Image Management Service

FAQs

 Issue
 02

 Date
 2025-05-06





Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2025. All rights reserved.

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

Trademarks and Permissions

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

Notice

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

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

Contents

1 Image Consulting 1	
1.1 How Do I Select an Image?	1
1.2 Are There Any Public Images Already Containing Certain Applications, Such as OpenVPN or PyTorch?	
1.3 What Are the Differences Between Images and Backups?	
1.4 Can I Tailor an Image?	5
1.5 How Can I Back Up the Current Status of an ECS for Restoration in the Case of a System Fault?	5
1.6 How Can I Apply a Private Image to an Existing ECS?	
1.7 Can I Import Data from a Data Disk Image to a Data Disk?	
1.8 Can I Use Private Images of Other Accounts?	5
2 Image Creation	7
2.1 General Creation FAQ	7
2.2 Full-ECS Image FAQ	Э
2.3 How Can I Use a Backup to Create an EVS Disk or ECS?	1
2.4 Is There Any Difference Between the Image Created from a CSBS/CBR Backup and That Created from an ECS?	
2.5 Why Can't I Find an ISO Image When I Want to Use It to Create an ECS or Change the OS of an ECS?	
2.6 How Do I Create a Full-ECS Image Using an ECS That Has a Spanned Volume?12	2
2.7 Why Is Sysprep Required for Creating a Private Image from a Windows ECS?	2
2.8 How Do I Handle the Startup Failure of a Windows ECS Created from a Windows Image Generalized by Sysprep?	
2.9 What Do I Do If I Cannot Create an Image in ZVHD2 Format Using an API?14	4
2.10 What Will the System Do to an Image File When I Use the File to Register a Private Image?14	4
2.11 How Do I Configure an ECS, a BMS, or an Image File Before I Use It to Create an Image?16	5
2.12 What Do I Do If a Windows Image File Is Not Pre-Configured When I Use It to Register a Private Image?	8
2.13 What Do I Do If a Linux Image File Is Not Pre-Configured When I Use It to Register a Private Image?	
2.14 How Do I Make a System Disk Image Support Fast ECS Creation?	4
3 OS	5
3.1 How Do I Select an OS?	
3.2 How Is BIOS Different from UEFI?	ŝ
3.3 What Do I Do If I Cannot Change the Resolution of a Windows OS Booted in UEFI Mode?27	7

3.4 How Do I Delete Redundant Network Connections from a Windows ECS?	31
3.5 What Do I Do If an ECS Starts Slowly?	
3.6 Why Can't I Find My Private Image When I Want to Use It to Create an ECS or Change the OS of ECS?	
3.7 What Do I Do If CentOS Linux Is No Longer Maintained?	
3.8 How Do I Enable NIC Multi-Queue for an Image?	35
3.9 How Do I Configure an ECS to Use DHCPv6?	
3.10 How Do I Clean Up the Disk Space of a Windows ECS?	57
4 Driver Installation	<mark>68</mark>
4.1 Must I Install Guest OS Drivers on an ECS?	68
4.2 Why Do I Need to Install and Update VirtIO Drivers for Windows?	
4.3 Why Did I Fail to Install Guest OS Drivers on a Windows ECS?	
4.4 How Do I Install VirtIO Drivers in Windows?	
4.5 How Do I Install Native KVM Drivers in Linux?	
4.6 How Do I Install Native Xen and KVM Drivers?	
5 Image Sharing	
5.1 General Sharing FAQ	
5.2 How Will I Be Billed for Sharing an Image?	
5.3 What Are the Differences Between Sharing Images and Replicating Images?	
6 Image Import	
6.1 Can I Use Images in Formats not Described in This Document?	
6.2 What Are the Impacts If I Do Not Pre-configure an ECS Used to Create a Private Image?	
6.3 How Do I Import an OVF or OVA File to the Cloud Platform?	
6.4 What Do I Do If I Chose the Wrong OS or System Disk Capacity When Registering a Private Image	
6.5 Why Did My VHD Upload Fail? Why Does the System Say the System Disk in the VHD Image File	ls
Larger Than What I Specified on the Management Console?	84
7 Image Export	85
7.1 Can I Download My Private Images to a Local PC?	
7.2 Can I Use the System Disk Image of an ECS on a BMS After I Export It from the Cloud Platform?	
7.3 Why Is the Image Size in an OBS Bucket Different from That Displayed in IMS?	
7.4 Can I Download a Public Image to My PC?	
7.5 What Are the Differences Between Standard and Fast Import/Export?7.6 Why the Export Option Is Unavailable for My Image?	
8 Image Replication	
8.1 Image Replication	
8.2 How Do I Limit Agency Permissions for Cross-Region Image Replication?	
9 Image Deletion	
10 Image Encryption	<mark>95</mark>
11 Accounts and Permissions	96

11.1 What Do I Do If I Enabled EPS But Now I Cannot Find Private Images in My Enterprise Project?96
11.2 What Do I Do If I Cannot Create an Image from a CSBS Backup or BMS Using a Sub-account with the Allow_all Permission After EPS Is Enabled?
12 Cloud-Init
12.1 Cloud-Init Installation FAQ
12.2 What Can I Do with a Cloud-Init ECS?
12.3 What Do I Do If Installed NetworkManager and Now I Can't Inject the Key or Password Using Cloud-Init?
12.4 How Do I Install growpart for SUSE 11 SP4?
13 ECS Creation104
13.1 Can I Change the Image of a Purchased ECS?
13.2 Can I Change the Specifications Defined by a Private Image When I Use the Image to Create an ECS?
13.3 Can I Specify the System Disk Capacity When I Create an ECS Using an Image?
13.4 What Do I Do If a Partition Is Not Found During the Startup of an ECS Created from an Imported Private Image?
13.5 What Do I Do If the Disks of a CentOS ECS Created from an Image Cannot Be Found?107
13.6 What Do I Do If I Enabled Automatic Configuration During Image Registration for an ECS Created from a Windows Image and Now It Won't Start?
13.7 What Do I Do If an Exception Occurs When I Start an ECS Created from an Image Using UEFI Boot?
14 Billing
14.1 What Are the Billing Standards for IMS?

Image Consulting

1.1 How Do I Select an Image?

When creating an ECS or a BMS, you can select an image based on the following factors:

- Image Type
- Image Fee
- **OS**

Image Type

Images are classified into public, private, and shared images.

lmag e Type	Description	Availability/ Security	Billing
Publi c	A public image is a standard, widely used image. It contains an OS and preinstalled public applications and is available to all users. Public images are very stable and their OS and any included software have been officially authorized for use.	High	Free
Privat e	A private image contains an OS or service data, preinstalled public applications, and a user's personal applications. Private images are only available to the users who created them.	Medium	Free

lmag e Type	Description	Availability/ Security	Billing
Share d	A shared image is a private image another user has shared with you.	Low	Free

You can select an image as follows:

- If you only need an OS, select a public image.
- If you want to clone an existing ECS, select a private image.
- If you want to use an image shared by others, select a shared image.

Image Fee

You may be billed for using an image. For details, see **What Are the Billing Standards for IMS?**

OS

When selecting an OS, consider the following factors:

• Architecture types

Applicable Memory	Constraints
Smaller than 4 GB	 If the instance memory is greater than 4 GB, a 32-bit OS cannot be used. A 32-bit OS only allows addressing within a 4 GB memory range. An OS with more than 4 GB memory
4 GB or larger	cannot be accessed. If your application requires more than 4 GB of memory or the memory may need to be expanded to more than 4 GB, use a 64-bit OS.
	Memory Smaller than 4 GB

• OS types

OS Type Applicable Scenario		Constraints	
Linux	 High-performance server applications (for example, Web) and common programming languages such as PHP and Python 	The system disk must be at least 40 GB, and there must be at least 512 MB of memory.	
	 Databases such as MySQL. (You need to install the database.) 		

1.2 Are There Any Public Images Already Containing Certain Applications, Such as OpenVPN or PyTorch?

No public image contains applications like OpenVPN or PyTorch.

You need to manually install OpenVPN or PyTorch or search for eligible images in KooGallery.

1.3 What Are the Differences Between Images and Backups?

CBR and Image Management Service (IMS) have some complementary functions and can be used together in certain scenarios. Like CBR, IMS can also be used to back up ECSs.

Differences Between Backups and Images

Table 1-2 lists the differences between them.

Item	CBR	IMS
Concept	A backup contains the status, configuration, and data of a cloud server or disk stored at a specific time point for recovery in case of a fault. It is used to ensure data security and improve availability.	An image provides all information required for starting a cloud server. It is used to create a cloud server and deploy software environments in batches. A system disk image contains an OS and pre-installed application software for running services. A data disk image contains service data. A full-ECS image contains data of the system disk and data disks.

Table 1-2 Differences between backups and images

ltem	CBR	IMS
Usage method	 Data storage location: Unlike server or disk data, backups are stored in OBS. Deleting a disk will not clear its backups. Operation object: A server or disk can be backed up at a given point in time. CBR supports automatic backup and automatic deletion by configuring backup policies. Usage: Backups can be used to restore data to the original server or disk, or to create a new disk or full- ECS image. Support exporting to a local PC: No 	 Data storage location: Unlike server or disk data, backups are stored in OBS. If a server or disk that is created using an image is deleted, the image will not be cleared. Operation object: The system disk and data disks of a server can be used to create private images. You can also create private image files. Usage: System disk images or full-ECS images can be used to create new servers, and data disk images can be used to create new disks for service migration. Support exporting to a local PC: Yes However, full-ECS images cannot be exported to a local PC.
Application scenarios	 Data backup and restoration Rapid service deployment and migration 	 Server migration to the cloud or between clouds Deploying a specific software environment Deploying software environments in batches Backing up server operating environments
Advantages	Supports automatic backup. Data on a server or disk at a certain time point can be retained periodically or quantitatively.	Supports system disk backup. You can import the data disk image of a local server or a server provided by another cloud platform to IMS and then use the image to create an EVS disk.

NOTE

Although backups and images are stored in OBS, you cannot view backup and image data in OBS, because they do not occupy your resources. Backup fees are charged according to the CBR billing standards, and image storage fees are charged according to the OBS billing standards.

Relationship Between Backups and Images

- 1. You can use an ECS backup to create a full-ECS image.
- 2. Before creating a full-ECS image for an ECS, you need to back up the target ECS.
- 3. A backup is compressed when it is used to create an image, so the size of the generated image may be smaller than the backup size.

1.4 Can I Tailor an Image?

Yes, but you'd better not.

When you import an external image file, you are advised to import the image that contains the official OS release version. Do not tailor or highly customize the release version. Otherwise, problems may occur.

OS vendors do not always update OS release versions regularly. Some versions are no longer maintained, and these deprecated versions no longer receive security patches. Ensure that you read the update notifications from OS vendors and update your OS so that it runs properly.

1.5 How Can I Back Up the Current Status of an ECS for Restoration in the Case of a System Fault?

You can back up the ECS in any of the following ways:

- (Recommended) Use CBR to create a scheduled backup for the ECS. If the ECS fails, select a backup you want to use to restore the ECS, create a full-ECS image from the backup, and use the image to create a new ECS or to reinstall the OS of the ECS.
- Create a system disk image from the ECS. If the ECS fails, use the system disk image to create a new ECS or to reinstall the OS.
- Create a snapshot for the system disk of the ECS. If the ECS fails, you can restore the system disk data from the snapshot.

1.6 How Can I Apply a Private Image to an Existing ECS?

• You can use the image to change the ECS OS. When you change the OS, select that image. For details about how to change the OS, see Changing the OS.

1.7 Can I Import Data from a Data Disk Image to a Data Disk?

No.

A data disk image can only be used to apply for a new disk and its data cannot be imported to a disk. To import the data, perform the following operations:

- 1. Use the data disk image to create a temporary disk.
- 2. Attach the temporary disk to the ECS where the target disk is located.
- 3. Copy data from the temporary disk to the target disk. Then, delete the temporary disk.

1.8 Can I Use Private Images of Other Accounts?

Yes.

You can use the private images other accounts shared with you.

For details, see Sharing Specified Images.

2 Image Creation

2.1 General Creation FAQ

How Can I Use an ECS to Quickly Provision Identical ECSs?

If you have an ECS with applications deployed, you can use the ECS to create a private image and then use the image to create identical ECSs. In this way, you do not need to deploy applications repeatedly.

• Creating an ECS from an Image

How Many Private Images Can I Create Under an Account?

You can create up to 50 private images under a given account within a region. If you need more, submit a service ticket to increase your quota.

Do I Have to Stop the ECS Before Using It to Create a Private Image?

No. You can create an image from a running ECS. But the data that was written into the ECS during the image creation will not be included in the image.

Where Can I Check the Image Creation Progress? How Long Does It Take to Create an Image?

Log in to the management console. Choose **Compute** > **Image Management Service** and click the **Private Images** tab. Check the image creation progress in the **Status** column.

Creating an image may take a period of time because it requires to install Xen and KVM drivers, load the OS kernel, and configure GRUB boot. In addition, the network speed, image file type, and disk size all have impacts on how long the entire process takes.

Can I Use a Private Image of an IAM User Under My Account to Create an ECS?

Yes.

Private images created by an IAM user are visible to the account that the IAM user belongs to as well as all other IAM users (if any) under this account.

- You can use a system disk image or full-ECS image to create an ECS.
- You can use a data disk image to create an EVS disk.

In addition, private images created by an account are visible to IAM users under this account.

What Is a Full-ECS Image?

A full-ECS image contains the OS, applications, and service data of an ECS. Generally, a full-ECS image is used to migrate all data of an ECS. For example:

- Sharing an ECS with other users
- Migrating an ECS from one region to another, for example, from CN East-Shanghai1 to CN South-Guangzhou
- Migrating data from an old ECS to a new one

Why Do I Have to Select a Vault When Creating a Full-ECS Image? Do I Need to Pay for the Vault?

When creating a full-ECS image from a CBR backup, you must select a vault. The vault is where your image and backup are stored. You need to pay for the vault.

When creating a full-ECS image from a CSBS backup, the space used to store the image and CSBS backup is not open to users but still need to be billed.

So, no matter which backup type you select, you need to pay for the storage. Selecting a vault does not mean that you need to pay extra fees.

Where Can I Check the Data Disk Details of a Full-ECS Image?

To check data disk details, you need to go to the CSBS or CBR console, depending on where the full-ECS image is created from. That is because only system disk information (**Disk Capacity**) is displayed in the image list and image details after a full-ECS image is created.

The following describes how to view the data disk details in CBR:

- 1. In the private image list, click the full-ECS image name. Image details are displayed.
- 2. Locate **Source** and click the backup ID following it.

The CBR backup details page is displayed.

Figure 2-1 Image details

Name	cbr_switch_test	ID	ecc65532-7431-49e0-971a-b1912f4e1e2b
Image Type	Full-ECS image	Disk Capacity (GB)	
OS	Ubuntu 16.04 server 64bit	Status	📀 Normal
Minimum Memory	1 GB	Maximum Memory	4096 GB
Released ⑦	No	Created	2023/01/14 09:05:42 GMT+08:00
Completed	2023/01/14 09:05:50 GMT+08:00	Source	CBR backup manualbk_vault-d9cd-aaa0
Description		OS Type	Linux
NIC Multi-Queue	Supported	Enterprise Project	default

3. Click the **Disk Backup** tab. Details about the system disk and data disks are displayed.

Figure 2-2 Disk backup

Details Disk Backup Share List Replication History				
Name	Status	Disk Name	Used As	Disk Capacity (GB)
manualbk_3c43_ecs-37389	 Available 	ecs-373896-volume-0000	System Disk	40
manualbk_3c43_ecs-37389	Available	ecs-373896-volume-0001	Data Disk	100

What Are the Restrictions on Using a Full-ECS Image?

- A full-ECS image cannot be exported. You are advised to create images for the system disk and data disks separately and then export the images.
- A full-ECS image is shareable only when it is created from a CBR backup or from an ECS that has never had a CSBS backup.
- A full-ECS image cannot be replicated within the same region.

2.2 Full-ECS Image FAQ

What Is a Full-ECS Image?

A full-ECS image contains the OS, applications, and service data of an ECS. Generally, a full-ECS image is used to migrate all data of an ECS. For example:

- Sharing an ECS with other users
- Migrating an ECS from one region to another, for example, from CN East-Shanghai1 to CN South-Guangzhou
- Migrating data from an old ECS to a new one

Why Do I Have to Select a Vault When Creating a Full-ECS Image? Do I Need to Pay for the Vault?

When creating a full-ECS image from a CBR backup, you must select a vault. The vault is where your image and backup are stored. You need to pay for the vault.

When creating a full-ECS image from a CSBS backup, the space used to store the image and CSBS backup is not open to users but still need to be billed.

So, no matter which backup type you select, you need to pay for the storage. Selecting a vault does not mean that you need to pay extra fees.

Where Can I Check the Data Disk Details of a Full-ECS Image?

To check data disk details, you need to go to the CSBS or CBR console, depending on where the full-ECS image is created from. That is because only system disk information (**Disk Capacity**) is displayed in the image list and image details after a full-ECS image is created.

The following describes how to view the data disk details in CBR:

- 1. In the private image list, click the full-ECS image name. Image details are displayed.
- 2. Locate **Source** and click the backup ID following it.

The CBR backup details page is displayed.

Figure 2-3 Image details

Name	cbr_switch_test	ID	ecc65532-7431-49e0-971a-b1912f4e1e2b
Image Type	Full-ECS image	Disk Capacity (GB)	
OS	Ubuntu 16.04 server 64bit	Status	📀 Normal
Minimum Memory	1 G8	Maximum Memory	4096 GB
Released 🕜	No	Created	2023/01/14 09:05:42 GMT+08:00
Completed	2023/01/14 09:05:50 GMT+08:00	Source	CBR backup manualbk_vault-d9cd-aaa0
Description	-	OS Type	Linux
NIC Multi-Queue	Supported	Enterprise Project	default

3. Click the **Disk Backup** tab. Details about the system disk and data disks are displayed.

Figure 2-4 Disk backup

Details Disk Backup S	Share List Replication Histo	гу		
Name	Status	Disk Name	Used As	Disk Capacity (GB)
manualbk_3c43_ecs-37389	Available	ecs-373896-volume-0000	System Disk	40
manualbk_3c43_ecs-37389	< Available	ecs-373896-volume-0001	Data Disk	100

What Are the Restrictions on Using a Full-ECS Image?

- A full-ECS image cannot be exported. You are advised to create images for the system disk and data disks separately and then export the images.
- A full-ECS image is shareable only when it is created from a CBR backup or from an ECS that has never had a CSBS backup.
- A full-ECS image cannot be replicated within the same region.

2.3 How Can I Use a Backup to Create an EVS Disk or ECS?

You can use CSBS backups to create ECSs and use VBS backups to create EVS disks.

• CSBS backups cannot be directly used to create ECSs. You need to use a backup to create a private image and then use the private image to create ECSs.

For details about how to create a private image from a CSBS backup, see **Creating a Full-ECS Image from a CSBS Backup**. For details about how to use a private image to create ECSs, see **Creating an ECS from an Image**.

• VBS backups can be directly used to create EVS disks. For details, see Using a Backup to Create a Disk.

2.4 Is There Any Difference Between the Image Created from a CSBS/CBR Backup and That Created from an ECS?

No.

You can create a full-ECS image from an ECS, a CSBS backup, or a CBR backup.

When you create a full-ECS image from an ECS, the system first creates a backup for the ECS and then uses the backup to create an image. Therefore, the image is essentially created from an ECS backup no matter you use an ECS or a CSBS/CBR backup.

2.5 Why Can't I Find an ISO Image When I Want to Use It to Create an ECS or Change the OS of an ECS?

- An ISO image is only an intermediate product and is not available on the ECS console. The final product is a system disk image that has an OS and drivers installed.
- You are not advised to use a temporary ECS as a normal ECS because it has limited functionality. For example, disks cannot be attached to it.

For details about how to create a private image using an ISO file, see:

• Creating a Linux System Disk Image from an ISO File

2.6 How Do I Create a Full-ECS Image Using an ECS That Has a Spanned Volume?

You are not advised to use a Windows ECS that has a spanned volume to create a full-ECS image. If you create such an image and then use it to create new ECSs, data may be lost.

If an ECS has a spanned volume, back up data in the spanned volume and then delete the volume. Use the ECS to create a full-ECS image. You can then use the full-ECS image to create an ECS and use the backup to create a spanned volume for the new ECS if necessary.

NOTE

If a Linux ECS has a volume group or a logical volume consisting of multiple physical volumes, to ensure you do not lose any data, back up data in the volume group or logical volume and delete the volume group or logical volume before using this ECS to create a full-ECS image.

2.7 Why Is Sysprep Required for Creating a Private Image from a Windows ECS?

Why Is Sysprep Required?

Sysprep is used to generalize images. It removes server-specific information, like the security identifier (SID), from an image so that ECSs created from this image can have unique SIDs in a domain. If your windows ECSs do not need to be joined to a domain, Sysprep is not required.

Before running Sysprep, ensure that Windows is activated.

Restrictions on Running Sysprep

Sysprep can only be used to configure new installations of Windows and not to reconfigure an existing installation. You can run Sysprep as many times as required to build and to configure your installation of Windows. However, you can reset Windows activation only up to three times.

NOTE

In the Windows command line, run the following command to check how many times you can run Sysprep:

slmgr /dlv

In the displayed **Windows Script Host** dialog box, if the value of **Remaining Windows** rearm count is **0**, you cannot run Sysprep.

2.8 How Do I Handle the Startup Failure of a Windows ECS Created from a Windows Image Generalized by Sysprep?

Symptom

1. When you start the ECS, information similar to the following is displayed.

Figure 2-5 Message displayed



Then, the following information is displayed in the dialog box:

Windows could not parse or process the unattend answer file for pass [specialize]. A component or setting specified in the answer file does not exist. The error was detected while processing settings for component [Microsoft-Windows-Shell-Setup].

- 2. Click **OK**. The following information is displayed in the dialog box: The computer accidentally restarts or encounters an error. Windows installation cannot continue. Click OK to restart the computer and restart the installation.
- 3. Open **setupact.log** in **C:\Windows\Panther**. The log contains the following information.

Figure 2-6 Viewing ECS logs



Solution

- 1. Create an ECS from a public image. (You are advised to use a public image to create another ECS because Sysprep can be executed only for certain times.)
- 2. Create an **Unattend.xml** file or modify the **Unattend.xml** file provided by the system.
 - If you create an **Unattend.xml** file, ensure that the created file is used when you run Sysprep. For details about the file, visit:
 - https://docs.microsoft.com/en-us/windows-hardware/ manufacture/desktop/update-windows-settings-and-scriptscreate-your-own-answer-file-sxs
 - https://docs.microsoft.com/en-us/windows-hardware/ manufacture/desktop/sysprep--system-preparation--overview

 If you modify the Unattend.xml file (in the C:\Program Files\Cloudbase Solutions\Cloudbase-Init\conf directory), delete the RunSynchronous part from the file.

Figure 2-7	Deleting the	RunSynchronous	snippet
------------	--------------	----------------	---------



3. Run Sysprep. For details, see **Running Sysprep**.

NOTICE

If you use the **Unattend.xml** file created by yourself, check the **Unattend.xml** path when running Sysprep to ensure that the newly created **Unattend.xml** file is used.

4. Create an image from the ECS where Sysprep has been executed.

2.9 What Do I Do If I Cannot Create an Image in ZVHD2 Format Using an API?

Symptom

When you create a ZVHD2 image using an API, the image is created in the ZVHD format.

Solution

Check whether your token contains the **op_gated_lld** role (**op_gated_lld** is the OBT tag, which can be viewed in the body of the response message of the API used to obtain a user token). The ZVHD2 image has the lazy loading feature. If the current environment does not support this feature or this feature is in the OBT phase, the ZVHD2 image will fail to be created.

Contact the customer service to ensure that the current environment supports lazy loading. Obtain a new token and use the new token to create an image.

2.10 What Will the System Do to an Image File When I Use the File to Register a Private Image?

When you register a private image using an image file, if you enable automatic configuration, the system will perform the following operations.

Linux

- Check whether any PV drivers exist. If yes, the system deletes them.
- Modify the **grub** and **syslinux** configuration files to add OS kernel boot parameters and change disk partition names to UUID.
- Change disk partition names in the /etc/fstab file to UUID.
- Check whether the initrd file has Xen and IDE drivers. If no, the system will load the Xen and IDE drivers.
- Modify the X Window configuration file **/etc/X11/xorg.conf** to prevent display failures.
- Delete services of VMware tools.
- Record the latest automatic modification made to the image into /var/log/ rainbow_modification_record.log.
- Copy built-in VirtIO drivers to initrd or initramfs. For details, see External Image File Formats and Supported OSs.

