

Huawei Cloud EulerOS

User Guide

Issue 08
Date 2024-03-30



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

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

Trademarks and Permissions



HUAWEI and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

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

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

Contents

1 Overview.....	1
2 Selecting Huawei Cloud EulerOS as the Public Image When Creating an ECS.....	2
3 Changing an OS to Huawei Cloud EulerOS.....	3
4 Migrating an OS to Huawei Cloud EulerOS.....	7
4.1 Using x2hce-ca to Evaluate the Compatibility with Huawei Cloud EulerOS.....	7
4.1.1 Overview of x2hce-ca.....	7
4.1.2 Constraints.....	8
4.1.3 Installing x2hce-ca.....	8
4.1.4 Evaluating Software Compatibility.....	9
4.2 Migrating an OS to Huawei Cloud EulerOS 2.0.....	11
4.2.1 Constraints.....	11
4.2.2 Migrating an OS to Huawei Cloud EulerOS.....	12
4.2.3 Conflicting Packages.....	18
4.3 Migrating an OS to Huawei Cloud EulerOS 1.1.....	26
4.3.1 Constraints.....	26
4.3.2 Migration Operations.....	26
4.3.3 Appendix: Description of the .conf File.....	30
5 Upgrading Huawei Cloud EulerOS and RPM Packages.....	32
5.1 Upgrade Overview.....	32
5.2 Using dnf or yum to Upgrade.....	33
5.3 Upgrade Using OSMT.....	36
5.3.1 Overview.....	36
5.3.2 Constraints.....	36
5.3.3 Version Upgrade and Rollback.....	37
5.3.4 Updating RPM Packages.....	39
5.3.4.1 Preparations.....	39
5.3.4.2 Manual Update Using osmt update.....	43
5.3.4.3 Automatic Update Using osmt-agent.....	44
5.3.5 Follow-up Operations.....	45
5.3.6 Rolling Back RPM Packages.....	45
5.4 Appendixes.....	45
5.4.1 OSMT Command Help Information.....	45

5.4.2 Description of the /etc/osmt/osmt.conf File.....	49
6 Security Updates for Huawei Cloud EulerOS.....	52
6.1 Security Updates Overview.....	52
6.2 About CVE.....	52
6.3 Yum Command Parameters.....	53
6.4 Querying Security Updates.....	54
6.5 Checking for Security Updates.....	55
6.6 Installing Security Updates.....	55
7 Obtaining the openEuler Extended Software Packages.....	57
8 Creating a Docker Image and Starting a Container.....	61
9 Tools.....	65
9.1 BiSheng Compiler.....	65
9.2 Workload Accelerator.....	66
9.2.1 Overview.....	66
9.2.2 Installing Workload Accelerator.....	67
9.2.3 Static Acceleration.....	68
9.2.4 Dynamic Acceleration.....	72
9.2.5 Configuration File.....	77
9.3 Pod Bandwidth Management Tool.....	80
10 Kernel Functions and Interfaces.....	84
10.1 OOM Process Control Policy.....	84
10.2 Multi-level Memory Reclamation Policy.....	87
10.3 Multi-level Hybrid Scheduling of Kernel CPU cgroups.....	90
10.4 Kernel Exception Events Analysis.....	92
11 xGPU.....	99
11.1 Overview.....	99
11.2 Installing and Using xGPU.....	101
11.3 GPU Compute Scheduling Examples.....	108
12 Configuring Repo Sources and Installing Software for Huawei Cloud EulerOS.....	112
13 Change History.....	115

1 Overview

You can use Huawei Cloud EulerOS in the following ways:

- Selecting Huawei Cloud EulerOS as the public image when creating an ECS
- Changing an OS to Huawei Cloud EulerOS

If the ECS OS needs to be changed but the other ECS details (NICs, disks, and VPNs) need to be kept the same, as long as the software is loosely coupled to the OS, you can change the OS to Huawei Cloud EulerOS with only a few additional changes needed.

- Migrating an OS to Huawei Cloud EulerOS

If you want to change an ECS OS but retain other ECS details (NICs, disks, and VPNs) and the OS component parameter settings, you can migrate the OS to Huawei Cloud EulerOS.

NOTE


The OS can be migrated only to Huawei Cloud EulerOS 2.0 standard edition or Huawei Cloud EulerOS 1.1 CentOS compatible edition.

Table 1-1 Differences between OS change and OS migration

Item	OS Change	OS Migration
Data backup	<ul style="list-style-type: none">• Data in all partitions of the system disk will be lost, so back up the system disk data before you change the OS.• Data in data disks remains unchanged.	<ul style="list-style-type: none">• System disk data is not lost, but you are still advised to back up the system disk data, in case there are any system software exceptions.• Data in data disks remains unchanged.
Custom settings	After the OS is changed, custom settings such as DNS and hostname will be reset and need to be reconfigured.	After the OS is changed, custom settings such as DNS and hostname do not need to be reconfigured.

2 Selecting Huawei Cloud EulerOS as the Public Image When Creating an ECS


Procedure

1. Log in to the [management console](#).
2. Hover over  and choose **Compute > Elastic Cloud Server**.
3. Click **Buy ECS**.

The page for purchasing ECSs is displayed.

4. Select Huawei Cloud EulerOS as the public image.

On the **Configure Basic Settings** page, select **Public Image**, and select **Huawei Cloud EulerOS** and an image version. For details about the other steps to purchase an ECS, see [Purchasing an ECS](#).



NOTE

The Huawei Cloud EulerOS public image can only be used on ECSs of certain flavors. For details, see [Supported Instance Types](#).

3 Changing an OS to Huawei Cloud EulerOS

Constraints

- If you want to change the OS of a yearly/monthly ECS, the system disk capacity may be insufficient for the new image. If this is the case, you need to detach the system disk from the ECS and expand its capacity before changing the OS.
- There must be at least one EVS disk.
- Changing from BIOS to UEFI is not supported.

Important Notes

- After the OS is changed, the original OS will not be retained, and the original system disk and data on the system disk will be deleted, including data on the system partition and other partitions. Back up the data before changing the OS. For details, see [Backing Up ECS Data](#).
- Changing the OS will not affect data in data disks.
- After the OS is changed, the IP and MAC addresses remain unchanged.
- The system will automatically start after the OS is changed.
- After the OS is changed, the system disk type cannot be changed.
- After you change the OS, you need to deploy services in the new OS.
- After the OS is changed, custom settings such as DNS and hostname will be reset and need to be reconfigured.

Billing Rules


- After the OS of a pay-per-use ECS is changed, the price may increase if the new OS needs more disk capacity.

Prerequisites

- The source system has a system disk attached.
- If a password is used to log in to the source OS but a key pair is required to log in to the target OS, the key pair has to be available.

- If a private image is required for changing the OS, create the private image by following the instructions provided in [Image Management Service User Guide](#).
 - If the image of a specific instance is required, ensure that a private image has been created from that instance.
 - If a local image file is required, ensure that this image file has been imported into the cloud platform and registered as a private image.
 - If a private image from another region is required, ensure that this image has been replicated to the target region.
 - If a private image from another account is required, ensure that this image has been shared with you.

Procedure

1. Log in to the management console.
2. Hover over  and choose **Compute > Elastic Cloud Server**.
3. Locate the row containing the target ECS. Choose **More > Manage Image/Disk > Change OS** in the **Operation** column.
Before changing the OS, stop the ECS or select **Automatically stop the ECSs and change their OSs**.
4. Change the **Image Type** and **Image**.

NOTE

For a yearly/monthly ECS, if the system disk size is smaller than the image size, you must detach the system disk, expand its capacity, and attach it to the ECS before changing the OS.

For details about how to expand the system disk capacity, see [Disk Capacity Expansion](#).

Figure 3-1 Changing OS

Change OS ×

The ECSs must be stopped before their OSs can be changed. If you have not stopped the ECSs, select the following option to automatically stop the ECSs.

Automatically stop the ECSs and change their OSs

Specifications c6.large.2 | 2 vCPUs | 4 GB

Current Image CentOS 8.0 64bit
System Disk Capacity: 40 GB OS Architecture: 64-bit

Select Image

Image Type Public Image Private image Shared image Marketplace image

Image CentOS CentOS 7.6 64bit C

Login Mode Password Key pair

OK Cancel

5. Set **Login Mode**.

Use the key pair if it is required for logging in to the new OS.

6. Click **OK**.

7. On the **Change ECS OS** page, confirm the specifications, read and agree to the agreement or disclaimer, and click **Submit Application**.

The task status changes to **Changing OS**. When **Changing OS** disappears, the change is successful.

NOTE

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

Follow-up Operations

- If the OS before and after the change is Linux, and auto-mount on startup has been enabled for the data disk, the auto-mounting information will be lost after the OS is changed. You will need to update the **/etc/fstab** configuration:
 - a. Write the new partition information into **/etc/fstab**.
It is a best practice to back up the **/etc/fstab** file before writing data into it.
To enable auto-mounting on startup, see [Initializing a Linux Data Disk \(fdisk\)](#).
 - b. Mount the partition to use the data disk.
mount *diskname mountpoint*
 - c. Check that the partition was mounted successfully.

df -TH

- If the OS change fails, repeat steps [2](#) to [7](#) to retry.
- If the retry fails, contact customer service for manual recovery.

4 Migrating an OS to Huawei Cloud EulerOS

4.1 Using x2hce-ca to Evaluate the Compatibility with Huawei Cloud EulerOS

4.1.1 Overview of x2hce-ca

x2hce-ca is a free tool provided by Huawei Cloud to evaluate the compatibility between applications and the OS. You can use x2hce-ca to evaluate the compatibility before migrating an OS to Huawei Cloud EulerOS.

Table 4-1 x86 public images that support compatibility evaluation

OS Series	Source OS	Target OS
Huawei Cloud EulerOS	64-bit: Huawei Cloud EulerOS 1.1	Huawei Cloud EulerOS 2.0 Standard Edition, 64-bit
EulerOS	64-bit: EulerOS 2.10, 2.9, 2.5, 2.3, and 2.2	Huawei Cloud EulerOS 2.0 Standard Edition, 64-bit
CentOS	64-bit: CentOS 7.9, 7.8, 7.7, 7.6, 7.5, 7.4, 7.3, 7.2, 7.1, and 7.0 64-bit: CentOS 8: 8.3, 8.2, 8.1, and 8.0	Huawei Cloud EulerOS 2.0 Standard Edition, 64-bit
	64-bit: CentOS 7.9 and 7.6	Huawei Cloud EulerOS 1.1 CentOS-compatible Edition

Table 4-2 Arm public images that support compatibility evaluation

OS Series	Source OS	Target OS
EulerOS	64-bit: EulerOS 2.10, 2.9, 2.8, and 2.3	Huawei Cloud EulerOS 2.0 Standard Edition, 64-bit, Arm-compatible

4.1.2 Constraints

- Do not run x2hce-ca in the runtime environment because additional resource packages are created during its installation. The x2hce-ca tool can be installed and used only on Huawei Cloud EulerOS 2.0.
- The x2hce-ca tool can scan only files in .jar, .py, .pyc, .bin, .sh, .rpm, or .ko format. RPM files must be written in C, C++, Java, or Python.
- The x2hce-ca tool does not support rollback. If a task is interrupted unexpectedly, you can retry. The residual files in the `/opt/x2hce-ca/` directory do not affect the use of the tool.
- The x2hce-ca tool can run only in the system that meets the following hardware requirements:

Table 4-3 Hardware requirements for running x2hce-ca

Hardware Type	Description
Architecture	x86_64
CPU	Dual-core or higher
Memory	At least 8 GB available memory
Hard disk	20 GB or above

4.1.3 Installing x2hce-ca

1. Confirm that the repository is configured correctly.

Check whether the parameters in the `/etc/yum.repos.d/hce.repo` file are configured correctly. The correct configuration is as follows:

```
[base]
name=HCE $releasever base
baseurl=https://repo.huaweicloud.com/hce/$releasever/os/$basearch/
enabled=1
gpgcheck=1
gpgkey=https://repo.huaweicloud.com/hce/$releasever/os/RPM-GPG-KEY-HCE-2

[updates]
name=HCE $releasever updates
baseurl=https://repo.huaweicloud.com/hce/$releasever/updates/$basearch/
.....
```

2. Install x2hce-ca.

Run `yum install -y x2hce-ca-hce.x86_64` to install the x2hce-ca tool. After the installation is complete, the directories listed in [Table 4-4](#) are generated.

Table 4-4 Directories generated after x2hce-ca is installed

Directory	Description
/var/log/x2hce-ca	Stores log files of the x2hce-ca tool.
/var/log/aparser	Stores log files of the parser.
/opt/x2hce-ca/output	Stores the reports generated by the x2hce-ca tool.
/opt/x2hce-ca/scan	Stores the application software packages to be scanned.
/opt/x2hce-ca/update	Stores the update packages and the corresponding license files.
/etc/x2hce-ca/config	Stores static configuration files.
/etc/x2hce-ca/database_2.0.0.630	Stores database files.
/usr/local/x2hce-ca	Stores program files.

- Restart the OS or run **alias x2hce-ca="python /usr/local/x2hce-ca/x2hce-ca.py"** for the **x2hce-ca** command to take effect.

4.1.4 Evaluating Software Compatibility

Scanning Methods

The x2hce-ca tool can:

- Scan a single application package in a given directory on the source OS.
- Scan all application packages in a given directory on the source OS.

Procedure

- Log in as or switch to user **root**.
- Scan the application packages to check for compatibility.

x2hce-ca scan <option> [-os_name *Source OS name*] [-target_os_name *Target OS name*]

You can replace **<option>** with either of the following values:

- Dir_Name/App_Name: scans a single application software package in a single directory.

The following uses the x86 and Arm OS architectures as examples:

For example, if you want to scan application package

NetworkManager-1.18.8-1.el7.x86_64.rpm (x86) in the **/mnt/** directory, run the following command:

```
x2hce-ca scan /mnt/NetworkManager-1.18.8-1.el7.x86_64.rpm -  
os_name centos7.9 -target_os_name hce2.0
```

For example, if you want to scan application package

NetworkManager-1.18.8-1.el7.aarch64.rpm (Arm) in the **/mnt/** directory, run the following command:

```
x2hce-ca scan /mnt/NetworkManager-1.18.8-1.el7.aarch64.rpm -
os_name EulerOSV2.0SP8arm -target_os_name hce2.0arm -arch aarch64
```

NOTE

The default value of **-arch** is **x86_64**.

- **-b Dir_Name**: scans all application packages in a single directory.

For example, if you want to scan all application packages in **directory1**, run

```
x2hce-ca scan -b directory1 -os_name centos7.9 -target_os_name hce2.0
```

NOTE

A single-level directory can contain a maximum of 750 files of up to 900 MB. Excessive software packages may cause tool failures.

Parameter	Type	Description
-os_name	String	Source OS. This parameter is optional. The default value is centos7.9 . For example, if this parameter is set to centos8.2 , CentOS 8.2 is selected as the source OS.
-target_os_name	String	Target OS. This parameter is optional. The default value is hce2.0 . For example, if this parameter is set to -target_os_name hce1.1 , the target OS is Huawei Cloud EulerOS 1.1.

3. Analyze the evaluation results.

For example, after the x2hce-ca tool scans three RPM packages in the **/tmp/x2hce-ca_test** directory, it displays the following outputs:

```
[x2hce-ca@localhost ~]$ x2hce-ca scan -b /tmp/x2hce-ca_test/ -os_name centos7.9 -target_os_name hce1.1
2022-08-31 04:21:59.808 - USER_ID:1001 - INFO - Log save directory: /var/log/x2hce-ca
2022-08-31 04:21:59.812 - USER_ID:1001 - INFO - x2hce-ca scan /tmp/x2hce-ca_test/ -os_name centos7.9 -target_os_name hce1.1 -arch x86_64 -b
2022-08-31 04:21:59.849 - USER_ID:1001 - INFO - Start analyse /tmp/x2hce-ca_test/qt-4.8.7-8.el7.x86_64.rpm
2022-08-31 04:22:01.993 - USER_ID:1001 - INFO - Start scanning Jar Package...
2022-08-31 04:22:01.993 - USER_ID:1001 - INFO - No jars found
2022-08-31 04:22:02.681 - USER_ID:1001 - INFO - Start analyse /tmp/x2hce-ca_test/texinfo-5.1-5.el7.x86_64.rpm
2022-08-31 04:22:03.547 - USER_ID:1001 - INFO - Start scanning Jar Package...
2022-08-31 04:22:03.547 - USER_ID:1001 - INFO - No jars found
2022-08-31 04:22:04.294 - USER_ID:1001 - INFO - Start analyse /tmp/x2hce-ca_test/zip-3.0-11.el7.x86_64.rpm
2022-08-31 04:22:04.654 - USER_ID:1001 - INFO - Start scanning Jar Package...
2022-08-31 04:22:04.654 - USER_ID:1001 - INFO - No jars found
2022-08-31 04:22:04.906 - USER_ID:1001 - INFO - The excel report is saved: /opt/x2hce-ca/output/software/batch-20220831042159/batch-2022083104
2159.xls
2022-08-31 04:22:05.069 - USER_ID:1001 - INFO - Generate Success! The Archive file results are saved: /opt/x2hce-ca/output/software/batch-2022
0831042159.tar.gz
2022-08-31 04:22:05.069 - USER_ID:1001 - INFO - Analyse finished, total 3 items, 3 success, 0 failed
```

- Download the compatibility evaluation reports from the **/opt/x2hce-ca/output/software/** directory.

- If the software package is compatible with the target OS, the following message is displayed:

Software assessment report Report generated 2022/12/07 17:11:16

✔ Evaluation result: The package is installed and running normally, you can [click here](#) to find the openEuler missing package.

Software to be evaluated	apache-commons-cli-1.2-13.el7.noarch.rpm
Source operating system	centos7.9
Target operating system	hce1.1
System architecture	x86_64

If the software package is not compatible with the target OS, the following message is displayed:

Software assessment report Report generated 2022/12/08 09:12:44

⚠ Evaluation results: the dependent package is not compatible, the interface is compatible, the package cannot be installed, you can [click here](#) to query the openEuler missing package.

Software to be evaluated	libreoffice-impress-4.1.4.2-3.el7.x86_64.rpm
Source operating system	centos7.9
Target operating system	hce1.1
System architecture	x86_64

- For details about the compatible dependency packages and APIs, see the Excel file with the same name as the software package.

software	src_os	target_os	arch	assessment item	runtime	generate_time	dependent count	package_pend_count	external_interface_count	interface_pend_count	interface_compatibility_percent	dependent_compatibility_percent	src_pkg	target_pkg
apache-commons-cli-1.2-13.el7.noarch.rpm	centos7.9	hce1.1	x86_64	direct dependency, function interface	4.865	20221207171116	2	0	63	0	100%	100%	apache-commons-cli-1.2-13.el7.noarch.rpm	apache-commons-cli-1.2-13.hce1.noarch.rpm
apache-commons-cli-1.2-13.el7.noarch.rpm	centos7.9	hce1.1	x86_64	direct dependency, function interface	4.865	20221207171116	2	0	63	0	100%	100%	apache-commons-cli-1.2-13.el7.noarch.rpm	apache-commons-cli-1.2-13.hce1.noarch.rpm

- If application packages fail the check, view the Excel report in the `/opt/x2hce-ca/output/software/` directory.

4.2 Migrating an OS to Huawei Cloud EulerOS 2.0

4.2.1 Constraints

- Migration is only supported for **public images** for ECSs listed in the **Supported Instance Types**.
- During the OS migration, RPM packages need to be uninstalled, installed, and updated. As a result, the OS may restart unexpectedly. Back up the system disk before the OS migration. For details, see **Creating a Cloud Server Backup**.
- It is recommended that the remaining OS memory be greater than 128 MB, the remaining system disk space (required for running the migration tool) be greater than 5 GB, and the available space of the boot partition be greater than 200 MB.
- The name of a customized RPM package must be different from that of the OS RPM package. Otherwise, the custom RPM package will be deleted by the migration tool during the migration.
- After the OS migration, the system disk type cannot be changed.

- During OS migration, certain conflicting packages are generated in the system to be migrated. They will be automatically deleted after the OS migration. For details about the conflicting packages, see [Conflicting Packages](#).

4.2.2 Migrating an OS to Huawei Cloud EulerOS

Preparing Dependent Packages

Prepare the software packages required by the migration tool:

1. Remotely connect to the source OS.

Remotely log in to the ECS to be migrated and check that the ECS can be accessed from the Internet. For details, see [Login Overview](#).

2. Check that the migration tool can access the Huawei Cloud EulerOS repository to obtain the dependent packages.

The migration tool can access the Huawei Cloud EulerOS repository in the output of `curl https://repo.huaweicloud.com/hce/2.0/os/x86_64/` displays information similar to the following:

```
% Total % Received % Xferd Average Speed Time Time Time Current
Dload Upload Total Spent Left Speed
100 3417 0 3417 0 0 373 0 --:--:-- 0:00:09 --:--:-- 696
<!doctype html>
<html>
<head>
<meta charset="utf-8">
<title></title>
<link rel="stylesheet" href="/repository/static/css/style.css" type="text/css"/>
<style>
* {
font-family: 'Verdana', sans-serif;
margin: 0;
padding: 0;
-webkit-box-sizing: border-box;
-moz-box-sizing: border-box;
box-sizing: border-box;
}
.....
```

3. Configure the repository of the source OS and ensure that the migration tool can obtain dependent software from this repository.

The repository address depends on the OS.

4. Install dependent packages.

- a. Install Python software packages.

```
[root@localhost ~]# yum install -y python //Run this command in any directory.
```

- b. (Optional) Create a soft link.

NOTE

The following steps are only available for CentOS 8 and EulerOS 2.10 and 2.9.

- i. Install Python 3.0 software packages.

```
[root@localhost ~]# yum install -y python 3//Run this command in any directory.
```

- ii. Check whether the Python symbolic link exists.

- If the Python symbolic link does not exist, go to [4.b.iii](#).
- If the Python symbolic link exists but is not linked to Python 3.0, run the following command to delete the original link and then go to [4.b.iii](#).


```
[root@localhost]# unlink /usr/bin/python
```

- If the Python symbolic link exists and has been linked to Python 3.0, go to [Installing the Migration Tool and Checking Migration Conditions](#).

iii. Create a Python symbolic link.

```
[root@localhost]# python
-bash: /usr/bin/python: No such file or directory //Indicates that the Python symbolic link
does not exist.
[root@localhost]# cd /usr/bin/ //Switch to the /usr/bin directory.
[root@localhost bin]# ln -s python3 python //Create a Python symbolic link.
[root@localhost bin]# python
Python 3.6.8 (default, Apr 16 2020, 01:36:27)
[GCC 8.3.1 20191121 (Red Hat 8.3.1-5)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>>
//Press Ctrl+D to exit.
```

Installing the Migration Tool and Checking Migration Conditions

1. Download tool package **centos2hce2-*.rpm** of the latest version from the [Huawei Cloud open-source image site](#). You can also contact customer service to obtain the download link from O&M engineers.

