

Figure 5-1 One-click password reset plugin in the task manager

cloudResetPwdAgent		cloud reset password agent
cloudResetPwdUpdateAgent	13872	cloud reset password update agent

- If there is neither **cloudResetPwdAgent** nor **cloudResetPwdUpdateAgent**, the one-click password reset plugin has not been installed. In this case, perform [Step 2](#) to install the plugin.

Step 2 Install the one-click password reset plugin.

1. Download the plugin to the ECS.
 - a. Use the browser on the ECS to log in at <http://support.huawei.com/enterprise> and choose **Technical Support > Product Support > Enterprise Data Center > Cloud Computing > Huawei Cloud Stack > FusionSphere SIA**.
 - b. Download **FusionSphere_SIA-x.x.x-CloudResetPwdAgent_for_FusionSphere_OpenStack.zip** and decompress it to the ECS.

 **NOTE**

If the plug-in cannot be downloaded to the ECS, download the plug-in to the local PC and then upload it to the ECS.

1. Click [here](#) and select the latest version to go to the details page.
 2. Download **FusionSphere_SIA-x.x.x-CloudResetPwdAgent_for_FusionSphere_OpenStack.zip** and upload it to the ECS.
2. Go to the folder of the Windows OS, decompress **CloudResetPwdAgent.zip**, and go to the **CloudResetPwdAgent** folder.
3. Double-click **CloudResetPwdAgent.Windows > setup.bat** to install the one-click password reset plugin.
4. Double-click **CloudResetPwdUpdateAgent.Windows > setup.bat** to install the update program for the one-click password reset plugin.
5. Press **Ctrl+Shift+Esc**. In the task manager, check whether there are **cloudResetPwdAgent** and **cloudResetPwdUpdateAgent**.
 - a. If yes, the one-click password reset plugin is successfully installed.

Figure 5-2 One-click password reset plugin in the task manager

cloudResetPwdAgent		cloud reset password agent
cloudResetPwdUpdateAgent	13872	cloud reset password update agent

- b. If no, the installation failed. Rectify the fault by performing steps provided in [18.8.5 What Should I Do If the One-Click Password Reset Plugin Fails to Start?](#).

Step 3 Check whether the password reset plugin automatically starts upon ECS startup.

1. Restart the ECS.
2. Press **Ctrl+Shift+Esc**. In the task manager, check whether there are **cloudResetPwdAgent** and **cloudResetPwdUpdateAgent**.
 - a. If yes, the installation is successful. Go to [Step 3.3](#).
 - b. If no, the installation failed. Rectify the fault by performing steps provided in [18.8.5 What Should I Do If the One-Click Password Reset Plugin Fails to Start?](#).
3. Run the following command to check whether the IP address can be pinged:
ping 169.254.169.254

 **NOTE**

- 169.254.169.254 is a local IP address and is used to query the metadata of the ECS.
 - If yes, the ECS NIC is working properly. Go to [Step 3.4](#).
 - If no, contact technical support.
4. Stop the ECS, [log in to the ECS consol](#), and reset the password.

----End

Related Operations

Uninstalling the one-click password reset plugin

Step 1 Go to the **C:\CloudResetPwdUpdateAgent\bin** directory.

Step 2 Double-click **UninstallApp-NT.bat** to uninstall the plugin.

Step 3 Delete all files from **C:\CloudResetPwdUpdateAgent**.

Step 4 Go to the **C:\CloudResetPwdAgent\bin** directory.

Step 5 Double-click **UninstallApp-NT.bat**.

Step 6 Delete all files from **C:\CloudResetPwdAgent**.

----End

5.4.3 Installing the One-Click Password Reset Plugin for a Linux ECS

Before You Start

- After the one-click password reset plug-in is installed on the ECS, you need to restart the ECS to use the one-click password reset function.
- If you perform operations that have impact on the original system disk, including reinstalling an ECS OS, changing an ECS OS, or replacing the system disk, you must reinstall the one-click password reset plugin.
- The one-click password reset plugin **CloudResetPwdAgent** has been added as an open-source program to [GitHub Open-Source Platform](#) based on GNU General Public License v2.0. It is up to you to determine whether to install this plugin.

- When the password of an ECS is lost or expires, you can only manually reset the password. In this case, you cannot install and use this plugin to reset the password. For details about how to reset the password, see [8.5 Manually Resetting the Password for Logging In to a Linux ECS](#).
- This plugin can be automatically updated only when the ECS is bound with an EIP. Alternatively, you can manually download the upgrade package to install or update the plugin.
- If the one-click password reset plugin has been installed in the image used for creating the ECS, you do not need to reinstall this plugin. For details about how to verify whether this plugin has been installed, go to [Step 1](#).
- If the plugin cannot automatically start upon ECS startup, rectify the fault by referring to [18.8.5 What Should I Do If the One-Click Password Reset Plugin Fails to Start?](#).
- [Table 5-4](#) lists the Linux OSs that support this plugin. You cannot use this plugin in any Linux OS that is not listed here.

Table 5-4 Linux OSs that support the one-click password reset plugin

OS	OS Version Supported by the x86 Server	OS Version Supported by the ARM Server
CentOS	<ul style="list-style-type: none">• CentOS 8.0 64-bit• CentOS 7.7 64-bit• CentOS 7.3 64-bit• CentOS 7.2 64-bit• CentOS 7.0 64-bit• CentOS 7.1 64-bit• CentOS 6.9 64-bit• CentOS 6.8 64-bit• CentOS 6.8 32-bit• CentOS 6.6 32-bit• CentOS 6.6 64-bit• CentOS 6.5 64-bit• CentOS 6.4 64-bit• CentOS 6.3 64-bit• CentOS 7.6 for Hygon 64-bit	<ul style="list-style-type: none">• CentOS 7.4 64-bit• CentOS 7.5 64-bit• CentOS 7.6 64-bit• CentOS 8.0 64-bit

OS	OS Version Supported by the x86 Server	OS Version Supported by the ARM Server
Debian	<ul style="list-style-type: none"> • Debian 10.1 64-bit • Debian 10.0 64-bit • Debian 9.11 64-bit • Debian 9.9 64-bit • Debian 9.0 64-bit • Debian 8.8 64-bit • Debian 8.2 64-bit • Debian 7.5 64-bit • Debian 7.5 32-bit 	Debian GNU/Linux 10.2 64-bit
openSUSE	<ul style="list-style-type: none"> • openSUSE 42.2 64-bit • openSUSE 13.2 64-bit • openSUSE Leap 42.2 64-bit • openSUSE Leap 42.1 64-bit • openSUSE Leap 15.1 64-bit 	openSUSE Leap 15.0 64-bit
SUSE	<ul style="list-style-type: none"> • SUSE Linux Enterprise Server 15 SP1 64-bit • SUSE Linux Enterprise Server 12 SP2 64-bit • SUSE Linux Enterprise Server 12 SP1 64-bit • SUSE Linux Enterprise Server 11 SP4 64-bit • SUSE Linux Enterprise Server 12 SP5 64-bit 	<ul style="list-style-type: none"> • SUSE Linux Enterprise Server 12 SP4 64-bit • SUSE Linux Enterprise Server 15 64-bit
Ubuntu	<ul style="list-style-type: none"> • Ubuntu 19.04 64-bit • Ubuntu 18.04.3 64-bit • Ubuntu 16.10 32-bit • Ubuntu 16.04 32-bit • Ubuntu Server 19.10 64-bit • Ubuntu Server 16.04 64-bit • Ubuntu Server 14.04 64-bit • Ubuntu Server 14.04 32-bit 	Ubuntu Server 18.04 64-bit Ubuntu Server 19.04 64-bit

OS	OS Version Supported by the x86 Server	OS Version Supported by the ARM Server
EulerOS	<ul style="list-style-type: none">• EulerOS 2.5 64-bit• EulerOS 2.2 64-bit• EulerOS 2.9 64-bit	<ul style="list-style-type: none">• EulerOS 2.8 64-bit• EulerOS 2.9 64-bit
Fedora	<ul style="list-style-type: none">• Fedora 24 64-bit• Fedora 25 64-bit• Fedora Server 29 64-bit• Fedora Server 30 64-bit• Fedora Server 31 64-bit	Fedora Server 29 64-bit
Oracle	<ul style="list-style-type: none">• Oracle Linux 7.3 64-bit• Oracle Linux 6.9 64-bit• Oracle Linux 6.5 64-bit	N/A
China Standard Software	Neokylin Linux Advanced Server release 7.0 for Hygon (64 bit)	<ul style="list-style-type: none">• NeoKylin Server release 5.0 U2 64-bit• NeoKylin Linux Advanced Server release 7.0 U5 64-bit• NeoKylin Linux Advanced Server release 7.0 U6 64-bit• NeoKylin Linux Desktop 7.0 U5 64-bit

OS	OS Version Supported by the x86 Server	OS Version Supported by the ARM Server
Others	<ul style="list-style-type: none"> • uos V20 desktop • Kylin Desktop V10 • Kylin Server V10 	<ul style="list-style-type: none"> • Red Hat Enterprise Linux 7.5 64-bit • Red Hat Enterprise Linux 7.6 64-bit • Deepin GNU/Linux 15.5 64-bit • Red Flag Asianux Server 7.5.1804 desktop/server version • iSoft Server OS for Kunpeng • iSoft Server OS V5.1 • TongyuanOS 7.7 • UniKylin Linux release 3.3 • Kylin 4.0.2 • uos 20 desktop • uos V20 server • Kylin Linux Advanced Server V10 (Kylin V10 SP1 Server)

Prerequisites

- A Linux ECS must have more than 300 MB remaining space and data can be written to its root directory.
- ECSs created using SUSE 11 SP4 must have 4 GB or a larger memory.
- The ECS network connectivity is normal and DHCP is enabled.
- The outbound rules of the security group to which the ECS belongs must meet the following configuration requirements:
 - **Protocol:** TCP
 - **Port Range:** 80
 - **Remote End:** 169.254.0.0/16

If the default security group rule is used, you do not need to modify any configuration because the default security group rule meets all the preceding configuration requirements. In the default security group rule, **Protocol**, **Port Range**, and **Remote End** are set to **ANY**, **ANY**, and **0.0.0.0/16**, respectively.

 NOTE

When ARM servers are used, if your OS is Debian GNU, Linux 10.2 64bit, Unified OS 20 desktop 64bit, or Kylin 4.0.2 SP2 64bit, perform the following operations before installing the one-click password reset plugin:

1. Log in to the ECS and run the following command as the **root** user to create a lib64 directory:
mkdir -p /lib64
2. Run the following command to copy the file to the lib64 directory:
cp /lib/ld-linux-aarch64.so.1 /lib64

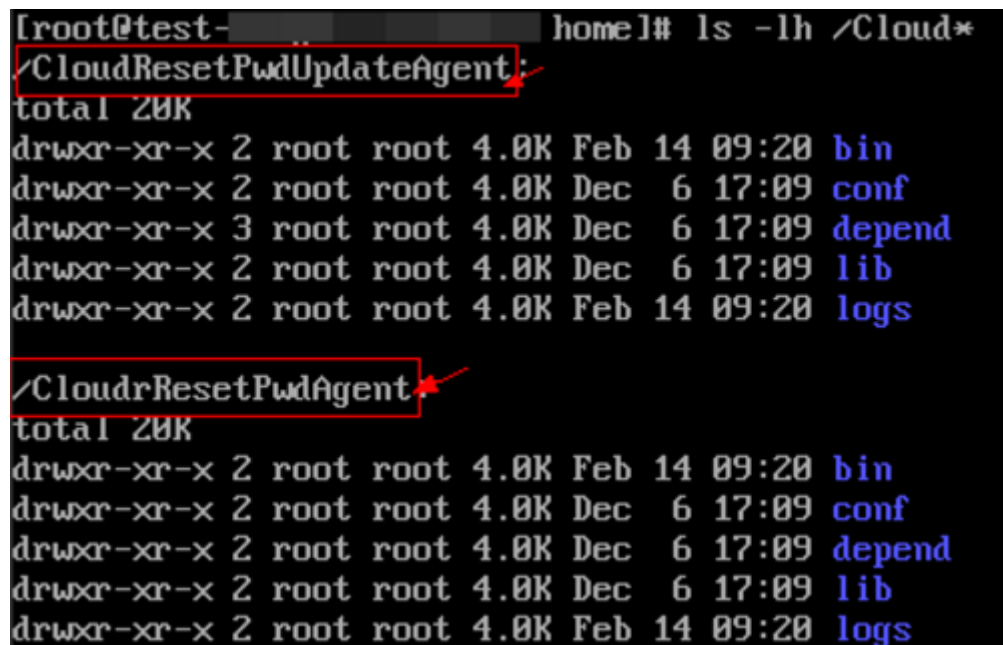
Procedure

Step 1 Check whether the one-click password reset plugin has been installed on the ECS.

1. Log in to the ECS as the **root** user. For details, see [6.1 Login Mode Overview](#).
2. Run the following command to check whether **CloudResetPwdAgent** and **CloudResetPwdUpdateAgent** have been installed on the ECS:

ls -lh /Cloud*

Figure 5-3 Command output



```
[root@test-... home]# ls -lh /Cloud*
/CloudResetPwdUpdateAgent:
total 20K
drwxr-xr-x 2 root root 4.0K Feb 14 09:20 bin
drwxr-xr-x 2 root root 4.0K Dec 6 17:09 conf
drwxr-xr-x 3 root root 4.0K Dec 6 17:09 depend
drwxr-xr-x 2 root root 4.0K Dec 6 17:09 lib
drwxr-xr-x 2 root root 4.0K Feb 14 09:20 logs

/CloudResetPwdAgent:
total 20K
drwxr-xr-x 2 root root 4.0K Feb 14 09:20 bin
drwxr-xr-x 2 root root 4.0K Dec 6 17:09 conf
drwxr-xr-x 2 root root 4.0K Dec 6 17:09 depend
drwxr-xr-x 2 root root 4.0K Dec 6 17:09 lib
drwxr-xr-x 2 root root 4.0K Feb 14 09:20 logs
```

3. If the information in [Figure 5-3](#) is displayed, the one-click password reset plugin has been installed on the ECS. In this case, you do not need to install this plugin again. Otherwise, the one-click password reset plugin has not been installed on the ECS. In this case, perform [Step 2](#) to install the plugin.

Step 2 Install the one-click password reset plugin.

1. Log in to the ECS. For details, see [6.1 Login Mode Overview](#).
2. Download the plugin to the ECS.
 - a. Use the browser on the ECS to log in at <http://support.huawei.com/enterprise> and choose **Support > Cloud Computing > Huawei Cloud Stack > FusionSphere SIA**.

- b. Download **FusionSphere_SIA-x.x.x-CloudResetPwdAgent_for_FusionSphere_OpenStack.zip** and decompress it to the ECS.

 NOTE

If the plug-in cannot be downloaded to the ECS, download the plug-in to the local PC and then upload it to the ECS.

1. Click [here](#) and select the latest version to go to the details page.
2. Download **FusionSphere_SIA-x.x.x-CloudResetPwdAgent_for_FusionSphere_OpenStack.zip** and upload it to the ECS.
3. Based on the ECS type (ARM or x86) and the number of bits of the ECS OS, go to the corresponding Linux folder, decompress **CloudResetPwdAgent.zip**, and go to the **CloudResetPwdAgent** folder.
4. Double-click **CloudResetPwdAgent.Linux > setup.sh** to install the one-click password reset plugin.
5. Double-click **CloudResetPwdUpdateAgent.Linux > setup.sh** to install the update program for the one-click password reset plugin.
6. Run the following command to decompress **CloudResetPwdAgent.zip**: There is no special requirement for the directory that stores the decompressed **CloudResetPwdAgent.zip**. Customize the directory.
unzip -o -d *Decompressed directory* CloudResetPwdAgent.zip
For example, if the directory for decompressing the package is **/home/linux/test**, run the following command:
unzip -o -d /home/linux/test CloudResetPwdAgent.zip
7. Run the following command to open the **CloudResetPwdUpdateAgent.Linux** file:
cd /home/linux/test
cd CloudResetPwdAgent/CloudResetPwdUpdateAgent.Linux
8. Run the following command to add the execute permission for the **setup.sh** file:
chmod +x setup.sh
9. Run the following command to install the plugin:
./setup.sh
10. Run the following commands to check whether the installation is successful:
service cloudResetPwdAgent status
service cloudResetPwdUpdateAgent status
If neither of the statuses of **CloudResetPwdAgent** and **CloudResetPwdUpdateAgent** is **unrecognized service**, the installation is successful, and go to [Step 3](#). If the installation fails, rectify the fault by performing steps provided in [18.8.5 What Should I Do If the One-Click Password Reset Plugin Fails to Start?](#).

- Step 3** Check whether the password reset plugin automatically starts upon ECS startup.
1. Restart the ECS.
 2. Run the following commands and check whether the statuses of **CloudResetPwdAgent** and **CloudResetPwdUpdateAgent** are **unrecognized service**.

service cloudResetPwdAgent status

service cloudResetPwdUpdateAgent status

- If yes, the plugin cannot automatically start upon ECS startup. Rectify the fault by performing steps provided in [18.8.5 What Should I Do If the One-Click Password Reset Plugin Fails to Start?](#).
 - If no, the plugin automatically starts upon ECS startup. Go to [Step 3.3](#).
3. Run the following command to check whether the IP address can be pinged:
ping 169.254.169.254

 **NOTE**

- 169.254.169.254 is a local IP address and is used to query the metadata of the ECS.
- If yes, the ECS NIC is working properly. Go to [Step 3.4](#).
 - If no, contact technical support.
4. Stop the ECS, [log in to the ECS console](#), and reset the password.

----End

Related Operations

Uninstalling the one-click password reset plugin

Step 1 Log in to an ECS. For details, see [6.1 Login Mode Overview](#).

Step 2 Run the following commands to open the bin file and delete **cloudResetPwdAgent**:

```
cd /CloudrResetPwdAgent/bin
```

```
sudo ./cloudResetPwdAgent.script remove
```

Step 3 Run the following commands to open the bin file and delete **cloudResetPwdUpdateAgent**:

```
cd /CloudResetPwdUpdateAgent/bin
```

```
sudo ./cloudResetPwdUpdateAgent.script remove
```

Step 4 Run the following commands to delete the password reset plugin:

```
sudo rm -rf /CloudrResetPwdAgent
```

```
sudo rm -rf /CloudResetPwdUpdateAgent
```

----End

5.5 Initializing EVS Data Disks

If you create an ECS with additional data disks, initialize the data disks after the ECS is created. For details, see [10.3 Initializing a Data Disk](#).

6 Logging In to an ECS

6.1 Login Mode Overview

You must log in to an ECS to configure and manage it. You can log in to only running ECSs.

You can log in to an ECS using VNC, SSH, or MSTSC.

- Logging in to an ECS using VNC
You can log in to both Windows and Linux ECSs using VNC by entering your username and password on UI.
- Logging in to an ECS using SSH
This method applies only to Linux ECSs that support the SSH protocol and have been connected. You can use a remote login tool (such as PuTTY) to log in to the ECS using a password or key.
- Logging in to an ECS using MSTSC
This method applies only to Windows ECSs bound with EIPs. You can run the **mstsc** command on a local server to log in to an ECS.

Select the login mode according to the ECS OS and login mode you selected when creating the ECS. For details, see [Table 6-1](#).

Table 6-1 Login mode overview

OS	Login Mode Selected When Creating an ECS	Available Login Method
Linux (Key pair or password authentication is supported when you log in to a	Key pair	For details about how to log in to a Linux ECS using an SSH key pair, see 6.2.1 Remotely Logging In to a Linux ECS Using a Key Pair (SSH) .

OS	Login Mode Selected When Creating an ECS	Available Login Method
Linux ECS. Available login options include SSH key pair, SSH password, and VNC.)	Password	<ul style="list-style-type: none"> For details about how to log in to a Linux ECS using SSH password, see 6.2.2 Remotely Logging In to a Linux ECS Using a Password (SSH). For details about how to remotely log in to a Linux ECS using VNC on the Web UI, see 6.2.3 Logging In to a Linux ECS Using VNC (Through the Console).
	No login mode configuration item displayed on the creation page	<p>The relationship between the authentication mode and login mode set during image creation is as follows:</p> <ul style="list-style-type: none"> Key authentication. For details about how to log in to a Linux ECS, see 6.2.1 Remotely Logging In to a Linux ECS Using a Key Pair (SSH). Password authentication. For details about how to log in to a Linux ECS, see 6.2.2 Remotely Logging In to a Linux ECS Using a Password (SSH) or 6.2.3 Logging In to a Linux ECS Using VNC (Through the Console).
Windows (Only password authentication is supported when you log in to an ECS. The login modes include VNC and MSTSC.)	Key pair	For details about how to obtain the password, see 6.3.1 Obtaining the Password for Logging In to a Windows ECS . Then, log in to a Windows ECS by referring to 6.3.3 Logging In to a Windows ECS Using a Password (MSTSC) or 6.3.2 Logging In to a Windows ECS Using VNC (Through the Console) .
	Password	<ul style="list-style-type: none"> For details about how to remotely log in to a Windows ECS using VNC on the Web UI, see 6.3.2 Logging In to a Windows ECS Using VNC (Through the Console). For details about how to log in to a Windows ECS using MSTSC, see 6.3.3 Logging In to a Windows ECS Using a Password (MSTSC).

OS	Login Mode Selected When Creating an ECS	Available Login Method
	No login mode configuration item displayed on the creation page	<p>The relationship between the authentication mode and login mode set during image creation is as follows:</p> <ul style="list-style-type: none"> • Key authentication. For details about how to obtain the password, see 6.3.1 Obtaining the Password for Logging In to a Windows ECS. Then, log in to a Windows ECS by referring to 6.3.3 Logging In to a Windows ECS Using a Password (MSTSC) or 6.3.2 Logging In to a Windows ECS Using VNC (Through the Console). • Password authentication. For details about how to log in to a Windows ECS, see 6.3.3 Logging In to a Windows ECS Using a Password (MSTSC) or 6.3.2 Logging In to a Windows ECS Using VNC (Through the Console).

6.2 Logging In to a Linux ECS

6.2.1 Remotely Logging In to a Linux ECS Using a Key Pair (SSH)

If your ECS uses the key pair authentication mode, you can log in to the ECS using SSH.

Prerequisites

- You have obtained the ECS key file.
- Both the network and OS of the ECS support the SSH protocol.
- You have configured the inbound rules of the security group to allow your local computer to access the ECS. For details, see [12.4 Configuring Security Group Rules](#).
- In the Region Type I scenario, if the subnet to which the ECS belongs is associated with a firewall, you need to allow the IP address of the login server and the login port to access the subnet.
- 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.

Procedure

Configuring the Login Permission in SSH Key Authentication Mode

Step 1 Log in to an ECS using VNC.

Step 2 Run the following command to open the `sshd_config` file.

```
sudo vi /etc/ssh/sshd_config
```

Step 3 Check whether the file contains the `RSAAuthentication` and `PubkeyAuthentication` configuration items.

- If yes, change the values of `RSAAuthentication` and `PubkeyAuthentication` to `yes`.
- If no, add `RSAAuthentication` and `PubkeyAuthentication` to the file and set their values to `yes`.

Step 4 Check whether the file contains the `PubkeyAcceptedKeyTypes` configuration item.

- If yes, check whether the configuration item contains `ssh-rsa`.
 - If yes, go to [Step 5](#).
 - If no, manually add `ssh-rsa` to `PubkeyAcceptedKeyTypes` and go to [Step 5](#).
- If no, go to [Step 5](#).

Step 5 Run the following command to restart the SSH service:

```
sudo service sshd restart
```

----End

Logging In to the Linux ECS from a Windows Computer

The following procedure uses PuTTY as an example. Before logging in to the ECS using PuTTY, make sure that the private key file has been converted to .ppk format.

Step 1 Check whether the private key file has been converted to .ppk format.

- If yes, go to [Step 6](#).
- If no, go to [Step 2](#).

Step 2 Check whether PuTTY and PuTTYgen have been installed.

- If yes, go to [Step 3](#).
- If no, download and install PuTTY and PuTTYgen.

NOTE


PuTTYgen is a private key generator, which is used to create a key pair that consists of a public key and a private key for login authentication. Ignore this step if PuTTY and PuTTYgen have been installed.

Step 3 Run PuTTYgen.

Step 4 In the **Actions** area, click **Load** and import the private key file that you stored when creating the ECS.

Set the file type to **All files (*.*)** when importing the private key file.

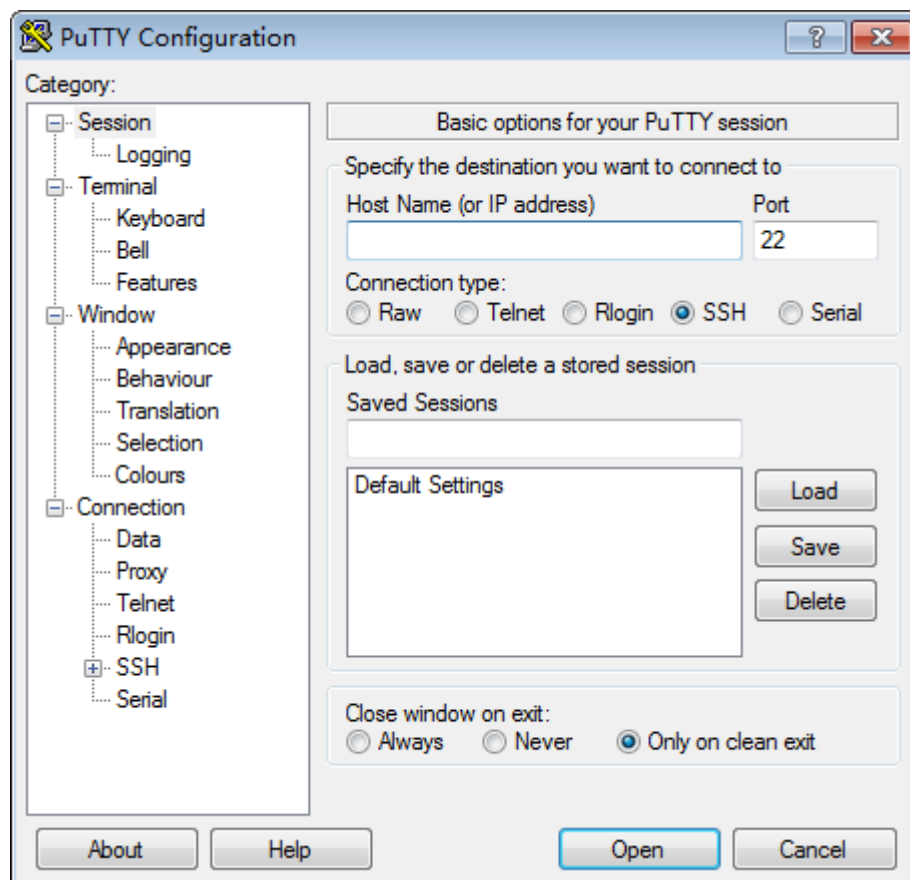
Figure 6-1 File Type



Step 5 Click **Save private key**. Save the converted private key, for example, **kp-123.ppk**, to the local computer.

Step 6 Double-click **PuTTY.EXE**. The **PuTTY Configuration** page is displayed.

Figure 6-2 PuTTY Configuration



Step 7 Choose **Connection > Data** and enter **root** in **Auto-login username**.

Step 8 Choose **Connection > SSH > Auth**. In configuration item **Private key file for authentication**, click **Browse** and select the private key file or the private key converted in [Step 5](#).

Step 9 Choose **Session** and enter the elastic IP address of the ECS under **Host Name (or IP address)**.

Step 10 Click **Open**.

Log in to the ECS.

----End

Logging In to the Linux ECS from a Linux Computer

The following procedure uses 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. Run the following command to change operation permission for the private key file to read-only:

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

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

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

```
ssh -i /path/kp-123 Default username@elastic IP address
```

For example, if the default username is **linux**, run the following command:

```
ssh -i /path/kp-123 linux@elasticIP address
```

In the preceding command,

path is the path where the private key file is stored and *EIP* is the EIP bound to the ECS.

6.2.2 Remotely Logging In to a Linux ECS Using a Password (SSH)

You can use SSH and a password to log in to an ECS in the following cases:

- The ECS uses password authentication.
- The ECS uses key authentication. The **root** user password is reset by using the key.
- The authentication mode cannot be set during ECS creation. You have obtained the **root** password for the image used to create the ECS.

Prerequisites

- Both the network and OS of the ECS support the SSH protocol.
- You have configured the inbound rules of the security group to allow your local computer to access the ECS. For details, see [12.4 Configuring Security Group Rules](#).
- In the Region Type I scenario, if the subnet to which the ECS belongs is associated with a firewall, you need to allow the IP address of the login server and the login port to access the subnet.
- 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.

Procedure

- Step 1** Perform operations according to the login mode you selected when applying for the ECS.

Table 6-2 Login method

Login Mode Selected When Creating an ECS	SSH Password Login Process
Key pair	<ol style="list-style-type: none"> 1. Reset the root user password. <ol style="list-style-type: none"> a. Use the SSH key to log in to the Linux ECS. For details, see 6.2.1 Remotely Logging In to a Linux ECS Using a Key Pair (SSH). b. Run the following command to set the user root password: sudo passwd root c. Enter the new password as prompted and press Enter. d. Confirm the password and press Enter. e. Verify that the information displayed is similar to the following, indicating that the password has been reset: passwd: all authentication tokens updated successfully. 2. Perform Step 2 to Step 3.
Password	Perform Step 2 to Step 3 .
No login mode configuration item displayed on the creation page	Perform Step 2 to Step 3 .

Step 2 Check whether the login permission in SSH password is enabled.

1. Log in to an ECS using VNC.
2. Run the following command to open the **sshd_config** file.
sudo vi /etc/ssh/sshd_config
3. Change the value of **PasswordAuthentication** in the **sshd_config** file to **yes**.

 **NOTE**

For the ECSs running the SUSE or openSUSE OSs, ensure that the values of the following configuration items in **/etc/ssh/sshd_config** are all **yes**.

- PermitRootLogin
- PasswordAuthenticatio
- ChallengeResponseAuthentication
- UsePAM

Step 3 Log in to the ECS.

- Windows computer

The following uses PuTTY as an example to describe how to log in to the ECS:

- a. Run PuTTY.
- b. Choose **Session** and enter the elastic IP address of the ECS under **Host Name (or IP address)**.

- c. Click **Open**.
- d. Enter the username **root** and the password set during ECS creation to log in to the ECS.

 **NOTE**

If the image has no Cloud-Init installed, enter the password of the **root** user configured during image creation.

- Linux computer

On the CLI of your computer, run the following command to log in to the ECS:

```
ssh EIP bound to the ECS
```

----End

6.2.3 Logging In to a Linux ECS Using VNC (Through the Console)

VNC is a capability provided by UI for remotely logging in to an ECS. Only password authentication is supported. You can use VNC and a password to log in to an ECS in the following cases:

- The ECS uses password authentication.
- The ECS uses key authentication. The **root** user password is reset by using the key.
- The authentication mode cannot be set during ECS creation. You have obtained the **root** password for the image used to create the ECS.

 **NOTE**

- If the ECS resides on an Arm host, you can only use the mouse embedded in the browser after logging in to the ECS using VNC.
- If a user has logged in to the ECS using VNC, other users cannot log in to it by default. You can modify related configuration items to allow other users to log in to the ECS in preemption mode. For details, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > Logging In to an ECS Using VNC in Preemption Mode**.

Prerequisites

- The remote login function is implemented using a system port (default port: 8002). Therefore, before attempting to log in remotely, ensure that the port to be used is not blocked by the firewall.
- If the client OS uses a local proxy and the firewall port cannot be configured on the local proxy, disable the proxy mode and then try logging in remotely.
- Only the ECS creator can remotely log in to an ECS by default. Other users or administrators do not have the login permission. To solve this issue, log in to a FusionSphere OpenStack controller node and modify the permission. For details, see "Product Management" > "Resource Pools" > "FusionSphere OpenStack" > "Compute" > "Modifying the Permission for Remotely Logging In to an ECS Using VNC" in *HUAWEI CLOUD Stack 8.1.0 O&M Guide*.

Procedure

- Step 1** Perform operations according to the login mode you selected when applying for the ECS.

Table 6-3 Login method

Login Mode Selected When Creating an ECS	VNC Login Process
Key pair	<ol style="list-style-type: none"> 1. Reset the root user password. <ol style="list-style-type: none"> a. Use the SSH key to log in to the Linux ECS. For details, see 6.2.1 Remotely Logging In to a Linux ECS Using a Key Pair (SSH). b. Run the following command to set the user root password: sudo passwd root c. Enter the new password as prompted and press Enter. d. Confirm the password and press Enter. e. Verify that the information displayed is similar to the following, indicating that the password has been reset: passwd: all authentication tokens updated successfully. 2. Perform Step 2 to Step 7.
Password	Perform Step 2 to Step 7 .
No login mode configuration item displayed on the creation page	Perform Step 2 to Step 7 .

- Step 2** Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 3** In the search box above the upper right corner of the ECS list, enter the target ECS name, IP address, ID, or CPU vendor, and click the search icon to search for the ECS. You can also search for the ECS by tag.
- Step 4** Locate the row containing the ECS and click **Remote Login** in the **Operation** column.
The **Configure Keyboard Layout for Remote Login** dialog box is displayed.
- Step 5** Select the English keyboard and click **Remote Login**.
- Step 6** (Optional) If the system displays "Press CTRL+ALT+DELETE to log on", click **Send CtrlAltDel** in the upper right corner of the remote login page to log in to the ECS.

Figure 6-3 Send CtrlAltDel



- Step 7** Enter the password you set during ECS creation or the password you reset in [Step 1.1](#) to log in.

 **NOTE**

If the image has no Cloud-Init installed, enter the password of the **root** user configured during image creation.

----End

6.3 Logging In to a Windows ECS

6.3.1 Obtaining the Password for Logging In to a Windows ECS

Context

Password authentication is required to log in to a Windows ECS. If you used key pair authentication when creating the ECS, you must obtain the administrator password generated when the ECS was initially installed. The administrator user is **Administrator** or the user configured using Cloudbase-Init. This password is randomly generated, offering high security.

Prerequisites

You have obtained the ECS key file.

Procedure

- Step 1** Obtain the private key file (.pem file) used when you created the ECS.
- Step 2** Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 3** In the ECS list, select the ECS for which the password is to be obtained.
- Step 4** In the **Operation** column, choose **More > Change Settings > Get Password**.
- Step 5** 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.
- Step 6** Click **Get Password** to obtain a random password.

----End

Follow Up

To ensure security, it is recommended that you clear the password record in the system after obtaining the password. For details, see section [8.2 Deleting the Initial Password for Logging In to a Windows ECS](#).

6.3.2 Logging In to a Windows ECS Using VNC (Through the Console)

VNC is a capability provided by UI for remotely logging in to an ECS. Only password authentication is supported. You can use VNC and a password to log in to an ECS in the following cases:

- The ECS uses password authentication.
- The ECS uses key authentication. The password is obtained by using the key.
- The authentication mode cannot be set during ECS creation. You have obtained the preset user password for the image used to create the ECS.

NOTE

- If the ECS resides on an Arm host, you can only use the mouse embedded in the browser after logging in to the ECS using VNC.
- If a user has logged in to the ECS using VNC, other users cannot log in to it by default. You can modify related configuration items to allow other users to log in to the ECS in preemption mode. For details, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > Logging In to an ECS Using VNC in Preemption Mode**.

Prerequisites

- The remote login function is implemented using a system port (default port: 8002). Therefore, before attempting to log in remotely, ensure that the port to be used is not blocked by the firewall.
- If the client OS uses a local proxy and the firewall port cannot be configured on the local proxy, disable the proxy mode and then try logging in remotely.
- If your ECS uses the password authentication mode, you have obtained the password for logging in to the Windows ECS. For details, see [6.3.1 Obtaining the Password for Logging In to a Windows ECS](#).
- Only the ECS creator can remotely log in to an ECS by default. Other users or administrators do not have the login permission. To solve this issue, log in to a FusionSphere OpenStack controller node and modify the permission. For details, see "Product Management" > "Resource Pools" > "FusionSphere OpenStack" > "Compute" > "Modifying the Permission for Remotely Logging In to an ECS Using VNC" in *HUAWEI CLOUD Stack 8.1.0 O&M Guide*.

Procedure

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 2** In the search box above the upper-right corner of the ECS list, enter the ECS name, IP address, or ID, and click the search icon to search for the ECS. You can also search for the ECS by tag.
- Step 3** Locate the row containing the ECS and click **Remote Login** in the **Operation** column.

The **Configure Keyboard Layout for Remote Login** dialog box is displayed.

Step 4 Select the English keyboard and click **Remote Login**.

Step 5 (Optional) If the system displays "Press CTRL+ALT+DELETE to log on", click **Send CtrlAltDel** in the upper right corner of the remote login page to log in to the ECS.

Figure 6-4 Send CtrlAltDel



Step 6 Enter the password you set during ECS creation or the password you obtained in [6.3.1 Obtaining the Password for Logging In to a Windows ECS](#) to log in to the ECS.

 **NOTE**

If the image has no Cloudbase-Init installed, enter the password configured during image creation.

----End

6.3.3 Logging In to a Windows ECS Using a Password (MSTSC)

If a Windows ECS uses any of the following authentication methods, you can use MSTSC (remote login tool) to log in to the ECS from a local server:

- The ECS uses password authentication.
- The ECS uses key authentication. The password is obtained by using the key.
- The authentication mode cannot be set during ECS creation. You have obtained the preset user password for the image used to create the ECS.

Context

Remote Desktop Protocol (RDP) is disabled on Windows ECSs by default. Before using MSTSC to log in to a Windows ECS, use VNC to log in to it and enable RDP.

Prerequisites

- If your ECS uses the password authentication mode, you have obtained the password for logging in to the Windows ECS. For details, see [6.3.1 Obtaining the Password for Logging In to a Windows ECS](#).
- You have bound an EIP to the ECS. If no EIP has been bound to the ECS, bind an EIP by referring to "Applying for an EIP" and "Binding an EIP" in **Operation Help Center > Network > Elastic IP > User Guide**.
- You have configured the inbound rules of the security group to allow your local computer to access the ECS. For details, see [12.4 Configuring Security Group Rules](#).
- In the Region Type I scenario, if the subnet to which the ECS belongs is associated with a firewall, you need to allow the IP address of the login server and the login port to access the subnet.
- 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.

Procedure

Step 1 Check whether RDP is enabled on the Windows ECS.

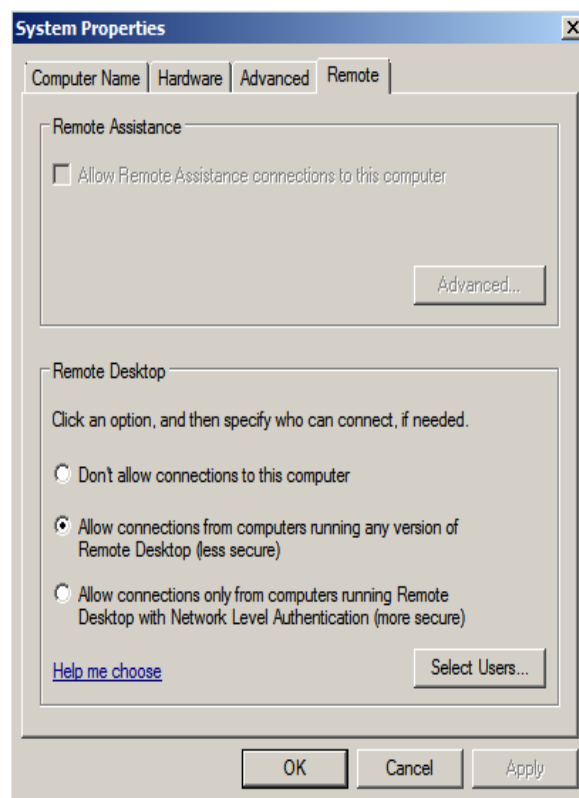
1. Log in to the Windows ECS using VNC.

For details, see [6.3.2 Logging In to a Windows ECS Using VNC \(Through the Console\)](#).

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 6-5 System Properties



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

Step 2 On your local computer, use MSTSC to log in to the Windows ECS.

1. Click **Start** in the task bar.
2. In the **Search programs and files** box, enter **mstsc**.
3. Log in to the ECS according to the prompts.

To ensure system security, change the login password after you log in to the ECS for the first time.

Step 3 (Optional) After logging in to the ECS using Remote Desktop Protocol (RDP), handle the issue that local files larger than 2 GB cannot be copied to a remote Windows ECS.

Perform this step only if you need to use the RDP clipboard. This issue occurs due to Windows OS limitations. For details, see <https://support.microsoft.com/en-us/help/2258090/copying-files-larger-than-2-gb-over-a-remote-desktop-services-or-terminal-services-session-by-using-clipboard-redirection-copy-and-paste-fails-silently>.

----End

7 Managing an ECS

7.1 Basic Operations

This section describes basic management operations on an ECS, including viewing details, changing the name, managing the tag, managing the life cycle, and changing the time zone of an ECS.

7.1.1 Viewing ECS Details

This section describes how to view details about an ECS, including its name, creator, EVS system disk, EVS data disks, VPCs, NICs, and security groups.

Procedure

Step 1 Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 In the search box above the upper right corner of the ECS list, enter the target ECS name, IP address, ID, or CPU vendor, and click the search icon to search for the ECS. You can also search for the ECS by tag.

Step 3 Click the name of the ECS to be queried.

The page providing details about the ECS is displayed.

Step 4 View the ECS details.



You can view ECS details by its EVS disks, NICs, security groups, EIPs, monitoring metrics, tags, ECS snapshots, CD-ROM drive, and ECS group.

----End

7.1.2 Changing the ECS Name

This section describes how to change the name of an ECS.

Procedure

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 2** In the search box above the upper right corner of the ECS list, enter the target ECS name, IP address, ID, or CPU vendor, and click the search icon to search for the ECS. You can also search for the ECS by tag.
- Step 3** Click the name of the ECS.
The page providing details about the ECS is displayed.
- Step 4** Click  next to the ECS name, change it, and click  to save the change.
----End

7.1.3 Adding and Managing ECS Tags

A tag identifies an ECS. Adding tags to an ECS facilitates ECS identification and management.

Context

A tag can be manually added or automatically generated.

- **Manually adding a tag:** You can add a tag when creating an ECS or after creating it (add the tag on the page providing details about the ECS).
- **Automatically generating a tag:** When you create an ECS, the system automatically generates a built-in tag for your ECS. The key of the built-in tag is the VPC ID of your ECS. The tag is not displayed by default.

A tag consists of a key and a value. The key and value come from the Tag Management Service. To add or modify a tag, contact the administrator to choose **Console > Mgmt & Deployment > Tag Management** in the menu bar on the upper part of the page. You can click the **Tag** tab on the ECS details page, and then query, add, delete, or modify tags.

Procedure

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 2** Go to the tag management page in either of the following ways:
Method 1: On the **Elastic Cloud Server** page, locate the row that contains the target ECS, choose **More > Change Settings > Add/Edit Tag** in the **Operation** column.
Method 2: Click the name of the target ECS. On the displayed ECS details page, click the **Tag** tab.
- Step 3** Add, delete, modify, and view ECS tags.
 - **Viewing tags**
You can view details of ECS tags, including the number of tags and the key and value of each tag.

 **NOTE**

When you create an ECS, the built-in tag automatically generated by the system for your ECS is unavailable on the UI.

- Adding a tag

In the upper left corner of the page, click **Add/Edit**. In the **Add/Edit Tag** dialog box, click **Add Tag**, select a key and value, and then click **OK**.

 **NOTE**

When multiple tags are added to an ECS, each tag must have a unique key.

- Modifying a tag

In the upper left corner of the page, click **Add/Edit**. In the **Add/Edit Tag** dialog box, change the key and value of the tag and click **OK**.

- Deleting a tag

In the upper left corner of the page, click **Add/Edit**. In the **Add/Edit Tag** dialog box, click **Delete** of the tag and click **OK**.

----End

7.1.4 Querying ECSs by Filters

You can query ECSs by running status, ECS name, private IP address, EIP, ID, or tag.

Procedure

Querying ECSs by name, private IP address, EIP, ID, or CPU vendor

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 2** In the search box above the upper right corner of the ECS list, select the search item, enter the ECS name, private IP address, EIP, ID, or CPU vendor, and click the search icon to search for the ECS.

 **NOTE**

When you query ECSs by name or IP address, fuzzy query is supported. If you enter part of an IP address, such as **8.3**, all ECSs whose IP addresses contain **8.3** will be displayed in the query result, for example, 192.168.8.3 and 192.168.35.2.

----End

Querying ECSs by tag

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 2** Click **Search by Tag** above the upper right corner of the ECS list to expand the search area.
- Step 3** Query ECSs by tag.
1. Select a key.
 2. Select a value.

3. Click **Add Tag**.

Step 4 Click **Search**.

The system searches for the target ECS based on the key and value of the tag.

 **NOTE**

If you need to query ECSs by multiple tags, repeat **Step 3**. When multiple tags are of the AND relationship, the system filters ECSs that meet the requirements according to all tags.

----End

7.1.5 Exporting ECS Details

You can export the details of some or all ECSs to a local directory in **.xlsx** format. The exported details include the ID, name, running status, flavor, image (version), IP address, EIP, AZ, and usage duration.

Procedure

Step 1 Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 Export an ECS.

- If you want to export the details of all ECSs in the current resource set, click the export icon in the upper right corner of the ECS list.
- If you want to export the details of some ECSs, select target ECSs on the same page and click **Export**. If you select multiple pages, the information about the ECS on the last page is exported.

 **NOTE**

- If the system displays a message indicating that you do not have the permission, contact the administrator to change your permissions. The procedure is as follows:
 1. Log in to ManageOne as an administrator and choose **System** on the top menu bar.
 2. In the navigation pane, choose **Role Management**.
 3. Click **Modify** in the **Operation** column of the corresponding role.
 4. In the **Project Permissions** area, select **Query ECS** under **Fine-grained permissions of ECS** and click **OK**.
- If you click the export icon in the upper right corner of the list of ECSs filtered by an IP address, all ECSs in the current resource set are exported. You can export all ECSs and filter them in the exported Excel file, or adjust the number of ECSs displayed on a single page to display all filtered ECSs on one page. Then select the ECSs to be exported, and click **Export** in the upper left corner of the page.
- Exporting all ECSs may take a long time. You cannot switch to other pages from the ECS list page during the export. To ensure normal operations, select **Open New Page** and perform other operations on the new page.

----End

7.1.6 Changing the Time Zone for an ECS

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 on an ECS to the local time or to another time zone in your network.

Procedure

- For a Linux ECS

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.

Step 1 Log in to the ECS.

Step 2 Run the following command to switch to user **root**:

```
su - root
```

Step 3 Run the following command to view the time zones supported by the ECS:

```
ls /usr/share/zoneinfo/
```

In the terminal display, the **/usr/share/zoneinfo** directory contains a hierarchy of time zone data files. Use the directory structure shown in **Figure 7-1** to obtain your desired time zone file.

Figure 7-1 Time zones supported by the ECS

```
[root@PEK***** ~]# ls /usr/share/zoneinfo/
Africa      Australia  Cuba      Etc      GMT-0     Indian    Kwajalein Navajo    posix     ROK       UTC
America     Brazil    EET       Europe   GMT+0     Iran      Libya     NZ        posixrules Singapore WET
Antarctica  Canada    Egypt     GB       Greenwich iso3166.tab MET       NZ-CHAT   PRC       Turkey    W-SU
Arctic      CET       Eire      GB-Eire  Hongkong  Israel    Mexico    Pacific   PST8PDT   UCT       zone.tab
Asia        Chile     EST       GMT      HST       Jamaica   MST       Poland    right     Universal Zulu
Atlantic    CST6CDT  EST5EDT  GMT0     Iceland   Japan     MST7MDT   Portugal  ROC       US
```

The directory structure shown in **Figure 7-1** 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 Singapore, the directory in which the time zone file is stored is **/usr/share/zoneinfo/Singapore**.
- 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**.

Step 4 Set the target time zone.

1. Run the following command to open the **/etc/sysconfig/clock** file:

```
vim /etc/sysconfig/clock
```

2. 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 Singapore, change the **ZONE** entry value to **Singapore**:
ZONE="Singapore"
- If the target time zone is for Paris, change the **ZONE** entry value to **Paris**:
ZONE="Europe/Paris"

Step 5 Press **Esc**, and run the **:wq** command to save and exit the **/etc/sysconfig/clock** file.

```
:wq
```

Step 6 Run the following command to check whether the **/etc/localtime** file is available on the ECS:

```
ls /etc/localtime
```

- If yes, go to [Step 7](#).
- If no, go to [Step 8](#).

Step 7 Run the following command to delete the existing **/etc/localtime** file:

```
rm /etc/localtime
```

Step 8 Run the following command to create a symbolic link between **/etc/localtime** and your time zone file so that the ECS can find this time zone file when it references the local time:

```
ln -sf /usr/share/zoneinfo/Singapore /etc/localtime
```

Step 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
```

Step 10 Log in to the ECS again and run the following command as user **root** to check whether the time zone has been changed successfully:

```
ls -lh /etc/localtime
```

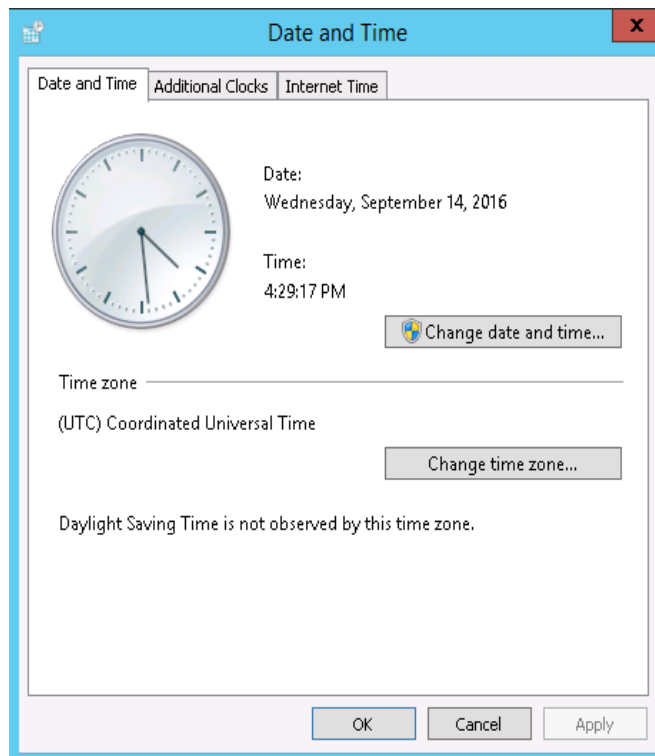
```
----End
```

- For a Windows ECS

Step 1 Click the time display on the far right side of the task bar located at the bottom of your screen. In the displayed dialog box, click **Change date and time settings**.

The **Date and Time** is displayed.

Figure 7-2 Date and Time



Step 2 Click **Change time zone**.

The **Time Zone Settings** page is displayed.

Step 3 In the **Set the time zone** pane, choose the target time zone from the **Time zone** drop-down list.

Step 4 Click **OK**.

----End

7.2 Life Cycle

7.2.1 Managing the Life Cycle of an ECS

ECS lifecycle management operations include starting, stopping, restarting, and deleting ECSs.

Context

- When many ECSs are started or stopped at the same time, heavy workload runs on the host. It is recommended that you start or stop only a few ECSs at a time so that no adverse impact will be caused on services of other ECSs.
- You can forcibly restart or stop a frozen ECS (that remains in the **Restarting** or **Stopping** state for a long time) at any time.

 **NOTE**

- Stopping or restarting an ECS will interrupt services running on the ECS. Exercise caution when performing this operation.
- Forcibly stopping or restarting an ECS will cause any unsaved data on the ECS to be lost. Exercise caution when performing this operation.
- Deleting an ECS will delete all snapshots created for it. Exercise caution when performing this operation.

Procedure

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 2** In the search box above the upper right corner of the ECS list, enter the target ECS name, IP address, ID, or CPU vendor, and click the search icon to search for the ECS. You can also search for the ECS by tag.
- Step 3** If you want to manage one ECS, choose **More > Change Status** in the **Operation** column and select the required operation, for example, to stop, delete, or restart the ECS. To manage ECSs in batches, select the target ECSs in the ECS list, click **Operation** above the ECS list, and choose **Start, Stop, Restart, Delete, or Extend**.

 **NOTE**

- If the ECS is stopped and the host group where the ECS resides is configured with tags for releasing resources upon ECS shutdown, the ECS releases the occupied resources after it is shut down. The resources include vCPU, memory, GPU, NPU, USB, and volume connections. For details, see step "Configure custom tags" in **Operation Help Center > Operation > Compute Services > Elastic Cloud Server (ECS) > Configuration Before ECS Creation > Creating a Host Group**.
- When Arm servers are used, NIC connections can be released when general computing-plus ECSs are shut down.
- Resources will not be released for disk-intensive or ultra-high I/O ECSs when they are shut down.
- Only ECSs that are in stopped or running state can be deleted.

- Step 4** Confirm the displayed information. [Table 7-1](#) describes the ECS statuses.

Table 7-1 ECS statuses

Status	Status Attribute	Description	Corresponding API Status
Creating	Intermediate	The ECS has been created but is not running.	BUILD/ BUILDING
Starting	Intermediate	The ECS is between the Stopped and Running states.	SHUTOFF
Running	Stable	The ECS is running properly. An ECS in this state can provide services.	ACTIVE

Status	Status Attribute	Description	Corresponding API Status
Stopping	Intermediate	The ECS instance is between the Running and Stopped states.	ACTIVE
Stopped	Stable	The ECS has been properly stopped. An ECS in this state cannot provide services.	SHUTOFF
Restarting	Intermediate	The ECS is being restarted.	REBOOT
Resizing	Intermediate	The ECS has received a change request and has started to perform the change operation.	RESIZE
Verifying resizing	Intermediate	The ECS is verifying the modified configuration.	VERIFY_RESIZE
Deleting	Intermediate	The ECS is being deleted. If the ECS remains in this state for a long time, exceptions may have occurred. In such an event, contact the administrator.	ACTIVE/ SHUTOFF/ REBOOT/ RESIZE/ VERIFR_RESIZE / / HARD_REBOOT/ REVERT_RESIZE/ ERROR
Deleted	Intermediate	The ECS has been deleted. An ECS in this state cannot provide services and will be promptly cleared from the system.	DELETED
Faulty	Stable	An exception has occurred on the ECS. An ECS in this state cannot provide services. Contact the administrator.	ERROR
Reinstalling OS	Intermediate	The ECS has received a request to reinstall the OS and has begun the installation.	SHUTOFF
Reinstalling OS failed	Stable	The ECS received a request to reinstall the OS, but due to exceptions, the installation failed. An ECS in this state cannot provide services. Contact customer service.	SHUTOFF
Changing OS	Intermediate	The ECS received a request to change the OS and has begun implementing the changes.	SHUTOFF

Status	Status Attribute	Description	Corresponding API Status
OS Change failed	Stable	The ECS has received a request to change the OS, but due to exceptions, the change failed to be carried out. An ECS in this state cannot provide services. Contact the administrator.	SHUTOFF
Forcibly restarting	Intermediate	The ECS is being forcibly restarted.	HARD_REBOOT
Rolling back resizing	Intermediate	The ECS is rolling back resizing.	REVERT_RESIZE

NOTE

If an ECS continues to remain in an intermediate state for over 30 minutes, exceptions may have occurred. In this event, contact the administrator.

----End

7.2.2 Deleting an ECS

This section describes how to delete an ECS. When deleting an ECS, you can choose to delete only the ECS, or to release the EIP bound to the ECS and the data disks attached to the ECS.

Constraints

- ECSs configured with Volume High Availability (VHA), Cloud Server High Availability (CSHA), Cloud Server Disaster Recovery (CSDR), or VHA+CSDR protection can only be soft-deleted or permanently deleted after DR protection is first canceled.
- ECSs in MRS clusters cannot be soft-deleted or permanently deleted.

Procedure

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 2** Perform either of the following operations:
 - Select all ECSs to be deleted, click **Operation** above the ECS list, and choose **Delete**. You can use this method to delete one or multiple ECSs.
 - In the **Operation** column of the ECS to be deleted, choose **More > Change Status > Delete**. You can use this method to delete only one ECS at a time.
- Step 3** On the displayed page, select **Permanently Delete** or not.
 - Selecting **Permanently Delete**

If you select **Permanently Delete**, you can also choose whether to release the EIP bound to and the data disks attached to the ECS. Then, click **OK**.


A permanently deleted ECS cannot be restored. Exercise caution when performing this operation.

 **NOTE**

- If the ECS is running, it cannot be permanently deleted. You need to stop it before.
 - When deleting an ECS that shares disks with other ECSs, if you choose to also delete its data disks, the ECS will be deleted but the shared disk will be retained.
- Not selecting **Permanently Delete**

Do not select **Permanently Delete**, and click **OK**.

If the ECS is running, it will be shut down before being deleted and put into the recycle bin. After being restored from the recycle bin, the ECS will be

automatically started. On the top navigation bar, click  and choose **Mgmt & Deployment > Recycle Bin**.

The ECS remains in the frozen state for 24 hours in the recycle bin. An ECS in the frozen period cannot be permanently deleted, but can be restored. After the frozen period ends, the ECS can be deleted or restored.

----End

7.2.3 Changing the Validity Period of an ECS

Context

If your ECS is about to expire or has expired, you can extend the validity period of your ECS so that it continues running properly.

 **NOTE**

- If you select **Never expires** as **Expired On** for an ECS, the validity period of the ECS cannot be extended.
- If the ECS has expired, you can only **Delete** or **Extend** the ECS.

Procedure

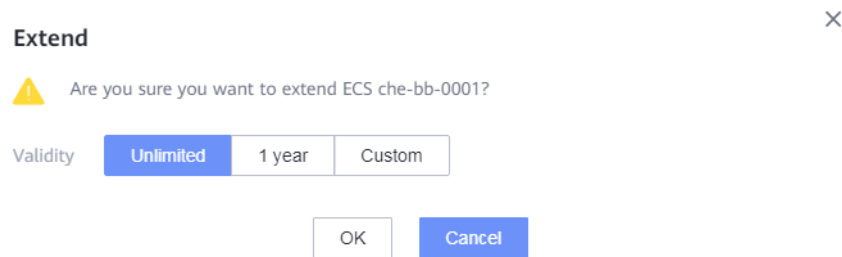
Step 1 Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 In the **Operation** column of the ECS to be extended, choose **More > Change Status > Extend**.

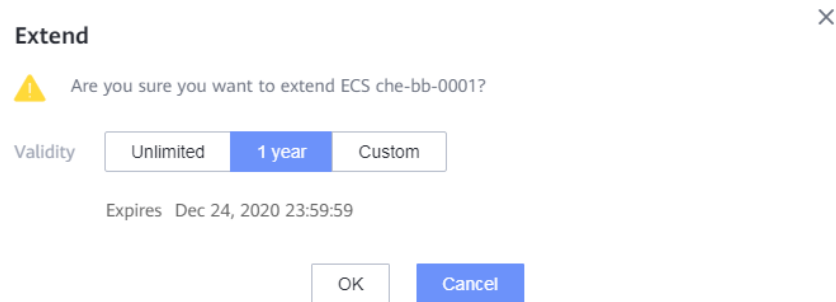
The **Extend** dialog box is displayed.

Step 3 Select **Validity**.


- Select **Unlimited** and click **OK**.



- Select **1 year** and click **OK**.



- Select **Custom**, click the calendar icon on the right, select a specific date to extend the validity period of an ECS, and click **OK**.



----End

7.3 Creating a Private Image Using an Existing ECS

You can create a private image from an existing ECS.

For details, see "Creating a Private Image" in [Operation Help Center > Compute > Image Management Service > User Guide](#).

7.4 Modifying the DR or Backup Function of an ECS

ECS backup, DR, and snapshot functions protect data integrity and service continuity. Before requesting a backup or DR service, set **Same Storage** to **Yes** for the ECS.

For existing ECSs, if **Same Storage** is set to **Yes**, you can change it to **No** at any time except for DR ECSs. If **Same Storage** is set to **No**, it can be changed to **Yes** only when specific conditions are met. For detailed operation procedure and constraints, see this section.

Procedure

Step 1 Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 Locate the row that contains the target ECS. In the **Operation** column, choose **More > Disk > Same storage**.

 **NOTE**

If **Same storage** is dimmed, **Boot Device** of the ECS flavor has been set to **Local Disk**.

Step 3 Set **Current Status** to **Yes** or **No**.

- If this parameter is set to **No**, the ECS does not support DR or backup. Whether it supports ECS snapshots depends on whether all disks of the ECS reside in the same storage backend. You can check whether ECS snapshots are supported on the ECS snapshot page.

 **NOTE**

Current Status cannot be set to **No** for ECSs configured with VHA, CSHA, CSDR, or VHA+CSDR protection.

- If this parameter is set to **Yes**, the ECS supports DR, backup, and ECS snapshot. **Current Status** can be set to **Yes** only for ECSs meeting the following conditions:
 - **Boot Device** of the ECS flavor must be **Cloud Disk**.
 - If the ECS only has the system disk, make sure that the storage backend where the system disk resides has the storage tag configured. For details about the configuration, visit **Operation Help Center > Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > Disk FAQs > (Optional) Creating a Disk Type**.
 - If the ECS has both system and data disks, these disks must reside in the same storage backend, and the backend storage must be configured with a storage tag. For details about the configuration, visit **Operation Help Center > Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > Disk FAQs > (Optional) Creating a Disk Type**.

----End

7.5 Cloning an ECS

Cloning an ECS is to copy its system and data disks and by doing so create a clone of it. (Not all ECSs can be cloned.) For security purposes, you may need to log in to the new ECS created by cloning a Linux ECS to delete the original password or key.

Context

- When you clone an existing ECS to create an ECS, the new ECS has the same attributes and parameter settings as the existing ECS except for the ID, MAC address, IP address, VIP, EIP, and password or key (the password or key is reset during cloning).

- Cloning includes online clone and offline clone. Online clone means that an ECS to be cloned stays in the **Running** state. Offline clone means that an ECS is in the **Stopped** state.
- Even if you have reset the password or key when cloning a Linux ECS, the new ECS may retain the original password or key. In this case, you can log in to the new ECS using the original password or key, or the new password or key. For security purposes, check whether you need to log in to the new ECS to delete the original password or key by referring to [Table 7-2](#).

Table 7-2 Authentication modes for the source ECS and the new ECS created by cloning the source ECS

ECS OS	Source ECS Authentication Mode	Clone ECS Authentication Mode	Whether to Retain Original Password or Key	Whether You Need to Delete Original Password or Key
Windows	Password or key	Password or key	No	No
Linux	Password	Password	No	No
	Password	Key	Yes	Yes
	Key	Password	Yes	Yes
	Key	Key	Yes	Yes

Constraints

Table 7-3 Limitation description

Restriction Type	Description
ECS Type	<ul style="list-style-type: none"> • Ultra-high I/O and USB-passthrough ECSs cannot be cloned. • If Boot Device is set to Local Disk in the flavor of an ECS, the ECS cannot be cloned. • ECSs whose virtualization type is KVM support online and offline clone. • ECSs configured with CSHA, CSDR, or VHA+CSDR protection cannot be cloned.
ECS running status	An ECS to be cloned has not expired and is in the Stopped or Running state.

Restriction Type	Description
Disks	<ul style="list-style-type: none">ECSs configured with shared disks cannot be cloned. ECSs with heterogeneous storage (non-Huawei SAN storage or non-Huawei Distributed Block Storage) or all disks (system and data disks) whose the backend storage SN is different support offline clone only. The backend storage where the system and data disks of the ECS reside must support snapshot and LUN copy. For OceanStor V3/V5 series storage, the administrator needs to activate the snapshot and LUN copy features when requesting a license for the disk array. For Dorado V3 series storage (software version: V3R1C21 or later), the snapshot feature needs to be activated. For storage systems of other vendors, also make sure that the snapshot and LUN copy features are activated as per the allocated license. Otherwise, the cloning will fail.When the backend storage is SAN storage, the default cloning rate is 10 to 20 MB per second. If the disks of an ECS contain large amounts of data, cloning it may take a long time. You can accelerate the cloning speed by increasing the LUN copy speed on the storage backend. For details, visit Operation Help Center and choose Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > How Do I Set the Cloning Speed of an ECS?
ECS Group	<ul style="list-style-type: none">If the source ECS is added to an ECS group, the cloned ECS is also added to the ECS group and complies with the ECS policy.If the policy of the ECS group is affinity or anti-affinity and no host meets the policy, the cloning fails.
Account	<ul style="list-style-type: none">The role to which your account belongs, for example, VDC operator, must have the permission to clone ECSs.If your account balance is insufficient, the system will display an error message after you submit a request for cloning an ECS. In this case, top up your account, and then clone your ECS again.

Procedure

Cloning an ECS

Step 1 Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 Locate the row containing the target ECS. In the **Operation** column, click **More**, and choose **Clone**.

The **Clone ECS** page is displayed.

Step 3 Set the ECS name.

- When cloning an ECS, you can customize the suffix.
- When you batch-clone ECSs, the system automatically numbers them, for example, ecs-0001, ecs-0002, and so on. The default start number is 0001, and the maximum number is 9999.

To start with a specific number, click **Change Suffix Start Number** to customize the value. For example, if you set the value to 1126, the ECS names will be xxx-1126, xxx-1127, and so on.

 **NOTE**

If you want to create a Windows ECS that needs to be added to a domain, or if you require that the host name of the ECS (that is, the computer name shown in the ECS OS) must be unique, set the ECS name by following the instructions provided in **Operation Help Center > Operation > Compute Services > Elastic Cloud Server (ECS) > ECS Host Name > Rules for Configuring ECS Names (Unique Host Names)**.

Step 4 Set the running status of the ECS. This parameter is available if the virtualization type of the AZ where the ECS resides is KVM.

- **Stopped:** A newly cloned ECS stays in the **Stopped** state.

 **NOTE**

- If **ECS Initial Status** is set to **Stopped** and the host group where the new ECS resides is configured with tags for releasing resources upon ECS shutdown, the new ECS does not occupy any of the following resources: vCPU, memory, GPU, NPU, USB, and volume connections. For details, see step "Configure custom tags" in **Operation Help Center > Operation > Compute Services > Elastic Cloud Server (ECS) > Configuration Before ECS Creation > Creating a Host Group**.
- When Arm servers are used, NIC connections can be released when general computing-plus ECSs are shut down.
- Resources will not be released for disk-intensive or ultra-high I/O ECSs when they are shut down.

- **Running:** A newly cloned ECS stays in the **Running** state.

Step 5 Determine whether to select **Joint Windows Domain**. This parameter is available if the virtualization type of the AZ where the ECS resides is KVM, the ECS is running a Windows OS, and the domain information has been configured for the corresponding ECS product.

The administrator can perform unified authentication for ECSs added to the same domain. The following functions will be available for ECSs added to a domain: manage compute resources, reduce network management complexity and costs, enhance security, and support account roaming and folder redirection. Resources can be shared among ECSs in the same domain. For more information about domain servers and their functions, click [here](#).

Specify whether to add an ECS to a Windows domain. You can select a domain from the drop-down list. Available options are those defined by the administrator during product creation. If the ECS image is a [static injection image](#) or does not have Cloudbase-Init installed, it cannot be added to a domain.

Step 6 Select a validity period and set the quantity of ECSs created from cloning.

 **NOTE**

- By default, a maximum of 30 ECSs can be cloned at a time.

Step 7 Configure the network. Select **Reconfigure** to reconfigure the network for the cloned ECS.

Table 7-4 Parameters

Parameter	Description	Example Value
Resource Set	<p>Select the current resource set or another resource set from the drop-down list. You can view the current resource set on the top menu bar of the page. Assume that the current resource set is Resource Set A and another resource set available is Resource Set B.</p> <ul style="list-style-type: none">• When you select the current resource set, VPCs available will be those in Resource Set A.• If you select Resource Set B, VPCs available will be those in Resource Set B. By selecting Resource Set B, you create ECSs in Resource Set A by using the network resources of Resource Set B. With other configurations including security groups, you enable these ECSs to communicate with all those in the VPCs that belong to Resource Set B, allowing ECSs of different projects to share the same VPCs. <p>NOTE This parameter is available when VPC sharing is enabled on Service OM and the shared VPC permission is configured for the resource set on ManageOne. Otherwise, this parameter is not displayed. By default, this function is disabled. For details, visit Operation Help Center and choose Operation > Network Services > Virtual Private Cloud (VPC) > Shared VPC Best Practices.</p>	project_02
VPC	<p>A VPC provides a network, including subnets and security groups, for an ECS.</p> <p>You can select an existing VPC, or click Create VPC to create one.</p>	N/A

Parameter	Description	Example Value
NIC	<p>Includes primary and extension NICs. You can add a maximum of 15 extension NICs to an ECS.</p> <ul style="list-style-type: none"> If you select VPC Subnet, all subnets in the VPC are available for you to choose from. In this case, the NIC supports layer 3 communication, allowing the ECS to communicate with networks (for example, the public network or other VPCs) beyond the VPC. If you select Intra-Project Subnet, all project-level subnets in the project are available for you to choose from. All NICs configured with the same subnet can communicate with each other at layer 2 on the project level. Layer 2 communication is supported within the same VPC and between different VPCs. <p>The network type for Primary NIC must be VPC Subnet. Otherwise, you cannot access the Internet through the allocated EIP, and you cannot access Object Storage Service (OBS) or a security service. Set the network type for Extension NIC as required.</p> <p>The IP address is automatically assigned by the system.</p> <p>NOTE</p> <ul style="list-style-type: none"> If the selected subnet has only an IPv4 address segment, the NIC will only have one IPv4 address. If the selected subnet has both IPv4 and IPv6 address segments, the NIC will have one IPv4 address and one IPv6 address. If the selected subnet does not have DHCP enabled and the selected image does not support static IP address injection, after an ECS is created, you need to manually configure IP addresses for the ECS. Otherwise, the NIC cannot be reached. For details, see 18.6.1 Configuring a Static IP Address for an ECS. Security Group: When cloning an ECS, you can select multiple security groups. Multiple security groups may affect the ECS network performance. You are advised to select a maximum of five security groups. Each NIC can be configured with different security groups but must be configured with at least one available security group. 	<p>VPC Subnet subnet-c869(192.168.0.0/24) default</p>

Step 8 Specify whether to set a new password or key for the new ECS for login authentication.

- If you select **No**, the password or key of the new ECS is the same as that of the source ECS. In this case, go to [Step 10](#).
- If you select **Yes**, set a new password or key for the new ECS. Then, go to [Step 9](#).

NOTE

- For Windows ECSs, the image password cannot be used as the ECS password. Therefore, this parameter is not displayed on the page.
- For Linux ECSs, if the **Set New Password or Key** parameter is not displayed on the page, the image used to create the source ECS is not installed with Cloud-Init. When cloning an ECS of this type, you cannot set a new password for the new ECS. The login password of the new ECS is the same as that of the source ECS.

Step 9 Select **Password** or **Key pair** as the login mode, and enter the password or select the key pair.

Step 10 Click **OK**.

NOTE

- After cloning, the system does not automatically allocate an EIP or a virtual IP address to the clone ECS, that is, the new ECS generated. If an EIP or a virtual IP address is bound to the original ECS, you need to manually allocate an EIP or a virtual IP address to the new ECS.
- If the source ECS has a snapshot, the snapshot will not be cloned during cloning. Therefore, you need to request a new snapshot for the cloned ECS.

Step 11 Check whether you need to log in to the new ECS to delete the original password or key by referring to [Table 7-2](#).

- If there is no need to delete the original password or key, cloning an ECS is complete.
- If there is the need to delete the original password, log in to the new ECS and perform the following operations. For details, see [6.1 Login Mode Overview](#).
 - a. Run the following command to delete the original password of the **root** user.

passwd -d root

```
root@heyan-cloudinit2:~# passwd -d root
passwd: password expiry information changed.
root@heyan-cloudinit2:~#
```

- b. Run the following command to set a new password for the **root** user.
passwd root
- c. Enter a new password twice as prompted. If the following information is displayed, resetting the password is successful.

