# Alibaba Cloud

Elastic Compute Service Images

Document Version: 20220712

(-) Alibaba Cloud

# Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

- You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloudauthorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
- 2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.
- 3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
- 4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).
- 5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.
- 6. Please directly contact Alibaba Cloud for any errors of this document.

# **Document conventions**

Style	Description	Example
<u>↑</u> Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	Danger:  Resetting will result in the loss of user configuration data.
<u> </u>	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	Warning: Restarting will cause business interruption. About 10 minutes are required to restart an instance.
Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	Notice:  If the weight is set to 0, the server no longer receives new requests.
? Note	A note indicates supplemental instructions, best practices, tips, and other content.	Note: You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click Settings> Network> Set network type.
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click OK.
Courier font	Courier font is used for commands	Run the cd /d C:/window command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	bae log listinstanceid  Instance_ID
[] or [a b]	This format is used for an optional value, where only one item can be selected.	ipconfig [-all -t]
{} or {a b}	This format is used for a required value, where only one item can be selected.	switch {active stand}

# **Table of Contents**

1.1	Image overview	- 09
2.	Select an image	- 11
3.	Find an image	- 16
4.	Public image	<del>-</del> 19
	4.1. Overview	- 19
	4.2	- 26
5.	Alibaba Cloud Linux	- 82
	5.1. Overview	- 82
	5.2. Release notes	- 88
	5.3. Use images	- 88
	5.3.1. Use Alibaba Cloud Linux 3 images in an on-premises e	- 88
	5.3.2. Use Alibaba Cloud Linux 2 images in an on-premises e	- 93
	5.3.3. Use Alibaba Cloud Linux 2	- 97
	5.3.4. Install and enable curl of a later version	100
	5.4. Security	100
	5.4.1. CIS Alibaba Cloud Linux benchmarks	100
	5.4.2. Subscribe to announcements for Alibaba Cloud Linux 2	104
	5.4.3. Use YUM to perform security updates (Alibaba Cloud L	105
	5.4.4. Use YUM to perform security updates (Alibaba Cloud L	110
	5.4.5. Manage Alibaba Cloud Linux security update notificati	114
	5.4.6. Use confidential containers based on Inclavare Contain	116
	5.4.7. Use Inclavare Containers to implement remote attestati	121
	5.5. Kernel Live Patching	133
	5.5.1. Overview of the Kernel Live Patching feature	133
	5.5.2. Operations related to kernel hotfixes	134
	5.5.3. Use livepatch-mgr	138

5.6. Features and interfaces supported by Alibaba Cloud Linux	145
5.6.1. Overview	145
5.6.2. Memcg Exstat feature	150
5.6.3. Memcg global minimum watermark rating	153
5.6.4. Memcg QoS feature of the cgroup v1 interface	155
5.6.5. Memcg backend asynchronous reclaim	156
5.6.6. THP reclaim	159
5.6.7. TCP-RT configurations	162
5.6.8. Change the TCP TIME-WAIT timeout period	170
5.6.9. Enable the cgroup writeback feature	171
5.6.10. Configure the weight-based throttling feature of blk-i	174
5.6.11. Enhance the monitoring of block I/O throttling	179
5.6.12. Use the JBD2 optimization interface	181
5.6.13. Create a hard link across project quotas	182
5.6.14. Track I/O latency	183
5.6.15. Detect I/O hangs of file systems and block layers	185
5.6.16. Enable the PSI feature for cgroup v1	188
5.6.17. Enable the CPU burst feature for cgroup v1	189
5.6.18. Group identity feature	191
5.6.19. Statistical status change of anonymous pages in the	197
5.6.20. Enhance the visibility of container resources	198
5.6.21. SMC-R	201
5.6.22. UKFEF	209
5.7. Operating System Migration Tool	213
5.7.1. Migrate from CentOS 8 to Alibaba Cloud Linux 3	213
5.7.2. Migrate from CentOS 7 to Alibaba Cloud Linux 2	218
5.8. Known issues in Alibaba Cloud Linux	220
5.9. Knowledge base	222

	5.10. Alibaba Cloud Linux FAQ	223
6	.Custom image	226
	6.1. Overview	226
	6.2. Create custom image	227
	6.2.1. Create a custom image from a snapshot	228
	6.2.2. Create a custom image from an instance	230
	6.2.3. Use Packer to create a custom image	233
	6.2.4. Use Packer to create and import an on-premises image	235
	6.2.5. Use OOS to create a custom image	241
	6.3. Import images	242
	6.3.1. Image import procedure	242
	6.3.2. Use the image compliance tool	243
	6.3.3. Instructions for importing images	248
	6.3.4. Install cloud-init	252
	6.3.5. Install the virtio driver	264
	6.3.6. Customize Linux images	271
	6.3.7. Convert the image file format	277
	6.3.8. Import custom images	278
	6.4. Update a custom image	283
	6.5. Copy a custom image	285
	6.6. Share or unshare a custom image	288
	6.7. Export a custom image	290
	6.8. Modify a custom image	292
	6.9. Edit the tags of a custom image	292
	6.10. Delete a custom image	293
7.	Community images	295
	7.1. Overview	295
	7.2. Publish a community image	296

7.3. Search for a community image	296
7.4. Use a community image	297
7.5. Unpublish a community image	298
8.Alibaba Cloud Marketplace images	300
9.Image family	302
9.1. Overview	302
9.2. Set an image family	303
9.3. Create an ECS instance from an image family	306
9.4. Set the state of an image to Deprecated or restore an im	308
10.Image Builder	311
10.1. Overview	311
10.2. Commands supported by Image Builder	313
10.3. Use Image Builder	314
10.3.1. Create an image component	315
10.3.2. Create an image template	315
10.3.3. View and execute an image template	318
10.3.4. Delete an image component	319
10.3.5. Delete an image template	320
10.4. Best practices for Image Builder	320
11.Change the operating system	325
12.End of support for operating systems	326
13.FAQ	333
13.1. Image FAQ	333
13.2. Manually update virtio drivers of Windows instances	368
13.3. Options for dealing with CentOS Linux EOL	381
13.4. Change CentOS 8 repository addresses	384
13.5. Change the CentOS 6 source address	386
13.6. End-of-life migration plan for CentOS Linux	389

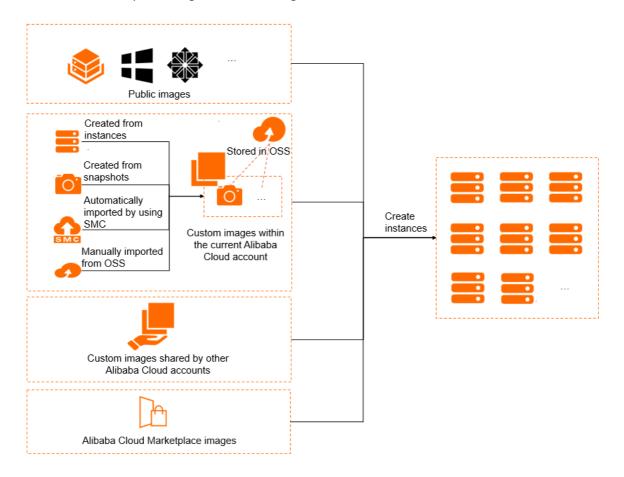
13.7. Security updates after the end of support and maintenan	392
13.8. Share encrypted custom images	393
13.9. Adapt Linux custom images to NVMe-based system disks	
13.10. Upgrade the SUSE Linux Enterprise Server public image	411
13.11. Linux time and time zones	414
13.12. Compatibility of AMD instances	415
13.13. Common image formats	416
13.14. Install GRUB on a Linux server	417
13.15. Manage Windows Server Semi-Annual Channel images a	420
13.16. Collect kdump information after the operating system of	425
13.17. Red Hat Cloud Access and Windows licenses	426
13.17.1. Red Hat Cloud Access	426
13.17.1.1. Overview	426
13.17.1.2. Step 1: Enroll in the Red Hat Cloud Access progra	428
13.17.1.3. Step 2: Create a RHEL image	429
13.17.1.4. Step 3: (Optional) Convert the image to a suppor	437
13.17.1.5. Step 4: Import the image to Alibaba Cloud ECS	437
13.17.1.6. Step 5: Create an ECS instance from the RHEL im	440
13.17.2. Windows licenses	441
13.17.2.1. Microsoft licenses and Software Assurance	441
13.17.2.2. Apply for License Mobility	443

# 1.Image overview

An Elastic Compute Service (ECS) image provides the information required to create an ECS instance. You must select an image when you create an ECS instance. An image is a copy of data from one or more disks. An instance image can contain data from only the system disk or from both system and data disks.

#### Image types

ECS images are classified into public images, custom images, shared images, community images, and Alibaba Cloud Market place images based on image sources.



#### • Public images

Public images provided by Alibaba Cloud are licensed, secure, and stable. Public images for Windows Server operating systems and mainstream Linux operating systems are provided. For more information, see Overview.

#### Cust om images

Custom images are created from instances or snapshots or imported from your computer. For more information, see Overview.

#### Shared images

Shared images are shared to you by other Alibaba Cloud accounts. Images shared to an account do not count against the image quota for the account. The account is not charged for images shared to it. For more information, see Share or unshare a custom image.

Alibaba Cloud Market place images

Alibaba Cloud Market place images are classified into the following types based on their providers:

- o Images provided by Alibaba Cloud
- Images provided by independent software vendors (ISVs) that are licensed by Alibaba Cloud
   Market place

An Alibaba Cloud Marketplace image contains an operating system and pre-installed software. The operating system and pre-installed software are thoroughly tested by the ISV and Alibaba Cloud to ensure that the image is safe to use. For more information, see Alibaba Cloud Marketplace images.

• Community images

Community images are publicly available. Custom images can be published as community images for other users to obtain and use. For more information, see Overview.

#### **Image prices**

You may be charged for using images. The following table describes the billing methods for different image types.

For more information, see Images and the "FAQ about the commercial availability of images" section of the Image FAQ topic.

## Select and find images

You can select an image based on the region, image type, image price, operating system, and built-in software. For more information, see Select an image.

You can find an image based on its type, name, ID, or the ID of the snapshot from which the image was created. Then, you can use the image to create instances or perform other operations. For more information, see Find an image.

## Replace the image of an instance

After you create an instance, you can replace its operating system by replacing its system disk.

- You can replace the image of the instance with a public image when you replace the system disk. For more information, see Replace the operating system of an instance by using a public image.
- You can also replace the image of the instance with a non-public image such as a custom, shared, or Alibaba Cloud Market place image when you replace the system disk. For more information, see Replace the operating system of an instance by using a non-public image.

#### Limits

For information about the limits and quotas of images, see the "Image limits" section of the Limits topic.

## **API** operations

You can call API operations to manage images. For more information, see the "Images" section of the List of operations by function topic.

# 2. Select an image

This topic describes how to select an appropriate image from multiple image types and operating systems to suit your business needs. You must select an image when you create an Elastic Compute Service (ECS) instance.

When you select an image, you must consider the following factors:

- Region
- Image type
- Image fee
- Operating system
- Built-in software (such as MySQL)

#### Region

An image is tied to a region and can be used to create instances only within the region. For example, if you want to create an instance in the China (Beijing) region, you can use images only in the China (Beijing) region. For more information about regions, see Regions and zones.

If you want to use an image that belongs to Region A to create instance in Region B, you must first copy the image to Region B. For more information, see Copy a custom image.

#### Image type

ECS images are classified into public images, custom images, shared images, community images, and Alibaba Cloud Market place images based on image sources. For more information, see Image types.

# Image fee

You may be charged for images that you use. For more information, see Image prices.

## Operating system

When you select an operating system, you must consider the following factors:

• Operating system architecture: 32-bit or 64-bit

Operating system architecture	Applicable memory	Limits
32-bit	Up to 4 GiB memory	<ul> <li>If the memory of an instance type is larger than 4 GiB, you cannot use a 32-bit operating system.</li> <li>A 32-bit Windows operating system supports up to four vCPUs.</li> </ul>
64-bit	Up to 4 GiB memory	If you want to use a memory of at least 4 GiB for your applications, use a 64-bit operating system.

• Operating system type: Windows, Linux, or UNIX-like

Operating system type	Connection method	Feature	Scenario
Windows	Remote Desktop Connection	A Windows public image is installed with a genuine activated system.	<ul> <li>Applicable to programs developed based on Windows architectures such as .NET programs.</li> <li>Supports SQL Server and other databases (The databases need to be manually installed).</li> </ul>
Linux and UNIX-like	SSH	<ul> <li>A common, stable, and secure server-side operating system.</li> <li>An open source operating system that provides fast deployment and easy source code compilation.</li> </ul>	<ul> <li>Typically used for server applications, such as highperformance web servers, and supports common programming languages such as PHP and Python.</li> <li>Supports MySQL and other databases (The databases need to be manually installed).</li> </ul>

Alibaba Cloud provides a list of public images that run Windows, Linux, or UNIX-like operating systems. For more information, see Overview.

#### Considerations for Windows

We recommend that you use a recent version of Windows. More recent versions of Windows have fewer vulnerabilities than earlier versions. IIS 7.5 provides more features and a more convenient console than IIS 6.

Read the following considerations and select a suitable hardware configuration and Windows version based on your needs:

- o Instance types that have only one vCPU and 1 GiB of memory do not support MySQL databases.
- Windows instances must have at least 2 GiB of memory to build websites and deploy web environments.
- To ensure service availability, we recommend that you select an instance type that has at least 2 GiB of memory when you use Windows 2012.
- You must select an instance type that has at least 2 GiB of memory if you want to use Windows
   Server 2016 or Windows Server 2019. If your selected instance type has less than 2 GiB of memory,
   Windows Server 2016 or Windows Server 2019 may not be displayed in the public image list on the
   buy page.
- Alibaba Cloud no longer provides technical support for Windows Server 2003 images. For more information, see Offline announcement of Windows Server 2003 images.
- Considerations for Linux and UNIX-like distributions

Alibaba Cloud Linux and UNIX-like public images contain the following distributions:

#### o Alibaba Cloud Linux

Alibaba Cloud Linux is an operating system that provides a safe, stable, and high-performance runtime environment for applications on ECS instances. Alibaba Cloud Linux 2 supports various cloud scenarios and instance types (excluding instances of the classic network type and non-I/O optimized instances). For more information, see Overview.

- o Red Hat series
  - Cent OS
  - Red Hat

The following table compares CentOS with Red Hat.

Operating system	Soft ware pack age form at	Packag e manag er	Fee	Feature	Relationship
CentOS	.rpm	yum	Free of charge	<ul> <li>Has stable but less frequent patch updates than Red Hat.</li> <li>Supports online and timely updates.</li> </ul>	<ul> <li>CentOS is a free version derived from the source code of Red Hat.</li> <li>They can use the same RPM package.</li> </ul>
Red Hat			Paid	Stable with enterprise- level technical support.	■ They can run the same commands.

- Debian series
  - Debian
  - Ubuntu

The following table compares Debian with Ubuntu.

Operating system	Softwar e package format	Package manager	Feature	Relationship
Debian		aptitude	Stable.	
Ubuntu	.deb	apt-get	<ul><li>User-friendly system configuration.</li><li>Timely software updates.</li><li>Easy to use and learn.</li></ul>	Ubuntu is built on the Debian architecture and infrastructure. Ubuntu is the enhanced version of Debian.

- SUSE series
  - SUSE Linux
  - OpenSUSE

The following table compares openSUSE with SUSE Linux.

Operating system	Feature	Relationship
OpenSUSE  SUSE Linux	<ul> <li>openSUSE is the community edition of SUSE Linux. SUSE Linux Enterprise is the enterprise edition of SUSE Linux.</li> <li>SUSE Linux Enterprise is more mature and stable, but its official distribution contains fewer software features than openSUSE.</li> <li>openSUSE provides advanced software versions, better extensibility (desktop and server installation available), and free updates (official technical support available).</li> <li>SUSE Linux Enterprise is more suitable for work and production environments, whereas openSUSE is more suitable for personal entertainment and professional purposes.</li> </ul>	<ul> <li>As of version 10.2, SUSE Linux is officially renamed to openSUSE.</li> <li>openSUSE uses the same kernel as SUSE Linux.</li> </ul>

#### CoreOS

CoreOS is an open source lightweight operating system based on the Linux kernel and designed to provide infrastructure for clustered deployments. CoreOS provides automation, ease of application deployment, security, reliability, and scalability. CoreOS provides the underlying functionalities required to deploy applications inside software containers in conjunction with a set of built-in tools for service discovery and configuration sharing.

#### o FreeBSD

FreeBSD is a UNIX-like operating system for a variety of platforms that focuses on features, speed, and stability. FreeBSD provides advanced networking, performance, security, and compatibility features that are still missing in other operating systems, even some of the best commercial ones. For more information, see FreeBSD Documentation.

#### Built-in software

Alibaba Cloud Market place images are typically pre-installed with a runtime environment or software applications. You can purchase appropriate images to create ECS instances based on your needs. For more information, see Alibaba Cloud Market place images.

#### What to do next

- Use an image to create an ECS instance. For more information, see Create an instance by using the wizard.
- Use an image to change the operating system of an ECS instance. For more information, see Change

the operating system.

# 3.Find an image

You can find an image based on its type, name, ID, or snapshot ID. After you find an image, you can use it to create an instance or perform other operations. This topic describes how to find an image.

#### Methods

You can use one of the following methods to find an image:

• Find an image by using the Elastic Compute Service (ECS) console

#### Examples:

- Example 1: Find a Windows public image in the China (Beijing) region.
- Example 2: Find a shared image whose name contains MySQL in the China (Hangzhou) region.
- Example 3: Find a custom image that corresponds to a snapshot ID of s-123\*\*\*\* in the China (Shenzhen) region.
- Example 4: Find a community image. For more information, see Search for a community image.
- Find an image by calling an API operation

## Find an image by using the Elastic Compute Service (ECS) console

You can find the image that you want to manage on the Images page in the ECS console.

- 1.
- 2.
- 3.
- 4. Select an image type.
- 5. (Optional)(Optional)Select an image family.
  - Note You can select an image family only for custom images. By default, all images are displayed.
- 6. Select a filter condition from the drop-down list.
  - ? Note
    - For custom images, public images, and shared images, you can select the image name, image ID, or snapshot ID as the search item.
    - For Alibaba Cloud Market place images, you can select the image ID, instance ID, or image product name as the search item.
    - For community images, you can select the image name, image ID, or owner ID as the search item.
- 7. Enter a value in the search box.
- 8. Click the \( \text{\text{Q}} \) icon to search for an image.

# Example 1

Example 1: Find a Windows public image in the China (Beijing) region.

Perform the following operations on the Images page.

- 1. In the left part of the top navigation bar, select China (Beijing) from the drop-down list.
- 2. Click the Public Images tab.
- 3. Select Image Name from the drop-down list.
- 4. Enter win in the search box.

#### Example 2

Example 2: Find a shared image whose name contains mysql in the China (Hangzhou) region.

Perform the following operations on the Images page.

- 1. In the left part of the top navigation bar, select China (Hangzhou) from the drop-down list.
- 2. Click the **Shared Images** tab.
- 3. Select Image Name from the drop-down list.
- 4. Enter MySQL in the search box.
- 5. Click the \( \text{\tint{\text{\tin}\text{\texi\text{\texi}\text{\text{\text{\text{\text{\text{\text{\text{\text{\text{\texi}\text{\text{\text{\text{\text{\text{\text{\text{\text{\text{\text{\texi}\tex

#### Example 3

Example 3: Find a custom image that corresponds to a snapshot ID of s-123\*\*\*\* in the China (Shenzhen) region.

Perform the following operations on the Images page.

- 1. In the left part of the top navigation bar, select China (Shenzhen) from the drop-down list.
- 2. Click the Custom Images tab.
- 3. Select Snapshot ID from the drop-down list.
- 4. Enter *s-123\*\*\*\** in the search box.

# Find an image by calling an API operation

You can also find an image by using OpenAPI Explorer or Alibaba Cloud CLI to call the DescribeImages operation. This section describes how to use OpenAPI Explorer to call the DescribeImages operation to find an image.

- 1. Go to the OpenAPI Explorer console.
- 2. Select a region from the **RegionId** drop-down list.
- 3. (Optional)Configure parameters such as ImageName and ImageId.

- Note Image IDs comply with the following naming rules:
  - IDs of public images are named after the operating system version numbers, architectures, languages, and release dates of the images. For example, the ID of the Windows Server 2008 R2 Enterprise 64-bit (English) public image is win2008r2 64 ent sp1 en-us 40G alibase 20190318.vhd.
  - IDs of custom images, community images, and Alibaba Cloud Marketplace images start with m.
  - IDs of shared images are the same as those of the custom images that were shared.
- 4. Click Submit Request.
- 5. Click the Debugging Result tab.

If the image is found, its information such as the image ID, description, and operating system is displayed on the **Debugging Result** tab. For more information, see **DescribeImages**.

#### What's next

After you find an image, you can perform the following operations:

- Create an ECS instance. For more information, see Create an instance by using the wizard.
- Share the image. For more information, see Share or unshare a custom image.
- Copy the image. For more information, see Copy a custom image.
- Export the image. For more information, see Export a custom image.
- Delete the image. For more information, see Delete a custom image.
- Modify the description of the image. For more information, see Modify a custom image.

#### Related information

Describelmages

# 4.Public image 4.1. Overview

Public images provided by Alibaba Cloud are fully licensed to provide a secure and stable operating environment for applications on Elastic Compute Service (ECS) instances. This topic describes the types of public images such as Alibaba Cloud Linux images, third-party images, and open source images.

## Types of public images

The following table describes two types of public images provided by Alibaba Cloud. You can use public images to create ECS instances free of charge, except Windows Server, Red Hat Enterprise Linux, and SUSE Linux Enterprise Server public images. For more information, see Image prices.

Туре	Description	Technical support
Alibaba Cloud Linux images	Alibaba Cloud Linux images are custom, proprietary operating systems provided by Alibaba Cloud to create ECS instances. Alibaba Cloud Linux images are fully tested to ensure their security, stability, and normal startup and operation.	Alibaba Cloud provides technical support to help resolve the issues that you may encounter when you use Alibaba Cloud Linux images.
Third-party images and open source images	Third-party and open source images are fully tested and released by Alibaba Cloud to ensure their security, stability, and normal startup and operation. The following third-party public images are available:  • Windows: Windows Server  • Linux: Anolis OS, Ubuntu, CentOS, Red Hat Enterprise Linux, Debian, openSUSE, SUSE Linux Enterprise Server, FreeBSD, CoreOS, Fedora CoreOS, Fedora, Rocky Linux, and AlmaLinux	We recommend that you contact corresponding operating system vendors or open source communities for technical support. Alibaba Cloud also provides information about image- and system-related issues.

# Alibaba Cloud Linux images

Alibaba Cloud Linux images are Linux public images developed by Alibaba Cloud. The following table describes the versions of Alibaba Cloud Linux images.

Operating system	Operating system version	Description
------------------	--------------------------	-------------

Operating system	Operating system version	Description
Alibaba Cloud Linux 3	<ul> <li>Alibaba Cloud Linux 3.2104 64-bit</li> <li>Alibaba Cloud Linux 3.2104 64-bit (Quick Start)</li> <li>Alibaba Cloud Linux 3.2104 64-bit for ARM</li> <li>Alibaba Cloud Linux 3.2104 64-bit (UEFI)</li> </ul>	The operating system supports various Alibaba Cloud instance types including ECS Bare Metal Instance types. By default, Alibaba Cloud Linux 3 is equipped with Alibaba Cloud CLI and other software packages.  If you want to replace other Linux distributions with Alibaba Cloud Linux 3, you can select Public Image and select Alibaba Cloud Linux 3 when you create an ECS instance or when you replace the system disk of an ECS instance.  For more information, see Alibaba Cloud Linux overview and Alibaba Cloud Linux 3镜像发布记录.
Alibaba Cloud Linux 2	<ul> <li>Alibaba Cloud Linux 2.1903 LTS 64-bit</li> <li>Alibaba Cloud Linux 2.1903 64-bit (Quick Start)</li> <li>Alibaba Cloud Linux 2.1903 64-bit (UEFI)</li> <li>Alibaba Cloud Linux 2.1903 64-bit (Trusted)</li> </ul>	The operating system supports various Alibaba Cloud instance types including ECS Bare Metal Instance types. By default, Alibaba Cloud Linux 2 is equipped with Alibaba Cloud CLI and other software packages.  If you want to replace other Linux distributions with Alibaba Cloud Linux 2, you can select Public Image and select Alibaba Cloud Linux 2 when you create an ECS instance or when you replace the system disk of an ECS instance.  For more information, see Overview and Alibaba Cloud Linux 2镜像发布记录.

# Third-party images and open source images

Alibaba Cloud releases and updates public images of third-party and open source image vendors on a regular basis. For more information, see . Alternatively, you can select a region in the ECS console and view all the public images available in that region on the Public Images tab. For more information, see Find an image.

The following tables describe the versions of third-party and open source public images for Windows and Linux provided by Alibaba Cloud.

#### • Windows images

Operating system	Operating system version
	<ul> <li>Windows Server 2019 Datacenter with Containers 64-bit (Simplified Chinese)</li> </ul>
Windows Server 2019	Windows Server 2019 Datacenter with Containers 64-bit (English)
	<ul> <li>Windows Server 2019 Datacenter 64-bit (Simplified Chinese)</li> </ul>
	Windows Server 2019 Datacenter 64-bit (English)

Operating system	Operating system version
Windows Server 2016	<ul> <li>Windows Server 2016 Datacenter 64-bit (Simplified Chinese)</li> <li>Windows Server 2016 Datacenter 64-bit (English)</li> </ul>
Windows Server 2012	<ul> <li>Windows Server 2012 R2 Datacenter 64-bit (Simplified Chinese)</li> <li>Windows Server 2012 R2 Datacenter 64-bit (English)</li> </ul>
Windows Server Version **** (Semi-Annual Channel)	<ul> <li>Windows Server Version **** Datacenter 64-bit (Simplified Chinese)</li> <li>Windows Server Version **** Datacenter 64-bit (English)</li> <li>The asterisks (****) indicate the latest version number of the Semi-Annual Channel release.</li> </ul>



## • Linux images

Operating system	Operating system version
	Anolis OS 8.4 ANCK 64-bit
	Anolis OS 8.4 RHCK 64-bit
	Anolis OS 8.4 ANCK 64-bit for ARM
	Anolis OS 8.4 RHCK 64-bit for ARM
	Anolis OS 8.2 ANCK 64-bit
	Anolis OS 8.2 RHCK 64-bit
	Anolis OS 8.2 ANCK 64-bit for ARM
	o Anolis OS 8.2 RHCK 64-bit for ARM
Anolis OS	o Anolis OS 7.9 ANCK 64-bit
	Anolis OS 7.9 RHCK 64-bit
	Anolis OS 7.9 ANCK 64-bit for ARM
	Anolis OS 7.9 RHCK 64-bit for ARM
	Anolis OS 7.7 ANCK 64-bit
	Anolis OS 7.7 RHCK 64-bit
	Anolis OS 7.7 ANCK 64-bit for ARM
	Anolis OS 7.7 RHCK 64-bit for ARM
	Cent OS Stream 8 64-bit
	○ CentOS 8.5 64-bit
	o Cent OS 8.4 64-bit
	• CentOS 8.4 64-bit for ARM
	• CentOS 8.4 64-bit (UEFI)
	<ul> <li>Cent OS 8.3 64-bit</li> </ul>
	<ul> <li>CentOS 8.3 64-bit for ARM</li> </ul>

Operating system	<ul> <li>Cent OS 8.3 64-bit (UEFI)</li> <li>Operating system version</li> <li>Cent OS 8.2 64-bit</li> </ul>
	<ul> <li>CentOS 8.2 64-bit (AMD-compatible)</li> </ul>
	o CentOS 8.1 64-bit
	Cent OS 8.0 64-bit
	o CentOS 7.9 64-bit
	<ul> <li>CentOS 7.9 64-bit (AMD-compatible)</li> </ul>
	○ CentOS 7.9 64-bit (UEFI)
	o Cent OS 7.8 64-bit
	CentOS 7.8 64-bit (Trusted)
	<ul> <li>CentOS 7.8 64-bit (AMD-compatible)</li> </ul>
	o Cent OS 7.7 64-bit
	○ CentOS 7.6 64-bit
	o CentOS 7.5 64-bit
	o Cent OS 7.4 64-bit
CentOS	○ CentOS 7.3 64-bit
	○ Cent OS 7.2 64-bit
	○ Cent OS 6.10 64-bit
	○ Cent OS 6.9 64-bit
	Cent OS 6.8 32-bit
	? Note
	• The CentOS project has announced the end of life (EOL) for CentOS Linux. CentOS Linux public images provided by Alibaba Cloud are derived from the repositories provided by the CentOS project. When CentOS Linux reaches its EOL, Alibaba Cloud stops providing support for CentOS Linux. For information about how to deal with the situation, see Options for dealing with CentOS Linux EOL.
	<ul> <li>CentOS 6 and CentOS 8 have reached its end of life (EOL) and is no longer maintained by the Linux community. If you want to continue to use specific installation packages of CentOS 6 and CentOS 8 during the transition period, you must change CentOS 6 and CentOS 8 repository addresses. For more information, see Change the CentOS 6 source address and Change CentOS 8 repository addresses.</li> </ul>
	<ul> <li>To use a 32-bit operating system, select an instance type that has 4 GiB or less of memory. For more information, see Select an image.</li> </ul>

Operating system	Operating system version
CoreOS	<ul> <li>CoreOS 2345.3.0 64-bit</li> <li>CoreOS 2303.4.0 64-bit</li> <li>CoreOS 2303.3.0 64-bit</li> <li>CoreOS 2247.6.0 64-bit</li> <li>CoreOS 2023.4.0 64-bit</li> <li>CoreOS 1745.7.0 64-bit</li> </ul> Note
Debian	<ul> <li>Debian 11.2 64-bit</li> <li>Debian 11.1 64-bit</li> <li>Debian 10.11 64-bit</li> <li>Debian 10.10 64-bit</li> <li>Debian 10.10 64-bit (UEFI)</li> <li>Debian 10.9 64-bit (UEFI)</li> <li>Debian 10.9 64-bit for ARM</li> <li>Debian 10.7 64-bit</li> <li>Debian 10.7 64-bit</li> <li>Debian 10.5 64-bit</li> <li>Debian 10.5 64-bit</li> <li>Debian 10.4 64-bit</li> <li>Debian 10.3 64-bit</li> <li>Debian 10.3 64-bit</li> <li>Debian 9.13 64-bit</li> <li>Debian 9.13 64-bit</li> <li>Debian 9.14 64-bit</li> <li>Debian 9.15 64-bit</li> <li>Debian 9.16 64-bit</li> <li>Debian 9.16 64-bit</li> <li>Debian 9.11 64-bit</li> <li>Debian 9.11 64-bit</li> <li>Debian 9.8 64-bit</li> <li>Debian 8.9 64-bit</li> <li>Debian 8.9 64-bit</li> </ul>
FreeBSD	<ul> <li>FreeBSD 12.1 64-bit</li> <li>FreeBSD 11.4 64-bit</li> <li>FreeBSD 11.3 64-bit</li> <li>FreeBSD 11.2 64-bit</li> </ul>

Operating system	Operating system version
OpenSUSE	<ul> <li>OpenSUSE 15.3 64-bit</li> <li>OpenSUSE 15.2 64-bit</li> <li>OpenSUSE 15.1 64-bit</li> <li>OpenSUSE 42.3 64-bit</li> </ul>
Red Hat	<ul> <li>Red Hat Enterprise Linux 8.5 64-bit</li> <li>Red Hat Enterprise Linux 8.4 64-bit</li> <li>Red Hat Enterprise Linux 8.2 64-bit</li> <li>Red Hat Enterprise Linux 8.1 64-bit</li> <li>Red Hat Enterprise Linux 7.9 64-bit</li> <li>Red Hat Enterprise Linux 7.9 64-bit</li> <li>Red Hat Enterprise Linux 7.8 64-bit</li> <li>Red Hat Enterprise Linux 7.8 64-bit</li> <li>Red Hat Enterprise Linux 7.6 64-bit</li> <li>Red Hat Enterprise Linux 7.5 64-bit</li> <li>Red Hat Enterprise Linux 7.4 64-bit</li> <li>Red Hat Enterprise Linux 7.4 64-bit</li> <li>Red Hat Enterprise Linux 6.10 64-bit</li> <li>Red Hat Enterprise Linux 6.10 64-bit</li> <li>Red Hat Enterprise Linux 6.9 64-bit</li> <li>Red Hat Enterprise Linux 6.9 64-bit</li> <li>Red Hat Enterprise Linux 6.9 64-bit</li> </ul>

Operating system	Operating system version
SUSE Linux	<ul> <li>SUSE Linux Enterprise Server 15 SP3 64-bit</li> <li>SUSE Linux Enterprise Server 15 SP2 64-bit</li> <li>SUSE Linux Enterprise Server 15 SP1 64-bit</li> <li>SUSE Linux Enterprise Server 12 SP5 64-bit</li> <li>SUSE Linux Enterprise Server 12 SP4 64-bit</li> <li>SUSE Linux Enterprise Server 12 SP3 64-bit</li> <li>SUSE Linux Enterprise Server 12 SP2 64-bit</li> <li>SUSE Linux Enterprise Server 12 SP2 64-bit</li> <li>SUSE Linux Enterprise Server 11 SP4 64-bit</li> </ul>
	Note Starting from October 30, 2021, SUSE Linux Enterprise Server public images are charged. For more information, see SUSE Linux Enterprise Server will be commercialized on October 30, 2021. The SUSE Linux Enterprise Server images used by existing instances can be upgraded to commercial versions to continue to receive image updates and technical support. For more information, see Upgrade the SUSE Linux Enterprise Server public images of existing instances to the commercial edition.
Ubuntu	<ul> <li>Ubuntu 20.04 64-bit</li> <li>Ubuntu 20.04 64-bit (AMD-compatible)</li> <li>Ubuntu 20.04 64-bit for ARM</li> <li>Ubuntu 20.04 64-bit (UEFI)</li> <li>Ubuntu 18.04 64-bit</li> <li>Ubuntu: 18.04 64-bit (AMD-compatible)</li> <li>Ubuntu 18.04 64-bit for ARM</li> <li>Ubuntu 18.04 64-bit (UEFI)</li> <li>Ubuntu 16.04 64-bit</li> <li>Ubuntu 16.04 32-bit</li> <li>Ubuntu 14.04 64-bit</li> <li>Ubuntu 14.04 32-bit</li> </ul>
	Note To use a 32-bit operating system, select an instance type that has 4 GiB or less of memory. For more information, see Select an image.

Operating system	Operating system version
Fedora CoreOS	<ul> <li>Fedora CoreOS 34.20210529.3.0_3</li> <li>Fedora CoreOS 33.20210217.3.0_3</li> <li>When you use these images, take note of the following items:</li> <li>When you create an instance or replace the system disk of an instance, you can use only a key pair as logon credentials. You can use only the key pair that you initially configured for your instance to log on to the instance. You cannot change or unbind the key pair.</li> <li>If you have created an instance or replaced the system disk of an instance, you cannot reset the password of the instance.</li> </ul>
Fedora	<ul><li>Fedora 34 64-bit</li><li>Fedora 33 64-bit</li></ul>
Rocky Linux	Rocky Linux 8.5
AlmaLinux	AlmaLinux 8.5

# 4.2.

•

•

<ul> <li>Anolis OS</li> <li>AlmaLinux</li> <li>Cent OS</li> <li>CoreOS</li> <li>Debian</li> <li>Fedora CoreOS</li> <li>Fedora</li> <li>FreeBSD</li> <li>OpenSUSE</li> <li>Red Hat Enterprise Linux</li> <li>Rocky Linux</li> <li>SUSE Linux Enterprise Server</li> <li>Ubuntu</li> </ul>	<ul> <li>Windows Server 2022</li> <li>Windows Server 2019</li> <li>Windows Server 2016</li> <li>Windows Server 2012</li> <li>Windows Server 2008</li> <li>Windows Server Version 2004</li> <li>Windows Server Version 1909</li> <li>Windows Server Version 1903</li> <li>Windows Server Version 1809</li> </ul>

# Anolis OS

Anolis OS 8.4 RHCK	anolisos_8_4_x64_20G_rhck_alibase_ 20220518.vhd	2022-05-18	•
Anolis OS 8.4 ANCK	anolisos_8_4_x64_20G_anck_alibase _20220518.vhd	2022-05-18	•
Anolis OS 7.9 RHCK	anolisos_7_9_x64_20G_rhck_alibase_ 20220518.vhd	2022-05-18	•
Anolis OS 8.4 ANCK	anolisos_8_4_x64_20G_anck_alibase _20220413.vhd	2022-04-13	•
Anolis OS 8.4 RHCK	anolisos_8_4_x64_20G_rhck_alibase_ 20220413.vhd	2022-04-13	•
Anolis OS 8.2 ANCK	anolisos_8_2_x64_20G_anck_alibase _20220413.vhd	2022-04-13	•
Anolis OS 8.2 RHCK	anolisos_8_2_x64_20G_rhck_alibase_ 20220415.vhd	2022-04-15	•
Anolis OS 7.9 RHCK	anolisos_7_9_x64_20G_rhck_alibase_ 20220414.vhd	2022-04-14	•

Anolis OS 7.9 RHCK	anolisos_7_9_x64_20G_rhck_alibase_ 20220316.vhd	2022-03-16	•
Anolis OS 7.9 ANCK	anolisos_7_9_x64_20G_anck_alibase _20220316.vhd	2022-03-16	•
Anolis OS 8.4 RHCK	anolisos_8_4_x64_20G_rhck_alibase_ 20220119.vhd	2022-01-19	<ul><li>Linux time and time zones</li></ul>
Anolis OS 7.9 ANCK	anolisos_7_9_x64_20G_anck_alibase _20220111.vhd	2022-01-11	•
Anolis OS 7.9 RHCK	anolisos_7_9_x64_20G_rhck_alibase_ 20220110.vhd	2022-01-10	•
Anolis OS 7.9 RHCK	anolisos_7_9_x64_20G_rhck_alibase_ 20211210.vhd	2021-12-10	•
Anolis OS 7.7 ANCK	anolisos_7_7_x64_20G_anck_alibase _20211118.vhd	2021-11-18	•
Anolis OS 7.7 RHCK	anolisos_7_7_x64_20G_rhck_alibase_ 20211118.vhd	2021-11-18	•

Anolis OS 8.4 ANCK	anolisos_8_4_x64_20G_anck_alibase _20211008.vhd	2021-10-08	
Anolis OS 8.4 RHCK	anolisos_8_4_x64_20G_rhck_alibase_ 20211008.vhd	2021-10-08	

Anolis OS 8.2 ANCK	anolisos_8_2_x64_20G_anck_alibase _20211008.vhd	2021-10-08	
Anolis OS 8.2 RHCK	anolisos_8_2_x64_20G_rhck_alibase_ 20211008.vhd	2021-10-08	
Anolis OS 8.2 RHCK	anolisos_8_2_x64_20G_rhck_alibase_ 20210529.vhd	2021-05-29	•

# AlmaLinux

AlmaLinux 8.6	almalinux_8_6_x64_20G_alibase_202 20524.vhd	2022-05-24	•

AlmaLinux 8.5	almalinux_8_5_x64_20G_alibase_202 20426.vhd	2022-04-26	•
AlmaLinux 8.5	almalinux_8_5_x64_20G_alibase_202 20322.vhd	2022-03-22	• 0 0
AlmaLinux 8.5	almalinux_8_5_x64_20G_alibase_202 20208.vhd	2022-02-08	• 0 0
AlmaLinux 8.5	almalinux_8_5_x64_20G_alibase_202 20110.vhd	2022-01-10	• • • •
AlmaLinux 8.5	almalinux_8_5_x64_20G_alibase_202 11126.vhd	2021-11-26	•

# CentOS

CentOS Stream 8	centos_stream_8_x64_20G_alibase_ 20220524.vhd	2022-05-24	•
Cent OS 7.9	centos_7_9_x64_20G_alibase_20220 524.vhd	2022-05-24	•
Cent OS 8.5	centos_8_5_x64_20G_alibase_20220 428.vhd	2022-04-28	•
CentOS Stream 8	centos_stream_8_x64_20G_alibase_ 20220426.vhd	2022-04-26	•
Cent OS 7.9	centos_7_9_x64_20G_alibase_20220 426.vhd	2022-04-26	•
Cent OS 8.5	centos_8_5_x64_20G_alibase_20220 325.vhd	2022-03-25	•

CentOS Stream 8	centos_stream_8_x64_20G_alibase_ 20220322.vhd	2022-03-22	•
Cent OS 7.9	centos_7_9_x64_20G_alibase_20220 322.vhd	2022-03-22	•
Cent OS 8.5	centos_8_5_x64_20G_alibase_20220 303.vhd	2022-03-03	o     Linux time and time zones
CentOS Stream 8	centos_stream_8_x64_20G_alibase_ 20220208.vhd	2022-02-08	• • • • • • • • • • • • • • • • • • •
Cent OS 7.9	centos_7_9_x64_20G_alibase_20220 208.vhd	2022-02-08	• • • • • • • • • • • • • • • • • • •
Cent OS 8.5	centos_8_5_x64_20G_alibase_20211 228.vhd	2021-12-28	• • • •
CentOS Stream 8	centos_stream_8_x64_20G_alibase_ 20211227.vhd	2021-12-27	•
Cent OS 7.9	centos_7_9_x64_20G_alibase_20211 227.vhd	2021-12-27	• • • • • • • • • • • • • • • • • • •
CentOS Stream 8	centos_stream_8_x64_20G_alibase_ 20211130.vhd	2021-11-30	• • • • • • • • • • • • • • • • • • •
Cent OS 7.8	centos_7_8_x64_20G_alibase_20211 130.vhd	2021-11-30	•

CentOS 7.7	centos_7_7_x64_20G_alibase_20211 130.vhd	2021-11-30	•
CentOS 7.6	centos_7_6_x64_20G_alibase_20211 130.vhd	2021-11-30	•
CentOS 7.5	centos_7_5_x64_20G_alibase_20211 130.vhd	2021-11-30	•
CentOS 8.5	centos_8_5_x64_20G_alibase_20211 1129.vhd	2021-11-29	• • • • • • • • • • • • • • • • • • • •
CentOS 7.9	centos_7_9_x64_20G_alibase_20211 123.vhd	2021-11-23	• • • • •
CentOS Stream 8	centos_stream_8_x64_20G_alibase_ 20211115.vhd	2021-11-15	•
CentOS 7.5	centos_7_5_x64_20G_alibase_20211 101.vhd	2021-11-01	• • • • • • • • • • • • • • • • • • •
Cent OS 7.6	centos_7_6_x64_20G_alibase_20211 030.vhd	2021-10-30	• • • • • • • • • • • • • • • • • • • •

Cent OS 7.8	centos_7_8_x64_20G_alibase_20211 028.vhd	2021-10-28	• • • • • • • • • • • • • • • • • • •
Cent OS 7.7	centos_7_7_x64_20G_alibase_20211 028.vhd	2021-10-28	• • • • • • • • • • • • • • • • • • •
CentOS 8.4	centos_8_4_x64_20G_alibase_20211 027.vhd	2021-10-27	•
CentOS 7.9	centos_7_9_x64_20G_alibase_20211 027.vhd	2021-10-27	•
Cent OS 8.4	centos_8_4_x64_20G_alibase_20210 927.vhd	2021-09-27	•
CentOS 7.9	centos_7_9_x64_20G_alibase_20210 927.vhd	2021-09-27	• • • • • • • • • • • • • • • • • • •
Cent OS 8.4	centos_8_4_x64_20G_alibase_20210 824.vhd	2021-08-24	• • • •
Cent OS 7.9	centos_7_9_x64_20G_alibase_20210 824.vhd	2021-08-24	• • • •
Cent OS 8.3	centos_8_3_x64_20G_alibase_20210 723.vhd	2021-07-23	• • • • • • • • • • • • • • • • • • • •

CentOS 8.2	centos_8_2_x64_20G_alibase_20210 712.vhd	2021-07-12	• • • • • • • • • • • • • • • • • • • •
CentOS 8.1	centos_8_1_x64_20G_alibase_20210 712.vhd	2021-07-12	•
CentOS 8.0	centos_8_0_x64_20G_alibase_20210 712.vhd	2021-07-12	• • • • • • • • • • • • • • • • • • •
Cent OS 8.4	centos_8_4_x64_20G_alibase_20210 623.vhd	2021-06-23	•
CentOS 7.9	centos_7_9_x64_20G_alibase_20210 623.vhd	2021-06-23	•
Cent OS 8.3	centos_8_3_x64_20G_alibase_20210 521.vhd	2021-05-21	•
CentOS 7.9	centos_7_9_x64_20G_alibase_20210 521.vhd	2021-05-21	•
Cent OS 8.3	centos_8_3_x64_20G_alibase_20210 420.vhd	2021-04-20	•
CentOS 7.9	centos_7_9_x64_20G_alibase_20210 420.vhd	2021-04-20	•
Cent OS 8.3	centos_8_3_x64_20G_alibase_20210 318.vhd	2021-03-18	• • • • • • • • • • • • • • • • • • •

CentOS 7.9	centos_7_9_x64_20G_alibase_20210 318.vhd	2021-03-18	• • • • • • •
CentOS 7.9	centos_7_9_x64_20G_alibase_20210 128.vhd	2021-01-28	• • • • • • • • • • • • • • • • • • •
Cent OS 7.9	centos_7_9_x64_20G_alibase_20201 228.vhd	2020-12-28	•
Cent OS 8.2	centos_8_2_x64_20G_alibase_20201 120.vhd	2020-11-20	• • • • • • • • • • • • • • • • • • •
CentOS 7.9	centos_7_9_x64_20G_alibase_20201 120.vhd	2020-11-20	•
CentOS 6.10	centos_6_10_x64_20G_alibase_2020 1120.vhd	2020-11-20	•
Cent OS 8.2	centos_8_2_x64_20G_alibase_20200 914.vhd	2020-09-14	•
CentOS 7.8	centos_7_8_x64_20G_alibase_20200 914.vhd	2020-09-14	•
CentOS 6.10	centos_6_10_x64_20G_alibase_2020 0914.vhd	2020-09-14	•
CentOS 8.2	centos_8_2_x64_20G_alibase_20200 824.vhd	2020-08-24	•
Cent OS 7.8	centos_7_8_x64_20G_alibase_20200 817.vhd	2020-08-17	•
Cent OS 6.10	centos_6_10_x64_20G_alibase_2020 0817.vhd	2020-08-17	•

Cent OS 8.2	centos_8_2_x64_20G_alibase_20200 717.vhd	2020-07-17	•
Cent OS 7.8	centos_7_8_x64_20G_alibase_20200 717.vhd	2020-07-17	•
CentOS 6.10	centos_6_10_x64_20G_alibase_2020 0717.vhd	2020-07-17	•
Cent OS 7.8	centos_7_8_x64_20G_alibase_20200 622.vhd	2020-06-22	•
Cent OS 8.2	centos_8_2_x64_20G_alibase_20200 616.vhd	2020-06-16	•
Cent OS 6.10	centos_6_10_x64_20G_alibase_2020 0616.vhd	2020-06-16	•
Cent OS 8.1	centos_8_1_x64_20G_alibase_20200 519.vhd	2020-05-19	•
Cent OS 7.8	centos_7_8_x64_20G_alibase_20200 519.vhd	2020-05-19	•
Cent OS 6.10	centos_6_10_x64_20G_alibase_2020 0519.vhd	2020-05-19	•
Cent OS 8.1	centos_8_1_x64_20G_alibase_20200 426.vhd	2020-04-26	•
Cent OS 7.7	centos_7_7_x64_20G_alibase_20200 426.vhd	2020-04-26	•
Cent OS 6.10	centos_6_10_x64_20G_alibase_2020 0426.vhd	2020-04-26	•
Cent OS 8.1	centos_8_1_x64_20G_alibase_20200 329.vhd	2020-03-29	•

CentOS 7.7	centos_7_7_x64_20G_alibase_20200 329.vhd	2020-03-29	• • o  Note Install cloudinit
CentOS 6.10	centos_6_10_x64_20G_alibase_2020 0319.vhd	2020-03-19	•
Cent OS 8.0	centos_8_0_x64_20G_alibase_20200 218.vhd	2020-02-18	● ○ ○ ○ ○ ○ ○ ○ ○ □知问题
Cent OS 7.7	centos_7_7_x64_20G_alibase_20200 220.vhd	2020-02-20	•
Cent OS 6.10	centos_6_10_x64_20G_alibase_2020 0214.vhd	2020-02-14	•
Cent OS 6.10	centos_6_10_x64_20G_alibase_2020 0103.vhd	2020-01-03	•
CentOS 7.7	centos_7_7_x64_20G_alibase_20191 225.vhd	2019-12-25	•
Cent OS 8.0	centos_8_0_x64_20G_alibase_20191 225.vhd	2019-12-25	• • o  • Note Install cloudinit
Cent OS 6.10	centos_6_10_x64_20G_alibase_2019 1223.vhd	2019-12-25	•
Cent OS 7.7	centos_7_7_64_20G_alibase_201910 08.vhd	2019-10-8	•

Cent OS 7.6	centos_7_06_64_20G_alibase_20190 711.vhd	2019-7-11	•
Cent OS 6.10	centos_6_10_64_20G_alibase_20190 709.vhd	2019-7-9	•
Cent OS 6.10	centos_6_10_64_20G_alibase_20190 621.vhd	2019-6-21	•
Cent OS 7.6	centos_7_06_64_20G_alibase_20190 619.vhd	2019-6-19	• • • • • • • • • • • • • • • • • • •
CentOS 7.6	centos_7_06_64_20G_alibase_20190 218.vhd	2019-2-18	•
CentOS 7.6	centos_7_05_64_20G_alibase_20181 212.vhd	2018-12-12	•
Cent OS 7.5	centos_7_05_64_20G_alibase_20181 210.vhd	2018-12-10	• • • • • • • • • • • • • • • • • • •

# CoreOS

#### ? Note

CoreOS 2345_3.0	coreos_2345_3.0_x64_30G_alibase_2 0200519.vhd	2020-05-19	•
CoreOS 2345_3.0	coreos_2345_3.0_x64_30G_alibase_2 0200423.vhd	2020-04-23	•
CoreOS 2345_3.0	coreos_2345_3.0_x64_30G_alibase_2 0200325.vhd	2020-03-25	•

CoreOS 2303_4.0	coreos_2303_4.0_x64_30G_alibase_2 0200217.vhd	2020-02-17	•
CoreOS 2303_3.0	coreos_2303_3_x64_30G_alibase_20 191223.vhd	2019-12-23	•

#### Debian

Debian 11.3	debian_11_3_x64_20G_alibase_2022 0531.vhd	2022-05-31	• • • • • • • Linux time and time zones • • • • • • • • • • • • • • • • • • •
Debian 10.12	debian_10_12_x64_20G_alibase_202 20524.vhd	2022-05-24	• o • Linux time and time zones •
Debian 11.3	debian_11_3_x64_20G_alibase_2022 0510.vhd	2022-05-10	•
Debian 10.12	debian_10_12_x64_20G_alibase_202 20428.vhd	2022-04-28	• • • • • • • • • • • • • • • • • • •
Debian 9.13	debian_9_13_x64_20G_alibase_2022 0428.vhd	2022-04-28	• • • • • • • • • • • • • • • • • • •
Debian 10.12	debian_10_12_x64_20G_alibase_202 20328.vhd	2022-03-28	•
Debian 9.13	debian_9_13_x64_20G_alibase_2022 0322.vhd	2022-03-22	• • o o

Debian 11.2	debian_11_2_x64_20G_alibase_2022 0308.vhd	2022-03-08	• • • • • • • • • • • • • • • • • • •
Debian 11.2	debian_11_2_x64_20G_alibase_2022 0208.vhd	2022-02-08	•
Debian 10.11	debian_10_11_x64_20G_alibase_202 20208.vhd	2022-02-08	•
Debian 9.13	debian_9_13_x64_20G_alibase_2022 0208.vhd	2022-02-08	• • • • • • • • • • • • • • • • • • •
Debian 11.2	debian_11_2_x64_20G_alibase_2021 1227.vhd	2021-12-27	• • • • • • • • • • • • • • • • • • •
Debian 10.11	debian_10_11_x64_20G_alibase_202 11227.vhd	2021-12-27	• • • • • • • • • • • • • • • • • • •
Debian 9.13	debian_9_13_x64_20G_alibase_2021 1227.vhd	2021-12-27	• • • • • • • • • • • • • • • • • • •
Debian 11.1	debian_11_1_x64_20G_alibase_2021 1130.vhd	2021-11-30	• • • • • • • • • • • • • • • • • • •
Debian 11.0	debian_11_0_x64_20G_alibase_2021 1130.vhd	2021-11-30	• • • • • • • • • • • • • • • • • • •

Debian 10.11	debian_10_11_x64_20G_alibase_202 11123.vhd	2021-11-23	• • • • • • • • • • • • • • • • • • •
Debian 9.13	debian_9_13_x64_20G_alibase_2021 1123.vhd	2021-11-23	• • • • •
Debian 11.1	debian_11_1_x64_20G_alibase_2021 1027.vhd	2021-10-27	•
Debian 10.11	debian_10_11_x64_20G_alibase_202 11027.vhd	2021-10-27	•
Debian 9.13	debian_9_13_x64_20G_alibase_2021 1027.vhd	2021-10-27	•
Debian 9.8	debian_9_8_x64_20G_alibase_20210 928.vhd	2021-09-28	•
Debian 11.0	debian_11_0_x64_20G_alibase_2021 0927.vhd	2021-09-27	•
Debian 10.10	debian_10_10_x64_20G_alibase_202 10927.vhd	2021-09-27	•
Debian 9.13	debian_9_13_x64_20G_alibase_2021 0927.vhd	2021-09-27	•
Debian 9.6	debian_9_6_x64_20G_alibase_20210 927.vhd	2021-09-27	• • • • • • • • • • • • • • • • • • • •
Debian 10.10	debian_10_10_x64_20G_alibase_202 10824.vhd	2021-08-24	• 0 0

Debian 9.13	debian_9_13_x64_20G_alibase_2021 0824.vhd	2021-08-24	•
Debian 10.10	debian_10_10_x64_20G_alibase_202 10729.vhd	2021-07-29	• • • •
Debian 9.13	debian_9_13_x64_20G_alibase_2021 0729.vhd	2021-07-29	• • • • •
Debian 10.10	debian_10_10_x64_20G_alibase_202 10623.vhd	2021-06-23	•
Debian 9.13	debian_9_13_x64_20G_alibase_2021 0623.vhd	2021-06-23	•
Debian 10.9	debian_10_9_x64_20G_alibase_2021 0521.vhd	2021-05-21	•
Debian 9.13	debian_9_13_x64_20G_alibase_2021 0521.vhd	2021-05-21	•
Debian 10.9	debian_10_9_x86_20G_alibase_2021 0420.vhd	2021-04-20	•
Debian 9.13	debian_9_13_x64_20G_alibase_2021 0420.vhd	2021-04-20	•
Debian 10.9	debian_10_9_x86_20G_alibase_2021 0318.vhd	2021-03-18	•
Debian 9.13	debian_9_13_x64_20G_alibase_2021 0318.vhd	2021-03-18	•
Debian 10.7	debian_10_7_x64_20G_alibase_2021 0128.vhd	2021-01-28	• • • • • • • • • • • • • • • • • • •

Debian 9.13	debian_9_13_x64_20G_alibase_2021 0128.vhd	2021-01-28	• • • •
Debian 10.7	debian_10_7_x64_20G_alibase_2020 1228.vhd	2020-12-28	•
Debian 9.13	debian_9_13_x64_20G_alibase_2020 1228.vhd	2020-12-28	•
Debian 10.6	debian_10_6_x64_20G_alibase_2020 1120.vhd	2020-11-20	• • • •
Debian 9.13	debian_9_13_x64_20G_alibase_2020 1120.vhd	2020-11-20	• • • • • • • • • • • • • • • • • • •
Debian 10.5	debian_10_5_x64_20G_alibase_2020 0914.vhd	2020-09-14	•
Debian 9.13	debian_9_13_x64_20G_alibase_2020 0914.vhd	2020-09-14	•
Debian 10.5	debian_10_5_x64_20G_alibase_2020 0819.vhd	2020-08-19	• • • • • • • • • • • • • • • • • • •
Debian 9.13	debian_9_13_x64_20G_alibase_2020 0819.vhd	2020-08-19	• • • • • • • • • • • • • • • • • • •
Debian 9.13	debian_9_13_x64_20G_alibase_2020 0730.vhd	2020-07-30	• • • • • • • • • • • • • • • • • • •
Debian 10.4	debian_10_4_x64_20G_alibase_2020 0717.vhd	2020-07-17	•

Debian 10.4	debian_10_4_x64_20G_alibase_2020 0622.vhd	2020-06-22	•
Debian 9.12	debian_9_12_x64_20G_alibase_2020 0622.vhd	2020-06-22	•
Debian 10.4	debian_10_4_x64_20G_alibase_2020 0519.vhd	2020-05-19	•
Debian 9.12	debian_9_12_x64_20G_alibase_2020 0519.vhd	2020-05-19	•
Debian 10.3	debian_10_3_x64_20G_alibase_2020 0426.vhd	2020-04-26	•
Debian 9.12	debian_9_12_x64_20G_alibase_2020 0426.vhd	2020-04-26	•
Debian 10.3	debian_10_3_x64_20G_alibase_2020 0329.vhd	2020-03-29	• • • • • • • • • • • • • • • • • • •
Debian 9.12	debian_9_12_x64_20G_alibase_2020 0324.vhd	2020-03-24	•
Debian 9.12	debian_9_12_x64_20G_alibase_2020 0220.vhd	2020-02-20	•
Debian 10.3	debian_10_3_x64_20G_alibase_2020 0218.vhd	2020-02-18	•
Debian 9.11	debian_9_11_x64_20G_alibase_2019 1225.vhd	2019-12-25	•
Debian 10.2	debian_10_2_x64_20G_alibase_2019 1223.vhd	2019-12-24	• • o • Note Install cloudinit

Debian 9.9	debian_9_09_64_20G_alibase_20190 702.vhd	2019-7-2	•
Debian 9.9	debian_9_09_64_20G_alibase_20190 510.vhd	2019-5-10	•
Debian 8.11	11_64_20G_alibase_20190311.vhd	2019-3-11	• • • •
Debian 9.8	debian_9_08_64_20G_alibase_20190 225.vhd	2019-2-25	•
Debian 9.6	debian_9_06_64_20G_alibase_20190 103.vhd	2019-1-3	•
Debian 9.6	debian_9_06_64_20G_alibase_20181 212.vhd	2018-12-12	<ul><li> ○</li><li> ○</li></ul>

# Fedora CoreOS

Fedora CoreOS 34.2021052 9.3.0_3	fcos_34_20210529_3_0_3_x64_20G_ alibase_20210824.vhd	2021-08-24	• • • • • • • • • • • • • • • • • • •
Fedora CoreOS 34.2021052 9.3.0_3	fcos_34.20210529.3.0_3_x64_20G_ali base_20210623.vhd	2021-06-23	• • • • •

Fedora CoreOS 33.2021021 7.3.0_3	fcos_33.20210217.3.0_3_x86_20G_ali base_20210331.vhd	2021-03-31	• • • • • • • •

## Fedora

fedora_35_x64_20G_alibase_202205 31.vhd	2022-05-31	• • • • • • • • • • • • • • • • • • •
fedora_35_x64_20G_alibase_202205 10.vhd	2022-05-10	• • • • • • • • • • • • • • • • • • •
fedora_35_x64_20G_alibase_202203 22.vhd	2022-03-22	•
fedora_34_1_x64_20G_alibase_2022 0322.vhd	2022-03-22	•
fedora_34_1_x64_20G_alibase_2022 0208.vhd	2022-02-08	• • • • • • • • • • • • • • • • • • •
fedora_34_1_x64_20G_alibase_2021 1028.vhd	2021-10-28	•
fedora_34_1_x64_20G_alibase_2021 0927.vhd	2021-09-27	•
fedora_34_1_x64_20G_alibase_2021 0824.vhd	2021-08-24	• • • •
fedora_34_1_x64_20G_alibase_2021 0623.vhd	2021-06-23	•

fedora_33_0_x64_20G_alibase_2021 0318.vhd	2021-03-18	• • • •
fedora_33_0_x64_20G_alibase_2021 0128.vhd	2021-01-28	• • • • • • •

#### FreeBSD

FreeBSD 12.3	freebsd_12_3_x64_30G_alibase_202 20331.vhd	2022-03-31	
FreeBSD 13.0	freebsd_13_0_x86_30G_alibase_202 20324.vhd	2022-03-24	•
FreeBSD 11.4	freebsd_11_4_x64_30G_alibase_202 10319.vhd	2021-03-19	•
FreeBSD 12.1	freebsd_12_1_x64_30G_alibase_202 01022.vhd	2020-10-22	•
FreeBSD 11.3	freebsd_11_3_x64_30G_alibase_202 00803.vhd	2020-08-03	•
FreeBSD 11.3	freebsd_11_3_x64_20G_alibase_202 00420.vhd	2020-04-20	•
FreeBSD 11.2	freebsd_11_02_64_30G_alibase_201 90806.vhd	2019-8-6	• • • • • • • • • • • • • • • • • • •

# OpenSUSE

Opensuse 15.3	opensuse_15_3_x64_20G_alibase_20 220524.vhd	2022-05-24	•

Opensuse 15.3	opensuse_15_3_x64_20G_alibase_20 220426.vhd	2022-04-26	•
Opensuse 15.3	opensuse_15_3_x64_20G_alibase_20 220311.vhd	2022-03-11	• • • •
Opensuse 15.3	opensuse_15_3_x64_20G_alibase_20 220208.vhd	2022-02-08	• • • • • • • • • • • • • • • • • • •
Opensuse 15.3	opensuse_15_3_x64_20G_alibase_20 211227.vhd	2021-12-27	• • • • • • • • • • • • • • • • • • •
Opensuse 15.3	opensuse_15_3_x64_20G_alibase_20 211123.vhd	2021-11-23	• • • • • • • • • • • • • • • • • • •
Opensuse 15.3	opensuse_15_3_x64_20G_alibase_20 211027.vhd	2021-10-27	•
Opensuse 15.3	opensuse_15_3_x64_20G_alibase_20 210927.vhd	2021-09-27	•
Opensuse 15.3	opensuse_15_3_x64_20G_alibase_20 210824.vhd	2021-08-24	•
Opensuse 15.2	opensuse_15_2_x64_20G_alibase_20 210729.vhd	2021-07-29	• • • • •
Opensuse 15.2	opensuse_15_2_x64_20G_alibase_20 210623.vhd	2021-06-23	•
Opensuse 15.2	opensuse_15_2_x64_20G_alibase_20 210521.vhd	2021-05-21	•

Opensuse 15.2	opensuse_15_2_x64_20G_alibase_20 210420.vhd	2021-04-20	•
Opensuse 15.2	opensuse_15_2_x64_20G_alibase_20 210318.vhd	2021-03-18	• • • • • • • • • • • • • • • • • • •
Opensuse 15.2	opensuse_15_2_x64_20G_alibase_20 210128.vhd	2021-01-28	• • • • • • • • • • • • • • • • • • •
Opensuse 15.2	opensuse_15_2_x64_20G_alibase_20 201228.vhd	2020-12-28	•
Opensuse 15.2	opensuse_15_2_x64_20G_alibase_20 201120.vhd	2020-11-20	•
Opensuse 15.2	opensuse_15_2_x64_20G_alibase_20 200914.vhd	2020-09-14	•
Opensuse 15.2	opensuse_15_2_x64_20G_alibase_20 200818.vhd	2020-08-18	•
Opensuse 15.2	opensuse_15_2_x64_20G_alibase_20 200717.vhd	2020-07-17	•
Opensuse 15.1	opensuse_15_1_x64_20G_alibase_20 200623.vhd	2020-06-23	•
Opensuse 15.1	opensuse_15_1_x64_20G_alibase_20 200520.vhd	2020-05-20	•
Opensuse 15.1	opensuse_15_1_x64_20G_alibase_20 200426.vhd	2020-04-26	• • • • • • • • • • • • • • • • • • •
Opensuse 15.1	opensuse_15_1_x64_20G_alibase_20 200331.vhd	2020-03-31	•

Opensuse 15.1	opensuse_15_1_x64_20G_alibase_20 200222.vhd	2020-02-22	•
Opensuse 15.1	opensuse_15_1_x64_20G_alibase_20 191219.vhd	2019-12-19	• • • •

# **Red Hat Enterprise Linux**



•

• Which instance families do Red Hat Enterprise Linux (RHEL) images support?

•

•

•

•

•

•

•

•

•

•

•

# **Rocky Linux**

Rocky Linux 8.6	rockylinux_8_6_x64_20G_alibase_202 20524.vhd	2022-05-24	Linux time and time zones
Rocky Linux 8.5	rockylinux_8_5_x64_20G_alibase_202 20426.vhd	2022-04-26	•

Rocky Linux 8.5	rockylinux_8_5_x64_20G_alibase_202 20322.vhd	2022-03-22	• • • •
Rocky Linux 8.5	rockylinux_8_5_x64_20G_alibase_202 20214.vhd	2022-02-14	• • • •
Rocky Linux 8.5	rockylinux_8_5_x64_20G_alibase_202 11227.vhd	2021-12-27	• • • •
Rocky Linux 8.5	rockylinux_8_5_x64_20G_alibase_202 11126.vhd	2021-11-26	•

## SUSE Linux Enterprise Server



- SUSE Linux Enterprise Server 15 SP3
- SUSE Linux Enterprise Server 15 SP2
- SUSE Linux Enterprise Server 15 SP1
- SUSE Linux Enterprise Server 12 SP5
- SUSE Linux Enterprise Server 12 SP4
- SUSE Linux Enterprise Server 12 SP3

#### Ubuntu

Ubuntu20.0	ubuntu_20_04_x64_20G_alibase_202 20524.vhd	2022-05-24	•
Ubuntu18.0	ubuntu_18_04_x64_20G_alibase_202 20524.vhd	2022-05-24	•
Ubunt u20.0 4	ubuntu_20_04_x64_20G_alibase_202 20428.vhd	2022-04-28	•
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 20428.vhd	2022-04-28	•

Ubuntu20.0 4	ubuntu_20_04_x64_20G_alibase_202 20331.vhd	2022-03-31	• • • • •
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 20322.vhd	2022-03-22	• • • • • • • • • • • • • • • • • • •
Ubuntu20.0 4	ubuntu_20_04_x64_20G_alibase_202 20215.vhd	2022-02-15	• o • Linux time and time zones •
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 20208.vhd	2022-02-08	• • • • • • • • • • • • • • • • • • •
Ubuntu20.0 4	ubuntu_20_04_x64_20G_alibase_202 20110.vhd	2022-01-10	• • • • • • • • Linux time and time zones
Ubuntu20.0 4	ubuntu_20_04_x64_20G_alibase_202 11227.vhd	2021-12-27	• • • • • • • • • • • • • • • • • • •
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 11227.vhd	2021-12-27	• • • • • • • • • • • • • • • • • • •
Ubuntu20.0 4	ubuntu_20_04_x64_20G_alibase_202 11123.vhd	2021-11-23	• • • •
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 11123.vhd	2021-11-23	• • • • •

Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_202 11028.vhd	2021-10-28	•
Ubuntu20.0 4	ubuntu_20_04_x64_20G_alibase_202 11027.vhd	2021-10-27	•
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 11027.vhd	2021-10-27	•
Ubuntu20.0 4	ubuntu_20_04_x64_20G_alibase_202 10927.vhd	2021-09-27	•
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 10927.vhd	2021-09-27	•
Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_202 10927.vhd	2021-09-27	•
Ubuntu20.0 4	ubuntu_20_04_x64_20G_alibase_202 10824.vhd	2021-08-24	• • • • •
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 10824.vhd	2021-08-24	• • • •
Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_202 10824.vhd	2021-08-24	• • • •
Ubuntu20.0 4	ubuntu_20_04_x64_20G_alibase_202 10623.vhd	2021-06-23	•
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 10623.vhd	2021-06-23	•
Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_202 10623.vhd	2021-06-23	•

Ubuntu20.0	ubuntu_20_04_x64_20G_alibase_202 10521.vhd	2021-05-21	•
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 10521.vhd	2021-05-21	•
Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_202 10521.vhd	2021-05-21	•
Ubuntu20.0 4	ubuntu_20_04_x64_20G_alibase_202 10420.vhd	2021-04-20	•
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 10420.vhd	2021-04-20	•
Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_202 10420.vhd	2021-04-20	•
Ubuntu20.0 4	ubuntu_20_04_x64_20G_alibase_202 10318.vhd	2021-03-18	•
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 10318.vhd	2021-03-18	•
Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_202 10318.vhd	2021-03-18	• • • •
Ubuntu20.0 4	ubuntu_20_04_x64_20G_alibase_202 10128.vhd	2021-01-28	• • • • • • • • • • • • • • • • • • •
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 10128.vhd	2021-01-28	• • • • • • • • • • • • • • • • • • •
Ubuntu20.0 4	ubuntu_20_04_x64_20G_alibase_202 01228.vhd	2020-12-28	•

Ubuntu18.0	ubuntu_18_04_x64_20G_alibase_202 01228.vhd	2020-12-28	•
Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_202 01228.vhd	2020-12-28	•
Ubuntu20.0 4	ubuntu_20_04_x64_20G_alibase_202 01120.vhd	2020-11-20	•
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 01120.vhd	2020-11-20	•
Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_202 01120.vhd	2020-11-20	•
Ubuntu20.0 4	ubuntu_20_04_x64_20G_alibase_202 00914.vhd	2020-09-14	• • • • • • • • • • • • • • • • • • •
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 00914.vhd	2020-09-14	• • • • • • • • • • • • • • • • • • •
Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_202 00914.vhd	2020-09-14	•
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 00817.vhd	2020-08-17	•
Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_202 00817.vhd	2020-08-17	•
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 00717.vhd	2020-07-17	•
Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_202 00717.vhd	2020-07-17	•
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 00618.vhd	2020-06-18	•

Ubuntu16.0	ubuntu_16_04_x64_20G_alibase_202 00618.vhd	2020-06-18	•
Ubuntu20.0 4	ubuntu_20_04_x64_20G_alibase_202 00522.vhd	2020-05-22	•
Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_202 00522.vhd	2020-05-22	•
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 00521.vhd	2020-05-21	•
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 00426.vhd	2020-04-26	•
Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_202 00426.vhd	2020-04-26	•
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 00329.vhd	2020-03-29	• • • •
Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_202 00319.vhd	2020-03-19	•
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_202 00220.vhd	2020-02-20	•
Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_202 00220.vhd	2020-02-20	•
Ubuntu18.0 4	ubuntu_18_04_x64_20G_alibase_201 91225.vhd	2019-12-25	•
Ubuntu16.0 4	ubuntu_16_04_x64_20G_alibase_201 91225.vhd	2019-12-25	•
Ubuntu 18.04	ubuntu_18_04_64_20G_alibase_2019 0624.vhd	2019-6-24	•

Ubuntu 16.04	ubuntu_16_04_64_20G_alibase_2019 0620.vhd	2019-6-20	•
Ubuntu 16.04	ubuntu_16_04_64_20G_alibase_2019 0513.vhd	2019-5-13	•
Ubuntu 18.04	ubuntu_18_04_64_20G_alibase_2019 0509.vhd	2019-5-9	• • • •
Ubuntu 16.04	ubuntu_16_04_64_20G_alibase_2019 0301.vhd	2019-3-1	•
Ubuntu 18.04	ubuntu_18_04_64_20G_alibase_2019 0223.vhd	2019-2-23	•
Ubuntu 18.04	ubuntu_18_04_64_20G_alibase_2018 1212.vhd	2018-12-12	• • • • • • • • • • • • • • • • • • •
	anolisos_8_4_arm64_20G_anck_aliba se_20220519.vhd	2022-05-19	•
	anolisos_8_4_arm64_20G_rhck_aliba se_20220519.vhd	2022-05-19	•
	anolisos_7_9_arm64_20G_rhck_aliba se_20220518.vhd	2022-05-18	•

anolisos_8_4_arm64_20G_anck_aliba se_20220413.vhd	2022-04-13	•
anolisos_8_4_arm64_20G_rhck_aliba se_20220415.vhd	2022-04-15	•
anolisos_8_2_arm64_20G_anck_aliba se_20220413.vhd	2022-04-13	•
anolisos_8_2_arm64_20G_rhck_aliba se_20220415.vhd	2022-04-15	•
anolisos_7_9_arm64_20G_rhck_aliba se_20220413.vhd	2022-04-13	•
anolisos_7_9_arm64_20G_rhck_aliba se_20220316.vhd	2022-03-16	•
anolisos_7_9_arm64_20G_anck_aliba se_20220316.vhd	2022-03-16	•
debian_11_2_arm64_20G_alibase_20 220308.vhd	2022-03-08	•

anolisos_7_9_arm64_20G_anck_aliba se_20220111.vhd	2022-01-11	•
anolisos_7_9_arm64_20G_rhck_aliba se_20220110.vhd	2022-01-10	•
anolisos_7_9_arm64_20G_rhck_aliba se_20211210.vhd	2021-12-10	•
anolisos_7_7_arm64_20G_anck_aliba se_20211118.vhd	2021-11-18	•
anolisos_7_7_arm64_20G_rhck_aliba se_20211118.vhd	2021-11-18	•
ubuntu_20_04_arm64_20G_alibase_2 0211027.vhd	2021-10-27	•

anolisos_8_4_arm64_20G_anck_aliba se_20211008.vhd	2021-10-08	
anolisos_8_4_arm64_20G_rhck_aliba se_20211008.vhd	2021-10-08	

anolisos_8_2_arm64_20G_anck_aliba se_20211008.vhd	2021-10-08	
anolisos_8_2_arm64_20G_rhck_aliba se_20211008.vhd	2021-10-08	
centos_8_4_arm64_20G_alibase_202 10927.vhd	2021-09-27	• • • •
ubuntu_20_04_arm64_20G_alibase_2 0210927.vhd	2021-09-27	•

	centos_8_4_arm64_20G_alibase_202 10824.vhd	2021-08-24	• • • • • • • • • • • • • • • • • • • •
	centos_8_3_arm64_20G_alibase_202 10707.vhd	2021-07-07	•
	centos_8_4_arm64_20G_alibase_202 10706.vhd	2021-07-06	•
	centos_7_9_arm64_20G_alibase_202 10706.vhd	2021-07-06	•
	ubuntu_18_04_arm64_20G_alibase_2 0210706.vhd	2021-07-06	•
	debian_10_9_arm64_20G_alibase_20 210706.vhd	2021-07-06	•
Debian 10.7	debian_10_7_amd_x64_20G_alibase_ 20201231.vhd	2020-12-31	•
Ubuntu 20.04	ubuntu_20_04_amd_x64_20G_alibas e_20200820.vhd	2020-08-20	•
CentOS 7.8	centos_7_8_amd_x64_20G_alibase_2 0200806.vhd	2020-08-06	•

Cent OS 8.2	centos_8_2_amd_x64_20G_alibase_2 0200804.vhd	2020-08-04	•
Ubuntu 18.04	ubuntu_18_04_amd_x64_20G_alibas e_20200804.vhd	2020-08-04	•
	centos_8_5_uefi_x64_20G_alibase_2 0220328.vhd	2022-03-28	•
	ubuntu_20_04_uefi_x64_20G_alibase _20220324.vhd	2022-03-24	<ul><li> o</li><li> Linux time and time zones</li><li> o</li></ul>
	anolisos_8_4_x64_20G_rhck_uefi_ali base_20220328.vhd	2022-03-28	• • • • • • • • • •
	anolisos_8_4_x64_20G_anck_uefi_ali base_20220328.vhd	2022-03-28	• • • • • •
	ubuntu_18_04_uefi_x64_20G_alibase _20220330.vhd	2022-03-30	• • • • • • • • • • • • • • • • • • •
	debian_10_12_uefi_x64_20G_alibase _20220330.vhd	2022-03-30	•
	debian_10_10_uefi_x64_20G_alibase _20210611.vhd	2021-06-11	•

debian_9_13_uefi_x64_20G_alibase_ 20210611.vhd	2021-06-11	•
centos_7_9_uefi_x64_20G_alibase_2 0210611.vhd	2021-06-11	•
centos_8_4_uefi_x64_20G_alibase_2 0210611.vhd	2021-06-11	•
centos_8_3_uefi_x64_20G_alibase_2 0210611.vhd	2021-06-11	•
ubuntu_20_04_uefi_x64_20G_alibase _20210609.vhd	2021-06-09	•
ubuntu_18_04_uefi_x64_20G_alibase _20210609.vhd	2021-06-09	•

•

•

#### Windows Server 2022

•	2022-06-16	•
•	2022-05-18	•

•	2022-04-18	•
•	2022-03-11	•
•	2022-03-11	•

## Windows Server 2019

•	2022-06-16	•
•	2022-06-16	•
•	2022-05-18	•
•	2022-05-18	•
•	2022-04-18	•
•	2022-04-18	•
•	2022-03-11	•
•	2022-03-11	•
•	2021-12-17	•

•	2021-12-17	•
•	2021-10-15	•
•	2021-10-15	•
•	2021-09-16	•
•	2021-09-16	•
•	2021-08-17	•
•	2021-08-17	•
•	2021-07-16	•
•	2021-07-15	•
•	2021-06-18	•
•	2021-06-18	•
•	2021-05-16	•
•	2021-05-16	•
•	2021-04-16	•

•	2021-04-16
•	2021-03-16
•	2021-03-16
•	2021-02-20
•	2021-02-20
•	2021-01-14
•	2021-01-14
•	2020-12-17
•	2020-12-17
•	2020-11-20
•	2020-11-20
•	2020-10-15
•	2020-10-15
•	2020-09-14

•	2020-09-14	•
•	2020-08-14	•
•	2020-08-14	•
•	2020-07-23	
•	2020-07-23	
•	2020-06-15	•
•	2020-06-15	•
•	2020-05-16	•
•	2020-05-16	•
•	2020-04-16	•
•	2020-04-16	•
•	2020-03-14	•

•	2020-03-14	•
•	2020-02-25	•
•	2020-02-13	•
•	2020-01-16	•
•	2019-12-20	
•	2019-10-12	
•	2019-8-16	•
•	2019-7-18	•
•	2019-5-28	
•	2019-3-18	

# Windows Server 2016

•	2022-06-16	•
•	2022-05-18	•

•	2022-04-18	•
•	2022-03-11	•
•	2021-12-17	•
•	2021-11-15	•
•	2021-10-15	•
•	2021-09-16	•
•	2021-08-17	•
•	2021-07-15	•
•	2021-06-18	•
•	2021-05-16	•
•	2021-04-16	•
•	2021-03-16	•
•	2021-02-20	•
•	2021-01-14	•

•	2020-12-17	•
•	2020-11-20	•
•	2020-10-15	•
•	2020-09-14	•
•	2020-08-14	•
•	2020-07-23	
•	2020-06-15	•
•	2020-05-16	•
•	2020-04-16	•
•	2020-03-14	•
•	2020-02-13	•
•	2020-01-16	•
•	2019-12-24	

•	2019-10-12	
•	2019-8-16	•
•	2019-7-18	•
•	2019-5-23	
•	2019-3-18	
•	2018-12-20	•

# Windows Server 2012

•	2022-06-16	•
•	2022-05-18	•
•	2022-04-18	•
•	2022-03-11	•
•	2021-12-17	•
•	2021-11-15	•

•	2021-10-15
•	2021-09-16
•	2021-08-17
•	2021-07-15
•	2021-06-18
•	2021-05-16
•	2021-04-16
•	2021-03-16
•	2021-02-20
•	2021-01-14
•	2020-12-17
•	2020-11-20
•	2020-10-15

•	2020-09-14	•
•	2020-08-14	•
•	2020-07-23	
•	2020-06-15	•
•	2020-05-16	•
•	2020-04-16	•
•	2020-03-14	•
•	2020-02-13	•
•	2020-01-16	•
•	2019-12-18	
•	2019-10-12	
•	2019-8-16	•
•	2019-7-18	•

•	2019-5-23	
•	2019-3-18	
•	2018-12-20	•

# Windows Server 2008

? Note			
	•	2020-01-16	•
	•	2019-12-20	
	•	2019-10-12	
	•	2019-8-16	•
	•	2019-8-16	•
	•	2019-7-18	•
	•	2019-5-17	•
	•	2019-5-15	•
	•	2019-3-18	

		•
•	2018-12-20	•
		•

# Windows Server Version 2004

•	2021-12-17	•
•	2021-12-17	•
•	2021-11-15	•
•	2021-11-15	•
•	2021-10-15	•
•	2021-10-15	•
•	2021-09-16	•
•	2021-09-16	•
•	2021-08-17	•
•	2021-08-17	•
•	2021-07-16	•

•	2021-07-15
•	2021-06-18
•	2021-06-18
•	2021-05-16
•	2021-05-16
•	2021-04-16
•	2021-04-16
•	2021-03-16
•	2021-03-16
•	2021-02-20
•	2021-02-20
•	2021-01-14
•	2021-01-14
•	2020-12-17

•	2020-12-17	•
•	2020-11-20	•
•	2020-11-20	•
•	2020-10-15	•
•	2020-10-15	•
•	2020-09-14	•
•	2020-09-14	•
•	2020-08-14	•
•	2020-08-14	•

# Windows Server Version 1909

•	2020-07-23	
•	2020-07-23	
•	2020-06-15	•
•	2020-06-15	•

•	2020-05-16	•
•	2020-05-16	•
•	2020-04-16	•
•	2020-04-16	•
•	2020-03-15	•
•	2020-02-13	•
•	2020-01-16	•
•	2019-12-19	

# Windows Server Version 1903

•	2019-10-12	
•	2019-8-16	•
•	2019-7-18	•

# Windows Server Version 1809

•	2019-5-28	
•	2019-3-18	
•	2018-12-22	•

# 5.Alibaba Cloud Linux 5.1. Overview

Alibaba Cloud Linux is a Linux distribution developed by Alibaba Cloud. It is built on the Anolis OS developed by the OpenAnolis community (OpenAnolis) and is fully compatible with the Cent OS and Red Hat Enterprise Linux (RHEL) ecosystems. Alibaba Cloud Linux offers a safe, stable, and high-performance customized runtime environment for cloud-based applications and is optimized for Alibaba Cloud infrastructure to improve user experience. Alibaba Cloud Linux 2 and Alibaba Cloud Linux 3 have been released

For more information, visit the Alibaba Cloud Linux 2 product page.

#### Use scenarios

Alibaba Cloud Linux is suitable for the following Elastic Compute Service (ECS) instance families and workloads:

- A variety of instance families, including ECS Bare Metal Instance families. For more information, see Instance family. Take note of the following items:
  - Alibaba Cloud Linux supports instances that have up to 768 vCPUs.
  - Alibaba Cloud Linux supports instances that have 0.5 GiB to 12,288 GiB of memory.
  - Alibaba Cloud Linux does not support non-I/O optimized instances.
- A variety of workloads in cloud environments, such as databases, cloud native containers, data analytics, web applications, and other workloads in the production environment.

### Billing

Alibaba Cloud Linux images are provided free of charge. However, you must pay for other resources that are consumed when you create instances from the images, such as vCPUs, memory, storage, public bandwidth, and snapshots. For information about the billing details, see Overview.

#### **Advantages**

Compared with other Linux distributions, Alibaba Cloud Linux has the following advantages:

- Alibaba Cloud provides long-term free software maintenance and technical support for Alibaba Cloud Linux. For information about the end of life (EOL) plans of different Alibaba Cloud Linux releases, see EOL plans of Alibaba Cloud Linux.
- Alibaba Cloud Linux is optimized for integration with Alibaba Cloud infrastructure and increases system startup speeds and runtime performance. It has long-tested and refined from its extensive use in Alibaba and Alibaba Cloud services and proved its outstanding performance and stability.
- Alibaba Cloud Linux is equipped with an updated Linux kernel, user-mode packages, and toolkits, and provides the latest enhanced operating system features from the Linux community to support cloudbased application environments.
  - Alibaba Cloud Linux 3 is powered by Linux kernel 5.10 LTS and uses GCC 10.2, Binutils 2.35, and glibc 2.32 by default.
  - Alibaba Cloud Linux 2 is powered by Linux kernel 4.19 LTS and provides devtoolset support for GCC 7.3.1, GCC 8.2.1, and GCC 9.1.1. For information about how to use devtoolset, see the "Use experimental software packages" section in Use experimental software packages.

- Alibaba Cloud Linux is compatible with the Cent OS and RHEL ecosystems. Most Cent OS and RHEL software can run on Alibaba Cloud Linux with little or no modifications.
  - Alibaba Cloud Linux 3 is built on Anolis OS 8 and compatible with the Cent OS 8 and RHEL 8 ecosystems.
  - Alibaba Cloud Linux 2 is built on Anolis OS 7 and compatible with the Cent OS 7 and RHEL 7 ecosystems.
- Alibaba Cloud Linux provides policies for monitoring and fixing security vulnerabilities and ensures system security. For more information, see Security updates.

#### **Features**

- Alibaba Cloud Linux is distributed with the latest version of the Alibaba Cloud kernel.
  - Alibaba Cloud Linux 2 is developed based on Linux kernel 4.19 LTS, and Alibaba Cloud Linux 3 is developed based on Linux kernel 5.10 LTS. Alibaba Cloud Linux is optimized to provide features for cloud-based scenarios, improve performance, and fix bugs. For more information, see Alibaba Cloud Linux 2镜像发布记录 and Alibaba Cloud Linux 3镜像发布记录.
  - Alibaba Cloud Linux provides customized and optimized kernel startup parameters and system configuration parameters for the ECS instance environment.
  - Alibaba Cloud Linux provides kdump, which is a kernel dumping mechanism used when the operating system fails. You can enable or disable this feature without the need to restart the operating system.
  - Alibaba Cloud Linux provides Kernel Live Patching (KLP). For more information, see Overview of the Kernel Live Patching feature.
- Alibaba Cloud Linux has software pre-installed or updated.
  - User-mode packages keep compatible with new CentOS versions. The user-mode packages of CentOS 7 can be used on Alibaba Cloud Linux 2. The user-mode packages of CentOS 8 can be used on Alibaba Cloud Linux 3.
  - o Alibaba Cloud Linux is pre-installed with Alibaba Cloud CLI.
  - Fixes for common vulnerabilities and exposures (CVEs) are continuously updated until the EOL of Alibaba Cloud Linux. For more information, see Alibaba Cloud Linux 2.1903 Security Advisories and Alibaba Cloud Linux 3 Security Advisories. Alibaba Cloud Linux provides solutions to automatically fix vulnerabilities. For more information, see Use YUM to perform security updates (Alibaba Cloud Linux 2).
- Alibaba Cloud Linux accelerates the startup process, improves runtime performance, and enhances system stability.
  - Alibaba Cloud Linux optimizes the startup speed of ECS instances. Tests have proven that Alibaba Cloud Linux can reduce the startup time by up to 60% compared with other operating systems.
  - Alibaba Cloud Linux optimizes the scheduling, memory, and I/O subsystems. In some open source benchmark tests, the performance of Alibaba Cloud Linux exceeds that of other operating systems by up to 30%.
  - Alibaba Cloud Linux provides enhanced system stability and can reduce downtime by 50% compared with other operating systems.

# Alibaba Cloud Linux images

Release	lmage	ID of the latest image version	lmage description
	Alibaba Cloud Linux 3.2104 64-bit	aliyun_3_x64_20G_aliba se_20210910.vhd	The default standard Alibaba Cloud Linux image that supports the x86 architecture.
Alibaba Cloud Linux 3.2104 64-bit (Quick Start)			This image is an image customized based on the Alibaba Cloud kernel. It supports quick boot (Qboot) and starts instances from the kernel. Instances created from this image start faster than those created from other images and share the same default runtime as those created from the standard image. This image has the following characteristics:
	aliyun_3_x64_20G_qbo ot_alibase_20211214.v hd	<ul> <li>The first-time startup of instances is accelerated, whereas subsequent startup of instances remains at normal speed.</li> <li>The memory initialization process is optimized and accelerated, devices that take</li> </ul>	
		time to start such as mice are modularized, and the kernel boot is sped up.  The cloud-init service is replaced with the latest AliyunInit service to reduce the amount of time required to initialize the operating system.	
Linux 3			<ul> <li>The kernel startup parameter o this image cannot be modified.</li> </ul>
		This image does not support the following instance families: GPU-accelerated compute- optimized instance families, vGPU-accelerated instance families, FPGA-accelerated compute-optimized instance families, NPU-accelerated	
		compute-optimized instance families, heterogeneous instance families, Super Computing Cluster (SCC) instance families, and security- enhanced instance families.	

Release	lmage	ID of the latest image version	Image description
	Alibaba Cloud Linux 3.2104 64-bit (UEFI)	aliyun_3_x64_20G_uefi_ alibase_20210910.vhd	This image is developed based on the default x86-based standard image and supports only the Unified Extensible Firmware Interface (UEFI) boot mode.
	Alibaba Cloud Linux 3.2104 64 for ARM	aliyun_3_arm64_20G_al ibase_20210910.vhd	The default standard Alibaba Cloud Linux image that supports the ARM architecture.
	Alibaba Cloud Linux 2.1903 LTS 64-bit	aliyun_2_1903_x64_20G _alibase_20210726.vhd	The default standard Alibaba Cloud Linux image that supports the x86 architecture.

Release	lmage	ID of the latest image version	lmage description
Alibaba Cloud Linux 2			This image is an image customized based on the Alibaba Cloud kernel. It supports quick boot (Qboot) and starts instances from the kernel. Instances created from this image start faster than those created from other images and share the same default runtime as those created from the standard image. This image has the following characteristics:  • The first-time startup of
			instances is accelerated, whereas subsequent startup of instances remains at normal speed.
			<ul> <li>The memory initialization process is optimized and accelerated, devices that take time to start such as mice are modularized, and the kernel boot is sped up.</li> </ul>
	Alibaba Cloud Linux 2.1903 LTS 64-bit (Quick Start)	aliyun_2_1903_x64_20G _qboot_alibase_20210 726.vhd	<ul> <li>The cloud-init service is replaced with the latest AliyunInit service to reduce the amount of time required to initialize the operating system.</li> </ul>
			• The kernel startup parameter of this image cannot be modified.
			• This image does not support the following instance families: GPU-accelerated compute-optimized instance families, vGPU-accelerated instance families, FPGA-accelerated compute-optimized instance families, NPU-accelerated compute-optimized instance families, heterogeneous instance families, Super Computing Cluster (SCC) instance families, and security-enhanced instance families.

Release	lmage	ID of the latest image version	lmage description
	Alibaba Cloud Linux 2.1903 LTS 64-bit (UEFI)	aliyun_2_1903_x64_20G _uefi_alibase_2021072 6.vhd	This image is developed based on the default x86-based standard image and supports only the UEFI boot mode.
	Alibaba Cloud Linux 2.1903 64-bit (Trusted)	aliyun_2_1903_x64_20G _secured_alibase_2021 0325.vhd	This image is developed based on the default x86-based standard image and supports Alibaba Cloud trusted instances.

# Obtain Alibaba Cloud Linux images

You can use one of the following methods to obtain and use Alibaba Cloud Linux images:

- ECS instances
  - When you create an ECS instance, select **Public Image** and then select Alibaba Cloud Linux and its version. For more information, see Create an instance by using the wizard.
  - Change the operating system of an existing ECS instance to Alibaba Cloud Linux by replacing the system disk. For more information, see Replace the operating system of an instance by using a public image.
- On-premises environments such as a virtualization environment based on Kernel-based Virtual Machine (KVM)

Download and install an Alibaba Cloud Linux image, and then restart the system. For more information, see Use Alibaba Cloud Linux 2 images in an on-premises environment or Use Alibaba Cloud Linux 3 images in an on-premises environment.