The asterisk (*) indicates the version of the migration tool. In this example, **centos2hce2-1.0.0-0.0.30.hce2.x86_64.rpm** is used. Replace with the actual installation package name during your operation.

```
[root@localhost test]# wget https://repo.huaweicloud.com/hce/2.0/updates/x86_64/Packages/
centos2hce2-1.0.0-0.0.30.hce2.x86_64.rpm //Download the centos2hce2-*.rpm package.
[root@localhost test]# ls //Check whether the download is successful.
centos2hce2-1.0.0-0.0.30.hce2.x86_64.rpm
```

2. Install the migration tool.

```
[root@localhost test]# rpm -ivh centos2hce2-1.0.0-0.0.30.hce2.x86_64.rpm --nodeps
warning: centos2hce2-1.0.0-0.0.30.hce2.x86_64.rpm: Header V4 RSA/SHA256 Signature, key ID
a8def926: NOKEY
Verifying... ##### [100%]
Preparing... ##### [100%]
Updating / installing...
1:centos2hce2-1.0.0-0.0.6.hce2 ##### [100%]
```

3. Configure the backup directory for the system disk of the source OS.

Before the OS migration, the migration tool automatically backs up all data of the system software to the backup directory.

You can run **vim /etc/centos2hce2.conf** to configure the **backup_dir** field in the **centos2hce2.conf** file. The default value of **backup_dir** is **/mnt/sdb/.osbak**.

```
# backup dir
backup_dir = "/mnt/sdb/.osbak" #Change the backup directory.
```

NOTE

- To prevent system data from being lost during the migration, configure a backup directory.
 - During OS migration, the migration tool checks the space of the backup directory. To prevent check failures caused by insufficient space, use an independent data disk (for example, /dev/sdb/ mounted to /mnt/sdb/) as the backup directory.
 - Do not use the tmpfs file system (such as /dev and /run) as the backup directory. If such tmpfs file system is used as the backup directory, files in the file system will be lost after the system is restarted.
4. Configure the migration parameters.

a. Set Web migration.

To perform a Web migration, the system needs to download the RPM package. The network cannot be disconnected during the download.

To do so, you need to modify the **centos2hce2.conf** file as follows:

```
[repo_relation]
.....
# default yum source, val: web or iso
default_yum_source = 'web'
.....
# if web as source, web link config as follow
web_link_dir = "https://repo.huaweicloud.com/hce/2.0/os/x86_64/;https://repo.huaweicloud.com/hce/2.0/updates/x86_64/"
```

Table 4-5 Parameter settings for a Web migration

Parameter	Description
default_yum_source	Set this parameter to web .
web_link_dir	Specifies the source address of base repo and updates repo of the Huawei Cloud EulerOS. Separate multiple repo sources with semicolons (;). This is an example: https://repo.huaweicloud.com/hce/2.0/os/x86_64/;https://repo.huaweicloud.com/hce/2.0/updates/x86_64/.

b. Configure the **isclose_modules** parameter, which is required only for CentOS 8.

CentOS 8 allows you to batch install RPM packages as a module, but Huawei Cloud EulerOS does not. Therefore, you need to disable the module function before performing an OS migration.

- **yes** (default): The system closes its modules before the migration.
- **no**: The system does not close its modules before the migration. If any module is enabled, the migration is interrupted.

```
[system]
# whether close modules, if value is no, system may be not migrate
isclose_modules = "yes"
```

 **NOTE**

- You can run **dnf module list** to view all running modules.
- You can run **dnf module list | grep '\[e\]** to view enabled modules.

5. Run **centos2hce2.py --check all** to check whether the current OS can be migrated.
 - If the message "Environment check passed!" is displayed, continue with the migration.
 - If the message "call migration failed" is displayed, perform step 6 to handle the exceptions.
6. Install software dependencies for the migration tool.

Run **centos2hce2.py --install all** to back up the files in the system, install software dependencies, and complete preprocessing.

If the following information is displayed, the software dependencies have been installed and preprocessing has been completed. You need to perform step 5 again to check the environment.

```
2022-08-19 03:12:58,373-INFO-centos2hce2.py-[line:832]: Dependency packages already exist!  
2022-08-19 03:12:58,373-INFO-centos2hce2.py-[line:891]: migrate install depend options finished
```

7. (Optional) Repeat the backup.

Run **centos2hce2.py --backup force** to back up the files in the system to the path configured in step 3.

NOTE

The software dependencies installed in step 6 are also backed up after this command is executed.

Migrating an OS to Huawei Cloud EulerOS

1. Run **centos2hce2.py --upgrade all** to migrate an OS to Huawei Cloud EulerOS.

If **migrate success** is displayed, the OS migration was successful. If the migration fails, perform a rollback by referring to 1.

```
[root@localhost ~]# centos2hce2.py --upgrade all  
2022-08-19 03:19:21,090-INFO-centos2hce2.py-[line:1233]: start migration  
2022-08-19 03:19:21,075-INFO-centos2hce2.py-[line:425]: config sut succeed  
2022-08-19 03:19:21,096-INFO-centos2hce2.py-[line:901]: SELinux service switches to permissive mode and has been temporarily closed!  
[ INFO ] - [initramfs]: set command line value done  
2022-08-19 03:30:35,117-INFO-centos2hce2.py-[line:1032]: migrate success  
2022-08-19 03:30:35,467-INFO-centos2hce2.py-[line:989]: python link is /usr/bin/python3.9  
2022-08-19 03:30:35,482-INFO-centos2hce2.py-[line:994]: create python link succeed  
2022-08-19 03:30:35,482-INFO-centos2hce2.py-[line:1044]: migrate excute finished
```

NOTE

- The migration command cannot be executed in the background in Linux.
- The **--simple_name** parameter can be added so that the abbreviation of Huawei Cloud EulerOS is displayed in the grub menu after the migration.

2. Reboot the server by running **reboot**. If the **reboot** command does not respond, run **reboot -f** instead.

After the system reboot, run **cat /etc/hce-release** to view the OS information and run **uname -a** to view the OS kernel information.

If Huawei Cloud EulerOS is displayed, the OS migration was successful. Otherwise, the migration failed. [Contact technical support](#).

```
[root@localhost ~]# cat /etc/hce-release  
Huawei Cloud EulerOS release 2.0 (West Lake)  
[root@localhost ~]# uname -a  
Linux localhost.localdomain 5.10.0-60.18.0.50.h425_1.hce2.x86_64 #1 SMP Thu Aug 18 16:31:04 UTC 2022 x86_64 x86_64 x86_64 GNU/Linux
```

NOTE

After the OS is migrated to Huawei Cloud EulerOS, the name of the original OS is still displayed on the console. You can manually update the OS name on the console.

3. Delete the files of the source OS's components.

After the OS migration, the target OS's components replace the source OS's components. However, the files of the source OS's components are still stored in the system. You need to run **centos2hce2.py --precommit upgrade** to delete such files.

If the message "upgrade precommit success" is displayed, the files have been successfully deleted.

```
2022-08-19 03:52:32,871-INFO-centos2hce2.py-[line:1112]: remove geolite2-city-20180605-1.el8.noarch succeed
2022-08-19 03:52:32,984-INFO-centos2hce2.py-[line:1112]: remove subscription-manager-rhsm-certificates-1.26.16-1.el8.0.1.x86_64 succeed
2022-08-19 03:52:33,543-INFO-centos2hce2.py-[line:1112]: remove cockpit-ws-211.3-1.el8.x86_64 succeed
2022-08-19 03:52:33,543-INFO-centos2hce2.py-[line:1113]: handle clean rpms finished
2022-08-19 03:52:37,082-INFO-centos2hce2.py-[line:1124]: remove useless link /usr/lib/gcc/x86_64-linux-gnu/10.3.1/32/libasan.a
2022-08-19 03:52:37,906-INFO-centos2hce2.py-[line:1124]: remove useless link /usr/lib/gcc/x86_64-linux-gnu/10.3.1/32/libatomic.a
2022-08-19 03:52:37,910-INFO-centos2hce2.py-[line:1124]: remove useless link /usr/lib/gcc/x86_64-linux-gnu/10.3.1/32/libgomp.so
2022-08-19 03:52:37,915-INFO-centos2hce2.py-[line:1124]: remove useless link /usr/lib/gcc/x86_64-linux-gnu/10.3.1/32/libitm.a
2022-08-19 03:52:37,919-INFO-centos2hce2.py-[line:1124]: remove useless link /usr/lib/gcc/x86_64-linux-gnu/10.3.1/32/libquadmath.a
2022-08-19 03:52:37,923-INFO-centos2hce2.py-[line:1124]: remove useless link /usr/lib/gcc/x86_64-linux-gnu/10.3.1/32/libubsan.a
2022-08-19 03:52:37,927-INFO-centos2hce2.py-[line:1124]: remove useless link /usr/share/doc/e2fsprogs/RELEASE-NOTES
2022-08-19 03:52:37,929-INFO-centos2hce2.py-[line:1126]: remove useless link finished
2022-08-19 03:52:37,930-INFO-centos2hce2.py-[line:1132]: clean system finished and migrate succeed
2022-08-19 03:52:37,930-INFO-centos2hce2.py-[line:1200]: upgrade precommit success
```

NOTE

The deletion can be performed for multiple times.

4. (Optional) Modify Cloud-Init configurations.

- Skip this step if Cloud-Init is running normally in the source OS and Cloud-Init is an RPM package.
- If Cloud-Init is running normally in the source OS and Cloud-Init is a file (for example, a CentOS 7 file) other than an RPM package, modify `/etc/cloud/cloud.cfg` as follows:

a. Enable remote login using the password for user `root` and allow SSH access to `root`.

Set `disable_root` to `0` to keep `root` enabled. Set `ssh_pwauth` to `1` to allow remote login using a password. Set `lock_passwd` to `False` to not lock the password.

```
users:
- name: root
  lock_passwd: False

disable_root: 0
ssh_pwauth: 1
```

b. Run `/usr/bin/cloud-init init --local`.

If there are no errors and the Cloud-Init version is displayed, Cloud-Init has been correctly configured.

```
[root@localhost ~]# /usr/bin/cloud-init init --local
Cloud-init v. 21.4 running 'init-local' at Fri, 22 Jul 2022 07:43:21 +0000. Up 602150.81 seconds.
[root@localhost ~]#
```

5. (Optional) If SELinux service is disabled during the migration but needs to be enabled after the migration, run `centos2hce2.py --precommit upg-selinux` to enable the SELinux service. This command is executed twice. After each execution, the system restarts.

a. Run `centos2hce2.py --precommit upg-selinux`.

```
[root@localhost ~]# centos2hce2.py --precommit upg-selinux
2022-08-21 23:46:23,891-INFO-centos2hce2.py-[line:1239]: precommit migration
2022-08-21 23:46:23,891-INFO-centos2hce2.py-[line:1149]: begin to set selinux
2022-08-21 23:46:23,892-INFO-centos2hce2.py-[line:1157]: grub path is /boot/grub2/grub.cfg
2022-08-21 23:46:23,895-INFO-centos2hce2.py-[line:1162]: sed selinux succeed
2022-08-21 23:46:23,897-INFO-centos2hce2.py-[line:1167]: create autorelabel file succeed
2022-08-21 23:46:23,901-INFO-centos2hce2.py-[line:1172]: modify selinux config succeed
2022-08-21 23:46:23,901-INFO-centos2hce2.py-[line:1174]: create phase 1 flag file succeed
2022-08-21 23:46:23,901-INFO-centos2hce2.py-[line:1184]: selinux has been set, please reboot now
2022-08-21 23:46:23,901-INFO-centos2hce2.py-[line:1206]: upgrade precommit selinux success
[root@localhost ~]# reboot
```

b. After the system is restarted, run `centos2hce2.py --precommit upg-selinux` again.

```
[root@localhost ~]# centos2hce2.py --precommit upg-selinux
2022-08-21 23:57:07,576-INFO-centos2hce2.py-[line:1239]: precommit migration
2022-08-21 23:57:07,576-INFO-centos2hce2.py-[line:1176]: now begin to set selinux phase 2
2022-08-21 23:57:07,580-INFO-centos2hce2.py-[line:1181]: modify selinux config succeed
2022-08-21 23:57:07,580-INFO-centos2hce2.py-[line:1183]: create phase 2 flag file succeed
```

```
2022-08-21 23:57:07,580-INFO-centos2hce2.py-[line:1184]: selinux has been set, please reboot now
2022-08-21 23:57:07,580-INFO-centos2hce2.py-[line:1206]: upgrade precommit selinux success
[root@localhost ~]# reboot
```

c. After the second restart, run **getenforce** to check the SELinux status. If it is **Enforcing**, SELinux has been enabled.

```
[root@localhost ~]# getenforce
Enforcing
```

6. (Optional) After the migration is complete, delete the source OS data.

After the migration, the system data of the source OS is still stored in the new system and occupies a large amount of memory. You can run **centos2hce2.py --commit all** to clear the data.

The system will automatically delete the system data of the source OS, including the system data in the backup directory mentioned in step 3.

NOTICE

After the command is executed, the OS cannot be rolled back.

```
[root@localhost ~]# centos2hce2.py --commit all
2022-08-22 04:45:32,601-INFO-centos2hce2.py-[line:1242]: commit migration
```

Rolling Back the OS

1. Roll back the OS if needed.
 - a. Run **centos2hce2.py --rollback all** to roll back the system. After the rollback, run **reboot** to restart the system.

```
[root@localhost ~]# centos2hce2.py --rollback all
2022-08-22 04:03:12,107-INFO-centos2hce2.py-[line:1236]: Start rollback
[ INFO ] - [sut]: ===== Pre Rollback Stage =====
[ INFO ] - [sut]: skip system version check during rollback
[ INFO ] - [sut]: start to do rollback before reboot
[ INFO ] - [sut]: set dracut module for rollback.
[ INFO ] - [sut]: reboot aborted.You need reboot as soon as possible
[root@localhost ~]# reboot
```

- b. Run **centos2hce2.py --precommit rollback** to restore the environment.

```
[root@localhost ~]# centos2hce2.py --precommit rollback
2022-08-22 04:36:13,902-INFO-centos2hce2.py-[line:1239]: precommit migration
2022-08-22 04:36:13,904-INFO-centos2hce2.py-[line:1071]: /opt/migrate//rsync_backup is not exists, skip it
2022-08-22 04:36:13,905-INFO-centos2hce2.py-[line:483]: sut not backup no need rollback
2022-08-22 04:36:17,996-INFO-centos2hce2.py-[line:1194]: rollback precommit success
```

2. (Optional) If SELinux has been enabled before the migration, the SELinux service will be automatically disabled during the migration. If necessary, manually enable the SELinux status after the rollback.

- a. Run **centos2hce2.py --precommit rbk-selinux**.

```
[root@localhost ~]# centos2hce2.py --precommit rbk-selinux
2022-09-05 03:58:37,015-INFO-centos2hce2.py-[line:1401]: precommit migration
2022-09-05 03:58:37,047-INFO-centos2hce2.py-[line:1319]: now begin to set selinux
2022-09-05 03:58:37,051-INFO-centos2hce2.py-[line:1324]: modify selinux config succeed
2022-09-05 03:58:37,051-INFO-centos2hce2.py-[line:1325]: selinux has been set, please reboot now
2022-09-05 03:58:37,051-INFO-centos2hce2.py-[line:1340]: set rollback selinux succeed
2022-09-05 03:58:37,051-INFO-centos2hce2.py-[line:1365]: upgrade precommit selinux success
```

- b. Run **reboot** to restart the system.

```
[root@localhost ~]# reboot
```
- c. After the system is restarted, you can see that SELinux is enabled.

```
[root@localhost ~]# getenforce  
Enforcing
```

3. Clear data from the OS.

Run **centos2hce2.py --commit all** to clear the data.

The system will automatically delete the system data of the source and target OSs, including the system data in the backup directory mentioned in step 3.

```
[root@localhost ~]# centos2hce2.py --commit all  
2022-08-22 04:45:32,601-INFO-centos2hce2.py-[line:1242]: commit migration
```

4.2.3 Conflicting Packages

Table 4-6 Conflicting packages in CentOS 8 series

CentOS Version	Conflicting Packages
CentOS 8.0	rust-doc intel-gpu-tools netcf-libs redhat-rpm-config asciidoc gnuplot-common perf tigervnc-icons libpq-devel paratype-pt-sans-caption-fonts scala-apidoc java-11-openjdk-devel java-11-openjdk-headless java-1.8.0-openjdk-headless dovecot systemd-journal-remote pcp-manager pcp-webapi libguestfs-java-devel libguestfs-javadoc icedtea-web-javadoc systemtap-runtime-java java-1.8.0-openjdk-accessibility java-1.8.0-openjdk-demo ant tigervnc-server-applet java-atk-wrapper java-11-openjdk-guava20 javapackages-tools jboss-jaxrs-2.0-api maven-shared-utils tagsoup cdi-api libbase geronimo-annotation pentaho-reporting-flow-engine maven-resolver-api apache-commons-codec maven-lib jansi-native maven-wagon-provider-api libguestfs-java apache-commons-cli istack-commons-tools jline plexus-cipher istack-commons-runtime jcl-over-slf4j apache-commons-io maven-resolver-spi maven-wagon-file httpcomponents-core icedtea-web glassfish-el-api aopalliance hawtjni-runtime plexus-containers-component-annotations flute jboss-annotations-1.2-api liblayout java-1.8.0-openjdk-postgresql-jdbc mariadb-java-client plexus-sec-dispatcher google-guice libformula jdeparser ant-lib maven-wagon-http-shared jboss-logging plexus-classworlds slf4j librepository ongres-scam-client sisu-plexus libfonts plexus-interpolation java-1.8.0-openjdk-src plexus-utils scala-swing maven-wagon-http ongres-scam maven-resolver-impl libloader httpcomponents-client atinject apache-commons-logging maven-resolver-connector-basic jansi jsoup maven-resolver-util jboss-interceptors-1.2-api libreoffice-ure byteman sac apache-commons-lang3 libserializer scala maven-resolver-transport-wagon jboss-logging-tools sisu-inject libreoffice-core java-1.8.0-openjdk-devel

CentOS Version	Conflicting Packages
CentOS 8.1	kernel-rpm-macros intel-gpu-tools netcf-libs redhat-rpm-config asciidoc gnuplot-common perf tigervnc-icons libpq-devel paratype-pt-sans-caption-fonts java-1.8.0-openjdk-headless java-11-openjdk-headless java-11-openjdk-devel pcp-pmda-rpm pcp-pmda-podman scala-apidoc libguestfs-java-devel libguestfs-javadoc icedtea-web-javadoc systemtap-runtime-java java-1.8.0-openjdk-accessibility java-1.8.0-openjdk-demo ant tigervnc-server-applet java-atk-wrapper java-11-openjdk jansi- native hawtjni-runtime ongres-scram jboss-annotations-1.2-api liblayout atinject plexus-utils istack-commons-tools jline apache-commons-io ongres-scram-client maven-shared-utils maven-resolver-impl libfonts jsoup apache-commons-codec glassfish-el-api jdeparser maven-resolver-util scala-swing tagsoup google-guice istack-commons-runtime jcl-over-slf4j pentaho-reporting-flow-engine maven-resolver-api maven- resolver-connector-basic libloader slf4j apache-commons-cli maven-wagon-provider-api maven-resolver-transport-wagon byteman httpcomponents-client jna java-1.8.0-openjdk-devel maven-lib libreoffice-core java-1.8.0-openjdk-src javapackages- tools plexus-cipher cdi-api jboss-logging sisu-inject httpcomponents-core guava20 sac libbase jboss-jaxrs-2.0-api java-1.8.0-openjdk libserializer plexus-containers-component- annotations jboss-interceptors-1.2-api jboss-logging-tools libguestfs-java ant-lib libreoffice-ure maven-resolver-spi maven- wagon-file jansi maven-wagon-http-shared apache-commons- lang3 postgresql-jdbc mariadb-java-client plexus-sec-dispatcher sisu-plexus scala plexus-classworlds flute maven-wagon-http icedtea-web libformula plexus-interpolation aopalliance geronimo-annotation librepository apache-commons-logging

CentOS Version	Conflicting Packages
CentOS 8.2	<p>python-psycopg2-doc exiv2 llvm-googletest adwaita-qt llvm-static rust-doc intel-gpu-tools netcf-libs flatpak-session-helper asciidoc perf tigervnc-icons paratype-pt-sans-caption-fonts java-1.8.0-openjdk-headless java-11-openjdk-devel java-11-openjdk-headless scala-apidoc libguestfs-java-devel libguestfs-javadoc icedtea-web-javadoc systemtap-runtime-java java-1.8.0-openjdk-accessibility java-1.8.0-openjdk-demo ant tigervnc-server-applet java-atk-wrapper java-11-openjdk jboss-annotations-1.2-api cdi-api ongres-scrum maven-resolver-util apache-commons-codec istack-commons-tools icedtea-web plexus-classworlds plexus-utils maven-wagon-http-shared atinject javapackages-tools istack-commons-runtime jline geronimo-annotation jansi jdeparser byteman liblayout maven-resolver-transport-wagon jmc-core ant-lib libreoffice-core jansi-native jcl-over-slf4j slf4j ee4j-parent libfonts maven-wagon-http jboss-logging jboss-interceptors-1.2-api tagsoup httpcomponents-client plexus-containers-component-annotations apache-commons-lang3 jaf java-1.8.0-openjdk-src jsoup guava20 flute apache-commons-cli libbase ongres-scrum-client jboss-logging-tools plexus-interpolation libloader librepository libreoffice-ure scala-swing jboss-jaxrs-2.0-api maven-resolver-spi maven-lib apache-commons-io hawtjni-runtime google-guice aopalliance libguestfs-java postgresql-jdbc jna glassfish-el-api maven-resolver-impl java-1.8.0-openjdk-directory-maven-plugin mariadb-java-client httpcomponents-core maven-wagon-file maven-wagon-provider-api owasp-java-encoder libserializer maven-shared-utils plexus-cipher java-1.8.0-openjdk-devel plexus-sec-dispatcher pentaho-reporting-flow-engine maven-resolver-api sac scala libformula sisu-inject apache-commons-logging maven-resolver-connector-basic sisu-plexus</p>