```
root@heyan-cloudinit2:~# passwd root
Enter new UNIX password:
Retype new UNIX password:
passwd: password updated successfully
root@heyan-cloudinit2:~# _
```

- If there is the need to delete the original key, log in to the new ECS and perform the following operations. For details, see [6.1 Login Mode Overview](#).

- a. Run the following command to open the file containing the public key.
vi /root/.ssh/authorized_keys
- b. Move the cursor to the row containing the target key, and then run the **dd** command to delete the key.
- c. Run the **:wq!** command to save the file.
- d. Run the **service sshd restart** command to restart the SSH service.

----End

7.6 Managing the Watchdog Status of an ECS

The watchdog function provides a heartbeat mechanism used to monitor the health status of ECSs. You can enable or disable watchdog for an ECS. The watchdog status for an ECS can be enabled, disabled, or not configured.

Prerequisites

- When x86 servers and KVM virtualization are used, before enabling the watchdog function, ensure that an IPMI-compliant dog-kicking program has been installed on the image file used to create the ECS. Otherwise, the ECS will restart repeatedly.
- In ARM scenarios, before enabling the watchdog function, ensure that a 6300ESB-compliant dog-kicking program has been installed on the image file used to create the ECS. Otherwise, the ECS will restart repeatedly.

Procedure

NOTE

- After the status of the watchdog is changed, the ECS will automatically restart. Perform this operation at a proper time to avoid impact on your services.
- If the ECS fails to restart, forcibly restart it.

Step 1 Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 Locate the row that contains the target ECS, and choose **More > Change Settings > Watchdog** in the **Operation** column.

Step 3 Change the status of watchdog.

- Toggle the watchdog status to **ON**, and click **OK**.
The status of watchdog will change to **Enable**.

NOTE

In ARM scenarios, if **Watchdog** is set to **Enable**, the **Watchdog Alarm Policy** parameter is displayed on the page. If the 6300ESB watchdog does not detect watchdog information in the specified time, an alarm will be generated. The ECS will determine, based on this alarm policy, whether to get restarted.

- Toggle the watchdog status to **OFF**, and click **OK**.
The status of watchdog will change to **Disable**.

----End

7.7 Managing the HA Status of an ECS

When HA is enabled, an ECS is automatically rebuilt on another host whenever the ECS or its host becomes faulty, ensuring service continuity. The HA status for an ECS can be enabled, disabled, or not configured.

Prerequisites

To support HA, ECSs must meet the following requirements: the global HA function is enabled, the HA function of the host group where the ECS resides is enabled or not configured, and the HA function of the ECS is enabled.

NOTE

- For details about how to enable the global HA function, see "Product Management" > "Resource Pool" > "FusionSphere OpenStack" > "Compute" > "Configuring the VM HA Function" in *HUAWEI CLOUD Stack 8.1.0 O&M Guide*.
- To check whether the HA function of the host group is enabled, log in to Service OM and check it in the **Custom Tag** area on the **Configuration** tab page of the host group details page. If a custom tag whose tag name is **_ha_enabled** and tag value is **False** exists, the HA function of the host group is disabled. If the tag does not exist or its value is **True**, the HA function of the host group is enabled.
You are advised not to enable this function for the management host group. Otherwise, services may be affected.
- Resources need to be reserved for HA. Otherwise, the ECS HA function may fail. To ensure that the ECS HA function works properly, you need to clear host exceptions and resource insufficiency alarms in a timely manner.

Procedure

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
 - Step 2** Locate the row that contains the target ECS, and choose **More > Change Settings > HA** in the **Operation** column.
 - Step 3** Change the HA status of the ECS.
 - Toggle the HA status to **ON**, and click **OK**.
The HA status will change to **Enable**.
 - Toggle the HA status to **OFF**, and click **OK**.
The HA status will change to **Disable**.
- End

7.8 ECS Snapshot

You can create an ECS snapshot that includes the system disk and data disks of an ECS whose virtualization type is KVM at a specific point in time. (Only EVS disks support ECS snapshot. Local disk and USB devices do not support this feature.) This snapshot records the disk data, quantity, and mount points at the time when the snapshot was taken. When necessary, you can use the ECS snapshot to roll

back the ECS. After you use an ECS snapshot to roll back the ECS, the system and data disks of the ECS roll back to their statuses at the time when the snapshot was taken, including disk data, quantity, and mount points. Any changes on the system or data disks later than the snapshot will be canceled. For example, new disks will be deleted, and detached disks will be attached again to their original mount points.

Notes

- Disk-intensive ECSs, intermediate-state ECSs, and faulty ECSs do not support ECS snapshots.
- ECSs with shared disks attached do not support ECS snapshots. If the disks (system and data disks) of an ECS are not in the same storage array or non-Huawei storage (Huawei Distributed Block Storage, OceanStor V3/V5 Series, and OceanStor Dorado V3 Series), the ECS does not support ECS snapshots. If the ECS from which the disk is being detached does not support ECS snapshots, try again after the disk is detached.
- USB passthrough ECSs support ECS snapshots, but USB devices do not support this function.
- If ECS snapshots have been created for an ECS whose flavor you want to change, you cannot change the flavor to one that does not support ECS snapshots, for example, disk-intensive ECS.
- An ECS configured with CSDR, CSHA, or VHA does not support rollback from a snapshot.
- If ECS snapshots have been created for an ECS, disks can be detached from the ECS. However, the detached disks cannot be attached to another ECS or deleted.
- Deleting an ECS will delete all snapshots created for it. Exercise caution when performing this operation.
- If ECS snapshots have been created for an ECS, the OS of the ECS cannot be reinstalled or changed.
- If resource quotas and fees have been configured on your platform, ensure that they are sufficient. Otherwise, you cannot create an ECS snapshot.
- The ECS snapshots automatically created by the system during full-ECS image creation cannot be rolled back or manually deleted. After the image is created, it is automatically deleted.

Procedure

Creating and deleting an ECS snapshot

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 2** Click the name of an ECS.
The ECS details page is displayed.
- Step 3** Click the **ECS Snapshot** tab.
- Step 4** Choose an operation as required.

- Creating an ECS snapshot
 - a. Click **Apply for Snapshot**.
The **ECS Snapshot** dialog box is displayed.
 - b. Enter a snapshot name and click **OK**.
When the snapshot that you have applied for is displayed in the snapshot list, the snapshot is successfully created. If the snapshot is not displayed in the snapshot list after a long time, go to the upper part of the ECS list, and click the fault icon next to the application status to view the failure cause.

 **NOTE**

- You can also locate the row that contains the target ECS in the ECS list, choose **More > Apply for Snapshot** in the **Operation** column.
 - Before creating a snapshot for an ECS, click the **EVS** tab to check whether all its data disks have been attached. If any of them has not been attached, creating an ECS from the snapshot will fail.
 - To view details about a child snapshot corresponding to the ECS snapshot, click the icon on the left of the ECS snapshot name. On the displayed details page, click the child snapshot name to access the child snapshot details page.
- Deleting an ECS snapshot
Click **Delete**, and in the displayed dialog box, click **OK**.

 **NOTE**

The ECS snapshots automatically created by the system during full-ECS image creation cannot be deleted. After the image is created, it is automatically deleted.

----End

Rolling Back an ECS from a Snapshot

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 2** Click the ECS name.
The ECS details page is displayed.
- Step 3** Click the **ECS Snapshot** tab, locate the row containing the target snapshot, and click **Rollback**.

NOTICE

- If the ECS is in the **Running** state, the system will shut down it during the rollback, which will interrupt your services. Exercise caution when performing this operation. After the rollback is complete, the ECS automatically starts.
- After you use an ECS snapshot to roll back the ECS, the system disk and data disks of the ECS roll back to their statuses at the time for creating the ECS snapshot, including disk data, quantity, and attachment points. Any changes on the system disk or data disks later than the ECS snapshot will be canceled. For example, new disks will be deleted, and detached disks will be attached again to their original attachment points.
- When creating a full-ECS image, the system automatically creates an ECS snapshot. You cannot roll back the ECS using the snapshot.

Step 4 Click **OK**.

----End

7.9 Creating a CD-ROM Drive and Attaching ISO/UVP VMTools

After creating a CD-ROM drive, you can perform the following operations:

- Attach the CD-ROM drive of the local PC to an ECS using the CD-ROM drive, and copy the content stored in the CD-ROM drive to the ECS whose virtualization type is KVM.
- Attach the local ISO file to an ECS using the CD-ROM drive, and install the plugin or software, for example, install UVP VMTools.
- Attach the UVP VMTools installation package using the CD-ROM drive to install or upgrade UVP VMTools.
 - In the HUAWEI CLOUD Stack 8.1.0 installation scenario, upload the UVP VMTools installation package. For details, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > How Do I Obtain and Upload the UVP VMTools Installation Package?**
 - If HUAWEI CLOUD Stack is upgraded from an earlier version to 8.1.0, the UVP VMTools installation package already exists on the physical host where the ECS is located.
- If UVP VMTools has been installed in the ECS image, UVP VMTools can be automatically upgraded using the CD-ROM drive. If the CD-ROM drive is not used and UVP VMTools has been installed in the image, UVP VMTools cannot be automatically upgraded. You need to manually upgrade it.

This section describes how to create a CD-ROM drive.

Context

To create a CD-ROM drive and attach the CD-ROM drive of the local PC, local ISO file, or UVP VMTools installation package, perform the following steps:

- Step 1** Shut down the ECS, and create a CD-ROM drive. If the ECS has a CD-ROM drive, go to the next step.
- Step 2** If you want to attach the CD-ROM drive of the local PC or an ISO file to the CD-ROM drive, download and install Microsoft .NET Framework 4.5 to 4.8 from the Microsoft official website on the local PC. If the UVP VMTools installation package is mounted to the physical host where the ECS is located, you do not need to download the client plugin.
- Step 3** Attach the CD-ROM drive of the local PC, local ISO file, or UVP VMTools installation package.
- Step 4** For a Linux ECS, you also need to log in to the ECS to perform the mounting operation. This step is not required for a Windows ECS.

----End

For management VMs and all service VMs, you can perform necessary operations on Service OM. For details, see "Resource Management" > "Service OM Resource Management" > "Compute Resource Management" > "VM Configuration Management" > "Creating a CD-ROM Drive for a VM and Mounting an ISO File to the VM" in *HUAWEI CLOUD Stack 8.1.0 O&M Guide*. For service VMs, you can also perform the necessary operations on the ECS details page on ManageOne Operation Portal. This section provides details about these operations.

Notes

- The virtualization type of the ECS must be **KVM**.
- An ECS, whether it is running or has stopped, can have a local CD-ROM drive, ISO file, or UVP VMTools installation package attached. Before attaching them, you need to create a CD-ROM drive.
- A CD-ROM drive can be created only when no task is being executed on the ECS and the ECS is stopped. Only one CD-ROM drive can be created for an ECS.
- Before attaching the local CD-ROM drive or an ISO file, install Microsoft .NET Framework 4.5 to 4.8 on the local PC. Otherwise, the attaching function will be unavailable. If the antivirus software blocks the plugin installation, disable it.
- Do not attach the local CD-ROM drive or an ISO file during UVP VMTools upgrade.
- An ECS that has a CD-ROM drive created and is idle supports all of the following features: live migration, cold migration, VM HA, online flavor change, offline flavor change (that is, the ECS needs to restart during the flavor change), and creating ECS snapshots and private images.
- An ECS that has a CD-ROM drive created and is occupied does not support live or cold migration, offline flavor change, or host group change, but supports online flavor change and creation of private images. It supports creation of ECS snapshots, but the snapshots will not contain information of the attached local CD-ROM drive, ISO image file, and UVP VMTools installation package. HA is supported. However, after HA is enabled for an ECS, CD-ROM drives and ISO files may not still be mounted to the ECS.
- The total number of CD-ROM drives and ISO files attached to all ECSs on a physical server cannot exceed 32. The UVP VMTools installation package is

not affected. If the total number has reached 32, you need to detach some of the CD-ROM drives or ISO files.

- Before attaching a local CD-ROM drive or an ISO file, ensure that the local PC can communicate with the ManageOne platform and port 8887 of the local PC is not occupied.
- Local PC operating systems needed for client plugin installation (ensure that the local PC runs one of the following OSs):
 - Windows 7 32-bit or 64-bit
 - Windows 7 SP1 32-bit or 64-bit
 - Windows 8 32-bit or 64-bit
 - Windows 8.1 32-bit or 64-bit
 - Windows 10 32-bit or 64-bit
 - Windows Server 2008 R2 64-bit
 - Windows Server 2012 64-bit
 - Windows Server 2012 R2 64-bit
- UVP VMTools can be installed on all Linux ECSs, but only some Windows ECSs are supported. For details about the supported OSs, see the following:
 - Carrier users: Click [here](#), search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)**, and obtain the latest document. Search for **UVP Tools** to view the support status.
 - Enterprise users: Click [here](#), search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)**, and obtain the latest document. Search for **UVP Tools** to view the support status.

Creating a CD-ROM Drive

Step 1 Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 If the ECS is running, stop the ECS first.

Step 3 Create a CD-ROM drive in either of the following ways:

Method 1: On the **Elastic Cloud Server** page, locate the row that contains the target ECS, choose **More > CD-ROM Drive And ISO > Create CD-ROM Drive** in the **Operation** column, and then click **OK**.

Method 2: Click the name of the target ECS. On the displayed ECS details page, click **CD-ROM Drive** and then **Create CD-ROM Drive**, and click **OK**.

----End

Attaching the CD-ROM Drive of the Local PC or an ISO File

If you want to attach the CD-ROM drive of the local PC or an ISO file to the CD-ROM drive, download and install Microsoft .NET Framework 4.5 to 4.8 from the Microsoft official website on the local PC. Ensure that the local PC runs an OS listed in [Local PC operating systems needed for client plugin installation](#).

Step 1 Go to the attachment page in either of the following ways:

Method 1: On the **Elastic Cloud Server** page, locate the row that contains the target ECS, and choose **More > CD-ROM Drive And ISO > Attach ISO** in the **Operation** column.

Method 2: Click the name of the target ECS. On the displayed ECS details page, click **CD-ROM Drive** and then **Attach ISO**.

Step 2 If the system prompts you to install the client plugin, download the plugin to the local PC.

Step 3 Install the plugin according to the installation wizard.

Step 4 Open the task manager of the OS, select the **Process** tab, and check whether the **FSPortalClientPluginKvm.exe** process exists.

- If yes, the installation is successful.
- If no, perform the installation again.

Step 5 Go to the **Attach ISO** page again. If the **External Protocol Request** dialog box is displayed, click **Start**.

 **NOTE**

If the system displays a message indicating that no plugin is detected, rectify the fault by referring to [18.1.3 Common Methods to Troubleshoot Local Client Plugin Used for Attaching ISO](#).

Step 6 Select an object to be mounted.

- Select **CD-ROM(ISO)** and select the local CD-ROM drive to be attached from the drop-down list.
- Select **File(ISO)**, click **Browse**, and select the local ISO file to be attached.

Step 7 If the system displays a success message, the operation is successful.

Step 8 If the ECS is in the **Stopped** state, start the ECS and determine whether a manual attaching operation is required based on the ECS.

1. For a Windows ECS, no further action is required.
2. For a Linux ECS, you also need to log in to the ECS to attach the CD-ROM drive or ISO file. Take CentOS as an example. Perform the following steps:
 - a. Log in to the ECS. For details, see [6.1 Login Mode Overview](#).
 - b. Run the following command to switch to the **root** user:
sudo
 - c. Run the following command to query the name of the attached device:
ls -lh /dev
In the command output, **/dev/sr*** indicates the CD-ROM drive. Assume that the device name is **/dev/sr0**.
 - d. Check the mount point. If no mount point exists, run the following command to create one, for example, **/mnt/vmtools**:
mkdir -p /mnt/vmtools
 - e. Run the following command to attach the local CD-ROM drive or an ISO file:
mount /dev/sr0 /mnt/vmtools

----End

UVP VMTools

On the **CD-ROM Drive** tab page, in the **Installed version** column of the component table, **Latest version in current environment** is the latest UVP VMTools. If you want to install UVP VMTools of this version for an ECS or upgrade UVP VMTools to this version, attach the UVP VMTools installation package to the ECS and install or upgrade UVP VMTools.

NOTE

- During the installation or upgrade of UVP VMTools, you need to remotely log in to the ECS. By default, only the ECS creator can remotely log in to the ECS. Other users or administrators do not have the login permission, and therefore cannot install or upgrade UVP VMTools. To solve this issue, log in to a FusionSphere OpenStack controller node and modify the permission. For details, see .
- If UVP VMTools 2.5.0 or later has been installed on the ECS, it can be automatically upgraded with UVP VMTools on the controller node. You can also manually upgrade it. UVP VMTools earlier than 2.5.0 cannot be upgraded. Contact technical support.

Step 1 On the **CD-ROM Drive** tab page of the ECS details page, check whether the CD-ROM drive is idle. In the **Operation** column on the right, click **Attach UVP VMTools Package**.

The ISO attaching dialog box is displayed.

Step 2 Click **OK**. If a dialog box is displayed, prompting you to remotely log in to the ECS and install or upgrade UVP VMTools, the mounting is successful.

Step 3 Click **Remote Login**, and then install or upgrade UVP VMTools. You can also click **Install UVP VMTools** or **Upgrade UVP VMTools** in the **Operation** column of a component to install or upgrade UVP VMTools.

For details about how to install UVP VMTools, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > Installing or Upgrading UVP VMTools for an Existing ECS > Installing UVP VMTools**. For details about how to upgrade UVP VMTools, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > Installing or Upgrading UVP VMTools for an Existing ECS > Upgrading UVP VMTools**.

----End

Related Operations

Deleting a CD-ROM drive

Step 1 If the CD-ROM drive is occupied, detach ISO or UVP VMTools on the **CD-ROM Drive** tab page before deleting it. Ensure that the CD-ROM drive is idle.

Step 2 Before deleting the CD-ROM drive, stop the ECS first, if the ECS is not stopped already.

Step 3 On the **CD-ROM Drive** tab page of the ECS details page, click **Delete**. Alternatively, on the **Elastic Cloud Server** page, locate the row that contains the ECS, choose **More > CD-ROM Drive And ISO > Delete CD-ROM Drive** in the **Operation** column.

A confirmation dialog box is displayed.

 **NOTE**

- If UVP VMTools has been installed in the image used for creating the ECS, UVP VMTools can be automatically upgraded using the CD-ROM drive. After the CD-ROM drive is deleted, UVP VMTools will need to be upgraded manually.
- UVP VMTools collects internal monitoring metrics of ECSs to monitor their running status and supports communication between ECSs and physical hosts. UVP VMTools improves disk and network I/O performance for Windows ECSs, and reports alarms when Linux ECSs become faulty.

Step 4 Click **OK**.

----End

Detaching the local CD-ROM drive/ISO file/UVP VMTools installation package

On the **CD-ROM Drive** tab page of the ECS details page, click **Detach ISO** or **Detach UVP VMTools Package**.

8 Passwords and Key Pairs

8.1 Overview

The ECS password is very important. Keep the password secure. You can obtain, clear, and reset the ECS password manually or with one click.

- After obtaining the initial password of a Windows ECS, delete the password. For details, see [8.2 Deleting the Initial Password for Logging In to a Windows ECS](#).
- For an ECS on which the one-click password reset plugin has been installed, you can reset the password of the ECS with one click. For details, see [8.3 Resetting the ECS Password with One Click \(Windows and Linux\)](#).
- For an ECS on which the one-click password reset plugin has not been installed, you need to manually reset the password of the ECS. For details about how to reset the password of a Windows ECS, see [8.4 Manually Resetting the Password for Logging In to a Windows ECS](#). For details about how to reset the password of a Linux ECS, see [8.5 Manually Resetting the Password for Logging In to a Linux ECS](#).

[Table 8-1](#) shows the ECS password complexity requirements.

Table 8-1 Password complexity requirements

Parameter	Requirement	Example Value
Password	<ul style="list-style-type: none">• Consists of 8 to 26 characters.• Contains at least three of the following character types:<ul style="list-style-type: none">- Uppercase letters- Lowercase letters- Digits- Special characters: !@%-_=[{}];,./\$#&*?• Cannot contain the username or the username spelled backwards.• Cannot contain more than two characters in the same sequence as they appear in the username. (This requirement applies only to Windows ECSs.)	test12\$@

8.2 Deleting the Initial Password for Logging In to a Windows ECS

To ensure the security of a Windows ECS that uses a key pair for authentication, it is recommended that you clear the initial password for logging in to the ECS stored in the system after obtaining the initial password. Deleting the initial password does not affect ECS operation or login.

Prerequisites

Once deleted, the password cannot be retrieved. Record the password before deleting it.

Procedure

Step 1 Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 In the ECS list, select the ECS whose password you want to delete.

Step 3 In the **Operation** column, choose **More > Change Settings** and click **Delete Password**.

The system displays a message, asking you whether you want to delete the password.

Step 4 Click **Delete** to delete the password.

----End

8.3 Resetting the ECS Password with One Click (Windows and Linux)

When the password of an ECS on which the one-click password reset plugin has been installed is lost or expires, you can reset the password of the ECS with one click.

NOTE

If the plugin cannot automatically start upon ECS startup, rectify the fault by referring to [18.8.5 What Should I Do If the One-Click Password Reset Plugin Fails to Start?](#).

Constraints

- Log in to DMK and check whether the value of `is_supported_reset_password` in ECS UI is **true**. If the value is **false**, the password resetting entry is shielded. For details about how to change the value, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > Modifying the Password Reset Configuration Item**.
- Only the ECS creator can change the passwords using the one-click password reset plugin by default. To allow a VDC administrator to change the passwords of all ECSs in the VDC to which the VDC administrator belongs and its lower-level VDCs, log in to a FusionSphere OpenStack management node and modify the permission. For details, see "Product Management > Resource Pools" > "FusionSphere OpenStack" > "Compute" > "Modifying the Permission of a VDC Administrator to Reset the Password of an ECS" in *HUAWEI CLOUD Stack 8.1.0 O&M Guide*.
- If you are not the actual user of this ECS when resetting the password, exercise caution when deciding to perform this operation, avoiding impact on the ECS user.
- If your account corresponds to a customized role, the role must have the permission to reset the password.
- You have installed the one-click password reset plugin before the password of your ECS is lost or expires. For details about how to check whether this plugin has been installed and how to install the plugin, see [5.4.1 Overview](#).
- If the power status of an ECS is stopped during ECS creation, a newly obtained ECS stays in the **Stopped** state. In this case, the one-click password reset plugin does not take effect. You need to start the ECS and then reset the password.
- ECSs created using SUSE 11 SP4 must have 4 GB or a larger memory.

Procedure

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 2** Ensure that the target ECS is stopped.
- Step 3** Locate the row that contains the target ECS, and in the **Operation** column, choose **More > Change Settings > Reset Password**.

Step 4 Enter a new password for the ECS and confirm the new password.

The complexity of the new password must meet the requirements listed in [Table 8-1](#).

Step 5 Click **OK**.

 **NOTE**

A new password takes effect after the ECS is restarted.

----End

8.4 Manually Resetting the Password for Logging In to a Windows ECS

If you need to change the initial password and reset a lost or expired password, see this section.

Resetting the Password After the Initial Connection

Step 1 Log in to the Windows ECS remotely. For details, see [6.1 Login Mode Overview](#).

Step 2 Press **Win+R** to start the **Open** dialog box.

Step 3 Enter **cmd**, and click **OK**.

The command line interface (CLI) is displayed.

Step 4 Run the following command to change the password (the complexity of the new password must meet the requirements listed in [Table 8-1](#)):

```
net user Administrator New password
```

----End

Resetting a Lost or Expired Password

 **NOTE**

The password resetting method provided in this section is for reference only and may not apply to all scenarios. If the password resetting fails, contact the OS vendor to provide other methods.

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 elastic IP address to the temporary ECS and configured the apt-get source.
- Use either of the following methods to install **ntfs-3g** and **chntpw** software packages in the temporary ECS:

Method One:

Run the following command to install the **ntfs-3g** and **chntpw** software packages:

sudo apt-get install ntfs-3g chntpw

Method Two:

Download the corresponding **ntfs-3g** and **chntpw** software packages according to the temporary ECS OS.

Visit <https://www.tuxera.com/community/open-source-ntfs-3g> to obtain the **ntfs-3g** software package.

Visit <https://pkgs.org/download/chntpw> to obtain the **chntpw** software package.

Procedure

Step 1 Stop the original ECS, detach its system disk from, and attach the system disk to a temporary ECS.

1. Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
2. Stop the original Windows ECS, switch to the details page, and click the **EVS** tab.
3. Locate the row containing the system disk and click **Detach** to detach the system disk from the ECS.
4. On the page showing the details of the temporary ECS, click the **EVS** tab.
5. Click **Attach Disk**. In the displayed dialog box, select the system disk detached in [Step 1.3](#) and attach it to the temporary ECS.

Step 2 Log in to the temporary ECS remotely and attach the system disk.

1. Run the following command to view the directory of the system disk detached from the original Windows ECS and now attached to the temporary ECS:

```
fdisk -l
```

2. Run the following command to attach the file system of the detached system disk to the temporary ECS:

```
mount -t ntfs-3g /dev/Query result obtained in Step 2.1 /mnt/
```

For example, if the query result obtained in [Step 2.1](#) is **xvde2**, run the following command:

```
mount -t ntfs-3g /dev/xvde2 /mnt/
```

Step 3 Change the password and clear the original password.

1. Run the following command to back up the SAM file:

```
cp /mnt/Windows/System32/config/SAM /mnt/Windows/System32/config/SAM.bak
```
2. Run the following command to change the password of a specified user:

```
chntpw -u Administrator /mnt/Windows/System32/config/SAM
```
3. Enter **1** and **y** as prompted, and press **Enter**

The password has been reset if the following information is displayed:

```
Select: [q] > 1  
Password cleared!
```

```
Hives that have changed:  
#Name  
0<SAM>
```

```
Write hive files? (y/n) [n] : y
0<SAM> - OK
```

Step 4 Stop the temporary ECS, detach the system disk of the original Windows ECS, and attach the system disk to the original Windows ECS.

1. Stop the temporary ECS, switch to the details page, and click the **EVS** tab.
2. Click **Detach** to detach the data disk temporarily attached in [Step 1.5](#).
3. On the page showing the details of the original Windows ECS, click the **EVS** tab.
4. Click **Attach Disk**. In the displayed dialog box, select the data disk detached in [Step 4.2](#) and mount point `/dev/sda`.

Step 5 Start the original Windows ECS and set a new login password.

1. Click **Start** to start the original Windows ECS. After the status becomes **Running**, click **Remote Login** in the **Operation** column.
2. Click **Start**. Enter **CMD** in the search box and press **Enter**.
3. Run the following command to change the password (the new password must meet the requirements listed in [Table 8-1](#)):

```
net user Administrator New password
```

----End

8.5 Manually Resetting the Password for Logging In to a Linux ECS

Logging In to an ECS and Changing the Password of the root User

Step 1 Use the existing key file to log in to the Linux ECS as user **root** in key pair authentication mode.

Step 2 Run the following command to reset the password of user **root**:

```
passwd
```

Replace **passwd** with **passwd username** if you want to reset the passwords of other users.

Step 3 Enter the new password as prompted. The new password must meet the requirements listed in [Table 8-1](#).

```
New password:
Retype new password:
```

The password has been reset if the following information is displayed:
passwd: password updated successfully

----End

Resetting a Lost or Expired Password

NOTE

The password resetting method provided in this section is for reference only and may not apply to all scenarios. If the password resetting fails, contact the OS vendor to provide other methods.

Prerequisites

- Log in to DMK and check whether the value of **is_supported_reset_password** in ECS UI is **true**. If the value is **false**, the password resetting entry is shielded. For details about how to change the value, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > Modifying the Password Reset Configuration Item**.
- You have prepared a temporary Linux ECS that resides in the same AZ as the target ECS.
- An EIP has been bound to the temporary ECS.

Procedure

Step 1 Download the script used to reset the password and upload the script to the temporary ECS.

1. Log in to ManageOne, and choose **Computing > Elastic Cloud Server**.
2. Locate the row that contains the ECS (the original Linux ECS) whose password needs to be reset, click **More > Reset Password** in the **Operation** column. Download the password resetting script as instructed.
3. Use a remote connection tool, such as WinSCP, to upload the **changepasswd.sh** script obtained from [Step 1.2](#) to the temporary ECS. To download WinSCP, visit <http://winscp.net/>.

Step 2 Stop the original Linux ECS, detach the system disk, and attach the system disk to the temporary ECS.

1. Stop the original Linux ECS, switch to the details page, and click the **EVS** tab.

NOTE

Do not forcibly stop the original ECS. Otherwise, the password resetting may fail.

2. Locate the row containing the system disk to be detached and click **Detach** to detach the system disk from the ECS.
3. On the page showing the details of the temporary ECS, click the **EVS** tab.
4. Click **Attach Disk**. In the displayed dialog box, select the system disk detached in [Step 2.2](#) and attach it to the temporary ECS.

Step 3 Remotely log in to the temporary ECS, upload the script to it, and reset the password.

1. Locate the row containing the temporary ECS and click **Remote Login** in the **Operation** column.
2. 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
3. Run the following commands in the directory where the 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. For CentOS 7.2, you can run the **yum install lvm2** command to install **lvm2**.

NOTICE

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.

4. Enter the new password and the directory obtained in [Step 3.2](#) as prompted.
The password is changed if the following information is displayed:
set password success.

Step 4 Stop the temporary ECS, detach the system disk of the original Linux ECS, attach the system disk to the original Linux ECS, and restart the ECS.

1. Stop the temporary ECS, switch to the details page, and click the **EVS** tab.
2. Click **Detach** to detach the data disk temporarily attached in [Step 2.4](#).
3. On the page showing the details of the original Linux ECS, click the **EVS** tab.
4. Click **Attach Disk**. In the displayed dialog box, select the data disk detached in [Step 4.2](#) and mount point **/dev/sda**.
5. Restart the original Linux ECS.

----End

8.6 Creating a Key Pair

If you plan to use a key pair on your created ECS for login authentication, you need to create a key pair or import an existing private key. For details, see this section.

Context

To ensure system security, you are recommended to use the key pair authentication mode to authorize the user who attempts to log in to an ECS. Therefore, you must use an existing key pair or create a new one for remote login authentication.

- Creating a key pair

If no key pair is available, create one. You can use either of the following methods:

- Create a key pair using UI. After the creation, the public key is automatically stored in the system, and the private key is manually stored in a local directory. For details, see [Step 1](#) to [Step 6](#). This method is recommended.

- Create a key pair using **puttygen.exe**. After the creation, both the public key and private key are stored locally. For details, see [Step 1](#) to [Step 3](#).
- Using an existing key pair
If a key pair exists locally, import the key pair on ManageOne Operation Portal. For details, see [Step 1](#) to [Step 4](#).

 **NOTE**

If the public key of the existing key pair is stored by clicking **Save public key** of **puttygen.exe**, the public key cannot be imported. If this key pair must be used for remote authentication, see [18.6.4 What Should I Do If a Public Key Fails to Be Imported to ManageOne After a Key Pair Is Created Using PuTTYgen?](#).

Procedure

(Recommended) Using UI to creating a key pair

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 2** In the navigation pane, choose **Key Pair**.
- Step 3** In the right pane of the page, click **Create Key Pair**.
- Step 4** Enter the key name.
- Step 5** Click **OK**.
- Step 6** In the displayed dialog box, click **OK**.

 **NOTE**

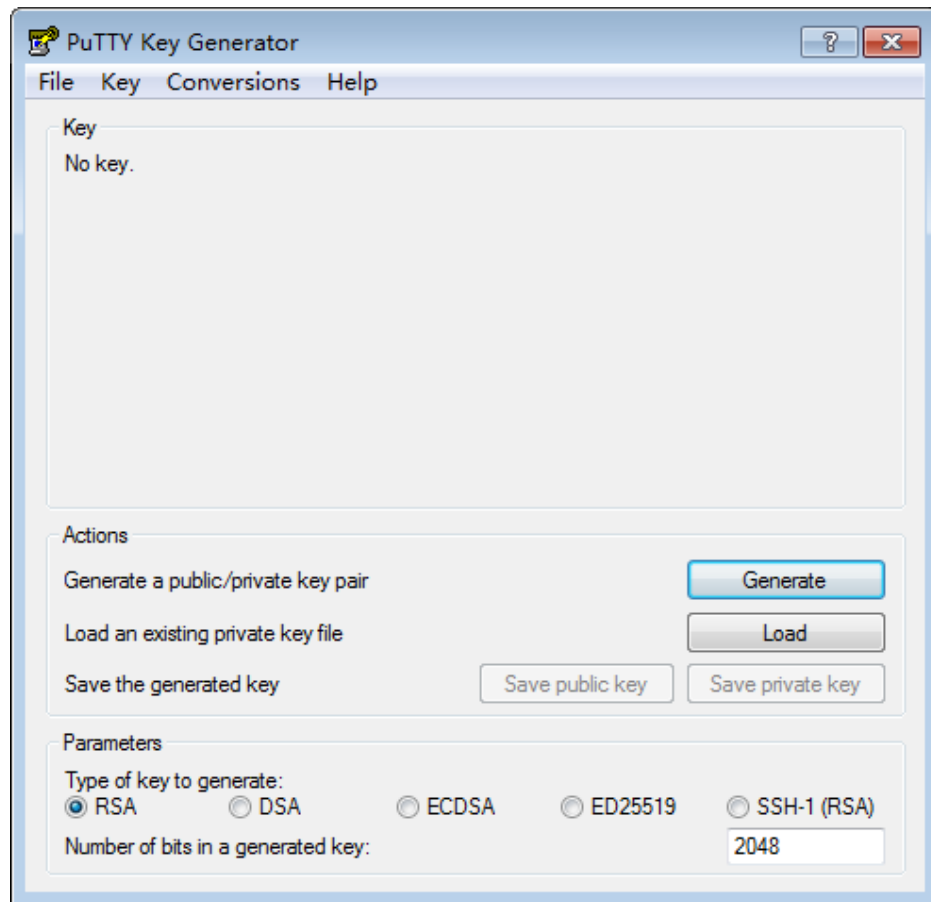
You can view and save the private key according to the prompts. To ensure ECS security, you are limited to downloading the private key only once.

----End

Using a tool to create a key pair

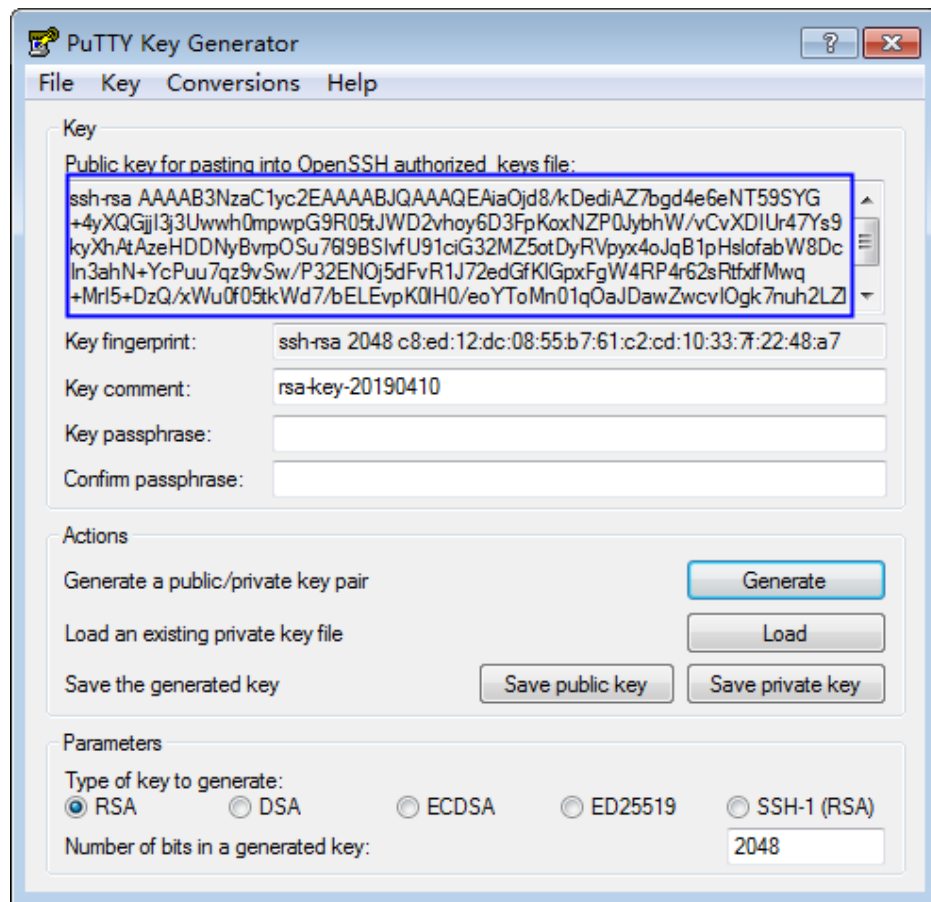
- Step 1** Obtain the public and private keys.
 1. Double-click **puttygen.exe**. The **PuTTY Key Generator** page is displayed. The following uses 0.70 (64-bit) as an example.

Figure 8-1 PuTTY Key Generator



2. In the **Parameters** area, set **Type of key to generate** to **RSA** and **Number of bits in a generated key** to **1024** (not recommended), **2048**, or **4096**.
3. Click **Generate**, and move the pointer randomly over the blank area. The key generator automatically generates a key pair that consists of a public key and a private key. The public key is shown in the blue box in [Figure 8-2](#).

Figure 8-2 Obtaining the public and private keys



Step 2 Copy the public key content to a .txt file and save the file in a local directory.

NOTE

Do not save the public key by clicking **Save public key**. Storing a public key by clicking **Save public key** of **puttygen.exe** will change the format of the public key content. Such a key cannot be imported.

Step 3 Save the private key.

When you are required to log in to a Linux ECS using PuTTY, you must use the .ppk private key. To save the private key in .ppk format, perform the following operations:

1. On the **PuTTY Key Generator** page, choose **File > Save private key**.
2. Save the converted private key, for example, **kp-123.ppk**, in a local directory.

----End

Importing the public key

Step 1 Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 In the navigation pane, choose **Key Pair**.

Step 3 On the **Key Pair** page, click **Import Public Key**.

Step 4 Use either of the following methods to import the public key:

 **NOTE**

The public key contains a maximum of 1024 characters, and the encryption mode is RSA 1024 (not recommended), RSA 2048, or RSA 4096. Otherwise, the import fails.

- Selecting a file
 - a. Select the locally stored public key.

 **NOTE**

When importing a public key, ensure that the public key is imported. Otherwise, importing the public key will fail.

- b. Click **OK**.
- Copying the public key content
 - a. Copy the content of the public key in .txt file into the **Public Key Content** text box.
 - b. Click **OK**.

After the public key is imported, you can change its name.

----End

Helpful Links

- [18.6.3 What Should I Do If a Public Key Cannot Be Imported?](#)
- [18.6.4 What Should I Do If a Public Key Fails to Be Imported to ManageOne After a Key Pair Is Created Using PuTTYgen?](#)

9 ECS Flavors

9.1 Changing the Flavor of an ECS

This section describes how to modify an ECS flavor. If your ECS flavor cannot meet service requirements, modify it, including the number of vCPUs and memory size. If an ECS flavor degrades, the performance of ECSs using the flavor will be adversely affected. You can modify an ECS flavor online or offline.

Online Flavor Change

Online flavor change allows you to change the flavor of an ECS without interrupting services running on the ECS. Online flavor change must meet the constraints listed in [Table 9-1](#). Otherwise, only offline flavor change can be performed. During online flavor change, the flavor that meets all conditions is displayed on the page. You can use any of them.

Table 9-1 Constraints on online flavor change

Category	Description
Scope	You can add CPUs, increase the memory size, and modify CPU QoS based on the selected flavor. However, the ECS type cannot be changed. CPU QoS specifies the relative importance of ECSs and defines the upper limit and lower limit of the clock speed of the physical CPU occupied by ECSs during resource contention.
Permission	You must have the permission to change the ECS flavor online.
Virtualization type and physical server type	<ul style="list-style-type: none">• Currently, only ECSs whose virtualization type is KVM support online flavor change.• x86 servers support online flavor change, but Arm servers do not support it.

Category	Description
Global online change switch	To allow online flavor change, log in to the FusionSphere OpenStack web client, choose Configuration > OpenStack > Nova , and set Online flavor change to On . After it is changed from OFF to On , you need to forcibly restart the ECS to support online flavor change.
ECS type	<ul style="list-style-type: none"> • Ultra-high I/O ECSs, general computing-plus ECSs (x86), general computing-plus ECSs (Arm), AI-accelerated ECSs, and GPU-accelerated ECSs do not support online flavor change. • ECSs in MRS clusters do not support online flavor change.
Image	<ul style="list-style-type: none"> • The ECS OS must support both Hot-adding vCPUs and Hot-adding memory. If the OS table contains the Maximum Number of sockets column, the maximum socket specifications must be greater than or equal to 32. For details, see <i>FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)</i>. <ul style="list-style-type: none"> – For carrier users, click here. – For enterprise users, click here. • UVP VMTools must have been installed in the ECS image. If UVP VMTools has not been installed in the image, manually install UVP VMTools before you can perform online flavor change. For details about how to check whether UVP VMTools is installed on an ECS, see Related Operations. For details about how to install UVP VMTools, visit Operation Help Center and choose Operation > Compute Services > Elastic Cloud Server (ECS) > Installing or Upgrading UVP VMTools for an Existing ECS. • If the ECS is managed by the VMware platform to HUAWEI CLOUD Stack, you need to modify the image attributes before online flavor change. Otherwise, the ECS flavor cannot be changed. For details, click here.

Category	Description
vCPU	<ul style="list-style-type: none">• The maximum number of vCPUs of the ECS cannot exceed the maximum number supported by the OS. For details, see <i>FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)</i>.• If the current ECS flavor is A and the target flavor to change to is B, the difference between the quantities of vCPUs contained in these two flavors must meet certain requirements. For details, see Related Operations. Only flavors that meet these requirements are displayed on the portal. To create a flavor, contact the administrator to create one on Service OM. For details about how to create a flavor, see visit Operation Help Center and choose Operation > Compute Services > Elastic Cloud Server (ECS) > Configuration Before ECS Creation > Creating a Flavor. For details about how to set vCPUs, see Related Operations.• The number of vCPUs in the selected flavor cannot exceed the total number of vCPUs available on the physical host. Otherwise, the flavor may fail to be changed. For details about how to view the quantity of vCPUs available on a physical host, visit Operation Help Center and choose Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > How Do I Check the Number of vCPUs Available on a Physical Host?
Memory/ Hugepage memory	<ul style="list-style-type: none">• The initial ECS memory cannot be less than 4 GB. After flavor change, the memory must be less than 8 times the original memory and cannot exceed the maximum memory capacity supported by the OS. For details, see <i>FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)</i>.• For ECSs using hugepage memory, the memory increase must be multiples of the hugepage memory.
Number of online flavor changes	<p>The maximum number of online flavor changes allowed for ECSs varies depending on the OS, initial number of vCPUs, memory size, and network attributes.</p> <p>Each time the flavor is changed online, the number of allowed changes displayed on the GUI is decreased by one. After an ECS is provisioned, the number of allowed changes is close to the upper limit. After the ECS is forcibly restarted, the number of allowed changes is restored to the upper limit. The result displayed on the GUI prevails.</p>

Changing the ECS Flavor Offline

Changing the ECS flavor offline means that the ECS will need to restart for the new flavor to take effect, which results in service interruptions. When changing

the ECS flavor offline, you can add or remove CPUs and increase or decrease the memory size. In addition, the ECS type can be changed.

 **NOTE**

- ECSs configured with CSHA, CSDR, or VHA+CSDR protection do not support flavor change.
- ECSs in MRS clusters do not support offline flavor change.
- If the ECS is managed by the VMware platform to HUAWEI CLOUD Stack, you need to modify the image attributes before offline flavor change. Otherwise, the ECS flavor cannot be changed. For details, click [here](#).
- The number of vCPUs in the selected flavor cannot exceed the total number of vCPUs available on the physical host. Otherwise, the flavor may fail to be changed. For details about how to view the quantity of vCPUs available on a physical host, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > How Do I Check the Number of vCPUs Available on a Physical Host?**
- When changing the flavor, ensure that the selected flavor meets the following requirements. Otherwise, an error message will be displayed after the application is submitted, and the ECS flavor cannot be changed.
 - For an ECS whose **Boot Device** is set to **Cloud Disk**, the memory size of the selected flavor must be greater than or equal to the minimum memory size set during image registration.
 - For an ECS whose **Boot Device** is set to **Local Disk**, the **Memory** and **Root Disk (GB)** of the selected flavor must be greater than or equal to the minimum memory size and minimum disk set during image registration.

Procedure

Step 1 Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 Locate the row that contains the target ECS. In the **Operation** column, choose **More > Change Settings > Change Flavor**.

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

 **NOTE**

If the ECS is in the **Stopped** status, perform [Step 4](#) to [Step 5](#). If the ECS is in the **Running** status, perform [Step 3](#) to [Step 5](#).

Step 3 Set **Change Method**.

- Online change means that the ECS flavor will be changed without the need to restart the ECS.
- Offline change means that the ECS will automatically restart during the flavor change, which leads to service interruptions.

 **NOTE**

If this setting is not configurable, the conditions for online flavor change are not met, and the ECS supports offline flavor change only. For details, see [Online Flavor Change](#).

Step 4 Select the new flavor.

 **NOTE**

In x86 scenarios, all available flavors are x86 flavors. In Arm scenarios, all available flavors are Arm flavors.

Step 5 Click **OK**.

 **NOTE**

For an ECS booted from a local disk and configured with data disks, if the ECS is migrated to a new host after the flavor is changed but the ECS is in the ERROR state because it fails to connect to the data disks, run the expansion command to forcibly roll back the migration process. For details, see .

----End

Related Operations

Deciding on the Quantity of vCPUs for the New Flavor During Online Flavor Change

Step 1 On the ECS details page, view the ECS flavor.

Step 2 Log in to Service OM as an O&M administrator and go to the **Elastic Cloud Server** page. For details, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > How Do I Log In to Service OM?**

Step 3 Click **Flavors**.

The flavor details page is displayed.

Step 4 Locate the target flavor in the flavor list based on the flavor information obtained in [Step 1](#) and click the flavor name.

Step 5 Click **Configuration**.

Step 6 Check whether the **Custom Tag** area contains the tags **hw:cpu_cores** and **hw:cpu_threads**.

- If yes, the number of vCPUs to be added must be an integral multiple of the product of the two tag values.
- If no, the number of vCPUs to be added must be an integer such as 1, 2, 3...

NOTICE

- The recommended value of **hw:cpu_threads** is 1 or 2.
- The product of **hw:cpu_cores** and **hw:cpu_threads** must be the nth power of 2, for example, 2, 4, 8, or 16.
- The number of vCPUs before the flavor change must be an integer multiple of the product of **hw:cpu_cores** and **hw:cpu_threads**.

For example, if the ECS currently has 4 vCPUs, and the values of **hw:cpu_cores** and **hw:cpu_threads** are 1 and 2, respectively, then the number of vCPUs to be added must be an integral multiple of 2. The total number of vCPUs after the

flavor change is as follows: 4 + 2, 4 + 2 x 2, 4 + 2 x 3, and so on. If these two tags do not exist, the number of vCPUs after flavor change is 5, 6, 7, and so on.

----End

Checking Whether UVP VMTools Is Installed in a Windows OS

Step 1 Log in to the ECS. For details, see [6.1 Login Mode Overview](#).

Step 2 Perform either of the following depending on the bit version of your OS:

- For a 32-bit OS, check whether `C:\Program Files\virtio\` exists. If yes, UVP VMTools has been installed. Otherwise, VMTools is not installed.
- For a 64-bit OS, check whether `C:\Program Files (x86)\virtio\` exists. If yes, UVP VMTools has been installed. Otherwise, VMTools is not installed.

Step 3 Check whether `vm-agent-daemon` and `VMTools Daemon Service` can be found in the Task Manager.

- If yes, UVP VMTools has started.
- If no, UVP VMTools has not started. To use UVP VMTools, start the program first.

----End

Checking Whether UVP VMTools Is Installed in a Linux OS

Step 1 Log in to the ECS. For details, see [6.1 Login Mode Overview](#).

Step 2 Run the following command to check whether the `/etc/vmtools/` directory exists. If the command is executed successfully, UVP VMTools has been installed. If a message is displayed indicating that the directory does not exist, UVP VMTools is not installed.

```
cd /etc/vmtools/
```

Step 3 Switch to the `root` user and run `service vm-agent status`. If the following information is displayed, UVP VMTools has started.

```
[root@localhost ~]# service vm-agent status
server (pid 1584 976) is running ...
[ OK ]
```

----End

10 EVS Disk

10.1 Applying for a Data Disk

A data disk can be created together with an instance. During the instance creation, you can configure parameters such as the disk type, capacity, and sharing attribute. After the instance is created, the disk will be automatically attached to the instance. This section describes how to apply for and add a blank data disk to an instance separately.

Context

ECSs do not support the merging of EVS disk spaces. Each EVS disk is independent, and the spaces of multiple EVS disks cannot be merged through formatting. You are advised to plan the number and capacity of EVS disks before disk creation. It is not recommended that logical volumes managed by the LVM be created on the disks, because snapshots are created for independent EVS disks and creating such logical volumes will generate differential data after a snapshot rollback.

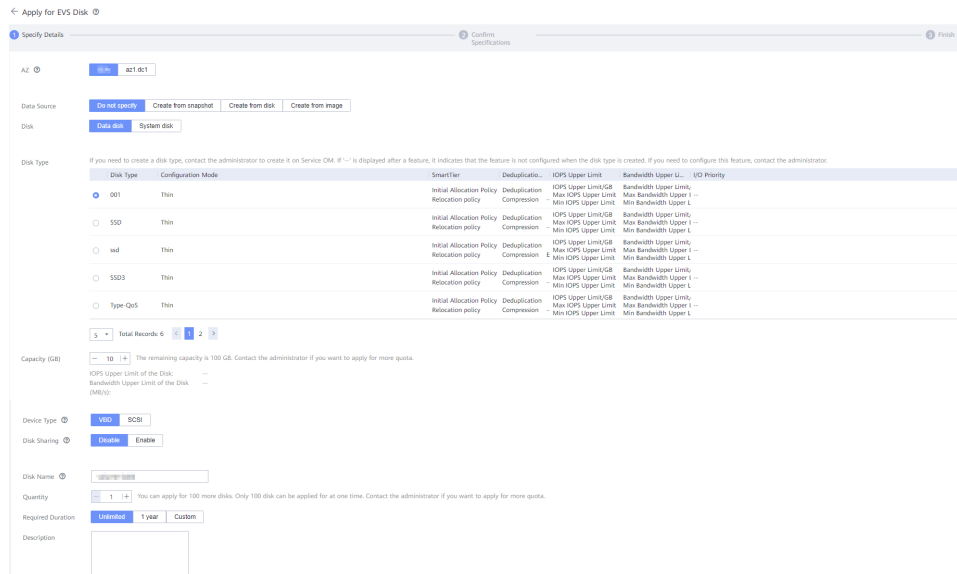
Procedure

- Step 1** Log in to the EVS console as a VDC administrator or VDC operator. For details, see [18.7.1 Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).
- Step 2** Click **Apply for EVS Disk**.
- Step 3** In the **Select Service** dialog box, select the target service and then click **Apply Now**.

NOTE

Services are created by administrators based on operation requirements. When creating a service, the administrator can lock service parameters (for example, apply for a fixed disk type or apply for resources in a specified AZ), specify the service publishing scope (for example, visible to only specific VDCs), and set the product approval process. For details about how to create more services, see "Creating a Common Service" in *HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide*.

- Step 4** In the **Apply for EVS Disk** dialog box, set the parameters as prompted and based on [Table 10-1](#).



NOTE

If the service you selected has parameters **AZ**, **Capacity (GB)**, **Disk Sharing**, or **Disk Type** configured, the parameter values configured for the service will be displayed for the EVS disk you apply for.

Table 10-1 Parameters required when applying for an EVS disk

Parameter	Description	Example Value
AZ	Specifies the availability zone (AZ) where an EVS disk is to be created. NOTE <ul style="list-style-type: none"> EVS disks can be attached to instances only in the same AZ. If an AZ is bound to a tenant upon the tenant creation, the resource pools in this AZ are used as dedicated storage pools for this tenant and invisible to other tenants not bound to the AZ. Services of this tenant can run on dedicated physical devices without interference from other tenants. 	az1.dc1
Data Source	Select Do not specify to apply for a blank EVS disk that does not contain data.	Do not specify
Disk	Specifies whether to apply for a data disk or system disk. This section describes how to apply for a data disk.	Data disk

Parameter	Description	Example Value
Disk Type	Select a disk type. You can select the disk type created in "Storage Services" > "Elastic Volume Service (EVS for ECS)" > "Configuration Before Applying for an EVS Disk" > "(Optional) Creating a Disk Type" in <i>HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide</i> . You can also select the disk type created during automatic installation and deployment using HUAWEI CLOUD Stack Deploy. The name of the automatically created disk type is the value of business_volume_type in deployment parameter summary file <i>xxx_export_all_v2_EN.xlsx</i> .	-
Capacity (GB)	EVS disk capacity. The EVS disk capacity can neither exceed the total capacity quota of EVS disks nor the capacity quota of the current disk type.	10 GB
Device Type	<ul style="list-style-type: none">• VBD indicates a VBD EVS disk.• SCSI indicates a SCSI EVS disk. SCSI-type disks allow the ECS OS to directly access the underlying storage media and support SCSI commands more advanced than VBD-type disks. For details about how to use SCSI disks, see Device Type. NOTE Device Type can be set only to VBD for the selected AZ using a heterogeneous storage device.	VBD
Disk Sharing	<ul style="list-style-type: none">• Disable indicates that the new disk will be a non-shared EVS disk.• Enable indicates that the new disk will be a shared EVS disk. Such a disk can be attached to multiple ECSs. NOTE If the backend storage device is a heterogeneous storage device, no shared EVS disk can be created.	Enable

Parameter	Description	Example Value
Disk Name	<p>The disk name can contain only letters, digits, underscores (_), and hyphens (-). When applying for a single EVS disk, ensure that the disk name contains less than or equal to 63 characters. When applying for EVS disks in batches, ensure that the disk name contains less than or equal to 58 characters.</p> <ul style="list-style-type: none"> If you apply for a single EVS disk, the value of this parameter is used as the name of the EVS disk. If you apply for multiple EVS disks in batches, the value of this parameter is used as the prefix of the names of the EVS disks. The name of each EVS disk resembles <i>Disk name-A four-digit number</i>. <p>NOTE For example, if you apply for two EVS disks and set Disk Name to volume, the names of the two EVS disks will be volume-0001 and volume-0002.</p>	volume-0001
Quantity	<p>Specifies the number of EVS disks that you apply for. The default value is 1, which means that you apply for one EVS disk.</p> <p>By default, a maximum of 100 EVS disks can be created at a time. The number of EVS disks that can be applied for in a batch varies with the current EVS disk quantity quota.</p>	1
Required Duration	<p>Specifies the validity period of the EVS disk that you apply for.</p> <ul style="list-style-type: none"> If you select Unlimited, the new EVS disk has no expiration date. If you select 1 year, the validity period of the new EVS disk will be one year, which is subject to the expiration date displayed on the console. If you select Custom, specify an expiration date for the EVS disk that you apply for. 	Unlimited
Description	<p>Describes the EVS disk that you apply for.</p> <p>The length cannot exceed 63 characters.</p>	-

Step 5 Click **Next**.

Step 6 Confirm that the information is correct and click **Add to Cart** or **Apply Now**.

If the configuration is incorrect, click **Back**.

- If you click **Add to Cart**, go to [Step 7](#).
- If you click **Apply Now**, go to [Step 8](#).

Step 7 Submit an application for the service in the shopping cart.

1. Click the shopping cart in the upper right corner of the page.
2. Select the service and click **Apply Now**.
3. Enter the order information and click **OK** to submit the application.

Step 8 If the application for an EVS disk requires approval, contact the administrator for approval. Otherwise, skip this step.

 **NOTE**

The application for an EVS disk can be rejected by the administrator. If you have entered incorrect configuration when applying for an EVS disk, you can contact the administrator to reject the application, correct the configuration, and submit the application again.

Step 9 On the **Elastic Volume Service** page, view the status of the EVS disk. After the EVS disk is created and its status changes to **Available**, the EVS disk is successfully created.

----End

Follow-up Procedure

If you want to use the new EVS disk, [attach it to an ECS](#) first.

10.2 Attaching an EVS Disk

A created EVS disk can be used by an instance only after being attached to the instance. The EVS disk created together with an instance is automatically attached, requiring no manual attaching operations.

Restrictions

- The ECS supports the attaching of disks in VBD and SCSI modes.
- Regardless if a shared EVS disk or non-shared EVS disk is attached to an instance, the EVS disk and the instance must be in the same AZ.
- Data disks can only be attached to ECSs as data disks. System disks can be attached to ECSs as system disks or data disks.
- An EVS disk cannot be attached to an instance that has expired.
- An EVS disk cannot be attached to an instance that has been soft deleted.
- When a disk is attached to an ECS configured with the disaster recovery (DR) service (CSDR/CSHA/VHA), you must ensure that the disk is created using the same backend storage as the existing disk on the ECS.
- An EVS disk with snapshots of a VM can be attached only to the VM and cannot be attached to any other VM.
- Neither shared EVS disks nor SCSI EVS disks can be attached to an ECS that has the CSHA service configured.
- If the ECS uses the Windows operating system and the administrator set **Disk Device Type** to **ide** when registering the image, shut down the ECS before attaching the EVS disk to the ECS.

- If the ECS to which the EVS disk belongs has not been created, the EVS disk cannot be attached to another ECS.

Context

- If you want to attach a SCSI EVS disk to an ECS, check whether the ECS supports SCSI disks. For details, see [Requirements and Restrictions on Using SCSI EVS Disks](#).
- ECSs to which SCSI shared disks are attached must be selected from the same anti-affinity ECS group. For details about the reasons, see [SCSI Reservation](#).
- You can attach the disks in the following ways:
 - EVS console: To attach multiple EVS disks to different ECSs, perform related operations on the EVS console.
 - ECS console: To attach multiple EVS disks to one instance, perform related operations on the ECS console.

Operations on the ECS Console

Step 1 Log in to the EVS console as a VDC administrator or VDC operator. For details, see [18.7.1 Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 Try either of the following to go to the **Attach Disk** page:

Method 1: On the **Elastic Cloud Server** page, locate the row that contains the target ECS, click **More** in the **Operation** column, and select **Attach Disk**.

Method 2: Click the name of the target ECS. On the displayed ECS details page, choose **EVS > Attach Disk**.

Step 3 On the **Attach Disk** page, determine whether to specify an attachment point.

- If you select **No**, the system automatically assigns an address.
Select **No** and select the disk to be attached. A data disk can only be attached to an ECS as a data disk. A system disk can be attached to an ECS as a system disk or a data disk.
- If you select **Yes**, you need to specify the device address.
Select **Yes** and select the disk to be attached and the attachment point.

NOTE

- Make sure that the ECS has no ongoing disk attaching task. Otherwise, disk attaching may fail.
- The device addresses displayed on the page are classified into the following types based on **Device Address** of images used by the ECS: a-x:x:x, a:x, and xxx:xx:xx.x. If the disk is successfully attached and the ECS is running properly, you can view the details about the disk on the ECS details page, including **Device Name** and **Device Address**. **Device Name**: indicates the actual device name of the disk on the ECS. **Device Address** indicates the address allocated by the system to the disk. If VMTools is not installed on the ECS or the version of VMTools is too early, only the device address is displayed on the page. For details about how to query the attachment point of a disk on an ECS based on the device address, see [Operation Help Center > Elastic Cloud Server > User Guide > Best Practices > Manually Viewing the Disk Mount Point](#).

If you want to apply for a new EVS disk, click **Apply for Disk**. For details, see section [10.1 Applying for a Data Disk](#).


Step 4 Click **OK**.

 **NOTE**

If a disk is attached to a stopped ECS, the device name is not displayed on the page after the disk is attached. Instead, it is displayed several minutes after the ECS is started.

----End

Operations on the EVS Console

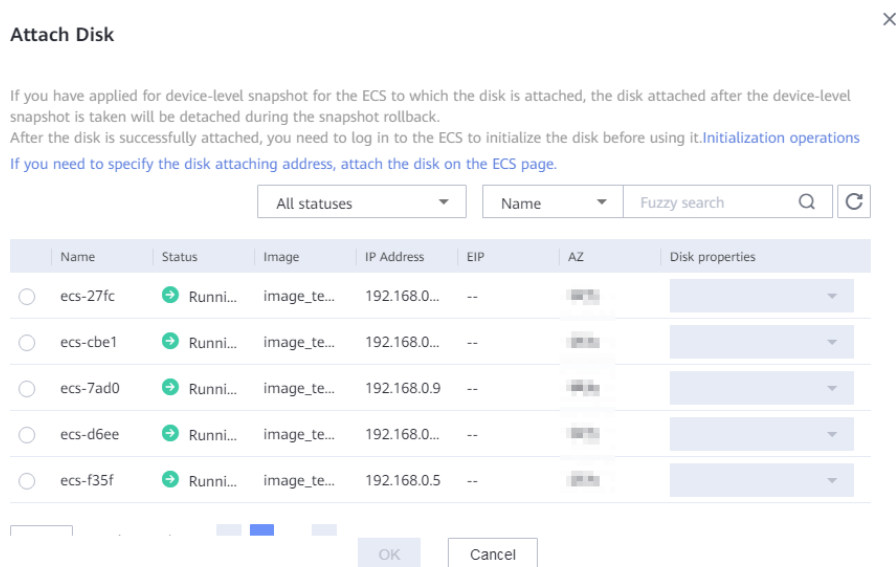
Step 1 Click  in the upper left corner, select a region and resource set, and choose **Storage > Elastic Volume Service**. The EVS console is displayed.

Step 2 On the **Elastic Volume Service** page, locate the row that contains the target EVS disk, and click **Attach** in the **Operation** column.

The disk can be either a shared disk or a non-shared disk.

Step 3 In the **Attach Disk** dialog box displayed, select the instance to which the disk is to be attached. A data disk can be attached to an instance only as a data disk. A system disk can be attached to an instance as a data disk or a system disk.

If you are attaching a shared disk, you can select multiple ECSs. A shared disk can be attached to a maximum of 16 ECSs by default.



 **NOTE**

To specify an address, you need to attach the disk on the ECS console. For details, see [Operations on the ECS Console](#).

Step 4 Click **OK**.

On the EVS disk list page, if the disk status is **In-use**, the EVS disk has been successfully attached to an instance.

----End

Follow-up Procedure

After the EVS disk is attached to the instance, perform subsequent operations by referring to [Table 10-2](#).

Table 10-2 Operations after the EVS disk is attached

EVS Disk	Instance OS	Follow-up Procedure
New blank data disk (Data Source is Do not specify)	Windows	To initialize the data disk of a Windows ECS, see Initializing a Windows Data Disk .
	Linux	<ul style="list-style-type: none"> When the data disk capacity is less than 2 TB, see Initializing a Linux Data Disk (fdisk). When the data disk capacity is greater than and equal to 2 TB, see Initializing a Linux Data Disk (parted).
New blank system disk (Data Source is Do not specify)	Windows/ Linux	The system disk cannot be used to directly start an ECS. If you need to use it to start the ECS, copy the backup data of a normal system disk to the system disk.
Data disk created from a snapshot, an image, a backup, or an existing disk	Windows	<p>Check whether the source disk corresponding to snapshots, images, existing disks, and backups has been initialized.</p> <ul style="list-style-type: none"> If yes, you can directly use the EVS disk without performing any other operations. If no, initialize the disk before using it. The operations are the same as those for a new blank data disk.

EVS Disk	Instance OS	Follow-up Procedure
	Linux	<p>Check whether the source disk corresponding to snapshots, images, existing disks, and backups has been initialized.</p> <ul style="list-style-type: none"> If yes, log in to the instance and run the mount command (mount <i>partition mount path</i>) to mount the partition. Configure automatic mounting upon system startup. Then, the EVS disk can be used properly. For details about how to set automatic mounting upon system startup, see Setting Automatic Disk Attachment Upon Instance Start. If no, initialize the disk before using it. The operations are the same as those for a new blank data disk.
System disk created from an image, a snapshot, an existing disk, or a backup	Windows/ Linux	You can directly use the system disk without performing any other operations.

10.3 Initializing a Data Disk

10.3.1 Initialization Overview

After attaching an EVS disk to an instance, you need to log in to the instance to partition and initialize the disk so that the data disk can be used for the instance. [Table 10-3](#) provides common partition styles.

Table 10-3 Disk partition style

Disk Partition Style	Maximum Disk Capacity Supported	Maximum Number of Partitions Supported	Partitioning Tool Supported
Main Boot Record (MBR)	2 TB	<ul style="list-style-type: none"> Four primary partitions Three primary partitions and one extended partition 	<ul style="list-style-type: none"> For Linux OS fdisk or parted For Windows OS Disk management

Disk Partition Style	Maximum Disk Capacity Supported	Maximum Number of Partitions Supported	Partitioning Tool Supported
GUID Partition Table (GPT)	18 EB NOTE 1 EB = 1,048,576 TB	Unlimited	<ul style="list-style-type: none">• For Linux OS parted• For Windows OS Disk management

10.3.2 Initializing a Windows Data Disk

A data disk attached to an instance or created together with an instance can be used by the instance only after being initialized. This section uses the Windows Server 2008 R2 Enterprise OS as an example. Specific initialization operations vary with OSs.

Prerequisites

- For details about how to log in to an ECS, see **Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS.**
- A disk has been attached to an instance and has not been initialized.

Context

Initializing a data disk is a high-risk operation. If the data disk contains useful data, perform the operations described in [18.7.11 Applying for a Snapshot](#) or [18.7.12 Creating a Backup](#) for the data disk first.

Procedure

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

The **Server Manager** page is displayed.

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

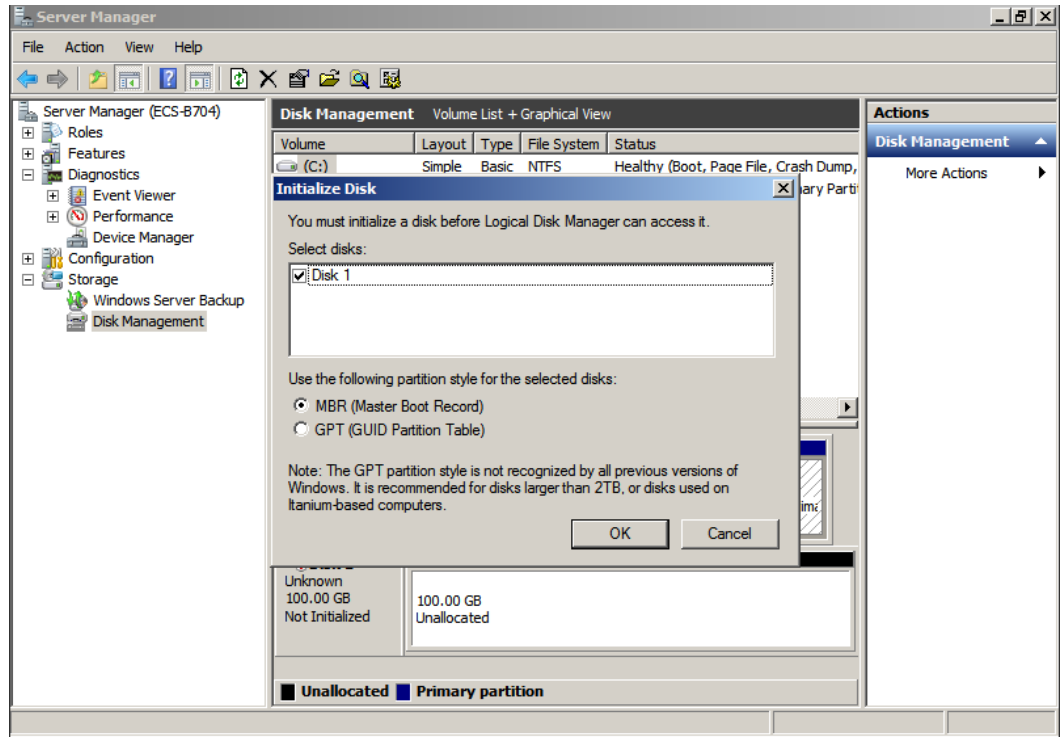
Step 3 If the disk to be initialized in the disk list is in **Offline** state, right-click in the disk area and choose **Online** from the shortcut menu.

Then, the disk status changes from **Offline** to **Uninitialized**.

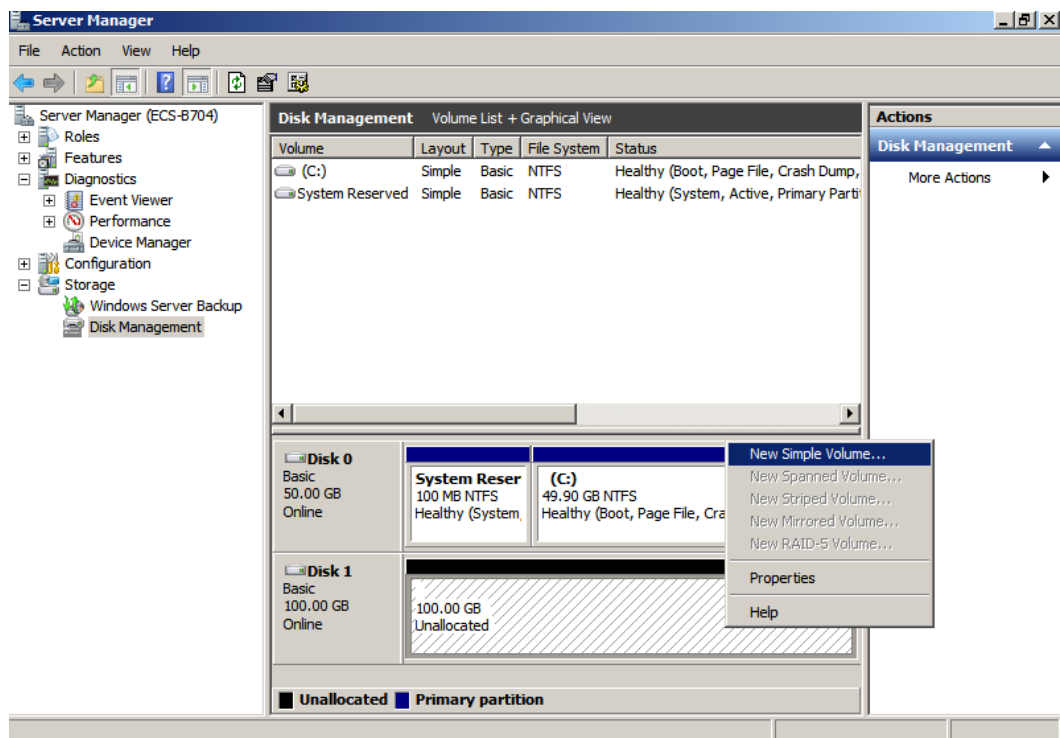
Step 4 Right-click in the disk area and choose **Initialize Disk** from the shortcut menu. In the displayed **Initialize Disk** dialog box, select **MBR (Master Boot Record)** and click **OK**.

NOTE

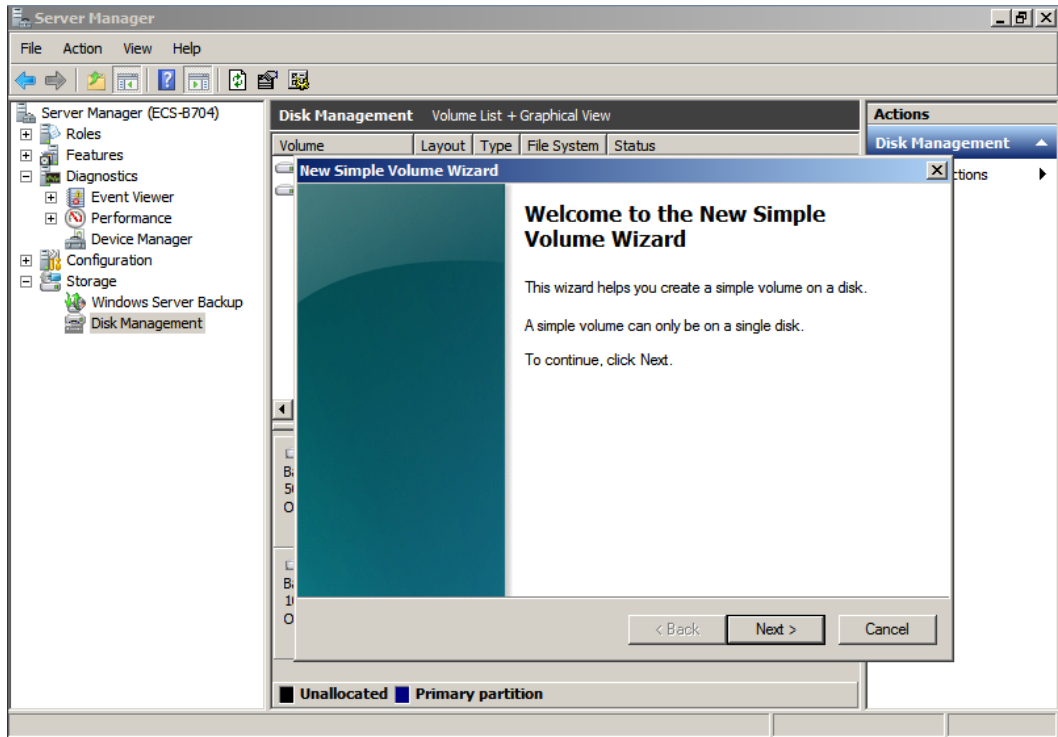
If the data disk to be initialized is larger than 2 TB, select **GPT (GUID Partition Table)** in the dialog box.



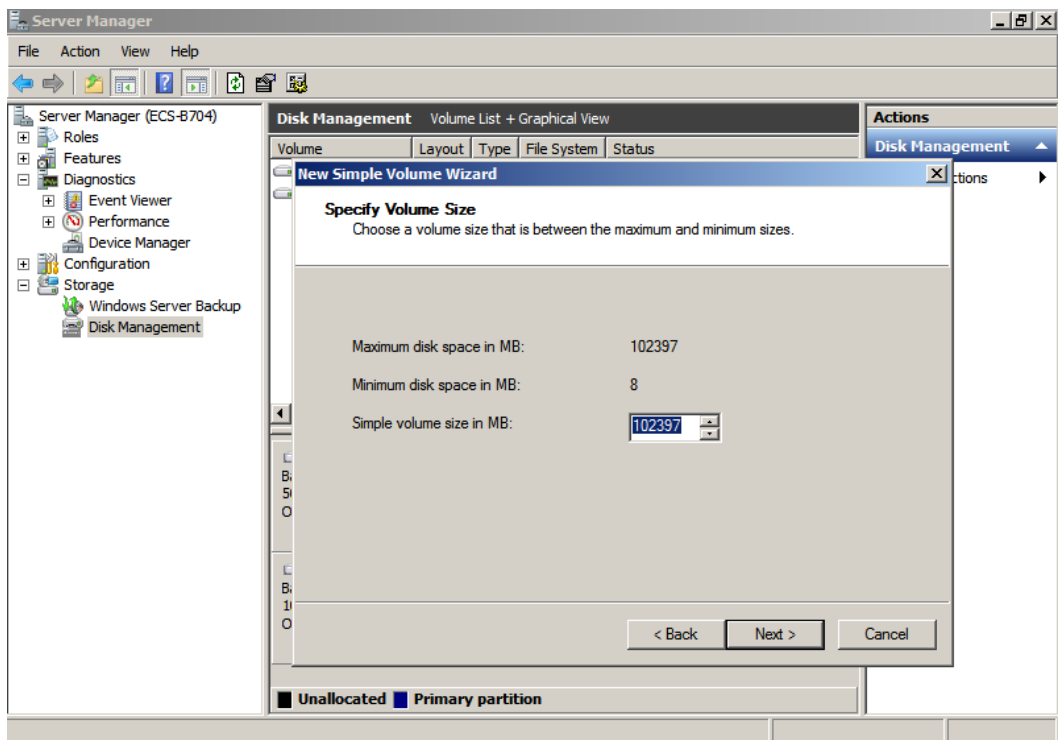
Step 5 Right-click at the unallocated disk space and choose **New Simple Volume** from the shortcut menu.




Step 6 On the displayed **New Simple Volume Wizard** page, click **Next**.



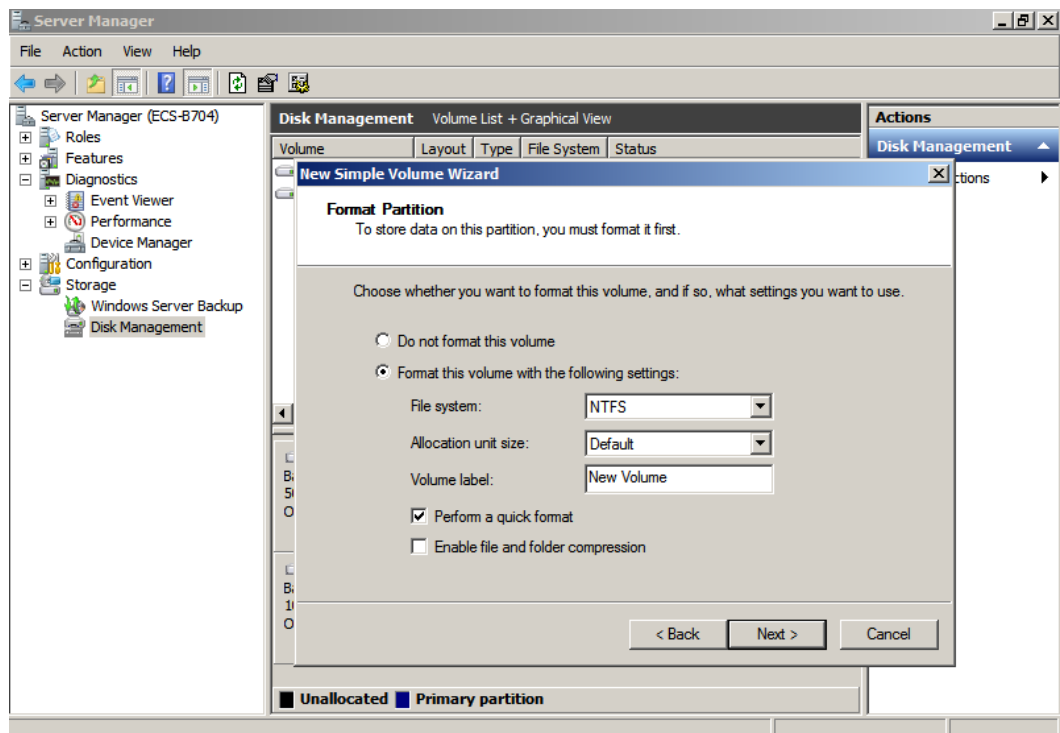
Step 7 Specify the simple volume size as required (the default value is the maximum) and click **Next**.




Step 8 Assign the driver letter and click **Next**.



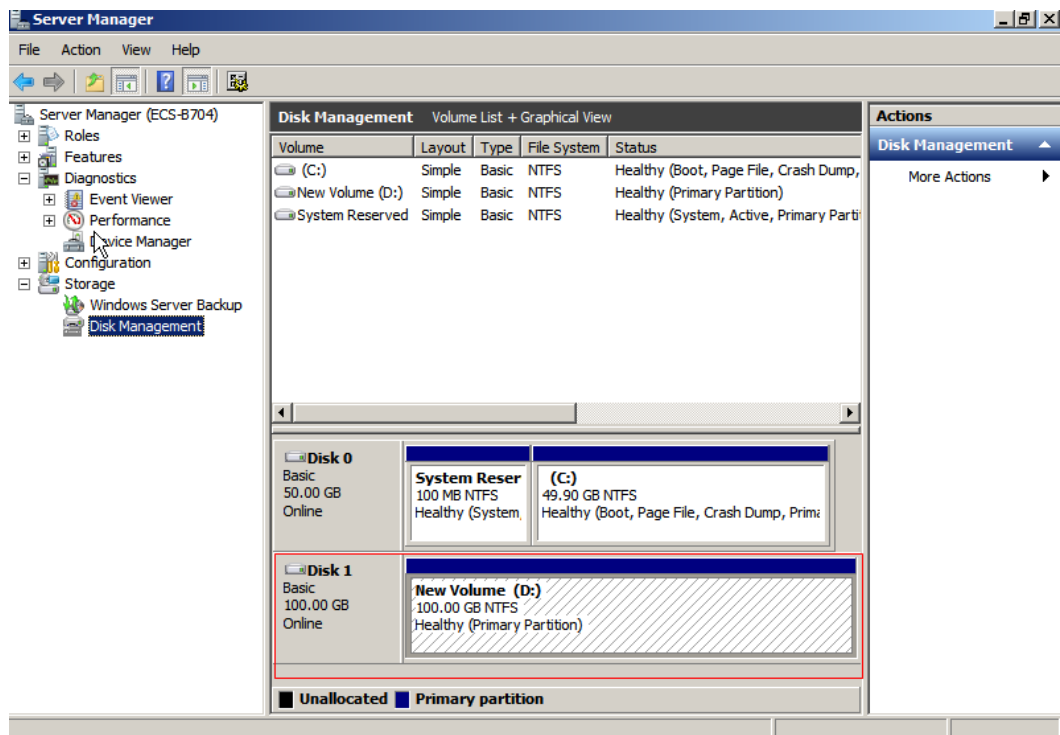
Step 9 Select **Format this volume with the following settings**, set parameters based on the actual requirements, and select **Perform a quick format**. Then click **Next**.



Step 10 Click **Finish**.



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



----End

10.3.3 Initializing a Linux Data Disk (fdisk)

A data disk attached to an instance or created together with an instance can be used by the instance only after being initialized. This section uses an instance running CentOS 7.0 (64-bit) as an example, and uses the fdisk partition tool to set up partitions for the data disk. Initialization operations vary with operating systems.

Prerequisites

- For details about how to log in to an ECS, see [Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS](#).
- A disk has been attached to an instance and has not been initialized.

Context

Both the fdisk and parted can be used to partition a Linux data disk. For a disk larger than 2 TB, only parted can be used because fdisk cannot partition such a large disk. For details, see [Using the parted tool](#).

Creating Partitions and Mounting a Disk

The following example shows how to create a new primary partition on a new data disk that has been attached to an instance. The primary partition will be created using fdisk, and MBR is the default partition style. Furthermore, the partition will be formatted using the ext4 file system, mounted on the `/mnt/sdc` directory, and set to be automatically mounted upon a system start.

Step 1 Run the following command to view information about the added data disk:

fdisk -l

Information similar to the following is displayed: (In the command output, the server contains two disks. `/dev/xvda` is the system disk, and `/dev/xvdb` is the added data disk.)

NOTE

If you do not log in to the ECS and run the `umount` command but directly detach the `/dev/xvdb` or `/dev/vdb` EVS disk on the management console, the disk name in the ECS may encounter a release delay. When you attach the disk to the server again, the mount point displayed on the management console may be inconsistent with that in the server. For example, device name `/dev/sdb` or `/dev/vdb` is selected for attachment, but `/dev/xvdc` or `/dev/vdc` may be displayed as the disk name in the OS. This issue does not adversely affect services.

```
[root@ecs-b656 test]# fdisk -l
```

```
Disk /dev/xvda: 42.9 GB, 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: 0x000cc4ad
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvda1	*	2048	2050047	1024000	83	Linux
/dev/xvda2		2050048	22530047	10240000	83	Linux
/dev/xvda3		22530048	24578047	1024000	83	Linux

```
/dev/xvda4 24578048 83886079 29654016 5 Extended
/dev/xvda5 24580096 26628095 1024000 82 Linux swap / Solaris
```

Disk /dev/xvdb: 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

NOTE

The capacity displayed here is inconsistent with the capacity of the EVS disk applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 Bytes and 1 GB = 1,000,000,000 Bytes.

Step 2 Run the following command to allocate partitions for the added data disk using `fdisk`:

fdisk *Newly added data disk*

In this example, `/dev/xvdb` is the newly added data disk.

fdisk /dev/xvdb

Information similar to the following is displayed:

```
[root@ecs-b656 test]# fdisk /dev/xvdb
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 0xb00005bd.
Command (m for help):
```

Step 3 Enter `n` and press **Enter**.

Entering `n` creates a partition.

There are two types of disk partitions:

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

```
Command (m for help): n
Partition type:
 p primary (0 primary, 0 extended, 4 free)
 e extended
```

Step 4 Enter `p` and press **Enter**.

The following describes how to create a primary partition.

Information similar to the following is displayed: (**Partition number** indicates the serial number of the primary partition. The value can be **1** to **4**.)

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

Step 5 Enter the primary partition number **1** and press **Enter**.

For example, select **1** as the partition number.