#### Release notes

- Alibaba Cloud Linux 3镜像发布记录
- Alibaba Cloud Linux 2镜像发布记录

# Security updates

For information about security updates for CVEs for Alibaba Cloud Linux, see the following documents:

- Alibaba Cloud Linux 3 Security Advisories
- Alibaba Cloud Linux 2.1903 Security Advisories

### Technical support

Alibaba Cloud provides the following technical support for Alibaba Cloud Linux:

- Security updates and vulnerability fixes are provided free of charge based on the lifecycle of Alibaba Cloud Linux releases. For information about the lifecycle of different Alibaba Cloud Linux releases, see EOL plans of Alibaba Cloud Linux. You can to obtain free technical support.
- Alibaba Cloud Linux repositories are frequently updated and base images are updated every month. Updates cover new features, security updates, and vulnerability fixes.
- Security updates are provided from YUM repositories. You can run the **yum update** command to update an image to the latest version.

# 5.2. Release notes

# 5.3. Use images

# 5.3.1. Use Alibaba Cloud Linux 3 images in an onpremises environment

Alibaba Cloud Linux 3 images are available in various formats and have cloud-init built in. This topic describes how to use Alibaba Cloud Linux 3 images in an on-premises environment.

#### Context

Alibaba Cloud Linux 3 images can run only on kernel-based virtual machines (KVMs). Alibaba Cloud Linux 3 images cannot directly start KVMs. You must configure a boot image.

In this topic, the on-premises operating system is CentOS. Alibaba Cloud Linux 3 is used to create a KVM, and cloud-init is used to initialize the system settings of the KVM. The NoCloud data source is used to create on-premises configuration files. After the configuration files are attached to the KVM as virtual disks, the KVM can be started. For more information about cloud-init, see Alibaba Cloud (AliYun) on the official cloud-init website.

# Step 1: Download an Alibaba Cloud Linux 3 image to your computer

You can download an Alibaba Cloud Linux 3 image to your computer from the Alibaba Cloud Linux 3 On-premise Image page. Alibaba Cloud Linux 3 images are available in the VHD or QCOW2 format.

# Step 2: Obtain the seed.img boot image from your computer

You must configure the network, account, and YUM repository for the boot image. Typically, the image is named seed.img. You can change the name of the image, but we recommend that you do not.

Notice The seed.img image contains only the configuration files that are required to start cloud-init. The image does not contain Alibaba Cloud Linux 3 system files.

You can use one of the following methods to obtain the seed.img image:

• Go to the Alibaba Cloud Linux 3 On-premise Image page and click *seed.img* to download the seed.img image.

The configurations in this image cannot be modified, which makes it less ideal for some scenarios. Make sure that you are familiar with the image before you use it.

- Perform the following operations to manually generate the seed.img image based on the NoCloud data source.
  - 1. In an on-premises directory, create two configuration files named meta-data and user-data
    - i. Create a directory named seed and go to the directory.

```
mkdir seed
cd seed/
```

ii. Create the meta-data configuration file.

The following example shows the content of the configuration file. You can modify the content.

```
#cloud-config
#vim:syntax=yaml
local-hostname: alinux-host
```

iii. Create the user-data configuration file.

The following example shows the content of the configuration file. You can modify the content.

```
#cloud-config
        #vim:syntax=yaml
        # Create an account named alinux who is authorized to run sudo commands.
          - default
          - name: alinux
           sudo: ['ALL=(ALL) ALL']
            plain text passwd: aliyun
            lock passwd: false
        # Create a YUM repository for Alibaba Cloud Linux 3.
        yum repos:
            alinux3-module:
                name: alinux3-module
                baseurl: https://mirrors.aliyun.com/alinux/$releasever/module/$basearch/
                enabled: 1
                gpgcheck: 1
                qpqkey: https://mirrors.aliyun.com/alinux/$releasever/RPM-GPG-KEY-ALINUX-3
            alinux3-updates:
                name: alinux3-updates
                baseurl: https://mirrors.aliyun.com/alinux/$releasever/updates/$basearch/
                enabled: 1
                gpgcheck: 1
                qpgkey: https://mirrors.aliyun.com/alinux/$releasever/RPM-GPG-KEY-ALINUX-3
            alinux3-plus:
                name: alinux3-plus
                baseurl: https://mirrors.aliyun.com/alinux/$releasever/plus/$basearch/
                gpgcheck: 1
                gpgkey: https://mirrors.aliyun.com/alinux/$releasever/RPM-GPG-KEY-ALINUX-3
            alinux3-powertools:
                name: alinux3-powertools
                baseurl: https://mirrors.aliyun.com/alinux/$releasever/powertools/$basearch
                gpgcheck: 1
                enabled: 1
                gpgkey: https://mirrors.aliyun.com/alinux/$releasever/RPM-GPG-KEY-ALINUX-3
            alinux3-os:
                name: alinux3-os
                baseurl: https://mirrors.aliyun.com/alinux/$releasever/os/$basearch/
                gpgcheck: 1
                enabled: 1
                gpgkey: https://mirrors.aliyun.com/alinux/$releasever/RPM-GPG-KEY-ALINUX-3
2. Install the cloud-utils software package on your computer.
```

```
yum install -y cloud-utils
```

3. In the seed directory, run the following command to generate the seed.img image:

```
cloud-localds seed.img user-data meta-data
```

### Step 3: Start the KVM

Use one of the following methods to start the KVM. Then, use the account information in the user-data configuration file to log on to the KVM.

- Use libvirt to start the KVM.
  - i. Create a configuration file in the XML format on your computer. In this example, the file is named alinux3.xml and contains the following content. You can modify this file based on your requirements.

```
<domain type='kvm'>
   <name>alinux3</name>
    <memory>1048576/memory><!-- Set the memory size to 1 GB. -->
    <vcpu>1</vcpu>
    cpu mode='host-passthrough'><!-- Start the KVM on the server of the corresponding
architecture. -->
   </cpu>
    <0s>
       <type arch='x86_64'>hvm</type><!-- If the architecture of the server is ARM 6 \,
4-bit, set arch to aarch64. -->
       <boot dev='hd'/>
   </09>
   <clock sync="localtime"/>
   <on poweroff>destroy</on poweroff>
   <on reboot>restart</on reboot>
   <on_crash>restart</on_crash>
    <devices>
       <emulator>/usr/bin/qemu-kvm</emulator><!-- Configure a KVM path based on the</pre>
operating system. For example, the KVM path for Ubuntu is /usr/bin/kvm. -->
       <disk type='file' device='disk'><!-- Specify the type parameter based on the</pre>
image format. Set type to qcow2 if the image is in the QCOW2 format, and set type to
vpc if the image is in the VHD format. -->
           <driver name='qemu' type='qcow2' cache='none' dataplane='on' io='native'/</pre>
> <!-- If you want to create a snapshot in the QCOW2 format, you must disable datapla
ne. -->
           <source file='path'/> <!-- Enter the absolute path of the Alibaba Cloud {\tt L}
inux 3 image. -->
           <target dev='vda' bus='virtio'/>
        <!-- Add the information of the seed.img image. -->
        <disk type='file' device='disk'>
           <driver name='qemu' type='raw'/>
            <source file='/path/to/your/seed.img'/> <!-- Enter the absolute path of t</pre>
he seed.img image -->
           <target dev='vdb' bus='virtio'/>
        </disk>
        <interface type='network'>
           <source network='default'/>
            <model type='virtio'/>
        </interface>
        <console type='pty'>
           <target type='virtio' port='0'/>
        </console>
        <video>
            <model type='cirrus' vram='9216' heads='1'/>
            <alias name='video0'/>
        </video>
        <input type='tablet' bus='usb'/>
        <input type='mouse' bus='ps2'/>
        <graphics type='vnc' port='-1' autoport='yes'/>
    </devices>
</domain>
```

ii. Run the virsh command to start the KVM. The following code shows sample commands.

**? Note** By default, libvirt is started by a common user. Make sure that common users have the permissions to manage image files and their paths.

```
virsh define alinux3.xml
virsh start KVMName  # Enter the name of the KVM.
```

• Run the qemu-kvm command to start the KVM. You must append the following parameter information to the command. Set the file parameter to the actual absolute path of the seed.img image.

```
-drive file=/path/to/your/seed.img,if=virtio,format=raw
```

For more information about how to configure and manage virtualization, see Configuring and managing virtualization.

• Use the graphical interface (virt-manager) to start the KVM. Before you start the KVM, find the configuration file of the KVM on your computer and add the absolute path of the seed.img image to the configuration file.

# 5.3.2. Use Alibaba Cloud Linux 2 images in an onpremises environment

Alibaba Cloud Linux 2 images are available in various formats and have cloud-init built in. This topic describes how to use Alibaba Cloud Linux 2 images in an on-premises environment.

#### Context

Alibaba Cloud Linux 2 images can run only on kernel-based virtual machines (KVMs). Alibaba Cloud Linux 2 images cannot directly start KVMs. You must configure a boot image. In this topic, the on-premises operating system is CentOS. Alibaba Cloud Linux 2 is used to create a KVM, and cloud-init is used to initialize the system settings of the KVM. For more information about cloud-init, see Alibaba Cloud (AliYun) on the official cloud-init website. The NoCloud data source is used to create on-premises configuration files. After the configuration files are attached to the KVM as virtual disks, the KVM can be started.

This topic is intended for users who are familiar with KVMs.

#### Step 1: Download an Alibaba Cloud Linux 2 image to your computer

You can download an Alibaba Cloud Linux 2 image to your computer from the Alibaba Cloud Linux 2 Onpremise Image page. Alibaba Cloud Linux 2 images are available in the VHD or QCOW2 format.

#### Step 2: Obtain the seed.img boot image from your computer

You must configure the network, account, and YUM repository for the boot image. In most cases, the image is named seed.img. You can change the name of the image, but we recommend that you do not.

Notice The seed.img image contains only the configuration files that are required to start cloud-init. The image does not contain Alibaba Cloud Linux 2 system files.

You can use one of the following methods to obtain the seed.img image:

• Go to the Alibaba Cloud Linux 2 On-premise Image page and click *seed.img* to download the seed.img image.

The configurations in this image cannot be modified, which makes it less ideal for some scenarios. Make sure that you are familiar with the image before you use it.

- Perform the following operations to manually generate the seed.img image based on the NoCloud data source:
  - 1. In an on-premises directory, create two configuration files named meta-data and user-data.
    - i. Create a directory named seed and go to the directory.

```
mkdir seed cd seed/
```

ii. Create the meta-data configuration file.

The following example shows the content of the configuration file. You can modify the content.

```
#cloud-config
#vim:syntax=yaml
local-hostname: alinux-host
```

iii. Create the user-data configuration file.

The following example shows the content of the configuration file. You can modify the content.

```
#cloud-config
#vim:syntax=yaml
# Create a user named alinux who is authorized to run sudo commands.
  - default
  - name: alinux
   sudo: ['ALL=(ALL)
   plain text passwd: aliyun
    lock passwd: false
# Create a YUM repository for Alibaba Cloud Linux 2.
yum repos:
   base:
        baseurl: https://mirrors.aliyun.com/alinux/$releasever/os/$basearch/
       enabled: true
       gpgcheck: true
        gpgkey: https://mirrors.aliyun.com/alinux/RPM-GPG-KEY-ALIYUN
        name: Aliyun Linux - $releasever - Base - mirrors.aliyun.com
    updates:
        baseurl: https://mirrors.aliyun.com/alinux/$releasever/updates/$basearch/
        enabled: true
        gpgcheck: true
        gpgkey: https://mirrors.aliyun.com/alinux/RPM-GPG-KEY-ALIYUN
        name: Aliyun Linux - $releasever - Updates - mirrors.aliyun.com
    extras:
        baseurl: https://mirrors.aliyun.com/alinux/$releasever/extras/$basearch/
        enabled: true
        gpgcheck: true
        gpgkey: https://mirrors.aliyun.com/alinux/RPM-GPG-KEY-ALIYUN
        name: Aliyun Linux - $releasever - Extras - mirrors.aliyun.com
    plus:
       baseurl: https://mirrors.aliyun.com/alinux/$releasever/plus/$basearch/
        enabled: true
        gpgcheck: true
        gpgkey: https://mirrors.aliyun.com/alinux/RPM-GPG-KEY-ALIYUN
        name: Aliyun Linux - $releasever - Plus - mirrors.aliyun.com
```

2. Install the cloud-utils software package on your computer.

```
yum install -y cloud-utils
```

3. In the seed directory, run the following command to generate the seed.img image:

```
cloud-localds seed.img user-data meta-data
```

#### Step 3: Start the KVM

Use one of the following methods to start the KVM. Then, use the account information in the user-data configuration file to log on to the KVM.

- Use libvirt to start the KVM.
  - i. Create a configuration file in the XML format on your computer. The sample file is named alinu x2.xml and contains the following content. You can modify this file based on your requirements.

```
<domain type='kvm'>
    <name>alinux2</name>
    <memory>1048576/memory> <!-- 1 GB memory -->
    <vcpu>1</vcpu>
    <os>
        <type arch='x86 64'>hvm</type>
       <boot dev='hd'/>
    </os>
    <clock sync="localtime"/>
    <on poweroff>destroy</on poweroff>
    <on reboot>restart</on reboot>
    <on crash>restart</on crash>
    <devices>
        <emulator>/usr/bin/qemu-kvm/emulator><!-- Configure a KVM path based on the</pre>
operating system. For example, the KVM path for Ubuntu is /usr/bin/kvm. -->
       <disk type='file' device='disk'><!-- Specify the type parameter based on the</pre>
image format. Set type to gcow2 if the image is in the QCOW2 format, and set type to
vpc if the image is in the VHD format. -->
            -- If you want to create a snapshot in the QCOW2 format, you must disable
dataplane. -->
            <source file='path'/> <!-- Enter the absolute path of the Alibaba Cloud {\tt L}
inux 2 image. -->
            <target dev='vda' bus='virtio'/>
        </disk>
        -- Add the information of seed.img. -->
        <disk type='file' device='disk'>
           <driver name='qemu' type='raw'/>
            <source file='/path/to/your/seed.img'/> <!-- Enter the absolute path of s</pre>
eed.img. -->.
            <target dev='vdb' bus='virtio'/>
        </disk>
        <interface type='network'>
            <source network='default'/>
            <model type='virtio'/>
        </interface>
        <console type='pty'>
            <target type='virtio' port='0'/>
        </console>
        <video>
            <model type='cirrus' vram='9216' heads='1'/>
            <alias name='video0'/>
        </video>
        <input type='tablet' bus='usb'/>
        <input type='mouse' bus='ps2'/>
        <graphics type='vnc' port='-1' autoport='yes'/>
    </devices>
</domain>
```

#### ii. Run the virsh command to start the KVM.

**? Note** By default, libvirt is started by a common user. Make sure that common users have the permissions to manage image files and their paths.

```
virsh define alinux2.xml
virsh start KVMName  # Enter the name of the KVM.
```

• Run the qemu-kvm command to start the KVM. You must append the following parameter information to the command. Set the file parameter to the actual absolute path of the seed.img image.

```
-drive file=/path/to/your/seed.img,if=virtio,format=raw
```

For more information about how to use the libvirt and qemu-kvm commands, see Installing Virtualization Packages Manually.

• Use the graphical interface (virt-manager) to start the KVM. Before you start the KVM, find the configuration file of the KVM on your computer and add the absolute path of the seed.img image to the configuration file.

# 5.3.3. Use Alibaba Cloud Linux 2

This topic describes how to view system or kernel parameters, enable or disable kdump, and obtain the Debuginfo packages, source code packages, and experimental software packages in Alibaba Cloud Linux 2 operating systems.

# View or modify parameters

• View or modify system parameters

You can run the *sysctl* command to view or modify the runtime system parameters of Alibaba Cloud Linux 2. Alibaba Cloud Linux 2 has the following kernel configuration parameters updated in the /etc/sysctl.d/50-aliyun.conf file.

System parameter	Description
<pre>kernel.hung_task_timeout_ secs = 240</pre>	Increases the kernel hung_task timeout period in seconds to prevent frequent hung_task prompts.
kernel.panic_on_oops = 1	Throws the kernel panic exception when an Oops error occurs in the kernel. System failure details are automatically captured if kdump is configured.
<pre>kernel.watchdog_thresh = 50</pre>	Sets the thresholds for events such as high resolution timer (hrtimer), non-maskable interrupt (NMI), soft lockup, and hard lockup to prevent potential kernel false positives.
<pre>kernel.hardlockup_panic = 1</pre>	Throws the kernel panic exception when a hard lockup error occurs in the kernel. System failure details are automatically captured if kdump is configured.

View kernel parameters

You can run the cat /proc/cmdline command to view the runtime kernel parameters of Alibaba Cloud Linux 2. Alibaba Cloud Linux 2 has the following kernel parameters updated.

Kernel parameter	Description
crashkernel=0M-2G:0M,2G-8 G:192M,8G-:256M	Reserves memory space for kdump.
cryptomgr.notests	Disables crypto self-check during kernel startup to accelerate system startup.
cgroup.memory=nokmem	Disables the kernel memory statistics feature of the memory cgroup to prevent potential kernel instability.
rcupdate.rcu_cpu_stall_ti meout=300	Increases the timeout threshold of RCU CPU Stall Detector to 300 seconds to prevent kernel false positives.

# Enable or disable kdump

Alibaba Cloud Linux 2 provides the kdump service. After this service is enabled, kernel errors can be captured to help you analyze kernel failures.

**? Note** If the memory of the selected instance type does not exceed 2 GiB, the kdump service cannot be used.

• Run the following commands in sequence to enable the kdump service:

```
# Enable the kdump service.
sudo systemctl enable kdump.service
# Restart the kdump service.
sudo systemctl restart kdump.service
```

• Run the following commands in sequence to return the memory address space reserved by the kdump service to the operating system and disable the kdump service:

```
# Change the configurations in the /sys/kernel/kexec_crash_size file.
sudo sh -c 'echo 0 > /sys/kernel/kexec_crash_size'
# Disable the kdump service.
sudo systemctl disable kdump.service
# Stop the kdump service.
sudo systemctl stop kdump.service
```

**? Note** After the memory address space reserved by the kdump service is returned to the operating system, you must restart the operating system before you can re-enable the kdump service.

# Obtain the Debuginfo package and the source code package

• Run the following commands in sequence to obtain the Debuginfo package:

```
# Install yum-utils.
sudo yum install -y yum-utils
# Install the Debuginfo package. In the command, packageName specifies the name of the so
ftware package that you want to install.
sudo debuginfo-install -y <packageName>
```

• Run the following commands in sequence to obtain the source code package:

```
# Install the source code.
sudo yum install -y alinux-release-source
# Install yum-utils.
sudo yum install -y yum-utils
# Install the source code package. In the command, sourcePackageName specifies the name o
f the software package that you want to install.
sudo yumdownloader --source <sourcePackageName>
```

# Use experimental software packages

Experimental software packages are provided by Alibaba Cloud, but are not fully tested. Alibaba Cloud does not guarantee the quality of these packages. Alibaba Cloud Linux 2 provides the following types of experimental packages:

Experimental software packages that serve regular purposes

```
Golang 1.12Golang 1.13
```

Run the following commands in sequence to install an experimental software package:

```
# Enable support for the YUM repository.
sudo yum install -y alinux-release-experimentals
# Install an experimental software package that serves regular purposes. In the command,
packageName specifies the name of the software package that you want to install.
sudo yum install -y <packageName>
```

• Development kits that support SCL plug-ins

```
    Development kit based on GCC-7.3.1 : devtoolset-7
    Development kit based on GCC-8.2.1 : devtoolset-8
    Development kit based on GCC-9.1.1 : devtoolset-9
```

Run the following commands in sequence to install an experimental software package:

```
# Install scl-utils.
sudo yum install -y scl-utils
# Enable support for the YUM repository.
sudo yum install -y alinux-release-experimentals
# Install the software packages that you need from the YUM repository. The following samp
le commands are run to install all development kits that support SCL plug-ins:
sudo yum install -y devtoolset-7-gcc devtoolset-7-gdb devtoolset-7-binutils devtoolset-7-
make
sudo yum install -y devtoolset-8-gcc devtoolset-8-gdb devtoolset-8-binutils devtoolset-8-
make
sudo yum install -y devtoolset-9-gcc devtoolset-9-gdb devtoolset-9-binutils devtoolset-9-
make
```

After the software packages are installed, you can use the later versions of GNU Compiler Collection (GCC) and related tools. Sample commands:

```
# Specify a repository name to view an existing SCL. In this example, the devtoolset-7 re
pository is used.
scl -1 devtoolset-7
# Run the related SCL software.
scl enable devtoolset-7 'gcc --version'
```

# 5.3.4. Install and enable curl of a later version

The default version of curl for Alibaba Cloud Linux 2 is 7.29.0. If your business requires a later version of curl, you can install and enable curl-7.61.1 by performing the operations described in this topic.

#### Context

To avoid incompatibility issues, Alibaba Cloud Linux 2 does not support curl upgrades within the system but provides a way to install and enable the version of curl-7.61.1.

#### **Procedure**

1. Run the following command to install the Alibaba Cloud experimentals repository:

```
sudo yum install -y alinux-release-experimentals
```

- 2. Run the following commands in sequence to install curl and its dependency:
  - i. Inst all curl.

```
sudo yum install -y httpd24-curl
```

ii. Install the curl dependency httpd24-libcurl-devel.

```
sudo yum install -y httpd24-libcurl-devel
```

3. Run the following command to enable curl:

```
source /opt/rh/httpd24/enable
```

4. Run the following command to check the version of curl:

```
curl --version
```

The command output shown in the following figure indicates that curl-7.61.1 is installed.

```
[root@test ~| # curl --version
curl 7.61.1 (x86_64-koji-linux-gnu) libcurl/7.61.1 NSS/3.44 zlib/1.2.7 libssh2/1.8.0 nghttp2/1.7.1
Release-Date: 2018-09-05
Protocols: dict file ftp ftps gopher http https imap imaps ldap ldaps pop3 pop3s rtsp scp sftp smb smbs smtp smtps telnet tftp
Features: AsynchDNS IPv6 Largefile GSS-API Kerberos SPNEGO NTLM NTLM_WB SSL libz HTTP2 UnixSockets HTTPS-proxy
```

# 5.4. Security

# 5.4.1. CIS Alibaba Cloud Linux benchmarks

Alibaba Cloud Linux security benchmarks have been certified by the Center for Internet Security (CIS). The CIS has published CIS Alibaba Cloud Linux benchmarks. This topic describes CIS Alibaba Cloud Linux benchmarks and how to download them.

#### Context

The CIS is a renowned nonprofit organization that has been operating as an online community for decades to safeguard against cyber threats. It works with public service sectors, businesses, and academia from around the world to develop security best practice solutions for cyber defense, called benchmarks. For more information, visit the official CIS website or see CIS in Wikipedia. The CIS has published benchmarks for most Linux distributions such as Alibaba Cloud Linux, CentOS, and Ubuntu. CIS benchmarks are well recognized as important security standards for cloud-based operating systems.

The CIS certification for Alibaba Cloud Linux 2 security benchmarks was completed on August 16, 2019. Then, the CIS published *CIS Aliyun Linux 2 Benchmark v1.0.0*. The CIS certification for Alibaba Cloud Linux 3 security benchmarks was completed on February 6, 2022. Then, the CIS published *CIS Alibaba Cloud Linux 3 Benchmark v1.0.0*. Alibaba Cloud Linux is the first CIS certified operating system in China.

#### Introduction to CIS Alibaba Cloud Linux benchmarks

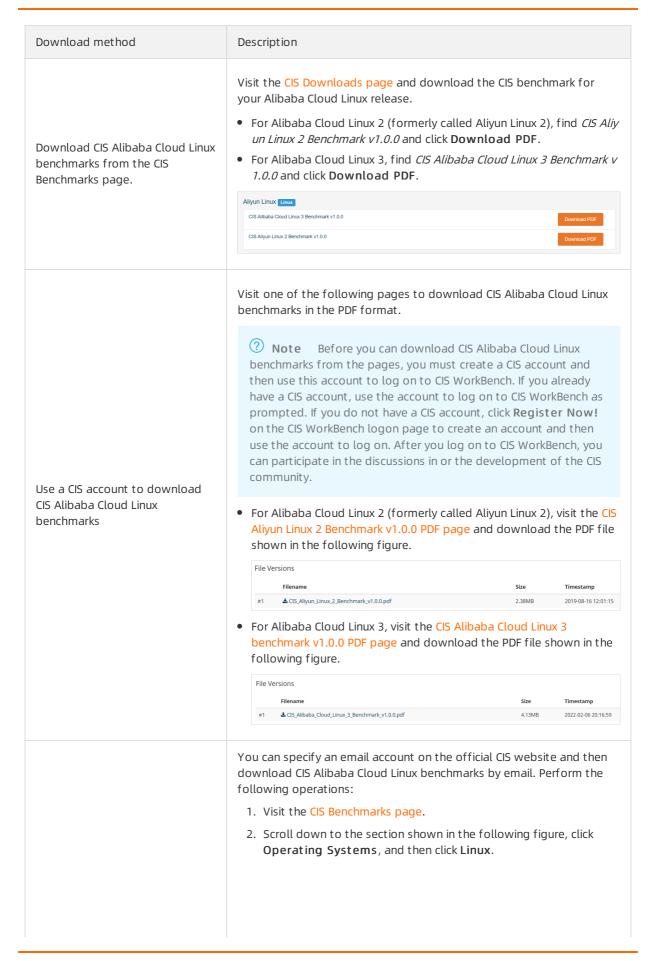
Take note of the following items about CIS Alibaba Cloud Linux benchmarks:

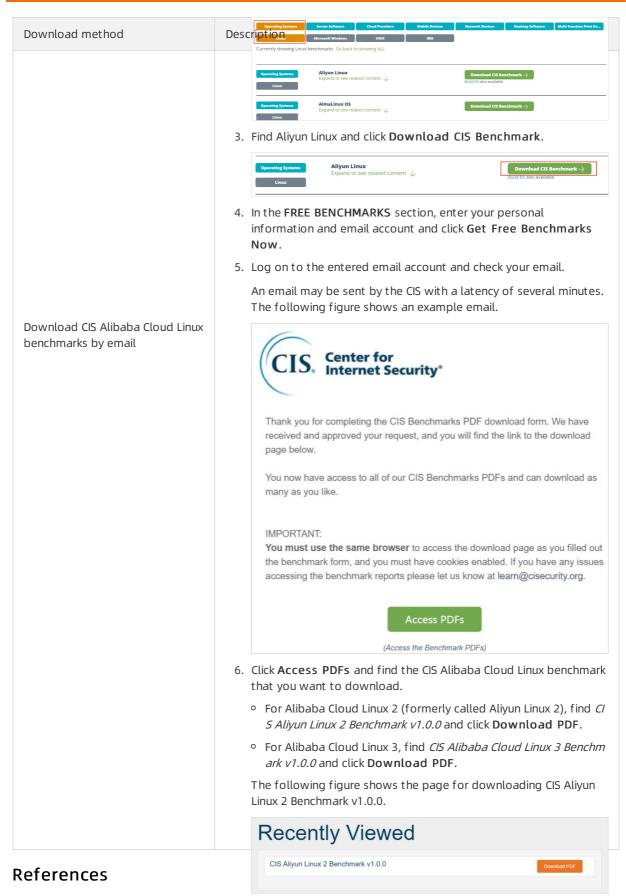
- The CIS benchmark for each Alibaba Cloud Linux release provides specific security recommendations for more than 200 items in the following areas: initial setup, services, network configuration, logging and auditing, access, authentication, authorization, and system maintenance.
  - o CIS Aliyun Linux 2 Benchmark v1.0.0 includes 204 items.
  - o CIS Alibaba Cloud Linux 3 Benchmark v1.0.0 includes 266 items.
- Each item is described from the following aspects: profile applicability, scoring information, description, rationale, impact, audit, remediation, and references. In CIS Alibaba Cloud Linux benchmarks, items are classified into different profile levels: Level 1 and Level 2.
  - Level 1 items cover basic security configurations that are easy to implement and have minimal impact on performance.
  - Level 2 it ems are suitable for high-security environments and may incur performance overheads.
- Items in CIS Aliyun Linux 2 Benchmark are categorized as Scored and Not Scored based on scoring information. Items in CIS Alibaba Cloud Linux 3 Benchmark are categorized as and Manual based on scoring information.
  - o Scored and Automated: If an item is categorized as Scored or Automated, the compliance of this item with the provided security recommendation affects the benchmark score. If this item complies with the provided security recommendation, the benchmark score is increased. If this item does not comply with the provided security recommendation, the benchmark score is decreased.
  - o Not Scored and Manual: If an item is categorized as Not Scored or Manual, the compliance of this item with the provided security recommendation does not affect the benchmark score. The benchmark score does not change with this item regardless of whether the item complies with the provided security recommendation.

Note When you use CIS Alibaba Cloud Linux benchmarks, you can select benchmark items based on your business requirements to improve the security of your Alibaba Cloud Linux operating systems. You can also refer to CIS Alibaba Cloud Linux benchmarks to improve the security of other compatible operating systems.

#### Download CIS Alibaba Cloud Linux benchmarks

The following table describes the methods that you can use to download CIS Alibaba Cloud Linux benchmarks.





The CIS provides instructions on how to use CIS Alibaba Cloud Linux benchmarks.

- CIS Aliyun Linux 2 Benchmark v1.0.0
- CIS Alibaba Cloud Linux 3 Benchmark v1.0.0

# 5.4.2. Subscribe to announcements for Alibaba Cloud Linux 2 CVEs

Fixes for Common Vulnerabilities and Exposures (CVEs) are continuously released until the end of life (EOL) of Alibaba Cloud Linux 2. You can use browser plug-ins to subscribe to announcements for Alibaba Cloud Linux 2 CVEs and check for CVE updates.

#### Context

This topic describes how to install the RSS Feed Reader plug-in to subscribe to updates for Alibaba Cloud Linux 2 CVEs. Google Chrome is used in the example. You can also select other browsers to subscribe to updates for Alibaba Cloud Linux 2 CVEs.

# Alibaba Cloud Linux 2 security advisories

To protect the system security of the Alibaba Cloud Linux 2 distribution, the Alibaba Cloud operating system team stays up to data on software issues and CVEs discovered by the Linux industry and community. The team promptly releases updates to YUM repositories to enhance security and fix software defects and security vulnerabilities. For more information, see Alibaba Cloud Linux 2.1903 Security Advisories.

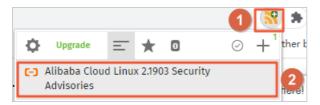
# (Optional) Install the RSS Feed Reader plug-in

If your Google Chrome is not installed with the RSS Feed Reader plug-in, install the plug-in first.

- 1. Open Google Chrome on your computer.
- 2. In the top bookmarks bar, click **Apps** and click **Web Store**.
- 3. In the upper-left corner of the chrome web store page, enter rss feed reader and search.
- 4. Click RSS Feed Reader to go to the details page.
- 5. On the details page of RSS Feed Reader, click Add to Chrome and click Add extension. When Add to Chrome changes to Remove from Chrome, the plug-in is added to extensions.

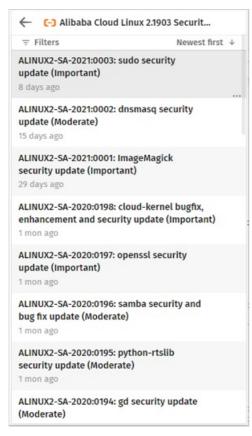
#### Subscribe to announcements for Alibaba Cloud Linux 2 CVEs

- 1. Subscribe to updates for Alibaba Cloud Linux 2 CVEs.
  - i. Access Alibaba Cloud Linux 2.1903 Security Advisories in Google Chrome.
  - ii. In the upper-right corner of the page, click the RSSI image.
  - iii. In the dialog box that appears, click Subscribe.When Subscribe changes to Following, you are subscribed to updates for Alibaba Cloud Linux 2 CVEs.
- 2. View updates for Alibaba Cloud Linux 2 CVEs by using RSS Feed Reader.



- i. In the upper-right corner of Google Chrome, click the RSS Feed Reader plug-in, which is tagged
   ① in the preceding figure.
- ii. Click **Alibaba Cloud Linux 2.1903 Security Advisories**, which is tagged ② in the preceding figure.

The update records are displayed as shown in the following figure.



# 5.4.3. Use YUM to perform security updates (Alibaba Cloud Linux 3)

This topic describes how to use YUM to query, check for, and install security updates for Alibaba Cloud Linux 3.

#### **Prerequisites**

An Elastic Compute Service (ECS) instance that runs Alibaba Cloud Linux 3 is created. For more information, see Creation method overview.

#### Background

To ensure system security, Alibaba Cloud Linux 3 stays up-to-date on Common Vulnerabilities and Exposures (CVE) through community-based, industry-supported efforts. Alibaba Cloud Linux 3 updates software packages including kernel packages, fixes software defects and security vulnerabilities, and enhances security features in a timely manner. For information about Alibaba Cloud Linux 3 security updates, see Alibaba Cloud Linux 3 Security Advisories.

Alibaba Cloud Linux 3 security updates are classified into the following severity levels based on the Common Vulnerability Scoring System 3 (CVSS3) for CVE:

- Critical: High-risk vulnerabilities exist and the security update is required.
- Important: Relatively high-risk vulnerabilities exist and the security update is strongly recommended.
- Moderate: Medium-risk vulnerabilities exist and the security update is recommended.
- Low: Low-risk vulnerabilities exist and the security update is optional.

# Query security updates

You can run the following command to query security updates:

```
yum updateinfo <command> [option]
```

The following table describes the variables of the command.

Variable	Valid value
<command/>	<ul> <li>list: queries the list of available security updates.</li> <li>info <update_id>: queries the details of a specific security update. Set <update_i d=""> to an advisory ID in Alibaba Cloud Linux 3 security advisories.</update_i></update_id></li> </ul>
	•sec-severity= <sevs> : specifies the security update severity levels. Set<sevs> to one or more security update severity levels.</sevs></sevs>
	Note If you want to specify multiple security update severity levels, separate them with commas (,). Security update severity levels are case-sensitive.
	Format description:
[option]	<ul> <li>Specify a single security update severity level in thesec-severity=<severity level="">         format. Example:sec-severity=Moderate .</severity></li> </ul>
	<ul> <li>Specify multiple security update severity levels in thesec-severity={<severity level="">,<severity level="">}. Example:sec-severity={Moderate, Important} .</severity></severity></li> </ul>
cve=< <i>CVE ID&gt;</i> : specifies the CVE IDs. Set <cve id=""> to one or more CVE can obtain CVE IDs from Alibaba Cloud Linux 3 security advisories.</cve>	

This section provides examples of the commands that you can run to query security updates.

- Run the yum updateinfo --help command to obtain the help information about the command.
- Run the following command to guery information about all available security updates:

```
yum updateinfo
```

#### A command output similar to the following one is returned:

```
Last metadata expiration check: 0:06:42 ago on Wed 02 Jun 2021 03:05:30 AM EDT.

Updates Information Summary: available

3 Security notice(s)

2 Important Security notice(s)

1 Moderate Security notice(s)
```

• Run the following command to guery the list of available security updates:

```
yum updateinfo list
```

A command output similar to the following one is returned:

```
Last metadata expiration check: 0:09:05 ago on Wed 02 Jun 2021 03:05:30 AM EDT.

ALINUX3-SA-2021:0008 Moderate/Sec. gnutls-3.6.14-7.1.al8.x86_64

ALINUX3-SA-2021:0029 Important/Sec. gnutls-3.6.14-8.1.al8.x86_64

ALINUX3-SA-2021:0028 Important/Sec. libldb-2.1.3-3.1.al8.x86_64

ALINUX3-SA-2021:0029 Important/Sec. nettle-3.4.1-4.1.al8.x86_64
```

• Run the yum updateinfo info <update\_id> command to query the details of a specific security update.

Example command with <update\_id> set to ALINUX3-SA-2021:0008 :

```
yum updateinfo info ALINUX3-SA-2021:0008
```

#### A command output similar to the following one is returned:

```
Last metadata expiration check: 0:11:58 ago on Wed 02 Jun 2021 03:05:30 AM EDT.

ALINUX3-SA-2021:0008: gnutls security and bug fix update (Moderate)

Update ID: ALINUX3-SA-2021:0008

Type: security

Updated: 1969-12-31 19:00:00

CVEs: CVE-2020-24659

Description: Package updates are available for Alibaba Cloud Linux 3 that fix the following vulnerabilities:

:
: CVE-2020-24659:
: An issue was discovered in GnuTLS before 3.6.15. A server can trigger a NULL pointer dereference in a TLS 1.3 client if a no_renegotiation alert is sent with unexpect ed timing, and then an invalid second handshake occurs. The crash happens in the applicat ion's error handling path, where the gnutls_deinit function is called after detecting a handshake failure.

: Severity: Moderate
```

• Run the following command to query the security updates of a specific severity level:

```
yum updateinfo list --sec-severity=Moderate
```

A command output similar to the following one is returned:

```
Last metadata expiration check: 0:05:25 ago on Mon 07 Jun 2021 09:08:25 AM EDT. ALINUX3-SA-2021:0008 Moderate/Sec. gnutls-3.6.14-7.1.al8.x86_64
```

# Check for security updates

You can run the yum check-update --security command to check for security updates available for the system. You can append --sec-severity=<SEVS> to the command to check for security updates of specific severity levels. Set <SEVS> to one or more severity levels.

**? Note** If you want to specify multiple severity levels, separate them with commas (,). Security update severity levels are case-sensitive.

The following examples demonstrate how to check for security updates:

• Example 1: Run the following command to check for all security updates:

```
yum check-update --security
```

A command output similar to the following one is returned:

```
Last metadata expiration check: 0:08:41 ago on Wed 02 Jun 2021 05:24:55 PM CST.
nss.x86 64
                              3.53.1-17.1.al8
                                                            alinux3-updates
nss-softokn.x86_64
nss-softokn-freebl.x86_64
                               3.53.1-17.1.al8
                                                             alinux3-updates
                             3.53.1-17.1.al8
                                                            alinux3-updates
nss-sysinit.x86 64
                              3.53.1-17.1.al8
                                                            alinux3-updates
nss-util.x86 64
                              3.53.1-17.1.al8
                                                            alinux3-updates
perl-Errno.x86 64
                               1.28-417.2.al8
                                                            alinux3-updates
perl-IO.x86 64
                               1.38-417.2.al8
                                                             alinux3-updates
```

• Example 2: Run the following command to check for security updates of the Critical and Important security levels:

```
yum check-update --security --sec-severity={Critical,Important}
```

A command output similar to the following one is returned:

```
Last metadata expiration check: 0:10:23 ago on Wed 02 Jun 2021 05:24:55 PM CST.
                              3.6.14-8.2.al8
gnutls.x86 64
                                                       alinux3-updates
nss.x86 64
                               3.53.1-17.1.al8
                                                       alinux3-updates
nss-softokn.x86_64
                                                       alinux3-updates
                              3.53.1-17.1.al8
nss-softokn-freebl.x86_64
                              3.53.1-17.1.al8
                                                       alinux3-updates
nss-sysinit.x86_64
                              3.53.1-17.1.al8
                                                       alinux3-updates
                               3.53.1-17.1.al8
nss-util.x86 64
                                                       alinux3-updates
nss-util.x86_64
perl-Errno.x86_64
                              1.28-417.2.al8
                                                       alinux3-updates
perl-IO.x86 64
                               1.38-417.2.al8
                                                        alinux3-updates
```

# Install security updates

You can run the yum upgrade command to install the security updates of specific severity levels or for specific CVEs.

• You can run the yum upgrade --security command to install security updates. You can append --sec-severity=<SEVS> to the command to install security updates of specific severity levels. Set < SEVS> to one or more severity levels.

**Note** If you want to specify multiple severity levels, separate them with commas (,). Security update severity levels are case-sensitive.

#### Example:

Run the following command to install security updates of the Critical and Important security levels:

```
yum upgrade --security --sec-severity={Critical,Important}
```

#### A command output similar to the following one is returned:

- You can run the yum upgrade --cve=<*CVE ID>* command to install security updates for specific CVEs. --cve=<*CVE ID>* specifies the CVE IDs. Set <*CVE ID>* to one or more CVE IDs.
  - **?** Note If you want to specify multiple CVE IDs, separate them with commas (,). CVE IDs are case-sensitive.

#### Example:

Run the following command to install the security update for CVE-2020-24659:

```
yum upgrade --cve=CVE-2020-24659
```

#### A command output similar to the following one is returned:

# 5.4.4. Use YUM to perform security updates (Alibaba Cloud Linux 2)

This topic describes how to use YUM to query, check for, and install security updates for Alibaba Cloud Linux 2.

## Background

To ensure system security, Alibaba Cloud Linux 2 stays up-to-date on Common Vulnerabilities and Exposures (CVEs) through a community-based, industry-supported efforts. Alibaba Cloud Linux 2 updates software packages including the kernel packages, and fixes software defects and security vulnerabilities in a timely manner. For information about Alibaba Cloud Linux 2 security updates, see Alibaba Cloud Linux 2.1903 Security Advisories.

Alibaba Cloud Linux 2 security updates are classified into the following severity levels based on the Common Vulnerability Scoring System 3 (CVSS3) for CVEs:

- Critical: High-risk vulnerabilities exist and the security update is required.
- Important: Relatively high-risk vulnerabilities exist and the security update is strongly recommended.
- Moderate: Medium-risk vulnerabilities exist and the security update is recommended.
- Low: Low-risk vulnerabilities exist and the security update is optional.

## Query security updates

You can run the following command to guery security updates:

```
yum updateinfo <command> [option]
```

The following table describes the variables of the command.

Variable	Valid value		
command	<ul> <li>list : queries the list of available security updates.</li> <li>info <update_id> : queries details about a specific security update. The value of <update_id> is an advisory ID in Alibaba Cloud Linux 2 security advisories.</update_id></update_id></li> </ul>		
option	•sec-severity= <sevs> orsecseverity=<sevs> : specifies the security update severity levels. <sevs> specifies one or more security update severity levels.</sevs></sevs></sevs>		
	Notice The values of <sevs> are case-sensitive.</sevs>		
	•cve= <cves> : specifies one or more CVE IDs. You can obtain CVE IDs from Alibaba Cloud Linux 2 security advisories.</cves>		

This section provides examples of the commands that you can run to query security updates.

- Run the yum updateinfo --help to obtain the help information about the command.
- Run the yum updateinfo command to query information about all available security updates. Sample command and command output:

```
# yum updateinfo
Loaded plugins: fastestmirror
Determining fastest mirrors
| 3.1 kB 00:00:00
extras
| 2.5 kB 00:00:00
plus
| 2.5 kB 00:00:00
updates
| 2.9 kB 00:00:00
(1/6): extras/2.1903/x86_64/primary_db
| 149 kB 00:00:00
(2/6): base/2.1903/x86_64/group_gz
| 101 kB 00:00:00
(3/6): updates/2.1903/x86_64/updateinfo
| 81 kB 00:00:00
(4/6): plus/2.1903/x86 64/primary db
| 1.5 MB 00:00:00
(5/6): base/2.1903/x86_64/primary_db
| 4.9 MB 00:00:00
(6/6): updates/2.1903/x86_64/primary_db
| 6.1 MB 00:00:00
Updates Information Summary: updates
   17 Security notice(s)
        7 Important Security notice(s)
        6 Moderate Security notice(s)
        4 Low Security notice(s)
updateinfo summary done
```

• Run the yum updateinfo list command to query the list of available security updates. Sample command and command output:

```
# yum updateinfo list
Loaded plugins: fastestmirror
Loading mirror speeds from cached hostfile
ALINUX2-SA-2019:0055 Moderate/Sec. binutils-2.27-41.base.1.al7.x86_64
ALINUX2-SA-2019:0058 Low/Sec. curl-7.29.0-54.1.al7.x86_64
ALINUX2-SA-2019:0059 Low/Sec. elfutils-default-yama-scope-0.176-2.1.al7.n
...
```

• Run the yum updateinfo info <update\_id> command to query the details about a specific security update. Sample command and command output:

```
# yum updateinfo info ALINUX2-SA-2020:0005
Loaded plugins: fastestmirror
Loading mirror speeds from cached hostfile
 ALINUX2-SA-2020:0005: nss, nss-softokn, nss-util security update (Important)
_____
 Update ID: ALINUX2-SA-2020:0005
   Release : Alibaba Cloud Linux 2.1903
     Type : security
    Status : stable
    Issued : 2020-01-03
      CVEs : CVE-2019-11729
           : CVE-2019-11745
Description: Package updates are available for Alibaba Cloud Linux 2.1903 that fix
          : the following vulnerabilities:
           : CVE-2019-11729:
           : Empty or malformed p256-ECDH public keys may
           : trigger a segmentation fault due values being
           : improperly sanitized before being copied into
           : memory and used. This vulnerability affects
           : Firefox ESR < 60.8, Firefox < 68, and Thunderbird
           : < 60.8.
           : CVE-2019-11745:
           : When encrypting with a block cipher, if a call to
           : NSC EncryptUpdate was made with data smaller than
           : the block size, a small out of bounds write could
           : occur. This could have caused heap corruption and
           : a potentially exploitable crash. This
           : vulnerability affects Thunderbird < 68.3, Firefox
           : ESR < 68.3, and Firefox < 71.
  Severity : Important
updateinfo info done
```

## Check for security updates

You can run the yum check-update --security command to check for security updates available for the system. You can append --secseverity=<SEVS> to the command to check for security updates of specific severity levels. Set <SEVS> to one or more severity levels.

Notice If you want to specify multiple security update severity levels, separate them with commas (,). The values of <SEVS> are case-sensitive.

The following examples demonstrate how to check for security updates:

• Example 1

```
# yum check-update --security |grep available
49 package(s) needed for security, out of 183 available
```

• Example 2

```
# yum check-update --security --secseverity=Critical,Important |grep available
30 package(s) needed for security, out of 183 available
```

## Install security updates

You can use the yum upgrade command to install security updates in one of the following ways:

• You can run the yum upgrade --security command to install security updates. You can append secseverity=<SEVS> to the command to install security updates of specific severity levels. Set <SE vs> to one or more severity levels.

Notice If you want to specify multiple security update severity levels, separate them with commas (,). The values of <SEVS> are case-sensitive.

#### Sample command and command output:

• You can run the yum upgrade -cves=<CVES> command to install security updates for specific CVEs. Set <CVES> to one or more CVE IDs.

Notice If you want to specify multiple CVE IDs, separate them with commas (,). The values of <CVES> are case-sensitive.

Sample command and command output:

```
# yum upgrade --cve=CVE-2019-11729, CVE-2019-11745
Loaded plugins: fastestmirror
Loading mirror speeds from cached hostfile
[snipped]
Dependencies Resolved
Package
                                                Arch
                                                                                    Vers
                                        Repository
Updating:
                                                                             854 k
.0-7.1.al7
                                        updates
 nss-softokn
                                               x86 64
                                                                                     3.44
                                                                              330 k
.0-8.1.a17
                                         updates
nss-softokn-freebl
                                               x86 64
                                                                                     3.44
.0-8.1.al7
                                                                              225 k
                                         updates
nss-sysinit
                                                x86 64
                                                                                     3.44
.0-7.1.al7
                                         updates
                                                                              65 k
                                                                                    3.44
nss-tools
                                               x86 64
.0-7.1.al7
                                                                              528 k
                                         updates
nss-util
                                                x86 64
                                                                                     3.44
.0-4.1.al7
                                                                              79 k
                                         updates
Updating for dependencies:
nspr
                                                x86_64
                                                                                     4.21
.0-1.1.al7
                                         updates
                                                                              127 k
Transaction Summary
Upgrade 6 Packages (+1 Dependent package)
Total download size: 2.2 M
Is this ok [y/d/N]:
```

Note The output of the man yum command shows that the yum upgrade command is equivalent to the yum update --obsoletes command. The yum upgrade command is also equivalent to the yum update command because obsoletes is enabled in the /etc/yum.conf configuration file by default.

# 5.4.5. Manage Alibaba Cloud Linux security update notifications

By default, the update-motd service is enabled in Alibaba Cloud Linux operating systems. If the system detects an available security update, a notification is displayed when you log on to an Elastic Compute Service (ECS) instance. This topic describes how to manage these security updates for Alibaba Cloud Linux.

#### Context

To ensure system security, Alibaba Cloud Linux stays up to date on Common Vulnerabilities and Exposures (CVEs) through a community-based, industry-supported effort. Alibaba Cloud Linux updates software packages (including the kernel), fixes software defects and security vulnerabilities, and enhances security features in a timely manner. For information about Alibaba Cloud Linux security updates, see Alibaba Cloud Linux 2.1903 Security Advisories and Alibaba Cloud Linux 3 Security Advisories.

To better notify you of security updates, the update-motd service is installed and enabled in Alibaba
Cloud Linux operating systems. The update-motd service periodically runs the yum check-update
command for Alibaba Cloud Linux 2 or the dnf updateinfo command for Alibaba Cloud Linux 3 to
check and obtain available security updates. When you log on to an instance by using a VNC
management terminal, the notification of security updates is displayed in the Message of the Day
(MOTD) of the operating system. Examples:

• Notification for Alibaba Cloud Linux 2:

```
3 package(s) needed for security, out of 3 available
Run "yum update" to apply all updates.
```

Notification for Alibaba Cloud Linux 3:

```
Updates Information Summary: available

8 Security notice(s)

1 Critical Security notice(s)

1 Important Security notice(s)

6 Moderate Security notice(s)

Run "dnf upgrade-minimal --security" to apply all updates.
```

Description of the previous examples:

- The notification for Alibaba Cloud Linux 2 indicates that three security vulnerabilities need to be addressed by security updates for the Alibaba Cloud Linux 2 operating system.
- The notification for Alibaba Cloud Linux 3 indicates that one critical security vulnerability, one important security vulnerability, and six moderate security vulnerabilities need to be addressed by security updates for the Alibaba Cloud Linux 3 operating system.

Alibaba Cloud Linux security updates are classified into the following severity levels based on the Common Vulnerability Scoring System 3 (CVSS3) for CVEs:

- Critical: High-risk vulnerabilities exist and the security update is strongly recommended.
- Important: Relatively high-risk vulnerabilities exist and the security update is strongly recommended.
- Moderate: Medium-risk vulnerabilities exist and the security update is recommended.
- Low: Low-risk vulnerabilities exist and the security update is optional.

For more information, see the "Check for security updates" section in Use YUM to perform security updates (Alibaba Cloud Linux 2) Or Use YUM to perform security updates (Alibaba Cloud Linux 3).

#### Recommended operations

You can update all or specific vulnerabilities based on your business requirements.

Operation	Description
-----------	-------------

Operation	Description
Full update	You can run the following command to install all security updates based on the operating system notification:  yum upgradesecurity
Partial update	You can install specific security updates based on your business requirements. For more information, see the "Install security updates" section in Use YUM to perform security updates (Alibaba Cloud Linux 2) or Use YUM to perform security updates (Alibaba Cloud Linux 3).

## Operations

You can run the systemctl command to manage the update-motd service.

```
Start the update-motd service.

systemctl start update-motd

Stop the update-motd service.

systemctl stop update-motd

Restart the update-motd service.

systemctl restart update-motd

Check the state of the update-motd service.

systemctl status update-motd service.
```

## 5.4.6. Use confidential containers based on Inclavare Containers in a Docker container

Inclavare Containers is the first open source container runtime intended for confidential computing in the industry. Inclavare Containers allows you to launch protected containers within a hardware-based trusted execution environment (TEE) to prevent untrusted entities such as untrusted cloud service providers (CSPs) from accessing sensitive data. You can deploy confidential containers based on Inclavare Containers within a Software Guard Extensions (SGX) encrypted computing environment. This topic describes how to deploy and use confidential containers based on Inclavare Containers in a Docker container.

#### **Prerequisites**

• A security-enhanced Elastic Compute Service (ECS) instance that uses an Alibaba Cloud Linux image is created.

When you create a security-enhanced instance, you must select a specific image to use the security features. For more information, see Create security-enhanced instances. For example, you can select the ecs.g7t instance type and the Alibaba Cloud Linux 3.2104 64-bit (UEFI) public image to create a security-enhanced instance.

• An SGX encrypted computing environment is created in the security-enhanced instance. For more information, see Build an SGX encrypted computing environment.

#### Context

You can use conventional methods to deploy and use confidential containers in a Docker container. However, you must have expertise in confidential computing and how to use Intel SGX SDKs to develop and create images. Inclavare Containers helps you streamline the processes and provides you with easy access to confidential computing. Inclavare Containers is compatible with different types of enclave runtimes and provides a consistent user experience across standard and confidential containers. For more information, visit inclavare-containers.

#### **Procedure**

To use a confidential container in a Docker container that runs on the Alibaba Cloud Linux operating system, perform the following steps:

- Step 1: Build a runtime environment to run a confidential container
- Step 2: Configure the OCI-compliant runtime rune for a Docker container
- Step 3: Create an image for the confidential container
- Step 4: Use the confidential container

## Step 1: Build a runtime environment to run a confidential container

1. Log on to the security-enhanced instance.

For more information, see Connection methodsGuidelines on instance connection.

2. Inst all rune.

rune is a command-line tool that conforms to the Open Container Initiative (OCI) runtime specification. rune is used to create and run enclaves. For more information, visit rune and runtimecontainer.

Run one of the following commands to install rune based on your Alibaba Cloud Linux operating system.

• If you are using the Alibaba Cloud Linux 2.1903 LTS 64-bit (UEFI) operating system, run the following command to install rune:

```
yum-config-manager --add-repo https://mirrors.openanolis.org/inclavare-containers/ali
nux2-repo && \
   rpm --import https://mirrors.openanolis.org/inclavare-containers/alinux2-repo/RPM-G
PG-KEY-rpm-sign && \
   yum install -y rune
```

• If you are using the Alibaba Cloud Linux 3.2104 64-bit (UEFI) operating system, run the following command to install rune:

```
yum-config-manager --add-repo https://mirrors.openanolis.org/inclavare-containers/ali
nux3-repo && \
   rpm --import https://mirrors.openanolis.org/inclavare-containers/alinux3-repo/RPM-G
PG-KEY-rpm-sign && \
   yum install -y rune
```

3. Run the following command to check whether rune is installed:

```
which rune
```

The following example command output indicates that rune is installed:

```
/usr/local/bin/rune
```

4. Run the following command to install the Occlum software stack.

Occlum is an enclave runtime supported by Inclavare Containers. Inclavare Containers must work with an enclave runtime to run confidential containers. For more information, visit occlum.

```
yum install -y occlum-pal
```

5. Run the following command to check whether Occlum is installed:

```
ls /opt/occlum/build/lib
```

The following example command output indicates that Occlum is installed:

```
libocclum-pal.so.0.21.0
```

## Step 2: Configure the OCI-compliant runtime rune for a Docker container

1. Install Docker Community Edition (Docker-CE).

Perform one of the following operations to install Docker-CE based on your Alibaba Cloud Linux operating system.

- Note You must install Docker-CE when you use the following methods to install Docker.
- To install Docker-CE on an instance that runs the Alibaba Cloud Linux 2.1903 LTS 64-bit (UEFI) operating system, perform the operations described in Deploy and use Docker on Alibaba Cloud Linux 2 instances.
- To install Docker-CE on an instance that runs the Alibaba Cloud Linux 3.2104 64-bit (UEFI) operating system, perform the operations described in Deploy and use Docker on Alibaba Cloud Linux 3 instances.
- 2. Configure the OCI-compliant runtime rune for the Docker container.
  - i. Run the following command to add the rune configurations to the Docker configuration file. Example Docker configuration file: /etc/docker/daemon.json.

```
mkdir -p /etc/docker && \
cat << EOF > /etc/docker/daemon.json
{
    "runtimes": {
        "rune": {
            "path": "/usr/local/bin/rune",
            "runtimeArgs": []
        }
    }
    EOF
```

ii. Run the following command to restart the Docker daemon:

```
systemctl restart docker
```

iii. Run the following command to check whether the rune container runtime is configured:

```
docker info grep rune
```

■ If you are using the Alibaba Cloud Linux 2.1903 LTS 64-bit (UEFI) operating system, the following example command output indicates that the rune container runtime is installed:

```
Runtimes: rune io.containerd.runc.v2 io.containerd.runtime.v1.linux runc
```

• If you are using the Alibaba Cloud Linux 3.2104 64-bit (UEFI) operating system, the following example command output indicates that the rune container runtime is installed:

```
Runtimes: io.containerd.runc.v2 io.containerd.runtime.v1.linux runc rune
```

## Step 3: Create an image for the confidential container

1. Run the following command to start the Occlum Docker image.

This operation can ensure that the subsequent steps to create an image for the confidential container are performed in the Occlum Docker image. For more information about the Occlum Docker image, see occlum 0.21.0.

```
mkdir "$HOME/rune_workdir" && \
docker run -it --privileged \
    -v /dev/sgx_enclave:/dev/sgx/enclave \
    -v /dev/sgx_provision:/dev/sgx/provision \
    -v "$HOME/rune_workdir":/root/rune_workdir \
    occlum/occlum:0.21.0-ubuntu18.04
```

2. Run the following command to write the Hello World code for testing:

```
cd /root/rune_workdir && \
cat << EOF > hello_world.c
#include <stdio.h>
#include <unistd.h>
void main(void)
{
    while (1) {
        printf("Hello World!\n");
        fflush(stdout);
        sleep(5);
    }
}
EOF
```

- 3. Create and package the Hello World trusted application.
  - i. Run the following command to compile the program by using the Occlum toolchain:

```
occlum-gcc -o hello_world hello_world.c
```

ii. Run the following command to initialize an Occlum instance:

```
occlum new occlum_instance
```

iii. Run the following commands to generate an Occlum FS image and an Occlum SGX enclave:

```
cd occlum_instance && \
   cp ../hello_world image/bin/ && \
   openssl genrsa -aes128 -out occlum_key.pem -3 3072 && \
   occlum build --sign-key occlum_key.pem
```

After you run the preceding commands, set a password based on the following message. The password is used to encrypt and protect the signature private key.

```
Enter pass phrase for occlum_key.pem:
Verifying - Enter pass phrase for occlum_key.pem:
```

iv. Run the following command to configure Production Enclave.

By default, Occlum generates Debug Enclave. You must run the following command to set Enclave to Production Enclave in the production environment. For more information about the differences between Debug Enclave and Production Enclave, see debugging-intel-sgx-enclaves-in-windows-737361.pdf.

```
sed -i 's/"debuggable": true/"debuggable": false/g' Occlum.json
```

After you run the preceding command to modify the Occlum.json configuration file, you must run the following command to generate Production Enclave:

```
occlum build --sign-key occlum_key.pem
```

v. Run the following command to package the Hello World trusted application and copy the trusted application to the *rune\_workdir* directory:

```
occlum package occlum_instance.tar.gz && \
   cp occlum_instance.tar.gz /root/rune_workdir
```

- 4. Create a confidential container image.
  - i. Run the following command to exit the Occlum Docker image:

```
exit
```

ii. Run the following command to write Dockerfile:

```
cd "$HOME/rune_workdir" && \
cat << EOF >Dockerfile
FROM scratch
ADD occlum_instance.tar.gz /
ENTRYPOINT ["/bin/hello_world"]
EOF
```

iii. Run the following command to create a confidential container image:

```
docker build . -t occlum-app
```

## Step 4: Use the confidential container

Run the following command to run the confidential container and check the operation logs of the confidential container:

```
docker run -it --rm --runtime=rune \
   -e ENCLAVE_TYPE=intelSgx \
   -e ENCLAVE_RUNTIME_PATH=/opt/occlum/build/lib/libocclum-pal.so.0.21.0 \
   -e ENCLAVE_RUNTIME_ARGS=occlum_instance \
   -e ENCLAVE_RUNTIME_LOGLEVEL="off" \
   -e OCCLUM_RELEASE_ENCLAVE=1 \
   occlum-app
```

Parameters in the preceding command:

- ENCLAVE\_TYPE: specifies the hardware type of the enclave. In this example, this parameter is set to intelSgx.
- ENCLAVE\_PATH: specifies the path to start the enclave runtime PAL.
- ENCLAVE\_ARGS: specifies the specific parameters of the enclave runtime PAL. Separate multiple parameters with commas (,).
- ENCLAVE\_RUNTIME\_LOGLEVEL: specifies the log level of the enclave runtime.
- OCCLUM\_RELEASE\_ENCLAVE: specifies the enclave type. When this parameter is set to 0, the enclaves of the Debug Enclave type are queried. When this parameter is set to 1, the enclaves of the Production Enclave type are queried.

A command output similar to the following one is returned. If the Inclavare Containers environment is installed and the confidential container runs normally, <code>Hello World!</code> is displayed in the container log every 5 seconds.

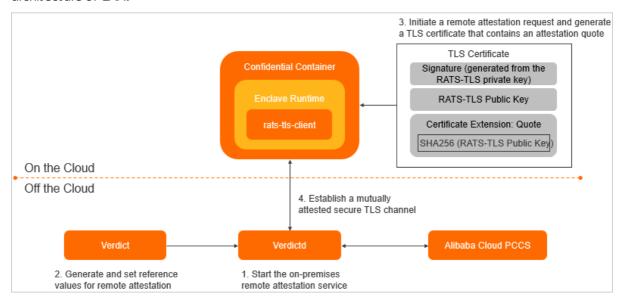
```
Hello World!
Hello World!
Hello World!
```

## 5.4.7. Use Inclavare Containers to implement remote attestation

Inclavare Containers implements a universal and cross-platform remote attestation architecture named Enclave Attestation Architecture (EAA). EAA can prove that sensitive workloads run in a hardware-based trusted execution environment (TEE). This topic describes how to use Inclavare Containers to implement remote attestation.

#### Context

EAA uses a Transport Layer Security (TLS) certificate that contains a quote of a hardware-based TEE as the root of trust. This ensures that the rats-tls-client remote attestation client and the on-premises Verdictd service communicate in a hardware-based TEE. The following figure shows the workflow and architecture of EAA.



To check whether the workloads of confidential containers run in a hardware-based TEE, you can start rats-tls-client in the cloud to send a request to Verdictd for verification. Procedure:

- 1. Start Verdictd in the trusted environment on the user side. Verdictd is an on-premises verifier.
- 2. Verdict generates and sets reference values for remote attestation, and uploads the values to the Verdict runtime environment.
- 3. rats-tls-client initiates a remote attestation request to Verdictd for verification and sends the information of the hardware-based TEE and sensitive data from a confidential container to Verdictd.
- 4. Verdictd compares the metric values against the reference values to determine whether workloads in the cloud are running in the expected trusted environment. Then, the entire remote attestation process is completed.

The following table describes the components involved in the process.

Component	Role	Description
Confidential container	Attester	Confidential containers are used to run rats-tls- client in the enclave runtime. The client obtains the metric values of the programs running in the HW-TEE (confidential container) and initiates remote attestation requests to Verdictd by using RATS-TLS. For more information, visit rats-tls.
Verdictd	On-premises verifier	Verdictd is a service that runs in the trusted environment on the user side. Verdictd uses Alibaba Cloud Provisioning Certificate Caching Service (PCCS) and Verdict to check the metric values of the programs running in the HW-TEE to complete the entire remote attestation process. For more information, visit Verdictd.

Component	Role	Description
Verdict	On-premises reference value provider	Verdict is a program that runs in the trusted environment on the user side. It is also the configuration program for Verdictd. It is used to configure Verdictd-approved reference values for the metric values of a program running in the HW-TEE to determine whether the program runs as expected or that the program has not been tampered with. For more information, visit Verdict.
Alibaba Cloud PCCS	Remote attestation service	The Alibaba Cloud Software Guard Extensions (SGX) remote attestation service is fully compatible with the remote attestation service for Intel® SGX Elliptic Curve Digital Signature Algorithm (ECDSA) and Intel® SGX SDK. Therefore, the vSGX instances provided by Alibaba Cloud, which are instances of the g7t, c7t, and r7t instance families, can gain trust from remote providers or producers through remote attestation. For more information, see Intel® SGX ECDSA Remote Attestation Service and Intel® SGX SDK.

In this topic, the following resources are used:

- An Elastic Compute Service (ECS) security-enhanced instance that uses an Alibaba Cloud Linux (UEFI) public image simulates rats-tls-client in the cloud.
- An ECS instance that uses a CentOS 8.2 public image simulates the on-promises host of Verdictd.

#### Precautions:

- Alibaba Cloud Linux 2.1903 LTS 64-bit (UEFI) and Alibaba Cloud Linux 3.2104 64-bit (UEFI) public images support remote attestation by using Inclavare Containers. Specific commands may vary based on the versions of operating systems.
- The Alibaba Cloud SGX remote attestation service is supported only in regions inside the Chinese mainland. For more information about regions, see Regions and zones.

#### **Procedure**

Perform the following steps to use Inclavare Containers to implement remote attestation:

- Step 1: Make preparations
- Step 2: Install a confidential container
- Step 3: Install and start Verdictd
- Step 4: Configure reference values for remote attestation
- Step 5: Verify whether remote attestation is implemented

## Step 1: Make preparations

1. Create a security-enhanced instance that uses an Alibaba Cloud Linux (UEFI) public image.

The instance simulates rats-tls-client in the cloud. When you create a security-enhanced instance, you must select an image that supports security features. For more information about how to create a security-enhanced instance, see Create security-enhanced instances.

Note When you create an instance, you must select a virtual private cloud (VPC) and a vSwitch for the instance and assign a public IP address to enable Internet access for the instance.

2. Build an SGX encrypted computing environment in the security-enhanced instance.

For more information, see Build an SGX encrypted computing environment.

3. Create an instance that uses a CentOS 8.2 public image.

The instance simulates the on-promises host of Verdictd. For information about how to create an ECS instance, see Create an instance by using the wizard. For more information about Verdictd, visit Verdictd.

- Note When you create an instance, you must select a VPC and a vSwitch for the instance and assign a public IP address to enable Internet access for the instance.
- 4. Add a rule to the security group of the CentOS 8.2 instance to allow inbound traffic on port 1111. For more information, see Add a security group rule.

## Step 2: Install a confidential container

- 1. Connect to the security-enhanced instance that uses an Alibaba Cloud Linux (UEFI) public image. For more information, see Connection methodsGuidelines on instance connection.
- 2. Run the following command to install the rats-tls-sgx confidential container.

The rats-tls-sgx installation package provides sample programs and library files for establishing secure communication channels based on the hardware-based TEE. For more information, visit rats-tls.

• Run the following command if the instance runs the Alibaba Cloud Linux 3.2104 64-bit (UEFI) operating system:

```
yum-config-manager --add-repo https://mirrors.openanolis.org/inclavare-containers/ali
nux3-repo && \
   rpm --import https://mirrors.openanolis.org/inclavare-containers/alinux3-repo/RPM-G
PG-KEY-rpm-sign && \
   yum install -y rats-tls-sgx
```

• Run the following command if the instance runs the Alibaba Cloud Linux 2.1903 LTS 64-bit (UEFI) operating system:

```
yum-config-manager --add-repo https://mirrors.openanolis.org/inclavare-containers/ali
nux2-repo && \
   rpm --import https://mirrors.openanolis.org/inclavare-containers/alinux2-repo/RPM-G
PG-KEY-rpm-sign && \
   yum install -y rats-tls-sgx
```

3. Run the following command to check whether rats-tls-sgx is installed:

```
ls /usr/share/rats-tls/samples
```

The following command output indicates that rats-tls-sgx is installed:

```
rats-tls-client rats-tls-server sgx_stub_enclave.signed.so
```

### Step 3: Install and start Verdictd

- Connect to the instance that uses a CentOS 8.2 public image.
   For more information, see Connection methodsGuidelines on instance connection.
- 2. Change the repository address of CentOS 8.
- 3. Run the following command to install the SGX Platform Software (PSW).

Verdictd must rely on the dynamic library provided by the SGX PSW to verify the TLS certificate that contains SGX information. For more information, see SGX PSW.

```
yum install -y yum-utils && \
    wget -c https://download.01.org/intel-sgx/sgx-linux/2.13/distro/centos8.2-server/sg
x_rpm_local_repo.tgz && \
    tar xzf sgx_rpm_local_repo.tgz && \
    yum-config-manager --add-repo sgx_rpm_local_repo && \
    yum makecache && rm -f sgx_rpm_local_repo.tgz && \
    yum install --nogpgcheck -y libsgx-dcap-quote-verify \
    libsgx-dcap-default-qpl libsgx-dcap-ql \
    libsgx-uae-service
```

4. Configure the public endpoint of Alibaba Cloud PCCS.

The Alibaba Cloud SGX remote attestation service is regionally deployed. For optimal stability, we recommend that you access this service in the region where the vSGX instance is deployed. You must manually modify the <code>/etc/sgx\_default\_qcnl.conf</code> file to adapt to the Alibaba Cloud SGX remote attestation service that is deployed in the region where the vSGX instance is deployed.

- **Note** The Alibaba Cloud SGX remote attestation service is supported only in regions inside the Chinese mainland. For more information about regions, see Regions and zones.
- If a public IP address is assigned to the vSGX instance, you must modify /etc/sgx\_default\_qcnl.c onf as described in the following code. You can use the vi or vim editor to modify the file. The example command for using the vi editor is vi /etc/sgx\_default\_qcnl.conf . The example command for using the vim editor is vim /etc/sgx\_default\_qcnl.conf .

```
# PCCS server address
PCCS_URL=https://sgx-dcap-server.<Region-ID>.aliyuncs.com/sgx/certification/v3/
# To accept insecure HTTPS cert, set this option to FALSE
USE_SECURE_CERT=TRUE
```

<Region-ID> is a variable. You must replace it with the ID of the region where the vSGX instance is deployed. For example, if the instance is deployed in the China (Hangzhou) region, you must replace <Region-ID> with cn-hangzhou .

If the vSGX instance resides within a VPC and has only an internal IP address, you must modify /et c/sgx\_default\_qcnl.conf as described in the following code. You can use the vi or vim editor to modify the file. The example command for using the vi editor is vi /etc/sgx\_default\_qcnl.conf

```
. The example command for using the vim editor is  \mbox{ vim /etc/sgx\_default\_qcnl.conf }.
```

```
# PCCS server address
PCCS_URL=https://sgx-dcap-server-vpc.<Region-ID>.aliyuncs.com/sgx/certification/v3/
# To accept insecure HTTPS cert, set this option to FALSE
USE_SECURE_CERT=TRUE
```

<Region-ID> is a variable. You must replace it with the ID of the region where the vSGX instance is deployed. For example, if the instance is deployed in the China (Hangzhou) region, you must replace <Region-ID> with cn-hangzhou .

5. Run the following command to install the Verdictd software stack.

Verdict d must rely on the library files of rats-tls-host to verify the TLS certificate that contains SGX information.

```
yum-config-manager --add-repo https://mirrors.openanolis.cn/inclavare-containers/rpm-re
po/ && \
    rpm --import https://mirrors.openanolis.cn/inclavare-containers/rpm-repo/RPM-GPG-KE
Y-rpm-sign && \
    yum install -y rats-tls-host verdictd
```

6. Run the following command to check whether Verdictd is installed:

```
which verdictd verdict
```

The following command output indicates that Verdictd is installed:

```
/usr/local/bin/verdictd
/usr/local/bin/verdict
```

7. Run the following command to check whether rats-tls-host is installed:

```
ls /usr/share/rats-tls/samples
```

The following command output indicates that rats-tls-host is installed:

```
rats-tls-client rats-tls-server
```

8. Run the following command to start Verdictd:

```
\label{lem:condition} \begin{tabular}{ll} verdictd --client-api 127.0.0.1:10001 --listen 0.0.0.0:1111 --mutual --attester nullatt ester --verifier sgx\_ecdsa --tls openssl --crypto openssl \\ \end{tabular}
```

Note After Verdictd is started, Verdictd listens on 0.0.0.0:1111 for remote attestation requests from rats-tls-client and listens on 127.0.0.1:10001 for requests to configure policies from Verdict.

In the subsequent steps, you must make sure that Verdictd continues to run and the ECS instance remains connected.

## Step 4: Configure reference values for remote attestation

1. Create a terminal to connect to the ECS instance that uses a Cent OS 8.2 public image.

For more information, see Connection methodsGuidelines on instance connection.

- Note When you create a terminal, you must make sure that the terminal that has Verdictd started is enabled or connected. Two terminals are connected to the ECS instance that uses a CentOS 8.2 public image.
- 2. Use Verdict to configure the reference values of SGX OPA for Verdictd so that only specific SGX enclaves can be started.

Open Policy Agent (OPA) is an open source, general-purpose policy engine that enables unified policy enforcement across the entire stack. For more information, visit OPA.

i. Run the following command to export the reference value file of SGX OPA:

```
verdict --client-api 127.0.0.1:10001 --export-opa-reference sgxData
```

The sgxData file contains the following content:

```
cat sgxData
{
    "mrEnclave": [],
    "mrSigner": [],
    "productId": 0,
    "svn": 0
}
```

Parameters in the sgxData reference value file:

- mrEnclave: The mrEnclave value of the SGX enclave is added to the mrEnclave reference value array.
- mrSigner: The mrSigner value of the SGX enclave is added to the mrSigner reference value array.
- productId: This parameter ensures that the productId value of the SGX enclave is greater than or equal to the reference value of productId.
- svn: This parameter ensures that the security version number (SVN) of the SGX enclave is greater than or equal to the reference value of svn.
- ii. Obtain the reference values of SGX OPA.

In actual business environments, you are the one who deploys applications in the cloud. Therefore, you must obtain reference values of the applications based on your trusted applications. The following steps demonstrate how to obtain the reference values of SGX OPA. In the example, the reference values of sgx\_stub\_enclave.signed.so deployed on the cloud are obtained.

- a. Run the following command to download the application deployed on the cloud.
  - Run the following command if the instance runs the Alibaba Cloud Linux 3.2104 64-bit (UEFI) operating system:

```
wget https://mirrors.openanolis.cn/inclavare-containers/alinux3-repo/rats-tls
-sgx-0.6.5-1.al8.x86_64.rpm
```

Run the following command if the instance runs the Alibaba Cloud Linux 2.1903 LTS 64-bit (UEFI) operating system:

```
\label{lem:wget} $$ wget $ https://mirrors.openanolis.cn/inclavare-containers/alinux2-repo/rats-tls -sgx-0.6.5-1.al7.x86_64.rpm $$
```

- b. Run the following command to decompress the downloaded RPM package.
  - Run the following command if the instance runs the Alibaba Cloud Linux 3.2104 64-bit (UEFI) operating system:

```
rpm2cpio rats-tls-sgx-0.6.5-1.al8.x86_64.rpm cpio -div
```

■ Run the following command if the instance runs the Alibaba Cloud Linux 2.1903 LTS 64-bit (UEFI) operating system:

```
rpm2cpio rats-tls-sgx-0.6.5-1.al7.x86_64.rpm cpio -div
```

c. Run the following command to download and install sgx\_sign:

```
export SGX_SDK_VERSION=2.14 && \
    export SGX_SDK_RELEASE_NUMBER=2.14.100.2 && \
    wget -c https://download.01.org/intel-sgx/sgx-linux/$SGX_SDK_VERSION/distro
/centos8.2-server/sgx_linux_x64_sdk_$SGX_SDK_RELEASE_NUMBER.bin && \
    chmod +x sgx_linux_x64_sdk_$SGX_SDK_RELEASE_NUMBER.bin && \
    echo -e 'no\n/opt/intel\n' ./sgx_linux_x64_sdk_$SGX_SDK_RELEASE_NUMBER.bin
&& \
    rm -f sgx_linux_x64_sdk_$SGX_SDK_RELEASE_NUMBER.bin
```

d. Run the following command to use sgx\_sign to dump the metadata file of the enclave and the enclave signature structure (SIGSTRUCT):

```
/opt/intel/sgxsdk/bin/x64/sgx_sign dump -enclave usr/share/rats-tls/samples/sgx _stub_enclave.signed.so-dumpfile metadata.txt -cssfile sigstruct.bin
```

e. Run the following command to obtain the Base64-encoded mrEnclave value:

```
dd skip=960 count=32 if=sigstruct.bin of=mrenclave.binary bs=1 && \ cat mrenclave.binary base64
```

■ A command output similar to the following one is returned if the instance runs the Alibaba Cloud Linux 3.2104 64-bit (UEFI) operating system:

```
zqs/chEBBNrInHlSCDO+0eV/pnlAB5qieF3M0hKfekA=
```

■ A command output similar to the following one is returned if the instance runs the Alibaba Cloud Linux 2.1903 LTS 64-bit (UEFI) operating system:

```
+PLc9QPrEZ6ad9e3BgihgOINOVPYvIj0v/ixt2Kbwyw=
```

f. Run the following command to obtain the Base64-encoded mrSigner value:

```
grep -A 2 -i "mrsigner->value" metadata.txt awk -F ":" 'NR==2,NR==3 {print $1}
' xargs echo -n sed 's/[[:space:]]//g;s/0x//g' xxd -r -p base64

g9cZ533qyhRw9rr2Kk13QwPImdtpAg+ccO4d/AjHzp4=
```

g. Run the following command to obtain the productId value:

```
dd skip=1024 count=2 if=sigstruct.bin of=productId.binary bs=1 && \
    od -An -tx2 productId.binary
```

A command output similar to the following one is returned for rats-tls-sgx-0.6.5-1.al8.x86\_64.rpm:

```
0000
```

h. Run the following command to obtain the svn value:

```
dd skip=1026 count=2 if=sigstruct.bin of=svn.binary bs=1 && \
    od -An -tx2 svn.binary
```

A command output similar to the following one is returned for rats-tls-sgx-0.6.5-1.al8.x86\_64.rpm:

```
0000
```

iii. Run the following command to generate an evidence file based on the reference values of SGX OPA.

Generate an evidence file based on the reference values of SGX OPA. The evidence file is a JSON file that contains the information of stub\_enclave evidence. It is used to test whether the updated sqxData reference value file can work normally.

Run the following command if the instance runs the Alibaba Cloud Linux 3.2104 64-bit (UEFI) operating system:

```
cat <<- EOF >evidence
{
    "mrEnclave": "zqs/chEBBNrInHlSCDO+0eV/pnlAB5qieF3M0hKfekA=",
    "mrSigner": "g9cZ533qyhRw9rr2Kk13QwPImdtpAg+ccO4d/AjHzp4=",
    "productId": 0,
    "svn": 0
}
EOF
```

■ Run the following command if the instance runs the Alibaba Cloud Linux 2.1903 LTS 64-bit (UEFI) operating system:

```
cat <<- EOF >evidence
{
    "mrEnclave": "+PLc9QPrEZ6ad9e3BgihgOINOVPYvIj0v/ixt2Kbwyw=",
    "mrSigner": "g9cZ533qyhRw9rr2Kk13QwPImdtpAg+ccO4d/AjHzp4=",
    "productId": 0,
    "svn": 0
}
EOF
```

- iv. Run the following command to update the local reference value file of SGX OPA.

  Update the *sgxData* file based on the obtained reference values of SGX OPA.
  - Run the following command if the instance runs the Alibaba Cloud Linux 3.2104 64-bit (UEFI) operating system:

```
cat <<- EOF >./sgxData
{
    "mrEnclave": ["zqs/chEBBNrInHlSCDO+0eV/pnlAB5qieF3M0hKfekA="],
    "mrSigner": ["g9cZ533qyhRw9rr2Kk13QwPImdtpAg+ccO4d/AjHzp4="],
    "productId": 0,
    "svn": 0
}
EOF
```

■ Run the following command if the instance runs the Alibaba Cloud Linux 2.1903 LTS 64-bit (UEFI) operating system:

```
cat <<- EOF >./sgxData
{
    "mrEnclave": ["+PLc9QPrEZ6ad9e3BgihgOINOVPYvIj0v/ixt2Kbwyw="],
    "mrSigner": ["g9cZ533qyhRw9rr2Kk13QwPImdtpAg+ccO4d/AjHzp4="],
    "productId": 0,
    "svn": 0
}
EOF
```

v. (Optional)Run the following command to test the updated local reference value file.

If you want to test the updated local reference value file, such as the format of the new reference value file and the added reference values, perform this step.

```
verdict --client-api 127.0.0.1:10001 --test-opa-local-reference sgxPolicy.rego ./sg
xData ./evidence
```

#### Parameters in the command:

- --test-opa-local-reference : the local reference value file.
- sgxPolicy.rego : the name of the OPA policy file used for the test. This parameter must be set to sgxPolicy.rego in the command.
- ./sgxData : the file to be tested, which is the updated reference value file.
- ./evidence : a JSON file generated by running a shell script that contains the stub\_encla
   ve evidence information. This file is used to test whether the updated sgxData reference
   value file can work normally.

If the test succeeds, a command output similar to the following one is returned.

■ Alibaba Cloud Linux 3.2104 64-bit (UEFI):

```
[root@test ~]# verdict --client-api 127.0.0.1:10001 --test-opa-local-reference sgxPolicy.rego ./sgxData ./evidence [2022-02-25707:23:16Z INFO verdict] Connect to Verdictd with addr: 127.0.0.1:10001 [2022-02-25707:23:16Z INFO verdict::opa OPA Test local reference: policy name: sgxPolicy.rego, reference file:./s gxData, input file:./evidence [2022-02-25707:23:16Z INFO verdict::opa TestOpa status is: "{\"allow\":true,\"parseInfo\":{\"mrEnclave\":[\"zqs/chEBBNrInHISCDO+0eV/pnlAB5qieF3M0hKfekA=\"]],\"mrSigner\":[\"g9cZ5 33qyhRW9rr2Kk13QwPImdtpAg+ccO4d/AjHzp4=\",\"g9cZ533qyhRW9rr2Kk13QwPImdtpAg+ccO4d/AjHzp4=\",\"g9cZ533qyhRW9rr2Kk13QwPImdtpAg+ccO4d/AjHzp4=\"]],\"productId\":[0,0],\"svn\":[0,0]}}"
```

■ Alibaba Cloud Linux 2.1903 LTS 64-bit (UEFI):

```
tverdict --client-api 127.0.0.1:10001 --test-opa-l
ocal-reference sgxPolicy.rego ./sgxData ./evidence
|2022-03-04T12:18:57Z INFO verdict| Connect to Verdictd with addr: 127.0.0.1:10001
|2022-03-04T12:18:57Z INFO verdict::opa| OPA Test local reference: policy name: sg
xPolicy.rego, reference file:./sgxData, input file:./evidence
|2022-03-04T12:18:57Z INFO verdict::opa| TestOpa status is: "{\"allow\":true,\"par
seInfo\":{\"mrEnclave\":[\"+PLc9QPrEZ6ad9e3BgihgOINOVPYvIj0v/ixt2Kbwyw=\",[\"+PLc9Q
PrEZ6ad9e3BgihgOINOVPYvIj0v/ixt2Kbwyw=\"]],\"mrSigner\":[\"g9cZ533qyhRw9rr2Kk13QwPI
mdtpAg+cc04d/AjHzp4=\",[\"g9cZ533qyhRw9rr2Kk13QwPImdtpAg+cc04d/AjHzp4=\"]],\"produc
tId\":[0,0],\"svn\":[0,0]}}"
```

vi. Run the following command to upload the reference value file.

After the reference value file is updated, you can upload the latest local reference value file to Verdictd.

```
verdict --client-api 127.0.0.1:10001 --set-opa-reference sgxData ./sgxData
```

A command output similar to the following one is returned.

```
[root@test ~]# verdict --client-api 127.0.0.1:10001 --set-opa-reference sgxData ./sgxData [2022-02-25T07:24:05Z INFO verdict] Connect to Verdictd with addr: 127.0.0.1:10001 [2022-02-25T07:24:05Z INFO verdict::opa] set_opa_reference status is: "OK"
```

You can run the **cat /opt/verdictd/opa/sgxData** command to check whether the reference value file is updated. Sample command output:

■ Alibaba Cloud Linux 3.2104 64-bit (UEFI):

```
cat /opt/verdictd/opa/sgxData
{
    "mrEnclave": ["zqs/chEBBNrInHlSCDO+0eV/pnlAB5qieF3M0hKfekA="],
    "mrSigner": ["g9cZ533qyhRw9rr2Kk13QwPImdtpAg+ccO4d/AjHzp4="],
    "productId": 0,
    "svn": 0
}
```

■ Alibaba Cloud Linux 2.1903 LTS 64-bit (UEFI):

```
cat /opt/verdictd/opa/sgxData
{
    "mrEnclave": ["+PLc9QPrEZ6ad9e3BgihgOINOVPYvIj0v/ixt2Kbwyw="],
    "mrSigner": ["g9cZ533qyhRw9rr2Kk13QwPImdtpAg+ccO4d/AjHzp4="],
    "productId": 0,
    "svn": 0
}
```

## Step 5: Verify whether remote attestation is implemented

1. Connect to the security-enhanced instance that uses an Alibaba Cloud Linux (UEFI) public image. For more information, see Connection methodsGuidelines on instance connection.

```
? Note You must make sure that the terminal that has Verdictd started is enabled or connected.
```

2. Run the following command to allow rats-tls-client to initiate a remote attestation request to Verdictd to check whether the workloads in the cloud are running in a hardware-based TEE:

```
cd /usr/share/rats-tls/samples/ && \
./rats-tls-client -a sgx_ecdsa -v nullverifier -t openssl -c openssl -i <Public IP addr
ess of the instance on which Verdictd is deployed> -p 1111 --mutual --verdictd
```

<Public IP address of the instance on which Verdictd is deployed> is a variable. Replace it with the public IP address of the instance on which Verdictd is deployed.

The following command output indicates that remote attestation is implemented.

## 5.5. Kernel Live Patching

## 5.5.1. Overview of the Kernel Live Patching feature

Alibaba Cloud Linux provides the Kernel Live Patching (KLP) feature for fixing the common vulnerabilities and exposures (CVEs) and critical bugs of a kernel. You can update hot fixes for the operating system kernel without restarting the server. This ensures the stability and security of the kernel. This topic describes the KLP feature and its benefits and limits.

#### Introduction

You can use KLP to update hotfixes for the kernel of Alibaba Cloud Linux in a timely manner. KLP consists of the following components:

- RPM package: an RPM package that contains the kernel module (.ko file) and description file of a hotfix. The kernel loads the kernel module in the hotfix to fix kernel bugs.
- kpatch utility: a command-line utility used to manage kernel modules in hotfixes.
- kpatch service: a KLP systemd service. This service loads the kernel module of each hotfix during the operating system initialization and is used to fix kernel bugs.

#### **Benefits**

KLP can update hotfixes for CVEs or critical bugs of a kernel in a smooth and quick manner without compromising server security and stability. You do not need to restart servers or other business-related task processes, wait until time-consuming tasks are completed, log off, or migrate business.

#### Limits

The following limits apply to KLP:

- For Alibaba Cloud Linux operating systems,
  - KLP applies to Alibaba Cloud Linux 2.1903 whose kernel version is kernel-4.19.24-9.al7.x86\_64 or later.
  - KLP applies to Alibaba Cloud Linux 3.1903 whose kernel version is **5.10.23-4.al8.x86\_64** or later.
- For each updated kernel version of Alibaba Cloud Linux operating systems, Alibaba Cloud provides one-year KLP support. After the one-year period ends, you must upgrade the kernels of the operating systems to the latest version.
- Not all CVEs or critical bugs can be fixed by using KLP. KLP is intended to reduce server restarts due to
  hotfix updates but cannot prevent server restarts all the time. KLP applies to CVEs of high and critical
  severity levels as well as critical bugs.
- KLP is not a general solution to upgrade kernels. It is applicable to updating hotfixes for CVEs or critical bugs only when it is inconvenient to immediately restart servers.

• During the process of updating hotfixes or after the hotfixes take effect, you cannot use SystemTap or kprobe to test or track functions involved in the hotfixes. Otherwise, the hotfixes become invalid.

### Related operations

For information about how to obtain, enable, or disable kernel hotfixes for Alibaba Cloud Linux, see Operations related to kernel hotfixes.

## 5.5.2. Operations related to kernel hotfixes

This topic describes how to obtain kernel hotfixes, enable or disable kernel hotfixes, and disable the kpatch service on Elastic Compute Service (ECS) instances that run Alibaba Cloud Linux operating systems.

#### Context

The following sections describe the operations related to kernel hotfixes:

- For information about how to obtain a kernel hotfix and view its details, see the Obtain a kernel hotfix section.
- For information about how to enable a kernel hotfix in an operating system, see the Enable a kernel hotfix section.
- For information about how to disable a kernel hotfix in an operating system when the hotfix is abnormal, see the Disable a kernel hotfix section.
- If the Kernel Live Patching (KLP) feature is enabled in the operating system, but you do not want the operating system to load all the kernel modules of your hotfixes when you restart your server, you can disable the kpatch service. For more information, see the Disable the kpatch service section.

#### Obtain a kernel hotfix

Each kernel hotfix is released by using an RPM package. You can use one of the following methods to obtain the installation package of a hotfix:

- Use the CVE announcement platform of Alibaba Cloud Linux
  - i. Access the common vulnerabilities and exposures (CVE) announcement platform.

All the released hotfixes are displayed on CVE announcement platforms. You can visit the following websites of CVE announcement platforms for Alibaba Cloud Linux distributions:

- Alibaba Cloud Linux 2.1903 Security Advisories
- Alibaba Cloud Linux 3 Security Advisories
- ii. Click **Advisory ID** to re-sort the information in the list. Arrange the hotfixes in descending order of time.

Advisory IDs starting with HOTFIX are kernel hotfixes, as shown in the following figure.

Updated Date	Affected Packages	Advisory ID	Severity	CVE ID(s)
2021-07-28	kernel-hotfix-5928799-5.al8	HOTFIX-SA-2021:0026	High	• CVE-2021-22555
2021-07-28	kernel-hotfix-5928799-4.al8	HOTFIX-SA-2021:0025	High	• CVE-2021-22555
2021-07-28	kernel-hotfix-5956925-6.al8	HOTFIX-SA-2021:0024	High	• CVE-2021-33909
2021-07-28	kernel-hotfix-5956925-5.al8	HOTFIX-SA-2021:0023	High	• CVE-2021-33909
2021-07-28	kernel-hotfix-5956925-4.al8	HOTFIX-SA-2021:0022	High	• CVE-2021-33909

134 > Document Version: 20220712

The following table describes the information on each CVE announcement platform.

Column	Description	
	The package name of a kernel hotfix. The name is prefixed by kernel -hotfix- and suffixed by the minor version number of Alibaba Cloud Linux operating systems.	
	You can determine whether a hotfix is applicable to your Alibaba Cloud Linux operating system based on the suffix of the package name of the hotfix. Example:	
Afforted Dackages	kernel-hotfix-5928799-5.al8 is displayed in the Affected Packages column, which indicates that the hotfix is applicable to operating systems whose minor version number is 5.al8.	
Affected Packages	■ After you run the uname -r command in your Alibaba Cloud Linux 3 operating system, 5.10.23-5.a18.x86_64 is returned.	
	The minor version number of all Alibaba Cloud Linux 3 operating systems is 5.al8. As such, the kernel-hotfix-5928799-5.al8 hotfix is applicable to your Alibaba Cloud Linux 3.	
	Note The version number of all Alibaba Cloud Linux 2 operating systems is al7.	
	The release sequential number of a hotfix. Hotfixes are classified into CVE hotfixes and Bugfix hotfixes. The following section describes the	
Advisory ID	naming conventions of hotfixes:  CVE hotfixes are prefixed by HOTFIX-SA	
Auvisory ID	■ Bugfix hotfixes are prefixed by HOTFIX-BA	
	You can click the <b>Advisory ID</b> of a hotfix to view its details and download the RPM package of the hotfix.	
CVE ID(s)	The ID of the CVE to be fixed. For a Bugfix hotfix, the CVE ID(s) column is empty.	

iii. Click the release sequential number of a hotfix in the **Advisory ID** column to go to the details page of the hotfix.