NOTE

For the following image files, the system does not copy built-in VirtIO drivers after **Enable automatic configuration** is selected:

- Image files whose **/usr** directory is an independent partition
- Fedora 29 64bit, Fedora 30 64bit, and CentOS 8.0 64bit image files that use the XFS file system
- SUSE 12 SP4 64bit image files that use the ext4 file system

2.11 How Do I Configure an ECS, a BMS, or an Image File Before I Use It to Create an Image?

ECS or Image File Configurations

Table 2-1	ECS	configurations
-----------	-----	----------------

OS	Configuration
Linux	 Configuring DHCP (Optional) Installing and Configuring Cloud-Init Deleting Files from the Network Rule Directory Changing Disk Identifiers in the GRUB File to UUID Changing Disk Identifiers in the fstab File to UUID Installing Native Xen and KVM Drivers on a Xen ECS Installing Native KVM Drivers on a KVM ECS Detaching Data Disks from an ECS NOTE Before creating a Linux private image, ensure that the /etc/ fstab file does not contain data disk configuration. Otherwise, the cloud servers created from this image cannot start. If any data disks are attached, comment out or delete the configuration of these data disks from the /etc/fstab file.

Table 2-2	Image file	configurations
-----------	------------	----------------

OS	Configuration
Linux	 Deleting Files from the Network Rule Directory Configuring DHCP Installing Native KVM Drivers on a KVM ECS Changing Disk Identifiers in the GRUB File to UUID Changing Disk Identifiers in the fstab File to UUID
	 Detaching Data Disks from an ECS NOTE Before creating a Linux private image, ensure that the /etc/ fstab file does not contain data disk configuration. Otherwise, the ECSs created from this image cannot start. If any data disks are attached, comment out or delete the configuration of these data disks from the /etc/fstab file. (Optional) Installing and Configuring Cloud-Init

- When registering an external image file as a private image, you are advised to perform the preceding operations on the VM where the external image file is located.
- When registering a Windows external image file as a private image, if the Guest OS drivers are installed, the cloud platform will check the image file after you select **Enable automatic configuration**. If the GuestOS drivers are not installed, the cloud platform will try to install them.

BMS or Image File Configurations

Table 2-3 BMS	configurations
---------------	----------------

OS	Configuration	Reference
Windows	 Install the bms-network-config package. Install Cloudbase-Init. Delete residual files from the OS. 	Creating a Private Image from a BMS
Linux	 Install software in the bms- network-config package. Install Cloud-Init. Delete residual files from the OS. 	Creating a Private Image from a BMS

Table 2-4 intage file configurations	Table	2-4	Image	file	configurations
--------------------------------------	-------	-----	-------	------	----------------

OS	Configuration	Reference
Windows	 Install drivers for x86 V5 BMSs. Install Cloudbase-Init. 	Private Image Creation Guide
	 Install Cloudbase-Init. Install software in the bms- network-config package. 	
	 (Optional) Install the SDI iNIC driver. 	
	 (Optional) Install the one-click password reset plug-in. 	
	• Set the Windows time zone.	
	• Set the virtual memory.	
	• (Optional) Configure automatic Windows update.	
	Configure SID.	

OS	Configuration	Reference
Linux	 Install and configure Cloud-Init. Modify the hardware device driver that boots the OS. 	Private Image Creation Guide
	 Install software in the bms- network-config package. 	
	 (Optional) Install the SDI iNIC driver. 	
	 (Optional) Install the Hi1822 NIC driver. 	
	• (Optional) Install the IB driver.	
	 (Optional) Install drivers for x86 V5 BMSs. 	
	• (Optional) Install the UltraPath software.	
	• (Optional) Install the one-click password reset plug-in.	
	• Perform security configuration.	
	• Configure remote login to the BMS.	
	 Configure automatic root partition expansion. 	

2.12 What Do I Do If a Windows Image File Is Not Pre-Configured When I Use It to Register a Private Image?

If an image file is not configured as instructed in **Preparing an Image File** before it is exported from the original platform, you can use it to create an ECS, configure the ECS, and use the ECS to create a private image. **Figure 2-8** shows the process.

An ECS can run properly only after Xen Guest OS drivers (PV drivers) and KVM Guest OS drivers (VirtIO drivers) are installed on it. Without these drivers, the performance of this ECS will be affected and some functions will be unavailable. Ensure that the drivers have been installed for the image file before it is exported from the original platform. Otherwise, the ECSs created from the image will fail to start.

- To install PV drivers, see Installing PV Drivers.
- To install VirtIO drivers, see Installing VirtIO Drivers.





Step 1: Upload the Image File

Upload the image file to an OBS bucket. For details, see **Uploading an Image** File.

Step 2 Register the Image File as a Private Image

On the management console, select the uploaded image file and register it as a private image. For details, see **Registering an Image File as a Private Image**.

Step 3: Create an ECS

- 1. Access the IMS console.
 - a. Log in to the management console.
 - b. Under **Compute**, click **Image Management Service**. The IMS console is displayed.
- 2. Click the **Private Images** tab.
- 3. Locate the row that contains the private image and click **Apply for Server** in the **Operation** column.
- 4. Set parameters as prompted to create an ECS. Note that:
 - You are advised to select Pay-per-use for Billing Mode because the ECS will be deleted when you are finished with it.
 - Bind an EIP to the ECS so that you can upload installation packages to the ECS or download them from the ECS.
 - You must add inbound rules for the ECS security groups to ensure that the ECS can be accessed.

 If the image file has Cloudbase-Init installed, set a password and log in to the ECS using the password when prompted. If Cloudbase-Init is not installed, use the password or certificate contained in the image file to log in the ECS.

For details, see **Purchasing an ECS**.

- 5. Check the ECS to see if the private image used to create the ECS has been pre-configured.
 - a. Check whether the ECS can be successfully started. If it can, Guest OS drivers have been installed for the image file on the original platform or the drivers have been automatically installed for the private image on the cloud platform. If the ECS cannot start up, install Guest OS drivers for the image file on the original platform and go back to **Step 1: Upload the Image File**.
 - b. Check whether you can log in to the ECS using your configured password or key. If you can, Cloudbase-Init has been installed. If you cannot, use the password or key contained in the image file to log in to the ECS and install Cloudbase-Init as instructed in **Installing and Configuring Cloudbase-Init**.
 - c. Check whether DHCP is configured by referring to 2 in **Step 4: Configure the ECS**.
 - d. Use MSTSC to log in to the ECS. If the login fails, enable remote desktop connection by referring to **3** in **Step 4: Configure the ECS**.

If the ECS meets the preceding requirements, the private image has been preconfigured. Skip **Step 4: Configure the ECS** and **Step 5: Create a Private Image from the ECS**.

Step 4: Configure the ECS

Remotely log in to the ECS created in **Step 3: Create an ECS** to configure it.

- 1. Log in to the ECS.
- 2. Check whether DHCP is configured. If the ECS is configured with a static IP address, change its IP address assignment mode to DHCP as instructed in **Configuring DHCP**.
- 3. Enable remote desktop connection for the ECS as needed. For details about how to enable this function, see **Enabling Remote Desktop Connection**.
- 4. (Optional) Configure value-added functions.
 - Install and configure Cloudbase-Init. For details, see Installing and Configuring Cloudbase-Init.
 - Enable NIC multi-queue. For details, see **How Do I Enable NIC Multi-Queue for an Image?**
 - Configure an IPv6 address. For details, see How Do I Configure an ECS to Use DHCPv6?

Step 5: Create a Private Image from the ECS

Create a private image from the ECS. For details, see **Creating a Windows System Disk Image from a Cloud Server**.

(Optional) Clear the Environment

After the image registration is complete, delete the image file as well as the intermediate private image and ECS to avoid generating any additional charges.

- Delete the image registered in **Step 2 Register the Image File as a Private Image**.
- Delete the ECS created in Step 3: Create an ECS.
- Delete the image file from the OBS bucket.

2.13 What Do I Do If a Linux Image File Is Not Pre-Configured When I Use It to Register a Private Image?

If an image file is not configured as instructed in **Preparing an Image File** before it is exported from the original platform, you can use it to create an ECS, configure the ECS, and use the ECS to create a private image. **Figure 2-9** shows the process.

A CAUTION

An ECS can run properly only after Xen and KVM drivers are installed on it. If no such drivers are installed, the performance of the ECS will be affected and some functions will be unavailable. Ensure that KVM drivers have been installed for the image file before it is exported from the original platform. Otherwise, the ECSs created from the image will fail to start.

- For Xen, install native Xen and KVM drivers. For details, see How Do I Install Native Xen and KVM Drivers?
- For KVM, install native KVM drivers. For details, see Installing Native KVM Drivers.





Step 1: Upload the Image File

Upload the image file to an OBS bucket. For details, see **Uploading an Image** File.

Step 2 Register the Image File as a Private Image

On the management console, select the uploaded image file and register it as a private image. For details, see **Registering an Image File as a Private Image**.

Step 3: Create an ECS

Create an ECS from the private image.

- 1. Access the IMS console.
 - a. Log in to the management console.
 - b. Under Compute, click Image Management Service.

The IMS console is displayed.

- 2. Click the **Private Images** tab.
- 3. Locate the row that contains the private image and click **Apply for Server** in the **Operation** column.
- 4. Set parameters as prompted to create an ECS. Pay attention to the following:
 - You are advised to select **Pay-per-use** for **Billing Mode** because the ECS will be deleted when you finish with it.
 - You must add inbound rules for security groups of the ECS to ensure that the ECS can be accessed.

 If Cloud-Init has been installed in the image file, set a login password as prompted. If Cloud-Init is not installed, use the password or certificate contained in the image file to log in.

For details, see Purchasing an ECS.

- 5. Check the ECS to see if the private image used to create the ECS has been pre-configured.
 - a. Check whether the ECS can be successfully started. If the start succeeds, Xen and KVM drivers have been installed for the external image file on the original platform or the drivers have been automatically installed for the private image on the cloud platform. If the start failed, install Xen and KVM drivers for the image file and start from **Step 1: Upload the Image File** again.
 - b. Check whether you can log in to the ECS using your configured password or key. If you can, Cloud-Init has been installed. If you cannot, use the password or key in the image file to log in to the ECS and install Cloud-Init as instructed in Installing Cloud-Init.
 - c. Check the network configuration by referring to **Step 4: Configure the ECS**.

If the ECS meets the preceding requirements, the private image has been preconfigured. Skip **Step 4: Configure the ECS** and **Step 5: Create a Private Image from the ECS**.

Step 4: Configure the ECS

Remotely log in to the ECS created in Step 3: Create an ECS to configure it.

- 1. Log in to the ECS.
- 2. Configure the network.
 - Run the ifconfig command to check whether the private IP address of the ECS is the same as that displayed on the console. If they are inconsistent, delete files from the network rule directory as instructed in Deleting Files from the Network Rule Directory.
 - Check whether DHCP is configured. If the ECS is configured with a static IP address, change its IP address assignment to DHCP as instructed in Configuring DHCP.
 - Run the service sshd status command to check whether SSH is enabled.
 If it is disabled, run the service sshd start command to enable it. Ensure that your firewall (for example, Linux iptables) allows SSH access.
- 3. Configure a file system.
 - Change disk identifiers in the GRUB file to UUID. For details, see
 Changing Disk Identifiers in the GRUB File to UUID.
 - Change disk identifiers in the fstab file to UUID. For details, see
 Changing Disk Identifiers in the fstab File to UUID.
 - Clear the automatic mount configuration of non-system disks in the /etc/ fstab file. For details, see Detaching Data Disks from an ECS.
- 4. (Optional) Configure value-added functions.
 - Install and configure Cloud-Init. For details, see Installing and Configuring Cloud-Init.

- Enable NIC multi-queue. For details, see How Do I Enable NIC Multi-Queue for an Image?
- Configure an IPv6 address. For details, see How Do I Configure an ECS to Use DHCPv6?

Step 5: Create a Private Image from the ECS

Create a private image from the ECS. For details, see **Creating a System Disk Image from a Cloud Server**.

(Optional) Clear the Environment

After the image registration is complete, delete the image file as well as the intermediate private image and ECS to avoid generating any additional charges.

- Delete the image registered in Step 2 Register the Image File as a Private Image.
- Delete the ECS created in Step 3: Create an ECS.
- Delete the image file from the OBS bucket.

2.14 How Do I Make a System Disk Image Support Fast ECS Creation?

Scenarios

Fast Create greatly reduces the time required to create ECSs from a system disk image. Currently, this feature is supported by all newly created system disk images by default, but some existing system disk images may not support this feature. You can make them support it through image replication.

For example, if image A does not support fast ECS creation, you can replicate it to generate image copy_A that supports fast ECS creation.

Constraints

Full-ECS images and ISO images cannot be configured using this method.

Check Whether an Image Supports Fast ECS Creation

- 1. Access the IMS console.
 - a. Log in to the management console.
 - b. Under **Compute**, click **Image Management Service**. The IMS console is displayed.
- 2. Click the **Private Images** tab to display the image list.
- 3. Click the name of the target image.
- 4. On the displayed image details page, check the value of **Fast ECS Creation**.

Name	newCloudinit-Image	ID	e587771f-9920-456f-a987-b00792c4d326
Image Type	System disk image	Disk Capacity (GB)	40
OS	Windows Server 2016 Datacenter 64bit	Status	🥑 Normal
Minimum Memory	Unlimited	Image Size	7.21 GB
Maximum Memory	Unlimited	Released	No
Created	2023/01/28 17:30:38 GMT+08:00	Completed	2023/01/28 17:38:58 GMT+08:00
Source	ECS	Description	
Encrypted	No	OS Type	Windows
NIC Multi-Queue	Supported	Enterprise Project	default
Fast ECS Creation	Supported		

Figure 2-10 Image details

Configure an Image to Make It Support Fast ECS Creation

1. Locate the system disk image, click **More** in the **Operation** column, and select **Replicate** from the drop-down list.

The **Replicate Image** dialog box is displayed.

- 2. Set **Replication Mode** to **Within Region** and set other parameters based on **Replicating Images Within a Region**.
- 3. Wait for the replication to complete. Then, the new image can be used to quickly create ECSs.

3 os

3.1 How Do I Select an OS?

• Linux

Used for development platforms or services that run Linux. CentOS and Ubuntu are provided. CentOS is recommended.

The system disk must be at least 40 GB, and there must be at least 512 MB of memory.

• If your servers require more than 4 GB of memory, select a 64-bit OS because 4 GB is the maximum memory a 32-bit OS can access.

3.2 How Is BIOS Different from UEFI?

Boot Mode	Description	Highlight
BIOS	Basic Input Output System (BIOS) stores important basic input/output programs of ECSs, system settings, self-test programs upon system startup, and automatic startup programs.	Provides basic settings and control for ECSs.
UEFI	Unified Extensible Firmware Interface (UEFI) is a specification that defines a software interface between an OS and platform firmware. UEFI can be used to automatically load an OS from a pre-boot operating environment.	Boots up or recovers from sleep state faster.

Table 3-1 Differences between the UEFI and BIOS boot modes

3.3 What Do I Do If I Cannot Change the Resolution of a Windows OS Booted in UEFI Mode?

Symptoms

The resolution of a Windows OS booted in UEFI mode cannot be changed in display settings.

Possible Causes

By default, x86 Windows cloud servers on Huawei Cloud are booted in BIOS mode. If users migrate Windows servers from other vendors or a local IDC, the servers may be booted in UEFI mode. The resolution of servers booted in UEFI mode cannot be changed in display settings.

Solutions

1. Log in to the Windows cloud server using VNC. Open the **Start** menu, click the **Power** button, and choose **Restart**. During the restart, press **Esc** on your keyboard quickly and repeatedly to enter the BIOS configuration.



2. Use the up and down arrow keys on the keyboard to select **Device Manager** and press **Enter**.

	Device Manager	
Devices List Driver Health Manager RAM Disk Configuration OUMF Platform Configuration iSCSI Configuration Network Device List Press ESC to exit.		List all the Driver Health instances to manage
†∔=Move Highlight <ente< td=""><td>r>=Select Entry Es</td><td>c=Exit</td></ente<>	r>=Select Entry Es	c=Exit

3. Use the up and down arrow keys to select **OVMF Platform Configuration** and press **Enter**.

	Device Manager	
Devices List Driver Health Manager RAM Disk Configuration JUMF Platform Configuration iSCSI Configuration Network Device List Press ESC to exit.		Change various OVMF platform settings.
†↓=Move Highlight	<enter≻=select entry<="" td=""><td>Esc=Exit</td></enter≻=select>	Esc=Exit

4. Use the up and down arrow keys to select **Change Preferred** (it is selected by default).

OVMF Settings			
Preferred Resolution at Next Boot Change Preferred Resolution for Next Boot Commit Changes and Exit Discard Changes and Exit		You can specify a new preference for the Graphics Console here. The list is filtered against the video RAM size.	
	9=Reset to Defaults Enter>=Select Entry	F10=Save Esc=Exit	

5. In the **Change Preferred** line, use the up and down arrow keys to change the resolution (for example, the resolution is changed from 800x600 to 1280x760), and press **Enter**.

OVMF Settings			
Preferred Resolution at Next Boot Change Preferred Resolution for Next Boot Commit Changes and Exit Discard Changes and Exit		You can specify a new preference for the Graphics Console here. The list is filtered against the video RAM size.	
	9=Reset to Defaults Enter>=Select Entry	F10=Save Esc=Exit	
		Configuration changed	

6. Use the up and down arrow keys to select **Commit Changes and Exit** and press **Enter** to save the settings and exit. Then, press **Esc** to exit the BIOS configuration.

7. Use the up and down arrow keys to select **Reset** and press **Enter** to restart the cloud server.

NOTE

You must select **Reset**. If you select **Continue**, you need to restart the cloud server again to make the BIOS configuration take effect.

OpenStack Nova pc-i440fx-2.8 0.0.0		2.00 GHz 8192 MB RAM
Select Language	<standard english=""></standard>	Reset
 Device Manager Boot Manager Boot Maintenance Manager 		
Continue Reset		
↑↓=Move Highlight <e< td=""><td>nter>=Select Entry</td><td></td></e<>	nter>=Select Entry	

8. Log in to the cloud server and check the desktop resolution.

After logging in to the cloud server, right-click the desktop and choose **Display settings** from the shortcut menu. The desktop resolution is 1280x760 which is the same as that configured in BIOS.

3.4 How Do I Delete Redundant Network Connections from a Windows ECS?

Method 1

1. Press **Win+R**. In the displayed dialog box, enter **regedit** and press **Enter** to open the registry editor.

NOTE

Modifying a registry may cause a system startup failure. So, back up the registry before modifying it.

2. Open the following registry key:

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT \CurrentVersion\NetworkList\Profiles

Click each item under **Profiles** and query the **Data** column of **ProfileName** in the right pane.

- 3. Double-click **ProfileName** and set **Value Data** to the name of a new network.
- 4. Restart the ECS for the change to take effect.
Method 2

1. Press **Win+R**. In the displayed dialog box, enter **regedit** and press **Enter** to open the registry editor.

NOTE

Modifying a registry may cause a system startup failure. So, back up the registry before modifying it.

- 2. Open the following registry keys:
 - HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT \CurrentVersion\NetworkList\Profiles HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT \CurrentVersion\NetworkList\Signatures\Unmanaged
- 3. Delete the directories shown in the following figure.

Figure 3-1 Registry directory



3.5 What Do I Do If an ECS Starts Slowly?

Symptom

If an ECS starts slowly, you can change the default timeout duration to speed up the startup.

Solution

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

sudo su

 Run the following command to query the version of the GRUB file: rpm -qa | grep grub

Figure 3-2 Querying the GRUB file version

[root@]#	rpm	-qa	grep	grub
grub2-2.02-0.44.el7.centos.x8					

- 4. Set **timeout** in the GRUB file to **0**.
 - If the GRUB file version is earlier than 2: Open **/boot/grub/grub.cfg** or **/boot/grub/menu.lst** and set **timeout** to **0**.
 - If the GRUB file version is 2:
 - Open /boot/grub2/grub.cfg and set the value of timeout to 0.

Figure 3-3 Modifying the timeout duration



3.6 Why Can't I Find My Private Image When I Want to Use It to Create an ECS or Change the OS of an ECS?

When you create an ECS or change the OS of an existing ECS, some of your private images are not shown. One possible cause is that the x86 and Arm architectures are incompatible with each other, or that there is an incompatibility issue between UEFI and BIOS boot modes.

- If a private image is created from a x86 ECS, this image will be invisible to you when you create an Arm (Kunpeng) ECS or change the OS of an Arm (Kunpeng) ECS, and vice versa.
- If you use an external image file to create a private image and select the x86 architecture, this image will be invisible to you when you create an Arm (Kunpeng) ECS or change the OS of an Arm (Kunpeng) ECS, and vice versa.
- If a private image is created from an ECS in BIOS boot mode, this image will be invisible to you when you create an ECS in UEFI boot mode or change the OS of an ECS in UEFI boot mode, and vice versa.
- If you use an external image file to create a private image and select the BIOS boot mode, this image will be invisible to you when you create an ECS in UEFI boot mode or change the OS of an ECS in UEFI boot mode, and vice versa.

3.7 What Do I Do If CentOS Linux Is No Longer Maintained?

CentOS has planned to stop maintaining CentOS Linux. Huawei Cloud will stop providing CentOS Linux public images. This section describes the impacts and tells you how to address the situation.

Background

On December 8, 2020, CentOS announced its plan to stop maintaining CentOS Linux and launched CentOS Stream. For more information, see CentOS Project shifts focus to CentOS Stream.

CentOS Linux 8 ended on December 31, 2021, and CentOS Linux 7 will end on June 30, 2024. CentOS Linux 9 and later versions will not be released, and patches will no longer be updated.

Impacts

CenterOS Linux users will be affected as follows:

- After December 31, 2021, CentOS Linux 8 users will not be able to obtain any maintenance or support services, including problem fixing and function updates.
- After June 30, 2024, CentOS Linux 7 users will not be able to obtain any maintenance or support services, including problem fixing and function updates.

Huawei Cloud users will be affected as follows:

- CentOS Linux 8 public images will continue for a certain time. ECSs created from CentOS Linux 8 images will not be affected, but the images will no longer be updated.
- Huawei Cloud will synchronize with CentOS for the support of CentOS Linux. After December 31, 2021, support services will no longer be available for CentOS 8. The support for CentOS 7 will continue until June 30, 2024.

Solution

You can change the OS so that the services originally running in CentOS Linux can continue to run in other OSs.

For details about how to change an OS, see Changing the OS.

ltem	Precautions for Changing an OS
Data backup	• Data in all partitions of the system disk will be cleared, so you are advised to back up the system disk data prior to an OS change.
	Data in data disks remains unchanged.
Custom settings	After the OS is changed, custom settings such as DNS and hostname will be reset and need to be reconfigured.

Table 3-2 Precautions for changing an OS

Table 3-3 Available OSs

OS	Description	Intended User
Rocky Linux	Rocky Linux is a community-driven enterprise-class OS. It is a downstream release of Red Hat Enterprise Linux (RHEL). Rocky Linux is fully compatible with and as stable as CentOS.	Individuals or enterprises that want to continue to use free images in an open source community
AlmaLinux	AlmaLinux is an open-source, community-driven Linux distribution developed by the CloudLinux team. It fills the gap left by the discontinuation of the CentOS Linux stable release. AlmaLinux can remain 1:1 compatibility with upstream RHEL. You can use it for OS changes without stopping your server.	Individuals or enterprises that want to continue to use free images in an open source community
Ubuntu	They are Linux distributions that differ in use and compatibilities.	Individuals or enterprises that can afford the OS change costs

3.8 How Do I Enable NIC Multi-Queue for an Image?

Scenarios

Network I/O bandwidth can keep increasing to the point where a single vCPU cannot process all of the NIC interrupts. NIC multi-queue allows multiple vCPUs to process NIC interrupts, thereby improving network PPS and I/O performance.

ECSs Supporting NIC Multi-Queue

NIC multi-queue can only be enabled on an ECS with the specifications and virtualization type described in this section.

• For details about the ECS specifications that support NIC multi-queue, see ECS Types.

NOTE

If there are more than 1 NIC queue, NIC multi-queue is supported.

• Only KVM ECSs support NIC multi-queue.

Operation Instructions

If an ECS has the required specifications and virtualization type, perform the following operations to enable NIC multi-queue:

NOTE

Currently, NIC multi-queue is only available for Ethernet NICs.

- 1. Register the External Image File as a Private Image.
- 2. Enable NIC Multi-Queue for the Image.
- 3. Create an ECS from the Private Image.
- 4. Run a Script to Enable NIC Multi-Queue on the ECS.

Register the External Image File as a Private Image

For details, see **Registering an Image File as a Private Image (Linux)**. After the image file is imported, view the value of **NIC Multi-Queue** on the page providing details about the image.

- If the value is Supported, go to Create an ECS from the Private Image.
- If the value is Not supported, go to Enable NIC Multi-Queue for the Image.

Enable NIC Multi-Queue for the Image

Use any of the following methods to enable NIC multi-queue for an image:

Method 1:

- 1. Access the IMS console.
 - a. Log in to the management console.
 - b. Under **Compute**, click **Image Management Service**. The IMS console is displayed.
- 2. On the displayed **Private Images** page, locate the row that contains the image and click **Modify** in the **Operation** column.
- 3. Enable NIC multi-queue for the image.

Method 2:

- 1. Access the IMS console.
 - a. Log in to the management console.
 - b. Under **Compute**, click **Image Management Service**. The IMS console is displayed.
- 2. On the displayed **Private Images** page, click the name of the image.