Information similar to the following is displayed: (**First sector** indicates the first sector number. The value can be **2048** to **20971519**, and the default value is **2048**.)

```
Partition number (1-4, default 1): 1
First sector (2048-20971519, default 2048):
```

Step 6 Press **Enter**.

The default start sector number 2048 is used as an example.

Information similar to the following is displayed: (**Last sector** indicates the last sector number. The value can be from **2048** to **20971519**, and the default value is **20971519**.)

```
First sector (2048-20971519, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
```

Step 7 Press **Enter**.

The default last sector number 20971519 is used as an example.

Information similar to the following is displayed, indicating that a primary partition is created for a 10 GB data disk.

```
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
Using default value 20971519
Partition 1 of type Linux and of size 10 GiB is set
Command (m for help):
```

Step 8 Enter **p** and press **Enter** to view the details about the created partition.

Information similar to the following is displayed: (Details about the **/dev/xvdb1** partition are displayed.)

```
Command (m for help): p
Disk /dev/xvdb: 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: 0xb00005bd
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvdb1		2048	20971519	10484736	83	Linux

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

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

Information similar to the following is displayed: (The partition is successfully created.)

```
Command (m for help): w
The partition table has been altered!
Calling ioctl() to re-read partition table.
Syncing disks.
```

Step 10 Run the following command to synchronize the new partition table to the data disk:

```
partprobe
```

Step 11 Run the following command to set the format for the file system of the newly created partition:

```
mkfs -t File system format /dev/xvdb1
```

For example, run the following command to set the **ext4** file system for the **/dev/xvdb1** partition:

```
mkfs -t ext4 /dev/xvdb1
```

Information similar to the following is displayed:

```
[root@ecs-b656 test]# mkfs -t ext4 /dev/xvdb1
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
655360 inodes, 2621184 blocks
131059 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2151677952
80 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

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

NOTE

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

Step 12 Run the following command to create a mount directory:

```
mkdir Mount directory
```

/mnt/sdc is used in this example.

```
mkdir /mnt/sdc
```

Step 13 Run the following command to mount the new partition to the mount directory created in [Step 12](#):

```
mount /dev/xvdb1 Mount directory
```

For example, run the following command to mount the newly created partition on **/mnt/sdc**:

```
mount /dev/xvdb1 /mnt/sdc
```

Step 14 Run the following command to view the mount result:

```
df -TH
```

Information similar to the following is displayed. The newly created **/dev/xvdb1** partition has been mounted on **/mnt/sdc**.

```
[root@ecs-b656 test]# df -TH
Filesystem Type Size Used Avail Use% Mounted on
/dev/xvda2 xfs 11G 7.4G 3.2G 71% /
devtmpfs devtmpfs 4.1G 0 4.1G 0% /dev
tmpfs tmpfs 4.1G 82k 4.1G 1% /dev/shm
tmpfs tmpfs 4.1G 9.2M 4.1G 1% /run
tmpfs tmpfs 4.1G 0 4.1G 0% /sys/fs/cgroup
/dev/xvda3 xfs 1.1G 39M 1.1G 4% /home
```

```
/dev/xvda1  xfs  1.1G 131M 915M 13% /boot
/dev/xvdb1  ext4  11G  38M 9.9G  1% /mnt/sdc
```

----End

Setting Automatic Disk Attachment Upon Instance Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify `/dev/xvdb1` in `/etc/fstab`. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in `/etc/fstab` to automatically attach the disk at a system start.

NOTE

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

```
blkid Disk partition
```

For example, run the following command to query the UUID of `/dev/xvdb1`:

```
blkid /dev/xvdb1
```

Information similar to the following is displayed: (The UUID of `/dev/xvdb1` is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the `fstab` file using the vi editor:

```
vi /etc/fstab
```

Step 3 Press `i` to enter the editing mode.

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

```
UUID=xxx attachment directory file system defaults 0 2
```

Assuming that the file system is `ext4` and the attachment directory is `/mnt/sdc`.
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the `fstab` file before detaching the disk. Otherwise, you may fail to access the OS after the disk is detached.

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

The system saves the configuration and exits the vi editor.

----End

10.3.4 Initializing a Linux Data Disk (parted)

A data disk attached to an instance or created together with an instance can be used by the instance only after being initialized. This section uses an instance running CentOS 7.0 (64-bit) as an example, and uses the parted partition tool to set up partitions for the data disk. Initialization operations vary with operating systems.

Prerequisites

- For details about how to log in to an ECS, see [Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS](#).
- A disk has been attached to an instance and has not been initialized.

Mounting Partitions to a Disk

The following example shows how to create a partition on a new data disk that has been attached to an instance. The partition will be created using parted and GPT is the default partition style. Furthermore, the partition will be formatted using the ext4 file system, mounted on the `/mnt/sdc` directory, and set to be automatically mounted upon a system start.

Step 1 Run the following command to view information about the added data disk:

lsblk

Information similar to the following is displayed:

NOTE

If you do not log in to the ECS and run the `umount` command but directly detach the `/dev/xvdb` or `/dev/vdb` EVS disk on the management console, the disk name in the ECS may encounter a release delay. When you attach the disk to the server again, the mount point displayed on the management console may be inconsistent with that in the server. For example, device name `/dev/sdb` or `/dev/vdb` is selected for attachment, but `/dev/xvdc` or `/dev/vdc` may be displayed as the disk name in the OS. This issue does not adversely affect services.

```
[root@ecs-centos-70 linux]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0 0 40G 0 disk
├─xvda1 202:1 0 4G 0 part [SWAP]
└─xvda2 202:2 0 36G 0 part /
xvdb 202:16 0 10G 0 disk
```

The command output indicates that the server contains two disks. `/dev/xvda` is the system disk and `/dev/xvdb` is the new data disk.

Step 2 Run the following command to enter parted to partition the added data disk:

parted *Added data disk*

In this example, `/dev/xvdb` is the newly added data disk.

parted `/dev/xvdb`

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# parted /dev/xvdb
GNU Parted 3.1
```

```
Using /dev/xvdb  
Welcome to GNU Parted! Type 'help' to view a list of commands.
```

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/xvdb: unrecognised disk label  
Model: Xen Virtual Block Device (xvd)  
Disk /dev/xvdb: 10.7GB  
Sector size (logical/physical): 512B/512B  
Partition Table: unknown  
Disk Flags:
```

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

 **NOTE**

The capacity displayed here is inconsistent with the capacity of the EVS disk applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 Bytes and 1 GB = 1,000,000,000 Bytes.

Step 4 Run the following command to set the disk partition style:

```
mklabel Disk partition style
```

The disk partition styles include MBR and GPT. For example, run the following command to set the partition style to GPT:

```
mklabel gpt
```

NOTICE

If you change the disk partition style after the disk has been used, the original data on the disk will be cleared. Therefore, select a proper disk partition style when initializing the disk.

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: Xen Virtual Block Device (xvd)  
Disk /dev/xvdb: 20971520s  
Sector size (logical/physical): 512B/512B  
Partition Table: gpt  
Disk Flags:  
  
Number Start End Size File system Name Flags
```

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

Step 7 Enter **mkpart opt 2048s 100%** and press **Enter**.

In the command, **opt** is the name of the new partition, **2048s** indicates the start of the partition, and **100%** indicates the end of the partition. You can plan the number and capacity of disk partitions based on 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? Cancel
```

If the preceding warning message is displayed, enter **Cancel** to stop the partitioning. Then, find the first sector with the best disk performance and use that value to partition the disk.

Step 8 Enter **p** and press **Enter** to view the details about the created partition.

Information similar to the following is displayed:

```
(parted) p  
Model: Xen Virtual Block Device (xvd)  
Disk /dev/xvdb: 20971520s  
Sector size (logical/physical): 512B/512B  
Partition Table: gpt  
Disk Flags:  
  
Number Start End Size File system Name Flags  
1 2048s 20969471s 20967424s opt
```

Details about the **/dev/xvdb1** partition are displayed.

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

Step 10 Run the following command to view the disk partition information:

lsblk

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# lsblk  
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT  
xvda 202:0 0 40G 0 disk  
├─xvda1 202:1 0 4G 0 part [SWAP]  
└─xvda2 202:2 0 36G 0 part /  
xvdb 202:16 0 100G 0 disk  
└─xvdb1 202:17 0 100G 0 part
```

In the command output, **/dev/xvdb1** is the partition you created.

Step 11 Run the following command to set the format for the file system of the newly created partition:

NOTICE

The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

mkfs -t *File system format* /dev/xvdb1

For example, run the following command to set the **ext4** file system for the **/dev/xvdb1** partition:

mkfs -t ext4 /dev/xvdb1

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# mkfs -t ext4 /dev/xvdb1
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
655360 inodes, 2620928 blocks
131046 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2151677925
80 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

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.

Step 12 Run the following command to create a mount point:

```
mkdir Mount point
```

For example, run the following command to create the **/mnt/sdc** mount point:

```
mkdir /mnt/sdc
```

Step 13 Run the following command to mount the new partition to the mount point created in [Step 12](#):

```
mount /dev/xvdb1 Mount point
```

For example, run the following command to mount the newly created partition on **/mnt/sdc**:

```
mount /dev/xvdb1 /mnt/sdc
```

Step 14 Run the following command to view the mount result:

```
df -TH
```

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda2      xfs       39G   4.0G 35G  11% /
devtmpfs        devtmpfs  946M   0 946M  0% /dev
tmpfs           tmpfs     954M   0 954M  0% /dev/shm
tmpfs           tmpfs     954M   9.1M 945M  1% /run
tmpfs           tmpfs     954M   0 954M  0% /sys/fs/cgroup
/dev/xvdb1      ext4      11G   38M 101G  1% /mnt/sdc
```

The newly created **/dev/xvdb1** is mounted on **/mnt/sdc**.

----End

Setting Automatic Disk Attachment at a System Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by

referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify `/dev/xvdb1` in `/etc/fstab`. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in `/etc/fstab` to automatically attach the disk at a system start.

NOTE

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

```
blkid Disk partition
```

For example, run the following command to query the UUID of `/dev/xvdb1`:

```
blkid /dev/xvdb1
```

Information similar to the following is displayed: (The UUID of `/dev/xvdb1` is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1  
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the `fstab` file using the vi editor:

```
vi /etc/fstab
```

Step 3 Press `i` to enter the editing mode.

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

```
UUID=xxx attachment directory file system defaults 0 2
```

Assuming that the file system is `ext4` and the attachment directory is `/mnt/sdc`.
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the `fstab` file before detaching the disk. Otherwise, you may fail to access the OS after the disk is detached.

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

The system saves the configuration and exits the vi editor.

----End

10.4 Expanding EVS Disk Capacity

10.4.1 Overview

You can expand EVS disk capacity if the disk capacity becomes insufficient. For system disks, you can only expand the capacity of existing disks. For data disks,

you can either expand the capacity of existing disks or add data disks by referring to [add new disks](#) and attach them to instances.

After expanding the capacity of a disk, you need to create partitions for the new capacity or create partitions to replace the original disk partitions.

- **Creating partitions for the new capacity**
When using this method, you need to create partitions, without unmounting any existing partitions. This operation does not interrupt ongoing services and has minor impacts on services. This method is recommended for system disks or disks carrying services that cannot be interrupted. If the MBR partition style is used, the disk capacity must be less than 2 TB and the number of partitions does not exceed the upper limit after the expansion.
- **Creating partitions to replace existing ones**
If the MBR partition style is used and the number of disk partitions has reached the upper limit, new partitions cannot be created. In this situation, you need to unmount existing partitions first and then create new ones to replace them. This operation does not delete data in existing partitions, but services must be suspended during the operation, affecting ongoing services.

[Table 10-4](#) describes the post-expansion operations.

Table 10-4 Post-expansion operations for an EVS disk

OS	Disk Capacity After Expansion	Current Disk Partition Style	Post-Expansion Operation
Linux	≤ 2 TB	MBR	Use the fdisk or parted tool to create partitions for the new capacity. Operations After Expanding the Capacity of a Disk in Linux (Adding Partitions Using fdisk) Operations After Expanding the Capacity of a Disk in Linux (Adding Partitions Using parted)
			Use the fdisk or parted tool to create new partitions to replace existing ones. This operation interrupts ongoing user services and is not recommended after system disk capacity expansion. Operations After Expanding the Capacity of a Disk in Linux (Replacing the Original Partitions Using fdisk) Operations After Expanding the Capacity of a Disk in Linux (Replacing the Original Partitions Using parted)

OS	Disk Capacity After Expansion	Current Disk Partition Style	Post-Expansion Operation
		GPT	Use the parted tool to add partitions for the new capacity or create partitions to replace the original partitions. Operations After Expanding the Capacity of a Disk in Linux (Adding Partitions Using parted) Operations After Expanding the Capacity of a Disk in Linux (Replacing the Original Partitions Using parted)
	> 2 TB	MBR	Use the parted tool to change the partition style from MBR to GPT. However, this operation will clear disk data.
		GPT	Use the parted tool to add partitions for the new capacity or create partitions to replace the original partitions. Operations After Expanding the Capacity of a Disk in Linux (Adding Partitions Using parted) Operations After Expanding the Capacity of a Disk in Linux (Replacing the Original Partitions Using parted)
Windows	-	MBR	Allocate the additional disk space to existing partitions. Operations After Expanding the Capacity of a Disk in Windows
		GPT	

10.4.2 Expanding Disk Capacity Online

This section describes how to expand the capacity of an EVS disk attached to an instance.

Restrictions

- When you expand the capacity of a disk online, the instance to which the disk is attached must be in the **Running** or **Stopped** state.
- Shared EVS disks do not support online capacity expansion, that is, the capacity of a shared EVS disk can be expanded only when the disk is in the **Available** state.
- The capacity of a disk configured with the DR service (CSHA/CSDR/VHA) cannot be expanded.

- When the backend storage is Huawei SAN storage (OceanStor V3/V5 series or OceanStor Dorado V3 series/6.x) or heterogeneous storage, if the EVS disk has snapshots, capacity expansion is not supported. When the backend storage is Huawei Distributed Block Storage, capacity expansion can be performed for an EVS disk with snapshots.
- If backend storage of the disk is heterogeneous storage, online capacity expansion is not supported while offline capacity expansion is supported.
- Capacity expansion cannot be performed when the disk status is **Reserved** or **Maintenance**.
- For the OSs that support online expansion, see [Table 10-5](#).

Table 10-5 OSs

OS	Version
CentOS	7.3 64-bit
	7.2 64-bit
	6.8 64-bit
	6.7 64-bit
	6.5 64-bit
Debian	8.6.0 64-bit
	8.5.0 64-bit
Fedora	25 64-bit
	24 64-bit
SUSE	SUSE Linux Enterprise Server 12 SP2 (64-bit)
	SUSE Linux Enterprise Server 12 SP1 (64-bit)
	SUSE Linux Enterprise Server 11 SP4 (64-bit)
	SUSE Linux Enterprise Server 12 (64-bit)
OpenSUSE	42.2 64-bit
	42.1 64-bit
Oracle Linux Server release	7.3 64-bit
	7.2 64-bit
	6.8 64-bit
	6.7 64-bit
Ubuntu Server	16.04 64-bit
	14.04 64-bit
	14.04.4 64-bit

OS	Version
Windows	Windows Server 2008 R2 Enterprise 64-bit
	Windows Server 2012 R2 Standard 64-bit
	Windows Server 2016 R2 Standard 64-bit
Red Hat Enterprise Linux (RHEL)	7.3 64-bit
	6.8 64-bit

Procedure

Step 1 [Log in to EVS Console as a VDC administrator or VDC operator.](#)

Step 2 In the EVS disk list, locate the row that contains the target disk, click **More** in the **Operation** column, and choose **Expand Capacity**.

The **Expand Capacity** page is displayed.

 **NOTE**

To view information about the instance to which the EVS disk is attached, click the instance name in the **Attaching Information** column.

Step 3 Set **Added Capacity (GB)** as prompted and click **Next**.

Step 4 On the **Resource Details** page, confirm the EVS disk specifications.

- If you do not need to modify the specifications, click **Apply Now** to start the EVS disk capacity expansion.
- If you need to modify the specifications, click **Previous** to modify parameters.

After the expansion is submitted, the disk list page is displayed.

Step 5 If the EVS disk whose capacity is to be expanded requires approval, contact the administrator for approval. Otherwise, skip this step.

On the **Elastic Volume Service** page, view the capacity of the EVS disk. If the disk capacity has increased, the expansion is successful.

----End

Follow-up Procedure

After you have expanded the capacity of an EVS disk, perform follow-up operations for the additional capacity.

 **NOTE**

If the instance to which the EVS disk is attached is stopped during capacity expansion, power on the instance before capacity expansion.

- For Windows, see [10.4.4 Operations After Expanding Disk Capacity in Windows](#).
- For Linux, see [Table 10-4](#).

10.4.3 Expanding Disk Capacity Offline

This section describes how to expand the capacity of an EVS disk not attached to any instance.

Restrictions

- The capacity of a disk configured with the DR service (CSHA/CSDR/VHA) cannot be expanded.
- When the backend storage is Huawei SAN storage (OceanStor V3/V5 series or OceanStor Dorado V3 series/6.x) or heterogeneous storage, if the EVS disk has snapshots, capacity expansion is not supported. When the backend storage is Huawei Distributed Block Storage, capacity expansion can be performed for an EVS disk with snapshots.
- Capacity expansion cannot be performed when the disk status is **Reserved** or **Maintenance**.

Procedure

Step 1 [Log in to EVS Console as a VDC administrator or VDC operator](#).

Step 2 In the EVS disk list, locate the row that contains the target disk, click **More** in the **Operation** column, and choose **Expand Capacity**.

The **Expand Capacity** page is displayed.

NOTE

To view information about the instance to which the EVS disk is attached, click the instance name in the **Attaching Information** column.

Step 3 Set **Added Capacity (GB)** as prompted and click **Next**.

Step 4 On the **Resource Details** page, confirm the EVS disk specifications.

- If you do not need to modify the specifications, click **Apply Now** to start the EVS disk capacity expansion.
- If you need to modify the specifications, click **Previous** to modify parameters.

After the expansion is submitted, the disk list page is displayed.

Step 5 If the EVS disk whose capacity is to be expanded requires approval, contact the administrator for approval. Otherwise, skip this step.

On the **Elastic Volume Service** page, view the capacity of the EVS disk. If the disk capacity has increased, the expansion is successful.

----End

Follow-up Procedure

After you have expanded the capacity of an EVS disk, perform follow-up operations on the additional volume.

1. Attach the expanded EVS disk to an instance. For details, see section [10.2 Attaching an EVS Disk](#).

2. Operations after attaching:
 - For Windows, see [Operations After Expanding Disk Capacity in Windows](#).
 - For Linux, see [Table 1 Post-expansion operations for an EVS disk](#).

10.4.4 Operations After Expanding Disk Capacity in Windows

This section describes how to perform follow-up operations for the additional disk space after you have expanded the capacity of an EVS disk and attached the EVS disk to an instance. The method for allocating the extended space of an EVS disk varies depending on the in-use server OS. This section uses Windows Server 2008 R2 Enterprise as the example OS to describe how to perform post-expansion operations for an EVS disk. For other Windows OSs, see the corresponding OS product documents.

Prerequisites

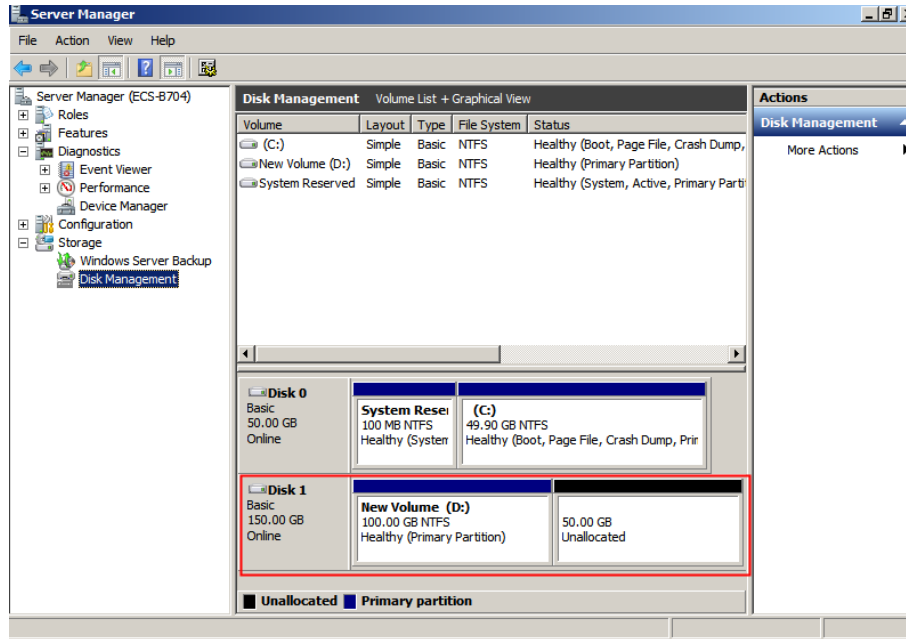
- For details about how to log in to an ECS, see [Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS](#).
- A disk has been attached to an instance and has not been initialized.

Context

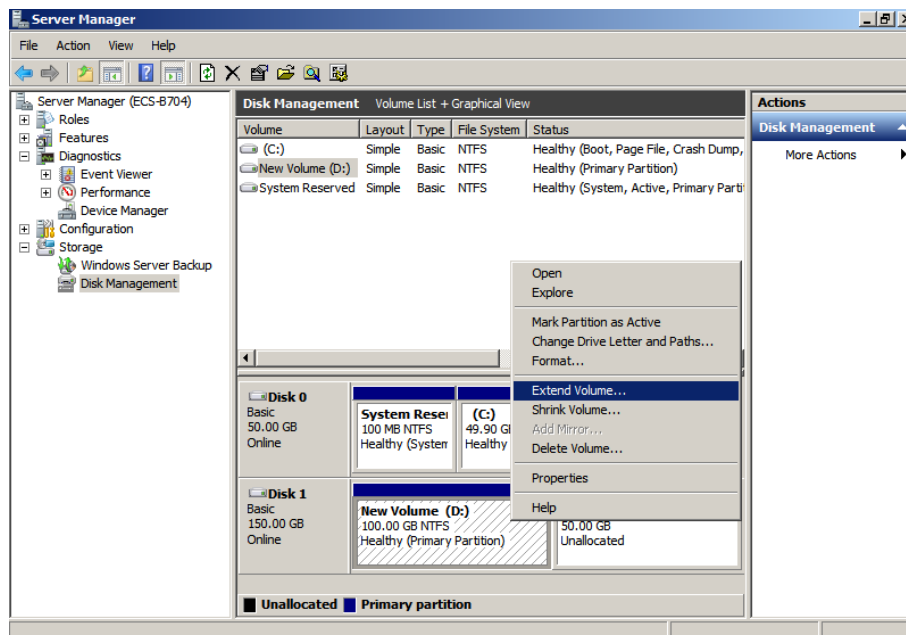
After the capacity expansion has succeeded, the new EVS disk space needs to be allocated to an existing partition or a new partition. This section describes how to allocate new EVS disk space to an existing partition.

Procedure

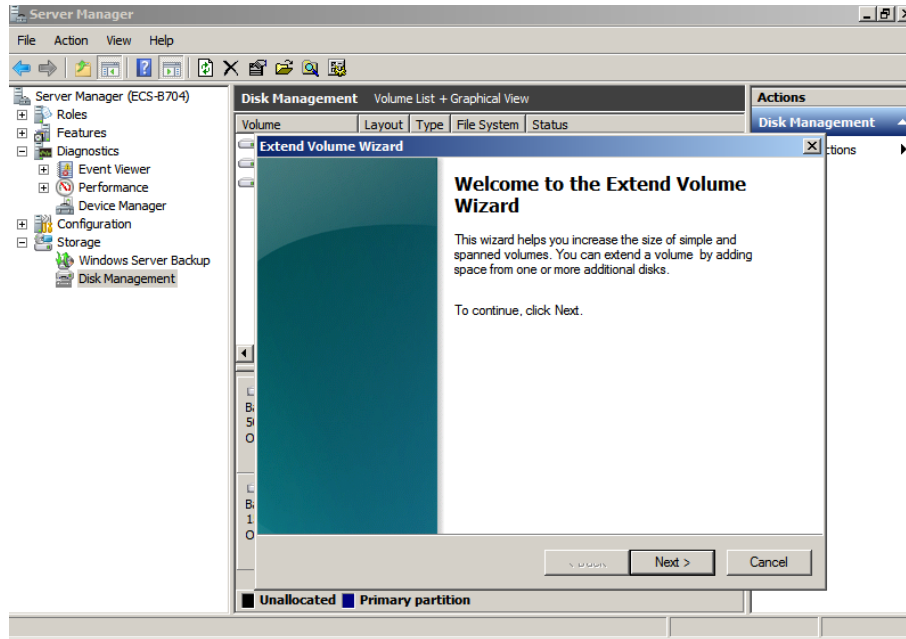
- Step 1** In desktop, right-click **Computer** and choose **Manage** from the shortcut menu.
The **Server Manager** window is displayed.
- Step 2** In the navigation tree, choose **Storage > Disk Management**.
The **Disk Management** page is displayed.
- Step 3** On the **Disk Management** page, select the disk and partition that needs to be extended. The current partition size and unallocated disk space are displayed.



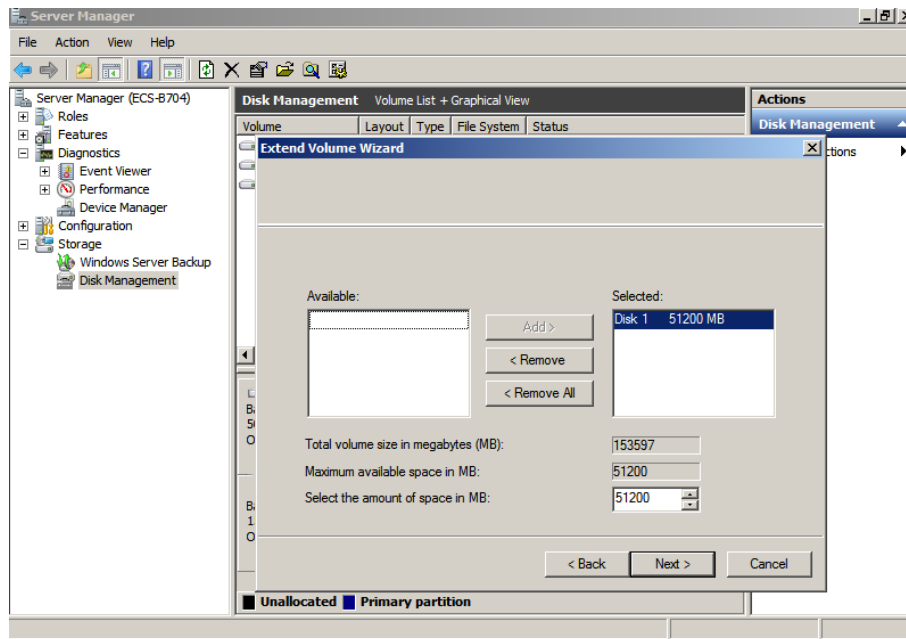
Step 4 Right-click the selected disk and choose **Extend Volume**.



Step 5 On the displayed **Extend Volume Wizard** page, click **Next**.

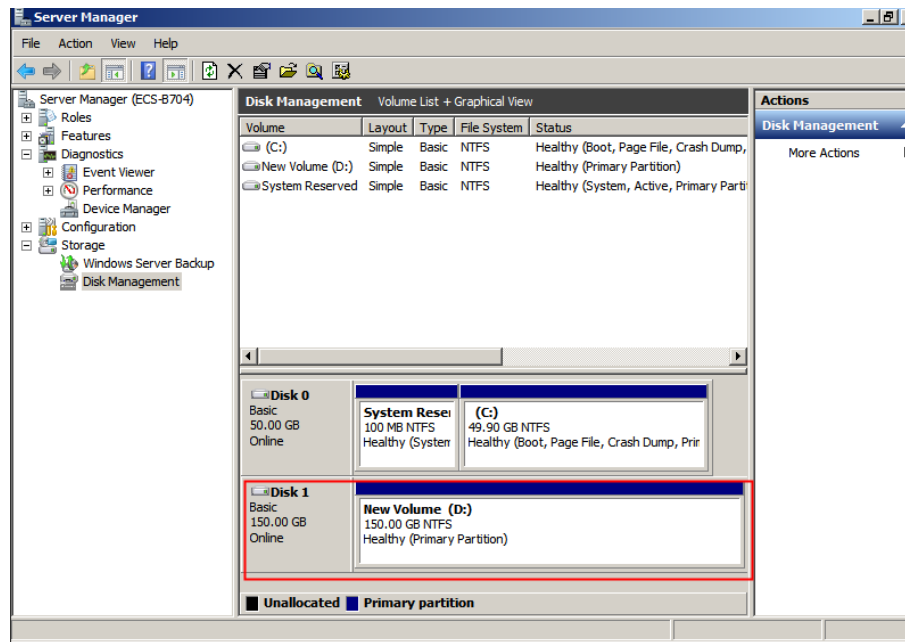


Step 6 In the text box to the right of **Select the amount of space in MB**, enter the extended capacity and click **Next**.



Step 7 Click **Finish** to complete the wizard.

After the expansion is successful, the disk capacity is greater than the original capacity.



----End

10.4.5 Operations After Expanding Disk Capacity in Linux (Adding Partitions Using fdisk)

After the capacity of an EVS disk is expanded and the disk is attached to an instance, the additional space needs to be allocated to a new partition. This section uses CentOS 7.0 (64-bit) as an example to describe how to use `fdisk` to create partitions for a disk after capacity expansion.

Prerequisites

- For details about how to log in to an ECS, see **Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS.**
- A disk has been attached to an instance and has not been initialized.

Procedure

The steps below describe how to allocate the additional space to a new partition and mount the partition on `/opt`. The disk is partitioned using MBR, the disk capacity is less than 2 TB, and the number of partitions is lower than the upper limit. You can use either the `fdisk` or `parted` tool to partition the disk space. This section uses the `fdisk` tool as an example.

Step 1 Run the following command to query and view the disk information:

```
fdisk -l
```

Information similar to the following is displayed: (`/dev/xvda` is the system disk.)

```
[root@ecs-bab9 test]# 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: 0x000cc4ad
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvda1	*	2048	2050047	1024000	83	Linux
/dev/xvda2		2050048	22530047	10240000	83	Linux
/dev/xvda3		22530048	24578047	1024000	83	Linux
/dev/xvda4		24578048	83886079	29654016	5	Extended
/dev/xvda5		24580096	26628095	1024000	82	Linux swap / Solaris

NOTE

The capacity displayed here is inconsistent with the capacity of the EVS disk applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 bytes and 1 GB = 1,000,000,000 bytes.

Step 2 Run the following command to enter the fdisk (`/dev/xvda` is used as an example):

```
fdisk /dev/xvda
```

Information similar to the following is displayed:

```
[root@ecs-bab9 test]# 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):
```

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

Because the system disk has five existing partitions, the system automatically creates the sixth one.

Information similar to the following is displayed:

```
Command (m for help): n
All primary partitions are in use
Adding logical partition 6
First sector (26630144-83886079, default 26630144):
```

Step 4 Enter the new partition's first sector number, for example the default value, and press **Enter**.

The first sector number must be greater than the last sector numbers of existing partitions.

Information similar to the following is displayed:

```
First sector (26630144-83886079, default 26630144):
Using default value 26630144
Last sector, +sectors or +size{K,M,G} (26630144-83886079, default 83886079):
```

Step 5 Enter the new partition's last sector number and press **Enter**.

The default last sector number is used as an example.

Information similar to the following is displayed:

```
Last sector, +sectors or +size{K,M,G} (26630144-83886079, default 83886079):
Using default value 83886079
Partition 6 of type Linux and of size 27.3 GiB is set
```

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

Step 6 Enter **p** and press **Enter** to view the created partition.

Information similar to the following is displayed:

```
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: 0x000cc4ad

   Device Boot      Start         End      Blocks   Id  System
/dev/xvda1 *         2048        2050047    1024000   83  Linux
/dev/xvda2          2050048        22530047    10240000   83  Linux
/dev/xvda3          22530048        24578047     1024000   83  Linux
/dev/xvda4          24578048        83886079    29654016    5  Extended
/dev/xvda5          24580096        26628095     1024000   82  Linux swap / Solaris
/dev/xvda6          26630144        83886079    28627968   83  Linux
```

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

Step 7 Enter **w** and press **Enter**.

The partition result is written into the partition table, and the partitioning is complete.

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.

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.
```

Step 8 Run the following command to synchronize the new partition table to the OS:

partprobe

Step 9 Run the following command to set the format for the file system of the new partition:

In this example, the ext4 file system is used.

mkfs -t ext4 /dev/xvda6

Information similar to the following is displayed:

```
[root@ecs-bab9 test]# mkfs -t ext4 /dev/xvda6
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
1790544 inodes, 7156992 blocks
357849 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2155872256
219 block groups
32768 blocks per group, 32768 fragments per group
8176 inodes per group
```

```
Superblock backups stored on blocks:  
32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,  
4096000
```

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

The formatting takes a while, and you need to observe the system running status. Once **done** is displayed in the command output, the formatting is complete.

Step 10 Run the following command to mount the new partition to the space-demanding directory, for example **/opt**:

```
mount /dev/xvda6 /opt
```

Information similar to the following is displayed:

```
[root@ecs-bab9 test]# mount /dev/xvda6 /opt  
[root@ecs-bab9 test]#
```

NOTE

If the new partition is mounted to a directory that is not empty, the subdirectories and files in the directory will be hidden. In this situation, you are advised to mount the new partition to an empty directory or a new directory. If the new partition must be mounted to a directory that is not empty, move the subdirectories and files in this directory to another directory temporarily. After the partition is mounted, move the subdirectories and files back.

Step 11 Run the following command to view the mount result:

```
df -TH
```

Information similar to the following is displayed:

```
[root@ecs-bab9 test]# df -TH  
Filesystem Type Size Used Avail Use% Mounted on  
/dev/xvda2 xfs 11G 7.4G 3.2G 71% /  
devtmpfs devtmpfs 4.1G 0 4.1G 0% /dev  
tmpfs tmpfs 4.1G 82k 4.1G 1% /dev/shm  
tmpfs tmpfs 4.1G 9.2M 4.1G 1% /run  
tmpfs tmpfs 4.1G 0 4.1G 0% /sys/fs/cgroup  
/dev/xvda3 xfs 1.1G 39M 1.1G 4% /home  
/dev/xvda1 xfs 1.1G 131M 915M 13% /boot  
/dev/xvda6 ext4 29G 47M 28G 1% /opt
```

```
----End
```

Setting Automatic Disk Mounting at a System Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify **/dev/xvdb1** in **/etc/fstab**. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in **/etc/fstab** to automatically attach the disk at a system start.

NOTE

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

```
blkid Disk partition
```

For example, run the following command to query the UUID of `/dev/xvdb1`:

```
blkid /dev/xvdb1
```

Information similar to the following is displayed: (The UUID of `/dev/xvdb1` is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the `fstab` file using the vi editor:

```
vi /etc/fstab
```

Step 3 Press `i` to enter the editing mode.

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

```
UUID=xxx attachment directory file system defaults 0 2
```

Assuming that the file system is `ext4` and the attachment directory is `/mnt/sdc`.

```
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2
```

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the `fstab` file before detaching the disk. Otherwise, you may fail to access the OS after the disk is detached.

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

The system saves the configuration and exits the vi editor.

```
----End
```

10.4.6 Operations After Expanding Disk Capacity in Linux (Adding Partitions Using parted)

This section describes how to perform follow-up operations for the additional disk space after you have expanded the capacity of an EVS disk and attached the EVS disk to an instance. This section uses EulerOS 2.0 (64-bit) to describe how to allocate the additional disk space to a partition using `parted`.

Prerequisites

- For details about how to log in to an ECS, see [Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS](#).
- A disk has been attached to an instance and has not been initialized.

Procedure

The steps below describe how to allocate the additional space to a new partition and mount the partition on `/opt`. The disk is partitioned using MBR, the disk

capacity is less than 2 TB, and the number of partitions is lower than the upper limit. You can use either the `fdisk` or `parted` tool to partition the disk space. This section uses the `parted` tool as an example.

Step 1 Run the following command to view the disk partition information:

lsblk

If the following information is displayed, the total capacity of the current system disk `dev/xvda` is 80 GB. 40 GB of the disk has been allocated to partitions, and the remaining 40 GB is additional disk space and have not been allocated to any partition.

```
[root@ecs-1120 linux]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0 0 80G 0 disk
├─xvda1 202:1 0 40G 0 part /
xvdb 202:16 0 250G 0 disk
├─xvdb1 202:17 0 100G 0 part
├─xvdb2 202:18 0 50G 0 part
xvdc 202:32 0 40G 0 disk
├─xvdc1 202:33 0 8G 0 part
└─xvdc2 202:34 0 32G 0 part
```

Step 2 Run the following command to create a partition for the additional system disk space:

parted System disk

parted /dev/xvda

The following information is displayed:

```
[root@ecs-1120 linux]# parted /dev/xvda
GNU Parted 3.1
Using /dev/xvda
Welcome to GNU Parted! Type 'help' to view a list of commands.
```

Step 3 Enter `unit s` and press **Enter** to set the measurement unit of the disk to sector numbers.

Step 4 Enter `p` and press **Enter** to view the current disk partition style.

Partition Table specifies the partition style for existing disks. `msdos` indicates that the disk partition style is MBR, and `gpt` indicates that the disk partition style is GPT.

The following information is displayed:

```
(parted) unit s
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvda: 167772160s
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number Start End      Size   Type  File system  Flags
 1      2048s 83886079s 83884032s primary ext4
```

Step 5 Enter `mkpart` and press **Enter** to create a partition.

Step 6 Enter `p` and press **Enter** to create a primary partition. Creating a primary partition is used as an example.

Set the file system format to `ext4`.

Set the first sector and last sector to **83886080** and **167772159** for the **dev/xvda2** partition, respectively.

Set the parameters as required.

The following information is displayed:

```
(parted) mkpart  
Partition type? primary/extended? p  
File system type? [ext2]? ext4  
Start? 83886080  
End? 167772159
```

NOTE

The file system type may fail to be configured in this step. You can reconfigure the file system format according to [Step 9](#).

Step 7 Enter **p** and press **Enter** to view the created partition.

Information similar to the following is displayed:

```
(parted) p  
Model: Xen Virtual Block Device (xvd)  
Disk /dev/xvda: 167772160s  
Sector size (logical/physical): 512B/512B  
Partition Table: msdos  
Disk Flags:  
  
Number Start End Size Type File system Flags  
1 2048s 83886079s 83884032s primary ext4  
2 83886080s 167772159s 83886080s primary
```

Step 8 Enter **q** and press **Enter** to exit parted.

Step 9 Run the following command to set the format for the file system of the new partition:

In this example, the ext4 file system is used.

```
mkfs -t ext4 /dev/xvda2
```

Information similar to the following is displayed:

Wait for the formatting. If **done** is displayed in the command output, the formatting has been complete.

```
[[root@ecs-1120 linux]# mkfs -t ext4 /dev/xvda2  
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  
2621440 inodes, 10485760 blocks  
524288 blocks (5.00%) reserved for the super user  
First data block=0  
Maximum filesystem blocks=2157969408  
320 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  
  
Allocating group tables: done  
Writing inode tables: done  
Creating journal (32768 blocks): done  
Writing superblocks and filesystem accounting information: done
```

Step 10 Run the following command to mount the new partition to the space-demanding directory, for example `/opt`:

```
mount /dev/xvda2 /opt
```

 **NOTE**

If the new partition is mounted to a directory that is not empty, the subdirectories and files in the directory will be hidden. In this situation, you are advised to mount the new partition to an empty directory or a new directory. If the new partition must be mounted to a directory that is not empty, move the subdirectories and files in this directory to another directory temporarily. After the partition is mounted, move the subdirectories and files back.

Step 11 Run the following command to view the mount result:

df -TH

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda1      ext4      43G   8.3G   33G   21% /
devtmpfs        devtmpfs  885M    0 885M    0% /dev
tmpfs           tmpfs     894M    0 894M    0% /dev/shm
tmpfs           tmpfs     894M   18M 877M    2% /run
tmpfs           tmpfs     894M    0 894M    0% /sys/fs/cgroup
tmpfs           tmpfs     179M    0 179M    0% /run/user/2000
tmpfs           tmpfs     179M    0 179M    0% /run/user/0
tmpfs           tmpfs     179M    0 179M    0% /run/user/1001
/dev/xvda2      ext4      43G    51M   40G    1% /opt
```

----End

Setting Automatic Disk Mounting at a System Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify `/dev/xvdb1` in `/etc/fstab`. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in `/etc/fstab` to automatically attach the disk at a system start.

 **NOTE**

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

```
blkid Disk partition
```

For example, run the following command to query the UUID of `/dev/xvdb1`:

```
blkid /dev/xvdb1
```

Information similar to the following is displayed: (The UUID of `/dev/xvdb1` is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the **fstab** file using the vi editor:

```
vi /etc/fstab
```

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

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

```
UUID=xxx attachment directory file system defaults 0 2
```

Assuming that the file system is **ext4** and the attachment directory is **/mnt/sdc**.

```
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2
```

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the **fstab** file before detaching the disk. Otherwise, you may fail to access the OS after the disk is detached.

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

The system saves the configuration and exits the vi editor.

```
----End
```

10.4.7 Operations After Expanding Disk Capacity in Linux (Replacing Original Partitions Using fdisk)

This section describes how to perform follow-up operations for the additional disk space after you have expanded the capacity of an EVS disk and attached the EVS disk to an instance.

NOTICE

For details, contact the OS vendor of the target VM. This section uses EulerOS 2.0 (64-bit) to describe how to allocate the additional disk space to a partition using fdisk.

Prerequisites

- For details about how to log in to an ECS, see **Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS**.
- A disk has been attached to an instance and has not been initialized.

Procedure

The following example shows how to create new partitions to replace existing ones on a disk that has been attached to an instance. In this example, the disk partition is **/dev/xvdb1**, the partition style is MBR, and the mounting directory is **/mnt/sdc**. After capacity expansion, the disk capacity is less than 2 TB. You can use either the fdisk or parted tool to partition the disk space. This section uses the fdisk tool as an example.

Step 1 Run the following command to view the disk partition information:

fdisk -l

Information similar to the following is displayed: (In the command output, the server contains two disks. **/dev/xvda** is the system disk, and **/dev/xvdb** is the data disk.)

```
[root@ecs-b656 test]# fdisk -l

Disk /dev/xvda: 42.9 GB, 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: 0x000cc4ad

   Device Boot      Start         End      Blocks   Id  System
/dev/xvda1 *        2048     2050047     1024000    83  Linux
/dev/xvda2          2050048     22530047     10240000    83  Linux
/dev/xvda3          22530048     24578047      1024000    83  Linux
/dev/xvda4          24578048     83886079     29654016     5  Extended
/dev/xvda5          24580096     26628095      1024000    82  Linux swap / Solaris

Disk /dev/xvdb: 24.7 GB, 24696061952 bytes, 48234496 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: 0xb00005bd

   Device Boot      Start         End      Blocks   Id  System
/dev/xvdb1          2048     20971519     10484736    83  Linux
```

In the command output, parameter **Disk label type** indicates the disk partition style. Value **dos** indicates the MBR partition style, and value **gpt** indicates the GPT partition style.

View the **/dev/xvdb** capacity and check whether the additional space is included.

NOTE

The capacity displayed here is inconsistent with the capacity of the EVS disk applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 Bytes and 1 GB = 1,000,000,000 Bytes.

- If the additional space is displayed, go to [Step 2](#).
- If the new capacity is not displayed in the command output, run the following command to update the disk capacity:

```
echo 1 > /sys/class/scsi_device/%d:%d:%d:%d/device/rescan &
```

In the command, **%d:%d:%d:%d** indicates a folder in the **/sys/class/scsi_device/** directory and can be obtained using **ll /sys/class/scsi_device/**.

Information similar to the following is displayed: (**2:0:0:0** indicates the folder to be obtained.)

```
cs-xen-02:/sys/class/scsi_device # ll /sys/class/scsi_device/
total 0
lrwxrwxrwx 1 root root 0 Sep 26 11:37 2:0:0:0-> ../../devices/xen/vscsi-2064/host2/target2:0:0/2:0:0:0/
scsi_device/2:0:0:0
```

Example command:

```
echo 1 > /sys/class/scsi_device/2:0:0:0/device/rescan &
```

After the update is complete, run the **fdisk -l** command to view the disk partition information.

Step 2 Run the following command to unmount the disk:

```
umount /mnt/sdc
```

Step 3 Run the following command and enter **d** and then **1** to delete existing partition **/dev/xvdb1**. The partition to be deleted varies between scenarios.

```
fdisk /dev/xvdb
```

The command output is as follows:

```
[root@ecs-b656 test]# fdisk /dev/xvdb
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): d
Selected partition 1
Partition 1 is deleted

Command (m for help):
```

 **NOTE**

Deleting partitions will not cause data losses on the data disk.

Step 4 Enter **n** and press **Enter** to create a partition.

Entering **n** creates a partition.

There are two types of disk partitions:

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

```
Command (m for help): n
Partition type:
  p primary (0 primary, 0 extended, 4 free)
  e extended
```

Step 5 Ensure that the entered partition type is the same as that of the partition to be replaced. In this example, creating a primary partition is used. Therefore, enter **p** and press **Enter** to create a primary partition.

Information similar to the following is displayed: (**Partition number** indicates the serial number of the primary partition. The value can be **1** to **4**.)

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

Step 6 Ensure that entered partition number is the same as that of the partition to be replaced. In this example, the partition number **1** is used. Therefore, enter **1** and press **Enter**.

Information similar to the following is displayed: (**First sector** indicates the first sector number. The value can be **2048** to **20971519**, and the default value is **2048**.)

```
Partition number (1-4, default 1): 1
First sector (2048-20971519, default 2048):
```

Step 7 Press Enter.

The default start sector number is used as an example.

Information similar to the following is displayed: (**Last sector** indicates the last sector number. The value can be **2048** to **20971519**, and the default value is **20971519**.)

```
First sector (2048-20971519, default 2048):  
Using default value 2048  
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
```

Step 8 Press Enter.

The default last sector number is used as an example.

Information similar to the following is displayed, indicating that the partition is successfully created.

```
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):  
Using default value 20971519  
Partition 1 of type Linux and of size 10 GiB is set  
Command (m for help):
```

Step 9 Enter p and press Enter to view the details about the created partition.

Information similar to the following is displayed: (Details about the **/dev/xvdb1** partition are displayed.)

```
Command (m for help): p  
  
Disk /dev/xvdb: 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: 0xb00005bd
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvdb1		2048	20971519	10484736	83	Linux

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

Step 10 Enter w and press Enter to write the changes into the partition table.

Information similar to the following is displayed: (The partition is successfully created.)

```
Command (m for help): w  
The partition table has been altered!  
  
Calling ioctl() to re-read partition table.  
Syncing disks.
```

If the following error is displayed when you write partition results to the partition table, the new partition table will be updated upon the next OS restart.

```
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.
```

Step 11 Run the following commands to check and adjust the size of the file system on `/dev/xvdb1`:

e2fsck -f /dev/xvdb1

Information similar to the following is displayed:

```
[root@ecs-b656 test]# e2fsck -f /dev/xvdb1
e2fsck 1.42.9 (28-Dec-2013)
Pass 1: Checking inodes, blocks, and sizes
Pass 2: Checking directory structure
Pass 3: Checking directory connectivity
Pass 4: Checking reference counts
Pass 5: Checking group summary information
/dev/xvdb1: 11/655360 files (0.0% non-contiguous), 83137/2621184 blocks
```

resize2fs /dev/xvdb1

Information similar to the following is displayed:

```
[root@ecs-b656 test]# resize2fs /dev/xvdb1
resize2fs 1.42.9 (28-Dec-2013)
Resizing the filesystem on /dev/xvdb1 to 6029056 (4k) blocks.
The filesystem on /dev/xvdb1 is now 6029056 blocks long.
```

----End

10.4.8 Operations After Expanding Disk Capacity in Linux (Replacing Original Partitions Using parted)

This section describes how to perform follow-up operations for the additional disk space after you have expanded the capacity of an EVS disk and attached the EVS disk to an instance.

NOTICE

For details, contact the OS vendor of the target VM. This section uses EulerOS 2.0 (64-bit) to describe how to allocate the additional disk space to a partition using parted.

Prerequisites

- For details about how to log in to an ECS, see **Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS**.
- A disk has been attached to an instance and has not been initialized.

Procedure

The following example shows how to create a partition to replace the `/dev/xvdc1` partition mounted on `/mnt/sdc`. `/dev/xvdc1` is the only partition of the `/dev/xvdc` disk attached to an instance and uses the GPT partition style. There are two disks attached to the instance. During the partition creation, services will be interrupted.

NOTICE

After the disk capacity has been expanded, the additional space is added to the end of the disk. When the disk has multiple partitions, only the partition at the end of the disk can be expanded.

Step 1 Run the following command to view the disk partition information:

lsblk

Information similar to the following is displayed:

```
[root@ecs-1120 sdc]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0 0 80G 0 disk
├─xvda1 202:1 0 40G 0 part /
└─xvda2 202:2 0 40G 0 part /opt
xvdb 202:16 0 350G 0 disk
├─xvdb1 202:17 0 100G 0 part
└─xvdb2 202:18 0 200G 0 part
xvdc 202:32 0 60G 0 disk
└─xvdc1 202:33 0 10G 0 part /mnt/sdc
```

Indicates that the total capacity of the current data disk **/dev/xvdc** is 60 GB and the allocated partition capacity is 10 GB. The last partition is **/dev/xvdc1**, which is attached to the **/mnt/sdc** directory.

View the **/dev/xvdc** capacity and check whether the additional space is included.

- If the additional space is included, go to [Step 2](#).
- If the new capacity is not displayed in the command output, run the following command to update the disk capacity:

```
echo 1 > /sys/class/scsi_device/%d:%d:%d:%d/device/rescan &
```

In the command, **%d:%d:%d:%d** indicates a folder in the **/sys/class/scsi_device/** directory and can be obtained using **ll /sys/class/scsi_device/**.

Information similar to the following is displayed: (**2:0:0:0** indicates the folder to be obtained.)

```
cs-xen-02:/sys/class/scsi_device # ll /sys/class/scsi_device/
total 0
lrwxrwxrwx 1 root root 0 Sep 26 11:37 2:0:0:0-> ../../devices/xen/vscsi-2064/host2/target2:0:0/2:0:0:0/
scsi_device/2:0:0:0
```

Example command:

```
echo 1 > /sys/class/scsi_device/2:0:0:0/device/rescan &
```

After the update is complete, run the **fdisk -l** command to view the disk partition information.

Step 2 Run the following command to unmount the disk partition:

```
umount /mnt/sdc
```

Step 3 Run the following command to view the unmount result:

lsblk

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# umount /mnt/sdc
[root@ecs-1120 linux]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0 0 80G 0 disk
├─xvda1 202:1 0 40G 0 part /
```

```
└─xvda2 202:2 0 40G 0 part /opt
xvdb 202:16 0 350G 0 disk
└─xvdb1 202:17 0 100G 0 part
└─xvdb2 202:18 0 200G 0 part
xvdc 202:32 0 60G 0 disk
└─xvdc1 202:33 0 10G 0 part
```

Step 4 Run the following command to enter parted to allocate the additional space of the data disk to a partition:

```
parted Data disk
```

In this example, `/dev/xvdc` is the data disk.

```
parted /dev/xvdc
```

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# parted /dev/xvdc
GNU Parted 3.1
Using /dev/xvdc
Welcome to GNU Parted! Type 'help' to view a list of commands.
```

Step 5 Enter **unit s** and press **Enter** to set the measurement unit of the disk to sector numbers.

Step 6 Enter **p** and press **Enter** to view the current disk partition style.

Information similar to the following is displayed:

```
(parted) unit s
(parted) p
Error: The backup GPT table is not at the end of the disk, as it should be.
This might mean that another operating system believes the disk is smaller.
Fix, by moving the backup to the end (and removing the old backup)?
Fix/Ignore/Cancel? Fix
Warning: Not all of the space available to /dev/xvdb appears to be used,
you can fix the GPT to use all of the space (an extra 104857600 blocks)
or continue with the current setting?
Fix/Ignore? Fix
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdc: 125829120s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start      End          Size         File system Name  Flags
 1     2048s    20969471s  20967424s   ext4   opt
```

Partition Table specifies the partition style for existing disks. **msdos** indicates that the disk partition style is MBR, and **gpt** indicates that the disk partition style is GPT.

If the preceding information is displayed, enter **Fix** to rectify the disk exception. Then take note of the first and last sectors of the `/dev/xvdc1` partition. These values will be used during the partition recreation. In this example, the partition's first sector is **2048**, and its last sector is **20969471**.

The `/dev/xvdc1` partition number is **1**. Therefore, enter **rm 1** and press **Enter** to delete the partition.

Step 7 Enter **p** and press **Enter** to check whether the `/dev/xvdc1` partition has been deleted.

```
(parted) rm 1
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdc: 125829120s
```

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

Number Start End Size File system Name Flags
```

Step 8 Enter **mkpart opt 2048s 100%** and press **Enter** to create a partition.

opt indicates the name of the created partition, **2048s** indicates the first sector recorded in [Step 6](#), and **100%** indicates that all capacity of the disk is allocated to one partition, which must be greater than or equal to the last sector recorded in [Step 6](#).

Information similar to the following is displayed:

```
(parted) mkpart opt 2048s 100%
Warning: You requested a partition from 2048s to 125829199s (sectors 2048..125829199).
The closest location we can manage is 2048s to 125829036s (sectors 2048..125829036).
Is this still acceptable to you?
Yes/No? Yes
```

Enter **Yes** as prompted to set the last sector.

If the following warning message is displayed, enter **Cancel** to stop the partitioning. Then, find the first sector with the best disk performance and use that value to partition the disk. The warning message will not be displayed if the first sector with the best disk performance has been entered. In this example, **2048s** is one of such first sectors. Therefore, the system does not display the warning message.

```
Warning: The resulting partition is not properly aligned for best performance.
Ignore/Cancel? Cancel
```

NOTE

Data will be lost if the following operations are performed:

- Select a first sector which is inconsistent with that of the original partition.
- Select a last sector which is smaller than that of the original partition.

Step 9 Enter **p** and press **Enter** to check whether the **/dev/xvdc1** partition has been recreated.

```
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 125829120s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start End Size File system Name Flags
1 2048s 125829086s 125827039s ext4 opt
```

The **/dev/xvdc1** partition has been recreated.

Step 10 Enter **q** and press **Enter** to exit parted.

Step 11 Run the following command to view the disk partition information after the partition expansion:

lsblk

If information similar to the following is displayed, the total capacity of **/dev/xvdc** is 60 GB. The new 50 GB of the capacity is allocated to the **/dev/xvdc1** partition, and the partition is mounted on the **/mnt/sdc** directory. Skip [Step 12](#), [Step 14](#), and [Step 15](#).

```
[root@ecs-1120 sdc]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0 0 80G 0 disk
├─xvda1 202:1 0 40G 0 part /
└─xvda2 202:2 0 40G 0 part /opt
xvdb 202:16 0 350G 0 disk
├─xvdb1 202:17 0 100G 0 part
└─xvdb2 202:18 0 200G 0 part
xvdc 202:32 0 60G 0 disk
└─xvdc1 202:33 0 60G 0 part /mnt/sdc
```

- Step 12** Run the following command to check the correctness of the file system on **/dev/xvdc1**:

```
e2fsck -f /dev/xvdc1
```

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# e2fsck -f /dev/xvdb2
e2fsck 1.42.9 (28-Dec-2013)
Pass 1: Checking inodes, blocks, and sizes
Pass 2: Checking directory structure
Pass 3: Checking directory connectivity
Pass 4: Checking reference counts
Pass 5: Checking group summary information
/dev/xvdc1: 11/655360 files (0.0% non-contiguous), 83137/2620928 blocks
```

- Step 13** Run the following command to expand the size of the file system on **/dev/xvdc1**:

```
resize2fs /dev/xvdc1
```

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# resize2fs /dev/xvdc1
resize2fs 1.42.9 (28-Dec-2013)
Resizing the filesystem on /dev/xvdc1 to 15728379 (4k) blocks.
The filesystem on /dev/xvdc1 is now 15728379 blocks long.
```

- Step 14** Run the following command to view the disk partition information after the partition expansion:

```
lsblk
```

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0 0 80G 0 disk
├─xvda1 202:1 0 40G 0 part /
└─xvda2 202:2 0 40G 0 part /opt
xvdb 202:16 0 350G 0 disk
├─xvdb1 202:17 0 100G 0 part
└─xvdb2 202:18 0 200G 0 part
xvdc 202:32 0 60G 0 disk
└─xvdc1 202:33 0 60G 0 part
```

In the command output, the total capacity of the **/dev/xvdc** disk is 60 GB, in which the additional 50 GB has been allocated to the **dev/xvdc1** partition.

- Step 15** Run the following command to mount the created partition to the **/mnt/sdc** directory:

```
mount /dev/xvdc1 /mnt/sdc
```

- Step 16** Run the following command to view the attachment result of **/dev/xvdc1**:

```
df -TH
```

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# mount /dev/xvdc1 /mnt/sdc
[root@ecs-1120 linux]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda1     ext4      43G   8.3G   33G   21% /
devtmpfs       devtmpfs  885M   0   885M   0% /dev
tmpfs          tmpfs     894M   0   894M   0% /dev/shm
tmpfs          tmpfs     894M   18M   877M   2% /run
tmpfs          tmpfs     894M   0   894M   0% /sys/fs/cgroup
tmpfs          tmpfs     179M   0   179M   0% /run/user/2000
tmpfs          tmpfs     179M   0   179M   0% /run/user/0
tmpfs          tmpfs     179M   0   179M   0% /run/user/1001
/dev/xvda2     ext4      43G   51M   40G   1% /opt
/dev/xvdc1     ext4      64G   55M   60G   1% /mnt/sdc
```

The `/dev/xvdc1` partition has been mounted on the `/mnt/sdc` directory.

----End

10.4.9 Adding a Data Disk

You can expand disk capacity by adding data disks. This section describes how to add data disks to expand the available capacity of an instance.

Procedure

- Step 1** Apply for a data disk as required by services. For details, see section [10.1 Applying for a Data Disk](#).
- Step 2** Attach the data disk to the instance requiring capacity expansion. For details, see [10.2 Attaching an EVS Disk](#).
- Step 3** After the attachment is complete, initialize the data disk by referring to section [10.3 Initializing a Data Disk](#).

----End

10.5 Releasing an EVS Disk

10.5.1 Detaching an EVS Disk

This section describes how to detach an EVS disk from an ECS.

The following describes how to detach a data disk. The operations for detaching a system disk are the same.

You can detach a data disk in the following ways:

- ECS console: To detach multiple data disks from one ECS, perform related operations on the ECS console.
- EVS console: To detach a shared disk from multiple ECSs, perform related operations on the EVS console. In other scenarios, you can perform related operations on either console.

Prerequisites

If automatic attachment upon instance start has been configured for the Linux operating system by referring to [Setting Automatic Disk Attachment Upon](#)

Instance Start, comment out or delete the added content in the `/etc/fstab` file before detaching the disk. Otherwise, you may fail to access the operating system after the disk is detached.

Restrictions

- ECSs of the KVM virtualization type support online data disk detachment, namely, you can detach a data disk from an ECS in **Running** state when the following conditions are met:
 - The ECS is running steadily. You are not advised to detach a disk immediately after the ECS is started. Otherwise, the operation may fail.
 - For a VBD disk, **Disk Device Type** should be set to **scsi** or **virtio** when the operating system of the ECS is registered with Service OM.
 - The operating system of the ECS supports online disk detaching. For details, see the latest SIA compatibility document. Search for **Hot-plugging out disks** in the document to check whether the operating system supports online disk detaching.

NOTE

Obtain *FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)*:

- x86
 - Carrier users: [Click here](#). Search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)**.
 - Enterprise users: [Click here](#). Search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)**.
- Arm
 - Carrier users: [Click here](#). Search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (ARM)**.
 - Enterprise users: [Click here](#). Search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (ARM)**.
- ECSs of the Xen virtualization type support online data disk detachment. However, ECSs of the VMware virtualization type do not support online data disk detachment.
- System disks cannot be detached online.
- Before detaching a disk online from an instance running Windows, log in to the instance to perform the offline operation and confirm that UVP VMTools has been installed on the ECS and is running properly. At the same time, ensure that this disk is not being read and written. Otherwise, the disk will fail to be detached.
- Before detaching a disk online from an instance running Linux, log in to the instance, run the **umount** command to cancel the relationship between the disk and the file system, and confirm that the disk is not being read and written. Otherwise, the disk will fail to be detached.

ECS Console


Step 1 Log in to ManageOne as a VDC administrator or VDC operator using a browser.

URL in non-B2B scenarios: <https://Domain name of ManageOne Operation Portal>, for example, <https://console.demo.com>.

URL in B2B scenarios: **https://Domain name of ManageOne Tenant Portal**, for example, **https://tenant.demo.com**.

You can log in using a password or a USB key.

- Login using a password: Enter the username and password.
The password is that of the VDC administrator or VDC operator.
- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.

Step 2 Click  in the upper left corner, select a region and resource set, and choose **Storage > Elastic Volume Service**. The EVS console is displayed.

Step 3 In the ECS list, locate the ECS from which you want to detach the disk.

Step 4 Click the name of the ECS.

The page providing details about the ECS is displayed.


Step 5 On the **EVS** tab, select the data disk to be detached and click **Detach**.

Step 6 Click **OK**.

Step 7 After the detaching is complete, the EVS disk will no longer be displayed on the **EVS** tab of the ECS.

----End

Operations on EVS Console

Step 1 Click  in the upper left corner, select a region and resource set, and choose **Storage > Elastic Volume Service**. The EVS console is displayed.

Step 2 Click **Detach** in the row where the EVS disk to be detached is located.

If you are attempting to detach a shared disk, select the instance where the shared disk resides. You can select multiple instances.

Step 3 Click **OK**.

- If the disk being detached is a non-shared disk, the detachment is successful after the EVS disk state changes to **Available**.
- If a shared disk is detached, **Status** of the EVS disk becomes **Available** only after it is detached from all instances.

----End

Follow-up Procedure

After the EVS disk is detached, you can delete it by referring to [delete the disk](#) to release space.

10.5.2 Deleting an EVS Disk

If an EVS disk is no longer needed, you can soft delete it or delete it permanently. The soft deleted EVS disk is moved to the recycle bin and can be restore. A permanently deleted EVS disk cannot be restored.

The following describes how to delete a data disk. The operations for deleting a system disk are the same.

Restrictions

- If a disk has been attached to an instance, the disk cannot be deleted.
- If a disk has snapshots, the disk can be deleted only when the snapshot status is **Available** or **Error**.
- You can delete a disk only when the disk status is **Available**, **Error**, **Restoration failed**, or **Rollback failed**, and no VM snapshot has been created for the ECS where the disk resides.
- Disks configured with the DR service (CSDR/CSHA/VHA) cannot be deleted.
- If an EVS disk has a snapshot, the EVS disk can be soft deleted only when the snapshot is in the **Available** or **Error** state.
- When an EVS disk is permanently deleted, all snapshots of the EVS disk are also deleted.
- A shared disk to be deleted must have been detached from all instances.
- If the ECS to which the EVS disk belongs has not been created, the EVS disk cannot be deleted.
- Local disks can be used as data disks or system disks for an ECS. When a local disk is used as the system disk or a data disk, its life cycle starts and ends with the ECS, and cannot be manually detached or deleted.

Procedure

Step 1 Log in to EVS Console as a VDC administrator or VDC operator. For details, see [18.7.1 Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 On the **Elastic Volume Service** page, locate the row that contains the target data disk, click **More** in the **Operation** column, and choose **Delete**.

To delete multiple data disks at a time, select in front of the data disks on the **Elastic Volume Service** page.


Step 3 Delete the data disk.

- If you want to delete an EVS disk permanently, select **By selecting this option, the disk will be permanently deleted instead of being deleted to the recycle bin.** and click **OK**.

When an EVS disk is permanently deleted, all snapshots of the EVS disk are also deleted and cannot be restored.

- If you want to delete a data disk but require that the deleted data disk can be restored later, click **OK** to soft delete the data disk.

Delete Disk
✕



Are you sure you want to delete the following disk?
This operation will delete the disk to the recycle bin. You can choose 'Console > Recycle Bin' to restore or permanently delete the disk.
If a disk is permanently deleted, all the snapshots of the disk will be deleted. Exercise caution when you perform this operation.

Name	Capacity (GB)	Snapshots	Disk Type	Created	Status
██████████	50	0	FC_volume_t...	Dec 25, 2019 15:29:48	✔ Available

By selecting this option, the disk will be permanently deleted instead of being deleted to the recycle bin.

OK
Cancel



Step 4 If deleting a data disk requires approval, contact the administrator for approval. Otherwise, skip this step.

- After a data disk is permanently deleted, it is no longer displayed in the EVS disk list.
- After a data disk is soft deleted, the status of the data disk changes to **Soft deleted**.

----End

Follow-up Procedure

A deleted EVS disk is stored in the recycle bin and still occupies storage space.

- To restore data on a disk that has been **Soft deleted**, click  in the upper left corner of the page and choose **Recycle Bin**. On the **Recycle Bin** page, select the target EVS disk and click **Restore**.
- To permanently delete an EVS disk that has been **Soft deleted**, click  in the upper left corner of the page and choose **Recycle Bin**. On the **Recycle Bin** page, select the target EVS disk and click **Delete Permanently**.

11 ECS Group

11.1 Creating an ECS Group

An ECS group is a logical group with affinity, anti-affinity, weak affinity, or weak anti-affinity rules configured. The policy of the ECS group is used to schedule ECSs. If **Anti-affinity** is specified, ECSs in the ECS group will be created on different hosts. If **Affinity** is specified, ECSs in the ECS group will be created on the same host. If your ECSs need to maintain the affinity policy with other ECSs, create ECS groups by performing operations provided in this section and add the created ECSs to appropriate ECS groups.

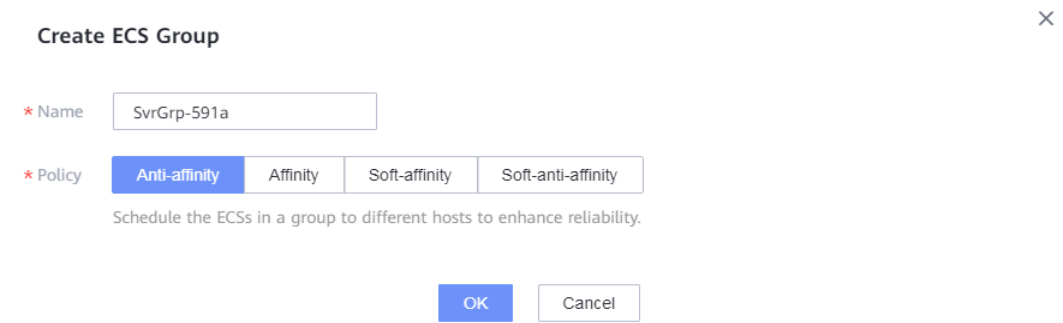
Context

When creating an ECS, you can add the ECS to an ECS group. Some existing ECSs can be added to ECS groups or moved to different groups. For details, see [11.2 Adding an ECS to or Removing an ECS from an ECS Group](#).

Procedure

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 2** In the navigation pane on the left, choose **ECS Group**.
- Step 3** click **Create ECS Group**.
- Step 4** Enter an ECS group name and specify a policy. The **Anti-affinity** policy is specified by default.

Figure 11-1 Creating an ECS group



When ECSs are being created, they are scheduled based on the policy of the ECS group to which they are added.

- If the policy of the ECS group is anti-affinity, the ECSs in a group are scheduled to different physical hosts to improve reliability.
- If the policy of the ECS group is affinity, the ECSs in a group are scheduled to the same physical host to improve performance and facilitate management.
- If the policy of the ECS group is weak anti-affinity, the ECSs in a group are scheduled to different physical hosts whenever possible. If the number of physical hosts is limited, the ECSs in a group may be scheduled to the same host.
- If the policy of the ECS group is weak affinity, the ECSs in a group are scheduled to the same physical host whenever possible. If the physical host resources are insufficient, the ECSs in a group may be scheduled to different hosts.

Step 5 Click **OK**.

Step 6 The ECS group is created.

 **NOTE**

When creating the ECS, expand **Advanced Settings** and select the target ECS group.

----End

11.2 Adding an ECS to or Removing an ECS from an ECS Group

Context


An ECS group is a logical group with affinity, anti-affinity, weak affinity, or weak anti-affinity rules configured. ECSs added to an ECS group will be scheduled to the same or different hosts according to the affinity rules of the group.

- If the policy of the ECS group is anti-affinity, the ECSs in a group are scheduled to different physical hosts to improve reliability.
- If the policy of the ECS group is affinity, the ECSs in a group are scheduled to the same physical host to improve performance and facilitate management.

- If the policy of the ECS group is weak anti-affinity, the ECSs in a group are scheduled to different physical hosts whenever possible. If the number of physical hosts is limited, the ECSs in a group may be scheduled to the same host.
- If the policy of the ECS group is weak affinity, the ECSs in a group are scheduled to the same physical host whenever possible. If the physical host resources are insufficient, the ECSs in a group may be scheduled to different hosts.

If you need to add an ECS to an ECS group, see [Adding an ECS to an ECS Group](#). If you need to remove an ECS from an ECS group, see [Removing an ECS from an ECS Group](#).

 **NOTE**

- For an ECS group with an affinity or anti-affinity policy configured, migrating a member of this group by specifying the destination host may invalidate the affinity or anti-affinity policy of the group.
- For an ECS group with an affinity policy configured, cross-AZ ECS migration invalidates the affinity policy of the ECS group.
- For an ECS group whose affinity or anti-affinity policy has been invalidated, the  icon indicating an exception is displayed in the **Policy** column. You can restore the affinity or anti-affinity policy of this ECS group by re-migrating the previously migrated ECS or removing this ECS from the group.

Adding an ECS to an ECS Group

An ECS can be added to only one ECS group in either of the following ways:

- You can add an ECS to an ECS group during ECS creation.
- You can add an ECS to an ECS group after the ECS is created.

Adding an ECS to an ECS group during ECS creation

For an ECS whose virtualization type is KVM, you can expand **Advanced Settings** and select the target ECS group during ECS creation.

 **NOTE**

- If the policy of the ECS group is affinity or anti-affinity, the ECS will fail to be created when existing hosts or resources are insufficient to fulfill the affinity or anti-affinity rules of the ECS group.
- ECSs whose virtualization type is KVM can be added to an ECS group configured with any affinity rule. ECSs of other virtualization types cannot be added to ECS groups.

Adding an ECS to an ECS group after the ECS is created

An ECS whose virtualization type is KVM can be added to an ECS group or moved to a different group after its creation. Adding an ECS to a group may cause the ECS to migrate. The restrictions on adding the created ECS to an ECS group are as follows:

- You must have the permission to add ECSs to ECS groups.
- General-purpose and general computing-plus ECSs can be added to ECS groups while they are in the **Running** or **Stopped** state. GPU-accelerated (computing-accelerated) ECSs or an ECS whose **Boot Device** is **Local Disk** can

be added to ECS groups only when the ECS is in the **Stopped** state. Other types of ECSs cannot be added to ECS groups.

- ECSs that have ongoing tasks cannot be added to ECS groups.
- If you add an ECS to a group that already contains ECSs and the new ECS is not in compliance with the affinity rule of the existing ECSs, the newly added ECS will need to migrate. If the ECS does not support migration, it cannot be added to the ECS group.
- If the target ECS group contains ECSs that have ongoing tasks, wait until the tasks are complete, and then proceed with the subsequent operation.
- The ECS group members cannot be added or removed in batches.

Step 1 Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 There are two places where you can add an ECS to an ECS group:

- ECS details page
 - a. Click the name of an ECS.
The ECS details page is displayed.
 - b. In the **ECS Group** area, click **Add to ECS Group**.
The **Add to ECS Group** dialog box is displayed.
 - c. Filter an ECS group by selecting the policy type or ECS group name, and click **OK**.
If information about the target ECS group is displayed on the ECS group tab page, the ECS group is successfully added. If the ECS group fails to be added, check whether the ECS meets the conditions for adding an ECS group by referring to [Added Restrictions](#).
- ECS group
 - a. In the navigation pane on the left, choose **ECS Group**.
 - b. Click **Add ECS** in the **Operation** column of the corresponding ECS group.
The **Add ECS** dialog box is displayed.
 - c. Filter an ECS by selecting the ECS status or name, and click **OK**.
If the target ECS is displayed in the ECS group, the ECS is successfully added to the ECS group. If the ECS group fails to be added, check whether the ECS meets the conditions for adding an ECS group by referring to [Added Restrictions](#).

----End

Removing an ECS from an ECS Group

Removing an ECS from an ECS group does not cause the ECS to migrate. You can remove the ECS from the ECS details page or ECS group list, but cannot remove multiple ECSs in batches.

- ECS details page
 - a. Click the name of an ECS.
The ECS details page is displayed.

- b. In the **ECS Group** area, click **Remove from ECS Group**.
- c. In the displayed dialog box, click **OK**.

If no ECS group information is displayed on the ECS group tab page, the ECS group is removed successfully.

- ECS group
 - a. In the navigation pane on the left, choose **ECS Group**.
 - b. Click the icon on the left of the corresponding ECS group to view details about the ECS group.
 - c. Click **Remove from ECS Group** in the **Operation** column of the corresponding ECS group.
 - d. Click **OK**.

If the ECS is not displayed in the ECS group, the ECS group is removed successfully.

Related Operations

For details about how to create an ECS group, see [11.1 Creating an ECS Group](#).

12 Network and Security

A VPC provides network resources, including subnets, security groups, network ACLs, virtual private networks (VPNs), and elastic IP addresses (EIPs), for ECSs. You can create security groups, configure security group policies, create network ACLs, add network ACL rules, associate subnets with network ACLs, and create VPC peering connections to enable communication or isolation between ECSs in the same VPC or between VPCs.

12.1 Configuring Intra-VPC Communication and Security Policy for ECSs

By default, all subnets in a VPC can communicate with each other. The IP address of each ECS belongs to a subnet in the VPC. You can associate a subnet with a network ACL and configure network ACL rules to enable or disable communication between subnets.

To enable communication between ECSs, you must ensure that the security groups and subnets to which the ECSs belong can communicate with each other.

- For details about how to create a security group and configure a security group rule, see "Security Group" in **Operation Help Center > Network > Virtual Private Cloud > User Guide**.
- For details about how to create a network ACL and configure a network ACL rule, see "Applying for a Network ACL" and "Managing a Network ACL" in **Operation Help Center > Network > Network ACL > User Guide**.

12.2 Enabling Communication Between an ECS and the Internet and Configuring Security Policies

To enable your ECS to communicate with the Internet, you can perform any of the following operations:

- Binding an EIP
An EIP is a static, public IP address. You can bind an EIP to an ECS in your subnet or unbind it from the ECS. An EIP enables Internet connectivity for an ECS in your VPC through a fixed public IP address.

For details, see "Creation" > "Binding an EIP" in **Operation Help Center > Network > Elastic IP > User Guide**.

Public IP addresses are scarce resources. An ECS not bound with a public IP address can access the Internet through an ECS bound with a public IP address in the same subnet. For details, see **18.6.7 Accessing the Internet Using an ECS Without a Public IP Address**.

- Some ECSs not only require services provided by the system but also need to access the Internet to obtain information or download software. The system allows users to bind EIPs to virtual NICs (ports) of ECSs to enable the ECSs to access the Internet. However, assigning a public IP address to each ECS consumes already-limited IPv4 addresses, incurs additional costs, and may increase the attack surface for a virtual environment. Therefore, enabling multiple ECSs to share one public IP address is a preferable and more feasible method. This can be done using .

For details about how to configure NAT, see "NAT Gateway" in **Operation Help Center > Network > Virtual Private Cloud > User Guide**.

- Configuring a VPN

A VPN establishes an encrypted communication tunnel between a remote user and a VPC, enabling the remote user to use service resources in the VPC through the VPN.

For details about how to configure a VPN, see "Creation" > "Applying for a VPN" in **Operation Help Center > Network > Virtual Private Network > User Guide**.

12.3 Modifying NIC Configurations

You can add a NIC to an ECS, delete a NIC from an ECS, and change the security group to which a NIC belongs.

Adding a NIC

Step 1 Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see **18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?**

Step 2 Go to the page for adding a NIC in either of the following ways:

Method 1: On the **Elastic Cloud Server** page, locate the row that contains the target ECS, choose **More > NIC > Add NIC** in the **Operation** column.

Method 2: Click the name of the target ECS. On the displayed ECS details page, click **NIC** and then **Add NIC**.

Step 3 Configure the NIC information.

- Type
 - The type of the network to which a NIC belongs can be **VPC Subnet** and **Intra-Project Subnet**.
 - If you select **VPC Subnet**, the **Subnet** option provides all subnets in the VPC. In this case, the NIC supports layer 3 communication, allowing the

ECS to communicate with networks (for example, the public network or other VPCs) beyond the VPC.

- If you select **Intra-Project Subnet**, the **Subnet** option provides all intra-project subnets in the project. If the ECS needs to communicate with ECSs in other VPCs at layer 2, select **Intra-Project Subnet**.
- Security Group
Select the security group for this NIC. You can select multiple security groups. In such a case, the access rules of all the selected security groups apply to the ECS.
- Subnet
Select the subnet where the NIC resides.

 **NOTE**

If the selected subnet does not have DHCP enabled and the ECS image does not support static IP address injection, after a NIC is added, you need to manually configure an IP address for it within the ECS. Otherwise, the NIC cannot be reached. For details, see [18.6.1 Configuring a Static IP Address for an ECS](#).

- IP Address
Select **Automatically Assign** or **Manually Assign** for the NIC parameter.
 - IPv4 address
You can select an available IP address from the selected subnet to use as the NIC IP address. Alternatively, the NIC IP address can be automatically assigned by the system.
 - IPv6 address
If IPv6 is enabled and the selected subnet has both an IPv4 address segment and an IPv6 address segment, this parameter is displayed. You can select an available IP address from the selected subnet to use as the NIC IP address. Alternatively, the NIC IP address can be automatically assigned by the system.

Step 4 Click **OK**.

 **NOTE**

- An ECS can have only one primary NIC, and the number of extension NICs cannot exceed 15.
- When ARM servers are used, if the ECS bus type is Virtio, the total number of NICs, disks (EVS disks and pass-through disks), GPU cards, and NPU cards cannot exceed 24.

----End

Follow-up Operations

Task 1: Activating a NIC

Some OSs cannot identify newly added NICs. In this case, you must manually activate the NICs. The following uses Ubuntu and a NIC that resides on a DHCP-enabled subnet as an example to describe how to activate the NIC. Required operations may vary depending on the OS. For additional information, see the documentation for your OS.

Step 1 Find the target ECS and click **Remote Login** in the **Operation** column.

Log in to the ECS.

Step 2 Run the following command to query the NIC name:

```
ifconfig -a
```

In this example, the NIC name is **eth2**.

Step 3 Run the following command to switch to the target directory:

```
cd /etc/network
```

Step 4 Run the following command to open the **interfaces** file:

```
vi interfaces
```

Step 5 Add the following information to the **interfaces** file:

```
auto eth2
```

```
iface eth2 inet dhcp
```

Step 6 Run the following command to save and exit the **interfaces** file:

```
:wq
```

Step 7 Run the **ifup ethX** or **/etc/init.d/networking restart** command to activate the newly added NIC.

X in the preceding command indicates the NIC name obtained in [Step 2](#), for example, **ifup eth2**.

Step 8 Run the following command to check whether the NIC name obtained in [Step 2](#) is displayed in the command output:

```
ifconfig
```

- If yes, the newly added NIC has been activated, and no further action is required.
- If no, the newly added NIC has failed to be activated. Go to [Step 9](#).

Step 9 Log in to ManageOne Operation Portal. Locate the row that contains the target ECS, click **More** in the **Operation** column, and then click **Restart**.

Step 10 Run the **ifconfig** command, and check whether the NIC name obtained in [Step 2](#) is displayed in the command output.

- If yes, no further action is required.
- If no, contact the system administrator.

----End

Task 2: Configuring Routes

If the static route function is not enabled for the selected subnet when you add a NIC, the NIC uses the default route of the ECS. If the static route function is enabled, the NIC uses the static route configured for the subnet during subnet creation. You can view the route information of the subnet on the subnet details page.

To modify the route information of a NIC, log in to the ECS first.

Task 3: Other Configurations

If an ECS uses a NIC to provide external services or communicate with other components, configure a security group and network ACL for the NIC to ensure proper security control.

Deleting a NIC

CAUTION

Deleting a NIC may cause the following problems. Exercise caution when performing this operation.

- The ECS uses the NIC to provide external services. Deleting the NIC may interrupt ECS access due to IP address changes.
- The ECS provides services in cluster or HA mode. Deleting the NIC may cause service interruption due to IP address changes.
- For most OSs, after a NIC is deleted, the custom route information of the NIC is also deleted, which may cause the loss of network connectivity.
- After a NIC is deleted, components associated with the MAC address of the NIC, such as the software license bound to the MAC address, cannot run properly.

Step 1 Click the ECS name to go to the ECS details page.

Step 2 Click the **NIC** tab.

The NIC details page of the ECS is displayed.

Step 3 Click **Delete** in the row of the target NIC. In the dialog box that is displayed, click **OK**.

NOTE

You are not allowed to delete the ECS primary NIC. By default, the ECS primary NIC is the first NIC displayed in the NIC list.

Step 4 If the NIC still exists after the NIC tab page is refreshed several minutes later, shut down the ECS and delete the NIC again.


----End

Changing the security group to which a NIC belongs

NOTICE

If you want to change the security group to which a NIC belongs, ensure that the inbound and outbound rules of the destination security group have been correctly configured. Otherwise, ECS communication may be interrupted after the security group is changed.

Step 1 Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID, and click . Alternatively, click **Search by Tag** above the upper right corner of the ECS list and search for an ECS by tag key and value.

Step 3 Click the name of the ECS.

The page providing details about the ECS is displayed.

Step 4 Click the **NIC** tab.

Step 5 Click **Change Security Group** in the row of the target NIC.


Step 6 In the displayed **Change Security Group** dialog box, select the target security group.

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

 **NOTE**

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

Step 7 Click **OK**.

The security group has been changed. Click  on the left of the NIC IP address to view the changed security group. Click the security group name to view details about the security group rules.

----End

Managing the private IP address of a NIC

 **CAUTION**

Changing the private IP address of a NIC may cause the following problems. Exercise caution when performing this operation.

- The ECS uses the NIC to provide external services. Changing the IP address of the NIC may interrupt ECS access.
- The ECS provides services in cluster or HA mode. Changing the IP address of the NIC may cause service interruption.

Step 1 On the ECS details page, click the **NIC** tab.

Step 2 Locate the row that contains the NIC whose private IP address needs to be changed, click **Manage IP Address**.

 **NOTE**

- If the **Manage IP Address** option is unavailable, you do not have the permission to manage IP addresses. Contact the administrator to change your permissions.
- If your ECS uses a shared VPC, you cannot redirect to the VPC details page by clicking **Manage IP Address**.

To change the private IP address of a NIC (prerequisite being that you have the permission to do so), manually switch the resource set in the upper left corner of the page, choose **Network > Virtual Private Cloud** on the console, locate the VPC to which the ECS belongs, and go to the VPC details page. Click the subnet where the ECS resides. On the **In-use IP Address** tab, perform [Step 3](#) to [Step 5](#).

Step 3 On the **In-use IP Address** tab page, locate the row that contains the IP address of the NIC, and click **Modify**.

The **Modify IP Address** dialog box is displayed.

Step 4 Change the private IP address of the NIC to another IP address within the IP address range of the subnet. Click **OK**.

Step 5 Perform the following operations based on whether DHCP is enabled for the subnet:

- If DHCP is enabled for the subnet, you need to restart the ECS to obtain a new IP address after changing the private IP address, or you can log in to the ECS and manually change the IP address.
- If DHCP is not enabled for the subnet, you need to log in to the ECS and manually change the IP address. For details, see [18.6.1 Configuring a Static IP Address for an ECS](#).

 **NOTE**

If you do not restart the ECS or manually change the IP address of the ECS, the ECS NIC will be unreachable.

----End

12.4 Configuring Security Group Rules

You can configure inbound and outbound access rules of a security group to function as access control policies for ECSs in the security group. This section describes how to configure security group rules.

Context

The default security group policy allows data exchange only within a security group. The ECSs in a security group cannot be accessed from an external network. Therefore, if you need to access an ECS in a security group externally, for example, using remote login, you must configure the inbound rules of the security group.

By default, a security group allows all outbound data packets from ECSs in the security group. If you need to limit access from ECSs in a security group to the Internet, configure outbound access rules of the security group.

 **NOTE**

When an ECS is added to or removed from a security group, applications running on the ECS may be interrupted. Therefore, exercise caution when performing this operation.

Procedure

Adding inbound rules

Step 1 Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 In the ECS list, click the name of the ECS whose security group rules you want to modify.

The page providing details about the ECS is displayed.

Step 3 Click the **Security Groups** tab and click the icon before the security group to view the security group rules.

Step 4 Click the security group ID.

The system automatically switches to the **Security Group** page.

Step 5 Click the security group whose rules are to be modified. Then, click **Delete** in the **Operation** column of the **Inbound** row to delete the inbound rule.

Step 6 Click **Add Rule** to add an inbound rule for the security group.

 **NOTE**

Assume that you have permission to remotely log in to the ECS.

- **Protocol and Port Range**
 - To remotely access a Windows ECS, set **Protocol** to **TCP** and **Port Range** to **3389**.
 - To remotely access a Linux ECS, set **Protocol** to **TCP** and **Port Range** to **22**.
- **Type:** The options include **IPv4** and **IPv6**. This parameter must be set when IPv6 is enabled on the platform.
- **Source:** Select **IP Address Range** and enter an IP address range, or select **Security Group** and set a security group.

Table 12-1 Parameter description

Rule Type	IP Address Range	Security Group
IPv4	If Source is set to IP Address Range , the configured IPv4 address segment or IP address can be used to access the ECSs in the set security group.	If Source is set to Security Group , the IPv4 addresses of all ECSs in the selected security group can be used to access the ECSs in the set security group.

Rule Type	IP Address Range	Security Group
IPv6	If Source is set to IP Address Range , the configured IPv6 address segment or IP address can be used to access the ECSs in the set security group.	If Source is set to Security Group , the IPv6 addresses of all ECSs in the selected security group can be used to access the ECSs in the set security group.

 **NOTE**

If the source IP address is set to 0.0.0.0/0 or 0:0:0:0:0:0/0 by default, all IP addresses can be used to access ECSs in the set security group.

Step 7 Click **OK**.

----End

Adding outbound rules

Step 1 On the **Outbound** tab page of the security group details page, click **Add Rule** to add an outbound access rule for the security group.

- **Protocol:** Customize a protocol. It specifies the protocol through which ECSs in the security group communicate with the Internet.
- **Port Range/ICMP Type:** Customize the value. It specifies the access destination port of the ECS in the security group.
- **Type:** The options include **IPv4** and **IPv6**. This parameter must be set when IPv6 is enabled on the platform.
- **Destination:** Customize an IP address range or security group. It specifies a network segment or security groups that ECSs in the security group are allowed to access.

Table 12-2 Parameter description

Rule Type	IP Address Range	Security Group
IPv4	If Destination is set to IP Address Range , the configured IPv4 address segment or IP address can be used to access the ECSs in the set security group.	If Destination is set to Security Group , the IPv4 addresses of all ECSs in the selected security group can be used to access the ECSs in the set security group.
IPv6	If Destination is set to IP Address Range , the configured IPv6 address segment or IP address can be used to access the ECSs in the set security group.	If Destination is set to Security Group , the IPv6 addresses of all ECSs in the selected security group can be used to access the ECSs in the set security group.

Step 2 Click **OK**.

----End

12.5 Changing the EIP

You can assign an EIP, bind an EIP to an ECS, unbind an EIP from an ECS, and release an EIP.

Context

An EIP is a static, public IP address. You can bind an EIP to an ECS in your subnet or unbind it from the ECS. An EIP enables Internet connectivity for an ECS in your VPC through a fixed public IP address.

NOTE

If the system displays a message indicating that you do not have the permission, contact the administrator to change your permissions.

Procedure

Binding an EIP

Step 1 Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 Go to the page for binding an EIP in either of the following ways:

Method 1: On the **Elastic Cloud Server** page, locate the row that contains the target ECS, choose **More** > **NIC** > **Bind EIP** in the **Operation** column.

Method 2: Click the name of the target ECS. On the displayed ECS details page, click **EIP** and then **Bind EIP**.

Step 3 Select the NIC to which you want to bind the EIP and the available EIP, and click **OK**.

NOTE

- One NIC can be bound to only one EIP. To bind another EIP, add a new NIC on the **NIC** tab page of the ECS details page.
- If no EIP is available, click **View EIP** in the lower left corner of the **Bind EIP** page, or click **View EIP** on the **EIP** tab page of the ECS details page. For details about how to apply for and bind an EIP, see "Creation" in **Operation Help Center** > **Network** > **Elastic IP** > **User Guide**.

----End

Unbinding an EIP

CAUTION

If an ECS provides services externally, unbinding an EIP from the ECS may interrupt external services. Exercise caution when performing this operation.

Step 1 On the **EIP** tab page of the ECS details page, click **Unbind** on the right of the NIC to be unbound.

The **Unbind EIP** dialog box is displayed.

Step 2 Click **OK**.

----End

12.6 Binding a Floating Private IP Address (Virtual IP Address)

A floating private IP address provides a second IP address for a NIC. This IP address can be bound to the NIC using the NIC subinterface or multi-IP address function. This enables more flexible network functions, such as the floating IP address for HA hot standby and VIP for LVS.

Procedure

Step 1 Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 In the ECS list, click the name of the target ECS.


The page providing details about the ECS is displayed.

Step 3 Click the **NIC** tab.

Step 4 If your ECS uses a shared VPC, you cannot redirect to the VPC details page by clicking **Manage IP Address**. Perform the following operations based on whether the VPC to which the ECS resides is a shared VPC:

- If yes, go to [Step 5](#).
- If no, click **Manage IP Address**. The page providing details about the subnet where the NIC resides is displayed. In this case, go to [Step 9](#).

Step 5 View the VPC to which the ECS belongs, and click the icon on the left of the NIC to view the subnet where the NIC resides.

Step 6 Click  in the upper left corner of the page and choose **Network > Virtual Private Cloud**.

Step 7 In the upper left corner of the page, manually switch to the resource set where the shared VPC resides.

Step 8 Click the VPC to which the ECS belongs. On the **Subnet** tab of the VPC details page, click the subnet obtained in [Step 5](#).

Step 9 On the **Virtual IP Address** tab page, click **Assign Virtual IP Address**.

Step 10 Select a virtual IP address assignment mode, and configure the parameters as required.

- If the subnet is an IPv4 subnet, the system assigns or you specify an IPv4 address.

- If the subnet is an IPv4&IPv6 one and the IP address assignment mode is **Automatic**, an IPv4 or IPv6 address is assigned depending on the IP address type you select.
- If the subnet is an IPv4&IPv6 one and the IP address assignment mode is **Manual**, set the IP address type to **IPv4** or **IPv6**, and manually specify an IP address of the corresponding type.

Step 11 Locate the row that contains the obtained virtual IP address, and click **Bind to Cloud Server** in the **Operation** column.

The **Bind to Cloud Server** page is displayed.

Step 12 Select the target ECS and NIC, and click **OK**.

 **NOTE**

- If an ECS is deployed in cluster or HA mode to provide services, ensure that the correct ECS or NIC is selected. Otherwise, services may be unavailable.
- After a virtual IP address is bound to an ECS, if the virtual IP address cannot be pinged, check and configure the virtual IP address by referring to [18.6.8 What Do I Do If the Virtual IP Address Cannot Be Pinged After Being Bound to the ECS NIC?](#). After a virtual IP address is bound to an ECS NIC, you can use an EIP, VPN, Direct Connect (Region Type I), or VPC peering connection to access the virtual IP address. How to use the virtual IP address to deploy HA for multiple active and standby ECSs is determined by upper-layer applications.

----End

13 Operating Systems

After the ECS is successfully created, you can reinstall or change the OS of the ECS as required.

13.1 Reinstalling an ECS OS

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

Constraints

- The system disk quota must be greater than 0.
- Before the OS reinstallation, make sure that the original image is still available. For public and private images, make sure that the images are not deleted. For shared images, make sure that image sharing is not canceled.
- If the ECS image supports static IP address injection, the OS cannot be reinstalled on the ECS.
- Back up data before the reinstallation.
 - Make sure that all necessary data has been backed up. Reinstalling the OS clears the data in all partitions of the system disk, including the system partition. After the OS is reinstalled, the existing system disk snapshot is deleted, but the data disk snapshot is not affected.
 - For Windows, if the system disk and data disks use dynamic disks (volume types: simple, spanned, striped, mirrored, and RAID 5), reinstalling the operating system may damage the logical volumes or cause the OS to become faulty. In this case, back up important data first.
 - For Linux, if Logical Volume Manager is configured for the system disk and data disks and the disks belong to the same volume group, reinstalling the operating system may damage the logical volumes or cause the OS to become faulty. In this case, back up important data first.
- The OS cannot be reinstalled on an ECS with ECS snapshots.
- If **Disk Device Type** of the OS has been changed, you cannot reinstall the OS for the ECS.
- The OSs of ECSs configured with VHA, CSHA, CSDR, or VHA+CSDR protection, and ECSs in MRS clusters cannot be reinstalled.

Prerequisites

- The target ECS is in the **Stopped** or **Reinstallation failed** state.
- The target ECS has an EVS system disk attached.

Procedure

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 2** Locate the row containing the target ECS. In the **Operation** column, choose **More > OS and Image > Reinstall OS**.
- Step 3** Determine whether to select **Joint Windows Domain**. This parameter is available if the virtualization type of the AZ where the ECS resides is KVM, the ECS is running a Windows OS, and the domain information has been configured for the corresponding ECS product.

The administrator can perform unified authentication for ECSs added to the same domain. The following functions will be available for ECSs added to a domain: manage compute resources, reduce network management complexity and costs, enhance security, and support account roaming and folder redirection. Resources can be shared among ECSs in the same domain. For more information about domain servers and their functions, click [here](#).

An ECS can be added to a domain when the following requirements are met:

- The image is a non-static injection image and Cloudbase-Init is installed.
- The ECS name cannot contain Chinese characters, can contain a maximum of 15 characters, and cannot be the same as that of an existing ECS to be added to a domain.

Specify whether to add an ECS to a Windows domain. You can select a domain from the drop-down list. Available options are those defined by the administrator during product creation.

- Step 4** Set a key or new password. Select the image password as the ECS password or reset a new password or key for the ECS. If **Set Key or New Password** is set to **Yes**, you need to select **Key Pair** or **Password** for **Login Mode**.

NOTE

This parameter is displayed when all of the following conditions are met:

- The image must have the Cloud-Init (Linux OS) plugin installed, and Cloud-Init is selected during image registration.
- The image password can be used as the ECS password.
- The image password cannot be used as the ECS password in the Windows OS.

- Step 5** Enter the password or select a key based on the selected login mode.

NOTE

If Cloud-Init or Cloudbase-Init is not installed in the original image, you cannot select the login mode and configure the login password or key during the OS reinstallation.

- Step 6** Click **OK**.

After the application is submitted, the ECS status changes to **Reinstalling**. The reinstallation has been successfully completed when the ECS status changes to **Running**.

 **NOTE**

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

----End

Follow-up Procedure

If the OS reinstallation is unsuccessful, perform [Step 2](#) to [Step 6](#) again.

If the second reinstallation attempt is unsuccessful, contact the system administrator.

13.2 Changing the ECS OS

This section describes how to change the image and OS of an ECS.

Context

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

The platform allows changing between image types (public, private, and shared images) and between OSs (between Windows OSs, between Linux OSs, or between a Windows OS and a Linux OS).

 **NOTE**

- After the OS is changed, the original OS is not retained and the original system disk is deleted.
- Changing the OS clears the data in all partitions of the EVS system disk, including the system partition. Back up useful data in advance.

Constraints

- The EVS system disk quota must be greater than 0.
- After the OS is changed, the system disk capacity may increase because a different image is used.
- If the ECS image supports static IP address injection, the OS cannot be changed for the ECS.
- The OS cannot be changed on an ECS with ECS snapshots created.
- Back up data before changing.
 - Make sure that all necessary data has been backed up. Changing the OS clears the data in all partitions of the system disk, including the system partition. After the OS is changed, the existing system disk snapshot is deleted, but the data disk snapshot is not affected.
 - For Windows, if the system disk and data disks use dynamic disks (volume types: simple, spanned, striped, mirrored, and RAID 5), changing

- the operating system may damage the logical volumes or cause the OS to become faulty. In this case, back up important data first.
- For Linux, if Logical Volume Manager is configured for the system disk and data disks and the disks belong to the same volume group, changing the operating system may damage the logical volumes or cause the OS to become faulty. In this case, back up important data first.
 - You cannot switch between OSs of different disk device types. The OS displayed on the Web UI is the one that meets the restriction.
 - When you change the OS for an ECS residing on an x86 server, all OSs whose registered architecture is x86 are displayed on the Web UI. In the case of an ARM server, all OSs whose registered architecture is ARM are displayed.
 - During OS change, the system filters image files based on the ECS specifications.
 - For a flavor whose **Boot Device** is set to **Cloud Disk**, an image is displayed here only when its **Min Memory (MB)** specified during image registration is less than or equal to the selected memory.
 - For a flavor whose **Boot Device** is set to **Local Disk**, an image is displayed here only when its **Min Memory (MB)** specified during image registration and the minimum disk required by the image are less than or equal to the **Memory** and **Root Disk (GB)** of the selected flavor, respectively.
 - If you select **BIOS** or **UEFI** as the boot mode, the image files that use this boot mode will be displayed. If the boot mode parameter is not displayed, it indicates that all the image files use **BIOS** as the default boot mode.
 - The OSs of ECSs configured with VHA, CSHA, CSDR, or VHA+CSDR protection, and ECSs in MRS clusters cannot be changed.

Prerequisites

- The target ECS is in **Stopped** or **OS change failed** state.
- The target ECS has an EVS system disk attached.

Procedure

Step 1 Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 Locate the row containing the target ECS. In the **Operation** column, choose **More > OS and Image > Change OS**.

The **Change OS** page is displayed.

Step 3 Select the specifications (including **Image Type** and **Image**) of the ECS to be replaced. For some images, you need to configure **Boot Mode**.

For more details, see section [5.2 Applying for an ECS](#).

 **NOTE**

- If the ECS has NICs that are using static IP addresses, but the new image you selected to change the ECS OS does not support static IP address injection, after the OS change, you need to manually configure static IP addresses within the ECS. Otherwise, the NICs cannot be reached. For details, see [18.6.1 Configuring a Static IP Address for an ECS](#).
- If the ECS image does not support static IP address injection, the system filters the image files and displays only the images that do not support static IP address injection.

Step 4 Determine whether to select **Joint Windows Domain**. This parameter is available if the virtualization type of the AZ where the ECS resides is KVM, the ECS is running a Windows OS, and the domain information has been configured for the corresponding ECS product.

The administrator can perform unified authentication for ECSs added to the same domain. The following functions will be available for ECSs added to a domain: manage compute resources, reduce network management complexity and costs, enhance security, and support account roaming and folder redirection. Resources can be shared among ECSs in the same domain. For more information about domain servers and their functions, click [here](#).

An ECS can be added to a domain when the following requirements are met:

- The image is a non-static injection image and Cloudbase-Init is installed.
- The ECS name cannot contain Chinese characters, can contain a maximum of 15 characters, and cannot be the same as that of an existing ECS to be added to a domain.

Specify whether to add an ECS to a Windows domain. You can select a domain from the drop-down list. Available options are those defined by the administrator during product creation.

Step 5 Set a key or new password. Select the image password as the ECS password or reset a new password or key for the ECS. If **Set Key or New Password** is set to **Yes**, you need to select **Key Pair** or **Password** for **Login Mode**.

 **NOTE**

This parameter is displayed when all of the following conditions are met:

- The image must have the Cloud-Init (Linux OS) plugin installed, and Cloud-Init is selected during image registration.
- The image password can be used as the ECS password.
- The image password cannot be used as the ECS password in the Windows OS.

Step 6 Enter the password or select a key based on the selected login mode.

 **NOTE**

If the selected image does not have Cloud-Init installed, you cannot select the login mode and configure the login password or key during OS reinstallation.

Step 7 Click **OK**.

After the application is submitted, the ECS status changes to **Changing OS**. The OS changing is completed when the ECS status changes to **Running**.

 **NOTE**

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

----End

Follow-up Procedure

If the OS change is unsuccessful, perform [Step 2](#) to [Step 7](#) again to try changing the OS again.

If the second attempt also fails, contact technical support.

14 Monitoring Metrics

14.1 ECS Monitoring Metrics

You can customize alarm rules, monitored objects, and notification policies for ECSs to monitor their status adequately. In addition, you can use the ECS metrics, such as CPU usage, memory usage, disk usage, and disk read and write rates, to get details about the performance and resource usage of ECSs and adjust their flavors to better meet your requirements.

Supported ECS monitoring metrics are listed in [Table 14-1](#). You can view the monitoring metrics of an ECS by clicking the **Monitoring** tab on the ECS details page.

Table 14-1 ECS monitoring metrics

Metric	Description	Formula	Remarks
CPU Usage	Indicates the CPU usage (%) of an ECS.	Total time used by each CPU core on an ECS/(Monitoring period x Number of vCPU cores of an ECS)	N/A
Memory Usage	Indicates the memory usage (%) of an ECS.	Used memory of an ECS/Total memory of the ECS	This metric is unavailable if the image has no UVP VMTools installed.
Network Inbound Rate	Indicates the number of incoming bytes on an ECS per second.	Total number of incoming bytes on an ECS/Monitoring period	This metric is unavailable if the image has no UVP VMTools installed.

Metric	Description	Formula	Remarks
Network Outbound Rate	Indicates the number of outgoing bytes on an ECS per second.	Total number of outgoing bytes on an ECS/Monitoring period	This metric is unavailable if the image has no UVP VMTools installed.
EVS Disk Usage	Indicates the internal disk usage (%) of an ECS.	Used capacity of an ECS disk/Total capacity of the ECS disk	This metric is unavailable if the image has no UVP VMTools installed.
EVS Disk I/O Write	Indicates the disk read rate of an ECS (bytes/s).	Total number of bytes written to disks/Monitoring period	N/A
EVS Disk I/O Read	Indicates the disk write rate of an ECS (bytes/s).	Total number of bytes read from disks/Monitoring period	N/A
Disk Read Requests	Indicates the number of read requests sent to an ECS per second.	Total number of read requests sent to an ECS disk/Monitoring period	N/A
Disks Write Requests	Indicates the number of write requests sent to an ECS per second.	Total number of write requests sent to an ECS disk/Monitoring period	N/A
GPU Usage	Indicates the current CPU usage (%) of an ECS.	-	This metric is unavailable if TenantPMAgent is not installed.
GPU Memory Usage	Indicates the current memory usage (%) of an ECS.	-	This metric is unavailable if TenantPMAgent is not installed.
GPU Performance Status	Indicates the current GPU performance status of an ECS. This metric has no unit.	-	This metric is unavailable if TenantPMAgent is not installed.

Metric	Description	Formula	Remarks
Chip Memory Usage	Indicates the memory usage (%) of a chip.	Chip memory usage = (Total memory - Idle memory)/Total memory of the AI chip operating system	N/A
AI Core Usage	Indicates the AI core usage (%).	AI core usage = $(t_1 + t_2 + \dots + t_n)/3(s)$. <i>tn</i> indicates the interval between each BS delivery and completion.	N/A
AI CPU Usage	Indicates the AI CPU usage (%).	$(\text{Total number of AI CPU running cycles} - \text{Number of AI CPU idle cycles}) / \text{Total number of AI CPU running cycles}$	N/A
Control CPU Usage	Indicates the control CPU usage (%).	$(\text{Total number of running cycles of the control CPU} - \text{Number of idle cycles of the control CPU}) / \text{Total number of running cycles of the control CPU}$	N/A
Memory Bandwidth Usage	Indicates the DDR memory bandwidth usage (%).	Bandwidth of the DDR memory in a fixed period / Theoretical bandwidth of the DDR memory in a fixed period	N/A

14.2 Viewing ECS Running Status

Scenarios

The platform monitors the running status of ECSs. You can view the monitoring metrics of an ECS on the details page of the ECS.

Monitored data requires a period of time for transmission and display. The status of an ECS displayed on the monitoring page is the status obtained 5 to 10 minutes before. Therefore, wait 5 to 10 minutes, and then obtain the monitored data of a newly created ECS.

Context

- When an ECS is in the **Stopped**, **Faulty**, or **Deleted** state, its monitoring metrics cannot be viewed. The monitoring metrics can be viewed after the ECS starts or recovers.
- The system 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.

Procedure

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 2** In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, ID, or CPU vendor, and click the search icon to search for the ECS.
- Step 3** Click the name of the target ECS. The page providing details about the ECS is displayed.
- Step 4** Click the **Monitoring** tab to view the monitoring data.
- Step 5** In the monitoring area, specify the monitoring period.
- Fixed and customized time ranges are provided.
1. Fixed time ranges include **Last 1 hour**, **Last 3 hours**, **Last 12 hours**, **Last 24 hours**, **Last 7 days**, and **Last 30 days**.
 2. A customized time range can be specified within the last seven days.
- End

14.3 Viewing Information Through Metadata

ECS metadata is data about ECS data. It can be used to configure or manage running ECSs.

Context

[Table 14-2](#) lists the types of ECS metadata.

Table 14-2 ECS metadata types

API Type	Metadata Type	Description	Compatibility
OpenStack	Metadata	Queries ECS metadata.	Supported
	GET Password	Queries the password for logging in to an ECS.	Supported
	User Data	Queries ECS user data.	Supported

API Type	Metadata Type	Description	Compatibility
	POST Password	Stores the password for logging in to an ECS.	Supported
AWS EC2	ami-id	Queries the image ID of an ECS.	Supported
	ami-launch-index	Queries an ECS launching sequence.	Supported
	ami-manifest-path	Queries the path where an image list is stored.	Supported
	block-device-mapping	Queries the block device of an ECS.	Supported
	hostname	Queries the name of the host accommodating an ECS.	Supported
	instance-id	Queries an ECS ID.	Supported
	instance-type	Queries an ECS flavor name.	Supported
	local-ipv4	Queries the fixed IP address of an ECS.	Supported
	availability-zone	Queries the AZ accommodating an ECS.	Supported
	public-ipv4	Queries the floating IP address of an ECS.	Supported
	public-keys	Queries the public key of an ECS.	Supported
	reservation-id	Queries an ECS reservation ID.	Supported
	user-data	Queries ECS user data.	Supported
	instance-action	Queries ECS actions.	Not supported
kernel-id	Queries an ECS kernel image ID.	Not supported	
local-hostname	Queries the local name of an ECS.	Not supported	

API Type	Metadata Type	Description	Compatibility
	public-hostname	Queries the external name of an ECS.	Not supported
	ramdisk-id	Queries an ECS ramdisk image ID.	Not supported
	security-groups	Queries the security group to which an ECS belongs.	Not supported

The following describes the URI and methods of using the supported ECS metadata.

Precautions

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 to only 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 to only 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
```

Prerequisites

Security group rules in the outbound direction meet the following requirements:

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

NOTE

If you use the default security group rules in the outbound direction, the preceding requirements are met, and the metadata can be accessed. Default security group rules in the outbound direction are as follows:

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

Metadata (OpenStack)

Queries the metadata of an ECS.

- URI
`/169.254.169.254/openstack/{version}/meta_data.json`

- Method
Supports GET requests.

- Example
The following section describes how to use the tool cURL to view ECS metadata.

curl `http://169.254.169.254/openstack/latest/meta_data.json`

```
{
  "admin_pass": "sWs9YVAiytTs",
  "availability_zone": "manage-az",
  "files": [
    {
      "content_path": "/content/0000",
      "path": "/etc/litao.ini"
    }
  ],
  "hostname": "lt-test-01.novalocal",
  "launch_index": 0,
  "meta": {
    "test_key": "test_vaule"
  },
  "name": "lt-test-01",
  "public_keys": {
    "novakey": "ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQCr4Mk6fRqbPRXE3lq9NivYvrysz0/D
+gAWCgKCG46lU+x3aLRRtUwpSoX4W7FtRegHAp7EplhmW40vM+9HrKEZbQsaiyfe0VP/
cHZFKJLU9cnDJNMIb0WojLWDORnUnsfZL1tFaYyIcAdll6TuB92sj4Bg8xrYcN3nfNtSSfHHszvHlc0kys7AC
+eLL4NWlyeGDSkmsHS0vniP0mpRgpB2QShmx/ZEIQQ+YxMoL8z+A44+v/V+/
R8K7aJK3LbQ8Yu8vwky9M1OLG176s9pQnTmdlrKNWc4dYC8zNRxaFvyuUO9FD71OSEkmoZwkbhYHml
Vlw49d00lr63ok8mMij root@6B1AA2D2-3B57-11DA-8567-000000821800\n"
  },
  "random_seed":
"ptvKzBu5SNFrWKAiar62VT7KOGZS38T3VdpEjRDmBOWf0RGKvjNGXasAXz08KfzBPCq
+MupA4lg9mUyHWJiAcl7lvOQV7EWubJ4pxQn2tNxrRR7VdSR8WJUqeZV1FjxxgxjJkQtnWmrT57J
+SEaqzRP64ONYdikPnBLOlpnsj+gKMxXdmFL3tQ0ljC2vKdZdbjc8QkT84/dnbJpx/DaatkiLWVVtMlGhP72j/
0NH1JbAnv5EECS5z51h2YFnRxJOJFIScwnl5UuFRjsO82T+7usRjk0IXAvDIKcqQlZrV5WWXrWc1e
+yejjjVvAGad5PcjLeaGaEEBNo5eBJ9e4FZHLkkUV7naYGsLiOLjO3qZjvFM1gcaKsrv92/Ys2DbwHFFv0Xo/
uNA9Jx7MFIXELdNOFDP8nS4fVp52w6KOUvQ1iEhEBaOlQVKENJ6jfbSYCoYe4JGK/
```

```
UeGCaVhjZPpe8L9nVjn2QYfe6e/p1TxRk1MFe6LgrpQXR8SzEJ71+X6j+pRvP8Ui8W1M88esq
+ZqbjjlpAxR9ISmd/aTM7WGa7OLOxDsWQDBAH7J/86u4P1Zebs6vFgORM
+RUc2PFAOMSEQ2Ak0Ulhx9FrqoK9plol7x9m8aQAj0Sd89IMiCDTgVbk8DSZA4qT5lnYzqZYAvg9XWitCLK
eM+C05P7748=",
  "uuid": "32274de3-6efe-45b3-86a8-46c4335ffeb7"
}
```

Password (OpenStack)

Configures and queries the password for logging in to an ECS.

- URI
`/169.254.169.254/openstack/{version}/password`
- Method
Supports both GET and POST requests, where
 - The GET request is used to view the password.
 - The POST request is used to store the password. Exercise caution when sending a POST request.
- Examples
 - Example 1: Query the password for logging in to an ECS.
curl http://169.254.169.254/openstack/latest/password
 - Example 2: Store the password for logging in to an ECS.
 - i. Run the following command to temporarily disable the historical record function to prevent passwords from being leaked when running the **history** command:
set +o history
 - ii. Run the following command to save the password for logging in to the ECS:
curl -X POST http://169.254.169.254/openstack/latest/password -d "xxxxxxx"
xxxxxxx indicates the password queried in [Example 1](#).
 - iii. Run the following command to enable the historical record function:
set -o history

User Data (OpenStack)

Queries ECS user data. The value is configured when you create an ECS. It cannot be changed after the configuration.

- URI
`/169.254.169.254/openstack/{version}/user_data`
- Method
Supports GET requests.
- Example
curl http://169.254.169.254/openstack/latest/user_data

```
ICAgICAgDQoiQSBjbG91ZCBkb2VzIG5vdCBrbm93IHdoeSBpdCBtb3ZlcyBpbjBqdXN0IHN1Y2ggYSBkaXJlY
3Rpb24gYW5kIGF0IHN1Y2ggYSBzcGVlZC4uLkI0IGZlZWxzIGFuIGltcHVsc2lvbi4uLnRoXMGaXMgdGhIIH
BsYWNIHRvIGdviG5vdy4gQnV0IHRoZSBza3kga25vd3MgdGhIIHJlYXNvbnMgYW5kIHRoZSBwYXR0ZXJl
cyBiZWphbmQgYWxslGNsb3VkcycgYW5kIHlvdSB3aWxslGtub3csIHRvbywg2hIbI5b3UgbGlmdCB5b3
```



```
Vyc2VsZiBoaWdoIGVub3VnaCB0byBzZWUgYmV5b25kIGhvcml6b25zLiINCg0KLVJpY2hhcmQgQmFjaA=
```

User Data (AWS EC2)

Queries ECS user data. The value is configured when you create an ECS. It cannot be changed after the configuration.

- URI
/169.254.169.254/{version}/meta-data/user-data
- Method
Supports GET requests.

- Example

curl http://169.254.169.254/latest/meta-data/user-data

```
ICAgICAgDQoiQSBjbG91ZCBkb2VzIG5vdCBrbm93IHdoeSBpdCBtb3ZlcyBpbjBqdXN0IHN1Y2ggYSBkaXJlY3Rpb24gYW5kIGF0IHN1Y2ggYSBzcgVlZC4uLk0IGZlZWxzIGFuIGltcHVsc2lubi4uLnRoaXMgaXMgdGhIIHBsYWNIHRvIGdvlG5vdy4gQnV0IHRob2ZSBza3kga25vd3MgdGhIIHJlYXNvbnMgYW5kIHRob2ZSBwYXR0ZXJucyBiZWphbmQgYWxslGNsb3VkcycgYW5kIHlvdSB3aWxslGtub3csIHRvbywg2hlbiB5b3UgbGlmdCB5b3Vyc2VsZiBoaWdoIGVub3VnaCB0byBzZWUgYmV5b25kIGhvcml6b25zLiINCg0KLVJpY2hhcmQgQmFjaA=
```

Ami ID (AWS EC2)

Queries an ECS image ID.

- URI
/169.254.169.254/{version}/meta-data/ami-id
- Method
Supports GET requests.

- Example

curl http://169.254.169.254/latest/meta-data/ami-id

```
ami-00000003
```

Ami Launch Index (AWS EC2)

Queries an ECS launching sequence. The value of the first launched ECS is 0.

- URI
/169.254.169.254/{version}/meta-data/ami-launch-index
- Method
Supports GET requests.

- Example

curl http://169.254.169.254/latest/meta-data/ami-launch-index

Ami Manifest Path (AWS EC2)

Queries the path where an image list is stored.

- URI
/169.254.169.254/{version}/meta-data/ami-manifest-path

- Method
Supports GET requests.
- Example
curl http://169.254.169.254/latest/meta-data/ami-manifest-path
FIXME

Block Device Mapping (AWS EC2)

Queries the block device of an ECS.

- URI
/169.254.169.254/{version}/meta-data/block-device-mapping/ami
- Method
Supports GET requests.
- Example
curl http://169.254.169.254/latest/meta-data/block-device-mapping/ami
vda

Hostname (AWS EC2)

Queries the name of the host accommodating an ECS. The .novalocal suffix will be added later.

- URI
/169.254.169.254/{version}/meta-data/hostname
- Method
Supports GET requests.
- Example
curl http://169.254.169.254/latest/meta-data/hostname
vm-test.novalocal

Instance ID (AWS EC2)

Queries an ECS ID.

- URI
/169.254.169.254/{version}/meta-data/instance-id
- Method
Supports GET requests.
- Example
curl http://169.254.169.254/latest/meta-data/instance-id
i-00000001

Instance Type (AWS EC2)

Queries an ECS flavor name.

- URI

/169.254.169.254/{version}/meta-data/instance-type

- Method
Supports GET requests.
- Example
curl http://169.254.169.254/latest/meta-data/instance-type
flavor_test

Local IPv4 (AWS EC2)

Queries the fixed IP address of an ECS.

- URI
/169.254.169.254/{version}/meta-data/local-ipv4
- Method
Supports GET requests.
- Example
curl http://169.254.169.254/latest/meta-data/local-ipv4
192.168.111.120

Availability Zone (AWS EC2)

Queries the AZ accommodating an ECS.

- URI
/169.254.169.254/{version}/meta-data/placement/availability-zone
- Method
Supports GET requests.
- Example
curl http://169.254.169.254/latest/meta-data/placement/availability-zone
az1.dc1

Public IPv4 (AWS EC2)

Queries the floating IP address of an ECS.

- URI
/169.254.169.254/{version}/meta-data/public-ipv4
- Method
Supports GET requests.
- Example
curl http://169.254.169.254/latest/meta-data/public-ipv4
192.168.111.120

Public Keys (AWS EC2)

Queries the public key of an ECS.

- URI

/169.254.169.254/latest/meta-data/public-keys/0/openssh-key

- Method
Supports GET requests.
- Example

curl http://169.254.169.254/latest/meta-data/public-keys/0/openssh-key

```
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQADIA5Fw5k8Fgzajn1zJwLoV3+wMP+6CyvsSilc/hioggSnYu/AD0Yqm8vVO0kWlun1rFbdO+QUZKyVr/OPUjQSw4SRh4qsTKf/+eFoWTjplFvd1WCBZzS/WRenxIwR00KkczHSJro763+wYcwKieb4eKRxaQoQvoFgVjLBULXajH4eKoKTVNtMXAvPP9aMy2SLgsJNtMb9ArfziAiblQynq7UIflnN3VclzPeiWrqtzjyOp6CPUXnL0lVPTvbLe8sUteBsJZwLL6K4i+Y0lf3ryqnmQgC21yW4Dzu+kwk8FVT2MgWkCwiZd8gQ/+uJzrJFyMfUOBikIOBfuUENIJUuhABGenerated-by-Nova
```

Reservation ID (AWS EC2)

Queries an ECS reservation ID.

- URI
/169.254.169.254/{version}/meta-data/reservation-id
- Method
Supports GET requests.
- Example

curl http://169.254.169.254/latest/meta-data/reservation-id

```
r-kso0e196
```

15 Load Balancing

Elastic Load Balance (ELB) is a service that automatically distributes incoming traffic across multiple backend Elastic Cloud Servers (ECSs) based on specified forwarding policies. ELB can expand the access handling capability of application systems through traffic distribution and achieve a higher level of fault tolerance and performance. ELB also improves system availability by eliminating single points of failures (SPOFs).

- For details about how to create a load balancer, see "Creation" in **Operation Help Center > Network > Elastic Load Balance > User Guide**.
- For details about how to add an ECS to or delete an ECS from a load balancer, see "Management" > Managing Backend Cloud Servers" in **Operation Help Center > Network > Elastic Load Balance > User Guide**.

16 Best Practices

16.1 Creating an Application Allowing Access from External Networks

16.1.1 Overview

HUAWEI CLOUD Stack provides a wide range of cloud services, enabling you to quickly and easily deploy, run, and maintain your applications on the cloud. If you are planning to create a simple application allowing access from external networks and the application has relatively low requirements on compute, network, and storage performance and average requirements on database security, you may deploy your application by referring to the procedures described in this section.

16.1.2 Implementation Plan

Related Services

The application has low requirements on the performance of the server and general requirements on database security. Therefore, you can apply for an Elastic Cloud Server (ECS), and deploy the application and database on the ECS. The following services are recommended:

- Elastic Cloud Server (ECS)
An ECS is a computing server consisting of CPUs, memory, and Elastic Volume Service (EVS) disks that allow on-demand allocation and elastic scaling. After an ECS is created, you can use it like using your local computer or physical server.
- Virtual Private Cloud (VPC)
VPC enables you to build a logically isolated virtual network environment for ECSs. You can customize and manage the network environment. VPC improves the security of your resources and simplifies network deployment.
The IP address of an ECS belongs to a subnet in a VPC, so the VPC and subnet are required. In addition, the ECS needs to allow access from external

networks, so the elastic IP address (EIP) is required to ensure the security of the ECS. Security groups are used and security group rules are configured to perform access control on protocols, ports, and IP addresses, ensuring the security of the ECS.

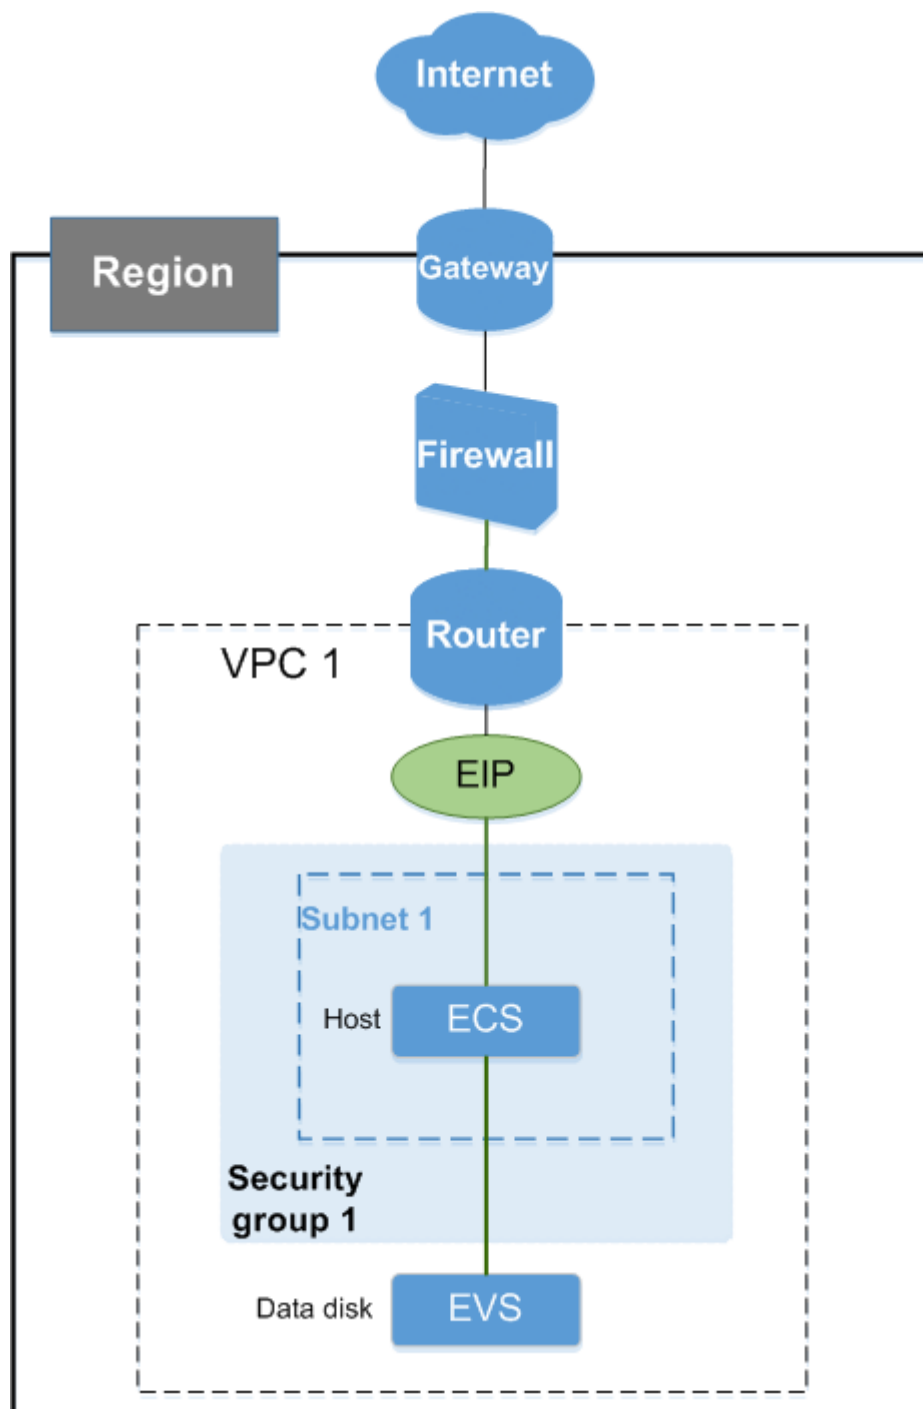
- Elastic Volume Service (EVS)

EVS is a virtual block storage service. Compared with traditional disks, EVS disks have higher data reliability, higher I/O throughput, and easier-to-use characteristics. EVS disks provide block storage for ECSs.

Networking

[Figure 16-1](#) shows the networking of this practice.

Figure 16-1 Networking



Deployment process

When applying for an ECS, you need to select a VPC, subnet, and security group for the ECS. Therefore, before creating an ECS, apply for a VPC and create a security group. When applying for an ECS, you can choose to assign an EIP to the ECS. After obtaining an ECS, you need to log in to the ECS and manually initialize data disks. After initializing data disks, you can start to deploy your applications.

Figure 16-2 shows the deployment process.

Figure 16-2 Deployment process

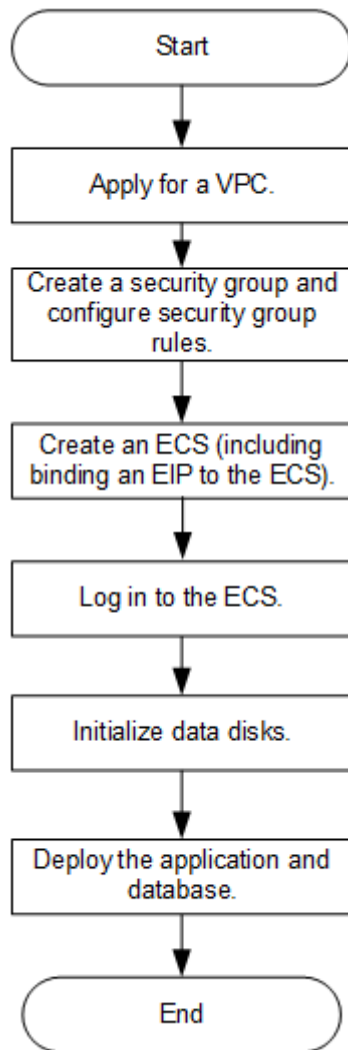


Table 16-1 Deployment process

Operation	Meaning and Purpose	Involved Feature
Requesting a VPC	A VPC will be created to provide an isolated virtual network for ECSs. You can configure and manage the virtual network as required.	VPC: VPC and subnet
Creating a security group and configuring a security group rule	A security group is a logically established group that provides VM-level security protection and access control. In this step, you need to create a security group and configure an inbound rule for the security group.	VPC: security group

Operation	Meaning and Purpose	Involved Feature
Applying for an ECS (including binding an EIP to the ECS)	In this step, you apply for an ECS on which the applications and database will be deployed. An EIP is a public IP address on the Internet. An ECS can be bound with an EIP and configured with security group rules to allow access from external networks.	<ul style="list-style-type: none"> • ECS • EVS • VPC: VPC, subnet, security group, and EIP
Remotely Logging in to the ECS and deploying the application	After creating an ECS, remotely log in to the ECS, and then perform operations including initializing data disks and deploying the application.	ECS
Initializing data disks	If you attach data disks when applying for an ECS, you need to manually initialize the data disks after the ECS is successfully provisioned.	<ul style="list-style-type: none"> • ECS • EVS
Deploying the application and database	After creating an ECS, remotely log in to the ECS and deploy the application and database.	ECS

16.1.3 Requesting and Configuring Services

16.1.3.1 Applying for a VPC

Context

This section describes how to quickly apply for a VPC and single-stack subnet.

Before performing the following operations, you need to correctly configure and allocate external networks as planned.

Procedure

Step 1 Log in to ManageOne as a VDC administrator or VDC operator using a browser.

URL in non-B2B scenarios: <https://Domain name of ManageOne Operation Portal>, for example, <https://console.demo.com>.

URL in B2B scenarios: <https://Domain name of ManageOne Tenant Portal>, for example, <https://tenant.demo.com>.


You can log in using a password or a USB key.

- Login using a password: Enter the username and password.

The password is that of the VDC administrator or VDC operator.

- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.



Step 2 Click  in the upper left corner, select a region, and choose **Network > Virtual Private Cloud**.

Step 3 Select a project in the region from the drop-down list in the upper left corner of the main menu.

Step 4 On the **Virtual Private Cloud** page, click **Apply for VPC**.

The **Select Service** dialog box is displayed.

Step 5 Select a service and click **Apply Now**.

The **Select VPC Configuration** page is displayed.

Step 6 Select **VPC with Internet Access (Optional)**, and click **Next**.

Step 7 Configure the VPC parameters by referring to [Table 16-2](#).

Table 16-2 VPC parameters

Parameter	Description	Example Value
Region and Project	The current region and project are displayed by default. To change them, use the selector in the upper left corner of the page.	az1.dc1 (test)
Name	Specifies the VPC name. The name can contain letters, digits, underscores (_), and hyphens (-).	VPC-001
External Network	Select an availability zone from the All AZs drop-down list and select the external network in the VPC.	az0.dc0 net-01
CIDR Block	Specifies the Classless Inter-Domain Routing (CIDR) block for the VPC. The CIDR block of a subnet can be the same as the CIDR block for the VPC (for a single subnet in the VPC) or a subset (for multiple subnets in the VPC).	192.168.0.0/16
Required Duration	Specifies the required duration of a VPC.	1 year

Step 8 Configure subnet 1 by referring to [Table 16-3](#). Subnet 1 is isolated from the extranet, which ensures high security. Cloud servers in subnet 1 cannot directly communicate with the extranet.

Table 16-3 Parameters for configuring subnet 1

Parameter	Description	Example Value
Subnet 1 Name	Specifies the subnet 1 name. The name consists of only letters, digits, underscores (_), and hyphens (-).	Subnet-f03c
Type	If you have deployed the dual stack (IPv4 & IPv6) in the system, you need to select a network type. If only the IPv4 service is deployed, configure subnet 1 parameters by referring to Table 16-4 . <ul style="list-style-type: none">IPv4IPv4&IPv6 To create a single-stack subnet, select IPv4 and set parameters based on Table 16-4 .	IPv4

Table 16-4 Parameters for configuring an IPv4 subnet

Parameter	Description	Example Value
Subnet 1 Segment	Specifies the IP address range of subnet 1.	192.168.0.0/24
Gateway	Specifies the gateway address of subnet 1.	192.168.0.1
DNS Server Address 1	Specifies the IP address of an associated DNS server. This parameter is optional.	192.168.71.3
DNS Server Address 2	Specifies the IP address of an associated DNS server. This parameter is optional. NOTE If the DNS server address is left blank, the subnet is not associated with any DNS server. When using only one DNS server address, enter it into DNS Server Address 1 . To add a DNS server address, click Add DNS Server Address .	192.168.72.3

Parameter	Description	Example Value
Static Route Switch	<p>If this switch is set to OFF, static routes will not be configured for cloud servers in the subnet. If this switch is set to ON, the configured static routes will be injected to those cloud servers by using the DHCP function of the subnet.</p> <ul style="list-style-type: none"> • Destination: indicates the destination IP address range of the static route. • Next Hop: indicates the next-hop IP address of the static route. <p>NOTE When you need to add more static routes for cloud servers in a subnet, click Add Static Route. You can configure up to five static routes at a time.</p>	<ul style="list-style-type: none"> • Destination: 10.10.0.0/24 • Next Hop: 192.168.20.2

Step 9 Disable the NAT function.

Step 10 Check the configuration of the new VPC.

- Click **Apply Now** to apply for the VPC.
- Alternatively, click **Add to Cart** and submit the application later.

----End

16.1.3.2 Creating a Security Group and Configuring Security Group Rules

Creating a Security Group


Step 1 Log in to ManageOne as a VDC operator using a browser.

URL in non-B2B scenarios: **https://Domain name of ManageOne Operation Portal**, for example, **https://console.demo.com**.

URL in B2B scenarios: **https://Domain name of ManageOne Tenant Portal**, for example, **https://tenant.demo.com**.

You can log in using a password or a USB key.

- Login using a password: Enter the username and password.
The password is that of the VDC operator.
- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.

Step 2 Click  on the left of the main menu, select a region, and choose **Network > Virtual Private Cloud**.

Step 3 In the navigation pane on the left, click **Security Group**.

Step 4 On the **Security Group** page, click **Create Security Group**.

Step 5 In the **Create Security Group** area, set parameters as prompted.

Table 16-5 Parameter description

Parameter	Description	Example Value
Name	Specifies the security group name.	sg-001
Description	Provides supplementary information about the security group. This parameter is optional.	-

Step 6 Click **OK**.

----End

Configuring a Security Group Rule

Step 1 Click the name of the security group. The details page is displayed.

Step 2 On the Inbound tab, click **Add Rule**, and set the required parameters.

 **NOTE**

The application needs to allow access from external networks. Therefore, permit access in the inbound direction.

Table 16-6 Parameters for adding a security group rule

Parameter	Description	Example Value
Protocol	<p>Specifies the network protocol. The value can be ANY, TCP, UDP, or ICMP.</p> <ul style="list-style-type: none"> • ANY means that this rule is effective for any protocol. • TCP: Transmission Control Protocol (TCP) is a transport layer protocol. It provides reliable data transmission and maintains a virtual connection between devices or services that communicate with each other. • UDP: indicates a transport layer protocol that is used to compress network data traffic into data packets. • ICMP: indicates a network layer protocol. It is used to transmit error report control messages, and the ping command is used for communication status check. 	TCP

Parameter	Description	Example Value
Port Range/ ICMP Type	<ul style="list-style-type: none"> When you select TCP or UDP for Protocol, this parameter is a port range. Its value ranges from 1 to 65535. When you select ANY for Protocol, this parameter is unconfigurable. When you select ICMP for Protocol, this parameter is the ICMP type. 	22 or 22-30
Type	<p>If you have deployed the dual stack (IPv4 & IPv6) in the system, you need to select a type of the security group rule.</p> <ul style="list-style-type: none"> If this parameter is set to IPv4, the traffic on the IPv4 network segment is allowed to pass. If this parameter is set to IPv6, traffic on the IPv6 network segment is allowed to pass. 	IPv4
Source/ Destination	<p>Specifies the source when you select the Inbound tab.</p> <p>Specifies the destination when you select the Outbound tab.</p> <p>The value can be an IP address range or a security group.</p> <ul style="list-style-type: none"> If you specify an IP address range as the value, select an IP address type, and enter an IP address range. <ul style="list-style-type: none"> Select IPv4 for Type. For example: xxx.xxx.xxx.xxx/32 (IP address) xxx.xxx.xxx.0/24 (subnet) 0.0.0.0/0 (any IP address) Select IPv6 for Type. For example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/64 (subnet) 0:0:0:0:0:0:0:0/0 (any address) If you specify a security group as the value, choose a source or destination security group from the drop-down list. 	0.0.0.0/0 default
Description	<p>Provides supplementary information about the rule. This parameter is optional.</p> <p>The description can contain a maximum of 64 characters and cannot contain angle brackets (<) or (>).</p>	-

Parameter	Description	Example Value
Operation	<ul style="list-style-type: none"> • Clone: specifies to copy an existing security group rule. • Delete: specifies to delete an existing security group rule. 	-
Adding a tag	Adds multiple rules to a security group. A maximum of 10 rules can be added at a time.	-


Step 3 Click **OK**.

----End

16.1.3.3 Creating an ECS

Step 1 Log in to ManageOne as a VDC operator using a browser.

- URL in non-B2B scenarios: **https://Address for accessing ManageOne Operation Portal**, for example, **https://console.demo.com**
URL in B2B scenarios: **https://Address for accessing ManageOne Tenant Portal**, for example, **https://tenant.demo.com**
- Username and password: username and password of the VDC operator

Step 2 Click  in the upper left corner, select a region, and choose **Computing > Elastic Cloud Server**.

The **Elastic Cloud Server** page is displayed.

Step 3 Select a resource set from the menu bar on the top of the page.

Step 4 Click **Apply for ECS**.

The **Select Service** page is displayed.

Step 5 Select a service and click **Apply Now**.

The **Apply for ECS** page is displayed.

Step 6 Configure basic information about the ECS to be created. For details, see [Table 16-7](#).

 **NOTE**

- When you select different services, the parameters to customize are different. The service you selected in [Step 5](#) determines whether **AZ**, **ECS Type**, **vCPUs**, **Memory**, **Image Type**, and **Image** can be customized. During the configuration, you can skip the parameters that cannot be customized.
- The screenshot is only an example. If the actual environment is different from the screenshot, use the actual environment.

Table 16-7 Parameter description

Parameter	Description	Example Value
AZ	A physical region where resources use independent power supplies and networks. AZs are physically isolated but interconnected through an internal network. To enhance application availability, create ECSs in different AZs.	kvm_az
Creation Method	Specifies the method for creating an ECS. <ul style="list-style-type: none">• New: Customize parameters to create an ECS.• Create from Template: Create an ECS using a full-ECS image or ECS backup as a template.	New
ECS Type	The platform provides various ECSs for you to select based on application scenarios.	General-purpose
Boot Mode	<ul style="list-style-type: none">• Basic Input/Output System (BIOS) is used to load the basic computer code to initialize hardware, check hardware functions, and boot the OS.• Unified Extensible Firmware Interface (UEFI) does not need a long sel-check as BIOS does, simplifying hardware initialization and OS boot. In addition, UEFI is easy to use because it supports graphical user interfaces (GUIs), various operation modes, and hardware driver insertion. NOTE <ul style="list-style-type: none">- Skip this parameter if it is not displayed.- In ARM scenarios, the ECS boot mode can only be UEFI and cannot be changed.	BIOS

Parameter	Description	Example Value
Image Type	<ul style="list-style-type: none"> • 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. To select a public image, set Image Type to Public Image and select a desired one from the Image drop-down lists. • Private Image A private image is an image available only to the user who created it using an existing ECS or external image file. It contains an OS, preinstalled public applications, and the user's private applications. Using a private image to create ECSs removes the need to configure multiple ECSs repeatedly. To select a private image, set Image Type to Private Image and select a desired one from the Image drop-down list. • Shared Image A shared image is a private image shared by another user. To select a shared image, set Image Type to Shared Image and select a desired one from the Image drop-down list. 	Public Image
Image	<ul style="list-style-type: none"> • Windows Used for development platforms or operating services that run Windows. An authorized license is included in the image. • Linux Used for development platforms or operating services that run Linux. 	Windows
Joint Windows Domain	<p>This parameter is available if the virtualization type of the selected AZ is KVM, the ECS is running a Windows OS, and the product selected in Step 5 has been configured with domain information. If the selected image uses a static IP address or does not have Cloudbase-Init installed, it cannot be added to a domain.</p> <p>Specify whether to add an ECS to a Windows domain. You can select a domain from the drop-down list. Available options are those defined by the administrator during product creation.</p>	-

Parameter	Description	Example Value
Same Storage	<p>If the new ECS needs to support backup or disaster recovery, select Yes. Otherwise, select No.</p> <p>If you select Yes, make sure that the system and data disks of the ECS reside in the same storage backend, and the storage backend is configured with the storage tag. Otherwise, the ECS cannot be provisioned.</p> <p>NOTE This parameter is available only when Boot Mode of the specified ECS flavor is set to Cloud Disk.</p>	No
System Disk	To ensure that the ECS runs properly, the minimum allowed capacity of the system disk is related to the selected image file.	10GB
Data Disk	<p>This parameter is displayed after you click Add Data Disk.</p> <p>Select a disk type and set the disk size. You can create multiple data disks for an ECS.</p>	40GB
Quantity	Set the number of ECSs to be created.	1

Step 7 Click **Next**. The **Network Configuration** page is displayed.

Step 8 Configure network information about the ECS. For details, see [Table 16-8](#).

Table 16-8 Parameter description

Parameter	Description	Example Value
Resource Set	<p>Select the current resource set or another resource set from the drop-down list. You can view the current resource set in the navigation bar at the top. You do not need to change the default resource set.</p> <p>NOTE This parameter is available when VPC sharing is enabled on Service OM and the shared VPC permission is configured for the resource set on ManageOne. Otherwise, this parameter is not displayed. By default, this function is disabled.</p>	project_02
VPC	Provides a network, including subnet and security group, for an ECS.	-

Parameter	Description	Example Value
NIC	<p>Includes primary and extension NICs.</p> <ul style="list-style-type: none">• If you select VPC Subnet, all subnets in the VPC are available for you to choose from. In this case, the NIC supports layer 3 communication, allowing the ECS to communicate with networks (for example, the public network or other VPCs) beyond the VPC.• If you select Intra-Project Subnet, all project-level subnets in the project are available for you to choose from. All NICs configured with the same subnet can communicate with each other at layer 2 on the project level. Layer 2 communication is supported within the same VPC and between different VPCs.	subnet-c869(192.168.0.0/24)
Security Group	Controls ECS access within a security group or between security groups by defining access rules. This enhances ECS security.	-
EIP	<p>A static public IP address bound to an ECS in a VPC. Using the EIP, the ECS provides services externally.</p> <p>The following options are provided:</p> <ul style="list-style-type: none">• Do Not Use: Without an EIP, the ECS cannot access the Internet and is used only in the private network or cluster.• Automatically Assign: The system automatically assigns an EIP for the ECS. The EIP provides exclusive bandwidth.• Specify: An existing EIP is assigned for the ECS. When using an existing EIP, you cannot create ECSs in batches.	Automatically Assign

Step 9 Click **Next**. The **Advanced Settings** page is displayed.

Step 10 Set the ECS name.

When you create ECSs in batches, the system automatically adds an incremental number to the end of each ECS name.

Step 11 Set the host name prefix of the ECS.

If this parameter is displayed, set it. The host name prefix and a suffix of 5 random characters (0-9 and a-z) form the ECS host name, that is, the computer name shown in the OS. It is in the format "Host Name Prefix-5 random characters".

Step 12 Set the power status of the ECS to **Running**.

- **Stopped:** A newly obtained ECS stays in the **Stopped** state.
- **Running:** A newly obtained ECS stays in the **Running** state.

Step 13 To add description for an ECS, such as the purpose of the ECS, enter the required information in the description text box.


Step 14 If **Set Key or New Password** is displayed, click **Yes**. You can customize the password or key pair for logging in to the ECS.

Step 15 Configure the login mode.

 **NOTE**

This password is used to log in to the ECS. Keep it secure.

Step 16 Retain the default values for other parameters and click **Next**. The **Confirmation** page is displayed.

- Check whether all configuration items are correct. If you need to modify a configuration item, click  next to the corresponding module.
- Confirm **Required Duration**.

Step 17 Click **Add to Cart** or **Apply Now**.

- **Add to Cart:** Add the configured ECS to the shopping cart, and submit the order after you confirm all the resources you need, including network and storage resources.
- **Apply Now:** Submit the task.

 **NOTE**


- If the ECS you requested needs administrator approval, it will be provisioned after your request is approved. Otherwise, the ECS will be provisioned immediately.
- If you create an ECS with additional data disks, initialize the data disks after the ECS is created.

----End

16.1.3.4 Logging In to an ECS

Step 1 Log in to ManageOne as a VDC operator using a browser.

- URL in non-B2B scenarios: **https://Address for accessing ManageOne Operation Portal**, for example, **https://console.demo.com**
URL in B2B scenarios: **https://Address for accessing ManageOne Tenant Portal**, for example, **https://tenant.demo.com**
- Username and password: username and password of the VDC operator

Step 2 Click  in the upper left corner, select a region, and choose **Computing > Elastic Cloud Server**.

The **Elastic Cloud Server** page is displayed.

Step 3 Select a resource set from the menu bar on the top of the page.

- Step 4** In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID, and click the search button to search for the ECS.
- Step 5** Locate the row containing the ECS and click **Remote Login** in the **Operation** column.
- The **Configure Keyboard Layout for Remote Login** is displayed.
- Step 6** Select the English keyboard and click **Remote Login**.
- Step 7** (Optional) If the system displays "Press CTRL+ALT+DELETE to log on", click **Send CtrlAltDel** in the upper right corner of the remote login page to log in to the ECS.

Figure 16-3 Send CtrlAltDel



- Step 8** Enter the password set in [16.1.3.3 Creating an ECS](#) and log in to the ECS.
- End

16.1.3.5 Initializing a Windows Data Disk

A data disk attached to an ECS or created together with an ECS must be initialized before it can become available. This section uses an instance running Windows Server 2008 R2 Enterprise as an example. Initialization operations vary with operating systems.

Prerequisites

- You have logged in to the ECS. For details, see [16.1.3.4 Logging In to an ECS](#).
- A disk has been attached to the ECS and has not been initialized.

Context

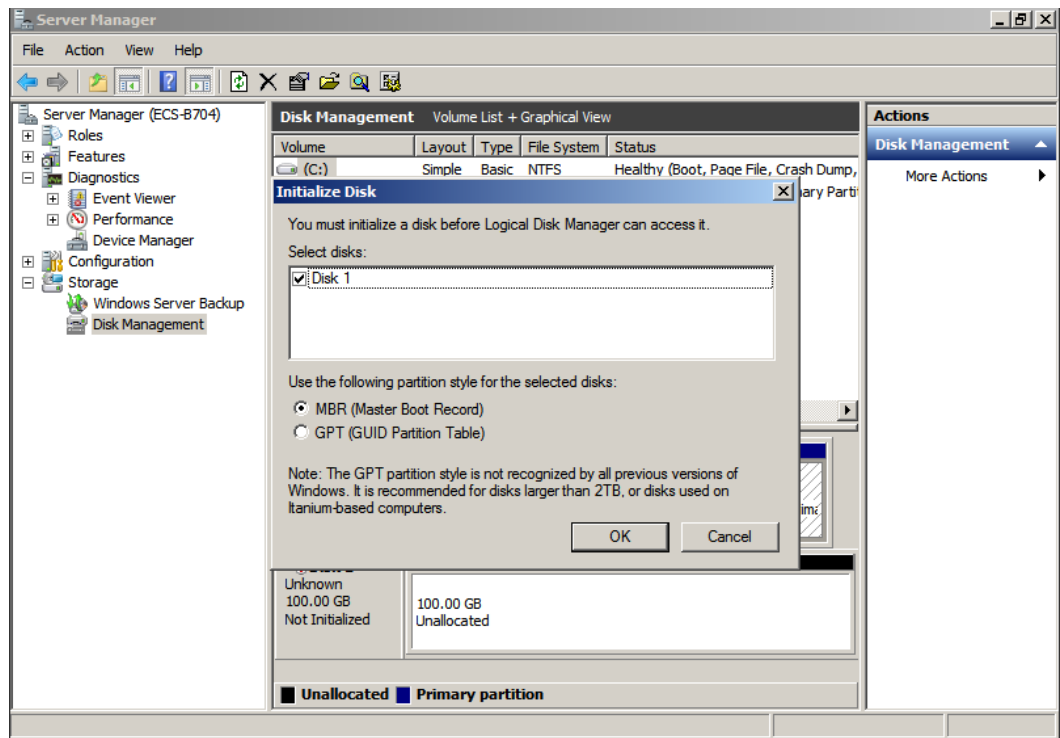
Initializing a data disk is highly risky. If there is useful data on the data disk to be initialized, create a snapshot or backup copy for the data disk before disk initialization.

Procedure

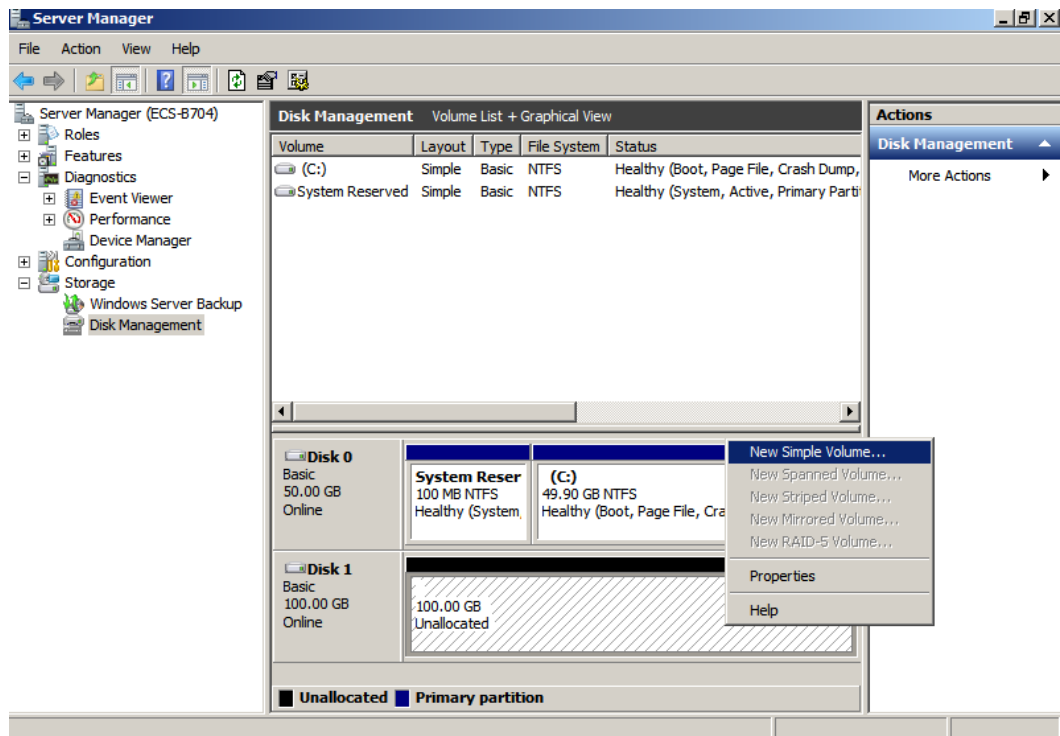
- Step 1** In desktop, right-click **Computer** and choose **Manage** from the shortcut menu.
- The **Server Manager** page is displayed.
- Step 2** In the navigation pane, choose **Storage > Disk Management**.
- Step 3** If the disk to be initialized in the disk list is in **Offline** state, right-click in the disk area and choose **Online** from the shortcut menu.
- Then, the disk status changes from **Offline** to **Uninitialized**.
- Step 4** Right-click in the disk area and choose **Initialize Disk** from the shortcut menu. In the displayed **Initialize Disk** dialog box, select **MBR (Master Boot Record)** and click **OK**.

NOTE

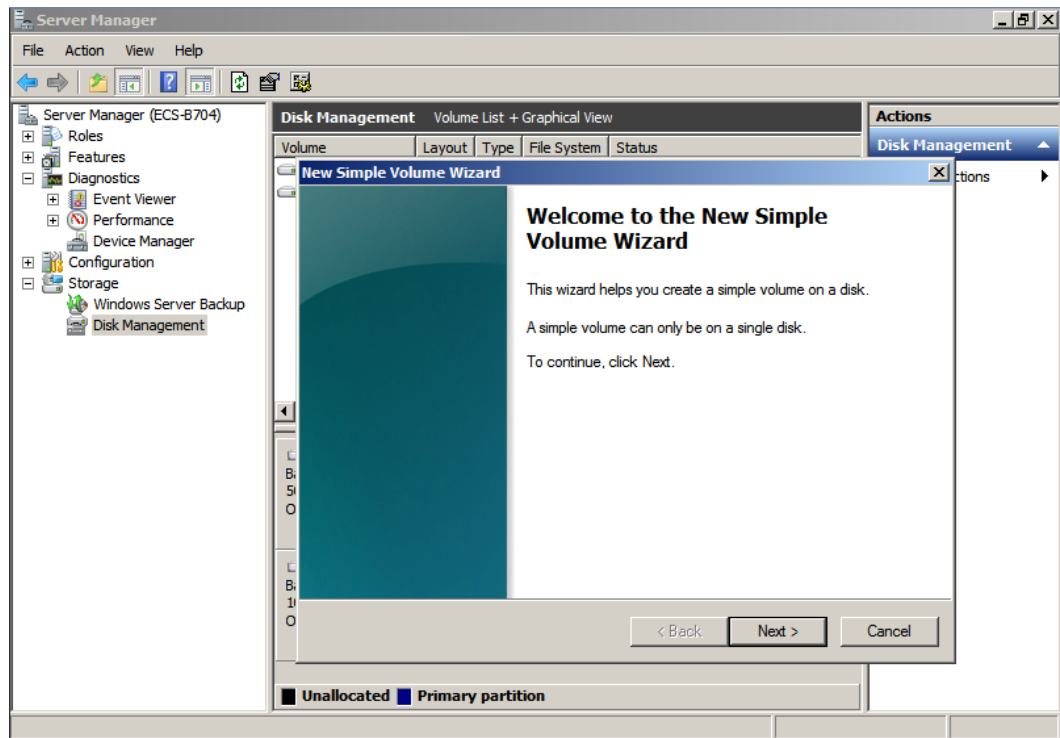
If the data disk to be initialized is larger than 2 TB, select **GPT (GUID Partition Table)** in the dialog box.



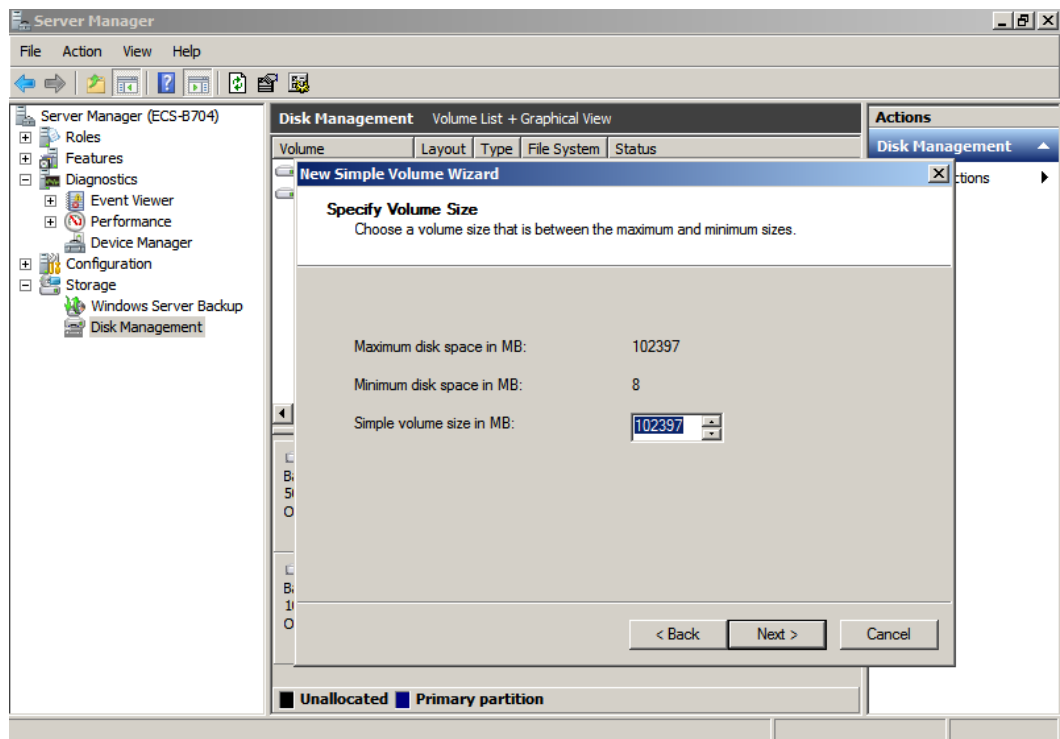
Step 5 Right-click at the unallocated disk space and choose **New Simple Volume** from the shortcut menu.



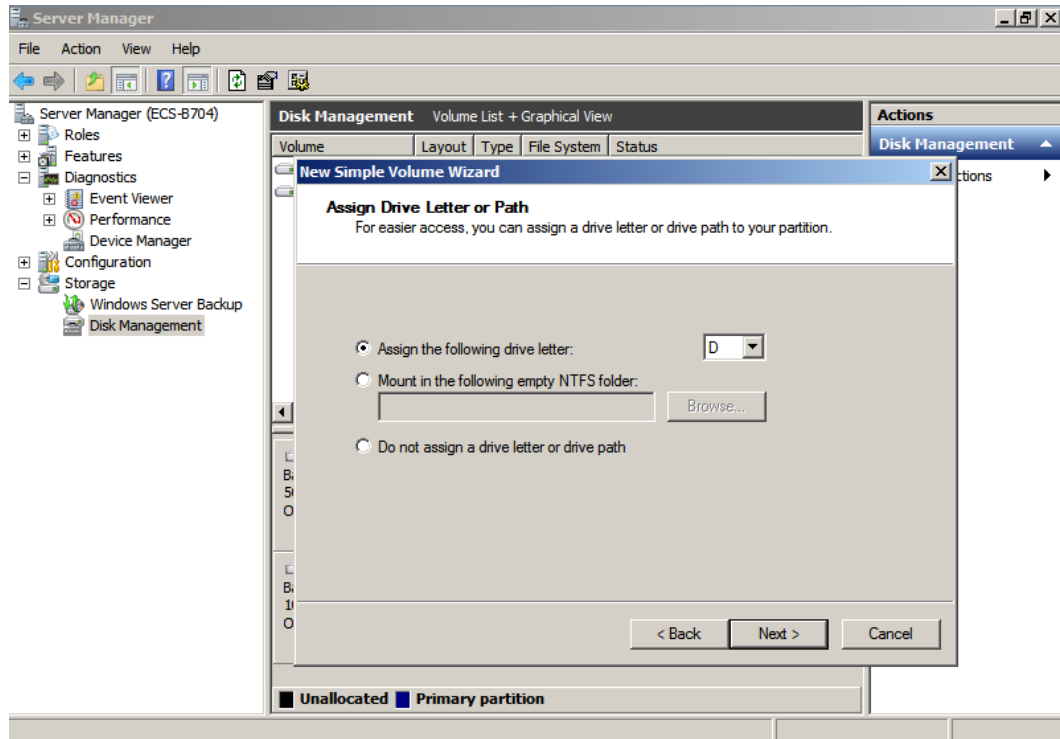
Step 6 On the displayed **New Simple Volume Wizard** page, click **Next**.