On the details page, you can view details about the hotfix and the name of the RPM package corresponding to the hotfix.

```
    Updated Packages
    x86_64
    kernel-hotfix-5928799-5.al8-1.0-20210720165816.al8.x86_64.rpm→ (download)
```

Each RPM package name is in the following format:  $kernel-hotfix-\{hotfix\_id\}-\{Minor\ version\ number\ of\ the\ operating\ system\}-\{Version\ number\ of\ the\ hotfix\}-\{Timestamp\ of\ the\ hotfix\}.\{Major\ version\ number\ of\ the\ operating\ system\}.\{System\ architecture\}.rpm$  .

The following section describes the RPM package name in the preceding figure:

■ 5928799: the ID of the hotfix.

- 5.al8: the minor version number of Alibaba Cloud Linux 3. You can use this hotfix only when the kernel version of your Alibaba Cloud Linux 3 is consistent with this version. You can run the uname -r command in your ECS instance to view the kernel version of the operating system.
- 1.0: the version number of the RPM package for the hotfix.
- 20210720165816: the time when the hotfix was created. The hotfix was created at 16:58:16 on July 20, 2021. An invalid value may be displayed in this field for some hotfixes of earlier versions.
- al8: the version number of Alibaba Cloud Linux 3. The version number of all Alibaba Cloud Linux 3 operating systems is al8.
  - Note The version number of all Alibaba Cloud Linux 2 operating systems is al7.
- x86\_64: the architecture of the operating system.
- Use a YUM reposit ory

You can run the **yum list** command to check the installation package of a kernel hotfix. For more information, see **Enable a kernel hotfix**.

#### Enable a kernel hotfix

- Connect to an Alibaba Cloud Linux ECS instance that requires hot fixes.
   For more information, see Connection methodsGuidelines on instance connection.
- 2. Run the following command to install the kpatch utility:

```
sudo yum -y install kpatch
```

- 3. Install a hotfix.
  - i. Run the following command to view the kernel version of the operating system:

```
sudo uname -r
```

In this example, the following command output indicates that the minor version number of the operating system is 5.al8.

5.10.23-5.al8.x86 64

 Run the yum list command to query the kernel hotfixes that are applicable to the operating system.

Command syntax:

```
sudo yum list | grep "kernel-hotfix" | grep "<Minor version number of the operating
system>"
```

In this example, the minor version number of the operating system is 5.al8. Run the following command:

```
sudo yum list | grep "kernel-hotfix" | grep "5.al8"
```

A command output similar to the following one is returned:

iii. Install the specified hotfix.

In this example,  $kernel-hotfix-5928799-5.al8.x86_64$  is used to demonstrate how to run the yum command to install a hotfix.

? Note When you use a YUM repository to install hotfixes, you do not need to include .rpm (the suffix of an RPM package) in the command.

```
sudo yum -y install kernel-hotfix-5928799-5.al8.x86_64
```

4. Run the following command to use the kpatch utility to check the status of the hotfix:

```
sudo kpatch list
```

A command output similar to the following one indicates that the kernel hotfix is installed and in effect:

```
Loaded patch modules:

kpatch_5928799 [enabled]

Installed patch modules:

kpatch_5928799 (5.10.23-5.al8.x86_64)
```

#### Disable a kernel hotfix

If errors exist in your hotfix (for example, if your hotfix does not take effect), you can perform the following steps to disable the kernel hotfix: In this example, the **yum** commands are used.

1. Run the following command to view the hotfix that can be disabled:

```
sudo yum list installed | grep kernel-hotfix
```

A command output similar to the following one is returned:

```
kernel-hotfix-5928799-5.al8.x86_64 1.0-20210720165816.al8 @alinux3-plus
```

2. Run the following command to delete the hotfix package in which errors exist.

In this example, the kernel-hotfix-5928799-5.al8.x86 64 hotfix is used.

```
sudo yum -y remove kernel-hotfix-5928799-5.al8.x86 64
```

3. Run the following command to check whether the hotfix in which errors exist has been deleted:

```
sudo kpatch list
```

A command output similar to the following one indicates that no hotfixes are installed and in effect:

```
Loaded patch modules:
Installed patch modules:
```

## Disable the kpatch service

If your hotfix is installed and in effect, but you do not want the operating system to load all the kernel modules of your hotfixes when you restart your server, you can perform the following steps to disable the kpatch service:

1. Run the following command to check whether the kpatch service is enabled:

```
sudo systemctl is-enabled kpatch.service
```

enabled in the command output indicates that the kpatch service is enabled.

2. Run the following command to disable the kpatch service:

```
sudo systemctl disable kpatch.service
```

A command output similar to the following one indicates that the kpatch service is disabled:

```
Removed symlink /etc/systemd/system/multi-user.target.wants/kpatch.service.
```

3. Run the following command to check the status of the kpatch service:

```
sudo systemctl status kpatch.service
```

A command output similar to the following one indicates that the kpatch service is disabled:

```
kpatch.service - "Apply kpatch kernel patches"
  Loaded: loaded (/usr/lib/systemd/system/kpatch.service; disabled; vendor preset: disabled)
```

Note You can run the sudo systemctl enable kpatch.service command to enable the kpatch service.

## 5.5.3. Use livepatch-mgr

livepatch-mgr is a kernel hotfix management tool provided by Alibaba Cloud Linux operating systems (OSs). You can use livepatch-mgr to view, install, or uninstall kernel hotfixes in Alibaba Cloud Linux OSs.

#### Context

To ensure system stability and security, Alibaba Cloud Linux distributions stay up-to-date on common vulnerabilities and exposures (CVEs) discovered by the Linux community and the industry and are updated with the latest versions of software packages (including kernels), bug fixes, and CVE fixes in a timely manner. For information about security updates for Alibaba Cloud Linux, see the following documents:

- Alibaba Cloud Linux 2.1903 Security Advisories
- Alibaba Cloud Linux 3 Security Advisories

You can choose based on the security updates to use hotfixes or coldfixes to update your software packages. This topic describes how to use livepatch-mgr.

## Install and view livepatch-mgr

- Connect to an Alibaba Cloud Linux Elastic Compute Service (ECS) instance that requires hot fixes.
   For information about how to connect to an ECS instance, see Connection methods Guidelines on instance connection.
- 2. Run the following command to install livepatch-mgr:

```
yum install livepatch-mgr -y
```

3. After livepatch-mgr is installed, run the following command to view the features supported by livepatch-mgr:

```
livepatch-mgr -h
```

The following command output is returned.

```
[root@test ~]# livepatch-mgr -h
usage: alinux-livepatch [-h] {list,update,load,unload} ...
 oositional arguments:
  {list,update,load,unload}
                        sub-command help
    list
                        list live patches
                        update all available live patches for the running
    update
                        kernel
    load
                        load all installed live patches for the running kernel
                        unload all available live patches for the running
    unload
                        kernel
 optional arguments:
                        show this help message and exit
```

livepatch-mgr provides the following commands:

- **update**: installs and applies applicable hotfixes to the instance. For more information, see the **Install and apply hotfixes** section.
- list: queries all hotfixes of the instance. For more information, see the Query hotfixes section.
- **load**: loads or applies hotfixes that are installed on the instance. For more information, see the **Load** or apply hotfixes section.
- **unload**: uninstalls hotfixes from the instance. For more information, see the **Uninstall hotfixes** section.

## Install and apply hotfixes

To install and apply all applicable hot fixes to the instance, run the following command:

```
livepatch-mgr update
```

livepatch-mgr provides the  $\ \ \, \text{--bugfix} \ \, , \ \ \text{--security} \ \, , \ \text{and} \ \ \, \text{-cve=}<\text{\it CVEs>} \ \,$  parameters for your use of commands in different scenarios.

• To install and apply applicable hotfixes for bugs, run the following command:

```
livepatch-mgr update --bugfix
```

• To install and apply applicable hotfixes for CVEs, run the following command:

```
livepatch-mgr update --security
```

• To install and apply hotfixes for specified CVEs, run the following command:

```
livepatch-mgr update -cve=<CVEs>
```

The *<CVEs>* variable specifies CVE IDs. CVE IDs are case-sensitive. If you specify multiple CVE IDs, separate them with commas (,). Example:

```
livepatch-mgr update -cve=CVE-2021-33909,CVE-2021-22555
```

Note You can obtain CVE IDs from Alibaba Cloud Linux 2.1903 Security Advisories or Alibaba Cloud Linux 3 Security Advisories.

In this example, the livepatch-mgr update command is run to install and apply all applicable hotfixes to the instance. The following command output is returned:

```
Plug-ins loaded: fastestmirror, langpacks, releasever-adapter, update-motd
Loading mirror speeds from cached hostfile
Resolving dependencies
--> Running transaction check
---> To install the kernel-hotfix-4121479-21.al7.x86 64.0.1.0-1.al7 software package
--> Dependency resolving complete
Dependencies resolved
Package
                               Architecture
                                                Version
                                                                    Source
ize
Software package being installed:
Transaction overview
Install 1 software package.
Total download size: 244 k
Installation package size: 1.0 M
Downloading packages:
Running transaction check
Running transaction test
Transaction test succeeded
Running transaction
 Installing : kernel-hotfix-4121479-21.al7-1.0-1.al7.x86_64
installing /var/khotfix/4.19.91-21.al7.x86 64/4121479/kpatch-4121479.ko (4.19.91-21.al7.x86
loading patch module: /var/khotfix/4.19.91-21.al7.x86 64/4121479/kpatch-4121479.ko
 Verifying : kernel-hotfix-4121479-21.al7-1.0-1.al7.x86 64
 kernel-hotfix-4121479-21.al7.x86 64 0:1.0-1.al7
Done.
```

## Query hotfixes

To query all hotfixes applicable to the instance, run the following command:

```
livepatch-mgr list
```

livepatch-mgr provides the --installed , --bugfix , --security , --running , and --available parameters for your use of commands in different scenarios.

• To query hotfixes installed on the instance, run the following command:

```
livepatch-mgr list --installed
```

• To query installed hot fixes for bugs on the instance, run the following command:

```
livepatch-mgr list --installed --bugfix
```

• To query installed hot fixes for CVEs on the instance, run the following command:

```
livepatch-mgr list --installed --security
```

• To query all hotfixes that are in effect on the instance, run the following command:

```
livepatch-mgr list --running
```

• To query hotfixes that are in effect for bugs on the instance, run the following command:

```
livepatch-mgr list --running --bugfix
```

• To query hotfixes that are in effect for CVEs on the instance, run the following command:

```
livepatch-mgr list --running --security
```

• To query all available hot fixes on the instance, run the following command:

```
livepatch-mgr list --available
```

• To query available hot fixes for bugs on the instance, run the following command:

```
livepatch-mgr list --available --bugfix
```

• To query available hot fixes for CVEs on the instance, run the following command:

```
livepatch-mgr list --available --security
```

In this example, the livepatch-mgr list command is run to query all hotfixes applicable to the instance. The following command output is returned. The queried hotfixes can be classified into the following categories:

- Loaded patch modules: hot fixes that are loaded or applied
- Installed patch modules: hotfixes that are installed
- Available and not installed patch modules: hotfixes that are available but not installed

```
Loaded patch modules:
              5692820 fix trace buffer out bound
CVE-2021-33909 5956925 seq file: disallow extremely large seq buffer allocations;
              4375449 ovl: fix dentry leak in ovl_get_redirect;
               5260815 ext4: Avoid freeing inodes on dirty list;
               5000697 mm, compaction: make capture control handling safe wrt interrupts;
CVE-2021-22555 5928799 netfilter: x tables: fix compat match/target pad out-of-bound write
              4121479 fix XDP socket queue descriptor entries overflow;
Installed patch modules:
              5692820 fix trace buffer out bound
CVE-2021-33909 5956925 seq_file: disallow extremely large seq buffer allocations;
              4375449 ovl: fix dentry leak in ovl get redirect;
               5260815 ext4: Avoid freeing inodes on dirty list;
              5000697 mm, compaction: make capture control handling safe wrt interrupts;
CVE-2021-22555 5928799 netfilter: x tables: fix compat match/target pad out-of-bound write
               4121479 fix XDP socket queue descriptor entries overflow;
Available and not installed patch modules:
{\tt CVE-2020-14386~kernel-hotfix-3768348-21.al7.x86\_64~Using~tp\_reserve~to~calculate~netoff~can}
overflow as tp_rese...(more)
              kernel-hotfix-5902278-21.al7.x86_64 We notice that hung task happens in a c
onner but practical ...(more)
```

## Load or apply hotfixes

To load all hotfixes that are installed on the instance, run the following commands:

```
livepatch-mgr load
```

livepatch-mgrtool provides the --security , --bugfix , and load <hotfix package> parameters for your use of commands in different scenarios.

• To load installed hot fixes for CVEs on the instance, run the following command:

```
livepatch-mgr load --security
```

• To load installed hot fixes for bugs on the instance, run the following command:

```
livepatch-mgr load --bugfix
```

• To load a specific hotfix that is installed on the instance, run the following command:

```
livepatch-mgr load <hotfix package>
```

The <hotfix package> value specifies the hotfix that you want to load. You can set the value to the absolute path or name of the kernel module (.ko file) in a hotfix. Examples:

• To specify a hotfix package by using the absolute path of the kernel module (.ko file) in a hotfix, run the following command:

```
livepatch-mgr load /var/khotfix/4.19.91-21.al7.x86_64/5956925/kpatch-5956925.ko
```

#### Sample command output:

```
loading kernel-hotfix-5956925 loading patch module: /var/khotfix/4.19.91-23.al7.x86_64/5956925/kpatch-5956925.ko
```

• To specify a hotfix package by using the name of the kernel module without the .ko suffix in a hotfix, run the following command:

```
livepatch-mgr load kpatch-5260815
```

#### Sample command output:

```
loading kernel-hotfix-5260815
loading patch module: /var/khotfix/4.19.91-23.a17.x86_64/5260815/kpatch-5260815.ko
```

• To specify a hotfix package by using the name of the kernel module with the .ko suffix in a hotfix, run the following command:

```
livepatch-mgr load kpatch-4121479.ko
```

#### Sample command output:

```
loading kernel-hotfix-4121479 loading patch module: /var/khotfix/4.19.91-23.al7.x86_64/4121479/kpatch-4121479.ko
```

In this example. the <code>livepatch-mgr load</code> command is run to load all hotfixes that are installed on the instance. The following command output is returned:

```
loading kernel-hotfix-5956925-23.x86_64
loading patch module: /var/khotfix/4.19.91-23.al7.x86_64/5956925/kpatch-5956925.ko
loading kernel-hotfix-5902278-23.x86_64
loading patch module: /var/khotfix/4.19.91-23.al7.x86_64/5902278/kpatch-5902278.ko
loading kernel-hotfix-5260815-23.x86_64
loading patch module: /var/khotfix/4.19.91-23.al7.x86_64/5260815/kpatch-5260815.ko
loading kernel-hotfix-5928799-23.x86_64
loading patch module: /var/khotfix/4.19.91-23.al7.x86_64/5928799/kpatch-5928799.ko
loading kernel-hotfix-5577959-23.x86_64
loading patch module: /var/khotfix/4.19.91-23.al7.x86_64/5577959/kpatch-5577959.ko
```

#### Uninstall hotfixes

To uninstall all hot fixes from the instance, run the following command:

```
livepatch-mgr unload
```

livepatch-mgr provides the --security , --bugfix , and unload <hotfix package> parameters for your use of commands in different scenarios.

• To uninstall hotfixes for CVEs from the instance, run the following command:

```
livepatch-mgr unload --security
```

• To uninstall hotfixes for bugs from the instance, run the following command:

```
livepatch-mgr unload --bugfix
```

• To uninstall a specific hotfix from the instance, run the following command:

```
livepatch-mgr unload <hotfix package>
```

The <hotfix package> value specifies the hotfix that you want to uninstall. You can set the value to the absolute path or name of the kernel module (.ko file) in a hotfix. Examples:

• To specify a hotfix package by using the absolute path of the kernel module (.ko file) in a hotfix, run the following command:

```
livepatch-mgr unload /var/khotfix/4.19.91-21.a17.x86_64/5928799/kpatch-5928799.ko
```

#### Sample command output:

```
unoading kernel-hotfix-5928799
disabling patch module: kpatch_5928799
unloading patch module: kpatch_5928799
```

• To specify a hotfix package by using the name of the kernel module without the .ko suffix in a hotfix, run the following command:

```
livepatch-mgr unload kpatch_5692820
```

#### Sample command output:

```
unoading kernel-hotfix-5692820
disabling patch module: kpatch_5692820
unloading patch module: kpatch_5692820
```

• To specify a hotfix package by using the name of the kernel module with the .ko suffix in a hotfix, run the following command:

```
livepatch-mgr unload kpatch_5000697.ko
```

#### Sample command output:

```
unoading kernel-hotfix-5000697
disabling patch module: kpatch_5000697
unloading patch module: kpatch_5000697
```

In this example, the livepatch-mgr unload command is run to uninstall all hotfixes from an instance. The following command output is returned:

```
unloading kernel-hotfix-5956925-23.x86_64
disabling patch module: kpatch_5956925
unloading patch module: kpatch_5956925
unloading kernel-hotfix-5902278-23.x86_64
disabling patch module: kpatch_5902278
unloading patch module: kpatch_5902278
unloading kernel-hotfix-5260815-23.x86_64
disabling patch module: kpatch_5260815
unloading patch module: kpatch_5260815
unloading patch module: kpatch_5260815
unloading kernel-hotfix-5928799-23.x86_64
disabling patch module: kpatch_5928799
unloading patch module: kpatch_5928799
unloading kernel-hotfix-5577959-23.x86_64
disabling patch module: kpatch_5577959
```

# 5.6. Features and interfacessupported by Alibaba Cloud Linux 2

# 5.6.1. Overview

This topic describes the kernel features and interfaces supported by Alibaba Cloud Linux. This topic is intended for users who are familiar with the Linux kernel and want to use Linux kernel features.

## Memory

References	Supported version	Description
Memcg Exstat feature	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.91- 18.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	Alibaba Cloud Linux 2 provides the Memcg Exstat feature. Unlike community versions of the Linux kernel, Alibaba Cloud Linux 2 kernel supports the following cgroup v1 interfaces for the Memcg Exstat feature: memory.events, memory.events.local, and memory.stat. Statistics on the latency caused by changes in the global minimum watermark and backend asynchronous reclaim are provided.

References	Supported version	Description
Memcg global minimum watermark rating	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.91- 18.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	Alibaba Cloud Linux 2 provides the memcg global minimum watermark rating feature. The global minimum watermark of resource-consuming tasks is increased to trigger direct memory reclaim. The global minimum watermark of latency-sensitive tasks is decreased to prevent direct memory reclaim. When a resource-consuming task requests a large amount of memory, the global minimum watermark is increased to throttle resources used for the task for a short period to prevent direct memory reclaim for latency-sensitive tasks. After a specified amount of memory is reclaimed by means of global kswapd backend reclaim, the resource-consuming task stops being throttled.
Memcg QoS feature of the cgroup v1 interface	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.91- 18.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	The memcg quality of service (QoS) feature can be used to control locks and limits on memory usage in a memcg. In community versions of the Linux kernel, this feature is supported only by the cgroup v2 interface. In Alibaba Cloud Linux 2, the memcg QoS feature is also supported by the cgroup v1 interface.
Memcg backend asynchronous reclaim	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.81- 17.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	Alibaba Cloud Linux 2 provides the backend asynchronous reclaim feature for memcgs. This feature differs from the global kswapd kernel thread because it uses the workqueue mechanism instead of a corresponding memcg kswapd kernel thread.
THP reclaim	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.91- 24.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	Alibaba Cloud Linux 2 provides the Transparent Huge Page (THP) reclaim feature. You can use the THP reclaim feature to fix memory usage issues caused by THPs, such as out of memory (OOM) errors.

# Network

	References	Supported version	Description
--	------------	-------------------	-------------

References	Supported version	Description
SMC-R	<ul> <li>Alibaba Cloud Linux 3 kernel version 5.10.60- 9.al8.x86_64 and later</li> </ul>	Alibaba Cloud Linux 3 is tuned to support Shared Memory Communications over RDMA (SMC-R). SMC-R is based on Alibaba Cloud eRDMA and can transparently replace TCP in applications without loss of functionality. SMC-R enables direct, high-speed, low-latency, and memory-to-memory communications and provides higher performance than TCP in various scenarios such as in-memory databases, remote procedure calls (RPCs), and transmission of large files.
TCP-RT configurations	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.91- 21.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	Alibaba Cloud Linux 2 supports TCP-layer service monitoring (TCP-RT). TCP-RT is a TRACE method. TCP-RT allows you to configure event tracking in a kernel-based TCP stack to identify a request and response when a single connection carries only one concurrent request and response. You can then obtain information such as the amount of time taken to receive the request in the TCP stack and the amount of time taken to process the request. In addition, TCP-RT supports statistical analysis in the kernel system and generates statistical information of specified connections on a regular basis.
Change the TCP TIME-WAIT timeout period	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.43- 13.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	In Linux kernels, TCP/IP connections remain in the TIME-WAIT state for 60 seconds. The length of this period cannot be changed. However, in specific scenarios such as heavy TCP loads, network performance can be improved if this period is shortened. To resolve this issue, Alibaba Cloud Linux 2 provides a kernel interface to change the period of time that TCP/IP connections remain in the TIME-WAIT state before the system closes them.

# Storage

References	Supported version	Description
Enable the cgroup writeback feature	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.36- 12.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	Alibaba Cloud Linux 2 provides the cgroup writeback feature for the cgroup v1 interface. This feature allows you to limit the buffered I/O rate when you use the cgroup v1 interface.

References	Supported version	Description
Configure the weight- based throttling feature of blk-iocost	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.81- 17.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	Alibaba Cloud Linux 2 provides the weight- based throttling feature (blk-iocost) based on the cost model. This feature is an improvement to the weight-based disk throttling feature of the I/O subsystem (blkcg) within a kernel.
Enhance the monitoring of block I/O throttling	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.81- 17.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	Alibaba Cloud Linux 2 provides interfaces to better monitor Linux block I/O throttling.
Use the JBD2 optimization interface	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.81- 17.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	JBD2 is the kernel thread of ext4 file systems. JBD2 often experiences the shadow (BH_Shadow) state when it is being used, which can affect the system performance. To resolve this issue, Alibaba Cloud Linux 2 provides an interface to optimize JBD2.
Create a hard link across project quotas	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.91- 18.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	By default, ext4 file systems have constraints that forbid hard links to be created across project quotas. However, some scenarios require hard links to be created. Alibaba Cloud Linux 2 provides a custom interface to bypass the constraints of ext4 file systems and create hard links across project quotas.
Track I/O latency	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.91- 18.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	Alibaba Cloud Linux 2 optimizes the /proc/diskstats interface that provides raw data for the I/O latency analysis tool iostat. Alibaba Cloud Linux 2 can calculate the durations of read, write, and special I/O (discard) operations on the device side. In addition, Alibaba Cloud Linux 2 provides the bcc tool to track I/O latency.
Detect I/O hangs of file systems and block layers	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.91- 18.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	An I/O hang occurs when the system becomes unstable or even goes down due to time-consuming I/O requests. Alibaba Cloud Linux 2 extends the core data structure and provides the feature to identify and detect I/O hangs at low system overheads.

# Monitoring

References	Supported version	Description
UKFEF	<ul> <li>Alibaba Cloud Linux 3 kernel version 5.10.60- 9.al8.x86_64 and later</li> </ul>	Alibaba Cloud Linux 3 supports Unified Kernel Fault Event Framework (UKFEF) in the kernel of the 5.10.60-9.al8.x86_64 version. UKFEF is used to collect the exception events on Alibaba Cloud Linux 3 that may cause risks and generate event reports in a unified format.
Enable the PSI feature for cgroup v1	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.81- 17.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	Alibaba Cloud Linux 2 provides the Pressure Stall Information (PSI) feature for the cgroup v1 interface. This feature allows you to monitor CPUs, memory, and I/O performance.

# Scheduling

References	Supported version	Description
Enable the CPU burst feature for cgroup v1	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.91- 22.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	Alibaba Cloud Linux 2 provides the CPU burst feature for the cgroup v1 interface. The CPU burst feature allows CPU-throttled containers to burst their CPU utilization for higher performance and lower latency.
Group identity feature	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.91- 24.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	Alibaba Cloud Linux 2 provides the group identity feature. This feature allows you to configure different identities for CPU cgroups to prioritize process tasks in the cgroups.

# Compatibility

References	Supported version	Description
Statistical status change of anonymous pages in the /proc/meminfo file	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.91- 22.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	In Alibaba Cloud Linux 2 kernel version 4.19.91-22.al7.x86_64 and later, the statistical  Active (anon) and Inactive (anon) states of anonymous pages in the /proc/meminfo file are changed. If your application uses these statistics, this change causes inconsistency between data queried before and after the change.

# Others

References Supported version	Description
------------------------------	-------------

References	Supported version	Description
Enhance the visibility of container resources	<ul> <li>Alibaba Cloud Linux 2 kernel version 4.19.91- 22.2.al7.x86_64 and later</li> <li>All Alibaba Cloud Linux 3 kernel versions</li> </ul>	Alibaba Cloud Linux 2 provides kernel interfaces for the container resource visualization feature to enhance the visibility of container resources.

# 5.6.2. Memcg Exstat feature

This topic describes the Memcg Exstat (Extend/Extra) feature provided by Alibaba Cloud Linux 2 kernel version 4.19.91-18.a17 or later.

#### Context

Compared with the Memcg Exstat feature provided by community versions of the Linux kernel, the Memcg Exstat feature that is available with the Alibaba Cloud Linux 2 kernel provides the following additional memcg statistical items:

- Interfaces supported by the cgroup v1 interface: memory.events, memory.events.local, and memory.stat
- Statistics on latency caused by changes in the global minimum watermark
- Statistics on latency caused by backend asynchronous reclaim

#### **Features**

Alibaba Cloud Linux 2 kernel features are implemented based on kernel interfaces. This topic describes the implementation methods for each feature.

Feature	Description
Statistics on memory events	The cgroup v2 interface of community versions of the Linux kernel supports the memory.events and memory.events.local interfaces to show the number of times that specific memory events have occurred in a memcg. For more information, visit cgroup-v2.rst.  The cgroup v1 interface of the Alibaba Cloud Linux 2 kernel supports the memory.events and memory.events.local interfaces.
	Note The interface files are not stored in the memcg root directory.

Feature	Description	
Statistics on memory workingsets	The cgroup v2 interface of community versions of the Linux kernel supports the memory.stat interface. The interface file contains the following statistical items: workingset refault , workingset activate , and workingset nodereclaim . For more information, visit cgroup-v2.rst.	
	The cgroup v1 interface of the Alibaba Cloud Linux 2 kernel supports the memory.stat interface. The interface file contains the following statistical items: workingset refault , workingset activate , workingset nodereclaim , and workingset restore . The following content is the official description of workingset restore :	
	<pre>+ workingset_restore + Number of restored pages which have been detected as an active + workingset before they got reclaimed.</pre>	
Statistics on latency caused by changes in the global minimum watermark	Alibaba Cloud Linux 2 provides the memcg global minimum watermark rating feature. For more information, see Memcg global minimum watermark rating.	
	Alibaba Cloud Linux 2 provides statistics on the time of throttling that results from the offset global minimum watermark specified in the memcg.exstat interface being exceeded. The statistical item is wmark_min_throttled_ms .	
	Note This statistical item is recursive to the parent group, and the interface file is not stored in the memcg root directory.	
	Statistics description:	
	• Interface file: memory.exstat	
	• Statistical item: wmark_min_throttled_ms	
	• Unit: ms	

Feature	Description	
Statistics on latency caused by backend asynchronous reclaim	Alibaba Cloud Linux 2 provides the backend asynchronous reclaim feature memcg kswapd. For more information, see Memcg backend asynchronous reclaim.  Alibaba Cloud Linux 2 provides statistics on latency caused by backend	
	asynchronous reclaim in the memcg.exstat interface. These statistics include the block time and actual working time of the reclaim. The statistical item is <a href="https://www.wmark_reclaim_work_ms">wmark_reclaim_work_ms</a> .	
	Note This statistical item is recursive to the parent group, and the interface file is not stored in the memcg root directory.	
	Statistics description:	
	Interface file: memory.exstat	
	• Statisticalitem: wmark_reclaim_work_ms	
	• Unit: ms	

## **Examples**

Create a test file at the /sys/fs/cgroup/memory mount point of the memcg. Make sure that the test file path contains the memory.events, memory.events.local, and memory.exstat interface files. You must also make sure that the memory.stat interface file contains the workingset statistical item.

1. Run the following command to create a test file:

```
mkdir /sys/fs/cgroup/memory/test
```

- 2. Query the memory.events, memory.events.local, and memory.exstat interfaces.
  - i. Run the following command to query the memory.events interface:

```
cat /sys/fs/cgroup/memory/test/memory.events
```

#### Example out put:

```
low 0
high 0
max 0
oom 0
oom_kill 0
```

ii. Run the following command to query the memory.events.local interface:

```
cat /sys/fs/cgroup/memory/test/memory.events.local
```

#### Example out put:

```
low 0
high 0
max 0
oom 0
oom_kill 0
```

iii. Run the following command to query the memory.exstat interface:

```
cat /sys/fs/cgroup/memory/test/memory.exstat
```

#### Example out put:

```
wmark_min_throttled_ms 0
wmark_reclaim_work_ms 0
```

3. Run the following command to check whether the memory.stat interface file contains the workin gset statistical item:

```
cat /sys/fs/cgroup/memory/test/memory.stat | grep workingset
```

A command output similar to the following one indicates that the interface file contains the work ingset statistical item:

```
workingset_refault 0
workingset_activate 0
workingset_restore 0
workingset_nodereclaim 0
total_workingset_refault 0
total_workingset_activate 0
total_workingset_restore 0
total_workingset_nodereclaim 0
```

# 5.6.3. Memcg global minimum watermark rating

This topic describes the memcg global minimum watermark rating feature provided by Alibaba Cloud Linux 2 kernel version 4.19.91-18.a17 and later.

#### Context

Global memory reclaim has a great impact on system performance in Linux kernels. When latency-sensitive tasks and resource-consuming tasks are deployed together, resource-consuming tasks request a large amount of memory instantly, which causes the free memory of the system to reach the global minimum watermark (global wmark\_min). As a result, direct memory reclaim is enabled for all system tasks, which causes performance jitters for latency-sensitive tasks. Both global kswapd backend reclaim and memory backend reclaim cannot resolve this problem.

In this case, Alibaba Cloud Linux 2 provides the memcg global minimum watermark rating feature. The global minimum watermark of resource-consuming tasks is increased to trigger direct memory reclaim. The global minimum watermark of latency-sensitive tasks is decreased to avoid direct memory reclaim. When a resource-consuming task requests a large amount of memory, an increase in the global minimum watermark throttles the resources used by the task for a short period to avoid direct memory reclaim for latency-sensitive tasks. After a certain amount of memory is reclaimed through global kswapd backend reclaim, the resource-consuming task is not throttled.

# Interface description

The interface that implements the memory global minimum watermark rating feature is memory.wmark\_min\_adj. The value of this interface indicates the percent of adjustment over the global minimum watermark. Valid values: -25 to 50.

• This interface inherits a value of 0 from its parent group when the interface is created. Therefore, the

default value is 0.

• A negative value is a percent of adjustment over the [0, wmark\_min] range, where wmark\_min is the value of global wmark min. Example:

```
memory.wmark min adj=-25, memcg WMARK MIN is "WMARK MIN + (WMARK MIN - 0) * (-25%)"
```

- **? Note** A negative value also indicates a decrease of the global minimum watermark. This can increase the memcg quality of service (QoS) of latency-sensitive tasks.
- A positive value is a percent of adjustment over the [wmark\_min, wmark\_low] range, where wmark min and wmark\_low are values of global wmark\_min and global wmark\_low. Example:

```
memory.wmark_min_adj=50, memog WMARK_MIN is "WMARK_MIN + (WMARK_LOW - WMARK_MIN) * 50%"
```

- **? Note** A positive value also indicates an increase of the global minimum watermark. This can decrease the memcg QoS of resource-consuming tasks.
- When the offset global minimum watermark is triggered, throttling is performed, and the throttling time is linearly proportional to the excessive memory usage. Valid values of the throttling time: 1 to 1000. Unit: ms.
  - **?** Note This interface file is not stored in the memcg root directory.

#### **Precautions**

A multi-level memorg contains <code>effective memory.wmark\_min\_adj</code> , which is the final effective value of memory.wmark\_min\_adj. The values of memory.wmark\_min\_adj at all levels are traversed to obtain the maximum value. Intermediate nodes that have 0 as the default value are excluded. The following hierarchy provides an example:

The following table describes the mapping between the values of memory.wmark\_min\_adj at each level and the final effective value.

Level	Value at each level	Final effective value
A	-10	-10
В	-25	-10
С	0	0
D	50	50

Level	Value at each level	Final effective value
E	-25	-10
F	50	50

#### ? Note

- The value displayed in the output of the cat /sys/fs/cgroup/memory/<Memcg path>/memory
  .wmark\_min\_adj command is the final effective value. In the command, <memcg path>
  indicates the root path of the memcg.
- We recommend that you use the global minimum watermark rating feature together with global wmark\_min. For example, you can set global wmark\_min to 2 GB or more in /proc/sys/vm/min\_free\_kbytes.

### Configuration examples

Example 1: Configure global minimum watermark rating for the memcg of latency-sensitive tasks.

- 1. Run the mkdir /sys/fs/cgroup/memory/test-lc command to create a test file.
- 2. Run the echo -25 > /sys/fs/cgroup/memory/test-lc/memory.wmark\_min\_adj command to write the value of -25 to the interface to increase the memorg QoS of latency-sensitive tasks.

Example 2: Configure global minimum watermark rating for the memcg of resource-consuming tasks.

- 1. Run the mkdir /sys/fs/cgroup/memory/test-be command to create a test file.
- 2. Run the echo 25 > /sys/fs/cgroup/memory/test-be/memory.wmark\_min\_adj command to write the value of 25 to the interface to decrease the memcg QoS of resource-consuming tasks.

# 5.6.4. Memcg QoS feature of the cgroup v1 interface

The memory control group (memcg) quality of service (QoS) feature can be used to control locks and limits on memory usage in a memcg. In community versions of the Linux kernel, this feature is supported only by the  $cgroup \ v2$  interface. In Alibaba Cloud Linux 2 kernel version 4.19.91-18.a17 and later, this feature is also supported by the  $cgroup \ v1$  interface.

## **Background**

In the Alibaba Cloud Linux 2 kernel, the memog QoS feature is enabled in the cgroup v1 interface by default. For more information about memog QoS , see Documentation/admin-guide/cgroup-v2.rst . You can obtain the kernel document from the Debuginfo package and source code package of Alibaba Cloud Linux 2. For more information, see Use Alibaba Cloud Linux 2.

#### **Precautions**

When you use the memcg QoS feature of the  $\ensuremath{\mathtt{cgroup}}\ \ensuremath{\mathtt{v1}}\ \ensuremath{\mathtt{interface}},$  we recommend that you place your tasks in a memcg leaf node such as  $\ensuremath{\mathtt{/sys/fs/cgroup/memory/<intermediate}}\ \ensuremath{\mathtt{node>/<leaf}}\ \ensuremath{\mathtt{node>/tasks}}.$ 

# Interface description

The following table describes the interfaces that implement the memcg QoS feature of the cgroup v1 interface in the Alibaba Cloud Linux 2 kernel.

Interface	Description
memory.min	Absolutely locks the memory. Memory locked by this interface cannot be reclaimed. You can perform the following read and write operations on this interface:  Read the size of locked memory from this interface.  Write the size of locked memory to this interface.
memory.low	Relatively locks the memory. Memory locked by this interface may be partially reclaimed if the system has no other memory to reclaim. You can perform the following read and write operations on this interface:  Read the size of locked memory from this interface.  Write the size of locked memory to this interface.
memory.high	Limits the memory usage. You can perform the following read and write operations on this interface:  Read the usage limits on the memory from this interface.  Write the usage limits on the memory to this interface.

## **Examples**

Create a test memcg in the memcg mount directory such as /sys/fs/cgroup/memory/. Make sure that the memcg contains the memory.min, memory.low, and memory.high interfaces.

#### Example command:

```
mkdir /sys/fs/cgroup/memory/test
ls /sys/fs/cgroup/memory/test | grep -E "memory.(min|low|high)"
```

#### Example output:

memory.high
memory.low
memory.min

# 5.6.5. Memcg backend asynchronous reclaim

The memory control group (memcg) backend asynchronous reclaim feature is provided by Alibaba Cloud Linux 2 kernel version 4.19.81-17.a17 and later. This topic describes the interfaces that implement the memcg backend asynchronous reclaim feature.

#### Context

In community versions of the Linux kernel, the system allocates memory and triggers memcg-level direct memory reclaim when the memory reaches a limit set by a memcg. Direct memory reclaim is synchronous reclaim that occurs in the context of memory allocation and affects the performance of the current process.

To resolve this problem, Alibaba Cloud Linux 2 provides the backend asynchronous reclaim feature for memcgs. This feature differs from the global kswapd kernel thread in that this feature uses the workqueue mechanism instead of creating a corresponding memcg kswapd kernel thread. Additionally, four memcg control interfaces are added in each of the cgroup v1 and cgroup v2 interfaces.

#### **Precautions**

- Memory allocation of the current memorg may recursively trigger the backend asynchronous reclaim of the parent group.
- When backend asynchronous reclaim is triggered, it starts from the memcg on which the feature is triggered and is performed in sequence down the hierarchy.
- When memory.high is configured and the value of memory.high is smaller than that of memory.limit\_in\_bytes, the values of memory.wmark\_high and memory.wmark\_low are calculated based on memory.high instead of memory.limit\_in\_bytes.

## Interface description

Interface	Description
memory.wmark_ratio	Specifies whether to enable the memcg backend asynchronous reclaim feature and sets the memcg memory watermark where asynchronous reclaim starts. Unit: percent of the memcg memory limit. Valid values: 0 to 100.  The default value is 0, which indicates that the memcg backend asynchronous reclaim feature is disabled.
	<ul> <li>When the value is not 0, the memcg backend asynchronous reclaim feature is enabled and the corresponding watermark is set.</li> </ul>
memory.wmark_high	A read-only interface.
	<ul> <li>When the memcg memory usage exceeds the value of this interface, backend asynchronous reclaim is started.</li> </ul>
	• The value of this interface is calculated based on the following formula: (memory.limit_in_bytes × memory.wmark_ratio/100) .
	<ul> <li>When the memcg backend asynchronous reclaim feature is disabled, the default value of memory.wmark_high is the maximum value and the backend asynchronous reclaim feature is never triggered.</li> </ul>
	This interface file is not stored in the memcg root directory.

Interface	Description	
memory.wmark_low	<ul> <li>A read-only interface.</li> <li>When the memcg memory usage is less than the value of this interface, backend asynchronous reclaim is ended.</li> <li>The value of this interface is calculated based on the following formula:         <pre>memory.wmark_high - memory.limit_in_bytes × memory.wmark_sca le_factor/10000</pre> </li> <li>This interface file is not stored in the memcg root directory.</li> </ul>	
memory.wmark_scale_facto r	Controls the difference between values of memory.wmark_high and memory.wmark_low. Unit: 0.01 percent of the memory memory limit. Valid values: 1 to 1000.  This interface inherits the value of its parent group when the interface is created. The inherited value is 50, which indicates 0.50% of the memory memory limit. This is also the default value.  This interface file is not stored in the memory root directory.	

## Configuration examples

1. Create a test file.

mkdir /sys/fs/cgroup/memory/test/

2. Specify the value of memory.limit\_in\_bytes.

In this example, the value is 1 G.

```
echo 1G > /sys/fs/cgroup/memory/test/memory.limit_in_bytes
```

3. Configure memory.wmark\_ratio.

In this example, the memcg memory watermark where asynchronous reclaim starts is set to 95% of the memcg memory limit.

```
echo 95 > /sys/fs/cgroup/memory/test/memory.wmark_ratio
```

4. Run the following command to check the value of memory.wmark\_scale\_factor:

```
cat /sys/fs/cgroup/memory/test/memory.wmark_scale_factor
```

The default value is 0.50% of the memcg memory limit. Example output: 50.

5. Query the values of memory.wmark\_high and memory.wmark\_low.

If the following values are displayed for memory.wmark\_high and memory.wmark\_low in the command outputs, the configurations are correct.

• Run the following command:

```
cat /sys/fs/cgroup/memory/test/memory.wmark_high
```

Example out put: 1020051456 .

• Run the following command:

```
cat /sys/fs/cgroup/memory/test/memory.wmark_low
```

Example out put: 1014685696 .

## 5.6.6. THP reclaim

Alibaba Cloud Linux 2 provides the Transparent Huge Page (THP) reclaim feature in kernel versions 4.19.91-24.al7 and later. You can use the THP reclaim feature to fix memory usage issues caused by THP, such as out of memory (OOM) errors. This topic describes the THP reclaim feature and provides examples on how to configure the operations of the feature.

#### Context

In Linux operating systems, memory is managed in blocks known as pages. THPs are blocks of memory that come in 2 MiB and 1 GiB sizes, while the size of a regular memory page (subpage) is 4 KiB. Applications consume large amounts of memory, which incurs high address translation overheads. When applications request memory, the kernel dynamically allocates THPs to reduce Translation Lookaside Buffer (TLB) misses and improve application performance.

However, THPs may cause memory bloat. When the THP feature is enabled, the kernel allocates 2 MiB blocks of memory as THPs, which is equivalent to 512 subpages. This leads to high memory fragmentation, which in turn leads to memory overcommitment.

Memory bloat may lead to OOM errors. For example, when an application that only requests 8 KiB of memory (2 subpages) is assigned a 2 MiB THP, the remaining 510 subpages are zero subpages, which represents a significant waste of resident set size (RSS) and will eventually result in an OOM error.

To fix memory issues caused by THPs, Alibaba Cloud Linux 2 provides the THP reclaim feature for memory control groups (memcgs). THP reclaim is a mechanism that splits THPs and reclaims zero subpages to avoid OOM errors caused by memory bloat. However, using this feature may cause memory performance degradation.

#### Operations

The following table describes the operations of the THP reclaim feature.

Operation	Description
memory.thp_reclaim	<ul> <li>Enables or disables the THP reclaim feature. Valid values:</li> <li>reclaim: enables the THP reclaim feature.</li> <li>swap: reserved for future use.</li> <li>disable: disables the THP reclaim feature.</li> </ul> Default value: disable.
memory.thp_reclaim_stat	<ul> <li>Queries the state of the THP reclaim feature. Description of response parameters:</li> <li>queue_length: the queue length of each node. When the THP reclaim feature is enabled, THPs are added to a reclaim queue.</li> <li>split_hugepage: the total number of THPs of each node that have been split.</li> <li>reclaim_subpage: the total number of zero subpages of each node that have been reclaimed.</li> <li>The values of these parameters are listed for each NUMA node in ascending order of node IDs (node0, node1) from left to right.</li> </ul>

Operation	Description
memory.thp_reclaim_ctrl	Controls how the THP reclaim feature is triggered. Description of request parameters:  • threshold: The upper limit on the number of zero subpages in a THP. If the number of zero subpages exceeds the specified limit, the THP reclaim feature is triggered immediately. The default value of the threshold parameter is 16.  • reclaim: manually triggers the THP reclaim feature.
/sys/kernel/mm/transparent	The global operation. If you do not want to configure each memory control group, you can use this operation. Valid values:  • memcg: the default value. This value specifies that each memory control group uses its own configuration that is configured by using memory.thp_reclaim.
_hugepage/reclaim	<ul> <li>reclaim: forcibly enables the THP reclaim feature for all memory control groups.</li> <li>swap: reserved for future use.</li> <li>disable: forcibly disables the THP reclaim feature for all memory control groups.</li> </ul>

## Configure the operations

This example describes how to configure the operations that implement the THP reclaim feature. In the example, a memory control group named *test* is used.

1. Run the following command to create a memory control group named test:

```
mkdir /sys/fs/cgroup/memory/test/
```

2. Run the following command to enable the THP reclaim feature for test:

```
echo reclaim > /sys/fs/cgroup/memory/test/memory.thp_reclaim
```

3. Run the following command to check whether the THP reclaim feature is enabled for test:

```
\verb|cat /sys/fs/cgroup/memory/test/memory.thp_reclaim| \\
```

The settings that take effect are enclosed in a pair of brackets ( [] ). [reclaim] shown in the following command output indicates that the THP feature is enabled for *test*.

```
[root@test ~]# cat /sys/fs/cgroup/memory/test/memory.thp_reclaim
[reclaim] swap disable
```

4. Run the following command to forcibly enable the THP reclaim feature by using the global operation:

```
echo reclaim > /sys/kernel/mm/transparent hugepage/reclaim
```

If you want to forcibly disable the THP reclaim feature, run the following command:

```
echo disable > /sys/kernel/mm/transparent_hugepage/reclaim
```

Note When the value of the global operation is reclaim or disable, the configurations of the global operation take precedence over the operation. However, the configurations for memory.thp\_reclaim in individual memory control groups are not affected.

5. Run the following command to use the memory.thp\_reclaim\_ctrl operation to set the upper limit on the number of zero subpages in a THP for *test*:

```
echo "threshold 32" > /sys/fs/cgroup/memory/test/memory.thp_reclaim_ctrl
```

In this case, if the number of zero subpages in a THP exceeds 32, zero subpages reclaim is triggered.

6. Manually trigger the zero subpage reclaim feature.

The THP reclaim feature reclaims excess zero subpages if the number of zero subpages exceeds the value of threshold . If you want to set the memory.thp\_reclaim\_ctrl operation to reclaim , take note of the following items:

Note You can only write reclaim into the /sys/fs/cgroup/memory/test/memory.thp\_reclaim\_ctrl file to trigger the zero subpage reclaim feature, and you cannot use the cat command to view the memory.thp\_reclaim\_ctrl operation configurations.

• Run the following command to trigger the zero subpage reclaim feature for the current memory control group:

```
echo "reclaim 1" > /sys/fs/cgroup/memory/test/memory.thp_reclaim_ctrl
```

• Run the following command to recursively trigger the zero subpage reclaim feature for the current memory control group and its sub control groups:

```
echo "reclaim 2" > /sys/fs/cgroup/memory/test/memory.thp_reclaim_ctrl
```

The zero subpage reclaim feature can not only be manually triggered by setting memory.thp\_reclaim\_ctrl to reclaim , but also be automatically triggered by memory reclaim.

- When an OOM error occurs, the zero subpage reclaim feature is triggered.
- When the backend asynchronous reclaim feature is triggered, the zero subpage reclaim feature is triggered. For more information about the backend asynchronous reclaim feature, see Memcg backend asynchronous reclaim.
- 7. Run the following command to view the state of the THP reclaim feature for *test*:

```
cat /sys/fs/cgroup/memory/test/memory.thp_reclaim_stat
```

A command output similar to the following one is displayed:

```
queue_length 14
split_hugepage 523
reclaim_subpage 256207
```

# Sample C code

This section provides sample C code that is used for applications to request THPs. You can see different results of the applications to request THPs when the THP reclaim feature is enabled or disabled.

1. Run the following command to set the value of memory.limit\_in\_bytes to 1 GiB:

```
echo 1G > /sys/fs/cgroup/memory/test/memory.limit_in_bytes
```

Then, run the following command to query the value of memory.limit in bytes.

```
cat /sys/fs/cgroup/memory/test/memory.limit_in_bytes
```

A command output similar to the following one is displayed:

```
[root@test ~]# echo 1G > /sys/fs/cgroup/memory/test/memory.limit_in_bytes
[root@test ~]# cat /sys/fs/cgroup/memory/test/memory.limit_in_bytes
1073741824
[root@test ~]#
```

2. Run the following command to disable the memog backend asynchronous reclaim feature:

For more information about the memcg backend asynchronous reclaim feature, see Memcg backend asynchronous reclaim.

```
echo 0 > /sys/fs/cgroup/memory/test/memory.wmark_ratio
```

3. Run the following C code when the THP reclaim feature is enabled and when the THP reclaim feature is disabled:

```
// The application requests 1 GiB of memory (512 THPs). Of these THPs, 10 contain zero
#define HUGEPAGE SIZE 4096 * 512
int main()
       int i, thp = 512;
       char *addr;
       posix memalign((void **)&addr, HUGEPAGE SIZE, HUGEPAGE SIZE * thp);
        for (i = 0; i < 10; i++) {
                memset(addr, 0xc, HUGEPAGE SIZE >> 1);
                addr += HUGEPAGE SIZE;
        }
        for (; i < thp; i++) {
                memset(addr, 0xc, HUGEPAGE SIZE);
                addr += HUGEPAGE SIZE;
        }
       pause();
        return 0:
}
```

The results vary based on whether the THP reclaim feature is enabled.

- When the THP reclaim feature is disabled, OOM errors occur when the last THP is requested.
- When the THP reclaim feature is enabled, the THPs allocated for applications are split into subpages and zero subpages are reclaimed. In this case, OOM errors do not occur.

# 5.6.7. TCP-RT configurations

TCP-layer service monitoring (TCP-RT) is supported on Alibaba Cloud Linux 2 with the kernel of version 4.19.81-17.al7 and later.

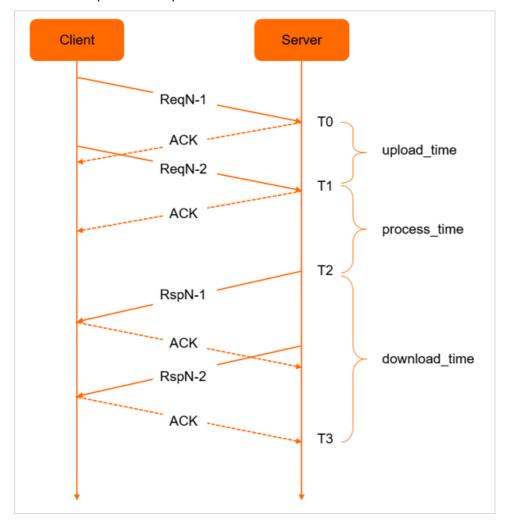
#### Context

TCP-RT is a TRACE method. TCP-RT allows you to configure event tracking in a kernel-based TCP stack to identify a request and response when a single connection carries only one concurrent request and response. You can then obtain information such as the amount of time taken to receive the request in the TCP stack and the amount of time taken to process the request. In addition, TCP-RT supports statistical analysis in the kernel system and generates statistical information of specified connections on a regular basis.

TCP-RT applies to services such as HTTP/1.1 web services, MySQL database services, and Redis services that have only one concurrent request and response on a single connection.

#### Principle

The following figure shows different phases when TCP-RT is used for a connection that carries one concurrent request and response.



In the preceding figure, the Nth request sent by the client to the server is displayed as ReqN. The request is composed of the ReqN-1 and ReqN-2 data packets. The time when the server receives the first packet is recorded as T0 and the time when the second packet is received is recorded as T1. The server processes the received request and sends the RspN-1 and RspN-2 response packets to the client. The time when the server sends the first response packet is recorded as T2. When the client receives a response packet, it sends an acknowledgement (ACK) message to the server. The time when the server receives the final ACK message is recorded as T3.

TCP-RT helps obtain the following information based on the recorded points in time:

- upload time: the amount of time taken to upload the request.
- process\_time: the amount of time taken by the server to process the request. It is a period that starts when the server receives the final data packet of the request and ends when the server sends the first response packet to the client.
- download\_time: the amount of time taken to download the data. It is a period that starts when the
  server sends the first response packet to the client and ends when the server receives the final ACK
  message from the client. This time period is essential for downloading large amounts of response
  data.

# Information output

TCP-RT collects and generates the parameters related to the TCP service in kernel mode. The following table describes the output modes and output time of files.

File type	Output mode	Output time
Log file	Log files are the rt-network-log* files generated in the /sys/kernel/debug/tcp-rt directory by using debugfs. Log files have the following features:  The suffix of the file name is the serial number of the CPU core. For example, the log files generated by a 32-core server are named from rt-network-log0 to rt-network-log31.  Each log file can be up to 2 MB in size. When a log file exceeds this size limit, previous data is deleted.  Log files generated by using debugfs are one-time files. After the data of a one-time file is read for the first time, the data is deleted.	<ul> <li>Time 1: For a TCP connection, information about a task is generated when the task is complete and the next task starts. A task consists of a request and response.</li> <li>Time 2: Information is generated when the connection is closed.</li> <li>After data is generated by TCP-RT, the application layer can immediately read the data.</li> </ul>
stats file	stats files are the rt-network-stats files generated in the /sys/kernl/debug/tcp-rt directory. These files are generated based on the server port- or client port-related statistics.	stats files are generated at regular intervals. By default, data is generated once a minute.

#### **Output formats**

This section describes the formats of generated log files and stats files.

- Note Definitions of task and TCP connection lifecycle:
  - A task consists of a request and response.
  - The TCP connection lifecycle includes multiple tasks.
- Log file format

In log files, each column of a record corresponds to a different set of information. The following figure shows an example of a log file.



The following section describes the parameters from left to right as recorded in the figure.

- The version number. In this example, the version number is V6.
- o The scenario record identifier. The identifiers are divided into five types: R, E, W, N, and P.
  - R: An R record is generated for the TCP connection when a request is received by the local server and a response is sent.
  - E: An E record is generated when the connection is closed.
  - W: A W record is generated when the connection is closed during data transmission.
  - N: An N record is generated when the connection is closed during data reception.
  - P: A P record is generated for the TCP connection when the local server sends a request to the peer server and a response is sent.
- The seconds portion of the start time of the task.
- o The microseconds portion of the start time of the task.
- The peer IP address of the TCP connection.
- The peer port of the TCP connection.
- o The local IP address of the TCP connection.
- The local port of the TCP connection.

The preceding parameters are followed by different parameters in different scenarios. The following table describes the parameters.

Scenario record identifier	Parameter
----------------------------	-----------

Scenario record identifier	Parameter
	An R record records the normal startup and shutdown of a task. Each TCP connection can have multiple R records.
	o The amount of data sent by the task. Unit: bytes.
	<ul> <li>The duration of the task, which is the interval between the time when the TCP service receives a request from the client and the time when the ACK message is received from the client. Unit: µs.</li> </ul>
	<ul> <li>The minimum TCP round trip time (RTT) for the task. Unit: μs.</li> </ul>
	• The number of TCP segments retransmitted by the task.
	<ul> <li>The sequence number of the task. After a TCP connection is established, the sequence number of the first task is 1.</li> </ul>
R	$\circ$ The latency of the task, which is the interval between the time when the final request segment is received and the time when the first response segment is sent . Unit: $\mu s$ .
	• The latency for sending request segments of the task, which is the interval between the time when the final request segment is received and the time when the final request segment is received. Unit:
	μs.
	• The amount of data received by the task. Unit: bytes.
	<ul> <li>Specifies whether unordered reception occurs during the task. A value of 1 indicates that unordered reception occurs, and a value of 0 indicates that unordered reception did not occur.</li> </ul>
	• The TCP maximum segment length (MSS) used during the task. Unit: bytes.
	A P record records the normal startup and shutdown of a task. Each TCP connection can have multiple P records. P is a new record type in V6 and indicates the request information from the client. This record is available only when the pports or pports_range option is configured.
	The amount of data sent by the task. Unit: bytes.
	<ul> <li>The duration of the task. The task starts when the local server sends data and ends when the local server receives the final response from the peer end. Unit: µs.</li> </ul>
	<ul> <li>The minimum TCP RTT for the task. Unit: μs.</li> </ul>
	The number of TCP segments retransmitted by the task.
Р	<ul> <li>The sequence number of the task. After a TCP connection is established, the sequence number of the first task is 1.</li> </ul>
	<ul> <li>The service time of the task. It is the period that starts when the request is sent and ends when the first response is received. Unit: μs.</li> </ul>
	<ul> <li>The amount of time during which the response for the task are received, which is from the time when the first response packet is received to the time when the final response packet is received. Unit: µs.</li> </ul>
	• The total size of the response packets received by the task. Unit: bytes.
	<ul> <li>Specifies whether unordered reception occurs during the task. A value of 1 indicates that unordered reception occurs, and a value of 0 indicates that unordered reception did not occur.</li> </ul>
	<ul> <li>The TCP MSS used during the task. Unit: bytes.</li> </ul>

Scenario record identifier	Parameter
	An E record records information of a closed TCP connection. Each closed TCP connection has an E record. TCP connections that have the pports or pports_range option configured also have this record.
	The sequence number of the final task.
	• The amount of data sent during the TCP connection lifecycle. Unit: bytes.
Е	<ul> <li>The amount of data for which the TCP connection has sent a response but has not received the ACK message. A value of 0 indicates that no such data exists. Unit: bytes.</li> </ul>
	<ul> <li>The amount of data received during the TCP connection lifecycle. Unit: bytes.</li> </ul>
	• The number of TCP segments retransmitted by the task.
	$\circ~$ The minimum TCP RTT in the TCP lifecycle. Unit: $\mu s.$
	An N record records information about a TCP connection that is closed when the task is requesting response segments. Each TCP connection may have one N record or none.
	The sequence number of the final task.
	<ul> <li>The duration of the final task. Only the receiving time is available because the connection is closed. Unit: μs.</li> </ul>
N	<ul> <li>The amount of data received during the TCP connection lifecycle. Unit: bytes.</li> </ul>
	<ul> <li>Specifies whether unordered reception occurs during the task. A value of 1 indicates that unordered reception occurs, and a value of 0 indicates that unordered reception did not occur.</li> </ul>
	The TCP MSS used during the task. Unit: bytes.
	A W record records information about a TCP connection that is closed when the task is responding to request segments. Each TCP connection may have one W record or none.
	The amount of data sent in the response of the final task. Unit: bytes.
	• The duration of the final task. The receiving time is incomplete because the connection is closed. Unit: µs.
	<ul> <li>The minimum TCP RTT for the final task. Unit: μs.</li> </ul>
W	The number of TCP segments retransmitted by the final task.
	The sequence number of the final task.
	<ul> <li>The latency of the final task. Unit: μs.</li> </ul>
	<ul> <li>The latency for sending request segments of the final task. Unit: μs.</li> </ul>
	<ul> <li>The amount of data for which the final task has sent a response but has not received the ACK message. A value of 0 indicates that no such data exists. Unit: bytes.</li> </ul>

Note If the connection was closed normally, the <a href="mailto:snd\_nxt">snd\_nxt</a> value is deducted from the amount of data sent by the final task. However, if the connection was closed in an abnormal manner, for example, if a reset packet is received after a task is complete and the connection is ended by responding to <a href="mailto:reset">reset</a>, the actual amount of data sent is one byte larger than the displayed value. Typically, parameter values related to the amount of data sent by a task are correct, and the margin of error is only one byte.

stats file format

The following section describes the parameters from left to right as recorded:

- The timestamp encoded at the time the record was generated.
- The reserved field. Valid value: all .
- o The port number.
- o The average value of the total time taken by tasks recorded in R records.
- The average latency of tasks recorded in R records.
- The number of packets lost per 1,000 packets.
- The average value of RTT. Unit: μs.
- The number of tasks closed per 1,000 tasks in the data sent by requests.
- o The average size of data sent by tasks.
- The average latency for sending requests of tasks.
- The average amount of data received by tasks.
- The number of tasks.

#### Operations

You can perform the following operations to use TCP-RP:

1. Load and configure a module.

You can use one of the following methods to load and configure a module:

o Configure the parameters when you load the module.

```
Example: modprobe tcp_rt lports=80 .
```

- Load the module and then configure the parameters.
  - a. Run the following command to load the module:

```
modprobe tcp_rt
```

b. After the module is loaded, run a command and configure the parameters in the /sys/modul e/tcp\_rt/parameters/ directory. The following table describes the parameters and corresponding command configurations.

Parameter	Description	Default value	Command configuration
stats	Specifies whether to generate stats files. A value of 0 indicates that stats files are generated, and a value of 1 indicates that stats files are not generated.	0	echo 0 > stats
stats_interval	The interval at which to generate stats files. Unit: seconds.	60	echo 60 > stats_inter
lports	The local server ports for collecting data. A maximum of six ports can be specified.	None	echo 80,800,8080 > lp orts
pports	The peer port of the TCP connection.	None	echo 80,800,8080 > pp orts
lports_range	The local server port range. Set a start port number and an end port number to specify a port range. Examples: port 80 to port 200, and port 1000 to port 2000.	None	echo 80,100,1000,2000 >lports_range
pports_range	The peer port range of the TCP connection. Set a start port number and an end port number to specify a port range. Examples: port 80 to port 200, and port 1000 to port 2000.	None	echo 80,100,1000,2000 >pports_range
log_buf_num	The number of buffers. The maximum size of a log file is log_buf_num * 25 6 k . This parameter can be configured only when you load the module.	8	<pre>modprobe tcp_rt log_b uf_num=10</pre>
stats_buf_num	The number of buffers. The maximum size of a stats file is stats_buf_n um * 16 k . This parameter can be configured only when you load the module.	8	<pre>modprobe tcp_rt stats _buf_num=10</pre>

## 2. Uninst all the module.

i. Run the following command to disable the top-rt module to ensure that no new connections use TCP-RT:

```
echo 1 > /sys/kernel/debug/tcp-rt/deactivate
```

ii. Run the following command to ensure that no connections are using the tcp-rt module:

```
If Used by in the tcp-rt module is displayed as 0 , no connections are using the tcp-rt module.
```

iii. When no connections are using the tcp-rt module, run the following command to uninstall the tcp-rt module:

```
rmmod tcp_rt
```

# 5.6.8. Change the TCP TIME-WAIT timeout period

In the Linux kernel, TCP/IP connections stay in the TIME-WAIT state for 60 seconds. This period cannot be changed. However, in some scenarios such as heavy TCP loads, network performance can be improved if this period is shortened. In this context, Alibaba Cloud Linux 2 provides a kernel interface to change the length of the TCP TIME-WAIT timeout period in the kernel of the 4.19.43-13.al7 or later version. This topic describes how to use the kernel interface.

#### Context

The TIME-WAIT state is a mechanism in TCP/IP stacks that keeps sockets open after an application shuts down the sockets. By default, this state lasts for 60 seconds to ensure complete data transmission between the server and the client. If a large number of connections are in the TIME-WAIT state, network performance may be compromised. To resolve this issue, Alibaba Cloud Linux 2 provides an interface to change the TIME-WAIT timeout period to improve network performance in high-concurrency scenarios. The value range of this interface is 1 to 600 seconds. The default value of the TIME-WAIT timeout period is 60 seconds.

#### Limits

A timeout period of less than 60 seconds may violate the TCP/IP quiet time restriction and may cause old data to be accepted as new or duplicated new data rejected as old. Therefore, we recommend that you adjust the TIME-WAIT timeout period based on the advice of Alibaba Cloud technicians. For more information about the TCP/IP quiet time, visit IETF RFC 793.

# Configuration methods

You can use one of the following methods to change the TIME-WAIT timeout period. In both methods, the [\$TIME VALUE] parameter specifies the new timeout period that you want to configure.

• Run the sysctl command to change the TIME-WAIT timeout period:

```
sysctl -w "net.ipv4.tcp_tw_timeout=[$TIME_VALUE]"
```

• Run the echo command as the root user to change the TIME-WAIT timeout period in the /proc/s ys/net/ipv4/tcp\_tw\_timeout interface.

```
echo [$TIME_VALUE] > /proc/sys/net/ipv4/tcp_tw_timeout
```

For example, in scenarios where a large number of short-lived connections exist, such as when NGINX is configured with a layer-7 proxy, we recommend that you change the TIME-WAIT timeout period to 5 seconds. Run one of the following commands to modify the timeout period:

```
Note You can run the netstat -ant | grep TIME_WAIT | wc -1 command to check
whether a large number of short-lived connections exist in the server.

sysctl -w "net.ipv4.tcp_tw_timeout=5"
echo 5 > /proc/sys/net/ipv4/tcp_tw_timeout
```

# 5.6.9. Enable the cgroup writeback feature

Alibaba Cloud Linux 2 supports the cgroup writeback feature for the cgroup v1 kernel interface in the 4.19.36-12.al7 kernel version and later. This feature allows you to limit the buffered I/O rate when you use the cgroup v1 kernel interface.

#### Context

cgroups (abbreviated from control groups) are a Linux kernel mechanism that organizes processes hierarchically and distributes system resources along the hierarchy in a controlled and configurable manner. Two versions of cgroups are available: cgroup v1 and cgroup v2. For more information, visit What are control groups. This topic describes how to enable the cgroup writeback feature for cgroup v1 to limit the buffered I/O rate of processes.

#### Limits

After you enable the cgroup writeback feature, check whether the mappings between memory subsystems (memcgs) and I/O subsystems (blkcgs) conform to the following rule. If the mappings conform to the rule, limit the buffered I/O rate of processes.

The cgroup writeback feature requires cooperation between memcgs and blkcgs to limit the buffered I/O rate. However, by default, the control subsystems of cgroup v1 do not work together. Memcgs and blkcgs must be associated based on the rule that each memcg is mapped to a unique blkcg. The mappings from memcgs to blkcgs can be one-to-one or many-to-one, but cannot be one-to-many or many-to-many.

For example, to limit the buffered I/O rate of Processes A and B, you must take note of the following items:

- When A and B belong to two different memcgs, each of the memcgs can be mapped to a unique blkcg. For example, A belongs to memcg1 and blkcg1 , and B belongs to memcg2 and blkcg0.
- When A and B belong to two different memcgs, these memcgs can also be mapped to the same blkcg. For example, A belongs to memcg1 and B belongs to memcg2, and both A and B belong to blkcg2.
- When A and B belong to the same memcg, the memcg can be mapped only to one blkcg. For example, both A and B belong to memcg0 and blkcg3.

To avoid exceptions, we recommend that you perform the following operations after you enable the cgroup writeback feature but before you limit the buffered I/O rate of a process: Configure the cgroup.procs interface for the corresponding blkcg, and write the ID of the process to the interface

to ensure that the memcgs mapped to this blkcg are not mapped to any other blkcgs. You can also use a tool to view the mappings between memcgs and blkcgs. For more information, see Verify the mappings from memcgs to blkcgs.

During O&M, processes may move to different cgroups. Based on the preceding rule, no exceptions occur if processes move between memcgs. However, If processes move between blkcgs, exceptions may occur. The following rule is defined in the code of the cgroup writeback feature to avoid exceptions: If a process in a running blkcg moves to a different blkcg, the memcg to which the process belongs is mapped to the root blkcg. Typically, no throttling threshold is set for the root blkcg. When the memcg is mapped to the root blkcg, the throttling feature no longer takes effect.

Notice Although the preceding rule is defined in the kernel code, we recommend that you do not move processes between blkcgs.

#### Enable cgroup writeback

By default, the cgroup writeback feature is disabled for the cgroup v1 interface. To enable this feature, perform the following steps:

1. Add the  $cgwb\_v1$  field to the grubby command and run the command to enable the cgroup writeback feature.

In this example, the kernel version is 4.19.36-12.a17.x86\_64. You must enter your actual kernel version during the operation. To query your kernel version, run the uname -a command.

```
sudo grubby --update-kernel="/boot/vmlinuz-4.19.36-12.al7.x86_64" --args="cgwb_v1"
```

2. Restart the system for the cgroup writeback feature to take effect.

```
sudo reboot
```

3. Run the following command to read the /proc/cmdline kernel file. Ensure that the parameters of the command include the cgwb\_v1 field. This way, the blkio.throttle.write\_bps\_device and blkio.throttle.write\_iops\_device interfaces in the corresponding blkcgs can limit the buffered I/O rate.

```
cat /proc/cmdline | grep cgwb_v1
```

# Verify the mappings from memcgs to blkcgs

Before you limit the buffered I/O rate of processes, you can use one of the following methods to check whether the mappings from memcgs to blkcgs are one-to-one or many-to-one.

Run the following command to view the mappings from memcgs to blkcgs:

```
sudo cat /sys/kernel/debug/bdi/bdi_wb_link
```

The following sample response shows that the mapping from the memcg to the blkcg is one-to-one.

```
memory <---> blkio
memcg1: 35 <---> blkcg1: 48
```

- Use the ftrace kernel monitoring tool.
  - i. Use the ftrace tool.

```
\label{lem:condition} $$ sudo bash -c "echo 1 > /sys/kernel/debug/tracing/events/writeback/insert_memcg_blkcg_link/enable"
```

ii. View the output information.

```
sudo cat /sys/kernel/debug/tracing/trace_pipe
```

The following sample response contains <code>memcg\_ino=35 blkcg\_ino=48</code> , which indicates that the mapping from the memcg to the blkcg is one-to-one.

```
<...>-1537 [006] .... 99.511327: insert_memcg_blkcg_link: memcg_ino=35 blkcg_ino=48 old_blkcg_ino=0
```

## Check whether cgroup writeback takes effect

In this example, two processes that generate I/O are simulated to check whether the cgroup writeback feature takes effect.

### ? Note

- The dd command responds quickly and the screen rolls too fast to be viewed. Run the state command to view the dd command output.
- The dd command writes data in sequence. The system performs sequential I/O refreshing every time 1 MB of output data is generated. Therefore, you must set the threshold for blkio.throttle.write\_bps\_device to a value of no less than 1 MB (1,048,576 bytes). If you set the threshold for blkio.throttle.write\_bps\_device to a value of less than 1 MB, I/O hangs may occur.
- 1. Simulate two processes that generate I/O, and set the cgroup.procs interface of the blkcg
  based on the preceding limits.

```
sudo mkdir /sys/fs/cgroup/blkio/blkcg1
sudo mkdir /sys/fs/cgroup/memory/memcg1
sudo bash -c "echo $$ > /sys/fs/cgroup/blkio/blkcg1/cgroup.procs"  # $$ specifies the process ID.
sudo bash -c "echo $$ > /sys/fs/cgroup/memory/memcg1/cgroup.procs"  # $$ specifies the process ID.
```

2. Use the blkio.throttle.write bps device interface in the blkcg to limit the buffered I/O rate.

```
sudo bash -c "echo 254:48 10485760 > /sys/fs/cgroup/blkio/blkcg1/blkio.throttle.write_b
ps_device"  # Set the writeback throttling threshold for the disk to 10 MB/s based on
the device number.
```

3. Use the dd command that does not have oflag set to sync to generate buffered I/O.

```
sudo dd if=/dev/zero of=/mnt/vdd/testfile bs=4k count=10000
```

4. Use the iostat tool to query the results. View the wmb/s column in the command output. If the value is 10 MB/s, the cgroup writeback feature takes effect.

```
iostat -xdm 1 vdd
```

# 5.6.10. Configure the weight-based throttling feature of blk-iocost

Alibaba Cloud Linux 2 provides the weight-based throttling feature (blk-iocost) based on the cost model in the 4.19.81-17.al7.x86\_64 kernel version and later. blk-iocost improves the weight-based disk throttling feature of the I/O subsystem (blkcg) within a kernel. Both cgroup v1 and cgroup v2 interfaces support blk-iocost in Alibaba Cloud Linux 2. This topic describes the interfaces that implement weight-based throttling.

# Interface description

Interface	Description	Configuration item

Interface	Description	Configuration item
cost.qos	A read/write interface whose file is stored only in the root group of the blkcg. The full name of the file is blkio.cost.qos in cgroup v1 and io.cost.qos in cgroup v2. This interface provides the blk-iocost feature and limits the I/O quality of service (QoS) rate based on the latency weight.  After blk-iocost is enabled, the kernel calculates the proportion of requests that exceed the read or write latency threshold (rlat wlat) out of all requests. When the proportion is greater than the read or write latency percentile (rpct wpct), the kernel considers the disk to be saturated and reduces the rate at which requests are sent to the disk. By default, the value of rlat wlat is set to 0, which indicates that the blk-iocost feature is disabled.	Each line of configuration in the interface file starts with the major number (MAJ) and minor number (MIN) of the disk in the MAJ:MIN format, followed by the following configuration items:  • enable: specifies whether to enable the blk-iocost controller to enable the blk-iocost controller to enable the blk-iocost is disabled. A value of 1 indicates that blk-iocost is disabled. A value of 1 indicates that blk-iocost is enabled.  • ctrl: the control mode. Valid values: auto and user. When ctrl is set to auto, the kernel automatically detects the disk category and uses built-in parameters. When ctrl is set to user, you must specify the following QoS control parameters:  • rpct: the read latency percentile. Valid values: 0 to 100.  • rlat: the read latency threshold. Unit: microseconds.  • wpct: the write latency percentile. Valid values: 0 to 100.  • wlat: the write latency threshold. Unit: microseconds.  • min: the minimum scaling percentage. Valid values: 1 to 10000.  • max: the maximum scaling percentage. Valid values: 1 to 10000.

Interface	Description	Configuration item
cost.model	A read/write interface whose file is stored only in the root group of the blkcg. The full name of the interface file is  blkio.cost.model in cgroup v1 and io.cost.model in cgroup v2. The interface is used to set the cost model.	Each line of configuration in the interface file starts with the major number (MAJ) and minor number (MIN) of the disk in the MAJ:MIN format, followed by the following configuration items:  • ctrl: the control mode. It specifies whether to allow the user to enter model parameters. Valid values: au to and user.  • model: the model parameter. Valid value: linear. You must define the following modeling parameters when model is set to linear:  • [r w]bps: the maximum sequential I/O throughput.  • [r w]seqiops: the maximum sequential input/output operations per second (IOPS).  • [r w]randiops: the maximum random IOPS.   **Onte You can use the tools/cgroup/iocost_coef_gen.py script in the kernel source code to generate the preceding parameters and then write the parameters to the cost.model interface file to configure the cost model.

Interface	Description	Configuration item
cost.weight	A read/write interface whose file is stored only in the sub-group of the blkcg. The full name of the interface file is  blkio.cost.weight in cgroup v1 and  io.cost.weight in cgroup v2. This interface is used to set the weight of a sub-group.  Default value: 100. Valid values: 1 to 10000. The interface can be used to set a weight for each disk or change the default weight of a sub-group.	<ul> <li><weight> : the default weight of the blkcg.</weight></li> <li>MAJ:MIN <weight> : the weight of the blkcg on the disk specified by MAJ: MIN.</weight></li> </ul>

#### **Precautions**

If ctrl is set to auto when you use the blk-iocost feature, you must set the rotational attribute to 0 for ultra disks, standard SSDs, enhanced SSDs, or local NVMe SSDs.

#### Example 1

Use *cost.qos* to enable the blk-iocost feature for the 254:48 disk. If more than 95% of requests have a latency (rlat|wlat) greater than 5 milliseconds, the disk is considered to be saturated. The kernel adjusts the rate at which requests are sent to the disk within the interval from 50% to 150% of the original rate. The following commands are used for the cgroup v1 and cgroup v2 interfaces:

• The command for cgroup v1:

```
echo "254:48 enable=1 ctrl=user rpct=95.00 rlat=5000 wpct=95.00 wlat=5000 min=50.00 max=1 50.00" > /sys/fs/cgroup/blkio/blkio.cost.qos
```

• The command for cgroup v2:

```
echo "254:48 enable=1 ctrl=user rpct=95.00 rlat=5000 wpct=95.00 wlat=5000 min=50.00 max=1 50.00" > /sys/fs/cgroup/io.cost.qos
```

## Example 2

Use  $_{\text{cost.model}}$  to configure a model on the  $_{254:48}$  disk based on the modeling parameters set when model is  $_{\text{linear}}$ . The following commands are used for the cgroup v1 and cgroup v2 interfaces:

• The command for cgroup v1:

```
echo "254:48 ctrl=user model=linear rbps=2706339840 rseqiops=89698 rrandiops=110036 wbps= 1063126016 wseqiops=135560 wrandiops=130734" > /sys/fs/cgroup/blkio/blkio.cost.model
```

• The command for cgroup v2:

echo "254:48 ctrl=user model=linear rbps=2706339840 rseqiops=89698 rrandiops=110036 wbps= 1063126016 wseqiops=135560 wrandiops=130734" > /sys/fs/cgroup/io.cost.model

### Example 3

Use cost.weight to change the default weight of blkcg1 to 50 and set the weight of blkcg1 on the 254:48 disk to 50. The following commands are used for the cgroup v1 and cgroup v2 interfaces:

• The command for cgroup v1:

```
echo "50" > /sys/fs/cgroup/blkio/blkcg1/blkio.cost.weight # Change the default weight to 50.
echo "254:48 50" > /sys/fs/cgroup/blkio/blkcg1/blkio.cost.weight # Set the weight on the 254:48 disk to 50.
```

• The command for cgroup v2:

```
echo "50" > /sys/fs/cgroup/cgl/io.cost.weight  # Change the default weight to 50. echo "254:48 50" > /sys/fs/cgroup/cgl/io.cost.weight  # Set the weight on the 254:48 di sk to 50.
```

### Common monitoring tools

• iocost monitor script

The tools/cgroup/iocost\_monitor.py script in the kernel source code uses the drgn debugger to obtain kernel parameters and then provide the I/O performance monitoring data. For more information about drgn, visit drgn. The script is used in the following manner:

Run the following command to monitor the I/O performance of the *vdd* disk.

```
./iocost_monitor.py vdd
```

A command output similar to the following one is returned:

```
      vdd RUN
      per=500.0ms
      cur_per=3930.839:v14620.321
      busy= +1
      vrate=6136.22%
      params=hdd

      active
      weight
      hweight%
      inflt%
      dbt
      delay usages%

      blkcg1
      * 50/ 50 9.09/ 9.09 0.00 0 0*000 009:009:009

      blkcg2
      * 500/ 500 90.91/ 90.91 0.00 0 0*000 089:091:092
```

• The blkio.cost.statcost.stat interface file of cgroup v1

The Alibaba Cloud Linux 2 kernel provides the blk-iocost interface file (blkio.cost.statcost.stat) of the cgroup v1 interface. This interface file records the QoS data of each controlled device. Run the following command to view the interface file:

```
cat /sys/fs/cgroup/blkio/blkcg1/blkio.cost.stat
```

#### Example out put:

```
254:48 is_active=1 active=50 inuse=50 hweight_active=5957 hweight_inuse=5957 vrate=159571
```

ftrace monitoring tool

178

The Alibaba Cloud Linux 2 kernel provides the ftrace tool related to blk-iocost for kernel-side analytics. The ftrace monitoring tool is used in the following manner:

i. Set the enable attribute to 1 to enable the ftrace tool.

```
echo 1 > /sys/kernel/debug/tracing/events/iocost/enable
```

ii. View the output information.

```
cat /sys/kernel/debug/tracing/trace_pipe
```

#### Example out put:

```
dd-1593 [008] d... 688.565349: iocost iocq activate: [vdd:/blkcq1] now=6890652
89:57986587662878 vrate=137438 period=22->22 vtime=0->57986365150756 weight=50/50 hwe
ight=65536/65536
   dd-1593 [008] d.s. 688.575374: iocost_ioc_vrate_adj: [vdd] vrate=137438->13743
8 busy=0 missed ppm=0:0 rq wait pct=0 lagging=1 shortages=0 surpluses=1
           [008] d.s. 688.608369: iocost ioc vrate adj: [vdd] vrate=137438->13743
8 busy=0 missed ppm=0:0 rq wait pct=0 lagging=1 shortages=0 surpluses=1
   dd-1594 [006] d... 688.620002: iocost_iocg_activate: [vdd:/blkcg2] now=6891199
46:57994099611644 vrate=137438 period=22->26 vtime=0->57993412421644 weight=250/250 h
weight=65536/65536
<idle>-0
           [008] d.s. 688.631367: iocost ioc vrate adj: [vdd] vrate=137438->13743
8 busy=0 missed ppm=0:0 rq wait pct=0 lagging=1 shortages=0 surpluses=1
<idle>-0 [008] d.s. 688.642368: iocost ioc vrate adj: [vdd] vrate=137438->13743
8 busy=0 missed_ppm=0:0 rq_wait_pct=0 lagging=1 shortages=0 surpluses=1
<idle>-0 [008] d.s. 688.653366: iocost ioc vrate adj: [vdd] vrate=137438->13743
8 busy=0 missed_ppm=0:0 rq_wait_pct=0 lagging=1 shortages=0 surpluses=1
<idle>-0 [008] d.s. 688.664366: iocost ioc vrate adj: [vdd] vrate=137438->13743
8 busy=0 missed ppm=0:0 rq wait pct=0 lagging=1 shortages=0 surpluses=1
```

# 5.6.11. Enhance the monitoring of block I/O throttling

Alibaba Cloud Linux 2 provides interfaces in the 4.19.81-17.al7 kernel version and later to better monitor Linux block I/O throttling. This topic describes how to use the interfaces.

#### Context

Linux block I/O throttling (bit/s or IOPS) is required in multiple scenarios, especially those where cgroup writeback is enabled. You can perform I/O throttling related operations in a more convenient manner if block I/O throttling is well monitored. In this context, Alibaba Cloud Linux 2 provides interfaces to enhance the monitoring of block I/O throttling.

# Interface description

Interface	Description
blkio.throttle.io_service_time	The total amount of time between request dispatch and request completion for I/O operations. Unit: nanoseconds.
blkio.throttle.io_wait_time	The total amount of time that the I/O operations wait in scheduler queues. Unit: nanoseconds.

Interface	Description
blkio.throttle.io_completed	The number of completed I/O operations. It is used to calculate the average latency of the block I/O throttling layer.
blkio.throttle.total_io_queued	The number of I/O operations that were throttled. The number of I/O operations that were throttled in the current cycle can be calculated based on periodic monitoring data and be used to analyze whether I/O latency is related to throttling.
blkio.throttle.total_bytes_queued	The total bytes of I/O that were throttled. Unit: bytes.

The path of the preceding interfaces is /sys/fs/cgroup/blkio/<cgroup>/, where <cgroup> is the control group.

## Example

You can obtain the average I/O latency of a disk by using the interface that enhances the monitoring of block I/O throttling. In this example, the average I/O write latency of the *vdd* disk is monitored at an interval of five seconds. Then, the average I/O latency of the *vdd* disk is calculated. The following table describes the relevant parameters.

Parameter	Description
write_wait_time <n></n>	Obtains the duration of throttling at the block I/O throttling layer.
write_service_time <n></n>	Obtains the total amount of time between request dispatch and request completion for I/O operations.
write_completed <n></n>	Obtains the number of completed I/O operations.

1. Obtain the monitoring data at the T1 time.

```
write_wait_time1 = `cat /sys/fs/cgroup/blkio/blkcg1/blkio.throttle.io_wait_time | grep
-w "254:48 Write" | awk '{print $3}'`
write_service_time1 = `cat /sys/fs/cgroup/blkio/blkcg1/blkio.throttle.io_service_time |
grep -w "254:48 Write" | awk '{print $3}'`
write_completed1 = `cat /sys/fs/cgroup/blkio/blkcg1/blkio.throttle.io_completed | grep
-w "254:48 Write" | awk '{print $3}'`
```

2. Wait five seconds and obtain the monitoring data at the T2 (T1 + 5s) time.

```
write_wait_time2 = `cat /sys/fs/cgroup/blkio/blkcgl/blkio.throttle.io_wait_time | grep
-w "254:48 Write" | awk '{print $3}'`
write_service_time2 = `cat /sys/fs/cgroup/blkio/blkcgl/blkio.throttle.io_service_time |
grep -w "254:48 Write" | awk '{print $3}'`
write_completed2 = `cat /sys/fs/cgroup/blkio/blkcgl/blkio.throttle.io_completed | grep
-w "254:48 Write" | awk '{print $3}'`
```

3. Calculate the average I/O latency during the five seconds based on the following formula:

Average I/O latency = (Total I/O duration at the T2 time - Total I/O duration at the T1 time)/(Number of completed I/O operations at the T2 time - Number of completed I/O operations at the T1 time).

```
avg_delay = `echo "((write_wait_time2 + write_service_time2) - (write_wait_time1+write_
service_time1)) / (write_completed2 - write_completed1)" | bc`
```

# 5.6.12. Use the JBD2 optimization interface

JBD2 is the kernel thread of the ext4 file system. It often experiences the shadow (BH\_Shadow) state during its use, which can affect the system performance. To solve this problem, Alibaba Cloud Linux 2 provides an interface in version 4.19.81-17.al7 and later of the kernel to optimize JBD2. This topic describes the interface.

## Context

Ext4 is one of the most common journaling file systems. JBD2 is the kernel thread of ext4 for updating journals and is a global resource for the ext4 file system. When the JBD2 kernel thread attempts to obtain access permissions from the cache, the cache page may be in the BH\_Shadow state. In this case, JBD2 waits an extended period of time for the cache page to write back to the disk, which can affect the system performance. To solve this problem, Alibaba Cloud Linux 2 provides the force\_copy kernel interface to optimize JBD2. The system forcibly copies the cache page to reduce the amount of time that JBD2 waits for the cache page in the BH\_Shadow state to write back to the disk. In addition, Alibaba Cloud Linux 2 provides the stats interface to help analyze quality of service (QoS) issues related to the file system.

# Interface description

Interface	Description
force_copy	The interface file is stored in /proc/fs/jbd2/ <device>-8/force_copy, where the device variable specifies the name of the Elastic Block Storage (EBS) device. After you enable the force_copy interface, the system forcibly copies data to reduce the waiting time of JBD2.</device>
	Notice The interface consumes memory to run.
stats	The interface file is stored in /proc/fs/jbd2/ <device>-8/stats. The interface helps determine whether QoS issues in the file system are caused by JBD2.</device>

# **Examples**

The following examples demonstrate how to implement the force copy and stats interfaces:

• By default, the force\_copy interface is disabled. You can set the value of the interface to 1 to enable the interface or set the value to 0 to disable the interface.

echo 1 > /proc/fs/jbd2/nvme0n1-8/force\_copy # Call the interface.

• Run the following command to query the stats interface:

cat /proc/fs/jbd2/nvme0n1-8/stats

A command output similar to the following one is returned:

337 336 65536 0 14837 1701504 16 0 20058 5 33082732 605 942 1000 1000

The following table describes the fields in the preceding sample output.

Field	Description
The first field	The ID of the transaction.
The second field	The number of transactions requested.
The third field	The maximum number of cached transactions.
The fourth field	The transaction wait time.
The fifth field	The latency of the transaction request.
The sixth field	The amount of time that the transaction ran.
The seventh field	The amount of time that the transaction was locked.
The eighth field	The amount of time it took to refresh the transaction.
The ninth field	The transaction logging time.
The tenth field	The average transaction commit time.
The eleventh field	The number of handles contained in the transaction.
The twelfth field	The number of blocks contained in the transaction.
The thirteenth field	The number of blocks recorded for the transaction.
The fourteenth field	The time constant of the kernel configuration, in Hertz.
The fifteenth field	The period of the time constant of the kernel configuration in milliseconds.

# 5.6.13. Create a hard link across project quotas

By default, the ext4 file system has constraints, and you are not allowed to create hard links across project quotas. However, some scenarios require hard links to be created. Alibaba Cloud Linux 2 provides a custom interface that can bypass the constraints of the ext4 file system to create hard links across project quotas. This topic describes the interface and provides an example on how to use the interface.

### Context

Linux distributions support the following disk quot a modes: user quota, group quota, and project quota. Project quotas provide more fine-grained disk quotas than user or group quotas. Project quotas identify directories and files within file systems by project ID. This topic describes how to create a hard link across project ID directories in an ext4 file system.

# Interface description

The default value of the /proc/sys/fs/hardlink\_cross\_projid interface is 0. In this case, hard links cannot be created across project quotas. If the /proc/sys/fs/hardlink\_cross\_projid interface is set to 1, you can bypass the constraints of the ext4 file system to create hard links across project quotas.

For more information about the interface, see <code>Documentation/sysctl/fs.txt</code>. You can obtain the kernel document from the Debuginfo package and the source code package of Alibaba Cloud Linux 2. For more information, see <code>Use Alibaba Cloud Linux 2</code>.

# Example

You can run the following command to query the value of the /proc/sys/fs/hardlink\_cross\_projid interface:

```
cat /proc/sys/fs/hardlink_cross_projid
```

A value of 0 is returned, which indicates that hard links cannot be created across project quotas.

You can change the value from 0 to 1 so that you can create hard link across project quotas.

```
echo 1 > /proc/sys/fs/hardlink cross projid
```

# 5.6.14. Track I/O latency

Alibaba Cloud Linux 2 optimizes the /proc/diskstats raw data source interface of the I/O latency analysis tool iostat. Alibaba Cloud Linux 2 can calculate the amount of time consumed by read, write, and special I/O (discard) operations on devices. In addition, Alibaba Cloud Linux 2 provides the bcc tool to track I/O latency. This topic describes the optimized /proc/diskstats interface and the bcc tool.

# Interface description

The /proc/diskstats interface in Alibaba Cloud Linux 2 allows you to query the I/O information on a disk and the amount of time consumed by read, write, and discard operations on a device.

Example: Query the /proc/diskstats interface by using the root account.

```
cat /proc/diskstats
```

Example output:

254 0 vda 6328 3156 565378 2223 1610 424 25160 4366 0 1358 5332 0 0 0 0 2205 3347 0

In the command output, the last three domains are new domains added in Alibaba Cloud Linux 2. The following table describes the three domains.

Domain	Description
The sixteenth domain	The amount of time consumed by read operations on the device. Unit: ms.
The seventeenth domain	The amount of time consumed by write operations on the device. Unit: ms.
The eighteenth domain	The amount of time consumed by discard operations on the device. Unit: ms.

Note For information about other domains, see the kernel document Documentation/iostats.txt. You can obtain the kernel document from the Debuginfo package and source code package of Alibaba Cloud Linux 2. For more information, see Use Alibaba Cloud Linux 2.

## bcc

Alibaba Cloud Linux 2 provides the bcc tool to track I/O latency. You must download the tool before you can use it. Run the following command to download the tool:

```
yum install -y bcc-tools
```

You can run one of the following commands to query the description of the bcc tool:

• Run the following command to query the description of the bcc tool:

/usr/share/bcc/tools/alibiolatency -h

Sample description:

```
usage: alibiolatency [-h] [-d DEVICE] [-i [DIS INTERVAL]]
                  [-t [AVG THRESHOLD TIME]] [-T [THRESHOLD TIME]] [-r]
Summarize block device I/O latency
optional arguments:
             show this help message and exit
 -h, --help
 -d DEVICE, --device DEVICE
                     inspect specified device
 -i [DIS_INTERVAL], --dis_interval [DIS_INTERVAL]
                      specify display interval
 -t [AVG THRESHOLD TIME], --avg threshold time [AVG THRESHOLD TIME]
                     display only when average request process time is
                      greater than this value
 -T [THRESHOLD TIME], --threshold time [THRESHOLD TIME]
                      dump request life cycle when single request process
                     time is greater than this value
 -r, --dump_raw
                     dump every io request life cycle
examples:
                          # summarize block I/O latency(default display interval is 2s
   ./alibiolatency
   ./alibiolatency -d sda3 # inspect specified device /dev/sda3
   ./alibiolatency -i 2  # specify display interval, 2s
   ./alibiolatency -t 10  # display only when average request process time is greater
than 10ms
   ./alibiolatency -T 20  # dump request life cycle when single request process time i
s greater than 20ms
   ./alibiolatency -r
                          # dump every io request life cycle
```

• Run the man command to query the description of the bcc tool:

man bcc-alibiolatency

# 5.6.15. Detect I/O hangs of file systems and block layers

An I/O hang occurs when the system becomes unstable or even goes down due to time-consuming I/O requests. To accurately detect I/O hangs, Alibaba Cloud Linux 2 extends the core data structure and provides a feature to locate and detect I/O hangs at low system overheads. This topic describes the interfaces for this feature and the usage examples of these interfaces.

