

# Elastic Compute Service

Quick Start for Linux

# Quick Start for Linux

## Overview

### Purpose of this document

This document describes how to quickly create an instance running Linux, connect to an instance remotely, and format the data disk. It is designed to be the one-stop process to walk you through the creation and purchase of instances, as well as remote login and quick environment deployment.

Elastic Compute Service (ECS) instance is sometimes called Alibaba Cloud Elastic Compute Service, Elastic Compute Service, or ECS. To avoid misunderstanding, this document uses the term **Elastic Compute Service (ECS) instance**.

### Target readers

This document is a quickstart reference for anyone who wants to get started with an ECS running Linux.

This document is intended for users that can purchase products and services online. For API users, refer to [API documentation](#).

### Quickstart flowchart

The following figure shows the process of creating and using an ECS instance.



Generally, follow the steps below to purchase and use an ECS instance:

1. Select configuration
2. Create an instance
3. Log on to an instance
4. Format a data disk

# Step 1. Select a configuration scheme

Alibaba Cloud recommends the following configuration schemes, which can meet the requirements of most users.

- **Entry level:** 1-core CPU, 1GB memory and 1Mbps bandwidth. This scheme is applicable to personal sites at their initial stage with relatively low traffic.
- **Advanced:** 1-core CPU, 2GB memory and 1Mbps bandwidth. This scheme is applicable to the websites with moderate traffic, simple development environments and codebase.
- **General:** 2-core CPU, 4GB memory and 1Mbps bandwidth. This scheme is applicable to enterprise operations, parallel computing applications and common data processing. It can meet the requirements of 90% cloud computing users.
- **Ideal:** 4-core CPU, 8GB memory and 1Mbps bandwidth. This scheme is applicable to services requiring higher computing performance, for example, enterprise operations, batch processing, distributed analysis and applications.

These recommended configuration schemes are only for reference when you start using an ECS. You can modify the configuration at any time (upgrade or downgrade) according to real usage situation.

For more information about web hosting, see [Web Application Hosting](#).

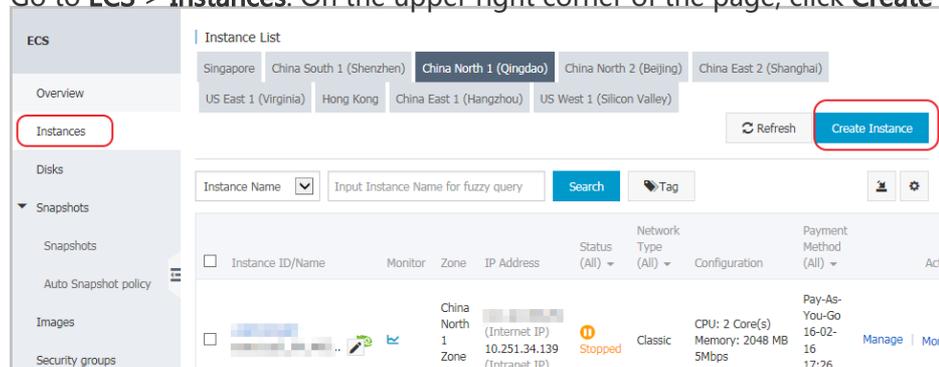
After confirming the configuration scheme, you can start to create an ECS instance.

# Step 2. Create an instance running Linux

Log on to the ECS Management Console.

Enter your account and password. If you do not have an account yet, click **Join Free** to sign up.

Go to **ECS > Instances**. On the upper right corner of the page, click **Create Instance**.



You can choose **Starter Package** or **Advanced Purchase**. Starter Package is mainly for new starters, while Advanced Purchase is for users who need to customize the configuration of the ECS instance.

Choose a pricing model: **Subscription** or **Pay-As-You-Go**.

- **Subscription:** A pre-paid and cost effective option for a longer term commitment (monthly or yearly pricing).

**Pay-As-You-Go:** A costlier hourly pricing option.

**Note:** For a Pay-As-You-Go instance, even if you stop it, your account will still be charged. If you decide not to use it any more, do not simply stop it, but **release it**.

Choose a region. See <https://intl.aliyun.com/why-alibaba-cloud> for the latest list of regions. For example, when you choose the Singapore region, you can either choose:

- **Random.** The instance will be put in a Random zone for the selected region unless otherwise specified.

Or select either **Singapore Zone A** or **Singapore Zone B**.