CentOS Version	Conflicting Packages
CentOS 8.3	netcf-libs rust-doc git-credential-libsecret texlive-context intel-gpu-tools flatpak-session-helper asciidoc perf tigervnc-icons paratype-pt-sans-caption-fonts java-1.8.0-openjdk-headless java-11-openjdk-devel java-11-openjdk-headless libguestfs-java-devel libguestfs-javadoc icedtea-web-javadoc systemtap-runtime-java java-1.8.0-openjdk-accessibility java-1.8.0-openjdk-demo ant tigervnc-server-applet java-atk-wrapper java-11-openjdk exiv2 llvm-googletest adwaita-qt llvm-static python-psycopg2-doc scala-apidoc libXau libappstream-glib jmc-core byteman libfonts jaf jcl-over-slf4j mariadb-java-client tagsoup libguestfs-java jsoup apache-commons-cli sisu-inject jansi-native jna apache-commons-lang3 flute librepository javapackages-tools cdi-api ongres-scam java-1.8.0-openjdk-devel sisu-plexus istack-commons-runtime jboss-logging guava20 java-1.8.0-openjdk-src maven-resolver-util geronimo-annotation hawtjni-runtime jboss-annotations-1.2-api ongres-scam-client maven-resolver-connector-basic slf4j sac apache-commons-codec atinject maven-wagon-http libreoffice-ure plexus-cipher jboss-interceptors-1.2-api jline pentaho-reporting-flow-engine httpcomponents-core liblayout istack-commons-tools jdeparser maven-wagon-provider-api ee4j-parent apache-commons-io maven-resolver-spi jboss-logging-tools plexus-sec-dispatcher plexus-containers-component-annotations jboss-jaxrs-2.0-api scala libbase libreoffice-core httpcomponents-client directory-maven-plugin java-1.8.0-openjdk libformula maven-wagon-file maven-shared-utils aopalliance glassfish-el-api owasp-java-encoder postgresql-jdbc libloader google-guice plexus-classworlds ant-lib maven-resolver-api plexus-interpolation java-1.8.0-openjdk-slowdebug maven-resolver-impl java-1.8.0-openjdk-headless-slowdebug prometheus-jmx-exporter maven-resolver-transport-wagon jolokia-jvm-agent maven-wagon-http-shared maven-lib jansi HdrHistogram apache-commons-logging plexus-utils icedtea-web libserializer scala-swing

CentOS Version	Conflicting Packages
CentOS 8.4	<p>python-psycopg2-doc anaconda-install-env-deps hwloc-gui python3-lit exiv2 cups-filters cups-filters-libs gutenprint adwaita-qt cups cups-lpd hplip-common hwloc-libs gutenprint-doc gutenprint-libs gutenprint-libs-ui hwloc foomatic-db-ppds foomatic-db python39-pip python39-setuptools python39-numpy python39-chardet python39-psutil python39-urllib3 python39-requests python39-wheel libasan6 paratype-pt-sans-caption-fonts python39-six python39-idna python39-ply python39-pyyaml python39-pycparser python39-lxml python39-pysocks rust-doc netcf-libs git-credential-libsecret texlive-context flatpak-session-helper asciidoc intel-gpu-tools tigervnc-icons jmc-core byteman libfonts jaf jcl-over-slf4j mariadb-java-client tagsoup libguestfs-java jsoup apache-commons-cli sisu-inject jansi-native jna apache-commons-lang3 flute librepository javapackages-tools cdi-api ongres-scam java-1.8.0-openjdk-devel sisu-plexus istack-commons-runtime jboss-logging guava20 java-1.8.0-openjdk-src maven-resolver-util geronimo-annotation hawtjni-runtime jboss-annotations-1.2-api ongres-scam-client maven-resolver-connector-basic slf4j sac apache-commons-codec atinject maven-wagon-http libreoffice-ure plexus-cipher jboss-interceptors-1.2-api jline pentaho-reporting-flow-engine httpcomponents-core liblayout istack-commons-tools jdeparser maven-wagon-provider-api ee4j-parent apache-commons-io maven-resolver-spi jboss-logging-tools plexus-sec-dispatcher plexus-containers-component-annotations jboss-jaxrs-2.0-api scala libbase libreoffice-core httpcomponents-client directory-maven-plugin java-1.8.0-openjdk libformula maven-wagon-file maven-shared-utils aopalliance glassfish-el-api owasp-java-encoder postgresql-jdbc libloader google-guice plexus-classworlds ant-lib maven-resolver-api plexus-interpolation java-1.8.0-openjdk-slowdebug maven-resolver-impl java-1.8.0-openjdk-headless-slowdebug prometheus-jmx-exporter maven-resolver-transport-wagon jolokia-jvm-agent maven-wagon-http-shared maven-lib jansi HdrHistogram apache-commons-logging plexus-utils icedtea-web libserializer scala-swing java-1.8.0-openjdk-headless java-11-openjdk-devel java-11-openjdk-headless libguestfs-java-devel libguestfs-javadoc icedtea-web-javadoc systemtap-runtime-java java-1.8.0-openjdk-accessibility java-1.8.0-openjdk-demo ant java-atk-wrapper java-11-openjdk scala-apidoc libappstream-glib PackageKit-gtk3-module gnome-software flatpak-libs PackageKit-glib PackageKit-gstreamer-plugin libpq-devel poppler perf</p>

CentOS Version	Conflicting Packages
CentOS 8.5	bluez python-psycopg2-doc perl-Devel-Peek OpenIPMI-libs anaconda-install-env-deps postfix-mysql perl-Devel-SelfStubber metacity bluez-libs libicu vte-profile qt5-qttools-examples exiv2 cups-filters cups-filters-libs gutenprint gnome-session cups cups-lpd hplip-common hwloc gnome-session-wayland-session gutenprint-doc gutenprint-libs gutenprint-libs-ui gnome-session-xsession foomatic-db-ppds foomatic-db gnome-classic-session gnome-shell-extension-apps-menu gnome-shell-extension-auto-move-windows gnome-shell-extension-drive-menu gnome-shell-extension-launch-new-instance gnome-shell-extension-native-window-placement gnome-shell-extension-places-menu gnome-shell-extension-screenshot-window-sizer gnome-shell-extension-user-theme gnome-shell-extension-window-list gnome-shell-extension-windowsNavigator gnome-shell-extension-workspace-indicator python39-six python39-idna python39-ply python39-pyyaml python39-pycparser python39-psutil python39-urllib3 python39-lxml python39-pysocks xorg-x11-server-Xwayland compat-hwloc1 bluez-obexd bluez-hid2hci netcf-libs git-credential-libsecret texlive-context flatpak-session-helper asciidoc intel-gpu-tools tigervnc-icons libasan6 paratype-pt-sans-caption-fonts pcp-pmda-podman jmc-core byteman libfonts jaf jcl-over-slf4j mariadb-java-client tagsoup libguestfs-java jsoup apache-commons-cli sisu-inject jansi-native jna apache-commons-lang3 flute librepository javapackages-tools cdi-api ongres-scrum java-1.8.0-openjdk-devel sisu-plexus istack-commons-runtime jboss-logging guava20 java-1.8.0-openjdk-src maven-resolver-util geronimo-annotation hawtjni-runtime jboss-annotations-1.2-api ongres-scrum-client maven-resolver-connector-basic slf4j sac apache-commons-codec atinject maven-wagon-http libreoffice-ure plexus-cipher jboss-interceptors-1.2-api jline pentaho-reporting-flow-engine httpcomponents-core liblayout istack-commons-tools jdeparser maven-wagon-provider-api ee4j-parent apache-commons-io maven-resolver-spi jboss-logging-tools plexus-sec-dispatcher plexus-containers-component-annotations jboss-jaxrs-2.0-api scala libbase libreoffice-core httpcomponents-client directory-maven-plugin java-1.8.0-openjdk libformula maven-wagon-file maven-shared-utils aopalliance glassfish-el-api owasp-java-encoder postgresql-jdbc libloader google-guice plexus-classworlds ant-lib maven-resolver-api plexus-interpolation java-1.8.0-openjdk-slowdebug maven-resolver-impl java-1.8.0-openjdk-headless-slowdebug prometheus-jmx-exporter maven-resolver-transport-wagon jolokia-jvm-agent maven-wagon-http-shared maven-lib jansi HdrHistogram apache-commons-logging plexus-utils icedtea-web libserializer scala-swing java-1.8.0-openjdk-headless java-11-openjdk-devel java-11-openjdk-headless libguestfs-java-devel libguestfs-javadoc icedtea-web-javadoc systemtap-runtime-java java-1.8.0-

CentOS Version	Conflicting Packages
	openjdk-accessibility java-1.8.0-openjdk-demo ant java-atk-wrapper java-11-openjdk scala-apidoc libappstream-glib PackageKit-gtk3-module gnome-software flatpak-libs PackageKit-glib PackageKit-gstreamer-plugin coreos-installer-bootinfra OpenIPMI rust cargo perf flatpak hplip-libs nautilus gutenprint-cups libgtop2 PackageKit libsane-hpaio PackageKit-command-not-found xorg-x11-drv-wacom-serial-support clutter clutter-gtk clutter-gst3 cheese-libs cheese gnome-initial-setup gnome-control-center clutter-gst2

Table 4-7 Conflicting packages in CentOS 7 series

CentOS Version	Conflicting Packages
CentOS 7.0	texlive-kpathsea-lib libdhash libref_array libbasicobjects qemu-kvm-tools texlive-dvipdfm-bin texlive-dvipdfm tomcat-servlet-3.0-api gnuplot-common postgresql-devel tigervnc-icons squid perf dovecot dovecot-mysql dovecot-pgsql dovecot-pigeonhole lvm2-cluster
CentOS 7.1	texlive-kpathsea-lib libdhash libref_array qemu-kvm-tools texlive-dvipdfm-bin tomcat-servlet-3.0-api gnuplot-common squid tigervnc-icons postgresql-devel perf dovecot dovecot-mysql dovecot-pgsql dovecot-pigeonhole lvm2-cluster texlive-dvipdfm libcacard
CentOS 7.2	texlive-kpathsea-lib libdhash qemu-kvm-tools rdma-ndd texlive-dvipdfm texlive-dvipdfm-bin dstat tomcat-servlet-3.0-api gnuplot-common perf squid tigervnc-icons tigervnc-icons postgresql-devel dovecot dovecot-pgsql dovecot-pigeonhole lvm2-cluster ipa-server-trust-ad
CentOS 7.3	spice-glib texlive-kpathsea-lib libdhash qemu-kvm-tools rdma-ndd texlive-dvipdfm texlive-dvipdfm-bin dstat tomcat-servlet-3.0-api gnuplot-common perf squid tigervnc-icons postgresql-devel dovecot dovecot-mysql dovecot-pgsql dovecot-pigeonhole lvm2-cluster pcp-pmda-kvm pcp-pmda-rpm spice-gtk3 vinagre ipa-server ipa-server-trust-ad
CentOS 7.4	spice-glib texlive-kpathsea-lib libdhash qemu-kvm-tools texlive-dvipdfm-bin texlive-dvipdfm dstat tomcat-servlet-3.0-api gnuplot-common perf squid tigervnc-icons postgresql-devel lvm2-cluster spice-gtk3 vinagre
CentOS 7.5	spice-glib texlive-kpathsea-lib qemu-kvm-tools texlive-dvipdfm-bin texlive-dvipdfm dstat tomcat-servlet-3.0-api gnuplot-common perf squid tigervnc-icons postgresql-devel lvm2-cluster spice-gtk3 vinagre

CentOS Version	Conflicting Packages
CentOS 7.6	shim-x64 spice-glib adwaita-gtk2-theme texlive-kpathsea-lib qemu-kvm-tools texlive-dvipdfm-bin texlive-dvipdfm dstat tomcat-servlet-3.0-api gnuplot-common cockpit-ws perf squid tigervnc-icons postgresql-devel java-11-openjdk-headless lvm2-cluster spice-gtk3 vinagre
CentOS 7.7	shim-x64 spice-glib openmpi adwaita-gtk2-theme exiv2 texlive-kpathsea-lib qemu-kvm-tools texlive-dvipdfm-bin texlive-dvipdfm dstat tomcat-servlet-3.0-api cockpit-ws gnuplot-common perf squid tigervnc-icons postgresql-devel java-11-openjdk-headless lvm2-cluster spice-gtk3 openmpi-devel vinagre
CentOS 7.8	shim-x64 spice-glib openmpi adwaita-gtk2-theme exiv2 texlive-kpathsea-lib qemu-kvm-tools texlive-dvipdfm-bin texlive-dvipdfm dstat tomcat-servlet-3.0-api cockpit-ws gnuplot-common perf squid tigervnc-icons postgresql-devel java-11-openjdk-headless lvm2-cluster spice-gtk3 openmpi-devel vinagre
CentOS 7.9	spice-glib openmpi adwaita-gtk2-theme exiv2 gnuplot-common texlive-kpathsea-lib perf qemu-kvm-tools texlive-dvipdfm-bin texlive-dvipdfm dstat tomcat-servlet-3.0-api cockpit-ws squid tigervnc-icons postgresql-devel java-11-openjdk-headless lvm2-cluster spice-gtk3 openmpi-devel

Table 4-8 Conflicting packages in Huawei Cloud EulerOS

Huawei Cloud EulerOS Version	Conflicting Packages
Huawei Cloud EulerOS 1.1	spice-glib openmpi exiv2 sg3_utils spice-gtk3 openmpi-devel kernel-hcek tomcat-servlet-3.0-api kernel-hcek-devel dstat gnuplot-common cockpit-ws perf squid postgresql-devel java-11-openjdk-headless lvm2-cluster fcoe-utils libblockdev udisks2 python-blivet device-mapper-multipath device-mapper-multipath-libs libblockdev-crypto libblockdev-fs libblockdev-loop libblockdev-mdraid libblockdev-nvdim libblockdev-part libblockdev-swap libblockdev-utils NetworkManager-team NetworkManager-bluetooth NetworkManager-wifi libstorage-uis-static kiwi-dlimage

Table 4-9 Conflicting packages in EulerOS

EulerOS Version	Conflicting Packages
EulerOS 2.9	euleros-release;euleros-latest-release;kiwi-systemdeps;python3-kiwi NetworkManager-team NetworkManager-bluetooth NetworkManager-wifi libstorage-uio-static kiwi-dlimage systemd-udev-compat
EulerOS 2.10	euleros-release;euleros-latest-release;kiwi-systemdeps;python3-kiwi NetworkManager-team NetworkManager-bluetooth NetworkManager-wifi libstorage-uio-static kiwi-dlimage systemd-udev-compat

4.3 Migrating an OS to Huawei Cloud EulerOS 1.1

4.3.1 Constraints

- Only a CentOS 7.9 OS with no GUI installed can be migrated to Huawei Cloud EulerOS 1.1.
- During the OS migration, RPM packages need to be uninstalled, installed, and updated. As a result, the OS may restart unexpectedly. Back up the system disk before the OS migration. For details, see [Creating a Cloud Server Backup](#).
- There should be at least 128 MB of available space in the memory and 1 GB on the system disk.

4.3.2 Migration Operations

Migrate CentOS 7.9 to Huawei Cloud EulerOS 1.1.

Preparing Dependent Packages

1. Remotely connect to the source OS.
Remotely log in to the ECS to be migrated and check that the ECS can be accessed from the Internet. For details, see [Login Overview](#).
2. Disable all repository configurations in `/etc/yum.repos.d` of CentOS. This ensures that the repositories of CentOS and Huawei Cloud EulerOS do not conflict.

```
root@ecs-2d53-g1 yum.repos.d]# ls
CentOS-Base.repo  CentOS-Debuginfo.repo  CentOS-Media.repo  CentOS-Vault.repo  epel.repo  epel-testing.repo
CentOS-CR.repo   CentOS-fasttrack.repo  CentOS-Sources.repo  CentOS-x86_64-kernel.repo  epel.repo.rpmnew
root@ecs-2d53-g1 yum.repos.d]#
```

Take `Centos_Base.repo` as an example. Add `enabled=0` under each item, as shown in the following figure.

```
[base]
name=CentOS-$releasever - Base
#mirrorlist=http://mirrorlist.centos.org/?release=$releasever&arch=$basearch&repo=os&infra=$infra
baseurl=https://repo.huaweicloud.com/centos/$releasever/os/$basearch/
gpgcheck=1
enabled=0
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-7

#released updates
[updates]
name=CentOS-$releasever - Updates
#mirrorlist=http://mirrorlist.centos.org/?release=$releasever&arch=$basearch&repo=updates&infra=$infra
baseurl=https://repo.huaweicloud.com/centos/$releasever/updates/$basearch/
gpgcheck=1
enabled=0
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-7

#additional packages that may be useful
[extras]
name=CentOS-$releasever - Extras
#mirrorlist=http://mirrorlist.centos.org/?release=$releasever&arch=$basearch&repo=extras&infra=$infra
baseurl=https://repo.huaweicloud.com/centos/$releasever/extras/$basearch/
gpgcheck=1
enabled=0
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-CentOS-7

#additional packages that extend functionality of existing packages
[centosplus]
```

3. Configure the repository of Huawei Cloud EulerOS.

Add the following content to **hce.repo** and then store it in the **/etc/yum.repos.d/** directory:

```
[centos7_everything]
name=centos7_everything
baseurl=https://repo.huaweicloud.com/hce/1.1/os/x86_64/
enable=1
gpgcheck=0
priority=1

#released updates
[updates]
name=hce1_updates
baseurl=https://repo.huaweicloud.com/hce/1.1/updates/x86_64/
gpgcheck=0
enabled=1
gpgkey=
```

4. Check whether CentOS 7.9 can access the repository of Huawei Cloud EulerOS.

Run the **curl https://repo.huaweicloud.com/hce/1.1/os/x86_64/** command to check whether the repository of Huawei Cloud EulerOS can be accessed. If information similar to the following is displayed, the repository can be accessed:

```
% Total % Received % Xferd Average Speed Time Time Time Current
Dload Upload Total Spent Left Speed
100 3417 0 3417 0 0 373 0 --:--:-- 0:00:09 --:--:-- 696
<!doctype html>
<html>
<head>
<meta charset="utf-8">
<title></title>
<link rel="stylesheet" href="/repository/static/css/style.css" type="text/css"/>
<style>
* {
font-family: 'Verdana', sans-serif;
margin: 0;
padding: 0;
-webkit-box-sizing: border-box;
-moz-box-sizing: border-box;
box-sizing: border-box;
}
.....
```

5. Install Python 3.

```
[root@localhost ~]# yum install -y python 3 //Run this command in any directory you want.
```

NOTE

If Python 3 has been installed on CentOS 7.9, skip this step.

6. Disable SELinux.

To ensure that system configuration files are consistent before and after the migration, SELinux needs to be disabled.

- a. Modify the `/etc/selinux/config` file by changing the value of **SELINUX** to **disabled**.

```
SELINUX=disabled
```

- b. Restart the OS to apply the changes.

Installing the Migration Tool

1. Download tool package **centos2hce1-*.rpm** from the Huawei Cloud open-source image site. Contact customer service to obtain the download link from O&M engineers.

The asterisk (*) indicates the version of the migration tool. In this example, **centos2hce1-1.0.0-0.0.2.x86_64.rpm** is used.

```
[root@localhost test]# wget https://repo.huaweicloud.com/hce/1.1/updates/x86_64/Packages/centos2hce1-1.0.0-0.0.2.x86_64.rpm //Download the centos2hce1-*.rpm package.
```

```
[root@localhost test]# ls //Check whether the download is successful.
centos2hce1-1.0.0-0.0.2.x86_64.rpm
```

2. Install the migration tool.

After the tool has been installed, the system automatically generates the `/etc/centos2hce1.conf` file.

```
[root@localhost ~]# rpm -ivh centos2hce1-1.0.0-0.0.2.x86_64.rpm
```

3. Configure the `centos2hce1.conf` file.

Configure the repository of Huawei Cloud EulerOS. It will be used for checking whether the repository can be accessed and updating RPM packages.

```
#iso as yum source link
[repo_info]
base_yum_url=https://repo.huaweicloud.com/hce/1.1/os/x86_64/

#iso as yum source
repositr_hce1_1 =
[base]
name=hceversion
baseurl=https://repo.huaweicloud.com/hce/1.1/os/x86_64/
gpgcheck=0
enabled=1
#released updates
[updates]
name=hce1_updates
baseurl=https://repo.huaweicloud.com/hce/1.1/updates/x86_64/
gpgcheck=0
enabled=1
gpgkey=
```

NOTE

To learn more about the parameters in the `centos2hce1.conf` file, see [Appendix: Description of the .conf File](#).

Migrating the OS

1. Back up the OS.

The migration to Huawei Cloud EulerOS 1.1 cannot be rolled back. Before performing the migration, you should back up CentOS including its system disk and data disks.

2. Run the **centos2hce1.py** command to migrate the OS.

The migration takes 20 minutes to 1 hour, depending on the number and size of RPM packages to be updated and the download speed of RPM packages from the repository. Reserve sufficient time for the migration based on your environment.

```
[root@localhost home]# centos2hce1.py
```

If the following information is displayed, the migration was complete. If the migration failed, use the backup to restore data.

```
Complete!
2023-04-18 15:57:22-centos2hce1.py [327]: redhat-lsb is replaced by system-lsb on hce1.1,centos-logs is replaced by hce-logs
"ggp-pubkey-f4a88eb5-53a7ff4b" gpg-pubkey-352c64e5-52ae6884package %(PACKAGE)s" is not installed
package ! is not installed
"grep-2.20-3.hce1cpackage CentOS is not installed
2023-04-18 15:57:27-centos2hce1.py [347]: remove left rpms
2023-04-18 15:57:27-centos2hce1.py [350]: remove rpm list: kernel-3.10.0-1160.el7.x86_64 kernel-devel-3.10.0-1160.el7.x86_64 ift
op-1.0-0.21.pre4.el7.x86_64 kernel-devel-3.10.0-1160.53.1.el7.x86_64 centos-release-7-9.2009.0.el7.centos.x86_64 kernel-3.10.0-1
160.53.1.el7.x86_64
2023-04-18 15:57:34-centos2hce1.py [357]: Removing yum cache
2023-04-18 15:57:41-centos2hce1.py [380]: Sync successfully, update grub.cfg.
[root@centos? home]#
```

NOTE

CentOS contains some RPM packages that are not provided by Huawei Cloud EulerOS 1.1. After you run the **centos2hce1.py** command to migrate the OS, these RPM packages are automatically deleted. If you want to retain them, run the **-s skip** command to migrate the OS.

3. (Optional) Delete unnecessary RPM packages.

The following two RPM packages are not used during the migration and do not affect how the system runs. You can just delete them.

```
[root@localhost home]# ll
total 24
-rw-r--r--. 1 root root 15972 Jul  1 06:31 hce-release-1.1-23.hce1c.x86_64.rpm
-rw-r--r--. 1 root root 5032 Jul  1 06:31 hce-repos-2.10-2.hce1c.x86_64.rpm
[root@localhost home]#
```

```
[root@localhost home]# rm hce-release-1.1-23.hce1c.x86_64.rpm hce-repos-2.10-2.hce1c.x86_64.rpm
rm: remove regular file 'hce-release-1.1-23.hce1c.x86_64.rpm'? y
rm: remove regular file 'hce-repos-2.10-2.hce1c.x86_64.rpm'? y
[root@localhost home]#
```

4. Run the **reboot** command to restart the OS.

5. Run the **cat /etc/os-release** command to check whether the migration was successful.

If the following Huawei Cloud EulerOS information is displayed, the migration was successful.

```
[root@localhost centos2hce1]# cat /etc/os-release
NAME="Huawei Cloud EulerOS"
VERSION="1.1 (x86_64)"
ID="hce"
VERSION_ID="1.1"
PRETTY_NAME="Huawei Cloud EulerOS 1.1 (x86_64)"
ANSI_COLOR="0;31"
[root@localhost centos2hce1]#
```

6. (Optional) Enable SELinux.
SELinux was disabled before the OS migration. Enable it if needed.
 - a. Modify the `/etc/selinux/config` file by changing the value of **SELINUX** to **enforcing**.
SELINUX=enforcing
 - b. Restart the OS to apply the changes.