# Interface description

Interface	Description
/sys/block/ <device>/queue/hang_threshold</device>	The interface allows you to query and modify the threshold for I/O hangs. Unit of the threshold: ms. Default value: 5000.
/sys/block/ <device>/hang</device>	The interface allows you to query the number of I/O operations that exceed the threshold for I/O hangs on the device.

Interface	Description
/sys/kernel/debug/block/ <device>/rq_hang</device>	The interface allows you to query details about I/O hangs.
/proc/ <pid>/wait_res</pid>	The interface allows you to query information about the resources for which the process is waiting.
/proc/ <pid>/task/<tid>/wait_res</tid></pid>	The interface allows you to query information about the resources for which the thread is waiting.

### The following table describes variables in the preceding interfaces.

Variable	Description
<device></device>	The name of the Elastic Block Storage (EBS) device.
<pid></pid>	The ID of the process.
<tid>&gt;</tid>	The ID of the thread.

# Example 1

You can call the  $\space{-10mm}\space{-10mm$ 

1. Change the threshold for I/O hangs of the *vdb* disk to 10,000 ms.

```
echo 10000 > /sys/block/vdb/queue/hang_threshold
```

2. View the change result.

```
cat /sys/block/vdb/queue/hang threshold
```

### Example output:

10000

# Example 2

You can call the /sys/block/<device>/hang interface to query the number of I/O operations that have I/O hangs on a disk. The *vdb* disk is used in this example.

The following command is used:

```
cat /sys/block/vdb/hang
```

### Example output:

```
0 1 \# The value on the left indicates the number of read operations that have I/O hangs. The value on the right indicates the number of write operations that have I/O hangs.
```

# Example 3

You can call the  $/sys/kernel/debug/block/<device>/rq_hang$  interface to query the details of I/O hangs. The vdb disk is used in this example.

The following command is used:

```
cat /sys/kernel/debug/block/vdb/rq_hang
```

### Example output:

```
ffff9e50162fc600 { .op=WRITE, .cmd_flags=SYNC, .rq_flags=STARTED|ELVPRIV|IO_STAT|STATS, .sta
te=in_flight, .tag=118, .internal_tag=67, .start_time_ns=1260981417094, .io_start_time_ns=1
260981436160, .current_time=1268458297417, .bio = ffff9e4907c31c00, .bio_pages = { ffffc859
60686740 }, .bio = ffff9e4907c31500, .bio_pages = { ffffc85960639000 }, .bio = ffff9e4907c3
0300, .bio_pages = { ffffc85960651700 }, .bio = ffff9e4907c31900, .bio_pages = { ffffc85960
608b00 }}
```

The preceding output shows the details of an I/O operation. io\_start\_time\_ns indicates the start time of the I/O request and this parameter has an assigned value. This indicates that the I/O request was not processed in a timely manner, which leads to prolonged I/O time.

# Example 4

You can call the /proc/<pid>/wait\_res interface to query information about the resources for which a process is waiting. In this example, the 577 process is used.

The following command is used:

```
cat /proc/577/wait_res
```

### Example output:

The following table describes the parameters in the example output.

Parameter	Description
Field 1	The types of the resources for which the process is waiting. A value of 1 indicates the cache page in the file system. A value of 2 indicates the block I/O layer.
Field 2	The addresses of the resources (cache page or block I/O layer) for which the process is waiting.
Field 3	The time at which the process began waiting for resources.
Field 4	The current time when the file is being read. The difference between Field 4 and Field 3 is the amount of time taken by the process to wait for the resources.

# 5.6.16. Enable the PSI feature for cgroup v1

In the Linux kernel, only the cgroup v2 interface supports the Pressure Stall Information (PSI) feature. Alibaba Cloud Linux 2 supports the PSI feature for the cgroup v1 interface in the 4.19.81-17.al7 kernel version and later to allow you to monitor the CPU, memory, and I/O performance. This topic describes how to enable the PSI feature for the cgroup v1 interface and query relevant information.

### Context

PSI is a kernel feature that can be used to monitor the CPU, memory, and I/O performance. For more information about the PSI feature, see the Documentation/accounting/psi.txt kernel document. This kernel document is contained in the Debuginfo package and source code package of Alibaba CloudLinux 2. For information about how to download the Debuginfo package and source code package, see Overview.

# Enable the PSI feature for the cgroup v1 interface

By default, the PSI feature is disabled for the cgroup v1 interface. You can perform the following steps to enable the PSI feature:

1. Run the grubby command to modify the startup parameter.

The default value of the <code>args</code> parameter is <code>"psi=1"</code>, which indicates that the PSI feature is enabled for the cgroup v2 interface. Change the value of the parameter to <code>"psi=1 psi\_v1=1"</code> to enable the PSI feature for the cgroup v1 interface. In this example, the kernel version is <code>4.19.81-17.al7.x86\_64</code>. You must use your actual kernel version when you the operation. To query the kernel version, run the <code>uname -a</code> command.

```
sudo grubby --update-kernel="/boot/vmlinuz-4.19.81-17.a17.x86_64" --args="psi=1 psi_v1=
1"
```

2. Restart the system to apply the change.

```
sudo reboot
```

# Check whether the PSI feature is enabled for cgroup v1

After the system restarts, you can run the following command to check whether the PSI feature is enabled for the cgroup v1 interface in /proc/cmdline of the kernel:

```
cat /proc/cmdline | grep "psi=1 psi_v1=1"
```

# Query the monitoring data of the CPU, memory, and I/O performance

After the PSI feature is enabled for the cgroup v1 interface, this feature monitors the CPU, memory, and I/O performance and transmits all the monitoring data to the cpuacct controller. You can query detailed monitoring data by running the following commands:

```
cat /sys/fs/cgroup/cpuacct/cpu.pressure
cat /sys/fs/cgroup/cpuacct/memory.pressure
cat /sys/fs/cgroup/cpuacct/io.pressure
```

# 5.6.17. Enable the CPU burst feature for cgroup

# **v**1

Alibaba Cloud Linux 2 provides the CPU burst feature for the cgroup v1 interface in kernel versions 4.19.91-22.al7 and later. The CPU burst feature allows CPU-throttled containers to burst their CPU utilization for higher performance and lower latency. This topic describes how to enable the CPU burst feature for the cgroup v1 interface and how to query CPU burst statistics.

### Context

In container services, containers are CPU-throttled by the CFS bandwidth controller of the kernel. You can enable the CPU burst feature for the cgroup v1 interface to allow containers to burst their CPU usage above the specified limit. This feature can ensure the quality of containers without reducing the deployment density of containers.

Note A CPU burst of a container means that the container utilizes CPU resources beyond the specified limit.

For more information about CPU burst, see the Alibaba Cloud Linux 2 kernel documentation. The Alibaba Cloud Linux 2 kernel documentation is stored in the Debuginfo package and source code package. The path of the documentation is *Documentation/scheduler/sched-bwc.rst*. For information about how to download the Debuginfo package and source code package, see Use Alibaba Cloud Linux 2.

## How the CPU burst feature works

The amount of CPU resources that can be allocated by the CPU burst feature depends on daily CPU resources accumulated. For example, a container accumulates idle CPU resources when it runs below the specified CPU limit. If the container needs to use greater amounts of CPU resources, it uses the CPU burst feature to consume the accumulated CPU resources to burst its CPU usage. To better understand the CPU burst feature, consider the following analogy:

Assume that you are entitled to an annual vacation of four days (CPU throttling limit). You can accumulate up to four outstanding vacation days (CPU burst limit) for future use. If you take only one day off each year in the first and second years, you are theoretically entitled to have six vacation days outstanding. In the third year, you are theoretically entitled to a vacation of ten days. However, you can accumulate only four outstanding vacation days (CPU burst limit) in the past two years and can actually take up to eight days off in the third year due to the limit on the maximum number of accumulated vacation days.

# Enable the CPU burst feature

1. To check whether the CPU burst feature is globally enabled for the cgroup v1 interface, run the following command:

```
cat /proc/sys/kernel/sched_cfs_bw_burst_enabled
```

Description about the command output:

• If the CPU burst feature is globally enabled, 1 is returned.

Note By default, the CPU burst feature is globally enabled for the cgroup v1 interface.

• If the CPU burst feature is globally disabled, a value other than 1 is returned. To globally enable the feature, run the following command:

```
echo 1 > /proc/sys/kernel/sched_cfs_bw_burst_enabled
```

2. Enable the CPU burst feature for a specific child cgroup.

By default, the CPU burst feature is disabled for all child cgroups in the cgroup v1 interface. Perform the following operations:

i. To go to the path of a child cgroup and check whether the CFS bandwidth controller has taken effect, run the following commands:

```
cd /sys/fs/cgroup/cpu/<Name of the child cgroup directory>/
cat cpu.cfs_period_us
cat cpu.cfs_quota_us
```

If the CFS bandwidth controller has taken effect, a positive integer is returned for  $_{\rm cpu.cfs\_qu}$  ota\_us in the command output. In this case, you can enable the CPU burst feature to allow containers to utilize more CPU resources than the specified limit.

ii. To check the default CPU burst settings, run the following command:

```
cat cpu.cfs_burst_us
```

If the CPU burst feature is disabled, a default value of 0 is returned in the command output.

iii. Configure cpu.cfs burst us to enable the CPU burst feature.

You can set cpu.cfs\_burst\_us to a positive integer to enable the CPU burst feature and use this integer as the CPU burst limit. The following operations provide an example of how to enable the CPU burst feature:

a. Configure the default cpu.cfs\_quota\_us and cpu.cfs\_period\_us files for the CFS bandwidth controller.

cpu.cfs\_period\_us specifies the length of a period during which CPU resources are utilized. cpu.cfs\_quota\_us specifies the maximum amount of CPU resources that each child cgroup can utilize during each period specified by cpu.cfs\_period\_us. Run the following commands to set cpu.cfs\_period\_us to 100 ms and set cpu.cfs\_quota\_us to 400 ms to ensure that each child cgroup continues to obtain four CPU resources, which is calculated by using the following formula: cpu.cfs quota us / cpu.cfs period us .

```
echo 400000 > cpu.cfs_quota_us
echo 100000 > cpu.cfs_period_us
```

b. Configure cpu.cfs burst us to enable the CPU burst feature.

```
echo 600000 > cpu.cfs_burst_us
```

# **Query CPU burst statistics**

190

1. To go to the path of a child cgroup, run the following command:

cd /sys/fs/cgroup/cpu/<Name of the child cgroup directory>/

2. To view CPU burst statistics, run the following command:

```
cat cpu.stat
```

Parameters in the command output:

- o nr burst: the number of CPU bursts triggered by the child cgroup.
- o burst time : the cumulative length of CPU burst time for the child cgroup.

# 5.6.18. Group identity feature

Alibaba Cloud Linux 2 with a kernel of the kernel-4.19.91-24.al7 version or later supports the group identity feature. You can use the group identity feature to configure different identities for CPU control groups (cgroups) to prioritize process tasks in the cgroups.

### Context

Notice Elastic Compute Service (ECS) instances that run Alibaba Cloud Linux 2 with a kernel of the kernel-4.19.91-25.1.al7 version can go down when the group identity feature is used. You must upgrade the kernel version to kernel-4.19.91-25.6.al7 or later. For more information, see the FAQ section of this topic.

When latency-sensitive tasks and computing tasks are deployed on the same instance, the Linux kernel scheduler must provide more scheduling opportunities to high-priority tasks to minimize scheduling latency and the impacts of low-priority tasks on kernel scheduling. For this scenario, Alibaba Cloud Linux 2 provides the group identity feature and adds interfaces used to configure scheduling priorities to CPU cgroups. Tasks with different priorities have the following characteristics:

- High-priority tasks have the minimum wakeup latency.
- Low-priority tasks do not affect the performance of high-priority tasks.
  - The wakeup of low-priority tasks does not affect the performance of high-priority tasks.
  - Low-priority tasks do not use the simultaneous multithreading (SMT) scheduler to share hardware and have no negative impacts on the performance of high-priority tasks.

### How the group identity feature works

The group identity feature allows you to configure identities for CPU cgroups to prioritize tasks in the cgroups. The group identity feature relies on a dual red-black tree architecture. A low-priority red-black tree is added based on the red-black tree of the Completely Fair Scheduler (CFS) scheduling queue to store low-priority tasks.

When the kernel schedules the tasks that have identities, the kernel processes the tasks based on their priorities. The following table describes the identities in descending order of priority.

Identity	Description
----------	-------------

Identity	Description
	Identifies a high-priority task. A high-priority task has more opportunities to preempt resources than a low-priority task.
	When the CFS schedules high-priority tasks, the following situations may occur:
	<ul> <li>If a high-priority task wakes up while a low-priority task is running, the high-priority task can unconditionally preempt resources from the low-priority task.</li> </ul>
ID_HIGHCLASS	• If a high-priority task wakes up while a normal-priority task is running and the virtual runtime (vruntime) of the high-priority task is less than that of the normal-priority task, the high-priority task can ignore the original scheduling policy and preempt resources. The original scheduling policy specifies that a task cannot preempt resources when its runtime on a CPU is less than the minimum runtime.
	<ul> <li>When tasks queue up to run, if a low- or normal-priority task is running, a high-priority task whose vruntime is less than that of the running task can ignore the original scheduling policy and preempt resources. The original scheduling policy specifies that a task cannot preempt resources when its runtime on a CPU is less than the minimum runtime.</li> </ul>
	Identifies a normal-priority task. A normal-priority task has more opportunities to preempt resources than a low-priority task.
	When the CFS schedules normal-priority tasks, the following situations may occur:
ID_NORMAL	<ul> <li>If a normal-priority task wakes up while a low-priority task is running, the normal-priority task can unconditionally preempt resources from the low-priority task.</li> </ul>
_	<ul> <li>When tasks queue up to run, if a low-priority task is running, a normal-priority task whose vruntime is less than that of the running task can ignore the original scheduling policy and preempt resources. The original scheduling policy specifies that a task cannot preempt resources when its runtime on a CPU is less than the minimum runtime.</li> </ul>
	Identifies a low-priority task.
ID INDEDGLACS	When the CFS schedules low-priority tasks, the following situations may occur:
ID_UNDERCLASS	If the peer SMT scheduler has run the <code>ID_SMT_EXPELLER</code> task, low-priority tasks cannot be scheduled to CPUs and are kicked out of the queue of tasks to be run.

The preceding identities apply based on the resource management policies of CPU cgroups.

- For tasks in cgroups of the same level, identity priorities take effect.
- For tasks in parent cgroups, identity priorities do not take effect. For tasks in child cgroups, identity priorities take effect.

• Resources are competed among tasks that have identities with the same priority in compliance with CFS policies. Note that the runtime of tasks identified by ID\_UNDERCLASS or ID\_NORMAL may not reach the minimum value.

#### Other identities

Identity	Description
ID_SMT_EXPELLER	Identifies the SMT expeller. The SMT expeller evicts the tasks that are identified by ID_UNDERCLASS from the peer CPU when the SMT scheduler runs.
ID_IDLE_SEEKER	Specifies that when a task wakes up, the task attempts to find idle CPUs within the limits of scheduler policies.
ID_IDLE_SAVER	Used with the sched_idle_saver_wmark kernel parameter. You can use sched_idle_saver_wmark to set a water mark for CPU idle time. When a task identified by ID_IDLE_SAVER wakes up, the task attempts to find only an idle CPU whose idle time exceeds the specified water mark.

### **Interfaces**

Interfaces used to configure identities

The group identity feature provides two interfaces for you to configure task identities: /sys/fs/cgr oup/cpu/\$cg/cpu.identity and /sys/fs/cgroup/cpu/\$cg/cpu.bvt\_warp\_ns . The \$cg variable indicates the child cgroup directory node where a task is located. Before you use the interfaces, take note of the following items:

- The cpu.bvt\_warp\_ns interface is a quick configuration interface. The written value of this interface can be converted to the value of identity.
- Both cpu.identity and cpu.bvt\_warp\_ns interfaces can be used to change the identities of cgroups.
- o After data is written to the cpu.identity interface, the last value written of the cpu.bvt\_warp\_ns interface is overwritten. This overwrite operation is not reflected in the cpu.bvt\_warp\_ns interface.
- After the data is written to the cpu.bvt\_warp\_ns interface, the last written value of the cpu.i
  dentity interface is overwritten. This overwrite operation is not reflected in the cpu.identity
  interface.
- You can use one of the interfaces to configure task identifies. We recommend that you do not configure both of the interfaces.
- If you are unfamiliar with operations related to the operating system kernel, we recommend that you do not use the cpu.identity interface.

The following table describes the interfaces.

Interface Description	
-----------------------	--

Interface	Description
	The default value is 0, which indicates the <code>ID_UNDERCLASS</code> identity.
	The interface is a 5-bit segment. Valid values of each bit of the interface:
	• If the interface is empty, it indicates the <code>ID_NORMAL</code> identity.
	• Bit 0: indicates the ID_UNDERCLASS identity.
	• Bit 1: indicates the ID_HIGHCLASS identity.
cpu.identity	• Bit 2: indicates the ID_SMT_EXPELLER identity.
	• Bit 3: indicates the ID_IDLE_SAVER identity.
	• Bit 4: indicates the ID_IDLE_SEEKER identity.
	other bits to 0 to obtain a binary value of 10010, which is converted to a decimal value of 18. Then, run the echo 18 > /sys/fs/cgro p/cpu/ \$cg /cpu.identity command to write 18 to cpu.identity  The default value is 0, which indicates the ID_NORMAL identity.
	Valid values:  • 2: indicates the ID_SMT_EXPELLER , ID_IDLE_SEEKER , and ID_HIGHCLASS identities. The value of the corresponding identity is 22.
only but warn ne	• 1: indicates the ID_HIGHCLASS and ID_IDLE_SEEKER identities. The value of the corresponding identity is 18.
cpu.bvt_warp_ns	<ul> <li>0: indicates the ID_NORMAL identity. The value of the corresponding identity is 0.</li> </ul>
	• -1: indicates the ID_UNDERCLASS and ID_IDLE_SAVER identities. The value of the corresponding identity is 9.

• Interfaces used to enable or disable scheduling features

You can run the following command to view the default settings of kernel scheduling features by using the sched\_features interface:

cat /sys/kernel/debug/sched\_features

The following table describes the scheduling features.

Scheduling feature	Description	Default value

Scheduling feature	Description	Default value
ID_IDLE_AVG	This feature is used with the <code>ID_ID</code> <code>LE_SAVER</code> identity to count the runtime of <code>ID_UNDERCLASS</code> tasks towards idle time. This ensures that no CPUs remain idle when only <code>ID_UNDERCLASS</code> tasks are running and prevents resource wastes.	ID_IDLE_AVG : indicates that the feature is enabled.
ID_RESCUE_EXPELLEE	This feature is used in load balancing scenarios. If tasks cannot find CPU resources available for use, CPUs that are evicting ID_UNDERCLASS tasks are load-balanced. This feature helps ID_UNDERCLASS tasks get out of the evicted state as soon as possible.	ID_RESCUE_EXPELLEE : indicates that the feature is enabled.
ID_EXPELLEE_NEVER_H	After this feature is enabled, when a task that is being evicted decides to migrate to another CPU, hot cache does not cause the migration request to be denied. This feature helps ID_UNDERCLASS tasks get out of the evicted state as soon as possible.	NO_ID_EXPELLEE_NEVER_HOT : indicates that the feature is disabled.
ID_LOOSE_EXPEL	After this feature is enabled, CPUs do not update their eviction states every time they select tasks but have the states automatically updated at the time specified by the sched_ex pel_update_interval kernel parameter. The configuration of this feature affects only state updates when CPUs select tasks, not the updates of IPI interrupts.	NO_ID_LOOSE_EXPEL : indicates that the feature is disabled.
ID_LAST_HIGHCLASS_S TAY	After this feature is enabled, the last ID_HIGHCLASS task that runs on a CPU cannot be migrated to another CPU.	<pre>ID_LAST_HIGHCLASS_STAY : indicates that the feature is enabled.</pre>

# • Interfaces used by sysctl to configure kernel parameters

Some capabilities of the group identity feature depend on the values of kernel parameters. The following table describes the parameters.

Kernel parameter	Description	Unit	Default value
------------------	-------------	------	------------------

Kernel parameter	Description	Unit	Default value
<pre>/proc/sys/kernel/sched_expel_u pdate_interval</pre>	The interval at which the eviction state is automatically updated when a CPU selects tasks. This parameter is valid only when the ID_LOOSE_EXPE feature is enabled.	ms	10
/proc/sys/kernel/sched_expel_i dle_balance_delay	The minimum idle balance interval when a CPU is evicting tasks. A value of -1 indicates that idle balance is not allowed.  If only ID_UNDERCLASS tasks exist on a CPU and the tasks are being evicted, the CPU is idle. Idle balance is performed on this CPU to improve load-balancing effects.  However, this may damage ID_UNDERCLASS tasks. You can set the sched_expel_idle_balance_delay parameter to alleviate this issue.	ms	-1
<pre>/proc/sys/kernel/sched_idle_sa ver_wmark</pre>	The water mark for CPU idle time. When an ID_IDLE_SAVER task wakes up, the task attempts to find an idle CPU whose idle time exceeds the specified water mark.	ns	0

# Information output

When you use the group identity feature, you can run the following command to view various parameters:

```
cat /proc/sched_debug
```

The following table describes the output parameters.

Parameter	Description
nr_high_running	The number of ID_HIGHCLASS tasks that are running on the current CPU.
nr_under_running	The number of ID_UNDERCLASS tasks that are running on the current CPU.
nr_expel_immune	The number of non- ID_UNDERCLASS tasks that are running on the current CPU.

Parameter	Description	
smt_expeller	Indicates whether ID_SMT_EXPELLER tasks are running on the current CPU. A value of 1 indicates that ID_SMT_EXPELLER tasks are running on the current CPU. A value of 0 indicates that no ID_SMT_EXPELLER tasks are running on the current CPU.	
on_expel	Indicates whether ID_SMT_EXPELLER tasks are running on the peer SMT scheduler. A value of 1 indicates that ID_SMT_EXPELLER tasks are running on the current CPU. A value of 0 indicates that no ID_SMT_EXPELLER tasks are running on the current CPU.	
high_exec_sum	The cumulative runtime of ID_HIGHCLASS tasks on the current CPU.	
under_exec_sum	The cumulative runtime of ID_UNDERCLASS tasks on the current CPU.	
h_nr_expel_immune	The number of non- ID_UNDERCLASS tasks that are running on cfs_rq .	
expel_start	The difference between the minimum vruntimes of the two red-black trees when the CPU starts to evict tasks.	
expel_spread	The cumulative difference between the minimum vruntimes of the two red-black trees caused by CPU eviction states.	
min_under_vruntime	The minimum vruntime of the low-priority red-black tree.	

## **FAQ**

How do I upgrade the kernel version from kernel-4.19.91-25.1.al7 to kernel-4.19.91-25.6.al7 or later? Solution:

1. Log on to the instance.

For more information, see Connect to a Linux instance by using a password or key.

2. Run the following command to query the kernel version:

```
uname -r
```

3. Run the following command to upgrade the kernel version:

```
yum update kernel
```

4. Run the following command to restart the instance to make the new kernel version take effect:

reboot

# 5.6.19. Statistical status change of anonymous pages in the /proc/meminfo file

In the 4.19.91-22.al7.x86\_64 and later kernel versions of Alibaba Cloud Linux 2 operating systems, the statistical Active (anon) and Inactive (anon) statuses of anonymous pages in the /proc/meminfo file are changed. If your application uses these statistics, this change causes inconsistency between the data that is queried before the change and the data after the change. This topic describes the change in the statistical status of anonymous pages in the /proc/meminfo file.

# Reason for the change

The reclaiming of anonymous pages in the latest kernel version is optimized in the Linux kernel community. The default storage location of anonymous pages is changed from the active least-recently used (LRU) list to the inactive LRU list. This affects the statistical Active (anon) and Inactive (anon) statuses of data in the /proc/meminfo file. The following items list the updates made in the community:

- New version used in the community: Linux Kernel 5.10
- Commit submitted to the community: b518154e59aab3ad0780a169c5cc84bd4ee4357e

Alibaba Cloud has incorporated this optimization item in the 4.19.91-22.al7.x86\_64 and later kernel versions of Alibaba Cloud Linux 2. The statistical Active (anon) and Inactive (anon) statuses of anonymous pages in the /proc/meminfo file will change.

**Note** You can run the **uname** -r command in operating systems to obtain kernel version information.

# Details of the change

By default, in the kernel versions earlier than  $4.19.91-22.al7.x86\_64$  of Alibaba Cloud Linux 2, the anonymous pages applied by applications are counted as Active (anon) in the proc/meminfo file, and are counted as Inactive (anon) only when the memory is insufficient.

By default, in the kernel version of 4.19.91-22.al7.x86\_64 or later of Alibaba Cloud Linux 2, the anonymous pages applied by applications are counted as Inactive (anon) in the /proc/meminfo file, and are counted as Active (anon) only after the anonymous pages are accessed. Therefore, if you use Active (anon) and Inactive (anon) in the /proc/meminfo file to indicate the statuses of anonymous pages, you need to adjust the configuration based on the change.

# Impact of the change and recommendations

Due to the statistical Active (anon) and Inactive (anon) status change of anonymous pages, the Active (anon) and Inactive (anon) data queried by an application before this change are inconsistent with the data queried after the change.

We recommend that you adjust the configurations of your applications based on the latest statistics of anonymous pages in your Alibaba Cloud Linux 2 operating system.

# 5.6.20. Enhance the visibility of container

### resources

198

Alibaba Cloud Linux 2 with a kernel of the 4.19.91-22.2.al7 or later version provides kernel interfaces for the container resource visualization feature to enhance the visibility of container resources. This topic describes these interfaces and their sample scenarios.

## Context

By default, the container resource visualization feature provided by Alibaba Cloud Linux 2 is disabled. After you enable this feature, you can run commands such as **top** and **free** in a container to read data from the following interfaces. When you run these commands, the resource information of the container rather than that of the Elastic Compute Service (ECS) instance on which the container resides is returned.

- /proc/cpuinfo
- /proc/meminfo
- /sys/devices/system/cpu/online

# **Interfaces**

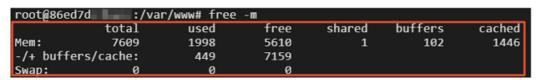
Interface	Description
/proc/sys/kernel/rich_container_enable	<ul> <li>Specifies whether the container resource visualization feature is enabled. Valid values:</li> <li>0: The container resource visualization feature is disabled.</li> <li>1: The container resource visualization feature is enabled.</li> <li>Default value: 0.</li> </ul>
/proc/sys/kernel/rich_container_source	<ul> <li>Specifies the data source of the cgroup interface.</li> <li>Valid values:</li> <li>0: uses the cgroup in which the current pointer resides as a data source.</li> <li>1: uses the cgroup in which the child reaper (Process 1 of the current PID namespace) resides as a data source.</li> <li>Default value: 0.</li> </ul>

Interface	Description
/proc/sys/kernel/rich_container_cpuinfo_source	Specifies the number of CPUs that are displayed in the /proc/cpuinfo and /sys/devices/system/cpu/online interfaces. Valid values:  • 0: uses the ratio of the value of Request to that of Limit ( Request/Limit ) in Kubernetes. This ratio equals the ratio of the value of quota to that of period ( quota/period ).  • 1: uses the CPU data source in the cpuset.cpus interface.  • 2: uses the integer to which the value obtained by dividing cpu.shares by /proc/sys/kern el/rich_container_cpuinfo_sharesbase is rounded up. For example, if the value obtained by dividing cpu.shares by /proc/sys/kernel/rich_container_cpuinfo_sharesb ase is 1.1, 2 is used as the final value. The final value cannot exceed the number of available CPUs.  Default value: 0.
/proc/sys/kernel/rich_container_cpuinfo_sharesbase	When the value of the /proc/sys/kernel/rich_container_cpuinfo_source interface is set to 2, you must use this interface as a part of the formula to obtain the final value. The value must be an integer greater than or equal to 2. Default value: 1024.

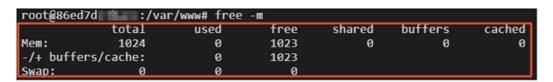
# **Examples**

In the examples, Docker is deployed in a Linux instance, and a container with 1 GB memory is created.

• If the container resource visualization feature is disabled, the value of the /proc/sys/kernel/rich\_cont ainer\_enable interface is set to 0. When you run the free -m command in the container, resource information is displayed, as shown in the following figure. It is the resource information of the Linux instance on which the container resides.



• If you have run the echo 1 > /proc/sys/kernel/rich\_container\_enable command on the Linux instance to enable the container resource visualization feature, when you run the free -m command in the container, resource information is displayed, as shown in the following figure. It is the resource information of the container.



# Special scenarios

In most cases, after the container resource visualization feature is enabled, the default values of the interfaces can meet your requirements. However, you must take note of the interface configurations in the following special scenarios:

For a pod in which the PID namespace is not shared, each container is an independent PID namespace. If you run the systemd command to start a privileged container and if the container process is Process 1, the task of collecting monitoring data is run in a child cgroup instead of the root cgroup of the container.

For example, when you log on to the container by using an SSH key pair and run the **cat** /**proc/cpuinfo** command, one of the following situations may occur:

- If the value of the /proc/sys/kernel/rich\_container\_source interface is set to 0 when you run the command, the current pointer resides in the child cgroup that is created by sshd.service in the container. In this case, error data is returned.
- If you run the echo 1 > /proc/sys/kernel/rich\_container\_source command, the cgroup in which the child reaper resides is used as a data source. In this case, valid data is returned.

# 5.6.21. SMC-R

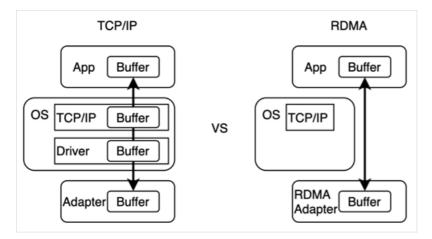
Alibaba Cloud Linux 3 is tuned to support Shared Memory Communications over RDMA (SMC-R). SMC-R is based on Alibaba Cloud eRDMA and can transparently replace TCP in applications without loss of functionality. SMC-R enables direct, high-speed, low-latency, and memory-to-memory communications and provides higher performance than TCP in various scenarios such as in-memory databases, remote procedure calls (RPCs), and transmission of large files. This topic describes SMC-R and how to work with it.

# **Background information**

SMC-R is based on Remote Direct Memory Access (RDMA). Before you can understand SMC-R, you must understand what RDMA and Alibaba Cloud eRDMA are.

RDMA is a remote direct memory access technology that bypasses kernel intervention. It is widely used in data- and compute-intensive scenarios and is suitable for a wide array of fields such as high-performance computing, machine learning, data centers, and mass storage. RDMA underlies the core services of Alibaba Cloud, such as Enhanced SSD (ESSD) and PolarDB, and has been tried and tested in critical business scenarios such as Double 11.

RDMA has the zero-copy and stack offload capabilities. It offloads protocol stacks to RDMA Network Interface Controllers (RNICs) and performs direct memory access without involving the kernel. By bypassing the operating system (OS) kernel stack, RDMA reduces CPU processing costs and delivers lower latency and higher throughput than traditional TCP networks. The following figure shows the differences between TCP/IP and RDMA stacks.



Previously, RDMA was used only with NICs and switches in some data centers and was complex to deploy. Alibaba Cloud eRDMA is a service that brings RDMA to the cloud for easy use, and allows Elastic Compute Service (ECS) users to use RDMA for data transmission without concerns about complex physical network environment configurations such as NICs and switches.

However, due to the great differences between the InfiniBand (IB) verbs interfaces used by RDMA and common POSIX socket interfaces, conventional applications must be significantly transformed before they can work with RDMA. Meanwhile, technical expertise is required to use IB verbs interfaces in an efficient manner.

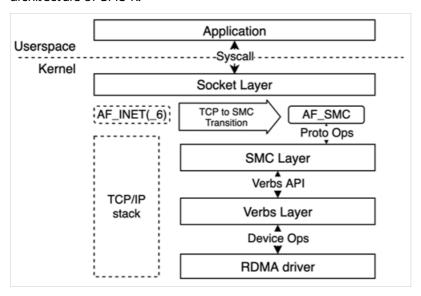
To make full use of eRMDA and deliver higher network performance, Alibaba Cloud Linux 3 provides optimized SMC-R and supports eRDMA. SMC-R provides a standard socket interface over RDMA to applications. SMC-R uses RDMA in an efficient manner and remains compatible with standard TCP applications to allow more applications to benefit from RDMA without modifications.

### **Architecture**

SMC-R architecture:

• Protocol hierarchy and transparent replacement

SMC-R is a reliable streaming transmission protocol that is fully compatible with sockets. SMC-R functions between the socket layer and the IB verbs layer in kernel space. It supports common socket interfaces and uses the IB verbs kernel-mode interface to help the RDMA driver transmit data. Alibaba Cloud Linux 3 provides a tool for the replacement of the protocol family at the socket layer to allow SMC-R to transparently replace TCP at the net namespace or application level. When you use Alibaba Cloud Linux 3, you can transition from TCP to SMC-R and achieve higher network performance based on RDMA without the need to modify network applications. The following figure shows the architecture of SMC-R.



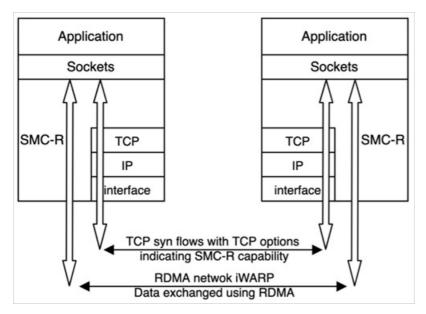
Automatic negotiation and secure fallback

SMC-R has the automatic negotiation capability and can dynamically fall back to TCP. To establish an RDMA connection, SMC-R first establishes in the kernel a TCP connection to the remote node. During the handshake, the local node uses specific TCP options to indicate its support for SMC-R and verifies that the remote node also supports SMC-R.

- o If both the local and remote nodes are confirmed to support SMC-R during the negotiation, SMC-R applies for required RDMA resources for user-mode network applications. The RDMA resources include the queue pair (QP) and completion queue (CQ) required to build an asynchronous RDMA communication model. At the same time, SMC-R creates a send buffer and a receive buffer and registers the receive buffer as the remote memory buffer (RMB) to which the remote node has direct access. Then, SMC-R initializes the RDMA connection. SMC-R encapsulates the access key (RKey) and start address of the RMB into a RT oken and notifies the remote node of the RT oken.
- If the local or remote node is found not to support SMC-R during the negotiation, the fallback-to-TCP mechanism is triggered and the local and remote nodes use the established TCP connection to transmit data and ensure network stability and reliability.

**Note** Note that SMC-R can fall back to a TCP stack only during connection negotiation. SMC-R cannot fall back to a TCP stack during data transmission.

The following figure shows the network flows for negotiation and data transmission.

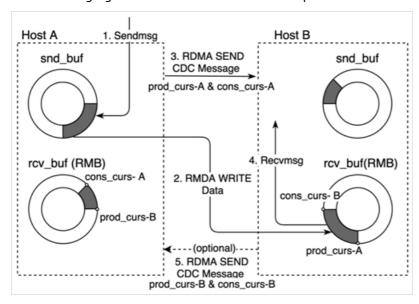


• Ring memory and data receiving and sending

SMC-R relies on efficient RDMA networks and a ring-shaped shared memory architecture for high-performance data transmission. After network applications transmit data to the kernel, SMC-R performs RDMA Write operations to write the data to the ring RMB of the remote node and performs RDMA Send and Receive operations to send and receive Connection Data Control (CDC) messages for updating and synchronizing the cursors in RMBs.

For the RMB on one side, the reading peer updates its consumer cursor to indicate the address of the next byte of data to be consumed. To prevent data loss, the writing peer does not write data to the RMB beyond the consumer cursor. Similarly, the writing peer updates its producer cursor to indicate the address of the next byte of data to be written. To ensure data correctness, the reading peer does not read data from the RMB beyond the producer cursor. Cursors are updated and synchronized between peers to manage and track writes and reads to ensure the security and reliability of data transmission.

The following figure shows the data transmission procedure.



## **Benefits**

SMC-R has the following benefits:

### • High performance

RDMA offloads data-plane features to RNICs and bypasses the kernel to directly access remote ring receive buffers. This enables SMC-R to have lower latency, higher throughput, and smaller CPU loads than traditional TCP stacks in specific scenarios.

- SMC-R protocol stacks are more lightweight than TCP stacks.
- o SMC-R uses RDMA for communication to lower latency and CPU loads and improve throughput.
- SMC-R has direct access to efficient, reliable remote ring buffers.

### • Transparent replacement

SMC-R is compatible with POSIX socket interfaces and can transparently replace TCP in new connection stacks at the net namespace or application level by using sysctl and user-mode tools at no additional costs for manually modifying or further developing applications.

- SMC-R exploits RDMA Reliable Connection (RC) transports at the underlying layer and is compatible with socket interfaces to provide reliable streaming transmission in place of TCP.
- SMC-R has the automatic negotiation and secure fallback-to-TCP mechanisms.
- SMC-R can transparently replace TCP at the net namespace or application level without loss of functionality.
- SMC-R is compatible with eRDMA Internet Wide-area RDMA Protocol (iWARP) and RDMA over Converged Ethernet (RoCE) networks at the underlying layer.

### Use scenarios

SMC-R is applicable to the following scenarios:

• Latency-sensitive data queries and processing

SMC-R is applicable to scenarios that involve high-performance data queries and processing and require high network performance, such as Redis, Memcached, and PostgreSQL. SMC-R allows applications to use it in place of TCP in a transparent and non-invasive manner and gain a 50% increase in QPS without the need for further development or adaptation.

• High-throughput data transmission

A large amount of bandwidth and CPU resources tend to be consumed when data is exchanged or transmitted at scale within a cluster. The efficient communication model used by Shared Memory Communications (SMC) enables SMC-R to deliver the same throughput at a lower CPU load than traditional TCP stacks. This way, computing resources are saved.

### Instructions

Alibaba Cloud Linux 3 provides a wide array of monitoring and maintenance tools for you to monitor the status of SMC-R and diagnose its issues. You can perform the following procedure to use SMC-R.

1. Load the SMC-R modules.

By default, SMC-R is compiled into kernel modules: smc and smc\_diag. You can manually load these modules in the system.

i. Run the following command to load the smc kernel module:

modprobe smc

ii. Run the following command to load the smc diag kernel module:

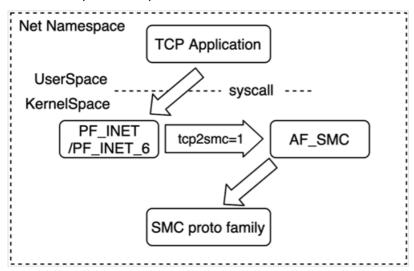
```
modprobe smc_diag
```

2. Enable transparent replacement.

Alibaba Cloud Linux 3 supports net namespace-level and application-level transparent replacement to allow SMC-R to transparently replace TCP for net namespaces or applications.

• Net namespace-level transparent replacement

Alibaba Cloud Linux 3 provides the net namespace-level transparent replacement feature to replace new TCP sockets with SMC-R sockets within a net namespace. The following figure shows the replacement procedure.



You can perform the following operations to configure transparent replacement for a net namespace:

a. Run the following command to enable transparent replacement for a net namespace:

```
sysctl net.smc.tcp2smc=1
```

b. Run applications.

If the operations in Step i are also performed at the peer, the local and remote nodes use SMC-R for data transmission. If the operations are not performed at the peer, the local and remote nodes fall back to TCP for data transmission. For more information about the negotiation process, see the Automatic negotiation and secure fallback section in this topic.

c. Run the following command to disable transparent replacement for the net namespace:

```
sysctl net.smc.tcp2smc=0
```

Application-level transparent replacement

Alibaba Cloud Linux 3 also provides the application-level transparent replacement feature to replace TCP with SMC-R for an application. This feature requires the SMC-R monitoring and maintenance toolkit smc-tools.

Note You can run the yum install smc-tools -y command to install the smc-tools toolkit. For more information about the smc-tools toolkit, see Step 3.

When you execute the <code>smc\_run</code> script from <code>smc-tools</code> to run applications, the <code>smc\_run</code> script uses the <code>LD\_PRELOAD</code> environment variable to define <code>libsmc-preload.so</code> in <code>smc-tools</code> as the dynamic library to be loaded first. <code>libsmc-preload.so</code> attempts to replace the new TCP sockets established for applications with SMC-R sockets. <code>smc\_run</code> command description:

```
Usage: smc_run [ OPTIONS ] COMMAND
Run COMMAND using SMC for TCP sockets
```

For example, to use SMC-R to run the *testApp* application in the current directory, run the following command:

```
smc_run ./testApp
```

Similar to net namespace-level transparent replacement, application-level transparent replacement requires the local and remote nodes to transparently replace TCP with SMC-R by using smc\_run before an SMC-R connection can be established for RDMA communication.

3. Use the SMC-R monitoring and maintenance tools to monitor SMC-R.

You can use smc-tools to track and diagnose SMC-R from multiple aspects. smc-tools includes the following tools:

- o smcr: shows information about SMC-R, such as information about linkgroups and devices.
- o smcss: shows information about active SMC-R sockets.
  - i. Run the following command to install smc-tools:

```
yum install smc-tools -y
```

### ii. Use smcr.

**smcr** is used to show information about SMC-R, such as information about linkgroups and devices.

### Command description:

```
Usage: smcr [ OPTIONS ] OBJECT {COMMAND | help}

OBJECT : { linkgroup | device }

    linkgroup

    Linkgroup(s) or link(s) used by SMC-R.

device

    One or more SMC-R devices.

OPTIONS : {-v[ersion] | -d[etails] | -dd[etails]}

-v, -version

    Print the version of the smcr utility and exit.

-d, -details

    Print detailed information.

-dd, -ddetails

    Print more detailed information.
```

### Usage example:

■ You can run the following command to view SMC-R device information:

```
smcr device
```

### Example command output:

```
Net-Dev IB-Dev IB-P IB-State Type Crit #Links PNET-ID eth0 erdma_01 1 ACTIVE 0x7ffd No 0
```

■ You can run the following command to view SMC-R linkgroup information:

```
smcr linkgroup
```

### Example command out put:

```
LG-ID LG-Role LG-Type VLAN #Conns PNET-ID 00000100 CLNT SINGLE 0 1
```

### iii. Use smcss.

smcss shows information about SMC-R sockets.

Command description:

```
Usage: smcss [ OPTIONS ]
OPTIONS :
       (none)
              displays a list of connecting, closing, or connected SMC sockets with
basic information.
       -a, --all
             displays all types of SMC sockets: listening, opening, closing, and c
onnected.
       -1, --listening
             shows listening sockets only. These are omitted by default.
       -d, --debug
             displays additional debug information, such as shutdown state.
              displays additional SMC-D specific information. Shows SMC-D sockets o
nly.
       -h, --help
             displays usage information.
       -R, --smcr
             displays additional SMC-R specific information. Shows SMC-R sockets o
nly.
       -v, --version
             displays program version.
       -W, --wide
              do not truncate IP addresses.
```

### Usage example:

You can run the following command to view detailed information about all SMC-R sockets.

```
smcss -a -R -d
```

### Example command output:

```
UID Inode Local Address
                                                 Peer Address
                Sndbuf Rcvbuf Peerbuf rxprod-Cursor rxcons-Cursor rxFlags t
ode Shutd Token
xprod-Cursor txcons-Cursor txFlags txprep-Cursor txsent-Cursor txfin-Cursor Role I
           Port Linkid GID
B-device
                                                                Peer-GID
             00000 1105985880 192.168.XX.XX:49080
ACTIVE
                                                  192.168.XX.XX:10003
                                                                         0000
SMCR <-> 00000001 00020000 00040000 00040000 0001:00026256 0001:00026256 00:00
0001:00026264 0001:00026256 00:00 0003:00006264 0003:00006264 0003:00006264 CLNT
erdma 012211
             01 01
                          0016:3e01:2211:0000:0000:0000:0000 0016:3e01:43b4
:0000:0000:0000:0000:0000
```

# 5.6.22. UKFEF

Alibaba Cloud Linux 3 supports Unified Kernel Fault Event Framework (UKFEF) in the kernel of the 5.10.60-9.al8.x86\_64 version. UKFEF is used to collect the exception events on Alibaba Cloud Linux 3 that may cause risks and generate event reports in a unified format. This topic describes the events collected by UKFEF, the methods used by UKFEF to generate event reports, and UKFEF control interfaces.

# **Background information**

An operating system may display some characteristics or send recognizable error messages before serious issues occur. During O&M, you can use such information to predict faults and take precautions. However, the information is generated in different formats and is distributed among different modules. As a result, you may encounter the following issues when you collect exception events on an operating system:

- Expertise on the operating system is required to parse exception events and determine their impacts.
- Exception events are displayed in a variety of formats, which increases the difficulty of automated O&M. Automated O&M matches the collected exception events against formats and then performs data cleansing to filter out unneeded data based on actual requirements.

To resolve the preceding issues, Alibaba Cloud Linux 3 provides UKFEF at the kernel layer. UKFEF collects information from a variety of exception events that may give rise to risks, determines the severity of the events, and then generates event reports in a unified format. These reports include the scenarios in which the issues occur and the recommended risk levels. This makes it easier to identify system exceptions during O&M. UKFEF also classifies known exception events and provides system risk reports that were unavailable in previous kernel versions.

UKFEF generates reports based on multiple dimensions such as the type, impact, and statistics of exception events. This allows you to efficiently diagnose exceptions during O&M. In addition, event reports are generated by using multiple methods to prevent data loss.

# **Event description**

The following table describes the event types and event levels classified by UKFEF and the methods for generating event reports.

Event information	Description
Event type	UKFEF collects the following common events on the operating system kernel:  • soft lockup  • Read-Copy Update (RCU) stall  • hung task  • global Out of Memory (OOM)  • cgroup OOM  • page allocation failure  • list corruption  • bad mm_struct  • I/O error  • EXT4-fs error  • Machine Check Exception (MCE)  • fatal signal  • warning  • panic

Event information	Description
	<ul> <li>UKFEF classifies exception events into three levels:</li> <li>Slight: The exception event does not affect the normal running of the operating system, but services deployed in the operating system may experience jitter. You can continue to observe for changes in the exception event.</li> </ul>
Event level	<ul> <li>Normal: The exception event may occur in the current application process. We recommend that you take measures such as terminating , restarting, or migrating the current application process.</li> </ul>
	Fatal: The exception event may cause fatal damages to the operating system.     We recommend that you immediately migrate your business.

Event information	Description	
	<ul><li>UKFEF generates event reports by using the following methods:</li><li>Use kernel logs to show the details of a single event. Example:</li></ul>	
	<pre>class Fault event[module:type]:messages. At cpu cpuid, task pid(cmdline). Total fault: cnt</pre>	
	<ul> <li>The details include the following parameters:</li> <li>class: The level of the exception event.</li> <li>module: The module to which the exception event belongs, including sched, mem, io, fs, net, and hardware. When an exception is caused by multiple modules, the value of this parameter is general.</li> <li>type: The type of the exception event.</li> <li>messages: The custom message of the exception event.</li> <li>cpuid: The ID of the CPU on which the exception event occurs.</li> <li>pid(cmdline): The pid and cmdline of the process corresponding to the exception event.</li> </ul>	
	<ul> <li>Note If the value of pid is -1, no corresponding process is running.</li> <li>cnt: The total number of the occurrences of exception events of the current type after system startup.</li> </ul>	
Event report	• Use the /proc/fault_events file to show the total number of the occurrences of exception events of all types. Example:	
	Total fault events: 0 Slight: 0 Normal: 0 Fatal: 0 soft lockup: 0 rcu stall: 0 hung task: 0 global oom: 0 cgroup oom: 0 page allocation failure: 0 list corruption: 0 bad mm_struct: 0 io error: 0 ext4 fs error: 0 mce: 0 fatal signal: 0 warning: 0 panic: 0	

# Control interfaces

Interface	Description
<pre>/proc/sys/kernel/fault_event_e nable</pre>	<ul><li>Specifies whether UKFEF is enabled. Valid values:</li><li>1: UKFEF is enabled.</li><li>0: UKFEF is disabled.</li></ul>
<pre>/proc/sys/kernel/fault_event_p rint</pre>	<ul> <li>Specifies whether UKFEF generates event reports. Valid values:</li> <li>1: UKFEF generates event reports.</li> <li>0: UKFEF does not generate event reports.</li> </ul>
/proc/sys/kernel/panic_on_fata l_event	Specifies whether to trigger the Panic mechanism of the operating system when a Fatal event occurs. Valid values:  1: triggers the Panic mechanism.  0: does not trigger the Panic mechanism.

# 5.7. Operating System Migration Tool

# 5.7.1. Migrate from CentOS 8 to Alibaba Cloud

# Linux 3

CentOS has announced the end of life (EOL) for CentOS Linux. You can use a migration tool provided by Alibaba Cloud Linux to smoothly migrate from CentOS 8 to Alibaba Cloud Linux 3. This topic describes how to use the migration tool.

### Limits

Alibaba Cloud Linux 3 uses Linux Kernel 5.10 LTS, and its GNU toolchain contains GCC 10.2, Binutils 2.35, and glibc 2.32. Cent OS 8 software is compatible with Alibaba Cloud Linux 3. When you use the migration tool to migrate from Cent OS 8 to Alibaba Cloud Linux 3, the migration tool determines which features and software may be limited after the migration. The following table describes the limits.

Feature or software	Supported	Migration affected	Description
GUI	No.	Yes.	Alibaba Cloud Linux 3 does not support GUIs. As a result, if your CentOS 8 operating system is configured with a GUI, you cannot use the migration tool to migrate to Alibaba Cloud Linux 3.
i686-based software package	No.	Yes.	Alibaba Cloud Linux 3 does not provide i686-based software packages. If your CentOS 8 operating system comes with i686-based software packages, you cannot use the migration tool to migrate to Alibaba Cloud Linux 3. To prevent system exceptions, you must delete the i686-based software packages before you perform a migration.

Feature or software	Supported	Migration affected	Description
epel-modular	No.	No.	The Extra Packages for Enterprise Linux (EPEL) repositories of CentOS 8 and Alibaba Cloud Linux 3 depend on different platforms. You can migrate from CentOS 8 to Alibaba Cloud Linux 3, but you cannot install the software packages that are based on epelmodular after the migration.
Compatibility with C/C++ programs	Not guaranteed.	No.	Alibaba Cloud Linux 3 comes with software packages that are optimized for better compatibility with C/C++ programs. The GNU toolchain of Alibaba Cloud Linux 3 contains GCC 10.2, Binutils 2.35, and glibc 2.32. After you migrate from CentOS 8 to Alibaba Cloud Linux 3, we recommend that you compile and verify your C/C++ programs by using the new GNU toolchain to prevent incompatibility.
Kernel module compatibility	No.	No.	Alibaba Cloud Linux 3 uses Linux Kernel 5.10 LTS. Alibaba Cloud Linux 3 is compatible with CentOS 8 software but is not fully compatible with kernel version 4.18 in CentOS 8. If your business uses the applications that depend on APIs, such as building out-of-tree kernel modules and using hardware resources, you must assess these applications beforehand. This way, you can determine whether a direct migration is feasible and make sure that the applications can run normally after the migration.
MySQL database	Only migrations to the same or later MySQL versions are supported.	<ul> <li>No if the MySQL version provided by Alibaba Cloud Linux 3 is the same as or later than that provided by CentOS 8</li> <li>Yes if the MySQL version provided by Alibaba Cloud Linux 3 is earlier than that provided by CentOS 8</li> </ul>	Alibaba Cloud Linux 3 supports MySQL. However, MySQL version downgrades are not supported. If the MySQL version provided by Alibaba Cloud Linux 3 is earlier than that provided by CentOS 8, you cannot use the migration tool to migrate from CentOS 8 to Alibaba Cloud Linux.

Feature or software	Supported	Migration affected	Description
Operating system hot migration	No.	No.	During the migration procedures, software packages are re-installed. This causes services to restart and results in brief service interruptions. To prevent impacts on services, we recommend that you back up data and isolate the source operating system environment before you use the migration tool to migrate from CentOS 8 to Alibaba Cloud Linux 3. For information about how to back up data, see Snapshot overview.

The migration tool consumes resources to download and install software packages during migration. Before you perform a migration on an Elastic Compute Service (ECS) instance, make sure that the following resources are available:

- Amount of memory space: 300 MiB
- Amounts of disk space available for the following directories:
  - /usr. 250 MiB/var. 1.5 GiB/boot: 50 MiB

**? Note** If your instance has a smaller number of vCPUs, little memory, and a small amount of storage space (system disk space and data disk space), we recommend that you stop services that are running on the instance before you migrate the instance to Alibaba Cloud Linux 3. This way, you can prevent migration failures caused by insufficient resources.

### Procedure

- Connect to the CentOS 8 ECS instance that you want to migrate.
   For more information, see Connection methodsGuidelines on instance connection.
- 2. Change the Cent OS 8 repository address.
- 3. Run the following commands in sequence to install the migration tool.
  - i. Run the following command to download the software repository of the migration tool:

```
\label{liminary} wget \ http://mirrors.aliyun.com/alinux/migration/alinux-migration.repo \ -O \ /etc/yum.repos.d/alinux-migration.repo
```

ii. Run the following command to install the centos8toalinux3 migration tool:

```
yum -y install centos8toalinux3
```

- 4. Run the following command to run the migration tool.
  - Note The migration tool can record software package information before and after the migration. To view the record files of post-migration software package information, run the centos8toalinux3.py -V command. For more information about how to view the record files of pre- and post-migration software package information, see the FAQ section.

```
centos8toalinux3.py
```

Description of migration results:

• The following command output indicates that the migration tool has terminated the migration because your operating system is configured with a GUI.

```
[root@localhost noarch]# centos&toalinux3.py
Checking if the tool is executed by root user
Checking required packages
Checking if there is GUI installed.
It seems that your system install a GUI module, but still now we don't support a GUI transfer.
```

• The following command output indicates that the migration is complete.

```
Complete!
gpg-pubkey-8483c65d-5ccc5b19 CentOS (CentOS Official Signing Key) <security@centos.org>
remove centos gpg-pubkey
Removing yum cache
Sync successfully, update grub.cfg.
Generating grub configuration file ...
done
Switch complete. Alinux recommends rebooting this system.
```

5. Run the following command to view post-migration operating system information:

```
cat /etc/os-release
```

The following command output indicates that the instance is migrated to Alibaba Cloud Linux 3.

6. Check configurations after the migration.

After the migration is complete, you must check whether the CentOS vault repository remains in the operating system. If the CentOS vault repository remains in the operating system, you must delete the respository to prevent the coexistence of Alibaba Cloud Linux 3 repository and CentOS 8 repository. Otherwise, when you run the **yum update** command, the CentOS package may be installed.

i. Run the following command to check whether the CentOS vault repository remains in the current operating system:

```
yum repolist
```

If a residual CentOS repo file remains in the the operating system, a message similar to the centOS-8.5.2111 - Base - mirrors.aliyun.com message is returned in the following command output:

```
# yum repolist
Repository ID
                                                                                   Rep
ository name
AppStream
                                                                           CentOS-8.5
.2111 - AppStream - mirrors.aliyun.com
alinux3-module
                                                                           alinux3-mo
dule
alinux3-os
                                                                           alinux3-os
alinux3-plus
                                                                           alinux3-pl
alinux3-powertools
                                                                           alinux3-po
wertools
alinux3-updates
                                                                           alinux3-up
dates
                                                                           CentOS-8.5
base
.2111 - Base - mirrors.aliyun.com
                                                                           Extra Pack
ages for Enterprise Linux 8 - x86 64
                                                                           CentOS-8.5
extras
.2111 - Extras - mirrors.aliyun.com
```

ii. Run the following command to check for CentOS repo files:

```
grep -rn "CentOS-8.5.2111" /etc/yum.repos.d/*.repo
```

#### A command output similar to the following one is returned:

```
# grep -rn "CentOS-8.5.2111" /etc/yum.repos.d/*.repo
/etc/yum.repos.d/Centos-vault-8.5.2111.repo:14:name=CentOS-8.5.2111 - Base - mirror
s.aliyun.com
/etc/yum.repos.d/Centos-vault-8.5.2111.repo:24:name=CentOS-8.5.2111 - Extras - mirr
ors.aliyun.com
/etc/yum.repos.d/Centos-vault-8.5.2111.repo:34:name=CentOS-8.5.2111 - Plus - mirror
s.aliyun.com
/etc/yum.repos.d/Centos-vault-8.5.2111.repo:43:name=CentOS-8.5.2111 - PowerTools -
mirrors.aliyun.com
/etc/yum.repos.d/Centos-vault-8.5.2111.repo:53:name=CentOS-8.5.2111 - AppStream - m
irrors.aliyun.com
```

iii. Run the following command to rename the found repo files:

```
mv /etc/yum.repos.d/Centos-vault-8.5.2111.repo /etc/yum.repos.d/Centos-vault-8.5.21
11.repo.bak
```

7. Run the following command to restart the instance to boot into the kernel of Alibaba Cloud Linux 3:

reboot

After the instance is restarted, run the **uname** -r command to view the kernel version information of Alibaba Cloud Linux 3. A command output similar to the following one is displayed.

#### **FAQ**

• How do I check for changes in the operating system that occurred during a migration from Cent OS 8 to Alibaba Cloud Linux 3?

The migration tool records software package information before and after the migration. After you run the **centos8toalinux3.py** -**V** command to migrate from CentOS 8 to Alibaba Cloud Linux 3, you can compare the files of pre- and post-migration software package information to find changes to your operating system. The record files are stored in different names in the /var/tmp directory.

- Record files of pre-migration software package information:
  - *\$(hostname)*-rpms-list-before.log
  - *\$(hostname)*-rpms-verified-before.log
- Record files of post-migration software package information:
  - *\$(hostname)*-rpms-list-after.log
  - *\$(hostname)*-rpms-verified-after.log
- What do I do if an exception occurs and causes a migration from CentOS 8 to Alibaba Cloud Linux 3 to fail?

Operating system migrations cannot be reversed. If an exception occurs during a migration, the operating system cannot be restored to its original state. We recommend that you create snapshots for your ECS instance beforehand. You can use these snapshots to restore disks when migration exceptions occur.

If a software package dependency causes a migration failure (for example, if an error is reported after the yum distro-sync command is run), save the error message and submit a ticket.

### 5.7.2. Migrate from CentOS 7 to Alibaba Cloud

### Linux 2

CentOS has announced the end of life (EOL) for CentOS Linux. You can use a migration tool provided by Alibaba Cloud Linux to smoothly migrate from CentOS 7 to Alibaba Cloud Linux 2. This topic describes how to use the migration tool.

#### Context

Before you use the migration tool to migrate away from CentOS 7, take note of the following items:

- Alibaba Cloud Linux 2 uses kernel version 4.19, which is incompatible with kernel version 3.10 of CentOS 7. Applications that must call the kernel API to program, such as to build out-of-tree kernel modules and to enable hardware, must be assessed beforehand to determine whether a direct migration is feasible and ensure that the applications can run normally after the migration.
- Migration procedures cannot be reversed, and software packages are automatically re-installed

during the procedures. Before you migrate your Elastic Compute Service (ECS) instances away from CentOS Linux, you must create snapshots for the instances to back up data. For more information about snapshots, see Snapshot overview.

• Alibaba Cloud Linux 2 does not support GUIs. As a result, the migration tool cannot be used to migrate away from Cent OS 7 operating systems (OSs) that are configured with a GUI.

#### **Procedure**

- Connect to the CentOS 7 ECS instance that you want to migrate.
   For more information, see Connection methodsGuidelines on instance connection.
- 2. Run the following commands in sequence to install the migration tool.
  - i. Run the following command to download the software repository of the migration tool:

```
wget http://mirrors.aliyun.com/alinux/migration/alinux-migration.repo -0 /etc/yum.r
epos.d/alinux-migration.repo
```

ii. Run the following command to install the centos7toalinux2 migration tool:

```
yum -y install centos7toalinux2
```

- 3. Run the following command to run the migration tool.
  - **Note** The migration tool can record software package information before and after the migration. To view the record files of post-migration software package information, run the centos7toalinux2.py -V command. For more information about how to view the record files of pre- and post-migration software package information, see the FAQ section.

```
centos7toalinux2.py
```

Description of migration results:

• The following command output indicates that the migration tool does not support your CentOS 7 OS and has terminated the migration because your OS is configured with a GUI.

```
[root@localhost ~]# centos7toalinux2.py
Checking if the tool is executed by root user
Checking required packages
Checking if there is GUI installed.
It seems that your system install a GUI module, but still now we don't support a GUI transfer.
```

• The following command output indicates that the migration is complete.

```
Sync successfully, update grub.cfg.

Senerating grub configuration file ...

Found linux image: /boot/vmlinuz-4.19.91-24.1.al7.x86_64

Found initrd image: /boot/initramfs-4.19.91-24.1.al7.x86_64.img

Found linux image: /boot/vmlinuz-3.10.0-1062.18.1.el7.x86_64

Found initrd image: /boot/initramfs-3.10.0-1062.18.1.el7.x86_64.img

Found linux image: /boot/vmlinuz-3.10.0-1062.el7.x86_64

Found initrd image: /boot/initramfs-3.10.0-1062.el7.x86_64.img

Found linux image: /boot/vmlinuz-0-rescue-20200426154603174201708213343640

Found initrd image: /boot/initramfs-0-rescue-20200426154603174201708213343640.img

done

Switch complete. Alinux recommends rebooting this system.
```

4. Run the following command to view post-migration OS information:

```
cat /etc/os-release
```

The following command output indicates that the instance is migrated to Alibaba Cloud Linux 2.

5. Run the following command to restart the instance to boot into the kernel of Alibaba Cloud Linux 2:

```
reboot
```

After the instance is restarted, run the **uname** -r command to view the kernel version information of Alibaba Cloud Linux 2: A command output similar to the following one is displayed.

#### **FAO**

• How do I check for changes in the OS that occurred during a migration from Cent OS 7 to Alibaba Cloud Linux 2?

The migration tool records software package information before and after the migration. After you run the **centos7toalinux2.py** -**V** command to migrate from CentOS 7 to Alibaba Cloud Linux 2, you can compare the record files of pre- and post-migration software package information to find changes to your OS. The record files are stored in different names in the /var/tmp directory.

- Record files of pre-migration software package information:
  - \$(hostname)-rpms-list-before.log
  - \$(hostname)-rpms-verified-before.log
- Record files of post-migration software package information:
  - *\$(hostname)*-rpms-list-after.log
  - *\$(hostname)*-rpms-verified-after.log
- What do I do if an exception occurs and causes a migration from CentOS 7 to Alibaba Cloud Linux 2 to fail?

OS migrations cannot be reversed. If an exception occurs during a migration, the OS cannot be reverted to its original state. We recommend that you create snapshots for your ECS instance to back up data so that you can use the snapshots to roll back disks in case of migration exceptions.

If a software package dependency causes a migration failure, for example, if an error is reported after the yum distro-sync command is run, you must save the error message and submit a ticket.

# 5.8. Known issues in Alibaba Cloud Linux

This topic describes known issues in Alibaba Cloud Linux images, the scope of these issues, and their corresponding solutions.

#### Known issues in Alibaba Cloud Linux 2:

- Performance issues may occur if you enable the CONFIG\_PARAVIRT\_SPINLOCK kernel feature
- The system may become unstable and performance issues may occur if you enable the Transparent Huge Page (THP) kernel feature
- A delegation conflict occurs in NFS v4.0
- Defects in NFS v4.1 or v4.2 cause applications not to exit
- System performance is affected after security vulnerabilities such as Meltdown and Spectre are fixed

# Performance issues may occur if you enable the CONFIG\_PARAVIRT\_SPINLOCK kernel feature

- Problem description: After you enable the CONFIG\_PARAVIRT\_SPINLOCK kernel feature, application
  performance is significantly affected if an Elastic Compute Service (ECS) instance has a large number
  of vCPUs and a large number of lock contentions exist in applications. For example, an NGINX
  application experiences a significant performance decrease when it processes short-lived
  connections.
- Solution: We recommend that you keep the CONFIG\_PARAVIRT\_SPINLOCK kernel feature disabled in Alibaba Cloud Linux 2. By default, this feature is disabled. If you are unsure about how to resolve this kernel issue, do not enable CONFIG\_PARAVIRT\_SPINLOCK.

# The system may become unstable and performance issues may occur if you enable the Transparent Huge Page (THP) kernel feature

- Problem description: After you enable the THP kernel feature by setting the transparent\_hugepage/enabled option to *always* in your production environment, the system becomes unstable and its performance deteriorates.
- Solution: For information about how to optimize THP-related performance, see THP-related performance optimization in Alibaba Cloud Linux 2.

#### A delegation conflict occurs in NFS v4.0

- Problem description: A delegation conflict occurs in Network File System (NFS) v4.0. For more information, see Delegation in NFS Version 4.
- Solution: We recommend that you disable the delegation feature when you use NFS v4.0. For information about how to disable this feature on the server side, see How to Select Different Versions of NFS on a Server.

#### Defects in NFS v4.1 or v4.2 cause applications not to exit

- Problem description: In NFS v4.1 or v4.2, if you use asynchronous I/O (AIO) in applications to send requests and then close the corresponding file descriptors before I/O responses are returned, a livelock may be triggered. This causes a failure in which the corresponding process exits.
- Solution: This problem is fixed in kernel versions 4.19.30-10.al7 and later. Application exit failures are not likely to occur. You can decide whether to upgrade the kernel to fix this issue. To upgrade the kernel, run the sudo yum update kernel -y command.

#### Notice ■

- The kernel upgrade may result in a system boot failure. Proceed with caution when you perform the operation.
- Before you upgrade the kernel, make sure that you have created snapshots or custom images to back up data. For more information, see Create a snapshot of a disk or Create a custom image from an instance.

# System performance is affected after security vulnerabilities such as Meltdown and Spectre are fixed

- Problem description: In the kernel of Alibaba Cloud Linux 2, the repair of important security vulnerabilities such as Meltdown or Spectre in processors is enabled by default. This affects system performance. Performance may be deteriorated during performance benchmark testing.
- Solution: Melt down and Spectre are two critical vulnerabilities discovered in Intel chips. Attackers can
  exploit these vulnerabilities to steal sensitive application data from the system memory. We
  recommend that you enable the repair feature. However, if you want to maximize system
  performance, you can disable the repair feature. For more information, see How to fix CPU
  vulnerabilities in the Alibaba Cloud Linux 2 system.

## 5.9. Knowledge base

This topic provides the troubleshooting information about and solutions to Alibaba Cloud Linux issues.

- Alibaba Cloud Linux 2
  - Solution to polkit memory leaks in Alibaba Cloud Linux 2
  - o systemd service error in Alibaba Cloud Linux 2
  - Prevent application performance jitter caused by cgroups
  - Causes of and solutions to the issue of OOM Killer being triggered
  - Identify the causes of high percentage of the slab\_unreclaimable memory in the Linux operating system
  - Solutions to memory fragmentation in Linux operating systems
  - Replace systemd-network settings with network-scripts settings for Alibaba Cloud Linux 2 ECS instances
  - How do I configure the kdump file in Alibaba Cloud Linux 2
  - Delay in system time synchronization after an ECS instance that runs an Alibaba Cloud Linux 2.1903 image is started or restarted
  - The yum command cannot be run on Alibaba Cloud Linux 2 ECS instances located in the classic network
  - A segfault error occurs when Alibaba Cloud Linux 2 ECS instances run Docker images of earlier versions
  - The route information cannot be queried or configured on Alibaba Cloud Linux 2 ECS instances
  - The application on the ECS instance occasionally suffers packet loss and the kernel Log (dmesg) contains the error message "kernel: nf\_conntrack: table full, dropping packet"
  - Alibaba Cloud Linux 2 ECS instances cannot create a large number of processes

- Influence of TCP congestion control algorithm BBR on network performance in Alibaba Cloud Linux 2
- The Buffer I/O write performance of the Ext4 file system in Alibaba Cloud Linux 2 ECS instances is lower than expected
- The return value of Send-Q is 0 when the ss command is run on an Alibaba Cloud Linux 2 ECS instance
- THP-related performance optimization in Alibaba Cloud Linux 2
- Alibaba Cloud Linux 3
  - Solution to the issue that the value of load is greater than 1 when the Alibaba Cloud Linux 3 operating system has no load
  - Third-party DNF or YUM repositories such as Docker-CE and EPEL cannot be used to install software on Alibaba Cloud Linux 3 ECS instances

### 5.10. Alibaba Cloud Linux FAQ

This topic provides answers to frequently asked questions about Alibaba Cloud Linux images.

- What are the differences between Alibaba Cloud Linux 2 and Alibaba Cloud Linux 3?
- What are the differences among Alibaba Cloud Linux images?
- How do I use Alibaba Cloud Linux images in Alibaba Cloud?
- Am I charged for using Alibaba Cloud Linux images in Alibaba Cloud ECS?
- Which ECS instance types do Alibaba Cloud Linux images support?
- Do Alibaba Cloud Linux images support 32-bit applications or libraries?
- Does Alibaba Cloud Linux provide a graphical user interface (GUI)?
- Can I use Alibaba Cloud Linux images in an on-premises environment?
- Can I view the source code of Alibaba Cloud Linux 2 components?
- Is Alibaba Cloud Linux 2 backward-compatible with the historical Aliyun Linux versions?
- Which third-party applications can run on Alibaba Cloud Linux 2?
- What are the advantages of Alibaba Cloud Linux 2 over other Linux operating systems?
- How does Alibaba Cloud Linux 2 protect data security?
- Does Alibaba Cloud Linux 2 support data encryption?
- How do I grant permissions to manage Alibaba Cloud Linux 2?

# What are the differences between Alibaba Cloud Linux 2 and Alibaba Cloud Linux 3?

Alibaba Cloud Linux 2, which was originally called Aliyun Linux 2, differs from Alibaba Cloud Linux 3 in kernel version, compiler, and lifecycle. The following table describes the differences.

Difference	Alibaba Cloud Linux 2	Alibaba Cloud Linux 3
Kernel version	4.19	5.10
Compiler	GCC 4.8.5	GCC 10.2
End-of-life date	2024-03-31	2031-04-30

#### What are the differences among Alibaba Cloud Linux images?

For information about the release notes of Alibaba Cloud Linux images, see Alibaba Cloud Linux 2镜像发布 记录 and Alibaba Cloud Linux 3镜像发布记录. For information about the relationships among Alibaba Cloud Linux images, see the "Images" section in Alibaba Cloud Linux images.

#### How do I use Alibaba Cloud Linux images in Alibaba Cloud?

Alibaba Cloud provides Alibaba Cloud Linux public images. When you create an Elastic Compute Service (ECS) instance, you can click Public Image and then select a version of Alibaba Cloud Linux images.

# Am I charged for using Alibaba Cloud Linux images in Alibaba Cloud ECS?

No, Alibaba Cloud Linux images are provided free of charge. You are charged only for the ECS instances that use the images.

#### Which ECS instance types do Alibaba Cloud Linux images support?

Alibaba Cloud Linux images support most ECS instance types, including ECS Bare Metal Instance types.

**Note** Alibaba Cloud Linux images do not support instances that run on the Xen virtual machine monitor.

# Do Alibaba Cloud Linux images support 32-bit applications or libraries?

No, Alibaba Cloud Linux images do not support 32-bit applications or libraries.

#### Does Alibaba Cloud Linux provide a graphical user interface (GUI)?

No, Alibaba Cloud Linux 2 and 3 images do not provide a graphical user interface (GUI).

# Can I use Alibaba Cloud Linux images in an on-premises environment?

Yes, you can use Alibaba Cloud Linux images in an on-premises environment. Alibaba Cloud Linux provides local images in the QCOW2 format. These images support only kernel-based virtual machines (KVMs). For more information, see Use Alibaba Cloud Linux 2 images in an on-premises environment or Use Alibaba Cloud Linux 3 images in an on-premises environment.

#### Can I view the source code of Alibaba Cloud Linux 2 components?

Yes, Alibaba Cloud Linux 2 is open source. You can use the yumdownloader tool or visit the official Alibaba Cloud download pages to download the source code package. You can also download the source code tree of the Alibaba Cloud Linux 2 kernel from GitHub. For more information, visit Github.

# Is Alibaba Cloud Linux 2 backward-compatible with the historical Aliyun Linux versions?

Yes, Alibaba Cloud Linux 2 is compatible with Aliyun Linux 17.01.

Note You may need to re-compile a compiled kernel module on Alibaba Cloud Linux 2 before the module can be used.

#### Which third-party applications can run on Alibaba Cloud Linux 2?

Alibaba Cloud Linux 2 is binary compatible with Cent OS 7.6.1810. Therefore, applications that can run on Cent OS can also run on Alibaba Cloud Linux 2.

# What are the advantages of Alibaba Cloud Linux 2 over other Linux operating systems?

Alibaba Cloud Linux 2 is binary compatible with CentOS 7.6.1810 and provides differentiated operating system features.

Compared with CentOS and Red Hat Enterprise Linux (RHEL), Alibaba Cloud Linux 2 has the following advantages:

- Frequent updates and updated Linux kernels, user-mode software, and toolkits are provided.
- Alibaba Cloud Linux 2 works out of the box and requires minimal configurations.
- Alibaba Cloud Linux 2 is optimized to work with cloud infrastructure and maximizes performance for users.
- Unlike RHEL, Alibaba Cloud Linux 2 does not have runtime charges. Different from Cent OS, Alibaba Cloud provides commercial support for Alibaba Cloud Linux 2.

#### How does Alibaba Cloud Linux 2 protect data security?

Alibaba Cloud Linux 2 is binary compatible with CentOS 7.6.1810 and RHEL 7.6 and complies with the RHEL safety specifications. Alibaba Cloud Linux 2 uses the following methods to protect your data:

- Uses industry-standard vulnerability scan and security test tools to perform periodical security scanning.
- Periodically assesses the CVE patch updates of CentOS 7 to fix operating system security vulnerabilities.
- Supports existing solutions of Alibaba Cloud for operating system security hardening.
- Uses the same mechanism as Cent OS 7 to release user security alerts and patch updates.

#### Does Alibaba Cloud Linux 2 support data encryption?

Yes, Alibaba Cloud Linux 2 retains the CentOS 7 data encryption toolkit and supports the data encryption solution coordinately implemented by Key Management Service (KMS) and CentOS 7.

#### How do I grant permissions to manage Alibaba Cloud Linux 2?

You can grant management permissions in Alibaba Cloud Linux 2 in the same manner as you would in CentOS 7. This means you can use the same commands to grant management permissions in Alibaba Cloud Linux 2 as you would in CentOS 7. The default permission configurations in images of these two operating systems are the same.

# 6.Custom image 6.1. Overview

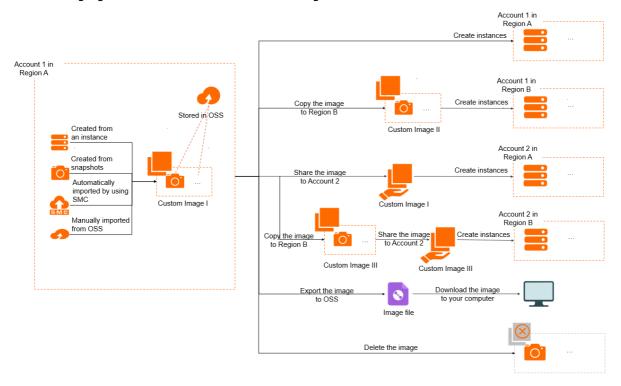
Custom images are created from Elastic Compute Service (ECS) instances or snapshots or imported from your computer. Custom images can contain applications and data. You can use custom images to create instances that have identical configurations. This eliminates the need to make repeated configurations.

#### Use a custom image

After you create or import a custom image, the image is in the Available state. You can use the image to create instances, share the image to other Alibaba Cloud accounts, copy the image to other regions, or export the image to an Object Storage Service (OSS) bucket. You can also delete the image when it is no longer needed.

**Note** Only the creator of a custom image can use, share, copy, and delete the image.

The following figure shows how to use a custom image.



#### **Operations**

The following table describes the operations that you can perform on custom images.

Operation	Description	References
Operation	Description	References

Operation	Description	References
Create an image	You can create a custom image from an instance or from a snapshot to copy the system environment. This eliminates the need to make repeated configurations. You can also use the custom image tool Packer to create a custom image.  If you have specified an image family when you created a custom image, you can deprecate or restore the image. To roll an image family back to a previous image that is not deprecated, you can deprecate the most recent image in the image family. If a custom image in an image family is deprecated, you can restore the image to use it as the most recent custom image in the image family.	<ul> <li>Create a custom image from an instance</li> <li>Create a custom image from a snapshot</li> <li>Use Packer to create a custom image</li> <li>Use Packer to create and import an on-premises image</li> <li>Set the state of an image to Deprecated or restore an image</li> </ul>
Import an image	You can import a custom image from your computer.	Image import procedure
Update an image	Operation Orchestration Service (OOS) provides public templates to automatically update custom images. To update a custom image, you need to only create a random or scheduled O&M task by selecting a source image and specifying required parameters (such as Cloud Assistant command-related parameters) in a public template. Then, the O&M task is automatically executed as defined in the template to update the image.	Update a custom image
Copy an image	If you want to use a custom image within another region, copy the image to that region. The image copy is assigned a unique ID and is independent of the copied image.	Copy a custom image
Share an image	You can share a custom image that you created to other Alibaba Cloud accounts. The Alibaba Cloud accounts can create instances from the shared image.	Share a custom image
Export an image	You can export a custom image that you created to an OSS bucket and then download the image to your computer.	Export a custom image
Modify image information	You can modify the name and description of a custom image.	Modify a custom image
Delete an image	You can delete a custom image that you no longer need.	Delete a custom image

# 6.2. Create custom image

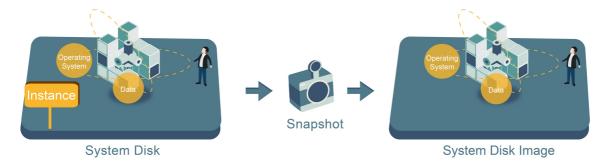
### 6.2.1. Create a custom image from a snapshot

You can create a custom image from a snapshot that contains the operating system and data of an Elastic Compute Service (ECS) instance. Then, you can use the custom image to create multiple identical instances.

#### **Prerequisites**

A system disk snapshot is created. For more information, see Create a snapshot of a disk.

#### Context



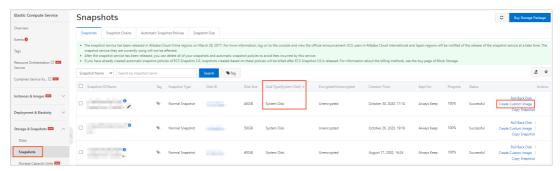
Before you create custom images from snapshots, take note of the following items:

- Notes on snapshots used to create custom images:
  - A custom image can be created from a system disk snapshot or from a system disk snapshot and one or more data disk snapshots. Data disk snapshots alone cannot be used to create custom images.
  - Both encrypted and unencrypted snapshots can be used to create custom images.
  - If the ECS instance from which a snapshot was created expires or is released, the custom image created from the snapshot and the ECS instance created from the image are not affected.
- Notes on custom images:
  - Custom images cannot be used across regions. To use custom images across regions, you can copy custom images from one region to another. For more information, see Copy a custom image.
  - The billing method of custom images is independent of that of the ECS instances from which the
    custom images were created and that of the ECS instances created from the custom images. For
    example, custom images created from subscription instances can be used to create pay-as-you-go
    instances.
- Notes on ECS instances created from custom images:
  - You can upgrade the configurations of ECS instances created from custom images, such as vCPUs, memory, bandwidth, and disks.
  - You can replace the operating systems of ECS instances created from custom images. The custom
    images are not affected and remain usable. For more information, see Replace the operating
    system of an instance by using a non-public image.
  - Network connection may fail to be established for some ECS Linux instances in virtual private clouds (VPCs) that are created from custom images. This may be caused by the configurations in /e tc/sysconfig/network. For more information, see How to solve unreachable network errors when a VPC-type instance is created from a custom image.
- Recommendations for data security:

Delete sensitive data from a snapshot in advance to enhance data security.

#### **Procedure**

- 1.
- 2. Use one of the following methods to find the system disk snapshot from which you want to create a custom image:
  - o On the Instances page
    - a.
    - b.
    - c. Find the instance. Click the instance ID or click Manage in the Actions column.
    - d. On the Instance Details page, click the **Snapshot** tab. Find the snapshot whose **Disk Type(All)** is **System Disk**. Click **Create Custom Image** in the **Actions** column.
  - o On the Snapshots page
    - a.
    - b.
    - c. Find the snapshot whose **Disk Type(All)** is **System Disk**. Click **Create Custom Image** in the Actions column.



3. In the Create Custom Image dialog box, configure parameters described in the following table.

Parameter	Description References	
System Snapshot ID	The snapshot must be a system disk snapshot.	N/A
Custom Image Name and Custom Image Description	Enter an image name in the Custom Image Name field and enter a description for the image in the Custom Image Description field.	
Image Family	Optional. Select an image family. Custom images that belong to an image family can be deprecated or restored to allow smooth image updates and rollbacks.	
Resource Group	Optional. Select a resource group to manage the custom image based on the usage and permissions.  Resource groups	

Parameter	Description	References
Add Data Disk Snapshot	<ul> <li>Optional. Select Add Data Disk Snapshot, click Add in the Snapshot Details section, and then select the ID of a data disk snapshot.</li> <li>If you do not select the ID of a data disk snapshot, an empty data disk is created with a default capacity of 5 GiB.</li> <li>If you select the ID of a data disk snapshot, the capacity of the data disk to be created is the same as the snapshot size.</li> </ul>	Snapshot overview
Tag	<ul> <li>Select a tag.</li> <li>Optional: In most scenarios, the tag is optional.</li> <li>Required: When you log on as a RAM user and a policy attached to the RAM user requires that specific tags be added, you must configure this parameter to specify the tags. Otherwise, an error is reported and indicates insufficient permissions.</li> </ul>	Create a resource with a specific tag

4. Click Create.

#### What's next

After the custom image is created, you can perform operations described in the following topics:

- Create an ECS instance by using a custom image
- Replace the operating system of an instance by using a non-public image

#### Related information

Createlmage

### 6.2.2. Create a custom image from an instance

After you create an Elastic Compute Service (ECS) instance, you can customize the instance by performing operations such as installing software and deploying application environments and then create a custom image from the instance. Instances created from the custom image contain all of the customized items, which eliminates the need to configure these items for each new instance.

#### **Prerequisites**

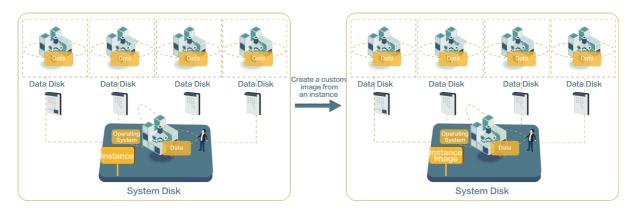
Sensitive data is deleted from the instance to enhance data security.

Before you create a custom image from a Linux instance, the following items are checked:

- Network configurations of the instance. The network configurations are checked based on your Linux distribution. For more information, see How to solve unreachable network errors when a VPC-type instance is created from a custom image.
- Free space on the system disk. The system disk must have sufficient free space.

#### Context

When you create a custom image from an instance, a snapshot is automatically created for each disk on the instance. All of the snapshots together constitute a complete custom image, as shown in the following figure.



Before you create a custom image from an instance, take note of the following items:

- When you create a custom image from an instance, you do not need to stop the instance.
  - Note When you create a custom image from an instance that is in the Running state, real-time data may not be saved to the image. To ensure data integrity, we recommend that you stop the instance before you use it to create a custom image. If you do not want to stop your instance, you can create a snapshot-consistent group from the instance and then use the snapshot-consistent group to create a custom image. For more information, see Create a snapshot-consistent group.
- While a custom image is being created from an instance, the state of the instance must not be changed. For example, if you stop, start, or restart an instance while a custom image is being created from the instance, the image cannot be created.
- You cannot create images from expired subscription instances. However, you can create snapshots for the system disks of the instances and then use the snapshots to create custom images.
- You cannot create images from released instances. However, if you have stored a system disk snapshot of a released instance, you can use the snapshot to create a custom image.
- A custom image resides within the same region as the instance from which the image was created. For example, if an instance is located in the China (Hangzhou) region, the image created from the instance is also located in the China (Hangzhou) region. For information about how to use images across regions, see Copy a custom image.
- The amount of time it takes to create an image depends on the disk size of the instance from which the image is created.
- You cannot create images that contain snapshots of system disks and data disks from instances that use local SSDs. For more information, see Overview.

When you create a custom image from a Linux instance, take note the following additional items:

- Do not upload data disk information to the /etc/fstab file. Otherwise, instances created from the image cannot be started.
- Do not update the kernel or operating system version.
- Do not adjust the system disk partitions. Only system disks with a single root partition are supported.
- Do not modify critical system files such as /sbin, /bin, and /lib.
- Do not modify the default logon username root.

#### **Procedure**

- 1.
- 2.
- 3.
- 4. Find the instance from which you want to create a custom image. In the **Actions** column, choose **More > Disk and Image > Create Custom Image**.
- 5. In the **Create Custom Image** dialog box, configure the parameters described in the following table.

Parameter	Description	References
Custom Image Name and Custom Image Description	Enter a name and description for the custom image.	N/A
Image Family	Optional. You can select an image family when you create a custom image. Custom images that belong to an image family can be deprecated or restored to allow smooth image updates and rollbacks.	Overview
Resource Group	Optional. Select a resource group to manage the custom image based on usage and permissions.  Resource groups	
Tag	<ul> <li>Select one or more tags.</li> <li>Optional: In most scenarios, Tags is an optional parameter.</li> <li>Required: When you log on as a RAM user and a policy attached to the RAM user requires that tags be added, you must configure this parameter. Otherwise, an error is reported and indicates insufficient permissions.</li> </ul>	Create a resource with a specific tag

#### 6. Click Create.

**? Note** The image can be used only after snapshots are created for all disks of the instance. Wait until the snapshots are created.

#### What's next

After you create a custom image, you can perform the following operations:

- Create an ECS instance by using a custom image
- Share or unshare a custom image
- Export a custom image

#### Related information

- Create a custom image from a snapshot
- Modify a custom image

Createlmage

### 6.2.3. Use Packer to create a custom image

Packer is a lightweight open source tool for creating images and runs on commonly used operating systems such as Windows, Linux, and macOS. This topic describes how to install and use Packer to create a custom image.

#### Context

In this topic, a Linux instance is used. For information about how to install Packer in Windows, visit Install Packer.

#### Step 1: Install Packer

- Connect to a Linux Elastic Compute Service (ECS) instance.
   For more information, see Connect to a Linux instance by using a password.
- 2. Run the cd /usr/local/bin command to access the /usr/local/bin directory.
  - **?** Note The /usr/local/bin directory is a directory that stores environment variables. You can install Packer to this directory or a directory that is added to environment variables.
- 3. Run the wget https://releases.hashicorp.com/packer/1.1.1/packer\_1.1.1\_linux\_amd64.zip command to download the Packer package.

You can go to the Download Packer page to download other versions of Packer packages.

- 4. Run the unzip packer 1.1.1 linux amd64.zip command to decompress the Packer package.
- 5. Run the packer -v command to verify the installation status of Packer.
  - o If Packer is installed, the Packer version number is returned.
  - o If Packer is not properly installed, the command not found error message is returned.

#### Step 2: Define a Packer template

If you want to use Packer to create a custom image, you must create a template file in the JSON format. In the template file, you must specify the image builder and provisioner. For more information, visit Alicloud Image Builder and Provisioners. Packer provides a variety of provisioners that allow you to configure the content generation mode for custom images. In this example, a Shell provisioner is used to define a Packer template.

Create a JSON file named *alicloud* in the Linux instance and copy the following content to the file:

```
"variables": {
 "access key": "{{env `ALICLOUD ACCESS KEY`}}",
  "secret_key": "{{env `ALICLOUD_SECRET_KEY`}}"
},
"builders": [{
  "type": "alicloud-ecs",
  "access_key":"{{user `access_key`}}",
  "secret_key":"{{user `secret_key`}}",
  "region": "cn-beijing",
  "image name": "packer basic",
  "source_image":"centos_7_02_64_20G_alibase_20170818.vhd",
  "ssh username":"root",
  "instance_type":"ecs.nl.tiny",
  "internet_charge_type":"PayByTraffic",
  "io_optimized":"true"
}],
"provisioners": [{
 "type": "shell",
  "inline": [
    "sleep 30",
    "yum install redis.x86_64 -y"
} ]
```

#### The following table describes the parameters that you must specify.

Parameter	Description	
	Your AccessKey ID. For more information, see Obtain an AccessKey pair.	
Note To prevent disclosing the AccessKey pair of your Alibaba Cloud account, we recommend that you create a RAM user and use the credentials of the RAM user to create an AccessKey pair. For more information, see Create a RA user and Obtain an AccessKey pair.		
secret_key	Your AccessKey secret. For more information, see Obtain an AccessKey pair.	
region	The region of the temporary instance used to create the custom image.	
image_name	The name of the custom image.	
source_image	The name of the source image used to create the custom image. You can obtain the name from the public image list of Alibaba Cloud.	
instance_type	The instance type of the temporary instance used to create the custom image.	
internet_charge_ty pe	The billing method for network usage of the temporary instance used to create the custom image.	

Parameter	Description
provisioners	The type of the provisioner used to create the custom image. For more information, visit Provisioners.

#### Step 3: Use Packer to create a custom image

Perform the following operations to create a custom image by using the Packer template file that you specified:

- 1. Run the export ALICLOUD\_ACCESS\_KEY=<Your AccessKey ID> command to import your AccessKey ID.
- 2. Run the export ALICLOUD\_SECRET\_KEY=<Your AccessKey secret> command to import your AccessKey secret.
- 3. Run the packer build alicloud.json command to create a custom image.

A command output similar to the following one indicates that a custom image that has ApsaraDB for Redis deployed is created:

```
alicloud-ecs output will be in this color.
==> alicloud-ecs: Prevalidating alicloud image name...
alicloud-ecs: Found image ID: centos_7_02_64_20G_alibase_20170818.vhd
==> alicloud-ecs: Start creating temporary keypair: packer 59e44f40-c8d6-0ee3-7fd8-blba08ea
94h8
==> alicloud-ecs: Start creating alicloud vpc
==> alicloud-ecs: Provisioning with shell script: /var/folders/3q/w38xx js6cl6k5mwkrqsnw7w0
000gn/T/packer-shell257466182
alicloud-ecs: Loaded plugins: fastestmirror
alicloud-ecs: Total
                                                                 1.3 MB/s | 650 kB 00:00
alicloud-ecs: Running transaction check
==> alicloud-ecs: Deleting temporary keypair...
Build 'alicloud-ecs' finished.
==> Builds finished. The artifacts of successful builds are:
--> alicloud-ecs: Alicloud images were created:
cn-beijing: m-2ze12578be1oa4ovs6r9
```

#### What's next

Create an ECS instance by using a custom image

#### Related information

- packer-provider
- Packer Documentation

# 6.2.4. Use Packer to create and import an onpremises image

Packer is a lightweight open source tool for creating images and runs on commonly used operating systems such as Windows, Linux, and macOS. This topic describes how to create an on-premises image for CentOS 6.9 and upload the image to Alibaba Cloud. You can create a Packer template to create images for other operating systems.