**Note:**

- Different regions are not interconnected via Intranet.
- Once an instance is created, you cannot change its region.
- Some features, such as available zones, network types, instance generations, disk types, thresholds for triggering DDoS black holes and instance price, vary by region. Select an appropriate region to meet your business needs.

Select an instance type. Options include Generations and I/O optimized instance.

- **Generation I:** Uses Intel Xeon CPU and DDR3 memory.

**Generation II:** Uses Intel Haswell CPU and DDR4 memory, and delivers improved memory computing power. It is an I/O optimized instance by default and delivers improved storage performance when used with SSD Cloud Disks.

Alibaba Cloud strives to bring the latest technologies and configurations based on user demand. It is always advised to check <https://intl.aliyun.com/buy/ecs> for any option changes.

**Note:**

To use the Windows OS for website building and Web environment deployment, you must select the instance type with **at least 2GB** memory or higher. The instance type with 1 core and 1GB of memory does not support running MySQL.

Select a network type.

- **Classic:** Both public and private IP addresses are distributed in a unified way by Alibaba Cloud. It is easy to configure and convenient to use, and suitable for users who require ease of use and fast use of ECS.
- **VPC (Virtual Private Cloud).** VPC refers to the logically isolated private networks. You can customize network topology and Intranet IP address range. It also supports dedicated network connectivity with external environment. It is suitable for users familiar with network management. Network type is different for different ECS features, but both of these network access services are BGP lines.

Select a network billing type.

- **Fixed Bandwidth:** Charges are determined by the fixed Internet bandwidth provided to the instance. Users that want a fixed monthly cost would prefer this setting.

**Data Transfer:** Charges are determined by the amount of the data transferred to an instance (usually calculated by GB). To help prevent high charges from sudden traffic, you can specify bandwidth for the Instance. Users that don't mind fluctuating costs to meet the Instance Internet network demands would prefer this billing type.

If you select a bandwidth of **0MB**, no IP address will be assigned and the instance cannot access the public network.

Select an image.

- **Public Image:** It contains the basic operating system image maintained officially by Alibaba Cloud. On top of this, you will need to install the related software and configure the application environment based on your specific requirements.
- **Custom Image:** It is generated based on the user system snapshot, including the initial system environment, application environment and related software configuration. Selecting custom images to create ECS can save your time for repeated configuration.
- **Shared Image:** It is the custom image shared by another account. Alibaba Cloud does not guarantee the integrity and security of such shared images, and you shall bear any risks associated with using these shared images.

When selecting an OS:

**For Windows:**

- By default, the Windows 2003/2008 system allows remote connection of up to 2 sessions. For more connections, you need to buy Remote Desktop Licensing (RD Licensing) service from Microsoft separately at your own expenses.
- It is suitable for running Windows programs, such as .net.

- It supports SQL Server and other databases (must be independently installed).
- It can be managed through remote desktop login.

**For Linux:**

- It is the most popular server operating system built with robust security and stability.
- It is free and open-source, easy to build and compile the source code.
- It is connected to the ECS through SSH.
- It is generally used for high-performance Web server and other server applications, and supports common programming languages such as PHP and Python, as well as MySQL and other databases (must be independently installed).
- CentOS is recommended.

Select a storage. **System Disk** is required. You can add up to four **Data Disks** based on service requirements.

- **Ultra Cloud Disk:** It delivers data reliability of up to 99.9999999% with triple-copy distributed storage technology. It initially provides 1,000 random IOPS and 50MB/s throughput, with an increment of 6 IOPS and 0.1MB throughput per GB, up to 3,000 random reading/writing IOPS and 80MB/s throughput. For example, a 200GB Ultra Cloud Disk provides 2,200 random reading/writing IOPS and 70 Mbps reading/writing throughput.
- **SSD Cloud Disk:** As with cloud disks, it delivers a data reliability of up to 99.9999999% with triple-copy distributed storage technology and provides stable and high performance storage of highly random IO. It offers 30 random IOPS per GB, a maximum of 20,000 random IOPS for a single disk, 256MB/s of throughput. To get maximum disk throughput with your Instance and Disk Type, you need to select the IO optimized option.

Set a logon password for the instance. You have the option to set the login password now or later. Password can also be changed anytime within the ECS instance details.

Select the Subscription Time as either 1 Month(s) or 1 Year with Auto-renewal set (Only available on Subscription plans).

Set the instance name.

Select the Number of Instance to launch. You can easily launch up to 50 instances at the same time.