3. In the upper right corner of the displayed image details page, click **Modify**. In the displayed **Modify Image** dialog box, enable NIC multi-queue for the image.

Method 3: Add hw_vif_multiqueue_enabled to the image using an API.

- 1. Obtain a token. For details, see Authentication.
- 2. Call an API to update image information. For details, see **Updating Image Information (Native OpenStack API)**.
- Add X-Auth-Token to the request header. The value of X-Auth-Token is the token obtained in step 1.
- 4. Add **Content-Type** to the request header.

Set Content-Type to application/openstack-images-v2.1-json-patch.

The request URI is in the following format:

PATCH /v2/images/{*image_id*}

The request body is as follows:

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

Create an ECS from the Private Image

Γ

1

Use the registered private image to create an ECS. For details, see **Purchasing an ECS**. Note the following when setting the parameters:

- **Region**: Select the region where the private image is located.
- **Image**: Select **Private image** and then the desired image from the drop-down list.

Run a Script to Enable NIC Multi-Queue on the ECS

For Linux, a script is provided to automatically enable NIC multi-queue. If an ECS has multiple NICs, the script will automatically enable multi-queue for all of them.

1. Log in to the ECS and run the following command to check the number of queues supported by and enabled for a NIC:

ethtool -l N/C

Example output:

```
[root@localhost ~]# ethtool -l eth0 # Number of gueues used by NIC eth0
Channel parameters for eth0:
Pre-set maximums:
RX:
            0
TX:
            0
Other:
             0
               4 # The NIC supports a maximum of four queues.
Combined:
Current hardware settings:
RX:
            0
TX:
            0
Other:
             0
               1 # One queue has been enabled for the NIC.
Combined:
```

If the values of the two **Combined** fields are the same, NIC multi-queue has been enabled. No further action is required.

 Download the configuration script multi-queue-hw.
 wget https://ecs-instance-driver.obs.cn-north-1.myhuaweicloud.com/ multi-queue-hw
 The download address is https://ocs_instance_driver.obs.cn

The download address is https://ecs-instance-driver.obs.cnnorth-1.myhuaweicloud.com/multi-queue-hw.

3. Make the script executable.

chmod +x multi-queue-hw

4. Move the **multi-queue-hw** script into the **/etc/init.d** directory.

mv multi-queue-hw /etc/init.d

Enter **y** when prompted: mv: overwrite '/etc/init.d/multi-queue-hw'?

5. Run the script.

/etc/init.d/multi-queue-hw start

NIC multi-queue is enabled immediately after the script is executed but is disabled again if the ECS stops.

- 6. Enable NIC multi-queue at startup.
 - For CentOS, Red Hat, Fedora, EulerOS, SUSE, and openSUSE:

chkconfig multi-queue-hw on

– For Ubuntu:

update-rc.d multi-queue-hw defaults 90 10

For Debian:
 systemctl enable multi-queue-hw

3.9 How Do I Configure an ECS to Use DHCPv6?

Scenarios

IPv6 addresses are used to deal with IPv4 address exhaustion. If an ECS uses an IPv4 address, the ECS can run in dual-stack mode after IPv6 is enabled for it. Then, the ECS will have two IP addresses to access the intranet and Internet: an IPv4 address and an IPv6 address.

In some cases, an ECS cannot dynamically acquire an IPv6 address even if it meets all the requirements in **Constraints**. You need to configure the ECS to dynamically acquire IPv6 addresses. For public images:

- By default, dynamic IPv6 address assignment is enabled for Windows public images. You do not need to configure it. The operations in Windows Server 2012 and Windows Server 2008 are for your reference only.
- Before enabling dynamic IPv6 address assignment for a Linux public image, check whether IPv6 has been enabled and then whether dynamic IPv6 address assignment has been enabled. Currently, IPv6 is enabled for all Linux public images.

Constraints

• Ensure that IPv6 has been enabled on the subnet where the ECS works.

If IPv6 is not enabled, enable it by referring to **Enabling IPv6 for a Network Interface**. Once enabled, IPv6 cannot be disabled.

• Ensure that the ECS flavor supports IPv6.

The ECS flavors that support IPv6 vary depending on regions and AZs. Check whether an ECS flavor supports IPv6 after you select a region and AZ on the management console.

If the value of IPv6 is Yes for an ECS flavor, the flavor supports IPv6.

NOTE

AZ and Flavor determine whether IPv6 is supported.

After you select an AZ, if **IPv6** is not displayed or the value of **IPv6** is **No**, IPv6 is not supported by any or certain flavors in the AZ.

- Ensure that **Automatically-assigned IPv6 address** is selected during ECS creation.
- After the ECS is started, its hot-swappable NICs cannot automatically acquire IPv6 addresses.
- Only ECSs can work in dual-stack mode and BMSs cannot.
- Only one IPv6 address can be bound to a NIC.
- Check that the ECS network configuration is correct.

For details about how to check the network configuration, see **Checking the ECS Network Configuration**.

If the network configuration is incorrect, **submit a service ticket**.

Procedure

- Windows: Windows Server 2012 is used as an example to describe how to enable dynamic assignment of IPv6 addresses in Windows, as shown in Table 3-4.
- Linux: Dynamic assignment of IPv6 addresses can be enabled automatically (recommended) or manually, as shown in Table 3-4.

If a private image created from a CentOS 6.x or Debian ECS with automatic IPv6 address assignment enabled is used to create an ECS in an environment that does not support IPv6, the ECS may start slow because of IPv6 address assignment timeout. You can set the timeout duration for assigning IPv6 addresses by referring to **Setting the Timeout Duration for IPv6 Address Assignment**.

Table 3-4 Enabling dynamic assignment of IPv6 addresses for different OSs

OS	Auto/Manual	Reference
Windows Server 2012	Auto	Windows Server 2012
Windows Server 2008	Auto	Windows Server 2008

OS	Auto/Manual	Reference
Linux	Auto (recommended)	Linux (Automatically Enabling Dynamic Assignment of IPv6 Addresses)
Linux	Manual	Linux (Manually Enabling Dynamic Assignment of IPv6 Addresses)

Windows Server 2012

Step 1 Check whether IPv6 is enabled for the ECS.

Run the following command in the CMD window:

ipconfig

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

Figure 3-4 Querying the IPv6 address

Σ	Administrator: Windows PowerShell	– – X
Windows IP Configuration		^
Ethernet adapter Ethernet	2:	
Connection-specific DN IPv6 Address	S Suffix . : openstacklocal : s :	
IPv4 Address		
Default Gateway	an a	

• If only a link-local IPv6 address is displayed, IPv6 is enabled but dynamic IPv6 assignment is not enabled. Go to **Step 2**.

Figure 3-5 Link-local IPv6 address

Z	Administrator: Windows PowerShell	_ D X
Windows IP Configuration		^
Ethernet adapter Ethernet		
Link-local IPv6 Addres: IPv4 Address	an an an the part of the second s	
Subnet Mask Default Gateway		

• If neither an IPv6 address nor link-local IPv6 address is displayed, IPv6 is disabled. Go to **Step 3**.

Figure 3-6 IPv6 disabled

	Administrator: C:\Windows\system32\cmd.exe					
	C:\Users\Administrator>ipconfig					
	Windows IP Configuration					
	Ethernet adapter Local Area Connection					
l	Connection-specific DNS Suffix . :					
l	IPv4 Address					
	Subnet Mask					
	Connection-specific DNS Suffix . Link-local IPv6 Address IPv4 Address					

NOTE

By default, dynamic IPv6 address assignment is enabled for Windows public images, as shown in **Figure 3-4**. No additional configuration is required.

Step 2 Enable dynamic IPv6 address assignment.

- 1. Choose **Start** > **Control Panel**.
- 2. Click Network and Sharing Center.
- 3. Click the Ethernet connection.

Figure 3-7 Ethernet connection

¥.	Network	and Sharing Center
🔄 💮 ▾ ↑ 🕎 ኑ Control I	Panel ► All Control Panel Items ► Network a	and Sharing Center 🛛 🗸 🖒
Control Panel Home	View your basic network inform	nation and set up connections
Change adapter settings	View your active networks	
Change advanced sharing	Network	Access type: Internet
settings	Public network	Connections: U Ethernet 2

- 4. In the **Ethernet Status** dialog box, click **Properties** in the lower left corner.
- 5. Select Internet Protocol Version 6 (TCP/IPv6) and click OK.



Figure 3-8 Configuring dynamic IPv6 address assignment

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

Step 3 Enable and configure IPv6.

- 1. In the **Internet Protocol Version 6 (TCP/IPv6) Properties** dialog box, configure an IPv6 address and a DNS server address.
 - IPv6 address: IPv6 address allocated during ECS creation. Obtain the value from the ECS list on the console.
 - Subnet prefix length: 64
 - Preferred DNS server: 240c::6666 (recommended)

Internet Protocol Version 6 (TCP/IPv6) Properties				
General				
You can get IPv6 settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IPv6 settings.				
Obtain an IPv6 address autom	natically			
 Use the following IPv6 address 	5:			
IPv6 address:				
Subnet prefix length:	64			
Default gateway:				
Obtain DNS server address automatically				
Use the following DNS server addresses:				
Preferred DNS server:	240c::6666			
Alternate DNS server:				
Ualidate settings upon exit	Advanced			
	OK Cancel			

Figure 3-9 Configuring an IPv6 address and a DNS server address

- (Optional) Run the following command depending on your ECS OS.
 For Windows Server 2012, run the following command in PowerShell or CMD:
 Set-NetIPv6Protocol -RandomizeIdentifiers disabled
- 3. Perform **Step 1** to check whether dynamic IPv6 address assignment is enabled.

----End

Windows Server 2008

Step 1 Check whether IPv6 is enabled for the ECS.

Run the following command in the CMD window:

ipconfig

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



Windows IP Configuration Ethernet adapter Local Area Connection 3: Connection-specific DNS Suffix .: IPv6 Address Link-local IPv6 Address IPv4 Address Default Gateway	Administrator: C:\Windows\system32\cmd.exe							
Connection-specific DNS Suffix .: IPv6 Address Link-local IPv6 Address IPv4 Address Subnet Mask	Windows IP Configuration							
IPv6 Address Link-local IPv6 Address IPv4 Address Subnet Mask	Ethernet adapter Local Area Connection 3:							
Link-local IPv6 Address : IPv4 Address : Subnet Mask :								
Subnet Mask								
Nefault Gateway								
borauto automay	Default Gateway							

• If only a link-local IPv6 address is displayed, IPv6 is enabled but dynamic IPv6 assignment is not enabled. Go to **Step 2**.

Figure 3-11 Link-local IPv6 address



 If neither an IPv6 address nor link-local IPv6 address is displayed, IPv6 is disabled. Go to Step 3.

Figure 3-12 IPv6 disabled

a Administrator: C:\Windows\system32\cmd.exe						
C:\Users\Administrator>ipconfig Windows IP Configuration						
Ethernet adapter Local Area Connection						
Connection-specific DNS Suffix . :						
Link-local IPv6 Address						
IPv4 Address						
Subnet Mask						
Default Gateway						

NOTE

By default, dynamic IPv6 address assignment is enabled for Windows public images, as shown in **Figure 3-10**. No additional configuration is required.

Step 2 Enable dynamic IPv6 address assignment.

1. Choose **Start > Control Panel**.

- 2. Click Network and Sharing Center.
- 3. Click Change adapter settings.
- 4. Right-click the local network connection and choose **Properties**.
- 5. Select Internet Protocol Version 6 (TCP/IPv6) and click OK.

Figure 3-13 Configuring dynamic IPv6 address assignment

🖣 Local Area Connection 3 Properties	×	
Networking		
Connect using:		
🔮 Red Hat VirtIO Ethernet Adapter		
Configure	וב	
This connection uses the following items:	-	
 Client for Microsoft Networks QoS Packet Scheduler File and Printer Sharing for Microsoft Networks Internet Protocol Version 6 (TCP/IPv6) Internet Protocol Version 4 (TCP/IPv4) Internet Protocol Version 9 (TCP/IPv4) Link-Layer Topology Discovery Mapper I/O Driver Link-Layer Topology Discovery Responder 		
Install Uninstall Properties		
Description TCP/IP version 6. The latest version of the internet protocol that provides communication across diverse interconnected networks.		
OK Canc	el	

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

Step 3 Enable and configure IPv6.

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

上 Local Area Connection Properties 🏾 🔋 🗙		
General Authentication Advanced		
Connect using:		
Bealtek RTL8139 Family PCI Fast Et		
This connection uses the following items:		
 Client for Microsoft Networks Betwork Load Balancing File and Printer Sharing for Microsoft Networks Ternet Protocol (TCP/IP) 		
Install Uninstall Properties Description Allows your computer to access resources on a Microsoft network.		
 Show icon in notification area when connected Notify me when this connection has limited or no connectivity 		
OK Cancel		

Figure 3-14 Enabling and configuring IPv6

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

Figure 3-15 Adding the protocol



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

Select Ne	twork Protocol
- <u>}</u>	Click the Network Protocol that you want to install, then click OK. If you have an installation disk for this component, click Have Disk.
App Mic Net	rk Protocol: bleTalk Protocol rosoft TCP/IP persion 6 twork Monitor Driver /Link IPX/SPX/NetBIOS Compatible Transport Protocol iable Multicast Protocol
	s driver is digitally signed. I me why driver signing is important OK

(Optional) Run the following commands depending on your ECS OS.
 For Windows Server 2008, run the following command in PowerShell or CMD:

netsh interface ipv6 set global randomizeidentifiers=disable

Disable the local connection and then enable it again.

To disable the local connection, choose **Start > Control Panel > Network and Internet > Network and Sharing Center > Change Adapter Options**. Rightclick the local connection and choose **Disable** from the shortcut menu.

To enable the local connection, choose **Start > Control Panel > Network and Internet > Network and Sharing Center > Change Adapter Options**. Rightclick the local connection and choose **Enable** from the shortcut menu.

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

----End

Linux (Automatically Enabling Dynamic Assignment of IPv6 Addresses)

The **ipv6-setup**-*xxx* tool can be used to enable Linux OSs to automatically acquire IPv6 addresses. *xxx* indicates a tool, which can be rhel or debian.

You can also enable dynamic IPv6 address assignment by following the instructions in Linux (Manually Enabling Dynamic Assignment of IPv6 Addresses).

- When you run **ipv6-setup**-*xxx*, the network service will be automatically restarted. As a result, the network is temporarily disconnected.
- If a private image created from a CentOS 6.x or Debian ECS with automatic IPv6 address assignment enabled is used to create an ECS in an environment that does not support IPv6, the ECS may start slow because of IPv6 address assignment timeout. Set the timeout duration for assigning IPv6 addresses to 30s by referring to Setting the Timeout Duration for IPv6 Address Assignment and try to create a new private image again.
- **Step 1** Run the following command to check whether IPv6 is enabled for the ECS:

ip addr

• If only an IPv4 address is displayed, IPv6 is disabled. Enable it by referring to **Setting the Timeout Duration for IPv6 Address Assignment**.

Figure 3-17 IPv6 disabled



• If a link-local address (starting with fe80) is displayed, IPv6 is enabled but dynamic assignment of IPv6 addresses is not enabled.

Figure 3-18 IPv6 enabled



• If the following address is displayed, IPv6 is enabled and an IPv6 address has been assigned:

Figure 3-19 IPv6 enabled and an IPv6 address assigned

eth0: <broadcast,multicast,up,lower_up> mtu 1500 qdisc mq state UP group default qlen 1000</broadcast,multicast,up,lower_up>			
link/ether fa:16:3e:75:af:4c brd ff:ff:ff:ff:ff:ff			
inet brd scope global noprefixroute dynamic eth0			
valid 1ft 86395sec preferred 1ft 86395sec			
inet6 2407:c080:802: /128 scope global dynamic			
valid_lft 7496sec preferred_lft 7196sec			
<pre>inet6 fe80::f816:3eff: /64 scope link noprefixroute</pre>			
valid_lft forever preferred_lft forever			

NOTE

IPv6 is enabled for Linux public images by default, as shown in Figure 3-18.

- **Step 2** Enable IPv6 for the ECS.
 - Run the following command to check whether IPv6 is enabled for the kernel: sysctl -a | grep ipv6
 - If a command output is displayed, IPv6 is enabled.
 - If no information is displayed, IPv6 is disabled. Go to Step 2.2 to load the IPv6 module.

- Run the following command to load the IPv6 module: modprobe ipv6
- Add the following content to the /etc/sysctl.conf file: net.ipv6.conf.all.disable_ipv6=0
- 4. Save the configuration and exit. Then, run the following command to load the configuration:

sysctl -p

- Step 3 Enable dynamic IPv6 address assignment for the ECS.
 - 1. Download **ipv6-setup-rhel** or **ipv6-setup-debian** with a required version and upload it to the target ECS.

ipv6-setup-*xxx* modifies the configuration file of a NIC to enable dynamic IPv6 address assignment or adds such a configuration file for a NIC, and then restarts the NIC or network service. **Table 3-5** lists the paths for obtaining **ipv6-setup-rhel** and **ipv6-setup-debian**.

Series	Release Version	How to Obtain
RHEL	 CentOS 6/7 EulerOS 2.2/2.3 Fedora 25 	https://ecs-instance- driver.obs.cn- north-1.myhuaweicloud.com/ ipv6/ipv6-setup-rhel
Debian	– Ubuntu 16/18 – Debian 8/9/10	https://ecs-instance- driver.obs.cn- north-1.myhuaweicloud.com/ ipv6/ipv6-setup-debian

 Table 3-5 Download paths of ipv6-setup-rhel and ipv6-setup-debian

- Run the following command to make ipv6-setup-xxx executable: chmod +x ipv6-setup-xxx
- 3. Run the following command to enable dynamic IPv6 address assignment for a NIC:

./ipv6-setup-*xxx* --dev [*dev*]

Example:

./ipv6-setup-xxx --dev eth0

D NOTE

- To enable dynamic IPv6 address assignment for all NICs, run the **./ipv6-setup-***xxx* command.
- To learn how to use **ipv6-setup**-*xxx*, run the **./ipv6-setup**-*xxx* --**help** command.

----End

3 OS

Linux (Manually Enabling Dynamic Assignment of IPv6 Addresses)

If a private image created from a CentOS 6.x or Debian ECS with automatic IPv6 address assignment enabled is used to create an ECS in an environment that does not support IPv6, the ECS may start slow because of IPv6 address assignment timeout. Set the timeout duration for assigning IPv6 addresses to 30s by referring to **Setting the Timeout Duration for IPv6 Address Assignment** and try to create a new private image again.

Step 1 Run the following command to check whether IPv6 is enabled for the ECS:

ip addr

• If only an IPv4 address is displayed, IPv6 is disabled. Enable it by referring to **Step 2**.

Figure 3-20 IPv6 disabled



• If a link-local address (starting with fe80) is displayed, IPv6 is enabled but dynamic assignment of IPv6 addresses is not enabled.

Figure 3-21 IPv6 enabled



• If the following address is displayed, IPv6 is enabled and an IPv6 address has been assigned:

Figure 3-22 IPv6 enabled and an IPv6 address assigned



NOTE

IPv6 is enabled for Linux public images by default, as shown in Figure 3-21.

- **Step 2** Enable IPv6 for the ECS.
 - Run the following command to check whether IPv6 is enabled for the kernel: sysctl -a | grep ipv6
 - If a command output is displayed, IPv6 is enabled.
 - If no information is displayed, IPv6 is disabled. Go to Step 2.2 to load the IPv6 module.

- Run the following command to load the IPv6 module: modprobe ipv6
- Add the following content to the /etc/sysctl.conf file: net.ipv6.conf.all.disable_ipv6=0
- 4. Save the configuration and exit. Then, run the following command to load the configuration:

sysctl -p

- Step 3 Enable dynamic IPv6 address assignment for the ECS.
 - Ubuntu 18.04/20.04
 - a. Run the following command to access /etc/netplan/:
 cd /etc/netplan
 - b. Run the following command to list the configuration file:
 ls

Figure 3-23 Configuration file name



c. Run the following command to edit the configuration file **01-networkmanager-all.yaml**:

vi 01-network-manager-all.yaml

Append the following content to the configuration file **01-network-manager-all.yaml** (pay attention to the YAML file format and text indentation):
 ethernets:
 eth0:
 dhcp6: true

Figure 3-24 Edited configuration file

```
# Let NetworkManager manage all devices on thin system
network:
    version: 2
    renderer: NetworkManager
    ethernets:
    eth0:
        dhcp6: true
```

Save the changes and exit.

- e. Run the following command to make the changes take effect: **sudo netplan apply**
- Ubuntu 22.04
 - a. Run the following command to access /etc/netplan/:
 cd /etc/netplan
 - b. Run the following command to list the configuration file:
 ls

Figure 3-25 Configuration file name



c. Run the following command to edit the configuration file **01**-**netcfg.yaml**:

vi 01-netcfg.yaml

Append the following content to the configuration file **01-netcfg.yaml** (pay attention to the YAML file format and text indentation):
 ethernets:
 eth0:
 dhcp6: true

Figure 3-26 Edited configuration file



Save the changes and exit.

- e. Run the following command to make the changes take effect: **sudo netplan apply**
- f. Run the following command to edit /etc/NetworkManager/ NetworkManager.conf:

vi /etc/NetworkManager/NetworkManager.conf

g. Append the following content to the configuration file **NetworkManager.conf** (pay attention to the file format and indentation): [main]

plugins=ifupdown,keyfile dhcp=dhclient

[ifupdown] managed=true

[device] wifi.scan-rand-mac-address=no

Figure 3-27 Modification result

```
[main]
plugins=ifupdown,keyfile
dhcp=dhclient
[ifupdown]
managed=true
[device]
wifi.scan-rand-mac-address=no
```

- h. Run the following command for the configuration to take effect: systemctl restart NetworkManager
- Debian
 - a. Add the following content to the /etc/network/interfaces file:

auto lo iface lo inet loopback auto **eth0** iface **eth0** inet dhcp iface eth0 inet6 dhcp pre-up sleep 3

- Add configurations for each NIC to the /etc/network/interfaces file. The following uses eth1 as an example: auto eth1 iface eth1 inet dhcp iface eth1 inet6 dhcp pre-up sleep 3
- c. Run the following command to restart the network service:

service networking restart

If no IPv6 address is assigned after the NICs are brought down and up, you can run this command to restart the network.

- d. Perform **Step 1** to check whether dynamic IPv6 address assignment is enabled.
- CentOS, EulerOS, or Fedora
 - a. Open the configuration file **/etc/sysconfig/network-scripts/ifcfg-eth0** of the primary NIC.

Add the following configuration items to the file: IPV6INIT=yes DHCPV6C=yes

- b. Edit the **/etc/sysconfig/network** file to add or modify the following line: NETWORKING_IPV6=yes
- c. For an ECS running CentOS 6, you need to edit the configuration files of its extension NICs. For example, if the extension NIC is eth1, you need to edit **/etc/sysconfig/network-scripts/ifcfg-eth1**.

Add the following configuration items to the file: IPV6INIT=yes DHCPV6C=yes

In CentOS 6.3, dhcpv6-client requests are filtered by **ip6tables** by default. So, you also need to add a rule allowing the dhcpv6-client request to the **ip6tables** file.

i.

Run the following command to add the rule to **ip6tables**:

ip6tables -A INPUT -m state --state NEW -m udp -p udp --dport 546 -d fe80::/64 -j ACCEPT

ii. Run the following command to save the rule in **ip6tables**: service ip6tables save

Figure 3-28 Example command



- d. (Optional) For CentOS 7/CentOS 8, change the IPv6 link-local address mode of extension NICs to EUI64.
 - i. Run the following command to query the NIC information: **nmcli con**

Figure 3-29 Querying NIC information



ii. Run the following command to change the IPv6 link-local address mode of eth1 to EUI64:

nmcli con modify "Wired connection 1" ipv6.addr-gen-mode eui64

NOTE

The NIC information varies depending on the CentOS series. In the command, *Wired connection 1* needs to be replaced with the value in the **NAME** column of the queried NIC information.

iii. Run the following commands to bring eth1 down and up:

ifdown eth1

ifup eth1

- e. Restart the network service.
 - i. For CentOS 6, run the following command to restart the network service:

service network restart

ii. For CentOS 7/EulerOS/Fedora, run the following command to restart the network service:

systemctl restart NetworkManager

- f. Perform **Step 1** to check whether dynamic IPv6 address assignment is enabled.
- SUSE, openSUSE, or CoreOS

SUSE 11 SP4 does not support dynamic IPv6 address assignment. No additional configuration is required for SUSE 12 SP1 or SUSE 12 SP2.

No additional configuration is required for openSUSE 12.3r1 of SOSE 12.3r2.

No additional configuration is required for CoreOS 10.10.5.

----End

Checking the ECS Network Configuration

1. Run the following command to check whether the ECS network service is normal:

systemctl status NetworkManager

If the network service is normal, the command output shows **active** (running) and the service status is **enabled**.

Figure 3-30 Network service status



 Run the following command to check how the ECS NIC obtains an IP address: cat /etc/sysconfig/network-scripts/ifcfg-ethx

NOTE

- **ethx** needs to be replaced with a specific NIC, for example, **eth0**.
- This command takes CentOS 7 as an example.