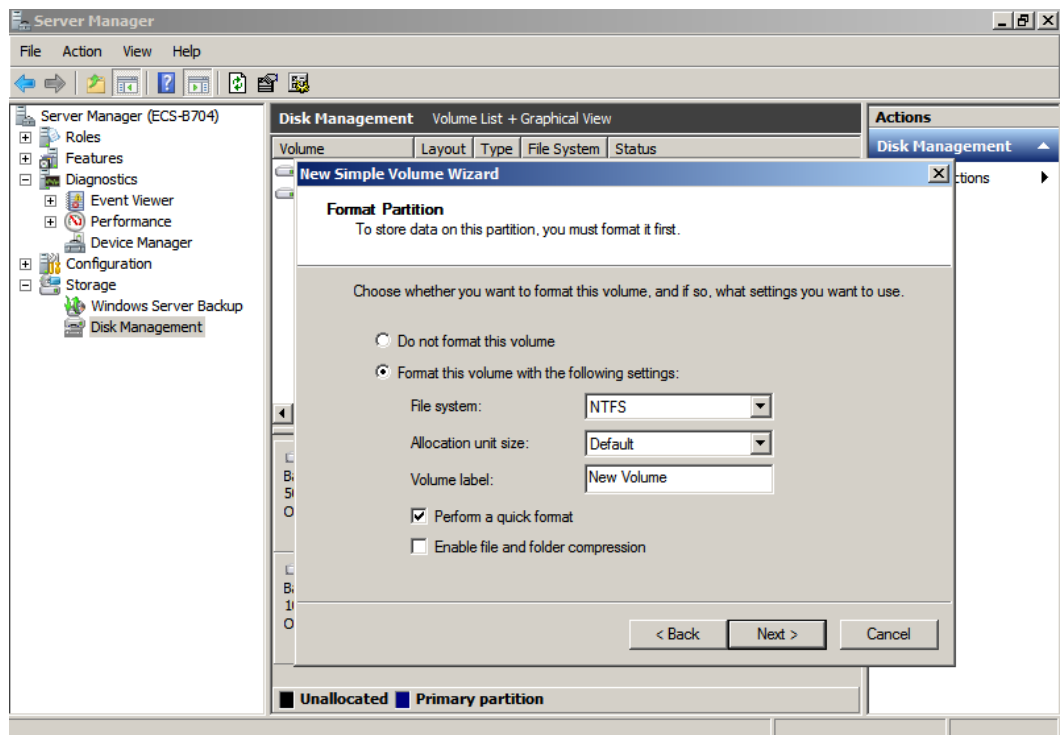
Step 7 Specify the simple volume size as required (the default value is the maximum) and click **Next**.



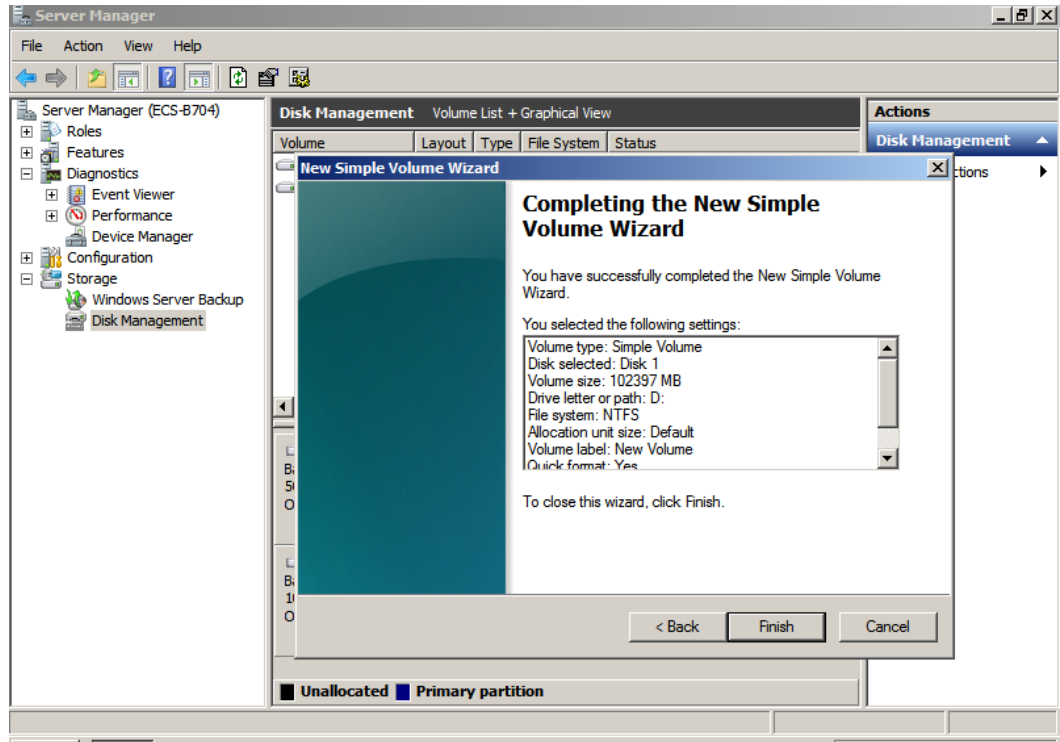
Step 8 Assign the driver letter and click **Next**.




Step 9 Select **Format this volume with the following settings**, set parameters based on the actual requirements, and select **Perform a quick format**. Then click **Next**.



Step 10 Click **Finish**.



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



----End

16.1.3.6 Initializing a Linux Data Disk (fdisk)

A data disk attached to an ECS or created together with an ECS must be initialized before it can become available. This section uses an instance running CentOS 7.0 64bit as an example, and uses the fdisk partition tool to set up partitions for the data disk. Initialization operations vary with operating systems.

Prerequisites

- You have logged in to the ECS. For details, see [17.1.3.4 Logging In to an ECS](#).
- A disk has been attached to the ECS and has not been initialized.

Context

Both the fdisk and parted can be used to partition a Linux data disk. For a disk larger than 2 TB, only parted can be used because fdisk cannot partition such a large disk. For details, see [16.1.3.7 Initializing a Linux Data Disk \(parted\)](#).

Creating Partitions and Mounting a Disk

The following example shows how to create a new primary partition on a new data disk that has been attached to an instance. The primary partition will be created using fdisk, and MBR is the default partition style. Furthermore, the partition will be formatted using the ext4 file system, mounted on the `/mnt/sdc` directory, and set to be automatically mounted upon a system start.

Step 1 Run the following command to view information about the added data disk:

fdisk -l

Information similar to the following is displayed: (In the command output, the server contains two disks. `/dev/xvda` is the system disk, and `/dev/xvdb` is the added data disk.)

NOTE

If you do not log in to the ECS and run the `umount` command but directly detach the `/dev/xvdb` or `/dev/vdb` EVS disk on the management console, the disk name in the ECS may encounter a release delay. When you attach the disk to the server again, the mount point displayed on the management console may be inconsistent with that in the server. For example, device name `/dev/sdb` or `/dev/vdb` is selected for attachment, but `/dev/xvdc` or `/dev/vdc` may be displayed as the disk name in the OS. This issue does not adversely affect services.

```
[root@ecs-b656 test]# fdisk -l
```

```
Disk /dev/xvda: 42.9 GB, 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: 0x000cc4ad
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvda1	*	2048	2050047	1024000	83	Linux
/dev/xvda2		2050048	22530047	10240000	83	Linux
/dev/xvda3		22530048	24578047	1024000	83	Linux
/dev/xvda4		24578048	83886079	29654016	5	Extended
/dev/xvda5		24580096	26628095	1024000	82	Linux swap / Solaris

```
Disk /dev/xvdb: 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
```

 **NOTE**

The capacity displayed here is inconsistent with the capacity of the EVS disk applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 Bytes and 1 GB = 1,000,000,000 Bytes.

Step 2 Run the following command to allocate partitions for the added data disk using `fdisk`:

fdisk *Newly added data disk*

In this example, `/dev/xvdb` is the newly added data disk.

fdisk /dev/xvdb

Information similar to the following is displayed:

```
[root@ecs-b656 test]# fdisk /dev/xvdb  
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 0xb00005bd.  
Command (m for help):
```

Step 3 Enter `n` and press **Enter**.

Entering `n` creates a partition.

There are two types of disk partitions:

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

```
Command (m for help): n  
Partition type:  
 p primary (0 primary, 0 extended, 4 free)  
 e extended
```

Step 4 Enter `p` and press **Enter**.

The following describes how to create a primary partition.

Information similar to the following is displayed: (**Partition number** indicates the serial number of the primary partition. The value can be `1` to `4`.)

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

Step 5 Enter the primary partition number `1` and press **Enter**.

For example, select `1` as the partition number.

Information similar to the following is displayed: (**First sector** indicates the first sector number. The value can be `2048` to `20971519`, and the default value is `2048`.)

```
Partition number (1-4, default 1): 1  
First sector (2048-20971519, default 2048):
```

Step 6 Press **Enter**.

The default start sector number 2048 is used as an example.

Information similar to the following is displayed: (**Last sector** indicates the last sector number. The value can be from **2048** to **20971519**, and the default value is **20971519**.)

```
First sector (2048-20971519, default 2048):  
Using default value 2048  
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
```

Step 7 Press **Enter**.

The default last sector number 20971519 is used as an example.

Information similar to the following is displayed, indicating that a primary partition is created for a 10 GB data disk.

```
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):  
Using default value 20971519  
Partition 1 of type Linux and of size 10 GiB is set  
Command (m for help):
```

Step 8 Enter **p** and press **Enter** to view the details about the created partition.

Information similar to the following is displayed: (Details about the **/dev/xvdb1** partition are displayed.)

```
Command (m for help): p  
  
Disk /dev/xvdb: 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: 0xb00005bd
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvdb1		2048	20971519	10484736	83	Linux

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

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

Information similar to the following is displayed: (The partition is successfully created.)

```
Command (m for help): w  
The partition table has been altered!  
  
Calling ioctl() to re-read partition table.  
Syncing disks.
```

Step 10 Run the following command to synchronize the new partition table to the data disk:

```
partprobe
```

Step 11 Run the following command to set the format for the file system of the newly created partition:

```
mkfs -t File system format /dev/xvdb1
```

For example, run the following command to set the **ext4** file system for the **/dev/xvdb1** partition:

mkfs -t ext4 /dev/xvdb1

Information similar to the following is displayed:

```
[root@ecs-b656 test]# mkfs -t ext4 /dev/xvdb1
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
655360 inodes, 2621184 blocks
131059 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2151677952
80 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

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

NOTE

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

Step 12 Run the following command to create a mount directory:

```
mkdir Mount directory
```

/mnt/sdc is used in this example.

```
mkdir /mnt/sdc
```

Step 13 Run the following command to mount the new partition to the mount directory created in [Step 12](#):

```
mount /dev/xvdb1 Mount directory
```

For example, run the following command to mount the newly created partition on **/mnt/sdc**:

```
mount /dev/xvdb1 /mnt/sdc
```

Step 14 Run the following command to view the mount result:

```
df -TH
```

Information similar to the following is displayed. The newly created **/dev/xvdb1** partition has been mounted on **/mnt/sdc**.

```
[root@ecs-b656 test]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda2     xfs       11G   7.4G  3.2G  71% /
devtmpfs       devtmpfs  4.1G   0  4.1G   0% /dev
tmpfs          tmpfs     4.1G   82k  4.1G   1% /dev/shm
tmpfs          tmpfs     4.1G   9.2M  4.1G   1% /run
tmpfs          tmpfs     4.1G   0  4.1G   0% /sys/fs/cgroup
/dev/xvda3     xfs       1.1G   39M   1.1G   4% /home
/dev/xvda1     xfs       1.1G  131M  915M  13% /boot
/dev/xvdb1     ext4      11G   38M  9.9G   1% /mnt/sdc
```

----End

Setting Automatic Disk Attachment Upon Instance Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify `/dev/xvdb1` in `/etc/fstab`. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in `/etc/fstab` to automatically attach the disk at a system start.

NOTE

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

```
blkid Disk partition
```

For example, run the following command to query the UUID of `/dev/xvdb1`:

```
blkid /dev/xvdb1
```

Information similar to the following is displayed: (The UUID of `/dev/xvdb1` is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1  
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the `fstab` file using the vi editor:

```
vi /etc/fstab
```

Step 3 Press `i` to enter the editing mode.

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

```
UUID=xxx attachment directory file system defaults 0 2
```

```
Assuming that the file system is ext4 and the attachment directory is /mnt/sdc.  
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2
```

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the `fstab` file before detaching the disk. Otherwise, you may fail to access the OS after the disk is detached.

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

The system saves the configuration and exits the vi editor.

----End

16.1.3.7 Initializing a Linux Data Disk (parted)

A data disk attached to an ECS or created together with an ECS must be initialized before it can become available. This section uses an instance running CentOS 7.0

64bit as an example, and uses the parted partition tool to set up partitions for the data disk. Initialization operations vary with operating systems.

Prerequisites

- You have logged in to the ECS. For details, see [17.1.3.4 Logging In to an ECS](#).
- A disk has been attached to the ECS and has not been initialized.

Creating Partitions and Attaching a Disk

The following example shows how to create a new primary partition on a new data disk that has been attached to an instance. The primary partition will be created using parted and GPT is the default partition style. Furthermore, the partition will be formatted using the ext4 file system, mounted on the `/mnt/sdc` directory, and set to be automatically mounted upon a system start.

Step 1 Run the following command to view information about the added data disk:

lsblk

Information similar to the following is displayed:

NOTE

If you do not log in to the ECS and run the `umount` command but directly detach the `/dev/xvdb` or `/dev/vdb` EVS disk on the management console, the disk name in the ECS may encounter a release delay. When you attach the disk to the server again, the mount point displayed on the management console may be inconsistent with that in the server. For example, device name `/dev/sdb` or `/dev/vdb` is selected for attachment, but `/dev/xvdc` or `/dev/vdc` may be displayed as the disk name in the OS. This issue does not adversely affect services.

```
[root@ecs-centos-70 linux]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0 0 40G 0 disk
├─xvda1 202:1 0 4G 0 part [SWAP]
└─xvda2 202:2 0 36G 0 part /
xvdb 202:16 0 10G 0 disk
```

The command output indicates that the server contains two disks. `/dev/xvda` is the system disk and `/dev/xvdb` is the new data disk.

Step 2 Run the following command to enter parted to partition the added data disk:

```
parted Added data disk
```

In this example, `/dev/xvdb` is the newly added data disk.

```
parted /dev/xvdb
```

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# parted /dev/xvdb
GNU Parted 3.1
Using /dev/xvdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
```

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/xvdb: unrecognised disk label
```



```
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 10.7GB
Sector size (logical/physical): 512B/512B
Partition Table: unknown
Disk Flags:
```

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

NOTE

The capacity displayed here is inconsistent with the capacity of the EVS disk applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 Bytes and 1 GB = 1,000,000,000 Bytes.

Step 4 Run the following command to set the disk partition style:

```
mklabel Disk partition style
```

The disk partition styles include MBR and GPT. For example, run the following command to set the partition style to GPT:

```
mklabel gpt
```

NOTICE

If you change the disk partition style after the disk has been used, the original data on the disk will be cleared. Therefore, select a proper disk partition style when initializing the disk.

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: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 20971520s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start End Size File system Name Flags
```

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

Step 7 Enter **mkpart opt 2048s 100%** and press **Enter**.

In the command, **opt** is the name of the new partition, **2048s** indicates the start of the partition, and **100%** indicates the end of the partition. You can plan the number and capacity of disk partitions based on 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? Cancel
```

If the preceding warning message is displayed, enter **Cancel** to stop the partitioning. Then, find the first sector with the best disk performance and use that value to partition the disk.

Step 8 Enter **p** and press **Enter** to view the details about the created partition.

Information similar to the following is displayed:

```
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 20971520s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name  Flags
  1      2048s  20969471s  20967424s                opt
```

Details about the **/dev/xvdb1** partition are displayed.

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

Step 10 Run the following command to view the disk partition information:

lsblk

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0  0  40G  0 disk
├─xvda1 202:1  0   4G  0 part [SWAP]
└─xvda2 202:2  0  36G  0 part /
xvdb  202:16  0 100G  0 disk
└─xvdb1 202:17  0 100G  0 part
```

In the command output, **/dev/xvdb1** is the partition you created.

Step 11 Run the following command to set the format for the file system of the newly created partition:

NOTICE

The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

mkfs -t *File system format* **/dev/xvdb1**

For example, run the following command to set the **ext4** file system for the **/dev/xvdb1** partition:

mkfs -t ext4 /dev/xvdb1

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# mkfs -t ext4 /dev/xvdb1
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
655360 inodes, 2620928 blocks
```

```
131046 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2151677925
80 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

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.

Step 12 Run the following command to create a mount point:

```
mkdir Mount point
```

For example, run the following command to create the **/mnt/sdc** mount point:

```
mkdir /mnt/sdc
```

Step 13 Run the following command to mount the new partition to the mount point created in [Step 12](#):

```
mount /dev/xvdb1 Mount point
```

For example, run the following command to mount the newly created partition on **/mnt/sdc**:

```
mount /dev/xvdb1 /mnt/sdc
```

Step 14 Run the following command to view the mount result:

```
df -TH
```

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda2     xfs       39G   4.0G  35G  11% /
devtmpfs       devtmpfs  946M   0  946M   0% /dev
tmpfs          tmpfs     954M   0  954M   0% /dev/shm
tmpfs          tmpfs     954M   9.1M  945M   1% /run
tmpfs          tmpfs     954M   0  954M   0% /sys/fs/cgroup
/dev/xvdb1     ext4      11G   38M  101G   1% /mnt/sdc
```

The newly created **/dev/xvdb1** is mounted on **/mnt/sdc**.

----End

Setting Automatic Disk Attachment at a System Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify **/dev/xvdb1** in **/etc/fstab**. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in **/etc/fstab** to automatically attach the disk at a system start.

 NOTE

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

```
blkid Disk partition
```

For example, run the following command to query the UUID of `/dev/xvdb1`:

```
blkid /dev/xvdb1
```

Information similar to the following is displayed: (The UUID of `/dev/xvdb1` is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1  
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the `fstab` file using the vi editor:

```
vi /etc/fstab
```

Step 3 Press `i` to enter the editing mode.

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

```
UUID=xxx attachment directory file system defaults 0 2
```

Assuming that the file system is `ext4` and the attachment directory is `/mnt/sdc`.

```
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2
```

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the `fstab` file before detaching the disk. Otherwise, you may fail to access the OS after the disk is detached.

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

The system saves the configuration and exits the vi editor.

```
----End
```

16.1.4 Deploying the Application and Database

After logging in to an ECS, you can deploy your application and database. You can deploy your application and database based on your application type and characteristics and the official application website.

16.2 Synchronizing the Clock of the Windows ECS

16.2.1 Overview

Network Time Protocol (NTP) is used for time synchronization among computers on a network. Time synchronization ensures that all ECSs have the same time, preventing adverse impact on running applications caused by different times.

This section describes how to use the NTP service to synchronize the time among Windows ECSs (Windows7 32bit is used as an example). To synchronize the time, perform the following operations:

1. Enable the NTP service. For details, see [16.2.2 Enabling the NTP Service](#).
2. Change the default NTP server address. For details, see [16.2.3 Changing the NTP Server Address](#).

16.2.2 Enabling the NTP Service

By default, the Windows Time service is enabled on the Windows OS. To use the NTP service, you need to enable the NTP service first. To enable the NTP service, perform the following operations:

1. Log in to the ECS.
2. Verify that the NTP service is enabled.


Procedure

Log in to the Windows ECS using VNC.

Step 1 Log in to ManageOne as a VDC administrator or operator using a browser.

URL in non-B2B scenarios: **https://Address for accessing ManageOne Operation Portal**, for example, **https://console.demo.com**

URL in B2B scenarios: **https://Address for accessing ManageOne Tenant Portal**, for example, **https://tenant.demo.com**

Step 2 Click  in the upper left corner of the page and select a region. Choose **Computing > Elastic Cloud Server**.

The **Elastic Cloud Server** page is displayed.

Step 3 In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID, and click the search button to search for the ECS.

Step 4 Locate the row containing the ECS and click **Remote Login** in the **Operation** column.

The **Configure Keyboard Layout for Remote Login** dialog box is displayed.

Step 5 Select the English keyboard and click **Remote Login**.

Step 6 (Optional) If the system displays "Press CTRL+ALT+DELETE to log on", click **Send CtrlAltDel** in the upper right corner of the remote login page to log in to the ECS.

Figure 16-4 Send CtrlAltDel




Step 7 Enter the password you set during ECS creation to log in to the ECS.

----End

Verify that the NTP service is enabled.

Step 1 Choose **Start**. In the **Search programs and files** text box, enter **services.msc**, and press **Enter**.

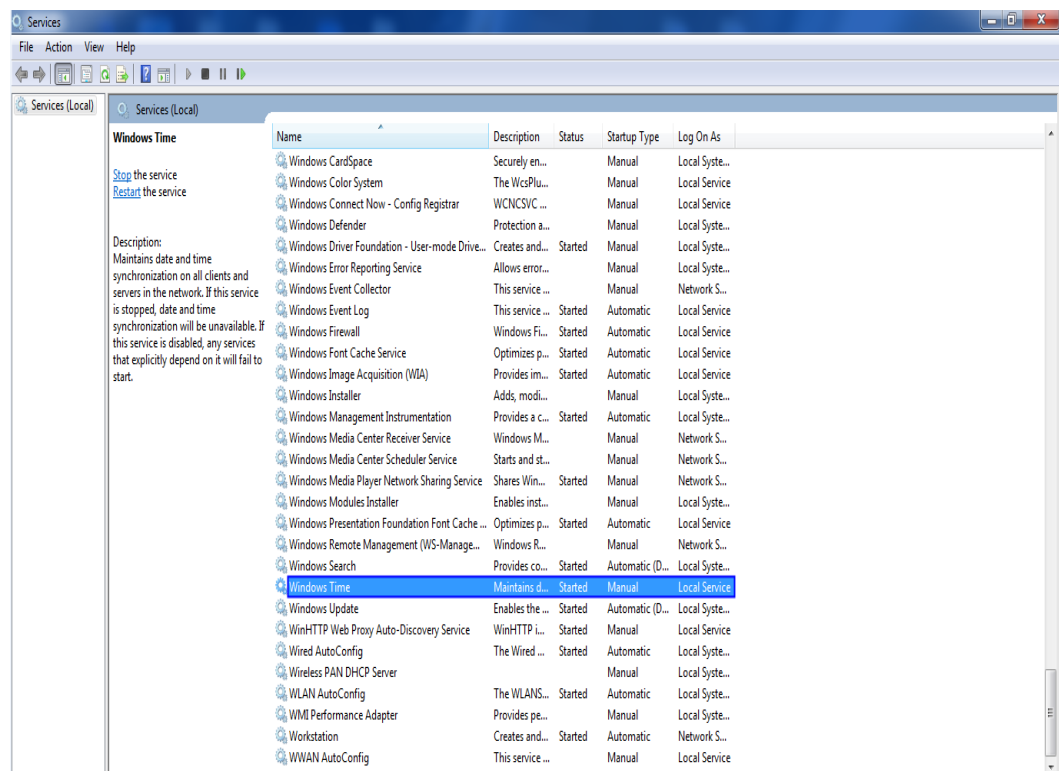


The **Services** page is displayed.

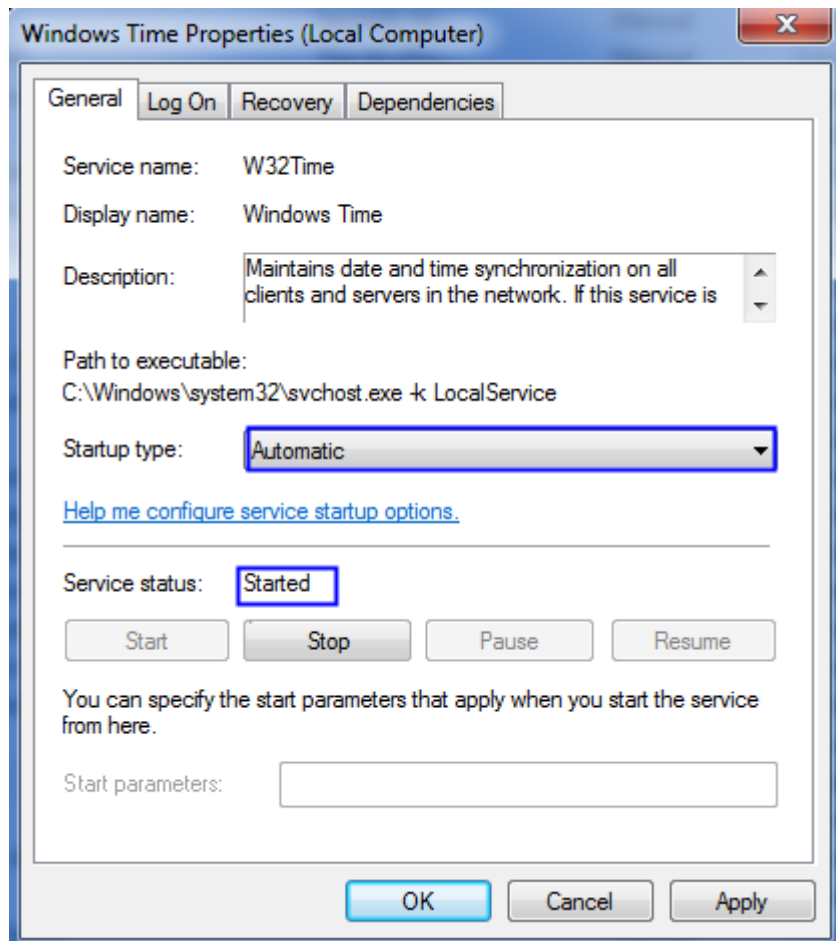
Step 2 Locate and double-click **Windows Time**.

The **Windows Time Properties (Local Computer)** dialog box is displayed.

Figure 16-5 Windows Time



Step 3 Choose **Automatic** as the startup type, and ensure that the service status is **Started**.



Step 4 Click **Apply**, and then click **OK**.

----End

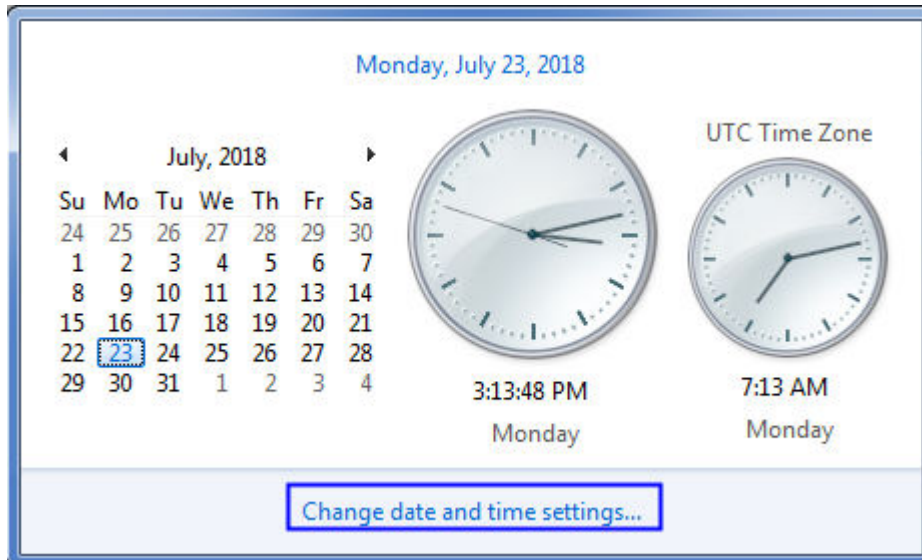
16.2.3 Changing the NTP Server Address

By default, Windows ECSs use **time.windows.com** as the NTP server. To ensure that all ECSs on a network have the same time, you can specify an NTP server for all ECSs.

Procedure

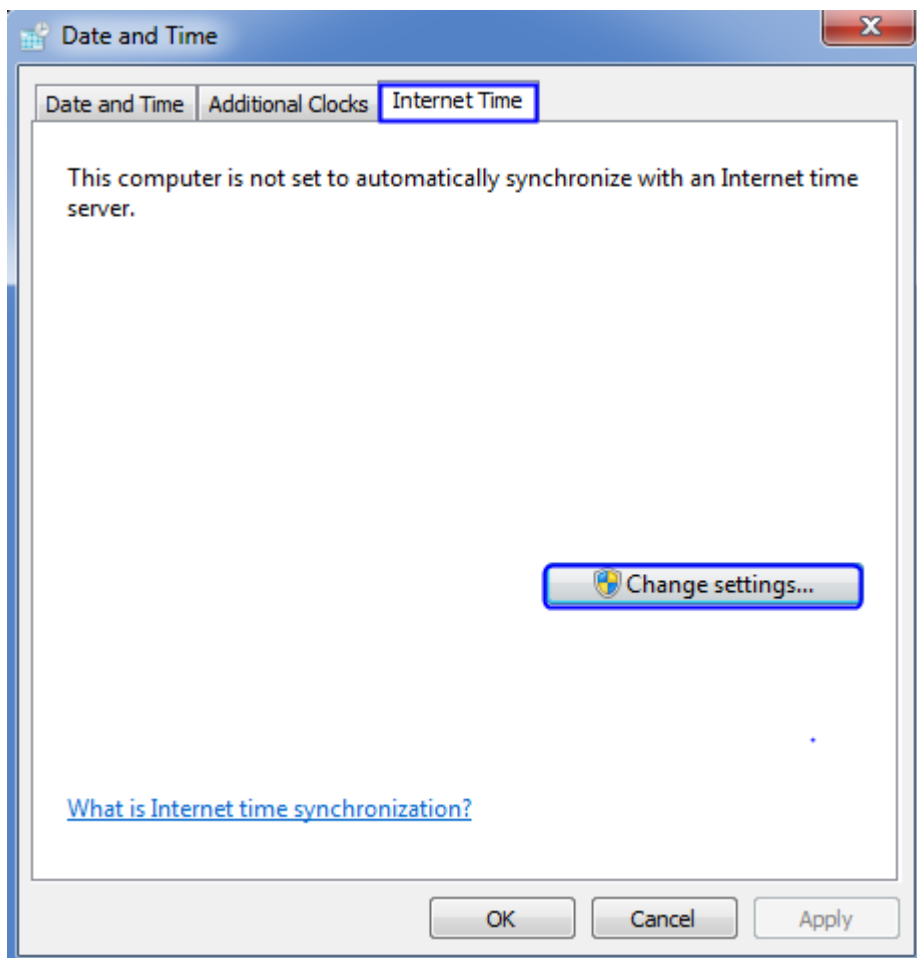
Step 1 Log in to the Windows ECS. For details, see [16.2.2 Enabling the NTP Service](#).

Step 2 On the toolbar, click **Date and Time**, and choose **Change date and time settings...**




Step 3 In the **Date and Time** dialog box, click the **Internet Time** tab, and then click **Change settings...**

The **Internet Time Settings** dialog box is displayed.



Step 4 Select **Synchronize with an Internet time server**, enter the server address, and click **Update now**.



----End

16.3 Manually Viewing the Disk Mount Point

On the ECS details page, you can view details about the attached disk, including **Device Name** and **Device Address**. **Device Name** indicates disk name on the ECS. **Device Address** indicates the address assigned by the ECS to the disk. If UVP VMTools is not installed on the ECS or the version of UVP VMTools is too early, only the device address is displayed. This section describes how to query the disk mount point on the ECS based on the device address when the device name is not displayed.

Solution Overview

Device addresses displayed on the page are divided into three types, as shown in the following table.

Table 16-9 Device address details

Device Address Format	Disk Device Type	Description	How to Obtain
a-x:x:x	The disk device type of the image used by the ECS is scsi .	a indicates the position of the SCSI controller where the disk is attached, and x:x:x indicates the disk address on the SCSI controller.	<ul style="list-style-type: none"> On a Linux ECS, run the lspci command to query the disk mount point. For details, see Linux ECSs. On a Windows ECS, view the position of the SCSI controller on Device Manager. Then, traverse all disks according to the disk addresses on the SCSI controller, compare the disk position information, and find the disk to be queried. For details, see Windows ECSs.
a:x	The disk device type of the image used by the ECS is ide .	a indicates the interface number, and x indicates the active and standby devices. For example, 0:0 indicates that the disk is located on the place representing the active device of IDE interface 0, and 0:1 indicates that the disk is located on the place representing the standby device of IDE interface 0.	<ul style="list-style-type: none"> If there are no ide disk devices in the images of Linux ECSs, device addresses in the a:x format are not displayed. On a Windows ECS, traverse all disks, view the disk location information, and find the disk to be queried. For details, see Disk Device Address: a:x.

Device Address Format	Disk Device Type	Description	How to Obtain
xxxx:xx:xx.x	The disk device type of the image used by the ECS is virtio .	The address format is domain:bus:slo t.function .	<ul style="list-style-type: none"> For a Linux ECS, run the ll /dev/disk/by-path/pci-xxxx:xx:xx.x command to query the disk mount point. For details, see Linux ECSs. On a Windows ECS, traverse all disks, view the disk location information, and find the disk to be queried. For details, see Windows ECSs.

Disk Device Address: a-x:x:x

- Linux ECSs

The following uses the device address **2-0:12:0** and CentOS 7.6 as examples. **2** indicates that the disk is mounted to the second SCSI controller, and **0:12:0** indicates the disk address on the SCSI controller.

- Log in to the ECS as the **root** user.
- Run the **lspci -D|grep "Virtio SCSI"** command to query the addresses of all SCSI controllers on the ECS. In the command output, view the address of the controller whose position is specified by **a** in the device address **a-x:x:x**.

Use **2-0:12:0** as an example, the address of the second SCSI controller is **0000:00:06.0**.

```
[root@host-88-97-3-44 ~]# lspci -D|grep "Virtio SCSI"
0000:00:05.0 SCSI storage controller: Red Hat, Inc. Virtio SCSI
0000:00:06.0 SCSI storage controller: Red Hat, Inc. Virtio SCSI
0000:00:07.0 SCSI storage controller: Red Hat, Inc. Virtio SCSI
0000:00:08.0 SCSI storage controller: Red Hat, Inc. Virtio SCSI
```

- Query the mount point name of the disk in the system based on **0000:00:06.0**, the obtained SCSI controller address, and **0:12:0**, the disk address on the SCSI controller.

lsscsi --verbose|grep "0000:00:06.0" -C 1|grep 0:12:0 -C 1

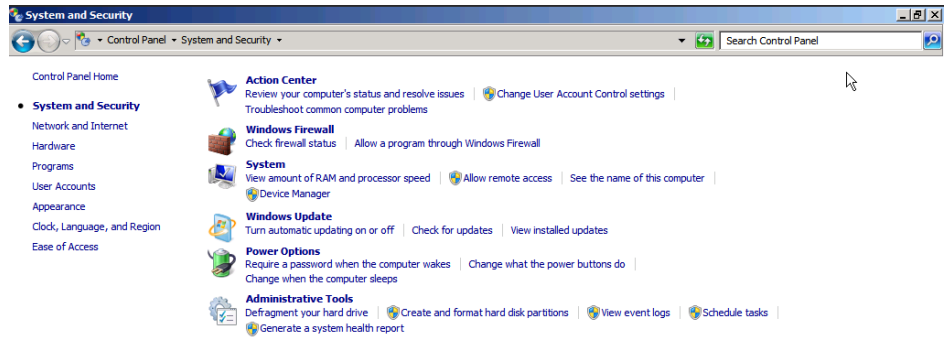
```
[root@host-88-97-3-44 ~]# lsscsi --verbose|grep "0000:00:06.0" -C 1|grep 0:12:0 -C 1
[3:0:12:0] disk HUAWEI XSG1 4303 /dev/sdb
dir: /sys/bus/scsi/devices/3:0:12:0 [/sys/devices/pci0000:00/0000:00:06.0/virtio3/host3/target3:0:12:0]
```

In the command output, the last column of the first line shows the disk mount point, that is, **/dev/sdb** in the figure.

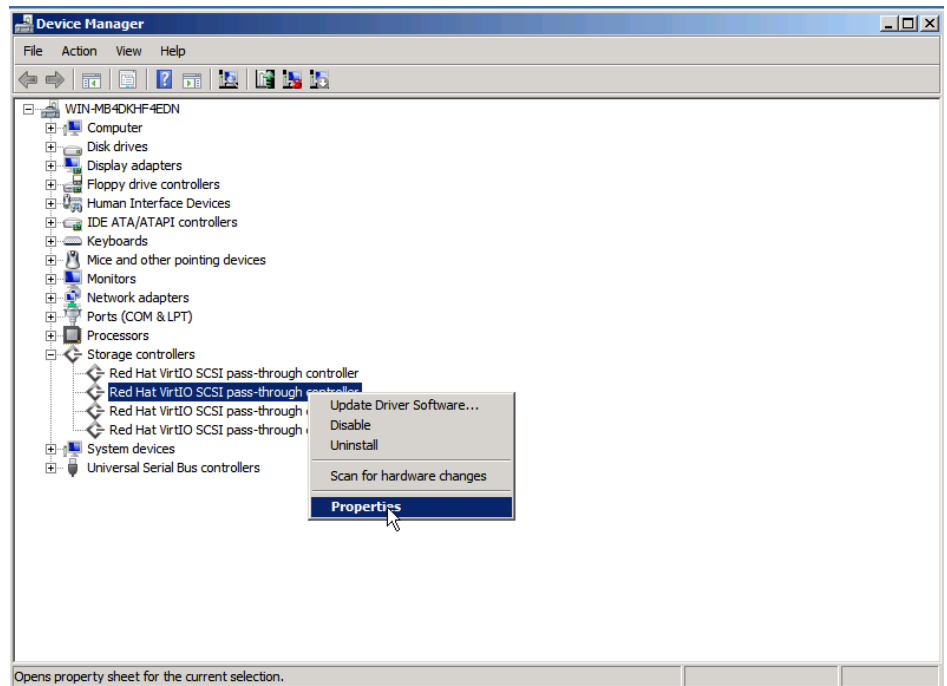
- Windows ECSs

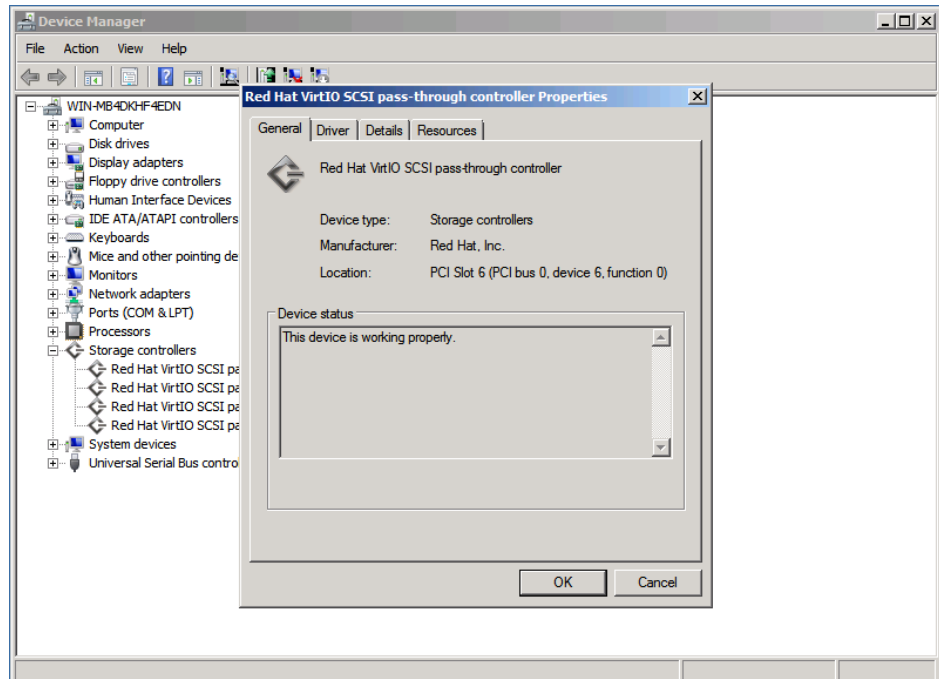
The following uses the device address **2-0:10:0** and Windows 2008 as examples. **2** indicates that the disk is mounted to the second SCSI controller, and **0:10:0** indicates the disk address on the SCSI controller.

- Log in to the ECS, and choose **Start > Control Panel > System and Security**.




- b. Under **Storage controllers** of Device Manager, view the properties of the second SCSI controller. **PCI Slot 6** on the **General** tab page indicates that the address of the second SCSI controller is **6**.





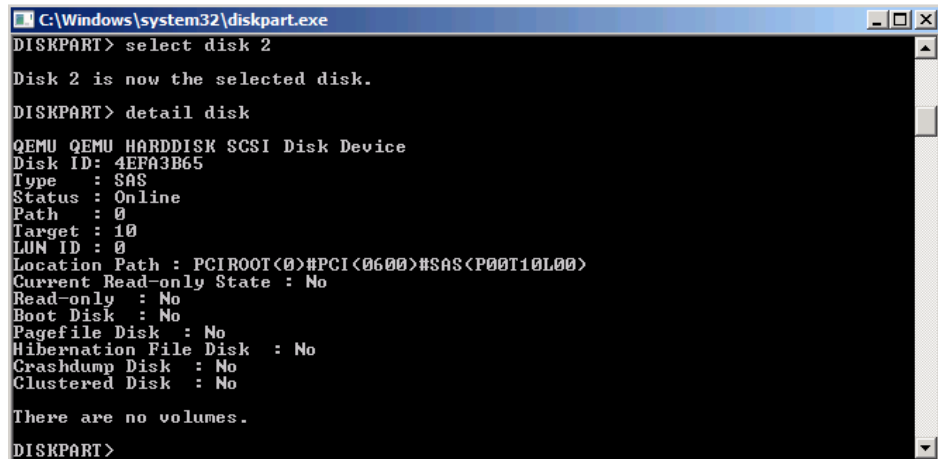
- c. Run the following commands in sequence to obtain the disk details. Disk 1 is used as an example.

```
select disk 1
detail disk
```



The path of disk 1 in the command output is **PCIROOT (0) #PCI (0500)#SAS (P00T01L00)**. According to the SCIS controller address 6 obtained in **b** and **0:10:0**, the disk address on the SCSI controller, the suffix of the path of the disk to be queried should be **PCI(0600)#SAS(P00T10L00)**. Therefore, disk 1 is not the one to be queried.

- d. Traverse disks. It is found that the location of disk 2 is **PCIROOT(0)#PCI(0600)#SAS(P00T10L00)**, which indicates that disk 2 is the one configured for the device in the system.



```
C:\Windows\system32\diskpart.exe
DISKPART> select disk 2
Disk 2 is now the selected disk.
DISKPART> detail disk
QEMU QEMU HARDDISK SCSI Disk Device
Disk ID: 4EFA3B65
Type : SAS
Status : Online
Path : 0
Target : 10
LUN ID : 0
Location Path : PCIROOT(0)#PCI(0600)#SAS(P00T10L00)
Current Read-only State : No
Read-only : No
Boot Disk : No
Pagefile Disk : No
Hibernation File Disk : No
Crashdump Disk : No
Clustered Disk : No
There are no volumes.
DISKPART>
```

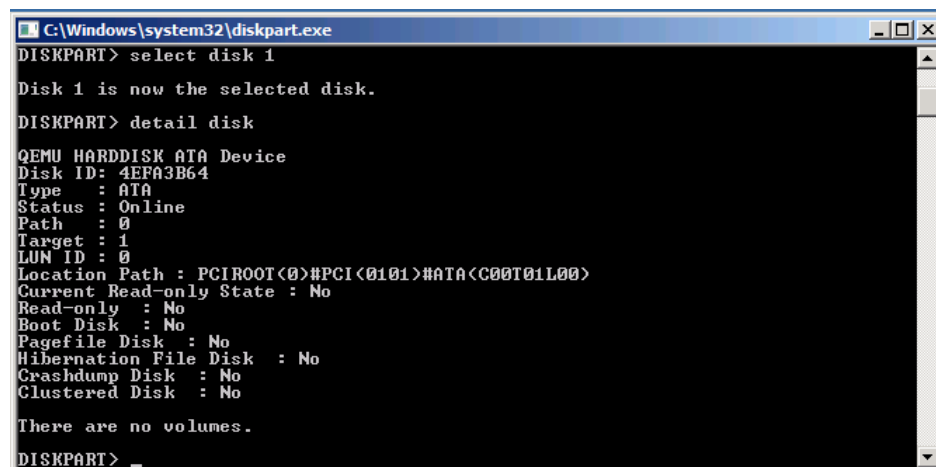
Disk Device Address: a:x

Windows ECSs

The following uses the device address **1:0** and Windows 2008 as examples.

1. Log in to the ECS and run the following commands in sequence to obtain the disk details. Disk 1 is used as an example.

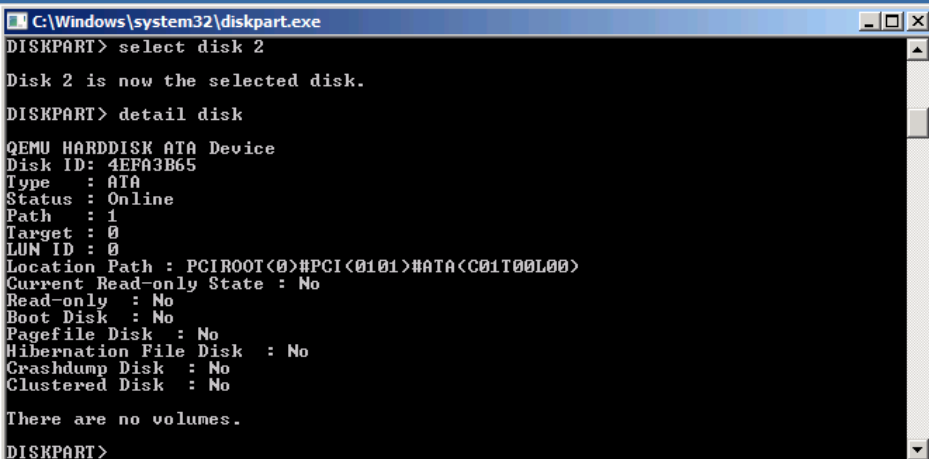
```
select disk 1
detail disk
```



```
C:\Windows\system32\diskpart.exe
DISKPART> select disk 1
Disk 1 is now the selected disk.
DISKPART> detail disk
QEMU HARDDISK ATA Device
Disk ID: 4EFA3B64
Type : ATA
Status : Online
Path : 0
Target : 1
LUN ID : 0
Location Path : PCIROOT(0)#PCI(0101)#ATA(C00T01L00)
Current Read-only State : No
Read-only : No
Boot Disk : No
Pagefile Disk : No
Hibernation File Disk : No
Crashdump Disk : No
Clustered Disk : No
There are no volumes.
DISKPART>
```

The path of disk 1 in the command output is **PCIROOT(0)#PCI(0101)#ATA(C00T01L00)**. Because the device address of the disk to be queried is **1:0**, the ATA part in the location information should contain **C01T00**. Therefore, disk 1 is not the one to be queried.

2. Traverse disks. It is found that the location of disk 2 is **PCIROOT(0)#PCI(0101)#ATA(C01T00L00)**, which indicates that disk 2 is the one configured for the device in the system



```
C:\Windows\system32\diskpart.exe
DISKPART> select disk 2
Disk 2 is now the selected disk.
DISKPART> detail disk
QEMU HARDDISK ATA Device
Disk ID: 4EFA3B65
Type : ATA
Status : Online
Path : 1
Target : 0
LUN ID : 0
Location Path : PCIROOT(0)#PCI(0101)#ATA(C01T00L00)
Current Read-only State : No
Read-only : No
Boot Disk : No
Pagefile Disk : No
Hibernation File Disk : No
Crashdump Disk : No
Clustered Disk : No
There are no volumes.
DISKPART>
```

Disk Device Address: xxxx:xx:xx.x

- **Linux ECSs**

The following uses the device address **0000:02:01.0** and CentOS 7.6 as examples.

- a. Log in to the ECS as the **root** user.
- b. Run the following command to view the device name:

```
ll /dev/disk/by-path/pci-0000:02:01.0
```

The command output displays the device name, that is, **/dev/vda** in the following figure.



```
root@host-88-97-3-44 ~# ll /dev/disk/by-path/pci-0000:02:01.0
lrwxrwxrwx. 1 root root 9 Jul 17 03:32 /dev/disk/by-path/pci-0000:02:01.0 -> ../../vda
root@host-88-97-3-44 ~#
```

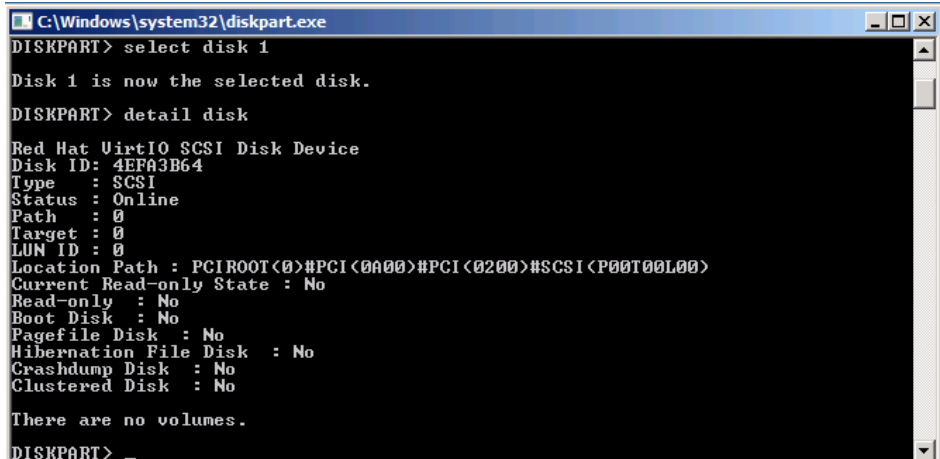
- **Windows ECSs**

The following uses the device address **0000:02:0d.0** and Windows 2008 as examples. The values of **domain** and **function** are always **0**, **bus** is **02**, and **slot** is **0d**.

- a. Log in to the ECS and run the following commands in sequence to obtain the disk details. Disk 1 is used as an example.

```
select disk 1
```

```
detail disk
```

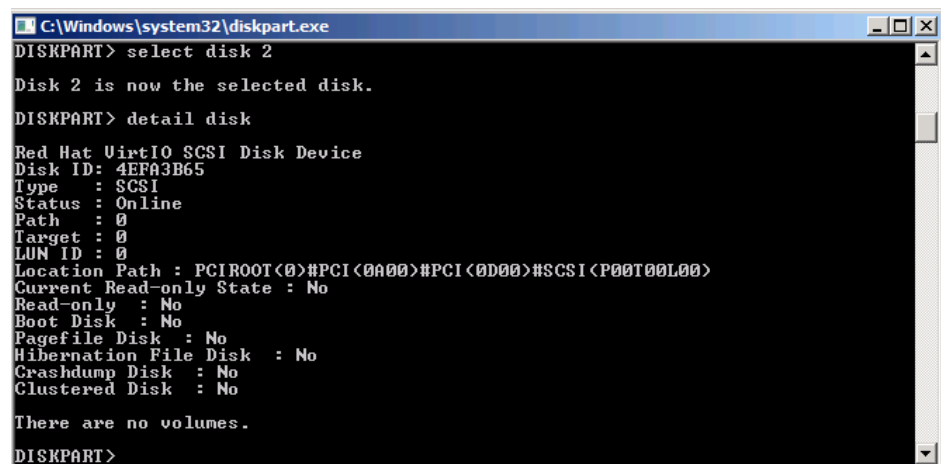


```
C:\Windows\system32\diskpart.exe
DISKPART> select disk 1
Disk 1 is now the selected disk.
DISKPART> detail disk
Red Hat VirtIO SCSI Disk Device
Disk ID: 4EFA3B64
Type : SCSI
Status : Online
Path : 0
Target : 0
LUN ID : 0
Location Path : PCIROOT(0)#PCI(0A00)#PCI(0200)#SCSI(P00T00L00)
Current Read-only State : No
Read-only : No
Boot Disk : No
Pagefile Disk : No
Hibernation File Disk : No
Crashdump Disk : No
Clustered Disk : No
There are no volumes.
DISKPART>
```

The following table describes the mapping between the first PCI in the location path and the bus in the address. The second PCI indicates the slot information. Therefore, the PCI part in the disk path for **0000:02:0d.0** should be **PCI(0A00)#PCI(0D00)** and disk 1 is not the one to be queried.

Bus ID	PCI Address ID
02	0A00
03	0B00

- b. Traverse all disks. The location path of disk 2 is **PCIROOT(0)#PCI(0A00)#PCI(0D00)#SCSI<P00T00L00>**. Therefore, disk 2 is the one configured for the device in the system.