Check the overview to make sure that the selected configuration details are correct. Then either **Add to Cart** (If you decide to continue shopping) or **Buy Now** if you want to confirm the purchase. Note down the instance details such as public and private IP details.

Confirm the order and make payment.

## Step 3. Log on to a Linux instance

You can log on to an ECS from a Windows, Linux or Mac OS X operating system. This section describes these logon methods respectively.

After an ECS instance is created, you can log on to the server through any one of the following methods:

- Using remote connection software: Common remote connection software include Putty and Xshell.
- **Management Terminal**: Whether you have purchased bandwidth or not when creating the instance, you can log on to the instance for management through the Management Terminal of the Management Console.
- Mobile phone: You can also connect to the instance through the remote desktop application on the mobile phone. The operations are simple, so we will not describe them again here.

### Log on to a Linux instance from Windows

The usage of different remote logon software is similar. This document uses Putty as an example to describe how to log on to an instance remotely. Putty, an installation-free tool, is available from <http://www.putty.org/>

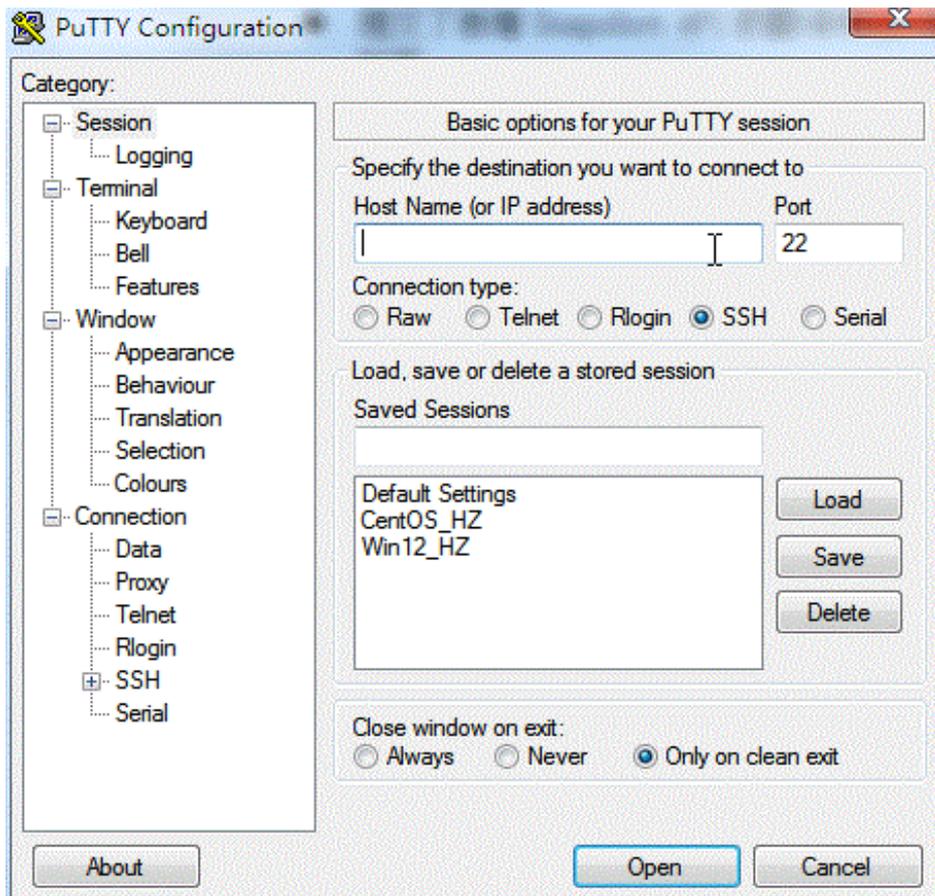
To log on to a Linux instance:

Start Putty.exe and access the main interface of Putty.