Figure 3-31 Method for ECS NIC to obtain an IP address



- If the value of **BOOTPROTO** is **dhcp**, the ECS NIC obtains an IP address using dynamic DHCP. Then, perform **3**.
- If the value of **BOOTPROTO** is **static**, the network of the ECS NIC is statically configured. Then, perform 4.
- 3. If the ECS NIC obtains an IP address using dynamic DHCP, run the following command to check whether the DHCP process is normal:

systemctl status NetworkManager

If the command output contains **dhclient**, the process is running properly.

Figure 3-32 Checking the DHCP process

 If the network of the ECS NIC is statically configured, run the following command to check whether the IP address configuration has taken effect: ip a

3 OS

If the IP address configuration in the command output is **forever**, the configuration has taken effect.

Figure 3-33 Checking the static configuration

[]] [] [] [] [] [] [] [] [] [
1: lo: <loopback,up,lower_up> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000</loopback,up,lower_up>
link/loopback 00:00:00:00:00 brd 00:00:00:00:00
inet 127.0.0.1/8 scope host lo
valid_lft forever preferred_lft forever
inet6 ::1/128 scope host
valid_lft forever preferred_lft forever
2: eth0: <broadcast,multicast,up,lower_up> mtu 1500 qdisc mq state UP group default qlen 1000</broadcast,multicast,up,lower_up>
link/ether fa:16:3e:fc:fc:14 brd ff:ff:ff:ff:ff
inet 10.0.0.24/24 hrd 10.0.0.255 scope global noprefixroute eth0
valid_lft forever preferred_lft forever
inet6 fe80::f816:3eff:fefc:fc14/64 scope link
valid_lft forever preferred_lft forever

Setting the Timeout Duration for IPv6 Address Assignment

After automatic IPv6 address assignment is configured on an ECS running CentOS 6.x or Debian, the ECS will be created as a private image. When this image is used to create an ECS in an environment that IPv6 is unavailable, the ECS may start slow because acquiring an IPv6 address times out. Before creating the private image, you can set the timeout duration for acquiring IPv6 addresses to 30s as follows:

- CentOS 6.*x*.
 - Run the following command to edit the dhclient.conf file:
 vi /etc/dhcp/dhclient.conf
 - b. Press i to enter editing mode and add the timeout attribute to the file. timeout 30;
 - c. Enter :wq to save the settings and exit.
- Debian 7.5:
 - a. Run the following command to edit the **networking** file: **vi /etc/init.d/networking**
 - b. Press i to enter editing mode and add the timeout attribute.

Figure 3-34 Modification 1

115 case "\$:	1" in
116 start)	
117	if init_is_upstart; then
118	exit 1
119	fi
120	process_options
121	check_ifstate
122	
123	if ["\$CONFIGURE_INTERFACES" = no]
124	then
125	log_action_msg "Not configuring network interfaces, see /etc/default/networking"
126	exit O
127	fi
128	set -f
129	exclusions=\$(process_exclusions)
130	log_action_begin_msg "Configuring network interfaces"
131 (1)	if <mark>/usr/bin/timeout 30</mark> ifup −a \$exclusions \$verbose && ifup_hotplug \$exclusions \$verbose
132	then
133	log_action_end_msg \$?
134	else
135	log_action_end_msg \$?
136	fi
137	;;
138	
139 stop)	
140	if init_is_upstart; then
141	exit 0
142	fi
143	check_network_file_systems
144	check_network_swap
	the entry have been and the set of the set o
144 145 146 147 (2) 148	check_network_swap log_action_begin_msg "Deconfiguring network interfaces" if/vusr/bin/timeout 30 ifdown -aexclude=1o \$verbose; then log_action_end_msg \$?

Figure 3-35 Modification 2



- Debian 8.2.0/8.8.0
 - Run the following command to edit the network-pre.conf file:
 vi /lib/systemd/system/networking.service.d/network-pre.conf
 - Press *i* to enter editing mode and add the timeout attribute to the file. [Service] TimeoutStartSec=30
- Debian 9.0
 - a. Run the following command to edit the **networking.service** file:

vi /etc/system/system/network-online.target.wants/ networking.service

b. Press i to enter editing mode and change **TimeoutStartSec=5min** to **TimeoutStartSec=30**.

3.10 How Do I Clean Up the Disk Space of a Windows ECS?

Scenarios

To clean up the disk space of a Windows ECS, you need to perform the following operations.

Disable Virtual Memory

Some disk space of an ECS serves as virtual memory which can be used when the ECS memory is exhausted. However, when the memory usage is high, frequent switching between the memory and virtual memory causes a large number of extra I/Os, deteriorating the I/O performance. You can disable virtual memory to free up disk space on Windows.

An ECS running Windows Server 2008 R2 Standard 64-bit is used as an example to describe how to disable virtual memory.

1. Right-click the **Computer** icon and choose **Properties**. In the displayed window, click **Advanced System Settings**.

- 2. In the **System Properties** dialog box, click the **Advanced** tab.
- 3. Click Settings in the Performance area.
- 4. In the displayed **Performance Options** dialog box, click the **Advanced** tab and then **Change**. The **Virtual Memory** dialog box is displayed.
- 5. Delete virtual memory.

In the **Virtual Memory** dialog box, deselect **Automatically manage paging file size for all drives**, select the paging file of the disk whose virtual memory is to be deleted, select **No paging file**, and click **Set**.

If the warning "If you disable the paging file or set the initial size to less than *xxx* megabytes and a system error occurs, Windows might not record details that could help identify the problem. Do you want to continue?" is displayed, set the initial size to *xxx* shown in the warning. The partition for storing paging files must have sufficient space. If the disk resources are insufficient, virtual memory will be insufficient. You only need to set virtual memory in one partition of the ECS.

6. Choose Start > Control Panel > Appearance and Personalization > Show hidden files and folders. The Folder Options dialog box is displayed. Click the View tab, deselect Hide protected operating system files (Recommended), and select Show hidden files, folders, and drives. When you deselect Hide protected operating system files (Recommended), a warning is displayed, as shown in Figure 3-36. Click Yes and then click Apply.

Figure 3-36 Warning



Figure 3-37 Showing hidden files

Folder Options	×		
General View Search			
Folder views You can apply the view (such as Details or Icons) that you are using for this folder to all folders of this type. Apply to Folders Reset Folders			
Advanced settings:			
Advanced settings. Display file size information in folder tips Display the full path in the title bar (Classic theme only) Hidden files and folders Ont show hidden files, folders, or drives Ont show hidden files, folders, and drives Hide empty drives in the Computer folder Hide extensions for known file types Hide protected operating system files (Recommended) Launch folder windows in a separate process Restore previous folder windows at logon Show drive letters Show encrypted or compressed NTFS files in color			
Restore Defaults			
OK Cancel Apply			

7. After file hiding is disabled, delete the hidden file **pagefile.sys** from disk C. After the OS is restarted, the available space increases.

Disable Hibernation

An ECS running Windows Server 2008 R2 Standard 64-bit is used as an example to describe how to disable hibernation.

Method 1: Delete the **Hiberfil.sys** file.

The **Hiberfil.sys** file is generated by Windows hibernation and occupies large system disk space. It saves memory data and sessions to disks so that the memory image file required by sessions can be quickly restored after the computer is restarted. Perform the following steps to delete the **Hiberfil.sys** file to free up some disk space:

- 1. Run cmd.exe as an administrator to open the command line interface (CLI).
- 2. Run the following command to disable hibernation (the **Hiberfil.sys** file will be automatically deleted):

powercfg -h off

Method 2: Disable hibernation.

 Choose Start > Control Panel > Appearance and Personalization > Personalization > Change screen saver and click Change power settings.

- 2. In the **Power Options** window, click **Change plan settings** on the right of **Balanced (recommended)**.
- 3. Set Turn off the display to Never.



Delete Internet Explorer Temporary Files

An ECS running Windows Server 2008 R2 Standard 64bit is used as an example to describe how to delete temporary files generated by Internet Explorer.

Internet Explorer stores all information about Internet access in a specified folder, including Internet Explorer cache files, cookies files, browsing history, visited websites, URLs in the address bar, and Internet Explorer forms/passwords. Perform the following steps to delete temporary files generated by Internet Explorer:

 Choose Start > Control Panel > Network and Internet > Internet Options. Click Delete browsing history and cookies to delete cookies and temporary files of Internet Explorer.



📴 Internet Properties 🔋 🔀			
General Security Privacy Content Connections Programs Advanced			
Home page			
To create home page tabs, type each address on its own line.			
https://www.huaweidoud.com/			
Use current Use default Use new tab			
Startup			
C Start with tabs from the last session			
Start with home page			
Tabs			
Change how webpages are displayed in tabs. Tabs			
Browsing history			
Delete temporary files, history, cookies, saved passwords, and web form information.			
Delete browsing history on exit			
Delete Settings			
Appearance			
Colors Languages Fonts Accessibility			
OK Cancel Apply			

 After the preceding step is performed, there may be residual files or records left. Open the C:\Windows\temp and C:\Users\Username\AppData\Local \Temp folders and delete all files from them. (The folders store visited websites and other temporary data.)

C:\Windows\Temp	~
C:\Users\Administrator\AppData\Local\Temp	~

Disable Data Execution Prevention (DEP)

DEP reserves a part of ECS memory for temporarily storing application data and another part of memory for temporarily storing application instructions. This protects ECSs from viruses and other security threats.

An ECS running Windows Server 2008 R2 Standard 64-bit is used as an example to describe how to disable DEP.

- 1. Right-click the **Computer** icon and choose **Properties**. In the displayed window, click **Advanced System Settings**.
- 2. In the **System Properties** dialog box, click the **Advanced** tab.

- 3. Click **Settings** in the **Performance** area.
- 4. In the **Performance Options** dialog box, click the **Data Execution Prevention** tab, select **Turn on DEP for essential Windows programs and services only**, and click **Apply**.

Performance Options	×		
Visual Effects Advanced Data Execution Prevention			
Data Execution Prevention (DEP) helps protect against damage from viruses and other security threats. <u>How does it work?</u>			
 Turn on DEP for essential Windows programs and services only 			
C Turn on DEP for all programs and services except those I select:			
	1		
Add Remove			
Your computer's processor supports hardware-based DEP.			
OK Cancel Apply			

Delete Redundant Application Files

The **C:\Windows\prefetch** folder stores the index files generated by applications. The files are used to improve system performance and speed up system startup and file reading. However, the number of files increases with time. Deleting redundant files can free up disk space on Windows.

Delete all files from the C:\Windows\prefetch folder to delete redundant application files.

Clean Up Disks

You can use **Disk Cleanup** of ECSs to delete temporary files, empty the recycle bin, and delete redundant system files and other files.

On the ECS, click Start. In the search box, enter Disk Cleanup. Click the displayed **Disk Cleanup** to scan for the space that can be freed up. After the scan is complete, confirm the files to be deleted and click **OK** to start disk cleanup.

If **Disk Cleanup** is unavailable, you need to install Desktop Experience first. An ECS running Windows Server 2008 R2 Standard 64bit is used as an example to describe how to install Desktop Experience.

- Choose Start > All Programs > Administrative Tools > Server Manager. In 1. the navigation pane, click Features.
- 2. Click Add Features.
- 3. In the Select Features dialog box, select Desktop Experience and click Next.

Add Features Wizard		×
Select Features		
Features Confirmation Progress Results	Select one or more features to install on this server. Features: Image:	Description: Desktop Experience includes features of Windows 7, such as Windows Media Player, desktop themes, and photo management. Desktop Experience does not enable any of the Windows 7 features; you must manually enable them. *
	< Previous Ne	ext > Install Cancel

Figure 3-38 Installing Desktop Experience

- In the Confirm Installation Selections dialog box, confirm the installation of 4. Desktop Experience, and click Install.
- 5. When the installation progress reaches 100% in the Results dialog box, the system prompts you to restart the server. Click **Close**, and then click **Yes** to restart the server.
- After the server is restarted, start the server manager. In the Function 6. Summary area, check whether Desktop Experience is installed.



Figure 3-39 Successful installation of Desktop Experience

7. Choose **Start** > **All Programs** > **Accessories** > **System Tools** > **Disk Cleanup** to start the disk cleanup tool.

Defragment and Optimize Drives

Disk defragmentation is a process in which system software or professional disk defragmentation software is used to organize the fragments generated during the long-term use of disks into the smallest number of contiguous fragments, improving the performance and speed of the server.

An ECS running Windows Server 2012 R2 Datacenter 64bit is used as an example to describe how to defragment and optimize drives.

1. Choose **Control Panel > All Control Panel Items > Administrative Tools**, and double-click **Defragment and Optimize Drives**.

🗟 l ⊋ 📗 = l	Shortcut Tools Application Tools	Administrative Tools	_ D X
File Home Share	View Manage Manage		~ (
🔄 🕘 🗉 🕇 🗟 « All	Control Panel Items Administrative Tools	🗸 🖒 Search Administ	rative Tools 🛛 🔎
⊿ 🔆 Favorites	Name	Date modified Type	Size
Desktop	퉬 Terminal Services	8/22/2013 23:39 File folder	
🐌 Downloads	Component Services	8/22/2013 14:57 Shortcut	
🔚 Recent places	🛃 Computer Management	8/22/2013 14:54 Shortcut	
	📸 Defragment and Optimize Drives	8/22/2013 14:47 Shortcut	
🔺 🌉 This PC	🛃 Event Viewer	8/22/2013 14:55 Shortcut	
🛛 隆 Desktop	🔝 iSCSI Initiator	8/22/2013 14:57 Shortcut	
🛛 📗 Documents	🛃 Local Security Policy	8/22/2013 14:54 Shortcut	1
🛛 🗼 Downloads	🔊 Microsoft Azure Services	11/22/2014 9:46 Shortcut	
🛛 🚺 Music	📷 ODBC Data Sources (32-bit)	8/22/2013 7:56 Shortcut	
🛛 📔 Pictures	📷 ODBC Data Sources (64-bit)	8/22/2013 14:59 Shortcut	
🛛 📔 Videos	Performance Monitor	8/22/2013 14:52 Shortcut	
🛛 📥 Local Disk (C:)	Resource Monitor	8/22/2013 14:52 Shortcut	
	🛃 Security Configuration Wizard	8/22/2013 14:45 Shortcut	
🛛 📬 Network	🔁 Server Manager	8/22/2013 14:55 Shortcut	-
	🔊 Services	8/22/2013 14:54 Shortcut	
	🛃 System Configuration	8/22/2013 14:53 Shortcut	
	🔁 System Information	8/22/2013 14:53 Shortcut	
	🔊 Task Scheduler	8/22/2013 14:55 Shortcut	-
[Mindows Firewall with Advanced Security	0/22/2012 14:45 Shortcut	>
24 items 1 item selected			

2. Select the disk partition to be optimized, and then click **Analyze** to analyze whether the disk partition needs to be optimized. If the analysis result indicates that optimization is necessary, click **Optimize**.

6		Optimize Drives		- 🗆 X
You can optimize your d optimized. Only drives o Status			or analyze them to find out	if they need to be
Drive	Media type	Last run	Current status	
📥 (C:)	Hard disk drive	8/16/2019 10:14	OK (0% fragmented)	
G System Reserved	Hard disk drive	Never run	OK (0% fragmented)	
Scheduled optimization			Analyze	Optimize
On				Change settings
	imized automatically.			change settings
				Close

Delete .dmp Files

When a blue screen of death (BSOD) occurs on a Windows ECS, the system automatically generates a BSOD error file with the suffix .dmp. A .dmp file is a system error file in Windows, such as, **memory.dmp** and **minixxxx.dmp**.

You can manually delete .dmp files from disk C to free up system disk space.

After the .dmp file is deleted, the BSOD cause on the ECS cannot be queried.

Clean Up Component Store

Windows Component Store stores all the files required for installing Windows. Updated installation files are also stored in Component Store, which causes the size of Component Store to increase as the update increases.

An ECS running Windows Server 2012 R2 Datacenter 64-bit is used as an example to describe how to clean up Component Store.

1. In Windows PowerShell, run the following command to delete the backup files generated during the Service Pack installation:

dism /online /cleanup-image /spsuperseded

```
PS C:\Users\Administrator≻ dism /online /cleanup-image /spsuperseded
Deployment Image Servicing and Management tool
Version: 6.3.9600.19408
Image Version: 6.3.9600.19397
Service Pack Cleanup cannot proceed: No Service Pack backup files were found.
The operation completed successfully.
```

2. Check the size of Component Store.

Dism.exe /Online /Cleanup-Image /AnalyzeComponentStore

PS C:\Users\Administrator> Dism.exe /Online /Cleanup-Image /AnalyzeComponentStore
Deployment Image Servicing and Management tool Version: 6.3.9600.19408
Image Version: 6.3.9600.19397
[=======]
Component Store (WinSxS) information:
Windows Explorer Reported Size of Component Store : 7.90 GB
Actual Size of Component Store : 7.75 GB
Shared with Windows : 4.12 GB Backups and Disabled Features : 3.33 GB Cache and Temporary Data : 297.92 MB
Date of Last Cleanup : 2019-08-16 11:00:48
Number of Reclaimable Packages : 3 Component Store Cleanup Recommended : Yes
The operation completed successfully.

3. Clean up Component Store.

Dism.exe /online /Cleanup-Image /StartComponentCleanup



Delete System Logs

System logs record hardware and software information, system problems, and system events. You can use the logs to locate error causes or track attacker actions. System logs include application logs, security logs, startup logs, and event forwarding logs. The **System32** > **LogFiles** folder on disk C of Windows stores the operation logs and event logs of Windows. Deleting the folder can free up the space of disk C.

An ECS running Windows Server 2008 R2 Standard 64-bit is used as an example to describe how to delete system logs.

1. Open the C:\Windows\System32\LogFiles folder, and delete all the files and folders from it.

C:\Windows\System32\LogFiles

 \sim

- 1. Choose **Start**, right-click **Computer**, and choose **Manage** from the shortcut menu.
- In the displayed window, choose Diagnostics > Event Viewer > Windows Logs and delete logs of Application, Security, Setup, System, and Forwarded Events.


4 Driver Installation

4.1 Must I Install Guest OS Drivers on an ECS?

Installing Guest OS drivers on an ECS improves your experience in using the ECS. In addition, it also ensures high reliability and stability of ECSs.

- Windows ECSs: Install PV and VirtIO drivers on ECSs.
- Linux ECSs: Install Xen PV and VirtIO drivers and add them to initrd.

4.2 Why Do I Need to Install and Update VirtIO Drivers for Windows?

Why Do I Need to Install VirtIO Drivers?

VirtIO drivers are paravirtualized drivers that provide high-performance disks and NICs for ECSs.

- Windows does not have VirtIO drivers installed by default.
- Public images have VirtIO drivers by default.
- You need to install VirtIO drivers for private images. For details, see Installing VirtIO Drivers.

Why Do I Need to Update VirtIO Drivers?

This ensures that known issues identified by the community can be eliminated from drivers as soon as possible.

What Do I Need to Do?

- Upgrade VirtIO drivers in Windows private images or running Windows ECSs.
- If you have any technical issues or questions, contact the customer service.

4.3 Why Did I Fail to Install Guest OS Drivers on a Windows ECS?

Possible causes:

- Your image file was exported from a VMware VM, and VMware Tools was not uninstalled or not completely uninstalled.
- You have downloaded Guest OS drivers of an incorrect version for your Windows ECS.
- The disk space available for installing Guest OS drivers is insufficient. Ensure that the disk where Guest OS drivers are installed has at least 300 MB space available.

4.4 How Do I Install VirtIO Drivers in Windows?

The installation only applies to KVM ECSs.

Before using an ECS or external image file to create a private image, ensure that VirtIO drivers have been installed in the OS so that ECSs created from this image can support KVM virtualization and the network performance can be improved.

For details, see Installing VirtIO Drivers.

4.5 How Do I Install Native KVM Drivers in Linux?

When optimizing a Linux private image, you need to install native KVM drivers on the ECS from which the image will be created. If you ECS already has native KVM drivers installed, you do not need to install the drivers again.

For details, see .

4.6 How Do I Install Native Xen and KVM Drivers?

Scenarios

Before using a Xen ECS to create a private image, you need to install native Xen and KVM drivers on the ECS.

This section describes how to install native Xen and KVM drivers.

If an ECS has no Xen drivers installed, the network performance of the ECS will be poor, and the security groups and firewall configured for the ECS will not take effect.

If an ECS has no KVM drivers installed, the NICs of the ECS may not be detected and the ECS will be unable to communicate with other resources.

Prerequisites

- The virtualization type of the ECS is Xen..
- The Linux ECS kernel version must be later than 2.6.24.
- Disable your antivirus and intrusion detection software. You can enable them after the driver installation is complete.

Procedure

Modify the configuration file depending on the OS.

• CentOS and EulerOS

Take CentOS 7.0 as an example. Modify the **/etc/dracut.conf** file. Add Xen PV and VirtIO drivers to **add_drivers**. Xen PV drivers include xen-blkfront and xen-netfront. VirtIO drivers include virtio_blk, virtio_scsi, virtio_net, virtio_pci, virtio_ring, and virtio. Separate driver names with spaces. Save and exit the **/etc/dracut.conf** file. Run the **dracut -f** command to regenerate initrd.

For details, see CentOS and EulerOS.

• Ubuntu and Debian

Modify the **/etc/initramfs-tools/modules** file. Add Xen PV and VirtIO drivers. Xen PV drivers include xen-blkfront and xen-netfront. VirtIO drivers include virtio_blk, virtio_scsi, virtio_net, virtio_pci, virtio_ring, and virtio. Separate driver names with spaces. Save and exit the **/etc/initramfs-tools/modules** file. Run the **update-initramfs -u** command to regenerate initrd.

For details, see **Ubuntu and Debian**.

- SUSE and openSUSE
 - If the OS version is earlier than SUSE 12 SP1 or openSUSE 13, modify the /etc/sysconfig/kernel file and add Xen PV and VirtIO drivers to INITRD_MODULES="". Xen PV drivers include xen_vnif, xen_vbd, and xen_platform_pci. VirtIO drivers include virtio_blk, virtio_scsi, virtio_net, virtio_pci, virtio_ring, and virtio. Separate driver names with spaces. Run the mkinitrd command to regenerate initrd.
 - If the OS version is SUSE 12 SP1, modify the /etc/dracut.conf file and add Xen PV and VirtIO drivers to add_drivers. Xen PV drivers include xen_vnif, xen_vbd, and xen_platform_pci. VirtIO drivers include virtio_blk, virtio_scsi, virtio_net, virtio_pci, virtio_ring, and virtio. Separate driver names with spaces. Run the dracut -f command to regenerate initrd.
 - If the OS version is later than SUSE 12 SP1 or openSUSE 13, modify the /etc/dracut.conf file and add Xen PV and VirtIO drivers to add_drivers. Xen PV drivers include xen-blkfront and xen-netfront. VirtIO drivers include virtio_blk, virtio_scsi, virtio_net, virtio_pci, virtio_ring, and virtio. Separate driver names with spaces. Save and exit the /etc/ dracut.conf file. Run the dracut -f command to regenerate initrd.

For details, see **SUSE and openSUSE**.

NOTE

For SUSE, run the following command to check whether xen-kmp (driver package for Xen PV) is installed:

rpm -qa |grep xen-kmp

If information similar to the following is displayed, xen-kmp is installed in the OS: xen-kmp-default-4.2.2 04 3.0.76 0.11-0.7.5

If xen-kmp is not installed, obtain it from the ISO file and install it.

If you add built-in drivers to the initrd or initramfs file by mistake, the ECS will not be affected.

CentOS and EulerOS

1. Run the following command to open the /etc/dracut.conf file:

vi /etc/dracut.conf

2. Press i to enter editing mode and add Xen PV and VirtIO drivers to add_drivers (the format varies depending on the OS). [root@CTU1000xxxxx ~]# vi /etc/dracut.conf # additional kernel modules to the default add_drivers+="xen-blkfront xen-netfront virtio_blk virtio_scsi virtio_net virtio_pci virtio_ring virtio"

- 3. Press **Esc**, enter **:wq**, and press **Enter**. The system saves the change and exits the **/etc/dracut.conf** file.
- 4. Run the following command to regenerate initrd:

dracut -f /boot/initramfs-2.6.32-573.8.1.el6.x86_64.img

If the virtual file system is not the default initramfs, run the **dracut -f** *Name of the initramfs or initrd file actually used* command. The actual initramfs or initrd file name can be obtained from the **grub.cfg** file, which can be **/boot/ grub/grub.cfg**, **/boot/grub2/grub.cfg**, or **/boot/grub/grub.conf** depending on the OS.