#### **Prerequisites**

- An AccessKey pair is created. For more information, see Obtain an AccessKey pair.
  - Note To prevent disclosing the AccessKey pair of your Alibaba Cloud account, we recommend that you create a RAM user and use the credentials of the RAM user to create an AccessKey pair. For more information, see Create a RAM user.
- The Object Storage Service (OSS) service is activated. For more information, see Activate OSS.

#### Context

In this topic, the following software versions are used. Operations may vary based on your software version.

- Operating system of the on-premises server: Ubuntu 16.04
- Operating system for which to create an image: CentOS 6.9

#### Example

In this example, CentOS 6.9 is used. Perform the following operations to create an on-premises image:

1. Run the egrep "(svm|vmx)" /proc/cpuinfo command to check whether your on-premises server or virtual machine supports Kernel-based Virtual Machine (KVM).

If KVM is supported, a command output similar to the following one is returned:

```
pat pse36 clflush dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx pdpe1gb rdtscp lm constant_tsc art arch_perfmon pebs bts rep_good nopl xtopology nonstop_tsc aperfmperf t sc_known_freq pni pclmulqdq dtes64 monitor ds_cpl vmx est tm2 ssse3 sdbg fma cx16 xtpr pdcm pcid sse4_1 sse4_2 x2apic movbe popcnt tsc_deadline_timer aes xsave avx f16c rdran d lahf_lm abm 3dnowprefetch epb intel_pt tpr_shadow vnmi flexpriority ept vpid fsgsbase tsc_adjust bmi1 avx2 smep bmi2 erms invpcid mpx rdseed adx smap clflushopt xsaveopt xsa vec xgetbv1 xsaves dtherm ida arat pln pts hwp hwp_notify hwp_act_window hwp_epp flags : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov
```

2. Run the following commands in sequence to install KVM if your on-premises server or virtual machine supports KVM:

```
sudo apt-get install qemu-kvm qemu virt-manager virt-viewer libvirt-bin bridge-uti
ls #Install KVM and related controls.

sudo virt-manager #Enable virt-manager.
```

If KVM is installed, you can create a virtual machine on the Graphical User Interface (GUI).

3. Inst all Packer.

For more information, see Use Packer to create a custom image.

4. Run the following commands in sequence to create a Packer template.

**Note** The following sample commands are based on the Cent OS 6.9 operating system. If you want to create on-premises images for other operating systems, create the *centos.json* configuration file. For more information, see the "Create a Packer template" in this topic.

```
cd /usr/local #Switch the directory.

wget https://github.com/hashicorp/packer-plugin-alicloud/blob/main/builder/examples
/local/centos.json #Download the centos.json configuration file that is released b
y Alibaba Cloud.

wget https://github.com/hashicorp/packer-plugin-alicloud/blob/main/builder/examples
/local/http/centos-6.8/ks.cfg #Download the ks.cfg configuration file that is rele
ased by Alibaba Cloud.

mkdir -p http/centos-6.9 #Create a directory.

mv ks.cfg http/centos-6.9/ #Move the ks.cfg file to the http/centos-6.9 directory.
```

5. Run the following commands in sequence to create and import an on-premises image:

```
export ALICLOUD_ACCESS_KEY=<Your AccessKey ID> #Import your AccessKey ID.

export ALICLOUD_SECRET_KEY=<Your AccessKey secret> #Import your AccessKey secret.

packer build centos.json #Create an on-premises image and import it to Alibaba Cloud.
```

#### A command output similar to the following one is returned:

```
qemu output will be in this color.
==> qemu: Downloading or copying ISO
    qemu: Downloading or copying: http://mirrors.aliyun.com/centos/6.9/isos/x86_64/Cen
tOS-6.9-x86_64-minimal.iso
....
==> qemu: Running post-processor: alicloud-import
    qemu (alicloud-import): Deleting import source https://oss-cn-beijing.aliyuncs.com
/packer/centos_x86_64
Build 'qemu' finished.
==> Builds finished. The artifacts of successful builds are:
--> qemu: Alicloud images were created:
cn-beijing: XXXXXXXXX
```

6. Wait a few minutes and go to the Images page in the ECS console to view the custom image in the specified region. In this example, the image is created in the China (Beijing) region.

#### Create a Packer template

In the preceding example, CentOS 6.9 is used to create an on-premises image. For more information, see the "Example" in this topic. The following JSON file shows how to create a Packer template based on CentOS 6.9:

```
{"variables": {
```

```
"box basename": "centos-6.9",
        "build timestamp": "{{isotime \"20060102150405\"}}",
        "cpus": "1",
        "disk_size": "4096",
        "git revision": " unknown git_revision__",
        "headless": "",
        "http proxy": "{{env `http proxy`}}",
        "https proxy": "{{env `https proxy`}}",
        "iso checksum type": "md5",
        "iso checksum": "af4axxxxxxxxxxxxxxx192a2",
        "iso name": "CentOS-6.9-x86 64-minimal.iso",
        "ks path": "centos-6.9/ks.cfg",
        "memory": "512",
        "metadata": "floppy/dummy metadata.json",
        "mirror": "http://mirrors.aliyun.com/centos",
        "mirror_directory": "6.9/isos/x86_64",
        "name": "centos-6.9",
        "no proxy": "{{env `no proxy`}}",
        "template": "centos-6.9-x86 64",
        "version": "2.1.TIMESTAMP"
     },
    "builders":[
       {
          "boot command": [
           "<tab> text ks=http://{{ .HTTPIP }}:{{ .HTTPPort }}/{{user `ks path`}}<enter><w
ait>"
          "boot wait": "10s",
          "disk size": "{{user `disk size`}}",
          "headless": "{{ user `headless` }}",
          "http directory": "http",
          "iso_checksum": "{{user `iso_checksum`}}",
          "iso checksum type": "{{user `iso checksum type`}}",
          "iso_url": "{{user `mirror`}}/{{user `mirror_directory`}}/{{user `iso_name`}}",
          "output directory": "packer-{{user `template`}}-qemu",
          "shutdown command": "echo 'vagrant'|sudo -S /sbin/halt -h -p",
          "ssh_password": "vagrant",
          "ssh port": 22,
          "ssh username": "root",
          "ssh wait timeout": "10000s",
          "type": "qemu",
          "vm name": "{{ user `template` }}.raw",
          "net device": "virtio-net",
          "disk interface": "virtio",
          "format": "raw"
        }
      ],
    "provisioners": [{
          "type": "shell",
          "inline": [
                "sleep 30",
                "yum install cloud-util cloud-init -y"
              ]
        }],
```

In this example, the QEMU Packer builder is used to create a virtual machine image. The following table describes the required parameters for the QEMU Packer builder. For more information about optional parameters, visit QEMU Builder.

Required parameters for the QEMU Packer builder

Parameter	Туре	Description
iso_checksum	string	The checksum for the ISO file of the operating system. Packer verifies this parameter before a virtual machine to which the ISO file is mounted is started. Make sure that you specify at least one of the iso_checksum and iso_checksum_url parameters. If you specify the iso_checksum parameter, the value of iso_checksum_url is ignored.
iso_checksum_typ e	string	The checksum type of the ISO file of the specified operating system. Valid values:  • none: If you set iso_checksum_type to none, the checksum process is ignored. This value is not recommended.  • md5  • sha1  • sha256  • sha512
iso_checksum_url	string	A URL that points to a GNU- or BSD-style checksum file that contains the ISO file checksum of an operating system. Make sure that you specify at least one of the iso_checksum and iso_checksum_url parameters. If you specify the iso_checksum parameter, the value of iso_checksum_url is ignored.
iso_url	string	<ul> <li>A URL that points to the ISO file and contains the installation image. This URL can be an HTTP URL or a file path:</li> <li>If the URL is an HTTP URL, Packer downloads the ISO file from the HTTP URL and caches the file for running it later.</li> <li>If the URL is a file path of an IMG or QCOW2 file, the QEMU Packer builder uses the file to create a custom image. If you specify the file path, set disk_image to true.</li> </ul>

Parameter	Туре	Description
headless	boolean	By default, Packer starts the GUI to build a QEMU virtual machine. If you set headless to <i>True</i> , a virtual machine without a console is started.

The provisioner used in the preceding example contains a Post-Processor module that enables automated upload of on-premises images to Alibaba Cloud. The following table describes the required parameters for the Packer provisioner. For more information about optional parameters, visit Alicloud Import Post-Processor.

#### Required parameters for the Packer provisioner

Parameter	Туре	Description
access_key	string	Your AccessKey ID.
secret_key	string	Your AccessKey secret.
region	string	The ID of the region to which you want to upload your on- premises image. In this example, the region ID is cn- beijing . For more information about regions, see Regions and zones.
image_name	string	<ul> <li>The name of your on-premises image.</li> <li>The name must be 2 to 128 characters in length.</li> <li>The name must start with a letter.</li> <li>The name can contain letters, digits, underscores (_), and hyphens (-).</li> <li>The name cannot start with http:// or https://.</li> </ul>
oss_bucket_name	string	The name of your OSS bucket. If you specify a bucket name that does not exist, Packer creates a bucket with the specified name when Packer uploads the image.
image_os_type	string	The type of the image. Valid values:  • linux  • windows
image_platform	string	The distribution of the image. In this example, CentOS is used.
image_architectur e	string	The architecture of the image. Valid values:  • i386  • x86_64
format	string	The format of the image. Valid values:  RAW  VHD

#### What's next

You can use the custom image uploaded to Alibaba Cloud to create Elastic Compute Service (ECS) instances. For more information, see Create an ECS instance by using a custom image.

#### Related information

#### References

- Packer Documentation
- packer
- opstools
- packer-provider
- Anaconda Kickstart

### 6.2.5. Use OOS to create a custom image

Operation Orchestration Service (OOS) provides a public template named ACS-ECS-CreateAndCopylmage to help you perform a set of operations. For example, you can use this template to send remote commands and create and copy custom images.

#### Context

You can view information such as the default process of executing this public template in the OOS console. For more information, see View the information of a public template.

#### **Procedure**

1. Go to the page for configuring a custom image creation task in the OOS console.

i.

ii.

iii.

- iv. In the top navigation bar, select a region.
- v. In the Create Image section, click Use OOS to Create.
- 2. On the Create or Update Images page, configure parameters.
  - i. Keep Create Image from Existing Instance selected.
  - ii. Complete the settings in the General Settings section.Enter a name and description for the custom image and set the Permissions parameter.
  - iii. Complete settings in the Image Distribution section.

Parameter	Description
Copy to Other Region	After a custom image is created in the current region, the image is automatically copied to one or more regions.
Shared Account ID	After a custom image is created, it is automatically shared to one or more Alibaba Cloud accounts.

#### iv. Select an instance.

OOS will create a custom image from the selected instance. When OOS creates a custom image, OOS generates snapshots for the disks attached to the instance. You are charged for the snapshots contained in the created custom image.

#### v. Enter a remote command.

The remote command that you entered is sent to the selected instance before OOS creates the custom image. You can enter a remote command to perform a specific operation on the instance, such as installing software or recording the execution time of O&M tasks.

#### vi. Configure advanced options.

Parameter	Description
Scaling Group Configuration ID	After a custom image is created, it is used to automatically update image parameters in the scaling configuration of the specified scaling group.
Tag	Add tags to the task for easy search.
Save Template	Save the settings of this task as a custom template for future reuse.

#### vii. Click Execute Now.

#### Result

When the task state changes to **Running**, you can click the execution ID to view the execution details such as output parameters and execution logs. When the task state changes to **Success**, the custom image is created and operations such as copying the image to other regions are performed as configured in the task.

### 6.3. Import images

### 6.3.1. Image import procedure

Importing a local image file is applicable to scenarios where your business is deployed on the cloud. You can use the imported custom image to create an ECS instance or replace the system disk. This topic describes how to import an image.

#### Linux image import procedure

The Linux image import procedure is as follows:

- 1. Complete the following preparations on the source server before creating a custom image:
  - i. Install an image compliance tool and use it to automatically check whether Linux settings meet the import requirements. For more information, see Use the image compliance tool.
  - ii. Install cloud-init to allow instances running the image to complete the initialization configuration. For more information, see Install cloud-init.
  - iii. Install VirtIO drivers to allow ECS instances created from the image to start. For more information, see Install the virtio driver.
- 2. Read Instructions for importing images.

- 3. Convert the image file format (except for RAW, VHD, and qcow2 image files). For more information, see Convert the image file format.
- 4. Import custom images.

If the operating system of the custom image is not supported by Alibaba Cloud and cloud-init cannot be installed, you can select Customized Linux when importing a custom image. For more information, see Customize Linux images.

#### Windows Image import procedure

The Windows image import procedure is as follows:

- 1. Install the virtio driver
- 2. Instructions for importing images
- 3. Convert the image file format
- 4. Import custom images

### 6.3.2. Use the image compliance tool

Before you import custom images, we recommend that you use the image compliance tool to check whether the Linux operating system configurations meet the import requirements. The tool checks whether imported custom images are valid images and can be used to create full-featured Elastic Compute Service (ECS) instances. This topic describes how to use the image compliance tool provided by Alibaba Cloud to check the validity of a custom Linux image. In this topic, example operations, parameters, and output items are provided.

#### Context

ECS allows you to create instances from imported custom images. You must import custom images to ECS before you can create instances from them. Custom images can be created based on on-premises servers, virtual machines (VMs), or cloud servers of other service providers. Custom images must meet specific requirements before they can be imported to ECS. For more information, see Instructions for importing images.

We recommend that you use the image compliance tool of ECS to reduce the amount of time required to create a custom image. The image compliance tool is designed to automatically validate configuration items in a Linux server environment to identify non-compliant items, generate detection reports in the TXT and JSON formats, and provide solutions if required.

The image compliance tool only supports Linux images, such as Ubuntu, CentOS, Debian, Red Hat, SUSE Linux Enterprise Server (SLES), openSUSE, FreeBSD, and CoreOS.

In this topic, a server that runs the CentOS 7.4 64-bit operating system is used to show how to use the image compliance tool.

#### **Procedure**

Perform the following operations to use the image compliance tool:

- 1. Log on to your server, VM, or cloud server.
- 2. Run the following command to download the image compliance tool to the working directory of your server, VM, or cloud server:

 $\label{localine} $$ wget $$ http://docs-aliyun.cn-hangzhou.oss.aliyun-inc.com/assets/attach/73848/cn_zh/155745 $$ 9863884/image\_check $$$ 

You can also click here to download the image compliance tool.

3. Run the image compliance tool with root privileges to ensure that it can read configuration files for which an access control list is set.

```
chmod +x image_check
sudo <Path of the image compliance tool>/image_check -p [Destination path]
```

In the preceding code example, <Path of the image compliance tool> is also the path in which the detection report is generated. Therefore, run the following command to start the image compliance tool:

```
sudo ./image_check
```

Note You can use \_\_p [Destination path] to specify the path in which to generate the detection report. If this parameter is not specified, the detection report is generated in the same path as the image compliance tool.

4. Wait for the image compliance tool to check the system configurations.

```
Begin check your system.....
The report is generating.
The infomation you need to input when you import your image to Alibaba Cloud website:
Current system: CentOS
Architecture: x86 64
System disk size: 42 GB
Check driver
                                                         [ OK ]
Check shadow file authority
                                                         [ OK ]
Check security
                                                         [ OK ]
                                                         [ OK ]
Check qemu-ga
Check network
                                                         [ OK
Check ssh
                                                         [ OK ]
Check firewall
                                                         [ OK ]
                                                         [ OK ]
Check filesystem
Check device id
                                                         [ OK
Check root account
                                                         [ OK ]
Check password
                                                         [ OK ]
Check partition table
                                                         [ OK ]
Check lib
                                                         [ OK
Check disk size
                                                         [ OK ]
Check disk use rate
                                                        [ OK ]
Check inode use rate
                                                        [ OK ]
16 items are OK.
0 items are failed.
0 items are warning.
The report is generated: /root/image check report 2019-05-10 13-28-21.txt
Please read the report to check the details.
```

#### 5. View the detection report.

The path of the detection report is displayed in the output of tool execution. In this example, the report is stored in the /root directory. The report is named in the format of  $image\_check\_report\_d$  ate\_time.txt Or  $image\_check\_report.json$ .

#### **Detection items**

The compliance tool detects the following configuration items to ensure that ECS instances created from your custom image are full-featured.

Detection item	Non-compliance issue	Suggestion
driver	The ECS instance cannot start normally.	Install the virtualization driver. For more information, see Install the virtio driver.
/etc/shadow	The password file cannot be modified. As a result, you cannot create an ECS instance from the custom image.	Do not run the chattr command to lock the /etc/shadow file.
SElinux	The ECS instance cannot start normally.	Do not start SELinux by modifying /etc/selinux/config.

Detection item	Non-compliance issue	Suggestion
qemu-ga	Some services required by ECS are unavailable, and the instance is not full-featured.	Uninstall qemu-ga.
network	Network features of the ECS instance are unstable.	Disable or delete Network Manager and enable the network service.
ssh	You cannot connect to the ECS instance in the console.	Enable the SSH service and do not set PermitRootLogin.
firewall	The system does not automatically configure your ECS instance environment.	Disable firewalls such as iptables, firewalld, IPFilter (IPF), IPFireWall (IPFW), or PacketFilter (PF).
file system	You cannot resize the disk.	<ul> <li>We recommend that you use the XFS, ext3, and ext4 file systems.</li> <li>The ext2, UFS, and UDF file systems are supported.</li> <li>Do not use the 64 bit feature for the ext4 file system.</li> <li>Note 64 bit is a feature of the ext4 file system. You can run the man ext4 command to view detailed descriptions.</li> </ul>
device id	The ECS instance cannot start normally.	Clean up the <i>fstab</i> file and remove device IDs that are not needed from the file to ensure that the device IDs in use appear in the output of the <b>blkid</b> command.
root	The username and password cannot be used to connect to the ECS instance.	Reserve the root account.
passwd	You cannot add users to the ECS instance.	Retain the passwd command or reinstall the password file.
Partitioned table	The ECS instance cannot start normally.	Use MBR partitioning.
/lib	The ECS instance cannot be automatically configured.	The /lib and /lib64 files cannot be stored in absolute paths. Modify the storage paths of the /lib and /lib64 files to their relative paths.

Detection item	Non-compliance issue	Suggestion
system disk	N/A	Increase the system disk capacity. The optimal system disk capacity is 40 GiB to 500 GiB. When you import images, configure the system disk capacity based on the virtual file size of images, instead of the size of images.
disk usage	Required drivers or services cannot be installed for the ECS instance.	Make sure that sufficient disk space is available.
inode usage	Required drivers or services cannot be installed for the ECS instance.	Make sure that sufficient inode resources are available.

The image compliance tool provides an OK , FAILED , or WARNING detection result based on detection items.

- OK: The detection items all comply with requirements.
- FAILED: The detection items do not comply with requirements. ECS instances created from the custom image cannot start normally. We recommend that you rectify the non-compliant items and create an image to improve instance startup efficiency.
- WARNING: The detection items do not comply with requirements. ECS instances created from the custom image can start normally, but ECS cannot use valid methods to configure your instance. You can immediately rectify the non-compliant items, or temporarily retain the items and create an image.

#### Output items

The image compliance tool generates detection reports in both the TXT and JSON formats in the destination path.

Note You can use \_\_p [Destination path] to specify the path in which to generate the detection report. If this parameter is not specified, the detection report is generated in the same path as the compliance tool.

• Reports in the TXT format are named image\_check\_report\_date\_time.txt. The reports include server configuration information and detection results. In the following example, a server that runs the CentOS 7.4 64-bit operating system is used:

```
The information you need to input when you import your image to Alibaba Cloud Website:

Current system is: CentOS #Server operating system

Architecture: x86_64 #System architecture

System disk size: 42 GB #Server system disk capacity

------

Check driver #Detection item name

Pass: kvm drive is exist #Detection result

Alibaba Cloud supports kvm virtualization technology

We strongly recommend installing kvm driver.
```

• Reports in the JSON format are named image\_check\_report.json. The reports include server configuration information and detection results. In the following example, a server that runs the Cent OS 7.4 64-bit operating system is used:

```
"platform": "CentOS", \\Server operating system
  "os_big_version": "7", \\Operating system version number (major)
  "os small version": "4", \\Operating system version number (minor)
  "architecture": "x86 64", \\System architecture
  "system disk size": "42", \\Server system disk capacity
  "version": "1.0.2", \\Compliance tool version
  "time": "2018-05-14 19-18-10", \Detection time
  "check_items": [{
     "name": "driver", \\Detection item name
     "result": "OK", \\Detection result
     "error code": "0", \\Error code
     "description": "Pass: kvm driver exists.", \\Description
      "comment": "Alibaba Cloud supports kvm virtualization technology. We strongly recom
mend installing kvm driver."
 } ]
}
```

#### What's next

View Instructions for importing images.

### 6.3.3. Instructions for importing images

To improve image the importing efficiency and availability of imported images, we recommend that you read the following instructions before you import an image. Make sure that the following operating types of images meet the specified requirements.

#### Windows Server images

ltem	Description
System-related item	<ul> <li>Verify the integrity of file systems.</li> <li>Do not modify critical system files.</li> <li>Make sure that the system disk has sufficient free space.</li> <li>Configure the system disk size based on the virtual disk size rather than the image size. The system disk size ranges from 40 GiB to 500 GiB.</li> <li>Disable the firewall.</li> <li>Allow traffic on Remote Desktop Protocol (RDP) port 3389.</li> <li>Make sure that the logon password for the administrator account meets the password complexity requirements. The password must be 8 to 30 characters in length and must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. Special characters include  () ` ~ ! @ # \$ % ^ &amp; *  + =   { } [ ] : ; ' &lt; &gt; , . ? / . The password cannot start with a forward slash (/).</li> <li>If an operating system uses the Unified Extensible Firmware Interface (UEFI) mode, the system disk can have only the boot partision (drive C). If the system disk has other partitions, the instance cannot be started or the disk cannot be resized.</li> </ul>

ltem	Description
ltems not supported	• ISO images. Before you import ISO images to Elastic Compute Service (ECS), you must use tools such as VirtualBox to convert the images to the RAW, VHD, or QCOW2 format on your computer.
	<ul> <li>Installation of the QEMU guest agent (GA) in images. If the QEMU GA is installed, some services that ECS requires may be unavailable.</li> </ul>
	<ul> <li>Installation of Community Edition virtio drivers in Windows Server operating systems. If Community Edition virtio drivers are installed, remove the read-only property of the following files:</li> </ul>
	<ul><li>C:\Windows\System32\drivers\netkvm.sys</li></ul>
	<ul> <li>C:\Windows\System32\drivers\balloon.sys</li> </ul>
	<ul><li>C:\Windows\System32\drivers\vioser.sys</li></ul>
	<ul><li>C:\Windows\System32\drivers\viostor.sys</li></ul>
	<ul> <li>C:\Windows\System32\drivers\pvpanic.sys</li> </ul>
	Multi-partition system disks.
Items supported	<ul> <li>Master boot record (MBR) and GUID Partition Table (GPT) partitions as well as New Technology File System (NTFS) file systems.</li> </ul>
	<ul> <li>Images in the RAW, QCOW2, or VHD format. If the image that you want to import is not in one of the preceding formats, you must convert the image to a supported format before you can import the image. For more information, see Convert the image file format.</li> </ul>
	• Images that run the following operating system versions:
	Microsoft Windows Server 2022
	Microsoft Windows Server 2019
	Microsoft Windows Server 2016
	Microsoft Windows Server 2012 R2
	Microsoft Windows Server 2012
	<ul> <li>Microsoft Windows Server 2008 R2</li> </ul>
	Microsoft Windows Server 2008
	<ul> <li>Windows Server 2003 Service Pack 1 (SP1) or later</li> </ul>

### Linux images

|--|

ltem	Description
System-related item	<ul> <li>Verify the integrity of file systems.</li> <li>Do not modify critical system files such as /sbin, /bin, and /lib*.</li> <li>Do not modify /etc/issue*. Otherwise, the distribution of the operating system cannot be identified by ECS and the instance operating system cannot be created.</li> <li>Do not modify /boot/grub/menu.lst. Otherwise, the ECS instance cannot be started.</li> <li>Do not modify /etc/fstab. Otherwise, abnormal partitions cannot be loaded and the ECS instance cannot be started.</li> <li>Do not set /etc/shadow to read-only. Otherwise, the password file cannot be modified and the instance operating system cannot be created.</li> <li>Do not modify /etc/selinux/config to enable SELinux. Otherwise, the instance operating system cannot be started. If you must enable SELinux, see Enable or disable SELinux.</li> <li>Make sure that the system disk has sufficient free space.</li> <li>Disable the firewall.</li> <li>Allow traffic on SSH port 22.</li> <li>Enable Dynamic Host Configuration Protocol (DHCP).</li> <li>Install Xen or KVM virtualization drivers. For more information, see Install the virtio driver.</li> <li>Install cloud-init to configure the hostname, NTP repositories, and YUM repositories. For more information, see Install cloud-init.</li> <li>Make sure that the logon password for the root account meets the password complexity requirements. The password must be 8 to 30 characters in length and must contain three of the following character types: lowercase letters, uppercase letters, digits, and special characters. Special characters include  () ` ~ ! @ # \$ % ^ &amp; * + =   ( ) [ ] : ; ' &lt; &gt; , . ? / .</li> <li>If an operating system uses the UEFI mode, the system disk can have only the root partition. If the system disk has other partitions, the instance cannot be started or the disk cannot be resized.</li> </ul>
Items not supported	<ul> <li>ISO images. Before you import ISO images to Elastic Compute Service (ECS), you must use tools such as VirtualBox to convert the images to the RAW, VHD, or QCOW2 format on your computer.</li> <li>Multiple network interfaces.</li> <li>IPv6 addresses.</li> <li>Adjustment of system disk partitions. Only disks with a single root partition are supported.</li> <li>Installation of the QEMU GA in images. If the QEMU GA is installed, some services that ECS requires may be unavailable.</li> <li>Creation of system disk partitions (the root partitions) across disks by using Logical Volume Manager (LVM), which makes ECS instances unable to start.</li> </ul>

Item	Description
	<ul> <li>Images in the RAW, QCOW2, or VHD format. If the image that you want to import is not in one of the preceding formats, you must convert the image to a supported format before you can import the image. For more information, see Convert the image file format.</li> <li>MBR and GPT partitions as well as XFS, ext3, and ext4 file systems.</li> <li>However, the ext4 file system cannot contain the 64bit feature, and the project and quota features cannot be used together. You can run the following command to view the list of features contained in an ext4 file system:</li> </ul>
	<ul><li>Images that run the following operating system versions:</li><li>Alibaba Cloud Linux</li></ul>
	∘ Anolis
Operation	o Cent OS 5/6/7/8
	CentOS Stream
	o Debian 6/7/8/9/10
	• FreeBSD
	o OpenSUSE 13/42/15
	o RedHat
	Red Hat Enterprise Linux (RHEL)
	SUSE Linux 10/11/12/15
	Ubuntu 10/12/13/14/16/18/20
	• CoreOS 681.2.0+
	AlmaLinux     Restautions
	Rocky Linux

#### Note:

If the image that you want to import uses the ARM64 architecture or if the image contains the CentOS Stream, Anolis, AlmaLinux, or Rocky Linux operating system, take note of the following items:

- To ensure that the password can be set or the key pair can be modified for an imported image, the image must meet the following requirements:
  - o The kernel of the operating system must support the <code>config\_fw\_cfg\_sysfs</code> feature. By default, the Linux community kernel 4.6 and later and CentOS kernel 3.10.0-826.el7 and later support this feature. You can run the <code>grep -nr config\_fw\_cfg\_sysfs /boot/config\_\$(uname -r)</code> command in the server corresponding to the image. If the command output contains

    <code>config\_fw\_cfg\_sysfs=y</code>, the kernel of this image supports the <code>config\_fw\_cfg\_sysfs</code> feature.
  - The operating system is installed with the latest version of Alibaba Cloud cloud-init. If the version of cloud-init is 19.1, the minor version must be 19.1.3 or later. If the version of cloud-init is 0.7.6a in some early versions of operating systems, the minor version must be 0.7.6a15 or later. For more information, see Install cloud-init.
  - o The operating system must support the SHA-512 encryption algorithm.

- To ensure that the imported image supports resizing of disks and file systems, the image must meet the following requirements:
  - The kernel version of the operating system must be later than 3.6.
  - The image must support the growpart command. To support this command, you must install the cloud-utils-growpart package. The methods of installing the package vary based on the operating systems. For more information, see Resize partitions and file systems of Linux system disks.
  - o The image must support the resize2fs command. To support this command, you must install the e2fsprogs package. By default, the package is installed on the operating system. If the package is not installed in the operating system, you must manually install it. For example, if you want to install the e2fsprogs package on a CentOS 7 operating system, run the yum install -y e2fsprogs command.
  - The operating system is installed with the latest version of Alibaba Cloud cloud-init. If the version of cloud-init is 19.1, the minor version must be 19.1.3 or later. If the version of cloud-init is 0.7.6a in some early versions of operating systems, the minor version must be 0.7.6a15 or later. For more information, see Install cloud-init.

#### Non-standard Linux images

Linux images that are not listed as ECS public images are considered as non-standard platform images. Such images are developed on a standard operating system platform, but their critical system configuration files, basic system environments, or applications do not comply with the standard operating system requirements of ECS. If you want to use a non-standard platform image, you must complete required configurations within the operating system contained in the image before and after the image is imported. The following section describes the non-standard platform image types that can be imported and how to configure these images:

- Others Linux: ECS identifies all images of this type as other Linux images. If you import an Others Linux image and then use it to create instances, ECS does not process the created instances. If DHCP is disabled in the Others Linux image, you must connect to the instances by using the remote connection feature in the ECS console and then manually configure IP addresses, routes, and passwords for the instances. If DHCP is enabled in the Others Linux image, networks are automatically configured for the instances when you log on to the instances.
- Customized Linux: Before a customized Linux image is imported, you can create instances from this image and configure networks and passwords for the instances in the same way as you do with instances created from standard ECS images. For more information, see Customize Linux images.

### 6.3.4. Install cloud-init

We recommend that you install cloud-init on a source server when you create a Linux custom image based on the server. This ensures that Elastic Compute Service (ECS) instances which run the custom image can initialize their system configurations. This topic describes how to install Alibaba Cloud cloud-init and the native cloud-init.

#### Context

cloud-init is open source software used by cloud platforms to initialize the system for Linux virtual machines. All major public cloud platforms such as Alibaba Cloud, Amazon Web Services (AWS), Microsoft Azure, and OpenStack support cloud-init. Alibaba Cloud cloud-init initializes system configurations on instance startup and executes user data scripts. The configurations include the Network Time Protocol (NTP), software repositories, hostnames, and SSH key pairs. For more information, see cloud-init Documentation.

By default, cloud-init is installed for all Alibaba Cloud public images. We recommend that you install Alibaba Cloud cloud-init on your Linux server based on which to create custom images in the following scenarios, so that instances created from the custom images can automatically initialize system configurations:

- Linux servers that you plan to migrate to the cloud are not installed with cloud-init.
  - **Note** Proceed with caution when you install Alibaba Cloud cloud-init on servers that you do not plan to migrate to the cloud.
- Linux servers are installed with cloud-init of versions earlier than 0.7.9.
- Alibaba Cloud ECS instances are not installed with cloud-init.

#### cloud-init versions

Different cloud platforms may support different versions of cloud-init. Select an appropriate version and configure an appropriate data source (datasource) for it. The latest version of Alibaba Cloud cloud-init is 19.1.10, and the data source is Aliyun. You can download other versions of cloud-init from the official cloud-init website.

Notice After cloud-init is installed, it automatically runs on system startup. If the installed version of cloud-init is not compatible with the operating system of the server or if the data source is not properly configured, cloud-init may not run normally and the instance may start slowly or may not start at all the next time you restart the instance. Therefore, you must select a later version of cloud-init and an appropriate data source such as Aliyum.

When you use cloud-init, take note of the differences among the following versions:

- 0.7.6a: the initial version of Alibaba Cloud cloud-init, which depends on Python 2.7. Some public images based on early versions of operating systems still support cloud-init 0.7.6a. If you want to install cloud-init 0.7.6a for your images, see the (Optional) Install Alibaba Cloud cloud-init 0.7.6a15 section.
  - Note The Python community no longer provides technical support for Python 2.7. To avoid potential risks associated with dependency libraries, we recommend that you use recent versions of cloud-init.
- 0.7.9 and earlier: initial versions of the native cloud-init, which are not applicable to initializing ECS instances and must be upgraded.
- 18: cloud-init 18 and later automatically initialize network configurations. The code for network configuration is BOOTPROTO=dhcp DEVICE=eth0 ONBOOT=yes STARTMODE=auto TYPE=Ethernet USERCTL=n

  o . If you want to customize network configurations after you install cloud-init, see the (Optional)

  Customize network configurations section.
- 19.1: Alibaba Cloud public images are being upgraded to have cloud-init 19.1 pre-installed, which

depends on Python 3.6.

#### Check the version of cloud-init

- 1. Log on to the source server.
- 2. Run the following command to check whether cloud-init is installed:

```
which cloud-init
```

If cloud-init is not installed, the output contains no path information. Install Alibaba Cloud cloud-init.

3. Run the following command to check the version of cloud-init:

```
cloud-init --version
```

If the returned version is earlier than 0.7.9, install Alibaba Cloud cloud-init 19.1.10.

4. Back up data on the server.

### (Recommended) Install Alibaba Cloud cloud-init 19.1.10

Perform the following operations to download cloud-init 19.1.10, whose data source is Aliyun:

1. Make sure that the python-pip dependency library is installed on the source server.

In the examples, python3-pip is used. Run the following commands to install the python-pip dependency library for specific Linux distributions:

• CentOS and Red Hat Enterprise Linux:

```
yum -y install python3-pip
```

Ubuntu and Debian:

```
apt-get -y install python3-pip
```

o openSUSE and SUSE:

```
zypper -n install python3-pip
```

2. Run the following command to download the Alibaba Cloud cloud-init installation package:

```
wget https://ecs-image-tools.oss-cn-hangzhou.aliyuncs.com/cloudinit/cloud-init-19.1.10.
tgz
```

3. Run the following command to decompress the cloud-init installation package to the current directory:

```
tar -zxvf cloud-init-19.1.10.tgz
```

4. Run the following commands to go to the *cloud-init* directory and install the dependency library:

```
cd ./cloud-init-19.1.10
pip3 install -r ./requirements.txt
```

5. Run the following command to go to the *tools* directory of the *cloud-init* file:

```
cd ./tools
```

6. Run the following command to execute the *deploy.sh* script to install cloud-init:

```
bash ./deploy.sh <issue> <major_version>
```

The following table describes the parameters and their values in the deploy.sh script.

Parameter	Description	Example
<issue></issue>	The operating system type. Valid values: <i>centos</i> , <i>redhat, rhel, debian, ubuntu, opensuse</i> , and <i>sles</i> . The parameter values are case-sensitive. <i>sles</i> stands for SUSE or SUSE Linux Enterprise Server.	centos
	The major version number of the operating system.	
<major_versio n&gt;</major_versio 	Note Alibaba Cloud cloud-init 19.1.10 cannot be installed on Ubuntu 14.	The major version number of CentOS 7.6 is <i>7</i> .

For example, if the current operating system is CentOS 7, you must run the bash ./deploy.sh cent os 7 command.

7. Check whether cloud-init is installed.

If Alibaba Cloud cloud-init is installed, "description": "success" is returned.

```
{
"status_code": 0,
"description": "success"
}
```

The following section provides sample shell scripts that are used to install Alibaba Cloud cloud-init for different Linux distributions. Adapt the script to your operating system.

Cent OS 7/8

```
# Check whether the python3-pip dependency library is installed. If not, install it.
if ! python3 -c 'import setuptools' >& /dev/null; then
   yum -y install python3-pip
fi
# Back up cloud-init of an earlier version.
test -d /etc/cloud && mv /etc/cloud /etc/cloud-old
# Download and decompress the Alibaba Cloud cloud-init installation package.
wget https://ecs-image-tools.oss-cn-hangzhou.aliyuncs.com/cloudinit/cloud-init-19.1.10.tg
z
tar -zxvf ./cloud-init-19.1.10.tgz
# Install cloud-init.
issue_major=$( cat /etc/redhat-release | grep -Eo '[0-9]+\.?[0-9]+' | head -1 | awk -F'.'
'{printf $1}')
bash ./cloud-init-*/tools/deploy.sh centos "$issue_major"
```

• Red Hat Enterprise Linux 7/8

```
# Check whether the python3-pip dependency library is installed. If not, install it.
if ! python3 -c 'import setuptools' >& /dev/null; then
   yum -y install python3-pip
fi
# Back up cloud-init of an earlier version.
test -d /etc/cloud && mv /etc/cloud /etc/cloud-old
# Download and decompress the Alibaba Cloud cloud-init installation package.
wget https://ecs-image-tools.oss-cn-hangzhou.aliyuncs.com/cloudinit/cloud-init-19.1.10.tg
z
tar -zxvf ./cloud-init-19.1.10.tgz
# Install cloud-init.
issue_major=$( cat /etc/os-release | grep VERSION_ID | grep -Eo '[0-9]+\.?[0-9]+' | head
-1 | awk -F'.' '{printf $1}')
bash ./cloud-init-*/tools/deploy.sh rhel "$issue_major"
```

#### Ubuntu 16/18/20

```
# Check whether the python3-pip dependency library is installed. If not, install it.
if ! python3 -c 'import setuptools' >& /dev/null; then
    apt-get install python36 python3-pip -y
fi
# Back up cloud-init of an earlier version.
test -d /etc/cloud && mv /etc/cloud /etc/cloud-old
# Download and decompress the Alibaba Cloud cloud-init installation package.
wget https://ecs-image-tools.oss-cn-hangzhou.aliyuncs.com/cloudinit/cloud-init-19.1.10.tg
z
tar -zxvf ./cloud-init-19.1.10.tgz
# Install cloud-init.
issue_major=$( cat /etc/os-release | grep VERSION_ID | grep -Eo '[0-9]+\.?[0-9]+' | head
-1 | awk -F'.' '{printf $1}')
bash ./cloud-init-*/tools/deploy.sh ubuntu "$issue_major"
```

#### • Debian 9/10

```
# Check whether the python3-pip dependency library is installed. If not, install it.
if ! python3 -c 'import setuptools' >& /dev/null; then
    apt-get -y install python3-pip
fi
# Back up cloud-init of an earlier version.
test -d /etc/cloud && mv /etc/cloud /etc/cloud-old
# Download and decompress the Alibaba Cloud cloud-init installation package.
wget https://ecs-image-tools.oss-cn-hangzhou.aliyuncs.com/cloudinit/cloud-init-19.1.10.tg
z
tar -zxvf ./cloud-init-19.1.10.tgz
# Install cloud-init.
issue_major=$( cat /etc/os-release | grep VERSION_ID | grep -Eo '[0-9]+\.?[0-9]+' | head
-1 | awk -F'.' '{printf $1}')
bash ./cloud-init-*/tools/deploy.sh debian "$issue_major"
```

#### • SUSE 12/15

```
# Check whether the python3-pip dependency library is installed. If not, install it.
if ! python3 -c 'import setuptools'>& /dev/null; then
    zypper -n install python3-pip
fi
# Back up cloud-init of an earlier version.
test -d /etc/cloud && mv /etc/cloud/etc/cloud-old
# Download and decompress the Alibaba Cloud cloud-init installation package.
wget https://ecs-image-tools.oss-cn-hangzhou.aliyuncs.com/cloudinit/cloud-init-19.1.10.tg
z
tar -zxvf ./cloud-init-19.1.10.tgz
# Install cloud-init.
issue_major=$( cat /etc/os-release | grep VERSION_ID | grep -Eo '[0-9]+\.?[0-9]+' | head
-1 | awk -F'.' '{printf $1}')
bash ./cloud-init-*/tools/deploy.sh sles "$issue_major"
```

#### OpenSUSE 15

```
# Check whether the python3-pip dependency library is installed. If not, install it.
if ! python3 -c 'import setuptools'>& /dev/null; then
    zypper -n install python3-pip
fi
# Back up cloud-init of an earlier version.
test -d /etc/cloud && mv /etc/cloud/etc/cloud-old
# Download and decompress the Alibaba Cloud cloud-init installation package.
wget https://ecs-image-tools.oss-cn-hangzhou.aliyuncs.com/cloudinit/cloud-init-19.1.10.tg
z
tar -zxvf ./cloud-init-19.1.10.tgz
# Install cloud-init.
issue_major=$( cat /etc/os-release | grep VERSION_ID | grep -Eo '[0-9]+\.?[0-9]+' | head
-1 | awk -F'.' '{printf $1}')
bash ./cloud-init-*/tools/deploy.sh opensuse"$issue_major"
```

# (Optional) Install Alibaba Cloud cloud-init 0.7.6a15

Some early operating systems such as CentOS 6, Debian 9, and SUSE Linux Enterprise Server 12 still support cloud-init 0.7.6a15.

Note By default, Alibaba Cloud public images of CentOS 6, Debian 9, and SUSE Linux Enterprise Server 12 are installed with <a href="cloud-init-0.7.6a15">cloud-init-0.7.6a15</a>. If you want to test installed cloud-init-0.7.6a15, first run the <a href="mailto:mv/etc/cloud/cloud.cfg/etc/cloud/cloud.cfg">mv/etc/cloud/cloud.cfg</a> / etc/cloud/cloud.cfg \_bak command to back up configuration files.

1. Run the following command to check whether the operating system version is CentOS 6, Debian 9, or SUSE Linux Enterprise Server 12:

```
cat /etc/issue
```

2. Run the following commands to download and decompress the Alibaba Cloud cloud-init 0.7.6a15 installation package:

```
wget https://ecs-image-tools.oss-cn-hangzhou.aliyuncs.com/cloud-init-0.7.6a15.tgz tar -zxvf cloud-init-0.7.6a15.tgz
```

3. Run the following command to go to the *tools* directory of the *cloud-init* file:

```
cd cloud-init-0.7.6a15/tools/
```

4. Run the following command to install cloud-init:

```
bash ./deploy.sh <issue> <major_version>
```

The following table describes the parameters and their values in the *deploy.sh* script.

Parameter	Description	Example
<issue></issue>	The operating system type. Valid values: <i>centos</i> , <i>debian</i> , and <i>sles</i> . The parameter values are casesensitive. <i>sles</i> stands for SUSE or SUSE Linux Enterprise Server.	centos
<major_versio< td=""><td>The major version number of the operating system.</td><td>The major version number of CentOS 6.5 is <i>6</i>.</td></major_versio<>	The major version number of the operating system.	The major version number of CentOS 6.5 is <i>6</i> .

For example, if the current operating system is CentOS 6, you must run the bash ./deploy.sh cent os 6 command.

## (Optional) Install the native cloud-init

1. Make sure that the Git, Python, and python-pip dependency libraries are installed on the source server.

In the examples, Git, Python 3.6, and python3-pip are used. Run the following commands to install Git, Python, and python-pip dependency libraries for specific Linux distributions:

o Cent OS and Red Hat Enterprise Linux:

```
yum -y install git python36 python3-pip
```

Ubunt u and Debian:

```
apt-get -y install git python36 python3-pip
```

openSUSE and SUSE:

```
zypper -n install git python36 python3-pip
```

2. Run the following command to download the cloud-init source code package from Git:

```
git clone https://git.launchpad.net/cloud-init
```

3. Run the following command to go to the cloud-init directory:

```
cd ./cloud-init
```

4. Run the following command to install all the dependency libraries:

```
pip3 install -r ./requirements.txt
```

5. Run the following command to install cloud-init:

```
python3 setup.py install
```

6. Modify the *cloud.cfg* configuration file.

i. Run the following command to open the configuration file:

```
vi /etc/cloud/cloud.cfg
```

```
# The top level settings are used as module
# and system configuration.
# A set of users which may be applied and/or used by various modules
# when a 'default' entry is found it will reference the 'default_user'
# from the distro configuration specified below
users:
- default
# If this is set, 'root' will not be able to ssh in and they
# will get a message to login instead as the default $user
disable_root: true
# This will cause the set-update hostname module to not operate (if true)
preserve_hostname: false
# Example datasource config
# datasource:
# Ec2:
# metadata_urls: [ 'blah.com' ]
# timeout: 5 # (defaults to 50 seconds)
# max_wait: 10 # (defaults to 120 seconds)
# The modules that run in the 'init' stage
cloud_init_modules:
```

ii. Change the content that precedes cloud init modules: to the following content:

```
# Example datasource config
\# The top level settings are used as module
# and system configuration.
# A set of users which may be applied and/or used by various modules
# when a 'default' entry is found it will reference the 'default user'
# from the distro configuration specified below
  - default
user:
   name: root
   lock passwd: False
# If this is set, 'root' will not be able to ssh in and they
# will get a message to login instead as the above $user
disable root: false
# This will cause the set+update hostname module to not operate (if true)
preserve_hostname: false
syslog fix perms: root:root
datasource_list: [ AliYun ]
# Example datasource config
datasource:
   AliYun:
       support_xen: false
       timeout: 5 # (defaults to 50 seconds)
       max wait: 60 # (defaults to 120 seconds)
      metadata urls: [ 'blah.com' ]
# The modules that run in the 'init' stage
cloud init modules:
```

# (Optional) Customize network configurations

1. After cloud-init is installed, open the /etc/cloud/cloud.cfg file.

```
vim /etc/cloud/cloud.cfg
```

2. Press the /key to enter the edit mode. Add the following disabled configuration before Example

```
datasource config :
 network:
   config: disabled
```

(?) Note After the configuration is added, you must manage the network configurations in the /etc/sysconfig/network-scripts/ directory on your own.

```
# and system configuration.
# A set of users which may be applied and/or used by various modules
# when a 'default' entry is found it will reference the 'default user'
# from the distro configuration specified below
users:
   - default
user:
   name: root
    lock passwd: False
# will get a message to login instead as the default $user
disable root: false
# This will cause the set+update hostname module to not operate (if true)
preserve_hostname: false
manage_etc_hosts: localhost
network:
  config: disabled
datasource_list: [ AliYun ]
# Example datasource config
```

3. Press the Esc key to exit the edit mode. Enter  $:_{\text{Wq}}$  and press the Enter key to save and close the file.

# (Optional) Keep the hostname and the /etc/hosts configuration file unchanged

By default, the host name in the cloud-init configuration file is not retained. If you do not want the host name and /etc/hosts configuration file of your instance to be modified, you can perform the following operations to modify the cloud-init configuration file.

1. After cloud-init is installed, open the /etc/cloud/cloud.cfg file.

```
vim /etc/cloud/cloud.cfg
2. Press the /key to enter the edit mode and change preserve hostname: false to preserve host
```

name: true .

```
# This will cause the set+update hostname module to not operate (if true)
preserve_hostname: true
manage_etc_hosts: localhost
```

3. Press the *Esc* key to exit the edit mode. Enter :wq and press the Enter key to save and close the file.

preserve\_hostname: true prevents cloud-init from modifying the instance hostname and the /etc/hosts configuration file.

### **Troubleshooting**

- After I use Python 3 to install cloud-init, the required dependency libraries are absent. What do I do? The dependency libraries that are absent may vary based on images. You can use pip to install the libraries and then install cloud-init again.
- cloud-init is running abnormally and an error message was returned. What do I do?

If the default software package manager such as **YUM** and the **pip** manager have been installed with different versions of dependency libraries, library version conflicts may occur and cause cloud-init to run abnormally. We recommend that you download dependency libraries based on the error message.

Error message	Cause	Troubleshooting command
no setuptools module in python	The python setuptools module is not installed.	Sample commands used to install setuptools for Python 3.6 on different operating systems:  • CentOS and Red Hat: yum -y instal 1 python3-pip  • Ubuntu and Debian: apt-get -y ins tall python3-pip  • openSUSE and SUSE: zypper -n ins tall python3-pip
<pre>File "/root/cloud- init/cloudinit/log.py", line 19, in <module>         import six    ImportError: No module named six )</module></pre>	The six dependency library is not installed.	pip3 install six
<pre>File "/root/cloud- init/cloudinit/url_helper.py", line 20, in <module>     import oauthlib.oauthl as oauthl     ImportError: No module named oauthlib.oauthl )</module></pre>	The oauthlib dependency library not installed.	pip3 install oauthlib

Error message	Cause	Troubleshooting command
No uninstalled dependency libraries specified	The error message is not mapped.	Run the following command to install all dependency libraries that are displayed in the <i>requirements.txt</i> file of cloud-init:
specified		pip3 install -r requirements.txt

• How do I make sure that a correct version of cloud init is installed?

After you install cloud-init, you can run the **cloud-init** --version command to check the version of cloud-init. If the returned version is the one that you want to install, the correct version of cloud-init is installed. For example, you want to install Alibaba Cloud cloud-init 19.1.10. After you install cloud-init and run the **cloud-init** --version command, the command output contains 19.1.10 , which indicates that cloud-init 19.1.10 is installed.

• cloud-init runs abnormally on an instance after I install a new version of Python 3 on the instance. What do I do?

If you install a new version of Python 3 such as Python 3.9 on your instance and create a symbolic link (such as ln -s /usr/bin/python3.9 /usr/bin/python3) for the Python 3 version to set it as the default version, cloud-init installed on the instance runs abnormally. For example, when you run the cloud-init --version command, an error is reported.

```
$cloud-init --version
Traceback (most recent call last):
   File "/usr/local/bin/cloud-init", line 33, in <module>
        sys.exit(load_entry_point('cloud-init==19.1.9', 'console_scripts', 'cloud-init')())
   File "/usr/local/bin/cloud-init", line 22, in importlib_load_entry_point
        for entry_point in distribution(dist_name).entry_points
   File "/usr/lib64/python3.9/importlib/metadata.py", line 524, in distribution
        return Distribution.from_name(distribution_name)
   File "/usr/lib64/python3.9/importlib/metadata.py", line 187, in from_name
        raise PackageNotFoundError(name)
importlib.metadata.PackageNotFoundError: cloud-init
```

You can use one of the following methods to resolve this issue:

o Method 1: Use the new version of Python 3 to re-install cloud-init.

- Method 2: Modify the cloud-init executable to change the Python interpreter path to the path to an earlier Python 3 version. Python 3.6 is used in the following example. Perform the following operation to change the Python interpreter path in the cloud-init executable:
  - a. Run the following command to open the cloud-init executable:

```
vim /usr/local/bin/cloud-init
```

b. Press the /key to enter the edit mode. At the beginning of the executable, change the content that follows #! to the path to Python 3.6.

Example of the new content on the line that starts with #!:

```
#!/usr/bin/python3.6
```

- c. Press the Esc key to exit the edit mode. Enter :wq and press the Enter key to save and close the file.
- After I install cloud-init on an instance, the user data that was specified in instance metadata when I created the instance does not or cannot run. What do I do?
  - i. Run the following command on the instance to check whether user data exists in instance metadata:

```
curl http://100.100.100.200/latest/user-data
```

#### Check results:

- If user data exists in instance metadata, the output contains the user data. You must perform other operations to further troubleshoot the problem.
- If user data does not exist in instance metadata, no output is returned.
- ii. Identify the reason why the user data does not or cannot run.
  - Check whether the user data is in a valid format.
    - User data must be in a valid format to be run by cloud-init. For example, the first line of the user data must start with #! For more information, see Manage the user data of Linux instances.
  - Check whether you selected **Security Hardening Mode** when you created the instance.
    - If you selected **Security Hardening Mode**, the components of cloud-init cannot be initialized, which affects instance configurations such as metadata and user data. For more information, see <u>View instance metadata</u>.
  - Check the execution result of the user data in the /var/log/cloud-init.log file of cloud-init and troubleshoot the problem based the returned error message.

Example of the user data execution result in the log file:

```
util.py[DEBUG]: Running command ['/var/lib/cloud/instance/scripts/part-001'] with
allowed return codes [0] (shell=False, capture=False)
```

• Check the standard output (stdout) and standard error (stderr) generated while the user data was running.

In systemd (a system initialization tool of Linux), user data is run by the cloud-final.service system service. You can run the following command to check the stdout and stderr generated while the user data was running. Then, you can troubleshoot the problem based on the command output.

```
journalctl -u cloud-final.service
```

If you still cannot identify the reason why the user data does not or cannot run after you perform the preceding operations, we recommend that you copy the user data to your computer and check whether the user data can be run on the computer.

#### What's next

• Linux servers that you plan to migrate to the cloud:

You can migrate the servers by using the Cloud Migration tool or by importing the custom images. For more information, see Migrate a server to Alibaba Cloud by using the Cloud Migration tool or Import custom images.

• Alibaba Cloud ECS instances that already run Linux custom images:

You can restart the system to check whether cloud-init is installed. If cloud-init is installed, the system configures the hostname, software repositories, and NTP. For example, if cloud-init is installed, the network configuration file shows the following information:

```
[root@iZbplios3psx4hoi******Z ~]# cat /etc/sysconfig/network-scripts/ifcfg-eth0
# Created by cloud-init on instance boot automatically, do not edit.
#
BOOTPROTO=dhcp
DEVICE=eth0
ONBOOT=yes
STARTMODE=auto
TYPE=Ethernet
USERCTL=no
```

#### Related information

• Official cloud-init website - Alibaba Cloud (AliYun)

# 6.3.5. Install the virtio driver

In some scenarios, you may want to create a custom image based on the operating system data of a source server, import the image to Alibaba Cloud, and then create an Elastic Compute Service (ECS) instance from the image. This source server can be a physical server, a virtual machine, or a cloud host. If the source server is not installed with the virtio driver, the created ECS instance may be unable to start. To prevent this issue, you must check whether the virtio driver is installed on the source server before you import a custom image to Alibaba Cloud.

#### Context

You can determine whether you need to manually install the virtio driver based on the operating system of your source server.

Operating system	Description
<ul> <li>Windows Server 2008</li> <li>Windows Server 2012</li> <li>Windows Server 2016</li> <li>Windows Server Version **** (Semi-Annual Channel)</li> <li>Windows Server 2019 and later</li> <li>CentOS 6, CentOS 7, CentOS 8, and later</li> <li>Ubuntu 12, Ubuntu 14, Ubuntu 16, Ubuntu 18, Ubuntu 20, and later</li> <li>Debian 7, Debian 8, Debian 9, Debian 10, and later</li> <li>SUSE 11, SUSE 12, SUSE 15, and later</li> </ul>	When you import a custom image created from a source server that runs one of the operating systems listed on the left, Alibaba Cloud automatically adds the virtio driver to the image. You do not need to manually install the virtio driver on the source server before you create the custom image.  For servers that are pre-installed with the virtio driver, if the driver is not contained in the temporary file systems, you must repair the file systems. For more information, see the Step 2: Repair the temporary file system section in this topic.
Other operating systems	If your source server runs a Linux-like operating system that is not contained in the preceding list, perform the following steps to install the virtio driver:  1. Step 1: Check whether the operating system kernel supports the virtio driver  2. Step 2: Repair the temporary file system  3. Step 3: Download the kernel installation package  4. Step 4: Compile the kernel

# Step 1: Check whether the operating system kernel supports the virtio driver

1. Run the following command to check whether the kernel of the current operating system supports the virtio driver:

grep -i virtio /boot/config-\$(uname -r)

```
[root@inplicementality ~]# grep -i virtio /boot/config-$(uname -r)
CONFIG_VIRTIO_VSOCKETS=m
CONFIG_VIRTIO_VSOCKETS_COMMON=m
CONFIG_VIRTIO_BLK=m
CONFIG_SCSI_VIRTIO=m
CONFIG_VIRTIO_NET=m
CONFIG_VIRTIO_CONSOLE=m
CONFIG_VIRTIO_CONSOLE=m
CONFIG_DRM_VIRTIO_SM
CONFIG_DRM_VIRTIO_SM
CONFIG_VIRTIO=m
# Virtio_drivers
CONFIG_VIRTIO_PCI=m
CONFIG_VIRTIO_PCI=m
CONFIG_VIRTIO_PCI_HEGACY=y
CONFIG_VIRTIO_BALLOON=m
CONFIG_VIRTIO_INPUT=m
# CONFIG_VIRTIO_INPUT=m
# CONFIG_VIRTIO_INPUT=m
# CONFIG_VIRTIO_INPUT=m
```

Check the CONFIG VIRTIO BLK and CONFIG VIRTIO NET parameters in the command output.

Check result	Description
The two parameters do not exist.	The virtio driver is not installed on the operating system and custom images created from the source server cannot be imported to Alibaba Cloud. You must compile the kernel and install the virtio driver on the source server. For more information, see the Step 3: Download the kernel installation package and Step 4: Compile the kernel sections in this topic.
The parameter values are both m.	Proceed to the next step to check whether the virtio driver is contained in the initramfs or initrd temporary file system.
The parameter values are both y.	The virtio driver is installed on the operating system and custom images created from the source server can be imported to Alibaba Cloud. For more information, see Instructions for importing images and Import custom images.

2. Run the following command to check whether the virtio driver is contained in the initramfs or initrd temporary file system:

#### Description of the command output:

- The preceding figure shows that initramfs contains the virtio\_blk driver and the virtio.ko, virtio\_pci.ko, and virtio\_ring.ko dependencies. You can import your custom images to Alibaba Cloud. For more information, see Instructions for importing images and Import custom images.
- If initramfs does not contain the virtio driver, you must repair the temporary file system before
  you can import custom images to Alibaba Cloud. For more information, see the Step 2: Repair the
  temporary file system section in this topic.

# Step 2: Repair the temporary file system

If the check result shows that the kernel of the operating system supports the virtio driver but the initramfs or initrd temporary file system does not contain the virtio driver, you must repair the temporary file system. Run one set of the following commands based on the operating system of your source server to repair the temporary file system:

Cent OS 8 or Red Hat 8

Cent OS 6, Cent OS 7, Red Hat 6, or Red Hat 7

Cent OS 5 or RedHat 5

Debian or Ubunt u

```
echo -e 'xen-blkfront\nvirtio_blk\nvirtio_pci\nvirtio_console' >> \
/etc/initramfs-tools/modules
mkinitramfs -o /boot/initrd.img-$(uname -r)
```

# Step 3: Download the kernel installation package

**Note** In this example, the *linux-4.4.24.tar.gz* kernel installation package is used. You must modify the commands based on the kernel version of your operating system.

1. Run the following command to install components required for compiling the kernel:

```
yum install -y ncurses-devel gcc make wget
```

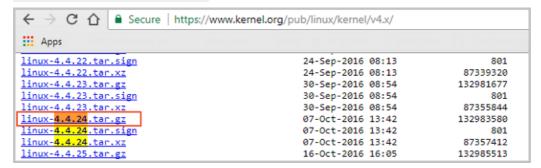
2. Run the following command to query the kernel version of your operating system. In this example, the kernel version is 4.4.24-2.a17.x86\_64.

```
uname -r
[root@iZbp1127hr3wi6p2cq9lnbZ ~]# uname -r
4.4.24-2.a17.x86_64
```

3. Go to the Index of pub/lunix/kernel/ page to query the download URL of the kernel version

repository.

In this example, the download URL of linux-4.4.24.tar.gz is https://www.kernel.org/pub/linux/kernel/v4.x/linux-4.4.24.tar.gz , as shown in the following figure.



4. Run the following command to switch the directory:

```
cd /usr/src/
```

5. Run the following command to download the installation package:

```
wget https://www.kernel.org/pub/linux/kernel/v4.x/linux-4.4.24.tar.gz
```

6. Run the following command to decompress the installation package:

```
tar -xzf linux-4.4.24.tar.gz
```

7. Run the following to create a symbolic link:

```
ln -s linux-4.4.24 linux
```

8. Run the following command to switch the directory:

```
cd /usr/src/linux
```

# Step 4: Compile the kernel

1. Run the following commands in sequence to compile the kernel:

```
make mrproper
symvers_path=$(find /usr/src/ -name "Module.symvers")
test -f $symvers_path && cp $symvers_path .
cp /boot/config-$(uname -r) ./.config
make menuconfig
```

Note If an error message is displayed when the make menuconfig command is run, you can troubleshoot the error based on the message. For example, if the ncurses-devel, bison, and flex components are not installed, you can run the following commands in sequence to install them:

```
yum install ncurses-devel
yum install bison
yum install flex
```

2. When the following page is displayed, start to make virtio-related configurations.

- **Note** If you select the configurations that contain asterisks (\*), the virtio driver is compiled into the kernel. If you select configurations that contain *m*, the virtio driver is compiled into a module and then the module is inserted into the kernel when the driver starts.
- i. Configure the Virtualization options.
  - a. Press the spacebar key to select **Virtualization** and press the *Enter* key to view the list of options.

b. Checkthat the Kernel-based Virtual Machine (KVM) support option is selected.

- ii. Configure the Processor type and features options.
  - a. Go back to the main menu, press the up and down arrow keys to select **Processor type** and features, and then press the *Enter* key to view the list of options.
  - b. Press the spacebar key to select **Paravirtualized guest support** and press the *Enter* key to view the list of options.

```
Arrow keys navigate the menu. <Enter> selects submenus --->. Highlighted letters are hotkeys. Pressing <Y> includes, <N> excludes, <N> modularizes features. Press <Esc><Esc> to exit, <?> for Help, </> for Search. Legend: [*] built-in [ ] excluded <M> module <> module capable

[*] Lickless System (Dynamic Ticks) in has unmet direct dependencies (EXPERIMENTAL && INET && SYSFS && CONFIGFS [*] High Resolution Timer Support
[*] symmetric multi-processing support
[*] support for extended (non-PC) X86 platforms
[] scaleMP vSMP
[4] single depth WCHM output
[M] Raravitriops debugging
[] maravitriops debugging
[] mitest
[] processor family (Generic-x86-64) --->
[8] Miximum number of CPUs (NEW)
```

c. Checkthat the **KVM paravirtualized clock** and **KVM Guest support** options are selected.

```
Arrow keys navigate the menu. <Enter> selects submenus ...>. Highlighted letters are hotkeys. Pressing <Y> includes, <N> excludes, <N> modularizes features. Press <Esc><Esc> to exit, <?> for Help, </> for Search. Legend: [*] built-in [] excluded <N> module <> module capable

--- Paravirtualized guest support
[*] Faravirtual steal time accounting
[*] Xen guest support
[*] inable Xen debug and tuning parameters in debugfs
[*] inable Xen debug checks. (NEW)

[*] ** WM paravirtualized clock

[*] ** WM paravirtualized clock
```

- iii. Configure the Device Drivers options.
  - a. Go back to the main menu, press the up and down arrow keys to select **Device Drivers**, and then press the *Enter* key to view the list of options.
  - b. Press the spacebar key to select **Block devices** and press the *Enter* key to view the list of options.
  - c. Checkthat the Virtio block driver option is selected.

- d. Go back to the list of **Device Drivers** options, press the up and down arrow keys to select **Network device support**, and then press the *Enter* key to view the list of options.
- e. Check that the Virtio network driver option is selected.

```
Arrow keys navigate the menu. <Enter> selects submenus --- (or empty submenus --- ). Highlighted letters are hotkeys excludes, <M> modularizes features. Press <Esc> to exit, <7> for Help, </> for Search. Legend: [*] built-in [ ] capable

--- Network device support

[*] Network core driver support

--- Bonding driver support

--- Bonding driver support

--- CPL (serial line load balancing) support

--- EQL (serial line load balancing) support

--- Intermediate Functional Block support

--- Henred team driver support

--- MM-C-VLAN support

--- MM-C-VLAN support

--- MM-C-VLAN support

--- MM-C-VLAN support (NEW)

--- M
```

- iv. Press the *Esc* key to exit the kernel configuration window and select **Yes** at the prompt to save the *.config* file.
- v. Check whether virtio-related configurations are complete.

For more information, see the Step 1: Check whether the operating system kernel supports the virtio driver section in this topic.

vi. (Optional)If virtio-related configurations are not complete, run the following commands in sequence to automatically edit the *.config* file:

```
make oldconfig
make prepare
make scripts
make
make install
```

vii. Run the following commands to check whether the virtio driver is installed:

```
find /lib/modules/"$(uname -r)"/ -name "virtio*" | grep -E "virtio*"
grep -E "virtio*" < /lib/modules/"$(uname -r)"/modules.builtin</pre>
```

If one of the command outputs contains virtio-related files such as virtio\_blk, virtio\_pci, and virtio\_console, the virtio driver is installed.

#### What's next

After the virtio driver is installed, you can perform the following operations:

- Migrate servers
- Import custom images

# 6.3.6. Customize Linux images

If the selected OS is not supported by Alibaba Cloud and cloud-init cannot be installed, you can select Customized Linux when you import a custom image. Alibaba Cloud regards customized Linux images as an unrecognized OS type. You must add a parsing script to the custom image before the import to automatically configure the instance when the instance is first started.

#### Limits

Customized Linux images have the following limits:

- The first partition must be writable.
- The type of the first partition must be FAT32, ext2, ext3, ext4, or UFS.
- The size of the virtual file of the customized Linux image must be larger than 5 GiB.

Customized Linux images have the following security requirements:

- No important vulnerabilities can be remotely exploited.
- When you log on to an instance for the first time by using a VNC management terminal of the ECS console, you must change the initial password (if there is one) before you can perform any other actions. For more information, see Connect to a Linux instance by using a password.
- Default SSH key pairs are not supported. The initial SSH key pair must be generated by Alibaba Cloud.

## Configuration method

Before you create and import a customized Linux image, you must perform the following operations:

1. Create the *aliyun\_custom\_image* directory in the root directory of the first partition of the server from which the image is created.

When the instance created from the customized Linux image is started for the first time, Alibaba Cloud writes instance configurations to the *os.conf* file in the *aliyun\_custom\_image* directory. If the *os.conf* file does not exist, Alibaba Cloud will create one.

2. Create a parsing script in the image to parse the *os.conf* file to implement instance configurations. For more information about how to write a script, see Considerations for the parsing script and Example of the parsing script.

### Example of the os.conf file

• os.conf file for instances in the classic network

```
hostname=<yourHostName>
password=<yourPassword>
eth0_ip_addr=10.0.0.2
eth0_mac_addr=00:xx:xx:xx:xx:23
eth0_netmask=255.255.255.0
eth0_gateway=10.0.0.1
eth0_route="10.0.0.0/8 10.0.0.1;172.16.0.0/12 10.0.0.1"
eth1_ip_addr=42.0.0.2
eth1_mac_addr=00:xx:xx:xx:xx:24
eth1_netmask=255.255.255.0
eth1_gateway=42.0.0.1
eth1_route="0.0.0.0/0 42.0.0.1"
dns_nameserver="7.7.7.7 8.8.8.8.8"
```

The following table describes the parameters in the preceding example.

Parameter	Description
hostname	The name of the host.
password	The password. It must be a Base64-encoded string.
eth0_ip_addr	The IP address of the eth0 NIC.
eth0_mac_addr	The MAC address of the eth0 NIC.
eth0_netmask	The network mask of the eth0 NIC.
eth0_gateway	The default gateway of the eth0 NIC.
eth0_route	The eth0 internal routes. By default, routes are separated with semicolons (;).
eth1_ip_addr	The IP address of the eth1 NIC.
eth1_mac_addr	The MAC address of the eth1 NIC.
eth1_netmask	The network mask of the eth1 NIC.
eth1_gateway	The default gateway of the eth1 NIC.
eth1_route	The eth1 Internet routes. By default, routes are separated with semicolons (;).

Parameter
-----------

dns_nameserver	The DNS address list. By default, addresses are separated with spaces.

#### os.conf file for instances in VPCs

```
hostname=<yourHostName>
password=<yourPassword>
eth0_ip_addr=10.0.0.2
eth0_mac_addr=00:xx:xx:xx:xx:23
eth0_netmask=255.255.255.0
eth0_gateway=10.0.0.1
eth0_route="0.0.0.0/0 10.0.0.1"
dns_nameserver="7.7.7.7 8.8.8.8.8"
```

#### The following table describes the parameters in the preceding example.

Parameter	Description
hostname	The name of the host.
password	The password. It must be a Base64-encoded string.
eth0_ip_addr	The IP address of the eth0 NIC.
eth0_mac_addr	The MAC address of the eth0 NIC.
eth0_netmask	The network mask of the eth0 NIC.
eth0_gateway	The default gateway of the eth0 NIC.
eth0_route	The eth0 internal routes. By default, routes are separated with semicolons (;).
dns_nameserver	The DNS address list. By default, addresses are separated with spaces.

# Considerations for the parsing script

Typically, when an instance is started for the first time, Alibaba Cloud writes instance configurations to the *os.conf* file. The os.conf file is in the *aliyun\_custom\_image* directory in the root directory of the partition. However, for a customized Linux image, you must create a predefined parsing script. The script reads the configurations from the *os.conf* file to configure the instance.

The parsing script must meet the conditions described in the following table.

Condition	Description	
Automatic start at system startup	Set the parsing script to run automatically at system startup by placing the script in the /etc/init.d/ directory.	

Condition	Description
Values for configuration items	As shown in the Example of the os.conf section, classic network-type and VPC-type instances differ in the number of configuration items and values of some configuration items.
Path for the configuration file	The device name allocated to the first partition for instances created from the customized Linux image varies depending on whether the instance is I/O optimized. We recommend that you use uuid or label in your parsing code to identify the device allocated for the first partition. The user password must be Base64-encoded in the parsing script.
National	The parsing script can determine the network type by checking whether eth1_route or eth1-related configuration items exist. The script parses and processes the instance accordingly based on the network type.  • VPC-type instances are configured with the default Internet route specified by
Network type	<ul> <li>the eth0_route parameter in the os.conf file.</li> <li>Classic network-type instances are configured with the default Internet route specified by the eth1_route parameter in the os.conf file, and with the default internal route specified by the eth0_route parameter.</li> </ul>
Configuration optimization	Configurations in the <i>os.conf</i> file are executed only once during the instance lifecycle. We recommend that you delete the <i>os.conf</i> file after the parsing script is executed. The parsing script will not execute the configurations in the <i>os.conf</i> file if the script does not read any configurations.
Customized image processing	When a custom image is created based on a customized Linux image, the automatic startup script is also included. Alibaba Cloud will write configurations to the <i>os.conf</i> file when the instance is started for the first time. Then, the parsing script will immediately execute the configurations upon detection.
Configuration change processing	When instance configurations are changed by using the Alibaba Cloud console or calling an API operation, Alibaba Cloud writes new configurations to the <i>os.conf</i> file. Then, the parsing script will run again to issue the changes.

# Example of the parsing script

A parsing script for CentOS is used in this example. You must change the script content based on your operating system and debug the script before you execute it.

```
if [[ -f $os conf file ]]; then
       . $os conf file
       return 0
   else
       return 1
    fi
}
cleanup() {
   # ensure $os conf file is deleted, to avoid repeating config system
   rm $os conf file >& /dev/null
   # ensure $os conf dir exists
   mkdir -p $os_conf_dir
config password() {
   if [[ -n $password ]]; then
       password=$(echo $password | base64 -d)
       if [[ $? == 0 && -n $password ]]; then
           echo "root:$password" | chpasswd
        fi
   fi
config hostname() {
   if [[ -n $hostname ]]; then
       sed -i "s/^HOSTNAME=. */HOSTNAME=$hostname/" /etc/sysconfig/network
       hostname $hostname
   fi
config dns() {
   if [[ -n $dns_nameserver ]]; then
       dns conf=/etc/resolv.conf
       sed -i '/^nameserver.*/d' $dns_conf
       for i in $dns nameserver; do
           echo "nameserver $i" >> $dns conf
       done
   fi
is_classic_network() {
   # vpc: eth0
    # classic: eth0 eth1
   grep -q 'eth1' $os conf file
config network() {
   /etc/init.d/network stop
   config interface eth0 ${eth0 ip addr} ${eth0 netmask} ${eth0 mac addr}
   config route eth0 "${eth0 route}"
   if is classic network; then
       config_interface eth1 ${eth1_ip_addr} ${eth1_netmask} ${eth1_mac_addr}
        config route eth1 "${eth1 route}"
   /etc/init.d/network start
config interface() {
   local interface=$1
   local ip=$2
```

```
local netmask=$3
   local mac=$4
   interface cfg="/etc/sysconfig/network-scripts/ifcfg-${interface}"
   cat << EOF > $interface cfg
DEVICE=$interface
IPADDR=$ip
NETMASK=$netmask
HWADDR=$mac
ONBOOT=yes
BOOTPROTO=static
EOF
config_default_gateway() {
   local gateway=$1
   sed -i "s/^GATEWAY=. */GATEWAY=$gateway/" /etc/sysconfig/network
}
config_route() {
   local interface=$1
   local route="$2"
   route conf=/etc/sysconfig/network-scripts/route-${interface}
   > $route conf
   echo $route | sed 's/;/\n/' | \
       while read line; do
          dst=$(echo $line | awk '{print $1}')
           gw=$(echo $line | awk '{print $2}')
           if ! grep -q "$dst" $route conf 2> /dev/null; then
              echo "$dst via $gw dev $interface" >> $route conf
           fi
           if [[ "$dst" == "0.0.0.0/0" ]]; then
              config default gateway $gw
           fi
       done
start() {
   if load os conf ; then
       config password
       config network
       config_hostname
       config dns
       cleanup
       return 0
   else
      echo "not load $os conf file"
       return 0
   fi
}
RETVAL=0
case "$1" in
   start)
      start
      RETVAL=$?
   ;;
```

```
echo "Usage: $U {start}"

RETVAL=3

;;
esac
exit $RETVAL
```

# 6.3.7. Convert the image file format

ECS can only import image files in the RAW, VHD, or qcow2 format. If you want to import image files in other formats, you must first convert their formats. This topic describes how to use the qemu-img tool to convert other image file formats to VHD or RAW.

#### Context

The qemu-img tool supports the following formats:

- Convert images from the RAW, qcow2, qcow1, VMDK, VDI, VHD (VPC), VHDX, or QED format to the VHD format.
- Convert between RAW and VHD.

**Note** The qemu-img tool cannot convert ISO images to the VHD or RAW format. To convert the image format, install the ISO image to VHD or RAW media, and then create an image in the specified format.

#### Windows

To install qemu-img and convert the image file format, perform the following steps:

- 1. Download and install gemu-img.
  - This example uses installation path *C:\Program Files\qemu*. Visit qemu-img to download the qemu-img tool.
- 2. Create an environment variable for gemu-img.
  - i. Choose **Start > Computer**, right-click Computer, and choose **Properties** from the shortcut menu.
  - ii. In the left-side navigation pane, click **Advanced System Settings**.
  - iii. In the **System Properties** dialog box that appears, click the **Advanced** tab and then click **Environment Variables**.
  - iv. In the Environment Variables dialog box that appears, find the Path variable from the System Variables section and then click Edit. If the Path variable does not exist, click New.
  - v. Add a system variable value.
    - In the Edit System Variable dialog box that appears, add C:\Program Files\qemu to the Variable Value field. Use semicolons (;) to separate different variable values.
    - In the **New System Variable** dialog box that appears, enter *Path* in the **Variable Name** field and enter *C:\Program Files\qemu* in the **Variable Value** field.
- 3. Check whether the environment variable is properly configured.
  - i. Open Command Prompt in Windows.
  - ii. Runthe qemu-img --help command.

If the command output is displayed, the environment variable is properly configured.

- 4. In Command Prompt, run the cd [Directory of the source image file] command to switch to a new file directory, for example, cd D:\ConvertImage .
- 5. Run the qemu-img convert -f qcow2 -O raw centos.qcow2 centos.raw command to convert the image file format.
  - The -f parameter is followed by the source image format.
  - The -O parameter (uppercase is required) is followed by the target image format, source file name, and target file name.

When the conversion is complete, the target file will appear in the directory where the source image file is located.

#### Linux

To install qemu-img and convert the image file format, perform the following steps:

1. Install the gemu-img tool.

Example:

o For Ubuntu, run the apt-get install qemu-utils command.

If the error Unable to locate package qemu-utils is returned during the installation on Ubuntu, run the following commands to install the qemu-img tool:

```
apt-get update #Update the package list.
apt-get install qemu-utils #Install the qemu-img tool.
```

• For CentOS, run the yum install qemu-img command.

Run pip install -r requirements.txt to install all of the dependent libraries based on the libraries contained in the requirements.txt file of cloud-init.

- 2. Run the qemu-img convert -f qcow2 -0 raw centos.qcow2 centos.raw command to convert the image file format.
  - The -f parameter is followed by the source image format.
  - The -O parameter (uppercase is required) is followed by the target image format, source file name, and target file name.

When the conversion is complete, the target file will appear in the directory where the source image file is located.

#### What's next

Import custom images

# 6.3.8. Import custom images

When you deploy services in the cloud, you can manually import image files of your own servers or use Elastic Compute Service (ECS) images. This topic describes the use scenarios and operation procedures of these two methods.

#### **Prerequisites**

The following requirements are met before you manually import a custom image:

- You fully understand the limits and requirements of images. For more information, see <u>Instructions for importing images</u>, <u>Customize Linux images</u> and <u>Convert the image file format</u>.
- Object Storage Service (OSS) is activated. For more information, see Activate OSS.
- If you are using a Resource Access Management (RAM) user, the AliyunECSImageImportDefaultRole role is assigned to the RAM user by the corresponding Alibaba Cloud account.

#### Context

The following table describes the scenarios in which you can use different features to import custom images.

Feature	Scenario	References
Auto- import	Migrate one or more servers such as servers in Internet data centers (IDCs), virtual machines, or cloud servers from other cloud platforms to Alibaba Cloud.	Auto-import
Manual import	<ul> <li>The required operating system is unavailable in Alibaba Cloud. Examples:</li> <li>The specified operating system type does not exist in Alibaba Cloud.</li> <li>The specified operating system version has been discontinued in Alibaba Cloud.</li> <li>The specified operating system is a custom operating system.</li> <li>The manual import feature allows you to select a license type to activate the source operating system and reduce your image costs in the cloud. For more information about license types, see the table that describes parameters for image import in the following section.</li> </ul>	Manual import

# Auto-import

Perform the following operations to use Server Migration Center (SMC) to generate an ECS image for your server.

- 1.
- 2.
- 3.
- 4. In the upper-right corner of the **Images** page, click **Online SMC Migration**. You are redirected to the SMC console.
- 5. Complete the operations as instructed.
  For more information, see Migration process. When your server is migrated, SMC generates a custom image for your server. The image name starts with IMAGE\_FROM\_SMC.

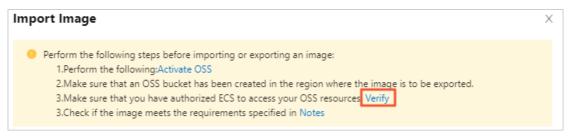
#### Manual import

Perform the following operations to manually import an on-premises image to Alibaba Cloud ECS.

1. Use a third-party OSS client or call an API operation to upload the custom image that you prepared from your computer to OSS.

For information about how to upload an image file larger than 5 GiB in size, see Multipart upload and resumable upload.

- 2.
- 3.
- 4.
- 5. Authorize ECS to access your OSS resources.
  - i. In the upper-right corner of the Images page, click Import Image.
  - ii. In the **Import Image** dialog box, click **Authorize** in Step 3 in the Info section, as shown in the following figure.



- iii. On the Cloud Resource Access Authorization page, AliyunECSImageImportDefaultRole and AliyunECSExportDefaultRole are displayed. Click Confirm Authorization Policy.
- 6. Import the custom image.
  - i. In the upper-right corner of the Images page, click Import Image again.
  - ii. In the Import Image dialog box, configure the parameters described in the following table.

Parameter	Required	Description	
Region of Image	Yes	This parameter is automatically set to the current region. To change the region, you can close the Import Image dialog box and select another region in the left part of the top navigation bar. Then, you can click Import Image to re-open the Import Image dialog box and continue to configure other parameters.	
	Yes	Paste the URL of the image object that you obtained from the OSS console. For more information, see Share objects.  Take note of the following items:	
		<ul> <li>You must use the same Alibaba Cloud account to obtain the URL of the image object and import the image.</li> </ul>	
OSS Object Address		■ The URL of the image object must be the public URL generated when you share the object in OSS. Internal endpoints provided by OSS cannot be used to import images. If you attempt to use an internal endpoint to import an image, an error will occur. For information about the internal endpoints of OSS, see Obtain the internal endpoint of OSS.	
Image Name	Yes	Enter a name for the custom image. The name must be 2 to 128 characters in length, and can contain letters, digits, periods (.), underscores (_), colons (:), and hyphens (-). It must start with a letter.	

Parameter	Required	Description	
Operating System/Platform	Yes	<ul> <li>Operating System:         Select the operating system type of your image. Valid values: Windows and Linux.         If your image is a non-standard platform image, select Linux.     </li> <li>Platform:         The valid values of Platform vary based on the Operating System value that you selected. Select the operating system platform of your image.         Valid values when Operating System is set to Windows: Windows Server 2003, Windows Server 2008, Windows Server 2012, Windows Server 2016, and Windows Server 2019.     </li> <li>Valid values when Operating System is set to Linux: CentOS, SUSE, openSUSE, Ubuntu, Debian, FreeBSD, CoreOS, Aliyun, Anolis, Red Hat, Customized Linux, and Others Linux. Submit a ticket to check whether Others Linux images can be imported.     </li> <li>If the operating system of your image is a distribution customized based on the Linux kernel, Submit a ticket.</li> </ul>	
SVCTAM		Select x86_64 for 64-bit operating systems, i386 for 32-bit operating systems, and arm64 for 64-bit ARM operating systems.	
System Disk Size (GiB)	No	Specify the system disk size. Unit: GiB. Valid values: 5 to 500.	
lmage Format	No	<ul> <li>Valid values: Auto-detect, RAW, QCOW2, and VHD.</li> <li>If you do not know the image format, you can select Auto-detect to allow the system to determine the image format.</li> <li>If you want to select an image format, we recommend that you select the QCOW2 or VHD format, which has a smaller transmission footprint.</li> <li>Note The ISO format is not supported and must be converted to the RAW, VHD, or QCOW2 format before the image can be imported. You can also use Packer to create and import an on-premises image. For more information, see Use Packer to create and import an on-premises image and Alicloud Image Builder parameters used to implement DevOps.</li> </ul>	

Parameter	Required	Description	
License Type	No	Select a license type to activate the source operating system after the image is imported. Default value: Auto. Valid values:  BYOL: the license that comes with the source operating system. If this value is selected, make sure that your license key can be used in Alibaba Cloud.  For information about Red Hat Cloud Access and Windows licenses, see Overview and Microsoft licenses and Software Assurance.  Aliyun: the license allocated by Alibaba Cloud. If this value is selected, a license allocated by Alibaba Cloud is applied based on your selected platform.  Auto: Alibaba Cloud automatically checks the source operating system and allocates a license. If this value is selected, the system automatically checks whether a license has been allocated by Alibaba Cloud for your specified operating system platform, and then performs an operation based on the check result.  If a license has been allocated by Alibaba Cloud for the specified operating system platform, the system allocates the license to the imported image.  If no license has been allocated by Alibaba Cloud for the specified operating system platform, the system switches the license type to BYOL.	
Description	No	Enter a description for the custom image.	
Add Data Disk Image		Select this option if you want to import an image that contains data from data disks. Supported data disk capacity ranges from 5 GiB to 1,000 GiB.	
Resource Group No		Specify a resource group for the custom image to simplify management.	
Tag No		Specify tags to add to the custom image. You can use tags to classify images for easy search and batch operations.	

- iii. Confirm that your configurations are correct and click **OK**. The system then creates a task to import the custom image.
- 7. (Optional)View the task progress in the image list of the destination region. Before the task is completed, you can find the imported custom image or cancel the task on the Task Logs page in the ECS console.

The amount of time it takes to import a custom image depends on the image size and the number of concurrent ongoing import tasks.

When you import an image, a snapshot is automatically generated. You can view the information of the snapshot on the <code>Snapshots</code> page in the ECS console. While the image import task is in progress, the state of the snapshot is displayed as <code>Failed</code>. When the image import task is completed, the snapshot enters the <code>Successful</code> state and the image enters the <code>Available</code> state. The snapshot size is equal to the size of the imported image file, regardless of the system disk size that you set when you import the image. Snapshots are billed based on the amount of storage space they occupy. For more information, see <code>Snapshots</code>.

#### What's next

Create an ECS instance by using a custom image

#### Related information

#### References

- Import Image
- Use Packer to create and import an on-premises image
- Change the operating system
- Image FAQ
- What is SMC?
- Overview
- Microsoft licenses and Software Assurance

# 6.4. Update a custom image

This topic describes how to update custom images by using Operation Orchestration Service (OOS). OOS provides public templates to update images automatically. To create an immediate or scheduled O&M task, you only need to select a source image and specify required parameters such as the Cloud Assistant command in a public template. The O&M task is then automatically executed based on the definitions in the template.

#### Context

You can use the ACS-ECS-UpdateImage template to sequentially execute the following tasks to update the image to a new custom image:

- 1. Check whether the name of the new custom image already exists and make sure that the name complies with the naming conventions.
- 2. Create and launch a temporary ECS instance based on parameters such as the instance type, source image ID, and security group ID that you have configured.
- 3. Check whether the Cloud Assistant client is installed on the temporary ECS instance. If not, install the Cloud Assistant client.
- 4. Run scripts by using Cloud Assistant to update the system environment of the instance.
  - **? Note** OOS calls Cloud Assistant API operations to run shell, bat, or PowerShell scripts to update the system environments of ECS instances. For more information, see **Overview**.
- 5. Stop the temporary ECS instance.
- 6. Create a custom image from the temporary ECS instance.

7. Release the temporary ECS instance.

#### Procedure

- 1.
- 2.
- 3.
- 4.
- 5. On the **Images** page, click the **Custom Image** tab.
- 6. Find the image that you want to update, click the icon in the **Actions** column, and then click **Update Image**.
- 7. In the Operation Orchestration Service console, perform the following operations:
  - i. Complete the configurations in the General Settings section, including New Image Name, Description, and Permissions.
  - ii. (Optional)Complete the Image Distribution settings. You can set whether to copy or share the image.
  - iii. Select an image. If you have selected an image in Step 6, the image is automatically selected in this step. Check whether the selected image is the one you want to update.
  - iv. Configure an intermediate instance. Set VPC, vSwitch, Security Group, and Instance Type to create the intermediate instance.
  - v. (Optional)Send remote commands. You can select a command type and enter the command content. The command is executed on the intermediate instance.
  - vi. (Optional)Complete the advanced settings. You can configure scaling groups, add tags, or save the current configurations as a template.
- 8. Click Execute Now, and confirm the prompt information in the Confirm Balance message.
- 9. On the **Create or Update Images** page, you can view the task status.

#### Result

If the O&M task is created and is in the **Running** state, the image is being updated. When Task Status changes to **Success**, the image is updated. You can click Details in the Actions column corresponding to the execution and view the ID of the new image in the Output Parameters section of the Basic Information tab.

**Note** To view details of the update process, click **Details** in the Actions column corresponding to the execution and click the **Logs** tab.

#### Related information

#### References

- •
- •
- Introduction to OOS

# 6.5. Copy a custom image

This topic describes how to copy a custom image from one region to another within your Alibaba Cloud account. This feature can be used to create identical Elastic Compute Service (ECS) instances and back up data of instances across regions.

#### Context

The following table describes a variety of scenarios in which custom images are copied.

Scenario	Procedure	API operation
Copy images across different regions within the same account	<ol> <li>Create a custom image from an instance</li> <li>Copy images to a different region</li> </ol>	<ol> <li>Createlmage</li> <li>Copylmage</li> </ol>
Copy images across regions and accounts	<ol> <li>Create a custom image from an instance</li> <li>Copy images to a different region</li> <li>Share or unshare a custom image</li> </ol>	<ol> <li>Createlmage</li> <li>Copylmage</li> <li>ModifylmageSharePermission</li> </ol>
Share images within the same region to another account	You can share images without the need to copy the images. For more information, see Share or unshare a custom image.	ModifylmageSharePermission
	Select Encrypt when you copy a custom image within the same region.  Note To copy a custom image within the same	
Change the encryption state of a custom image, encrypt a system disk, or encrypt operating system data	region, you must encrypt the image copy.  • You cannot cancel an ongoing copy task of an encrypted image.	Copylmage

Before you copy a custom image, take note of the following items:

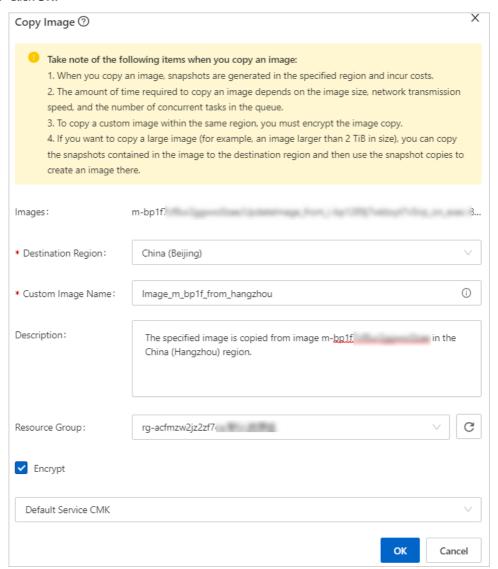
- You may be charged the following fees:
  - Fees for the use of storage capacity occupied by snapshot copies. For more information, see
     Overview.
    - **? Note** When an image is copied to a region, one or more corresponding snapshots are generated in the region. Then, a custom image is generated in the region based on these snapshots. These snapshots are retained permanently, regardless of the retention period of the snapshots of the source image.
  - Fees for cross-region data transfer. This data transfer is not billed for the moment. For the latest billing details, see the official Alibaba Cloud website for announcements.
- The amount of time it takes to copy an image depends on the image size, network transmission speed, and the number of concurrent tasks in the gueue.
  - **Note** To copy a large image such as an image larger than 2 TiB in size across regions, you can first copy the associated snapshots to the destination region and then create a custom image from these snapshots in the destination region. For more information, see Copy a snapshot and Create a custom image from a snapshot.
- The role authorization information, service authorization information, and instance user data of the source image are not copied to the image copy.

#### Procedure

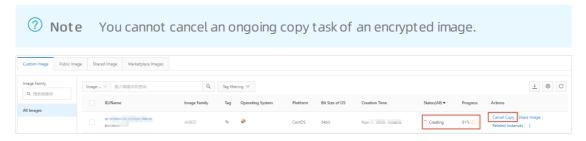
- 1.
- 2.
- 3.
- 4. On the **Images** page, click the **Custom Images** tab.
- 5. Find the custom image that you want to copy and click **Copy Image** in the **Actions** column.
- 6. In the **Copy Image** dialog box, specify parameters.
  - i. Confirm the ID of the custom image that you want to copy.
  - ii. Select a region from the **Destination Region** drop-down list, which can be the current region or another region.
  - iii. In the **Custom Image Name** field, enter an image name to display in the destination region. Example: **Image m bp1f from hangzhou**.
  - iv. (Optional)In the **Description** field, enter an image description to display in the destination region.
  - v. (Optional)Select a resource group from the Resource Group drop-down list.
  - vi. (Optional)Select Encrypt and then select Default Service CMK from the drop-down list.

    Default Service CMK: The default key generated by Key Management Service (KMS).
    - ? Note To copy a custom image within the same region, you must select Encrypt and select a key to encrypt the image.