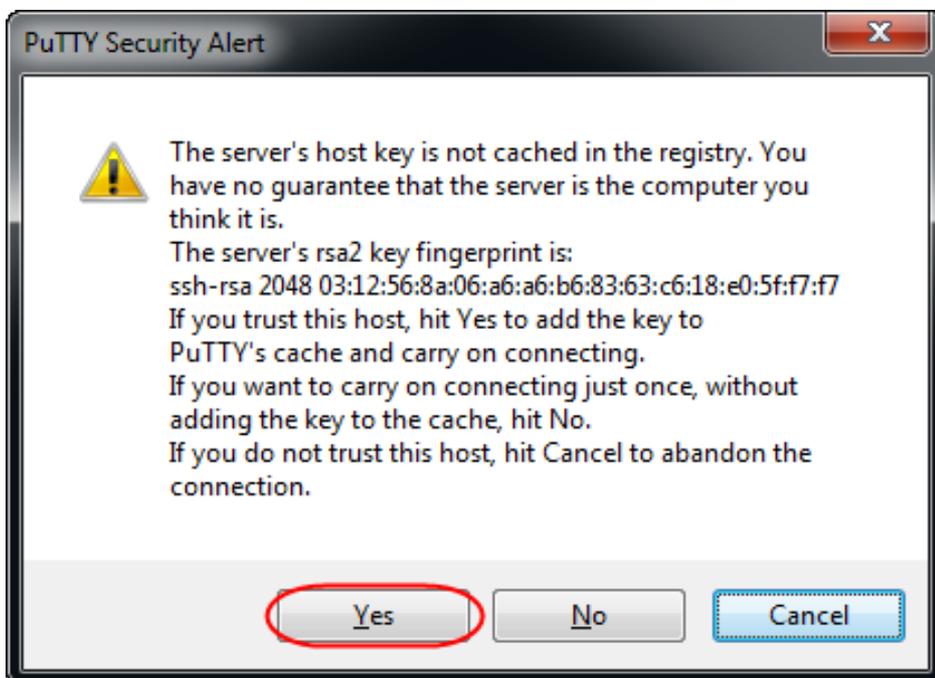
Enter the instance' s public IP address in **Host Name**.

- Use the default port **22**.
- Under **Connection type**, select **SSH**.
- Enter a session name in **Saved Session**, and click **Save**. You can directly access the session without entering the IP address next time.

Click **Open** to connect.

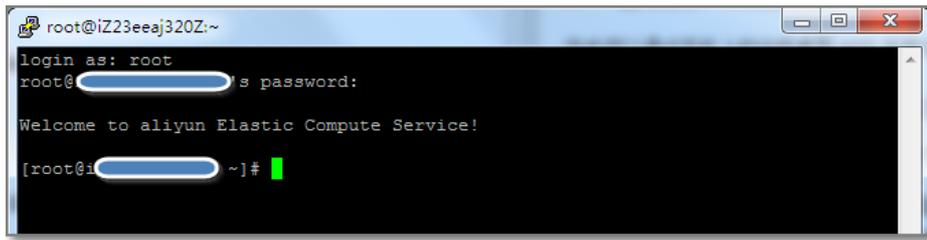


In the first connection, the following message will be displayed. Click **Yes**.



Enter your user name and password for the Linux ECS instance as prompted. The password

will not be displayed on the screen. Then press **Enter**.



Now, you are successfully connected to the instance and can work on it.

## Log on to a Linux instance from Linux or Mac OS X

Connect to the instance by using the SSH command directly. For example, `ssh root@instance' s public IP address`. Then, enter the password of the root user of the instance to log on.

## Log on to an instance through the Management Terminal

The management terminal is applicable to the following scenarios:

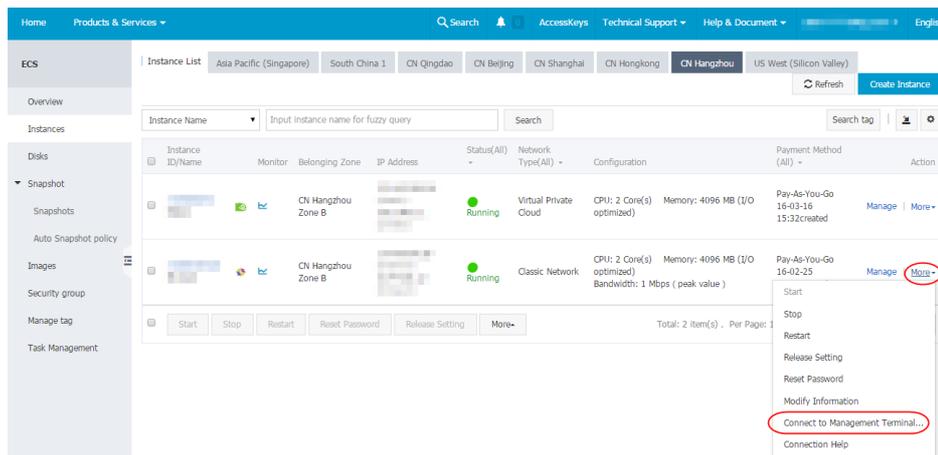
- If the instance boot speed is low and the self-check is initiated, you can view the progress on the Management Terminal.
- If remote connection fails due to a software setup error in the instance (for example, a firewall setup error), you can log on through the Management Terminal and reconfigure the firewall.
- If remote connection fails due to high CPU usage or bandwidth occupation by applications (for example, full occupation of the CPU/bandwidth by processes due to a Botnet attack), you can log on through the Management Terminal and end abnormal processes.