4.3.3 Appendix: Description of the .conf File

This section describes the important fields in the `.conf` file.

```
#rpm lists for os migration
[rpm_lists]
#origin system must need rpms
baserpms_list = "basesystem initscripts hce-logos plymouth grub2 grubby" //The RPM packages
required for the OS migration.

#old rpm and default conflict rpms //The conflicting RPM packages that may exist in the source OS during
the migration.
oldrpms_list = centos-backgrounds centos-release-cr desktop-backgrounds-basic \
centos-release-advanced-virtualization centos-release-ansible26 centos-release-ansible-27 \
centos-release-ansible-28 centos-release-ansible-29 centos-release-azure \
centos-release-ceph-jewel centos-release-ceph-luminous centos-release-ceph-nautilus \
centos-release-ceph-octopus centos-release-configmanagement centos-release-dotnet centos-release-fdio \
centos-release-gluster40 centos-release-gluster41 centos-release-gluster5 \
centos-release-gluster6 centos-release-gluster7 centos-release-gluster8 \
centos-release-gluster-legacy centos-release-messaging centos-release-nfs-ganesha28 \
centos-release-nfs-ganesha30 centos-release-nfv-common \
centos-release-nfv-openvswitch centos-release-openshift-origin centos-release-openstack-queens \
centos-release-openstack-rocky centos-release-openstack-stein centos-release-openstack-train \
centos-release-openstack-ussuri centos-release-opstools centos-release-ovirt42 centos-release-ovirt43 \
centos-release-ovirt44 centos-release-paas-common centos-release-qemu-ev centos-release-qpip-proton \
centos-release-rabbitmq-38 centos-release-samba411 centos-release-samba412 \
centos-release-scl centos-release-scl-rh centos-release-storage-common \
centos-release-virt-common centos-release-xen centos-release-xen-410 \
centos-release-xen-412 centos-release-xen-46 centos-release-xen-48 centos-release-xen-common \
python3-syspurpose python-oauth sl-logos yum-rhn-plugin centos-indexhtml \
libreport-centos libreport-web libreport-plugin-mantisbt libreport-plugin-rhtsupport \
libreport hunspell-en-US hunspell-en policycoreutils-gui libcanberra-gtk2 cups \
NetworkManager-libreswan-gnome plymouth-graphics-libs avahi cups-lpd pinentry-qt \
libsvg2-devel libcanberra-gtk3 gnome-themes-standard wodim gsettings-desktop-schemas-devel \
avahi-ui-gtk3 freerdp-libs pulseaudio-utils gstreamer1-plugins-bad-free-gtk ghostscript-cups \
setools-console libxkbcommon-x11 cups plymouth-plugin-two-step pulseaudio-module-x11 ImageMagick-c+
+ \
cups-devel policycoreutils-sandbox PackageKit-gstreamer-plugin gtk3-imodule-xim avahi-glib avahi-
autoipd \
mesa-libGLES foomatic libcanberra-devel plymouth-plugin-label PackageKit-gtk3-module colorl avahi-
gobject \
pinentry-qt4 avahi-ui-gtk3 plymouth-plugin-two-step ghostscript-cups ImageMagick-perl firewall-config \
plymouth-plugin-label redhat-redhat-lsb-corelib vim-X11 dbus-x11 pulseaudio PackageKit-command-not-
found libproxy-mozjs \
pinentry-gtk nm-connection-editor gtk2-imodule-xim wireshark-gnome pulseaudio-module-bluetooth
pidgin-sipe freerdp kmod-kvdo \
redhat-lsb-core

#The following list contains the same symbol as centos/redhat
dstrpms_list = "hce-release hce-repos"

[log_conf]
# migration tool log common dir
migrate_common_dir = "/var/log/migrate-tool/" //The path for storing logs.

# migration tool classification log dir
migrate_classification_dir = %(migrate_common_dir)s/centos2hce1/
```

```
#iso as yum source link
[repo_info]
base_yum_url=https://repo.huaweicloud.com/hce/1.1/os/x86_64/ //The base yum URL used for checking
the network connection.

#iso as yum source
repostr_hce1_1 = //The source path that provides the migration method.
[base]
name=hceversion
baseurl=https://repo.huaweicloud.com/hce/1.1/os/x86_64/ //The base yum URL used for obtaining the
RPM packages.
gpgcheck=0
enabled=1
gpgkey=

#released updates
[updates]
name=hce1_updates
baseurl=
gpgcheck=0
enabled=0
gpgkey=

#additional packages that may be useful
[extras]
name=hce1_extras
baseurl=
gpgcheck=0
enabled=0
gpgkey=

# plus packages provided by Huawei Linux dev team
[plus]
name=hce1_plus
baseurl=
gpgcheck=0
enabled=0
gpgkey=
```

5 Upgrading Huawei Cloud EulerOS and RPM Packages

5.1 Upgrade Overview

Huawei provides updates and maintenance for Huawei Cloud EulerOS and the RPM packages, including RPM packages deployed on Huawei Cloud EulerOS and those related to security updates for vulnerability fixing. To ensure system security, always install updates in a timely manner.

You can upgrade Huawei Cloud EulerOS using either `dnf`, `yum`, or `OSMT`:

- In Linux, you can use `dnf` or `yum` to upgrade or roll back RPM packages.
- `OSMT` is a software tool from Huawei Cloud that you can use to upgrade or roll back Huawei Cloud EulerOS and RPM packages. `OSMT` allows you to customize the upgrade scope and configure scheduled checks and delayed restarts.

Differences between the two methods are described in the following table.

Table 5-1 Differences between two methods

Item	dnf or yum	OSMT
RPM package upgrade	<ul style="list-style-type: none">• Upgrading all RPM packages, including those related to security updates for vulnerability fixing• Upgrading only RPM packages related to security updates	<ul style="list-style-type: none">• Upgrading all RPM packages, including those related to security updates for vulnerability fixing• Custom upgrades:<ul style="list-style-type: none">– Upgrading only RPM packages that do not need an OS restart– Upgrading only RPM packages that need an OS restart– Upgrading the RPM packages defined in a custom blacklist or whitelist– Upgrading RPM packages related to security updates– Fixing vulnerabilities– Upgrading the RPM packages related to new functions– Updating new RPM packages• Automatically updating RPM package and delaying restarts until specified times
OS version upgrade	Not supported	Upgrade of Huawei Cloud EulerOS 2.0 or later
Supported upgrade version	RPM packages of Huawei Cloud EulerOS 1.1 or later	RPM packages of Huawei Cloud EulerOS 2.0 or later
Rollback	Rollback to any historical update	Rollback only to the last update

5.2 Using dnf or yum to Upgrade

This section describes how to update or roll back the RPM packages of Huawei Cloud EulerOS 1.1 or later. The methods for using dnf and yum are the same. In this section, dnf is used. If you are upgrading to Huawei Cloud EulerOS 1.1, use yum.

NOTE

- Huawei Cloud EulerOS 2.0 and later support both yum and dnf.
- Huawei Cloud EulerOS 1.1 supports only yum.

Prerequisites

The dnf component has been installed and the dnf commands are available.

```
[root@localhost bin]# dnf
usage: dnf [options] COMMAND

List of Main Commands:

alias                List or create command aliases
autoremove           remove all unneeded packages that were originally installed as dependencies
check                check for problems in the packagedb
check-update         check for available package upgrades
clean                remove cached data
deplist              [deprecated, use repoquery --deplist] List package's dependencies and what packages
provide them
distro-sync          synchronize installed packages to the latest available versions
downgrade            Downgrade a package
.....
```

Context

Yum is a package management utility for Linux distributions, like CentOS. Yum allows for automatic updates and package management on RPM-based distributions. dnf is a rewrite of yum and improves on yum in several ways: improved performance, less memory usage, and better resolution of dependency conflicts. To ensure compatibility, Huawei Cloud EulerOS also supports yum commands.

Upgrade Procedure

1. Check what RPM package updates are available.

- Run **dnf list updates** to see the available updates.

```
[root@localhost bin]# dnf list updates
Last metadata expiration check: 6:49:11 ago on Tue 28 Jun 2022 01:55:35 PM CST.
hce-config.x86_64                3.0-66.hce2
hce-latest-release.x86_64       2.0-1656179342.2.0.2206.B032.hce2
irqbalance.x86_64              3:1.8.0-7.h9.hce2
kernel.x86_64                   5.10.0-60.18.0.50.h316_1.hce2
kernel-tools.x86_64            5.10.0-60.18.0.50.h316_1.hce2
kernel-tools-libs.x86_64       5.10.0-60.18.0.50.h316_1.hce2
kexec-tools.x86_64             2.0.23-4.h8.hce2
libcurl.x86_64                  7.79.1-2.h4.hce2
libssh.x86_64                   0.9.6-2.h3.hce2
libstdc++.x86_64                10.3.1-10.h10.hce2
libxml2.x86_64                  2.9.12-5.h5.hce2
openssh.x86_64                  8.8p1-2.h12.hce2
openssh-clients.x86_64         8.8p1-2.h12.hce2
openssh-server.x86_64          8.8p1-2.h12.hce2
Obsoleting Packages
dnf-data.noarch                 4.10.0-3.h6.hce2
dnf.noarch                      4.10.0-3.h5.hce2
dnf-data.noarch                 4.10.0-3.h6.hce2
dnf-data.noarch                 4.10.0-3.h5.hce2
```

- Run **dnf list updates --security** to list only the RPM packages related to security updates.

```
[root@localhost bin]# dnf list updates --security
Last metadata expiration check: 0:00:03 ago on Fri 08 Jul 2022 04:45:56 PM CST.
No security updates needed, but 2 updates available
```

2. Update the RPM packages.

- Run **dnf update** to update all the RPM packages, including those related to security updates for vulnerability fixing. The components' target versions are returned in the **Version** column.

```
[root@localhost bin]# dnf update
Last metadata expiration check: 7:12:18 ago on Tue 28 Jun 2022 01:55:35 PM CST.
Dependencies resolved.

=====
=====
Package                Arch   Version                               Repo   Size
=====
=====
Installing:
kernel                 x86_64 5.10.0-60.18.0.50.h316_1.hce2       hce2   47 M
Upgrading:
hce-config             x86_64 3.0-66.hce2                          hce2   13 k
hce-latest-release    x86_64 2.0-1656179342.2.0.2206.B032.hce2   hce2   5.2 k
kernel-tools          x86_64 5.10.0-60.18.0.50.h316_1.hce2       hce2   230 k
kernel-tools-libs     x86_64 5.10.0-60.18.0.50.h316_1.hce2       hce2   62 k
kexec-tools           x86_64 2.0.23-4.h8.hce2                     hce2   400 k
libcurl                x86_64 7.79.1-2.h4.hce2                      hce2   284 k
libssh                 x86_64 0.9.6-2.h3.hce2                       hce2   194 k
libstdc++              x86_64 10.3.1-10.h10.hce2                   hce2   535 k
libxml2                x86_64 2.9.12-5.h5.hce2                     hce2   659 k
logrotate              x86_64 3.18.1-1.h2.hce2                     hce2   60 k
mdadm                  x86_64 4.1-5.h2.hce2                         hce2   331 k
nftables               x86_64 1:1.0.0-1.h3.hce2                     hce2   303 k
perl                   x86_64 4:5.34.0-3.h5.hce2                   hce2   3.2 M
perl-libs              x86_64 4:5.34.0-3.h5.hce2                   hce2   1.8 M
Installing dependencies:
grub2-tools-efi        x86_64 1:2.06-3.h5.hce2                      hce2   472 k

Transaction Summary
=====
=====
Install 2 Packages
Upgrade 72 Packages
Total download size: 105 M
Is this ok [y/N]:
```

- Run **dnf update --security** to update only the RPM packages related to security updates.

```
[root@localhost bin]# dnf update --security
Last metadata expiration check: 7:15:16 ago on Tue 28 Jun 2022 01:55:35 PM CST.
No security updates needed, but 73 updates available Dependencies resolved.
Nothing to do.
Complete!
```

3. After the update is successful, check that services are running properly.

The issue that may occur during an update:

Based on the requirements of security regulations, the **chronyd** service will be disabled after Huawei Cloud EulerOS is upgraded from 2.0.2206 to a new version. If required, run **systemctl enable chronyd** to enable the service and run **systemctl start chronyd** to start the service.

Rollback Procedure

1. Run **dnf history** to query the IDs of historical operations.

```
[root@localhost ~]# dnf history
ID      | Command line          | Date and time   | Action(s)      | Altered
-----|-----|-----|-----|-----
5      | upgrade chrony        | 2022-10-09 11:38 | Upgrade        | 1
4      | history undo 3        | 2022-10-09 11:37 | Downgrade     | 1
3      | upgrade chrony        | 2022-10-09 11:36 | Upgrade        | 1
2      | install createrepo    | 2022-10-09 11:30 | Install        | 2
1      |                       | 2022-10-09 11:15 | Install        | 421 EE
[root@localhost ~]#
```

2. Run **dnf history undo <ID>** to roll back to the desired historical operation.

5.3 Upgrade Using OSMT

5.3.1 Overview

OSMT is a tool provided by Huawei Cloud to upgrade or roll back Huawei Cloud EulerOS and RPM packages. It also allows you to customize the upgrade scope, configure scheduled checks, perform a single upgrade at the specified time period, and schedule restarts for RPM packages.

- To upgrade or roll back the Huawei Cloud EulerOS, see [Version Upgrade and Rollback](#).
- To update or roll back only RPM packages, see [Updating RPM Packages](#).

NOTE

OSMT can only upgrade or roll back Huawei Cloud EulerOS 2.0 or later.

5.3.2 Constraints

- An upgrade or a rollback takes no more than 30 minutes, depending on the number and size of RPM packages to be updated and the download speed of RPM packages from the repository. Reserve sufficient time based on your environment.
- OSMT can only be used to update the RPM packages in the **base** and **updates** official source repositories. Ensure that these two repositories are correctly configured. You must run **systemctl restart osmt-agent** to restart the osmt-agent service after modifying the source repository.
- Modifying the OSMT configuration file using a method other than the **osmt config** command may lead to abnormal OSMT functions, so you are advised to run **osmt config** to modify the file.
- Upgrades must be performed as user **root**.
- Upgrading or rolling back the OS or RPM packages has the following requirements:
 - Memory: 512 MB
 - Root partition: 1.5 GB
 - Backup storage path (**store_path**): 8 GB
 - **/boot** partition in the OS: 100 MB

NOTE

- The required storage space varies depending on the upgrade scope and target version. During the upgrade, OSMT automatically estimates the space required for the upgrade. If the available space is insufficient, an error message is displayed.
- The upgrade and rollback impacts on the SELinux status are as follows:
 - An upgrade has no impact on the SELinux status. The SELinux status before and after the upgrade is the same.
 - If the SELinux status before a rollback is **enforcing**, after the rollback, its status automatically changes to **permissive**.

- To enable SELinux, manually change the SELinux status to **enforcing** and restart the OS.
- If the SELinux status before a rollback is **disabled**, the rollback has no impact on the SELinux status. The SELinux status before and after the rollback is the same.
- To enable SELinux, set the SELinux status to **permissive**, create the **.autorelabel** file in the root directory, restart the OS, change the SELinux status to **enforcing**, and restart the OS.
- OSMT checks the OS health status before the upgrade. If the check fails, resolve the issue based on the information provided. You can also manually perform the check by referring to [OSMT Command Help Information](#).
- The upgrade using OSMT depends on the DNF tool. To ensure the stability of the upgrade, OSMT will update the DNF tool and its dependent RPM packages to the latest version. For details about how to roll back the RPM packages, see [Rollback Procedure](#).
- If the system configuration (you can run **sysctl -a** to query system configuration) is modified after the RPM packages are updated, the upgrade cannot be performed using OSMT. You can run **sysctl -p** to update the system configuration. You can run **sysctl -p <file>** to specify the configuration file that takes effect. The **sysctl --system** command can be run on the configuration files in all system directories. Before running this command, confirm the kernel configuration files in all system directories.
- The kernel or kernel hot patch cannot have more than five versions. If the kernel or kernel hot patch has more than five versions, the OSMT check will fail. If this happens, uninstall unnecessary versions and perform the check again.
- If chrony and NTP coexist and chrony is in the active state, the OSMT check will fail. If this happens, stop the chrony service or uninstall either chrony or the NTP service and perform the upgrade again.

5.3.3 Version Upgrade and Rollback

This section describes how to upgrade or roll back the Huawei Cloud EulerOS.

During an OS upgrade or rollback, RPM packages will be updated to the versions of the target OS. The blacklist and whitelist configured in **osmt.conf** will not be applied.

Upgrading the OS Version

1. Confirm that the repository is configured correctly.

Check whether the parameters in the **/etc/yum.repos.d/hce.repo** file are configured correctly. The correct configuration is as follows:

```
[base]
name=HCE $releasever base
baseurl=https://repo.huaweicloud.com/hce/$releasever/os/$basearch/
enabled=1
gpgcheck=1
gpgkey=https://repo.huaweicloud.com/hce/$releasever/os/RPM-GPG-KEY-HCE-2

[updates]
name=HCE $releasever updates
baseurl=https://repo.huaweicloud.com/hce/$releasever/updates/$basearch/
.....
```

 NOTE

Incorrect configuration may result in OSMT upgrade failures or unexpected upgrades.

2. Update the OSMT version.

There is a mapping between OSMT versions and Huawei Cloud EulerOS versions. Huawei Cloud EulerOS installs by default the OSMT of the current OS. When upgrading the Huawei Cloud EulerOS, you also need to update the OSMT to the matched version.

Run **dnf update osmt -y --releasever** [*Target OS version*] to update the OSMT version. For example, to upgrade Huawei Cloud EulerOS 2.0 to 2.1, run **dnf update osmt -y --releasever 2.1** to update OSMT to the latest version.

 NOTE

You can also run **dnf install osmt -y --releasever** [*Target OS version*] to install OSMT if it is deleted by mistake. For example, you can run **dnf install osmt -y --releasever 2.1** to install the latest OSMT.

3. Upgrade the Huawei Cloud EulerOS version.

osmt update --releasever [*Target version*] **--reboot_config** [*Restart configuration*]

Choose an appropriate upgrade method. For more upgrade options, see [osmt update -h](#).

- To upgrade Huawei Cloud EulerOS 2.0 to 2.1, run the following command:

```
osmt update --releasever 2.1
```

The upgrade is only applied after a reboot.

- Upgrade Huawei Cloud EulerOS 2.0 to 2.1 and restart the OS immediately.

```
osmt update --releasever 2.1 --reboot_config always
```

- Upgrade Huawei Cloud EulerOS 2.0 to 2.1 and restart the OS at 2022-12-30 23:00:00.

```
osmt update --releasever 2.1 --reboot_config "2022-12-30 23:00:00"
```

4. Check whether the upgrade was successful.

Run **cat /etc/hce-latest** and view the **hceversion** field. If the **--releasever** value is the version you specified, the upgrade was successful.

5. (Optional) Delete backup files.

After verifying the OS functions, run **osmt remove** to delete the backup files.

 NOTE

The operation of deleting backup files cannot be undone. Ensure that no exception occurs after the upgrade before you run **osmt remove**.

Rolling Back the OS Version

1. Choose an appropriate rollback method.

- To roll back and not restart the OS, run the following command:

```
osmt rollback
```

- To roll back and restart the OS immediately, run the following command: (Then skip step 2.)

osmt rollback --reboot_config always

2. Run **reboot** to restart the OS.
The rollback is only applied after a restart.
3. Check whether the rollback was successful.
Run **cat /etc/hce-latest** and view the **hceversion** field. If the **hceversion** value is the source version, the rollback was successful.

5.3.4 Updating RPM Packages

5.3.4.1 Preparations

RPM packages can be updated manually (using **osmt update**) or automatically (using the background osmt-agent service). You need to perform the following operations for both manual and automatic updates.

1. Confirm that the repository is configured correctly.

Check whether the parameters in the **/etc/yum.repos.d/hce.repo** file are configured correctly. The correct configuration is as follows:

```
[base]
name=HCE $releasever base
baseurl=https://repo.huaweicloud.com/hce/$releasever/os/$basearch/
enabled=1
gpgcheck=1
gpgkey=https://repo.huaweicloud.com/hce/$releasever/os/RPM-GPG-KEY-HCE-2

[updates]
name=HCE $releasever updates
baseurl=https://repo.huaweicloud.com/hce/$releasever/updates/$basearch/
.....
```

 NOTE

Incorrect configuration may result in OSMT upgrade failures or unexpected upgrades.