5. Check whether native Xen and KVM drivers have been installed. If the virtual file system is initramfs, run the following commands:

lsinitrd /boot/initramfs-`uname -r`.img | grep xen

lsinitrd /boot/initramfs-`uname -r`.img | grep virtio

If the virtual file system is initrd, run the following commands:

lsinitrd /boot/initrd-`uname -r` | grep xen

lsinitrd /boot/initrd-`uname -r` | grep virtio

Assume that the virtual file system is initramfs. The command output will be:

[root@CTU10000xxxxx home]# lsinitrd /boot/initramfs-`uname -r`.img | grep xen -rwxr--r-- 1 root root 54888 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/ block/**xen-blkfront.ko** -rwxr--r-- 1 root root 45664 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/ drivers/net/**xen-netfront.ko**

[root@CTU10000xxxxx home]# lsinitrd /boot/initramfs-`uname -r`.img | grep virtio

-rwxrr 1 root root	23448 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
block/ virtio_blk.ko	

1 root	root	50704 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/

drivers/net/virtio_net.ko -rwxr--r-- 1 root root 28424 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/

scsi/**virtio_scsi.ko** drwxr-xr-x 2 root root 0 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86 64/kernel/drivers/

-rwxr--r-- 1

virtio

-rwxrr 1 root	root	14544 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
virtio/ virtio.ko		
-rwxrr 1 root	root	21040 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
virtio/ virtio_pci.ko		
-rwxrr 1 root	root	18016 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86 64/kernel/drivers/
virtio/ virtio_ring.ko		/ / /
· · · · · · · · · · · · · · · · · ·		

NOTE

If you add built-in drivers to the initrd or initramfs file by mistake, the ECS will not be affected but the drivers cannot be found by running the **lsinitrd** command. You can run the following commands to check whether built-in drivers are in the kernel:

```
cat /boot/config-`uname -r` | grep CONFIG_VIRTIO | grep y
cat /boot/config-`uname -r` | grep CONFIG_XEN | grep y
```

Ubuntu and Debian

1. Run the following command to open the **modules** file:

vi /etc/initramfs-tools/modules

 Press i to enter editing mode and add Xen PV and VirtIO drivers to the /etc/ initramfs-tools/modules file (the format varies depending on the OS). [root@CTU10000xxxxx ~]#vi /etc/initramfs-tools/modules

```
# Examples:
#
# raid1
# sd_mOd
xen-blkfront
xen-netfront
virtio_blk
virtio_scsi
virtio_net
virtio_pci
virtio_ring
virtio
```

- 3. Press **Esc**, enter **:wq**, and press **Enter**. The system saves the change and exits the **/etc/initramfs-tools/modules** file.
- 4. Run the following command to regenerate initrd:

update-initramfs -u

5. Run the following commands to check whether native Xen and KVM drivers have been installed:

lsinitramfs /boot/initrd.img-`uname -r` |grep xen

lsinitramfs /boot/initrd.img-`uname -r` |grep virtio

[root@ CTU10000xxxxx home]# lsinitramfs /boot/initrd.img-`uname -r` |grep xen lib/modules/3.5.0-23-generic/kernel/drivers/net/ethernet/qlogic/netxen lib/modules/3.5.0-23-generic/kernel/drivers/net/ethernet/qlogic/netxen/netxen_nic.ko lib/modules/3.5.0-23-generic/kernel/drivers/net/xen-netback lib/modules/3.5.0-23-generic/kernel/drivers/net/xen-netback/xen-netback.ko lib/modules/3.5.0-23-generic/kernel/drivers/block/xen-blkback lib/modules/3.5.0-23-generic/kernel/drivers/block/xen-blkback lib/modules/3.5.0-23-generic/kernel/drivers/block/xen-blkback.ko

[root@ CTU10000xxxxx home]# lsinitramfs /boot/initrd.img-`uname -r` |grep virtio lib/modules/3.5.0-23-generic/kernel/drivers/scsi/**virtio_scsi.ko**

D NOTE

If you add built-in drivers to the initrd or initramfs file by mistake, the ECS will not be affected but the drivers cannot be found by running the **lsinitrd** command. You can run the following commands to check whether built-in drivers are in the kernel:

[root@ CTU10000xxxxx home]**# cat /boot/config-`uname -r` | grep CONFIG_VIRTIO | grep y** CONFIG_VIRTIO_BLK=y CONFIG_VIRTIO_NET=y CONFIG_VIRTIO=y CONFIG_VIRTIO_RING=y CONFIG_VIRTIO_PCI=y CONFIG_VIRTIO_MMIO_CMDLINE_DEVICES=y [root@ CTU10000xxxxx home]**# cat /boot/config-`uname -r` | grep CONFIG_XEN | grep y** CONFIG_XEN_BLKDEV_FRONTEND=y CONFIG_XEN_NETDEV_FRONTEND=y

SUSE and openSUSE

If the OS version is earlier than SUSE 12 SP1 or openSUSE 13, modify the **/etc/ sysconfig/kernel** file to install the drivers. For details, see **scenario 1**.

If the OS version is SUSE 12 SP1, modify the **/etc/dracut.conf** file to install the drivers. For details, see **scenario 2**.

If the OS version is later than SUSE 12 SP1 or openSUSE 13, modify the **/etc/ dracut.conf** file to install the drivers. For details, see **scenario 3**.

Earlier than SUSE 12 SP1 or openSUSE 13:

NOTE

Before installing the drivers, run the following command to check whether xen-kmp (driver package for Xen PV) is installed in the OS:

rpm -qa |grep xen-kmp

If information similar to the following is displayed, xen-kmp is installed in the OS: xen-kmp-default-4.2.2 04 3.0.76 0.11-0.7.5

If xen-kmp is not installed, obtain it from the ISO package and install it first.

a. Run the following command to modify the **/etc/sysconfig/kernel** file:

vi /etc/sysconfig/kernel

b. Add Xen PV and VirtIO drivers after **INITRD_MODULES**= (the format varies depending on the OS).

SIA10000xxxxx:~ # vi /etc/sysconfig/kernel # (like drivers for scsi-controllers, for lvm or reiserfs)

(like universitor scal-controllers, for

INITRD_MODULES="ata_piix ata_generic xen_vnif xen_vbd xen_platform_pci virtio_blk virtio_scsi virtio_net virtio_pci virtio_ring virtio"

c. Run the **mkinitrd** command to regenerate initrd.

D NOTE

If the virtual file system is not the default initramfs or initrd, run the **dracut -f** *Name of the initramfs or initrd file actually used* command. The actual initramfs or initrd file name can be obtained from the **menu.lst** or **grub.cfg** file (/**boot**/ **grub/menu.lst**, /**boot/grub/grub.cfg**, or /**boot/grub2/grub.cfg**).

The following is an example initrd file of SUSE 11 SP4:

default 0 timeout 10 gfxmenu (hd0,0)/boot/message title sles11sp4_001_[_VMX_] root (hd0,0) kernel /boot/linux.vmx vga=0x314 splash=silent console=ttyS0,115200n8 console=tty0 net.ifnames=0 NON_PERSISTENT_DEVICE_NAMES=1 showopts initrd /boot/initrd.vmx title Failsafe_sles11sp4_001_[_VMX_] root (hd0,0) kernel /boot/linux.vmx vga=0x314 splash=silent ide=nodma apm=off noresume edd=off powersaved=off nohz=off highres=off processsor.max+cstate=1 nomodeset x11failsafe console=ttyS0,115200n8 console=tty0 net.ifnames=0 NON_PERSISTENT_DEVICE_NAMES=1 showopts initrd /boot/initrd.vmx

/boot/initrd.vmx in the **initrd** line is the initrd file actually used. So, the command used to regenerate initrd is **dracut -f /boot/initrd.vmx**. If **/ boot** is missing from the initrd file path, you need to add it when you run the **dracut -f** command. For example, if the file path is **/initramfs-***xxx*, the command should be **dracut -f /boot/initramfs-***xxx*.

d. Run the following commands to check whether Xen PV and KVM VirtIO drivers have been installed:

lsinitrd /boot/initrd-`uname -r` | grep xen

lsinitrd /boot/initrd-`uname -r` | grep virtio

SIA10000xxxxx:~ # lsinitrd /boot/initrd-`uname -r` | grep xen -rwxr--r-- 1 root root 42400 Jun 22 2012 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/block/ xen-blkfront.ko

-rwxr--r-- 1 root root 44200 Jun 22 2012 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/net/ xen-netfront.ko

SIA10000xxxxx:~ # lsinitrd /boot/initrd-`uname -r` | grep virtio -rwxr--r-- 1 root root 19248 Jun 22 2012 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/scsi/ virtio_scsi.ko

-rwxr--r-- 1 root root 23856 Jun 22 2012 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/block/ virtio_blk.ko

drwxr-xr-x 2 root root 0 Jul 12 14:53 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/**virtio** -rwxr--r-- 1 root root 15848 Jun 22 2012 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/virtio/ **virtio_ring.ko**

-rwxr--r-- 1 root root 20008 Jun 22 2012 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/virtio/ virtio_pci.ko

-rwxr--r-- 1 root root 12272 Jun 22 2012 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/virtio/ virtio.ko

-rwxr--r-- 1 root root 38208 Jun 22 2012 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/net/ virtio_net.ko

e. Restart the ECS.

f. Modify the /boot/grub/menu.lst file to add xen_platform_pci.dev_unplug=all and change the root settings.

Before the modification:

###Don't change this comment -YaST2 identifier: Original name: linux###
title SUSE Linux Enterprise Server 11SP4 - 3.0.76-0.11 (default)
root (hd0,0)
kernel /boot/vmlinuz-3.0.76-0.11-default root=UUID=4eb40294-4c6f-4384-bbb6b8795bbb1130 splash=silentcrashkernel=256M-:128M showopts vga=0x314
initrd /boot/initrd-3.0.76-0.11-default

After the modification:

###Don't change this comment -YaST2 identifier: Original name: linux###
title SUSE Linux Enterprise Server 11SP4 - 3.0.76-0.11 (default)
root (hd0,0)
kernel /boot/vmlinuz-3.0.76-0.11-default root=UUID=4eb40294-4c6f-4384-bbb6b8795bbb1130 splash=silentcrashkernel=256M-:128M showopts vga=0x314
xen_platform_pci.dev_unplug=all
initrd /boot/initrd-3.0.76-0.11-default

- Ensure that the root partition is in UUID format.
- xen_platform_pci.dev_unplug=all is used to shield QEMU devices.
- For SUSE 11 SP1 64bit to SUSE 11 SP4 64bit, add xen_platform_pci.dev_unplug=all to the menu.lst file. For SUSE 12 or later, QEMU device shield is enabled by default, and you do not need to configure it.
- g. Run the following commands to check whether Xen drivers exist in initrd:

lsinitrd /boot/initrd-`uname -r` | grep xen

lsinitrd /boot/initrd-`uname -r` | grep virtio

SIA10000xxxx:~ # lsinitrd /boot/initrd-`uname -r` | grep xen -rwxr--r-- 1 root root 42400 Jun 22 2012 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/block/ xen-blkfront.ko

-rwxr--r-- 1 root root 44200 Jun 22 2012 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/net/ xen-netfront.ko

SIA10000xxxxx:~ # lsinitrd /boot/initrd-`uname -r` | grep virtio

-rwxr--r-- 1 root root 19248 Jun 22 2012 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/scsi/ virtio_scsi.ko

-rwxr-r-- 1 root root 23856 Jun 22 2012 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/block/ virtio_blk.ko

drwxr-xr-x 2 root root 0 Jul 12 14:53 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/**virtio** -rwxr--r-- 1 root root 15848 Jun 22 2012 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/virtio/ **virtio** ring.ko

-rwxr--r-- 1 root root 20008 Jun 22 2012 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/virtio/ virtio_pci.ko

-rwxr--r-- 1 root root 12272 Jun 22 2012 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/virtio/ virtio.ko

-rwxr--r-- 1 root root 38208 Jun 22 2012 lib/modules/2.6.32-279.el6.x86_64/kernel/drivers/net/ virtio_net.ko

D NOTE

If you add built-in drivers to the initrd or initramfs file by mistake, the ECS will not be affected but the drivers cannot be found by running the **lsinitrd** command. You can run the following commands to check whether built-in drivers are in the kernel:

```
cat /boot/config-`uname -r` | grep CONFIG_VIRTIO | grep y
```

```
cat /boot/config-`uname -r` | grep CONFIG_XEN | grep y
```

- SUSE 12 SP1:
 - a. Run the following command to open the /etc/dracut.conf file:

vi /etc/dracut.conf

- b. Press i to enter editing mode and add Xen PV and VirtIO drivers to adddrivers (the format varies depending on the OS). [root@CTU10000xxxx ~]# vi /etc/dracut.conf # additional kernel modules to the default add_drivers+="ata_piix ata_generic xen_vnif xen_vbd xen_platform_pci virtio_blk virtio_scsi virtio_net virtio_pci virtio_ring virtio"
- c. Press **Esc**, enter **:wq**, and press **Enter**. The system saves the change and exits the **/etc/dracut.conf** file.
- d. Run the following command to regenerate initrd:

dracut -f /boot/initramfs-File name

If the virtual file system is not the default initramfs, run the **dracut -f** *Name of the initramfs or initrd file actually used* command. The actual initramfs or initrd file name can be obtained from the **grub.cfg** file, which can be **/boot/grub/grub.cfg**, **/boot/grub2/grub.cfg**, or **/boot/grub/ grub.conf** depending on the OS.

e. Check whether native Xen and KVM drivers have been installed. If the virtual file system is initramfs, run the following commands:

lsinitrd /boot/initramfs-`uname -r`.img | grep xen

lsinitrd /boot/initramfs-`uname -r`.img | grep virtio

If the virtual file system is initrd, run the following commands:

lsinitrd /boot/initrd-`uname -r` | grep xen

lsinitrd /boot/initrd-`uname -r` | grep virtio

• Later than SUSE 12 SP1 or openSUSE 13:

Take SUSE Linux Enterprise Server 12 SP2 (x86_64) as an example.

- Run the following command to open the /etc/dracut.conf file:
 vi /etc/dracut.conf
- b. Press i to enter editing mode and add Xen PV and VirtIO drivers to add_drivers (the format varies depending on the OS). [root@CTU10000xxxx ~]# vi /etc/dracut.conf # additional kernel modules to the default add_drivers+="ata_piix ata_generic xen-blkfront xen-netfront virtio_blk virtio_scsi virtio_net virtio_pci virtio_ring virtio"
- c. Press **Esc**, enter **:wq**, and press **Enter**. The system saves the change and exits the **/etc/dracut.conf** file.
- d. Run the following command to regenerate initrd:

dracut -f /boot/initramfs-File name

If the virtual file system is not the default initramfs, run the **dracut** -f *Name of the initramfs or initrd file actually used* command. The actual initramfs or initrd file name can be obtained from the **grub.cfg** file, which can be **/boot/grub/grub.cfg**, **/boot/grub2/grub.cfg**, or **/boot/grub/ grub.conf** depending on the OS.

e. Check whether native Xen and KVM drivers have been installed. If the virtual file system is initramfs, run the following commands:

lsinitrd /boot/initramfs-`uname -r`.img | grep xen

lsinitrd /boot/initramfs-`uname -r`.img | grep virtio

If the virtual file system is initrd, run the following commands:

lsinitrd /boot/initrd-`uname -r` | grep xen

lsinitrd /boot/initrd-`uname -r` | grep virtio

Assume that the virtual file system is initrd. The command output will be:

sluo-ecs-30dc:~ # lsinitrd /boot/initrd-`uname -r` | grep xen -rw-r--r-- 1 root root 69575 Oct 26 2016 lib/modules/4.4.21-69-default/kernel/drivers/block/**xenblkfront.ko**

-rw-r--r-- 1 root root 53415 Oct 26 2016 lib/modules/4.4.21-69-default/kernel/drivers/net/xennetfront.ko

drwxr-xr-x 2 root root 0 Sep 28 10:21 lib/modules/4.4.21-69-default/updates/pvdriver/**xen-hcall** -rwxr-xr-x 1 root root 8320 Sep 28 10:21 lib/modules/4.4.21-69-default/updates/pvdriver/xen-hcall/**xen-hcall.ko**

sluo-ecs-30dc:~ # lsinitrd /boot/initrd-`uname -r` | grep virtio -rw-r--r-- 1 root root 29335 Oct 26 2016 lib/modules/4.4.21-69-default/kernel/drivers/block/ virtio_blk.ko -rw-r--r-- 1 root root 57007 Oct 26 2016 lib/modules/4.4.21-69-default/kernel/drivers/net/ virtio net.ko -rw-r--r-- 1 root root 32415 Oct 26 2016 lib/modules/4.4.21-69-default/kernel/drivers/scsi/ virtio_scsi.ko

drwxr-xr-x 2 root root 0 Sep 28 10:21 lib/modules/4.4.21-69-default/kernel/drivers/virtio -rw-r--r-- 1 root root 19623 Oct 26 2016 lib/modules/4.4.21-69-default/kernel/drivers/virtio/ virtio.ko

-rw-r--r-- 1 root root 38943 Oct 26 2016 lib/modules/4.4.21-69-default/kernel/drivers/virtio/ virtio_pci.ko

-rw-r--r-- 1 root root 24431 Oct 26 2016 lib/modules/4.4.21-69-default/kernel/drivers/virtio/ virtio_ring.ko

NOTE

If you add built-in drivers to the initrd or initramfs file by mistake, the ECS will not be affected but the drivers cannot be found by running the **lsinitrd** command. You can run the following commands to check whether built-in drivers are in the kernel:

cat /boot/config-`uname -r` | grep CONFIG_VIRTIO | grep y cat /boot/config-`uname -r` | grep CONFIG_XEN | grep y

5 Image Sharing

5.1 General Sharing FAQ

How Many Tenants Can I Share an Image With?

A system disk image or data disk image can be shared with up to 128 tenants, but a full-ECS image can only be shared with up to 10 tenants.

How Many Images Can Be Shared with Me?

There is no limit.

Can I Share Images Between My Accounts at the Chinese Mainland Website and International Website?

Yes. You can share images between your accounts of the Chinese Mainland and International websites, but only in regions that are available on both websites. For example, you cannot share images in the CN North-Beijing1 region of the Chinese Mainland website with your account of the International website which does not have this region.

Do Shared Images Affect My Private Image Quota?

No.

I Shared an Image to an Account But the Account Did Not Accept or Reject the Image. Will My Image Sharing Quota Be Consumed?

No.

Where Can I View the Images Shared with Me?

Switch to the region where the shared image is, and choose **Service List** > **Compute** > **Image Management Service** > **Images Shared with Me**.

If you are a multi-project user, make clear which of your projects will receive the shared image. Switch to the region where the project is and select the project.

Then, choose Service List > Compute > Image Management Service > Images Shared with Me.

If the image is not accepted, a red dot is displayed on the **Images Shared with Me** tab page (as shown in Figure 5-1) and a message is displayed, asking you whether to accept the shared image. After the image is accepted, it is displayed in the list on the **Images Shared with Me** tab page.

Figure 5-1 Images Shared with Me

Publi	c Images	Private Images		Images Shared with Me
~	Tenant	has shared	3 ima	ages with you.

If I Want to Share a System Disk Image with Another Account, Should the Account Purchase an ECS in Advance?

No. The account can use the shared image to create ECSs.

Is There Any Restriction on the Region When I Create ECSs Using a Shared Image?

Yes. You can only create ECSs in the same region as the shared image.

Can I Share Images Shared with Me with Others?

You cannot directly share such images with other tenants. If you do need to do so, you can replicate a shared image as a private image and then share it.

Can I Use an Image I Have Shared with Others to Create an ECS?

Yes. After sharing an image with others, you can still use the image to create an ECS and use the created ECS to create a private image.

What Are the Risks of Creating ECSs Using a Shared Image?

The image owner can view, stop sharing, or delete the image at any time. After the shared image is deleted, you will be unable to use it to create a new ECS or change the OS of an existing ECS.

The cloud platform does not ensure the integrity or security of images shared by other accounts. You are advised to choose only images shared by trusted accounts.

What Are the Risks of Sharing Images?

Data, files, and software may be disclosed. Before sharing an image, you must take care to delete any sensitive data or important files from the image. The image recipient can use the shared image to create ECSs and use the created ECSs to create private images. If the created private images are shared with others, any data leakage that occurs can be quite widespread.

How Do I Replicate an Image Between Projects in a Region?

Image sharing allows you to replicate images between projects in the same region.

In the **Share Image** dialog box, enter target project IDs, and click **OK**, as instructed in **Sharing Specified Images**. After the image is shared, switch to the target projects. On the **Image Management Service** > **Images Shared with Me** page, you can view the shared image to be accepted.

Can I Specify a Region or an AZ for Sharing an Image?

No. When sharing an image, you can only specify a project ID. You cannot specify a region or an AZ. An image can only be shared within a given region. Once shared, it can be used in any AZ in that region.

Can I Restore My Data Disks from a Data Disk Image Shared by Another Account?

No. You can only use a shared image to create a new data disk but not to restore your existing data disks. However, you can use the new data disk to restore data. For details, see **Can I Import Data from a Data Disk Image to a Data Disk?**

Can I Share My ECS and Its EIP with Another Account?

You can share an ECS by sharing a full-ECS image but the EIP cannot be shared.

What Can I Do If I Want to Use a Rejected Image?

If you have rejected an image shared by another tenant, but now want to use it, two methods are available:

Method 1

Ask the image owner to share the image with you again. For details, see **Adding Tenants Who Can Use Shared Images**.

Method 2

Accept the rejected image again. For details, see Accepting Rejected Images.

5.2 How Will I Be Billed for Sharing an Image?

- If it is a system or data disk image, the OBS storage is free.
- If it is a full-ECS image, you will be billed for the space used to store the CBR backup.

NOTE

Full-ECS images created from a CBR backup can be shared but those created from a CSBS backup cannot.

5.3 What Are the Differences Between Sharing Images and Replicating Images?

• Sharing images:

You can only share images within a region with other users. To share an image across regions, replicate the image to the target region and then share it. For details, see **Sharing Images**.

- Replicating images:
 - In-region: You can convert encrypted and unencrypted images into each other or enable some advanced features (such as fast ECS creation from an image) by replicating an image within a region.

Scenario	Opera tion	Description	Helpful Link
Sharing	Share	The image is shared with another user in the same region. The target user can use the image (with the same ID as the source image) but the image is still owned by the user who shared it.	Sharing Specified Images
In-region replication under the same account	Replica te	This is used for conversion between encrypted and unencrypted images or for enabling advanced features (such as fast ECS creation from an image).	Replicating Images Within a Region

The following table describes the details.

6 Image Import

6.1 Can I Use Images in Formats not Described in This Document?

No. Currently, only VMDK, VHD, RAW, QCOW2, VHDX, QED, VDI, QCOW, ZVHD2, ISO, and ZVHD formats are supported.

Images in -flat.vmdk format and image file packages containing snapshot or delta volumes are not supported. You can use **qemu-img** to convert an image to one of the supported formats before uploading it to the cloud platform.

NOTE

For details about image file format conversion, see **Converting the Image Format Using qemu-img**.

6.2 What Are the Impacts If I Do Not Pre-configure an ECS Used to Create a Private Image?

Before using an ECS or external image file to create a private image, you need to pre-configure the ECS or the source VM of the image file. If you do not perform the pre-configuration, there are some potential impacts:

- 1. If you do not delete network rule files from the **udev** directory, those rules will be applied to newly created ECSs. Also, if you do not configure DHCP, the NICs of newly created ECSs will not start from eth0. You need to remotely log in to the new ECSs to resolve these issues.
- 2. When creating Linux ECSs:
 - Custom passwords cannot be injected.
 - Certificates cannot be injected.
 - Other custom configurations cannot be applied to new ECSs.
- 3. If you do not delete the automatic mount configuration from the **fstab** file, new ECSs may fail to start.

6.3 How Do I Import an OVF or OVA File to the Cloud Platform?

Scenarios

Open Virtualization Appliance (OVA) is a single file archive (with the .ova extension) of all of the files comprising an Open Virtualization Format (OVF) package. OVF is a package of the files required for defining and deploying VMs. An OVF package generally includes .ovf, .mf, and .vmdk files.

- An .ovf file is an XML descriptor that defines the metadata of a VM, such as the name and hardware requirements, and contains information about other files in the OVF package.
- An .mf file contains the SHA hash codes of all the files in the package and is used to prevent the image file from being tampered with.
- A .vmdk file is a virtual disk file that is used to create a disk image. An OVF package may contain multiple .vmdk files.

This section describes how to import OVF and OVA files to the cloud platform.

Procedure

Manually extract VMDK files from an OVF or OVA template and upload them to an OBS bucket. Then, you can select one of them from the bucket when you use an external file to create a system or data disk image.

NOTE

The following assumes that the OVF or OVA template contains only one VMDK file. If there are multiple VMDK files (for example, there are three VMDK files, one used as a system disk image file and the others as data disk image files), upload them to an OBS bucket and register them as a system disk image and data disk images, respectively.

- The source VM runs Linux.
 - If you choose to export an OVF template, upload the VMDK file generated in the folder to the cloud platform.
 - If you choose to export an OVA template and name it MyVm, perform the following operations:
 - i. Run the following command to check the OVA file:

file MyVm.ova

The command output is as follows:

MyVm.ova: POSIX tar archive (GNU)

MyVm.ova contains the following two files:

\$**tar tf MyVm.ova** MyVm.ovf MyVm.vmdk

ii. Run the following command to decompress MyVm.ova:

tar xvf MyVm.ova

The extracted folder contains the following files:

MyVm.ovf MyVm.vmdk

The .vmdk image file can be uploaded to the cloud platform.