#### vii. Click OK.



- 7. Check the result and manage the task.
  - i. In the upper-left corner of the top navigation bar, select the destination region from the drop-down list and view the copy progress of the custom image.
     When the copy task is completed, the progress reaches 100% and the image enters the Available state. An image is generated in the destination region and has a unique image ID.
  - ii. (Optional) If the **Progress** has not reached 100% and the **Status** is **Creating**, you can cancel the copy task by clicking **Cancel Copy**.



#### What's next

You can use the image copy to create an instance or replace the system disk of an instance as described in the following topics:

- Create an ECS instance by using a custom image
- Replace the operating system of an instance by using a non-public image

#### Related information

- Copylmage
- CancelCopyImage
- Copy a snapshot

# 6.6. Share or unshare a custom image

This topic describes how to share or unshare custom images. Shared images can be used to deploy Elastic Compute Service (ECS) instances across accounts. You can share your custom images to other Alibaba Cloud accounts. These Alibaba Cloud accounts can then use the shared images to create their own instances. You can also unshare shared custom images from other Alibaba Cloud accounts.

#### **Prerequisites**

Before you share a custom image, make sure that all sensitive data and files are removed from the image.

#### Context

Images shared to an account do not count against the image quota for the account. The account is not charged for images shared to it.

When an account creates their own instances by using an image that you shared, you do not incur charges. Only the account that used the image to create instances is charged. For more information about the billing of shared images, see the "Image prices" section in Image Overview.

The account can use the shared images but cannot delete the shared images. You must unshare a shared image before you can delete it. For more information, see Delete a custom image.

Before you share an image, take note of the following items:

- You can share only the custom images that are created within your account. You cannot share custom images that are created and shared by other accounts.
- Each custom image can be shared to up to 50 accounts.
- You can share images between the accounts on the China (aliyun.com), International (alibabacloud.com), and Japan (jp.alibabacloud.com) sites. However, the custom images that are created based on Alibaba Cloud Market place images cannot be shared.
- Custom images cannot be shared across regions. If you want to share a custom image across regions, you must copy the image to the destination region first. For more information, see Copy a custom image.
- ECS does not guarantee the integrity and security of shared images. Before you use shared images, make sure that the images come from trusted accounts. You must assume all risks.
- If you want to share an encrypted custom image, you must use Resource Access Management (RAM) to obtain permissions. For more information, see <a href="Share encrypted custom images">Share encrypted custom images</a>.

② **Note** If encryption is enabled for a disk attached to an ECS instance, the custom image created from the ECS instance is an encrypted custom image. For more information about the encryption feature, see 加密概述.

#### Share a custom image

You can perform the following operations to share your custom image to an Alibaba Cloud account.

- 1.
- 2.
- 3.
- 4. On the **Custom Image** tab, find the custom image that you want to share and click **Share Image** in the **Actions** Column.
- 5. In the **Share Image** dialog box, enter the ID of the account to which you want to share the image and select the option under the Account field. Then, click **Share Image**.

Images can be shared only to Alibaba Cloud accounts, not to RAM users. To obtain the ID of your Alibaba Cloud account, move the pointer over the profile picture in the upper-right corner of the Alibaba Cloud Management Console. If the account is identified as **Main Account** in the user information panel, the account ID is an Alibaba Cloud account ID.

After you share a custom image to an account, the account can view the shared image in the ECS console by choosing Instances & Images > Images and then clicking the Shared Image tab in the same region as that you previously selected. The account can perform the following operations:

- Create one or more ECS instances from the shared image.
  - When the account creates an ECS instance, it can click the Share Image tab and then select the shared image in the Image section. For more information, see Create an instance by using the wizard.
- Use the shared image to replace the system disk of an ECS instance.
   For more information, see Replace the operating system of an instance by using a non-public image.

#### Unshare a custom image

You can perform the following operations to unshare your custom image from an Alibaba Cloud account.

- $\bigcirc$  Notice After you unshare a custom image from an account, the following results occur:
  - The account is no longer able to query the image by using the ECS console or by calling an API operation.
  - The account is no longer able to create ECS instances or replace system disks by using the image.
  - The system disks of ECS instances that were created from the shared image cannot be reinitialized.
- 1.
- 2.
- 3.

- 4. On the **Custom Image** tab, find the custom image that you want to unshare and click **Share Image** in the **Actions** Column.
- 5. In the **Share Image** dialog box, find the ID of the account from which you want to unshare the image and click **Cancel Sharing** in the Actions column.

#### Related information

- ModifyImageSharePermission
- DescribeImageSharePermission

# 6.7. Export a custom image

You can export a created custom image to an Object Storage Service (OSS) bucket and download the image to local computers. This topic describes how to export custom images and what notes must be taken of.

#### **Prerequisites**

• An OSS bucket is available in the same region as the custom image to be exported.

If you do not have OSS buckets available in the region, create an OSS bucket. For information about how to create an OSS bucket, see Create buckets.

**Note** When you export a custom image, you are charged for OSS storage and the traffic generated for image download. For more information, see **Billing overview**.

- The custom image to be exported meets the following requirements:
  - It is not created based on an Alibaba Cloud Market place image.
  - o It does not contain a Windows Server operating system.
  - It does not contain snapshots of more than four data disks. The size of each data disk does not exceed 500 GiB.

**Note** If the capacity of a data disk exceeds 500 GiB, you must distribute the data on the disk to multiple data disks that are smaller than 500 GiB in size and then create and export custom images from these smaller data disks.

#### Context

290

Before you export a custom image, take note of the following points:

- The amount of time it takes to export a custom image depends on the size of the image and the number of concurrent tasks in the queue.
- If an exported custom image contains data disk snapshots, multiple objects appear in your OSS bucket.

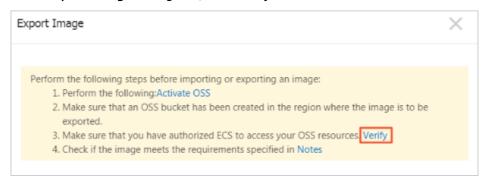
Objects whose names contain *system* are system disk snapshots. Objects whose names contain *data* are data disk snapshots. The identifier of a data disk snapshot is the mount point of the source data disk, such as xvdb and xvdc.

• To use the exported image to create identical Linux instances, make sure that the storage location and storage space division of files recorded in /etc/fstab are consistent with the exported data disk snapshot information.

- If the cloud disk does not contain data when the custom image is created, the decompressed image file also does not contain data.
- The exported image may not be normally started on instances that run operating systems different from the images, but you can mount the disk in the image to obtain data.

#### **Procedure**

- 1.
- 2.
- 3.
- 4. Authorize ECS to access OSS.
  - i. Find the image that you want to export. In the **Actions** column, click the icon, and then click **Export Image**.
  - ii. In the Export Image dialog box, click Verify in the notes.



iii. In the Cloud Resource Access Authorization dialog box, click Confirm Authorization Policy to allow ECS to access your OSS resources.

5.

- 6. On the **Custom Images** tab, find the image that you want to export and choose **More > Export Image** in the **Actions** column.
- 7. In the **Export Image** dialog box, configure the following parameters:
  - Image Format: Select a format in which to export the custom image. Valid values: RAW, VHD, QCOW2, VDI, and VMDK.
    - Note The feature of selecting image export formats is available only in some regions.
  - OSS Bucket Address: Select an OSS bucket that belongs to the same region as the custom image.
  - OSS Object Prefix: Set the prefix of the object name for the custom image. For example, if you set OSS Object Prefix to Demo, the exported image is named *Demo-<Automatically generated o* bject name>.
- 8. Click **OK** to export the custom image.

You can cancel an image export task at any time before the task is complete. Go to the Tasks page in the ECS console, find the corresponding task in the specified region, and cancel the task.

#### What's next

Download the custom image. For more information, see Download objects.

Note If you select the RAW image format, the default file name extension of the exported custom image is .raw.tar.gz, and the file name extension of the decompressed image is .raw. If your local computer runs a Mac OS X operating system, we recommend that you use GNU Tar to decompress the image.

#### Related information

- ExportImage
- CancelTask

# 6.8. Modify a custom image

This topic describes how to modify the name and description of a custom image.

#### **Prerequisites**

A cust om image is created or a cust om image creation request is submitted.

#### Procedure

- 1.
- 2.
- 3.
- 4. On the **Images** page, click the **Custom Image** tab. Find the custom image that you want to modify and modify its name and description.
  - o Modify the image name
    - a. Move the pointer over the name of the image and click the 🗾 icon.
    - b. In the dialog box that appears, enter a new image name in the **Image Name** field and click **OK**.
  - Modify the image description
    - a. In the Actions column, click the : icon and select Modify Description.
    - b. In the **Modify Description** dialog box, enter a new description and click **OK**.

#### Related information

ModifylmageAttribute

# 6.9. Edit the tags of a custom image

Tags can be used to identify resources with the same characteristics (such as custom images that belong to the same organization or that serve the same purpose) for easy search and management. This topic describes how to edit the tags of an existing custom image.

#### Context

For information about how to use tags, the resources that support tags, and the limits on tags, see

Overview and the "Tag limits" section of the Limits topic.

#### Procedure

- 1.
- 2.
- 3.
- 4. On the **Custom Image** tab, find the custom image whose tags you want to edit and move the pointer over the icon in the **Tag** column.
  - If you have not added tags to the custom image, click Add to add tags.
  - If you have added tags to the custom image, click **Edit** to edit the tags.
- 5. In the **Edit Tags** dialog box, click Available Tags to select existing tags or click Create to create tags. Then, click **OK**.

#### What's next

After tags are added to your custom images, you can filter the images by tag to perform different operations. For example, you can copy custom images that have a set of tags and update custom images that have a different set of tags.

## 6.10. Delete a custom image

You can delete a custom image that you no longer need.

#### Context

Take note of the following items before you delete a custom image:

- After a custom image is deleted, you are no longer able to create instances from the image.
- After a custom image is deleted:
  - The ECS instances that were created from the image can still be used and you continue to be charged for the instances.
  - The system disks of ECS instances that were created from the image cannot be re-initialized.
- Before you delete a shared image, you must unshare the image. After a shared image is deleted:
  - The accounts to which the image was shared are no longer able to query the image by using the ECS console or by calling an API operation.
  - The accounts to which the image was shared are no longer able to create ECS instances or replace the system disks by using the image.
  - The system disks of ECS instances that were created from the shared image cannot be reinitialized.
- If you delete an image that has been copied, the image copies are not affected. If you delete the copies of an image, the image is not affected.
- If you delete a custom image, snapshots contained in the image are not deleted. If you do not want to keep the snapshots, navigate to the Snapshots page and delete the snapshots.

#### **Procedure**

1.

- 2.
- 3.
- 4. On the Images page, click the Custom Image tab. Select one or more custom images that you want to delete and click Delete Image in the lower part of the page.
- 5. In the Delete Image dialog box, select Forcibly Delete.
- 6. Click OK.

#### **Related information**

• Deletelmage

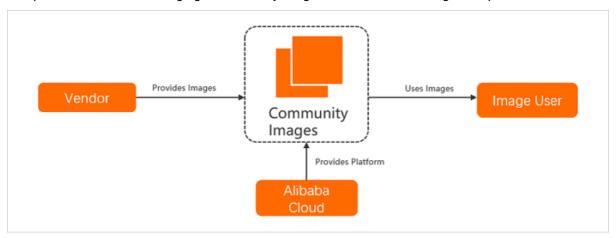
# 7. Community images

# 7.1. Overview

Community images are publicly available. Custom images can be published as community images for other users to obtain and use.

#### Introduction

The process related to managing a community image involves the following three parties.



- Alibaba Cloud: provides the platform where community images are published and managed by vendors.
- Vendor: provides custom images and publishes them as community images.
- Image user: uses community images.

#### Operations

The following table describes the operations supported by community images.

Operation	Description	References
Publish a community image	You can publish a custom image that is in the Available state as a community image.	Publish a community image
Search for a community image	You can search for a specific community image by specifying the name, ID, or owner ID of the community image. Then, you can use the community image to create Elastic Compute Service (ECS) instances.	Search for a community image
Use a community image	You can use a community image to quickly deploy ECS instances whose operating systems, applications, and data suit your business needs.	Use a community image
Unpublish a community image	You can unpublish a published community image when you want to update or stop maintaining the image.	Unpublish a community image

# 7.2. Publish a community image

You can publish custom images that are in the Available state as community images. This topic describes how to publish a custom image in the Elastic Compute Service (ECS) console.

#### **Prerequisites**

- The owner of each community image is responsible for maintaining the quality and updates of the image. Alibaba Cloud merely provides the platform where community images are published and managed. Therefore, you must fully understand Product Terms, Terms of Use, and Membership Agreement before you publish a community image.
- No sensitive words are contained in the name and description of the community image to be published, and no trojans, vulnerabilities, or weak passwords are contained in the community image.

#### Context

Before you publish a community image, take note of the following items:

- Encrypted images cannot be published as community images.
- You must sign *Product Terms, Terms of Use,* and *Membership Agreement* before you can publish a community image.
- Community images are publicly available. A community image is available to all Alibaba Cloud accounts in the region where it resides.
- Community images cannot be shared, exported, or copied.

#### **Procedure**

1. Go to the Images page.

i.

ii.

iii.

- 2. On the **Custom Images** tab, find the custom image that you want to publish as a community image. Move the pointer over the icon in the **Actions** column and select **Publish as Community Image**.
- 3. In the Publish As Community Image message, click OK.
  - If you have not signed Community Image Terms of Service, read the related terms and agreements, select You have signed and agreed to the following agreements:, and then click Publish.

Notice Read *Product Terms, Terms of Use* and *Membership Agreement* before you sign Community Image Terms of Service.

o If you have signed the Community Image Terms of Service, click OK.

After the community image is published, you can go to the **Images** page and click the **Community Images** tab to view the information of the published community image.

# 7.3. Search for a community image

You can search for a community image by image name, image ID, or owner ID. Then, you can use the community image to create Elastic Compute Service (ECS) instances. This topic describes how to search for a community image in the ECS console.

#### Procedure

- 1. Go to the Images page.
  - i.
  - ii.
  - iii.
- 2. Click the Community Images tab.

On the Community Images tab, you can view all community images located in the current region.

3. Search for a community image in the search box.

Perform the following operations.



- i. In the section tagged ①, select an option from the drop-down list. The following options are supported:
  - Image Name
  - Image ID
  - Owner ID
- ii. In the section tagged ②, enter a value that matches the selected option.
- iii. In the section tagged ③, click the 🔍 icon to search for the community image.

For example, you can select Image ID from the drop-down list and enter m-bp10pa9j0rio19ky\*\*\*\* in the search box. A similar result is displayed, as shown in the following figure.



#### What's next

You can use the community image to create ECS instances. For more information, see Use a community image.

# 7.4. Use a community image

You can use community images to deploy Elastic Compute Service (ECS) instances whose operating systems, applications, and data suit your business needs. This topic describes how to create an ECS instance by using a community image.

#### Context

Before you use a community image, take note of the following items:

- Community images are provided and updated by Alibaba Cloud users. Alibaba Cloud does not guarantee the quality of community images. You must fully understand and agree to the *Product Ter ms, Terms of Use,* and *Membership Agreement* before you use community images.
- If a community image is unpublished, disks of ECS instances created by using this community image cannot be re-initialized.

#### **Procedure**

1. Go to the Images page.

i.

ii.

iii.

2. Click the **Community Images** tab and find the community image that you want to use. For more information about how to find a community image, see Search for a community image.

3. Click Create Instance in the Actions column of the community image.



4. On the instance buy page, configure the parameters.

Take note of the following items:

- The region and the community image are specified by default.
  - If you know the ID and region of another community image, you can change the community image and the region.
- Read and select the *Product Terms, Terms of Use,* and *Membership Agreement*.
- Configure other parameters based on your business needs. For more information, see Create an instance by using the wizard.

# 7.5. Unpublish a community image

You can unpublish published community images if you want to update them or stop maintaining them. This topic describes how to unpublish a community image in the Elastic Compute Service (ECS) console.

#### Context

Before you unpublish a community image, take note of the following items:

- After a community image is unpublished, it is no longer available to other Alibaba Cloud accounts. If the community image was shared to other Alibaba Cloud accounts before it is unpublished, the share relationship continues.
- After a community image is unpublished, the system disks of ECS instances created by using this community image cannot be re-initialized.

#### **Procedure**

- 1. Go to the Images page.
  - i.
  - ii.
  - iii.
- 2. On the **Custom Image**s tab, find the published community image that you want to unpublished.

  Move the pointer over the icon in the **Actions** column and click **Unpublish Community Image**.
- 3. In the **Unpublish Community Image** message, confirm the image information and click **OK**. When the community image is unpublished, the **The community image is unpublished**. massage appears on the **Images** page.

# 8. Alibaba Cloud Marketplace images

You can purchase Alibaba Cloud Market place images to create Elastic Compute Service (ECS) instances that are equipped with pre-installed runtime environments or software applications. These images can be used for scenarios such as website building, application development, and visualized management. The created ECS instances can be used out-of-the-box, which is time-saving and convenient.

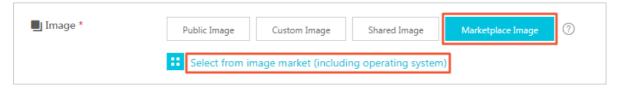


Alibaba Cloud Market place is a network platform established by Alibaba Cloud. On this network platform, independent software vendors (ISVs) provide their images and relevant services to users. Alibaba Cloud Market place images integrate specific software environments and features such as the PHP runtime environment and the control panel into the operating system. You can select images based on these environments and features when you create instances.

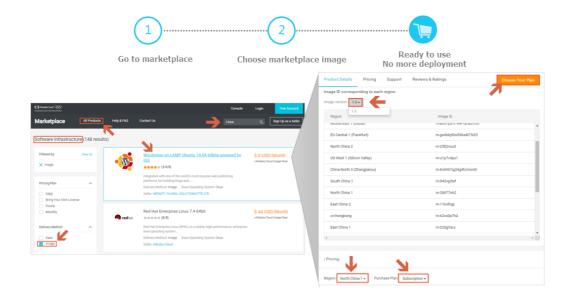
#### Create an instance from an Alibaba Cloud Marketplace image

You can use one of the following methods to create an instance from an Alibaba Cloud Marketplace image. For more information, see Create an instance by using the wizard.

Method 1: Select an Alibaba Cloud Market place image when you create an instance.



Method 2: Go to Alibaba Cloud Market place to purchase an image and then create an instance.



**Note** If you create an instance from an Alibaba Cloud Marketplace image, your license may become invalid. If you encounter this problem, contact the image vendor for technical support.

#### Change the operating system of an instance by using an Alibaba Cloud Marketplace image

If you have purchased an ECS instance and you want to replace the image of the instance with an Alibaba Cloud Market place image, you can replace the system disk.

On the Change Operating System page, select Market place Image for Image and then click Select from Image Market (Including Operating System). In the Market place Image dialog box, select the desired image. For more information, see Replace the operating system of an instance by using a non-public image.



# 9.1mage family 9.1. Overview

Alibaba Cloud provides the image family feature to update and roll back images. You can set an image family for an updated image. Then, when you use the image family to create Elastic Compute Service (ECS) instances, you do not need to obtain the ID of the updated image. The system automatically selects the latest available custom image in the image family. You can also set the state of the latest available custom image family to Deprecated to roll back the image.

#### **Features**

- When you call the Runinstances or Createinstance operation to create ECS instances, you can set the <code>ImageFamily</code> parameter to automatically select the latest available custom image in the image <code>family</code> instead of specifying the <code>ImageId</code> parameter.
- If you want to roll back an image in an image family, you can set the state of the image to Deprecated or delete the latest available custom image in the image family, and then create an ECS instance from the image family.
- You can restore an image from the Deprecated state to the Available state by using the ECS console or by calling the ModifylmageAttribute operation.

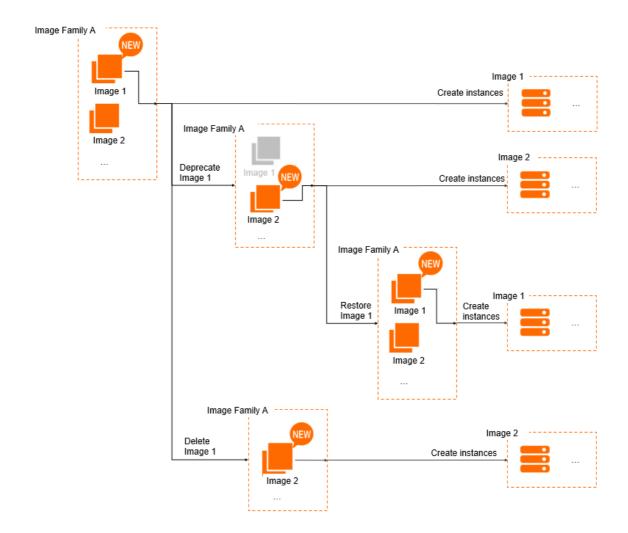
#### Limits

- You can set image families only for custom images.
- After you set an image family for a custom image, you can modify but cannot delete the image family.
- Image family names must be unique within the same account in the same region. Image family names can be the same in different regions.
- An image family can contain a maximum of 100 images.
- The number of image families within a single region cannot exceed 100.
- If a custom image is shared with other users, you must unshare the custom image before you put the custom image into the Deprecated state.
- Custom images in the Deprecated state can be used to create instances or replace system disks, but cannot be shared or copied.

#### **Usage notes**

- Before you use the image family feature, you must set an image family for your images by using the ECS console or by calling API operations. For more information, see Set an image family.
- You can create ECS instances from image families only by calling API operations. When you create an ECS instance from an image family, the system automatically selects the latest available custom image in the image family. For more information, see Create an ECS instance from an image family.
- You can set the state of an image in an image family to Deprecated or restore the image. When you set the state of an image to Deprecated, the image can be rolled back. For more information, see Set the state of an image to Deprecated or restore an image.
- You can call the Describe image From Family operation to query the latest image in the specified image family.

The following figure shows the procedure for using an image family.



#### Related information

- RunInstances
- CreateInstance
- Createlmage
- Describelmages
- ModifyImageAttribute

# 9.2. Set an image family

This topic describes how to set image families by using the Elastic Compute Service (ECS) console or by calling API operations.

#### **Prerequisites**

An ECS instance or a system disk snapshot is created. For more information, see Creation method overview or Create a snapshot of a disk.

#### Context

Image families cannot be separately created. You must set image families when you create custom images. If the image family that you set does not exist, the system automatically creates the image family. If the image family that you set already exists, the system assigns your custom images to the image family.

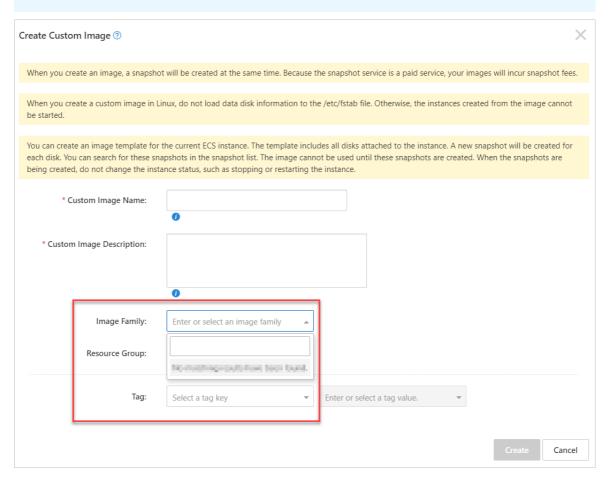
#### Set an image family by using the ECS console

- 1.
- 2.
- 3. Use one of the following methods to create a custom image:
  - In the left-side navigation pane, choose Instances & Images > Instances. For more information, see Create a custom image from an instance.
  - In the left-side navigation pane, choose **Storage & Snapshots** > **Snapshots**. For more information, see Create a custom image from a snapshot.

In the Create Custom Image dialog box, set Image Family for the custom image.

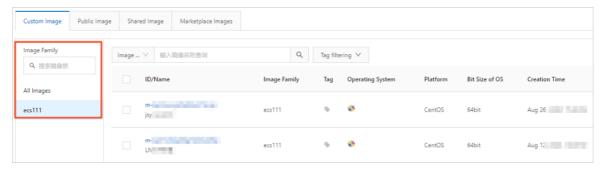
You can specify a name for the image family or select an existing image family.

**?** Note The name must be 2 to 128 characters in length. It must start with a letter and cannot start with http://, https://, acs:, or aliyun. It can contain letters, digits, colons (:), underscores (\_), and hyphens (-).



4. In the left-side navigation pane, choose **Instances & Images** > **Images** to view the created custom image and the corresponding image family.

In the **Image Family** section, enter an image family name or click the image family name for exact match.



#### Set an image family by calling an API operation

This section describes how to set an image family when you call the CreateImage operation to create a custom image. For more information about the operation, see CreateImage.

- 1. Log on to OpenAPI Explorer.
- 2. Call the CreateImage operation to create a custom image and set an image family.

In this example, create a custom image from an instance and set an image family named hangzhou -image . Specify the following API request parameters:

- RegionId: Specify the ID of the region. Example: cn-hangzhou .
- InstanceId: Specify the ID of the instance. Example: i-bp117mtqmllecpwl\*\*\*\* .
- ImageFamily: Specify the image family name. Example: hangzhou-image.

After the custom image is created, the ID of the custom image is displayed.

```
令

"RequestId": "46BC46F8-922C-4419-8B5D-E713F663E00A",
"ImageId": "m-bp1h""

]
```

3. Call the Describe Images operation to query the image family of the custom image created in the previous step. For more information about the operation, see Describe Images.

Specify the following API request parameters:

- RegionId: Specify the same region ID as in the previous step.
- Imageld: Enter the ID of the custom image created in the previous step.

The following figure shows a sample response. The value of the ImageFamily parameter is han gzhou-image .

```
},
                 "Status": "Available",
                 "Progress": "100%",
                 "Usage": "none",
                 "Architecture": "x86_64",
                  "ProductCode": "",
                  'IsCopied": false,
                 "ImageFamily": "hangzhou-image",
                  "IsSupportIoOptimized": true,
                 "IsSupportCloudinit": true,
                  "ImageName": "m-bp1h46wfpj ...,
                  "DiskDeviceMappings": {
"DiskDeviceMapping": [
-
                             "SnapshotId": "s-bp1h4",
                             "Type": "system",
                             "Progress": "",
```

#### What's next

If you want to modify an image family, you can call the ModifylmageAttribute operation. For more information, see ModifylmageAttribute.

# 9.3. Create an ECS instance from an image family

This topic describes how to create an Elastic Compute Service (ECS) instance from an image family. You can only call the Runinstances or Createinstance operation to create an ECS instance from an image family.

#### **Prerequisites**

One or more custom images are created, and the same image family is set for these custom images. For more information, see Set an image family. In this example, the image family named image-test is used.

#### Context

When you create an ECS instance from an image family, the system automatically selects the latest available custom image in the image family. The latest available custom image is the available custom image that is created at the latest time in the image family. For example, you have two available custom images in your image family. One was created on January 1, 2020, and the other was created on January 3, 2020. When you create an ECS instance from the image family, the system selects the custom image that created on January 3, 2020 by default.

For more information about the operations called to create ECS instances, see Runinstances or CreateInstance. In this topic, the Runinstances operation is called.

#### Call the Runinstances operation to create an ECS instance

- 1. Log on to OpenAPI Explorer.
- 2. Call the Runinstances operation to create an ECS instance.

  In this example, the image-test image family is used to create the instance. Specify the

> Document Version: 20220712

following API request parameters:

- RegionId: Specify the ID of the region. Example: cn-hangzhou .
- InstanceType: Specify the instance type of the instance. Example: ecs.g6.large .
- ImageFamily: Specify the image family name. Example: image-test .
- SecurityGroupId: Specify the ID of the security group. Example: sg-bp1i4c0xgqxadew2\*\*\*\* .
- VSwitchId: Specify the ID of the vSwitch. Example: vsw-bp1ddbrxdlrcbim46\*\*\*\* .

The following section shows a sample response. The ID of the created instance is displayed.

```
{
    "RequestId": "409D4604-84D0-4F16-B99E-809E2E72****",
    "InstanceIdSets": {
        "InstanceIdSet": [
            "i-bplenv7nl3mijm2t****"
        ]
    }
}
```

#### Call an API operation to verify the image information

- 1. Log on to OpenAPI Explorer.
- 2. Call the DescribeImageFromFamily operation to query the latest available custom image in the specified image family.

In this example, the latest image in the image-test image family is queried. Specify the following
API request parameters:

- RegionId: Specify the region ID of the image family.
- ImageFamily: Specify the image family name. Example: image-test .

The following figure shows a sample response. The ID of the latest available custom image is displayed.

```
},
"ImageVersion": "",
"OSTvpe": "linux",
"ImageId": "m-bp1e----------",
"CreationTime": "2020-03-17T06:19:19Z",
```

3. Call the DescribeInstances operation to query the image information of an instance.

Query the ID of the image based on the ID of the instance and determine whether the image is the latest available custom image in the image family. Specify the following API request parameters:

- RegionId: Select the region ID of the instance.
- InstanceIds: Specify the instance ID in the ["i-bplenv7nl3mijm2t\*\*\*\*"] format.

The following figure shows a sample response. The image ID of the instance is the same as the ID of the latest available custom image in the image family.

```
"OSName": "CentOS 8.0 64位",
"InstanceNetworkType": "vpc",
"InnerIpAddress": {
    "IpAddress": []
    },
    "ExpiredTime": "2099-12-31T15:597",
    "ImageId": "m-bp16",
    "EipAddress": {
        "AllocationId": "",
        "IpAddress": "",
        "InternetChargeType": ""
```

#### References

DescribeInstances

# 9.4. Set the state of an image to Deprecated or restore an image

Custom images can be put into the Deprecated state or be restored. You can set the state of an available custom image in an image family to Deprecated to roll back the image. You can also restore the latest custom image in the image family.

#### **Prerequisites**

Two custom images are available in the same image family. For more information, see Set an image family. In this example, the image family is named image-test and contains two images: image1 and image2 . image2 is the latest available custom image.

#### Context

The state of a custom image can be changed regardless of whether the custom image belongs to an image family. The state of a custom image can be changed from Available to Deprecated, and an image in the Deprecated state can be restored to the Available state. Operations vary based on the states of images:

- A custom image that has been shared must be unshared before the image can be put into the Deprecated state.
- Custom images in the Deprecated state can be used to create instances or replace system disks, but cannot be shared or copied.
- You can set the state of the latest available custom image in an image family to Deprecated to roll
  back the image. When you use the image family to create Elastic Compute Service (ECS) instances, the
  latest available custom image that is not in the Deprecated state is used. Custom images in the
  Deprecated state are not used.

Notice If you want to roll back a custom image in the image family to a previous version, you can put the latest available custom image into the Deprecated state. However, if the custom image that you want to roll back is the only image in the Available state within the image family, the image family does not contain available custom images that can be used to create instances after you put the custom image into the Deprecated state. Proceed with caution when you put a custom image into the Deprecated state.

• If you want to use the latest custom image that is in the Deprecated state in an image family, you can restore the image and use the image as the latest available custom image.

# Set the state of an image in an image family to Deprecated to roll back the image

- 1.
- 2.
- 3.
- 4. Set the state of the latest image named image2 in the image-test image family to 
  Deprecated.
  - i. In the **Image Family** section of the **Custom Images** tab, enter image—test in the search box for exact match.
  - ii. Click the iii icon and choose More > Deprecate Image in the Actions column that correspondsto image2.
  - iii. Deprecated is displayed in the Status column that corresponds to image2 .
- 5. Log on to OpenAPI Explorer.
- 6. Call the DescribeImageFromFamily operation to query the latest available custom image in the specified image family.

In this example, the latest image in the image-test image family is queried. Specify the following API request parameters:

- RegionId: Specify the region ID of the image family.
- ImageFamily: Specify the image family name. In this example, the name image-test is used.

The following figure shows a sample response. The name of the latest available image is displayed as image1, which indicates that the image in the image family is rolled back.

```
"OSName": "CentOS 7.2 64位",
"IsSupportIoOptimized": true,
"IsSupportCloudinit": true,
"ImageName": "image1",
"DiskDeviceMappings": {
"DiskDeviceMappings": [
```

#### Restore an image in an image family

- 1.
- 2.
- 3.
- 4. In the Image Family section of the Custom Images tab, enter in the search box for exact match.
- 5. Click the icon and choose More > **Restore** in the Actions column that corresponds to image2.
- 6. **Available** is displayed in the Status column that corresponds to image2.

image2 is restored to the latest available custom image. You can call the DescribeImageFromFamily operation to query the latest available custom image in the specified

image family.

7. Call the DescribeImageFromFamily operation to query the latest available custom image in the specified image family.

In this example, the latest image in the <u>image-test</u> image family is queried. Specify the following API request parameters:

- RegionId: Specify the region ID of the image family.
- ImageFamily: Specify the image family name. In this example, the name image-test is used.

The following figure shows a sample response. The name of the latest available custom image is displayed as image2.

```
"OSName": "CentOS 8.0 64\(\frac{1}{2}\)",

"IsSupportIoOptimized": true,

"IsSupportCloudinit": true,

"ImageName": "image2",

"DiskDeviceMappings": {

"DiskDeviceMapping": [
```

#### Related information

• ModifyImageAttribute

•

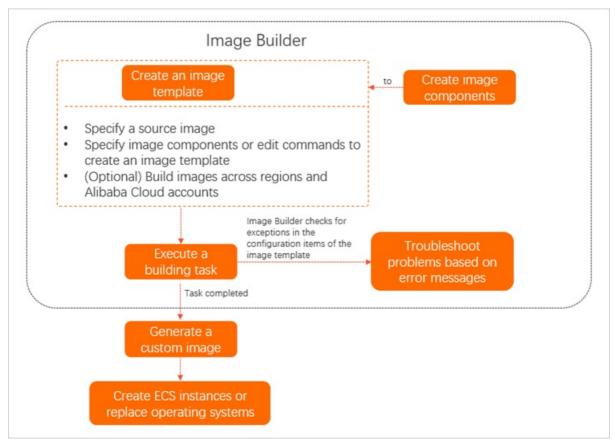
# 10.Image Builder

## 10.1. Overview

Image Builder is an all-in-one image customization service provided by Elastic Compute Service (ECS) and can be used to create images for ECS instances across regions and Alibaba Cloud accounts.

#### Introduction

Image Builder provides the image template feature that allows you to use image templates to create, distribute, and share images. The following figure shows the process of creating an image.



1. (Optional)Create image components.

An image component consists of one or more image template commands. You can reuse image components to create image templates. Image Builder allows you to use image template commands to create images in a standard and transparent manner. Image template commands are compatible with some Dockerfile commands. For more information, see Commands supported by Image Builder and Create an image component.

2. Create an image template.

Image template is a key feature of Image Builder. When you create an image template, you can specify a source image, image components, or image template commands to define image data. You can also specify regions in which to generate the image or specify Alibaba Cloud accounts to which to share the image. For more information, see Create an image template.

3. Execute an image creation task.

When you execute the image template, an image creation task is generated. When the task is executed, Image Builder checks for exceptions in the configuration items of the image template. For more information, see View and execute an image template.

4. View the created custom image.

You can use the created custom image to create ECS instances or replace the operating systems of ECS instances.

#### Billing

You are not charged for Image Builder. When you use an image template to create an image, an intermediate ECS instance is automatically created. You are charged for the intermediate ECS instance on a pay-as-you-go basis. For more information, see Pay-as-you-go.

**Note** You can accept the default configurations of the intermediate ECS instance. For example, an instance type that has a minimum number of vCPUs and a minimum size of memory is selected for the instance by default. You can also modify the configurations of the instance based on your business requirements.

#### **Scenarios**

Scenario	Description
Create images	You can use an image template to create an image based on a public image or an existing custom image.
Maintain server applications	You can use an image template to update applications on your server. You need only to configure the image information of the source server in an image template, add application update commands in the Image Template section, and then run the commands to update the applications.  Each image template that is created to update applications can be stored as an application version, and used to update or roll back applications based on your business requirements.
Manage images	To create or iterate an image in multiple regions or Alibaba Cloud accounts, you can use an image template to define image data and specify the regions to which to distribute the image or Alibaba Cloud accounts to which to share the image. Then, the system generates the image in the specified regions or shares the image to the specified Alibaba Cloud accounts.

For information about examples of how to use Image Builder, see Best practices for Image Builder.

#### **Benefits**

- Efficient
  - To facilitate the process of creating images, Image Builder is compatible with some Dockerfile commands. If you are familiar with Dockerfile commands, you can use Dockerfile to edit your commands and then import the Dockerfile template in the ECS console.

• Image Builder is an all-in-one service that can reduce the amount of time required to manually create ECS instances, perform O&M on ECS instances, and create, update, copy or share images.

#### Secure

When an image creation task is executed, Image Builder checks for exceptions in the parameters of the image template. For example, Image Builder checks whether the resource quotas are sufficient for the intermediate instance and whether image template commands run normally.

# 10.2. Commands supported by Image Builder

This topic describes the Dockerfile commands supported by Image Builder.

The commands supported by Image Builder are compatible with some Dockerfile commands. If you are familiar with Dockerfile commands, you can quickly get started with Image Builder. If you are not familiar with Dockerfile commands, you can learn about the commands supported by Image Builder from this topic. For more information about Dockerfile commands, such as the difference between and commands, visit Dockerfile reference.

The following table describes the commands supported by Image Builder.

Command	Format or example	Description
		Specifies the source image. The command is not supported by image components. After you configure a source image in an image template, the system generates a FROM command.
FROM	<ul><li>FROM IMAGE:<image id=""/></li><li>FROM IMAGE_FAMILY:<image family="" name=""/></li></ul>	Note If you manually configure a FROM command in an image template, the FROM command overwrites the command generated by the system after you configure a source image.
COMPONENT	COMPONENT ic- bp18hy47cqavewsb****	Specifies the image components. You can specify system components or custom components. The command is applicable only to image templates. When you create an image template, you can run the command to specify one or more image components to define the image template in a quick manner.
RESTART	RESTART	Restarts the instance. Do not add content to the command.
RUN	RUN echo hello;\ echo world;	Builds an image. The command can be separated into multiple lines. A line break (\) must be added to the end of each line except the last line.

Command	Format or example	Description
ENV	• ENV key value • ENV key1="value1" key2=" value2"	Configures system variables. The command is in the form of a key-value pair.  Note Double quotation marks (") must be included in the following command format: ENV key1="value1" key2="value2".
WORKDIR	<ul><li>WORKDIR /<path></path></li><li>WORKDIR <path1>/<path2></path2></path1></li></ul>	Configures working directories.
СОРУ	<ul><li>COPY {Network file address} <local directory=""></local></li><li>COPY <path file="" local="" of="" the=""> <local directory=""></local></path></li></ul>	Copies files.
USER	USER <username></username>	Configures users who execute commands.  ? Note The <username> variable must be set to an existing username in the system. Otherwise, an error is reported when the command is run.</username>
LABEL	LABEL user="username"\ date="2020-11-11" key="value"	Defines template metadata. The command can be separated into multiple lines. A line break (\) must be added to the end of each line except the last line.
CMD	<ul><li>CMD ["executable", "param 1", "param2"]</li><li>CMD command param1 param 2</li></ul>	Runs on instance startup.
ENT RYPOINT	<ul><li>ENTRYPOINT ["executable", "param1", "param2"]</li><li>CMD command param1 param</li></ul>	Runs on instance startup.

# 10.3. Use Image Builder

### 10.3.1. Create an image component

Image components store common commands that are used to create an image. The commands include application commands, configuration commands, and test commands. When you create an image template, you can select one or more image components in the Elastic Compute Service (ECS) console to define the image template. This topic describes how to create an image component in the ECS console.

#### Context

- Before you create an image component, take note of the following items:
  - Only custom image components can be created.
  - Image components can be used to create only Linux images.
  - When you create an image component, you must set Component Type to Build.
  - Each command used to create an image component cannot exceed 16 KB in size. A maximum of 127 commands can be configured to create an image component.
- For more information about the commands that the **Image Template** parameter of image components support, see Commands supported by Image Builder.
  - 1.
  - 2.
  - 3.
  - 4. On the Image Builder Management page, click the Image Components tab.

The first time that you use Image Builder, the Image Builder welcome page appears. Click Image Builder Management to go to the Image Builder Management page.

5. Click Create Image Component.

Configure the parameters to create an image component.

- **? Note** You can import a component template by using Dockerfile and preview the template after commands are configured.
- Set Name and Description.
- Set System Type to Linux.
- Set Component Type to Build.
- Set Tag. For more information, see Overview.
- In the Image Template section, set Command and Content.
- 6. Click Create.

After the image component is created, you can view the component on the **Image Components** tab. You can use the image component to create an image template. For more information, see Create an image template.

### 10.3.2. Create an image template

You can use image templates to customize image content and create images in different regions and accounts. This topic describes how to create an image template in the Elastic Compute Service (ECS) console.

#### Context

- Before you create an image template, take note of the following items:
  - Only custom image templates can be created.
  - You can configure only public, custom, or shared Linux images or image families as the source images of the custom image templates.
  - When you use an image template to create an image, you must create an intermediate instance to help create the image. You are charged for the intermediate instance on a pay-as-you-go basis.
     For more information, see Pay-as-you-go.
  - Each command used to create an image template cannot exceed 16 KB in size. A maximum of 127 commands can be specified.
- For information about the commands that can be used in Image Template, see Commands supported by Image Builder.
  - 1.
  - 2.
  - 3.
  - ${\bf 4. \ \ On \, the \, Image \, Builder \, Management \, \, page, \, click \, Create \, \, Image \, Template.}$ 
    - If you access Image Builder for the first time, the **Image Builder** welcome page appears. Click **Create Image Template**.
  - 5. Configure parameters described in the following tables to create an image template. Then, click **Create**.
    - Parameters in the Basic Information section

Parameter	Description
Template Name	The name of the image template. This parameter is required.
Description	The description of the image template.
Tag	The tags of the image template. You can add tags to the image template to manage the image template in a centralized manner. For more information, see Overview.

#### • Parameters in the Configure Source Image section

Parameter	Description
Import Template through Dockerfile	You can import the edited template content by using Dockerfile. The imported template content overwrites the <b>Source Image Type, Image,</b> and <b>Image Template</b> parameters that you specified in the ECS console. Proceed with caution.
Source Image Type	The type of the source image of the image template. Valid values: Images and Image Family.

Parameter	Description
Image or Image Family	This parameter is required and varies based on <b>Source Image Type</b> .
	Note Only Linux images are supported.
	If you set Source Image Type to Images, you can select Custom Image, Public Image, or Shared Image from the first Image drop-down list and then select an image ID or image version from the second Image drop-down list.
	If you set Source Image Type to Image Family, you must select an image family from the Image Family drop-down list.
Image Template	You must configure <b>Command</b> and <b>Content</b> for the image template. For information about supported commands, see the "Context" section in this topic. After you configure the parameters, you can click <b>Preview</b> to view the template content.
	Note You can select Custom Component or System Component from the Command drop-down list and then select the ID of a component from the Content drop-down list.

#### o Parameters in the **Configure Intermediate Instance** section

Parameter	Description
vSwitch	Select a vSwitch that is available in the current region. If you do not configure this parameter, the system creates a temporary virtual private cloud (VPC) and vSwitch. Make sure that the unused quota is sufficient in the region.
Instance Type	Select an instance type for the intermediate instance. If you do not configure this parameter, an instance type that provides the fewest vCPUs and memory resources is selected by default. This configuration is subject to the resource availability of instance types. For example, the ecs.g6.large instance type is selected by default. However, if available ecs.g6.large resources are insufficient, the ecs.g6.xlarge instance type is selected.
System Disk Size	Specify the system disk size for the intermediate instance. Unit: GiB. Valid values: 20 to 500. Default value: 40.
Public Bandwidth (Mbit/s)	Specify the public bandwidth for the intermediate instance.  Default value: 0. This value indicates that the intermediate instance cannot access the Internet. If your image building template content contains a command that requires access to the Internet, you must set a non-zero value for the public bandwidth.  Otherwise, the image building task fails.

Parameter	Description
	You can select whether to retain the intermediate instance.
Release if Image Building Fails	Note If you select No to retain the intermediate instance, you continue to be charged for the instance after the image building task fails. When the intermediate instance is no longer needed, you must manually release it in a timely manner. For more information, see Release an instance.

#### • Parameters in the Configure Image Distribution section

Parameter	Description	
Destination Image Name	Specify a name prefix for the image that you want to create. The system appends the ID of the image building task (ExecutionId) to the prefix to generate a full name of the image in the name prefix>_ <id building="" image="" of="" task="" the=""> format.</id>	
Current Region	The current region of the ECS console.	
Default Distribution Region	If you want to copy the image that you want to create to multiple regions, you can specify the distribution regions. Then, the image is distributed to the specified regions. You can specify up to 20 distribution regions or leave this parameter empty.	
Shared User ID	If you want to share the image that you want to create to other Alibaba Cloud accounts, you can specify the user IDs (UIDs) to share the image. Then, the image is shared to the specified Alibaba Cloud accounts within the same region as the image. You can specify up to 20 UIDs or leave this parameter empty.	
Execute Immediately	You can select whether to immediately execute the image building task. If you select No to not immediately execute the image building task, you can manually execute the task. For more information, see Execute an image building task.	

After the image template is created, you can view the image template on the **Image Builder Management** page.

## 10.3.3. View and execute an image template

After an image template is created, you can view its details. On the image template details page, you can manually execute image building tasks. This topic describes how to view the details of an image template and how to manually execute an image building task.

#### View an image template

- 1.
- 2.
- 3.

- 4. On the Image Templates tab of the Image Builder Management page, find the image template that you want to view. Then, click the image template ID in the Template Name column.
  - On the **Det ail** page, you can perform the following operations:
  - In the Basic Information section, view the basic information of the image template.
  - On the **Building History** tab, view historical image building tasks or execute image building tasks.
  - On the View Template tab, view the image template.

#### Execute an image building task

You can execute only a single image building task based on an image template at a time. You can cancel multiple image building tasks at the same time. When some image building tasks are canceled, other image building tasks are not affected.

- 1.
- 2.
- 3.
- 4. On the Image Templates tab of the Image Builder Management page, find the image template based on which you want to execute an image building task. Then, click the image template ID in the Template Name column.
- 5. On the **Building History** tab, click **Execute**.
  - To cancel an image building task that is in the Building state, click **Cancel** in the Actions column.
  - If the image building task is executed, the building task enters the Success state and an image ID is generated.
    - Note If Failed is displayed in the Status column corresponding to an image building task, you can view the details in the Execution Result column to identify and troubleshoot the issue.

### 10.3.4. Delete an image component

You cannot modify image components. If an exception occurs on a created image component, you can delete the component and create a new one. You can also delete the image components that are no longer needed. This topic describes how to create an image component in the Elastic Compute Service (ECS) console.

#### Context

Before you start the procedure, take note of the following items:

- Only custom image components can be deleted.
- After an image component is deleted, the image templates that have used the component are not affected.

#### **Procedure**

- 1.
- 2.
- 3.

- 4. On the Image Builder Management page, click the Image Components tab.
- 5. Find the image component that you want to delete. Click Delete in the Actions column.
- 6. Verify the name and ID of the image component that you want to delete. Then, click OK.

### 10.3.5. Delete an image template

You cannot modify image templates. If an exception occurs on a created image template, you can delete the template and create a new one. You can also delete the image templates that are no longer needed. This topic describes how to delete an image template in the Elastic Compute Service (ECS) console.

#### **Prerequisites**

Make sure that all image building tasks that are based on the image template are in the Success, Failed, or Canceled state. If an image building task that is based on the image template is in the Building state, you cannot delete the template. For more information about how to view the status of an image building task, see View an image template.

#### Context

After an image template is deleted, the status of custom images created from the image template are not affected.

#### **Procedure**

- 1.
- 2.
- 3
- 4. On the **Image Builder Management** page, find the image template that you want to delete. Then, click **Delete** in the **Actions** column.

# 10.4. Best practices for Image Builder

This topic describes how to create a custom image for the NGINX development environment by using Image Builder and how to create an Elastic Compute Service (ECS) instance from the custom image and test the environment.

#### Context

Image Builder is a service provided by Alibaba Cloud that integrates image customization, creation, maintenance, verification, sharing, and distribution. You can use Image Builder to create images and update or roll back applications within the images. For more information, see Overview.

In the examples, the following configurations are used:

- Region: China (Beijing)
- Public image: Cent OS 8.2 64-bit
- Image content in the image template: NGINX development environment

An image is created based on the preceding configurations, and then an instance is created from the image and is used to test whether the NGINX environment works normally.

#### Step 1: Create and execute an image template

1.

2.

3.

4. On the Image Builder Management page, click Create Image Template.

The first time you access Image Builder, the Image Builder welcome page appears. Click Create Image Template.

5. On the Create Image Template page, configure parameters and click Create.

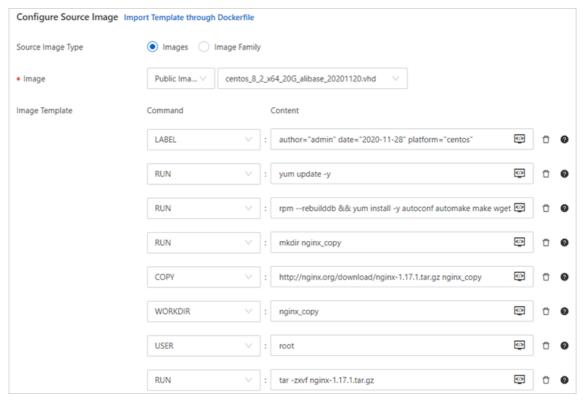
This section describes the parameters in the **Configure Source Image** section. For information about how to configure other parameters, see **Create an image template**.

- In the Basic Information section, configure the Template Name, Description, and Tag parameters.
- In the **Configure Source Image** section, click **Import Template through Dockerfile**. In the dialog box that appears, enter text in code editor and click **Parse**.

In this example, enter the following data in the code editor:

```
FROM IMAGE:centos_8_2_x64_20G_alibase_20201120.vhd
LABEL author="admin" date="2020-11-28" platform="centos"
RUN yum update -y
RUN rpm --rebuilddb && yum install -y autoconf automake make wget net-tools zlib zlib
-devel make gcc openssl-devel pcre pcre-devel tar
RUN mkdir nginx copy
COPY http://nginx.org/download/nginx-1.17.1.tar.gz nginx copy
WORKDIR nginx copy
USER root
RUN tar -zxvf nginx-1.17.1.tar.gz
ENV nginxversion nginx-1.17.1
WORKDIR $nginxversion
RUN ./configure --prefix=/usr/local/nginx && make && make install
WORKDIR /usr/local/nginx
ENTRYPOINT ["sbin/nginx","-c"]
CMD conf/nginx.conf
```

After the entered data is parsed, the **Image** and **Image Template** parameters are automatically configured, as shown in the following figure.



- In the **Configure Intermediate Instance** section, configure the Instance Type, System Disk Size, Public Bandwidth (Mbit/s) parameters.
  - Note By default, Public Bandwidth (Mbit/s) is set to 0. This value indicates that the intermediate instance cannot access the Internet. If your image content contains a command that requires access to the Internet, you must set Public Bandwidth to a non-zero value. Otherwise, the image creation task fails.
- In the **Configure Image Distribution** section, you can configure the created image to be distributed to specified regions or to be shared to specified Alibaba Cloud accounts.

In this example, the following configurations are used:

- **Destination Image Name**: Specify the name prefix of the image to be created. In this example, a name prefix of test-image is used. The system appends the ID of the image creation task to the prefix to generate a name for the image.
- Current Region: is set to Default distribution regions include China (Beijing).
- Default Distribution Region: Enter China (Hangzhou).
- Shared User ID: Specify an Alibaba Cloud account.
- Execute Immediately: Select No.
- 6. On the Image Templates tab of the Image Builder Management page, find the created image template. Click the template ID in the Template Name column to go to the Detail page of the image template.
- 7. On the **Building History** tab, click **Execute**.

After the image creation task is executed, the o icon is displayed in the Status column, and the corresponding image ID is generated in the Image ID column.

Note If Failed is displayed in the Status column, you can view the details of Execution Result to identify and troubleshoot the error.

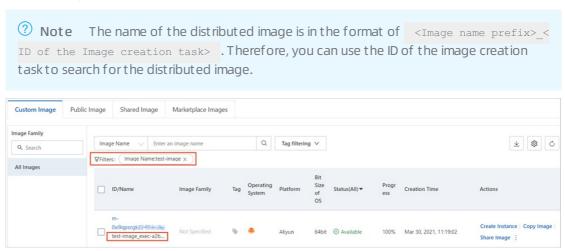
#### Step 2: View the created image

- 1. On the **Building History** tab of the **Detail** page of the image template, click the image ID to go to the details page of the created image.
- 2. Click **Object to Share** to view the Alibaba Cloud accounts to which the created image is shared.
- 3. Check the distribution of the created image.

i.

- ii. In the top navigation bar, select the region to which the image is distributed.In this example, select China (Hangzhou).
- iii. On the **Custom Image** tab, select **Image Name** from the drop-down list and enter the name of the distributed image in the search box to search for the image.

Enter the test-image name prefix. If the image is distributed to the selected region, you can find the image in the search result.



# Step 3: Create an instance from the image and test the NGINX environment

1

- 2. In the top navigation bar, select the region in which the image template is created. In this example, the image is created in the **China (Beijing)** region.
- 3. On the **Image Templates** tab, find the created image template and click the template ID to go to the **Det ail** page of the image template.
- 4. On the Building History tab, click an image ID to go to the details page of the image.
- 5. Click **Create Instance** in the Actions column corresponding to the image to create an instance. Create an instance based on your business needs. Take note of the following configuration items when you create the instance. For more information, see **Create an ECS instance** by using a custom

#### image.

- o Select Assign Public IPv4 Address in the Public IP Address section in the Networking step.
- Configure an inbound rule in a security group to allow traffic on port 80.
- 6. Obtain the public IP address of the instance.
  - i. After the instance is created, go back to the Instances page.
  - ii. Find the created instance and copy the public IP address in the IP Address column.
- 7. Open a browser on your computer, enter the  $\ensuremath{\texttt{<Public IP Address of the ECS instance>}}\$  in the address bar, and then press the Enter key.

If the following page appears, the NGINX development environment works normally.

### Welcome to nginx!

If you see this page, the nginx web server is successfully installed and working. Further configuration is required.

For online documentation and support please refer to <u>nginx.org</u>. Commercial support is available at <u>nginx.com</u>.

Thank you for using nginx.

# 11.Change the operating system

You can change the operating systems of instances, such as from Linux to Windows or from Ubuntu to CentOS. In regions outside mainland China, you can only change the operating system of an ECS instance between Linux editions or between Windows editions. You cannot change the operating system between Linux and Windows.

You can change the operating system of an instance by changing its system disk.

- Change the image of the system disk to a non-public image. For more information, see Replace the operating system of an instance by using a non-public image.
- Change the image of the system disk to a public image. For more information, see Replace the operating system of an instance by using a public image.

# 12.End of support for operating systems

Alibaba Cloud may stop providing technical support for services for a variety of reasons, such as end of service lifecycle, end of third-party support, or evolution of open source projects. We recommend that you stay up-to-date on technical support plans so that you can make informed and timely decisions on whether to update your services or software to general availability (GA) versions. When you use Elastic Compute Service (ECS), pay attention to the lifecycle plans of the operating systems contained in images that you used to create ECS instances. You can visit the official websites of the operating systems to obtain information about their lifecycle plans.

**? Note** When an operating system version reaches its end of life (EOL), Alibaba Cloud stops providing support for ECS instances that run the operating system version.

- Alibaba Cloud Linux 2 is an operating system provided by Alibaba Cloud. Alibaba Cloud provides technical support for issues that you encounter when you use Alibaba Cloud Linux 2 up until the operating system reaches its EOL.
- If you are using a third-party operating system, contact the operating system vendor or the open source community for technical support. Alibaba Cloud also provides information about image- and system-related issues. If a third-party vendor stops supporting or maintaining its operating system, Alibaba Cloud also stops providing corresponding technical assistance.

### **EOL plans of Alibaba Cloud Linux**

Version	End-of-support date
Alibaba Cloud Linux 3	2031-04-30
Alibaba Cloud Linux 2	2024-03-31

### EOL plans of third-party operating systems

The following tables describe the EOL plans of some third-party operating system versions.

#### Anolis OS

Version	End-of-support date
Anolis OS 8	2031-06-30
Anolis OS 7	2024-06-30

#### Cent OS

Version	End-of-update date	End-of-maintenance date
Cent OS Stream 8	2024-05-31	2024-05-31

Version	End-of-update date	End-of-maintenance date
Cent OS 8	2021-12-31	2021-12-31
Cent OS 7	2020-08-06	2024-06-30
CentOS 6	2017-05-10	2020-11-30
Cent OS 5	2007-04-12	2017-03-31

Note For information about what to do after CentOS Linux reaches EOL, see Options for dealing with CentOS Linux EOL.

### • Debian

Version	Codename	Release date	End-of- support date	EOL LTS	EOL ELTS
Debian 13	Trixie	To be decided	To be decided	To be decided	To be decided
Debian 12	Bookworm	To be decided	To be decided	To be decided	To be decided
Debian 11	Bullseye	2021-08-14	2024-07	2026-06	To be decided
Debian 10	Buster	2019-07-06	2022-07	2024-06	To be decided
Debian 9	Stretch	2017-06-17	2020-07-06	2022-06-30	To be decided
Debian 8	Jessie	2015-04-25	2018-06-17	2020-06-30	2022-06-30
Debian 7	Wheezy	2013-05-04	2016-04-25	2018-05-31	2020-06-30
Debian 6.0	Squeeze	2011-02-06	2014-05-31	2016-02-29	To be decided
Debian 5.0	Lenny	2009-02-14	2012-02-06	To be decided	To be decided

### • Ubuntu

Version	End-of-support date
Ubuntu 22.04 LTS	2027-04
Ubuntu 20.04 LTS	2030-04
Ubuntu 18.04 LTS	2028-04

Version	End-of-support date		
	2026-04		
Ubuntu 16.04 LTS	Note Ubuntu 16.04 LTS has reached the end of standard support. You can still obtain security updates with Extended Security Maintenance. For more information, see Security updates after the end of support and maintenance for Ubuntu 16.04 LTS.		
Ubuntu 14.04 LTS	2024-04		

### • Red Hat

Version	Release date	End date of Full Support <sup>①</sup>	End date of Maintenanc e Support 1 <sup>2</sup>	End date of Maintenanc e Support 2 <sup>3</sup>	End date of Extended Life <sup>@</sup>	Last minor release
Red Hat 4	2005-02-14	2009-03-31	2011-02-16	2012-02-29	2017-03-31	Red Hat 4.9
Red Hat 5	2007-03-15	2013-01-08	2014-01-31	2017-03-31	2020-11-30	Red Hat 5.11
Red Hat 6	2010-11-10	2016-05-10	2017-05-10	2020-11-30	2024-06-30	Red Hat 6.10
Red Hat 7	2014-06-10	2019-08-06	2020-08-06	2024-06-30	N/A	Red Hat 7.9
Red Hat 7 for ARM	2017-11-13	2020-11-30	N/A	N/A	N/A	Red Hat 7.9
Red Hat 7 (POWER9)	2017-11-13	2020-11-30	N/A	N/A	N/A	Red Hat 7.9
Red Hat 7 for System z (Structure A)	2018-04-10	2020-11-30	N/A	N/A	N/A	Red Hat 7.9
Red Hat 8	2019-05	2024-05	N/A	2029-05	To be decided	Red Hat 8.8

### Description:

 $\circ$   $^{\scriptsize \textcircled{1}}$ : During the Full Support phase, Red Hat provides updated installation images for minor releases.

- o <sup>2</sup>: During the Maintenance Support 1 phase, Red Hat may provide at their discretion hardware enablement in minor releases that do not require substantial software changes. No new software features are made available during this phase. During this phase, Red Hat provides at their discretion updated installation images for minor releases only when required due to installer changes.
- <sup>③</sup>: No features or hardware enablement is released during the Maintenance Support phase of Red Hat 8 or the Maintenance Support 2 phase of Red Hat 6 and Red Hat 7. Red Hat may provide minor releases with updated installation images during this phase.
- <sup>(4)</sup>: Red Hat provides limited technical support for products in the Extended Life phase. During this phase, Red Hat does not provide bugfixes, security fixes, hardware enablement, or root cause analyses but provides support only for existing installations.

For more information, see Red Hat Enterprise Linux Life Cycle.

#### FreeBSD

Version	Release date	End-of-maintenance date
FreeBSD 13.0	2021-04-13	3 months after FreeBSD 13.1 release
FreeBSD 12.3	2021-12-07	3 months after FreeBSD 12.4 release
FreeBSD 12.2	2020-10-27	2022-03-31
FreeBSD 12.1	2019-11-04	2021-01-31
FreeBSD 12.0	2018-12-11	2020-02-04
FreeBSD 11.4	2020-06-16	2021-09-30
FreeBSD 11.3	2019-07-09	2020-09-30
FreeBSD 11.2	2018-06-28	2019-10-31
FreeBSD 11.1	2017-07-26	2018-09-30
FreeBSD 11.0	2016-10-10	2017-11-30

### OpenSUSE

Version	End-of-maintenance date
OpenSUSE 15.3	2022-11-30
OpenSUSE 15.2	2022-01-04
OpenSUSE 15.1	2021-02-02
OpenSUSE 15.0	2019-12-03
OpenSUSE 42.3	2019-07-01

Version	End-of-maintenance date
OpenSUSE 42.2	2018-01-26
OpenSUSE 42.1	2017-05-17

### • SUSE Linux

Version	Release date	End data of general support	End date of Long Term Service Pack Support (LTSS)
SUSE Linux Enterprise Server 12	2014-10-27	2016-06-30	2019-07-01
SUSE Linux Enterprise Server 12 SP1	2015-12-15	2017-05-31	2020-05-31
SUSE Linux Enterprise Server 12 SP2	2016-11-08	2018-03-31	2021-03-31
SUSE Linux Enterprise Server 12 SP3	2017-09-07	2019-06-30	2022-06-30
SUSE Linux Enterprise Server 12 SP4	2018-12-12	2020-06-30	2023-06-30
SUSE Linux Enterprise Server 12 SP5	2019-12-09	2024-10-31	2027-10-31
SUSE Linux Enterprise Server 15	2018-07-16	2019-12-31	2022-12-31
SUSE Linux Enterprise Server 15 SP1	2019-06-24	2021-01-31	2024-06-31
SUSE Linux Enterprise Server 15 SP2	2020-07-21	2021-12-31	2024-12-31
SUSE Linux Enterprise Server 15 SP3	2021-06-22	6 months after SUSE Linux Enterprise Server 15 SP4 release	To be decided

### • Fedora

Version	End-of-maintenance date
Fedora 36	2023-05-16
Fedora 35	2022-11-15
Fedora 34	2022-05-17
Fedora 33	2021-11-30

Version	End-of-maintenance date
Fedora 32	2021-05-25
Fedora 31	2020-11-24
Fedora 30	2020-05-26
Fedora 29	2019-11-26
Fedora 28	2019-05-28
Fedora 27	2018-11-30

### • Rocky Linux

Version	Release date	End-of-maintenance date
Rocky Linux 8.5	2021-11-14	2031-11-14

### • AlmaLinux

Version	Release date	End-of-maintenance date
AlmaLinux 8.5	2021-11	2031-11

### • Fedora CoreOS

To be decided

### • Windows Server

Version	End-of-update date	End-of-maintenance date
Windows Server 2022	2026-10-13	2031-10-14
Windows Server 2019 Datacenter	2024-01-09	2029-01-09
Windows Server 2016 Datacenter	2022-01-11	2027-01-12
Windows Server 2012 Datacenter	2018-10-09	2023-10-10
Windows Server 2012 R2 Datacenter	2018-10-19	2023-10-10
Windows Server 2008	Stopped	Stopped
Windows Server 2003	Stopped	Stopped

For more information about EOL plans of third-party operating systems, visit the official websites.

- Cent OS
- Debian
- Ubuntu

- Red Hat
- FreeBSD
- OpenSUSE
- SUSE Linux
- Fedora
- Rocky Linux
- AlmaLinux
- Fedora CoreOS
- CoreOS
- Microsoft Windows Server

Elastic Compute Service Images·FAQ

### 13.FAQ 13.1. Image FAQ

This topic provides answers to frequently asked questions about Elastic Compute Service (ECS) images.

### Common FAQ

- Can I replace the selected image of an ECS instance?
- Do the system disks of ECS instances support Key Management Service (KMS) encryption? How do I use KMS encryption based on Terraform or Packer?
- What are the differences between snapshots and images? How are snapshots and images related?
- Which instance families do Red Hat Enterprise Linux (RHEL) images support?
- Do RHEL images provided by Alibaba Cloud include the official Red Hat subscriptions?
- How do I obtain technical support if I encounter issues when I use the RHEL operating system?
- How do Windows versions differ in terms of memory limits?

#### • FAQ about custom images

- Can I use a snapshot of a data disk to create a custom image?
- How do I view dat a disk usage?
- How do I unmount file systems and delete disk table data?
- How do I confirm that a data disk has been unmounted and that a new custom image can be created?
- Does a custom image still exist after the instance from which the image was created is released?
- When an instance expires or its data is deleted, are custom images that were created from the instance affected? Are instances created from the custom images affected?
- Can I replace the operating system of an instance created from a custom image? Can the custom image still be used after the operating system is replaced?
- Can I select a custom image with a different operating system when I replace the system disk of an instance?
- Can I use a custom image to overwrite the system disk data of an ECS instance?
- Can I upgrade the CPU, memory, bandwidth, and hard disks of an ECS instance that was created from a custom image?
- Can I use a cust om image across regions?
- Can a custom image created from a subscription instance be used to create a pay-as-you-go instance?
- I created an ECS instance from a custom image and specified a system disk capacity larger than that in the image. However, the system disk capacity of the new ECS instance is the same as that in the image. What do I do?
- Why must I comment out mount entries when I create a custom image or an instance?
- How do I configure and use a private Docker image registry?
- How do I clone an ECS instance?
- Some custom images cannot be used to create I/O optimized instances. What do I do?
- Where do I view the progress of an image import task? How long does it take to import an image?
- Where do I view the progress of an image creation task? How long does it take to create an image?

Images· FAQ Elastic Compute Service

### • FAQ about copying images

- When do I need to copy a custom image?
- Which images can be copied?
- Which regions support copying custom images?
- How long does it take to copy a custom image?
- How am I charged when I copy a custom image?
- What limits apply to the original image (image copied) and the new image (image copy) during an image copy process?
- How do I copy images within my Alibaba Cloud account to other regions within other Alibaba Cloud accounts?
- Can I copy a cust om image derived from an Alibaba Cloud Market place image across regions?
- How do I migrate data from regions outside the Chinese mainland to regions inside the Chinese mainland?

### • FAQ about sharing images

- How many images can be shared to me?
- How many users can an image be shared to?
- I have accounts on different Alibaba Cloud sites. Can I share images between these accounts?
- Do shared images count against my image quota?
- Do geographical limits apply when shared images are used to create instances?
- What are the risks of creating an instance from a shared image?
- What risks may arise if I share a custom image to other accounts?
- After an account shares an image to me, can I share this image to another account?
- After I share an image, can I still use this image to create an instance?
- Can an image created from Instance A in one region be used by Instance B in a different region?

### • FAQ about importing images

- Is Bring Your Own License (BYOL) supported when I import custom images?
- What types of licenses can I use when I import custom images?
- How are the imported BYOL images charged?
- How are BYOL licenses authenticated and subscribed in Alibaba Cloud when their subscriptions expire?

### • FAQ about exporting images

- I want to export an image to my computer for testing. What do I do?
- Can an image exported from ECS be started on other platforms?

### • FAQ about deleting images

- Can I delete a custom image after it has been used to create an ECS instance?
- Can I delete a custom image from my account after the image is shared to another account?
- What happens if I unshare Custom Image M to Account A?
- Why am I prompted with a message similar to "The specified image cannot be deleted because it is associated with instances" when I attempt to delete an image?
- FAQ about replacing images or operating systems

- When I replace a system disk, can I select an image that contains data disk snapshots?
- I want to replace the operating system of my ECS instance by using an existing image. What do I do?
- Can I use an image created from an instance in Account A to replace a system disk in Account B?

#### FAQ about image pricing

• I am creating an ECS instance. Why is the total instance cost displayed when I select a custom image higher than that displayed when I select a public image?

### • FAQ about the commercial availability of images

- What features do Alibaba Cloud Market place images provide?
- What are the benefits of Alibaba Cloud Market place images?
- What server environments and scenarios do Alibaba Cloud Marketplace images support?
- Are Alibaba Cloud Market place images safe?
- What do I do if I encounter an issue when I am installing or using an Alibaba Cloud Marketplace image?
- How do I purchase an Alibaba Cloud Market place image?
- How long can I use a purchased image?
- Can I request refunds for my Red Hat images or SUSE images that are no longer needed?
- Are Alibaba Cloud Market place images refundable?
- Are free Alibaba Cloud Marketplace images still available after Alibaba Cloud Marketplace images are commercially available?
- I bought an Alibaba Cloud Marketplace image in the China (Hangzhou) region. Can I use it to create an ECS instance or replace a system disk in the China (Beijing) region?
- I have an instance created from an Alibaba Cloud Market place image. Do I need to make further payments for the image when I renew the instance or upgrade the configurations of the instance?
- I have an ECS instance created from an Alibaba Cloud Market place image. After the instance is released, can I continue to use that image free of charge when I purchase a new ECS instance?
- I created an ECS instance from an Alibaba Cloud Market place image and then created a custom image from the instance. Do I need to pay for the custom image when I use it to create an ECS instance?
- If I copy an Alibaba Cloud Marketplace image that I bought to another region to create an ECS instance, do I need to pay for the image?
- I created an ECS instance from an Alibaba Cloud Market place image and then created a custom image from that instance. If I share the custom image to Account B, does Account B need to pay for the custom image when it uses this image to create an ECS instance?
- Is a fee charged if I replace a system disk by using an Alibaba Cloud Market place image or an image derived from an Alibaba Cloud Market place image?
- My ECS instance is using an Alibaba Cloud Marketplace image. Am I charged if I replace the system disk of the instance?
- How do I call an ECS API operation to use an Alibaba Cloud Market place image or a custom or shared image derived from an Alibaba Cloud Market place image to create an ECS instance or replace a system disk?

- If I do not purchase an Alibaba Cloud Market place image or an image derived from an Alibaba Cloud Market place image, is an error reported when I call an ECS API operation to use the image to create an ECS instance or replace a system disk?
- I have configured a scaling group for which the minimum number of instances is set to 10 and the maximum number of instances is set to 100. What do I do with Alibaba Cloud Market place images to ensure that ECS instances are created to meet my computing needs?
- Can I purchase multiple Alibaba Cloud Market place images at a time?
- If an image (such as jxsc000010 or jxsc000019) in use in an active scaling configuration no longer exists in Alibaba Cloud Marketplace, what do I do to ensure that ECS instances can continue to be created based on the scaling configuration in the corresponding scaling group?
- Can one product code support images in different regions?
- I bought 100 images that have the same product code. Can I use the images within all regions?
- After I select I/O Optimized, I cannot select Alibaba Cloud Market place images when I purchase an ECS instance. What is the cause and how do I resolve this issue?

#### • FAQ about subscription Alibaba Cloud Market place images

- What are yearly, monthly, and weekly subscription Alibaba Cloud Market place images?
- On which ECS instances can subscription images be used?
- How do I purchase a subscription image? Can I separately purchase it?
- How do I pay for subscription images?
- Can I use a subscription image after it expires? How do I continue to use it?
- After I purchase a subscription image, can I request a refund if I no longer want to use it?
- What can I expect when a refund is made?
- Can a subscription image be converted into a pay-as-you-go image?
- Can I replace a subscription image with an image of another type or vice versa? How is the fee calculated?
- Where do I view and manage the subscription images that I purchased?
- Is a fee charged for a custom image derived from a subscription image? How is the custom image affected if the subscription image expires?

### • FAQ about ECS instances and operating system images

- Will CentOS 8 public images be unpublished after CentOS 8 reaches its end of life (EOL)?
- Will my instances that were created from CentOS 8 images be affected after CentOS 8 reaches its EOL?
- Which operating systems are not supported by instance families that use AMD EPYCTM MILAN processors?
- In ECS Bare Metal instances, the kdump service cannot generate crash dump files. How do I resolve this issue?
- A softlockup occurs during the writeback process of the Linux operating system kernel. How do I resolve this issue?
- A softlockup error occurs when I delete cgroups in an ECS instance. What do I do?
- An ECS instance goes down. What do I do?
- Why am I unable to use an SSH key pair that uses the ssh-rsa signature algorithm to connect to an instance that runs Fedora 33 64-bit?
- How do I install patches and compile the kernel on FreeBSD?

Elastic Compute Service Images-FAQ

• Why does the load average become high on ECS instances that run Ubuntu operating systems of specific versions after the Server Guard (AliYunDun) process is started on the instances?

- Why am I unable to select a Windows operating system for ECS instances?
- Does Alibaba Cloud support Windows Server 2008 and Windows Server 2008 R2?
- The operating system of my instance is Windows Server. I am prompted with a message indicating that the operating system is not genuine. What do I do?
- Am I charged for the images used by ECS instances?
- Can I install or upgrade the operating systems of ECS instances on my own?
- Do operating systems have graphical interfaces?
- How do I choose an operating system?
- Do public images come with the FTP service?
- Which SUSE versions do Alibaba Cloud public images support?
- What service support is available for SUSE operating systems?
- If an image was manually created from an ECS instance, can I retrieve instance data after the instance is released on expiration?
- I have an ECS instance and I want to create another ECS instance from the image of this existing ECS instance. What do I do?
- I have purchased an ECS instance. How do I restore my shared image to the newly purchased instance?
- I have multiple Alibaba Cloud accounts. I want to transfer an instance from Account A to Account B or migrate the environment and applications of an instance within Account A to an instance within Account B. What do I do?
- How do I migrate data between ECS instances?
- Can ECS instances in different VPCs communicate with each other?
- How do I handle a Cent OS DNS resolution timeout?
- Why does ECS disable virtual memory and leave swap partitions unconfigured by default?
- How do I enable the kdump service in a public image?
- How do I obtain the dump file for RHEL images?
- How do I enable the kdump service on an instance that uses a Red Hat 8.1 image?
- Red Hat 8.1 and Red Hat 8.2 images start slowly on ECS bare metal instances. How do I solve this issue?
- After I use a CentOS 7 or Windows ECS instance for an extended period of time without restarting it, the instance is disconnected from the network, the network is no longer available, or the public or private IP address of the instance cannot be pinged. What do I do?
- The "UNEXPECTED INCONSISTENCY; RUN fsck MANUALLY." error is reported when an ECS instance starts. What do I do?
- How do Lupgrade RHEL 7 to RHEL 8?
- How do I enable a program to distinguish ECS instances when I deploy my business?
- After I used a Fedora CoreOS image to create an instance of a specific instance family, I found that the number of CPUs in the Iscpu command output is only half the number of vCPUs of the selected instance type. Why?

### Can I replace the selected image of an ECS instance?

Yes, you can replace the image of your ECS instance by selecting Replace Operating System in the ECS console. Be aware that if you replace the image of an instance, data stored on the system disk of the instance is lost. We recommend that you back up your data before you replace the image. For more information, see Change the operating system.

# Do the system disks of ECS instances support Key Management Service (KMS) encryption? How do I use KMS encryption based on Terraform or Packer?

- Yes, the system disks of ECS instances can be encrypted by using your own keys (BYOK) and customer master keys (CMKs) stored in KMS. For more information, see 加密概述.
- Packer-based encryption will be supported soon.
- In Terraform, you can set the encrypted parameter to enable or disable KMS encryption. For more information, see <a href="mailto:alicloud\_disks">alicloud\_disks</a>.

## What are the differences between snapshots and images? How are snapshots and images related?

Images and snapshots differ in the following ways:

- Images can be used to create ECS instances, whereas snapshots cannot.
- A snapshot can be a data backup of the system disk or a data disk of an ECS instance, whereas an image must contain the system disk data of an ECS instance.
- Snapshots can be used only to restore data of disks on existing instances, whereas images can be used to create instances or replace the system disks of instances.
- Images and snapshots suit different scenarios. The following section describes some scenarios for which snapshots and custom images are suitable:

#### Snapshots:

- Back up data on a regular basis. You can use automatic snapshot policies to automatically create snapshots to back up data on a daily, weekly, or monthly basis.
- o Temporarily back up dat a. Examples:
  - You can manually create a snapshot to back up the system data before a temporary system change such as system update or application release.
  - You can create a snapshot to back up data before you resize a system disk.
  - To migrate data from a disk, you can create a snapshot for the disk and use the snapshot to create a disk.

### Cust om images:

- Back up systems that do not change in a short term, such as released or updated application systems.
- Create ECS instances. For example, you can use a custom image to create an ECS instance that hosts multiple applications.
- Migrate systems and data. For example, you can migrate ECS instances from the classic network to virtual private clouds (VPCs).
- o Restore systems across regions and zones.

Snapshots and images have the following relationships:

- When you create a custom image from an instance, ECS creates a snapshot for each disk of the instance. The created custom image consists of the snapshots of all the disks of the instance. For more information, see Create a custom image from an instance.
- You can also create custom images from system disk snapshots. For more information, see Create a custom image from a snapshot.

## Which instance families do Red Hat Enterprise Linux (RHEL) images support?

RHEL images support the following instance families. For more information, see Instance family.

- ecs.r6 (supported only by RHEL 7.7 and later)
- ecs.c6 (supported only by RHEL 7.7 and later)
- ecs.g6 (supported only by RHEL 7.7 and later)
- ecs.r5
- ecs.c5
- ecs.q5
- ecs.re4
- ecs.t5

We recommend that you do not use RHEL images for the following instance types. This is because these instance types have less than 1.5 GiB of memory and cannot meet the minimum requirements of RHEL images.

- o ecs.t5-lc2m1.nano
- o ecs.t5-lc1m1.small
- ecs.hfc5
- ecs.hfg5
- ecs.i2
- ecs.sn1ne
- ecs.sn2ne
- ecs.se1ne
- ecs.sn1
- ecs.sn2
- ecs.se1

For more information, see the following topics:

- RHEL certification
- Which Instance Types are RHEL Guest supported with on Aliyun
- Red Hat Enterprise Linux technology capabilities and limits

## Do RHEL images provided by Alibaba Cloud include the official Red Hat subscriptions?

RHEL images provided by Alibaba Cloud include the official Red Hat subscriptions. When you purchase RHEL images, you can obtain Red Hat updates from the update source provided by Alibaba Cloud. Unlike RHEL images purchased from other channels, RHEL images provided by Alibaba Cloud eliminate the need for a separate account to obtain updates from the update source provided by Red Hat.

### How do I obtain technical support if I encounter issues when I use the RHEL operating system?

Traditionally, you have to log on to the Red Hat system and submit an after-sales request for technical support. Now, you can submit a ticket for technical support. Alibaba Cloud after-sales engineers will help you resolve the issues. If Alibaba Cloud is unable to resolve the issues related to the RHEL operating system, Alibaba Cloud will submit the issues to Red Hat for assistance.

### How do Windows versions differ in terms of memory limits?

Different Windows versions have different upper limits on memory. For more information, see Memory Limits for Windows and Windows Server Releases. Choose appropriate instance types based on the memory limits of your Windows version.

### Can I use a snapshot of a data disk to create a custom image?

No, you cannot use data disk snapshots to create custom images. Only system disk snapshots can be used to create custom images.

However, you can add a data disk snapshot when you use a system disk snapshot to create a custom image. For more information, see Create a custom image from a snapshot.

### How do I view data disk usage?

You can run the **df** command to view data disk usage and the locations where file systems are mounted. Example: **df** -**lh**.

You can run the fdisk command to view the partition information of a data disk. Example: fdisk -l.

### How do I unmount file systems and delete disk table data?

If /dev/hda5 is mounted to /mnt/hda5, you can run one of the following commands to unmount the file system:

```
umount /dev/hda5

umount /mnt/hda5

umount /dev/hda5 /mnt/hda5
```

/etc/fstab is an important configuration file in Linux systems. It contains details about the file systems and storage devices that are automatically mounted on system startup.

If you do not want a partition to be automatically mounted on instance startup, you must delete the corresponding statement from the /etc/fstab file. For example, after the following statement is deleted from the /etc/fstab file, xvdb1 is not mounted on system startup:

```
/dev/xvdb1 /leejd ext4 defaults 0 0
```

The following table describes other important configuration files in Linux systems.

Configuration file Description Risk of modifying the file	
---	--

Configuration file	Description	Risk of modifying the file
/etc/issue*, /etc/*-release, /etc/*_version	The system distribution configuration file	Modifications made to /etc/issue* cause failures to recognize system distributions and to create the system.
/boot/grub/menu.lst, /boot/grub/grub.conf	The system boot configuration file	Modifications made to /boot/grub/menu.lst cause kernel load and system boot failures.
/etc/fstab	The configuration file for mounting partitions on system startup	Modifications made to /etc/fstab cause partition load and system boot failures.
/etc/shadow	The system password- related configuration file	Changes of /etc/shadow to read-only cause failures to modify password files and to create the system.
/etc/selinux/config	The system security policy configuration file	Modifications made to /etc/selinux/config to enable SELinux cause system boot failures.

### How do I confirm that a data disk has been unmounted and that a new custom image can be created?

- 1. Confirm that the statement used to automatically mount the partitions of the data disk has been deleted from the /etc/fstab file.
- 2. Run the **mount** command to view the mount information of all devices. Confirm that the information of the corresponding data disk partitions is not displayed in the command output.

## Does a custom image still exist after the instance from which the image was created is released?

Yes, the custom image still exists after the instance from which the image was created is released. After an instance is manually or automatically released, the custom images created from the instance are not affected.

### When an instance expires or its data is deleted, are custom images that were created from the instance affected? Are instances created from the custom images affected?

No, the custom images and the instances created from the custom images are not affected. After an instance is manually released or is automatically released when it expires, the custom images created from the instance and the instances created from the custom images are not affected.

# Can I replace the operating system of an instance created from a custom image? Can the custom image still be used after the operating system is replaced?

Yes, you can replace the operating system of an instance created from a custom image. The custom image can still be used after the operating system is replaced.

## Can I select a custom image with a different operating system when I replace the system disk of an instance?

Yes, you can select a custom image with a different operating system when you replace the system disk of an instance. For more information, see Replace the operating system of an instance by using a non-public image.

**Note** When a custom image is used to replace a system disk, all data on the original system disk is overwritten.

### Can I use a custom image to overwrite the system disk data of an ECS instance?

Yes, you can use a custom image to overwrite the system disk data of an ECS instance. For more information, see Replace the operating system of an instance by using a non-public image.

**Note** When you use a custom image to replace the image of an instance, the custom image overwrites all data on the system disk of the instance.

## Can I upgrade the CPU, memory, bandwidth, and hard disks of an ECS instance that was created from a custom image?

Yes, you can upgrade the CPU, memory, bandwidth, and hard disks of an ECS instance that was created from a custom image. For more information, see Overview of instance configuration changes.

### Can I use a custom image across regions?

No, custom images cannot be used across regions. For example, a custom image created from an instance located in the China (Hangzhou) region cannot be used to create an instance in the China (Shanghai) region.

If you want to use a custom image across regions, you can copy the image to the destination region. For more information, see Copy a custom image.

## Can a custom image created from a subscription instance be used to create a pay-as-you-go instance?

Yes, a custom image created from a subscription instance can be used to create a pay-as-you-go instance. The use of custom images has nothing to do with the billing methods of instances.

I created an ECS instance from a custom image and specified a system disk capacity larger than that in the image. However, the system disk capacity of the new ECS instance is the same as that in the image. What do I do?

The system disk capacity of an instance created from a custom image may fail to be expanded due to one of the following reasons: The cloud-init service is not installed, the cloud-init service fails, or the file systems do not support the capacity expansion.

You can manually expand the system disk capacity.

## Why must I comment out mount entries when I create a custom image or an instance?

When you create an instance from a custom image, the following situations may cause failures to mount disks:

- The created instance does not have data disks.
- Data disks are new disks that are not formatted or partitioned.
- The mount entries for the data disks are not commented out in the /etc/fstab file of the custom image.

The following example shows the failure to mount a data disk. In this example, a data disk of an instance that was created from a custom image is not partitioned, and the mount entry for this data disk is not commented out in the /etc/fstab file of the custom image.

1. A data disk of the instance is not partitioned, as shown in the following figure.

```
[root@test ~]# fdisk -1
Disk /dev/xvda: 21.5 GB, 21474836480 bytes
255 heads, 63 sectors/track, 2610 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00078f9c
    Device Boot
                    Start
                                  End
                                            Blocks
                                                     Id System
/dev/xvdal *
                                  2611
                                          20970496
                                                     83 Linux
                         1
Disk /dev/xvdb: 10.7 GB, 10737418240 bytes
255 heads, 63 sectors/track, 1305 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00000000
```

2. In the instance, the mount entry for the data disk is not commented out in the /etc/fstab file, as shown in the following figure.

```
[root@test ~]# cat /etc/fstab
# /etc/fstab
# Created by anaconda on Thu Aug 14 21:16:42 2014
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
UUID= / ext4 defaults, barrier=0 1 1
                      /dev/shm
                                             tmpfs
                                                     defaults
                                                                    0 0
tmpfs
                      /dev/pts
                                             devpts gid=5, mode=620 0 0
devpts
                      /sys
                                                     defaults
                                                                    0 0
sysfs
                                             sysfs
                                                     defaults
                                                                    0 0
proc
                       proc
                                             proc
/dev/xvdb1
                                                     defaults
                                                                    0 0
                       alidata
                                             ext4
          ~]#
[root@test
```

3. When the instance starts, the system attempts to mount the data disk based on the configurations in the /etc/fstab file. The data disk cannot be mounted because the data disk is not partitioned, as shown in the following figure.

```
[root@test ~]# mount -a
mount: special device /dev/xvdb1 does not exist
[root@test ~]#
```

You do not need to comment out the mount entries for the data disks of an instance if these disks are created along the instance from snapshots of partitioned and formatted data disks.

If you have further questions, submit a ticket.

### How do I configure and use a private Docker image registry?

Image management is at the core of Docker. To allow organizations to internally share images, Docker has created the open source docker-registry project on GitHub to provide Docker Registry. Docker Registry is a tool used to manage private Docker image registries.

Start Docker Registry that supports Alibaba Cloud Object Storage Service (OSS). You can download Docker Registry from GitHub, install Docker Registry, and then run the pip install Docker Registry-driver-alioss command to install the OSS driver.