2. Run **dnf update osmt -y** to update OSMT.
3. Configure the **/etc/osmt/osmt.conf** file.

OSMT updates RPM packages based on the **osmt.conf** file. Configure the file as required.

```
[auto]
#if auto_upgrade is True, the osmt-agent will auto upgrade rpms use osmt.conf and reboot between
time interval we specified
#the value of cycle_time means the osmt-agent will check upgrade every cycle_time seconds, default
86400s(1 day)
#When a configuration item has a line break, you need to leave a space or tab at the beginning of
the line
auto_upgrade = False
cycle_time = 3600
minimal_interval = 3600
auto_upgrade_window = "22:00-05:00"
auto_upgrade_interval = 1

[Package]
There are three rules of filters, all enabled by default. Severity will be effect only when the types
contain security, it is the subtype of security.
# The following are the three rules:
# 1. white list has the highest priority, if whitelist is configured then ignore other rules and filter out
the whitelist packages from the full list of packages to be upgrade
# 2. Filter the update range by types, when the types contain security, further filter the severity of
security updates severity, only upgrade the severity level of security.
```

```
# 3. Filter blacklist to remove packages in blacklist from types filter results, and packages which
depend on packages in blacklist will also be removed.
# filters must contain at least one types rule, if the types rule is empty, the -a option will not upgrade
any packages (by default all 3 filters are enabled).

filters = "types, blacklist"
whitelist = ""

# types include: security, bugfix, enhancement, newpackage, unknown
# if types is empty, no package will be upgrade
# types = security, bugfix, enhancement, newpackage, unknown
types = "security"
# severity is the subtype of security, include: low, moderate, important, critical
severity = "important, critical"
blacklist = ""
# The rpm package that requires a system reboot to take effect after the upgrade
need_reboot_rpms = "kernel, kernel-debug, kernel-debuginfo, kernel-debuginfo-common, kernel-
devel, kernel-headers, kernel-ori, kernel-tools, kernel-tools-libs, glibc, glibc-utils, glibc-static, glibc-
headers, glibc-devel, glibc-common, dbus, dbus-python, dbus-libs, dbus-glib-devel, dbus-glib, dbus-
devel, systemd, systemd-devel, systemd-libs, systemd-python, systemd-sysv, grub2, grub2-efi, grub2-
tools, openssl, openssl-devel, openssl-libs, gnutls, gnutls-dane, gnutls-devel, gnutls-utils, linux-
firmware, openssh, openssh-server, openssh-clients, openssh-keycat, openssh-askpass, python-
libs, python, grub2-pc, grub2-common, grub2-tools-minimal, grub2-pc-modules, grub2-tools-extra, grub2-
efi-x64, grub2-efi-x64-cdboot, kernel-cross-headers, kernel-source, glibc-all-langpacks, dbus-
common, dbus-daemon, dbus-tools, systemd-container, systemd-pam, systemd-udev, grub2-efi-
aa64, grub2-efi-aa64-cdboot, grub2-efi-aa64-modules, openssl-perl, openssl-pkcs, kernel-tools-libs-
devel, glibc-debugutils, glibc-locale-source, systemd-help, grub2-efi-ia32-modules, grub2-efi-x64-
modules, grub2-tools-efi, grub2-help, openssl-pkcs11, grub2-efi-ia32-cdboot, osmt"
preinstalled_only = False

[backup]
store_path = /var/log
backup_dir = /etc,/usr,/boot,/var,/run
exclude_dir =
recover_service =

#the minimum resources required(MB)
[resource_needed]
#min_req_boot_space = 100
#min_req_backup_space = 8192
#min_req_root_space = 1536
#min_req_memory = 512

[cmdline]
cmdline_value =

[conflict]
#conflict_rpm = test1,test2

[strategy]
timeout_action = "stop"
timeout_action_before = 0

[check]
daemon_whitelist = "sysstat-collect.service, sysstat-summary.service, systemd-tmpfiles-clean.service"
```

Table 5-2 Major configuration items in **osmt.conf**

Configuration Item	Description
[auto]	<ul style="list-style-type: none">• auto_upgrade: specifies the RPM package update method. The default value is False.<ul style="list-style-type: none">- True: RPM packages can be updated either manually or automatically.- False: RPM packages can only be updated manually.• If auto_upgrade is set to True, the following parameters are available:<ul style="list-style-type: none">- cycle_time: defines the interval of checking for available updates, in seconds. The default value is 3600.- minimal_interval: defines the minimum interval (in seconds) between the start time and end time specified using osmt update -b. The default value is 3600.- auto_upgrade_window: defines the start time and end time of automatic updates using the osmt-agent service. The value is in the format of "HH:MM-HH:MM". If the end time is smaller than the start time, the update period covers two dates. For example, 22:00-05:00 indicates an update period from 22:00 on the current day to 05:00 on the next day.- auto_upgrade_interval: defines the minimum interval between two automatic updates, in days.• If auto_upgrade is set to False, only the following parameters are available, and any other [auto] parameter configured will not take effect.<ul style="list-style-type: none">- cycle_time: defines the interval of checking for available updates, in seconds. The default value is 3600.- minimal_interval: defines the minimum interval (in seconds) between the start time and end time specified using osmt update -b. The default value is 3600.• motd_setup: specifies whether to enable the login prompt. The default value is True.<ul style="list-style-type: none">- True: Enable the login prompt.- False: Disable the login prompt. After the setting, the login prompt is deleted immediately and will not be generated again. If the option is enabled again, you need to run the osmt update -s command or any upgrade command to trigger the generation again.

Configuration Item	Description
[Package]	<ul style="list-style-type: none">• filters: specifies the update scope. The value can be types, blacklist, or whitelist. For example, specifying blacklist will not update the RPM packages in the blacklist.<ul style="list-style-type: none">– types: the type of RPM packages to be updated.– blacklist: the RPM package blacklist. Packages added to the blacklist will not be updated. If an RPM package depends on a package in the blacklist, this RPM package will also not be updated.– whitelist: the RPM package whitelist. If the whitelist and blacklist are not configured, all RPM packages will be updated. The whitelist has a higher priority than the blacklist, meaning that if an RPM package is added to both lists, it will be updated.NOTE The blacklist and whitelist specified in the command are applied, and those defined in the configuration file will not.• need_reboot_rpms: lists the RPM packages that need an OS restart. Automatic updates using osmt-agent will not update packages in need_reboot_rpms. To update these packages during automatic updates, you must run osmt update --auto --reboot_config always or osmt update --auto --reboot_config "Restart time".• preinstalled_only: If this parameter is set to True, only the RPM packages in /etc/osmt/preinstalled.list are required to be upgraded.
[backup]	<ul style="list-style-type: none">• store_path: the directory under which the backup directory is created. During the upgrade, OSMT creates the .osbak directory under store_path. If .osbak already exists, run the osmt remove command to delete it first.• backup_dir: the directories that need to be backed up. Directories /etc, /usr, /boot, /var, and /run are backed up by default during an update and cannot be removed.• recover_service: OSMT checks whether the status of each service in the list is the same before and after the upgrade. If the status of a service is changed, OSMT will restore the service status. NOTE The path in [backup] must be an absolute path.
[cmdline]	cmdline_value : determines the startup items after an upgrade. Configure correct startup items to ensure that the OS can be started properly. By default, the default startup items of Huawei Cloud EulerOS are used.

Configuration Item	Description
[conflict]	conflict_rpm : specifies the RPM packages that will be deleted if there is a conflict during the upgrade.

 NOTE

You are advised not to modify other configuration items. For details, see [Description of the /etc/osmt/osmt.conf File](#).

5.3.4.2 Manual Update Using osmt update

You can manually update RPM packages in the following ways:

- Update the RPM packages using the **filters** field in the configuration file.
osmt update --auto --reboot_config [*Restart configuration*]

Table 5-3 Restart parameter description

Value	Description
never	Does not restart the OS after the update. If the reboot parameter is not configured or its value is set to never , the OS will also not restart after the update. This way, RPM packages in the need_reboot_rpms list will not be updated. To update them, run the following commands to set filters to whitelist and add the packages to the whitelist: (The update is only applied only after a restart.) osmt config -k filters -v "whitelist" osmt config -k whitelist -v "rpm1, rpm2, rpm3"
always	Updates the RPM packages (including the packages in need_reboot_rpms) and restart the OS immediately after the update.
<specific time>	Updates the RPM packages (including the packages in need_reboot_rpms) and restart the OS at the specified time, for example "2020-02-02 2:02:02".

- Update the RPM packages using the whitelist and blacklist.
osmt update --pkgs [*rpm1 rpm2 rpm3 ...*] **--exclude_pkgs** [*rpm4 rpm5 rpm6 ...*] **--reboot_config** [*Restart configuration*]
- pkgs**: (optional) specifies the whitelisted packages to be updated. Multiple packages are separated by spaces.
For example, run the following command to update the **hce-logos**, **hce-lsb**, and **tomcat** whitelisted packages:
osmt update --pkgs hce-logos hce-lsb tomcat

- **--exclude_pkgs**: (optional) specifies the blacklisted packages that will not be updated. Multiple packages are separated by spaces.

For example, run the following command to not update the **ongres-scram** and **llvm-static** blacklisted packages:

```
osmt update --exclude_pkgs ongres-scram llvm-static
```

- **--reboot_config** [*Restart configuration*]: (optional) configures the restart method. The value can be **always**, **never**, or a specific restart time.
 - **always**: restarts the OS after the update if some package updates are only applied after a restart. If there are no such packages, the OS will not restart.
 - **never**: does not restart the OS after the update.
 - *<specific time>*: specifies a specific restart time. If some package updates are only applied after a restart, the OS will restart at the specified time after the update. The restart time is in the format of "2020-02-02 2:02:02". If there are no such packages, the OS will not restart.

NOTE

- If you update packages using the blacklist or whitelist, specify at least one of **--pkgs** and **--exclude_pkgs**.
- The blacklist and whitelist specified in the command are applied, and those defined in the configuration file will not.

5.3.4.3 Automatic Update Using osmt-agent

The osmt-agent service periodically checks whether there are available RPM package updates and updates them automatically. You can configure how often to check updates and when to perform the update.

1. Run the following command to ensure that the value of **auto_upgrade** in the **osmt.conf** file is **True**:

```
osmt config -k auto_upgrade -v True
```

2. Run **systemctl status osmt-agent.service** to check whether the osmt-agent service is started.

- If the **Active** value is **active (running)**, osmt-agent is started.
- Otherwise, run **systemctl start osmt-agent.service** to start osmt-agent.

```
● osmt-agent.service - osmt-agent - The agent that manages HCE OS.  
Loaded: loaded (/usr/lib/systemd/system/osmt-agent.service; enabled; vendor preset: disabled)  
Active: active (running) since Sat 2022-12-24 18:32:42 CST; 2 days ago  
Main PID: 1421 (python)  
Tasks: 3 (limit: 21113)  
Memory: 72.7M  
CGroup: /system.slice/osmt-agent.service  
├─ 1421 python /usr/bin/osmt_server  
├─ 1474 /bin/bash /usr/bin/osmt-agent  
└─ 32445 sleep 3600
```

3. Run the desired command to configure when or how often to perform the updates:

- To configure a time window for automatic updates:

```
osmt config -k auto_upgrade_window -v "auto_upgrade_window"
```


auto_upgrade_window: defines the start time and end time of automatic updates using the osmt-agent service. The value is in the format of "HH:MM-HH:MM".

If the end time is smaller than the start time, the update period covers two dates. For example, **22:00-05:00** indicates an update period from 22:00 on the current day to 05:00 on the next day.

For example, run the following command to configure a time window starting from 23:00 on the current day to 01:00 on the next day:

```
osmt config -k auto_upgrade_window -v "23:00-01:00"
```

- To configure the interval between two automatic updates:

```
osmt config -k auto_upgrade_interval -v auto_upgrade_interval
```

auto_upgrade_interval: defines the minimum interval between two automatic updates, in days.

For example, run the following command to configure automatic updates every other day:

```
osmt config -k auto_upgrade_interval -v 1
```

5.3.5 Follow-up Operations

1. After the update is successful, check that services are running properly. Then run **osmt remove** to delete the backup files when appropriate. Once deleted, the update cannot be rolled back.
2. Based on the requirements of security regulations, the chronyd service will be disabled after Huawei Cloud EulerOS is upgraded from 2.0.2206 to a new version. If required, run **systemctl enable chronyd** to enable the service and run **systemctl start chronyd** to start the service.

5.3.6 Rolling Back RPM Packages

Run **osmt rollback --reboot_config always** to roll back RPM packages. They can only be rolled backup to the last update.

In the command, **--reboot_config always** is optional. You must use it if there are RPM packages in **need_reboot_rpms** updated in the last update.

If **--reboot_config always** is not specified, you need to manually restart the OS so that the rollback of packages in **need_reboot_rpms** can be applied.

NOTE

Use the latest OSMT. You are advised not to roll back OSMT using OSMT.

5.4 Appendixes

5.4.1 OSMT Command Help Information

- Run **osmt -h** to display the OSMT help information.

```
[root@localhost SOURCES]# osmt -h
usage: osmt [-h] {update,rollback} ...
positional arguments:
  {update,rollback}
```

```

update      update os version or packages
rollback    rollback last update
remove      remove backup files in store path
config      modify config file by command line
job         handle upgrade task.
optional arguments:
-h, --help  show this help message and exit
    
```

Table 5-4 OSMT help parameters

Parameter	Description
update	Upgrades the OS or update the RPM packages.
rollback	Rolls back the OS or RPM packages.
remove	Deletes the backup files from the storage path.
config	Queries or modifies the configuration file.
job	Queries or manages the upgrade tasks.
-h, --help	(Optional) Provides help information of the osmt command.

- Run **osmt update -h** to display the help information about the OS or RPM package updates.

```

[root@localhost SOURCES]# osmt update -h
usage: osmt update [-h] [--nosignature] [-s] [--all] [--security] [--version] [-a] [-p PKGS [PKGS ...]] [-e EXCLUDE_PKGS [EXCLUDE_PKGS ...]] [-v RELEASEVER] [-r REBOOT_CONFIG] [-b BETWEEN] [-j] [-c]
optional arguments:
-h, --help            show this help message and exit
--nosignature         ignore the signature of packages
-s, --show            show updateinfo
--all                show all pkgs which can update, 'osmt update --show --all'
--security            show security pkgs which can update
--version            show all version can update to
-a, --auto            auto update use config file
-p PKGS [PKGS ...], --pkgs PKGS [PKGS ...]
                    specify the packages to upgrade
-e EXCLUDE_PKGS [EXCLUDE_PKGS ...], --exclude_pkgs EXCLUDE_PKGS [EXCLUDE_PKGS ...]
                    specify the packages not to upgrade
-v RELEASEVER, --releasever RELEASEVER
                    specify the release version to upgrade
-r REBOOT_CONFIG, --reboot_config REBOOT_CONFIG
                    you can choose between always, never or a specific time. 'always': reboot os after
                    update ends if need. 'never': never reboot os automatically. '<specific time>':
                    reboot at specified time, format like "2020-02-02 2:02:02".
-b BETWEEN, --between BETWEEN
                    start upgrade time and end upgrade time, format like: '2020-02-02
                    2:02:02','2020-02-02 4:02:02'
-j, --job            run upgrade in background.
-c, --check          check upgrade task.
-V, --verbose        show more log to screen
-o, --preinstalled-only
                    upgrade preinstalled packages only
-t, --retry          retry previous upgrade action
--nocheck            do not check before upgrade
    
```

Table 5-5 OSMT update parameters

Parameter	Description
-h, --help	Provides help information about the osmt update command.
--nosignature	Specifies not to filter the RPM packages to be updated by package signature.
-s,--show	Displays available upgrade or update information. <ul style="list-style-type: none">• --all: displays all RPM packages to be updated.• --security: displays the security packages to be updated.• --version: displays the version to be upgraded to.
-a,--auto	Specifies the RPM package update method. This parameter is mutually exclusive with -v , -p , and -e .
-p,--pkgs	Specifies the whitelisted RPM packages to be updated. This parameter is mutually exclusive with -v and -a .
-e,--exclude_pkgs	Specifies the blacklisted RPM packages that will not be updated. This parameter is mutually exclusive with -v and -a .
-v , --releasever	Specifies the target Huawei Cloud EulerOS version.
-r, --reboot_config	Specifies the restart configuration. <ul style="list-style-type: none">• always: restarts the OS after the update if some package updates are only applied after a restart. If there are no such packages, the OS will not restart.• never: does not restart the OS after the update.• <specific time>: specifies a specific restart time. If some package updates are only applied after a restart, the OS will restart at the specified time after the update. The restart time is in the format of "2020-02-02 2:02:02". If there are no such packages, the OS will not restart.
-b, --between	Specifies the start time and end time of automatic updates using osmt-agent. The value is in the format of "HH:MM-HH:MM". If the end time is smaller than the start time, the update period covers two dates. For example, 22:00-05:00 indicates an update period from 22:00 on the current day to 05:00 on the next day.
-j, --job	Performs the upgrade using background processes.

Parameter	Description
-c, --check	Checks the system status before the upgrade. This parameter is optional. The system also performs the check before executing the upgrade. You are advised to add -c to the update command to perform the check. For example, run osmt update -a -c to perform the check before running the osmt update -a update command.
-V, --verbose	(Optional) Displays detailed upgrade information.
-o, --preinstalled-only	(Optional) Only the RPM package list in /etc/osmt/preinstalled.list is upgraded. This parameter is valid only for version upgrade.
-t, --retry	(Optional) Retries the upgrade.
--nocheck	(Optional) No checks will be performed before the upgrade. The upgrade starts directly.

- Run **osmt rollback -h** to display the help information about the OS or RPM package rollbacks.

```
usage: osmt rollback [-h] [-r {never,always}]
optional arguments:
-h, --help            show this help message and exit
-r {never,always}, --reboot_config {never,always}
                    whether to reboot after rollback
-V, --verbose         show more log to screen
-t, --retry           retry previous upgrade action
--nocheck            do not check before rollback
```

Table 5-6 OSMT rollback parameters

Parameter	Description
-h, --help	Provides help information about the osmt rollback command.
-r, --reboot_config	Specifies the restart configuration.
-V, --verbose	(Optional) Displays detailed process logs.
-t, --retry	(Optional) Retries the rollback.
--nocheck	(Optional) No checks will be performed before the rollback. The rollback starts directly.

- Run **osmt config -h** to display the help information about configuration item modification or query.

```
usage: osmt config [-h] [-k {minimal_interval,cycle_time...}] [-v VALUE]
optional arguments:
-h, --help            show this help message and exit
-k {minimal_interval,cycle_time...} --key {minimal_interval,cycle_time...}
-v VALUE, --value VALUE
-V, --verbose         show more log to screen
```

Table 5-7 OSMT configuration parameters

Parameter	Description
-h, --help	Provides help information about the osmt config command.
-k, --key	Specifies the keys to be queried or modified.
-v, --value	Specifies the values of the keys to be modified.
-V, --verbose	(Optional) Displays detailed process logs.

NOTE

You are advised to run **osmt config** to modify the configuration file. Any modification made to the file using a method other than **osmt config** may lead to abnormal OSMT functions.

- Run **osmt job -h** to display the help information about task management.

usage: osmt job [-h] [-s] [-c] [-d DELAY] [-y]

optional arguments:

```
-h, --help      show this help message and exit
-s, --show      show task info.
-c, --cancel    cancel current task.
-d DELAY, --delay DELAY
                delay task
-y, --yes       never ask for yes.
-V, --verbose   show more log to screen
```

Table 5-8 OSMT task management parameters

Parameter	Description
-h, --help	Provides help information about the osmt job command.
-s, --show	Displays background task information.
-d DELAY, --delay DELAY	Allows for postponing the pending restart task. For example, to postpone the task for 1 hour and 50 minutes, set the value to " 1:50:00 ".
-c, --cancel	Cancels the current job.
-y, --yes	Considers by default that the user agrees the operation.
-V, --verbose	(Optional) Displays detailed process logs.

5.4.2 Description of the /etc/osmt/osmt.conf File

This section describes the OSMT configuration items that you are advised not to modify in the **osmt.conf** file.

```
[auto]
# if auto_upgrade is True, the osmt-agent will auto upgrade rpms use osmt.conf and reboot between time
```

```
interval we specified
# the value of cycle_time means the osmt-agent will check upgrade every cycle_time seconds, default
86400s(1 day)
# When a configuration item has a line break, you need to leave a space or tab at the beginning of the line
auto_upgrade = False
cycle_time = 3600
minimal_interval = 3600
auto_upgrade_window = "22:00-05:00"
auto_upgrade_interval = 1
[Package]
# There are three rules of filters, all enabled by default. Severity will be effect only when the types contain
security, it is the subtype of security.
# The following are the three rules:
# 1. white list has the highest priority, if whitelist is configured then ignore other rules and filter out the
whitelist packages from the full list of packages to be upgrade
# 2. Filter the update range by types, when the types contain security, further filter the severity of
security updates severity, only upgrade the severity level of security.
# 3. Filter blacklist to remove packages in blacklist from types filter results, and packages which depend
on packages in blacklist will also be removed.
# filters must contain at least one types rule, if the types rule is empty, the -a option will not upgrade any
packages (by default all 3 filters are enabled).

filters = "types, blacklist"
whitelist = ""

# types include: security, bugfix, enhancement, newpackage, unknown
# if types is empty, no package will be upgrade
# types = security, bugfix, enhancement, newpackage, unknown
types = "security"
# severity is the subtype of security, include: low, moderate, important, critical
severity = "important, critical"
blacklist = "mysql"
# RPM packages that only take effect after an OS restart
need_reboot_rpms = kernel,kernel-debug,glibc,glibc-utils,dbus,dbus-python...
preinstalled_only = False
[backup]
store_path = /var/log
backup_dir = /etc,/usr,/boot,/var,/run
exclude_dir =
recover_service =
[resource_needed]
#the minimum resources required(MB)
#min_req_boot_space = 100
#min_req_backup_space = 8192
#min_req_root_space = 1536
#min_req_memory = 512
[cmdline]
cmdline_value = crashkernel=512M resume=/dev/mapper/hce-swap rd.lvm.lv=hce/root rd.lvm.lv=hce/swap
crash_kexec_post_notifiers panic=3 nmi_watchdog=1 rd.shell=0
[conflict]
#conflict_rpm = test1,test2

[strategy]
timeout_action = "stop"
timeout_action_before = 0

[check]
daemon_whitelist=sysstat-collect.service, sysstat-summary.service, systemd-tmpfiles-clean.service
```

Table 5-9 Configuration items that should not be modified in **osmt.conf**

Configuration Item	Description
types	The parameter that defines the RPM package update scope, including five configuration items security , bugfix , enhancement , newpackage , and unknown . You are advised not to modify it unless in some special cases.
severity	The system upgrades security updates by default. You are advised not to modify it unless in some special cases.
[resource_needed]	The minimum resource required by the system to perform the update or update check. You are advised not to modify it unless in some special cases.

6 Security Updates for Huawei Cloud EulerOS

6.1 Security Updates Overview

This section describes how to query and install Huawei Cloud EulerOS security updates using yum or dnf.

The support for yum and dnf depends on the OS version. This section uses yum as an example.

NOTE

As a substitute for yum, dnf delivers better performance. The methods for using dnf and yum are the same.

- Huawei Cloud EulerOS 2.0 and later support both yum and dnf.
- Huawei Cloud EulerOS versions earlier than 2.0 support only yum.

6.2 About CVE

Common Vulnerabilities and Exposures (CVE) is a list of publicly disclosed vulnerabilities, each with a unique CVE serial number. To ensure Huawei Cloud EulerOS security, Huawei Cloud closely follows industry vulnerability warnings and fixes software vulnerabilities in a timely manner. You can view the security updates at Huawei Cloud EulerOS security advisories:

- [Huawei Cloud EulerOS 1.1 Security Advisories](#)
- [Huawei Cloud EulerOS 2.0 Security Advisories](#)

In accordance with the Common Vulnerability Scoring System (CVSS), Huawei Cloud EulerOS security updates are classified into the following levels:

- Critical (high risk, mandatory)
- Important (medium high risk, strongly recommended)
- Moderate (medium risk, recommended)
- Low (low risk, optional)

If you have installed Huawei Cloud EulerOS 1.1 or later (Huawei Cloud EulerOS 2.0 as an example), install security updates by performing operations in the subsequent sections.

 **NOTE**

The command outputs vary depending on the Huawei Cloud EulerOS version.