```
C:\Windows\system32\diskpart.exe
DISKPART> select disk 2
Disk 2 is now the selected disk.
DISKPART> detail disk
Red Hat VirtIO SCSI Disk Device
Disk ID: 4EFA3B65
Type : SCSI
Status : Online
Path : 0
Target : 0
LUN ID : 0
Location Path : PCIROOT(0)#PCI(0A00)#PCI(0D00)#SCSI<P00T00L00>
Current Read-only State : No
Read-only : No
Boot Disk : No
Pagefile Disk : No
Hibernation File Disk : No
Crashdump Disk : No
Clustered Disk : No
There are no volumes.
DISKPART>
```

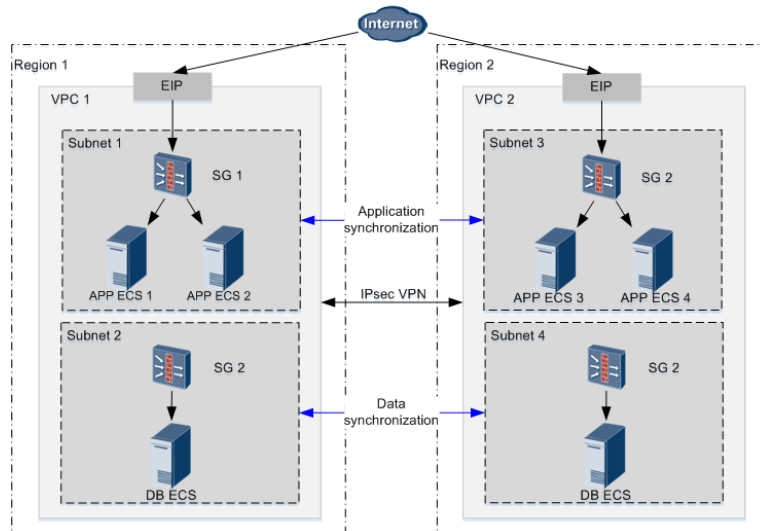
16.4 Using Direct Connect to Establish VPN Channels to Implement Cross-Region DR

Scenario

A customer applies for multiple ECSs in different regions to provide cross-region data synchronization and cross-region DR. [Figure 16-6](#) shows the networking diagram.

- The VPC in Region 1 is the production environment, and the VPC in Region 2 is the DR environment.
- Two ECSs are applied for in each region to run application services, called APP ECSs. External networks access the application services using public IP addresses.
- An ECS is applied for in each region to run the database, called DB ECS. The application services and database are deployed in different subnets.

Figure 16-6 Networking diagram



The following cloud services are used: ECS, VPC, SG, EIP, and VPN.

Requirement Analysis

The analysis based on the user requirements is as follows:

1. You can use VPNs to interconnect VPCs in different regions so that user data and services in these regions can interact with each other.
2. For security purposes, access control is performed, and the application services and databases in each partition need to be isolated from each other. Therefore, there are two subnets in the VPC. One subnet is used to deploy the APP ECSs, and the other subnet is used to deploy the DB ECS. The two subnets use different security groups to improve security.

Based on the preceding analysis, VPNs can be used to connect VPCs in different regions. After VPCs in the two regions are interconnected, data DR is implemented by using external technologies such as database/application replication and DNS.

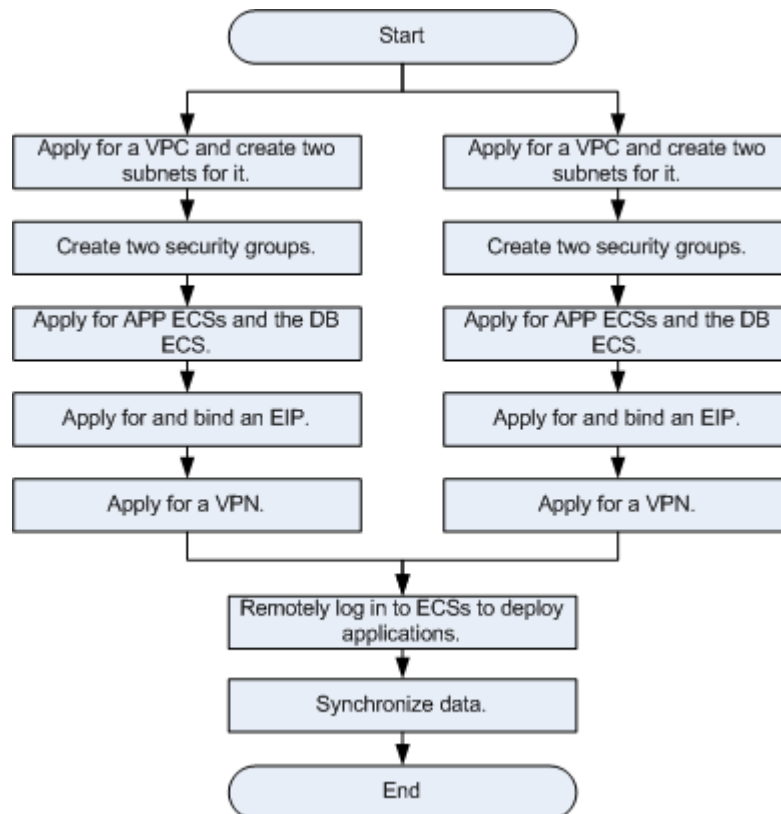
Configuration Plan

The application and database deployment policies vary depending on the application and database types. Therefore, you do not need to describe in detail how to deploy the application services and database on the available ECSs.

The configuration plan is as follows:

1. Create a VPC and create subnets for the VPC.
2. Create a security group and add rules to it.
3. Create ECSs in each subnet.
4. Apply for an EIP and bind it to the ECSs.
5. Create an IPsec VPN.

Figure 16-7 Configuration plan



Configuration Procedure

Creating a VPC


Step 1 Log in to ManageOne as a VDC administrator or VDC operator using a browser.

URL in non-B2B scenarios: <https://Domain name of ManageOne Operation Portal>, for example, <https://console.demo.com>.

URL in B2B scenarios: <https://Domain name of ManageOne Tenant Portal>, for example, <https://tenant.demo.com>.

You can log in using a password or a USB key.

- Login using a password: Enter the username and password.
The password is that of the VDC administrator or VDC operator.
- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.

Step 2 Click  in the upper left corner of the main menu, select a region, and choose **Network > Virtual Private Cloud**.

Step 3 Click **Apply for VPC**. The **Select Service** dialog box is displayed.

Step 4 Select a service and click **Apply Now**. The **VPC Configuration Wizard** page is displayed.

Step 5 Select **VPC with Internet Access (Optional)**, and click **Next**. Configure the VPC parameters, check your configuration, and click **Apply Now**.

For details, see "Creation" > "Applying for a VPC" > "(Optional) Applying for a VPC That Can Access the Internet" in **Operation Help Center > Network > Virtual Private Cloud > User Guide**.

Creating a subnet and a security group

Step 6 When you apply for a VPC in **Step 5**, subnet 1 is created. To create subnet 2 for the VPC created in **Step 5**, click the VPC name, click **Create Subnet** on the **Subnet** tab page, and configure parameters as prompted.


For details, see "VPC and Subnet" > "Subnet" > "Creating a Subnet" in **Operation Help Center > Network > Virtual Private Cloud > User Guide**.

Step 7 In the navigation pane on the left of the VPC home page, click **Security Group**. On the **Security Group** page, click **Create Security Group**, specify a name and enter a description, and click **OK**.

Step 8 Locate the row containing the newly created security group, and in the **Operation** column, click **Add Rule**. Click the **Inbound** or **Outbound** tab, and then add a rule as prompted.

For details, see "Security Group" > "Creating a Security Group" in **Operation Help Center > Network > Virtual Private Cloud > User Guide**.

Applying for an ECS

Step 9 Click  in the upper left corner, and choose **Computing > Elastic Cloud Server**. On the ECS console, click **Create ECS**. On the **Select Service** page, select a service and click **Apply Now**.

Step 10 Configure basic information about the ECS to be created, and set the number of ECSs to **2**.

Step 11 Click **Next**. On the **Network Configuration** page, configure network information about the ECS to be created, in which select the VPC obtained in **Step 5** for **VPC**.


Step 12 Click **Next** to enter the **Advanced Settings** page. Enable the HA function. When HA is enabled, an ECS is automatically rebuilt on another host whenever the ECS or its host becomes faulty, ensuring service continuity.

Step 13 Click **Next**. On the page for confirming specifications, confirm that the specifications are correct and click **Submit** to apply for two ECSs (ECS 1 and ECS 2) for the application services.

Step 14 Repeat **Step 9** to **Step 13** to apply for an ECS for the database.

For details, see **5 Creating an ECS**.

Applying for and binding an EIP

Step 15 Click  in the upper left corner, and choose **Network > Elastic IP**. On the **Elastic IP** page, click **Apply for EIP**. The **Select Service** dialog box is displayed. Select a service and click **Apply Now**.

Step 16 On the **Apply for EIP** page, configure the parameters of the basic information and bandwidth, and click **Apply Now**.


Step 17 On the **Elastic IP Address** page, locate the row that contains the target EIP, and click **Bind**.

Step 18 On the **Bind EIP** page, select the APP cloud resource that has been applied for.

Step 19 Click **OK**.

For details, see .

Applying for a VPN

Step 20 Click  in the upper left corner of the main menu, select a region, and choose **Network > Virtual Private Network**.

Step 21 In the navigation pane on the left, choose **Virtual Private Network > VPN Gateway**.

Step 22 On the **VPN Gateway** page, click **Create VPN Gateway** to display the **Create VPN Gateway** page.

Step 23 Set the parameters as prompted, and click **Create Now**.

Step 24 In the VPN connection details area, click **Apply for VPN Connection**. The **Select Service** dialog box is displayed.

Step 25 Select a service and click **Apply Now**.

Step 26 On the **Apply for VPN Connection** page, configure the parameters as prompted.

For details, see .

Remotely logging in to the ECSs to deploy services

Step 27 After applying for the APP ECSs and DB ECS, you can remotely log in to the ECSs and deploy applications and DB software. After the deployment is complete, you need to save the database configuration information in the template of the DB ECS to the application ECSs.

Synchronizing data

Step 28 After VPC interconnection in two regions is verified, data DR is implemented by using external technologies such as database/application replication and DNS.

----End

16.5 Distributing Traffic on APP ECSs

Scenarios

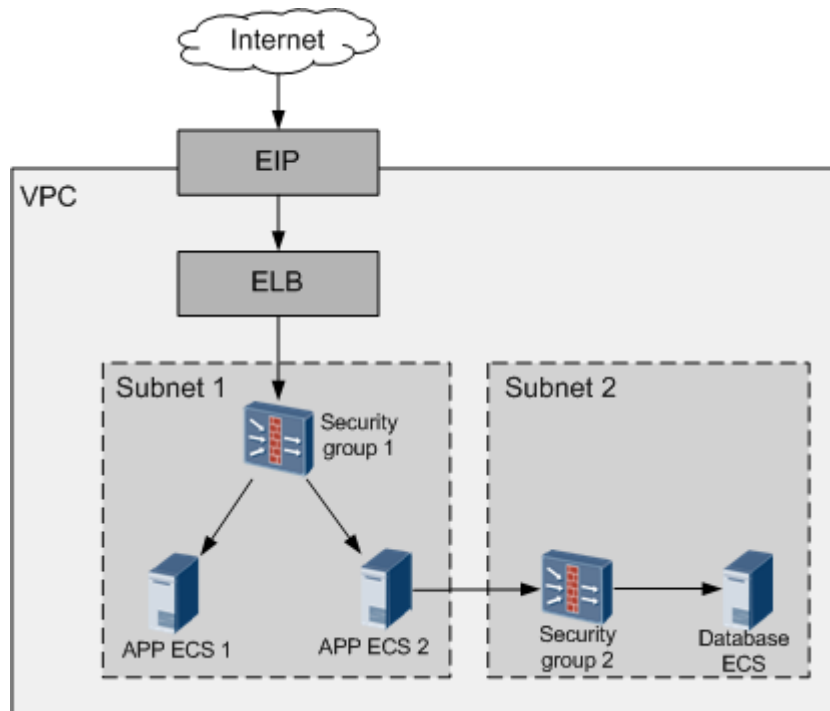
A user has applied for multiple Elastic Cloud Servers (ECSs) to run the application service and used the Elastic Load Balance (ELB) service to distribute access traffic to multiple ECSs. The networking is shown in [Figure 16-8](#).

- Apply for three ECSs, two running the application service and one running the database.
- Create secure and isolated environments for APP ECSs and the database ECS by applying for a VPC and subnets and creating security group rules. APP ECSs

and the database ECS are deployed in different subnets and security groups. The subnet where APP ECSs reside can connect to the Internet using the elastic IP address (EIP) and the subnet where the database ECS resides supports internal data access only. APP ECSs can access the database ECS.

- Bind an EIP to the load balancer to assign a public IP address to the load balancer and the external network can use the public IP address to access APP ECSs.

Figure 16-8 Service networking



Requirement Analysis

The analysis based on the user requirements is as follows:

1. The Virtual Private Cloud (VPC) service enables you to provision logically isolated, configurable, and manageable virtual networks for ECSs. You can select the IP address range, create multiple subnets, customize security groups, and configure route tables and gateways in the VPC. Based on the Internet access control requirements, you must deploy the APP and database ECSs in different subnets and configure independent security group rules for them to enhance their access control.
2. Elastic Load Balance (ELB) is a service that automatically distributes incoming traffic across multiple backend ECSs based on a specified forwarding policy. After you add a listener and backend ECS to the load balancer, the listener automatically checks the running status of the backend ECS. You need to apply for a load balancer for the subnet where the created APP ECSs reside and bind an EIP to the load balancer, which not only enables the APP ECS to access the Internet, but also improves the fault tolerance and increases the availability of your applications through traffic distribution.

Based on the preceding analysis, the ELB service can provide a unified external access entrance for external networks and distribute access traffic to multiple ECSs.

Configuration Roadmap

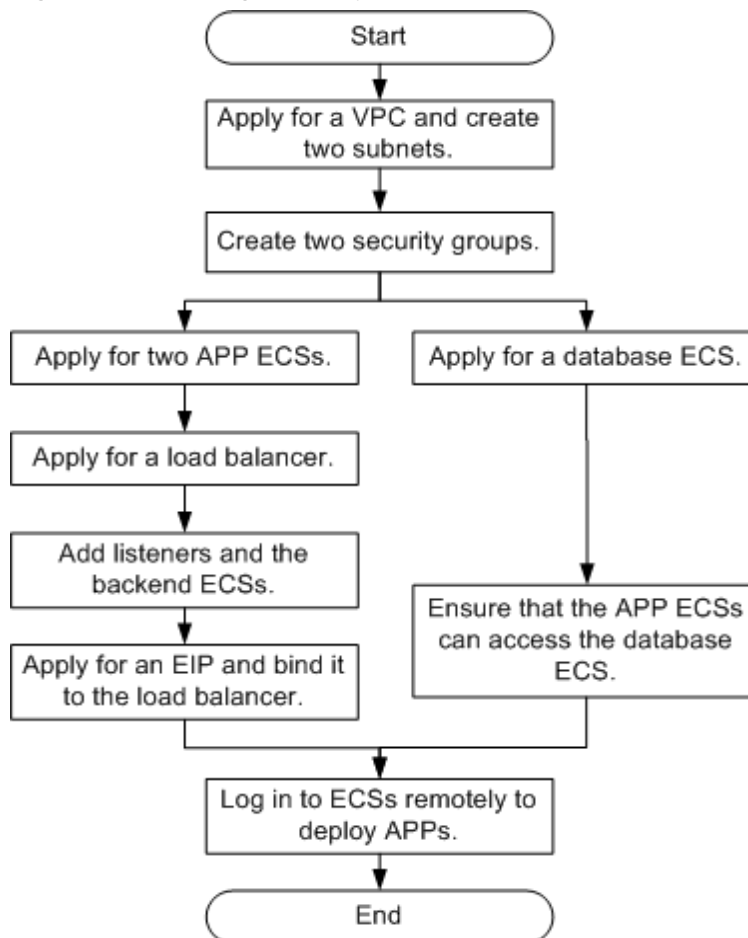
The APP and database deployment policies vary depending on their types. Therefore, this part does not provide the details about how to deploy applications and databases on the obtained ECSs.

The configuration roadmap is as follows:

1. Apply for a VPC and create two subnets and two security groups to ensure the privacy and security of the networks where APP ECSs and the database ECS reside.
2. Apply for three ECSs, two running the application service and one running the database. Select the created subnets separately when applying for the ECSs.
3. Apply for a load balancer for subnet 1 and add a listener and the APP ECSs to the load balancer to distribute traffic and perform health check for the APP ECSs.
4. Apply for an EIP and bind it to the load balancer to ensure that the APP ECSs can access the Internet.

Figure 16-9 shows the configuration process based on the preceding analysis.

Figure 16-9 Configuration process



Configuration Procedure

Apply for a VPC


Step 1 Log in to ManageOne as a VDC administrator or VDC operator using a browser.

URL in non-B2B scenarios: **https://Domain name of ManageOne Operation Portal**, for example, **https://console.demo.com**.

URL in B2B scenarios: **https://Domain name of ManageOne Tenant Portal**, for example, **https://tenant.demo.com**.

You can log in using a password or a USB key.

- Login using a password: Enter the username and password.
The password is that of the VDC administrator or VDC operator.
- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.

Step 2 Click  in the upper left corner, select a region, and choose **Network > Virtual Private Cloud**.

Step 3 Select a project in the region from the drop-down list in the upper left corner of the main menu.

Step 4 On the **Virtual Private Cloud** page, click **Apply for VPC**.

The **Select Service** dialog box is displayed.

Step 5 Select a service and click **Apply Now**.

The **VPC Configuration Wizard** page is displayed.


Step 6 Select **VPC with Internet Access (Optional)**, and click **Next**.

The **Apply for VPC** page is displayed.


Step 7 Configure the VPC parameters by referring to [Table 16-10](#).

Figure 16-10 VPC parameters in Region Type I

Basic Information

Region and Project  Select your region and then project from the top menu bar.

* Name

* External Network 

CIDR Block /

CIDR Block is mandatory if Subscribe to EI&PaaS Services is set to Yes.

Unavailable CIDR blocks: 10.0.0.0/24, 10.1.0.0/24, 10.2.0.0/24, 10.3.0.0/24, 10.4.0.0/24, 10.5.0.0/24, 10.6.0.0/24, 10.7.0.0/24, 10.8.0.0/24, 10.9.0.0/24, 10.10.0.0/24, 10.11.0.0/24, 10.12.0.0/24, 10.13.0.0/24, 10.14.0.0/24, 10.15.0.0/24, 10.16.0.0/24, 10.17.0.0/24, 10.18.0.0/24, 10.19.0.0/24, 10.20.0.0/24, 10.21.0.0/24, 10.22.0.0/24, 10.23.0.0/24, 10.24.0.0/24, 10.25.0.0/24, 10.26.0.0/24, 10.27.0.0/24, 10.28.0.0/24, 10.29.0.0/24, 10.30.0.0/24, 10.31.0.0/24, 172.16.0.0/16, 192.168.0.0/16, 100.64.0.0/10, 172.31.0.0/16, 192.168.0.0/16, 100.64.0.0/10, 172.31.0.0/16, 192.168.0.0/16, 100.64.0.0/10

* Required Duration

Table 16-10 VPC parameters

Parameter	Description	Example Value
Region and Project	The current region and project are displayed by default. To change them, use the selector in the upper left corner of the page.	az1.dc1 (test)
Name	Specifies the VPC name. The name can contain only letters, digits, underscores (_), hyphens (-), and dots (.).	VPC-001
External Network	Choose an AZ from the first drop-down list, and choose a VPC external network from the second drop-down list. If no external networks are available, contact the administrator to configure external networks as described in "Prerequisites".	az0.dc0 net-01
CIDR Block (available in Region Type I)	Specifies the Classless Inter-Domain Routing (CIDR) block for the VPC. The CIDR block of a subnet can be the same as the CIDR block for the VPC (for a single subnet in the VPC) or a subset (for multiple subnets).	192.168.0.0/16
Required Duration	Specifies the required duration for a VPC.	1 year

Step 8 Configure subnet 1 by referring to [Table 16-11](#). Subnet 1 is isolated from the Internet, which ensures high security. Cloud servers in subnet 1 cannot directly communicate with the Internet.

Table 16-11 Parameters for configuring subnet 1

Parameter	Description	Example Value
Subnet 1 Name	Specifies the name of subnet 1. The name can contain only letters, digits, underscores (_), hyphens (-), and dots (.).	Subnet-f03c
Type	When deploying IPv4 & IPv6 dual-stack services, you need to select a network type. When deploying single-stack services, configure subnet 1 parameters by referring to Table 16-12 . <ul style="list-style-type: none"> IPv4 Specifies the IPv4 subnet. You can configure parameters by referring to Table 16-12. IPv4&IPv6 Specifies that the subnet has both IPv4 and IPv6 address ranges. You can configure parameters by referring to Table 16-12 and Table 16-13. 	IPv4

Table 16-12 Parameters for configuring an IPv4 subnet

Parameter	Description	Example Value
Subnet 1 Segment	Specifies the IP address range of subnet 1.	192.168.0.0/24
Gateway	Specifies the gateway address of subnet 1.	192.168.0.1
Allocation Pool	<p>Specifies the range of IP addresses that can be automatically assigned to NICs if you choose to automatically assign an IP address when creating a cloud server or adding a NIC to a cloud server. This parameter is optional. The IP address range of the allocation pool must be within the subnet CIDR block.</p> <p>To reserve some IP addresses in a subnet so that they will not be automatically assigned to NICs, configure an allocation pool. When configuring the allocation pool, enter an IP address range that does not contain these IP addresses.</p> <p>NOTE</p> <ul style="list-style-type: none">• If a VPC requires two or more allocation pools, click Add Allocation Pool.• When creating a cloud server or adding a NIC to a cloud server, if you manually assign an IP address to the NIC, you can specify an IP address that belongs to the subnet but does not belong to the allocation pool.	192.168.0.2-192.168.0.221 192.168.0.225-192.168.0.251
DNS Server Address 1	Specifies the IP address of an associated DNS server. This parameter is optional.	192.168.71.3
DNS Server Address 2	<p>Specifies the IP address of an associated DNS server. This parameter is optional.</p> <p>NOTE</p> <p>If the DNS server address is left blank, the subnet is not associated with any DNS server.</p> <p>When using only one DNS server address, enter it into DNS Server Address 1.</p> <p>To add a DNS server address, click Add DNS Server Address.</p>	192.168.72.3

Parameter	Description	Example Value
Static Route Switch	<p>If this switch is set to OFF, static routes will not be configured for cloud servers in the subnet. If this switch is set to ON, the configured static routes will be injected to those cloud servers by using the DHCP function of the subnet.</p> <ul style="list-style-type: none"> • Destination: specifies the destination IP address range of the static route. • Next Hop: specifies the next-hop IP address of the static route. <p>NOTE When you need to add more static routes for cloud servers in a subnet, click Add Static Route. You can configure up to five static routes at a time.</p>	<ul style="list-style-type: none"> • Destination: 10.10.0.0/24 • Next Hop: 192.168.20.2

Table 16-13 Parameters for configuring an IPv6 subnet

Parameter	Description	Example Value
Advanced Settings	<p>You can use the default IPv6 subnet configuration or configure advanced settings. If you enable Advanced Settings, you need to configure additional parameters of the IPv6 subnet.</p>	enabled
Address Mode	<p>Only dhcpv6-stateful is available. In this mode, all IP addresses and other stateless configurations (such as DNS and NTP) are managed to implement fine-grained management and allocation of IP addresses.</p>	dhcpv6-stateful

Parameter	Description	Example Value
CIDR	<p>This parameter is optional. Enter a valid subnet CIDR block, which must be within one of the displayed available network segments. The text format of an IPv6 address is <code>xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx</code>. Each <code>x</code> is a hexadecimal digit, which represents a 4-bit binary number. Preliminary zeroes can be omitted during configuration.</p> <p>NOTE</p> <ul style="list-style-type: none"> If the IPv6 mode is dhcpv6-stateful, the prefix length ranges from 64 to 124 digits. In the displayed available network segments, <code>::</code> represents a string of consecutive 0 bits. For example, fc00:: in the available network segment fc00::/32 represents fc00:0000:0000:0000:0000:0000:0000:0000. fc00::/32 indicates that the available network segments are fc00:0000:0000:0000:0000:0000:0000:0000-fc00:0000:ffff:ffff:ffff:ffff:ffff:ffff. 	fc00:0000:00ff: f: 0000:0000:0 000:0000:00 00/64
Gateway	Specifies the gateway address of the subnet.	fc00:0000:00ff: f: 0000:0000:0 000:0000:00 01
DNS Server Address 1	Specifies the IP address of an associated DNS server. This parameter is optional.	-
DNS Server Address 2	<p>Specifies the IP address of an associated DNS server. This parameter is optional.</p> <p>NOTE</p> <p>If the DNS server address is left blank, the subnet is not associated with any DNS server.</p> <p>When using only one DNS server address, enter it into DNS Server Address 1.</p> <p>To add a DNS server address, click Add DNS Server Address.</p>	-

Step 9 Determine whether to enable NAT.

- If you will create a subnet that does not need to communicate with the extranet, disable the NAT function, and proceed to the next step.
- If you will create a subnet that needs to communicate with the extranet, enable the NAT function, and create and configure subnet 2 by making reference to [Table 16-14](#).

Table 16-14 NAT settings and subnet 2

Configuration Item	Description
Name	Specifies the name of the NAT gateway. The value contains 1 to 64 characters including only letters, digits, underscores (_), and hyphens (-).
External Network	Select the external network for the EIP. NOTE If no external networks are available, contact the administrator to configure external networks.
EIP	Select an EIP. This is the EIP of the NAT gateway bound to subnet 2. NOTE If no EIPs are available in the drop-down list, click View EIP and apply for a new EIP on the Elastic IP page.
Subnet 2 Name Subnet 2 Segment Gateway DNS Server Address	For details about how to configure subnet 2, see 8 .
Specifications	Specifies the number of connections supported by the NAT gateway. Table 16-15 shows the type details. Select an appropriate NAT gateway type based on your network planning.

Table 16-15 NAT gateway types in Region Type I

Type	Maximum Number of SNAT Connections	Number of New SNAT Connections per Second
Small	10,000	1,000
Medium	50,000	5,000
Large	200,000	10,000
Extra-large	1,000,000	30,000

Step 10 Check the configuration of the new VPC.

- Click **Apply Now** to apply for the VPC.
- Alternatively, click **Add to Cart** and submit the application later.

Create security groups

Step 11 In the navigation pane under **Network Console**, click **Security Group**.

Step 12 On the **Security Group** page, click **Create Security Group**.

Step 13 In the **Create Security Group** dialog box, set the parameters as prompted.

Figure 16-11 Creating a security group

Table 16-16 Parameter description

Parameter	Parameter Description	Example Value
Name	The security group name can be a maximum of 64 characters long, and can contain letters, digits, underscores (_), hyphens (-), and dots (.). The name cannot contain spaces. NOTE You can change the security group name after a security group is created. It is recommended that you use different names for different security groups.	sg-34d6
Description	The security group description can contain a maximum of 64 characters and cannot contain angle brackets (<) or (>).	-

Step 14 Click **OK**.

Step 15 On the **Security Group** page, click the target security group.

Step 16 Click the **Inbound** or **Outbound** tab.

Step 17 Add rules using **Add Rule** or **Fast Add Rule**.

- Click **Add Rule** and configure the parameters for adding a security group rule.

Table 16-17 Parameters for adding a security group rule

Parameter	Description	Example Value
Protocol	<p>Specifies the network protocol. The value can be ANY, TCP, UDP, or ICMP.</p> <ul style="list-style-type: none">- ANY means that this rule is effective for any protocol.- TCP: Transmission Control Protocol (TCP) is a transport layer protocol. It provides reliable data transmission and maintains a virtual connection between devices or services that communicate with each other.- UDP: indicates a transport layer protocol that is used to compress network data traffic into data packets.- ICMP: indicates a network layer protocol. It is used to transmit error report control messages, and the ping command is used for communication status check.	TCP
Port Range/ ICMP Type	<ul style="list-style-type: none">- When you select TCP or UDP for Protocol, this parameter is a port range. Its value ranges from 1 to 65535.- When you select ANY for Protocol, this parameter is unconfigurable.- When you select ICMP for Protocol, this parameter is the ICMP type.	22 or 22-30
Type	<p>If you have deployed the dual stack (IPv4 & IPv6) in the system, you need to select a type of the security group rule.</p> <ul style="list-style-type: none">- If this parameter is set to IPv4, the traffic on the IPv4 network segment is allowed to pass.- If this parameter is set to IPv6, traffic on the IPv6 network segment is allowed to pass.	IPv4

Parameter	Description	Example Value
Source/ Destination	<p>Specifies the source when you select the Inbound tab.</p> <p>Specifies the destination when you select the Outbound tab.</p> <p>The value can be an IP address range or a security group.</p> <ul style="list-style-type: none"> - If you specify an IP address range as the value, select an IP address type, and enter an IP address range. <ul style="list-style-type: none"> ▪ Select IPv4 for Type. For example: <ul style="list-style-type: none"> xxx.xxx.xxx.xxx/32 (IP address) xxx.xxx.xxx.0/24 (subnet) 0.0.0.0/0 (any IP address) ▪ Select IPv6 for Type. For example: <ul style="list-style-type: none"> xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/64 (subnet) 0:0:0:0:0:0:0:0/0 (any address) - If you specify a security group as the value, choose a source or destination security group from the drop-down list. 	0.0.0.0/0 default
Description	<p>Provides supplementary information about the rule. This parameter is optional.</p> <p>The description can contain a maximum of 64 characters and cannot contain angle brackets (<) or (>).</p>	-
Operation	<ul style="list-style-type: none"> - Clone: specifies to copy an existing security group rule. - Delete: specifies to delete an existing security group rule. 	-
Adding a tag	<p>Adds multiple rules to a security group. A maximum of 10 rules can be added at a time.</p>	-

- Click **Fast Add Rule** to add rules with the preset common port.

Table 16-18 Parameters for adding a rule quickly

Parameter	Description	Example Value
Type	<p>If you have deployed the dual stack (IPv4 & IPv6) in the system, you need to select a type of the security group rule.</p> <ul style="list-style-type: none"> - If this parameter is set to IPv4, the traffic on the IPv4 network segment is allowed to pass. - If this parameter is set to IPv6, traffic on the IPv6 network segment is allowed to pass. 	IPv4
Port	<p>Multiple common ports are preset for users. Users can select common ports based on network requirements.</p> <p>If the preset ports fail to meet the requirements, you can add a user-defined port to the TCP or UDP protocol.</p> <p>NOTE If you select multiple ports, the system adds multiple rules at a time.</p>	FTP(20-21)
Source/ Destination	<p>Specifies the source when you select the Inbound tab.</p> <p>Specifies the destination when you select the Outbound tab.</p> <p>The value can be an IP address range or a security group.</p> <ul style="list-style-type: none"> - If you specify an IP address range as the value, select an IP address type, and enter an IP address range. - If you specify a security group as the value, choose a source or destination security group from the drop-down list. 	0.0.0.0/0 default

 **NOTE**

Source and **Destination** can be set to **Security Group** or **IP Address Range**. The details are as follows:

- **IP Address Range:** This rule takes effect for the specified IP address range. **0.0.0.0/0** and **0:0:0:0:0:0:0:0/0** indicate that this rule takes effect for all IP addresses.
- **Security Group:** This rule takes effect for all cloud servers in the selected security group.

Step 18 Click **OK**.


Apply for ECSs

Step 19 Log in to ManageOne as a VDC operator using a browser.

- URL in non-B2B scenarios: **https://Address for accessing ManageOne Operation Portal**, for example, **https://console.demo.com**

URL in B2B scenarios: **https://Address for accessing ManageOne Tenant Portal**, for example, **https://tenant.demo.com**

- Username and password: username and password of the VDC operator.

Step 20 Click  in the upper left corner, select a region, and choose **Computing > Elastic Cloud Server**.

The **Elastic Cloud Server** page is displayed.

Step 21 Select a resource set from the menu bar on the top of the page.

Step 22 Click **Apply for ECS**.

The **Select Service** page is displayed.

Step 23 Select a service and click **Apply Now**.

The **Apply for ECS** page is displayed.

Step 24 Complete basic configurations for the ECS.

 **NOTE**

- Customizable settings vary depending on the product you select. The ECS you selected in **Step 23** determines whether **AZ**, **ECS Type**, **vCPUs**, **Memory**, **Image Type**, and **Image** can be customized. During the configuration, you can skip the parameters that cannot be customized.
- The screenshot is only an example. If the actual environment is different from the screenshot, use the actual environment.

Table 16-19 Parameter description

Parameter	Description	Example Value
AZ	Specifies a physical region where resources use independent power supplies and networks. AZs are physically isolated but interconnected through an internal network. To enhance application availability, create ECSs in different AZs.	kvm_az

Parameter	Description	Example Value
Creation Method	<p>Select a creation method for an ECS.</p> <ul style="list-style-type: none"> • New: Customize parameters to create an ECS. • Create from Template: Create an ECS using an ECS image or existing ECS backup as a template. To create an ECS using ECS backup, the CSBS service must be deployed on the platform. The ECS image and ECS backup displayed on the page are in the current region and support cross-AZ deployment, but do not support cross-region deployment. A full-ECS image and ECS backup that have expired or been deleted cannot be used to create an ECS. <p>When an ECS is requested, the OS and the boot mode of the OS cannot be changed. You can click Upgrade Flavor to modify the ECS type and flavor. The type and capacity of the system disk or data disk can be changed, but the capacity cannot be less than that of the source ECS disk. Other parameters can be customized as required.</p>	New
ECS Type	<p>The platform provides various ECSs for you to select based on application scenarios.</p> <p>The ECS type is determined by the ECS type tag selected during flavor creation. For details, visit Operation Help Center and choose Operation > Compute Services > Elastic Cloud Server (ECS) > Configuration Before ECS Creation > Creating a Flavor.</p>	General-purpose

Parameter	Description	Example Value
Boot Mode	<p>Specifies the ECS boot mode, which can be BIOS or UEFI.</p> <p>Basic Input/Output System (BIOS) is used to load the basic computer code to initialize hardware, check hardware functions, and boot the OS.</p> <p>Unified Extensible Firmware Interface (UEFI) does not need a long self-check as BIOS does, simplifying hardware initialization and OS boot. In addition, UEFI is easy to use because it supports graphical user interfaces (GUIs), various operation modes, and hardware driver insertion.</p> <p>NOTE</p> <p>This parameter is available only when the following requirements are met. Otherwise, this parameter is not available.</p> <ul style="list-style-type: none"> • The virtualization type of the selected AZ is KVM. • After you select vCPUs and memory (MB), the system filters image files based on the selected memory size. An image will be displayed here only if its Min Memory (MB) specified during image registration is smaller than the memory size of the selected flavor. <p>This parameter is available only if 1) at least one of the displayed image files is configured to use the UEFI boot mode; 2) UEFI boot is selected during image registration on Service OM. Otherwise, this parameter is unavailable, indicating that all image files use BIOS as the default boot mode.</p> <ul style="list-style-type: none"> • In Arm scenarios, the ECS boot mode can only be UEFI and cannot be changed. 	BIOS
Image Type	<ul style="list-style-type: none"> • 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. • Private Image Image available only to the user who created it using an existing ECS or external image file. It contains an OS, pre-installed public applications, and your private applications. Using a private image to create ECSs removes the need to configure multiple ECSs repeatedly. • Shared Image A shared image is a private image shared by another user. 	Public Image

Parameter	Description	Example Value
Image	<ul style="list-style-type: none"> ● Windows Used for development platforms or production workloads that run on Windows. It is recommended that the memory capacity be at least 1 GB. ECSs created using a Windows image support the installation of Internet Information Services (IIS) and SQL servers. ● Linux Used for development platforms or production workloads that run on Linux. <p>NOTE Select a 64-bit OS if the required memory capacity is 4 GB or larger. This is because 32-bit OSs allow addressing only within a 4 GB memory range.</p> <p>During ECS creation, the system filters image files based on the selected flavor.</p> <ul style="list-style-type: none"> ● For a flavor whose Boot Device is set to Cloud Disk, an image is displayed here only when its Min Memory (MB) specified during image registration is less than or equal to the selected memory. ● For a flavor whose Boot Device is set to Local Disk, an image is displayed here only when its Min Memory (MB) specified during image registration and the minimum disk required by the image are less than or equal to the Memory and Root Disk (GB) of the selected flavor, respectively. <p>If you select BIOS or UEFI as the boot mode, the image files that use this boot mode will be displayed. If the boot mode parameter is not displayed, it indicates that all the image files use BIOS as the default boot mode.</p> <p>NOTE If you select vGPU-accelerated for ECS Type and the driver is not installed in the image, install the GRID driver by referring to B Installing a GRID Driver on a vGPU-accelerated ECS after the ECS is created.</p>	CentOS

Parameter	Description	Example Value
Joint Windows Domain	<p>This parameter is available if the virtualization type of the selected AZ is KVM, the ECS is running a Windows OS, and the service selected in Step 23 has been configured with domain information. If the selected image is a static injection image or does not have Cloudbase-Init installed, it cannot be added to a domain.</p> <p>The administrator can perform unified authentication for ECSs added to the same domain. The following functions will be available for ECSs added to a domain: manage compute resources, reduce network management complexity and costs, enhance security, and support account roaming and folder redirection. Resources can be shared among ECSs in the same domain. For more information about domain servers and their functions, click here.</p> <p>Specify whether to add an ECS to a Windows domain. You can select a domain from the drop-down list. Available options are those defined by the administrator during product creation.</p>	-

Parameter	Description	Example Value
Same Storage	<p>This parameter is available only when Boot Device of the ECS flavor is set to Cloud Disk.</p> <ul style="list-style-type: none"> • If this parameter is set to Yes, the created ECS will support backup, DR, and ECS snapshot. • If this parameter is set to No, the created ECS will not support backup or DR. Whether it will support ECS snapshot depends on whether all disks of the ECS reside in the same storage backend. <p>If this parameter is set to Yes, the system selects the same storage backend configured with a storage tag to create the system and data disks of the ECS. The ECS can be provisioned successfully only when a storage backend that meets the requirements above is available in the environment. For details about how to configure the storage tag, visit Operation Help Center > Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > Disk FAQs > (Optional) Creating a Disk Type.</p> <p>After the ECS is provisioned, you can change the Same Storage setting if certain conditions are met. For details, see 7.4 Modifying the DR or Backup Function of an ECS.</p>	Yes
System Disk	<p>Select a disk type and set the disk size. To ensure that the ECS runs properly, the minimum allowed capacity of the system disk is related to the selected image file.</p> <p>If the system disk capacity cannot be changed on the web page, log in to DMK and change the value of is_supported_modify_sys_disk_size to true. For details, visit Operation Help Center and choose Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > Modifying Configuration Items on DMK.</p>	10GB

Parameter	Description	Example Value
Data Disk	<p>This parameter is displayed after you click Add Data Disk.</p> <p>Select a disk type and set the disk size. You can create multiple data disks for an ECS.</p> <p>NOTE</p> <ul style="list-style-type: none">• If Shared Disk is selected, the data disk can be attached to multiple ECSs. A shared disk can be attached to a maximum of 16 ECSs.• If you select Yes for Same Storage, ensure that the system disk and data disks of the ECS reside in the same storage backend, and the storage backend has the storage tag configured. Otherwise, the ECS cannot be provisioned.• If you select SCSI, transparent SCSI command transmission is supported. Therefore, when the VM HA function is used, lock protection is not supported, and the disk may be dual written. If SCSI is not displayed, log in to DMK and set the value of is_supported_volume_device_type in the ECS_UI configuration file to true. For details about how to change the value, visit Operation Help Center and choose Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > Modifying Configuration Items on DMK.	100GB
Quantity	<p>Set the number of ECSs to be created.</p> <p>NOTE</p> <p>By default, a maximum of 30 ECSs can be requested at a time.</p>	1

Step 25 Click **Next**. The **Network Configuration** page is displayed.

Step 26 Complete network configurations for the ECS.

Table 16-20 Parameter description

Parameter	Description	Example Value
Resource Set	<p>Select the current resource set or another resource set from the drop-down list. You can view the current resource set in the navigation bar at the top. Assume that the current resource set is Resource Set A and another resource set available is Resource Set B.</p> <ul style="list-style-type: none"> • When you select the current resource set, VPCs available will be those in Resource Set A. • If you select Resource Set B, VPCs available will be those in Resource Set B. By selecting Resource Set B, you create ECSs in Resource Set A by using the network resources of Resource Set B. With other configurations including security groups, you enable these ECSs to communicate with all those in the VPCs that belong to Resource Set B, allowing ECSs of different projects to share the same VPCs. <p>NOTE</p> <ul style="list-style-type: none"> • This parameter is available when VPC sharing is enabled on Service OM and the shared VPC permission is configured for the resource set on ManageOne. Otherwise, this parameter is not displayed. By default, this function is disabled. For details, visit Operation Help Center and choose Operation > Network Services > Virtual Private Cloud (VPC) > Shared VPC Best Practices. • If you select Resource Set B, security groups available will be those in Resource Set A, but the EIPs will be those that belong to Resource Set B. 	project_02
VPC	<p>Provides network functions for ECSs. The network functions include the subnet and security group.</p> <p>You can select an existing VPC, or click Create VPC to create one.</p>	-

Parameter	Description	Example Value
NIC	<p>Includes primary and extension NICs. You can add a maximum of 15 extension NICs to an ECS.</p> <ul style="list-style-type: none">• If you select VPC Subnet, all subnets in the VPC are available for you to choose from. In this case, the NIC supports layer 3 communication, allowing the ECS to communicate with networks (for example, the public network or other VPCs) beyond the VPC.• If you select Intra-Project Subnet, all project-level subnets in the project are available for you to choose from. All NICs configured with the same subnet can communicate with each other at layer 2 on the project level. Layer 2 communication is supported within the same VPC and between different VPCs. <p>The Primary NIC network type must be VPC Subnet. Otherwise, you cannot access the Internet through the allocated EIP, and you cannot access the Object Storage Service (OBS) or a security service. You can select the network type for an Extension NIC as required.</p> <p>You can choose to use the automatically assigned IP address or manually assign one. If you choose to manually assign IP addresses when creating ECSs in batches, you can set the value of Specify IP Address By to IP Address Range or Single IP Address. Then, click OK.</p> <p>NOTE</p> <ul style="list-style-type: none">• If the selected subnet has only an IPv4 address segment, the NIC will only have one IPv4 address. If the selected subnet has both IPv4 and IPv6 address segments, the NIC will have one IPv4 address and one IPv6 address.• If the selected subnet does not have DHCP enabled and the selected image does not support static IP address injection, after an ECS is created, you need to manually configure IP addresses for the ECS. Otherwise, the NIC cannot be reached. For details, see 18.6.1 Configuring a Static IP Address for an ECS.	VPC Subnet subnet-c869(192.168.0.0/24)

Parameter	Description	Example Value
Security Group	<p>A security group implements access control for ECSs within the security group, enhancing security protection on ECSs. This enhances ECS security.</p> <p>When creating an ECS, you can select multiple security groups. Multiple security groups may affect the ECS network performance. You are advised to select a maximum of five security groups. In such a case, the access rules of all the selected security groups apply to the ECS.</p> <p>NOTE</p> <p>If the image has Cloud-Init or Cloudbase-Init installed, you need to initialize the ECS after it is created. Before initializing an ECS, ensure that the security group rule in the outbound direction meets the following requirements:</p> <ul style="list-style-type: none"> • Protocol: TCP • Port Range: 80 • Remote End: 169.254.0.0/16 <p>If you use the default security group rules in the outbound direction, the preceding requirements are already met, and this parameter does not need to be set.</p> <p>If you want to be able to log in to an ECS in SSH mode, you need to configure the inbound rules of the security group to allow your local computer to access the ECS. For details, see 12.4 Configuring Security Group Rules.</p>	-

Parameter	Description	Example Value
EIP	<p>A static public IP address bound to an ECS in a VPC. Using the EIP, the ECS provides services externally.</p> <p>The following options are provided:</p> <ul style="list-style-type: none"> • Do Not Use Without an EIP, the ECS cannot access the Internet and is used only in the private network or cluster. • Automatically Assign: The system automatically assigns an EIP for the ECS. In the Region Type I scenario, you also need to select External Network, Subnet, Bandwidth Type, and Bandwidth to which the EIP belongs. <ul style="list-style-type: none"> – If Bandwidth Type is set to Dedicated Bandwidth, you need to configure the bandwidth of the EIP. The EIP occupies bandwidth exclusively. Select this mode when you desire stable and large bandwidth. – If Bandwidth Type is set to Shared Bandwidth, you need to set Bandwidth Name for the shared bandwidth. The EIP shares the bandwidth with other EIPs that are added to the shared bandwidth. If the shared bandwidth contains three EIPs and the peak bandwidth is 10 MB/s, the total traffic of the three EIPs cannot exceed 10 MB/s. Select this mode when your application does not have a high bandwidth requirement and there is a bandwidth cap. <p>NOTE If Automatically Assign is selected, ensure that the EIP quota is sufficient. Otherwise, ECS provisioning will fail.</p> <ul style="list-style-type: none"> • Specify: An existing EIP is assigned for the ECS. When using an existing EIP, you cannot create ECSs in batches. <p>NOTE If you are not authorized to use EIP, contact the administrator to change your permissions.</p>	Do Not Use

Step 27 Click **Next**. The **Advanced Settings** page is displayed.

1. Set the ECS name.

When you create ECSs in batches, the system automatically adds an incremental number to the end of the custom ECS name, for example,

ecs-0001, ecs-0002, and so on. The value ranges from 0001 to 9999 by default.

To start with a specific number, click **Change Suffix Start Number** to customize the value. For example, if you set the value to 1126, the ECS names will be xxx-1126, xxx-1127, and so on.

 **NOTE**

If you want to create a Windows ECS that needs to be added to a domain, or if you require that the host name of the ECS (that is, the computer name shown in the ECS OS) must be unique, set the ECS name by following the instructions provided in **Operation Help Center > Operation > Compute Services > Elastic Cloud Server (ECS) > ECS Host Name > Rules for Configuring ECS Names (Unique Host Names)**.

2. Set the host name prefix of the ECS.

The host name prefix of the ECS and a suffix of 5 random characters (0-9 and a-z) form the ECS host name, that is, the computer name shown in the OS. The value is in the format "Host Name Prefix-5 random characters".

- This parameter needs to be configured if this parameter is to be displayed. The system generates the host name prefix according to the ECS name you configured in [27.1](#) and by automatically filtering out unallowed characters based on the naming rules displayed on the page. You can change this prefix. The generated ECS host name is unique in the region or resource set. For details, see visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > ECS Host Name > Enabling ECS Host Name Uniqueness**.
- Skip this parameter if it is not displayed. The system generates host names for ECSs based on the default rules. The host names of different ECSs may be the same. For details, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > ECS Host Name > Default Rules for Generating ECS Host Names**.

3. Set the running status of the ECS. This parameter is available if the virtualization type of the AZ where the ECS resides is KVM.

- **Stopped:** A newly obtained ECS stays in the **Stopped** state.

 **NOTE**

- If **ECS Initial Status** is set to **Stopped** and the host group where the ECS resides is configured with tags for releasing resources upon ECS shutdown, the ECS does not occupy any of the following resources: vCPU, memory, GPU, NPU, USB, and volume connections. For details, see step "Configure custom tags" in **Operation Help Center > Operation > Compute Services > Elastic Cloud Server (ECS) > Configuration Before ECS Creation > Creating a Host Group**.
- When Arm servers are used, NIC connections can be released for general computing-plus ECSs when they are shut down.
- Resources will not be released for disk-intensive or ultra-high I/O ECSs when they are shut down.
- **Running:** A newly obtained ECS stays in the **Running** state.

4. Set a key or new password. Select the image password as the ECS password or reset a new password or key for the ECS.

 **NOTE**

This parameter is displayed only when the following conditions are met:

- The image selected in [Step 24](#) must have the Cloud-Init (Linux OS) plugin installed, and Cloud-Init is selected during image registration.
- The image selected in [Step 24](#) can use the image password as the ECS password.
- The image password cannot be used as the ECS password in the Windows OS.
- **No:** Use the password set during image creation as the ECS password.
- **Yes:** Set a key or new password for the ECS.

5. Select an authentication mode for the ECS.

 **NOTE**

The image selected in [Step 24](#) must have the Cloud-Init (Linux OS) or Cloudbase-Init (Windows OS) plugin installed, and Cloud-Init is selected during image registration.

- Key pair

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

 **NOTE**

- If you use an existing key pair, make sure that you have saved the key file to a local directory. Otherwise, logging in to the ECS will fail.
- Windows ECSs support only the password authentication mode. If the login mode is set to **Key pair**, you must use the key file used during ECS creation to obtain the password of user **Administrator** or Cloudbase-Init account generated during ECS installation for subsequent logins. For details, see [6.3.1 Obtaining the Password for Logging In to a Windows ECS](#).
- For Windows, if the selected image supports static IP address injection, the key pair authentication is dimmed. Only the password mode can be selected.

- Password

A password is used for ECS login authentication. If the ECS runs a Linux OS, you can use username **root** and its password to log in to the ECS. If the ECS runs a Windows OS, you can use username **Administrator** and its initial password to log in to the ECS.

When using password login authentication on an ECS whose virtualization type is KVM, you can select **Customize user** and customize a username and password to create a user.

- On a Windows ECS, this is a common user without administrator rights. To perform operations that require administrator rights, switch to the administrator role first.
 - On a Linux ECS, this is a common user without administrator rights. You can use this user to log in to the ECS over SSH. If you need to run a script or system command after login, enter the password of the **root** user to upgrade rights. After this user is created, logging in to the ECS as the **root** user over SSH is disabled by default. To enable the function, see [Enabling Login Using a Password over SSH](#).
6. If you need to configure advanced settings for the ECS, select **Configure**. Otherwise, go to [Step 28](#).

- **Watchdog:** Enable or disable watchdog for an ECS.
The watchdog function provides a heartbeat mechanism used to monitor the health status of ECSs. When an ECS does not work properly, an alarm is generated, and the system attempts to restart the ECS. If restarting the ECS is successful, the alarm is cleared.

NOTICE

- Before enabling the watchdog function in x86 scenarios, ensure that the watchdog program that complies with the standard IPMI watchdog has been installed on the image selected in [Step 24](#). Otherwise, the ECS may restart repeatedly.
 - Before enabling the watchdog function in Arm scenarios, ensure that the watchdog program that complies with the standard 6300ESB watchdog has been installed on the image selected in [Step 24](#). Otherwise, the ECS may restart repeatedly.
-
- **Watchdog Alarm Policy:** In Arm scenarios, this parameter is available if **Watchdog** is set to **Enable**. If the 6300ESB watchdog does not detect watchdog information in the specified time, an alarm will be generated. The ECS will determine, based on this alarm policy, whether to get restarted.
 - **HA:** Enable or disable the HA function for an ECS.
To support HA, ECSs must meet the following requirements: the global HA function is enabled, the HA function of the host group where the ECS resides is enabled or not configured, and the HA function of the ECS is enabled. When HA is enabled, an ECS is automatically rebuilt on another host whenever the ECS or its host becomes faulty, ensuring service continuity.

NOTE

- For details about how to enable the global HA function, see "Product Management" > "Resource Pool" > "FusionSphere OpenStack" > "Compute" > "Configuring the VM HA Function" in *HUAWEI CLOUD Stack 8.1.0 O&M Guide*.
 - To check whether the HA function of the host group is enabled, log in to Service OM and check it in the **Custom Tag** area on the **Configuration** tab page of the host group details page. If a custom tag whose tag name is **_ha_enabled** and tag value is **False** exists, the HA function of the host group is disabled. If the tag does not exist or its value is **True**, the HA function of the host group is enabled.
You are advised not to enable this function for the management host group. Otherwise, services may be affected.
 - Resources need to be reserved for HA. Otherwise, the ECS HA function may fail. To ensure that the ECS HA functions properly, you need to clear alarms such as host exceptions and insufficient resources in a timely manner.
- **CD-ROM Drive**
For ECSs whose virtualization type is KVM, you can select **Use** or **Not use** for **CD-ROM Drive**.

- Select **Use** if you want to attach the local CD-ROM drive or an ISO file to the ECS.
- If the CD-ROM drive is used and UVP VMTools is installed in the image you have selected, UVP VMTools will be automatically upgraded after the ECS is provisioned. If the CD-ROM drive is not used, even if UVP VMTools is installed in the image you have selected, UVP VMTools cannot be automatically upgraded after ECS provisioning. You will need to manually upgrade UVP VMTools.

NOTE

UVP VMTools collects internal monitoring metrics of ECSs to monitor their running status and supports communication between ECSs and physical hosts. UVP VMTools also provides the following functions:

- Improves disk I/O performance and network I/O performance for Windows ECSs.
- Reports alarms when faults occur on Linux ECSs.

– ECS Group

An ECS group is a logical group with affinity, anti-affinity, weak affinity, or weak anti-affinity rules configured. ECSs added to an ECS group will be scheduled to the same or different hosts according to the affinity rules of the group. For details about how to create an ECS group, see [11.1 Creating an ECS Group](#).

NOTE

- If the policy of the ECS group is affinity or anti-affinity, the ECS will fail to be created when existing hosts or resources are insufficient to fulfill the affinity or anti-affinity rules of the ECS group.
- ECSs whose virtualization type is KVM can be added to an ECS group configured with any affinity rule. ECSs of other virtualization types cannot be added to ECS groups.

– Tag

Specifies the ECS tags. This parameter is optional and helps you identify and manage your ECSs.

Click **Add Tag**, and select an existing key and value from the drop-down list box. The tags come from the Tag Management Service. To add or modify a tag, ask the administrator to do it after choosing **Console > Mgmt & Deployment > Tag Management** from the top menu bar.

NOTE

- When multiple tags are added to an ECS, each tag must have a unique key.
- During ECS creation, the system also automatically generates a built-in tag. The tag key is identical to the VPC ID and is invisible on the UI.

For details about tag management, see [7.1.3 Adding and Managing ECS Tags](#).

– File Injection


This parameter is optional. Enables the ECS to automatically inject a script file or other files into a specified directory when you create the ECS.

For details about file injection, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > Configuration Before ECS Creation > Creating a File Injection Script**.

– User Data Injection

Enables the ECS to automatically inject user data when the ECS starts for the first time. This configuration is optional. After this function is enabled, the ECS automatically injects the user data upon its first startup. For details about user data injection, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > Configuration Before ECS Creation > Creating a User Data Injection Script**.

Step 28 Click **Next** to go to the **Confirmation** page.

1. Check whether all settings are correct. If you need to modify a configuration item, click the  icon next it.
2. Confirm **Required Duration**.

 **NOTE**

Specifies the required duration for an ECS. It begins from the time when the ECS was created. You can use it within the use duration. When this duration expires, the ECS status becomes **Expired**.

If the ECS is running properly before the expiration, the ECS will still run properly and the system will not be shut down. In this case, you can only **Extend** and **Delete** the ECS.

Step 29 Click **Add to Cart** or **Apply Now**.

- **Add to Cart:** Add the configured ECS to the shopping cart, and submit the order after you confirm all the resources you need, including network and storage resources.
- **Apply Now:** Submit the task.

 **NOTE**

- If the ECS you requested needs administrator approval, it will be provisioned after your request is approved. Otherwise, the ECS will be provisioned immediately.
- If you create an ECS with additional data disks, initialize the data disks after the ECS is created. For details about how to initialize the data disks, see [10.3 Initializing a Data Disk](#).
- If your ECS is assigned both IPv4 and IPv6 addresses and runs CentOS 7.5 or Ubuntu Server 18.04.1, or the network communication is abnormal after the application is successful, visit **Operation Help Center** and choose **Operation > Compute Services > Image Management Service (IMS) > FAQs > Configuring a VM to Dynamically Obtain IPv6 Addresses**.

Step 30 Check whether the APP ECS can ping the database ECS.

- If yes, perform the following steps.
- If no, contact technical support for assistance.

Apply for a load balancer


Step 31 Log in to ManageOne as a VDC administrator or VDC operator using a browser.

URL in non-B2B scenarios: **https://Domain name of ManageOne Operation Portal**, for example, **https://console.demo.com**.

URL in B2B scenarios: **https://Domain name of ManageOne Tenant Portal**, for example, **https://tenant.demo.com**.

You can log in using a password or a USB key.

- Login using a password: Enter the username and password.
The password is that of the VDC administrator or VDC operator.
- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.

Step 32 Click  in the upper left corner of the page, select a region, and choose **Network > Elastic Load Balance**.

Step 33 Click **Apply for ELB** to enter the **Select Service** page.

Step 34 Select a service, and click **Apply Now**.

Step 35 On the **Apply for ELB** page, set the parameters according to the displayed tips. For details, see [Table 16-21](#).

Figure 16-12 Apply for ELB

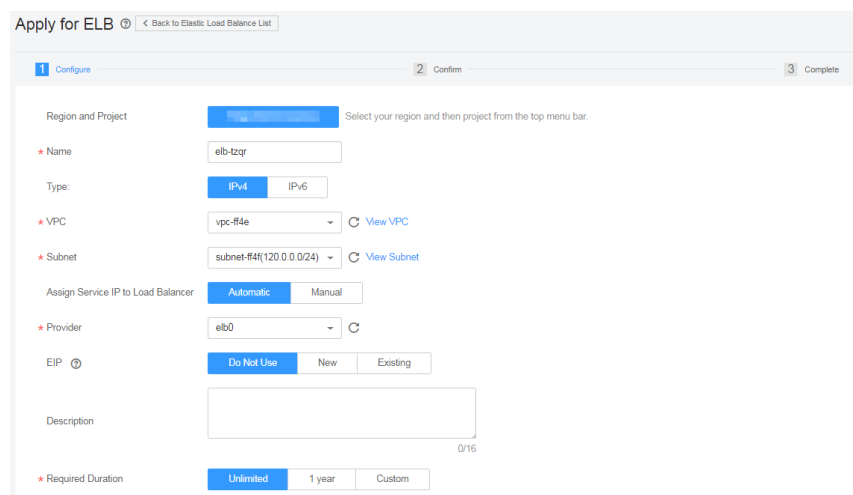


Table 16-21 Parameters for creating a load balancer

Parameter	Description	Example Value
Region and Project	The current region and project are displayed by default. To change them, use the selector in the upper left corner of the page.	cn-global-1 (vpc_project)
Name	Specifies the load balancer name. The parameter value consists of 1 to 64 characters, including letters, digits, and periods (.), underscores (_), or hyphens (-).	ulb-cjjw

Parameter		Description	Example Value
Type		<ul style="list-style-type: none"> IPv4: Load balancers support IPv4 addresses. If Type is set to IPv4, IPv4 addresses are used for load balancing. IPv6: Load balancers support IPv6 addresses. If Type is set to IPv6, IPv6 addresses are used for load balancing. 	IPv4
VPC		<p>Specifies the VPC to which the ELB instance belongs.</p> <p>The VPC provides network for load balancing. It also determines the scope of backend cloud servers that can be associated with this load balancer.</p> <p>You can select an existing VPC or click View VPC to create a new one. If no VPC is available, click Apply for VPC to create a new one.</p> <p>For more information about VPCs, see Operation Help Center > Network > Virtual Private Cloud > User Guide.</p>	vpc-123
Subnet		Specifies the subnet that contains the ELB service IP address. The value range includes all subnets that are bound to the VPC where the load balancer resides.	subnet-5dbb
Assign Service IP of Load Balancer		<p>Specifies the IP address used for accessing the load balancer. Each ELB must have a service IP address.</p> <p>The service IP address of a load balancer can be set in either of the following modes:</p> <ul style="list-style-type: none"> Automatic: An IP address that belongs to the selected subnet is automatically assigned as the ELB service IP address. Manual: Enter an IP address that belongs to the selected subnet to use as the ELB service IP address. 	Automatic

Parameter	Description	Example Value
EIP	<p>To receive access requests from external networks, bind an EIP to the load balancer. If no EIPs are available, create one.</p> <p>The following options are available:</p> <ul style="list-style-type: none"> ● Do Not Use: The load balancer cannot receive requests from external networks. ● New: Create an EIP. <ul style="list-style-type: none"> – External Network: Select the external network accessible to the load balancer. If no external network is available, plan, configure, and allocate external networks by following the instructions provided in Operation Help Center > Operation > VDC Tenant Modeling > Dividing External Networks (HUAWEI CLOUD Stack Scenario). – Subnet: Select the subnet of the external network to which the EIP is connected. This determines the range of the EIP. – Assignment Mode: When you select Automatic, the system automatically assigns an EIP from the Allocation Pools and you do not need to manually assign an EIP. The value of Allocation Pools is specified by users when they create subnets for an external network on Service OM. When you select Manual, you need to manually assign an EIP. The IP address can be any IP address of the subnet for the external network. <p>NOTE If the value of Allocation Pools is left blank on Service OM, an IP address can be assigned from all subnet CIDR blocks for the external network whether you choose to automatically or manually assign an EIP.</p> <ul style="list-style-type: none"> – Elastic IP Address: When Automatic is selected, an IP address that belongs to the selected subnet is automatically assigned as the EIP; 	Do Not Use

Parameter		Description	Example Value
		<p>when Manual is selected, enter an IP address that belongs to the selected subnet to use as the EIP.</p> <ul style="list-style-type: none"> - Bandwidth Type (available only in Region Type I): specifies the bandwidth type of the EIP to be bound to the ELB. The value can be Dedicated Bandwidth or Shared Bandwidth. You can select a proper bandwidth type as required. - Bandwidth Name (available only in Region Type I): specifies the bandwidth used by the backend cloud servers to access the external network using the EIP. The value can contain only letters, digits, underscores (_), and hyphens (-). When Bandwidth Type is set to Shared Bandwidth, you must select Bandwidth Name. If no option is available in the drop-down list box, click Manage shared bandwidth to create a shared bandwidth. - Bandwidth Size (Mbit/s) specifies the data transmission capability. This parameter is available when Bandwidth Type is set to Dedicated Bandwidth. - Bandwidth Description: Provides supplementary information about bandwidth. This parameter is available when Bandwidth Type is set to Dedicated Bandwidth. <ul style="list-style-type: none"> ● Existing: Select an existing EIP from the drop-down list to create the load balancer. 	
Description		<p>Provides supplementary information about the load balancer.</p> <p>NOTE The value can contain 1 to 64 characters and cannot contain angle brackets (<>).</p>	N/A
Required Duration		<p>Specifies the validity period of a load balancer. Required Duration can be set to Unlimited, 1 year, or Custom. Select the one that best meets your requirements.</p>	Unlimited

Step 36 Confirm the ELB configuration.

- Click **Add to Cart** and submit the application later.
- Click **Apply Now**. The system automatically applies for an instance.

Step 37 Click **Back to ELB List**. If the instance status changes to **Running**, the ELB instance is successfully created.

Add a listener and backend ECS

Step 38 On the **Elastic Load Balance** page, click the name of the newly created load balancer. In the **Listener** area, click **Add Listener**.

Step 39 Set parameters as prompted and click **OK**.

Figure 16-13 Add Listener

×

Add Listener

* Name:

* Protocol/Frontend Port:
The value ranges from 1 to 65535.

Connection limit:

* LB Algorithm:

Sticky Session Type:

* Cookie Name:
Ensure that the configured cookie name is the same as that configured for the applications running on the backend ECSs.

Description:
0/64

Health Check Configuration

* Health Check Mode:

* Interval (s):

* Timeout (s):

* Check Path:
The value can contain 1 to 80 characters.

* Maximum Retries: The value ranges from 1 to 10.

HTTP Method:

HTTP Status Code:

Step 40 In the **Backend ECS Group** area, click **Add Backend ECS Group** in the row where the target backend ECS group is located.

Step 41 Set parameters as prompted and click **OK**.

Bind an EIP.

Step 42 In the navigation pane on the left, choose **Elastic Load Balance**.

Step 43 Click the target load balancer name to go to the details page.

Step 44 Click **Bind** next to **EIP** and choose the EIP you want to bind to the load balancer from the drop-down list.

If no EIP is available in the drop-down list, click **View EIP** to apply for an EIP and bind it to the load balancer.

Step 45 Click .

 **NOTE**

To unbind an EIP, click **Unbind**. In the displayed dialog box, click **OK**.

----End

16.6 Creating and Attaching a Data Disk to a Database Host

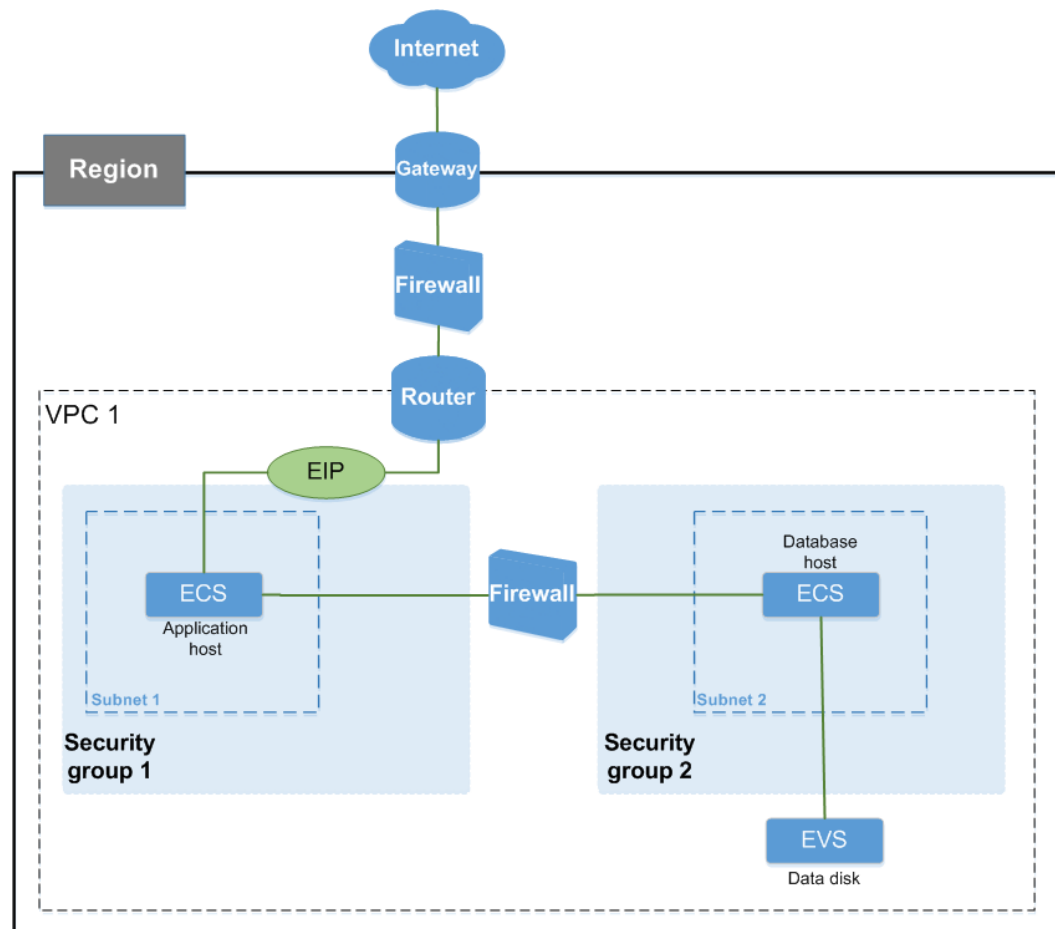
Scenario

A customer has applied for multiple Elastic Cloud Servers (ECSs) to provide a public network application. [Figure 16-14](#) shows the networking.

- An ECS is used to run the application service. External networks access the application service using public IP addresses.
- An ECS is used to run the database. The application service and the database are deployed on different cloud servers. Security groups are configured on the database side to allow traffic from database ports of security groups on the application side.

The customer has created the database host and system disk. A data disk needs to be added to the database host to store critical data in the database. In addition, an easy-to-use data protection mechanism is required for data in the data disk to ensure the security of critical data.

Figure 16-14 User service networking



Requirements

The analysis based on the user requirements is as follows:

1. Elastic Volume Service (EVS) is a virtual block storage service that provides block storage for ECSs. You can create EVS disks on the console and attach them to ECSs. The method for using EVS disks is the same as that for using hard disks on physical servers. With the EVS service, you can create and attach data disks to database hosts, meeting your storage requirements.
2. EVS disk snapshot is an important data recovery method that records the status of EVS disk data at a specific point in time. The snapshot created for an EVS disk at a certain point in time is independent from the life cycle of the EVS disk. The snapshot can be used to roll back and restore data of the EVS disk at the time when the snapshot was taken. You can use snapshots to back up important service data on the service host on a routine basis, mitigating data loss risks caused by misoperations, attacks or viruses to meet user requirements on data security.

The preceding analysis shows that you can add data disks to database hosts and create snapshots for the data disks to protect the data.

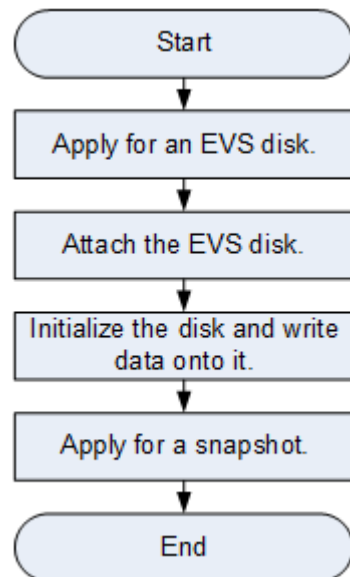
Configuration Plan

On the installation and deployment phase, backend storage was configured and the disk type for backend storage was created. The current task is to create a data disk for the database host, and then verify the availability. The configuration plan is as follows:

1. When the ECS was created for the database host, only a system disk was created, so a data disk needs to be created now.
2. A newly created data disk must be attached to a database host to work.
3. The system disk is created together with the database host and is automatically initialized by the system. Therefore, you need to perform initialization operations, such as partitioning and formatting, on only the attached data disk. After initialization, you can write test data into the data disk to check whether the data disk works properly.
4. Create a snapshot of the data disk of the database host for data backup to deal with data loss or inconsistency caused by misoperations, viruses, or hacker attacks.


Figure 16-15 shows the configuration flow based on the preceding analysis.

Figure 16-15 Configuration flow



Configuration Operations

Step 1 Apply for an EVS disk.

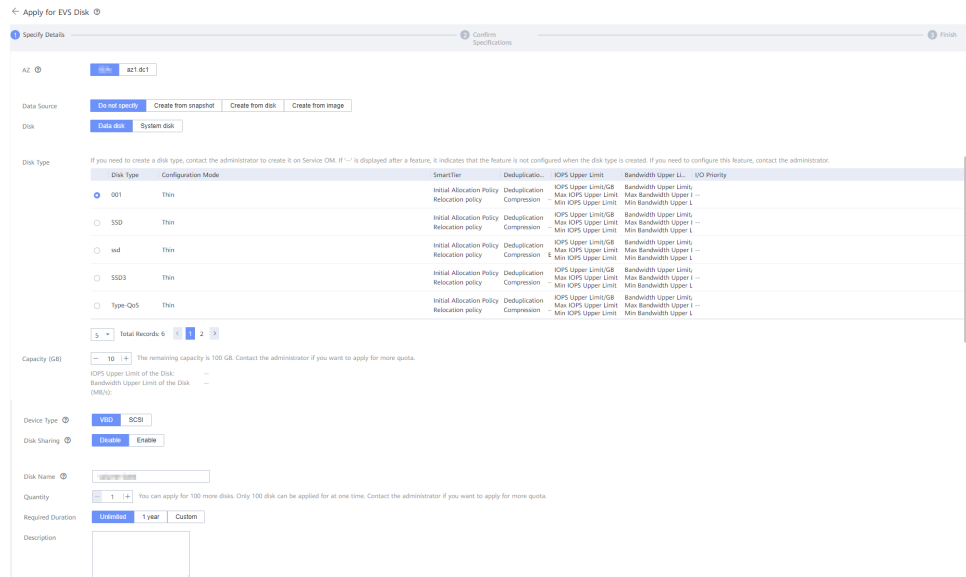
1. Log in to ManageOne as a VDC administrator or operator using a browser.
URL in non-B2B scenarios: **https://Address for accessing ManageOne Operation Portal**, for example, **https://console.demo.com**.
URL in B2B scenarios: **https://Address for accessing ManageOne Tenant Portal**, for example, **https://tenant.demo.com**.
2. Click  in the upper left corner and select a region. Choose **Storage > Elastic Volume Service**.

3. Select a project from the menu bar.
4. Click **Apply for EVS Disk**. In the **Select Service** dialog box, select the product to be applied for and click **Apply Now**.
5. Configure the data disk parameters.

[Table 16-22](#) lists the example data disk parameters.

Table 16-22 Example data disk parameters

Parameter	Description	Example value
AZ	This AZ must be the AZ where the database host resides.	az1.dc1
Data Source	If you select Do not specify , a blank EVS disk will be created, which contains no data.	Do not specify
Disk	Select the data disk to be applied for.	Data disk
Disk Type	Select a disk type that has been created on ManageOne Maintenance Portal.	-
Capacity	Determine the capacity of the EVS disk based on the service volume of the database host.	100GB
Device Type	It is recommended that you select SCSI . SCSI-type disks allow the ECS OS to directly access the underlying storage media and support advanced SCSI commands compared with VBD-type disks.	SCSI
Disk Sharing	In this example, the database host is deployed in the single-node mode, so select Disable .	Disable
Disk Name	It is recommended that you use the name of the database host as the prefix of the disk name to facilitate future management and maintenance.	<i>Database host name_volume_0001</i>
Quantity	Retain the default value 1 because only one data disk will be created.	1
Required Duration	Specifies the validity period of the EVS disk.	Not limited
Description	Provides additional information about the data disk to facilitate future management and maintenance.	Used for the database host



6. Click **Next**.
7. Confirm the application information, and click **Apply Now**.

Step 2 Attach the EVS disk.

1. Locate the row containing the EVS disk applied for in **Step 1**, and click **Attach** in the **Operation** column.
2. Select the database host to which the disk is to be attached.
If the state of the EVS disk is **In-use**, the EVS disk is successfully attached to the database host.

Step 3 Initialize and write data into the EVS disk.

NOTE

This section uses an ECS running CentOS 7.0 64-bit as an example, and uses the fdisk partition tool to set up partitions for the data disk. Initialization operations vary with operating systems.

1. Log in to the database host.
2. Run the following command to view information about the added data disk:

fdisk -l

Information similar to the following is displayed: (In the command output, the server contains two disks. **/dev/xvda** is the system disk, and **/dev/xvdb** is the added data disk.)

```
[root@ecs-b656 test]# fdisk -l
```

```
Disk /dev/xvda: 42.9 GB, 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: 0x000cc4ad
```

```
Device Boot      Start      End  Blocks  Id System
/dev/xvda1 *    2048    2050047  1024000  83 Linux
/dev/xvda2    2050048  22530047  10240000  83 Linux
```

```
/dev/xvda3 22530048 24578047 1024000 83 Linux
/dev/xvda4 24578048 83886079 29654016 5 Extended
/dev/xvda5 24580096 26628095 1024000 82 Linux swap / Solaris
```

```
Disk /dev/xvdb: 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
```

3. Run the following command to allocate partitions for the added data disk using fdisk:

```
fdisk Newly added data disk
```

In this example, **/dev/xvdb** is the newly added data disk.

```
fdisk /dev/xvdb
```

Information similar to the following is displayed.

```
[root@ecs-b656 test]# fdisk /dev/xvdb
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 0xb00005bd.
Command (m for help):
```

4. Enter **n** and press **Enter** to create a partition.

5. Enter **p** and press **Enter** to create a primary partition.

Information similar to the following is displayed. **Partition number** indicates the serial number of the primary partition. The value can be **1** to **4**.

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

6. Primary partition number **1** is used in this example. Enter **1** and press **Enter**.

Information similar to the following is displayed. **First sector** indicates the start cylinder number. The value can be **2048** to **20971519**, and the default value is **2048**.

```
Partition number (1-4, default 1): 1
First sector (2048-20971519, default 2048):
```

7. The default first sector **2048** is used in this example. Select **2048** and press **Enter**.

Information similar to the following is displayed. **Last sector** indicates the end cylinder number. The value can be **2048** to **20971519**, and the default value is **20971519**.

```
First sector (2048-20971519, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
```

8. The default last sector **20971519** is used in this example. Select **20971519** and press **Enter**.

Information similar to the following is displayed. A primary partition is created for a 10 GB data disk.

```
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
Using default value 20971519
Partition 1 of type Linux and of size 10 GiB is set
Command (m for help):
```

9. Enter **p** and press **Enter** to view the details about the created partition.

10. Enter **w** and press **Enter** to write the changes into the partition table.

Information similar to the following is displayed. The partition is successfully created.

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

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

11. Run the following command to synchronize the new partition table to the data disk:

```
partprobe
```

12. Run the following command to set the format for the file system of the newly created partition:

```
mkfs -t File system format /dev/xvdb1
```

For example, run the following command to set the **ext4** file system for the **/dev/xvdb1** partition:

```
mkfs -t ext4 /dev/xvdb1
```

Information similar to the following is displayed.

```
[root@ecs-b656 test]# mkfs -t ext4 /dev/xvdb1
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
655360 inodes, 2621184 blocks
131059 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2151677952
80 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

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

NOTE

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

13. Run the following command to create a mounting directory:

```
mkdir Mounting directory
```

/mnt/sdc is used in this example.

```
mkdir /mnt/sdc
```

14. Run the following command to mount the new partition on the directory created in [Step 3.13](#):

```
mount /dev/xvdb1 Mounting directory
```

For example, run the following command to mount the newly created partition on **/mnt/sdc**:

```
mount /dev/xvdb1 /mnt/sdc
```

15. Run the following command to view the mounting result:

```
df -TH
```

Information similar to the following is displayed. The newly created **/dev/xvdb1** partition has been mounted on **/mnt/sdc**.

```
[root@ecs-b656 test]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda2      xfs       11G   7.4G 3.2G  71% /
devtmpfs        devtmpfs  4.1G   0   4.1G   0% /dev
tmpfs           tmpfs     4.1G  82k  4.1G   1% /dev/shm
tmpfs           tmpfs     4.1G  9.2M  4.1G   1% /run
tmpfs           tmpfs     4.1G   0   4.1G   0% /sys/fs/cgroup
/dev/xvda3      xfs       1.1G  39M  1.1G   4% /home
/dev/xvda1      xfs       1.1G 131M  915M  13% /boot
/dev/xvdb1      ext4      11G  38M  9.9G   1% /mnt/sdc
```

16. Write data into **/mnt/sdc** to verify that the data disk works properly.

Step 4 Apply for a snapshot.

 **NOTE**

Currently, snapshots have to be created manually. If you require that data protection be performed regularly on the EVS disk, it is recommended that you use Volume Backup Service (VBS).

1. Select the EVS disk applied in **Step 1** and choose **More > Apply for Snapshot** in the **Operation** column to apply for a snapshot.
2. Configure the snapshot name and description, and click **Next**.
3. Confirm the application information, and click **Apply Now**.

If the status of the snapshot is **Available**, the snapshot is successfully created.

----End

17 Website Construction Tutorial

17.1 Building a Discuz Website

17.1.1 Overview

HUAWEI CLOUD Stack provides a wide range of services. A flexible combination of these services enables you to conveniently and rapidly deploy, run, and maintain various applications on the cloud. This section describes how to build a Discuz website using a combination of these services.

17.1.2 Implementation Plan

Requirement Analysis

To build a website, you need to perform the following tasks:

1. Register a domain name.
2. Apply for and configure servers.
3. Create a website.
4. File the website.

Tasks 1 and 4 are performed on a third-party website. Therefore, this tutorial describes tasks 2 and 3.

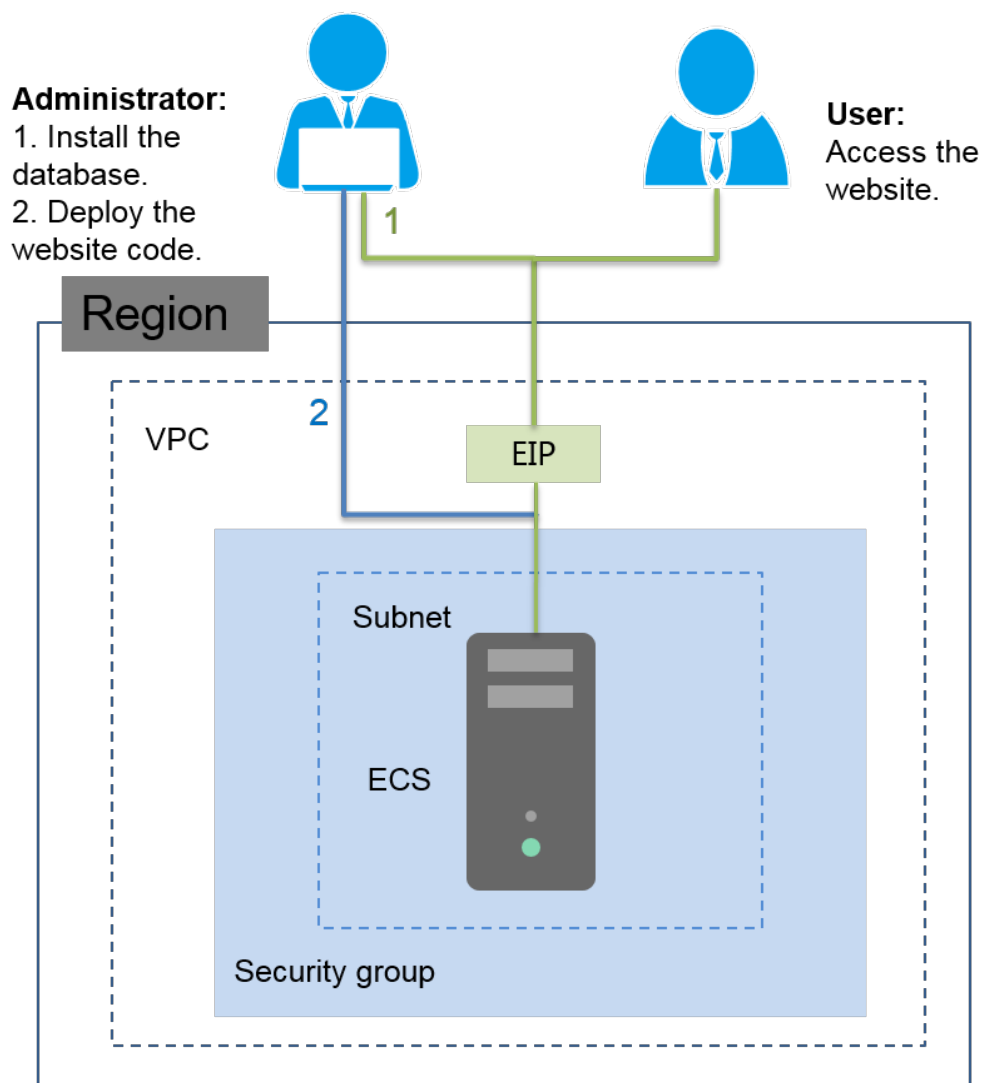
Usually, one Elastic Cloud Server (ECS) is sufficient to meet the service requirements of a small website. The ECS is used to host static and dynamic pages of the website and databases, handle user access, and generate page response. The ECS, Virtual Private Cloud (VPC), and Elastic Volume Service (EVS) services are required. If traffic increase occurs, you can use the Auto Scaling (AS) service so that ECSs are added at traffic peaks and removed at traffic lulls and use the Elastic Load Balance (ELB) service so that user access traffic is distributed to different ECSs. This tutorial mainly describes how to build a website using the ECS, VPC, and EVS services. [Table 17-1](#) lists the functions of each service and the operations related to each service.

Table 17-1 Service description

Service Name	Function	Related Operation
ECS	The ECS service is used to deploy databases, build websites, and store dynamic and static pages.	<ul style="list-style-type: none">• Creating an ECS• Logging in to an ECS
VPC	The VPC service provides a secure network environment for ECSs.	<ul style="list-style-type: none">• Creating a VPC• Creating a subnet• Creating a security group and adding a security group rule
EVS	The EVS service provides storage space for ECSs.	<ul style="list-style-type: none">• Creating an EVS disk• Attaching an EVS disk• Initializing EVS disks

Networking

Figure 17-1 Networking

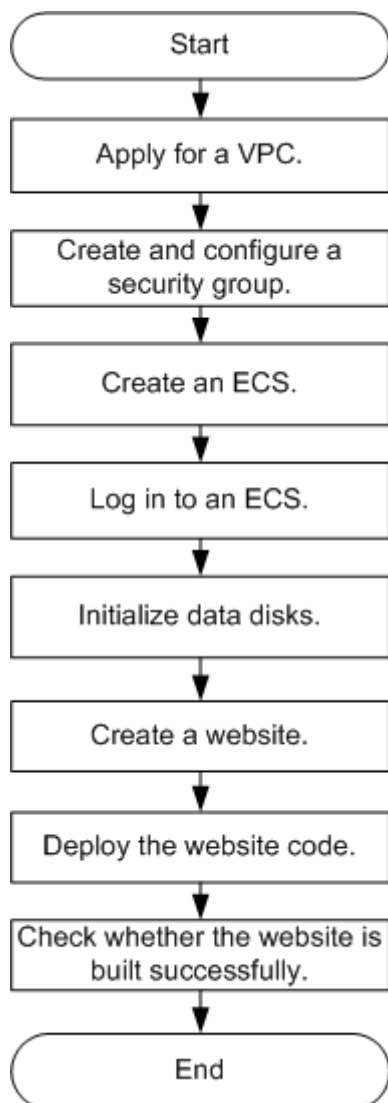


Configuration Plan

Before building a website, apply for an ECS. When applying for an ECS, you need to choose a VPC, subnet, security group, and EIP for the ECS. Therefore, before applying for an ECS, configure all these items, such as the network. After an ECS is created successfully, initialize data disks, install the database, deploy the website code, and check whether the website is built successfully.

The preceding analysis helps draw up the process for building a website, as shown in [Figure 17-2](#).

Figure 17-2 Process for building a website



17.1.3 Applying for and Configuring Services

17.1.3.1 Applying for a VPC

Context

This section describes how to quickly apply for a VPC and single-stack subnet.

Before performing the following operations, you need to correctly configure and allocate external networks as planned.

Procedure

Step 1 Log in to ManageOne as a VDC administrator or VDC operator using a browser.


URL in non-B2B scenarios: **https://Domain name of ManageOne Operation Portal**, for example, **https://console.demo.com**.

URL in B2B scenarios: **https://Domain name of ManageOne Tenant Portal**, for example, **https://tenant.demo.com**.

You can log in using a password or a USB key.

- Login using a password: Enter the username and password.
The password is that of the VDC administrator or VDC operator.
- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.



Step 2 Click  in the upper left corner, select a region, and choose **Network > Virtual Private Cloud**.

Step 3 Select a project in the region from the drop-down list in the upper left corner of the main menu.

Step 4 On the **Virtual Private Cloud** page, click **Apply for VPC**.

The **Select Service** dialog box is displayed.

Step 5 Select a service and click **Apply Now**.

The **Select VPC Configuration** page is displayed.

Step 6 Select **VPC with Internet Access (Optional)**, and click **Next**.

Step 7 Configure the VPC parameters by referring to [Table 17-2](#).

Table 17-2 VPC parameters

Parameter	Description	Example Value
Region and Project	The current region and project are displayed by default. To change them, use the selector in the upper left corner of the page.	az1.dc1(test)
Name	Specifies the VPC name. The name can contain letters, digits, underscores (_), and hyphens (-).	VPC-001
External Network	Select an availability zone from the All AZs drop-down list and select the external network in the VPC.	az0.dc0 net-01
CIDR Block	Specifies the Classless Inter-Domain Routing (CIDR) block for the VPC. The CIDR block of a subnet can be the same as the CIDR block for the VPC (for a single subnet in the VPC) or a subset (for multiple subnets in the VPC).	192.168.0.0/16
Required Duration	Specifies the required duration of a VPC.	1 year

Step 8 Configure subnet 1 by referring to [Table 17-3](#). Subnet 1 is isolated from the extranet, which ensures high security. Cloud servers in subnet 1 cannot directly communicate with the extranet.

Table 17-3 Parameters for configuring subnet 1

Parameter	Description	Example Value
Subnet 1 Name	Specifies the subnet 1 name. The name consists of only letters, digits, underscores (_), and hyphens (-).	Subnet-f03c
Type	If you have deployed the dual stack (IPv4 & IPv6) in the system, you need to select a network type. If only the IPv4 service is deployed, configure subnet 1 parameters by referring to Table 17-4 . <ul style="list-style-type: none">IPv4IPv4&IPv6 To create a single-stack subnet, select IPv4 and set parameters based on Table 17-4 .	IPv4

Table 17-4 Parameters for configuring an IPv4 subnet

Parameter	Description	Example Value
Subnet 1 Segment	Specifies the IP address range of subnet 1.	192.168.0.0/24
Gateway	Specifies the gateway address of subnet 1.	192.168.0.1
DNS Server Address 1	Specifies the IP address of an associated DNS server. This parameter is optional.	192.168.71.3
DNS Server Address 2	Specifies the IP address of an associated DNS server. This parameter is optional. NOTE If the DNS server address is left blank, the subnet is not associated with any DNS server. When using only one DNS server address, enter it into DNS Server Address 1 . To add a DNS server address, click Add DNS Server Address .	192.168.72.3

Parameter	Description	Example Value
Static Route Switch	<p>If this switch is set to OFF, static routes will not be configured for cloud servers in the subnet. If this switch is set to ON, the configured static routes will be injected to those cloud servers by using the DHCP function of the subnet.</p> <ul style="list-style-type: none"> • Destination: indicates the destination IP address range of the static route. • Next Hop: indicates the next-hop IP address of the static route. <p>NOTE When you need to add more static routes for cloud servers in a subnet, click Add Static Route. You can configure up to five static routes at a time.</p>	<ul style="list-style-type: none"> • Destination: 10.10.0.0/24 • Next Hop: 192.168.20.2

Step 9 Disable the NAT function.

Step 10 Check the configuration of the new VPC.

- Click **Apply Now** to apply for the VPC.
- Alternatively, click **Add to Cart** and submit the application later.

----End

17.1.3.2 Creating a Security Group and Configuring Security Group Rules

Creating a Security Group


Step 1 Log in to ManageOne as a VDC operator using a browser.

URL in non-B2B scenarios: <https://Domain name of ManageOne Operation Portal>, for example, <https://console.demo.com>.

URL in B2B scenarios: <https://Domain name of ManageOne Tenant Portal>, for example, <https://tenant.demo.com>.

You can log in using a password or a USB key.

- Login using a password: Enter the username and password.
The password is that of the VDC operator.
- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.

Step 2 Click  on the left of the main menu, select a region, and choose **Network > Virtual Private Cloud**.

Step 3 In the navigation pane on the left, click **Security Group**.

Step 4 On the **Security Group** page, click **Create Security Group**.

Step 5 In the **Create Security Group** area, set parameters as prompted.

Table 17-5 Parameter description

Parameter	Description	Example Value
Name	Specifies the security group name.	sg-001
Description	Provides supplementary information about the security group. This parameter is optional.	-

Step 6 Click **OK**.

----End

Configuring a Security Group Rule

Step 1 Click the name of the security group. The details page is displayed.

Step 2 On the Inbound tab, click **Add Rule**, and set the required parameters.

 **NOTE**

The application needs to allow access from external networks. Therefore, permit access in the inbound direction.

Table 17-6 Parameters for adding a security group rule

Parameter	Description	Example Value
Protocol	<p>Specifies the network protocol. The value can be ANY, TCP, UDP, or ICMP.</p> <ul style="list-style-type: none">• ANY means that this rule is effective for any protocol.• TCP: Transmission Control Protocol (TCP) is a transport layer protocol. It provides reliable data transmission and maintains a virtual connection between devices or services that communicate with each other.• UDP: indicates a transport layer protocol that is used to compress network data traffic into data packets.• ICMP: indicates a network layer protocol. It is used to transmit error report control messages, and the ping command is used for communication status check.	TCP

Parameter	Description	Example Value
Port Range/ ICMP Type	<ul style="list-style-type: none"> When you select TCP or UDP for Protocol, this parameter is a port range. Its value ranges from 1 to 65535. When you select ANY for Protocol, this parameter is unconfigurable. When you select ICMP for Protocol, this parameter is the ICMP type. 	22 or 22-30
Type	<p>If you have deployed the dual stack (IPv4 & IPv6) in the system, you need to select a type of the security group rule.</p> <ul style="list-style-type: none"> If this parameter is set to IPv4, the traffic on the IPv4 network segment is allowed to pass. If this parameter is set to IPv6, traffic on the IPv6 network segment is allowed to pass. 	IPv4
Source/ Destination	<p>Specifies the source when you select the Inbound tab.</p> <p>Specifies the destination when you select the Outbound tab.</p> <p>The value can be an IP address range or a security group.</p> <ul style="list-style-type: none"> If you specify an IP address range as the value, select an IP address type, and enter an IP address range. <ul style="list-style-type: none"> Select IPv4 for Type. For example: xxx.xxx.xxx.xxx/32 (IP address) xxx.xxx.xxx.0/24 (subnet) 0.0.0.0/0 (any IP address) Select IPv6 for Type. For example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/64 (subnet) 0:0:0:0:0:0:0:0/0 (any address) If you specify a security group as the value, choose a source or destination security group from the drop-down list. 	0.0.0.0/0 default
Description	<p>Provides supplementary information about the rule. This parameter is optional.</p> <p>The description can contain a maximum of 64 characters and cannot contain angle brackets (<) or (>).</p>	-

Parameter	Description	Example Value
Operation	<ul style="list-style-type: none"> • Clone: specifies to copy an existing security group rule. • Delete: specifies to delete an existing security group rule. 	-
Adding a tag	Adds multiple rules to a security group. A maximum of 10 rules can be added at a time.	-


Step 3 Click **OK**.

----End

17.1.3.3 Creating an ECS

Step 1 Log in to ManageOne as a VDC operator using a browser.

- URL in non-B2B scenarios: **https://Address for accessing ManageOne Operation Portal**, for example, **https://console.demo.com**
URL in B2B scenarios: **https://Address for accessing ManageOne Tenant Portal**, for example, **https://tenant.demo.com**
- Username and password: username and password of the VDC operator

Step 2 Click  in the upper left corner, select a region, and choose **Computing > Elastic Cloud Server**.

The **Elastic Cloud Server** page is displayed.

Step 3 Select a resource set from the menu bar on the top of the page.

Step 4 Click **Apply for ECS**.

The **Select Service** page is displayed.

Step 5 Select a service and click **Apply Now**.

The **Apply for ECS** page is displayed.

Step 6 Configure basic information about the ECS to be created. For details, see [Table 17-7](#).

NOTE

- When you select different services, the parameters to customize are different. The service you selected in [Step 5](#) determines whether **AZ**, **ECS Type**, **vCPUs**, **Memory**, **Image Type**, and **Image** can be customized. During the configuration, you can skip the parameters that cannot be customized.
- The screenshot is only an example. If the actual environment is different from the screenshot, use the actual environment.

Table 17-7 Parameter description

Parameter	Description	Example Value
AZ	A physical region where resources use independent power supplies and networks. AZs are physically isolated but interconnected through an internal network. To enhance application availability, create ECSs in different AZs.	kvm_az
Creation Method	Specifies the method for creating an ECS. <ul style="list-style-type: none"> • New: Customize parameters to create an ECS. • Create from Template: Create an ECS using a full-ECS image or ECS backup as a template. 	New
ECS Type	The platform provides various ECSs for you to select based on application scenarios.	General-purpose
Boot Mode	<ul style="list-style-type: none"> • Basic Input/Output System (BIOS) is used to load the basic computer code to initialize hardware, check hardware functions, and boot the OS. • Unified Extensible Firmware Interface (UEFI) does not need a long self-check as BIOS does, simplifying hardware initialization and OS boot. In addition, UEFI is easy to use because it supports graphical user interfaces (GUIs), various operation modes, and hardware driver insertion. <p>NOTE</p> <ul style="list-style-type: none"> • Skip this parameter if it is not displayed. • In ARM scenarios, the ECS boot mode can only be UEFI and cannot be changed. 	BIOS

Parameter	Description	Example Value
Image Type	<ul style="list-style-type: none">• 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. To select a public image, set Image Type to Public Image and select a desired one from the Image drop-down lists.• Private Image A private image is an image available only to the user who created it using an existing ECS or external image file. It contains an OS, preinstalled public applications, and the user's private applications. Using a private image to create ECSs removes the need to configure multiple ECSs repeatedly. To select a private image, set Image Type to Private Image and select a desired one from the Image drop-down list.• Shared Image A shared image is a private image shared by another user. To select a shared image, set Image Type to Shared Image and select a desired one from the Image drop-down list.	Public Image
Image	<ul style="list-style-type: none">• Windows Used for development platforms or operating services that run Windows. An authorized license is included in the image.• Linux Used for development platforms or operating services that run Linux.	CentOS
Same Storage	<p>If the new ECS needs to support backup or disaster recovery, select Yes. Otherwise, select No.</p> <p>If you select Yes, make sure that the system and data disks of the ECS reside in the same storage backend, and the storage backend is configured with the storage tag. Otherwise, the ECS cannot be provisioned.</p> <p>NOTE This parameter is available only when Boot Mode of the specified ECS flavor is set to Cloud Disk.</p>	No
System Disk	To ensure that the ECS runs properly, the minimum allowed capacity of the system disk is related to the selected image file.	16 GB

Parameter	Description	Example Value
Data Disk	This parameter is displayed after you click Add Data Disk . Select a disk type and set the disk size. You can create multiple data disks for an ECS.	40 GB
Quantity	Set the number of ECSs to be created.	1

Step 7 Click **Next**. The **Network Configuration** page is displayed.

Step 8 Configure network information about the ECS. For details, see [Table 17-8](#).

Table 17-8 Parameter description

Parameter	Description	Example Value
Resource Set	Select the current resource set or another resource set from the drop-down list. You can view the current resource set in the navigation bar at the top. You do not need to change the default resource set. NOTE This parameter is available when VPC sharing is enabled on Service OM and the shared VPC permission is configured for the resource set on ManageOne. Otherwise, this parameter is not displayed. By default, this function is disabled.	project_02
VPC	Provides a network, including subnet and security group, for an ECS.	-
NIC	Includes primary and extension NICs. <ul style="list-style-type: none"> If you select VPC Subnet, all subnets in the VPC are available for you to choose from. In this case, the NIC supports layer 3 communication, allowing the ECS to communicate with networks (for example, the public network or other VPCs) beyond the VPC. If you select Intra-Project Subnet, all project-level subnets in the project are available for you to choose from. All NICs configured with the same subnet can communicate with each other at layer 2 on the project level. Layer 2 communication is supported within the same VPC and between different VPCs. 	subnet-c869(192.168.0.0/24)

Parameter	Description	Example Value
Security Group	Controls ECS access within a security group or between security groups by defining access rules. This enhances ECS security.	-
EIP	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: <ul style="list-style-type: none">● Do Not Use: Without an EIP, the ECS cannot access the Internet and is used only in the private network or cluster.● Automatically Assign: The system automatically assigns an EIP for the ECS. The EIP provides exclusive bandwidth.● Specify: An existing EIP is assigned for the ECS. When using an existing EIP, you cannot create ECSs in batches.	Automatically Assign

Step 9 Click **Next**. The **Advanced Settings** page is displayed.

Step 10 Set the ECS name.

When you create ECSs in batches, the system automatically adds an incremental number to the end of each ECS name.

Step 11 Set the host name prefix of the ECS.

If this parameter is displayed, set it. The host name prefix and a suffix of 5 random characters (0-9 and a-z) form the ECS host name, that is, the computer name shown in the OS. It is in the format "Host Name Prefix-5 random characters".

Step 12 Set the power status of the ECS to **Running**.

- **Stopped:** A newly obtained ECS stays in the **Stopped** state.
- **Running:** A newly obtained ECS stays in the **Running** state.

Step 13 To add description for an ECS, such as the purpose of the ECS, enter the required information in the description text box.


Step 14 If **Set Key or New Password** is displayed, click **Yes**. You can customize the password or key pair for logging in to the ECS.

Step 15 Select **Password** for the login mode.

 **NOTE**

This password is used to log in to the ECS. Keep it secure.

Step 16 Retain the default values for other parameters and click **Next**. The **Confirmation** page is displayed.

1. Check whether all configuration items are correct. If you need to modify a configuration item, click  next to the corresponding module.
2. Confirm **Required Duration**.

Step 17 Click **Add to Cart** or **Apply Now**.

- **Add to Cart:** Add the configured ECS to the shopping cart, and submit the order after you confirm all the resources you need, including network and storage resources.
- **Apply Now:** Submit the task.

 **NOTE**


- If the ECS you requested needs administrator approval, it will be provisioned after your request is approved. Otherwise, the ECS will be provisioned immediately.
- If you create an ECS with additional data disks, initialize the data disks after the ECS is created.

----End

17.1.3.4 Logging In to an ECS

Step 1 Log in to ManageOne as a VDC operator using a browser.

- URL in non-B2B scenarios: **https://Address for accessing ManageOne Operation Portal**, for example, **https://console.demo.com**
URL in B2B scenarios: **https://Address for accessing ManageOne Tenant Portal**, for example, **https://tenant.demo.com**
- Username and password: username and password of the VDC operator

Step 2 Click  in the upper left corner, select a region, and choose **Computing > Elastic Cloud Server**.

The **Elastic Cloud Server** page is displayed.

Step 3 Select a resource set from the menu bar on the top of the page.

Step 4 In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID, and click the search button to search for the ECS.

Step 5 Locate the row containing the ECS and click **Remote Login** in the **Operation** column.

The **Configure Keyboard Layout for Remote Login** is displayed.

Step 6 Select the English keyboard and click **Remote Login**.

Step 7 (Optional) If the system displays "Press CTRL+ALT+DELETE to log on", click **Send CtrlAltDel** in the upper right corner of the remote login page to log in to the ECS.

Figure 17-3 Send CtrlAltDel



Step 8 Enter the password set in [16.1.3.3 Creating an ECS](#) and log in to the ECS.

----End

17.1.3.5 Initializing a Linux Data Disk (fdisk)

A data disk attached to an ECS or created together with an ECS must be initialized before it can become available. This section uses an instance running CentOS 7.0 64bit as an example, and uses the fdisk partition tool to set up partitions for the data disk. Initialization operations vary with operating systems.

Prerequisites

- You have logged in to the ECS. For details, see [17.1.3.4 Logging In to an ECS](#).
- A disk has been attached to the ECS and has not been initialized.

Context

Both the fdisk and parted can be used to partition a Linux data disk. For a disk larger than 2 TB, only parted can be used because fdisk cannot partition such a large disk. For details, see [17.1.3.6 Initializing a Linux Data Disk \(parted\)](#).

Creating Partitions and Mounting a Disk

The following example shows how to create a new primary partition on a new data disk that has been attached to an instance. The primary partition will be created using fdisk, and MBR is the default partition style. Furthermore, the partition will be formatted using the ext4 file system, mounted on the `/mnt/sdc` directory, and set to be automatically mounted upon a system start.

Step 1 Run the following command to view information about the added data disk:

fdisk -l

Information similar to the following is displayed: (In the command output, the server contains two disks. `/dev/xvda` is the system disk, and `/dev/xvdb` is the added data disk.)