1. Run Docker Registry.

```
docker run -e OSS_BUCKET=-e STORAGE_PATH=/docker/ -e OSS_KEY=-e OSS_SECRET=-p 5000:500
0 -d chrisjin/registry:ali_oss
```

2. Configure config.yml.

```
```local: &local
<<: *common
storage: alioss
storage_path: _env:STORAGE_PATH:/devregistry/
oss_bucket: _env:OSS_BUCKET[:default_value]
oss_accessid: _env:OSS_KEY[:your_access_id]
oss_accesskey: _env:OSS_SECRET[:your_access_key]```</pre>
```

3. Start Docker Registry.

DOCKER\_REGISTRY\_CONFIG=[your\_config\_path] gunicorn -k gevent -b 0.0.0.0:5000 -w 1 dock er\_registry.wi:application

If you have further questions, submit a ticket.

### How do I clone an ECS instance?

You can clone the environment and data of an existing ECS instance within your account to create ECS instances with identical environments and configurations. Use one of the following methods:

• Method 1: Manually clone an ECS instance within the same region by using the ECS console.

i.

- ii. Find the ECS instance that you want to clone and create a custom image from the instance. For more information, see Create a custom image from an instance.
- iii. Use the created custom image to create an ECS instance by following the instructions in Create an instance by using the wizard. When you create the ECS instance, take note of the following parameters:
  - Region: Select the region where the custom image is located.
  - Image: In the Image section, select Custom Image. Then, select the custom image that you created in the previous step from the drop-down list.
    - **Note** If the selected custom image contains one or more data disk snapshots, an equal number of data disks are automatically created from the snapshots. Each disk has the same size as the snapshot from which it is created. You can extend a data disk but cannot shrink it.
- Method 2: Manually clone an ECS instance across regions by using the ECS console.

i

- ii. Select the ECS instance that you want to clone and create snapshots for its system disk and data disks. For more information, see Create a snapshot of a disk.
  - ? Note To ensure data consistency, create snapshots only when the instance is in the Stopped state.
- iii. Copy the snapshots to a different region in which you want to create an instance. For more information, see Copy a snapshot.
- iv. Create a custom image from the copy of the system disk snapshot. In the Create Image dialog box, select **Add Data Disk Snapshot** and click **Add** to add one or more data disk snapshots to the image. For more information, see Create a custom image from a snapshot.
- v. Use the created custom image to create an ECS instance by following the instructions in Create an instance by using the wizard. When you create the ECS instance, take note of the following parameters:
  - Region: Select the region where the custom image is located.

- Image: In the Image section, select Custom Image. Then, select the custom image that you created in the previous step from the drop-down list.
  - Note If the selected custom image contains one or more data disk snapshots, an equal number of data disks are automatically created from the snapshots. Each disk has the same size as the snapshot from which it is created. You can extend a data disk but cannot shrink it.
- vi. Use the created custom image to create an ECS instance by following the instructions in Create an instance by using the wizard. When you create the ECS instance, take note of the following parameters:
  - Region: Select the region where the custom image is located.
  - Image: In the Image section, select Custom Image. Then, select the custom image that you created in the previous step from the drop-down list.
    - Note If the selected custom image contains one or more data disk snapshots, an equal number of data disks are automatically created from the snapshots. Each disk has the same size as the snapshot from which it is created. You can extend a data disk but cannot shrink it.
- Method 3: Automatically clone an ECS instance by using Operation Orchestration Service (OOS).
  - To clone an ECS instance across regions, access the ACS-ECS-CloneInstancesAcrossRegion public template in the OOS console. In the top navigation bar, select the region where the instance is located. Use the ACS-ECS-CloneInstancesAcrossRegion public template to clone the ECS instance to a different region.
  - To clone an ECS instance across zones within a region, access the ACS-ECS-CloneInstancesAcrossAZ public template in the OOS console. In the top navigation bar, select the region where the instance is located. Use the ACS-ECS-CloneInstancesAcrossAZ public template to clone the ECS instance from one zone to another.

### Some custom images cannot be used to create I/O optimized instances. What do I do?

Some custom images cannot be used to create I/O optimized instances. If you want to use such a custom image to create an I/O optimized instance, submit a ticket in which the image name is specified.

## Where do I view the progress of an image import task? How long does it take to import an image?

You can view the progress of an image import task on the Images page in the ECS console. It may take an extended period of time to import a custom image. The amount of time it takes to import an image depends on the image size and the number of concurrent import tasks in the queue.

## Where do I view the progress of an image creation task? How long does it take to create an image?

You can view the progress of an image creation task on the Images page in the ECS console. The amount of time it takes to create an image depends on the size of the disks that are used to create the image.

### When do I need to copy a custom image?

Each custom image can be used only within a single region and cannot be used across regions. To achieve the following goals, you can copy custom images:

- Deploy applications in ECS instances to multiple regions.
- Migrate ECS instances to other regions.
- Use custom images across regions.
- Change the encryption state of a custom image, encrypt a system disk, or encrypt operating system data.

For more information, see Copy a custom image.



- To copy a custom image within the same region, you must encrypt the image copy.
- You cannot cancel an ongoing copy task of an encrypted image.

### Which images can be copied?

Only custom images can be copied. Public images, Alibaba Cloud Market place images, and images shared by other accounts cannot be copied.

### Which regions support copying custom images?

All Alibaba Cloud regions support copying custom images.

### How long does it take to copy a custom image?

When you copy a custom image across regions, the image file is transmitted from one region to another. The amount of time it takes to copy a custom image depends on the network transmission speed and the number of transmission tasks in the queue.

To copy a large image such as an image larger than 2 TiB in size across regions, you can first copy the associated snapshots to the destination region and then create a custom image from these snapshots in the destination region. This procedure takes less time than the procedure to directly copy the image. For more information, see Copy a snapshot and Create a custom image from a snapshot. For information about snapshot fees, see Snapshots.

### How am I charged when I copy a custom image?

You must perform the following operations to copy a custom image:

- 1. Copy the snapshot from which the custom image was created from the source region to the destination region.
- 2. Create a custom image from the snapshot in the destination region.

You may be charged the following fees for the preceding operations:

- Fees for traffic between the two regions. Alibaba Cloud does not charge you for cross-region traffic. For the latest billing details, see the official Alibaba Cloud website for announcements.
- The new snapshot (snapshot copy) consumes snapshot storage space in the destination region. Snapshots are billed based on the storage space used. For more information, see Snapshots.

## What limits apply to the original image (image copied) and the new image (image copy) during an image copy process?

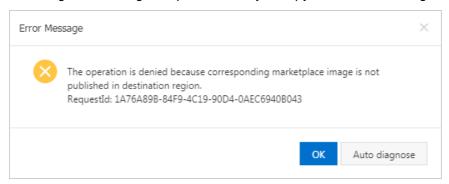
During an image copy process, the original image cannot be deleted, and the new image cannot be used to replace a system disk or create an ECS instance. The image copy process can be canceled.

## How do I copy images within my Alibaba Cloud account to other regions within other Alibaba Cloud accounts?

You must copy your own images to the destination regions and then share the images to the intended Alibaba Cloud accounts. After the images are shared, they are displayed in the shared image lists of those accounts.

## Can I copy a custom image derived from an Alibaba Cloud Marketplace image across regions?

If an Alibaba Cloud Market place image is available in the destination region, you can copy custom images derived from the Alibaba Cloud Market place image to the destination region. Otherwise, the following error message is reported when you copy such a custom image.



## How do I migrate data from regions outside the Chinese mainland to regions inside the Chinese mainland?

You can migrate data from regions outside the Chinese mainland to regions inside the Chinese mainland by copying images. For more information, see Copy a custom image.

### How many images can be shared to me?

A maximum of 100 images can be shared to you.

### How many users can an image be shared to?

An image can be shared to a maximum of 50 users.

### I have accounts on different Alibaba Cloud sites. Can I share images between these accounts?

Yes, images (except for custom images derived from Alibaba Cloud Marketplace images) can be shared between your accounts on the China site (aliyun.com), International site (alibabacloud.com), and Japan site (jp.alibabacloud.com).

### Do shared images count against my image quota?

No, shared images do not count against your image quota.

### Do geographical limits apply when shared images are used to create instances?

Yes, shared images can be used to create instances only in the regions where the shared images are located.

### What are the risks of creating an instance from a shared image?

The image owner can check how the image is shared and can delete the image. After a shared image is deleted by its owner, the system disks of the instances that use this image cannot be re-initialized.

Alibaba Cloud does not guarantee the integrity and security of images shared by other accounts. We recommend that you select only images shared by trusted accounts. After an ECS instance is created from a shared image, you must log on to the ECS instance to check the security and integrity of the shared image.

### What risks may arise if I share a custom image to other accounts?

If you share a custom image to other accounts, data and software may be leaked or stolen. Before you share a custom image to other accounts, check whether the image contains sensitive or important data. After the image is shared to other accounts, they can use the shared image to create ECS instances, which can then be used to create more custom images. During this process, data can be repeatedly spread. This gives rise to a risk of data being disclosed beyond your original intentions.

### After an account shares an image to me, can I share this image to another account?

No, only the owner of an image can share the image to other accounts.

### After I share an image, can I still use this image to create an instance?

Yes, after you share an image to another account, you can still use the image to create ECS instances and then create custom images from the instances.

## Can an image created from Instance A in one region be used by Instance B in a different region?

- If Instances A and B belong to the same account, you can copy the image to the region where
  Instance B is located and then apply the image to Instance B. For more information, see Copy a
  custom image.
- Instances A and B belong to different accounts, you can copy the image to the region where Instance B is located and then share the image to the account to which Instance B belongs. For more information, see Copy a custom image and Share or unshare a custom image.

## Is Bring Your Own License (BYOL) supported when I import custom images?

Yes, BYOL is supported when you import custom images. You can configure the license types by using the image import feature in the ECS console or by calling the ImportImage operation. For more information, see Import custom images and ImportImage.

### What types of licenses can I use when I import custom images?

When you import custom images, you can select one of the following license types:

#### Aliyun

Aliyun licenses are provided by Alibaba Cloud and include the licenses for Windows Server operating systems. If cloud-init is installed on the imported images, Alibaba Cloud uses KMS to activate the operating systems and provides Windows Server Update Services (WSUS).

#### BYOL

BYOL licenses are used in the following scenarios:

#### Microsoft

Microsoft BYOL licenses are used in the following scenarios:

■ BYOL implemented by using Software Assurance (SA)

BYOL can be implemented for software programs such as SQL Server and SharePoint that support License Mobility when ECS instances are created.

■ Windows operating system

Windows client access licenses (CALs) do not support License Mobility. Therefore, existing Windows licenses cannot be used within shared hardware environments. You must deploy Windows operating systems within a dedicated physical environment, which can be an Alibaba Cloud dedicated host or an ECS bare metal instance. For more information, see the Dedicated Host documentation and ECS Bare Metal Instance documentation.

For these kinds of ECS instances, Alibaba Cloud does not provide KMS, WSUS, or software technical support. You can contact Microsoft for software technical support.

■ No SA or no support for BYOL implemented by using SA

This scenario is similar to the Windows operating system scenario. You can reuse software licenses that you have purchased and downloaded, and deploy software programs in a dedicated hardware environment.

### o Redhat

Red Hat provides the Cloud Access program. If your Red Hat subscription to be migrated uses Bring Your Own Subscription (BYOS), you can sign up with Red Hat Cloud Access. .

#### Auto

Auto is the default value for License Type. When Auto is selected, a license type is automatically configured based on the operating system distribution to be imported.

- For operating systems such as Windows Server for which Alibaba Cloud has a signed licensing agreement and provides official licenses, the license type is Aliyun.
- For other operating systems such as noncommercial Linux images, the license type is BYOL. Alibaba Cloud does not provide software technical support for these operating systems.

### How are the imported BYOL images charged?

You are not charged for operating system components of BYOL images. This rule is applicable to newly created, renewed, or re-initialized ECS instances as well as ECS instances that have their configurations upgraded or downgraded.

## How are BYOL licenses authenticated and subscribed in Alibaba Cloud when their subscriptions expire?

You can change BYOL images to public images provided by Alibaba Cloud or Alibaba Cloud Market place images.

- For Windows Server operating systems, you can use the public images provided by Alibaba Cloud. For more information, see Overview.
- You can obtain SQL Server and Red Hat images in Alibaba Cloud Market place. For more information, see Alibaba Cloud Market place images.

### I want to export an image to my computer for testing. What do I do?

Alibaba Cloud custom images can be exported. For more information, see Export a custom image.

### Can an image exported from ECS be started on other platforms?

Images exported from ECS may not be started normally on other platforms. You can mount the disk files in the images to obtain data.

## Can I delete a custom image after it has been used to create an ECS instance?

You can select **Forcibly Delete** in the Delete Image dialog box to forcibly delete the image. However, after the image is deleted, the disks of the ECS instances created from the image cannot be reinitialized. For more information, see Re-initialize a system disk.

### Can I delete a custom image from my account after the image is shared to another account?

Yes, you can delete a custom image from your account after the image is shared to another account. However, after the shared image is deleted, the system disks of all ECS instances created from the image cannot be re-initialized. We recommend that you unshare the custom image before you delete it.

### What happens if I unshare Custom Image M to Account A?

If you unshare Custom Image M to Account A, Account A is unable to query Custom Image M by using the ECS console or by calling ECS API operations, and cannot use Custom Image M to create ECS instances or replace system disks. If Account A has created ECS instances from Custom Image M before the image is unshared, the system disks of these instances cannot be re-initialized.

# Why am I prompted with a message similar to "The specified image cannot be deleted because it is associated with instances" when I attempt to delete an image?

You may have created the image from a snapshot. To delete this image, you must select **Forcibly Delete**. After the image is forcibly deleted, instances created from it are still available, but their disks cannot be re-initialized. For more information, see <u>Delete a custom image</u>.

## When I replace a system disk, can I select an image that contains data disk snapshots?

No, you cannot select an image that contains data disk snapshots when you replace a system disk. If you want to use such an image to replace a system disk, we recommend that you use the image to create a pay-as-you-go instance and create a snapshot for the system disk of this new instance. You can then use the snapshot to create a custom image that contains only a system disk snapshot, and use the created custom image to replace the system disk.

### I want to replace the operating system of my ECS instance by using an existing image. What do I do?

For information about how to use an existing image to replace the operating system of an ECS instance, see Change the operating system.

**Note** We recommend that you create snapshots to back up data before you proceed.

### Can I use an image created from an instance in Account A to replace a system disk in Account B?

Yes, you can use an image created from an instance within Account A to replace a system disk within Account B. You can share the image to Account B and then replace the system disk. For more information, see Share or unshare a custom image.

Notice To use an image to replace a system disk, make sure that the image contains only a system disk snapshot.

### I am creating an ECS instance. Why is the total instance cost displayed when I select a custom image higher than that displayed when I select a public image?

This situation may occur in the following circumstances:

- The custom image contains data disk snapshots. When such an image is selected, the costs of the data disks cause the total cost of the instance to be higher than that of an instance created from a public image.
- The custom image was derived from a paid public image such as a Windows Server or RHEL image.

### What features do Alibaba Cloud Marketplace images provide?

A software environment such as the PHP, .NET, JAVA, or LAMP runtime environment and a variety of features such as the control panel and website building systems are pre-installed on the operating systems in Alibaba Cloud Market place images. You can use Alibaba Cloud Market place images to deploy runtime environments or software applications to ECS instances.

### What are the benefits of Alibaba Cloud Marketplace images?

You can use an Alibaba Cloud Market place image to create an ECS instance and deploy the preinstalled system environment or software of the image to the ECS instance. This eliminates the need to manually configure the environment or install software and enables you to create a ready-to-run runtime environment and conveniently build and manage services.

### What server environments and scenarios do Alibaba Cloud Marketplace images support?

Alibaba Cloud Market place provides hundreds of high-quality third-party images. These images not only support the deployment of runtime environments such as PHP, .NET, JAVA, LAMP, and Docker virtual containers, but can also meet personalized demands for website building, application development, and visual management.

### Are Alibaba Cloud Marketplace images safe?

All image service providers in Alibaba Cloud Market place have a wealth of experience in system maintenance and environment configuration. All images are made based on the official Alibaba Cloud operating systems that are installed with Alibaba Cloud Security. All images have passed strict security reviews and are safe to use.

## What do I do if I encounter an issue when I am installing or using an Alibaba Cloud Marketplace image?

You can view the service information on the buy page and contact the image service provider by TradeManager, phone, or email. They will answer your questions as soon as possible.

### How do I purchase an Alibaba Cloud Marketplace image?

You can purchase an Alibaba Cloud Market place image from Alibaba Cloud Market place or from the ECS instance buy page when you create an ECS instance.

### How long can I use a purchased image?

Theoretically, a purchased image can be used indefinitely. However, an image is a piece of software and has its own lifecycle. In addition, image providers provide services only over a limited period of time, which is described in the product details.

## Can I request refunds for my Red Hat images or SUSE images that are no longer needed?

No, you cannot request refunds for your Red Hat images or SUSE images that are no longer needed. By default, refunds are not available for ECS instances that use Red Hat or SUSE images. If you want to request a refund for an instance that uses a Red Hat or SUSE image, you can . However, the Red Hat or SUSE image is not refundable. After you receive the refund for the instance, the image can no longer be used by other instances.

### Are Alibaba Cloud Marketplace images refundable?

Alibaba Cloud Market place images support money-back guarantee within a period of time based on the Alibaba Cloud Market place rules. However, you are ineligible for a refund in the following situations:

- You have deployed the purchased image to an ECS instance within the money-back guarantee period.
- You have deployed the purchased image to an ECS instance before your application for a refund for this image is approved.
- You can receive refunds only for images that have not been used.

## Are free Alibaba Cloud Marketplace images still available after Alibaba Cloud Marketplace images are commercially available?

Free Alibaba Cloud Market place images are still available. However, you must purchase them at a price of USD 0.00 before you can use them.

I bought an Alibaba Cloud Marketplace image in the China (Hangzhou) region. Can I use it to create an ECS instance or replace a system disk in the China (Beijing) region?

No, Alibaba Cloud Marketplace images are region-specific. You can use an Alibaba Cloud Marketplace image that you purchased in a region to create ECS instances or replace system disks only within that region.

I have an instance created from an Alibaba Cloud Marketplace image. Do I need to make further payments for the image when I renew the instance or upgrade the configurations of the instance?

No, you do not need to make further payments for the image. After you purchase an Alibaba Cloud Market place image, you can use it on instances at no additional costs.

I have an ECS instance created from an Alibaba Cloud Marketplace image. After the instance is released, can I continue to use that image free of charge when I purchase a new ECS instance?

No, you cannot continue to use that image free of charge when you purchase an ECS instance.

I created an ECS instance from an Alibaba Cloud Marketplace image and then created a custom image from the instance. Do I need to pay for the custom image when I use it to create an ECS instance?

Yes, you must pay the original price of the Alibaba Cloud Market place image.

If I copy an Alibaba Cloud Marketplace image that I bought to another region to create an ECS instance, do I need to pay for the image?

Yes, you must pay the original price of the Alibaba Cloud Market place image.

I created an ECS instance from an Alibaba Cloud Marketplace image and then created a custom image from that instance. If I share the custom image to Account B, does Account B need to pay for the custom image when it uses this image to create an ECS instance?

Yes, Account B must pay the original price of the Alibaba Cloud Marketplace image.

Is a fee charged if I replace a system disk by using an Alibaba Cloud Marketplace image or an image derived from an Alibaba Cloud Marketplace image?

It depends. If the current image of your ECS instance is a different version of the replacement image, no fees are charged. Otherwise, a fee is charged.

My ECS instance is using an Alibaba Cloud Marketplace image. Am I charged if I replace the system disk of the instance?

No, you are not charged when you replace the system disk of the instance.

How do I call an ECS API operation to use an Alibaba Cloud Marketplace image or a custom or shared image derived from an Alibaba Cloud Marketplace image to create an ECS instance or replace a system disk?

1. Check whether the image in use is an Alibaba Cloud Market place image or an image derived from an Alibaba Cloud Market place image. You can call the Describelmages operation to query image information.

If the product ID ( ProductCode ) of your image is not empty, your image is an Alibaba Cloud Marketplace image or a custom or shared image derived from an Alibaba Cloud Marketplace image. For example, if the ProductCode of your image is abcd000111 , you can access the image at http://market.aliyun.com/products/123/abcd000111.html .

2. Select the version and region of the image and then purchase the image.

An image can be used only on ECS instances that are deployed within the same region in which the image was purchased. In addition, you can purchase only a single image at a time. If you want to create multiple ECS instances, you must purchase multiple images.

3. Use the image that you purchase to create an ECS instance or replace the system disk of an ECS instance.

If I do not purchase an Alibaba Cloud Marketplace image or an image derived from an Alibaba Cloud Marketplace image, is an error reported when I call an ECS API operation to use the image to create an ECS instance or replace a system disk?

Yes, an error is reported with the QuotaExceed.BuyImage error code.

I have configured a scaling group for which the minimum number of instances is set to 10 and the maximum number of instances is set to 100. What do I do with Alibaba Cloud Marketplace images to ensure that ECS instances are created to meet my computing needs?

If you want N instances that use the same image to be automatically created, you must purchase the image N times from Alibaba Cloud Market place in advance.

Can I purchase multiple Alibaba Cloud Marketplace images at a time?

No, you cannot purchase multiple Alibaba Cloud Market place images at a time.

If an image (such as jxsc000010 or jxsc000019) in use in an active scaling configuration no longer exists in Alibaba Cloud Marketplace, what do I do to ensure that ECS instances can continue to be created based on the scaling configuration in the corresponding scaling group?

We recommend that you select a suitable replacement image from Alibaba Cloud Market place to ensure that ECS instances are properly created in your scaling group.

### Can one product code support images in different regions?

Yes, one product code can support images in different regions if the regions already support the images.

## I bought 100 images that have the same product code. Can I use the images within all regions?

Alibaba Cloud Market place images are region-specific. If you want to use an image within a specific region, you must purchase the image in that region.

# After I select I/O Optimized, I cannot select Alibaba Cloud Marketplace images when I purchase an ECS instance. What is the cause and how do I resolve this issue?

View the details about and solution to this issue.

- Problem description: When you purchase an ECS instance on the official Alibaba Cloud website, you cannot select Alibaba Cloud Market place images.
- Cause: If you select I/O Optimized when you purchase an ECS instance, you cannot select Alibaba Cloud Market place images.
  - Compared with non-I/O optimized ECS instances, I/O optimized ECS instances provide better network capabilities between instances and disks to maximize the storage performance of standard SSDs. However, not all images support I/O optimized instances because the related optimization operations involve networks, storage, and internal drivers.
- Solution: When you purchase an I/O optimized instance, we recommend that you select an official standard image supported by the instance and then deploy your business environment on the instance.

If the issue persists, submit a ticket.

## What are yearly, monthly, and weekly subscription Alibaba Cloud Marketplace images?

Yearly, monthly, or weekly subscription Alibaba Cloud Market place images are images that are purchased from Alibaba Cloud Market place and billed on a subscription basis. These images are developed and maintained by image providers, who are responsible for both pre-sales consultation and after-sales services. In this topic, these images are collectively referred to as subscription images.

### On which ECS instances can subscription images be used?

A subscription image can be used only on a subscription instance with the same subscription duration.

### How do I purchase a subscription image? Can I separately purchase it?

No, you cannot separately purchase a subscription image.

You can use one of the following methods to purchase a subscription image:

• When you create an instance, set Billing Method to **Subscription**, select an **Alibaba Cloud Market place** image, and then specify a subscription duration by setting Duration.

**Note** In this case, you must pay for both the instance and the image. The instance is created on successful payment for both the image and instance.

• To use a subscription image on an existing subscription ECS instance, you can use the image to replace the operating system of the instance. In this case, you must select the image subscription duration based on the instance subscription duration. For more information, see Replace the operating system of an instance by using a non-public image.



Note In this case, you need to only pay for the image.

### How do I pay for subscription images?

Subscription images require payment upfront. The subscription duration of a subscription image must be the same as that of the subscription instance on which the image is used.

Image prices are set by the image providers.

### Can I use a subscription image after it expires? How do I continue to use it?

When a subscription image expires, it cannot be used unless it is renewed in a timely manner.

You cannot separately renew a subscription image. If you want to continue to use the image, you must renew the image together with the corresponding ECS instance. After the image is renewed, you can continue to use it.

### After I purchase a subscription image, can I request a refund if I no longer want to use it?

The image provider determines whether to make a refund. You can consult the image provider before you purchase the image.

### What can I expect when a refund is made?

If a refund is available, the image provider makes the refund based on your usage.

### Can a subscription image be converted into a pay-as-you-go image?

No, subscription images cannot be converted into pay-as-you-go images. This conversion feature is under development for release in the future. Stay updated on the official Alibaba Cloud website.

### Can I replace a subscription image with an image of another type or vice versa? How is the fee calculated?

Yes, you can replace images when you replace system disks of ECS instances. You can make the following replacements:

- Replace an image of another type (such as public image, custom image, or shared image) with a subscription image. After the image is replaced, the system calculates the actual cost based on the image cost and the remaining subscription duration of the ECS instance.
- Replace a subscription image with an image of another type (such as public image, custom image, or shared image). If the image provider allows for refunds, a refund is made based on your actual usage.
- Replace one subscription image with another subscription image. If a refund is available for the

replaced image after the replacement, the refund is made based on the refund policy. The actual cost of the replacement image is calculated based on the image price and the remaining subscription duration of the ECS instance.

## Where do I view and manage the subscription images that I purchased?

Log on to the ECS console. In the left-side navigation pane, choose Instances & Images > Images. Then, click the Market place Images tab to view and manage the subscription images that you purchased.

# Is a fee charged for a custom image derived from a subscription image? How is the custom image affected if the subscription image expires?

When you use a custom image derived from a subscription image to create an instance or replace a system disk, you are re-ordering the subscription image on Alibaba Cloud Marketplace. The custom image is not affected regardless of whether the subscription image expires.

## Will CentOS 8 public images be unpublished after CentOS 8 reaches its end of life (EOL)?

No, Alibaba Cloud will continue to provide CentOS 8 images. ECS instances that were created from CentOS 8 images are not affected. However, Alibaba Cloud will no longer provide updates for the CentOS 8 images because the CentOS community no longer provides updates for CentOS 8.

### Will my instances that were created from CentOS 8 images be affected after CentOS 8 reaches its EOL?

Instances that were created from CentOS 8 images are not affected. However, business deployed on the instances may be exposed to security risks because the CentOS community no longer provides support and security updates for CentOS 8. We recommend that you migrate your instances away from CentOS 8 at your earliest convenience. For more information, see Options for dealing with CentOS Linux EOL.

## Which operating systems are not supported by instance families that use AMD EPYC<sup>TM</sup> MILAN processors?

Ubuntu 18 or Debian 9 operating systems are not supported by instance families that use AMD EPYC MILAN processors. Ubuntu 18 and Debian 9 operating system kernels do not support AMD EPYC<sup>TM</sup> MILAN processors. If you use Ubuntu 18 or Debian 9 images to create instances of instance families that use AMD EPYC<sup>TM</sup> MILAN processors, the instances cannot be started. The following instance families use AMD EPYC<sup>TM</sup> MILAN processors:

- ecs.g7a
- ecs.r7a
- ecs.c7a
- ecs.ebmg7a
- ecs.ebmr7a
- ecs.ebmc7a

## In ECS Bare Metal instances, the kdump service cannot generate crash dump files. How do I resolve this issue?

For information about the cause of and solutions to this issue, see Solution to the issue that crash dump files cannot be generated on some ECS instances.

## A softlockup occurs during the writeback process of the Linux operating system kernel. How do I resolve this issue?

A softlockup occurs when some earlier versions of the Linux kernel write data back to file caches. For information about how to resolve this issue, see Solution to the issue that a softlockup occurs during the writeback process of the Linux operating system kernel.

### A softlockup error occurs when I delete cgroups in an ECS instance. What do I do?

For information about how to resolve this issue, see A softlockup error occurs when I delete cgroups in an ECS instance. What do I do?.

### An ECS instance goes down. What do I do?

- An ECS instance goes down and the "Out of memory and no killable processes" error message appears. What do I do?
- An ECS instance goes down and the "Objects remaining in kmalloc" message appears in an alert log. What do I do?
- An ECS instance goes down and the "RIP:get\_target\_pstate\_use\_performance" message appears in a log. What do I do?
- An ECS instance goes down and the "VFS: Unable to mount root fs on unknown-block" message appears in a log. What do I do?

# Why am I unable to use an SSH key pair that uses the ssh-rsa signature algorithm to connect to an instance that runs Fedora 33 64-bit?

When you use an SSH key pair that uses the ssh-rsa signature algorithm to log on to an ECS instance that runs Fedora 33 64-bit, you may be unable to connect to the instance. You can use one of the following methods to resolve this issue:

- Replace the SSH key pair that uses the ssh-rsa signature algorithm with an SSH key pair that uses other signature algorithms such as the ECDSA signature algorithm.
- Run the update-crypto-policies --set LEGACY command in the system to change POLICY to LEGACY. Then, you can continue to use the SSH key pair that uses the ssh-rsa signature algorithm.

### How do I install patches and compile the kernel on FreeBSD?

Alibaba Cloud FreeBSD public images already have their kernels patched to meet the startup requirements for instance families in Generation V or later. You can call the DescribeInstanceTypeFamilies operation with the Generation parameter set to query the instance families.

In the following situations, you can use the FreeBSD kernel source code to install patches and compile the kernel to solve the issue:

• If you use a FreeBSD image that is not provided by Alibaba Cloud or a custom image derived from

such a FreeBSD image to create an ECS instance of an instance family in Generation V or later, the instance may be unable to start.

• If you use a FreeBSD public image to create an ECS instance of an instance family in Generation V or later and use freebsd-update to update the kernel with new patches, the instance may be unable to start.

In this example, the FreeBSD kernel source code is used to demonstrate how to install patches and compile the kernel. In this example, FreeBSD 12.3 is used.

1. Download and decompress the FreeBSD kernel source code package.

```
wget https://mirrors.aliyun.com/freebsd/releases/amd64/12.3-RELEASE/src.txz -0 /src.txz
cd /
tar -zxvf /src.txz
```

2. Download patches.

In this example, 0001-virtio.patch is downloaded.

```
cd /usr/src/sys/dev/virtio/
wget https://ecs-image-tools.oss-cn-hangzhou.aliyuncs.com/0001-virtio.patch
patch -p4 < 0001-virtio.patch</pre>
```

3. Copy the kernel files and compile and install the kernel.

N in the make -j < N > command indicates the number of jobs that run in parallel. Set N based on your compiling environment. For example, for a single-vCPU environment, we recommend that you set -j < N > to  $-j \ge 0$ . This indicates that the ratio of the number of vCPUs to the N value is  $-1 \ge 0$ .

```
cd /usr/src/
cp ./sys/amd64/conf/GENERIC .
make -j2 buildworld KERNCONF=GENERIC
make -j2 buildkernel KERNCONF=GENERIC
make -j2 installkernel KERNCONF=GENERIC
```

4. After the kernel is compiled, delete the source code.

```
rm -rf /usr/src/*
rm -rf /usr/src/.*
```

Why does the load average become high on ECS instances that run Ubuntu operating systems of specific versions after the Server Guard (AliYunDun) process is started on the instances?

After the Server Guard process is started on ECS instances that run Ubuntu operating systems of specific versions such as Ubuntu 18.04, the load average of the instances becomes high. After the Server Guard process is terminated, the load average drops to a normal level.

For information about the causes of and solutions to this issue, see The system load is high after the Ubuntu 18.04 process is started on an ECS instance of the Server Guard version.

Why am I unable to select a Windows operating system for ECS instances?

When you create an ECS instance based on a Windows operating system, make sure that the instance memory size is greater than or equal to 1 GiB. For ECS instances that have less than 1 GiB of memory, you can select only Linux and Windows Server 1709 images.

### Does Alibaba Cloud support Windows Server 2008 and Windows Server 2008 R2?

## The operating system of my instance is Windows Server. I am prompted with a message indicating that the operating system is not genuine. What do I do?

Activate the Windows operating system. For more information, see How to use the KMS domain name to activate a Windows instance in a VPC.

#### Am I charged for the images used by ECS instances?

You are charged for the Windows Server, Red Hat, and SUSE public images. The fees vary based on instance types. Other public images are free of charge. For more information about the fees for other types of images, see Image overview.

### Can I install or upgrade the operating systems of ECS instances on my own?

No, you cannot install or upgrade the operating systems of ECS instances on your own. ECS instances must use images provided by Alibaba Cloud, which you cannot add or upgrade on your own. However, you can perform the following operations:

- Replace system disks and select new operating systems. For more information, see Change the
  operating system or Upgrade the operating systems of ECS instances in batches by replacing the
  system disks.
- Create ECS instances from custom images that were imported from your computers.
  - For information about how to import an image, see Instructions for importing images.
  - For information about how to create ECS instances from a custom image, see Create an ECS instance by using a custom image.
- Patch the operating systems.

#### Do operating systems have graphical interfaces?

Each Windows operating system except for Windows Server Semi-Annual Channel offers a management desktop. For information about how to use Windows Server Semi-Annual Channel operating systems, see Manage Windows Server Semi-Annual Channel images and instances.

Each Linux operating system offers a command line interface. You can install a graphical desktop.

#### How do I choose an operating system?

For information about how to choose an operating system, see Select an image.

#### Do public images come with the FTP service?

No, public images do not come with the FTP service. You must configure the FTP service on your own. For more information, see Manually build an FTP site on a Windows instance and Manually build an FTP site on a CentOS 7 instance.

#### Which SUSE versions do Alibaba Cloud public images support?

Alibaba Cloud public images support SUSE versions. For more information about the SUSE versions provided by Alibaba Cloud public images, see the "Alibaba Cloud Linux images" section in Overview.

#### What service support is available for SUSE operating systems?

SUSE Linux Enterprise Server (SLES) operating systems that are sold in Alibaba Cloud Market place are synchronized with SUSE update sources on a regular basis. For instances created from Alibaba Cloud SLES public images, the support for their operating systems is covered by the Alibaba Cloud enterpriselevel support service. If you have purchased the enterprise-level support service and encounter an issue when you use an SLES operating system, submit a ticket to contact Alibaba Cloud technical support personnel.

#### If an image was manually created from an ECS instance, can I retrieve instance data after the instance is released on expiration?

Yes, you can retrieve instance data by using one of the following methods:

- Create an instance from the created image. For more information, see Create an ECS instance by using a cust om image.
- Use the created image to replace the system disk of an existing instance. For more information, see Replace the operating system of an instance by using a non-public image.

Notice When you replace the system disk, take note of the following items:

- o All data on the replaced system disk is lost, and the disk is released.
- The replacement image must be in the same region as the existing instance.

#### I have an ECS instance and I want to create another ECS instance from the image of this existing ECS instance. What do I do?

You can create a custom image from the existing ECS instance and then use the custom image to create an ECS instance. For more information, see Create a custom image from an instance and Create an ECS instance by using a custom image.

#### I have purchased an ECS instance. How do I restore my shared image to the newly purchased instance?

Make sure that you have shared the image to the account to which the newly purchased instance belongs. Use one of the following methods to restore the image to the instance:

- If the shared image and the instance are located in the same region, replace the system disk of the instance and select the shared image for the new system disk. For more information, see Replace the operating system of an instance by using a non-public image.
- If the shared image and the instance are not located in the same region, copy the image to the region where the instance is located. Then, replace the system disk of the instance and select this image for the new system disk. For more information, see Copy a custom image and Replace the operating system of an instance by using a non-public image.

- Notice The following risks may arise when you replace the system disk of an instance:
  - The original system disk will be released. We recommend that you create a snapshot to back up your dat a in advance.
  - You must stop the instance before you can replace its system disk. When the instance is stopped, the services that run on the instance are interrupted.
  - After you replace the system disk, you must re-deploy the service environment on the new system disk. This may cause services on the instance to be interrupted for an extended period of time.
  - When the system disk is being replaced, a new system disk with a different disk ID is allocated to the instance. The snapshots of the original system disk cannot be used to roll back the new system disk.

#### I have multiple Alibaba Cloud accounts. I want to transfer an instance from Account A to Account B or migrate the environment and applications of an instance within Account A to an instance within Account B. What do I do?

You can perform the following steps:

- 1. Create a custom image from the instance within Account A. For more information, see Create a custom image from an instance.
- 2. Share the image to Account B. For more information, see Share or unshare a custom image.
- 3. Create an instance within Account B from the shared image. For more information, see Create an ECS instance by using a custom image.

#### How do I migrate data between ECS instances?

You can perform the following steps to migrate data from one ECS instance to another:

- 1. Create a custom image from the source ECS instance.
- 2. Copy or share the custom image.
  - If the source and destination instances are located within the same region and belong to the same account, go to the next step.
  - o If the source and destination instances are located within different regions but belong to the same account, copy the image to the region where the destination instance is located. For more information, see Copy a custom image.
  - If the source and destination instances are located within the same region but belong to different accounts, share the custom image to the account to which the destination instance belongs. For more information, see Share or unshare a custom image.
  - o If the source and destination instances are located within different regions and belong to different accounts, copy the image to the region where the destination instance is located and then share the image to the account to which the destination instance belongs. For more information, see Copy a custom image and Share or unshare a custom image.
- 3. Use the shared image to create an ECS instance or replace the image of the destination instance. For more information, see Create an ECS instance by using a custom image or Change the operating system.

**Note** If you want to replace the image of the destination instance, make sure that the original image does not contain data disk snapshots.

If the preceding steps are not applicable, see Migrate data between Alibaba Cloud VMs.

#### Can ECS instances in different VPCs communicate with each other?

Express Connect and Cloud Enterprise Network (CEN) can be used to allow ECS instances in different VPCs to communicate with each other. For more information, see Use CEN to enable inter-region network communication in CEN documentation.

#### How do I handle a CentOS DNS resolution timeout?

View the details about and solution to the CentOS DNS resolution timeout issue.

Cause

The DNS resolution mechanism of CentOS 6 and CentOS 7 has changed. A DNS resolution timeout may occur in CentOS 6 or CentOS 7 instances that were created before February 22, 2017 or that use custom images created before February 22, 2017.

Solution

You can perform the following steps to solve this issue:

- i. Download the fix\_dns.sh script.
- ii. Place the downloaded script in the /tmp directory of the CentOS system.
- iii. Run the bash /tmp/fix\_dns.sh command to execute the script.
- Script role

The script determines whether the /etc/resolv.conf file contains the options > single-request-representation. For more information, see resolv.conf configuration file.

The DNS resolution mechanism of Cent OS 6 and Cent OS 7 uses the same 5-tuple to send IPv4 and IPv6 DNS requests, for which the single-request-reopen option must be added. When two requests from the same port need to be handled after the option is added, the resolver closes the socket after the resolver sends the first request and then opens a new socket before the resolver sends the second request. The option takes effect immediately after it is added. You do not need to restart the instance.

- Script logic
  - i. Determine whether the operating system of the instance is CentOS.
    - If the operating system is not CentOS (for example, the operating system is Ubuntu or Debian), the script stops running.
    - If the operating system is CentOS, the script continues to run.
  - ii. Check the /etc/resolv.conf file for the options configuration.
    - The options configuration is unavailable.

By default, the Alibaba Cloud options configuration (options timeout:2 attempts:3 rotate single-request-reopen) is used.

options timeout:2 attempts:3 rotate single-request-reopen
; generated by /sbin/dhclient-script

Elastic Compute Service Images- FAQ

- The options configuration is available.
  - If the single-request-reopen option does not exist, append this option to the options configuration.
  - If the single-request-reopen option exists, the script stops running and the DNS nameserver configuration does not change.

### Why does ECS disable virtual memory and leave swap partitions unconfigured by default?

When physical memory is insufficient, the memory manager saves memory data that has been inactive for an extended period of time to a swap partition or virtual memory file. This mechanism helps increase the available memory.

However, if memory usage is already high and I/O performance is poor, the mechanism decreases the available memory instead. Alibaba Cloud ECS cloud disks use distributed file systems for storage and provide multiple strongly consistent replicas for each piece of data. This mechanism ensures the security of user data but deteriorates the storage and I/O performance of local disks by tripling the number of I/O operations.

Therefore, virtual memory is not enabled for Windows and swap partitions are not configured for Linux by default to avoid further decreasing I/O performance when system resources are insufficient.

#### How do I enable the kdump service in a public image?

By default, the kdump service is disabled in public images. If you want your instance to generate a core file when the instance is down so that you can analyze the downtime cause based on the file, you can perform the following steps to enable the kdump service. In the following example, the CentOS 7.2 public image is used: Operations may vary based on the version of your operating system.

- 1. Configure the directory in which to generate the core file.
  - i. Run the vim /etc/kdump.conf command to open the kdump configuration file.
  - ii. Run the path command to configure the directory in which to generate the core file. In this example, the directory is /var/crash, and the following path command is used:

```
path /var/crash
```

- iii. Save and close the /etc/kdump.conf file.
- 2. Enable the kdump service.

Use one of the following methods based on the operating system to enable the kdump service. In this example, the kdump service is enabled in CentOS 7.2 by using Method 1.

• Method 1: Run the following commands to enable the kdump service:

```
systemctl enable kdump.service
systemctl start kdump.service
```

• Method 2: Run the following commands to enable the kdump service:

```
chkconfig kdump on service kdump start
```

3. Run the following command to simulate the scenario where the instance is down:

```
echo c > /proc/sysrq-trigger
```

**Note** After the command is run, the instance is disconnected from the network. You must reconnect the instance to the network to perform the subsequent operations.

- 4. Analyze the core file.
  - i. Run the following command to install the Crash analysis tool:

```
yum install crash
```

ii. Download the debuginfo installation package.

Run the uname -r command to view the operating system kernel version and download the debuginf o installation package that matches the kernel version.

- kernel-debuginfo-common-x86\_64-<Kernel version>.rpm
- kernel-debuginfo-<Kernel version>.rpm

**Note** The download URLs for the debuginf o package vary based on the CentOS versions. You can find the download link that corresponds to your kernel version on the official CentOS website. For more information, see CentOS debuginf o packages.

In this example, the kernel version is 3.10.0-514.26.2.e17.x86\_64 . The following download commands are used:

```
\label{lem:model} wget \ http://debuginfo.centos.org/7/x86_64/kernel-debuginfo-common-x86_64-3.10.0-514\\ .26.2.el7.x86_64.rpm
```

wget http://debuginfo.centos.org/7/x86\_64/kernel-debuginfo-3.10.0-514.26.2.el7.x86\_
64.rpm

iii. Run the following commands to install the debuginfo package:

```
rpm -ivh kernel-debuginfo-common-x86_64-3.10.0-514.26.2.el7.x86_64.rpm
rpm -ivh kernel-debuginfo-3.10.0-514.26.2.el7.x86_64.rpm
```

iv. Run the following commands to use the Crash analysis tool to analyze the core file:

```
cd <core file directory>
crash /usr/lib/debug/lib/modules/<Kernel version>/vmlinux vmcore
```

In this example, the core file directory is  $\sqrt{var/crash/127.0.0.1-2019-07-08-15:52:25}$ , and the kernel version is 3.10.0-514.26.2.e17.x86\_64 . The following commands are used:

```
cd /var/crash/127.0.0.1-2019-07-08-15:52:25

crash /usr/lib/debug/lib/modules/3.10.0-514.26.2.el7.x86_64/vmlinux vmcore
```

#### How do I obtain the dump file for RHEL images?

Some RHEL images do not have the kdump service enabled by default. You can submit a ticket to obtain the dump file. For instances that have memory of greater than 16 GiB, you may be unable to obtain the dump file by submitting a ticket. Details on the ticket shall prevail.

### How do I enable the kdump service on an instance that uses a Red Hat 8.1 image?

By default, the kdump service is disabled in Red Hat 8.1 v1 images. To enable the kdump service, run the following commands:

```
systemctl enable kdump
systemctl start kdump
```

### Red Hat 8.1 and Red Hat 8.2 images start slowly on ECS bare metal instances. How do I solve this issue?

When a Red Hat 8.1 or Red Hat 8.2 image is used on an ECS bare metal instance, the image takes up to 2 minutes longer to start than a Red Hat 7 image. To solve this issue, you can change the kernel startup parameter settings from console=tty80 console=tty80, 115200n8 to console=tty0 console=tty80, 115200n8 in the /boot/grub2/grubenv file of the Red Hat 8.1 or Red Hat 8.2 image, and restart the instance for the change to take effect.

After I use a CentOS 7 or Windows ECS instance for an extended period of time without restarting it, the instance is disconnected from the network, the network is no longer available, or the public or private IP address of the instance cannot be pinged. What do I do?

For information about the cause of and solution to this issue, see Troubleshoot IP address faults in CentOS 7 instances and Windows instances.

## The "UNEXPECTED INCONSISTENCY; RUN fsck MANUALLY." error is reported when an ECS instance starts. What do I do?

This indicates that a file system error occurs due to data loss in the memory of the ECS instance, which may be caused by conditions such as a power outage. For more information about the issue and the solution to it, see How to solve the "UNEXPECTED INCONSISTENCY; RUN fsck MANUALLY" error returned when the ECS instance operating system fails to start.

#### How do I upgrade RHEL 7 to RHEL 8?

For information about how to upgrade RHEL 7 to RHEL 8, see Upgrading from RHEL 7 to RHEL 8.

How do I enable a program to distinguish ECS instances when I deploy my business?

If you want to enable your program to distinguish ECS instances when you deploy your business on a Linux operating system, you can run the systemd-machine-id-setup command to reset /etc/machine-id. This method requires you to re-generate the machine ID on each ECS instance. To prevent excess operations, we recommend that you use the /sys/class/dmi/id/product\_uuid file to distinguish ECS instances. The /sys/class/dmi/id/product\_uuids file is one of the unique external identifiers of an instance and can replace the /etc/machine-id file. This method is easy to use. You need to only obtain the UUID of the corresponding instance. You can run the following command to view the /sys/class/dmi/id/product\_uuid file of an ECS instance:

```
cat /sys/class/dmi/id/product_uuid
```

The following figure shows the output of an example command.

```
[root@test ~]# cat /sys/class/dmi/id/product_uuid
44CC9932-AF69-/
```

After I used a Fedora CoreOS image to create an instance of a specific instance family, I found that the number of CPUs in the lscpu command output is only half the number of vCPUs of the selected instance type. Why?

Assume that after you used a Fedora CoreOS image to create an instance of a specific instance family such as g5, you ran the lscpu command on the instance to view CPU information. In the command output, the total number of vCPUs indicated by the On-line CPU(s) list value is only half the number of vCPUs of the selected instance type. For example, for an instance type with two vCPUs used to create an instance, the On-line CPU(s) list value indicates only one vCPU, as shown in the following figure.

```
[core@iZbp1fjhzdqnsr47keavplZ ~]$ lscpu
Architecture: x86_64
CPU op-mode(s): 32-bit, 64-bit
Byte Order: Little Endian
Address sizes: 46 bits physical, 48 bits virtual
CPU(s): 2
On-line CPU(s) list: 0
Off-line CPU(s) list: 1
Thread(s) per core: 1
```

Note The On-line CPU(s) list value indicates vCPU serial numbers. In the preceding figure, only a single vCPU whose serial number is 0 is available.

This is because in Fedora CoreOS kernel, the mitigations boot parameter is set to auto, nosmt by default to mitigate all default mitigations and disable simultaneous multithreading (SMT) for vulnerable systems. To view the mitigations parameter that is set to auto, nosmt, you can run the cat /proc/cmdline command.

For more information about SMT, see Automatically disable SMT when needed to address vulnerabilities and Policy for disabling SMT.

## 13.2. Manually update virtio drivers of Windows instances

This topic describes how to manually update the virtio drivers of Windows Elastic Compute Service (ECS) instances.

#### **Procedure**

1. Create snapshots to back up data.

We recommend that you create snapshots for your Windows instance to back up instance data. If an exception occurs in the instance when the virtio driver is updated, you can use snapshots to roll back the data stored in the instance. For more information, see Create a snapshot of a disk.

2. Disable the third-party antivirus software.

To ensure that the update script or Cloud Assistant plug-in of the virtio driver runs normally, we recommend that you log on to your Windows instance and temporarily disable the third-party antivirus software installed in the instance. You do not need to disable the antivirus software that comes with the instance. If you confirm that your Windows instance does not need to update the virtio driver or that the virtio driver is updated, restart the third-party antivirus software in a timely manner.

- 3. Manually check or use Alibaba Cloud Security Center to check whether the virtio driver of your Windows instance needs to be updated.
  - You can use the vulnerability fixing feature of Security Center to check whether the virtio driver
    of your Windows instance needs to be updated. For more information, see Use Security Center to
    update virtio drivers.

Notice Security Center provides multiple editions, and different editions support different features. For more information, see Features. If you use Security Center Basic or other editions that support only the vulnerability detection feature but not the vulnerability fixing feature, you can use the vulnerability detection feature to view the information of your Windows instance for which you want to update the virtio driver, and then manually update the virtio driver.

- You can manually check whether the virtio driver of your Windows instance needs to be updated. For more information, see Manually check whether virtio drivers need to be updated.
- 4. Manually update the virtio driver.

You can use one of the following methods to manually update the virtio driver.

Notice When you manually update the virtio driver, the network interface controller (NIC) driver is also updated. The network of the ECS instance is interrupted for about 10 seconds. Pay attention to the impacts on business caused by network interruption.

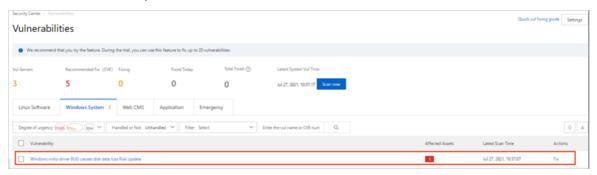
- If your Windows instance can access the Internet, we recommend that you use a script to update the virtio driver. For more information, see Method 1 for manual update: Use a script.
- If your Windows instance cannot access the Internet, we recommend that you upload an
  installation package of the virtio driver to manually update the driver. For more information, see
  Method 2 for manual update: Use an installation package.
- If you want to update the virtio drivers of multiple Windows instances that have access to the
  Internet or the internal networks within virtual private clouds (VPCs), we recommend that you use
  Cloud Assistant provided by Alibaba Cloud to batch update the drivers. For more information,
  see Method 3 for manual update: Use Cloud Assistant to batch update virtio drivers of multiple
  instances.

Note If the virtio driver cannot be updated by using a script, we recommend that you use a Cloud Assistant plug-in ( acs-plugin-manager.exe --exec --plugin=UpdateVirtIo ) to update the driver. If the update still fails, .

#### Use Security Center to update virtio drivers

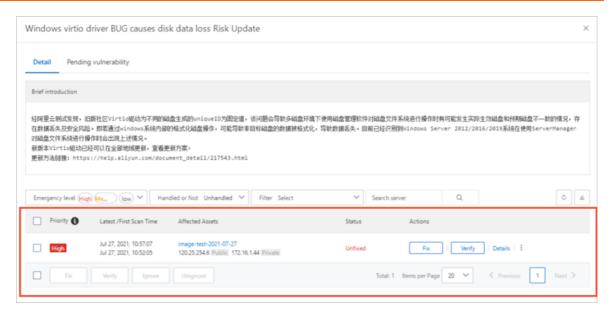
Alibaba Cloud Security Center provides the vulnerability detection and fixing feature to handle the possible impacts caused by virtio drivers of earlier versions. The following procedure describes how to use the vulnerability fixing feature of Security Center to check and update virtio drivers. For more information about the vulnerability fixing feature of Security Center, see Overview.

- 1.
- 2.
- 3.
- 4. Select **Unhandled** from the **Handled or Not Unhandled** drop-down list and check whether the **Windows virtio driver BUG causes disk data loss Risk Update** vulnerability announcement exists in the vulnerability announcement list.



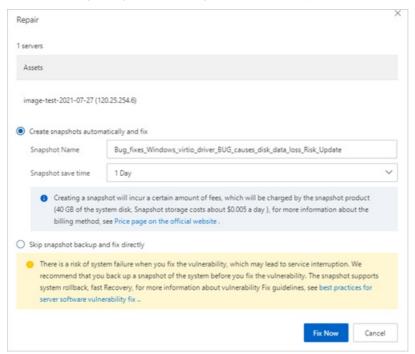
#### Description of the vulnerability announcement:

- If the vulnerability announcement does not exist, you do not have Windows instances whose virtio drivers need to be updated within your Alibaba Cloud account. You can skip the steps in this topic.
- If the vulnerability announcement exists, you can choose one of the following methods based on your Security Center edition to update your virtio drivers. For information about the vulnerability fixing feature of different editions of Security Center, see Vulnerability fixing.
  - If your Security Center edition supports the vulnerability detection feature but not the vulnerability fixing feature, you can manually update the virtio driver based on the methods for manually updating virtio drivers in this topic.
  - If your Security Center edition supports both the vulnerability detection and fixing features, you can follow this procedure to update the virtio driver.
- 5. Click Windows virtio driver BUG causes disk data loss Risk Update to go to the details panel of the vulnerability.
- 6. On the **Detail** tab, select all the Windows instances whose virtio drivers need to be updated and click **Fix** in the lower part of the panel.



7. In the **Repair** dialog box, click **Fix Now**.

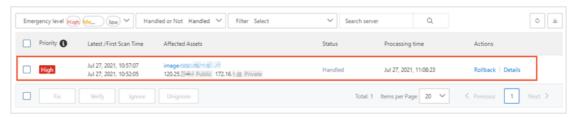
You can manually specify whether to create snapshots for ECS instances whose vulnerabilities are to be fixed. We recommend that you select **Create snapshots automatically and fix**. If you confirm that snapshots are already created for each instance before you fix the vulnerabilities, you can select **Skip snapshot backup and fix directly**.



- 8. After the vulnerabilities are fixed, click **Restart** in the **Actions** column to restart all the Windows instances whose vulnerabilities are fixed.
  - After the instances are restarted, **Reboot successful** is displayed in the **Status** column on the **Detail** tab.
- 9. View the details of handled vulnerabilities.
  - i. Go back to the Vulnerabilities page and click the Windows System tab.

ii. Select **Handled** from the **Handled** or **Not Handled** drop-down list and click **Windows** virtio driver BUG causes disk data loss Risk Update.

On the **Detail** tab, you can view the instances whose vulnerabilities are fixed, as shown in the following figure.



#### Manually check whether virtio drivers need to be updated

Before you update virtio drivers, you can run commands in your Windows instances to check whether the virtio drivers of your instances need to be updated. After you update virtio drivers based on the methods provided in this topic, you can also check whether the virtio drivers of your Windows instances are of the latest version.

- Connect to your Windows instance.
   For more information, see Connection methodsGuidelines on instance connection.
- 2. On the desktop, press Win+Rto open the Run dialog box.
- 3. Enter powershell to start PowerShell.
- 4. Run the following command to check the version of the virtio driver:

```
[System.Diagnostics.FileVersionInfo]:: GetVersionInfo("C:\Windows\System32\drivers\viostor.sys") \\
```

The following figure shows the command output. If the product version (**ProductVersion**) is earlier than 58017, we recommend that you follow the instructions provided in this topic to update your virtio driver.

5. Run the following command to view the unique identifier of the disk in the instance:

```
Get-Disk | Select-Object uniqueid
```

The following figure shows the command output. If the uniqueid value is 141463431303031 or a value that does not start with 000000, we recommend that you follow the instructions provided in this topic to update your virtio driver.

```
PS C:\Users\Administrator> Get-Disk | Select-Object uniqueid uniqueid uniqueid ------
```

#### Method 1 for manual update: Use a script

If your Windows instance can access the Internet, you can use a script to update the virtio driver of your instance.

Elastic Compute Service Images- FAQ

1. Connect to your Windows instance.

For more information, see Connection methods.

2. In the Windows instance, download the script used to update the virtio driver.

Click InstallVirtIo.ps1 to download the InstallVirtIo.ps1 script.

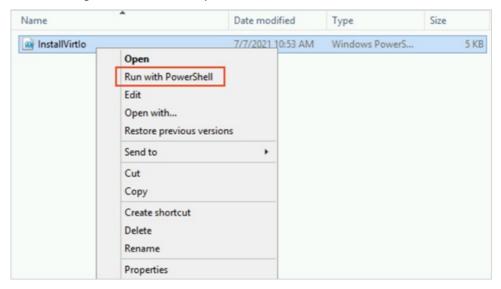
3. Run the InstallVirtIo.ps1 script to update the virtio driver.

In this example, *InstallVirtIo.ps1* is downloaded to the *C:\test* directory.

i. Open the C:\test folder.

Access the directory to which *InstallVirtIo.ps1* is downloaded.

ii. Select and right-click *InstallVirtIo.ps1* and then select **Run with PowerShell**.



Alternatively, you can press the *Shift* key and right-click a blank area in the folder window at the same time. Then, select **Open PowerShell window here**. Run the *InstallVirtIo.ps1* script in Windows PowerShell.



- If you are logged on to the Windows instance as a regular user, you must run this script as an administrator. If you are logged on as an administrator, you can directly run this script.
- If you are prompted that you are not allowed to run the script when you manually run the script, you must first modify the script execution policy of PowerShell. For more information, see FAQ.
- 4. After the script is run, restart the instance.

For more information, see Restart an instance. The updated virtio driver takes effect after your instance is restarted.

#### Method 2 for manual update: Use an installation package

If your Windows instance cannot access the Internet, you can use this method to update the virtio driver.

1. On your computer, download the virtio driver package provided by Alibaba Cloud.

Click Virtio Driver Package to download the package.

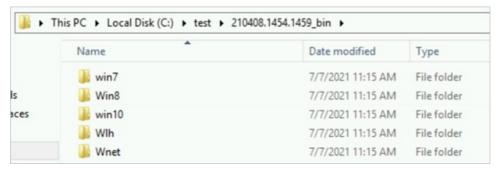
The downloaded package is named 210408.1454.1459 bin.zip.

2. Upload *210408.1454.1459\_bin.zip* to the Windows instance.

You can use one of the following methods to upload the package:

- If your computer runs a Windows operating system, use the Windows Remote Desktop Connection tool to connect to the instance from your computer and then upload the package.
- Build an FTP site on the instance and then use an FTP client to upload the package. For more information, see Manually build an FTP site on a Windows instance.
- 3. On the Windows instance, decompress *210408.1454.1459\_bin.zip* and open the *210408.1454.1459\_bin* folder.

In the 210408.1454.1459\_bin folder, you can find the subfolders that correspond to different Windows versions.



These subfolders contain the virtio driver installation files for the corresponding Windows versions. Take note of the following subfolders:

- win10: contains the driver installation files for Windows Server 2016, Windows Server 2019, and Windows 10.
- Win8: contains the driver installation files for Windows Server 2012 R2 and Windows 8.1.
- o win7: contains the drivers installation files for Windows Server 2008 R2 and Windows 7.

In each of these subfolders, driver installation files are put into the x86 or amd64 subfolder based on the architectures on which the drivers can be installed. The installation files in the amd64 subfolders are applicable to 64-bit operating systems with an AMD64 architecture. The installation files in the x86 subfolders are applicable to 32-bit operating systems with an x86 architecture.

- 4. Open an appropriate subfolder based on the operating system version of the Windows instance. In this example, open the amd64 subfolder in the *C:\test\210408.1454.1459\_bin\Win8\amd64* directory because the Windows instance runs a Windows Server 2012 R2 64-bit operating system.
- 5. After the *amd64* subfolder is opened, press the *Shift* key and right-click a blank area at the same time. Then, select **Open PowerShell window here**.



- 6. Run the following command in Windows PowerShell to install the virtio driver of the latest version.
  - **? Note** If you are logged on to the Windows instance as a regular user, you must run this command as an administrator. If you are logged on as an administrator, you can directly run this command.

```
pnputil -i -a *.inf
```

A command output shown in the following figure indicates that the virtio driver of the latest version is installed.

```
PS C:\Users\alienista = 20-14-p\210408.1454.1459_bin\win10\amd64> pnputil -i -a *.inf
Microsoft PnP Utility

Processing inf:
Successfully installed the driver.
Driver package added successfully.
Published name:

oem7.inf

Processing inf:
netkvm.inf
Successfully installed the driver.
Driver package added successfully.
Published name:
oem8.inf

Processing inf:
Successfully installed the driver.
Driver package added successfully.
Published name:
oem9.inf

Processing inf:
vioser.inf
Successfully installed the driver.
Driver package added successfully.
Published name:
oem10.inf

Processing inf:
viostor.inf
Successfully installed the driver.
Driver package added successfully.
Published name:
oem10.inf

Processing inf:
successfully installed the driver.
Driver package added successfully.
Published name:
oem10.inf

Total attempted:
5
Number successfully imported: 5
```

7. Restart the Windows instance.

For more information, see Restart an instance. The updated virtio driver takes effect after your instance is restarted.

Method 3 for manual update: Use Cloud Assistant to batch update virtio drivers of multiple instances

If you want to update virtio drivers of multiple Windows instances that have access to the Internet or the internal networks within VPCs, you can use Cloud Assistant provided by Alibaba Cloud to batch update the drivers of your instances.

**Note** To update virtio drivers by using Cloud Assistant, you must download the required installation package. The system downloads the package over the internal networks within VPCs. If the package cannot be downloaded over the internal networks, the system downloads the package over the Internet.

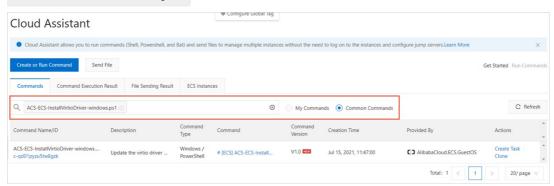
- 1.
- 2.
- 3. In the top navigation bar, select a region.

The selected region must be the region where your instances reside.

4. Use Cloud Assistant to batch update the virtio drivers.

You can use Cloud Assistant to manually set the command content to update the virtio drivers. You can also use the common commands of Cloud Assistant to update the drivers. This section describes different methods for updating virtio drivers by using Cloud Assistant.

- o (Recommended) Use common commands to batch update virtio drivers
  - a. On the Cloud Assistant page, click the Common Commands tab and enter allVirtioDriver-windows.ps1 in the search box.



- b. Click Create Task in the Actions column corresponding to the acs-Ecs-InstallVirtioDriv er-windows.ps1 common command.
- c. In the **Select Instances** section of the **Create Task** panel, select the IDs of the instances whose virtio drivers need to be updated.
- o Manually set the command content to batch update virtio drivers
  - a. On the Cloud Assistant page, click Create or Run Command.

Elast ic Comput e Service Images • FAQ

- b. In the Create Command panel, configure parameters.
  - In the **Command Information** section, configure the required parameters described in the following table and use default values for other parameters:

Parameter	Description		
Command Source	Select Enter Command Content.		
Command Name	Specify a name or use the default name.		
Command Type	Select PowerShell.		
Command	Use one of the following methods to update the virtio drivers:  Run the InstallVirtlo.ps1 script to update the virtio drivers.  You must download the InstallVirtlo.ps1 script to your computer and copy the entire script to the Command field.  Use a Cloud Assistant plug-in to update the virtio drivers.  Copy the following command to the Command field:		
	acs-plugin-manager.exeexec plugin=UpdateVirtIo		

■ In the **Select Instances** section, select the IDs of the instances whose virtio drivers need to be updated.

#### 5. Click Run.

You can view the execution results of the command on the **Command Execution Result** tab. For more information, see View execution results in the ECS console. The following figure shows a command output on one of the selected instances and indicates that the virtio driver of the instance is updated.

```
Published name: oem8.inf

Processing inf: pvpanic.inf
Successfully installed the driver.
Driver package added successfully.
Published name: oem9.inf

Processing inf: vioser.inf
Successfully installed the driver.
Driver package added successfully.
Published name: oem10.inf

Processing inf: viostor.inf
Successfully installed the driver.
Driver package added successfully.
Published name: oem11.inf

Total attempted: 5
Number successfully imported: 5
```

6. Batch restart the instances.

For more information, see Restart an instance. The updated virtio driver takes effect after your instance is restarted.

#### **FAO**

• What do I do if I am not allowed to run a script in PowerShell?

In some versions of Windows, PowerShell disables the script running feature by default to ensure security. You can modify the script execution policy of PowerShell to enable the script running feature. Perform the following operations:

i. Run the following command to make PowerShell run scripts:

```
Set-ExecutionPolicy Unrestricted
```

- ii. In the command line, enter A and press the *Enter* key to modify the script execution policy.
- iii. Manually run the script again.
- iv. To ensure system security, a security alert is returned in the command line of PowerShell. You must enter R in the command line and press the *Enter* key to run the script again.
- v. After the script is run, run the following command to restore the default settings of PowerShell to ensure system security:

```
Set-ExecutionPolicy Restricted
```

- vi. In the command line, enter A and press the *Enter* key to modify the script execution policy.
- What do I do if I am prompted that no digital signature is available for the NIC driver after the virtio driver is updated?

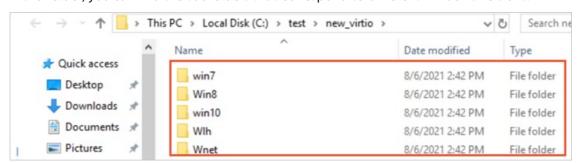
You can manually download the latest Windows driver package to replace the faulty NIC driver. Perform the following operations:

- i. Connect to the Windows instance.
   For more information, see Connect to a Windows instance by using a password.
- ii. Click Windows driver package to download the latest Windows driver package.

Note If you cannot download the driver package to your Windows instance, you can download the driver package to a physical host. Then, you can use Remote Desktop Connection or an FTP site to upload the driver package to the Windows instance. For information about how to build an FTP site, see Manually build an FTP site on a Windows instance.

After the package is downloaded, you can view the new\_virtio.zip driver package.

iii. Decompress the Windows driver package and open the new\_virtio folder.In the folder, you can find the subfolders that correspond to different Windows versions.



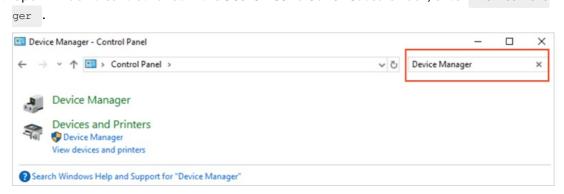
These subfolders contain the virtio driver installation files for the corresponding Windows versions. Take note of the following subfolders:

- win10: contains the driver installation files for Windows Server 2016, Windows Server 2019, and Windows 10.
- Win8: contains the driver installation files for Windows Server 2012 R2 and Windows 8.1.
- win7: contains the drivers installation files for Windows Server 2008 R2 and Windows 7.
- iv. Uninst all the NIC driver of the Windows instance.

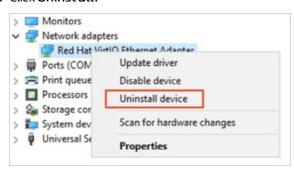
In this example, Windows Server 2019 64-bit is used. You can perform the similar steps for other Windows operating system versions.

Windows operating system versions.

a. Open Windows Control Panel. In the Search Control Panel search box, enter Device Mana



- b. Click Device Manager.
- c. Click Network adapters and right-click Red Hat Virt IO Ethernet Adapter.
- d. Click Uninst all.



e. In the Confirm Device Uninstall dialog box, select Delete the driver software for this device and click OK.

After the package is uninstalled, **Network adapters** is not displayed in Device Manager.

- v. Install the latest Windows NIC driver.
  - a. Open the Windows driver package file based on the version of your Windows operating system.
    - In this example, Windows Server 2019 64-bit is used, and the Windows driver package is stored in the *C:\test* folder. Therefore, the folder path is *C:\test\new\_virtio\win10\amd64*.
  - b. After the *amd64* subfolder is opened, press the *Shift* key and right-click a blank area at the same time. Then, select **Open PowerShell window here**.

c. In the PowerShell window, run the following command to install the latest Windows NIC driver:

```
pnputil -i -a netkvm.inf
```

The following figure shows that the NIC driver is installed.

```
Administrator: Windows PowerShell

PS C:\test\new_virtio\win10\amd64> pnputil -i -a netkvm.inf

Microsoft PnP Utility

Processing inf: netkvm.inf

Successfully installed the driver.

Driver package added successfully.

Published name: oem1.inf

Total attempted: 1

Number successfully imported: 1

PS C:\test\new_virtio\win10\amd64> __
```

vi. Return to or reopen Device Manager, right-click the hostname, and then click **Scan for** hardware changes.

In this example, the host name of the Windows instance is test, as shown in the following figure.



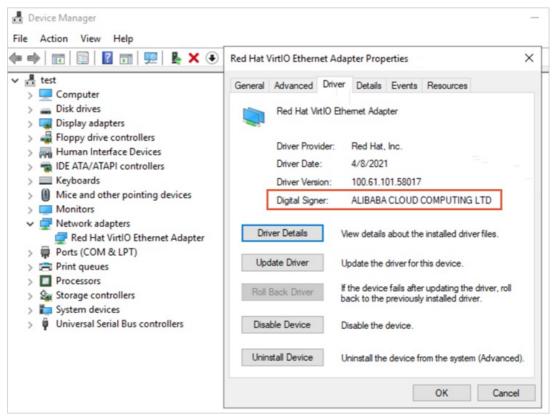
After the scanning is complete, you can view the added **Network adapters** in Device Manager.

- vii. View the information of the latest NIC driver.
  - a. Click Network adapters and right-click Red Hat Virt IO Ethernet Adapter.
  - b. Click Properties.

Elastic Compute Service Images FAQ

c. In the Red Hat VirtIO Ethernet Adapter Properties dialog box, click the Driver tab.

You can view the digital signature information of the latest NIC driver, as shown in the following figure.



viii. Restart the Windows instance.

For more information, see Restart an instance. After the instance is restarted, the new driver takes effect

## 13.3. Options for dealing with CentOS Linux EOL

Cent OS has announced the end of life (EOL) for Cent OS Linux. Alibaba Cloud Cent OS Linux public images are derived from the sources provided by Cent OS. When Cent OS Linux reaches its EOL, Alibaba Cloud stops providing support for the Cent OS Linux operating systems. This topic describes the impacts of the Cent OS Linux EOL and provides options for dealing with this change.

#### **Background information**

On December 8, 2020, Cent OS announced that they would be shifting focus from Cent OS Linux, the rebuild of Red Hat Enterprise Linux (RHEL), to Cent OS Stream. For more information, see Cent OS Project shifts focus to Cent OS Stream. The following section describes the EOL plans of different Cent OS Linux releases.

- Cent OS Linux 7, as a rebuild of RHEL 7, no longer receives updates from Cent OS as of August 6, 2020, and will reach its EOL on June 30, 2024.
- CentOS Linux 8, as a rebuild of RHEL 8, has had its lifecycle cut short and reached its EOL on December 31, 2021.

• CentOS Linux 9 and later versions are not provided. Instead, CentOS Stream is provided.

#### **Impacts**

If you use CentOS Linux, take note of the following items:

- If you use CentOS Linux 8, you no longer receive software maintenance or support such as vulnerability fixes and feature updates as of December 31, 2021.
- If you use Cent OS Linux 7, you will no longer receive software maintenance or support such as vulnerability fixes and feature updates as of June 30, 2024.

Since the lifecycle of CentOS Linux 8 has been cut short, CentOS provides the following options for CentOS Linux 8 users. If you use CentOS Linux 8 in a development or test environment, you can use CentOS Stream. If you use CentOS Linux 8 in a production environment or to deploy critical business, you can use RHEL. If you want to use CentOS Stream or RHEL, take note of the following items:

- Cent OS Stream is a rolling release distribution of Cent OS rather than a rebuild of RHEL. Therefore, system stability and compatibility cannot be guaranteed. If Cent OS Stream is used in a production environment, risks cannot be predicted.
- RHEL images are not provided for free, and you must subscribe to the images. We recommend that you evaluate the costs involved before you migrate to RHEL.
- CentOS Linux EOL affects the public images and support provided by Alibaba Cloud in the following aspects:
  - Alibaba Cloud will continue to provide CentOS 8 images. Elastic Compute Service (ECS) instances that were created from CentOS 8 are not affected. However, Alibaba Cloud will no longer update these images. For more information, see .
  - The support cycles provided by Alibaba Cloud for the CentOS operating systems are consistent
    with the lifecycles of the official CentOS operating systems. CentOS discontinues support for
    CentOS Linux 8 as of December 31, 2021, and will discontinue support for CentOS Linux 7 as of June
    30, 2024.

For more information about the CentOS project shifts, see FAQ - CentOS Project shifts focus to CentOS Stream.

#### Options

Alibaba Cloud provides public images that contain a variety of operating systems. You can select a suitable operating system to replace Cent OS Linux. The following table describes different operating systems and their intended users.

Operating system	Description	Intended user
------------------	-------------	---------------

Operating system	Description	Intended user	
Alibaba Cloud Linux	Alibaba Cloud Linux is an operating system provided by Alibaba Cloud. It is compatible with the RHEL and CentOS ecosystems, and is optimized for Alibaba Cloud. Alibaba Cloud provides Alibaba Cloud Linux 2 and Alibaba Cloud Linux 3 and provides free longterm support (LTS) for Alibaba Cloud Linux. For more information, see Overview or Alibaba Cloud Linux overview.  Alibaba Cloud Linux provides migration tools to migrate CentOS Linux to Alibaba Cloud Linux. If your instance runs a CentOS 7 or CentOS 8 operating system, you can use migration tools to migrate the operating system of your instance to Alibaba Cloud Linux. For more information, see the following topics:  • Migrate from CentOS 8 to Alibaba Cloud Linux 3  • Migrate from CentOS 7 to Alibaba Cloud Linux 2	Individuals and enterprises who require technical support and want to use free images.	
Anolis OS	Anolis OS is released by the OpenAnolis community. It is fully compatible with CentOS 8 and other mainstream Linux distributions. Anolis OS is an open source, independent Linux distribution and provides enterprise-level stability, security, reliability, and high performance. For more information, see OpenAnolis.	Individuals and enterprises who are used to using open source images and want to use free images.	
Red Hat Enterprise Linux	RHEL is an enterprise-level Linux operating system developed by Red Hat. RHEL images are not provided for free, and you must subscribe to the images. For more information, see Red Hat Enterprise Linux.	Individuals and enterprises who require technical support and want to migrate to an operating system that is similar to CentOS.	
CentOS Stream	CentOS Stream is a rolling release distribution provided by CentOS.  Individuals and enterprises who want to use an operating system that has rolling updates and that is similar to CentOS.		
Rocky Linux	Rocky Linux is a community-driven enterprise-level Linux operating system and a downstream rebuild of RHEL. Like CentOS, Rocky Linux offers stable releases for servers and is intended to be an alternative that is fully compatible with CentOS.		

Operating system	Description	Intended user
AlmaLinux	AlmaLinux is a stable, community-driven Linux distribution developed by the CloudLinux team. The AlmaLinux operating system is 1:1 binary compatible with RHEL. You can change the operating system of an instance to AlmaLinux without stopping the instance.	Individuals and enterprises who are used to using open source images and want to use free images.
Other Linux operating systems, such as Debian and Ubuntu	These operating systems vary in application compatibility and usage habits. Alibaba Cloud provides images that contain a variety of operating systems. For more information, see Select an image.	Individuals and enterprises who are capable of handling operating system changes.

#### References

• End of support for operating systems

•

Image FAQ

## 13.4. Change CentOS 8 repository addresses

Cent OS 8 has reached its end of life (EOL) and is no longer maintained by the Linux community. We recommend that you migrate to Anolis or Alibaba Cloud Linux. If you want to continue to use specific installation packages of Cent OS 8 during the transition period, change Cent OS 8 repository addresses as described in this topic.

#### Context

Cent OS 8 reached its EOL on December 31, 2021. In accordance with Linux community rules, all content has been removed from the following Cent OS 8 repository address:

http://mirror.centos.org/centos/8/ . All third-party image providers have removed CentOS 8 repositories for their images. Alibaba Cloud repositories at http://mirrors.cloud.aliyuncs.com and http://mirrors.aliyun.com can no longer be synchronized with CentOS 8 repositories. If you continue to use the default CentOS 8 repository configurations on Alibaba Cloud, an error is reported. The following figure shows an example error.

You can perform the following procedure to change the repository configurations of Elastic Compute Service (ECS) instances that run CentOS 8 based on network environments.

- YUM reposit ory for Cent OS 8
  - For ECS instances located in virtual private clouds (VPCs), change the YUM repository address to http://mirrors.cloud.aliyuncs.com/centos-vault/8.5.2111/ .
  - For ECS instances located in the classic network, change the repository address to http://mirrors.aliyuncs.com/centos-vault/8.5.2111/.

- Extra Packages for Enterprise Linux (EPEL) repository for CentOS 8
  - For ECS instances located in VPCs, change the EPEL repository address to http://mirrors.cloud.a liyuncs.com/epel-archive/8/.
  - For ECS instances located in the classic network, change the repository address to <a href="http://mirrors.com/epel-archive/8/">http://mirrors.com/epel-archive/8/</a>.

Note This topic describes how to change CentOS 8 repository addresses for ECS instances. If your server is not an ECS instance, make sure that the server can access the Internet and replace <a href="http://mirrors.cloud.aliyuncs.com">http://mirrors.cloud.aliyuncs.com</a> in CentOS 8 repository addresses with <a href="http://mirrors.aliyun.com">http://mirrors.aliyun.com</a>. For example, you can change the address of the YUM repository for CentOS 8 to <a href="http://mirrors.aliyun.com/centos-vault/8.5.2111/">http://mirrors.aliyun.com/centos-vault/8.5.2111/</a> and change the address of the EPEL repository for CentOS 8 to <a href="http://mirrors.aliyun.com/epel-archive/8/">http://mirrors.aliyun.com/epel-archive/8/</a>.

#### **Procedure**

Connect to an ECS instance that runs CentOS 8.
 For more information, see Connection methodsGuidelines on instance connection.

2. Run the following command to back up the repo files:

```
rename '.repo' '.repo.bak' /etc/yum.repos.d/*.repo
```

3. Run the following commands to download the latest repo files:

```
wget http://mirrors.cloud.aliyuncs.com/repo/Centos-vault-8.5.2111.repo -0 /etc/yum.repo
s.d/Centos-vault-8.5.2111.repo
wget http://mirrors.cloud.aliyuncs.com/repo/epel-archive-8.repo -0 /etc/yum.repos.d/epe
l-archive-8.repo
```

- Note If your server can access the Internet, replace <a href="http://mirrors.cloud.aliyuncs.co">http://mirrors.cloud.aliyuncs.co</a>
  in CentOS 8 repository addresses with <a href="http://mirrors.aliyun.com">http://mirrors.aliyun.com</a>.
- 4. Run the following commands to replace the addresses in the downloaded repo files:

```
sed -i 's/mirrors.cloud.aliyuncs.com/url_tmp/g' /etc/yum.repos.d/Centos-vault-8.5.2111 .repo && sed -i 's/mirrors.aliyun.com/mirrors.cloud.aliyuncs.com/g' /etc/yum.repos.d/C entos-vault-8.5.2111.repo && sed -i 's/url_tmp/mirrors.aliyun.com/g' /etc/yum.repos.d/C entos-vault-8.5.2111.repo sed -i 's/mirrors.aliyun.com/mirrors.cloud.aliyuncs.com/g' /etc/yum.repos.d/epel-archiv e-8.repo
```

5. Run the following command to re-create a cache:

```
yum clean all && yum makecache
```

After you change the addresses of the YUM and EPEL repositories for CentOS 8, you can use the yum install command to install the software packages that you need on your ECS instance.

#### Related information

- CentOS Product Specifications
- End of support for operating systems

## 13.5. Change the CentOS 6 source address

CentOS 6 has reached its end of life (EOL), and the Linux community no longer maintains this operating system version. We recommend that you upgrade your operating system to CentOS 7 or later. If you still require some of the installation packages of CentOS 6 during the transition, you can change the CentOS 6 source address for your images as instructed in this topic.

#### Context

CentOS 6 has reached its EOL on November 30, 2020. In accordance with the community rules of Linux, the content at the source address <a href="http://mirror.centos.org/centos-6/">http://mirror.centos.org/centos-6/</a> of CentOS 6 has been removed. All third-party image providers have already removed the CentOS 6 source address. The <a href="http://mirrors.cloud.aliyuncs.com">http://mirrors.cloud.aliyuncs.com</a> and <a href="http://mirrors.aliyun.com">http://mirrors.aliyun.com</a> source addresses of Alibaba Cloud cannot synchronize with the CentOS 6 source address. If you continue to use the default CentOS 6 source address on Alibaba Cloud, an error is reported. The following figure shows an example of the error.

```
http://mirrors.cloud.aliyuncs.com/centos/6/os/x86_64/repodata/repomd.xml: [Errno 14] PYCURL ERROR 22 - "The requested URL returned error: 404 Not Found"
Trying other mirror.
http://mirrors.cloud.aliyuncs.com/centos/6/extras/x86_64/repodata/repomd.xml: [Errno 14] PYCURL ERROR 22 - "The requested URL returned error: 404 Not Found"
Trying other mirror.
http://mirrors.cloud.aliyuncs.com/centos/6/updates/x86_64/repodata/repomd.xml: [Errno 14] PYCURL ERROR 22 - "The requested URL returned error: 404 Not Found"
```

You can perform the following procedure to change the source address configurations in the ECS instances that run CentOS 6 operating system based on the network environment.

- YUM reposit ory
  - For VPC-type instances, you must change the source address of YUM repository to <a href="http://mirrors.cloud.aliyuncs.com/centos-vault/6.10/">http://mirrors.cloud.aliyuncs.com/centos-vault/6.10/</a>.
- EPEL reposit ory
  - For VPC-type instances, you must change the source address of EPEL repository to <a href="http://mirrors.com/epel-archive/6/">http://mirrors.com/epel-archive/6/</a>.
  - For classic network-type instances, you must change the source address of EPEL repository to <a href="https://mirrors.aliyuncs.com/epel-archive/6/">http://mirrors.aliyuncs.com/epel-archive/6/</a>.

Note This topic describes how to change the Cent OS 6 source address in ECS instances. If your server is not an ECS instance, make sure that the server is accessible from the Internet and that the <a href="http://mirrors.cloud.aliyuncs.com">http://mirrors.cloud.aliyuncs.com</a> source address is replaced with <a href="http://mirrors.aliyun.com">http://mirrors.aliyun.com</a>. For example, if you change the source address of a YUM repository, replace the source address with <a href="http://mirrors.aliyun.com/centos-vault/6.10/">http://mirrors.aliyun.com/centos-vault/6.10/</a>. If you change the source address of an EPEL repository, replace the source address with <a href="http://mirrors.aliyun.com/epel-archive/6/">http://mirrors.aliyun.com/epel-archive/6/</a>.

#### **Procedure**

Connect to an ECS instance that runs CentOS 6.
 For more information, see Connection methodsGuidelines on instance connection.

Elastic Compute Service Images- FAQ

2. Run the following command to modify the CentOS-Base.repo file:

```
vim /etc/yum.repos.d/CentOS-Base.repo
```

- 3. Press the /key to enter the edit mode. Modify the following content to change the source address. Modify the file based on the network type of the instance.
  - VPC

```
[base]
name=CentOS-6.10
enabled=1
failovermethod=priority
baseurl=http://mirrors.cloud.aliyuncs.com/centos-vault/6.10/os/$basearch/
gpgcheck=1
gpgkey=http://mirrors.cloud.aliyuncs.com/centos-vault/RPM-GPG-KEY-CentOS-6
[updates]
name=CentOS-6.10
enabled=1
failovermethod=priority
baseurl=http://mirrors.cloud.aliyuncs.com/centos-vault/6.10/updates/$basearch/
gpgkey=http://mirrors.cloud.aliyuncs.com/centos-vault/RPM-GPG-KEY-CentOS-6
name=CentOS-6.10
enabled=1
failovermethod=priority
baseurl=http://mirrors.cloud.aliyuncs.com/centos-vault/6.10/extras/$basearch/
gpgcheck=1
gpgkey=http://mirrors.cloud.aliyuncs.com/centos-vault/RPM-GPG-KEY-CentOS-6
```

#### Classic network

```
[base]
name=CentOS-6.10
enabled=1
failovermethod=priority
baseurl=http://mirrors.aliyuncs.com/centos-vault/6.10/os/$basearch/
gpgcheck=1
gpgkey=http://mirrors.aliyuncs.com/centos-vault/RPM-GPG-KEY-CentOS-6
[updates]
name=CentOS-6.10
enabled=1
failovermethod=priority
baseurl=http://mirrors.aliyuncs.com/centos-vault/6.10/updates/$basearch/
gpgkey=http://mirrors.aliyuncs.comm/centos-vault/RPM-GPG-KEY-CentOS-6
[extras]
name=CentOS-6.10
enabled=1
failovermethod=priority
baseurl=http://mirrors.aliyuncs.com/centos-vault/6.10/extras/$basearch/
apacheck=1
gpgkey=http://mirrors.aliyuncs.com/centos-vault/RPM-GPG-KEY-CentOS-6
```

After you modify the file, press the *Esc* key to exit the edit mode and enter :wq to save and exit the file.

4. Run the following command to edit the epel.repo file:

```
vim /etc/yum.repos.d/epel.repo
```

- 5. Press the /key to enter the edit mode. Modify the following content to change the source address. Modify the file based on the network type of the instance.
  - VPC

```
[epel]
name=Extra Packages for Enterprise Linux 6 - $basearch
enabled=1
failovermethod=priority
baseurl=http://mirrors.cloud.aliyuncs.com/epel-archive/6/$basearch
gpgcheck=0
gpgkey=http://mirrors.cloud.aliyuncs.com/epel-archive/RPM-GPG-KEY-EPEL-6
```

o Classic network

```
[epel]
name=Extra Packages for Enterprise Linux 6 - $basearch
enabled=1
failovermethod=priority
baseurl=http://mirrors.aliyuncs.com/epel-archive/6/$basearch
gpgcheck=0
gpgkey=http://mirrors.aliyuncs.com/epel-archive/RPM-GPG-KEY-EPEL-6
```

After you modify the file, press the *Esc* key to exit the edit mode and enter :wq to save and exit the file.

#### What's next

After you change the source addresses of a YUM repository and an EPEL repository, you can use the **yum install** command to install the software packages that you need on the instance.

If you create an ECS instance by using a custom image, when you start the instance, cloud-init automatically initializes the source address configurations of the system. If you want to create a custom image from the ECS instance that has the source address changed and want to retain the original source address configurations, you must perform the following operations to modify the cloud-init configuration file of /etc/cloud/cloud.cfg in the ECS instance that has the source address changed before you create the custom image.

1. Run the following command to edit the /etc/cloud/cloud.cfg file:

```
vim /etc/cloud/cloud.cfg
```

2. Press the /key to enter the edit mode, and use # to comment out the - source-address module under cloud init modules: .

Elastic Compute Service Images·FAQ

The following figure shows the configuration information in the file.

```
# The modules that run in the 'init' stage
cloud_init_modules:
    - migrator
# - source-address
    - seed_random
    - bootcmd
    - write-files
    - growpart
    - resizefs
    - set_hostname
    - update_hostname
    - update_etc_hosts
    - rsyslog
    - users-groups
    - ssh
```

3. After you modify the file, press the *Esc* key to exit the edit mode and enter :wq to save and exit the file.

#### Related information

- CentOS Product Specifications
- End of support for operating systems

## 13.6. End-of-life migration plan for CentOS Linux

The CentOS Project announced the end of life (EOL) for CentOS Linux. If you are a CentOS Linux user, you may be concerned about how to update, maintain, and migrate away from CentOS Linux after EOL. To address these concerns, Alibaba Cloud provides several plans for migrating away from CentOS Linux. This topic describes the plans and some success stories about using these plans.

#### Context

On December 08, 2020, the Cent OS Project announced that it would discontinue Cent OS Linux. For more information, see Cent OS Project shifts focus to Cent OS Stream.

The CentOS Project is shifting focus from CentOS Linux, a rebuild of Red Hat Enterprise Linux (RHEL), to CentOS Stream. This shift affects current CentOS Linux versions in the following ways:

- Cent OS 7 continues to be supported and is scheduled to reach EOL on June 30, 2024.
- CentOS 8 has its lifecycle cut short to December 31, 2021.

#### Impacts of CentOS Linux EOL

According to CentOS Project announcements and as known by Alibaba Cloud, the changes in EOL of CentOS Linux have the following impacts:

- If you are using Cent OS 8, you no longer receive software maintenance or technical support such as bug fixes and feature updates as of January 1, 2022.
- If you are using CentOS 7, you will no longer receive software maintenance or technical support such as bug fixes and feature updates as of July 1, 2024.

#### Migration plans

Alibaba Cloud provides comprehensive plans in collaboration with the Anolis community (OpenAnolis) for you to migrate from CentOS Linux to Anolis OS or Alibaba Cloud Linux, which involve prior migration assessment, migration tools, and long-term support from Alibaba Cloud.

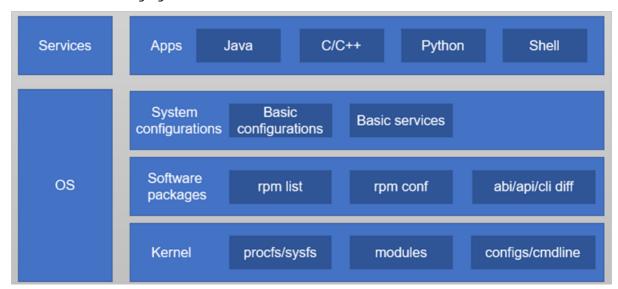
Alibaba Cloud migration plans have the following advantages over other open source migration plans: Alibaba Cloud migration plans incorporate systemized solutions and are backed by support tool platforms and long-term support from Alibaba Cloud. Alibaba Cloud migration plans are fully compatible with the CentOS ecosystem, provide differential technical benefits, and have been verified in a wide array of scenarios and user success practices.

#### Plan details

Each migration plan consists of prior migration assessment and migration implementation.

#### Prior migration assessment

In prior migration assessment, the migration assessment tool scans and retrieves the information of source CentOS Linux operating systems and compares the information with that of destination operating systems to identify migration risks. Both operating systems and services are assessed, as shown in the following figure.



- Prior migration assessment of services: The migration assessment tool recognizes application types and provides migration solutions based on the application types.
- Prior migration assessment of operating systems: The migration assessment tool finds the differences in the kernels, software packages, and system configurations between source and destination operating systems to use as the basis for migration.

#### Migration implementation

Migration can be implemented in new deployment or in-place migration mode.

New deployment

In new deployment mode, a destination operating system is installed and all applications and configurations are re-deployed. This mode is suitable for scenarios such as cluster scaling and node replacement.

• In-place migration

In in-place migration mode, a Cent OS Linux operating system is upgraded in place to Anolis OS or Alibaba Cloud Linux, and all applications, system configurations, and application configurations are retained without the need to re-deploy them or migrate data. This mode is suitable for services that are not affected by underlying operating system changes, such as Java programs. Java Development Kit (JDK) adapts to underlying operating systems. If the pre-migration and post-migration JDK versions are the same, services can run normally without the need to reconfigure or redeploy the services. For example, assume that a Hadoop multi-node cluster is set up based on java-1.8.0-openjdk on Cent OS 8. In an in-place migration, after the cluster is migrated to Anolis OS 8, the configurations of the head node are retained, data storage on data nodes is not affected, and the cluster can work normally.

#### Migration guidelines

Alibaba Cloud provides a specific migration guide for each Cent OS Linux version.

Notice Before you use the migration tool to migrate away from CentOS Linux, note that migration procedures cannot be reversed and that software packages are automatically reinstalled during migration. Before you migrate your Elastic Compute Service (ECS) instances away from CentOS 7 or CentOS 8, we recommend that you create snapshots for the instances to back up data. For more information about snapshots, see Snapshot overview.

- For information about how to migrate from CentOS 7 to Anolis OS 7, see Migrate from CentOS 7 to Anolis OS 7.
- For information about how to migrate from Cent OS 8 to Anolis OS 8, see Migrate from Cent OS 8 to Anolis OS 8.
- For information about how to migrate from CentOS 7 to Alibaba Cloud Linux 2, see Migrate from CentOS 7 to Alibaba Cloud Linux 2.

Note Alibaba Cloud Linux 2 uses kernel version 4.19, which is incompatible with kernel version 3.10 in Cent OS 7. If your business uses the applications that depend on APIs, such as building out-of-tree kernel modules and enabling hardware resources, you must assess these applications beforehand. This way, you can determine whether your applications can run normally after you perform a direct migration.

• For information about how to migrate from Cent OS 8 to Alibaba Cloud Linux 3, see Migrate from Cent OS 8 to Alibaba Cloud Linux 3.

Note Alibaba Cloud Linux 3 uses kernel version 5.10, which is incompatible with kernel version 4.18 of CentOS 8. If your business uses the applications that depend on APIs, such as building out-of-tree kernel modules and enabling hardware resources, you must assess these applications beforehand. This way, you can determine whether your applications can run normally after you perform a direct migration.

#### Success stories

- OpenAnolis helped China Unicom experiment with replacing Cent OS in its core business
- Zhejiang Mobile put Loongson 3A5000 chips and UOS in end-to-end commercial use for the first time in China

# 13.7. Security updates after the end of support and maintenance for Ubuntu 16.04 LTS

In April 2021, Ubuntu 16.04 LTS (Xenial Xerus) reached the end of its standard five-year security maintenance window. Accordingly, Alibaba Cloud no longer provides updates for Ubuntu 16.04 LTS public images. Ubuntu 16.04 LTS transitions into the extended security maintenance (ESM) support phase. You can perform operations described in this topic and use the Ubuntu Advantage for Infrastructure (UA-I) subscription to obtain updates for the ESM phase.

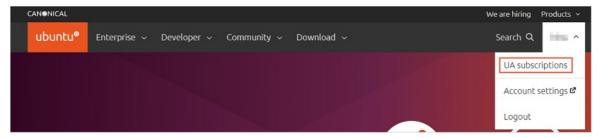
#### Context

While Ubuntu 16.04 LTS is in the ESM phase, you can use ESM in the UA-I subscription to continue to obtain security updates for the Ubuntu base operating systems, critical software packages, and infrastructure components. For more information, see <a href="Ubuntu Advantage for Infrastructure">Ubuntu Advantage for Infrastructure</a> and <a href="Extended Security Maintenance">Extended Security Maintenance</a>.

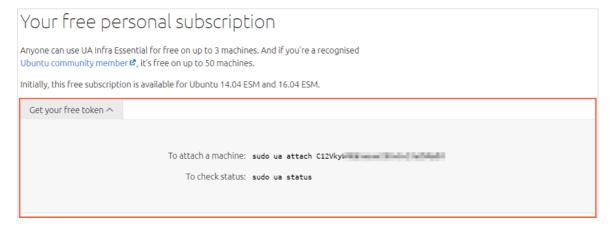
If you want to continue to obtain security updates for the operating system in an Ubuntu 16.04 instance, perform the following steps.

#### Step 1: Subscribe to the UA-I service on the Ubuntu official website

- 1. On your computer, enter <a href="https://ubuntu.com">https://ubuntu.com</a> in a browser to go to the Ubuntu official website.
- 2. In the top navigation bar, click  ${\bf Sign\ in}.$ 
  - You must use an Ubuntu account to log on to the Ubuntu official website. If you do not have an Ubuntu account, you must create one as prompted and then log on to the Ubuntu official website.
- 3. In the top navigation bar, click the username and select **UA** subscriptions.



4. In the Your free personal subscription section, click Get your free token.



You must save the command in the To attach a machine line.

### Step 2: Add UA-I to the Ubuntu 16.04 instance and perform security updates

1. Connect to the Ubuntu 16.04 instance.

For more information, see Connection methodsGuidelines on instance connection.

- 2. Run the following commands in sequence to install the latest Ubuntu Advantage (UA) client.
  - i. Upgrade software packages.

```
sudo apt update
```

ii. Install the UA client.

```
sudo apt install ubuntu-advantage-tools
```

3. Run the command in the To attach a machine line that you saved in Step 1.

Run the command in the following format. Replace <token> with the value that you saved.

```
sudo ua attach <token>
```

A command output similar to the following one indicates that ESM is enabled.

SERVICE cis	ENTITLED yes		DESCRIPTION Center for Internet Security Audit Tools
esm-infra	yes	enabled	UA Infra: Extended Security Maintenance (ESM)
fips fips-updates livepatch	yes yes yes		NIST-certified core packages NIST-certified core packages with priority security updates Canonical Livepatch service

4. (Optional) Run the following command to enable ESM.

If you are not sure whether ESM is enabled, you can run the following command to enable ESM:

```
sudo ua enable esm-infra
```

- 5. Run the following commands in sequence to upgrade software packages and update security patches:
  - i. Upgrade software packages.