NOTE

You can import an image file in the VHD, VMDK, QCOW2, RAW, VHDX, QCOW, VDI, QED, ZVHD, or ZVHD2 format to create a private image.

For details, see Creating a System Disk Image from an Image File or Creating a Data Disk Image from an Image File.

6.4 What Do I Do If I Chose the Wrong OS or System Disk Capacity When Registering a Private Image?

If you selected the wrong OS, the ECS creation may fail when the image is used.

If the system disk capacity you configured is less than the one in the image file, image registration will fail.

In such cases, delete the image and create a new one using the correct settings.

6.5 Why Did My VHD Upload Fail? Why Does the System Say the System Disk in the VHD Image File Is Larger Than What I Specified on the Management Console?

The possible causes may be:

- 1. Too small a value was specified when registering the image. Check the system disk capacity in the VHD image file. Specify a value at least this large when you use the VHD image file to register an image.
- 2. The VHD's actual disk size is larger than its virtual size. This can happen if the VHD image file was generated using **qemu-img** or a similar tool. For details, see https://bugs.launchpad.net/qemu/+bug/1490611.

Run the following command to check the VHD image file information:

[xxxx@xxxx test]\$ **qemu-img info 2g.vhd** image: 2g.vhd file format: vpc virtual size: 2.0G (2147991552 bytes) disk size: 8.0K cluster_size: 2097152

The virtual size is always an integer number of GBs. As a result, if an actual size is, like in the example here, **2147991552 bytes** (**2.0004 GB**), the virtual size will be only **2 GB**. In this example, you need to specify an integer larger than the actual size 2.0004 GB. The system disk capacity on the management console can only be an integer value, so you enter an integer larger than 2 GB.

7 Image Export

7.1 Can I Download My Private Images to a Local PC?

Yes. You can download private images as instructed in **Exporting an Image**.

Currently, only images in VMDK, VHD, QCOW2, or ZVHD format can be exported.

The default format of a private image is ZVHD2. Images exported in different formats may vary in size.

For more information, see **Exporting an Image**.

7.2 Can I Use the System Disk Image of an ECS on a BMS After I Export It from the Cloud Platform?

No.

The system disk image of an ECS is a VM file that contains a system running environment and does not have an installation boot program. So, it cannot be used on a BMS.

For how to create a BMS private image, see Bare Metal Server User Guide.

You can use either of the following methods to create a BMS private image:

- Creating a Private Image from a BMS
- Creating a Private Image from an External Image File

7.3 Why Is the Image Size in an OBS Bucket Different from That Displayed in IMS?

Symptom

When a private image is exported to an OBS bucket, the image size in the bucket is different from that displayed in IMS.

For example, the size of a private image is 1.04 GB on the IMS console, but in an OBS bucket, the size is 2.91 GB.

Cause Analysis

The default format of a private image is ZVHD2, but it may be stored in a different format (VMDK, VHD, QCOW2, or ZVHD) in an OBS bucket after it is exported. The format conversion may lead to size changes.

For more information, see **Exporting an Image**.

7.4 Can I Download a Public Image to My PC?

You cannot directly download a public image. However, you can use the public image to create an ECS, use the ECS to create a private image, export the private image to your OBS bucket, and then download the image.

Helpful links:

- Creating a System Disk Image from a Cloud Server
- Exporting an Image

NOTE

- Ubuntu public images and the private images created from these public images cannot be exported.
- However, if an Ubuntu private image is created from an external image file, this private image can be exported.

7.5 What Are the Differences Between Standard and Fast Import/Export?

Item	Description	Helpful Link
Standar d import	Import an external image file to the Huawei Cloud management console for creating a private image.	 Creating a System Disk Image from an Image File
	External image files in the following formats can be imported: VMDK, VHD, QCOW2, RAW, VHDX, QED, VDI, QCOW, ZVHD2, and ZVHD.	• Creating a Data Disk Image from an Image File
	Maximum file size: 128 GB	
	During the import, operations such as driver injection will be performed in the background. Therefore, the import takes a longer time than fast import.	

ltem	Description	Helpful Link
Fast import	 When importing an external image file in the RAW or ZVHD2 format to the Huawei Cloud management console, you can select Enable Fast Create. The import is fast because the system does not perform any operations like driver injection. Verify that: The image file converted to the RAW format has been optimized as required and a bitmap file has been generated for it. The image file converted to the ZVHD2 format has been optimized as required. 	Fast Import of an Image File
Standar d export	You can export private images to OBS buckets and download them to your local PC for further use on other cloud platforms. Maximum file size: 128 GB (If an image file is larger than 128 GB, use fast export.) You can specify the format of the exported image file. Currently, only QCOW2, VMDK, VHD, and ZVHD are supported.	Exporting an Image
Fast export	On the Export Image page, select Enable following Fast Export . You cannot specify the format of the exported image file. After the export is complete, you can use a tool to convert the exported image to your desired format. The file size is not limited. Encrypted images do not support fast export.	Exporting an Image

7.6 Why the Export Option Is Unavailable for My Image?

Some images cannot be exported. So, the **Export** option is not provided for them in the **Operation** column. The following images cannot be exported:

• Public images

- Full-ECS images
- ISO images

8 Image Replication

8.1 Image Replication

When Do I Need to Replicate an Image?

You can replicate an image within a region or across regions.

• In-region replication

This is used for conversion between encrypted and unencrypted images or for enabling advanced features (such as fast ECS creation). For details, see **Replicating Images Within a Region**.

• Cross-region replication

This is used for replicating a private image in the current region to the same account in another region. You can use this private image to deploy the same application environment in the two regions. For details, see **Replicating Images Across Regions**.

What Images Can I Replicate?

Only private images can be replicated. The restrictions on image replication are as follows:

• In-region replication

Private images created from ISO files or full-ECS images cannot be replicated within a region.

Cross-region replication

Encrypted images cannot be replicated across regions.

Can I Replicate Images Across Regions Between the Chinese Mainland, International, and European Websites?

The Chinese Mainland, International, and European websites are accessed by using different accounts. You cannot replicate images between these websites through cross-region replication. However, there is a way to do the replication.

For example, if you want to replicate a system disk image from the CN-Hong Kong region of the International website to the CN East-Shanghai1 region of the Chinese Mainland website, perform the following operations:

1. Log in to the console of the International website, export the image from the CN-Hong Kong region to an OBS bucket in that region, and download the image to a local directory. Assume that the image file is exported as **Windows Server 2012.qcow2**.

For details, see **Exporting an Image**.

2. Log in to the console of the Chinese Mainland website, switch to the CN East-Shanghai1 region, and upload the **Windows Server 2012.qcow2** image file to an OBS bucket in that region.

For details, see Uploading an Image File.

3. Register the image file as a private image

For details, see **Registering an Image File as a Private Image**.

How Long Does It Take to Replicate an Image?

The time required for replicating an image depends on the network transmission speed and the number of tasks in the queue.

How Will I Be Billed for Replicating Images?

• In-region replication

The replicas of system disk and data disk images are stored in OBS buckets for free.

NOTE

Full-ECS images cannot be replicated within the same region.

Cross-region replication

System disk and data disk images: You are billed for the OBS traffic for replicating images cross regions.

Full-ECS images: You are billed for the CSBS or CBR traffic generated by replicating the images cross regions and for the space used to store the CSBS or CBR backups associated with the replicas.

NOTE

The process of cross-region full-ECS image replication is as follows:

- 1. The system replicates the backup of a full-ECS image to the destination region.
- 2. In the destination region, the system automatically generates a full-ECS image from the backup.

Is There a Limit on the Image Size During Cross-Region Replication?

Yes. An image to be replicated cannot be larger than 128 GB.

To replicate an image larger than 128 GB across regions, you can export the image to an OBS bucket, import the image to an OBS bucket in the target region, and register it as a private image.

How Many Cross-Region Image Replication Tasks Can I Run Concurrently?

Five.

How Do I Replicate an Image Between Projects in a Region?

You can use image sharing to replicate images between projects in the same region.

In the **Share Image** dialog box, enter target project IDs, and click **OK**, as instructed in **Sharing Specified Images**. After the image is shared, switch to the target projects. On the **Image Management Service** > **Images Shared with Me** page, you can view the shared image to be accepted.

Figure 8-1 Accepting or rejecting shared images

Public Images	rivate Images Images Shared	with Me				
∧ Tenant hwstaff_p	b_CBUInfo_laaS has shared 4 images w	rith you.				Accept All
Accept Re	ect					
Name		ID	O S Type	05	Disk Capacity (GB)	Source Tenant
no-del-master	qiujiandong-00612884	26171a5	Linux	CentOS 7.6 64bit	40	Account facility and another and
v no-del-node1-	glujiandong-00612884	752fb62	Linux	CentOS 7.6 64bit	40	Account
docker-kubec	-heim	6e0e036	Linux	CentOS 7.6 64bit	40	Account mental and mental and
no-del-kind-qi	ijiandong-00612884	1bdc9da	Linux	CentOS 7.6 64bit	40	Account and an international and

What Are the Regions That Support Cross-Region Replication of System and Data Disk Images?

Destination Region Source Region	CN North- Beijing1	CN North- Beijing2	CN North- Beijing4	CN North- Ulanqab1	CN East- Shanghai1	CN East- Shanghai2	CN South- Guangzhou	CN Southwest- Guiyang1	CN-Hong Kong	AP-Bangkok	AP-Singapore	AF- Johannesburg	LA-Mexico City1	LA-Sao Paulo1
CN North-Beijing1		1	1	1	V	V	V	V	x	x	x	V	x	x
CN North-Beijing2	1		1	1	1	1	V	1	x	x	x	x	x	x
CN North-Beijing4	1	1		1	1	1	V	V	x	x	x	x	x	x
CN North-Ulanqab1	1	1	1		1	1	1	1	x	x	x	x	x	x
CN East-Shanghai1	1	1	1	1		1	V	V	x	x	x	x	x	x
CN East-Shanghai2	1	1	1	1	1		V	V	x	x	x	x	x	x
CN South-Guangzhou	1	1	1	1	1	1		1	x	x	x	x	x	x
CN Southwest- Guiyang1	1	1	1	1	1	V	V	-	x	x	x	x	x	x
CN-Hong Kong	x	x	x	x	x	x	x	x		x	x	x	x	x
AP-Bangkok	x	x	x	x	x	x	x	x	x	-	x	x	x	x
AP-Singapore	x	x	x	x	x	x	x	x	x	x	-	x	x	x
AF-Johannesburg	V	x	x	x	x	x	x	x	x	x	x		x	x
LA-Mexico City1	x	x	x	x	x	x	x	. ×	x	x	x	x		x
LA-Sao Paulo1	x	x	x	x	×	x	x	x	x	x	x	x	x	

Figure 8-2 Cross-region replication of system and data disk images

D NOTE

- \checkmark indicates replication is supported.
- × indicates replication is not supported. You can export the desired image and then import it to the target region.
- - indicates that the source and destination regions are the same.
- Region adjustment may affect cross-region image replication of users that are registered after the adjustment. The users that were registered before the adjustment are not affected. You can see on the console if cross-region can be performed for a given region.
- Regions not listed in the figure above do not support cross-region replication of system or data disk images.

What Are the Regions That Support Cross-Region Full-ECS Image Replication?

Cross-region replication of full-ECS images is only available for certain regions. You can see on the console if it is available for a given region.

If a full-ECS image cannot be replicated to a different region, you can use it to create an ECS, use the ECS to create a system disk image and a data disk image, and replicate the images to the destination region.

A full-ECS image created using an ECS backup can be replicated from the region where they reside to another region, but the replicated full-ECS image cannot be replicated across regions again.

8.2 How Do I Limit Agency Permissions for Cross-Region Image Replication?

Scenarios

During cross-region image replication, an agency is required to verify cloud service permissions in the destination region. So, you need to create a cloud service agency before the replication. In the past, you might be required to assign the Tenant Administrator or IMS Administrator system-defined role to an agency, but the two roles provide more excessive permissions than the agency actually needs.

You can use fine-grained system-defined policies to limit the agency permissions.

Procedure

- 1. Log in to the console.
- 2. Hover the mouse pointer over the username in the upper right corner and select **Identity and Access Management** from the drop-down list.
- 3. In the navigation pane, choose Agencies.
- 4. Click the name of the agency you used for cross-region replication.
- 5. On the **Permissions** tab page, if the **Tenant Administrator** or **IMS Administrator** role is displayed, proceed with the following steps to limit the agency permissions.
- 6. Click **Authorize**. Select system-defined policies based on the image type. For details, see **Table 8-1**.

- 7. Click **Next** and set the minimum authorization scope.
- 8. Click OK.
- 9. Click **Finish** to return back to the permission list.
- 10. In the list, select the **Tenant Administrator** and **IMS Administrator** roles and click **Delete**.

Table 8-1 Permissions required for cross-region image replication

Scenario	System-defined Policy
Cross-region replication of a system or data disk image	IMS CrossCopyAgencyPolicy
Cross-region replication of a full-ECS image	IMS CrossCopyAgencyCBRPolicy

9 Image Deletion

Will a Private Image Be Automatically Deleted If I Delete or Unsubscribe from the ECS Used to Create the Image?

No. Private images created using ECSs are stored in OBS buckets. Deleting or unsubscribing from the ECSs does not affect the images.

Can I Delete an Image I Shared with Others?

• If an image is shared with a project, you can delete the image without requiring any participation from the image recipients. After you delete the image, the image recipients cannot use it any longer. Inform the recipients to back up their data before you delete the image.

How Do I Delete a Shared Image? Does the Deletion Affect Any ECS or EVS Disk Created from It?

Reject this image on the **Images Shared with Me** tab page. This does not affect any ECS or EVS disk created from it.

10 Image Encryption

How Can I Change an Unencrypted Image to an Encrypted One?

You can select an encryption key when you replicate the image. Then, the system will generate an encrypted version of the unencrypted image.

What Are the Restrictions on Using an Encrypted Image?

- An encrypted image cannot be shared with others.
- The key used for encrypting an image cannot be changed.

How Do I Check the Encryption Status of a Disk in an Encrypted Full-ECS Image?

- 1. Access the IMS console.
- 2. On the **Private Images** tab page, click an image name to go to its details page.
- 3. Click the **Backups** tab.
- 4. On the **Backup Details** page, check the disk encryption status in the **Encrypted** column.

1 1 Accounts and Permissions

11.1 What Do I Do If I Enabled EPS But Now I Cannot Find Private Images in My Enterprise Project?

Scenarios

If you cannot find the private images on the **Enterprise Project Management Service** page, add the private images to their associated enterprise project.

Procedure

- 1. Log in to the management console.
- 2. Under Compute, click Image Management Service.
- 3. Click the **Private Images** tab.
- 4. Locate the row that contains the image, click **More** in the **Operation** column, and select **Allocate to Enterprise Project**.
- 5. In the displayed dialog box, select the target enterprise project.

11.2 What Do I Do If I Cannot Create an Image from a CSBS Backup or BMS Using a Sub-account with the Allow_all Permission After EPS Is Enabled?

When an enterprise project subaccount is used to create an image, the system displays a message indicating that CSBS or BMS is not supported by EPS.

This is because CSBS and BMS are not interconnected with EPS regionally or globally. The global resource viewing permission must be granted to the subaccount in IAM. For example, you can view resources of other cloud services if you have the Tenant Guest permission.

For details, see Assigning Permissions to an IAM User.

12 Cloud-Init

12.1 Cloud-Init Installation FAQ

You are advised to install Cloud-Init on the ECS that will be used to create a private image so that new ECSs created from the private image support custom configurations (for example, changing the ECS login password).

For details about how to install and configure Cloud-Init, see **Installing and Configuring Cloud-Init**.

The following describes common problems you may encounter when installing Cloud-Init and their solutions.

Ubuntu 16.04/CentOS 7: Failed to Set Cloud-Init Automatic Start

Symptom

After Cloud-Init is installed, you run the following command to configure Cloud-Init automatic start:

systemctl enable cloud-init-local.service cloud-init.service cloudconfig.service cloud-final.service

Information similar to the following is displayed:

Figure 12-1 Failed to enable Cloud-Init to start automatically

root@ecs-wjq-ubuntu14:~# systemctl enable cloud-init-local.service	cloud-init.se
rvice cloud-config.service cloud-final.service	
Failed to execute operation: Unit file is masked	
root@ecs_wig_ubuntu14.~#	

- Solution
 - a. Run the following command to roll back the configuration:

systemctl unmask cloud-init-local.service cloud-init.service cloudconfig.service cloud-final.service

- Run the following command to configure automatic start again: systemctl enable cloud-init-local.service cloud-init.service cloudconfig.service cloud-final.service
- c. Run the following command to check the Cloud-Init status:

systemctl status cloud-init-local.service cloud-init.service cloud-config.service cloud-final.service

As shown in the following figures, **failed** is displayed and all services are in the **inactive** state.

Figure 12-2 Checking Cloud-Init status (1)

<pre>root@ecs-wjq-ubuntui4:~# systemct1 status cloud-init-local.service • cloud-init-local.service - Initial cloud-init job (pre-networking) Loaded: loaded (/lib/systemd/system/cloud-init-local.service; enabled; vendo Active: failed (Result: exit-code) since Fri 2018-08-17 07:12:20 UTC; 1min 2 Process: 4418 ExecStart=/usr/bin/cloud-init initlocal (code=exited, status Main PID: 4418 (code=exited, status=203/EXEC)</pre>
Aug 17 07:12:20 ecs-wjq-ubuntu14 systemd[1]: Starting Initial cloud-init job (p Aug 17 07:12:20 ecs-wjq-ubuntu14 systemd[1]: cloud-init-local.service: Main pro Aug 17 07:12:20 ecs-wjq-ubuntu14 systemd[1]: Failed to start Initial cloud-init Aug 17 07:12:20 ecs-wjq-ubuntu14 systemd[1]: cloud-init-local.service: Unit ent Aug 17 07:12:20 ecs-wjq-ubuntu14 systemd[1]: cloud-init-local.service: Failed w Ines 1-11/11 (END)

Figure 12-3 Checking Cloud-Init status (2)

Loaded:	it-local.service – I loaded (/lib/system				ed; vendor
preset: er					
Active:	failed (Result: exi	t-code) since	Fri 2018-08	-17 07:12:20 UT	C; 59s ago
Process:	4418 ExecStart=/usr.	/bin/cloud-ini	t initlo	cal (code=exite	d, status=
203/EXEC)					
Main PID:	4418 (code=exited,	status=203/EXE	C)		
Aug 17 07:1 e-networkin	12:20 ecs-wjq-ubuntu ng)	14 systemd[1]:	Starting I	nitial cloud-in	it job (pr
Aug 17 07:1	12:20 ecs-wjq-ubuntu		cloud-init	-local.service:	Main proc
	, code=exited, statu				
	12:20 ecs-wjq-ubuntu	14 systemd[1]:	Failed to	start Initial c.	loud-init
job (pre-ne	etworking).				
Aug 17 07:1	12:20 ecs-wjq-ubuntu	14 systemd[1]:	cloud-init	-local.service:	Unit ente
red failed	state.				
	12:20 ecs-wjq-ubuntu 'exit-code'.	14 systemd[1]:	cloud-init	-local.service:	Failed wi

This is because the address that the system uses to access Cloud-Init is redirected to **/usr/bin/**, but the actual installation path is **/usr/local/bin**.

- Run the following command to copy Cloud-Init to the usr/bin directory:
 cp /usr/local/cloud-init /usr/bin/
- e. Run the following command to restart Cloud-Init:

systemctl restart cloud-init-local.service cloud-init.service cloudconfig.service cloud-final.service

Figure 12-4 Restarting Cloud-Init

					tial cloud−init jo system/cloud−init−.		
					Fri 2018-08-17 07		
					in/cloud-init init		
				(code=exited, sta			
ug	17	07::	18:01	ecs-wjq-ubuntu14	cloud-init[4491]:	[CLOUDINIT]	util.py[DEBUG]:
ug	17	07:1	18:01	ecs-wjq-ubuntu14	cloud-init[4491]:	[CLOUDINIT]	util.py[DEBUG]:
ug.	17	07:1	18:01	ecs-wjq-ubuntu14	cloud-init[4491]:	[CLOUDINIT]	initpy[DEB
1g	17	07::	18:01	ecs-wjq-ubuntu14	cloud-init[4491]:	[CLOUDINIT]	util.py[DEBUG]:
lg	17	07:1	18:01	ecs-wjq-ubuntu14	cloud-init[4491]:	[CLOUDINIT]	util.py[DEBUG]:
18	17	07:1	18:01	ecs-wjq-ubuntu14	cloud-init[4491]:	[CLOUDINIT]	util.py[DEBUG]:
1g	17	07::	18:01	ecs-wjq-ubuntu14	cloud-init[4491]:	[CLOUDINIT]	cloud-init[DEBU
1g	17	07::	18:01	ecs-wjq-ubuntu14	cloud-init[4491]:	[CLOUDINIT]	util.py[DEBUG]:
1g	17	07:1	18:01	ecs-wjq-ubuntu14	cloud-init[4491]:	[CLOUDINIT]	util.py[DEBUG]:
19	17	07:1	18:01	ecs-wjq-ubuntu14	cloud-init[4491]:	[CLOUDINIT]	util.py[DEBUG]:

 Run the following command to check the Cloud-Init status: systemctl status cloud-init-local.service cloud-init.service cloudconfig.service cloud-final.service

Ubuntu 14.04: chkconfig and systemctl Not Installed

- Symptom chkconfig is not installed.
 - Solution
 Run the following commands to install chkconfig:
 apt-get update
 apt-get install sysv-rc-conf
 cp /usr/sbin/sysv-rc-conf /usr/sbin/chkconfig
 Run the following command to query the Cloud-Init version:
 cloud-init -v
 Information similar to the following is displayed:
 -bash:/usr/bin/cloud-init:not found this command
 Run the following command to copy Cloud-Init to the usr/bin directory:
 cp /usr/local/bin/cloud-init /usr/bin/

Debian 9.5: Failed to Query the Cloud-Init Version and Configure Automatic Start

1. Run the following command to query the Cloud-Init version:

cloud-init -v

Information similar to the following is displayed: -bash:/usr/bin/cloud-init:not found this command

Run the **cp /usr/local/bin/cloud-init /usr/bin/** command to copy Cloud-Init to the **usr/bin** directory.

2. Run the **cloud-init init --local** command.

Information similar to the following is displayed:

Figure 12-5 Information returned when Cloud-Init automatic start successfully set

robuess-debian=3:rtmp/Lbu0p=init/= cloud=init# cloud=ini# cloud=init# cloud=init# clo

The compilation fails because GCC is not installed.

To solve this issue:

Run the following command to install GCC. Then, install Cloud-Init again.

yum -y install gcc

3. After Cloud-Init is installed, run the following command to configure Cloud-Init automatic start:

systemctl enable cloud-init-local.service cloud-init.service cloudconfig.service cloud-final.service

Information similar to the following is displayed.

Figure 12-6 Prompt indicating the failure to configure Cloud-Init automatic start

Failed to enable unit: Unit file /etc/systemd/system/cloud-init-local.service is masked.