6.3 Yum Command Parameters

Command format: yum <command> [option]

Table 6-1 Major <command> parameters

Parameter	Description
help	Displays help information.
updateinfo	Displays summary of package updates.
upgrade	Installs package updates.
check-update	Checks for available package updates.

Table 6-2 Major [option] parameters

Parameter	Description
-h, --help, --help-cmd	Displays command help information.
--security	Displays available security updates.
--advisory ADVISORY	Specifies specific advisories. This parameter can be used together with updateinfo , upgrade , and check - update . Multiple packages are separated by commas (,).
--cve CVES	Specifies specific CVEs. This parameter can be used together with updateinfo , upgrade , and check - update . Multiple packages are separated by commas (,).
--sec-severity {Critical,Important,Moderate,Low}	Specifies specific security levels. This parameter can be used together with updateinfo , upgrade , and check - update . Values in the brackets can be any combination of security update levels.

 NOTE

Use the **yum --help** command to obtain more information.

6.4 Querying Security Updates

Command format: `yum updateinfo <command> [option]`

- Run the **yum updateinfo** command to query all available security updates.

```
[root@localhost ~]# yum updateinfo
Last metadata expiration check: 0:03:05 ago on Thu 08 Sep 2022 05:30:23 PM CST.
Updates Information Summary: available
12 Security notice(s)
4 Critical Security notice(s)
6 Important Security notice(s)
2 Moderate Security notice(s)
```

- Major `<command>` parameters include:

- **list**: lists the available security updates.

```
[root@localhost ~]# yum updateinfo list
Last metadata expiration check: 0:03:32 ago on Thu 08 Sep 2022 05:30:23 PM CST.
HCE2-SA-2022-0006 Critical/Sec. curl-7.79.1-2.h6.hce2.x86_64
HCE2-SA-2022-0011 Moderate/Sec. gnupg2-2.2.32-1.h6.hce2.x86_64
HCE2-SA-2022-0002 Important/Sec. kernel-5.10.0-60.18.0.50.h425_2.hce2.x86_64
```

- **info <SA ID>**: queries the security updates of a specific advisory.

```
[root@localhost ~]# yum updateinfo info HCE2-SA-2022-0029
Last metadata expiration check: 5:09:15 ago on Tue 13 Sep 2022 09:43:13 AM CST.
```

```
=====
An update for python3 is now available for HCE 2.0
=====
```

```
Update ID: HCE2-SA-2022-0029
```

```
Type: security
```

```
Updated: 2022-09-08 22:08:34
```

```
CVEs: CVE-2021-28861
```

```
Description: Security Fix(es):
```

```
: Python 3.x through 3.10 has an open redirection vulnerability in lib/http/server.py due to no protection against multiple (/) at the beginning of URI path which may lead to information disclosure. (CVE-2021-28861)
```

```
Severity: Important
```

- Major `[option]` parameters include:

- **--sec-severity={Critical,Important,Moderate,Low}**: queries security updates of a specific level. Values in the brackets can be any combination of security update levels.

In the following example, **--sec-severity=Critical** is used to query critical security updates.

```
[root@localhost ~]# yum updateinfo list --sec-severity=Critical
Last metadata expiration check: 0:10:15 ago on Thu 08 Sep 2022 05:30:23 PM CST.
HCE2-SA-2022-0006 Critical/Sec. curl-7.79.1-2.h6.hce2.x86_64
HCE2-SA-2022-0003 Critical/Sec. libarchive-3.5.2-1.h2.hce2.x86_64
HCE2-SA-2022-0006 Critical/Sec. libcurl-7.79.1-2.h6.hce2.x86_64
...
```

In the following example, **--sec-severity={Critical,Moderate}** is used to query critical and moderate security updates.

```
[root@localhost ~]# yum updateinfo list --sec-severity={Critical,Moderate}
Last metadata expiration check: 0:11:07 ago on Thu 08 Sep 2022 05:30:23 PM CST.
HCE2-SA-2022-0006 Critical/Sec. curl-7.79.1-2.h6.hce2.x86_64
HCE2-SA-2022-0011 Moderate/Sec. gnupg2-2.2.32-1.h6.hce2.x86_64
HCE2-SA-2022-0003 Critical/Sec. libarchive-3.5.2-1.h2.hce2.x86_64
...
```

- **--cve=<CVE ID>**: queries security updates of a specific CVE.

```
[root@localhost ~]# yum updateinfo info --cve=CVE-2021-28861
Last metadata expiration check: 5:10:38 ago on Tue 13 Sep 2022 09:43:13 AM CST.
=====
An update for python3 is now available for HCE 2.0
=====
Update ID: HCE2-SA-2022-0029
Type: security
Updated: 2022-09-08 22:08:34
CVEs: CVE-2021-28861
Description: Security Fix(es):
: Python 3.x through 3.10 has an open redirection vulnerability in lib/http/server.py due to no
protection against multiple (/) at the beginning of URI path which may lead to information
disclosure. (CVE-2021-28861)
Severity: Important
```

NOTE

Use the `yum updateinfo --help` command to obtain more information.

6.5 Checking for Security Updates

- Run the `yum check-update --security` command to check for available security updates in the OS.

```
[root@localhost ~]# yum check-update --security
Last metadata expiration check: 0:11:39 ago on Thu 08 Sep 2022 05:30:23 PM CST.
curl.x86_64 7.79.1-2.h6.hce2 hce2
gnupg2.x86_64 2.2.32-1.h6.hce2 hce2
kernel.x86_64 5.10.0-60.18.0.50.h425_2.hce2 hce2
unbound-libs.x86_64 1.13.2-3.h2.hce2 hce2
```

- Run the `yum check-update --sec-severity={Critical,Important,Moderate,Low}` command to check for security updates of the specified level.

Values in the brackets can be any combination of security update levels.

```
[root@localhost ~]# yum check-update --sec-severity=Moderate
Last metadata expiration check: 0:23:57 ago on Thu 08 Sep 2022 05:30:23 PM CST.
gnupg2.x86_64 2.2.32-1.h6.hce2 hce2
python3-unbound.x86_64 1.13.2-3.h2.hce2 hce2
unbound-libs.x86_64 1.13.2-3.h2.hce2 hce2
```

6.6 Installing Security Updates

- Run the `yum upgrade --security` command to install all security updates.

```
[root@localhost ~]# yum upgrade --security
Last metadata expiration check: 5:21:24 ago on Tue 13 Sep 2022 09:43:13 AM CST.
Dependencies resolved.
=====
Package Arch Version Repository Size
=====
Installing:
Kernel x86_64 5.10.0-60.18.0.50.h498_2.hce2 hce2 49 M
Upgrading:
Curl x86_64 7.79.1-2.h6.hce2 hce2 147 k
...
Transaction Summary
=====
Install 1 Package
Upgrade 22 Packages
Total download size: 69 M
Is this ok [y/N]:
```

- Run the **yum upgrade --sec-severity={Critical,Important,Moderate,Low}** command to install security updates of the specified level.

Values in the brackets can be any combination of security update levels.

```
[root@localhost ~]# yum upgrade --sec-severity=Moderate
Last metadata expiration check: 0:32:27 ago on Thu 08 Sep 2022 05:30:23 PM CST.
Dependencies resolved.
=====
Package                Architecture    Version          Repository        Size
=====
Upgrading:
gnupg2                 x86_64         2.2.32-1.h6.hce2 hce2              2.2 M
python3-unbound        x86_64         1.13.2-3.h2.hce2 hce2              96 k
unbound-libs          x86_64         1.13.2-3.h2.hce2 hce2              505 k
Transaction Summary
=====
Upgrade 3 Packages
Total download size: 2.8 M
Is this ok [y/N]:
```

- Run the **yum upgrade --advisory =<SA ID>** command to install security updates of a specific advisory.

Multiple packages are separated by commas (,).

```
root@localhost ~]# yum upgrade --advisory=HCE2-SA-2022-0033
Last metadata expiration check: 1:48:44 ago on Tue 13 Sep 2022 09:43:13 AM CST.
Dependencies resolved.
=====
Package                Architecture    Version          Repository        Size
=====
Upgrading:
python3-rpm            x86_64         4.17.0-8.h13.hce2 hce2              87 k
rpm                    x86_64         4.17.0-8.h13.hce2 hce2              498 k
rpm-libs               x86_64         4.17.0-8.h13.hce2 hce2              376 k
Transaction Summary
=====
Upgrade 3 Packages
Total download size: 962 k
Is this ok [y/N]:
```

- Run the **yum upgrade --cve=<CVE ID>** command to install security updates of a specific CVE.

Multiple packages are separated by commas (,).

```
[root@localhost ~]# yum upgrade --cve=CVE-2021-28861
Last metadata expiration check: 5:16:36 ago on Tue 13 Sep 2022 09:43:13 AM CST.
Dependencies resolved.
=====
Package                Architecture    Version          Repository        Size
=====
Upgrading:
python3                x86_64         3.9.9-7.h10.hce2 hce2              8.0 M
python3-unversioned-command x86_64         3.9.9-7.h10.hce2 hce2              3.9 k
Transaction Summary
=====
Upgrade 2 Packages
Total download size: 8.0 M
Is this ok [y/N]:
```

7 Obtaining the openEuler Extended Software Packages

By default, Huawei Cloud EulerOS does not load the openEuler repository to avoid conflicts with Huawei Cloud EulerOS software packages.

Huawei Cloud EulerOS 2.0 is compatible only with openEuler 22.03 LTS. This section describes how you can obtain the extended software packages of openEuler 22.03 LTS.

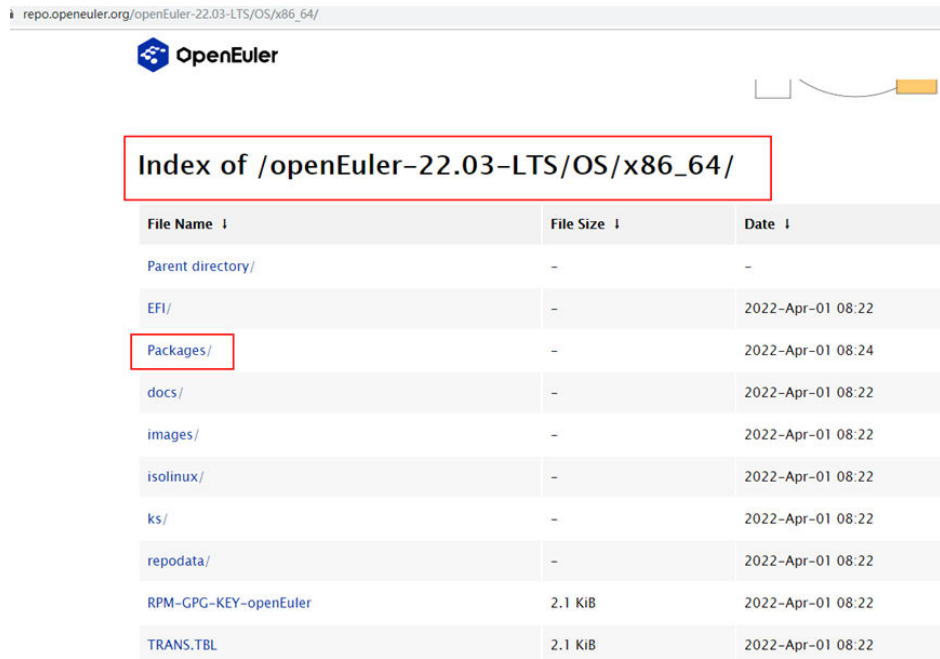
How to Obtain	Application Scenario	How to Install RPM Packages	How to Back Up and Restore Repository Files
Run the wget Command to Download the RPM Packages	Installing a few RPM packages	Manually download and install the RPM packages.	N/A
Batch Downloading RPM Packages Using the Repository File	Installing multiple RPM packages	The RPM packages are automatically installed. You do not need to download them again.	Back up the repository files in the <code>/etc/yum.repos.d</code> directory first and then delete them. Restore the repository files after the RPM packages are installed.

Run the wget Command to Download the RPM Packages

You can run the `wget` command to download the RPM packages. `hadoop-3.1-common-3.1.4-4.oe2203.noarch.rpm` is used as an example.

1. [Sign in to the openEuler community.](#)

- In the **OS/everything** directory, select the **aarch64/** or **x86_64/ System architecture directory** and open the **Packages/** directory.



- Search for the required RPM package, for example, **hadoop-3.1-common-3.1.4-4.oe2203.noarch.rpm**.

h2-javadoc-1.4.196-2.oe2203.noarch.rpm	149.2 KiB	2022-Apr-01 07:56
hadoop-3.1-client-3.1.4-4.oe2203.noarch.rpm	80.4 MiB	2022-Apr-01 07:52
hadoop-3.1-common-3.1.4-4.oe2203.noarch.rpm	29.2 MiB	2022-Apr-01 07:48
hadoop-3.1-common-native-3.1.4-4.oe2203.x86_64.rpm	65.3 KiB	2022-Apr-01 07:49
hadoop-3.1-devel-3.1.4-4.oe2203.x86_64.rpm	16.1 KiB	2022-Apr-01 07:48

- Right-click the RPM package, copy the download link, and run the **wget** command to download the RPM package.

```
[root@ecs-zty ~]# wget https://repo.openeuler.org/openEuler-22.03-LTS/everything/x86_64/Packages/hadoop-3.1-common-3.1.4-4.oe2203.noarch.rpm
--2023-05-18 21:31:10-- https://repo.openeuler.org/openEuler-22.03-LTS/everything/x86_64/Packages/hadoop-3.1-common-3.1.4-4.oe2203.noarch.rpm
Resolving repo.openeuler.org (repo.openeuler.org)... 49.0.230.196
Connecting to repo.openeuler.org (repo.openeuler.org)|49.0.230.196|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 30580053 (29M) [application/x-redhat-package-manager]
Saving to: 'hadoop-3.1-common-3.1.4-4.oe2203.noarch.rpm'

hadoop-3.1-common-3.1.4-4.oe2203.noarch.rpm [=====] 29.16M 243KB/s in 2m 15s
2023-05-18 21:33:25 (221 KB/s) - 'hadoop-3.1-common-3.1.4-4.oe2203.noarch.rpm' saved [30580053/30580053]
```

- If the following information is displayed, the package is downloaded.

```
[root@ecs-zty ~]# ls
hadoop-3.1-common-3.1.4-4.oe2203.noarch.rpm
[root@ecs-zty ~]#
```

- Run the **rpm -ivh hadoop-3.1-common-3.1.4-4.oe2203.noarch.rpm** command to install the RPM package. If the following information display is displayed, the package is installed.

If other packages are required during the installation, repeat the above steps to install the dependent packages.

```
root@ecs-zty ~]# rpm -ivh hadoop-3.1-common-3.1.4-4.oe2203.noarch.rpm
warning: hadoop-3.1-common-3.1.4-4.oe2203.noarch.rpm: Header U3 RSA/SHA1 Signature, key ID b25e7f66: NOKEY
error: Failed dependencies:
  apache-zookeeper is needed by hadoop-3.1-common-3.1.4-4.oe2203.noarch
  leveldb is needed by hadoop-3.1-common-3.1.4-4.oe2203.noarch
  protobuf2-java is needed by hadoop-3.1-common-3.1.4-4.oe2203.noarch
root@ecs-zty ~]# rpm -ivh hadoop-3.1-common-3.1.4-4.oe2203.noarch.rpm --force
warning: hadoop-3.1-common-3.1.4-4.oe2203.noarch.rpm: Header U3 RSA/SHA1 Signature, key ID b25e7f66: NOKEY
error: Failed dependencies:
  apache-zookeeper is needed by hadoop-3.1-common-3.1.4-4.oe2203.noarch
  leveldb is needed by hadoop-3.1-common-3.1.4-4.oe2203.noarch
  protobuf2-java is needed by hadoop-3.1-common-3.1.4-4.oe2203.noarch
root@ecs-zty ~]# rpm -ivh hadoop-3.1-common-3.1.4-4.oe2203.noarch.rpm --nodeps
warning: hadoop-3.1-common-3.1.4-4.oe2203.noarch.rpm: Header U3 RSA/SHA1 Signature, key ID b25e7f66: NOKEY
Preparing...
Updating / installing...
 1:hadoop-3.1-common-3.1.4-4.oe2203
root@ecs-zty ~]#
```

Batch Downloading RPM Packages Using the Repository File

Download the RPM packages under the **openEuler-22.03-LTS/everything/x86_64** directory and run the **yum** commands to install them.

1. Ensure that your VM can access <https://repo.openeuler.org/openEuler-22.03-LTS/>.
2. Configure a yum repository.

Go to the **/etc/yum.repos.d** directory, create an **openEuler.repo** file, and copy and paste the following content to **openEuler.repo**.

NOTE

As the **openEuler.repo** file conflicts with the Huawei Cloud EulerOS repository file, back up the Huawei Cloud EulerOS repository file in the **/etc/yum.repos.d** directory and delete this file, and then create an **openEuler.repo** file.

```
[openeuler]
name=openeuler
baseurl=https://repo.openeuler.org/openEuler-22.03-LTS/OS/x86_64/
gpgcheck=1
enabled=1
priority=3
gpgkey=https://repo.openeuler.org/openEuler-22.03-LTS/OS/x86_64/RPM-GPG-KEY-openeuler
[everything]
name=everything
baseurl=https://repo.openeuler.org/openEuler-22.03-LTS/everything/x86_64
gpgcheck=1
enabled=1
priority=3
gpgkey=https://repo.openeuler.org/openEuler-22.03-LTS/everything/x86_64/RPM-GPG-KEY-openeuler
```

3. Run the **yum clean all** command to delete the cache of the yum repository.
4. Run **yum makecache** to connect to the newly configured repository. If the following information is displayed, the repository is connected.

```
root@ecs-zty ~]# yum clean all
12 files removed
root@ecs-zty ~]# yum makecache
openeuler                235 kB/s | 3.4 MB  00:14
everything                229 kB/s | 1.6 MB  01:12
HCE 2.0 base              59 MB/s | 6.1 MB  00:00
HCE 2.0 updates           1.8 MB/s | 7.8 MB  00:04
Last metadata expiration check: 0:00:02 ago on Thu 18 May 2023 09:55:03 PM CST.
Metadata cache created.
```

5. Install the RPM packages. **hadoop-3.1-common** is used as an example here.
 - a. Run the **yum list** command to check whether the package exists.

```
root@ecs-zty ~]# yum list | grep hadoop-3.1-common
hadoop-3.1-common.noarch      3.1.4-4.oe2203      @System
hadoop-3.1-common-native.x86_64 3.1.4-4.oe2203      everything
root@ecs-zty ~]#
```

- b. Run the **yum -y install hadoop-3.1-common** command to install the package. If the following information is displayed, the package is installed.

```
Installed:
alsa-lib-1.2.5.1-1.oe2203.x86_64
cairo-1.17.4-1.oe2203.x86_64
dejavu-fonts-2.37-1.oe2203.noarch
fonts-filesystem-2.0.5-2.oe2203.noarch
gdk-pixbuf2-2.42.6-1.oe2203.x86_64
graphite2-1.3.14-5.oe2203.x86_64
gtk2-2.24.33-4.oe2203.x86_64
harfbuzz-2.8.2-1.oe2203.x86_64
java-1.8.0-openjdk-1:1.8.0.312.b07-11.oe2203.x86_64
javapackages-filesystem-5.3.0-6.oe2203.noarch
leveldb-1.20-4.oe2203.x86_64
libXau-1.0.9-2.oe2203.x86_64
libXcursor-1.2.0-1.oe2203.x86_64
libXext-1.3.4-2.oe2203.x86_64
libXft-2.3.4-1.oe2203.x86_64
libXinerama-1.1.4-5.oe2203.x86_64
libXrender-0.9.10-10.oe2203.x86_64
libdatrie-0.2.13-1.oe2203.x86_64
libjpeg-turbo-2.1.1-1.oe2203.x86_64
libtiff-4.3.0-9.oe2203.x86_64
lksctp-tools-1.0.19-1.oe2203.x86_64
nspr-4.32.0-1.oe2203.x86_64
nss-help-3.72.0-2.oe2203.x86_64
nss-util-3.72.0-2.oe2203.x86_64
pixman-0.40.0-1.oe2203.x86_64
tzdata-java-2021e-2.oe2203.noarch
xorg-x11-fonts-otf-7.5-24.oe2203.noarch
atk-2.36.0-1.oe2203.x86_64
copy-jdk-configs-4.0-1.oe2203.noarch
fontconfig-2.13.94-1.oe2203.x86_64
fribidi-1.0.10-1.oe2203.x86_64
gdk-pixbuf2-modules-2.42.6-1.oe2203.x86_64
gtk-update-icon-cache-3.24.30-5.oe2203.x86_64
hadoop-3.1.1-common-3.1.4-4.oe2203.noarch
hicolor-icon-theme-0.17-4.oe2203.noarch
java-1.8.0-openjdk-headless-1:1.8.0.312.b07-11.oe2203.x86_64
jbigkit-libs-2.1-17.oe2203.x86_64
libX11-1.7.2-2.oe2203.x86_64
libXcomposite-0.4.5-1.oe2203.x86_64
libXdamage-1.1.5-1.oe2203.x86_64
libXfixes-6.0.0-1.oe2203.x86_64
libXi-1.0-1.oe2203.x86_64
libXrandr-1.5.2-1.oe2203.x86_64
libXtst-1.2.3-10.oe2203.x86_64
libfontenc-1.1.4-2.oe2203.x86_64
libthai-0.1.28-4.oe2203.x86_64
libxcb-1.14-1.oe2203.x86_64
lua-posix-35.0-1.oe2203.x86_64
nss-3.72.0-2.oe2203.x86_64
nss-softoken-3.72.0-2.oe2203.x86_64
pango-1.49.3-2.oe2203.x86_64
protobuf2-java-2.5.0-3.oe2203.x86_64
xorg-x11-font-utis-1:7.5-43.oe2203.x86_64
zookeeper-3.6.2-2.4.oe2203.noarch

Complete!
```

6. Restore the repository file.

After installing the required openEuler packages, delete the openEuler.repo file and restore the repository file deleted in step 2.

8 Creating a Docker Image and Starting a Container

This section describes how you can create a Docker image for Huawei Cloud EulerOS and start the container on Huawei Cloud EulerOS.

Constraints

- The version of Huawei Cloud EulerOS running the container image must be the same as that of the created container image.

Creating an Image Archive File

1. Ensure that the repository address of Huawei Cloud EulerOS is correctly configured.

Check whether the repository address in the `/etc/yum.repos.d/hce.repo` file is correctly configured as follows:

```
[base]
name=HCE $releasever base
baseurl=https://repo.huaweicloud.com/hce/$releasever/os/$basearch/
enabled=1
gpgcheck=1
gpgkey=https://repo.huaweicloud.com/hce/$releasever/os/RPM-GPG-KEY-HCE-2

[updates]
name=HCE $releasever updates
baseurl=https://repo.huaweicloud.com/hce/$releasever/updates/$basearch/
.....
```

2. Create a temporary directory as the root system file of the Docker image and install the software package in the temporary directory.

```
rm -rf /tmp/docker_rootfs
mkdir -p /tmp/docker_rootfs
yum --setopt=install_weak_deps=False --installroot /tmp/docker_rootfs --releasever 2.0 install
bash yum coreutils security-tool procps-ng vim-minimal tar findutils filesystem hce-repos hce-
rootfiles cronie -y
```

CAUTION

- Replace **releasever** in the command with the Huawei Cloud EulerOS version number.
- You can also install other required software packages here, but ensure that the space in the **/tmp** directory is sufficient.