```
sudo apt update
```

ii. Update security patches.

```
sudo apt upgrade
```

### 13.8. Share encrypted custom images

If the encryption feature is enabled for disks attached to an Elastic Compute Service (ECS) instance, custom images created from the instance are encrypted. This topic describes how to share encrypted custom images.

#### Context

For more information about the encryption feature, see 加密概述. To share encrypted custom images, you must use Resource Access Management (RAM) to create a RAM role named

AliyunECSShareEncryptImageDefaultRole and attach specific policies to the role.

Note Encrypted custom images can be shared only in the China (Beijing), China (Shanghai), China (Hong Kong), and Singapore (Singapore) regions.

#### **Procedure**

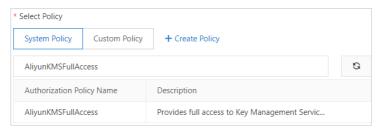
1. Make preparations.

Before you share encrypted custom images with an account, you must know the ID of its Alibaba Cloud account. To obtain the ID of the Alibaba Cloud account, move the pointer over the profile picture in the upper-right corner of the Alibaba Cloud Management Console. If the account is identified as Main Account in the user information panel, the account ID is an Alibaba Cloud account ID.

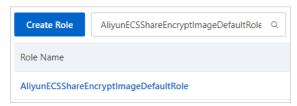
In this example, the ID of the Alibaba Cloud account with which the images are shared is 125\*\*\*\*.

- 2. Log on to the RAM console by using the Alibaba Cloud account of the image owner.
- 3. In the left-side navigation pane, choose **Identities > Roles**.
- 4. Click Create Role. In the Create Role panel, perform the following operations:
  - i. In the Select Trusted Entity section, select Alibaba Cloud Account and click Next.
  - ii. In the Configure Role step, enter AliyunECSShareEncryptImageDefaultRole in the RAM Role Name field, select Current Alibaba Cloud Account in the Select Trusted Alibaba Cloud Account section, and then click OK.
  - iii. In the Finish step, click Add Permissions to RAM Role.
  - iv. In the Add Permissions panel, click System Policy in the Select Policy section and enter liyunKMSFullAccess in the search box.

Use the default configurations for other parameters.



- v. Add the AliyunKMSFullAccess policy and click **OK**. Then, click **Complete**.
- 5. On the Roles page, enter AliyunECSShareEncryptImageDefaultRole in the search box next to Create Role. Then, click the role name to go to the role details page.



6. Modify the trust policy on the role details page.

- i. Click the Trust Policy Management tab.
- ii. Click Edit Trust Policy and replace the default trust policy with the following policy:

<UID> is a variable. Replace it with the ID of the Alibaba Cloud account with which the images are shared. In this example, the ID of the Alibaba Cloud account with which the images are shared is 125\*\*\*\* . The following code shows the modified trust policy:

If you want to share encrypted custom images with multiple Alibaba Cloud accounts, you must attach the trust policy to the accounts. Example policy:

iii. Click OK.

#### What's next

After the AliyunECSShareEncryptImageDefaultRole role is created and configured, you can use the Alibaba Cloud account of the image owner to share encrypted custom images with Alibaba Cloud accounts to which the trust policy is attached. For more information, see Share or unshare a custom image.

## 13.9. Adapt Linux custom images to NVMe-based system disks

Some Elastic Compute Service (ECS) instance families, such as ecs.g7se, ecs.c7se, and ecs.r7se, support disks being attached based on the Non-Volatile Memory Express (NVMe) protocol. If you want to use a custom image to create stable instances of such instance families, make sure that the custom image was created based on an operating system that contains required NVMe-related settings.

#### Context

For information about instance families that support disks being attached based on the NVMe protocol, see Instance family.

Before you configure NVMe-related settings on an instance and then use the instance to create a custom image, we recommend that you create snapshots for the system disk and data disks of the instance to back up data. For more information, see Snapshot overview.

This topic describes how to configure NVMe-related settings in the following operating systems:

- Cent OS
- Debian
- Ubuntu
- SUSE Linux Enterprise Server (SLES)

Cloud Assistant of Alibaba Cloud provides a plug-in that can be used to automatically configure NVMerelated settings in supported operating systems. For more information, see Use the Cloud Assistant plug-in to automatically configure NVMe-related settings.

#### CentOS

- Connect to the CentOS instance from which you want to create a custom image.
   For more information, see Connection methodsGuidelines on instance connection.
- 2. Configure NVMe-related settings based on the configurations of your operating system.

Run the following command to check whether the kernel has loaded the NVMe driver:

```
cat /boot/config-`uname -r` | grep -i nvme | grep -v "^#"
```

A command output similar to the following one is returned. Check whether the value of CONFIG\_BLK\_DEV\_NVME is y, which indicates that the operating system can be started on instance families that support the NVMe protocol. If yes, skip Steps ii and iii and proceed to Step iv. If not, perform the following steps in sequence.

```
[root@test ~]# cat /boot/config-`uname -r` | grep -i nvme | grep -v "^#"
CONFIG NVME CORE=m
CONFIG_BLK_DEV_NVME=m
CONFIG_NVME_MULTIPATH=y
CONFIG_NVME_FABRICS=m
CONFIG_NVME_FABRICS=m
CONFIG_NVME_TOP=m
CONFIG_NVME_TCP=m
CONFIG_NVME_TARGET=m
CONFIG_NVME_TARGET_LOOP=m
CONFIG_NVME_TARGET_FC=m
CONFIG_NVME_TARGET_FC=m
CONFIG_NVME_TARGET_FC=m
CONFIG_NVME_TARGET_FC=m
CONFIG_NVME_TARGET_FCLOOP=m
CONFIG_NVME_TARGET_TCP=m
CONFIG_NVME_TARGET_TCP=m
CONFIG_NVME_TARGET_TCP=m
CONFIG_RTC_NVMEM=y
CONFIG_NVMEM=y
```

ii. Run the following command to check whether the initial RAM file system (initramfs) contains the NVMe driver:

```
lsinitrd /boot/initramfs-`uname -r`.img | grep -i nvme | awk '{print $NF}'
```

Check whether a command output similar to the following one is returned, which indicates that the operating system can be started on instance families that support the NVMe protocol. If yes, skip Step iii and proceed to Step iv. If not, perform the following steps in sequence.

iii. Run the following commands in sequence to configure the initramfs to support the NVMe driver:

```
mkdir -p /etc/dracut.conf.d

echo 'add_drivers+=" nvme nvme-core nvme-fabrics nvme-fc nvme-rdma nvme-loop nvmet
nvmet-fc nvme-tcp "' >/etc/dracut.conf.d/nvme.conf

dracut -v -f

Note If the dracut tool is not installed on your operating system, run the yum -y
```

install dracut command to install the tool.

- iv. Add NVMe-related timeout parameters in GRand Unified Bootloader (GRUB).
  - a. Run the following command to open the grub file:

```
vi /etc/default/grub
```

b. Press the /key to enter the edit mode. On the GRUB\_CMDLINE\_LINUX= line, add the nvme core.io\_timeout and nvme\_core.admin\_timeout parameters and set them both to 4294967295,

as shown in the following figure.

- **Note** If the grub file already contains the preceding parameter settings by default, you do not need to add the parameters again.
- c. Press the *Esc* key to exit the edit mode. Then, enter and close the file.
- v. Apply the GRUB configurations.

Run one of the following commands based on the boot mode of your operating system:

■ Legacy boot mode

```
grub2-mkconfig -o /boot/grub2/grub.cfg
```

■ Unified Extensible Firmware Interface (UEFI) boot mode

```
grub2-mkconfig -o /boot/efi/EFI/centos/grub.cfg
```

vi. Run the following command again to check NVMe-related information:

```
lsinitrd /boot/initramfs-`uname -r`.img | grep -i nvme | awk '{print $NF}'
```

The following command output indicates that the operating system can be started on instance families that support the NVMe protocol.

```
[root@test ~]# lsinitrd /boot/initramfs-`uname -r`.img | grep -i nvme | awk '{print $NF}' usr/lib/modules/4.18.0-193.28.1.el8_2.x86_64/kernel/drivers/nvme
usr/lib/modules/4.18.0-193.28.1.el8_2.x86_64/kernel/drivers/nvme/host
usr/lib/modules/4.18.0-193.28.1.el8_2.x86_64/kernel/drivers/nvme/host/nvme-core.ko.xz
usr/lib/modules/4.18.0-193.28.1.el8_2.x86_64/kernel/drivers/nvme/host/nvme-fabrics.ko.xz
usr/lib/modules/4.18.0-193.28.1.el8_2.x86_64/kernel/drivers/nvme/host/nvme-fc.ko.xz
usr/lib/modules/4.18.0-193.28.1.el8_2.x86_64/kernel/drivers/nvme/host/nvme-core.ko.xz
usr/lib/modules/4.18.0-193.28.1.el8_2.x86_64/kernel/drivers/nvme/host/nvme-fc.ko.xz
usr/lib/modules/4.18.0-193.28.1.el8_2.x86_64/kernel/drivers/nvme/host/nvme-tcp.ko.xz
usr/lib/modules/4.18.0-193.28.1.el8_2.x86_64/kernel/drivers/nvme/target
usr/lib/modules/4.18.0-193.28.1.el8_2.x86_64/kernel/drivers/nvme/target/nvme-loop.ko.xz
usr/lib/modules/4.18.0-193.28.1.el8_2.x86_64/kernel/drivers/nvme/target/nvme-tcp.ko.xz
usr/lib/modules/4.18.0-193.28.1.el8_2.x86_64/kernel/drivers/nvme/target/nvme-tcp.ko.xz
usr/lib/modules/4.18.0-193.28.1.el8_2.x86_64/kernel/drivers/nvme/target/nvme-tcp.ko.xz
usr/lib/modules/4.18.0-193.28.1.el8_2.x86_64/kernel/drivers/nvme/target/nvme-tcp.ko.xz
```

3. Create a custom image from the CentOS instance.

After you configure NVMe-related settings in the operating system of your CentOS instance, you can create a custom image from the instance. For more information about how to create a custom image from an instance, see Create a custom image from an instance.

Note For a custom image that you created in an on-premises environment, you can use the image import feature to import the image to Alibaba Cloud. For more information, see Linux image import procedure.

It takes time to create a custom image. Wait until the custom image is created and enters the **Available** state. Then, proceed to the next step.

4. Create an instance from the custom image that you created.

For more information about how to create an instance from a custom image, see Create an ECS instance by using a custom image. When you create an instance from the custom image, select an instance type that suits your needs and supports disks being attached based on the NVMe protocol.

5. Connect to the new instance.

For more information, see Connection methodsGuidelines on instance connection.

- 6. Run the following commands in sequence on the new instance to check whether the NVMe-related settings work as expected.
  - i. Run the following command to check the information of the disks that support the NVMe protocol:

```
lsblk
```

ii. Run the following command to check whether the /proc/cmdline file contains NVMe-related timeout parameter settings:

```
cat /proc/cmdline
```

The following command output indicates that the custom image is adapted to the NVMe-based system disk.

#### Debian

- 1. Connect to the Debian instance from which you want to create a custom image.

  For more information, see Connection methodsGuidelines on instance connection.
- 2. Configure NVMe-related settings.

i. Run the following command to view the NVMe driver contained in the initrd:
 By default, the initrd in each Debian operating system loads the NVMe driver.

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

A command similar to the following one is returned.

```
root@test:~# lsinitramfs /boot/initrd.img-`uname -r` | grep -i nvme
usr/lib/modules/4.19.0-13-amd64/kernel/drivers/nvme
usr/lib/modules/4.19.0-13-amd64/kernel/drivers/nvme/host
usr/lib/modules/4.19.0-13-amd64/kernel/drivers/nvme/host/nvme-core.ko
usr/lib/modules/4.19.0-13-amd64/kernel/drivers/nvme/host/nvme-fabrics.ko
usr/lib/modules/4.19.0-13-amd64/kernel/drivers/nvme/host/nvme-rdma.ko
usr/lib/modules/4.19.0-13-amd64/kernel/drivers/nvme/host/nvme.ko
usr/lib/modules/4.19.0-13-amd64/kernel/drivers/nvme/target
usr/lib/modules/4.19.0-13-amd64/kernel/drivers/nvme/target/nvmet-fc.ko
usr/lib/modules/4.19.0-13-amd64/kernel/drivers/nvme/target/nvmet-rdma.ko
usr/lib/modules/4.19.0-13-amd64/kernel/drivers/nvme/target/nvmet-rdma.ko
usr/lib/modules/4.19.0-13-amd64/kernel/drivers/nvme/target/nvmet-rdma.ko
usr/lib/modules/4.19.0-13-amd64/kernel/drivers/nvme/target/nvmet.ko
root@test:~#
```

- ii. Add NVMe-related timeout parameters in GRUB.
  - a. Run the following command to open the /etc/default/grub file:

```
vi /etc/default/grub
```

- b. Press the /key to enter the edit mode. On the GRUB\_CMDLINE\_LINUX= line, add the nwme \_core.multipath, nwme\_core.io\_timeout, and nwme\_core.admin\_timeout parameters and set nvme\_core.multipath to n and nvme\_core.io\_timeout and nvme\_core.admin\_timeout both to 4294967295,
  - as shown in the following figure.

```
GRUB_CHROUT=0
GRUB_CHROUT=1
GRUB_DISTRIBUTOR=1
GRUB_DISTRIBUTOR=1
GRUB_CHROUT=1
GRUB_C
```

- **Note** If the grub file already contains the preceding parameter settings by default, you do not need to add the parameters again.
- c. Press the *Esc* key to exit the edit mode. Then, enter :wq and press the *Enter* key to save and close the file.

Elastic Compute Service Images FAQ

iii. Apply the GRUB configurations.

Run one of the following commands based on the boot mode of your operating system:

Legacy boot mode

```
grub-mkconfig -o /boot/grub/grub.cfg
```

■ UEFI boot mode

```
grub-mkconfig -o /boot/efi/EFI/debian/grub.cfg
```

Note Alternatively, you can run the following command to execute the script provided by Debian to apply the GRUB configurations. The boot mode of the operating system is irrelevant to the script. You can use the script regardless of whether the boot mode of your operating system is Legacy or UEFI.

update-grub2

3. Create a custom image from the CentOS instance.

After you configure NVMe-related settings in the operating system of your CentOS instance, you can create a custom image from the instance. For more information about how to create a custom image from an instance, see Create a custom image from an instance.

Note For a custom image that you created in an on-premises environment, you can use the image import feature to import the image to Alibaba Cloud. For more information, see Linux image import procedure.

It takes time to create a custom image. Wait until the custom image is created and enters the **Available** state. Then, proceed to the next step.

4. Create an instance from the custom image that you created.

For more information about how to create an instance from a custom image, see Create an ECS instance by using a custom image. When you create an instance from the custom image, select an instance type that suits your needs and supports disks being attached based on the NVMe protocol.

5. Connect to the new instance.

For more information, see Connection methodsGuidelines on instance connection.

- 6. Run the following commands in sequence on the new instance to check whether the NVMe-related settings work as expected.
  - i. Run the following command to check the information of the disks that support the NVMe protocol:

lsblk

ii. Run the following command to check whether the /proc/cmdline file contains NVMe-related timeout parameter settings:

```
cat /proc/cmdline
```

The following command output indicates that the custom image is adapted to the NVMe-based system disk.

#### Ubuntu

- Connect to the Ubuntu instance from which you want to create a custom image.
   For more information, see Connection methodsGuidelines on instance connection.
- 2. Configure NVMe-related settings.
  - i. Run the following command to check whether the initrd contains the NVMe driver:
     By default, the initrd in each Ubuntu operating system loads the NVMe driver.

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

A command similar to the following one is returned.

```
root@test:~# lsinitramfs /boot/initrd.img-`uname -r` | grep -i nvme
lib/modules/4.15.0-166-generic/kernel/drivers/
lib/modules/4.15.0-166-generic/kernel/drivers/
  /host
  e/host/nvme-fabrics.ko
e/host/nvme-core.ko
e/host/nvme-fc.ko
e/host/nvme-rdma.ko
e/host/nvme.ko
lib/modules/4.15.0-166-generic/kernel/drivers/
lib/modules/4.15.0-166-generic/kernel/drivers/
lib/modules/4.15.0-166-generic/kernel/drivers/r
lib/modules/4.15.0-166-generic/kernel/drivers/
lib/modules/4.15.0-166-generic/kernel/drivers/
  /target
lib/modules/4.15.0-166-generic/kernel/drivers/
  //target/nvmet-fc.ko
e/target/nvme-loop.ko
//target/nvmet.ko
lib/modules/4.15.0-166-generic/kernel/drivers/
lib/modules/4.15.0-166-generic/kernel/drivers/
  e-loop.ko
lib/modules/4.15.0-166-generic/kernel/drivers/
  /target/m
  et.ko
  /target/nvm
lib/modules/4.15.0-166-generic/kernel/drivers/
  t-rdma.ko
root@test:~#
```

- ii. Add NVMe-related timeout parameters in GRUB.
  - a. Run the following command to open the /etc/default/grub file:

```
vi /etc/default/grub
```

b. Press the /key to enter the edit mode. On the GRUB\_CMDLINE\_LINUX= line, add the nvme \_core.multipath, nvme\_core.io\_timeout, and nvme\_core.admin\_timeout parameters and set nvme\_core.multipath to n and nvme\_core.io\_timeout and nvme\_core.admin\_timeout both to 4294967295,

as shown in the following figure.

```
GRUB_DEFAULT=0
GRUB_TIMEOUT_STYLE=hidden
GRUB_TIMEOUT_STYLE=hidden
GRUB_DISTRIBUTOR="15b_release -1 -s 2> /dev/null || echo Debian'
GRUB_DISTRIBUTOR="15b_release -1 -s 2> /dev/null || echo Debian'
GRUB_CHOLINE_LINXX_OFFAULT= quiet_splash'
GRUB_CHOLINE_LINXX_OFFAULT= quiet_splash'
GRUB_CHOLINE_LINXX_OFFAULT= quiet_splash'
GRUB_CHOLINE_LINXX_OFFAULT= quiet_splash'
```

- **Note** If the grub file already contains the preceding parameter settings by default, you do not need to add the parameters again.
- c. Press the *Esc* key to exit the edit mode. Then, enter :wq and press the *Enter* key to save and close the file.
- iii. Apply the GRUB configurations.

Run one of the following commands based on the boot mode of your operating system:

■ Legacy boot mode

```
grub-mkconfig -o /boot/grub/grub.cfg
```

■ UEFI boot mode

```
grub-mkconfig -o /boot/efi/EFI/ubuntu/grub.cfg
```

**Note** Alternatively, you can run the following command to execute the script provided by Ubuntu to apply the GRUB configurations. The boot mode of the operating system is irrelevant to the script. You can use the script regardless of whether the boot mode of your operating system is Legacy or UEFI.

```
update-grub2
```

3. Create a custom image from the CentOS instance.

After you configure NVMe-related settings in the operating system of your CentOS instance, you can create a custom image from the instance. For more information about how to create a custom image from an instance, see Create a custom image from an instance.

Note For a custom image that you created in an on-premises environment, you can use the image import feature to import the image to Alibaba Cloud. For more information, see Linux image import procedure.

It takes time to create a custom image. Wait until the custom image is created and enters the **Available** state. Then, proceed to the next step.

4. Create an instance from the custom image that you created.

For more information about how to create an instance from a custom image, see Create an ECS instance by using a custom image. When you create an instance from the custom image, select an instance type that suits your needs and supports disks being attached based on the NVMe protocol.

5. Connect to the new instance.

For more information, see Connection methodsGuidelines on instance connection.

- 6. Run the following commands in sequence on the new instance to check whether the NVMe-related settings work as expected.
  - i. Run the following command to check the information of the disks that support the NVMe protocol:

```
lsblk
```

ii. Run the following command to check whether the /proc/cmdline file contains NVMe-related timeout parameter settings:

```
cat /proc/cmdline
```

The following command output indicates that the custom image is adapted to the NVMe-based system disk.

#### SUSE Linux Enterprise Server (SLES)

- Connect to the SLES instance from which you want to create a custom image.
   For more information, see Connection methodsGuidelines on instance connection.
- 2. Configure NVMe-related settings based on the configurations of your operating system.

i. Run the following command to check whether the kernel has loaded the NVMe driver:

```
cat /boot/config-`uname -r` | grep -i nvme | grep -v "^#"
```

A command output similar to the following one is returned. Check whether the value of G\_BLK\_DEV\_NVME is y, which indicates that the operating system can be started on instance families that support the NVMe protocol. If yes, skip Steps ii and iii and proceed to Step iv. If not, perform the following steps in sequence.

```
sles125:~ # cat /boot/config-`uname -r` | grep -i nvme | grep -v "^#"
CONFIG_NVME_CORE=m
CONFIG_BLK_DEV_NVME=m
CONFIG_NVME_MULTIPATH=y
CONFIG_NVME_FABRICS=m
CONFIG_NVME_FC=m
CONFIG_NVME_FC=m
CONFIG_NVME_TCP=m
CONFIG_NVME_TARGET=m
CONFIG_NVME_TARGET_LOOP=m
CONFIG_NVME_TARGET_RDMA=m
CONFIG_NVME_TARGET_FC=m
CONFIG_NVME_TARGET_FC=m
CONFIG_NVME_TARGET_FC=m
CONFIG_NVME_TARGET_FC=m
CONFIG_NVME_TARGET_FCLOOP=m
CONFIG_NVME_TARGET_FCLOOP=m
CONFIG_NVME_TARGET_TCP=m
CONFIG_NVME_TARGET_TCP=m
```

ii. Run the following command to check whether the initrd contains the NVMe driver:

```
lsinitrd /boot/initrd-`uname -r` | grep -i nvme | awk '{print $NF}'
```

Check whether a command output similar to the following one is returned, which indicates that the operating system can be started on instance families that support the NVMe protocol. If yes, skip Step iii and proceed to Step iv. If not, perform the following steps in sequence.

```
sles125:~ # lsinitrd /boot/initrd-`uname -r` | grep -i nvme | awk '{print $NF}' lib/modules/4.12.14-122.80-default/kernel/drivers/nvme lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/host lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/host/nvme-core.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/host/nvme-fabrics.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/host/nvme-fc.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/host/nvme-ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/host/nvme-rdma.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/host/nvme-tcp.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/target lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/target/nvme-loop.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/target/nvmet-fc.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/target/nvmet-fc.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/target/nvmet-fc.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/target/nvmet-fc.ko
```

iii. Run the following commands in sequence to configure the initrd to support the NVMe driver:

```
mkdir -p /etc/dracut.conf.d

echo 'add_drivers+=" nvme nvme-core nvme-fabrics nvme-fc nvme-rdma nvme-loop nvmet
nvmet-fc nvme-tcp "' >/etc/dracut.conf.d/nvme.conf

dracut -v -f

? Note If the dracut tool is not installed on your operating system, run the zypper
```

- Note If the dracut tool is not installed on your operating system, run the zypper install -y dracut command to install the tool.
- iv. Add NVMe-related timeout parameters in GRUB.
  - a. Run the following command to open the grub file:

```
vi /etc/default/grub
```

b. Press the /key to enter the edit mode. On the GRUB\_CMDLINE\_LINUX= line, add the nvme core.io\_timeout and nvme\_core.admin\_timeout parameters and set them both to 4294967295,

as shown in the following figure.

```
GRUB_DISTRIBUTOR=

GRUB_DEFAULT=saved

GRUB_HIDDEN_TIMEOUT=0

GRUB_HIDDEN_TIMEOUT_QUIET=true

GRUB_TIMEOUT=1

GRUB_CMDLINE_LINUX_DEFAULT="mitigations=auto splash=silent quiet showopts"

GRUB_CMDLINE_LINUX_=" net.ifnames=0 console=tty8 console=tty50,115200n8 nvme_core.io_timeout=4294967295 nvme_core.admin_timeout=4294967295"
```

- **? Note** If the grub file already contains the preceding parameter settings by default, you do not need to add the parameters again.
- c. Press the *Esc* key to exit the edit mode. Then, enter :wq and press the *Enter* key to save and close the file.
- v. Apply the GRUB configurations.

Run one of the following commands based on the boot mode of your operating system:

Legacy boot mode

```
grub2-mkconfig -o /boot/grub2/grub.cfg
```

■ UEFI boot mode

```
grub2-mkconfig -o /boot/efi/EFI/sles/grub.cfg
```

vi. Run the following command again to check NVMe-related information:

```
lsinitrd /boot/initrd-`uname -r` | grep -i nvme | awk '{print $NF}'
```

The following command output indicates that the operating system can be started on instance families that support the NVMe protocol.

```
sles125:~ # lsinitrd /boot/initrd-`uname -r` | grep -i nvme | awk '{print $NF}' lib/modules/4.12.14-122.80-default/kernel/drivers/nvme lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/host lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/host/nvme-core.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/host/nvme-fabrics.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/host/nvme-fc.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/host/nvme.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/host/nvme-rdma.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/host/nvme-tcp.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/target lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/target/nvme-loop.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/target/nvmet-fc.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/target/nvmet-fc.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/target/nvmet-fc.ko lib/modules/4.12.14-122.80-default/kernel/drivers/nvme/target/nvmet-fc.ko
```

3. Create a custom image from the CentOS instance.

After you configure NVMe-related settings in the operating system of your CentOS instance, you can create a custom image from the instance. For more information about how to create a custom image from an instance, see Create a custom image from an instance.

**?** Note For a custom image that you created in an on-premises environment, you can use the image import feature to import the image to Alibaba Cloud. For more information, see Linux image import procedure.

It takes time to create a custom image. Wait until the custom image is created and enters the **Available** state. Then, proceed to the next step.

4. Create an instance from the custom image that you created.

For more information about how to create an instance from a custom image, see Create an ECS instance by using a custom image. When you create an instance from the custom image, select an instance type that suits your needs and supports disks being attached based on the NVMe protocol.

5. Connect to the new instance.

For more information, see Connection methodsGuidelines on instance connection.

- 6. Run the following commands in sequence on the new instance to check whether the NVMe-related settings work as expected.
  - i. Run the following command to check the information of the disks that support the NVMe protocol:

```
lsblk
```

ii. Run the following command to check whether the /proc/cmdline file contains NVMe-related timeout parameter settings:

```
cat /proc/cmdline
```

The following command output indicates that the custom image is adapted to the NVMe-based system disk.

## Use the Cloud Assistant plug-in to automatically configure NVMerelated settings

Cloud Assistant provides the ecs\_nvme\_config plug-in to help you configure NVMe-related settings in operating systems. The ecs\_nvme\_config plug-in supports the following operations systems:

- Alibaba Cloud Linux
- Anolis OS
- CentOS/CentOS Stream
- Debian
- Ubunt u
- OpenSUSE
- SUSE Linux Enterprise Server
- Red Hat Enterprise Linux
- Fedora
- Rocky Linux
- AlmaLinux

The ecs name config plug-in does not support the following operating system versions:

- CentOS and Red Hat Enterprise Linux: versions earlier than CentOS 6.6 and Red Hat Enterprise Linux 6.6
- Debian: versions earlier than Debian 9
- Ubuntu: versions earlier than Ubuntu 16
- openSUSE: versions earlier than openSUSE 42
- SLES: versions earlier than SLES 11.4
  - 1. Connect to the instance from which you want to create a custom image.

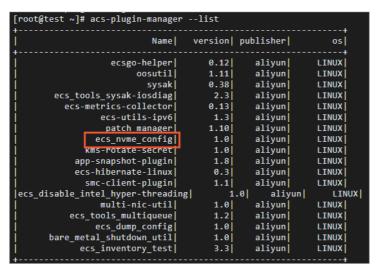
For more information, see Connection methodsGuidelines on instance connection.

2. Run the following command to check whether the Cloud Assistant client is installed on the instance and whether the installed Cloud Assistant client provides the ecs\_nvme\_config plug-in.

```
acs-plugin-manager --list
```

The following command output indicates that the Cloud Assistant client is installed. If the Cloud Assistant client is not installed on the instance, the command output contains no information about the client. In this case, you must install the Cloud Assistant client first. For more information,

#### see Install the Cloud Assistant client.



If you find in the command output a list of Cloud Assistant plug-ins that exclude the ecs\_nvme\_co nfig plug-in, the plug-in is unavailable in the current region and zone. In this case, you must use a method applicable to your operating system to manually configure NVMe-related settings.

- Cent OS
- o Debian
- o Ubuntu
- SUSE Linux Enterprise Server (SLES)

If the Cloud Assistant client is installed and provides the ecs\_nvme\_config plug-in on your instance, proceed to the next step.

- 3. Use the ecs nyme config plug-in to configure NVMe-related settings.
  - i. Run the following command to use the ecs\_nvme\_config plug-in to check whether the instance has the NVMe module and whether the instance supports the NVMe module:

```
acs-plugin-manager --exec --plugin ecs_nvme_config --params --check
```

Command output description:

If the following SUCCESS message appears in the command output, proceed to create a custom image without the need to configure NVMe-related settings:

```
[SUCCESS] Summary: Your image can Runnig on nvme instance
```

- If an [ERROR] message appears in the command output, proceed to the next step.
- ii. Configure the NVMe driver and relevant parameters based on the error message in the command output.

Example error message:

```
[ERROR] 1.initrd/initramfs not has nvme module, Please run acs-plugin-manager --ex ec --plugin ecs_nvme_config --params -f/--fix to enable nvme;
```

Run the following command to configure NVMe-related settings based on the error message:

```
acs-plugin-manager --exec --plugin ecs_nvme_config --params --fix
```

iii. After the NVMe-related settings are configured, run the following command to restart your instance:

reboot

iv. After the instance is restarted, run the following command to check whether the NVMe module is configured:

```
acs-plugin-manager --exec --plugin ecs_nvme_config --params --check
```

#### Example success output:

```
[OK] 1.initrd/initramfs already contain nvme module;

[OK] 2.fstab file looks fine and does not contain any device names;

[OK] 3.The nvme parameters already included.

[SUCCESS] Summary: Your image can Runnig on nvme instance
```

4. Create a custom image from the Cent OS instance.

After you configure NVMe-related settings in the operating system of your CentOS instance, you can create a custom image from the instance. For more information about how to create a custom image from an instance, see Create a custom image from an instance.

**Note** For a custom image that you created in an on-premises environment, you can use the image import feature to import the image to Alibaba Cloud. For more information, see Linux image import procedure.

It takes time to create a custom image. Wait until the custom image is created and enters the **Available** state. Then, proceed to the next step.

5. Create an instance from the custom image that you created.

For more information about how to create an instance from a custom image, see Create an ECS instance by using a custom image. When you create an instance from the custom image, select an instance type that suits your needs and supports disks being attached based on the NVMe protocol.

6. Connect to the new instance.

For more information, see Connection methodsGuidelines on instance connection.

- 7. Run the following commands in sequence on the new instance to check whether the NVMe-related settings work as expected.
  - i. Run the following command to check the information of the disks that support the NVMe protocol:

lsblk

ii. Run the following command to check whether the /proc/cmdline file contains NVMe-related timeout parameter settings:

```
cat /proc/cmdline
```

The following command output indicates that the custom image is adapted to the NVMe-based system disk.

# 13.10. Upgrade the SUSE Linux Enterprise Server public images of existing instances to the commercial edition

Starting from October 30, 2021, Alibaba Cloud SUSE Linux Enterprise Server (SLES) public images are available for commercial use and charge fees. As of April 1, 2022, if you have Elastic Compute Service (ECS) instances that use free SLES images, you must upgrade the SLES images of these instances to the commercial edition as soon as possible before you can obtain image updates and technical support for these images. You can perform the steps described in this topic to upgrade your SLES images to the commercial edition.

#### Context

Existing instances that run SLES operating systems are instances that were created from free SLES public images or created from your free custom SLES images. You can upgrade the SLES images of these instances to the commercial edition in the ECS console. After you upgrade the SLES images of instances to the commercial edition, you can continue to obtain image updates and technical support for these images. If you do not upgrade the SLES images of instances to the commercial edition, the instances can continue to provide services. As of April 1, 2022, image updates and technical support for SUSE are unavailable.

For the announcement on commercial availability of SLES public images, see SUSE Linux Enterprise Server will be commercialized on October 30, 2021.

#### **Precautions**

- If you upgrade an SLES image to the commercial edition, you are charged for the image based on the price at the time of upgrade.
- If you do not want to continue using SLES images, we recommend that you use a free Alibaba Cloud Linux operating system to replace SLES.
- SLES images become available for commercial use but SLES for SAP images have not
- The SLES images of existing instances that reside in the classic network cannot be upgraded to the commercial edition. If you have instances that reside in the classic network and the SLES images of these instances must be upgraded to the commercial edition, migrate the instances to virtual private

clouds (VPCs) and then upgrade the SLES images. For information about how to migrate instances from the classic network to VPCs, see Migrate ECS instances from the classic network to a VPC. If you encounter issues, submit a ticket.

#### **Procedure**

If you have instances created from free SLES images that need to be upgraded to the commercial edition, you can receive the notification about the commercial availability of SLES images on the Instances page or on the Images page when you log on to the ECS console. You can upgrade the SLES images as needed.

Before you upgrade your SLES images, you must use your Alibaba Cloud account to authorize Alibaba Cloud to run the register command after you pay for the upgrade order to obtain the SLES update repository. Perform the following operations:

- 1. Use your Alibaba Cloud account to log on to the Resource Access Management (RAM) console.
- 2. Create a RAM role named **AliyunECSAssistantRunCommandRole** and change the trust policy of the role to the following content.

**?** Note The role name must be AliyunECSAssistant RunCommand Role. Otherwise, the register command cannot be run.

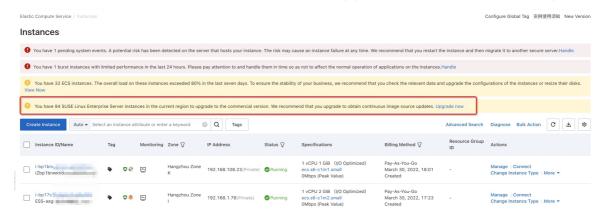
For more information about how to create a RAM role and change its trust policy, see Attach an instance RAM role to an ECS instance.

3. Attach the **AliyunECSAssist ant RunCommandRolePolicy** policy to the role. The policy contains the following content:

For more information about how to attach the policy to the RAM role, see Control access to resources by using RAM users.

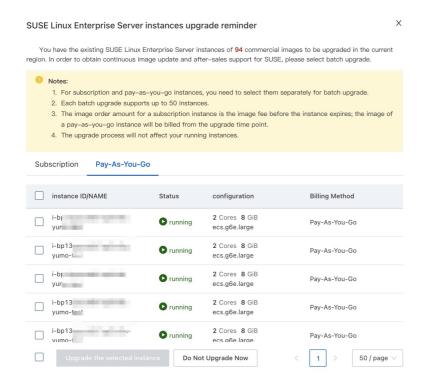
4.

In the left-side navigation pane, choose Instances & Images > Instances.
 In the section above the instance list, you can view the upgrade notification about SLES images.



- 6. Click Upgrade Now.
- 7. Upgrade the SLES images to the commercial edition as prompted.

You can also choose not to upgrade the images to the commercial edition.



After you place an upgrade order and complete the payment, Alibaba Cloud is authorized to call the RunCommand operation to run the register command. The SLES images can be upgraded and SUSE operating systems can be registered. The registration process takes several seconds. We recommend that you wait and then go to the Instances page, refresh the page and confirm the number of remaining instances whose SLES images must be upgraded, and then upgrade the images.

After SLES images are upgraded, you can obtain the SLES update repository and technical support for SUSE.

#### 13.11. Linux time and time zones

This topic describes the terms and operations related to time and time zones in Linux operating systems, and describes the changes in clock configurations in some Alibaba Cloud Linux public images.

#### Context

The following table describes the terms related to time and time zones in Linux operating systems.

Term	Description	
clock	<ul> <li>Linux has the following clocks:</li> <li>Hardware clock, also known as real-time clock (RTC): the battery-powered clock on the physical motherboard, which can be set in Linux BIOS.</li> <li>System clock: the clock in the Linux kernel. When a Linux operating system starts, the hardware clock first reads the settings of the system clock. After that, the system clock runs independently of the physical hardware. All commands and functions in Linux operating systems use the system clock.</li> </ul>	
time standard	<ul> <li>Linux has the following time standards:</li> <li>localtime: the time standard based on the local time zone where the operating system resides.</li> <li>Coordinated Universal Time (UTC): a global time standard that is independent of time zones.</li> </ul>	

#### Time-related operations

• Run the following command to view the time settings in the Linux operating system:

```
timedatectl status
```

A command output similar to the following one is returned. You can view the local time, UTC time, RTC time, and time zone of the Linux operating system.

```
root@test:"# timedatectl status

Local time: Fri 2022-01-21 17:56:58 CST

Universal time: Fri 2022-01-21 09:56:58 UTC

RTC time: Fri 2022-01-21 09:56:58

Time zone: Asia/Shanghai (CST, +0800)

System clock synchronized: yes

NTP service: active

RTC in local TZ: no

root@test: #
```

- Run the timedatectl command to set the time standard of the hardware clock.
  - o Set the time standard of the hardware clock to local time:

```
timedatectl set-local-rtc 1
```

If the time standard of the hardware clock is localtime, you can find that the value of  $_{\rm RTC}$  in  $_{\rm loc}$  al  $_{\rm TZ}$  in the time settings is  $_{\rm yes}$  after you run the timedatectl status command.

Elastic Compute Service Images·FAQ

Set the time standard of the hardware clock to UTC:

```
timedatectl set-local-rtc 0
```

If the time standard of the hardware clock is UTC, you can find that the value of RTC in local TZ in the time settings is no after you run the timedatectl status command.

**Note** The /etc/adjtime configuration file is used to synchronize the system clock. When you run the preceding commands, the /etc/adjtime file is automatically updated.

## Public images in which the Alibaba Cloud RTC uses the UTC time standard

By default, RTC in Alibaba Cloud public images uses the localtime standard. Note that Alibaba Cloud has changed the localtime standard used by RTC by default to the UTC time standard. In the following public images and their later versions, RTC uses the UTC time standard. For more information about image releases, see the following topics:

•

- Alibaba Cloud Linux 3镜像发布记录
- Release notes for Alibaba Cloud Linux 2

### 13.12. Compatibility of AMD instances

This topic describes the versions of the operating system images that are compatible with AMD instances.

#### **Background information**

AMD instances are new products released by Alibaba Cloud, which are powered by the AMD EPYC TM Zen 2 x86 microprocessor (code: Rome). AMD instances provide predictable and consistent high performance and reduce virtualization overhead. The following AMD instance families are available:

- g6a, general purpose instance family
- c6a, compute optimized instance family
- r6a, memory optimized instance family

For more information, see Instance family.

AMD instances are compatible with different operating system versions. The AMD Zen architecture was released in 2017. Some new features of the processor on the AMD Zen architecture are not supported by earlier operating system versions. You can determine the operating system version that you need to use based on your applications. The following table lists the recommended versions for mainstream operating systems.

#### Operating system versions that support AMD Rome processors

The following table lists the operating system versions that support AMD Rome processors.

Operating system	Operating system version	Official documentation	
------------------	--------------------------	------------------------	--

Operating system	Operating system version	Official documentation
Windows	<ul><li>Windows Server 2012 R2</li><li>Windows Server 2016</li><li>Windows Server 2019</li></ul>	Windows Server support and installation instructions for the AMD Rome family of processors
Linux	<ul> <li>Alibaba Cloud Linux 2</li> <li>Cent OS 7.6/7.7/7.8/7.9/8.0/8.1/8.2/8.3</li> <li>Ubuntu 16.04/18.04/20.04</li> </ul>	<ul> <li>Red Hat Enterprise Linux: AMD CPUs and Supported Red Hat Enterprise Linux (RHEL) Versions</li> <li>Ubuntu: AMD EPYC Rome support in Ubuntu Server</li> </ul>

#### Related information

•

## 13.13. Common image formats

ECS allows you to import images in multiple formats. This topic compares the QCOW2, VHD, and RAW formats supported by ECS.

The following table describes the common image formats.

Format	Description	Feature
QCOW2	QCOW2 is a virtual machine image format implemented by QEMU, which can represent a fixed-size block device-type disk in the form of a file.	<ul> <li>Occupies a smaller amount of disk space.</li> <li>Supports copy-on-write (CoW). Image files reflect only the underlying disk changes.</li> <li>Supports snapshots and can contain multiple snapshots.</li> <li>Supports compression and encryption. You can compress data by using zlib and encrypt data by using AES.</li> </ul>

Format	Description	Feature
VHD	VHD is a virtual disk file format provided by Microsoft. Files in the VHD format can be compressed into a single file and stored on the file system of a physical host. It includes the file systems required to start ECS instances.	<ul> <li>Simple maintenance. You can partition, format, compress, and delete physical partitions without affecting them.</li> <li>Easy backup. You can back up only the created VHD files. You can also use a backup tool to back up the entire physical partition where the VHD files reside.</li> <li>Easy migration. If you want to use a VHD file on multiple computers, you can detach the VHD file, copy it to the destination computer, and then attach it.</li> <li>Availability for system deployment. You can use the Imagex tool to release captured images to VHD virtual disks or use Windows Deployment Services (WDS) servers to deploy the system to VHD virtual disks.</li> </ul>
RAW	RAW files can be directly read and written by ECS instances. RAW is a format that provides the optimal performance and does not allow the space to dynamically grow.	<ul> <li>Simple addressing and high access efficiency.</li> <li>You can convert RAW to other formats by using format conversion tools.</li> <li>RAW files can be mounted by physical hosts. You can perform data transfer between RAW files and hosts without the need to start virtual machines.</li> </ul>

#### 13.14. Install GRUB on a Linux server

To use Server Migration Center (SMC) to migrate a Linux server to Alibaba Cloud, if the system boot program GRand Unified Bootloader (GRUB) is not installed on the server, you must install GRUB v2.02 or later. If the server runs an early Linux distribution such as CentOS 5 or Debian 7, if the GRUB version is earlier than 2.02, or if the "Do Grub Failed" error message is displayed in the log file, you must upgrade GRUB to v2.02 or later.

#### Context

This topic describes how to install GRUB on a Linux server. GRUB v2.02 is used in the example. The procedure to install GRUB of other versions is similar. The source code package may vary with GRUB versions. For more information, visit Index of /gnu/grub.

#### Procedure

- 1. Log on to the Linux server.
- 2. Run the following commands to check the paths of the grub, grub-install, and grub-mkconfig files of the current GRUB version:

which grub-install

```
which grub-mkconfig
```

If the outputs of preceding commands indicate that one or more of these paths do not exist,
 GRUB is not installed on the server or the corresponding files are missing. You must perform operations described in the following section to install GRUB.

• If you can view the paths of all the files, run the following commands to back up the grub, grub-install, and grub-mkconfig files by renaming them. When you install a new GRUB version, the new version overwrites the current version.

```
mv /sbin/grub /sbin/grub-old
mv /sbin/grub-install /sbin/grub-install-old
mv /sbin/grub-mkconfig /sbin/grub-mkconfig-old
```

- **Note** After you use SMC to migrate the server, you can restore the files by changing their names back to the original ones.
- 3. Install the GRUB dependencies including bison, gcc, and make.

```
yum install -y bison gcc make
```

- 4. Perform the following operations to install flex:
  - i. Check whether the tools folder exists. If the folder does not exist, create it.

```
test -d /root/tools || mkdir -p /root/tools
```

ii. Go to the tools folder and download the flex installation package.

```
cd /root/tools
wget https://github.com/westes/flex/releases/download/v2.6.4/flex-2.6.4.tar.gz
```

iii. Decompress the flex installation package.

```
tar xzf flex-2.6.4.tar.gz
```

iv. Go to the directory to which the flex installation package is decompressed, and create a folder named build.

```
cd flex-2.6.4 mkdir -p build
```

v. Go to the build folder, and compile and install flex.

```
cd build
../configure

make && make install
```

vi. Create the symbolic link.

```
ln -s /usr/local/bin/flex /usr/bin/flex
```

5. Perform the following operations to install GRUB.

You must use GRUB v2.02 or later for earlier distributions of operating systems such as CentOS 5, Red Hat Enterprise Linux 5, Debian 7, Amazon Linux, and Oracle Linux.

i. Check whether the tools folder exists. If the folder does not exist, create it.

```
test -d /root/tools || mkdir -p /root/tools
```

ii. Go to the tools folder and download the GRUB v2.02 installation package.

```
cd /root/tools
wget https://alpha.gnu.org/gnu/grub/grub-2.02~rc1.tar.gz
```

iii. Decompress the GRUB v2.02 installation package.

```
tar xzf grub-2.02~rc1.tar.gz
```

iv. Go to the directory to which the GRUB v2.02 installation package is decompressed, and create a folder named build.

```
cd grub-2.02~rc1
mkdir -p build
```

v. Go to the build folder, and compile and install GRUB.

```
cd build
../configure

sed -i -e "s/-Werror//" ./grub-core/Makefile
sed -i -e "s/-Werror//" . /Makefile
make && make install
```

vi. Create the symbolic links.

```
ln -s /usr/local/sbin/grub-install /sbin/grub-install
ln -s /usr/local/sbin/grub-mkconfig /sbin/grub-mkconfig
```

```
Note If the —Werror error is reported during compilation, find the makefile compile file, remove the —Werror option from the file, and then try again.
```

6. Run the following command to check whether GRUB v2.02 is installed or GRUB is upgraded to v2.02:

```
grub-install --version
```

#### What's next

- If GRUB v2.02 is installed or GRUB is upgraded to v2.02, you can use SMC to migrate the server to Alibaba Cloud. For more information, see Migration process.
- Optional. After the sever is migrated, if you want to use the previous version of GRUB, run the following commands to restore GRUB to its previous version:

```
rm /sbin/grub-install
rm /sbin/grub-mkconfig
rm /boot/grub/grub.cfg
mv /sbin/grub-old /sbin/grub
mv /sbin/grub-install-old /sbin/grub-install
```

## 13.15. Manage Windows Server Semi-Annual Channel images and instances

This topic describes how to manage an Elastic Compute Service (ECS) instance that is created from a Windows Server Semi-Annual Channel image.

#### Context

Windows Server Semi-Annual Channel runs in Server Core mode and has no graphical user interface (GUI). Windows Server Semi-Annual Channel offers advantages such as support for remote management, lower requirements for hardware, and reduced need for updates. Windows Server Semi-Annual Channel instances exclude Resource Manager, Control Panel, and Windows Explorer and do not support the \\*.msc command line option such as devmgmt.msc. You can manage these instances by using tools such as Sconfig, Server Manager, PowerShell, and Windows Admin Center.

When you create an instance, you can view the following Windows Server Semi-Annual Channel images in the public image list:

- Windows Server Version 1809 Dat acenter Edition
- Windows Server Version 1709 Dat acenter Edition
- Windows Server Version 1903 Dat acenter Edition
- Windows Server Version 1909 Dat acenter Edition
- Windows Server Version 2004 Dat acent er Edit ion

Windows Server Semi-Annual Channel runs in Server Core mode. We recommend that you use advanced management tools such as PowerShell and Windows Admin Center. For more information, see Manage a Server Core server in Microsoft Docs.

#### Manage an instance by using PowerShell

PowerShell runs on .NET Framework and uses object-oriented scripts. This allows you to manage Windows instances in the same manner as you would with SSH. For example, assume that the public IP address of your Windows instance is 172.16.1XX.183. You can perform the following steps to manage your instance by using PowerShell.

- 1. Connect to the Windows instance. For more information, see Connect to a Windows instance by using a username and password.
- 2. Enter PowerShell on the command line to start PowerShell.
- 3. Run the following commands in PowerShell of the instance:

```
Enable-PSRemoting -Force

Set-NetFirewallRule -Name "WINRM-HTTP-In-TCP-PUBLIC" -Remot
eAddress Any
```

- 4. Add rules to a security group to which the instance belongs to allow access over HTTP port 5985 and HTTPS port 5986. For more information about how to add rules to a security group, see Add a security group rule.
- 5. Enter PowerShell on the command line on your computer to start PowerShell.
- 6. Run the following command in PowerShell on your computer:

```
Set-Item WSMan:localhost\client\trustedhosts -value
172.16.1XX.183 -Force

Note 172.16.1XX.183 indicates that only your instance is trusted. You can use * to indicate that all computers are trusted.
```

7. Run Enter-PSSession '172.16.1XX.183' -Credential: 'administrator' in PowerShell and enter the password of the instance as prompted.

Now you can manage your Windows instances on your computer.

#### Install Windows Admin Center

Windows Admin Center is a browser-based GUI management tool. It can replace existing server management tools like Server Manager and Microsoft Management Console (MMC) to manage Server Core servers. For example, assume that the public IP address of your Windows instance is 172.16.1XX.183. You can use one of the following methods to install Windows Admin Center:

- Use commands
  - i. Connect to the Windows instance. For more information, see Connect to a Windows instance by using a username and password.
  - ii. Add rules to a security group to which the instance belongs to allow access over HTTP port 5985 and HTTPS port 5986. For more information, see Add a security group rule.
  - iii. Enter PowerShell on the command line to start PowerShell.
  - iv. Run the following commands in PowerShell of the instance:

```
Enable-PSRemoting -Force

Set-NetFirewallRule -Name "WINRM-HTTP-In-TCP-PUBLIC"
-RemoteAddress Any
```

v. Run the following commands to download Windows Admin Center:

```
wget -Uri

http://download.microsoft.com/download/E/8/A/E8A26016

-25A4-49EE-8200-E4BCBF292C4A/HonoluluTechnicalPreview1802.msi

-UseBasicParsing -OutFile c:\HonoluluTechnicalPreview

1802.msi

msiexec /i c:\HonoluluTechnicalPreview1802.msi /qn /L

*v log.txt SME_PORT=443

SSL_CERTIFICATE_OPTION=generate
```

vi. Run the *cat log.txt* command to check the download progress.

When information similar to the following content is displayed in the log file, Windows Admin Center is installed:

```
MSI (s) (14:44) [09:48:37:885]: Product: Project

'Honolulu' (Technical Preview) -- Installation co

mpleted successfully.

MSI (s) (14:44) [09:48:37:885]: Product installed

by Windows Installer. Product

name: Project 'Honolulu' (Technical Preview). Pro

duct version: 1.1.10326.0. Product

language: 1033. Manufacturer: Microsoft Corporati

on. Installation success or error

status: 0.
```

- Use a browser
  - o Prerequisit es

PowerShell is configured and can be used to manage instances. For more information, see Manage an instance by using PowerShell section. You must install Windows Admin Center on your computer.

- Procedure
  - a. Download Windows Admin Center and install it. For more information, see .
  - b. Access https://localhost/after you install Windows Admin Center.
  - c. Click Add. In the dialog box that appears, add the IP address of the instance.

Now you can use Windows Admin Center to manage instances in Microsoft Edge or Chrome.

#### **FAQ**

Question 1: How do I copy files to a Windows Server Semi-Annual Channel instance?

If the files to be copied are stored on your computer, and Windows Admin Center is installed, or PowerShell is configured to manage instances, you can use one of the following methods to copy files to the instance:

- Use a Remote Desktop (RDP) application
  - i. Connect to the Windows instance. For more information, see Connect to a Windows instance by using a username and password.
  - ii. Copy files from your computer.
  - iii. Enter notepad on the command line of the instance.
  - iv. Choose File > Open . In the dialog box that appears, right-click the destination directory and select Paste.
- Use PowerShell
  - i. Start the Windows instance.
  - ii. Enter PowerShell on the command line on your computer to start PowerShell.
  - iii. Use PowerShell to manage the instance. For more information, see Manage an instance by using PowerShell.
  - iv. Run the following commands on your computer:

```
$session = New-PSSession -ComputerName

172.16.1XX.183

Copy-Item -ToSession $session -Path C:\1.txt -Des

tination c:\2.txt
```

**?** Note The path to the file to be copied (original file) is C:1.txt on your computer. The path to the file copy (new file) on the Windows instance is C:1.txt.

- Use Windows Admin Center
  - i. Start the Windows instance.
  - ii. Configure Windows Admin Center. For more information, see Install Windows Admin Center.
  - iii. Start Windows Admin Center and click the Windows instance. Click File, select the file, and then click Upload.

Question 2: How do I stop or restart a Windows Server Semi-Annual Channel instance by connecting to the Windows instance?

- Use an RDP application
  - i. Connect to the Windows instance. For more information, see Connect to a Windows instance by using a username and password.
  - ii. Enter *sconfig* on the command line, select 13 to restart the instance or 14 to stop the instance, and then press the Enter key.
- Use PowerShell
  - i. Connect to the Windows instance. For more information, see Connect to a Windows instance by using a username and password.
  - ii. Enter PowerShell on the command line to start PowerShell.
  - iii. Run one of the following commands to restart or stop the instance:

- Use PowerShell
  - i. Start the Windows instance.
  - ii. Enter PowerShell on the command line on your computer to start PowerShell.
  - iii. Use PowerShell to manage the instance. For more information, see Manage an instance by using PowerShell.
  - iv. Run the following commands to restart or stop the instance:

```
Enter-PsSession -ComputerName 172.16.1XX.183

Restart-Computer -Force # Restarts the instance.

Stop-Computer -Force # Stops the instance.
```

- Use Windows Admin Center
  - i. Start the Windows instance.
  - ii. Configure Windows Admin Center. For more information, see Install Windows Admin Center.
  - iii. Start Windows Admin Center and select the Windows instance. In the left-side navigation pane, click **Overview**. On the Overview page, click **Restart** or **Shutdown**.

Question 3: How do I install the IIS service?

- Use an RDP application
  - i. Connect to the Windows instance. For more information, see Connect to a Windows instance by using a username and password.
  - ii. Enter PowerShell on the command line to start PowerShell.
  - iii. Run the following commands to install IIS:

```
Import-Module ServerManager

Add-WindowsFeature Web-Server, Web-CGI, Web-Mgmt-Cons
ole
```

- Use PowerShell
  - i. Start the Windows instance.
  - ii. Enter PowerShell on the command line on your computer to start PowerShell.
  - iii. Use PowerShell to manage the instance. For more information, see Manage an instance by using PowerShell.
  - iv. Run the following PowerShell commands on your computer:

```
Enter-PsSession -ComputerName 172.16.1XX.183

Import-Module ServerManager

Add-WindowsFeature Web-Server, Web-CGI, Web-Mgmt-Cons
ole
```

- Use Windows Admin Center
  - i. Start the Windows instance.
  - ii. Configure Windows Admin Center. For more information, see the Install Windows Admin Center section.
  - iii. Start Windows Admin Center and select the Windows instance. In the left-side navigation pane, click Roles & Features. On the Roles & Features page, click Web Server. Select suitable features based on your needs and click Yes.

Question 4: How do I reopen a command line window that I accidentally closed during an RDP session?

If a command line window is closed by accident during an RDP session, the remote application displays a black screen and operations cannot be performed. In this case, you can perform the following steps:

- 1. Press Ctrl+Alt+End if an MSTSC connection is used. In other cases, press Ctrl+Alt+Del.
- 2. Select Task Manager and press the Enter key.
- 3. In Task Manager, choose File > Run new task. Enter cmd and click OK.

#### Related information

#### References

- Windows Server servicing channels: LTSC and SAC
- Introducing Windows Server, version 1709
- Windows Admin Center
- About Remote Troubleshooting

## 13.16. Collect kdump information after the operating system of an instance breaks down

When the operating system of an Elastic Compute Service (ECS) instance breaks down, a kdump file is generated. You can collect and submit kdump information for Alibaba Cloud technical support personnel to identify and analyze the issue. This topic describes how to collect kdump information after the operating system of an instance breaks down.

#### Context

Alibaba Cloud technical support personnel can review the kdump information that you submit to analyze the cause of the operating system breakdown and notify you of the cause.

**Note** You can also check kdump files by using the Windows Debugger (WinDbg) tool for Windows instances or the crash tool for Linux instances.

#### Linux instance

- Connect to the Linux instance.
   For more information, see Connect to a Linux instance by using a password or key.
- 2. Run the following command to collect kdump information:

```
acs-plugin-manager --exec --plugin=ecs_dump_collector --params="-c"
```

A command output similar to the following one indicates that kdump information is collected.

```
[root@iZwz9bjyi0fk ~]# acs-plugin-manager --exec --plugin=ecs_dump_collector --params="-c"
[INFO] query kdump status
[INFO] Detected kdump service installed in your system
[INFO] now kdump service's status is: active
[INFO] kdump status is: Kdump is operational
check core path: /var/crash/127.0.0.1-2022-05-20-11:52:16/vmcore
uploading dumps: /var/crash/127.0.0.1-2022-05-20-11:52:16/vmcore
upload dumps finished
```

#### Windows instance

- Connect to the Windows instance.
   For more information, see Connect to a Windows instance by using a password or key.
- 2. In the lower-left corner of the Windows desktop, enter PowerShell in the search box to find Windows PowerShell and then open Windows PowerShell.
- 3. Run the following command in PowerShell to go to the *assist* directory:

cd C:\ProgramData\aliyun\assist

4. Run the following command to view the information of files in the assist directory:

ls

5. Run the following command to access the latest version of the directory. Example: 2.1.3.282

```
cd 2.1.3.282
```

6. Run the following command to collect kdump information:

```
acs-plugin-manager.exe --exec --plugin=ecs_dump_collector_win --params="-c"
```

A command output similar to the following one indicates that kdump information is collected.

```
C:\ProgramData\aliyun\assist\2.1.3.282\acs-plugin-manager.exe --exec --plugin=ecs_dump_collector_win --params="-c" [06/01/2022 15:23:32] [info] Crash status is ON. check core path: c:/windows/MEMORY.DMP uploading dumps: c:/windows/MEMORY.DMP upload dumps finished
```

#### **FAQ**

If the error message shown in the following figure is displayed when you attempt to collect kdump information for an instance, it indicates that the dump service is not enabled in the operating system.

```
[INFO] Detected kdump service installed in your system
[INFO] now kdump service's status is: unknown
[INFO] kdump service is NOT enabled in system, please run --enable
[INFO] kdump status is: Kdump is not operational
```

#### Solution:

- 1. Run one of the following commands to enable the dump service.
  - o For Linux operating systems, run the acs-plugin-manager --exec --plugin=ecs\_dump\_collector --params="--enable" command.
  - For Windows operating systems, run the acs-plugin-manager.exe --exec --plugin=ecs\_dump\_co llector win --params="--enable" command.
- 2. Restart the instance in the ECS console for the new settings to take effect.

For more information, see Restart an instance.

## 13.17. Red Hat Cloud Access and Windows licenses

#### 13.17.1. Red Hat Cloud Access

#### 13.17.1.1. Overview

This topic describes how to import Red Hat Enterprise Linux (RHEL) virtual machine (VM) images to Alibaba Cloud and use the imported images to create Elastic Compute Service (ECS) instances in the ECS console. Alibaba Cloud is a Red Hat certified cloud and service provider. The procedure described in this topic is applicable to RHEL 6.9, 6.10, 7.4 and later minor versions, and 8.0 and later minor versions.

#### Requirements for ECS instances and RHEL custom images

The RHEL images to be imported to Alibaba Cloud ECS must meet the following requirements:

- Images can be imported over a single network interface and cannot be imported from IPv6 addresses.
- The sshd service is enabled for remote access.
- Firewalls are disabled and SSH port 22 is enabled so that you can connect to the instances after they are created.

We recommend that you use security groups to control inbound or outbound traffic for one or more ECS instances. For more information, see Overview.

• The operating system is RHEL 6.9, 6.10, 7.4 or a later minor version, or 8.0 or a later minor version, which includes the latest cloud-init version and security updates.

For more information about RHEL versions, see Overview.

- One of the following file systems is used on the root partition: xfs, ext3, and ext4.
- Images only in the RAW, QCOW2, or VHD format can be imported to ECS. We recommend that you import images in the QCOW2 or VHD format, which have a smaller transmission footprint. If you want to import images in other formats, you can use the qemu-img tool to convert the images to the supported formats.
- The /etc/issue\*file is not modified. Otherwise, the distribution of the guest operating system cannot be properly identified, and you may not be able to create ECS instances.
- The GRand Unified Bootloader (GRUB) configuration file is correctly configured. Otherwise, the system may not be able to start.
- The /etc/fstab file is correctly configured. Otherwise, exceptions that prevent the loading of partitions may occur, which can cause the system to fail to start.
- The default permissions on the /etc/shadow file must be used. Otherwise, the password file cannot be modified and the system cannot start.
- A single root partition exists on the system disk and the system disk has sufficient space for the image files.
- Critical system files are correctly configured, such as /sbin, /bin, and /lib\*.

For more information about the requirements for ECS instances and RHEL custom images, see Instructions for importing images.

#### **Prerequisites**

Before you start, make sure that the following requirements are met:

- You have both a Red Hat account and an Alibaba Cloud account.
  - If you do not have an Alibaba Cloud account, create one first. For more information, see Sign up with Alibaba Cloud.
  - If you do not have a Red Hat account, create a Red Hat portal account on the official Red Hat website.
- You have a physical RHEL 7 server as an offline environment to tailor the downloaded virtual image to the Alibaba Cloud image requirements. For more information, see <a href="Step 2">Step 2</a>: Create a RHEL image.
- If you want to create an ECS instance within a Chinese mainland region, you must sign up with Alibaba Cloud by using your real name.

#### **Procedure**

The following figure shows how to import a RHEL VM image and create a RHEL instance from the imported image in the ECS console.



#### **Procedure**

- 1. Step 1: Enroll in the Red Hat Cloud Access program
- 2. Step 2: Create a RHEL image
- 3. Step 3: (Optional) Convert the image to a supported format
- 4. Step 4: Import the image to Alibaba Cloud ECS
- 5. Step 5: Create an ECS instance from the RHEL image

#### 13.17.1.2. Step 1: Enroll in the Red Hat Cloud Access

#### program

This topic describes how to enroll in the Red Hat Cloud Access program to run your Red Hat product subscriptions on Alibaba Cloud.

#### **Procedure**

- 1. Log on to the official Red Hat website by using your Red Hat account.

  Click Red Hat website to visit the official Red Hat website.
- 2. Go to the Red Hat Cloud Access page and click Get started.

Elastic Compute Service Images-FAQ

Click Access page to go to the Red Hat Cloud Access page.

3. Click **Subscription eligibility** in the Top Resources section to verify whether you are eligible for Red Hat Cloud Access.

Click Verifying eligibility for Cloud Access to review subscription eligibility guidelines.

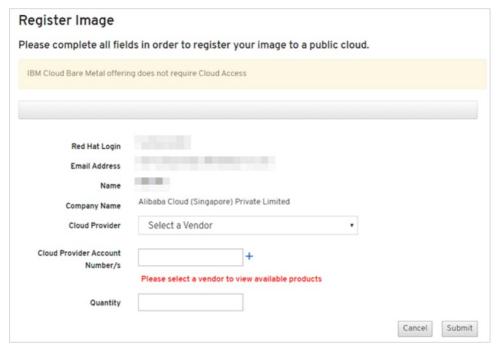
4. Click the Enable and use your own images link in the Enable and use your subscriptions in the cloud section.

Click Enroll with Use your own image to enable and use your own images.

- 5. In the **Register Image** dialog box, set the following parameters:
  - o Cloud Provider: Select Alibaba Cloud.
  - Cloud Provider Account: Enter the ID of your Alibaba Cloud account.

You can choose **Account Management > Basic Information** in the ECS console to view your account ID.

• Quantity: Specify the number of images.



6. Click Submit to register the images with Alibaba Cloud ECS.

#### 13.17.1.3. Step 2: Create a RHEL image

This topic describes how to migrate physical hosts from on-premises data centers to Alibaba Cloud or migrate virtual machines (VMs) from classic virtualization platforms to Alibaba Cloud. In this topic, a Red Hat Enterprise Linux (RHEL) 7 server is used.

#### Context

You can migrate physical servers or VMs to Alibaba Cloud.

• Migrate physical servers to Alibaba Cloud ECS (P2V and V2V)

You can use Server Migration Center (SMC) to migrate the operating system images of physical servers to Alibaba Cloud Elastic Compute Service (ECS). After the images are imported to ECS, you can use the imported images to create instances.

Migrate VMs to Alibaba Cloud ECS (V2V)

You can create VM images supported by Alibaba Cloud on on-premises physical servers. You can perform the operations described in this topic or use SMC to migrate VMs to Alibaba Cloud ECS.

For more information about the migration procedure in SMC, see Migration process.

#### Migrate a VM to Alibaba Cloud ECS

Perform the following steps to migrate a VM to Alibaba Cloud ECS:

- 1. Step 1: Prepare a physical RHEL 7 server
- 2. Step 2: Create a VM
- 3. Step 3: Configure the guest operating system
- 4. Step 4: Create a VM image

Note If a VM that meets the requirements described in Step 1: Prepare a physical RHEL 7 server is already running on the RHEL 7 server, you can start from Step 3: Configure the guest operating system.

#### Step 1: Prepare a physical RHEL 7 server

- 1. Log on to the physical RHEL 7 server.
- 2. Make sure that all the packages listed in the following table are installed on the RHEL server.

You can run the yum info Package name command to view the installed packages.

Package	Repository	Description
libvirt	rhel-7-server- rpms	An open source API, daemon, and management tool for managing platform virtualization.
virt-manager	rhel-7-server- rpms	A desktop tool for managing VMs.
libguestfs	rhel-7-server- rpms	A set of tools for accessing and modifying VM disk images.
libguestfs-tools	rhel-7-server- rpms	A set of system management tools for VMs, including guestfish.
qemu-img	rhel-7-server- rpms	A tool for converting VM disk images from one format to another.

#### Step 2: Create a VM

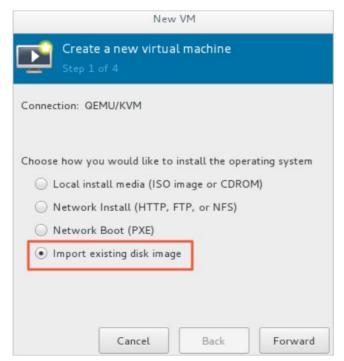
- 1. Click Red Hat to download the latest Red Hat Enterprise Linux 7.5 Update KVM Guest Image or Red Hat Enterprise Linux 7.4 Update KVM Guest Image.
  - Note We recommend that you download RHEL 7.5 that has more comprehensive features than RHEL 7.4.
- 2. Save the image to the RHEL physical server prepared in Step 1.

Elastic Compute Service Images·FAQ

- 3. Move the image to the /var/lib/libvirt/images directory.
- 4. In the Virtual Machine Manager window, click the 🔝 icon to create a VM.

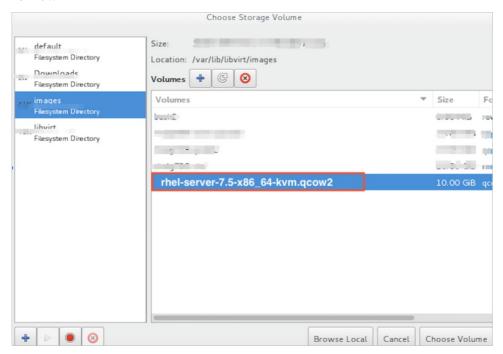


- 5. Complete the configurations in the **Create a virtual machine** wizard.
  - i. Select **Import existing disk image** to install the guest operating system. Then, click **Forward**.

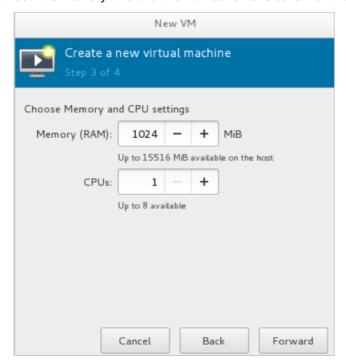


Images· FAQ Elastic Compute Service

ii. Browse the /var/lib/libvirt/images directory and select the KVM guest image in the QCOW2 format.

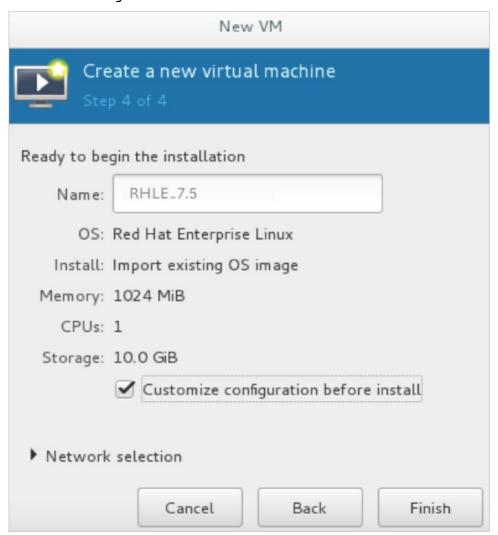


iii. Set the memory size and the number of CPU cores for the new VM.



Elast ic Compute Service Images• FAQ

iv. Check the configurations of the new VM, specify a name for the VM, and then select **Customize configuration before install**.



- v. In the left-side pane of the Custom configuration dialog box, click NIC and set **Device model** to **virtio**.
- vi. Click Begin Installation.
  - The VM begins to be created. During this process, the VM console appears and boots to a logon prompt.
- vii. When the logon prompt appears, shut down the VM.
- 6. Run the following command to set a password for the root user of the VM:

```
virt-customize -a /var/lib/libvirt/images/rhel-server-7.5-x86_64-kvm.qcow2 --root-passw
ord password>
```

In this example, the password for the root user of the VM is set to redhat! . Sample command:

```
# virt-customize -a /var/lib/libvirt/images/rhel-server-7.5-x86_64-kvm.qcow2 --root-pa
ssword password:redhat!
[ 0.0] Examining the guest ...
[ 103.0] Setting a random seed
[ 103.0] Setting passwords
[ 112.0] Finishing off
```

7. Start the RHEL VM and log on as the root user.

# Step 3: Configure the guest operating system

After you log on to the VM, you must configure the guest operating system.

1. Make sure that the guest operating system of the RHEL VM meets the specified requirements.

For more information, see the "Requirements for ECS instances and RHEL custom images" section in Overview.

The following table describes the requirements for the guest operating system.

ltem	Description	
/etc/shadow	/etc/shadow is unlocked by running the chattr command.	
qemu-ga	qemu-ga is uninstalled.	
Network	The network manager is disabled or uninstalled, and network services are enabled.	
ssh	The SSH service is enabled and PermitRootLogin is not set.	
Firewall	Firewalls are disabled.	
File system	The xfs, ext3, or ext4 file system is used. Do not use the 64-bit ext4 file system.	
passwd	passwd is installed.	
Partition	The Master Boot Record (MBR) or GUID Partition Table (GPT) partitioning format is used.	
Logical Volume Manager (LVM)	LVM is not used.	
/lib	Relative paths can be used only for /lib and /lib64.	
cloud-init	cloud-init is installed in the guest operating system.	
	Note By default, cloud-init disables password-based logon. If you want to enable password-based logon, see Step 6 of the " (Optional) Install the native cloud-init" section in Install cloud-init. This step shows you how to modify the content that precedes cloud_init_modules in the /etc/cloud/cloud.cfg configuration file.	

2. Configure the ntp.conf file.

Perform the following steps to configure the ntp.conf file that is located in the /etc/ntp.conf directory:

i. Run the following command to back up the ntp.conf file as ntp.conf.bk:

```
cp /etc/ntp.conf /etc/ntp.conf.bk
```

ii. Edit the /etc/ntp.conf file by replacing its content with the following content:

```
driftfile /var/lib/ntp/drift
pidfile /var/run/ntpd.pid
logfile /var/log/ntp.log
# Access Control Support
restrict default kod nomodify notrap nopeer noquery
restrict -6 default kod nomodify notrap nopeer noquery
restrict 127.0.0.1
restrict 192.168.0.0 mask 255.255.0.0 nomodify notrap nopeer noquery
restrict 172.16.0.0 mask 255.240.0.0 nomodify notrap nopeer noquery
restrict 100.64.XX.XX mask 255.192.0.0 nomodify notrap nopeer noquery
restrict 10.0.0.0 mask 255.0.0.0 nomodify notrap nopeer noquery
# local clock
server 127.127.1.0
fudge 127.127.1.0 stratum 10
restrict ntpl.aliyun.com nomodify notrap nopeer noquery
restrict ntpl.cloud.aliyuncs.com nomodify notrap nopeer noquery
restrict ntp10.cloud.aliyuncs.com nomodify notrap nopeer noquery
restrict ntpl1.cloud.aliyuncs.com nomodify notrap nopeer noquery
restrict ntp12.cloud.aliyuncs.com nomodify notrap nopeer noquery
restrict ntp2.aliyun.com nomodify notrap nopeer noquery
restrict ntp2.cloud.aliyuncs.com nomodify notrap nopeer noquery
restrict ntp3.aliyun.com nomodify notrap nopeer noquery
restrict ntp3.cloud.aliyuncs.com nomodify notrap nopeer noquery
restrict ntp4.aliyun.com nomodify notrap nopeer noquery
restrict ntp4.cloud.aliyuncs.com nomodify notrap nopeer noquery
restrict ntp5.aliyun.com nomodify notrap nopeer noquery
restrict ntp5.cloud.aliyuncs.com nomodify notrap nopeer noquery
restrict ntp6.aliyun.com nomodify notrap nopeer noquery
restrict ntp6.cloud.aliyuncs.com nomodify notrap nopeer noquery
restrict ntp7.cloud.aliyuncs.com nomodify notrap nopeer noquery
restrict ntp8.cloud.aliyuncs.com nomodify notrap nopeer noquery
restrict ntp9.cloud.aliyuncs.com nomodify notrap nopeer noquery
server ntpl.aliyun.com iburst minpoll 4 maxpoll 10
server ntpl.cloud.aliyuncs.com iburst minpoll 4 maxpoll 10
server ntp10.cloud.aliyuncs.com iburst minpoll 4 maxpoll 10
server ntpl1.cloud.aliyuncs.com iburst minpoll 4 maxpoll 10
server ntp12.cloud.aliyuncs.com iburst minpoll 4 maxpoll 10
server ntp2.aliyun.com iburst minpoll 4 maxpoll 10
server ntp2.cloud.aliyuncs.com iburst minpoll 4 maxpoll 10
server ntp3.aliyun.com iburst minpoll 4 maxpoll 10
server ntp3.cloud.aliyuncs.com iburst minpoll 4 maxpoll 10
server ntp4.aliyun.com iburst minpoll 4 maxpoll 10
server ntp4.cloud.aliyuncs.com iburst minpoll 4 maxpoll 10
server ntp5.aliyun.com iburst minpoll 4 maxpoll 10
server ntp5.cloud.aliyuncs.com iburst minpoll 4 maxpoll 10
server ntp6.aliyun.com iburst minpoll 4 maxpoll 10
server ntp6.cloud.aliyuncs.com iburst minpoll 4 maxpoll 10
server ntp7.cloud.aliyuncs.com iburst minpoll 4 maxpoll 10
server ntp8.cloud.aliyuncs.com iburst minpoll 4 maxpoll 10
server ntp9.cloud.aliyuncs.com iburst minpoll 4 maxpoll 10
```

After you make the modification, press the  $\mathit{Esc}$  key, enter  $:_{\mathtt{wq}}$ , and press the Enter key to save the modification and exit.

Elastic Compute Service Images-FAQ

You can create a VM image after you complete the configurations for the guest operating system. You can also use SMC provided by Alibaba Cloud to create and import VM images. For more information, see What is SMC?

### Step 4: Create a VM image

- 1. On the menu bar of the RHEL VM, choose View > Snapshots.
- 2. Click the + icon to create a snapshot.

After the snapshot is created, a VM image is created. The VM image is in the QCOW2 format.

### What's next

Step 3: (Optional) Convert the image to a supported format

# 13.17.1.4. Step 3: (Optional) Convert the image to a

# supported format

Images only in the RAW, QCOW2, or VHD format can be imported to Elastic Compute Service (ECS). If you want to import images in other formats, you must first convert the images to supported formats. This topic describes how to use the qemu-img tool to convert images in other formats to the QCOW2, RAW, or VHD format.

### Context

The gemu-img tool supports the conversion of images in the following formats:

- Convert images from the RAW, qcow2, qcow1, VMDK, VDI, VHD (VPC), VHDX, or QED format to the VHD format.
- Convert between RAW and VHD.

For more information, see Convert the image file format

### **Procedure**

1. Run the following command to install the gemu-img tool:

```
yum install qemu-img
```

2. Run the following command to convert the format of the image:

```
qemu-img convert -f VHDX -O vhd rhel.vhdx rhel.vhd
```

Parameter description:

- o -f: This parameter is followed by the original image format.
- -O: This parameter must be in uppercase and is followed by the new image format, the name of the original image, and the name of the new image.

After the conversion is complete, the new image appears in the directory of the original image.

# 13.17.1.5. Step 4: Import the image to Alibaba Cloud ECS

To import an image to Alibaba Cloud Elastic Compute Service (ECS), you must first upload the image to an Object Storage Service (OSS) bucket. This topic describes how to import the Red Hat Enterprise Linux (RHEL) image to Alibaba Cloud ECS.

### **Prerequisites**

- The OSS bucket, the image, and instances to use the image are located in the same region and belong to the same account.
- OSS is activated. For more information, see Activate OSS.
- If you are using a RAM user, the AliyunECSImageImportDefaultRole role is attached to the RAM user by the corresponding Alibaba Cloud account.

### Context

OSS is a secure, cost-effective, and highly durable service that enables you to store large amounts of data in Alibaba Cloud. For more information, see What is OSS?

### **Procedure**

- 1. Step 1: Create an OSS bucket
- 2. Step 2: Upload the image to the OSS bucket
- 3. Step 3: Import the image to ECS

### Step 1: Create an OSS bucket

- 1. Log on to the OSS console.
- 2. In the left-side navigation pane, click **Buckets**.
- 3. On the **Buckets** page, click **Create Bucket**.
- 4. In the Create Bucket panel, configure the following parameters. Then, click OK.
  - Bucket Name: Enter a valid name for the bucket. In this example, enter importrhel.
  - Region: Select a region from the drop-down list.
    - **Note** The OSS bucket and ECS instances to use the image must be in the same region.
  - Storage Class: Select the default storage class. In this example, select Standard.
  - Access Control List (ACL): Select the default ACL. In this example, select Private.

For information about how to configure other parameters, see Create buckets.

### Step 2: Upload the image to the OSS bucket

If the image file is greater than 5 GiB in size, you must upload the image by calling an API operation or using ossutil.

- For more information about the API operation, see Multipart upload.
- For more information about ossutil, see Overview.

In this example, the image file is less than 5 GiB in size. You can upload it to an OSS bucket by using the OSS console.

- 1. Log on to the OSS console.
- 2. In the left-side navigation pane, click **Buckets**.

Elastic Compute Service Images-FAQ

3. In the bucket list of the Buckets page, click the bucket name in the Bucket Name column. In this example, click importrhel.

- 4. On the Files page, click Create Folder.
- 5. In the Create Folder panel, enter a valid folder name. In this example, enter rhel/. Then, click OK.
- 6. Click the folder name. Then, click **Upload**.
- 7. On the **Upload** page, configure the following parameters:
  - Upload To: Click Current.
  - o File ACL: Click Inherited from Bucket.
  - Files to Upload: You can click Select Files to select a local file or click Select Folders to select a directory. You can also drag the required local file or directory to the Files to Upload section.

For information about how to configure other parameters, see Simple upload.

8. Click Upload.

You can view the upload progress of the file in the Task List panel. When Uploaded is displayed in the Status column corresponding to the file, the image file is uploaded and is displayed in the *rhel/* directory.

9. Click the image name. In the View Details panel, click Copy File URL below the URL field.



- The file URL is required when you import the image.
- If File ACL is set to Private, the URL contains the Expires and Signature information.

## Step 3: Import the image to ECS

After you upload the image to the OSS bucket, you can perform the following operations to import the image to ECS in the ECS console. You can also call the <a href="majorithmage">importimage</a> operation to import the image to ECS.

**Note** We recommend that you use an Alibaba Cloud account to import the image. If you are using a RAM user, you must attach AliyunECSImageImportDefaultRole to the RAM user.

- 1.
- 2.
- 3. In the top navigation bar, select a region.
  - Note Instances to use the image must be created in the region where the OSS bucket is located.
- 4.
- 5. Import the custom image.
  - i. In the upper-right corner of the Images page, click Import Image.

ii. In the Import Image dialog box, configure the following parameters:

- OSS Object Address: Paste the file URL that you copied in Step 2.
- Image Name: Enter a valid name for the image. You can use this name to search for the image when you create an ECS instance.
- Operating System/Platform: Select Linux from the first drop-down list and select RedHat from the second drop-down list.
- System Architecture: Select x86\_64.
- System Disk (GiB): Enter a value ranging from 40 to 500 as the size of the system disk.
- Image Format : Select VHD.
- License Type: Select BYOL.

For information about how to configure other parameters, see Import custom images.

- iii. Confirm that your configurations are correct and click **OK**. The system creates a task to import the custom image.
- 6. (Optional)

### What's next

Step 5: Create an ECS instance from the RHEL image

# 13.17.1.6. Step 5: Create an ECS instance from the RHEL

# image

After the Red Hat Enterprise Linux (RHEL) image is imported to Elastic Compute Service (ECS), you can use the image to create ECS instances. This topic describes how to use the RHEL image to create an instance in the ECS console. You can also call the RunInstances operation to create an instance to use the RHEL image.

#### **Procedure**

- 1.
- 2.
- 3.
- 4. On the Instances page, click Create Instance.
- 5. Perform the following operations in the Basic Configurations step:
  - i. Select a billing method.
    - For more information about billing methods of instances, see Overview.
  - ii. Select a region and a zone.
    - Select the region where the imported image is located and select a zone. For more information, see Regions and zones.
  - iii. Select an instance type and specify the number of instances to be created.
    - Select an instance type that is supported by Red Hat. For more information, visit Which Instance Types are RHEL Guest supported with on Aliyun.

Elastic Compute Service Images-FAQ

iv. Select an image.

In the Image section, click **Custom Image** and select the imported RHEL image from the drop-down list.

v. Configure storage.

Select the disk category of the system and data disks and specify the size of the disks. The size of the system disk cannot be smaller than that of the image.

- vi. (Optional) Configure the snapshot service.
- vii. Click Next.
- 6. (Optional) Complete configurations in the Network, System Configurations (Optional), and Grouping (Optional) steps. Then, confirm the order.

For more information about the steps and parameters in the wizard, see Create an instance by using the wizard.

#### Result

After the preceding operations are performed, go to the Instances page to view the state of the instance. If the RHEL instance is created, **Running** is displayed in the Status column corresponding to the instance.

### What's next

You can use Workbench, Virtual Network Console (VNC), or a third-party client to connect to the instance. For more information, see Connection methodsGuidelines on instance connection.

### Related information

RunInstances

# 13.17.2. Windows licenses

# 13.17.2.1. Microsoft licenses and Software Assurance

When you want to deploy your business in a Microsoft software-based environment, Alibaba Cloud can provide you with Microsoft licenses.

You can use one of the following methods to use Microsoft licenses in Alibaba Cloud:

• Use Microsoft licenses provided by Alibaba Cloud

Purchase Microsoft licenses from Alibaba Cloud. These licenses are included in Alibaba Cloud public images for Windows Server or Alibaba Cloud Market place images with pre-installed SQL Server databases.

• Bring your own license

If you have purchased a Microsoft license, apply for License Mobility by using Microsoft Software Assurance to bring your own license (BYOL) to Alibaba Cloud.

# Use Microsoft licenses provided by Alibaba Cloud

Alibaba Cloud is licensed to provide Microsoft licenses by using the Services Provider License Agreement (SPLA). Alibaba Cloud Elastic Compute Service (ECS) and ApsaraDB RDS provide software licensed by Microsoft, such as Windows Server images and ApsaraDB RDS for SQL Server databases. Client Access Licenses (CALs) are not required for Alibaba Cloud ECS instances that run Windows Server images. By default, each Windows Server instance has two connections that are used only for management purposes. If you need more connections or want to use the connections for non-management purposes, an extra Remote Desktop Services CAL (RDS CAL) is required.

You can select an appropriate Windows Server image or database type when you create an ECS instance. For more information, see Create an instance by using the wizard. The following table lists the versions of Windows Server images and SQL Server databases provided by Alibaba Cloud.

Windows Server operating system	SQL Server database
<ul> <li>Windows Server 2019 Datacenter with Containers 64-bit (Simplified Chinese)</li> <li>Windows Server 2019 Datacenter with Containers 64-bit (English)</li> <li>Windows Server 2019 Datacenter 64-bit (Simplified Chinese)</li> <li>Windows Server 2019 Datacenter 64-bit (English)</li> <li>Windows Server 2016 Datacenter 64-bit (Simplified Chinese)</li> <li>Windows Server 2016 Datacenter 64-bit (English)</li> <li>Windows Server 2012 R2 Datacenter 64-bit (Simplified Chinese)</li> <li>Windows Server 2012 R2 Datacenter 64-bit (English)</li> <li>Windows Server Version **** Datacenter 64-bit (Simplified Chinese)</li> <li>Windows Server Version **** Datacenter 64-bit (English)</li> <li>For more information about Windows Server versions, see .</li> </ul>	<ul> <li>Windows Server 2016 SQL Server 2014 Enterprise Edition</li> <li>Windows Server 2016 SQL Server 2014 Standard Edition</li> <li>Windows Server 2012 R2 SQL Server 2012 SP3 Enterprise Edition</li> <li>Windows Server 2012 R2 SQL Server 2012 SP3 Standard Edition</li> </ul>



? Note

# Bring your own license

You can continue to use your own purchased Microsoft licenses by bringing them to Alibaba Cloud. Check whether your Microsoft license provides Software Assurance. For more information, visit the official Microsoft website. You can perform one of the following operations based on whether your Microsoft license provides Software Assurance:

Note You can import a BYOL custom image to Alibaba Cloud and use this image to create an ECS instance. Set License Type to BYOL when you import a BYOL image. For more information, see Import custom images.

• Microsoft licenses without Software Assurance

If your Microsoft license does not provide Software Assurance, you can use ECS bare metal instances or dedicated hosts. The resources offered by these instances or hosts are exclusive for your use, so you can use your own Microsoft licenses even though these licenses do not provide Software Assurance. For more information, see Overview and What is DDH?.

Notice To meet compliance requirements, your license must be obtained before October 1, 2019, and the operating system versions must be released before October 1, 2019.

All Microsoft software covered by the Microsoft Product Terms can be deployed on ECS bare metal instances or dedicated hosts. Examples include:

- Microsoft Windows Server
- Microsoft Exchange Server
- o Microsoft SharePoint Server
- o Microsoft SQL Server
- o Microsoft System Center
- Microsoft Dynamics
- MSDN
- o Microsoft Remote Desktop Services (user CALs)
- Microsoft licenses with Software Assurance

You can migrate your Microsoft license to Alibaba Cloud and continue to use it by using the Software Assurance and Microsoft Volume Licensing programs. For software covered by the license terms of valid Microsoft products, you can apply for License Mobility through Software Assurance before you migrate services to ECS. For more information, see Apply for License Mobility and What is SMC?.

The following list describes eligible Microsoft applications and services:

- Microsoft Exchange Server
- Microsoft SharePoint Server
- o Microsoft SQL Server
- Microsoft System Center
- Microsoft Dynamics
- Microsoft BizTalk Server
- Microsoft Remote Desktop Services (user CALs)
- o Microsoft Skype for Business Server

### Statement of rights and responsibilities

Microsoft offers a comprehensive set of applications and services that are performance-oriented, secure, and stable. Select the solution that fits your requirements, and always comply with Microsoft licensing policy. This document is for your reference only and must not be construed as legal advice. For questions about the authorization and reserved rights of Microsoft applications and services, consult your legal team, a law firm, Microsoft, or Microsoft partners.

# 13.17.2.2. Apply for License Mobility

This topic describes how to apply for License Mobility through Software Assurance. If you plan to use an Elastic Compute Service (ECS) dedicated host or purchase a license provided by Alibaba Cloud, skip this topic.

### Context

As an Authorized Mobility Partner of Microsoft, Alibaba Cloud accepts the licenses of services covered in active Software Assurance contracts. You can send an application to Microsoft to migrate your license and continue to use it in Alibaba Cloud as a Bring Your Own License (BYOL) license.

If you want to continue using your existing licenses through License Mobility, make sure that the following requirements are met:

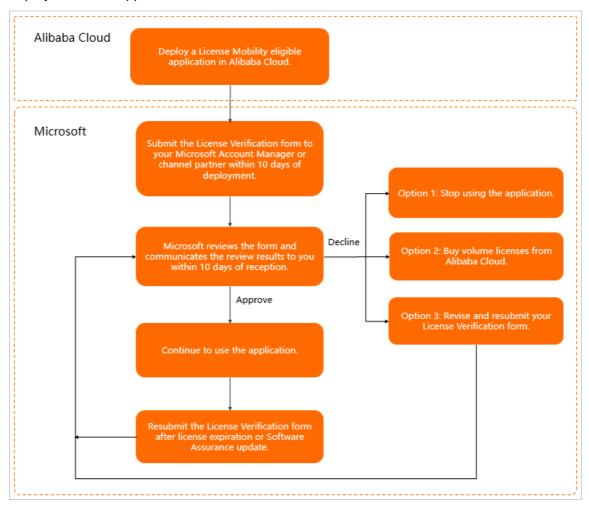
- Your licenses have active Microsoft Software Assurance. For more information, see Software Assurance.
- The Microsoft services must support License Mobility agreements. The following list provides some Microsoft services that are eligible for License Mobility:
  - o Exchange Server
  - o SharePoint Server
  - o SQL Server Standard Edition
  - o SQL Server Enterprise Edition
  - Skype for Business Server
  - o System Center Server
  - o Project Server
  - o Visual Studio Team Foundation Server
  - o BizTalk Server
  - o Remote Desktop Services

For information about a full list of eligible Microsoft services, visit Licensing terms.

- Eligible Volume Licensing programs include the Enterprise Agreement, Enterprise Subscription Agreement, and Microsoft Open Value Agreement, where Software Assurance is included, as well as other Volume Licensing programs where Software Assurance is an option, such as Microsoft Open License or Select Plus.
- The licenses cannot be applied to Windows client operating systems such as Windows Server or desktop applications such as Microsoft Office.

### Procedure

In accordance with Microsoft License Mobility terms, you can deploy application environments before Microsoft confirms your application. Within 10 days of the deployment, you must complete and submit the License Mobility Verification Form for review. The following figure shows how to apply for the deployment of an application.



- 1. Deploy an application that is eligible for License Mobility in Alibaba Cloud.
- 2. Submit the license migration application to Microsoft.

For information about detailed operations, visit License Mobility through Software Assurance. Specify Alibaba Cloud as the Authorized Mobility Partner on the verification form and provide the following information:

- o Authorized Mobility Partner name: Alibaba Cloud Computing Ltd.
- Authorized Mobility Partner website address: www.alibaba-inc.com
- Authorized Mobility Partner email address: microsoft@alibaba-inc.com
- 3. Wait for Microsoft to review the application and reply.

### What's next

In accordance with the Microsoft Product Terms, you can perform the following operations on licenses that have License Mobility benefits:

• Reallocate these licenses to other ECS instances in the same Alibaba Cloud zone. For more information, see Regions and zones.

• Reallocate these licenses to other ECS instances in different Alibaba Cloud zones. The interval between the consecutive reallocations cannot be less than 90 days. For more information, see Licensing terms.