To solve this issue:

a. Run the following command to roll back the configuration:

systemctl unmask cloud-init-local.service cloud-init.service cloudconfig.service cloud-final.service

b. Run the following command to set automatic start again:

systemctl enable cloud-init-local.service cloud-init.service cloudconfig.service cloud-final.service

c. Run the following command to restart Cloud-Init:

systemctl restart cloud-init-local.service cloud-init.service cloudconfig.service cloud-final.service

Run the **systemctl status** command to check the Cloud-Init status. Information similar to the following is displayed:

Figure 12-7 Verifying the service status

Loaded: loaded (/lik-gystem/cloud-init-local.service: enabled: uendor preset: enabled) ncture: active (stied) since Mno 2018-08-20 02:40:37 UTC: 6s ago Process: 1002 ExecStart=/usr>hin/cloud-init initlocal (code=exited, status=0/SUCCESS) Tasks: 0 (linit: 4915) CGroup: regetted, status=0/SUCCESS) Tasks: 0 (linit: 4915) Tasks: 0 (linit: 4915) Tasks: 0 (linit: 4915) Tasks: 0 (linit: 4915) Tasks: 0 (linit: 4915) CGROUP: Tasks: 0 (linit: 4915) Tasks: 0 (🛚 cloud-init-local.service – Initial cloud-init job (pre-networking)
Process: 1002 Execetited, status=0/SUCCESS) Hain PDI: 1002 (code=exited, status=0/SUCCESS) Tasks: 0 (linit: 4915) CGroup: /system.slice/cloud-init10021: CLOUDINIT) util.py(DEBUG1: Running conmand ['blkid', '-tLABEL=config=2', '-odevi hug 20 02:40:37 ecs-debian=9 cloud-init10021: CLOUDINIT) util.py(DEBUG1: Reading from /proc/nounts (uletFalse) Hug 20 02:40:37 ecs-debian=9 cloud-init10021: CLOUDINIT) util.py(DEBUG1: Reading from /proc/nounts (uletFalse) Hug 20 02:40:37 ecs-debian=9 cloud-init10021: CLOUDINIT) util.py(DEBUG1: Reading from /proc/nounts (uletFalse) Hug 20 02:40:37 ecs-debian=9 cloud-init10021: CLOUDINIT) util.py(DEBUG1: Reading from /proc/nounts (uletFalse) Hug 20 02:40:37 ecs-debian=9 cloud-init10021: CLOUDINIT) util.py(DEBUG1: Reading from /proc/nounts (uletFalse) Hug 20 02:40:37 ecs-debian=9 cloud-init10021: CLOUDINIT) util.py(DEBUG1: Reading from /proc/nounts (uletFalse) Hug 20 02:40:37 ecs-debian=9 cloud-init10021: CLOUDINIT) til.py(DEBUG1: Read High cloud / init' took 0.104 seconds (0.10) Hug 20 02:40:37 ecs-debian=9 cloud-init10021: CLOUDINIT) util.py(DEBUG1: Read High cloud / init' took 0.104 seconds (0.10) Hug 20 02:40:37 ecs-debian=9 cloud-init10221: CLOUDINIT) util.py(DEBUG1: cloud - init' took 0.104 seconds (0.10) Hug 20 02:40:37 ecs-debian=9 cloud-init10021: CLOUDINIT) util.py(DEBUG1: cloud - init' took 0.104 seconds (0.10) Hug 20 02:40:37 ecs-debian=9 cloud-init10021: CLOUDINIT) util.py(DEBUG1: cloud - init' took 0.104 seconds (0.10) Hug 20 02:40:37 ecs-debian=9 cloud-init10021: CLOUDINIT) til.py(DEBUG1: cloud - init' took 0.104 seconds (0.10) Hug 20 02:40:37 ecs-debian=9 cloud-init10021: CLOUDINIT) til.py(DEBUG1: cloud - init' took 0.104 seconds (0.10) Hug 20 02:40:40 ecs-debian=9 cloud-init10051: CLOUDINIT) helpers.py(DEBUG1: cloud - init' not cloud - ini	
<pre>Main PID: 1002 (code=exited, status=0/SUCCESS) Tasks: 0 (lini: 4915) CGroup: /system.slice/cloud-init-local.service Mug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.py(DEBUG): Running command ['blkid', '-tLABEL=config-2', '-devi 10 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.py(DEBUG): Seeing if we can get any data from (class 'cloud 10 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.py(DEBUG): Read 1947 bytes from /proc/mounts 10 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.py(DEBUG): Read 1947 bytes from /proc/mounts 10 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.py(DEBUG): Read 1947 bytes from /proc/mounts 10 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.py(DEBUG): Read 1947 bytes from /proc/mutils 10 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.py(DEBUG): Read 14 bytes from /proc/mutils 10 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.py(DEBUG): Read 14 bytes from /proc/mutils 10 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.py(DEBUG): Read 14 bytes from /proc/mutils 10 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.py(DEBUG): Read 14 bytes from /proc/mutils 10 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.py(DEBUG): Read 14 bytes from /proc/mutils 10 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.py(DEBUG): Read 14 bytes from /proc/mutils 10 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) tall.py(DEBUG): endi-init mode 'init' took 0.104 seconds (0.10) 10 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) tall.py(DEBUG): endi-init mode 'init' took 0.104 seconds (0.10) 10 20 02:40:40 ecs-debian-9 cloud-init.service: enabled) 10 20 20:40:40 ecs-debian-9 cloud-init init (code=exited, status=0/SUCCESS) 10 asks: 0 (linit: 4915) 10 20 20:40:40 ecs-debian-9 cloud-init(10051): CLOUDINIT) helpers, py(DEBUG): config-ca-certs already ran (freq=once-per-instanc 10 20 02:40:40 ecs-debian-9 cloud-init(10051): CLOUDINIT) helpers, py(DEBUG): config-ca-c</pre>	
Tasks: 0 (lini: 4915) CGroup: /system.slice/cloud-init-local.service hug 20 02:40:37 ecs-debian-9 cloud-init10021: CLOUDINIT1 util.pyDEENGI: Running command ['blkid', '-tLABEL=config-2', '-odevi hug 20 02:40:37 ecs-debian-9 cloud-init10021: CLOUDINIT1 util.pyDEENGI: Reading from /proc/monts hug 20 02:40:37 ecs-debian-9 cloud-init10021: CLOUDINIT1 util.pyDEENGI: Reading from /proc/monts hug 20 02:40:37 ecs-debian-9 cloud-init10021: CLOUDINIT1 util.pyDEENGI: Reading from /proc/monts hug 20 02:40:37 ecs-debian-9 cloud-init10021: CLOUDINIT1 util.pyDEENGI: Reading from /proc/monts hug 20 02:40:37 ecs-debian-9 cloud-init10021: CLOUDINIT1 util.pyDEENGI: Reading from /proc/monts hug 20 02:40:37 ecs-debian-9 cloud-init10021: CLOUDINIT1 util.pyDEENGI: Reading from /proc/monts' /dw/pts', 'opts': hug 20 02:40:37 ecs-debian-9 cloud-init10021: CLOUDINIT1 util.pyDEENGI: Reading from /proc/mptinte hug 20 02:40:37 ecs-debian-9 cloud-init10021: CLOUDINIT1 util.pyDEENGI: Reading from /proc/mptinte hug 20 02:40:37 ecs-debian-9 cloud-init10021: CLOUDINIT1 util.pyDEENGI: cloud-init mode 'init' took 0.104 seconds (0.10) hug 20 02:40:37 ecs-debian-9 cloud-init10021: CLOUDINIT1 util.pyDEENGI: cloud-init mode 'init' took 0.404 seconds (0.10) hug 20 02:40:37 ecs-debian-9 cloud-init.service: nabled: uendor preset: enabled) for ecs-debian-9 cloud-init.service: nabled: uendor preset: enabled) for ecs-debian-9 cloud-init.service frocess: 0 (linit: 415) CGroup: /system.slice/cloud-init119051: CLOUDINIT1 belpers.pyDEENG1: config-ca-certs already ran (freq=once-per-instance hug 20 02:40:40 ecs-debian-9 cloud-init119051: CLOUDINIT1 belpers.pyDEENG1: Running module rsyslog (Gnodule 'cloudinit.config.ca hug 20 02:40:40 ecs-debian-9 cloud-init119051: CLOUDINIT1 belpers.pyDEENG1: Running module rsyslog (Gnodule 'cloudinit.config.ca hug 20 02:40:40 ecs-debian-9 cloud-init119051: CLOUDINIT1 belpers.pyDEENG1: Running module rsyslog (Gnodule 'cloudinit.config.ca hug 20 02:40:40 ecs-debian-9 cloud-init119051: CLOUDINIT1 belpers.pyDEE	
<pre>CGroup: /system.slice/cloud-init-local.service hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.pu[DEBUG): Running command ['blkid', '-tAREL=config-2', '-devi hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.pu[DEBUG): Seeing if we can get any data from <class 'cloud<br="">hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.pu[DEBUG): Read 1947 bytes from /proc/mounts (quiet=False) hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.pu[DEBUG): Read 1947 bytes from /proc/mounts hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.pu[DEBUG]: No local datasource found hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.pu[DEBUG]: No local datasource found hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.pu[DEBUG]: Read 14 bytes from /proc/uptime hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.pu[DEBUG]: Read 14 bytes from /proc/uptime hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.pu[DEBUG]: Read 14 bytes from /proc/uptime hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.pu[DEBUG]: Read 14 bytes from /proc/uptime hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.pu[DEBUG]: Read 14 bytes from /proc/uptime hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util.pu[DEBUG]: enabled) hctive: active (exited) since fmo 2010-00-20:24:04:00 UTC: 3s ago Process: 1096 ExceStart=/usr/bin/cloud-init init (code=exited, status=0/SUCESS) fasks: 0 (linit: 4915) CGroup: /system.slice/cloud-init(10961: CLOUDINIT) helpers.pu[DEBUG]: config-ca-certs already ran (freq=once-per-instanc hug 20 02:40:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) helpers.pu[DEBUG]: config-ca-certs already ran (freq=once-per-instanc hug 20 02:40:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) helpers.pu[DEBUG]: config-ca-certs already ran (freq=once-per-instanc hug 20 02:40:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) helpers.pu[DEBUG]: config-ca-certs already ran (freq=once-per-instanc hug 20 02:40:40 ecs-debian-9</class></pre>	
<pre>hug 20 02:40:37 ecs-debian-9 cloud-init110021: [CLOUDINIT] util.pup[DEBUG]: Running command ['blkid', '-tARBEL=config-2', '-devi hug 20 02:40:37 ecs-debian-9 cloud-init10021: [CLOUDINIT]initpup[DEBUG]: Seeing if we can get any data from (class 'cloudi hug 20 02:40:37 ecs-debian-9 cloud-init10021: [CLOUDINIT] util.pup[DEBUG]: Reading from /proc/nounts hug 20 02:40:37 ecs-debian-9 cloud-init10021: [CLOUDINIT] util.pup[DEBUG]: Reading from /proc/nounts hug 20 02:40:37 ecs-debian-9 cloud-init10021: [CLOUDINIT] util.pup[DEBUG]: Reading from /proc/nounts hug 20 02:40:37 ecs-debian-9 cloud-init10021: [CLOUDINIT] util.pup[DEBUG]: Reading from /proc/numts (must '/dw/pts', 'opts': hug 20 02:40:37 ecs-debian-9 cloud-init10021: [CLOUDINIT] util.pup[DEBUG]: Reading from /proc/numts [must '/dw/pts', 'opts': hug 20 02:40:37 ecs-debian-9 cloud-init10021: [CLOUDINIT] util.pup[DEBUG]: Reading from /proc/numts [must '/dw/pts', 'opts': hug 20 02:40:37 ecs-debian-9 cloud-init10021: [CLOUDINIT] util.pup[DEBUG]: cloud-init mode 'init' took 0.104 seconds (0.10) hug 20 02:40:37 ecs-debian-9 cloud-init10021: [CLOUDINIT] util.pup[DEBUG]: cloud-init mode 'init' took 0.104 seconds (0.10) hug 20 02:40:37 ecs-debian-9 cloud-init1 service: nabled: vendor preset: enabled) for ecs-debian-9 system/cloud-init.service: nabled: vendor preset: enabled) for ecs-debian-9 cloud-init10021: CLOUDINIT] util.pup[DEBUG]: config-ca-certs already ran (freq=once-per-instance hug 20 02:40:40 ecs-debian-9 cloud-init10951: (CLOUDINIT] blpers.pup[DEBUG]: Running module rsuslog (Gnodule 'cloudinit.config.ca hug 20 02:40:40 ecs-debian-9 cloud-init10951: (CLOUDINIT] blpers.pup[DEBUG]: Running module rsuslog (Gnodule 'cloudinit.config.ca hug 20 02:40:40 ecs-debian-9 cloud-init10951: (CLOUDINIT] blpers.pup[DEBUG]: Running module rsuslog (Gnodule 'cloudinit.config.ca hug 20 02:40:40 ecs-debian-9 cloud-init10951: (CLOUDINIT] blpers.pup[DEBUG]: Running module rsuslog (Gnodule 'cloudinit.config.co hug 20 02:40:40 ecs-debian-9 cloud-init10951: (CLOUDINIT] blpers.p</pre>	
<pre>hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT)inityUDEBNG: Seaing from yproc/nomts (multi-false) hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util, pyUDEBNG: Reading from yproc/nomts (multi-false) hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util, pyUDEBNG: Fetched ('devyts': ('nomtpoint': 'devyts', 'opts': hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util, pyUDEBNG: Reading from yproc/nomts hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) cloud-init(DEBNG: No local datasource found hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util, pyUDEBNG: Reading from yproc/nytime (multi-false) hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util, pyUDEBNG: ecad-init node 'init' took 0.104 seconds (0.10) hug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util, pyUDEBNG: ecad-init node 'init' took 0.104 seconds (0.10) hug 20 02:40:37 ecs-debian-9 cloud-init(service: cnabled: vendor preset: enabled) hug 20 02:40:37 ecs-debian-9 cloud-init service: nabled: vendor preset: enabled) het loaded (/it/systend/systend/systenc/loud-init cambicd: vendor preset: enabled) het loaded (/it/systend/systenc/loud-init enabled: vendor preset: enabled) het setvice (exited) since fion 2010-00-20 02:40:40 UTC: 3s ago Process: 1096 ExectState/susten/suclud-init init Code=exited, status=0/SUCCESS) Hain TID: 1096 (code=exited, status=0/SUCCESS) Tasks: 0 (linit: 4915) CGroup: /aysten.slice/cloud-init(1996): (CLOUDINIT) helpers.py(DEBNG): config-ca-certs already ran (freq=once-per-instance hug 20 02:48:40 ecs-debian-9 cloud-init(1996): (CLOUDINIT) helpers.py(DEBNG): Raming module rsysing (cnodule 'cloudinit.config. Aug 20 20:48:40 ecs-debian-9 cloud-init(1996): (CLOUDINIT) helpers.py(DEBNG): Raming module rsysing (cnodule 'cloudinit.config. Aug 20 02:48:40 ecs-debian-9 cloud-init(1996): (CLOUDINIT) helpers.py(DEBNG): Raming module rsysing (cnodule 'cloudinit.config. Aug 20 02:48:40 ecs-debian-9 cloud-init(1996): (CLOUDINIT) helpers.py(DEBNG): Raming mod</pre>	Group: /system.slice/cloud-init-local.service
 cloud-init.service - Initial cloud-init jab (metadata service crauler) Loaded: loaded (/lix-gustend/susten/cloud-init.service: enabled; uendor preset: enabled) Active: active fexied) since Mno 2018-08-20 20:248:40 UTC: 3s ago Process: 1096 Excestent=/sur/bin/cloud-init init (code=exited, status=0/SUCCESS) Tasks: 0 (linit: 4915) CGroup: /system.slice/cloud-init.service Aug 20 02:48:40 UTC: Sa ago Process: 1096 Excestent=/status=0/SUCCESS) Tasks: 0 (linit: 4915) CGroup: /system.slice/cloud-init.service Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: (CLOUDINIT) helpers,putDEBUG1: config-ca-certs already ran (freq=once-per-instanc Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: (CLOUDINIT) helpers,putDEBUG1: config-ca-certs already ran (freq=once-per-instanc Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: (CLOUDINIT) helpers,putDEBUG1: config-ca-certs already ran (freq=once-per-instanc Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: (CLOUDINIT) helpers,putDEBUG1: config-ca-certs already ran (freq=once-per-instanc Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: (CLOUDINIT) helpers,putDEBUG1: config-ca-certs-gorups already ran (freq=once-per-instance Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: (CLOUDINIT) helpers,putDEBUG1: config-ca-certs-gorups already ran (freq=once-per-instance Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: (CLOUDINIT) tolpers,putDEBUG1: config-ca-certs-gorups already ran (freq=once-per-instance Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: (CLOUDINIT) tolpers,putDEBUG1: Rand IT sproc-uptime Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: (CLOUDINIT) tut1, putDEBUG1: Rand IT sproc-uptime Aug 20 02:48:40 ecs-debian-9 sugtent(11: started lnitial cloud-init job (netadata service crawler). cloud-config.service - Apply the settings specified in cloud-config Loaded: loaded (/il/vsytend/sytem/cloud-config.service: enabled)	Aug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT)initputDEBUG1: Seeing if we can get any data from class 'cloudi Aug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util, putDEBUG1: Reading from /proc/nounts Aug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util, putDEBUG1: Reading from /proc/nounts Aug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util, putDEBUG1: Fetched ('devyts': ('nountpoint': '/devyts', 'opts': Aug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) cloud-init(DEBUG1: Fetched ('devyts': ('nountpoint': '/devyts', 'opts': Aug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) cloud-init(DEBUG1: Reading from /proc/vuptine (upietFalse) Aug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util, putDEBUG1: Reading from /proc/vuptine Aug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util, putDEBUG1: Reading from /proc/vuptine Aug 20 02:40:37 ecs-debian-9 cloud-init(10021: CLOUDINIT) util, putDEBUG1: ecod-init for August
Laaded: loaded (/lik-gustend/susten/cloud-init.service: enabled: uendor preset: enabled) Active: active fexited) since Mno 2018-08-20 02:43640 UTC: 35 ago Process: 1096 Ecode-serited, status=0/SUCCESS) Tasks: 0 (linit: 4915) CGroup: /system.slice/cloud-initservice Aug 20 02:48:40 ecs-debian-9 cloud-init110961: [CLOUDINIT] helpers.pufDEBUG1: config-ca-certs already ran (freq=once-per-instanc Aug 20 02:48:40 ecs-debian-9 cloud-init110961: [CLOUDINIT] helpers.pufDEBUG1: config-ca-certs already ran (freq=once-per-instanc Aug 20 02:48:40 ecs-debian-9 cloud-init110961: [CLOUDINIT] helpers.pufDEBUG1: config-ca-certs already ran (freq=once-per-instanc Aug 20 02:48:40 ecs-debian-9 cloud-init10961: [CLOUDINIT] helpers.pufDEBUG1: config-ca-certs already ran (freq=once-per-instanc Aug 20 02:48:40 ecs-debian-9 cloud-init10961: [CLOUDINIT] helpers.pufDEBUG1: config-ca-certs already ran (freq=once-per-instance Aug 20 02:48:40 ecs-debian-9 cloud-init10961: [CLOUDINIT] helpers.pufDEBUG1: config-ca-certs already ran (freq=once-per-instance Aug 20 02:48:40 ecs-debian-9 cloud-init10961: [CLOUDINIT] helpers.pufDEBUG1: config-ca-certs-already ran (freq=once-per-instance Aug 20 02:48:40 ecs-debian-9 cloud-init10961: [CLOUDINIT] helpers.pufDEBUG1: config-ca-certs-groups already ran (freq=once-per-instance Aug 20 02:48:40 ecs-debian-9 cloud-init10961: [CLOUDINIT] totad-init10EBUG1: Ran 13 nodules with 0 failures Aug 20 02:48:40 ecs-debian-9 cloud-init10961: [CLOUDINIT] tot1, pufDEBUG1: cloud-init 0 proc-vuptime Aug 20 02:48:40 ecs-debian-9 cloud-init10961: [CLOUDINIT] tot1, pufDEBUG1: cloud-init 0 init 0 koz 2.657 seconds (2.66) Aug 20 02:48:40 ecs-debian-9 sustend[11: Started Initial cloud-init job (netadata service crawler). ■ cloud-config.service - Apply the settings specified in cloud-config uendor preset: enabled) Active: active (exited) since Mon 2018-08-02 02:48:10 UTC: 2s ago Process: 1140 Excestart=xusrbin:cloud-config.service: enabled: uendor preset: enabled) Active: active (exited) since Mon 2018-08-02 02	Aug 20 02:48:37 ecs-debian-9 systemd[1]: Started Initial cloud-init job (pre-networking).
<pre>hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) stages.pujDEBUG): Remaining module rsuslog (Kondule 'cloudinit.config. hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) helpers.putDEBUG): config-rsuslog already ran (freq=once-per-instance hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) stages.pujDEBUG): config-rsuslog already ran (freq=once-per-instance hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) helpers.putDEBUG): config-rsus-groups already ran (freq=once-per-instance hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) helpers.putDEBUG): config-users-groups already ran (freq=once-per-instance hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) cloud-init(DEBUG): Randing from proc-uptime (quietFalse) hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT] util.putDEBUG): cloud-init hole 'init' took 2.657 seconds (2.66) hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT] util.putDEBUG): cloud-init hole 'init' took 2.657 seconds (2.66) hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT] util.putDEBUG): cloud-init hole 'init' took 2.657 seconds (2.66) hug 20 02:48:40 ecs-debian-9 cloud-init(1995): [CLOUDINIT] util.putDEBUG): cloud-init hole 'init' took 2.657 seconds (2.66) hug 20 02:48:40 ecs-debian-9 cloud-init(1995): [cloud-config Loaded: loaded (/i)Fusutend/suster/cloud-config.service: enabled' uendor preset: enabled) Active: active (exited) since Mon 2018-08-20 02:48:41 UTC: 2s ago Process: 1140 Excestort=/susrbin:cloud-indits.putenesconfig (code=exited, status=0/SUCCESS) hain PID: 1140 (code=exited, status=0/SUCCESS) hain PID: 1140 (code=exited, status=0/SUCCESS)</pre>	Loaded: loaded (/lik/system/system/sloud-init.service: enabled: vendor preset: enabled) Active: active (exited) since Mon 2018-00-20 02:48:40 UTC: 3s ago Process: 1096 ExecStart=/usr/bin/sloud-init init (code=exited, status=0/SUCCESS) Main PID: 1096 (code=exited, status=0/SUCCESS) Tasks: 0 (linit: 415)
<pre>Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) helpers/pyUDEBUG1: config-rsyslog already ran (freq=once-per-instaince Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) helpers/pyUDEBUG1: Running module users-groups (cmodule 'cloudinit.con Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) helpers/pyUDEBUG1: Running modules with 0 failures Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) totil/pyUDEBUG1: Run 13 modules with 0 failures Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) util/pyUDEBUG1: Run 13 modules with 0 failures Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) util/pyUDEBUG1: Read 14 bytes from /proc/uptime Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) util/pyUDEBUG1: Read 14 bytes from /proc/uptime Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) util/pyUDEBUG1: Read 14 bytes from /proc/uptime Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) util/pyUDEBUG1: Read 14 bytes from /proc/uptime Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) util/pyUDEBUG1: Read 14 bytes from /proc/uptime Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) util/pyUDEBUG1: cloud-init Aug 2.557 seconds (2.66) Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) util/pyUDEBUG1: cloud-init Aug 2.557 seconds (2.66) Aug 20 02:48:40 ecs-debian-9 cloud-init/py the settings specified in cloud-config Loaded: loaded (/ilk/system/dsystem/cloud-config service: enabled' uendor preset: enabled' Active: active (exited) since Mon 2018-08-20 02:48:41 UTC; 2s ag0 Process: 1140 Excestart=xusr/bin/cloud-init Auglesnode=config (code=exited, status=0/SUCCESS) Main PID: 1140 (code=exited, status=0/SUCCESS) Main PID: 1140 (code=exited, status=0/SUCCESS) Alar S0 (linit: 4515)</pre>	
<pre>hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) stages.pujDEBNG:: Running nodule users-groups (cnodule 'cloudinit.com hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) helpers.pujDEBNG:: config-users-groups (cnodule 'cloudinit.com hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) helpers.pujDEBNG:: config-users-groups (cnodule 'cloudinit.com hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) til, pujDEBNG:: Reading from /proc/uptime hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) util, pujDEBNG:: Reading from /proc/uptime hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) util, pujDEBNG: cloud-init hok 2.657 seconds (2.66) hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT] util, pujDEBNG: cloud-init hok 2.657 seconds (2.66) hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [cloud-init job (netadata service crawler).</pre> I cloud-config.service - Apply the settings specified in cloud-config Loaded: loaded (/i)Fvystend/systew/cloud-config.service: enabled; vendor preset: enabled) Active: active (exited) since Mno 2018-08-20 02:44:11 UTC: 2s ago Process: 1140 ExceStart=/surs-functoud-loud-Exnode=config (code=exited, status=0/SUCCESS) Main PID: 1140 (code=exited, status=0/SUCCESS)	
<pre>hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) helpers.ju/DEBUG1: contig-users-groups already ran (freq=once-per-ins hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) helpers.ju/DEBUG1: Ran 13 modules uith 0 failures hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) uit] pyIDEBUG1: Read 14 bytes from /proc/uptime hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) uit] pyIDEBUG1: Read 14 bytes from /proc/uptime hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) uit] pyIDEBUG1: Read 14 bytes from /proc/uptime hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT) uit] pyIDEBUG1: cloud-init mode 'init' cook 2.657 seconds (2.66) hug 20 02:48:40 ecs-debian-9 systemd(11: Started Initial cloud-init job (metadata service crawler).</pre> I cloud-config.service - Apply the settings specified in cloud-config Loaded: loaded (/ilb/systemd/system/cloud-config.service: enabled; vendor preset: enabled) Artive: active (exited) since Mon 2018-08-20 02:48:41 UTC; 2s ag0 Process: 1140 Excestart=/susre/binceloud-load=14 UTC; 2s ag0 Hain PID: 1140 (code=exited, status=0/SUCCESS) Hain PID: 1140 (code=exited, status=0/SUCCESS)	
<pre>hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT] cloud-init(DEBUG): Reading from yonc-uptime (quiet=False) hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT] util, pujDEBUG): Reading from yonc-uptime (quiet=False) hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT] util, pujDEBUG): cloud-init node 'init' took 2.657 seconds (2.66) hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT] util, pujDEBUG): cloud-init node 'init' took 2.657 seconds (2.66) hug 20 02:48:40 ecs-debian-9 cloud-init(1096): [CLOUDINIT] util, pujDEBUG): cloud-init node 'init' took 2.657 seconds (2.66) hug 20 02:48:40 ecs-debian-9 systemd[1]: Started Initial cloud-init job (netadata service crawler).</pre> I cloud-config.service - Apply the settings specified in cloud-config Loaded: loaded (/i)/systemd/system/cloud-config.service: enabled; vendor preset: enabled) Active: active (exited) since Mon 2018-08-20 02:48:41 UTC; 2s ago Process: 1140 Excestart=/surs/bin/cloud-init nodulesnode=config (code=exited, status=0/SUCCESS) Main PID: 1140 (code=exited, status=0/SUCCESS)	
 Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) util,pijDEBNG1: Read 14 bytes from /procruptime Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) util,pijDEBNG1: cloud-init mode 'init' took 2.657 seconds (2.66) Aug 20 02:48:40 ecs-debian-9 systemd(11: Started Initial cloud-init job (metadata service crawler). ■ cloud-config.service - Apply the settings specified in cloud-config Loaded: loaded (/ih/systemd/system/cloud-config.service: enabled) Active: active (exited) since Mon 2018-08-20 02:48:41 UTC: 2s ago Process: 1140 Excestart-zusr>hinc:loud-loudels: -mode=config (code=exited, status=0/SUCCESS) Main PID: 1140 (code=exited, status=0/SUCCESS) 	Aug 20 02:48:40 ecs-debian-9 cloud-init[1096]: [CLOUDINIT] cloud-init[DEBUG]: Ran 13 modules with 0 failures
Aug 20 02:48:40 ecs-debian-9 cloud-init(10961: CLOUDINIT) util.jú[DEBUG]: cloud-init mode 'init' took 2.657 seconds (2.66) Aug 20 02:48:40 ecs-debian-9 systemd[11: Started Initial cloud-init job (metadata service creaver). ■ cloud-config .service - Apply the settings specified in cloud-config Loaded: loaded (/ilb/systemd/system/cloud-config.service; enabled; vendor preset: enabled) Active: active (exited) since Non 2018-08-20 02:48:41 UTC; 2s ago Process: 1140 Exectort=/sury-Dinz-cloud-init modulesmode=config (code=exited, status=0/SUCCESS) Main PID: 1140 (code=exited, status=0/SUCCESS) Tasks: 0 (linit: 915)	
Aug 20 02:48:40 ecs-debian-9 systemd[11: Started Initial cloud-init job (metadata service crawler). ■ cloud-config.service - Apply the settings specified in cloud-config Loaded (:livesystemd'system/cloud-config.service: enabled; vendor preset: enabled) Active: active (exited) since Mon 2018-08-20 02:48:41 UTC; 2s ago Process: 1140 Excestort=/surs/bin/cloud-init modulesnode=config (code=exited, status=0/SUCCESS) Main PID: 1140 (code=exited, status=0/SUCCESS) Tasks: 0 (linit: 915)	
<pre>cloud-config.service - Apply the settings specified in cloud-config Loaded: loaded (/lib/system/system/cloud-config.service; enabled; vendor preset: enabled) Active: active (exited) since Mon 2018-08-20 02:46:41 UTC; 2s ago Process: 1140 Exectort=/sury-bin-cloud-init modulesmode=config (code=exited, status=0/SUCCESS) Main PID: 1140 (code=exited, status=0/SUCCESS) Tasks: 0 (linit: 915)</pre>	
Loaded: loaded (/lik/system/system/sloud-config.service: enabled; vendor preset: enabled) Active: active (exited) since Mon 2018-08-20 02:48:41 UTC; 2s ago Process: 1140 ExceStart=/usr/bin/sloud-init modulesmode=config (code=exited, status=0/SUCCESS) Main PID: 1140 (code=exited, status=0/SUCCESS) Tasks: 0 (linit: 915)	Aug 20 02:48:40 ecs-debian-9 systemdill: Started Initial cloud-init job (metadata service crawler).
	Loaded: loaded (/lix/syttem/cygstem/cfoud-config.service: enabled; vendor preset: enabled) Active: active (exited) since Mon 2018-08-20 02:48:41 UTC; 2s ago Process: 1140 ExectStart=/usr/bln/cloud-init modulesmode-config (code=exited, status=0/SUCCESS) Main PID: 1140 (code=exited, status=0/SUCCESS) Tasks: 0 (linit: 415)