To log on to an instance through the Management Terminal:

Log on to the ECS Management Console.

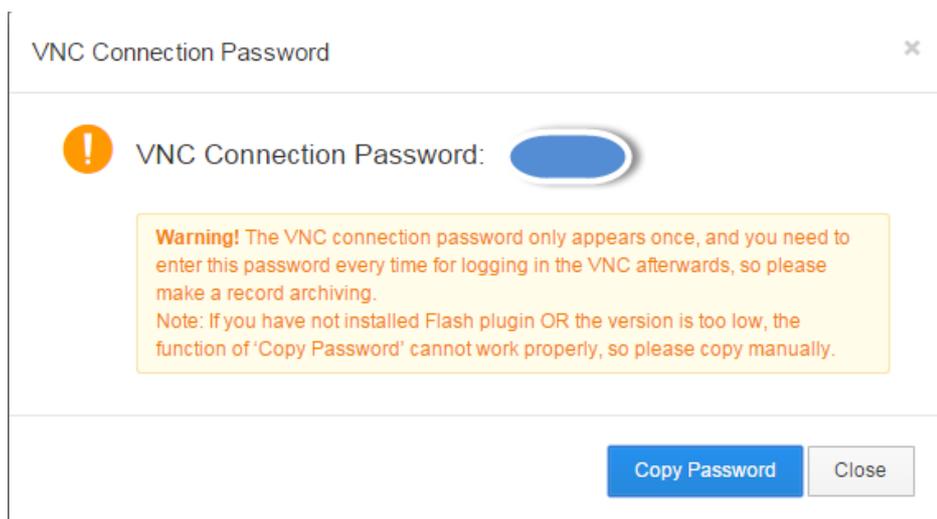
Go to the instance to be connected.

Click **More > Connect to Management Terminal** on the right.

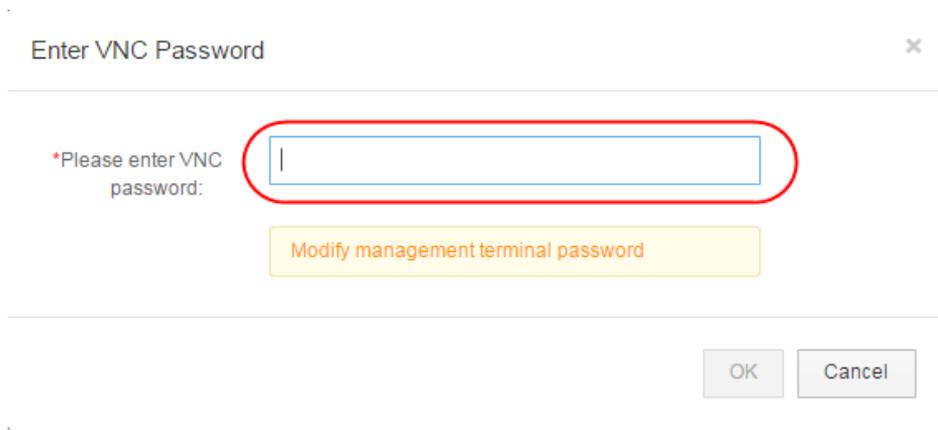


At your first logon, the password will be prompted. The prompt appears only once. You need to enter this password each time when you log on later. **Make sure you note down the password.**

If you forgot the password, or prefer to use a password of your own, click **Change the Password of the Management Terminal** in the upper right corner.



Click **Send Remote Command** in the upper left corner, and then click **Connect Management Terminal**. Enter the password to connect to the instance.



Enter the user name "root" and the password when creating an instance to log on.

```
CentOS release 6.5 (Final)
Kernel 2.6.32-431.23.3.el6.x86_64 on an x86_64
login: _
```

A persisting black screen after you enter the password indicates that the system is sleeping. In this case, you can press any key to wake up the system.

## Step 4. Format and attach a data disk

If you select a data disk when creating an instance, you need to format the data disk before using it.