3. Go to the temporary directory of the chroot file system.

```
chroot /tmp/docker_rootfs
```

```
[root@localhost tmp]# chroot /tmp/docker_rootfs
[root@localhost /]#
[root@localhost /]#
```

a. Execute **HCE security-tool.sh** to disable unnecessary services.

```
export EULEROSECURITY=0
echo "export TMOUT=300" >> /etc/bashrc
/usr/sbin/security-tool.sh -d / -c /etc/hce_security/hwsecurity/hce_security_install.conf -
u /etc/hce_security/usr-security.conf -l /var/log/hce-security.log -s
```

During the execution, it is normal if the following errors are displayed:

- The service file was not found. The service is not started in the chroot file system.
- The **/etc/sysconfig/init** file for booting the system was not found. The tool disables services during system startup. The image rootfs is not involved in system startup.
- The **/proc/sys/kernel/sysrq** file was not found. This file is used for calling after the system is started and does not exist in the chroot file system.

```
[root@localhost /]# /usr/sbin/security-tool.sh -d / -c /etc/hce_security/hwsecurity/hce
security_install.conf -u /etc/hce_security/usr-security.conf -l /var/log/hce-security.lo
g -s
Failed to disable unit, unit avahi-daemon.service does not exist.
Failed to disable unit, unit avahi-daemon.socket does not exist.
Failed to disable unit, unit snmpd.service does not exist.
Failed to disable unit, unit squid.service does not exist.
Failed to disable unit, unit smb.service does not exist.
Failed to disable unit, unit vsftpd.service does not exist.
Failed to disable unit, unit tftp.service does not exist.
Failed to disable unit, unit nfs-server.service does not exist.
System has not been booted with systemd as init system (PID 1). Can't operate.
Failed to connect to bus: Host is down
System has not been booted with systemd as init system (PID 1). Can't operate.
Failed to connect to bus: Host is down
Failed to disable unit, unit rpcbind.service does not exist.
Failed to disable unit, unit slapd.service does not exist.
Failed to disable unit, unit dhcpd.service does not exist.
System has not been booted with systemd as init system (PID 1). Can't operate.
Failed to connect to bus: Host is down
Failed to enable unit, unit haveged.service does not exist.
System has not been booted with systemd as init system (PID 1). Can't operate.
Failed to connect to bus: Host is down
Failed to disable unit, unit postfix.service does not exist.
System has not been booted with systemd as init system (PID 1). Can't operate.
Failed to connect to bus: Host is down
Failed to disable unit, unit chronyd.service does not exist.
System has not been booted with systemd as init system (PID 1). Can't operate.
Failed to connect to bus: Host is down
System has not been booted with systemd as init system (PID 1). Can't operate.
Failed to connect to bus: Host is down
tail: cannot open '/etc/sysconfig/init' for reading: No such file or directory
package at is not installed
[security-tool.sh:839] [error] package at does not exist
cronie-1.5.7-1.r2.hce2.x86_64
sysctl: cannot stat /proc/sys/kernel/sysrq: No such file or directory
Removed /etc/systemd/system/multi-user.target.wants/hce-security.service.
```

b. Uninstall the security-tool, cronie, and systemd software packages and their dependent software packages.

- ```
cp -af /etc/pam.d /etc/pam.d.bak
rm -f /etc/yum/protected.d/sudo.conf /etc/yum/protected.d/systemd.conf
yum remove -y security-tool cronie systemd
rpm -e --nodeps logrotate crontabs
rm -rf /etc/pam.d
mv /etc/pam.d.bak /etc/pam.d
sh -c 'shopt -s globstar; for f in $(ls /**/*.rpm); do rm -f $f; done'
[-d /var/lib/dnf] && rm -rf /var/lib/dnf/*
[-d /var/lib/rpm] && rm -rf /var/lib/rpm/*_db.*
```
- c. Remove the **/boot** directory.  

```
rm -rf /boot
```
  - d. Set the container image language to en\_US.  

```
cd /usr/lib/locale;rm -rf $(ls | grep -v en_US | grep -vw C.utf8)
rm -rf /usr/share/locale/*
```
  - e. Remove shared files **man**, **doc**, **info**, and **mime**.  

```
rm -rf /usr/share/{man,doc,info,mime}
```
  - f. Remove the cached log files.  

```
rm -rf /etc/ld.so.cache
[-d /var/cache/ldconfig] && rm -rf /var/cache/ldconfig/*
[-d /var/cache/dnf] && rm -rf /var/cache/dnf/*
[-d /var/log] && rm -rf /var/log/*.log
```
  - g. Remove the Java security certificate.  

```
rm -rf /etc/pki/ca-trust/extracted/java/cacerts /etc/pki/java/cacerts
```
  - h. Remove **/etc/machine-id**.  

```
rm -rf /etc/machine-id
```
  - i. Remove **/etc/mtab**.  

```
rm -rf /etc/mtab
```
4. Exit from the chroot file system.  

```
exit
```
  5. Compress the temporary directory and generate the Docker image archive file **hce-docker.x86\_64.tar.xz**.  
The archive path is **/tmp/docker\_rootfs/hce-docker.x86\_64.tar.xz**.  

```
pushd /tmp/docker_rootfs/
tar cvf hce-docker.x86_64.tar .
xz hce-docker.x86_64.tar
popd
```
  6. Convert the image archive file to a file with layer information.  
You cannot run the **docker load** command to load the image because the generated image archive file does not contain layer information. You need to run the **docker import** command to generate an image from the image archive file, and then run the **docker save** command to save the image as an image file with layer information. In this way, you can run the **docker load** command to load the image. In the following example, the image archive file is **docker\_save.tar.xz** and the image is named **my\_image**.  

```
docker import hce-docker.x86_64.tar.xz my_image:v1
docker save -o docker_save.tar.xz my_image:v1
```

You can run the **docker load** command to load the **docker\_save.tar.xz** file.

```
docker load -i docker_save.tar.xz
```

## Starting a Container Using an Image Archive File

1. Ensure that the repository address of Huawei Cloud EulerOS is correctly configured.

Check whether the repository address in the **/etc/yum.repos.d/hce.repo** file is correctly configured as follows:

```
[base]
name=HCE $releasever base
baseurl=https://repo.huaweicloud.com/hce/$releasever/os/$basearch/
enabled=1
gpgcheck=1
gpgkey=https://repo.huaweicloud.com/hce/$releasever/os/RPM-GPG-KEY-HCE-2

[updates]
name=HCE $releasever updates
baseurl=https://repo.huaweicloud.com/hce/$releasever/updates/$basearch/
.....
```

2. Install the Docker software package.

```
yum install docker -y
```

3. Use the image archive file to create a container image.

```
mv /tmp/docker_rootfs/hce-docker.x86_64.tar.xz .
docker import hce-docker.x86_64.tar.xz
```

Run the **docker images** command to check the container image ID. In this example, the container image ID is **6cfefae3a541**.

```
[root@localhost ~]# mv /tmp/docker_rootfs/hce-docker.x86_64.tar.xz .
[root@localhost ~]# docker import hce-docker.x86_64.tar.xz
sha256:6cfefae3a5410e3b48208f8a9e0a28fc08fa1b3a62ad39b27196a742969d5bfc
[root@localhost ~]#
[root@localhost ~]# docker images
REPOSITORY TAG IMAGE ID CREATED SIZE
<none> <none> 6cfefae3a541 6 seconds ago 181MB
```

#### NOTE

To create an image, you can run the following command to specify the **REPOSITORY** and **TAG** parameters:

```
docker import [OPTIONS] file|URL|- [REPOSITORY[:TAG]]
```

4. Run the image bash script in the container.

If the shell view changes after you run the following command, you have started the bash for the container:

```
docker run -it 6cfefae3a541 bash
```

```
[root@localhost ~]# docker run -it 6cfefae3a541 bash
[root@5c639b63fc0e ~]#
[root@5c639b63fc0e ~]#
```

# 9 Tools

## 9.1 BiSheng Compiler

BiSheng compiler is a high-performance, high-reliability, and easy-to-expand compiler developed by Huawei. BiSheng compiler has introduced multiple compilation technologies and supports programming languages C, C++, and Fortran.

### Constraints

- BiSheng compiler can be used only with Huawei Cloud EulerOS 2.0 of the x86 architecture.
- The native Clang-compiled languages of Huawei Cloud EulerOS 2.0 cannot work with the Clang-compiled languages of the BiSheng compiler. If you have installed the native Clang-compiled language package and need to use it, do not install the BiSheng compiler anymore.

If you have installed the BiSheng compiler but want to use the native Clang-compiled language, run **rpm -e bisheng-compiler** to delete the BiSheng compiler, and then open a new terminal to use the native Clang-compiled language.

### Installing BiSheng Compiler

1. Confirm that the repository is configured correctly.

Check whether the parameters in the **/etc/yum.repos.d/hce.repo** file are configured correctly. The correct configuration is as follows:

```
[base]
name=HCE $releasever base
baseurl=https://repo.huaweicloud.com/hce/$releasever/os/$basearch/
enabled=1
gpgcheck=1
gpgkey=https://repo.huaweicloud.com/hce/$releasever/os/RPM-GPG-KEY-HCE-2

[updates]
name=HCE $releasever updates
baseurl=https://repo.huaweicloud.com/hce/$releasever/updates/$basearch/
.....
```

2. Run **yum install bisheng-compiler** to install BiSheng compiler.

3. Run **source /usr/local/bisheng-compiler/env.sh** to import environment variables.  
If you have opened a new terminal, import the environment variables into the new terminal.
4. Check whether BiSheng compiler has been installed.  
Run **clang -v** to view the version number. If the command output contains the BiSheng compiler version, the BiSheng compiler has been successfully installed.

## Using BiSheng Compiler

1. Compile and run C and C++ programs.  

```
clang [command line flags] hello.c -o hello.o
./hello.o
clang++ [command line flags] hello.cpp -o hello.o
./hello.o
```
2. Compile and run Fortran programs.  

```
flang [command line flags] hello.f90 -o hello.o
./hello.o
```
3. Specify a linker.  
Specify the LLVM lld for BiSheng compiler. If you do not specify the LLVM lld, the default linker ld will be used.  

```
clang [command line flags] -fuse-ld=lld hello.c -o hello.o
./hello.o
```

## 9.2 Workload Accelerator

### 9.2.1 Overview

Workload Accelerator is a tool provided by Huawei Cloud for application optimization.

It works in two ways:

- Static acceleration  
Static acceleration collects the PMU monitoring information on the CPU when an application is running, and statically builds a new high-performance binary based on the collected information. It may require only adjustments to compiler parameters, without involving any other code modifications. There are two optimization methods in static acceleration.
  - **Use the native BOLT tool:** Only fixed parameter combinations can be used to optimize applications.
  - **Run the hce-wae-auto commands:** Different parameter combinations can be generated based on the user-defined parameter range for application optimization.
- Dynamic acceleration  
Dynamic acceleration directly accelerates the application process without interrupting services.

**Table 9-1** Advantages and disadvantages of the optimization methods

| Optimization Method  | Pros                                                                                                                                                                                                                                                                                                                                           | Cons                                                                                           |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|
| Static acceleration  | Applications are optimized on the basis of binary executable files, and the program code does not need to be modified.                                                                                                                                                                                                                         | After the optimization, you need to restart the applications.                                  |
| Dynamic acceleration | Application processes are optimized directly. Applications do not need to be restarted. Optimization results can be saved using application snapshots. In addition, binary file source tracing can be ensured, and application processes can be continuously optimized for iteration until the performance improvement reaches the bottleneck. | Data can only be collected using instrumentation, and optimization can be performed only once. |

## Constraints

- Workload Accelerator can be used only with Huawei Cloud EulerOS 2.0 of the x86 architecture.
- Only the **root** user can use Workload Accelerator.

## Process for Using Workload Accelerator to Optimize Applications

1. [Install the Workload Accelerator](#).
2. Optimize the application through [static acceleration](#) or [dynamic acceleration](#).

### 9.2.2 Installing Workload Accelerator

1. Confirm that the repository is configured correctly.

Check whether the parameters in the `/etc/yum.repos.d/hce.repo` file are configured correctly. The correct configuration is as follows:

```
[base]
name=HCE $releasever base
baseurl=https://repo.huaweicloud.com/hce/$releasever/os/$basearch/
enabled=1
gpgcheck=1
gpgkey=https://repo.huaweicloud.com/hce/$releasever/os/RPM-GPG-KEY-HCE-2

[updates]
name=HCE $releasever updates
baseurl=https://repo.huaweicloud.com/hce/$releasever/updates/$basearch/
.....
```

2. Run `yum install hce-wae` to install Workload Accelerator.
3. Check whether Workload Accelerator has been installed.

Run **llvm-bolt**. Workload Accelerator has been installed if the following information is displayed.

```
[root@localhost sdb]# llvm-bolt --help | more
OVERVIEW: BOLT - Binary Optimization and Layout Tool

USAGE: llvm-bolt [options] <executable> <executable>

OPTIONS:

BOLT generic options:
--bolt-id=<string> - add any string to tag this execution in the output binary via bolt info section
--data=<string> - <data file>
--data2=<string> - <data file>
--deterministic-debuginfo - disables parallel execution of tasks that may produce non-deterministic debug info
--dwarf-output-path=<string> - Path to where .dwo files or dwp file will be written out to.
--dymo-stats - print execution info based on profile
--enable-bat - write BOLT Address Translation tables
--hot-data - hot data symbols support (relocation mode)
--hot-functions-at-end - if reorder-functions is used, order functions putting hottest last
```

## 9.2.3 Static Acceleration

### Preparations

1. Run the following command to check whether the binary file to be optimized can be located again:

```
readelf -a application | grep .rela.text
```

- If the binary file contains **.rela.text**, the file can be located again. In this case, the application can be optimized.

```
[root@hce2 shard]# readelf -a gemini-redis-server | grep .rela.text
[15] .rela.text RELA 0000000000000000 0750a490
Relocation section '.rela.text' at offset 0x750a490 contains 241690 entries:
[root@hce2 shard]#
```

- If the binary file does not contain **.rela.text**, to allow BOLT to re-arrange functions (in addition to re-arranging code within functions) in your program, you need to add **--emit-relocs** or **-q** to the final link step of your application.

2. Collect the logs when the application is running.

After deploying and preheating the application, you can run **llvm-bolt -instrument -o -instrumentation-file** to configure the way how the logs are collected.

For example, to collect logs every 30 seconds after the **test.so** file is executed, run the following command to save the logs to the **test.log** file:

```
llvm-bolt tests.so -instrument -o testd.so -instrumentation-file=test.log -instrumentation-sleep-time=30 -instrumentation-no-counters-clear
```

- **instrument -o**: indicates the new dynamic library file generated after the log collection way is configured. In this example, the newly generated dynamic library is **testd.so**.
- **instrumentation-file**: indicates the name of the file where the logs are saved. In this example, the log file is **test.log**.
- **instrumentation-sleep-time**: indicates the interval for collecting logs, in seconds. In this example, logs are collected every 30 seconds.
- **instrumentation-no-counters-clear**: indicates that the log counter information is not cleared after each log collection to ensure that the log context is continuous.

3. Run the application corresponding to **testd.so**. The run logs of the application are automatically saved in the **test.log** file.

Workload Accelerator will optimize the application based on the dynamic data in the **test.log** file.



## Procedure

- Use the native BOLT tool to optimize the application.

Once you have **test.log** ready, you can use it for optimizations with BOLT. For example, you can run the following command:

```
llvm-bolt <executable> -o <executable>.bolt -data=test.log -reorder-
blocks=ext-tsp -reorder-functions=hfsort -split-functions -split-all-cold -
split-eh -dyno-stats
```

### NOTE

The parameters in the command are for illustration purposes and are not the optimal combination of parameters.

- Run the **hce-wae-auto** command to optimize the application.

The **hce-wae-auto** command optimizes the application based on user-defined [configuration file](#).

The command format is **hce-wae-auto [-h] [-c <Path>] [-s <Keyword>] [-e <Pattern>] [-l] [--free] *Application***.

### NOTE

- After you run the **hce-wae-auto** command, the path of the generated parameter set is the same as the binary output path (`/data/hce-wae-auto/hce-wae-auto.data`) in the configuration file by default.
- After you run the **hce-wae-auto** command, a log file is automatically generated, and its default path is `/var/log/hce-wae-auto/hce-wae-auto.log`.

**Table 9-2** Parameter description

| Parameter | Value Required | Value Type | Mandatory | Description                                                                                                                                 |
|-----------|----------------|------------|-----------|---------------------------------------------------------------------------------------------------------------------------------------------|
| -h        | No             | /          | No        | Displays command parameter help information.                                                                                                |
| -c        | Yes            | String     | No        | Defines the configuration file path. The default path is <code>/etc/hce-wae-auto.conf</code> .                                              |
| -s        | Yes            | String     | No        | Defines the keyword for fuzzy search of the last parameter set.<br>If there is no saved parameter set, this parameter will not take effect. |

| Parameter   | Value Required | Value Type | Mandatory | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------|----------------|------------|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -e          | Yes            | String     | No        | Executes the selected parameter combination based on the content of the last parameter set. A sequence number must be specified. <ul style="list-style-type: none"> <li>A single number specifies the parameter combination corresponding to the sequence number.</li> <li>The colon (:) indicates the sequence number range. For example, 4:6 indicates the sequence number range from 4 to 6. :7 indicates the sequence number range is from 1 to 7. 5: indicates the sequence number range is from 5 to the last. : indicates all the sequence numbers.</li> <li>A single number can be used together with a sequence number range (:) and separated using a comma (,). For example, 1,4:6 indicates sequence numbers 1, 4, 5, and 6.</li> </ul> |
| -l          | No             | /          | No        | Generates and prints the parameter set based on the configuration file. The command is not automatically executed and the parameter set is not saved.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| --free      | No             | /          | No        | Indicates that the validity of parameters in the configuration file is not verified.<br>Workload Accelerator performs simple verification on the parameters in the configuration file to determine whether the parameters are used for optimization using BOLT. If you run the <b>--free</b> command, the verification is canceled and parameters that are not for optimization can be entered.                                                                                                                                                                                                                                                                                                                                                     |
| application | No             | /          | Yes       | Specifies the binary file to be executed. You can add a relative path or an absolute path, for example, <b>mysqld</b> , <b>/usr/bin/mysqld</b> , or <b>/bin/mysqld</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

#### NOTE

If the **-c**, **-s**, **-e** and **-l** parameters conflict, **-s** has the highest priority, followed by **-l** and then **-e**, and **-c** has the lowest priority.

The following table describes some example commands.

| Example Command                               | Description                                                                                                                                                                                                                                             |
|-----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>hce-wae-auto mysqld</b>                    | Reads the configuration from the default path, generates a parameter set based on configuration parameters, automatically executes the command, and automatically saves the parameter set after the command is executed.                                |
| <b>hce-wae-auto mysqld -c /etc/my.conf -l</b> | Reads configuration parameters from the <b>/etc/my.conf</b> path, generates a parameter set, prints the generated command, but does not execute the command or save the parameter set.                                                                  |
| <b>hce-wae-auto mysqld -c /etc/my.conf</b>    | Reads configuration parameters from the <b>/etc/my.conf</b> path, generates a parameter set, automatically executes the command, and automatically saves the parameter set after the command is executed.                                               |
| <b>hce-wae-auto mysqld -s align</b>           | Searches for the keyword <b>align</b> in fuzzy mode from the parameter set generated last time and prints the matched parameter combination.                                                                                                            |
| <b>hce-wae-auto mysqld -e 4:6</b>             | Selects the parameter combinations of No.4, No.5, and No.6 from the parameter set generated last time to generate a binary file.                                                                                                                        |
| <b>hce-wae-auto mysqld -e 1,4:6</b>           | Selects the parameter combinations of No.1, No.4, No.5, and No.6 from the parameter set generated last time to generate a binary file.                                                                                                                  |
| <b>hce-wae-auto mysqld -e :</b>               | Selects all parameter combinations from the parameter set generated last time to generate a binary file.                                                                                                                                                |
| <b>hce-wae-auto mysqld -e 4:6 -l</b>          | Selects the parameter combinations of No.4, No.5, and No.6 from the parameter set generated last time, prints the generated command, but does not execute the command.                                                                                  |
| <b>hce-wae-auto mysqld -e :-l</b>             | Selects all parameter combinations from the parameter set generated last time, prints the generated command, but does not execute the command.                                                                                                          |
| <b>hce-wae-auto mysqld --free</b>             | Reads the configuration from the default path but does not verify parameters. Generates a parameter set based on configuration parameters, automatically executes the command, and automatically saves the parameter set after the command is executed. |
| <b>hce-wae-auto -h</b>                        | Displays command help information.                                                                                                                                                                                                                      |

## 9.2.4 Dynamic Acceleration

### Preparations

Before dynamic acceleration, you need to perform two checks. Applications can be dynamically accelerated only when both conditions are met.

1. Run the following command to check whether the binary file to be optimized can be located again:

```
readelf -a application | grep .rela.text
```

- If the binary file contains **.rela.text**, the file can be located again. In this case, the application can be optimized.

```
[root@hce2 shard]# readelf -a gemini-redis-server | grep .rela.text
[15] .rela.text RELA 0000000000000000 0750a490
Relocation section '.rela.text' at offset 0x750a490 contains 241690 entries:
[root@hce2 shard]#
```

- If the binary file does not contain **.rela.text**, to allow BOLT to re-arrange functions (in addition to re-arranging code within functions) in your program, you need to add **--emit-relocs** or **-q** to the final link step of your application.

2. Run the **hce-wae --check** command to check whether an application supports dynamic acceleration.

If **3** is displayed, dynamic acceleration is supported.

```
[root@localhost ~]# hce-wae --check /data/apps/mysql-8.0.28/bin/mysqld
2023-09-14 20:41:21,968-INFO: Log level: INFO
2023-09-14 20:41:38,425-INFO: check /data/apps/mysql-8.0.28/bin/mysqld result: 3
```

### Procedure

In the following operations, the `mysqld` application in the `/data/apps/mysql-8.0.28/bin` directory is used as an example.

1. Generate an instrumentation application and run it.
  - a. Run the `/data/hce-wae/dbo/gen_instrumentation /data/apps/mysql-8.0.28/bin/mysqld` command to generate an instrumentation application.