CentOS 7/Fedora 28: Required C Compiler Not Installed

• Symptom

After Cloud-Init is successfully installed, you run the following command:

cloud-init init --local

The following information is displayed:

/usr/lib/python2.5/site-packages/Cheetah/Compiler.py:1532: UserWarning: You don't have the C version of NameMapper installed! I'm disabling Cheetah's useStackFrames option as it is painfully slow with the Python version of NameMapper. You should get a copy of Cheetah with the compiled C version of NameMapper. "\nYou don't have the C version of NameMapper installed!

Cause analysis

This alarm is generated because C version of NameMapper needs to be compiled when Cloud-Init is installed. However, GCC is not installed in the system, and the compilation cannot be performed. As a result, NameMapper is missing.

Solution

Run the following command to install GCC:

yum -y install gcc

Reinstall Cloud-Init.

CentOS 7/Fedora: Failed to Use the New Password to Log In to an ECS Created from an Image

Symptom

You cannot use a new password to log in to an ECS created from an image with Cloud-Init installed. After logging in to the ECS using the old password, you find that NICs of the ECS are not started.

Figure 1	2-8 NIC	not	started
----------	---------	-----	---------



Solution

Log in to the ECS used to create that image, open the DHCP configuration file **/etc/sysconfig/network-scripts/ifcfg-eth***X*, and comment out **HWADDR**.

12.2 What Can I Do with a Cloud-Init ECS?

Introduction to Cloud-Init

Cloud-Init is an open-source tool for cloud instance initialization. When creating ECSs from an image with Cloud-Init, you can use user data injection to customize initialization details (for example, an ECS login password) to the ECSs. You can also configure and manage a running ECS by querying and using metadata. If Cloud-Init is not installed, you cannot apply these custom configurations to your ECSs, and you will have to use the original password in the image file to log in to the ECSs.

Installation Methods

You are advised to install Cloud-Init or Cloudbase-Init on the ECS to be used to create a private image so that new ECSs created from this private image can be customized.

• For Linux, download and install Cloud-Init.

For how to install and configure Cloud-Init, see **Installing and Configuring Cloud-Init**.

12.3 What Do I Do If Installed NetworkManager and Now I Can't Inject the Key or Password Using Cloud-Init?

Cause

One likely possibility is that the version of Cloud-Init is incompatible with that of NetworkManager. In Debian 9.0 and later versions, NetworkManager is incompatible with Cloud-Init 0.7.9.

Solution

Uninstall the current Cloud-Init and install Cloud-Init 0.7.6 or an earlier version.

For details about how to install Cloud-Init, see Installing Cloud-Init.

12.4 How Do I Install growpart for SUSE 11 SP4?

Scenarios

For SUSE and openSUSE, growpart is an independent tool and is not included in a **cloud-*** package. You need to install it separately.

Procedure

1. Check whether Cloud-Init and growpart are installed.

rpm -qa | grep cloud-init

If cloud-init is installed, the command output should be similar to the following:

cloud-init-0.7.8-39.2

rpm -qa | grep growpart

If growpart is installed, the command output should be similar to the following:

growpart-0.29-8.1

2. If they are installed, uninstall them.

zypper remove cloud-init growpart

3. Delete residual files.

rm -fr /etc/cloud/*

rm -fr /var/lib/cloud/*

4. Install growpart.

zypper install http://download.opensuse.org/repositories/home:/garloff:/ OTC:/cloudinit/SLE_11_SP4/noarch/growpart-0.27-1.1.noarch.rpm 5. Install python-oauth.

zypper install http://download.opensuse.org/repositories/home:/garloff:/ OTC:/cloudinit/SLE_11_SP4/x86_64/python-oauth-1.0.1-35.1.x86_64.rpm

6. Install Cloud-Init.

zypper install http://download.opensuse.org/repositories/home:/garloff:/ OTC:/cloudinit/SLE_11_SP4/x86_64/cloud-init-0.7.6-27.23.1.x86_64.rpm

7. Check whether growpart, python-oauth, and Cloud-Init are installed successfully.

rpm -qa | grep growpart

If growpart is installed, the command output should be similar to the following:

growpart-0.27-1.1

rpm -qa | grep python-oauth

If python-oauth is installed, the command output should be similar to the following:

python-oauthlib-0.6.0-1.5 python-oauth-1.0.1-35.1

rpm -qa | grep cloud-init

If Cloud-Init is installed, the command output should be similar to the following:

cloud-init-0.7.6-27.19.1

- 8. Check the configurations.
 - chkconfig cloud-init-local on;chkconfig cloud-init on;chkconfig cloudconfig on;chkconfig cloud-final on

13_{ECS Creation}

13.1 Can I Change the Image of a Purchased ECS?

Yes.

If you have selected the wrong image or your service requirements have changed, you can change the image of your ECS.

You can change the image type (public, private, and shared images) and OS. For details, see **Changing the OS**.

13.2 Can I Change the Specifications Defined by a Private Image When I Use the Image to Create an ECS?

Yes. You can specify the CPU, memory, bandwidth, system and data disks of the ECS you are creating. The system disk must be smaller than 1,024 GB but no less than the system disk capacity in the image.

13.3 Can I Specify the System Disk Capacity When I Create an ECS Using an Image?

Yes, but you cannot create a system disk smaller than the original and the maximum allowed is 1,024 GB.

13.4 What Do I Do If a Partition Is Not Found During the Startup of an ECS Created from an Imported Private Image?

Cause

Disk partition IDs are changed after the cross-platform image import. As a result, no partition can be found at startup. In this case, you need to change disk partition IDs in the image to UUID.

Solution

openSUSE 13.2 is used as an example.

1. Check disk partition IDs.

ls -l /dev/disk/by-id/

An output similar to the following should be seen:

total 0

```
lrwxrwxrwx 1 root root 10 Jul 22 01:35 ata-QEMU_HARDDISK_QM00001 -> ../../xvda
lrwxrwxrwx 1 root root 11 Jul 22 01:35 ata-QEMU_HARDDISK_QM00001-part1 -> ../../xvda1
lrwxrwxrwx 1 root root 12 Jul 22 01:35 ata-QEMU_HARDDISK_QM00001-part10 -> ../../xvda10
lrwxrwxrwx 1 root root 11 Jul 22 01:35 ata-QEMU_HARDDISK_QM00001-part2 -> ../../xvda2
lrwxrwxrwx 1 root root 11 Jul 22 01:35 ata-QEMU_HARDDISK_QM00001-part5 -> ../../xvda5
lrwxrwxrwx 1 root root 11 Jul 22 01:35 ata-QEMU_HARDDISK_QM00001-part6 -> .././xvda6
lrwxrwxrwx 1 root root 11 Jul 22 01:35 ata-QEMU_HARDDISK_QM00001-part7 -> ../../xvda7
lrwxrwxrwx 1 root root 11 Jul 22 01:35 ata-QEMU_HARDDISK_QM00001-part8 -> ../../xvda8
lrwxrwxrwx 1 root root 11 Jul 22 01:35 ata-QEMU_HARDDISK_QM00001-part9 -> .././xvda9
lrwxrwxrwx 1 root root 10 Jul 22 01:35 ata-QEMU_HARDDISK_QM00005 -> ../../xvde
lrwxrwxrwx 1 root root 10 Jul 22 01:35 scsi-SATA QEMU HARDDISK QM00001 -> ../../xvda
lrwxrwxrwx 1 root root 11 Jul 22 01:35 scsi-SATA_QEMU_HARDDISK_QM00001-part1 -> ../../xvda1
lrwxrwxrwx 1 root root 12 Jul 22 01:35 scsi-SATA_QEMU_HARDDISK_QM00001-part10 -> ../../xvda10
lrwxrwxrwx 1 root root 11 Jul 22 01:35 scsi-SATA_QEMU_HARDDISK_QM00001-part2 -> ../../xvda2
lrwxrwxrwx 1 root root 11 Jul 22 01:35 scsi-SATA_QEMU_HARDDISK_QM00001-part5 -> .././xvda5
lrwxrwxrwx 1 root root 11 Jul 22 01:35 scsi-SATA_QEMU_HARDDISK_QM00001-part6 -> .././xvda6
lrwxrwxrwx 1 root root 11 Jul 22 01:35 scsi-SATA_QEMU_HARDDISK_QM00001-part7 -> ../../xvda7
Irwxrwxrwx 1 root root 11 Jul 22 01:35 scsi-SATA QEMU HARDDISK QM00001-part8 -> ../../xvda8
lrwxrwxrwx 1 root root 11 Jul 22 01:35 scsi-SATA_QEMU_HARDDISK_QM00001-part9 -> ../../xvda9
lrwxrwxrwx 1 root root 10 Jul 22 01:35 scsi-SATA_QEMU_HARDDISK_QM00005 -> ../../xvde
```

ata-QEMU_HARDDISK_*xxx* and **scsi-SATA_QEMU_HARDDISK**_*xxx* indicate that the ECS disks are simulated using Quick EMUlator (QEMU). The content on the left of -> are disk partition IDs, and that on the right of -> are partition names.

2. Check disk partition UUIDs.

ls -l /dev/disk/by-uuid/

An output similar to the following should be seen:

```
total 0
```

```
Irwxrwxrwx 1 root root 11 Jul 22 01:35 45ecd7a0-29da-4402-a017-4564a62308b8 -> ../../xvda5
Irwxrwxrwx 1 root root 11 Jul 22 01:35 55386c6a-9e32-41d4-af7a-e79596221f51 -> ../../xvda9
Irwxrwxrwx 1 root root 11 Jul 22 01:35 55f36660-9bac-478c-a701-7ecc5347f789 -> ../../xvda8
Irwxrwxrwx 1 root root 11 Jul 22 01:35 780f36bc-0ada-4c98-9a8d-44570d65333d -> ../../xvda1
Irwxrwxrwx 1 root root 11 Jul 22 01:35 b3b7c47f-6a91-45ef-80d6-275b1cc16e19 -> ../../xvda6
Irwxrwxrwx 1 root root 11 Jul 22 01:35 ea63b55d-3b6e-4dcd-8986-956b72bac3e9 -> ../../xvda7
Irwxrwxrwx 1 root root 12 Jul 22 01:35 eb3cc645-925e-4bc5-bedf-c2a6f3b65809 -> ../../xvda10
```

The content on the left of -> are disk partition UUIDs, and that on the right of -> are partition names. Based on the outputs in **1** and this step, you can obtain the mappings between the partition names, IDs, and UUIDs.

3. Open the **/etc/fstab** file to check partition names.

vi /etc/fstab

An example command output is as follows:

/dev/disk/by-id/scsi-SATA_QEMU_HARDDISK_QM00001-part5 / ext3 defaults,errors=panic 1 1 /dev/disk/by-id/scsi-SATA_QEMU_HARDDISK_QM00001-part1 /boot ext3 defaults,errors=panic 1 2 /dev/disk/by-id/scsi-SATA_QEMU_HARDDISK_QM00001-part6 /home ext3 nosuid,errors=panic 1 2 /dev/disk/by-id/scsi-SATA_QEMU_HARDDISK_QM00001-part10 /opt ext3 defaults,errors=panic 1 2 /dev/disk/by-id/scsi-SATA_QEMU_HARDDISK_QM00001-part7 /tmp ext3 nodev,nosuid,errors=panic 1 2 /dev/disk/by-id/scsi-SATA_QEMU_HARDDISK_QM00001-part7 /tmp ext3 nodev,nosuid,errors=panic 1 2 /dev/disk/by-id/scsi-SATA_QEMU_HARDDISK_QM00001-part9 /usr ext3 defaults,errors=panic 1 2 /dev/disk/by-id/scsi-SATA_QEMU_HARDDISK_QM00001-part8 /var ext3 nodev,nosuid,errors=panic 1 2 sysfs /sys sysfs noauto 0 0 proc /proc proc defaults 0 0 usbfs /proc/bus/usb usbfs noauto 0 0 devpts /dev/pts devpts mode=0620,gid=5 0 0 /dev/cdrom /media/ udf,iso9660 noexec,noauto,nouser,nodev,nosuid 1 2 tmpfs /dev/shm tmpfs noexec,nodev,nosuid 0 0

The values in the first column are the disk partition IDs.

4. Press i to enter editing mode. Change each disk partition ID to a UUID based on the outputs in 1 and 2.

The modified content is as follows.

UUID=45ecd7a0-29da-4402-a017-4564a62308b8 / ext3 defaults,errors=panic 1 1 UUID=780f36bc-0ada-4c98-9a8d-44570d65333d /boot ext3 defaults,errors=panic 1 2 UUID=b3b7c47f-6a91-45ef-80d6-275b1cc16e19 /home ext3 nosuid,errors=panic 1 2 UUID=eb3cc645-925e-4bc5-bedf-c2a6f3b65809 /opt ext3 defaults,errors=panic 1 2 UUID=ea63b55d-3b6e-4dcd-8986-956b72bac3e9 /tmp ext3 nodev,nosuid,errors=panic 1 2 UUID=55386c6a-9e32-41d4-af7a-e79596221f51 /usr ext3 defaults,errors=panic 1 2 UUID=55f36660-9bac-478c-a701-7ecc5347f789 /var ext3 nodev,nosuid,errors=panic 1 2 sysfs /sys sysfs noauto 0 0 proc /proc proc defaults 0 0 usbfs /proc/bus/usb usbfs noauto 0 0 devpts /dev/pts devpts mode=0620,gid=5 0 0 /dev/cdrom /media/ udf,iso9660 noexec,noauto,nouser,nodev,nosuid 1 2 tmpfs /dev/shm tmpfs noexec,nodev,nosuid 0 0

NOTE

Ensure that the UUIDs are correct, or the ECS will be unable to start up normally.

- 5. Press **Esc**, enter **:wq**, and press **Enter**. The system saves the configuration and exits the editor.
- 6. Check the partition names in the system boot configuration file.

The system boot configuration file varies depending on the OS. Confirm the boot configuration file of the current OS.

- GRUB file
 - /boot/grub/grub.conf
 - /boot/grub/menu.lst
 - /boot/grub/grub.cfg
 - /boot/grub2/grub.cfg
- Syslinux configuration file
 - /extlinux.conf

- /boot/syslinux/extlinux.conf
- /boot/extlinux/extlinux.conf
- /boot/syslinux/syslinux.cfg
- /syslinux/syslinux.cfg
- /syslinux.cfg

The boot file in this example is **/boot/grub/menu.lst**. Run the following command to check it:

vi /boot/grub/menu.lst

```
default 0
timeout 3
title xxx Server OS - xxxxxx
kernel /boot/vmlinuz-3.0.101-0.47.52-default root=/dev/disk/by-id/scsi-
SATA_QEMU_HARDDISK_QM00001-part5 resume= memmap=0x2000000$0x3E000000
nmi_watchdog=2 crashkernel=512M-:256M console=ttyS0,115200 console=tty0 xen_emul_unplug=all
initrd /boot/initrd-3.0.101-0.47.52-default
```

7. Press **i** to enter editing mode and change the partition names in the system boot configuration file.

Change each disk partition name in the **/boot/grub/menu.lst** file in 6 to **UUID=***UUID* of the disk partition based on the query results in 1 and 2.

```
default 0
timeout 3
title xxx Server OS - xxxxx
kernel /boot/vmlinuz-3.0.101-0.47.52-default root=UUID=45ecd7a0-29da-4402-a017-4564a62308b8
resume= memmap=0x2000000$0x3E00000 nmi_watchdog=2 crashkernel=512M-:256M
console=ttyS0,115200 console=tty0 xen_emul_unplug=all
initrd /boot/initrd-3.0.101-0.47.52-default
```

8. Press **Esc**, enter **:wq**, and press **Enter**. The system saves the configuration and exits the vi editor.

13.5 What Do I Do If the Disks of a CentOS ECS Created from an Image Cannot Be Found?

Symptom

When you started a CentOS ECS, the system cannot find disks. Generally, this is because the xen-blkfront.ko module was not loaded during the startup. You need to modify OS kernel startup parameters.

Figure 13-1 Startup screen

 DK] Started Show Plymouth Boot Screen. OK] Reached target Paths. OK] Reached target Basic System. Incut-initqueue[465]: Warning: Could not boot. Iracut-initqueue[465]: Warning: /dev/disk/by-uuid/545e232a-f59b-4576-af34-eccb82 Pea3d2 does not exist Starting Dracut Emergency Shell Varning: /dev/disk/by-uuid/545e232a-f59b-4576-af34-eccb829ea3d2 does not exist
Generating "/run/initramfs/rdsosreport.txt"
Entering emergency mode. Exit the shell to continue. Type "journalctl" to view system logs. You might want to save "/run/initramfs/rdsosreport.txt" to a USB stick or /boot after mounting them and attach it to a bug report.
lracut:/# _

Solution

Modify OS kernel boot parameters.

NOTE

These operations can only be performed after a normal OS startup. So, perform them in the source ECS of the image instead of the current ECS.

1. Run the following command to log in to the OS:

lsinitrd /boot/initramfs- `uname -r `img |grep -i xen

- If the command output contains xen-blkfront.ko, contact the customer service.
- If no command output is displayed, go to 2.
- 2. Back up the GRUB file.
 - If the ECS runs CentOS 6, run the following command:

cp /boot/grub/grub.conf /boot/grub/grub.conf.bak

- If the ECS runs CentOS 7, run the following command:

cp /boot/grub2/grub.cfg /boot/grub2/grub.cfg.bak

3. Use the **vi** editor to open the GRUB configuration file. Run the following command (CentOS 7 as an example):

vi /boot/grub2/grub.cfg

4. Add **xen_emul_unplug=all** to the default boot kernel.

NOTE

Search for the line that contains **root=UUID=** and add **xen_emul_unplug=all** to the end of the line.

```
menuentry 'CentOS Linux (3.10.0-229.el7.x86_64) 7 (Core) with debugging' --class centos --class gnu-
linux --class gnu --class os --unrestricted $menuentry_id_option 'gnulinux-3.10.0-229.el7.x86_64-
advanced-bf3cc825-7638-48d8-8222-cd2f412dd0de' {
```

```
load_video
set gfxpayload=keep
insmod gzio
insmod part_msdos
insmod ext2
set root='hd0,msdos1'
if [ x$feature_platform_search_hint = xy ]; then
search --no-floppy --fs-uuid --set=root --hint='hd0,msdos1' bf3cc825-7638-48d8-8222-
```

```
cd2f412dd0de
else
search --no-floppy --fs-uuid --set=root bf3cc825-7638-48d8-8222-cd2f412dd0de
fi
linux16 /boot/vmlinuz-3.10.0-229.el7.x86_64 root=UUID=bf3cc825-7638-48d8-8222-
cd2f412dd0de xen_emul_unplug=all ro crashkernel=auto rhgb quiet systemd.log_level=debug
systemd.log_target=kmsg
initrd16 /boot/initramfs-3.10.0-229.el7.x86_64.img
}
```

- 5. Press Esc, enter :wq, and press Enter to exit the vi editor.
- 6. Create an image using the ECS, upload and register the image on the cloud, and then use the image to create a new ECS.

13.6 What Do I Do If I Enabled Automatic Configuration During Image Registration for an ECS Created from a Windows Image and Now It Won't Start?

Cause

This may be caused by an issue with offline VirtIO driver injection.

Solution

When you inject VirtIO drivers into the image for a Windows ECS, note that:

- If the boot mode in the image file is UEFI, the VirtIO drivers cannot be injected offline.
- Disable Group Policy Object (GPO) because some policies may cause offline VirtIO driver injection to fail.
- Stop any installed antivirus software. They may cause offline VirtIO driver injection to fail.

To install VirtIO drivers, see **Preparations for Creating a Windows Private Image**.

13.7 What Do I Do If an Exception Occurs When I Start an ECS Created from an Image Using UEFI Boot?

Symptom

An ECS created from a private image booting to UEFI cannot start.

Possible Causes

The image is configured for UEFI boot, but the uefi attribute was not added to the image.

Solution

- 1. Delete the ECS that failed to start.
- Call the API to update the image attributes and change the value of hw_firmware_type to uefi.

API URI: PATCH /v2/cloudimages/{*image_id*}

For details about how to call the API, see **Updating Image Information**.

3. Use the updated image to create an ECS.

14 Billing

14.1 What Are the Billing Standards for IMS?

Image Type	Billing	
Private	 System disk image, ISO image, or data disk image: free OBS buckets are used in the background for free image storage. 	
	• Full-ECS image: charged for the storage on CSBS or CBR If the image is created from a CSBS or CBR backup, you will be billed for the space used to store the CSBS or CBR backup.	
Shared	• If it is a system or data disk image, the OBS storage is free.	
	• If it is a full-ECS image, you will be billed for the space used to store the CBR backup.	
	NOTE Full-ECS images created from a CBR backup can be shared but those created from a CSBS backup cannot.	

Billing of Different Image Types

Billing of Private Image Operations

Operation	Billing Item
Uploading an external image file to an OBS bucket	OBS storage (You are advised to delete unneeded image files.)

Operation	Billing Item
Creating an ECS from a system disk image or full-	EVS disks attached to the created ECS (The system disk image and full-ECS image are free of charge.)
ECS image	For example, when you use a full-ECS image (with a 40 GB system disk and 100 GB data disk) to create an ECS, you will be charged based on the total capacity of the system disk and data disk.
Replication of a system or data disk image within a region	None. OBS buckets are available for free to store the image copies.
Exporting images	OBS space for storing the exported images

Reference

- For OBS pricing details, see **OBS Pricing Details**.
- For CSBS pricing details, see CSBS Pricing Details.
- For CBR pricing details, see CBR Pricing Details.
- For EVS pricing details, see EVS Pricing Details.