NOTE

If you do not log in to the ECS and run the **umount** command but directly detach the `/dev/xvdb` or `/dev/vdb` EVS disk on the management console, the disk name in the ECS may encounter a release delay. When you attach the disk to the server again, the mount point displayed on the management console may be inconsistent with that in the server. For example, device name `/dev/sdb` or `/dev/vdb` is selected for attachment, but `/dev/xvdc` or `/dev/vdc` may be displayed as the disk name in the OS. This issue does not adversely affect services.

```
[root@ecs-b656 test]# fdisk -l
```

```
Disk /dev/xvda: 42.9 GB, 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: 0x000cc4ad
```

```
Device Boot      Start         End      Blocks   Id  System
/dev/xvda1    *          2048     2050047     1024000   83  Linux
```

```
/dev/xvda2 2050048 22530047 10240000 83 Linux
/dev/xvda3 22530048 24578047 10240000 83 Linux
/dev/xvda4 24578048 83886079 29654016 5 Extended
/dev/xvda5 24580096 26628095 10240000 82 Linux swap / Solaris
```

Disk /dev/xvdb: 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

NOTE

The capacity displayed here is inconsistent with the capacity of the EVS disk applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 Bytes and 1 GB = 1,000,000,000 Bytes.

Step 2 Run the following command to allocate partitions for the added data disk using `fdisk`:

fdisk *Newly added data disk*

In this example, `/dev/xvdb` is the newly added data disk.

fdisk /dev/xvdb

Information similar to the following is displayed:

```
[root@ecs-b656 test]# fdisk /dev/xvdb
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 0xb00005bd.
Command (m for help):
```

Step 3 Enter `n` and press **Enter**.

Entering `n` creates a partition.

There are two types of disk partitions:

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

```
Command (m for help): n
Partition type:
 p primary (0 primary, 0 extended, 4 free)
 e extended
```

Step 4 Enter `p` and press **Enter**.

The following describes how to create a primary partition.

Information similar to the following is displayed: (**Partition number** indicates the serial number of the primary partition. The value can be **1** to **4**.)

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

Step 5 Enter the primary partition number **1** and press **Enter**.

For example, select **1** as the partition number.

Information similar to the following is displayed: (**First sector** indicates the first sector number. The value can be **2048** to **20971519**, and the default value is **2048**.)

```
Partition number (1-4, default 1): 1
First sector (2048-20971519, default 2048):
```

Step 6 Press **Enter**.

The default start sector number 2048 is used as an example.

Information similar to the following is displayed: (**Last sector** indicates the last sector number. The value can be from **2048** to **20971519**, and the default value is **20971519**.)

```
First sector (2048-20971519, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
```

Step 7 Press **Enter**.

The default last sector number 20971519 is used as an example.

Information similar to the following is displayed, indicating that a primary partition is created for a 10 GB data disk.

```
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
Using default value 20971519
Partition 1 of type Linux and of size 10 GiB is set
Command (m for help):
```

Step 8 Enter **p** and press **Enter** to view the details about the created partition.

Information similar to the following is displayed: (Details about the **/dev/xvdb1** partition are displayed.)

```
Command (m for help): p

Disk /dev/xvdb: 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: 0xb00005bd
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvdb1		2048	20971519	10484736	83	Linux

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

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

Information similar to the following is displayed: (The partition is successfully created.)

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

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

Step 10 Run the following command to synchronize the new partition table to the data disk:

```
partprobe
```


Step 11 Run the following command to set the format for the file system of the newly created partition:

```
mkfs -t File system format /dev/xvdb1
```

For example, run the following command to set the **ext4** file system for the **/dev/xvdb1** partition:

```
mkfs -t ext4 /dev/xvdb1
```

Information similar to the following is displayed:

```
[root@ecs-b656 test]# mkfs -t ext4 /dev/xvdb1
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
655360 inodes, 2621184 blocks
131059 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2151677952
80 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

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

 **NOTE**

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

Step 12 Run the following command to create a mount directory:

```
mkdir Mount directory
```

/mnt/sdc is used in this example.

```
mkdir /mnt/sdc
```

Step 13 Run the following command to mount the new partition to the mount directory created in [Step 12](#):

```
mount /dev/xvdb1 Mount directory
```

For example, run the following command to mount the newly created partition on **/mnt/sdc**:

```
mount /dev/xvdb1 /mnt/sdc
```

Step 14 Run the following command to view the mount result:

```
df -TH
```

Information similar to the following is displayed. The newly created **/dev/xvdb1** partition has been mounted on **/mnt/sdc**.

```
[root@ecs-b656 test]# df -TH
Filesystem Type Size Used Avail Use% Mounted on
/dev/xvda2 xfs 11G 7.4G 3.2G 71% /
```

```
devtmpfs    devtmpfs 4.1G  0 4.1G  0% /dev
tmpfs       tmpfs    4.1G  82k 4.1G  1% /dev/shm
tmpfs       tmpfs    4.1G  9.2M 4.1G  1% /run
tmpfs       tmpfs    4.1G  0 4.1G  0% /sys/fs/cgroup
/dev/xvda3  xfs     1.1G  39M 1.1G  4% /home
/dev/xvda1  xfs     1.1G  131M 915M 13% /boot
/dev/xvdb1  ext4    11G  38M 9.9G  1% /mnt/sdc
```

----End

Setting Automatic Disk Attachment Upon Instance Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify `/dev/xvdb1` in `/etc/fstab`. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in `/etc/fstab` to automatically attach the disk at a system start.

NOTE

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

```
blkid Disk partition
```

For example, run the following command to query the UUID of `/dev/xvdb1`:

```
blkid /dev/xvdb1
```

Information similar to the following is displayed: (The UUID of `/dev/xvdb1` is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the `fstab` file using the vi editor:

```
vi /etc/fstab
```

Step 3 Press `i` to enter the editing mode.

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

```
UUID=xxx attachment directory file system defaults 0 2
```

Assuming that the file system is `ext4` and the attachment directory is `/mnt/sdc`.
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the `fstab` file before detaching the disk. Otherwise, you may fail to access the OS after the disk is detached.

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

The system saves the configuration and exits the vi editor.

----End

17.1.3.6 Initializing a Linux Data Disk (parted)

A data disk attached to an ECS or created together with an ECS must be initialized before it can become available. This section uses an instance running CentOS 7.0 64bit as an example, and uses the parted partition tool to set up partitions for the data disk. Initialization operations vary with operating systems.

Prerequisites

- You have logged in to the ECS. For details, see [17.1.3.4 Logging In to an ECS](#).
- A disk has been attached to the ECS and has not been initialized.

Creating Partitions and Attaching a Disk

The following example shows how to create a new primary partition on a new data disk that has been attached to an instance. The primary partition will be created using parted and GPT is the default partition style. Furthermore, the partition will be formatted using the ext4 file system, mounted on the **/mnt/sdc** directory, and set to be automatically mounted upon a system start.

Step 1 Run the following command to view information about the added data disk:

lsblk

Information similar to the following is displayed:

NOTE

If you do not log in to the ECS and run the **umount** command but directly detach the **/dev/xvdb** or **/dev/vdb** EVS disk on the management console, the disk name in the ECS may encounter a release delay. When you attach the disk to the server again, the mount point displayed on the management console may be inconsistent with that in the server. For example, device name **/dev/sdb** or **/dev/vdb** is selected for attachment, but **/dev/xvdc** or **/dev/vdc** may be displayed as the disk name in the OS. This issue does not adversely affect services.

```
[root@ecs-centos-70 linux]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0 0 40G 0 disk
├─xvda1 202:1 0 4G 0 part [SWAP]
└─xvda2 202:2 0 36G 0 part /
xvdb 202:16 0 10G 0 disk
```

The command output indicates that the server contains two disks. **/dev/xvda** is the system disk and **/dev/xvdb** is the new data disk.

Step 2 Run the following command to enter parted to partition the added data disk:

parted *Added data disk*

In this example, **/dev/xvdb** is the newly added data disk.

parted /dev/xvdb

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# parted /dev/xvdb
GNU Parted 3.1
Using /dev/xvdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
```

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/xvdb: unrecognised disk label
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 10.7GB
Sector size (logical/physical): 512B/512B
Partition Table: unknown
Disk Flags:
```

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

 **NOTE**

The capacity displayed here is inconsistent with the capacity of the EVS disk applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 Bytes and 1 GB = 1,000,000,000 Bytes.

Step 4 Run the following command to set the disk partition style:

```
mklabel Disk partition style
```

The disk partition styles include MBR and GPT. For example, run the following command to set the partition style to GPT:

```
mklabel gpt
```

NOTICE

If you change the disk partition style after the disk has been used, the original data on the disk will be cleared. Therefore, select a proper disk partition style when initializing the disk.

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: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 20971520s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start End Size File system Name Flags
```

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

Step 7 Enter **mkpart opt 2048s 100%** and press **Enter**.

In the command, **opt** is the name of the new partition, **2048s** indicates the start of the partition, and **100%** indicates the end of the partition. You can plan the number and capacity of disk partitions based on 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? Cancel
```

If the preceding warning message is displayed, enter **Cancel** to stop the partitioning. Then, find the first sector with the best disk performance and use that value to partition the disk.

Step 8 Enter **p** and press **Enter** to view the details about the created partition.

Information similar to the following is displayed:

```
(parted) p  
Model: Xen Virtual Block Device (xvd)  
Disk /dev/xvdb: 20971520s  
Sector size (logical/physical): 512B/512B  
Partition Table: gpt  
Disk Flags:  
  
Number Start End Size File system Name Flags  
1 2048s 20969471s 20967424s opt
```

Details about the **/dev/xvdb1** partition are displayed.

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

Step 10 Run the following command to view the disk partition information:

lsblk

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# lsblk  
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT  
xvda 202:0 0 40G 0 disk  
├─xvda1 202:1 0 4G 0 part [SWAP]  
└─xvda2 202:2 0 36G 0 part /  
xvdb 202:16 0 100G 0 disk  
└─xvdb1 202:17 0 100G 0 part
```

In the command output, **/dev/xvdb1** is the partition you created.

Step 11 Run the following command to set the format for the file system of the newly created partition:

NOTICE

The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

mkfs -t *File system format* **/dev/xvdb1**

For example, run the following command to set the **ext4** file system for the **/dev/xvdb1** partition:

mkfs -t ext4 /dev/xvdb1

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# mkfs -t ext4 /dev/xvdb1
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
655360 inodes, 2620928 blocks
131046 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2151677925
80 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

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.

Step 12 Run the following command to create a mount point:

```
mkdir Mount point
```

For example, run the following command to create the **/mnt/sdc** mount point:

```
mkdir /mnt/sdc
```

Step 13 Run the following command to mount the new partition to the mount point created in [Step 12](#):

```
mount /dev/xvdb1 Mount point
```

For example, run the following command to mount the newly created partition on **/mnt/sdc**:

```
mount /dev/xvdb1 /mnt/sdc
```

Step 14 Run the following command to view the mount result:

```
df -TH
```

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# df -TH
Filesystem Type Size Used Avail Use% Mounted on
/dev/xvda2 xfs 39G 4.0G 35G 11% /
devtmpfs devtmpfs 946M 0 946M 0% /dev
tmpfs tmpfs 954M 0 954M 0% /dev/shm
tmpfs tmpfs 954M 9.1M 945M 1% /run
tmpfs tmpfs 954M 0 954M 0% /sys/fs/cgroup
/dev/xvdb1 ext4 11G 38M 101G 1% /mnt/sdc
```

The newly created **/dev/xvdb1** is mounted on **/mnt/sdc**.

----End

Setting Automatic Disk Attachment at a System Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by

referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify `/dev/xvdb1` in `/etc/fstab`. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in `/etc/fstab` to automatically attach the disk at a system start.

 **NOTE**

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

blkid *Disk partition*

For example, run the following command to query the UUID of `/dev/xvdb1`:

blkid /dev/xvdb1

Information similar to the following is displayed: (The UUID of `/dev/xvdb1` is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the `fstab` file using the vi editor:

vi /etc/fstab

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

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

UUID=xxx attachment directory file system defaults 0 2

Assuming that the file system is **ext4** and the attachment directory is `/mnt/sdc`.
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the `fstab` file before detaching the disk. Otherwise, you may fail to access the OS after the disk is detached.

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

The system saves the configuration and exits the vi editor.

----End

17.1.4 Building a Discuz Website

17.1.4.1 Installing the Database

Installing the MySQL Database

Step 1 Log in to an ECS.

Step 2 Run the following command to install the MySQL database server, MySQL client, and libraries and files required for MySQL development:

```
yum install -y mysql-server mysql mysql-devel
```

The installation is successful if the following information is displayed:
Complete!

Step 3 Run the following command to start the MySQL service:

```
service mysqld start
```

The MySQL service starts if the following information is displayed:
Starting mysqld: [OK]

Step 4 Run the following command to set the account and password of the database administrator:

```
mysqladmin -u root password 'xxxxxx'
```

In the preceding command, *xxxxxx* indicates the user-defined password of the database administrator.

Step 5 Run the following command, and then enter the password to log in to the MySQL database:

```
mysql -u root -p  
Enter password
```

Step 6 Enter the administrator password (password of the **root** user) of the MySQL database that is set in step [Step 4](#) to log in to the database. The login is successful if the information "Welcome to the MySQL monitor." is displayed in the command output.

Welcome to the MySQL monitor. Commands end with ; or \g

Step 7 Run the following command to use the MySQL database:

```
use mysql
```

Step 8 Run the following command to display the user list:

```
select host,user from user;
```

Step 9 Run the following command to update the user list and allow all IP addresses to access the database:

```
update user set host='%' where user='root' LIMIT 1;
```

This operation is successful if information similar to the following is displayed:
Query OK, 1 row affected (0.00 sec)
Rows matched:1 Changed:1 Warnings:0

Step 10 Run the following command to forcibly update the permissions:

```
flush privileges;  
Query OK, 0 row affected (0.00 sec)
```


Step 11 Run the following command to exit the database:

```
quit
```

```
Bye
```

Step 12 Run the following command to restart the MySQL service:

```
service mysqld restart
```

```
Stopping mysqld;      [ OK ]  
Starting mysqld;     [ OK ]
```

Step 13 Run the following command to enable the MySQL service to automatically start up upon system boot:

```
chkconfig mysqld on
```

Step 14 Run the following command to disable the firewall:

```
service iptables stop
```

```
iptables: Setting chains to policy ACCEPT: filter [ OK ]  
iptables: Flushing firewall rules: [ OK ]  
iptables: Unloading modules: [ OK ]
```

Step 15 Run the following command to disable the firewall permanently upon server restart:

```
chkconfig iptables off
```

```
----End
```

17.1.4.2 Configuring the Web Environment

Step 1 Log in to an ECS.

Step 2 Run the following command to start the httpd service:

```
service httpd start
```

Step 3 Run the following command to enable the httpd service to automatically start up upon system boot:

```
chkconfig httpd on
```

Step 4 Run the following command to start the php-fpm service:

```
service php-fpm start
```

Step 5 Run the following command to enable the php-fpm service to automatically start up upon system boot:

```
chkconfig php-fpm on
```

Step 6 Run the following command to disable the firewall:

```
service iptables stop
```

Step 7 Run the following command to disable the firewall permanently upon server restart:

```
chkconfig iptables off
```

Step 8 Run the following command to start the MySQL service:

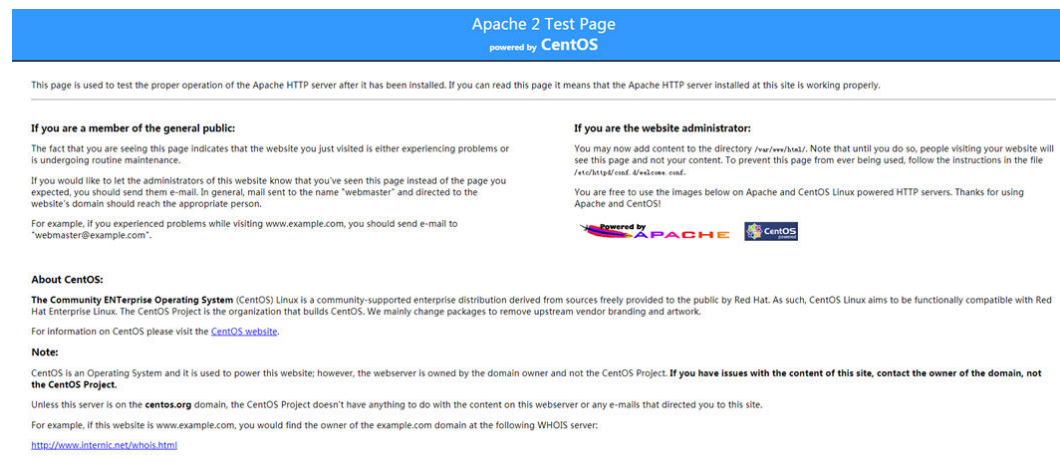
```
service mysqld start
```

Step 9 Run the following command to enable the MySQL service to automatically start up upon system boot:

```
chkconfig mysqld on
```

Step 10 In the address box of a browser, enter **http://elastic IP address** to access the default homepage of the server. The configuration is successful if information similar to **Figure 17-4** is displayed.

Figure 17-4 Proper running of the Apache HTTP server



----End

17.1.4.3 Deploying the Website Code

Step 1 On a local PC, decompress the **Discuz_X3.3_SC_UTF8.zip** file to the **Discuz_X3.3_SC_UTF8** folder.

Step 2 Use WinSCP to upload the files in the **update** folder of the **Discuz_X3.3_SC_UTF8** folder to the **/var/www/html** directory on the ECS.

Step 3 By default, only the **root** user has the write permission. Therefore, you need to log in to the ECS and run the following command to grant the write permission to other users:

```
chmod -R 777 /var/www/html
```

Step 4 In the address box of a browser, enter **http://elastic IP address**.

The Discuz installation page is displayed. Follow the installation wizard to install Discuz.

 NOTE

- Enter the private IP address of the ECS as the IP address of the MySQL database server.
- Enter the administrator password of the MySQL database that is set in [Step 4](#) as the password of the MySQL database.

----End

17.1.5 Checking Whether the Website Is Built

In the address box of a browser, enter **`http://elastic IP address/forum.php`** to visit the forum homepage. If the forum homepage is displayed, the website is built successfully.

18 FAQs

18.1 General FAQs

18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?

Scenarios

This section describes how to log in to ManageOne Operation or Tenant Portal to manage ECSs.

Procedure


Step 1 Log in to ManageOne as a VDC operator using a browser.

URL in non-B2B scenarios: **https://Domain name of ManageOne Operation Portal**, for example, **https://console.demo.com**.

URL in B2B scenarios: **https://Domain name of ManageOne Tenant Portal**, for example, **https://tenant.demo.com**.

You can log in using a password or a USB key.

- Login using a password: Enter the username and password.
The password is that of the VDC operator.
- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.

Step 2 Click  in the upper left corner of the page and select a region. Choose **Computing > Elastic Cloud Server**.

The **Elastic Cloud Server** page is displayed.

Step 3 Select a resource set from the menu bar on the top of the page.

----End

18.1.2 How Do I Handle Error Messages Displayed on ManageOne?

Symptom

An error message is displayed on ManageOne when an error occurs after you perform ECS-related operations.

Context

After you perform ECS-related operations on ManageOne, the system displays the request status on the ECS page. You can determine the request execution status based on the information displayed in the request status.

- If the operation request is successfully 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, find the error code and perform the corresponding operations listed in [Table 18-1](#).

Table 18-1 Error codes and solution suggestions

Error Code	Error Message on ManageOne	Solution
Ecs.0000	Request error. Try again later or contact the system administrator.	Adjust the request structure as required.
Ecs.0001	The maximum number of ECSs or EVS disks has been reached. Contact the system administrator and request a quota increase.	Contact the system administrator to apply for an increased ECS quota. NOTE When applying for increasing your ECS quota, first determine the number of target ECSs, CPU cores (vCPUs), and memory capacity (RAM) required.
Ecs.0003	You do not have the permission or your balance is insufficient.	Contact the system administrator to check the account information.

Error Code	Error Message on ManageOne	Solution
Ecs.0005	Invalid parameter. Refer to the FAQs or contact the system administrator.	Adjust the request structure as required.
Ecs.0010	The private IP address is in use. Select an available IP address and create the ECS again.	Use an idle IP address to create your ECS.
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.
Ecs.0012	The number of IP addresses in the subnet is insufficient. Release IP addresses in the subnet or select another subnet, and create the ECS again.	Release idle IP addresses on the target subnet or use a new subnet for creating ECSs.
Ecs.0013	Insufficient EIP quota. Contact the system administrator and request a quota increase.	Contact the system administrator to apply for an increased EIP quota.
Ecs.0015	Disks of this type are not applicable to the ECS.	Select a supported EVS disk and attach it to the ECS.
Ecs.0100	The ECS status does not meet requirements. Make the ECS in the required status and try again.	Change the ECS status to the required status and try again.

Error Code	Error Message on ManageOne	Solution
Ecs.0101	The system disk status does not meet the requirement.	This error code is an exception protection error code and does not occur in normal scenarios. If this error code is displayed, contact technical support.
Ecs.0103	The disk is unavailable.	Change the ECS status to the required status and try again. If the EVS disk is faulty, contact the system administrator for troubleshooting.
Ecs.0104	Insufficient number of ECS slots for attaching disks.	Detach an EVS disk from the ECS before attaching a new EVS disk.
Ecs.0105	No system disk found.	Reattach the EVS system disk to the ECS and perform the desired operation again.
Ecs.0107	The number of shared disks that can be attached to an ECS exceeds the maximum limit.	Detach an EVS disk from the ECS before attaching a new EVS disk.
Ecs.0201	Failed to create the port.	<ol style="list-style-type: none"> 1. On ManageOne Maintenance Portal, check whether ALM-73203 Component Fault exists. If yes, clear the alarm and try again. 2. On the VPC details page, check whether the VPC and subnet status is normal. If the network status is abnormal, rectify the network exception and try again. 3. If an IP address is specified when you apply for the ECS, check whether the IP address is occupied on the console. If the IP address is occupied, change another IP address and try again. 4. If the fault persists, use CloudMonitorXRay to view error details and contact technical support.
Ecs.0202	Failed to create the system disk.	<ol style="list-style-type: none"> 1. On ManageOne Maintenance Portal, check whether ALM-73203 Component Fault exists. If yes, clear the alarm and try again. 2. Use CloudMonitorXRay to view the error code reported by the EVS and rectify the fault based on the handling procedure of the EVS error code.

Error Code	Error Message on ManageOne	Solution
Ecs.0204	Failed to create the VM.	<ol style="list-style-type: none"> 1. On ManageOne Maintenance Portal, check whether ALM-73203 Component Fault exists. If yes, clear the alarm and try again. 2. If multiple NICs are not allowed to reside in the same subnet, check whether the NICs used for creating the VM belong to the same subnet. If yes, specify different subnets and create the VM again. 3. On Service OM, choose Resources > Compute Resource > Host Groups. Check whether the tag set of the host group contains tags in the VM flavor. Ensure that the tags are case sensitive. Otherwise, the host cannot be selected during VM creation. 4. On Service OM, check whether NUMA affinity is enabled or hugepage memory is configured in the VM flavor. If yes, log in to the FusionSphere OpenStack web client, and choose Configuration > Kernel Option > Host Group > Configuration > Hugepage Configuration to check whether hugepage memory is configured. Otherwise, the host cannot be selected during VM creation. 5. Check whether the CPU and memory resources are sufficient on the hypervisor page of the FusionSphere OpenStack Management Console. Otherwise, the host cannot be selected during VM creation. 6. If the fault persists, use CloudMonitorXRay to view error details and contact technical support.
Ecs.0205	Failed to attach the data disk.	<ol style="list-style-type: none"> 1. On ManageOne Maintenance Portal, check whether ALM-73203 Component Fault exists. If yes, clear the alarm and try again. 2. Use CloudMonitorXRay to view error details and contact technical support.

Error Code	Error Message on ManageOne	Solution
Ecs.0219	Failed to create the VM.	<ol style="list-style-type: none"> 1. On ManageOne Maintenance Portal, check whether ALM-73203 Component Fault exists. If yes, clear the alarm and try again. 2. If multiple NICs are not allowed to reside in the same subnet, check whether the NICs used for creating the VM belong to the same subnet. If yes, specify different subnets and create the VM again. 3. On Service OM, choose Resource > Compute Resource > Host Groups. Check whether the tag set of the host group contains tags in the VM flavor. Ensure that the tags are case sensitive. Otherwise, the host cannot be selected during VM creation. 4. On Service OM, check whether NUMA affinity is enabled or hugepage memory is configured in the VM flavor. If yes, log in to the FusionSphere OpenStack web client, and choose Configuration > Kernel Option > Host Group > Configuration > Hugepage Configuration to check whether hugepage memory is configured. Otherwise, the host cannot be selected during VM creation. 5. Check whether the CPU and memory resources are sufficient on the hypervisor page of the FusionSphere OpenStack Management Console. Otherwise, the host cannot be selected during VM creation. 6. If the fault persists, use CloudMonitorXRay to view error details and contact technical support.
Ecs.1300	Failed to create the data disk.	<ol style="list-style-type: none"> 1. On ManageOne Maintenance Portal, check whether ALM-73203 Component Fault exists. If yes, clear the alarm and try again. 2. Use CloudMonitorXRay to view the error code reported by the EVS and rectify the fault based on the handling procedure of the EVS error code.

Error Code	Error Message on ManageOne	Solution
Other error codes	Other error messages	Initiate the request again. Alternatively, record the returned error code and contact the system administrator for handling.

18.1.3 Common Methods to Troubleshoot Local Client Plugin Used for Attaching ISO

After the client plugin is installed, the system displays a message indicating that the plugin cannot be detected.

Step 1 Open the task manager on the local PC, and check whether the **FSPortalClientPluginKvm.exe** process exists on the **Process** tab.

FSPortalClientPluginKvm.exe	g0039...	00	11,592 K	FCPorta...
FSPortalClientPluginKvm.exe	g0039...	00	19,020 K	FCPorta...

Step 2 If the process does not exist, double-click **startup.bat** in the installation directory to start the plugin.

Step 3 Click **Attach ISO**, and check whether a CD-ROM drive or ISO file can be selected.

- If yes, no further action is required.
- If no, uninstall the plugin and reinstall it.

----End

Port 8887 is disabled or is occupied.

The installation and running of the plugin on the local PC rely on port 8887. Before using the client plugin, ensure that port 8887 on the local PC is not occupied or restricted. If a firewall exists between the local PC and the cloud environment, open port 8887 on the firewall.

Configuration items to be modified

If the plugin is running properly and port 8887 is enabled and not occupied, but the system still displays a message indicating that the plugin information cannot be detected, log in to DMK, change the value of **ecs_check_client_version** to **false**, and try again. For details about how to change the value, visit **Operation Help Center** and choose **Operation > Compute Services > Elastic Cloud Server (ECS) > FAQs > Modifying Configuration Items on DMK**.

Reinstallation of the client plugin fails with the error message "Java(TM) Platform SE binary"

If the plugin has been installed, when you install the plugin again, the installation fails and the message **Java(TM) Platform SE binary** is displayed. In this case,

open the task manager on the local PC and view all processes of the Java application on the **Process** tab.

- If only the **java.exe** process exists, manually stop this process and reinstall the plugin.
- If multiple Java application processes exist and you cannot tell them apart, perform the following:
 - Method 1: Manually clear the residual files in the original installation directory, change the installation directory of the plugin, and reinstall the plugin after 5 minutes. After the installation is complete, reload the **Attach ISO** page and perform the attach operation.
 - Method 2: Restart the local PC and try again.

The default client plugin cannot read data from or write data into logs and import the image file uploading progress to the logs. As a result, the plugin is unavailable.

If the user does not have the highest privileges on the OS or drive C, the default client plugin cannot read data from and write data into logs and import the image file uploading progress to the logs. As a result, the plugin function is unavailable. In this case, handle the issue using the following methods:

- Method 1: If the user does not have the highest privileges on the OS or drive C, it is recommended that the client plugin should be installed in the **C:\Users\Username** directory.
- Method 2: If the user does not have the highest privileges on the OS, grant the user the highest privileges.
- Method 3: If the user does not have the required permissions on drive C, it is recommended that the plugin be installed on a drive to which the user has access, for example, drive D.

18.1.4 What Is Quota?

The system limits the number and capacity of resources for tenants. You can log in to ManageOne as a VDC administrator and choose **System > VDC List** to view the quota.

If your resource quota is insufficient, contact the administrator to change the quota. For details, see visit **Operation Help Center** and choose **Operation > Operation Management > Managing Organizations > Managing Quotas**.

18.2 Image FAQs

18.2.1 What Should I Do If an Image Failed to Be Updated?

Symptom

During ECS creation, the image fails to be updated.

Solution

Check whether the flavor is updated. If yes, perform the following steps:

Step 1 Log in to Service OM.

Step 2 Choose **Services > Resource > Image Resource**.

Step 3 On the **Images** page, click **Image List**.

Step 4 On the right of the image registered by the user, click **Modify**. In the displayed dialog box, select **Instance Type**. The selected instance type matches the selected tag during the flavor creation.

Step 5 If the configuration is correct, contact technical support for assistance.

----End

18.2.2 What Is a Static Injection Image?

Concepts and Application Scenarios

A static injection image is a process of using configdrive and Cloud-Init (for Linux) or Cloudbase-Init (for Windows) to inject the IP address, username, password, and file of an ECS into the ECS. This operation applies to the following scenarios:

- If DHCP is not available, the ECS cannot automatically obtain an IP address. You can inject the preset IP address into the ECS using the static injection image.
- If DHCP is available and the ECS provides important functions, high reliability and stable running are required. If DHCP is faulty, the ECS cannot automatically obtain an IP address, affecting ECS services. In this case, you can use a static injection image to obtain an ECS and inject the planned IP address into the ECS without being affected by DHCP.

Constraints

ECSs created using static injection images do not support Windows automated domain joining.

Implementation Principle

The platform uses OpenStack to inject information into the ECS in standard mode configdrive. The preset NIC information, username, password, and file are injected into the ECS. During application and startup, the preset information is stored on a disk device mounted to the ECS. The disk device is injected into the ECS using Cloud-Init (for Linux) or Cloudbase-Init (for Windows).

18.3 ECS FAQs

18.3.1 What Is the cloudbase-init Account in Windows ECSs?

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.

Do not modify or delete this account or uninstall the Cloudbase-Init agent program. Otherwise, injecting the customized data for initializing the ECS generated using the Windows private image created based on this ECS will fail.

 **NOTE**

This account is unavailable in Linux ECSs.

18.3.2 What Should I Do When an ECS Remains in the Restarting or Stopping State for a Long Time?

If an ECS remains in the **Restarting** or **Stopping** state for over 30 minutes after being restarted, you can forcibly restart or stop the ECS as follows:

Step 1 Access the Elastic Cloud Server console as a VDC administrator or VDC operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 In the ECS list, locate the row that contains the target ECS, click **More** in the **Operation** column, and select **Restart** or **Stop**.

A dialog box is displayed to confirm whether you want to restart or stop the ECS.

Step 3 In the dialog box displayed, select **Forcibly restart** or **Forcibly stop**.

Step 4 Click **OK**.

----End

18.3.3 Can a Deleted ECS Be Provisioned Again?

It depends:

- If you soft delete an ECS, the ECS is removed to the recycle bin and remains in the **Frozen** state for a period of time. An ECS in the **Frozen** state cannot be permanently deleted. After the **Frozen** period, the ECS can be restored.
- If you select **Permanently Delete** when deleting an ECS, the ECS is deleted and cannot be restored.

18.3.4 How Can I Change the Static Host Name of a Linux ECS?

Symptom

The static host name of a Linux ECS is user defined and injected using Cloud-Init during the ECS creation. Although the host name can be changed by running the **hostname** command, the changed host name is restored after the ECS is restarted.

Solution

To make the changed host name take effect even after the ECS is stopped or restarted, save the changed name into configuration files.

The changed host name is assumed to be **new_hostname**.

Step 1 Modify the `/etc/hostname` configuration file.

1. Run the following command to edit the configuration file:
sudo vim /etc/hostname
2. Change the host name to the new one.
3. Run the following command to save and exit the configuration file:
:wq

Step 2 Modify the `/etc/sysconfig/network` configuration file.

1. Run the following command to edit the configuration file:
sudo vim /etc/sysconfig/network
2. Change the **HOSTNAME** value to the new host name.
HOSTNAME=Changed host name
An example is provided as follows:
HOSTNAME=new_hostname
3. Run the following command to save and exit the configuration file:
:wq

Step 3 Modify the `/etc/cloud/cloud.cfg` configuration file.

1. Run the following command to edit the configuration file:
sudo vim /etc/cloud/cloud.cfg
2. 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 host name 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 host name of the new ECS is the host name (**new_hostname**) used by the private image, and user-defined host names cannot be injected using Cloud-Init.
 - Method 2: Delete or comment out - **update_hostname**.
If you use method 2, the changed host name 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 host name permanently takes effect, and user-defined host names (such as **new_new_hostname**) can be injected using Cloud-Init.

Step 4 Run the following command to restart the ECS:

sudo reboot

Step 5 Run the following command to check whether the host name is changed:

sudo hostname

If the changed host name is displayed in the command output, the host name is changed and the new name permanently takes effect.

----End

18.3.5 What Restrictions Are Involved with Using ECSs?

- Do not upgrade the kernel and OS versions of an ECS.
- Do not uninstall the preinstalled performance optimization software from the ECS.
- Do not change the MAC address of the ECS NIC.

18.3.6 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 email systems, web systems, and Enterprise Resource Planning (ERP) systems.

18.3.7 How Long Does It Take to Provision an ECS?

If administrator approval is required, it depends on the response speed of the administrator.

If no administrator approval is required, an ECS can be provisioned within a few minutes in most cases. The time it takes to provision an ECS depends on the ECS flavor, available resources (such as EVS disks and EIPs), and system load.

18.3.8 What Functions Does the Delete Button Provide?

Delete means to delete an ECS. You can soft delete or permanently delete an ECS.

- If you soft delete an ECS, the ECS is removed to the recycle bin and remains in the **Frozen** state for a period of time. An ECS in the **Frozen** state cannot be restored or permanently deleted. After the **Frozen** period, the ECS can be restored or permanently deleted.
- If you select **Permanently Delete** when deleting an ECS, the ECS will be permanently deleted. You can choose to delete the associated EIP and data disks at the same time. If you do not delete them, they are reserved. If necessary, you can manually delete them later.

18.3.9 What Is a Deleted ECS?

Deleted is an intermediate state of the ECS. This status indicates that the ECS is deleted. An ECS in the **Deleted** state cannot provide services. If you soft delete an ECS, the ECS is removed to the recycle bin. If you permanently delete an ECS, the ECS is deleted from the system.

18.3.10 Why Does the Task Status Area Show an ECS Creation Failure But the ECS List Displays the Created ECS?

Symptom

After you created an ECS bound with an EIP, the ECS creation was successful but binding the EIP failed due to insufficient EIPs. Although the **Task Status** showed that the ECS creation failed, the ECS was displayed in the ECS list. The results of the ECS creation task were inconsistent.

Causes

- The ECS list displays details about created ECSs.
- The **Task Status** area shows the task status of the ECS, including the sub-tasks, such as creating the ECS resource and binding an EIP. Only when all subtasks have succeeded, the task status becomes **Succeeded**. Otherwise, the task status is **Failed**.

If the ECS is successfully created but EIP binding fails, the task fails. However, the ECS you created is temporarily displayed in the list. After the system rolls back, the ECS is removed.

18.4 EIP FAQs

18.4.1 Can Multiple EIPs Be Bound to an ECS?

Yes. However, this configuration is not recommended. To configure multiple EIPs, you must manually configure routing policies.

18.4.2 Will a NIC Added to an ECS Be Identified Automatically?

It is recommended that you log in to the ECS after adding a NIC to the ECS to check whether the OS identifies the NIC. If not, manually activate the NIC. For details, see section [12.3 Modifying NIC Configurations](#).

18.5 Login FAQs

18.5.1 What Should I Do After I Log In to an ECS Using VNC and Perform an Operation But the Page Does not Respond for a Long Time?

If your computer is running Windows 7 and you have logged in to the ECS using VNC through Internet Explorer 10 or 11, and the VNC page becomes unresponsive after long time of inactivity, click **AltGr** twice on the VNC page.

18.5.2 What Should I Do If I Cannot View Data After Logging In to an ECS Using VNC?

After you log in to a Linux ECS using VNC and perform a data viewing operation, such as running the `cat` command to view large files or playing videos, VNC may become unavailable due to the high memory usage of the browser. In this case, use another browser to log in to the ECS again.

18.5.3 Why Does a Blank Screen Appear While the System Displays a Message Indicating Successful Authentication After I Attempted to Log In to an ECS Using VNC?

Another user has logged in to this ECS using VNC.

Only one user can log in to an ECS using VNC at a time. If multiple users attempt to log in to an ECS at the same time, only the first user can log in to it successfully. For other users, the system displays a message indicating that the user is authenticated, but the screen turns blank. If this occurs, wait until the other user logs out of the ECS.

18.5.4 Why Was My Login to a Linux ECS with a Key File Unsuccessful?

Symptom

When the key file for creating a Linux ECS is used to attempt to log in to the ECS, login fails.

Possible Causes

Possible causes vary depending on the image used to create the Linux ECS.

- Cause 1: The image used to create the Linux ECS is a private image, on which the Cloud-Init tool is not installed.
- Cause 2: The Cloud-Init tool is installed on the image, but the key pair was not successfully obtained when the ECS was created.

Solution

- If the issue is a result of cause 1, proceed as follows:
If a private image is created without the Cloud-Init tool installed, the ECS configuration cannot be customized successfully. 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 configured when the private image was created. If the original image password or key pair has been lost, use the password resetting function available on the **Elastic Cloud Server** page to reset the password.
- If the issue is a result of cause 2, proceed as follows:
 - a. Select the target ECS and click **Restart**.

- b. Use the key file to log in to the ECS again and check whether the login is successful.
 - If yes, no further action is required.
 - If no, contact the system administrator.

18.5.5 Why Does the System Display a Message Indicating that the Password for Logging In to a Windows ECS Cannot Be Queried?

Symptom

Password authentication mode is required to log in to a Windows ECS. Therefore, a key file is required to obtain the initial password for logging in to the ECS. However, when I clicked **Get Password** (see section [6.3.1 Obtaining the Password for Logging In to a Windows ECS](#)), the system displayed a message indicating that the password failed to obtain. As a result, I failed to log in to the ECS.

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 the Cloudbase-Init tool is not installed.
- Cause 2: The Cloudbase-Init tool is installed on the image, but the key pair was not successfully 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 the Cloudbase-Init tool installed, the ECS configuration cannot be customized successfully. 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 original image password has been lost, use the password resetting function available on the **Elastic Cloud Server** page to reset the password.
- If the issue is a result of cause 2, proceed as follows:
 - a. Select the target ECS and click **Restart** to restart it.
 - b. After the restart is successful, choose **More > Change Settings > Get Password** in the **Operation** column to check whether the password can be obtained.
 - If yes, no further action is required.
 - If no, contact the system administrator.

18.5.6 What Should I Do If I Cannot Use MSTSC to Log In to an ECS Running Windows Server 2012?

Symptom

An ECS running Windows Server 2012 has the password authentication mode configured during ECS creation. When a user uses the initial password and MSTSC to log in to the ECS, the login fails and the system displays the message "You must change your password before logging on for the first time. Please update your password or contact your system administrator or technical support."

Possible Causes

The local computer used by the user is running the Windows 10 OS.

Due to limitations, the Windows 10 OS does not support remote logins to an ECS running the Windows Server 2012 OS using the initial password.

Solutions

- Solution 1
Use a local computer running the Windows 7 OS for remotely logging in to the ECS running the Windows Server 2012 OS.
- Solution 2
Retain the original local computer and change the initial login password.
 - a. Use VNC to log in to the ECS running the Windows Server 2012 OS for the first time.
 - b. Change the login password as prompted.
 - c. Use the changed password and MSTSC to log in to the ECS again.
- Solution 3:
Retain the original local computer and initial login password.
 - a. Choose **Start**. In the **Search programs and files** text box, enter **mstsc** and press **Enter**.
The **Remote Desktop Connection** page is displayed.
 - b. Enter the elastic IP address and click **Connect**. Then, use the username **administrator** and the login password configured during ECS creation to establish the connection.
The connection fails, and the system displays the message "You must change your password before logging on for the first time. Please update your password or contact your system administrator or technical support."
 - c. Click **Options** in the lower left corner of the **Remote Desktop Connection** page.
 - d. On the **General** tab, click **Save As** in the **Connection settings** pane and save the remote desktop file in .rdp format.
 - e. Use Notepad++ to open the .rdp file.
 - f. Add the following statement to the last line of the .rdp file and save the file.

enablecredsspsupport:i:0

- g. Double-click the edited .rdp file to set up the remote desktop connection.
- h. Click **Connect** to connect to the ECS running the Windows Server 2012 OS again.

18.6 Network and Security FAQs

18.6.1 Configuring a Static IP Address for an ECS

This section describes the common procedures for configuring a static IP address for an ECS. Some of the steps may vary depending on the OS type and version.

Windows

This section uses Windows 7 Professional as an example to describe how to configure a static IPv4 or IPv6 address. For other Windows OSs, see the corresponding guides on the official website.

- Step 1** Choose **Start > Control Panel > Network and Internet > Network and Sharing Center**.
- Step 2** In the **View your active networks** area, click the connection you want to modify.
- Step 3** Click **Properties**.
- Step 4** On the **Networking** tab, select **Internet Protocol Version 4 (TCP/IPv4)** or **Internet Protocol Version 6 (TCP/IPv6)** under **This connection uses the following items**, and then click **Properties**.
- Step 5** Select **Use the following IP address** or **Use the following IPv6 address**, set the required parameters, and click **OK**.

----End

SUSE Linux Enterprise Server Series

NOTE

Set the parameters according to your network settings.

- Step 1** Log in to the ECS as the root user.
- Step 2** Run the following command to check the network devices of the ECS. Record the name of the NIC to be configured.

ifconfig

- Step 3** Run the following command to open the NIC configuration file:

```
vi /etc/sysconfig/network/ifcfg-eth0
```

In the command above, **eth0** indicates the name of the NIC to be configured. Set this parameter according to the query result in [Step 2](#).

- Step 4** Modify the parameters in the configuration file according to your network settings.

```
BOOTPROTO='static' #Static IP address  
BROADCAST="" #Broadcast address  
IPADDR='192.168.83.247' #IP address  
NETMASK='255.255.255.0' #Subnet mask  
NETWORK='192.168.83.0' #Network address  
STARTMODE='auto' #Start the network immediately after ECS startup.
```

Step 5 Run the following command to set the gateway:

```
vi /etc/sysconfig/network/routes
```

```
default 192.168.83.2 - -
```

Step 6 Run the following command to restart the network:

```
service network restart
```

----End

Red Hat Series

Step 1 Log in to the ECS as the root user.

Step 2 Run the following command to view all configuration files of the ECS and confirm the name of the configuration file for the NIC to be configured, for example, **ifcfg-eth1**.

```
ls /etc/sysconfig/network-scripts/ifcfg-*
```

Step 3 Run the following command to open the NIC configuration file:

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

Step 4 In the configuration file, change the value of **BOOTPROTO** to **static**, and add information about the static IP address according to your network settings.

```
BROADCAST=192.168.1.255 #Broadcast address
```

```
IPADDR=192.168.1.33 #IP address
```

```
NETMASK=255.255.255.0 #Subnet mask
```

```
GATEWAY=192.168.1.1 #Gateway
```

Step 5 Press **Esc**, enter **:wq**, and then press **Enter** to save the change and exit.

Step 6 Run the following command to restart the network service:

```
service network restart
```


----End

18.6.2 Why Can I Remotely Connect to an ECS But Cannot Ping It?

To successfully ping an ECS, enable Internet Control Message Protocol (ICMP) in the security group rules for the ECS.

18.6.3 What Should I Do If a Public Key Cannot Be Imported?

If you use Internet Explorer 9 to access the management console, the public key may fail to import or the file injection function may become unavailable. In this case, perform the following steps to modify browser settings and then try again:

- Step 1** Click  in the upper right corner of the browser.
 - Step 2** Select **Internet Options**.
 - Step 3** Click the **Security** tab in the displayed dialog box.
 - Step 4** Click **Internet**.
 - Step 5** If the security level indicates **Custom**, click **Default Level** to restore to the default settings.
 - Step 6** Move the scroll bar to set the security level to **Medium** and click **Apply**.
 - Step 7** Click **Custom Level**.
 - Step 8** Set **Initialize and script ActiveX controls not marked as safe for scripting** to **Prompt**.
 - Step 9** Click **Yes**.
- End

18.6.4 What Should I Do If a Public Key Fails to Be Imported to ManageOne After a Key Pair Is Created Using PuTTYgen?

Symptom

When a key pair created using PuTTYgen is imported, the system displays a message indicating that importing the public key failed.

Possible Causes

The format of the public key content does not meet system requirements.

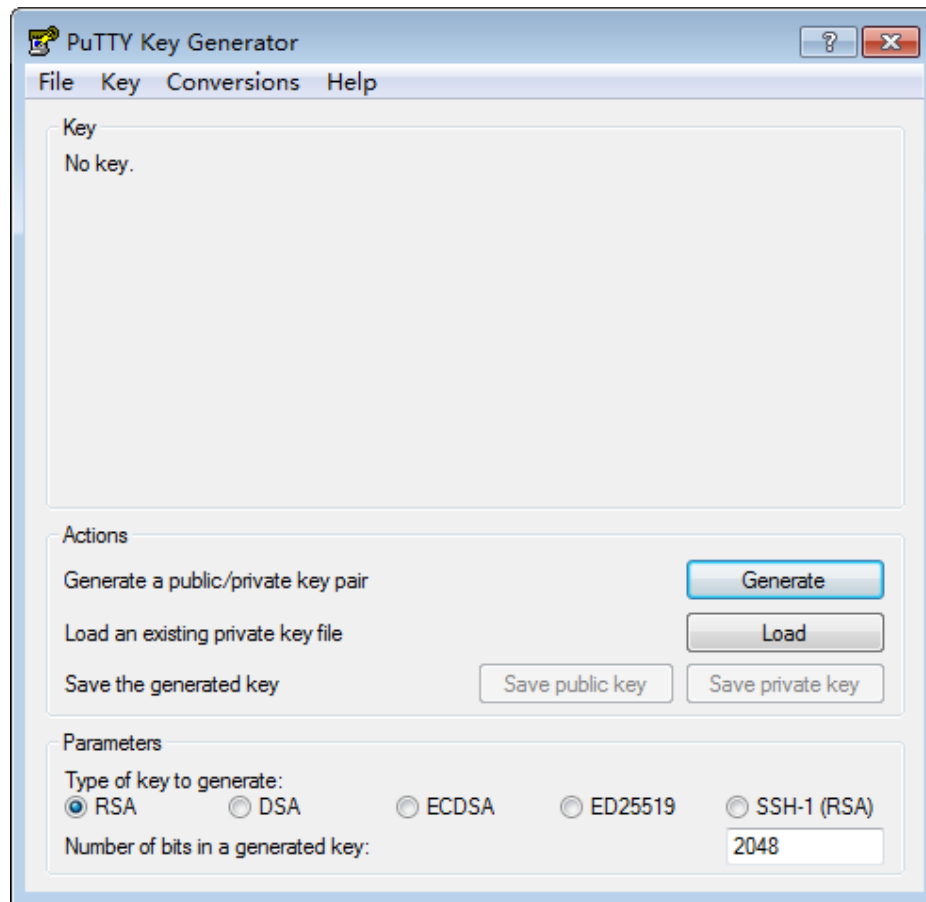
When you use PuTTYgen to create a key pair and store the public key by clicking **Save public key**, the format of the public key content will be changed. Such a key cannot be imported.

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.

- Step 1** Double-click **puttygen.exe**. The **PuTTY Key Generator** page is displayed. The following uses 0.70 (64-bit) as an example.

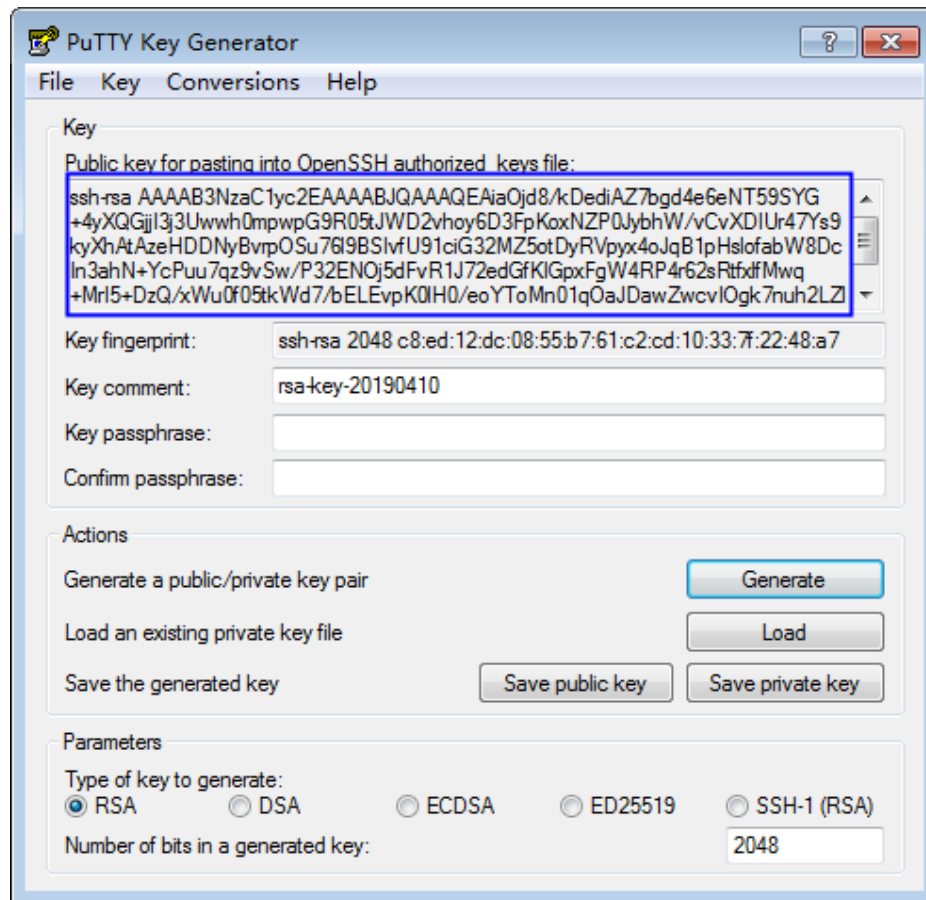
Figure 18-1 PuTTY Key Generator



Step 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 box in [Figure 18-2](#) is the public key with the format meeting system requirements.

Figure 18-2 Restoring the format of the public key content



Step 3 Copy the public key content to a .txt file and save the file in a local directory.

Step 4 Import the public key file.

1. Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
2. In the navigation pane, choose **Key Pair**.
3. On the **Key Pair** page, click **Import Public Key**.
4. Copy the public key content in the .txt file to **Public Key Content** and click **OK**.

----End

18.6.5 How Can I Change the MTU of a Linux ECS NIC?

Maximum transmission unit (MTU) specifies the largest packet of data that can be transmitted on a network and ranges from 1,280 to 8,888 in the unit of byte. This section uses a disk-intensive Linux ECS as an example to describe how to change the NIC MTU of ECSs running SUSE, CentOS, and Ubuntu OSs, respectively.

SUSE Linux OSs

The following operations use the SUSE Enterprise Linux Server 11 SP3 64bit OS as an example to describe how to change the MTU:

Step 1 Log in to the ECS as an administrator and run the **sudo su root** command to switch to user **root**.

Step 2 Run the **ifconfig** command to query the network ports of the ECS.

Step 3 Run the following command to open the **ifcfg-xxx** configuration file:

```
vi /etc/sysconfig/network/ifcfg-xxx
```

xxx specifies the network port queried in [Step 2](#), for example, eth0.

Step 4 Press **i** and run the following command to configure the MTU for the NIC:

```
MTU=8888
```



```
INTERFACETYPE=
## Type: integer
## Default: ""
##
## Set a specific Maximum Transmission Unit (MTU) for this interface.
##
MTU=8888
## Type: string
## Default: ""
##
## Set a specific link level address (LLADDR), a.k.a. MAC address or hardware
## address, for this interface (if the driver supports it).
##
LLADDR=
```

Step 5 Press **Esc** and run the **:wq!** command to save and exit the configuration file.

Step 6 Run the following command to restart the network:

```
service network restart
```

Step 7 Run the **ifconfig** command to check whether the MTU value has been changed.

```
----End
```

CentOS

The following operations use the CentOS 7.2 64bit OS as an example to describe how to change the MTU:

Step 1 Log in to the ECS as an administrator and run the **sudo su root** command to switch to user **root**.

Step 2 Run the **ifconfig** command to view the NIC that has a bound IP address.

Step 3 Run the following command to open the **ifcfg-xxx** configuration file:

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

xxx specifies the network port queried in [Step 2](#), for example, eth0.

Step 4 Press **i** and run the following command to configure the MTU for the NIC:

```
MTU=8888
```

```
DEVICE="eth0"  
BOOTPROTO="dhcp"  
ONBOOT="yes"  
STARTMODE="onboot"  
MTU=8888
```

Step 5 Press **Esc** and run the **:wq!** command to save and exit the configuration file.

Step 6 Run the following command to restart the network:

```
service network restart
```

Step 7 Run the **ifconfig** command to check whether the MTU value has been changed.

----End

Ubuntu Linux OSs

Step 1 Log in to the ECS as an administrator and run the **sudo su root** command to switch to user **root**.

Step 2 Run the following command to open the **interfaces** file:

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

Step 3 Press **i** and run the following command to configure the MTU for the NIC:

```
post-up /sbin/ifconfig/ eth0 mtu 8888
```

```
# This file describes the network interfaces available on your system  
# and how to activate them. For more information, see interfaces(5).  
  
# The loopback network interface  
auto eth0  
iface eth0 inet dhcp  
post-up /sbin/ifconfig/ eth0 mtu 8888
```

Step 4 Press **Esc** and run the **:wq!** command to save and exit the **interfaces** file.

Step 5 Run the following command to restart the network:

```
/etc/init.d/networking restart
```

Step 6 Run the **ifconfig** command to check whether the MTU value has been changed.

----End

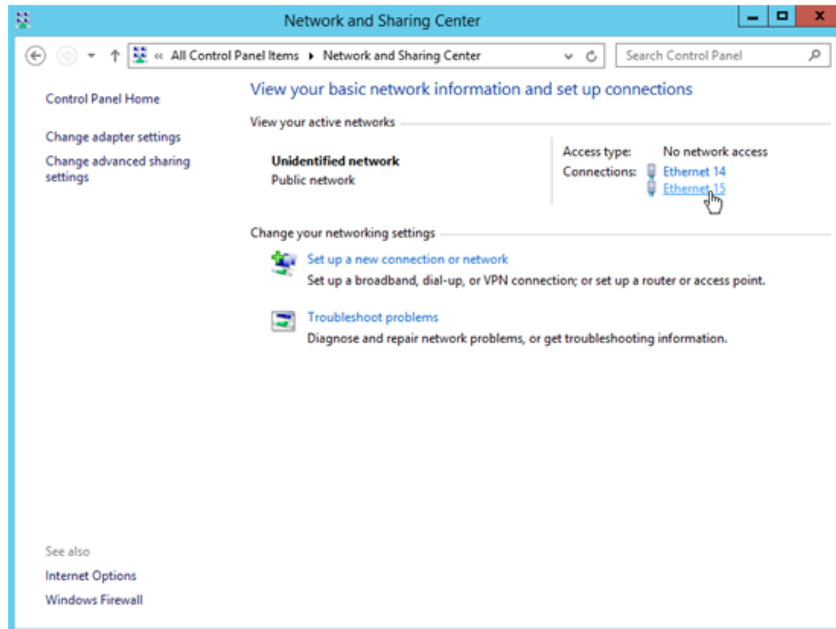
18.6.6 How Can I Change the MTU of a Windows ECS NIC?

Maximum transmission unit (MTU) specifies the largest packet of data that can be transmitted on a network and ranges from 1,280 to 8,888 in the unit of byte. This section uses an ECS running the Windows Server 2012 OS as an example to describe how to change the MTU of the ECS NIC.

Step 1 Enable Jumbo Packet on the NIC.

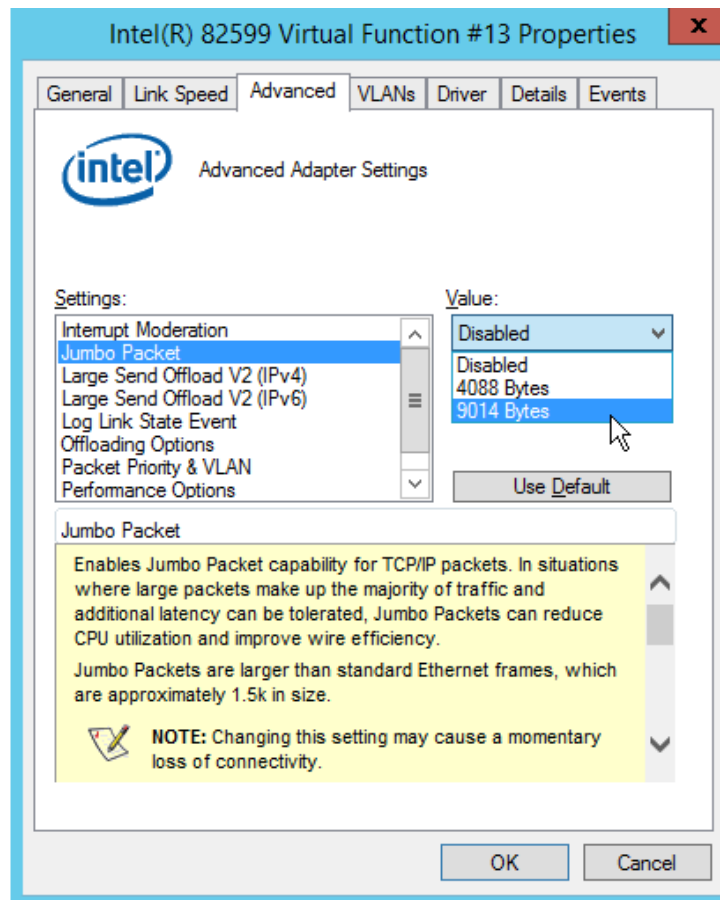
1. Click **Start** in the task bar and choose **Control Panel**.
2. Click **View network status and tasks** under **Network and Internet**.

Figure 18-3 Network and Sharing Center



3. In the **View your active networks** area, select the target NIC. Take the Ethernet 15 NIC shown in **Figure 18-3** as an example.
Click **Ethernet 15**.
The page showing the Ethernet 15 NIC status is displayed.
4. Click **Properties**.
The page showing the Ethernet 15 NIC properties is displayed.
5. Click **Configure**. In the dialog box that is displayed, click the **Advanced** tab.

Figure 18-4 Inter(R) 82599 Virtual Function #13 Properties

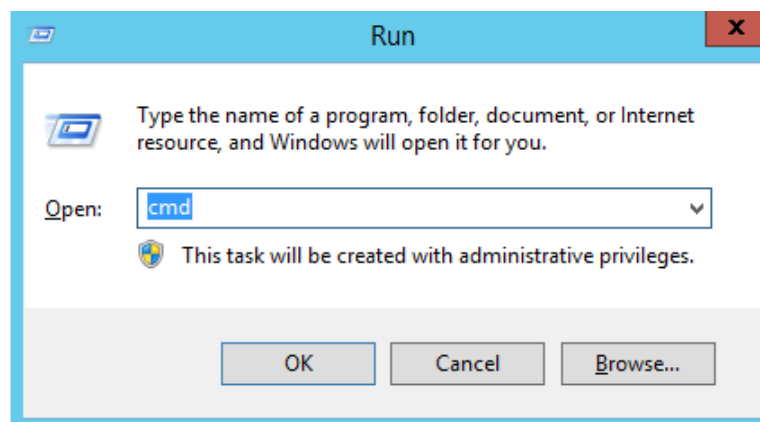


6. In the **Settings** area, select **Jumbo Packet**. In the **Value** area, select **9014 Bytes**.
7. Click **OK**.

Step 2 Change the MTU.

1. Press **Win+R** to open the **Run** dialog box.

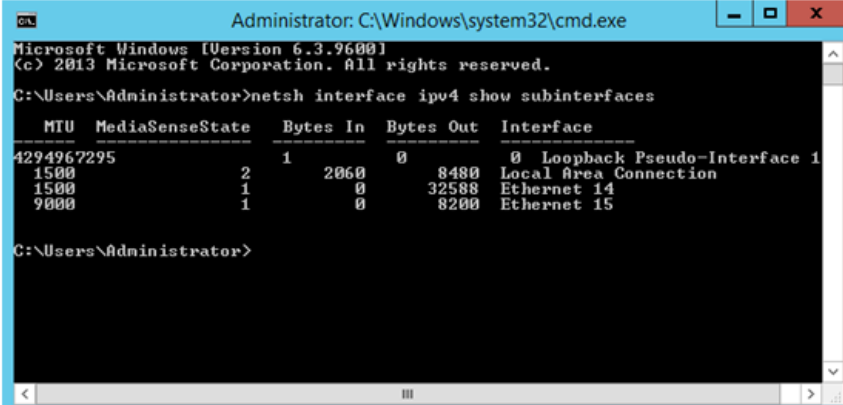
Figure 18-5 Run



2. In the **Open** text box, enter **cmd** and click **OK**.
3. Run the following command to view the NIC MTU:

netsh interface ipv4 show subinterfaces

4. Obtain the result. The MTU of the NIC with Jumbo Packet enabled is 9,000.

Figure 18-6 Obtaining the NIC MTU

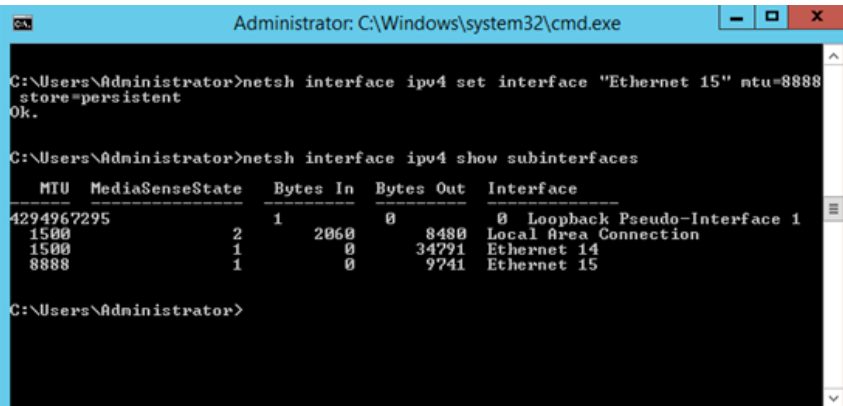
```
Administrator: C:\Windows\system32\cmd.exe
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>netsh interface ipv4 show subinterfaces

   MTU   MediaSenseState   Bytes In   Bytes Out   Interface
-----
4294967295      1           0           0   0 Loopback Pseudo-Interface 1
1500           2           2060          8480   Local Area Connection
1500           1           0          32588   Ethernet 14
9000           1           0           8200   Ethernet 15

C:\Users\Administrator>
```

5. Run the following command to change the NIC MTU:
netsh interface ipv4 set interface "NIC name" mtu=Changed MTU store=persistent
For example, if the MTU of the Ethernet 15 NIC is changed to 8888, run the following command:
netsh interface ipv4 set interface "Ethernet 15" mtu=8888 store=persistent
6. Run the following command to view the changed NIC MTU:
netsh interface ipv4 show subinterfaces

Figure 18-7 Obtaining the changed NIC MTU

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

C:\Users\Administrator>netsh interface ipv4 set interface "Ethernet 15" mtu=8888 store=persistent
Ok.

C:\Users\Administrator>netsh interface ipv4 show subinterfaces

   MTU   MediaSenseState   Bytes In   Bytes Out   Interface
-----
4294967295      1           0           0   0 Loopback Pseudo-Interface 1
1500           2           2060          8480   Local Area Connection
1500           1           0          34791   Ethernet 14
8888           1           0           9741   Ethernet 15

C:\Users\Administrator>
```

----End

18.6.7 Accessing the Internet Using an ECS Without a Public IP Address

Public IP addresses are scarce resources. This section describes how An ECS not bound with a public IP address can access the Internet through an ECS bound with a public IP address in the same subnet.

Context

To ensure platform security and conserve public IP address resources, public IP addresses are assigned only to specified ECSs. ECSs without public IP addresses cannot access the Internet directly. If these ECSs need to access the Internet (to perform a software upgrade or install a patch, for example), you can select an ECS with a bound public IP address to be an agent ECS to access the Internet.

Prerequisites

- An agent ECS with a bound public IP address is available.
In this example, the agent ECS runs CentOS 6.5.
- The agent ECS and the ECS that needs to access the Internet through the agent ECS are in the same subnet and same security group.

Procedure

- Step 1** Access the Elastic Cloud Server console as a VDC administrator or operator. For details, see [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)
- Step 2** In the search box above the upper right corner of the ECS list, enter the agent ECS name and click the search icon.
- Step 3** Click the name of the agent ECS. The page providing details about the ECS is displayed.
- Step 4** Click the **NIC** tab and expand the details. Then, set **Source/Destination Check** to **OFF**.
- Step 5** Log in to the agent ECS.
For details, see section [6.1 Login Mode Overview](#).
- Step 6** Run the following command to check whether the agent ECS can successfully connect to the Internet:
- ```
ping www.google.com
```
- The agent ECS can successfully connect to the Internet if information similar to the following is displayed:
- ```
64 bytes from 172.16.111.148: icmp_seq=1 ttl=51 time=9.34 ms
64 bytes from 172.16.111.148: icmp_seq=2 ttl=51 time=9.11 ms
64 bytes from 172.16.111.148: icmp_seq=3 ttl=51 time=8.99 ms
```
- Step 7** Run the following command to check whether IP forwarding is enabled on the agent ECS:
- ```
cat /proc/sys/net/ipv4/ip_forward
```
- If **0** (disabled) is displayed, go to [Step 8](#).
  - If **1** (enabled), go to [Step 14](#).
- Step 8** Run the following command to open the IP forwarding configuration file in the vi editor:
- ```
vi /etc/sysctl.conf
```

Step 9 Press **i** to enter editing mode.

Step 10 Set the value of the `net.ipv4.ip_forward` parameter to **1**.

Step 11 Press **Esc**, type `:wq`, and press **Enter**.

The system saves the configurations and exits the vi editor.

Step 12 Run the following command to effect the modification:

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

Step 13 Run the following command to delete the original iptables rule:

```
iptables - F
```

Step 14 Run the following command to configure source network address translation (SNAT) to enable ECSs in the same network segment as the agent ECS to access the Internet:

```
iptables -t nat -A POSTROUTING -o eth0 -s subnet/netmask-bits -j SNAT --to nat-instance-ip
```

For example, if the agent ECS is on 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:

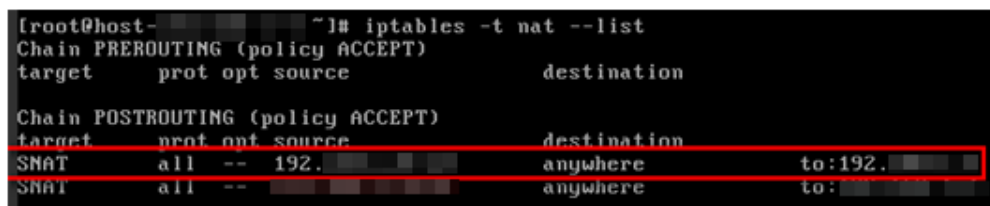
```
iptables -t nat -A POSTROUTING -o eth0 -s 192.168.125.0/24 -j SNAT --to 192.168.125.4
```

Step 15 Run the following command to check whether SNAT is successfully configured:

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

SNAT has been successfully configured if information similar to [Figure 18-8](#) is displayed.

Figure 18-8 Successful SNAT configuration



```
[root@host- ~]# iptables -t nat --list
Chain PREROUTING (policy ACCEPT)
target     prot opt source                destination

Chain POSTROUTING (policy ACCEPT)
target     prot opt source                destination
SNAT      all  --  192.168.125.0/24      anywhere             to:192.168.125.4
SNAT      all  --  anywhere             anywhere            to:192.168.125.4
```

Step 16 Add a route.

1. Log in to ManageOne as a VDC administrator or operator using a browser.
 - URL in non-B2B scenarios: **https://Address for accessing ManageOne Operation Portal**, for example, **https://console.demo.com**.
 - URL in B2B scenarios: **https://Address for accessing ManageOne Tenant Portal**, for example, **https://tenant.demo.com**.
 - Username and password: username and password of a VDC administrator or operator
2. Under **Network**, click **Virtual Private Cloud**.
3. Select a VPC to which a route is to be added and click **Route Table**. On the **Route Table** page, click **Add Route**.

4. Set route information as prompted.
 - **Destination:** indicates the destination network segment. The default value is **0.0.0.0/0**.
 - **Next Hop:** indicates the private IP address of the SNAT ECS.
You can obtain the private IP address of the ECS on the ECS page.

----End

18.6.8 What Do I Do If the Virtual IP Address Cannot Be Pinged After Being Bound to the ECS NIC?

Troubleshooting Method

1. Check whether the source/destination check function of the NIC is disabled and whether the virtual IP address has been bound to the NIC.
2. Check whether the sub-interface of the ECS NIC has been created properly.
3. Check whether the ECS security groups and the network ACL rules associated with the subnets used by the ECS NICs block traffic.

Troubleshooting Procedure

Step 1 Check whether the source/destination check function of the NIC is disabled and whether the virtual IP address has been bound to the NIC.

1. Log in to the management console.
2. Choose **Computing > Elastic Cloud Server**.
3. Click the target ECS name in the ECS list.
4. On the displayed ECS details page, click the **NIC** tab.
5. Ensure that **Source/Destination Check** is disabled.
6. Ensure that a virtual IP address has been bound to the ECS. On the NIC details page, click **Manage IP Address**. On the **Virtual IP Address** tab page of the subnet where the NIC resides, check whether the virtual IP address has been bound to the ECS.

If no virtual IP address has been bound to the ECS, bind a virtual IP address to the ECS on the **Virtual IP Address** tab page of the subnet where the NIC resides.

Step 2 Check whether the sub-interface of the ECS NIC has been created properly.

For a Linux ECS:

1. Run the following command on the ECS to check whether the NIC of the **ethX:X** type exists:
ifconfig


```
[root@scy ~]# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 192.168.1.2 netmask 255.255.255.0 broadcast 192.168.1.255
    inet6 fe80::f816:3eff:fe4d:5b98 prefixlen 64 scopeid 0x20<link>
    ether fa:16:3e:4d:5b:98 txqueuelen 1000 (Ethernet)
    RX packets 77399 bytes 5101164 (4.8 MiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 68798 bytes 8090922 (7.7 MiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

eth0:1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 192.168.1.137 netmask 255.255.255.0 broadcast 192.168.1.255
    ether fa:16:3e:4d:5b:98 txqueuelen 1000 (Ethernet)
```

The command output in the preceding figure contains the NIC of the **ethX:X** type. **192.168.1.137** is the virtual IP address in **1**.

- If yes, the sub-interface of the ECS NIC has been created properly.
 - If no, go to **2.2**.
2. If the command output does not contain the NIC of the **ethX:X** type, run the following command to switch to the **/etc/sysconfig/network-scripts** directory:

```
cd /etc/sysconfig/network-scripts
```

3. Run the following command to create and then modify the **ifcfg-eth0:1** file:
vi ifcfg-eth0:1

Add the following NIC information to the file:

```
BOOTPROTO=static
DEVICE=eth0:1
HWADDR=fa:16:3e:4d:5b:98
IPADDR=192.168.1.137
GATEWAY=192.168.1.1
NETMASK=255.255.255.0
ONBOOT=yes
ONPARENT=yes
```

4. Press **Esc**, enter **:wq!**, and save the file and exit.
5. Restart the ECS and run the **ifconfig** command to check whether the virtual IP address has been configured.

For a Windows ECS:

1. Run the following command on the ECS to check whether the sub-interface of the ECS NIC has been created properly:

```
ipconfig
```

- If the command output contains the virtual IP address bound to the NIC in **1**, the sub-interface has been created properly.
 - If the command output does not contain the virtual IP address bound to the NIC in **1**, go to **2.2**.
2. On the ECS, choose **Control panel > Network & Internet > Network and Sharing Center**.
 3. Under **View your basic network information and set up connections** and **View your active networks**, locate the current network, and click the current connection.

4. In the displayed dialog box, click **Properties**.
5. Double-click **Internet Protocol Version 4**.
6. Select **Use the following IP address**.
7. Click **Advanced**. The **Advanced TCP/IP Settings** dialog box is displayed.
8. On the **IP Settings** tab page, click **Add** in the **IP addresses** area.

 **NOTE**

If the ECS dynamically obtains an IP address using DHCP, you also need to add the IP address contained in the command output in 2.1 as the private IP address of the ECS.

Step 3 Check whether the ECS security groups and the network ACL rules associated with the subnets used by the ECS NICs block traffic.

1. On the ECS details page, click the **Security Groups** tab and confirm that required security group rules have been configured for the virtual IP address. If the required security group rules have not been configured, click the **Security Groups** tab, click the ID of the security group to display the details of the security group, and then reconfigure security group rules.
2. Go to the **Network ACL** page and check whether the network ACL rules associated with the subnets used by the ECS NICs block access to the virtual IP address.

----End

18.7 Disk FAQs

18.7.1 Logging In to the EVS Console as a VDC Administrator or VDC Operator

Procedure


Step 1 Log in to ManageOne as a VDC administrator or VDC operator using a browser.

URL in non-B2B scenarios: **https://Domain name of ManageOne Operation Portal**, for example, **https://console.demo.com**.

URL in B2B scenarios: **https://Domain name of ManageOne Tenant Portal**, for example, **https://tenant.demo.com**.

You can log in using a password or a USB key.

- Login using a password: Enter the username and password.
The password is that of the VDC administrator or VDC operator.
- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.

Step 2 Click  in the upper left corner, select a region and resource set, and choose **Storage > Elastic Volume Service**. The EVS console is displayed.

----End

18.7.2 How Can I Attach a New EVS Disk to an ECS?

Log in to ManageOne Tenant Portal or ManageOne Operation Portal, click **Elastic Volume Service** under **Storage**, create an EVS disk, and attach the disk to an ECS. The EVS disk becomes available to the ECS after you initialize the disk.

For details about how to create, attach, and initialize an EVS disk, see [10 EVS Disk](#).

18.7.3 Can Multiple EVS Disks Be Attached to a Single ECS?

The total number of system and data disks that can be attached to an ECS cannot exceed 60.

NOTE

- If you create an ECS earlier than FusionSphere Service 6.3.1, a maximum of 11 disks can be attached to your ECS.
- If the number of disks that can be attached to an ECS is less than the number that you specify, some drive letters have been pre-occupied by the system. For example, occupied by the system, CD-ROM drive, or pass-through disk. Use the actual number of disks that can be attached.

SCSI disks can be used as only data disks. A maximum of 59 SCSI disks can be attached. The VBD disk can act as either a system disk or a data disk. [Table 18-2](#) describes the relationship between the maximum number of disks that can be attached to an ECS and the **Disk Device Type** value you set during Service OM registration for the image used for ECS creation.

Table 18-2 Total number of VBD disks that can be attached

Disk Device Type	Total VBD Disks
ide	4 (x86)
	0 (ARM) NOTE When ARM servers are used, Disk Device Type cannot be set to ide during image registration.

Disk Device Type	Total VBD Disks
virtio	<p>24</p> <p>NOTE</p> <ul style="list-style-type: none"> When ARM servers are used, if the ECS bus type is Virtio, the total number of NICs cannot exceed 16, the total number of VBD disks cannot exceed 24, and the total number of NICs, disks (EVS disks and pass-through disks), and NPU cards cannot exceed 24. When x86 servers are used, if Boot Mode of the image is set to UEFI during ECS creation: During online disk attachment, the mount point must be between vda and vdp. A maximum of 16 disks (including system disks) can be attached. If the mount point exceeds vdp, for example, vdq, you must shut down the ECS, attach the target disk, and then start the ECS. Such a process is an offline disk attachment process.
scsi	60

18.7.4 Can I Attach an EVS Disk to Multiple Instances?

A non-shared EVS disk can be attached to only one instance.

A shared EVS disk can be attached to a maximum of 16 instances by default.

18.7.5 How Many States Does an EVS Disk Have?

Disks are managed using the EVS during the entire process from disk creation to release, ensuring optimal user experience of applications or sites hosted on them.

[Figure 18-9](#) shows the switching between different states of an EVS disk, and [Table 18-3](#) describes the meaning and supported operations of each state.

Figure 18-9 EVS disk status

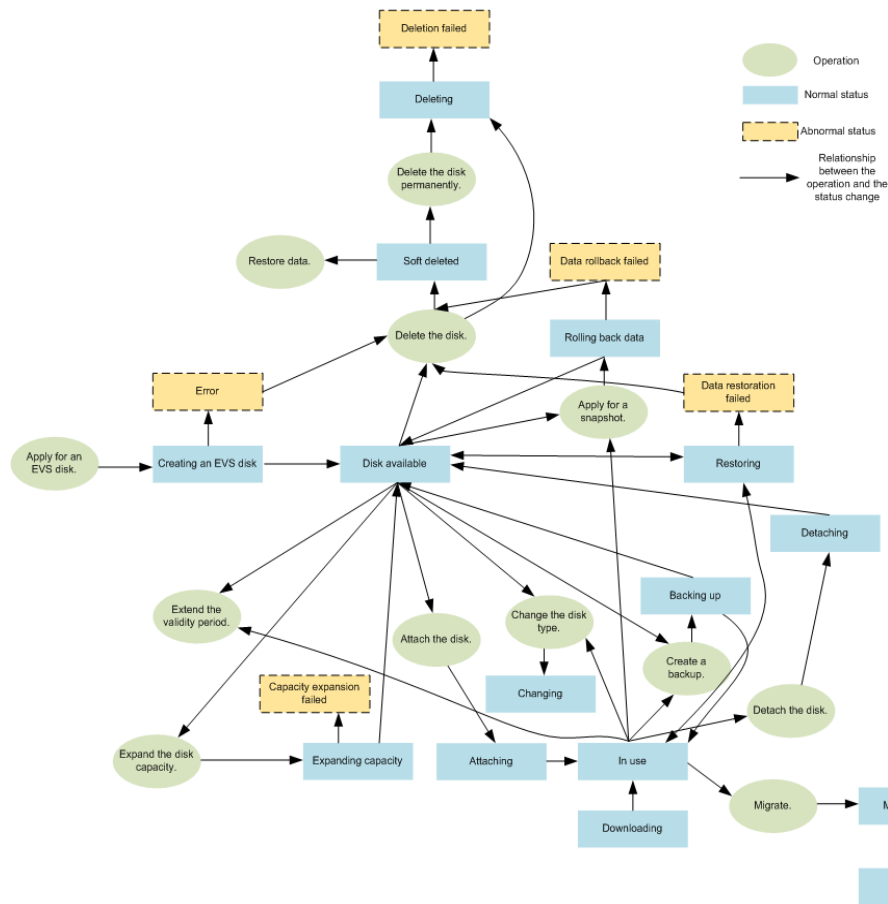


Table 18-3 EVS disk status description

EVS Disk Status	Status Description	Supported Operation
In-use	The EVS disk is attached to an instance and is in use.	<ul style="list-style-type: none"> Detach Expand capacity Create snapshot Create backup Extend validity period of EVS disks (except EVS disks whose expiration time is Unlimited) Change disk type
Downloading	Data is being downloaded from an image to the EVS disk. An EVS disk is in this status when you are creating an ECS.	-

EVS Disk Status	Status Description	Supported Operation
Available	The EVS disk is created and has not been attached to any server.	<ul style="list-style-type: none"> ● Attach ● Expand capacity ● Create snapshot ● Create backup ● Delete ● Extend validity period of EVS disks (except EVS disks whose expiration time is Unlimited) ● Change disk type
Creating	The EVS disk is being created.	-
Attaching	The EVS disk is being attached to an instance.	-
Detaching	The EVS disk is being detached from an instance.	-
Deleting	The EVS disk is being deleted.	-
Extending	The capacity of the EVS disk is being expanded.	-
Rolling back	The EVS disk data is being rolled back using the snapshot.	-
Backing up	The EVS disk is being backed up.	-
Restoring	The EVS disk data is being restored using the backup.	-
Retyping	The EVS disk type is being changed.	-
Soft deleted	The EVS disk has been soft deleted and moved to the recycle bin.	<ul style="list-style-type: none"> ● Restore ● Permanently delete
Error	<p>An error occurs when you are creating an EVS disk.</p> <p>You can delete the EVS disk and create it again.</p>	Delete
Deletion failed	<p>An error occurs when you are deleting an EVS disk.</p> <p>Contact the administrator.</p>	None

EVS Disk Status	Status Description	Supported Operation
Expansion failed	An error occurs when you are expanding the capacity of an EVS disk. The administrator will contact you and help you handle this error. Do not perform any operations on the disk before the administrator contact you. If you require that the error be handled as soon as possible, contact the administrator.	-
Rollback failed	An error occurs during an EVS disk rollback.	Delete
Restoration failed	An error occurs during EVS disk restoration. The administrator will contact you and help you handle this error. Do not perform any operations on the disk before the administrator contact you. If you require that the error be handled as soon as possible, contact the administrator.	Delete
Maintenance	The EVS disk is being migrated.	-
Reserved	Status of the source disk after the migration task is complete	-

18.7.6 Does an EVS Disk or Snapshot Generate Metering Information in All States?

EVS disks and snapshots do not generate metering information in the following states:

- EVS disk state: Creating, Error, Restoration failed, Downloading, Expansion failed, Deleting, Deletion failed, Rollingback, Rollback failed, or Retyping
- Snapshot state: Creating, Error, Deleting, or Deletion failed

18.7.7 Can I Change the EVS Disk Capacity?

EVS disk capacity can be expanded but cannot not be reduced at present.

For details about how to expand the EVS disk capacity, see [Expanding EVS Disk Capacity](#).

18.7.8 Will Data in an EVS Disk Be Lost When the EVS Disk Is Detached?

The data will not be lost.

To prevent data loss, you are advised to perform the following operations:

- Windows operating system: Before the detachment, log in to the instance and perform the offline operation.
- Linux operating system: Before the detachment, log in to the instance and run the **umount** command to unmount the disk partition.

18.7.9 Device Type

Definition

Device types of EVS disks are divided based on whether advanced SCSI commands are supported. The device type can be VBD or SCSI.

- VBD: EVS disks of this type support only basic SCSI read and write commands. They are used in common scenarios, for example, OA, tests, Linux clusters such as RHCS.
- SCSI: EVS disks of this type support transparent SCSI command transmission and allow the ECS operating system to directly access the underlying storage media. SCSI EVS disks support advanced SCSI commands (such as SCSI-3 persistent pre-lock) in addition to basic SCSI read and write commands. They can be used in cluster scenarios where data security is ensured by using the SCSI lock mechanism, such as the Windows MSCS cluster.

NOTE

For details about ECS operating systems supported and ECS software required by SCSI EVS disks, see [Requirements and Restrictions on Using SCSI EVS Disks](#).

Requirements and Restrictions on Using SCSI EVS Disks

- If the VM HA function, storage plane anti-split-brain function, and data disk protection function are enabled, only non-SCSI disks can be locked to prevent data disks from being written in dual-write mode. (For details about how to check whether the functions are enabled, see "Configuring the VM HA Function" in *HUAWEI CLOUD Stack 8.1.0 O&M Guide*). Lock protection is not supported on SCSI data disks. Lock protection is implemented based on SCSI commands, but SCSI disks support transparent SCSI command transmission. As a result, lock protection on SCSI disks may conflict with user operations on SCSI disks, resulting in task failure.
- When SCSI disks are attached to an ECS, check whether the ECS supports SCSI disks based on the following description:
 - Windows OS
 - i. Check whether the ECS operating system supports the SCSI feature. Obtain the ECS operating system version by referring to *FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)* and check whether the SCSI (virtio-*scsi*) or raw device

mapping feature is supported. Obtain the document by referring to the following note.

- ii. The Windows operating system must have UVP VMTools installed to support SCSI.

Generally, the administrator has installed the UVP VMTools when creating a public image. You do not need to install it manually.

If the Windows operating system supports the SCSI feature and UVP VMTools has been installed in the operating system, you can attach SCSI disks to the ECS.

– Linux OS

The SCSI feature of the Linux operating system is not provided by the UVP VMTools but the driver in the operating system. Therefore, you only need to check whether the current ECS operating system supports the SCSI feature.

The check method is as follows: Check the ECS operating system version by referring to FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud) and check whether the SCSI (virtio-scsi) or raw device mapping feature is supported.

If the Linux operating system supports the SCSI feature, you can attach SCSI disks to the ECS. Obtain the document by referring to the following note.

 NOTE

Obtain *FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)*:

- x86
 - Carrier users: [Click here](#). Search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)**.
 - Enterprise users: [Click here](#). Search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)**.
- Arm
 - Carrier users: [Click here](#). Search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (ARM)**.
 - Enterprise users: [Click here](#). Search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (ARM)**.

18.7.10 Shared Disk

In the traditional cluster architecture, multiple computing nodes need to access the same data. This ensures that when a single or multiple computing nodes are faulty, the HA cluster can continue providing services, which means that a faulty component will not cause service interruption. Therefore, important data files need to be stored on shared block storage, and shared block storage is centrally managed using the cluster file system. When multiple frontend computing nodes access data, the data will be the same on the multiple computing nodes.

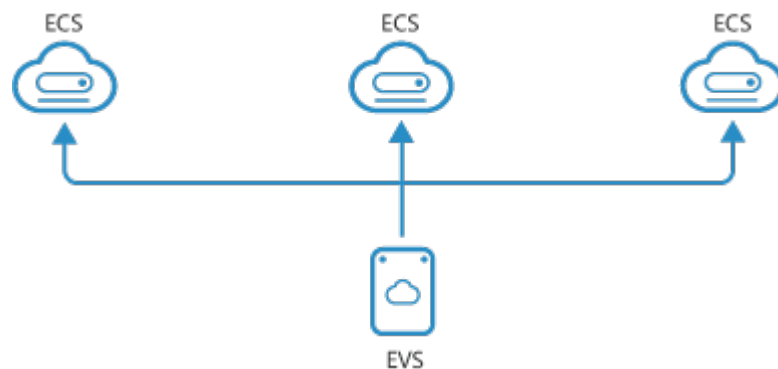
The shared disk is designed for the core service HA architecture of enterprise customers. The shared disk is suitable for scenarios that require shared block storage access in the share-everything architecture. The scenarios include the HA Oracle RAC database architecture for government, enterprise, and finance customers and the HA server cluster architecture.

Definition

Shared EVS disks are block storage devices that support concurrent read/write operations of multiple ECSs/BMSs. Shared EVS disks feature multiple attachments, high-concurrency, high-performance, and high-reliability. A shared EVS disk can be attached to a maximum of 16 ECSs/BMSs. A non-shared EVS disk can be attached to only one ECS/BMS. This document uses ECS as an example, as shown in [Figure 18-10](#).

Currently, shared EVS disks can be used as data disks only and cannot be used as system disks. Shared EVS disks of the VBD or SCSI type can be created.

Figure 18-10 Shared EVS disk



SCSI Reservation

Shared EVS disks of the VBD type do not support SCSI locks. SCSI shared EVS disks support SCSI reservation. If SCSI reservation is required for your applications, create SCSI shared EVS disks.

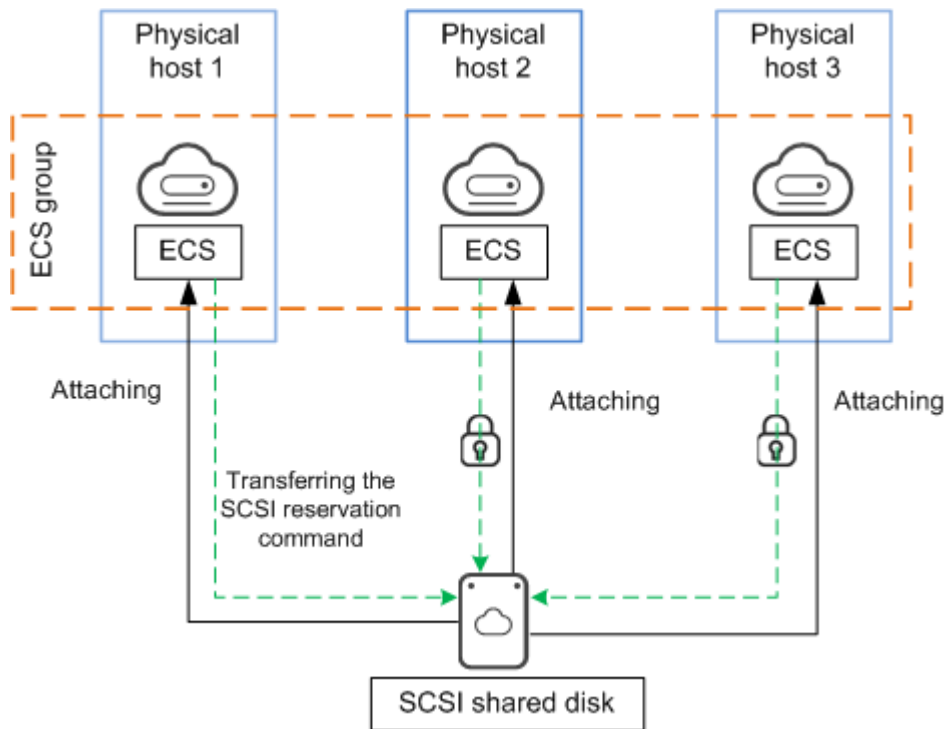
SCSI reservation is the basic mechanism for multiple hosts to use disks. In a shared storage environment, multiple service hosts may access a disk simultaneously. If multiple hosts perform the write operation on the disk at the same time, the disk does not know data from which host will be written first. To prevent this problem that may cause data damage, SCSI reservation is introduced.

SCSI reservation for an EVS disk cannot distinguish multiple ECSs on a single physical host, and SCSI reservation is supported only when ECSs are deployed on different physical hosts. If a SCSI shared EVS disk is attached to ECSs, use anti-affinity ECS groups with SCSI reservation so that SCSI reservation takes effect.

The ECSs in an anti-affinity ECS group will be created on different physical hosts, thereby improving service reliability. You can add an ECS to an ECS group only when creating the ECS. An existing ECS cannot be added to any ECS group.

[Figure 18-11](#) shows how SCSI reservation is implemented. When a SCSI shared disk is attached to multiple ECSs in an anti-affinity ECS group, if one of the ECSs sends a SCSI reservation command to the SCSI shared disk, the SCSI shared disk is locked for the other ECSs. In this case, the other ECSs cannot write data into the SCSI shared disk.

Figure 18-11 SCSI reservation implementation mechanism



NOTICE

If an ECS does not belong to any ECS group, the SCSI shared EVS disk should better not be attached to the ECS. Otherwise, SCSI reservations may not work properly, which puts your data at risk.

Precautions for Using the Shared EVS Disk

A shared EVS disk is essentially the disk that can be attached to multiple instances for use, which is similar to a physical disk in that the disk can be attached to multiple physical servers, and each server can read data from and write data into any space on the disk. If the data read and write rules, such as the read and write sequence and meaning, between these servers are not defined, data read and write interference between servers or other unpredictable errors may occur.

Shared EVS disks provide block storage devices whose data can be randomly read or written and allows shared access. Shared EVS disks do not provide the cluster file system. You need to install the cluster file system to manage shared EVS disks.

If a shared EVS disk is attached to multiple instances but is managed using a common file system, disk space allocation conflict will occur and data files will be inconsistent. The details are as follows:

- Disk space allocation conflict
Suppose that a shared EVS disk is attached to multiple instances. When a process on instance A writes files into the shared EVS disk, it checks the file system and available disk space. After files are written into the shared EVS disk, instance A will change its own space allocation records, but will not

change the space allocation records on the other instances. Therefore, when instance B attempts to write files to the shared EVS disk, it may allocate disk space addresses that have been allocated by instance A, resulting in disk space allocation conflict.

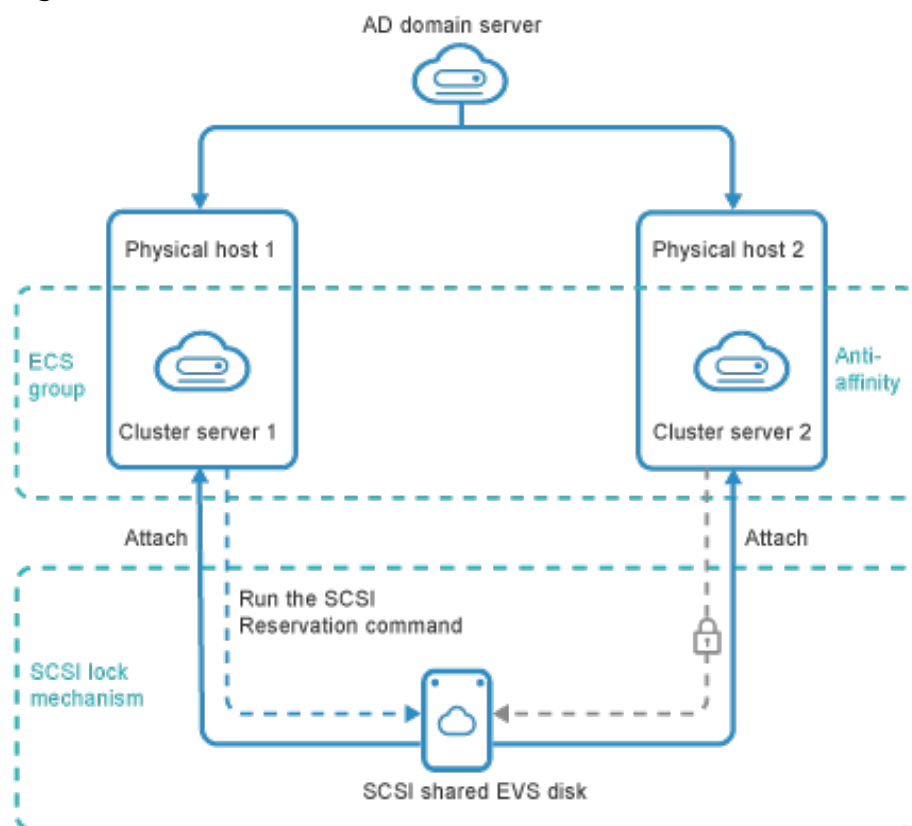
- Inconsistent data files
Suppose instance A reads data and records it in the cache. When another process on instance A accesses the same data, the process will read the data directly from the cache. If instance B changes the data, instance A will not know and will read the data from the cache. As a result, service data will be inconsistent on instance A and instance B.

The cluster management system is used to manage the shared EVS disks. If the cluster needs to use SCSI reservations, you need to apply for a shared EVS disk of SCSI type. Example enterprise applications include Windows MSCS (Microsoft Cluster Service) and Linux RHCS (Red Hat Cluster Suite). Windows MSCS and Linux RHCS are used as an example to describe how to use shared EVS disks.

- Windows MSCS

Figure 18-12 shows the Windows MSCS diagram. Multiple nodes in the cluster share the same storage. The cluster needs to use the SCSI reservation. Therefore, the shared EVS disk of the SCSI type is used. When a node in the cluster is faulty, services on the node are switched to another available node.

Figure 18-12 Windows MSCS

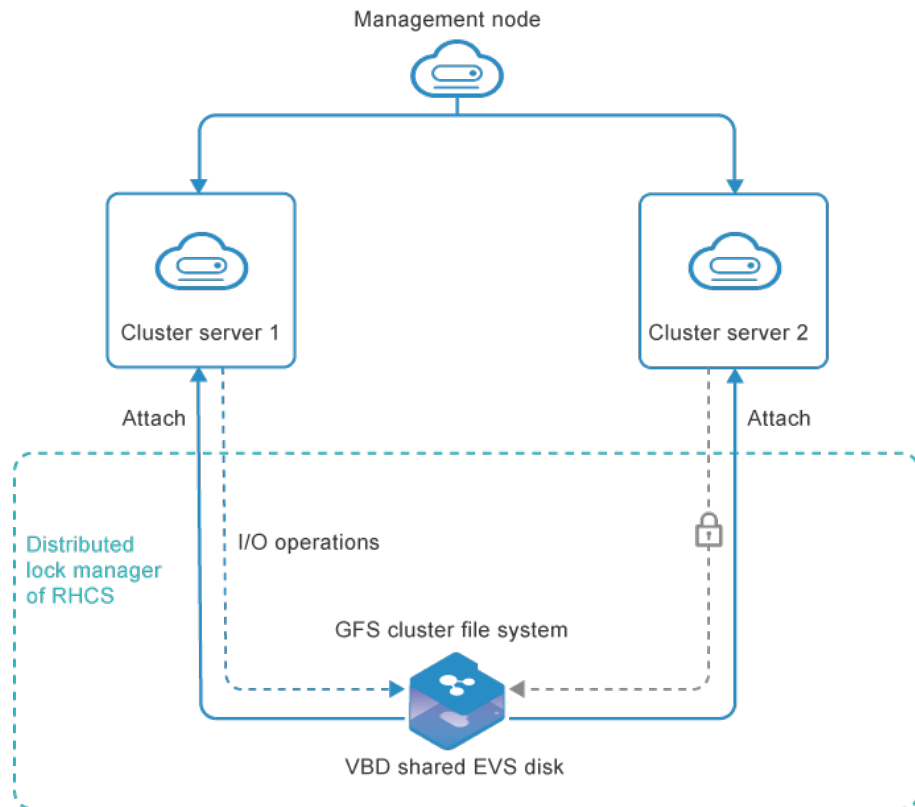


- Linux RHCS

Figure 18-13 shows the Linux RHCS diagram. Linux RHCS is a cluster suite that provides high availability, load balancing, and storage sharing. Linux RHCS provides a distributed lock manager. The GFS uses the lock mechanism

of the lock manager to ensure data consistency when multiple nodes share the same disk. Therefore, a shared EVS disk of VBD type without SCSI reservation can be used in the cluster.

Figure 18-13 Linux RHCS



18.7.11 Applying for a Snapshot

A snapshot can capture the data and status of a disk at a certain time point. If a service change or application software upgrade is required, you can create a snapshot for the disk in advance. If a fault occurs during the change or upgrade, you can use the snapshot to quickly restore disk data, ensuring service continuity and security. You can also use snapshots for routine backup of disk data.

Restrictions

- If backend storage is one of OceanStor V3/V5 or OceanStor Dorado V3 series, it is necessary for the administrator to import the HyperSnap license on the device in advance.
- Snapshots can be created only for disks in the **Available** or **In-use** state.
- A snapshot name cannot be the same as that of the prefix of the temporary snapshot created by the backup service (VBS/CSBS), the DR service (CSDR/CSHA/VHA), or the VM snapshot.
- Snapshots created using the EVS console consume the capacity quota instead of quantity quota of EVS disks.
- Temporary snapshots created by the backup service (VBS/CSBS) or the DR service (CSDR/CSHA/VHA) do not consume EVS disk quotas. Snapshots created using the VM snapshot function do not consume EVS disk quotas.

- Snapshots created using the EVS console, temporary snapshots created by DR and backup services, and snapshots created using the VM snapshot function consume backend storage capacity. If a large number of snapshots are created, contact the administrator to set the thin provisioning ratio of backend storage to a large value, preventing EVS disk provisioning failures caused by excessive snapshots.
- If backend storage of the disk is heterogeneous storage, snapshots can be created.
- In the VRM or VMware scenario, no snapshots can be created for shared EVS disks.
- If an EVS disk is created from data storage of the VIMS type in the VRM scenario and the EVS disk has been attached to an ECS, a snapshot can be created for the EVS disk only when the ECS is in the stopped state.
- In the VMware scenario, no snapshots can be created for an EVS disk that has been attached to an ECS.
- No snapshots can be created for disks that have expired.
- No snapshots can be created for disks that have been soft deleted.
- Snapshots cannot be created when the disk status is **Reserved** or **Maintenance**.
- If a task for creating a snapshot fails, the task is automatically deleted.

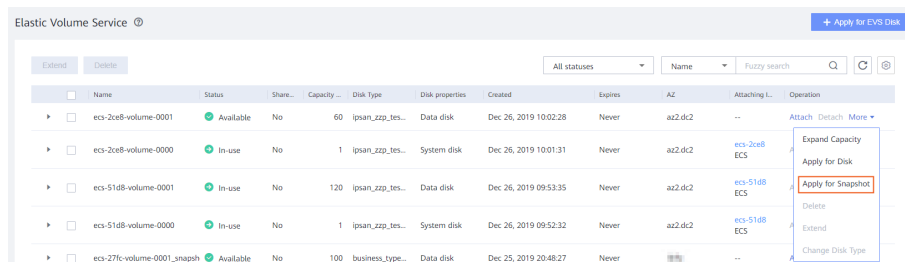
Procedure

Step 1 Log in to the EVS console as a VDC administrator or VDC operator. For details, see [18.7.1 Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 Use one of the following methods to display the **Apply for Snapshot** page.


- Method 1

In the EVS disk list, locate the row that contains the target disk, click **More** in the **Operation** column, and choose **Apply for Snapshot**.



- Method 2

In the navigation pane on the left, choose **Elastic Volume Service > Snapshot**. Then, click **Apply for Snapshot**.



Step 3 Configure the basic information about the snapshot, as shown in [Table 18-4](#).

Table 18-4 Parameters for creating a snapshot

Parameter	Description	Example Value
Name	The value contains only digits, letters, underscores (_), and hyphens (-) and cannot exceed 63 characters. NOTE The snapshot name cannot be the same as the name prefix of any temporary snapshot created by the backup service (VBS or CSBS) or the DR service (CSDR, CSHA, or VHA) or the name prefix of any VM snapshot. The prefixes include autobk_snapshot , manualbk_snapshot , and sys_snapshot .	snapshot-01
EVS Disk	<ul style="list-style-type: none">When you create a snapshot for the specified disk, the disk has been determined.When you create a snapshot on the snapshot page, click Select, select the target disk, and click OK.	volume-01
Description	Description of the created snapshot. The value contains a maximum of 63 bytes.	-

Step 4 Click **Next**.

Step 5 If you do not need to modify the specifications, click **Apply Now** to start the snapshot application.

If you need to modify the specifications, click **Back** to modify parameter settings.

Step 6 If creating a snapshot requires approval, contact the administrator for approval. If no, skip this step.

Return to the **Snapshot** page. If the snapshot status is **Available**, the snapshot is created successfully.

----End

18.7.12 Creating a Backup

To prevent data loss caused by misoperations or system faults, you can create a backup for an EVS disk to ensure data correctness and security.

You can create backups for EVS disks in the following ways:

- EVS console: To create backups for a single EVS disk, perform related operations on the EVS console.
- VBS console: To create backups for multiple EVS disks in batches, perform related operations on the VBS console.

Restrictions

- Only disks in the **Available** or **In-use** state can be backed up.

- If backend storage of the disk is heterogeneous storage, backups cannot be created.

EVS Console

Before creating a backup on the EVS console, ensure that the VBS service has been deployed in the environment and enabled by the administrator. Otherwise, this function is unavailable. For details about how to enable the VBS service, see "Storage Services" > "Elastic Volume Service (EVS for ECS)" > "FAQs" > "How Do I Enable the VBS?" in *HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide*.

- Step 1** Log in to the EVS console as a VDC administrator or VDC operator. For details, see [18.7.1 Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).
- Step 2** On the **Elastic Volume Service** page, locate the row that contains the target EVS disk, click **More** in the **Operation** column, and then choose **Create Backup**.
- Step 3** On the **Create VBS Backup** page, configure parameters as prompted.

For details, see **Operation Help Center > DR & Backup > User Guide > Creating a Periodic Backup Task > Creating a Backup Task**.

----End

VBS Console

For details about how to create a backup task on the VBS page, see **Operation Help Center > DR & Backup > User Guide > Creating a Periodic Backup Task > Creating a Backup Task**.

Follow-up Procedure

After an EVS disk is backed up successfully, a VBS backup of the disk is generated. You can view the details about the backup on **Operation Help Center > DR & Backup > User Guide > Managing Backups and Replicas > Viewing EVS Disk Backups and Replicas**. You can also view backup tree information on the **Elastic Volume Service** page.

18.8 OS FAQs

18.8.1 Can I Install or Upgrade the OS by Myself?

No. If you need to install or upgrade the OS, use the **Reinstall OS** or **Change OS** function to perform the operation.

18.8.2 Can the OS of an ECS Be Changed?

Yes. An ECS OS can be changed.

For instructions about how to change an ECS OS, see section [13.2 Changing the ECS OS](#).

18.8.3 Can I Select Other OSs 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 section [13.2 Changing the ECS OS](#).

18.8.4 How Can I Obtain Data Disk Information If Tools Is Deleted?

If Tools is uninstalled from a Linux ECS in a non-PVOPS system, data disks cannot be identified. In this case, you can obtain information about these data disks by creating a new ECS and attaching the original data disks to the original ECS to the new ECS. The procedure is as follows:

Create a new ECS.

For details, see [5.2 Applying for an ECS](#).

NOTE

The new ECS must belong to the same AZ and have the same parameter settings as the original ECS.

Step 1 (Optional) In the ECS list, locate the row containing the original ECS, choose **More > Change Status > Stop** in the **Operation** column. On the **Stop ECS** page, select **Forcibly stop** and **Yes** and click **OK** to forcibly stop the original ECS.

Manually refresh the ECS list. The original ECS is stopped once the **Status** value changes to **Stopped**.

NOTE

The ECSs running certain OSs support online data disk detaching. If your OS supports this feature, you can detach data disks from running ECSs.

For details how to detach data disks from running ECSs, see [10.5 Releasing an EVS Disk](#).

Step 2 Click the name of the original ECS. On the ECS details page, click the **EVS** tab to view information about the data disk attached to the ECS. You can click the data disk ID to go to the **EVS** page.

NOTE

If multiple data disks are attached to the original ECS and information about all data disks needs to be obtained, repeat [Step 3](#) to [Step 4](#) to detach and attach each data disk.

Step 3 Locate the data disk to be detached, click **Detach**, and click **OK** to detach the data disk from the original ECS.

Manually refresh the EVS list. The data disk is successfully detached from the original ECS once the **Status** value changes to **Available**.

Step 4 Locate the row that contains the detached data disk, click **Attach**, select the newly created ECS, and click **OK** to attach the target data disk to the newly created ECS.

Manually refresh the EVS list. The data disk is successfully attached to the new ECS once the **Status** value changes to **In-use**. After an EVS disk is successfully attached to a new ECS, you can log in to ManageOne to view the latest disk information.

----End

18.8.5 What Should I Do If the One-Click Password Reset Plugin Fails to Start?

For ECSs running the following OSs, if the installed one-click password reset plugin cannot automatically start upon ECS startup, perform the operations described in this section. Then, perform operations in [Follow-up Procedure](#) to check whether the one-click password reset plugin can automatically start upon ECS startup.

Table 18-5 Enabling the plugin to automatically start upon OS startup in x86 scenarios

OS	Procedure
CoreOS	For CoreOS, see CoreOS .
SUSE	<ul style="list-style-type: none">For SUSE 12, see SUSE 12.For other SUSE OSs (except SUSE 11 and 12), see SUSE (Except SUSE 11 and SUSE 12), Ubuntu, or Debian (Except Debian 8). <p>NOTE For SUSE 11, this problem will not occur.</p>
Red Hat 7, Oracle 7, CentOS 7, and CentOS 8.0	For Red Hat 7, Oracle 7, CentOS 7, and CentOS 8.0, see Red Hat 7, Oracle 7, CentOS 7, and CentOS 8.0 .
OpenSUSE	For openSUSE 13, see openSUSE 13 . For openSUSE Leap 15.1 64bit, see OpenSUSE 15.1 .
Debian	<ul style="list-style-type: none">For Debian 8, see Debian 8.For other Debian OSs, see SUSE (Except SUSE 11 and SUSE 12), Ubuntu, or Debian (Except Debian 8).
Ubuntu	For SUSE (except SUSE 11 and 12), Ubuntu, or Debian (except Debian 8), see SUSE (Except SUSE 11 and SUSE 12) , Ubuntu , or Debian (Except Debian 8) .
Fedora	<ul style="list-style-type: none">For Fedora 24 64bit and Fedora 25 64bit, see Fedora 24 64bit and Fedora 25 64bit.For Fedora Server 29/31 64bit, see Fedora 29/31 and Euler 2.8/2.9.For Fedora Server 30 64bit, see Fedora 30.
Other OSs	For other OSs, see Other OSs .

Table 18-6 Enabling the plugin to automatically start upon OS startup in Arm scenarios

OS	Procedure
SUSE	For SUSE Linux Enterprise Server 15 64bit, see SUSE Linux Enterprise Server 15 64-bit and openSUSE Leap 15.0 64-bit .
OpenSUSE	For openSUSE Leap 15.0 64bit, see SUSE Linux Enterprise Server 15 64-bit and openSUSE Leap 15.0 64-bit .
China Standard Software	<ul style="list-style-type: none"> • NeoKylin Server release 5.0 U2 64bit • NeoKylin Linux Advanced Server release 7.0 U5 64bit • NeoKylin Linux Advanced Server release 7.0 U6 64bit • NeoKylin Linux Desktop 7.0 U5 64bit For NeoKylin, see NeoKylin OSs .
Fedora	For Fedora 29, see Fedora 29/31 and Euler 2.8/2.9 .
Euler	For Euler 2.8 and Euler 2.9, see Fedora 29/31 and Euler 2.8/2.9 .
Deepin	For Deepin GNU/Linux 15.5, see Deepin GNU/Linux 15.5 .
Cent OS	For CentOS 8.0 64bit, see CentOS 8.0 (ARM) .
uos	<ul style="list-style-type: none"> • uos V20 server • uos V20 desktop (64-bit) • uos 20 server For details, see uos V20 server .
Kirin	<ul style="list-style-type: none"> • Kylin V10 SP1 Desktop (64-bit) • Kylin V10 SP1 Server (64-bit) • Kylin V10 Desktop (64-bit) For details, see Kylin V10 .
OpenEuler	For OpenEuler 20.03, see OpenEuler 20.03 .

CoreOS

Run the following commands to start the plugin:

```
cat >/etc/systemd/system/cloudResetPwdAgent.service <<EOT
```

```
[Unit]
```

```
Description=cloudResetPwdAgent service
```

```
Wants=local-fs.target
```

```
Requires=local-fs.target
[Service]
Type=simple
ExecStart=/CloudrResetPwdAgent/bin/cloudResetPwdAgent.script start
RemainAfterExit=yes
ExecStop=/CloudrResetPwdAgent/bin/cloudResetPwdAgent.script stop
KillMode=none
[Install]
WantedBy=multi-user.target
EOT
systemctl enable cloudResetPwdAgent.service
cat >/etc/systemd/system/cloudResetPwdUpdateAgent.service <<EOT
[Unit]
Description=cloudResetPwdUpdateAgent service
Wants=local-fs.target
Requires=local-fs.target
[Service]
Type=simple
ExecStart=/CloudResetPwdUpdateAgent/bin/
cloudResetPwdUpdateAgent.script start
RemainAfterExit=yes
ExecStop=/CloudResetPwdUpdateAgent/bin/
cloudResetPwdUpdateAgent.script stop
KillMode=none
[Install]
WantedBy=multi-user.target
EOT
systemctl enable cloudResetPwdUpdateAgent.service
```

SUSE 12

Step 1 Run the following command to open the **boot.local** file:

```
vi /etc/init.d/boot.local
```

Step 2 Press **i** to enter the editing mode, and add the following content to the end of the file:

```
/CloudrResetPwdAgent/bin/cloudResetPwdAgent.script start
```

```
/CloudResetPwdUpdateAgent/bin/cloudResetPwdUpdateAgent.script start
```

Step 3 Press **Esc**, enter **:wq**, and press **Enter** to save the configuration and exit.

----End

Red Hat 7, Oracle 7, CentOS 7, and CentOS 8.0

Step 1 Run the following commands to open the **rc.local** file:

```
chmod +x /etc/rc.d/rc.local
```

```
vi /etc/rc.d/rc.local
```

Step 2 Press **i** to enter the editing mode, and add the following content to the end of the file:

```
/CloudrResetPwdAgent/bin/cloudResetPwdAgent.script start
```

```
/CloudResetPwdUpdateAgent/bin/cloudResetPwdUpdateAgent.script start
```

Step 3 Press **Esc**, enter **:wq**, and press **Enter** to save the configuration and exit.

----End

openSUSE 13

Step 1 Run the following command to open the **boot.local** file:

```
vi /etc/init.d/boot.local
```

Step 2 Press **i** to enter the editing mode, and add the following content to the end of the file:

```
/CloudrResetPwdAgent/bin/cloudResetPwdAgent.script start
```

```
/CloudResetPwdUpdateAgent/bin/cloudResetPwdUpdateAgent.script start
```

Step 3 Press **Esc**, enter **:wq**, and press **Enter** to save the configuration and exit.

----End

OpenSUSE 15.1

Step 1 Run the **chkconfig** command to check the status of the one-click password reset plugin.

- If the status is **on**, no further action is required.
- If the status is **off**, perform [Step 2](#).

Step 2 Run the following commands:

```
/sbin/chkconfig cloudResetPwdUpdateAgent on
```

```
/sbin/chkconfig cloudResetPwdAgent on
```

----End

SUSE Linux Enterprise Server 15 64-bit and openSUSE Leap 15.0 64-bit

Step 1 Log in to an ECS as the **root** user and run the following command to check whether insserv is installed:

```
ls /sbin/insserv
```

- If the following information is displayed, insserv has been installed. In this case, perform [Step 2](#).

```
host-192-168-0-200:~ # ls /sbin/insserv
/sbin/insserv
```

- If the following information is displayed, insserv is not installed.

```
host-192-168-0-200:~ # ls /sbin/insserv
ls: cannot access '/sbin/insserv': No such file or directory
host-192-168-0-200:~ #
```

Run the following command to install insserv:

```
zypper install insserv-compat
```

Step 2 Run the following commands to enable the plugin to automatically start upon ECS startup:

```
/sbin/chkconfig cloudResetPwdUpdateAgent on
```

```
/sbin/chkconfig cloudResetPwdAgent on
```

```
----End
```

Debian 8

Step 1 Run the following command to open the **rc.local** file:

```
vi /etc/rc.local
```

Step 2 Press **i** to enter the editing mode, and add the following content to the end of the file:

```
/CloudrResetPwdAgent/bin/cloudResetPwdAgent.script start
```

```
/CloudResetPwdUpdateAgent/bin/cloudResetPwdUpdateAgent.script start
```

Step 3 Press **Esc**, enter **:wq**, and press **Enter** to save the configuration and exit.

```
----End
```

SUSE (Except SUSE 11 and SUSE 12), Ubuntu, or Debian (Except Debian 8)

Step 1 Run the following command to open the **rc** file:

```
vi /etc/init.d/rc
```

Step 2 Press **i** to enter the editing mode, and add the following content to the end of the file:

```
/CloudrResetPwdAgent/bin/cloudResetPwdAgent.script start
```

```
/CloudResetPwdUpdateAgent/bin/cloudResetPwdUpdateAgent.script start
```

Step 3 Press **Esc**, enter **:wq**, and press **Enter** to save the configuration and exit.

----End

Ubuntu 18.04.1 (ARM)

Step 1 Log in to the ECS and run the following command as the **root** user to create a **lib64** directory:

```
mkdir -p /lib64
```

Step 2 Run the following command to copy the file to the **lib64** directory:

```
cp /lib/ld-linux-aarch64.so.1 /lib64
```

Step 3 Run the following command to open the **rc** file:

```
vi /etc/init.d/rc
```

Step 4 Press **i** to enter the editing mode, and add the following content to the end of the file:

```
/CloudrResetPwdAgent/bin/cloudResetPwdAgent.script start
```

```
/CloudResetPwdUpdateAgent/bin/cloudResetPwdUpdateAgent.script start
```

Step 5 Press **Esc**, enter **:wq**, and press **Enter** to save the configuration and exit.

----End

Fedora 24 64bit and Fedora 25 64bit

Step 1 Run the following commands to open the **rc.local** file:

```
touch /etc/rc.d/rc.local
```

```
chmod +x /etc/rc.d/rc.local
```

```
vi /etc/rc.d/rc.local
```

Step 2 Press **i** to enter the editing mode, and add the following content to the end of the file:

```
/CloudrResetPwdAgent/bin/cloudResetPwdAgent.script start
```

```
/CloudResetPwdUpdateAgent/bin/cloudResetPwdUpdateAgent.script start
```

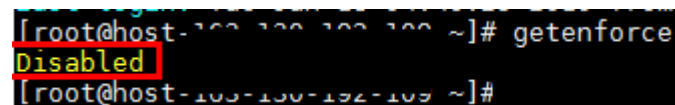
Step 3 Press **Esc**, enter **:wq**, and press **Enter** to save the configuration and exit.

----End

Fedora 29/31 and Euler 2.8/2.9

Log in to the ECS and run the following command as the **root** user to check whether **selinux** is enabled:

```
getenforce
```



```
[root@host-100-100-100-100 ~]# getenforce  
Disabled  
[root@host-100-100-100-100 ~]#
```

- If the returned value is **Disabled**, selinux is disabled. In this case, no further action is required.
- If the returned value is **Enforcing** or **Permissive**, run the following commands:

```
chcon -u system_u -r object_r -t initrc_exec_t /CloudrResetPwdAgent/bin/  
cloudResetPwdAgent.script  
  
chcon -u system_u -r object_r -t initrc_exec_t /  
CloudResetPwdUpdateAgent/bin/cloudResetPwdUpdateAgent.script
```

Fedora 30

Step 1 Run the following commands to open the **rc.local** file:

```
touch /etc/rc.d/rc.local  
  
chmod +x /etc/rc.d/rc.local  
  
vi /etc/rc.d/rc.local
```

Step 2 Press **i** to enter the editing mode, and add the following content to the end of the file:

```
#!/bin/bash  
  
/CloudrResetPwdAgent/bin/cloudResetPwdAgent.script start  
  
/CloudResetPwdUpdateAgent/bin/cloudResetPwdUpdateAgent.script start
```

Step 3 Run the following command:

```
chcon -u system_u -r object_r -t initrc_exec_t /etc/rc.d/rc.local  
  
----End
```

NeoKylin OSs

Log in to the ECS and run the following commands as the **root** user to configure the one-click password reset plugin upon system start:

```
/sbin/chkconfig cloudResetPwdUpdateAgent on  
  
/sbin/chkconfig cloudResetPwdAgent on
```

Deepin GNU/Linux 15.5

Step 1 Log in to the ECS and run the following command as the **root** user to create a lib64 directory:

```
mkdir -p /lib64
```

Step 2 Run the following command to copy the file to the lib64 directory:

```
cp /lib/ld-linux-aarch64.so.1 /lib64  
  
----End
```


Other OSs

Step 1 Run the following command to open the **rc.local** file:

```
vi /etc/rc.d/rc
```

Step 2 Press **i** to enter the editing mode, and add the following content to the end of the file:

```
/CloudrResetPwdAgent/bin/cloudResetPwdAgent.script start
```

```
/CloudResetPwdUpdateAgent/bin/cloudResetPwdUpdateAgent.script start
```

Step 3 Press **Esc**, enter **:wq**, and press **Enter** to save the configuration and exit.

----End

CentOS 8.0 (ARM)

Step 1 Log in to the ECS as the **root** user.

Step 2 Change the value of **SELINUX** in the **/etc/selinux/config** file to **disabled**.

```
SELINUX=disabled
```

Step 3 Restart an ECS.

Step 4 Run the following command. If **Disabled** is displayed, the configuration is successful.

```
getenforce
```

----End

uos V20 server

Step 1 Log in to the ECS and run the following command as the **root** user to create a **lib64** directory:

```
mkdir -p /lib64
```

Step 2 Run the following command to copy the file to the **lib64** directory:

```
cp /lib/ld-linux-aarch64.so.1 /lib64
```

Step 3 Run the following command to open the **rc** file:

```
vi /etc/init.d/rc
```

Step 4 Press **i** to enter the editing mode, and add the following content to the end of the file:

```
/CloudrResetPwdAgent/bin/cloudResetPwdAgent.script start
```

```
/CloudResetPwdUpdateAgent/bin/cloudResetPwdUpdateAgent.script start
```

Step 5 Press **Esc**, enter **:wq**, and press **Enter** to save the configuration and exit.

----End

OpenEuler 20.03

Log in to the ECS and run the following command as user **root** to check whether selinux is enabled:

getenforce

```
[root@host-192-168-192-192 ~]# getenforce  
Disabled  
[root@host-192-168-192-192 ~]#
```

- If **Disabled** is returned, selinux is disabled. In this case, no further action is required.
- If the returned value is **Enforcing** or **Permissive**, run the following commands:
chcon -u system_u -r object_r -t initrc_exec_t /CloudrResetPwdAgent/bin/cloudResetPwdAgent.script
chcon -u system_u -r object_r -t initrc_exec_t /CloudResetPwdUpdateAgent/bin/cloudResetPwdUpdateAgent.script

Kylin V10

Step 1 Log in to the ECS and run the following command as user **root** to create the **lib64** directory:

```
mkdir -p /lib64
```

Step 2 Run the following command to copy the file to the **lib64** directory:

```
cp /lib/ld-linux-aarch64.so.1 /lib64
```

----End

Follow-up Procedure

Windows ECSs

Step 1 Restart the ECS.

Step 2 Press **Ctrl+Shift+Esc**. In the task manager, check whether there are **cloudResetPwdAgent** and **cloudResetPwdUpdateAgent**.

- If yes, the installation is successful. Go to [Step 3](#).
- If no, the installation failed. Contact Huawei technical support.

Step 3 Run the following command to check whether the IP address can be pinged:

```
ping 169.254.169.254
```

- If yes, the ECS NIC is working properly. No further action is required.
- If no, contact technical support.

----End

Linux ECSs

Step 1 Restart the ECS.

Step 2 Run the following commands and check whether the statuses of **CloudResetPwdAgent** and **CloudResetPwdUpdateAgent** are **unrecognized service**.

service cloudResetPwdAgent status

service cloudResetPwdUpdateAgent status

- If yes, the plugin cannot automatically start upon ECS startup. Contact technical support.
- If no, the plugin automatically starts upon ECS startup. Go to [Step 3](#).

Step 3 Run the following command to check whether the IP address can be pinged:

ping 169.254.169.254

- If yes, the ECS NIC is working properly. No further action is required.
- If no, contact technical support.

----End

18.8.6 What Can I Do If the OS Reinstallation or Change Fails?

Procedure

Step 1 [18.1.1 How Do I Log In to ManageOne Operation or Tenant Portal?](#)

Step 2 On the ECS page, click the name of the instance whose OS fails to be changed or reinstalled to switch to the ECS instance details page.

Step 3 Obtain the instance ID from the ECS instance details page.

Step 4 Use PuTTY to log in to the FusionSphere OpenStack controller node using the IP address of the External OM plane.

Default username: **fsp**

See the default password of the account for logging in to the target node on the "Type A (Background)" sheet in *HUAWEI CLOUD Stack 8.1.0 Account List*.

NOTE

To obtain the IP address of the External OM plane, perform the following steps: Obtain the parameter information summary file *xxx_export_all_v2_EN.xlsx* exported from HUAWEI CLOUD Stack Deploy during software installation and search for **Reverse-Proxy** on the "Tool_generated_IP_Params" sheet.

Step 5 Run the following command to switch to the **root** user:

su - root

See the default password of the account for logging in to the target node on the "Type A (Background)" sheet in *HUAWEI CLOUD Stack 8.1.0 Account List*.

Step 6 Run the following command to disable user logout upon system timeout:

TMOUT=0

Step 7 Run the following command to import environment variables:

source set_env

Select **1** and enter the password as prompted.

For the default password, see the default password of the **dc_admin** account on the "Type B (FusionSphere OpenStack)" sheet of [HUAWEI CLOUD Stack 8.1.0 Account List](#).

Step 8 Run the following command to delete the VM intermediate state label:

```
nova meta ceacdafe-f179-4b77-b296-c505241b1111 delete op_svc_lockaction
```

In the preceding command, *ceacdafe-f179-4b77-b296-c505241b1111* is the value of **ID** obtained in [Step 3](#).

Information similar to the following is displayed:

```
A5597B6D-1624-2E8F-E911-E310A8EDB70F:/home/fsp # nova meta ceacdafe-f179-4b77-b296-c505241b1111 delete op_svc_lockaction
A5597B6D-1624-2E8F-E911-E310A8EDB70F:/home/fsp #
```

Step 9 Run the following command to query the original volume ID (**oldVolumeld**) of the VM:


```
nova show ceacdafe-f179-4b77-b296-c505241b1111 | grep oldVolumeld
```

In the preceding command, *ceacdafe-f179-4b77-b296-c505241b1111* is the value of **ID** obtained in [Step 3](#).

Information similar to the following is displayed. Record the **oldVolumeld** value in the command output.

```
A5597B6D-1624-2E8F-E911-E310A8EDB70F:/home/fsp # nova show ceacdafe-f179-4b77-b296-c505241b1111 | grep oldVolumeld
| metadata | {"cascaded.instance_extrainfo":
"stopped_release_resource:True,system_serial_number:ceacdafe-f179-4b77-b296-c505241b1111,max_mem:
4194304,max_cpu:254,cpu_num_for_one_plug:1,org_cpu:2,current_mem:
2048,xml_support_live_resize:True,pcibridge:2,org_mem:2048,iohang_timeout:720,current_cpu:
2,uefi_mode_sysinfo_fields:version_serial_uuid_family_asset,num_of_mem_plug:0", "metering.image_id":
"e23c54c2-93d5-4c9c-b66f-488043bbcb6a", "metering.imagetype": "gold", "metering.resourcetype": "1",
"metering.resourcetypecode": "wbl_2u2g.linux", "_ha_policy_type": "remote_rebuild", "server_expiry": "0",
"image_name": "cloudinit_image", "__instance_vwatchdog": "false", "metering.cloudServiceType":
"sys.service.type.ec2", "os_bit": "64", "vpc_id": "c61182cb-ed68-4f26-b635-ed6e98856bd1", "os_type":
"Linux", "charging_mode": "0", "oldVolumeld": "bcc8bfc1-62af-45a6-b96b-02952e3dc03c", "productId":
"fe74e736afde437cbc68dcf045584ff3"} |
```

Step 10 Attach an EVS disk to the ECS.

1. Perform [Step 1](#) to [Step 2](#). Click the **EVS** tab and click **Attach Disk**.
2. Copy the original volume ID (**oldVolumeld** value) obtained in [Step 9](#) to the right of **ID** and click .
3. Select **System Disk** for **Disk Properties** and click **OK**.

If the information about the attached disk is displayed on the **EVS** tab, the system disk is successfully attached. In this case, you can reinstall or change the OS.

----End

A Supported vGPU Types

Tesla T4 vGPU Types

Q-Series vGPU Types for Tesla T4

Required license edition: vDWS

The maximum number of available combined resolution pixels supported by these vGPU types is determined by their frame buffer size. You can choose to use a small number of high-resolution displays or a larger number of low-resolution displays with these vGPU types. The maximum number of displays per vGPU is based on a configuration in which all displays have the same resolution.

Table A-1 Q-series vGPU types for Tesla T4

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
T4-16Q	Virtual desktops	16384	1	66355200	7680 × 4320	2
					5120 × 2880 or lower	4
T4-8Q	Virtual desktops	8192	2	66355200	7680 × 4320	2
					5120 × 2880 or lower	4
T4-4Q	Virtual desktops	4096	4	58982400	7680 × 4320	1
					5120 × 2880 or lower	4

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
T4-2Q	Virtual desktops	2048	8	35389440	7680 × 4320	1
					5120 × 2880	2
					4096 × 2160 or lower	4
T4-1Q	Virtual desktops and virtual workstations	1024	16	17694720	5120 × 2880	1
					4096 × 2160	2
					3840 × 2160	2
					2560 × 1600 or lower	4

B-Series vGPU Types for Tesla T4

Required license edition: vPC or vDWS

The maximum number of available combined resolution pixels supported by these vGPU types is determined by their frame buffer size. You can choose to use a small number of high-resolution displays or a larger number of low-resolution displays with these vGPU types. The maximum number of displays per vGPU is based on a configuration in which all displays have the same resolution.

Table A-2 B-series vGPU types for Tesla T4

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
T4-2B	Virtual desktops	2048	8	17694720	5120 × 2880	1
					4096 × 2160	2
					3840 × 2160	2

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
					2560 × 1600 or lower	4
T4-2B4	Virtual desktops	2048	8	17694720	5120 × 2880	1
					4096 × 2160	2
					3840 × 2160	2
					2560 × 1600 or lower	4
T4-1B	Virtual desktops	1024	16	16384000	5120 × 2880	1
					4096 × 2160	1
					3840 × 2160	1
					2560 × 1600 or lower	4
T4-1B4	Virtual desktops	1024	16	16384000	5120 × 2880	1
					4096 × 2160	1
					3840 × 2160	1
					2560 × 1600 or lower	4

NOTE

- -1B and -1B4 vGPUs perform adequately when each vGPU has only two 2560 x 1600 virtual displays. If you want to use more than two 2560 x 1600 virtual displays per vGPU, use a vGPU with more frame buffers, such as a -2B or -2B4 vGPU.
- In this version, -1B4 and -2B4 vGPUs are not recommended because they may be removed in a future release. Use the following vGPU types, which provide equivalent functionality:
 - Replace -1B4 vGPU types with -1B vGPU types.
 - Replace -2B4 vGPU types with -2B vGPU types.

C-Series vGPU Types for Tesla T4

Required license edition: Virtual Compute Server (vCS) or vDWS

These vGPU types support a single display with a fixed maximum resolution.

Table A-3 C-series vGPU types for Tesla T4

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Maximum Display Resolution	Virtual Displays per vGPU
T4-16C	Training workloads	16384	1	4096 x 2160	1
T4-8C	Training workloads	8192	2	4096 x 2160	1
T4-4C	Training workloads	4096	4	4096 x 2160	1

NOTE

C-series vGPU types are NVIDIA vCS vGPU types, which are optimized for compute-intensive workloads. Therefore, C-series vGPUs support only a single display head and do not support Quadro graphics acceleration.

A-Series vGPU Types for Tesla T4

Required license edition: vApps

These vGPU types support a single display with a fixed maximum resolution.

Table A-4 A-series vGPU types for Tesla T4

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Maximum Display Resolution	Virtual Displays per vGPU
T4-16A	Virtual applications	16384	1	1280 × 1024	1
T4-8A	Virtual applications	8192	2	1280 × 1024	1

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Maximum Display Resolution	Virtual Displays per vGPU
T4-4A	Virtual applications	4096	4	1280 × 1024	1
T4-2A	Virtual applications	2048	8	1280 × 1024	1
T4-1A	Virtual applications	1024	16	1280 × 1024	1

 **NOTE**

A-series NVIDIA vGPUs support a single display with low resolution to be used as the console display for remote applications such as RDSH and Citrix Virtual Apps and Desktops. The maximum resolution and number of virtual display heads of A-series NVIDIA vGPUs apply only to the console display. The maximum resolution of each RDSH or Citrix Virtual Apps and Desktops session is determined by the remote solution and is not restricted by the maximum resolution of the vGPU. Similarly, the number of virtual display heads supported by each session is determined by the remote solution and is not restricted by the vGPU.

Quadro RTX 6000 vGPU Types

Q-Series vGPU Types for Quadro RTX 6000

Required license edition: vDWS

The maximum number of available combined resolution pixels supported by these vGPU types is determined by their frame buffer size. You can choose to use a small number of high-resolution displays or a larger number of low-resolution displays with these vGPU types. The maximum number of displays per vGPU is based on a configuration in which all displays have the same resolution.

Table A-5 Q-series vGPU types for Quadro RTX 6000

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
RTX6000-24Q	Virtual workstations	24576	1	6635520	7680 × 4320	2
					5120 × 2880 or lower	4
RTX6000-12Q	Virtual workstations	12288	2	6635520	7680 × 4320	2

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
					5120 × 2880 or lower	4
RTX6000-8Q	Virtual workstations	8192	3	66355200	7680 × 4320	2
					5120 × 2880 or lower	4
RTX6000-6Q	Virtual workstations	6144	4	58982400	7680 × 4320	1
					5120 × 2880 or lower	4
RTX6000-4Q	Virtual workstations	4096	6	58982400	7680 × 4320	1
					5120 × 2880 or lower	4
RTX6000-3Q	Virtual workstations	3072	8	35389440	7680 × 4320	1
					5120 × 2880	2
					4096 × 2160 or lower	4
RTX6000-2Q	Virtual workstations	2048	12	35389440	7680 × 4320	1
					5120 × 2880	2
					4096 × 2160 or lower	4
RTX6000-1Q	Virtual workstations	1024	24	17694720	5120 × 2880	1
					4096 × 2160	2

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
					3840 × 2160	2
					2560 × 1600 or lower	4

B-Series vGPU Types for Quadro RTX 6000

Required license edition: vPC or vDWS

The maximum number of available combined resolution pixels supported by these vGPU types is determined by their frame buffer size. You can choose to use a small number of high-resolution displays or a larger number of low-resolution displays with these vGPU types. The maximum number of displays per vGPU is based on a configuration in which all displays have the same resolution.

Table A-6 B-series vGPU types for Quadro RTX 6000

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
RTX6000-2B	Virtual desktops	2048	12	17694720	5120 x 2880	1
					4096 x 2160	2
					3840 × 2160	2
					2560 × 1600 or lower	4
RTX6000-1B	Virtual desktops	1024	24	16384000	5120 x 2880	1
					4096 x 2160	1
					3840 × 2160	1
					2560 × 1600 or lower	4

 **NOTE**

-1B vGPUs perform adequately when each vGPU has only two 2560 x 1600 virtual displays. If you want to use more than two 2560 x 1600 virtual displays per vGPU, use a vGPU with more frame buffers, such as a -2B vGPU.

C-Series vGPU Types for Quadro RTX 6000

Required license edition: vCS or vDWS

These vGPU types support a single display with a fixed maximum resolution.

Table A-7 C-series vGPU types for Quadro RTX 6000

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Maximum Display Resolution	Virtual Displays per vGPU
RTX6000-2 4C	Training workloads	24576	1	4096 x 2160	1
RTX6000-1 2C	Training workloads	12288	2	4096 x 2160	1
RTX6000-8 C	Training workloads	8192	3	4096 x 2160	1
RTX6000-6 C	Training workloads	6144	4	4096 x 2160	1
RTX6000-4 C	Inference workloads	4096	6	4096 x 2160	1

A-Series vGPU Types for Quadro RTX 6000

Required license edition: vApps

These vGPU types support a single display with a fixed maximum resolution.

Table A-8 A-series vGPU types for Quadro RTX 6000

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Maximum Display Resolution	Virtual Displays per vGPU
RTX6000-2 4A	Virtual applications	24576	1	1280 × 1024	1
RTX6000-1 2A	Virtual applications	12288	2	1280 × 1024	1

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Maximum Display Resolution	Virtual Displays per vGPU
RTX6000-8 A	Virtual applications	8192	3	1280 × 1024	1
RTX6000-6 A	Virtual applications	6144	4	1280 × 1024	1
RTX6000-4 A	Virtual applications	4096	6	1280 × 1024	1
RTX6000-3 A	Virtual applications	3072	8	1280 × 1024	1
RTX6000-2 A	Virtual applications	2048	12	1280 × 1024	1
RTX6000-1 A	Virtual applications	1024	24	1280 × 1024	1

Tesla V100 PCIe vGPU Types

Q-Series vGPU Types for Tesla V100 PCIe

Required license edition: vWS

The maximum number of available combined resolution pixels supported by these vGPU types is determined by their frame buffer size. You can choose to use a small number of high-resolution displays or a larger number of low-resolution displays with these vGPU types. The maximum number of displays per vGPU is based on a configuration in which all displays have the same resolution.

Table A-9 Q-series vGPU types for Tesla V100 PCIe

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
V100-16 Q	Virtual workstations	16384	1	6635520	7680 × 4320	2
					5120 × 2880 or lower	4
V100-8Q	Virtual workstations	8192	2	6635520	7680 × 4320	2
					5120 × 2880 or lower	4

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
V100-4Q	Virtual workstations	4096	4	58982400	7680 × 4320	1
					5120 × 2880 or lower	4
V100-2Q	Virtual workstations	2048	8	35389440	7680 × 4320	1
					5120 × 2880	2
					4096 × 2160 or lower	4
V100-1Q	Virtual desktops and virtual workstations	1024	16	17694720	5120 × 2880	1
					4096 × 2160	2
					3840 × 2160	2
					2560 × 1600 or lower	4

B-Series vGPU Types for Tesla V100 PCIe

Required license edition: vPC or vWS

The maximum number of available combined resolution pixels supported by these vGPU types is determined by their frame buffer size. You can choose to use a small number of high-resolution displays or a larger number of low-resolution displays with these vGPU types. The maximum number of displays per vGPU is based on a configuration in which all displays have the same resolution.

Table A-10 B-series vGPU types for Tesla V100 PCIe

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
V100-2B	Virtual desktops	2048	8	17694720	5120 × 2880	1

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
					4096 x 2160	2
					3840 x 2160	2
					2560 x 1600 or lower	4
V100-2B4	Virtual desktops	2048	8	17694720	5120 x 2880	1
					4096 x 2160	2
					3840 x 2160	2
					2560 x 1600 or lower	4
V100-1B	Virtual desktops	1024	16	16384000	5120 x 2880	1
					4096 x 2160	1
					3840 x 2160	1
					2560 x 1600 or lower	4
V100-1B4	Virtual desktops	1024	16	16384000	5120 x 2880	1
					4096 x 2160	1
					3840 x 2160	1
					2560 x 1600 or lower	4

 NOTE

- -1B and -1B4 vGPUs perform adequately when each vGPU has only two 2560 x 1600 virtual displays. If you want to use more than two 2560 x 1600 virtual displays per vGPU, use a vGPU with more frame buffers, such as a -2B or -2B4 vGPU.
- In this version, -1B4 and -2B4 vGPUs are not recommended because they may be removed in a future release. Use the following vGPU types, which provide equivalent functionality:
 - Replace -1B4 vGPU types with -1B vGPU types.
 - Replace -2B4 vGPU types with -2B vGPU types.

C-Series vGPU Types for Tesla V100 PCIe

Required license edition: Virtual Compute Server (vCS) or Quadro vDWS

These vGPU types support a single display with a fixed maximum resolution.

Table A-11 C-series vGPU types for Tesla V100 PCIe

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Maximum Display Resolution	Virtual Displays per vGPU
V100-16C	Training workloads	16384	1	4096 x 2160	1
V100-8C	Training workloads	8192	2	4096 x 2160	1
V100-4C	Training workloads	4096	4	4096 x 2160	1

 NOTE

C-series vGPU types are NVIDIA vCS vGPU types, which are optimized for compute-intensive workloads. Therefore, C-series vGPUs support only a single display head and do not support Quadro graphics acceleration.

A-Series vGPU Types for Tesla V100 PCIe

Required license edition: vApps

These vGPU types support a single display with a fixed maximum resolution.

Table A-12 A-series vGPU types for Tesla V100 PCIe

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Maximum Display Resolution	Virtual Displays per vGPU
V100-16A	Virtual applications	16384	1	1280 × 1024	1
V100-8A	Virtual applications	8192	2	1280 × 1024	1

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Maximum Display Resolution	Virtual Displays per vGPU
V100-4A	Virtual applications	4096	4	1280 × 1024	1
V100-2A	Virtual applications	2048	8	1280 × 1024	1
V100-1A	Virtual applications	1024	16	1280 × 1024	1

 **NOTE**

A-series NVIDIA vGPUs support a single display with low resolution to be used as the console display for remote applications such as RDSH and Citrix Virtual Apps and Desktops. The maximum resolution and number of virtual display heads of A-series NVIDIA vGPUs apply only to the console display. The maximum resolution of each RDSH or Citrix Virtual Apps and Desktops session is determined by the remote solution and is not restricted by the maximum resolution of the vGPU. Similarly, the number of virtual display heads supported by each session is determined by the remote solution and is not restricted by the vGPU.

Tesla V100 PCIe 32 GB vGPU Types

Q-Series vGPU Types for Tesla V100 PCIe 32 GB

Required license edition: vWS

The maximum number of available combined resolution pixels supported by these vGPU types is determined by their frame buffer size. You can choose to use a small number of high-resolution displays or a larger number of low-resolution displays with these vGPU types. The maximum number of displays per vGPU is based on a configuration in which all displays have the same resolution.

Table A-13 Q-series vGPU types for Tesla V100 PCIe 32 GB

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
V100D-32Q	Virtual workstations	32768	1	66355200	7680 × 4320	2
					5120 × 2880 or lower	4
V100D-16Q	Virtual workstations	16384	2	66355200	7680 × 4320	2

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
					5120 × 2880 or lower	4
V100D-8 Q	Virtual workstations	8192	4	66355200	7680 × 4320	2
					5120 × 2880 or lower	4
V100D-4 Q	Virtual workstations	4096	8	58982400	7680 × 4320	1
					5120 × 2880 or lower	4
V100D-2 Q	Virtual workstations	2048	16	35389440	7680 × 4320	1
					5120 × 2880	2
					4096 × 2160 or lower	4
V100D-1 Q	Virtual desktops and virtual workstations	1024	32	17694720	5120 × 2880	1
					4096 × 2160	2
					3840 × 2160	2
					2560 × 1600 or lower	4

B-Series vGPU Types for Tesla V100 PCIe 32 GB

Required license edition: vPC or vDWS

The maximum number of available combined resolution pixels supported by these vGPU types is determined by their frame buffer size. You can choose to use a small number of high-resolution displays or a larger number of low-resolution displays with these vGPU types. The maximum number of displays per vGPU is based on a configuration in which all displays have the same resolution.

Table A-14 B-series vGPU types for Tesla V100 PCIe 32 GB

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
V100D-2 B	Virtual desktops	2048	16	17694720	5120 x 2880	1
					4096 x 2160	2
					3840 x 2160	2
					2560 x 1600 or lower	4
V100D-2 B4	Virtual desktops	2048	16	17694720	5120 x 2880	1
					4096 x 2160	2
					3840 x 2160	2
					2560 x 1600 or lower	4
V100D-1 B	Virtual desktops	1024	32	16384000	5120 x 2880	1
					4096 x 2160	1
					3840 x 2160	1
					2560 x 1600 or lower	4
V100D-1 B4	Virtual desktops	1024	32	16384000	5120 x 2880	1
					4096 x 2160	1
					3840 x 2160	1
					2560 x 1600 or lower	4

 **NOTE**

- -1B and -1B4 vGPUs perform adequately when each vGPU has only two 2560 x 1600 virtual displays. If you want to use more than two 2560 x 1600 virtual displays per vGPU, use a vGPU with more frame buffers, such as a -2B or -2B4 vGPU.
- In this version, -1B4 and -2B4 vGPUs are not recommended because they may be removed in a future release. Use the following vGPU types, which provide equivalent functionality:
 - Replace -1B4 vGPU types with -1B vGPU types.
 - Replace -2B4 vGPU types with -2B vGPU types.

C-Series vGPU Types for Tesla V100 PCIe 32 GB

Required license edition: Virtual Compute Server (vCS) or Quadro vDWS

These vGPU types support a single display with a fixed maximum resolution.

Table A-15 C-series vGPU types for Tesla V100 PCIe 32 GB

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Maximum Display Resolution	Virtual Displays per vGPU
V100D-32C	Training workloads	32768	1	4096 x 2160	1
V100D-16C	Training workloads	16384	2	4096 x 2160	1
V100D-8C	Training workloads	8192	4	4096 x 2160	1
V100D-4C	Training workloads	4096	8	4096 x 2160	1

 **NOTE**

C-series vGPU types are NVIDIA vCS vGPU types, which are optimized for compute-intensive workloads. Therefore, C-series vGPUs support only a single display head and do not support Quadro graphics acceleration.

A-Series vGPU Types for Tesla V100 PCIe 32 GB

Required license edition: vApps

These vGPU types support a single display with a fixed maximum resolution.

Table A-16 A-series vGPU types for Tesla V100 PCIe 32 GB

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Maximum Display Resolution	Virtual Displays per vGPU
V100D-32A	Virtual applications	32768	1	1280 × 1024	1
V100D-16A	Virtual applications	16384	2	1280 × 1024	1
V100D-8A	Virtual applications	8192	4	1280 × 1024	1
V100D-4A	Virtual applications	4096	8	1280 × 1024	1
V100D-2A	Virtual applications	2048	16	1280 × 1024	1
V100D-1A	Virtual applications	1024	32	1280 × 1024	1

NOTE

A-series NVIDIA vGPUs support a single display with low resolution to be used as the console display for remote applications such as RDSH and Citrix Virtual Apps and Desktops. The maximum resolution and number of virtual display heads of A-series NVIDIA vGPUs apply only to the console display. The maximum resolution of each RDSH or Citrix Virtual Apps and Desktops session is determined by the remote solution and is not restricted by the maximum resolution of the vGPU. Similarly, the number of virtual display heads supported by each session is determined by the remote solution and is not restricted by the vGPU.

Tesla V100S PCIe 32 GB vGPU Types

Q-Series vGPU Types for Tesla V100S PCIe 32 GB

Required license edition: vWS

The maximum number of available combined resolution pixels supported by these vGPU types is determined by their frame buffer size. You can choose to use a small number of high-resolution displays or a larger number of low-resolution displays with these vGPU types. The maximum number of displays per vGPU is based on a configuration in which all displays have the same resolution.

Table A-17 Q-series vGPU types for Tesla V100S PCIe 32 GB

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
V100S-32 Q	Virtual workstations	32768	1	66355200	7680 × 4320	2
					5120 × 2880 or lower	4
V100S-16 Q	Virtual workstations	16384	2	66355200	7680 × 4320	2
					5120 × 2880 or lower	4
V100S-8 Q	Virtual workstations	8192	4	66355200	7680 × 4320	2
					5120 × 2880 or lower	4
V100S-4 Q	Virtual workstations	4096	8	58982400	7680 × 4320	1
					5120 × 2880 or lower	4
V100S-2 Q	Virtual workstations	2048	16	35389440	7680 × 4320	1
					5120 × 2880	2
					4096 × 2160 or lower	4
V100S-1 Q	Virtual desktops and virtual workstations	1024	32	17694720	5120 × 2880	1
					4096 × 2160	2
					3840 × 2160	2
					2560 × 1600 or lower	4

B-Series vGPU Types for Tesla V100S PCIe 32 GB

Required license edition: vPC or vDWS

The maximum number of available combined resolution pixels supported by these vGPU types is determined by their frame buffer size. You can choose to use a small number of high-resolution displays or a larger number of low-resolution displays with these vGPU types. The maximum number of displays per vGPU is based on a configuration in which all displays have the same resolution.

Table A-18 B-series vGPU types for Tesla V100S PCIe 32 GB

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Available Pixels	Display Resolution	Virtual Displays per vGPU
V100S-2B	Virtual desktops	2048	16	17694720	5120 x 2880	1
					4096 x 2160	2
					3840 x 2160	2
					2560 x 1600 or lower	4
V100S-1B	Virtual desktops	1024	32	16384000	5120 x 2880	1
					4096 x 2160	1
					3840 x 2160	1
					2560 x 1600 or lower	4

 **NOTE**

-1B vGPUs perform adequately when each vGPU has only two 2560 x 1600 virtual displays. If you want to use more than two 2560 x 1600 virtual displays per vGPU, use a vGPU with more frame buffers, such as a -2B vGPU.

C-Series vGPU Types for Tesla V100S PCIe 32 GB

Required license edition: Virtual Compute Server (vCS) or Quadro vDWS

These vGPU types support a single display with a fixed maximum resolution.

Table A-19 C-series vGPU types for Tesla V100S PCIe 32 GB

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Maximum Display Resolution	Virtual Displays per vGPU
V100S-32C	Training workloads	32768	1	4096 x 2160	1
V100S-16C	Training workloads	16384	2	4096 x 2160	1
V100S-8C	Training workloads	8192	4	4096 x 2160	1
V100S-4C	Training workloads	4096	8	4096 x 2160	1

NOTE

C-series vGPU types are NVIDIA vCS vGPU types, which are optimized for compute-intensive workloads. Therefore, C-series vGPUs support only a single display head and do not support Quadro graphics acceleration.

A-Series vGPU Types for Tesla V100S PCIe 32 GB

Required license edition: vApps

These vGPU types support a single display with a fixed maximum resolution.

Table A-20 A-series vGPU types for Tesla V100S PCIe 32 GB

vGPU Type	Expected Use Case	Frame Buffer (MB)	Maximum vGPUs per GPU	Maximum Display Resolution	Virtual Displays per vGPU
V100S-32A	Virtual applications	32768	1	1280 × 1024	1
V100S-16A	Virtual applications	16384	2	1280 × 1024	1
V100S-8A	Virtual applications	8192	4	1280 × 1024	1
V100S-4A	Virtual applications	4096	8	1280 × 1024	1
V100S-2A	Virtual applications	2048	16	1280 × 1024	1
V100S-1A	Virtual applications	1024	32	1280 × 1024	1

 **NOTE**

A-series NVIDIA vGPUs support a single display with low resolution to be used as the console display for remote applications such as RDSH and Citrix Virtual Apps and Desktops. The maximum resolution and number of virtual display heads of A-series NVIDIA vGPUs apply only to the console display. The maximum resolution of each RDSH or Citrix Virtual Apps and Desktops session is determined by the remote solution and is not restricted by the maximum resolution of the vGPU. Similarly, the number of virtual display heads supported by each session is determined by the remote solution and is not restricted by the vGPU.

B Installing a GRID Driver on a vGPU-accelerated ECS

Scenario

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](#)

NOTE

NVIDIA allows you to apply for a 90-day trial license.

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 required information.
For details about how to register an account and apply for a trial license, see the [official NVIDIA help page](#).

- Step 5** Enter the PAK obtained in step **Step 4** on the **Redeem Product Activation Keys** page and click **Redeem**.
- Step 6** Specify **Username** and **Password** and click **LOGIN**.
- Step 7** Log in at the official NVIDIA website as prompted and select **SOFTWARE DOWNLOADS**.
- Step 8** Download the GRID driver of the required version based on the VM OS.
- Step 9** Decompress the GRID driver installation package and install the driver that matches your ECS OS.
- Step 10** On the **SOFTWARE DOWNLOADS** page, click **ADDITIONAL SOFTWARE** to download the license software package.
- End

Deploying and Configuring the License Server

The following uses an ECS running CentOS 7.5 as an example to describe how to deploy and configure the license server on the ECS.

NOTE

- The target ECS must have at least 2 vCPUs and 4 GiB of memory.
- Ensure that the MAC address of the target ECS has been recorded.
- If the license server is used in the production environment, deploy it in high availability mode. For details, see [official NVIDIA documentation for license server high availability](#).

Step 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.

Step 2 Install the license server.

For details, see the [official NVIDIA documentation for installing the license server](#).

Step 3 Obtain the license file.

1. Log in to the [NVIDIA website](#) on a new tab and select **LICENSE SERVERS**. Click **CREATE SERVER**.
2. Enter the information about the license server.
 - Set **Server Name**, **Description**, and **MAC Address** (MAC address of the license server).
 - Select **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**.

Click **CREATE LICENSE SERVER**.

3. Download the license file.

Step 4 In the web browser, access the homepage of the license server management page using the link configured during the installation.

Step 5 Choose **License Server > License Management**, select the .bin license file to be uploaded, and click **Upload**.

----End

Installing the GRID Driver and Configuring the License

Install the GRID driver of a desired version, for example, on a GPU-accelerated Windows ECS.

NOTE

Microsoft remote login protocols do not support GPU 3D hardware acceleration. To use this function, install third-party desktop protocol-compliant software, such as VNC, PCoIP, or NICE DCV, and access the ECS through the client.

Step 1 Open the NVIDIA control panel on the Windows control panel.

Step 2 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, 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.

----End

C Supported Driver Versions and OSs

- [Table C-1](#) lists supported driver versions.

Table C-1 Driver versions for GPU passthrough ECSs

vGPU	OS	CUDA Toolkit	Driver
NVIDIA Tesla P4	Windows 10	11.4	471.41
	Windows Server 2016	11.4	471.41
	Windows Server 2019	11.4	471.41
	Ubuntu Desktop 18.04.1	11.4	471.41
	Ubuntu Server 18.04.1	11.4	471.41
	CentOS 7.6/7.7	11.4	470.57.02
	CentOS 8.0	11.4	470.57.02
NVIDIA Tesla P40	Windows 10	11.4	471.41
	Windows Server 2016	11.4	471.41
	Windows Server 2019	11.4	471.41
	Ubuntu Desktop 18.04.1	11.4	471.41
	Ubuntu Server 18.04.1	11.4	471.41
	CentOS 7.6/7.7	11.4	470.57.02
	CentOS 8.0	11.4	470.57.02

vGPU	OS	CUDA Toolkit	Driver
Tesla P100 PCIe 12 GB	Windows 10	11.4	471.41
	Windows Server 2016	11.4	471.41
	Windows Server 2019	11.4	471.41
	Ubuntu Desktop 18.04.1	11.4	471.41
	Ubuntu Server 18.04.1	11.4	471.41
	CentOS 7.6/7.7	11.4	470.57.02
	CentOS 8.0	11.4	470.57.02
Tesla V100 PCIe	Windows 10	11.4	471.41
	Windows Server 2016	11.4	471.41
	Windows Server 2019	11.4	471.41
	Ubuntu Desktop 18.04.1	11.4	470.57.02
	Ubuntu Server 18.04.1	11.4	470.57.02
	Ubuntu Server 20.04	11.4	470.57.02
	CentOS 7.6/7.7/7.8	11.4	470.57.02
	CentOS 8.0/8.1/8.2	11.4	470.57.02
Tesla V100 PCIe 32 GB	Windows 10	11.4	471.41
	Windows Server 2016	11.4	471.41
	Windows Server 2019	11.4	471.41
	Ubuntu Desktop 18.04.1	11.4	470.57.02
	Ubuntu Server 18.04.1	11.4	470.57.02
	Ubuntu Server 20.04	11.4	470.57.02

vGPU	OS	CUDA Toolkit	Driver
	CentOS 7.6/7.7/7.8	11.4	470.57.02
	CentOS 8.0/8.1/8.2	11.4	470.57.02
Tesla V100S PCIe 32GB	Windows 10	11.4	471.41
	Windows Server 2016	11.4	471.41
	Windows Server 2019	11.4	471.41
	Ubuntu Desktop 18.04.1	11.4	470.57.02
	Ubuntu Server 18.04.1	11.4	470.57.02
	Ubuntu Server 20.04	11.4	470.57.02
	CentOS 6.10	11.4	470.57.02
	CentOS 7.9	11.4	470.57.02
	Red Hat Enterprise Linux (RHEL) 6.10	11.4	470.57.02
	NeoKylin 7.6	11.4	470.57.02
Tesla T4	Windows 10	11.4	471.41
	Windows Server 2016	11.4	471.41
	Windows Server 2019	11.4	471.41
	Ubuntu Desktop 18.04.1	11.4	470.57.02
	Ubuntu Server 18.04.1	11.4	470.57.02
	Ubuntu Server 20.04	11.4	470.57.02
	CentOS 7.6/7.7/7.8	11.4	470.57.02
	CentOS 8.0/8.1/8.2	11.4	470.57.02

- [Table C-2](#) lists supported guest OSs.

Table C-2 Guest OSs for vGPU-accelerated ECSs

vGPU	Guest OS
Tesla T4	Windows 10, Windows Server 2016, and Windows Server 2019 Ubuntu Server 18.04.1, Ubuntu Desktop 18.04.1 Ubuntu Server 20.04
Quadro RTX 6000	Windows 10, Windows Server 2016, and Windows Server 2019 Ubuntu Server 18.04.1, Ubuntu Desktop 18.04.1 Ubuntu Server 20.04
Tesla V100 PCIe	Windows 10, Windows Server 2016, and Windows Server 2019 Ubuntu Server 18.04.1, Ubuntu Desktop 18.04.1 Ubuntu Server 20.04
Tesla V100 PCIe 32 GB	Windows 10, Windows Server 2016, and Windows Server 2019 Ubuntu Server 18.04.1, Ubuntu Desktop 18.04.1 Ubuntu Server 20.04
Tesla V100S PCIe 32 GB	Windows 10, Windows Server 2016, and Windows Server 2019 Ubuntu Server 18.04.1, Ubuntu Desktop 18.04.1 Ubuntu Server 20.04 CentOS 6.10 and CentOS 7.9 RHEL 6.10 NeoKylin 7.6