You can also configure multiple data disk partitions based on service requirements. It is recommended that you use the built-in system tool for partitioning.

**Note:** ECS supports only secondary partitioning of the **data disk**, but does not support secondary partitioning of the **system disk** (either in Windows or Linux OS). If you forcibly use a third-party tool to perform secondary partitioning on the system disk, it may lead to unknown risks, such as system crash and data loss.

The operation is applicable to only **Non-I/O optimized + SSD cloud disk Linux** (Redhat, CentOS, Debian, and Ubuntu) instances.

To format and attach a data disk:

1. Use a remote connection tool. Enter the user name **root** and password to log on to the instance.

Run the **fdisk -l** command to view the data disk. **Note:** Before the data disk is partitioned

and formatted, you cannot view the data disk by running the **df -h** command. In the following example, a 5GB data disk needs to be attached.

If you do not find `/dev/xvdb` after running the **fdisk -l** command, it indicates that your instance does not have a data disk. Therefore, attaching is not required. In this case, you can skip this chapter.

```
[root@xxxx ~]# fdisk -l

Disk /dev/xvda: 42.9 GB, 42949672960 bytes
255 heads, 63 sectors/track, 5221 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00078f9c

Device Boot Start End Blocks Id System
/dev/xvda1 * 1 5222 41940992 83 Linux

Disk /dev/xvdb: 5368 MB, 5368709120 bytes
255 heads, 63 sectors/track, 652 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00000000
```

3. Run the following command to partition the data disk.

```
fdisk /dev/xvdb
```

Enter **n**, **p**, **1** in sequence as prompted, press Enter twice, and then enter **wq**. The partitioning will begin.

```
[root@xxx ~]# fdisk /dev/xvdb
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel
Building a new DOS disklabel with disk identifier 0x33eb5059.
Changes will remain in memory only, until you decide to write them.
After that, of course, the previous content won't be recoverable.

Warning: invalid flag 0x0000 of partition table 4 will be corrected by w(rite)

WARNING: DOS-compatible mode is deprecated. It's strongly recommended to
switch off the mode (command 'c') and change display units to
sectors (command 'u').

Command (m for help): n
Command action
e extended
p primary partition (1-4)
p
```

```
Partition number (1-4): 1
First cylinder (1-652, default 1):
Using default value 1
Last cylinder, +cylinders or +size{K,M,G} (1-652, default 652):
Using default value 652
```

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

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

Run the **fdisk -l** command to view the new partition. The new partition xvdb1 is created, for example, `/dev/xvdb1` below.

```
[root@xxx ~]# fdisk -l

Disk /dev/xvda: 42.9 GB, 42949672960 bytes
255 heads, 63 sectors/track, 5221 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00078f9c

Device Boot Start End Blocks Id System
/dev/xvda1 * 1 5222 41940992 83 Linux

Disk /dev/xvdb: 5368 MB, 5368709120 bytes
255 heads, 63 sectors/track, 652 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x33eb5059

Device Boot Start End Blocks Id System
/dev/xvdb1 1 652 5237158+ 83 Linux
```

Run the following command to format the new partition. The formatting time depends on the size of the data disk. You can also choose other file formats, for example, ext4.

```
mkfs.ext3 /dev/xvdb1
```

7. Run the following command to write the new partition information.

```
echo '/dev/xvdb1 /mnt ext3 defaults 0 0'>> /etc/fstab
```

Upon completion, run the `cat /etc/fstab` command to view the information.

**Note:** Ubuntu 12.04 does not support barrier. Therefore, the correct command for the

system is as follows:

```
echo '/dev/xvdb1 /mnt ext3 barrier=0 0 0' >> /etc/fstab
```

To attach the data disk to a folder separately, for example, to store webpages separately, modify the `/mnt` part in the above command.

8. Run the **mount /dev/xvdb1 /mnt** command to attach the new partition. Then, run the **df -h** command to view the partition. If data disk information is displayed, the attaching is successful and the new partition can be used.

```
[root@xxx ~]# mount /dev/xvdb1 /mnt
[root@xxx ~]# df -h
Filesystem Size Used Avail Use% Mounted on
/dev/xvda1 40G 1.5G 36G 4% /
tmpfs 498M 0 498M 0% /dev/shm
/dev/xvdb1 5.0G 139M 4.6G 3% /mnt
```

**Note:** ECS does not support installation and deployment of virtualization software (for example, KVM, Xen, VMware, and so on).