Command format: `/data/hce-wae/dbo/gen_instrumentationApplication path`

```
[root@localhost ~]# ./data/hce-wae/dbo/gen_instrumentation /data/apps/mysql-8.0.28/bin/mysqld
BOLT-INFO: Target architecture: x86_64
BOLT-INFO: BOLT version: b93bff771fd4
BOLT-INFO: first alloc address is 0x400000
BOLT-INFO: creating new program header table at address 0x4200000, offset 0x3e00000
BOLT-WARNING: debug info will be stripped from the binary. Use -update-debug-sections to keep it.
BOLT-INFO: enabling relocation mode
BOLT-INFO: forcing -jump-tables=move for instrumentation
BOLT-INFO: enabling -align-macro-fusion=all since no profile was specified
BOLT-INFO: enabling lite mode
BOLT-WARNING: split function detected on input : _ZL28delete_dictionary_tablespacev.cold/1. The support is limited in re
location mode.
BOLT-INFO: 0 out of 69377 functions in the binary (0.0%) have non-empty execution profile
BOLT-INFO: the input contains 6637 (dynamic count : 0) opportunities for macro-fusion optimization that are going to be
fixed
BOLT-INSTRUMENTER: Number of indirect call site descriptors: 37513
BOLT-INSTRUMENTER: Number of indirect call target descriptors: 87025
BOLT-INSTRUMENTER: Number of function descriptors: 68951
BOLT-INSTRUMENTER: Number of branch counters: 686408
BOLT-INSTRUMENTER: Number of ST leaf node counters: 481684
BOLT-INSTRUMENTER: Number of direct call counters: 248150
BOLT-INSTRUMENTER: Total number of counters: 1416242
BOLT-INSTRUMENTER: Total size of counters: 11329936 bytes (static alloc memory)
BOLT-INSTRUMENTER: Total size of string table emitted: 7829557 bytes in file
BOLT-INSTRUMENTER: Total size of descriptors: 70240472 bytes in file
BOLT-INSTRUMENTER: Profile will be saved to file /data/hce-wae/dbo/tmp/perf.fdata
BOLT-INFO: 425568 instructions were shortened
BOLT-INFO: removed 11295 empty blocks
BOLT-INFO: UCE removed 143092 blocks and 8769293 bytes of code.
BOLT-INFO: SCTC: patched 0 tail calls (0 forward) tail calls (0 backward) from a total of 0 while removing 0 double jump
s and removing 0 basic blocks totalling 0 bytes of code. CTCs total execution count is 0 and the number of times CTCs ar
e taken is 0.
BOLT-INFO: output linked against instrumentation runtime library, lib entry point is 0xd0616e0
BOLT-INFO: clear procedure is 0xd05fe10
BOLT-INFO: setting _end to 0x3c83b18
BOLT-INFO: setting _end to 0x3c83b18
BOLT-INFO: patched build-id (flipped last bit)
[root@localhost ~]#
```

The `mysqld.inst` instrumentation file is generated in the current directory.

```
[root@localhost ~]# ls -al *.inst
-rwxrwxrwx. 1 root root 304995968 Sep 15 10:24 mysqld.inst
[root@localhost ~]#
```

- b. Run the instrumentation file to obtain the PID of the application process. In this example, the PID is 87042.

```
[root@localhost ~]# ./data/apps/mysql-8.0.28/bin/mysqld.inst -uroot &
[1] 87042
[root@localhost ~]# 2023-09-15T02:30:08.152217Z 0 [System] [MY-010116] [Server] /data/apps/mysql-8.0.28/bin/mysqld.inst (mysqld 8.0.28) start
ing as process 87042
2023-09-15T02:30:08.565656Z 1 [System] [MY-013576] [InnoDB] InnoDB initialization has started.
2023-09-15T02:30:11.547443Z 1 [System] [MY-013577] [InnoDB] InnoDB initialization has ended.
2023-09-15T02:30:13.032029Z 0 [System] [MY-010229] [Server] Starting XA crash recovery ...
2023-09-15T02:30:13.040024Z 0 [System] [MY-010232] [Server] XA crash recovery finished.
2023-09-15T02:30:13.377357Z 0 [Warning] [MY-010058] [Server] CA certificate ca.pem is self signed.
2023-09-15T02:30:13.377427Z 0 [System] [MY-013602] [Server] Channel mysql_main configured to support TLS. Encrypted connections are now suppo
rted for this channel.
2023-09-15T02:30:13.488591Z 0 [System] [MY-011323] [Server] X Plugin ready for connections. Bind-address: '::' port: 33060, socket: /tmp/mysql
x.sock
2023-09-15T02:30:13.488631Z 0 [System] [MY-010931] [Server] /data/apps/mysql-8.0.28/bin/mysqld.inst: ready for connections. Version: '8.0.28'
socket: '/tmp/mysql.sock' port: 3306 Source distribution.
```

2. Create a dynamic acceleration configuration file for `mysqld`.

Each application to be optimized must have a configuration file. Workload Accelerator dynamically accelerates the application based on the configuration file.

- a. Run the following command to copy the default configuration file `/data/hce-wae/config/mysqld.conf`:

```
[root@localhost]# cp /data/hce-wae/config/hce-wae-tmp.conf /data/hce-wae/config/mysqld.conf
```

- b. Set the `origin-exe` field in the configuration file.

`origin-exe` indicates the location of the application to be optimized. In this example, the location is `/data/apps/mysql-8.0.28/bin/mysqld`.

```
[root@localhost]# vim /data/hce-wae/config/mysqld.conf
```

```
Configuration for hce-wae
each config file should be configured for one mission, which is
one running process in the environment
[mission]
config the way to collect run-time data, can be defined in [perf, instrument]
log-type=instrument

config the way to hotpatch the optimized segments, can be defined in [mode1, mode2, mode3]
mode1 will hotpatch by dbo tools, other two types are currently not supported
hotpatch-type=mode1

config the location where the optimized executable file to be saved in
snapshot-path=/data/hce-wae/snapshot

let hce-wae tool know the path to the origin executable file
should be configured when log-type is 'instrument'
origin-exe=/data/apps/mysql-8.0.28/bin/mysqld

config stop strategy type, three strategies can be selected:
1. run-times=N stop accelerating after optimized N times
2. period=N stop accelerating after N seconds
3. condition="example condition" stop accelerating after satisfied the condition, currently not supported
[stop-strategy]
run-times=1
```

3. Use the configuration file and the PID to configure dynamic acceleration.

Command format: **hce-wae --conf [PID] [/path/to/config]**

```
[root@localhost ~]# hce-wae --conf 87042 /data/hce-wae/config/mysqld.conf
2023-09-15 10:33:29,478-INFO: Log level: INFO
2023-09-15 10:33:29,479-INFO: mission will stop after run 1 times
2023-09-15 10:33:29,479-INFO: record mission from config succeed
[root@localhost ~]#
```

4. Enable dynamic acceleration to optimize the instrumentation application.

Command format: **hce-wae --start [PID]**

```
[root@localhost ~]# hce-wae --start 87042
2023-09-15 10:43:18,421-INFO: Log level: INFO
2023-09-15 10:43:18,422-INFO: start dbo mission for /data/apps/mysql-8.0.28/runtime_output_directory/mysqld.inst ...
2023-09-15 10:43:18,422-INFO: extracting call sites, it may take several minutes ...
2023-09-15 10:44:00,414-INFO: preparing ...
start running dbo server to monitor process 87042
2023-09-15 10:44:00,419-INFO: run dbo for /data/apps/mysql-8.0.28/runtime_output_directory/mysqld.inst(pid: 87042) succeed
2023-09-15 10:44:00,419-INFO: starting ...
start optimizing
2023-09-15 10:44:00,422-INFO: run dbo for /data/apps/mysql-8.0.28/runtime_output_directory/mysqld.inst(pid: 87042) succeed
2023-09-15 10:44:00,422-INFO: mission started, target pid: 87042
2023-09-15 10:44:00,422-INFO: recording started mission info...
2023-09-15 10:44:00,423-INFO: start mission worker...
2023-09-15 10:44:00,424-INFO: mission worker for 87042 started
```

You can view the optimization status from the **--status** parameter. If the status is **Running**, the process is being optimized. If the status is **Finished**, the process has been optimized.

Command format: **hce-wae --status [PID]**

```
[root@localhost ~]# hce-wae --status 87042
2023-09-15 10:44:48,379-INFO: Log level: INFO
2023-09-15 10:44:48,384-INFO: status: {
 "status": "Running",
 "sub_status": "DataCollecting",
 "run_times": 0,
 "failed_code": "NoFailed"
}
[root@localhost ~]# hce-wae --status 87042
2023-09-15 10:45:22,066-INFO: Log level: INFO
2023-09-15 10:45:22,072-INFO: status: {
 "status": "Finished",
 "sub_status": "Done",
 "run_times": 1,
 "failed_code": "NoFailed"
}
```

5. Use the **--snapshot** parameter to generate an optimized **.dbo** binary snapshot file. In this example, the file is **mysqld.dbo**.

```
[root@localhost ~]# hce-wae --snapshot 87042
2023-09-15 10:46:00,433-INFO: Log level: INFO
2023-09-15 10:46:00,438-INFO: dbo snapshot for 87042 succeed
2023-09-15 10:46:00,438-INFO: snapshot generated, target_pid: 87042, path: /data/hce-wae/snapshot/
[root@localhost ~]# ll /data/hce-wae/snapshot/
total 71384
-rwxrwxrwx. 1 root root 148936512 Sep 15 10:46 mysqld.dbo
```

The default snapshot path is **/data/hce-wae/snapshot/**. You can change the path in the configuration file as needed. You can use this snapshot file to run the application without repeated optimization.

6. Terminate dynamic acceleration to stop application optimization.

Command format: **hce-wae --stop [PID]**

```
[root@localhost ~]# hce-wae --stop 87042
2023-09-15 10:46:20,867-INFO: Log level: INFO
stop dbo server successfully!
2023-09-15 10:46:20,871-INFO: dbo stop for 87042 succeed
2023-09-15 10:46:20,871-INFO: mission stopped, target_pid: 87042
```

## CLI for Dynamic Acceleration

A CLI is provided for dynamic acceleration. [Table 9-3](#) lists the supported commands.

**Figure 9-1** Startup page of dynamic acceleration

```
[root@localhost ~]# hce-wae --console

Welcome to hce-wae!

Version : 1.0.0
Release : 0.0.16.hce2
Architecture: x86_64

Type: 'h' or 'help' for help with commands
 'quit' to quit

hce-wae>
```

**Figure 9-2** Help page of dynamic acceleration

```

hce-wae> help
h, help | Show this help message and exit
list | List the process information of the current environment,
 | including PID, PPID, and command
check <Binary file path> | Check if the application support static or dynamic
 | acceleration. 0: both not support; 1: static only; 2:
 | dynamic only; 3: both are supported.
show <Pid> | Show all non-kernel processes which contain one of the
 | shared-object libraries that the target process
 | depends on.
conf <Pid> <Config path> | Config the target process by pid
conf <Process> | Set the Config of the target process
start <Pid> | Start dynamic acceleration for process by its pid
stop <Pid> | Stop dynamic acceleration for process by its pid
status <Pid> | Show the status of the dynamic accelerating process by
 | its pid
snapshot <Pid> | Make a snapshot for the dynamic accelerating process
 | by its pid
quit | Exit console
hce-wae>

```

**Table 9-3** CLI commands supported for dynamic acceleration

| Command | How to Use   | Description                                                                                                                                                                                                                                                                                                              |
|---------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| list    | list         | Obtains the process information in the current environment, including the PID, PPID, and command information.                                                                                                                                                                                                            |
| status  | status <PID> | Queries ongoing dynamic acceleration tasks and their status. The following information is returned: <ul style="list-style-type: none"> <li>• PID</li> <li>• Process name</li> <li>• Optimization count</li> <li>• Acceleration status (Initialized, DataCollecting, Optimizing, HotPatching, Done, or Failed)</li> </ul> |
| check   | check <PID>  | Checks whether the current process supports static or dynamic acceleration and displays the text result.                                                                                                                                                                                                                 |
| show    | show <PID>   | Queries the process dependencies.                                                                                                                                                                                                                                                                                        |
| conf    | conf <PID>   | Configures the dynamic acceleration capability for the target process. After you run the command, configuration options are displayed in sequence. Configure the options as prompted.                                                                                                                                    |



| Command  | How to Use        | Description                                |
|----------|-------------------|--------------------------------------------|
| start    | start <PID>       | Enables dynamic acceleration.              |
| stop     | stop <PID>        | Stops dynamic acceleration.                |
| snapshot | snapshot<br><PID> | Generates a dynamic acceleration snapshot. |
| quit     | quit              | Exits the CLI.                             |
| h/help   | h/help            | Shows help information.                    |

## 9.2.5 Configuration File

This section describes each configuration item in the [static acceleration configuration file](#) and [dynamic acceleration configuration file](#).

### Static Acceleration Configuration File

The following figure shows the default configuration of static acceleration. You can customize a configuration file to optimize applications.

```
Section 'binary' defined options associated with binary file
[binary]
Output path for generated binary files, absolute path is recommended
default: /data/hce-wae-auto/
binary_out_path = "/data/hce-wae-auto"

Threshold for the number of generated binary file, currently not in use
default value: 100
binary_num_threshold = 100

Name suffix for auto-generated binary files
default value: `blot.auto`
binary_file_suffix = "blot.auto"

Section 'parameter' defined option associated with user defined parameter collection
[parameter]
User defined parameter collection, parameters in this option will automatically be separated to different
combination parameter groups, which are used as the parameter input for `llvm-blot` respectively, and
generate different binary files.
For those parameters which can be assigned values, use `=` connect parameter name and parameter value.
Each line should contain one parameter and should not configure the same parameters in this option.
example:
-plt=all
parameters =
 align-blocks
 frame-opt
 align-functions=1

Section 'include' defined include filter for auto-generated parameter collections
[include]
options here should be the parameters combination that should be included for the auto-generated
parameter group.
Each line should contain one parameter, if the parameter is not assigned a value,
it must end with `=`
example:
frame-opt=none
align-blocks=

Section 'exclude' defined exclude filter for auto-generated parameter collections
[exclude]
Options here should be the parameters combination that should be excluded for the auto-generated
parameter group.
Options has the same formate rule with the 'include' section
frame-opt=none
```

Table 9-4 Configuration information

| Module | Description                                                                                      |
|--------|--------------------------------------------------------------------------------------------------|
| binary | Binary file attributes. For details about the binary parameters, see <a href="#">Table 9-5</a> . |

| Module    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| parameter | A collection of user-defined parameters. A parameter set is generated based on this collection of parameters. At least one parameter must be defined.<br>You can run the <b>llvm-bolt -h</b> command to view all parameters.                                                                                                                                                                                                                                                                                                                                                                                                        |
| include   | Parameters that must be contained in the collection of user-defined parameters. Multiple parameters are allowed. The logical relationship between parameters is <b>AND</b> .<br>The key of the configuration item is the parameter name, and the value is specified, for example, <b>frame-opt=none</b> .<br><b>NOTE</b> <ul style="list-style-type: none"> <li>If the value cannot be specified or does not need to be specified, end the key with an equal sign (=), for example, <b>frame-opt=</b>.</li> <li>If <b>include</b> and <b>exclude</b> contain the same parameter, <b>exclude</b> precedes <b>include</b>.</li> </ul> |
| exclude   | Parameters that do not need to be contained in the collection of user-defined parameters. Multiple parameters are allowed. The logical relationship between parameters is <b>AND</b> .<br><b>NOTE</b> <ul style="list-style-type: none"> <li>If the value cannot be specified or does not need to be specified, end the key with an equal sign (=), for example, <b>frame-opt=</b>.</li> <li>If <b>include</b> and <b>exclude</b> contain the same parameter, <b>exclude</b> precedes <b>include</b>.</li> </ul>                                                                                                                    |

**Table 9-5** binary parameters

| Item               | Value Type | Default Value        | Description                                                                  |
|--------------------|------------|----------------------|------------------------------------------------------------------------------|
| binary_out_path    | String     | "/data/hce-wae-auto" | Defines the path for saving the binary file that is automatically generated. |
| binary_file_suffix | String     | "blot.auto"          | Defines the suffix of the binary file name that is automatically generated.  |

## Example Static Acceleration Configuration File

```
[binary]
binary_out_path = "/data/llvm_auto"
binary_num_threshold = 1000
binary_file_suffix = "blot.auto"

[parameter]
parameters =
 --align-blocks # Parameter prefix -- can be added.
 frame-opt # If you do not need to specify the value of a parameter in the collection of user-
 # defined parameters, you do not need to end the key with an equal sign (=).
 align-functions=1 # If the value of a parameter in the collection of user-defined parameters is
 # specified, the value of the parameter is 1 in all parameter combinations in the generated parameter set.
```

```
[include]
align-blocks= # If the parameter value cannot be specified, the configuration item is still ended
 # with an equal sign (=).
[exclude]
frame-opt=none # The parameter and its value are specified. Parameter combinations whose
 # parameter is frame-opt and value is none are filtered out from the generated parameter set.
indirect-call-promotion= # The parameter is specified and is of the enumeration type. All parameter
 # combinations whose parameter is frame-opt are filtered out from the generated parameter set.
```

## Dynamic Acceleration Configuration File

The following figure shows the default configuration of dynamic acceleration. You can customize a configuration file to optimize applications.

```
Configuration for hce-wae
each config file should be configured for one mission, which is
one running process in the environment
[mission]
config the way to collect run-time data, can be defined in [perf, instrument]
log-type=instrument

config the way to hotpatch the optimized segments, can be defined in [mode1, mode2, mode3]
mode1 will hotpatch by dbx tools, other two types are currently not supported
hotpatch-type=mode1

config the location where the optimized executable file to be saved in
snapshot-path=/data/hce-wae/snapshot

let hce-wae tool know the path to the origin executable file
should be configured when log-type is 'instrument'
origin-exe=/path/to/origin/executable/file

config stop strategy type, three strategies can be selected:
1. run-times=N stop accelerating after optimized N times
2. period=N stop accelerating after N seconds
3. condition="example condition" stop accelerating after satisfied the condition, currently not supported
[stop-strategy]
run-times=10
```

**Table 9-6** Configuration information

| Module          | Description                                                                                                                                                 |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| [mission]       | Parameters to be configured for the application                                                                                                             |
| log-type        | Indicates the way to collect logs during runtime. Only the instrumentation is supported                                                                     |
| hotpatch-type   | Indicates the way to hotpatch the optimized segments. Only mode1 (ptrace) is supported.                                                                     |
| snapshot-path   | Indicates the path for storing the optimized binary snapshot file.                                                                                          |
| origin-exe      | Indicates the location of the original application. This parameter must be specified when logs are collected using instrumentation.                         |
| [stop-strategy] | Policy for stopping application optimization. Select one of the configurations.                                                                             |
| run-times       | Specifies how many times the application is optimized. Optimization stops when the number is reached. Currently, an application can be optimized only once. |
| period          | Specifies the optimization period, in seconds. Optimization stops when the period expires. The value ranges from <b>1</b> to <b>600</b> .                   |

| Module    | Description                                                                                                      |
|-----------|------------------------------------------------------------------------------------------------------------------|
| condition | Specifies the optimization condition. Optimization stops when this condition is met. This parameter is reserved. |

## 9.3 Pod Bandwidth Management Tool

In hybrid service deployments, the pod bandwidth management function schedules resources based on QoS levels to improve network bandwidth utilization. Huawei Cloud EulerOS provides `oncn-tbwm` for you to manage the pod bandwidth. You can run the `tbwmcli` commands to limit the network rate in packet sending and receiving.

### Prerequisites

Before using the bandwidth management tool, ensure that the virtual NIC `ifb0` is not in use, and load the `ifb` driver.

### Constraints

- Only Huawei Cloud EulerOS 2.0 x86 architecture supports the `tbwmcli` commands.
- Only the root user is allowed to run the `tbwmcli` commands.
- The `tbwmcli` commands can be used to enable the QoS function for only one NIC at a time.
- After the NIC is removed and then inserted, the original QoS rules will be lost. In this case, you need to manually reconfigure the QoS function.
- `tbwmcli` commands are not supported for `cgroup v2`.
- Upgrading the `oncn-tbwm` software package does not affect the enabling status before the upgrade. Uninstalling the `oncn-tbwm` software package will disable the QoS function for all devices.
- Only NIC names containing digits, letters, hyphens (-), and underscores (\_) can be identified.
- Bandwidth limiting may cause protocol stack memory overstock. If this happens, the transport layer protocol will perform backpressure. For UDP and other protocols that do not have backpressure mechanisms, packet loss, `ENOBUFFS`, and inaccurate traffic limiting may occur.
- Network rate limiting for packet receiving depends on the TCP backpressure capability. In scenarios where TCP is not used, network packets have been received by the target NIC, and network rate limiting is not supported.
- The `tbwmcli`, `tc`, and NIC commands cannot be used together. You can only run the `tbwmcli` commands for rate limiting. For example, if the `tc qdisc` rule has been configured for a NIC, enabling the QoS function for the NIC may fail.

### How to Use

1. Install the `oncn-tbwm` software package.

- a. Confirm that the repository is configured correctly.

Check whether the parameters in the `/etc/yum.repos.d/hce.repo` file are configured correctly. The correct configuration is as follows:

```
[base]
name=HCE $releasever base
baseurl=https://repo.huaweicloud.com/hce/$releasever/os/$basearch/
enabled=1
gpgcheck=1
gpgkey=https://repo.huaweicloud.com/hce/$releasever/os/RPM-GPG-KEY-HCE-2

[updates]
name=HCE $releasever updates
baseurl=https://repo.huaweicloud.com/hce/$releasever/updates/$basearch/
.....
```

- b. Run `yum install oncn-tbwm` to install the oncn-tbwm software package.  
c. Verify the oncn-tbwm software package.

- Run `tbwmcli -v`. If the following information is displayed, the installation is successful:

```
version: 1.0
```

- Check whether the following oncn-tbwm service components exist.

```
/usr/bin/tbwmcli
/usr/share/tbwmcli
/usr/share/tbwmcli/README.md
/usr/share/tbwmcli/bwm_prio_kern.o
/usr/share/tbwmcli/tbwm_tc.o
```

2. Run the `tbwmcli` command as required.

**Table 9-7** tbwmcli commands

| Command                                           | Description                                                                                                                                                                                                                                                                                         |
|---------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>tbwmcli -e ethx tbwmcli -d ethx egress</pre> | <p>Enables or disables the packet sending QoS for a NIC.</p> <p>Example: Enabling the packet sending QoS for eth0</p> <pre>tbwmcli -e eth0 enable eth0 egress success</pre> <p>Example: Disabling the packet sending QoS for eth0</p> <pre>tbwmcli -d eth0 egress disable eth0 egress success</pre> |

| Command                                                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>tbwmcli -i ethx online/<br/>offline</p> <p>tbwmcli -d ethx ingress</p> | <p>Enables or disables the packet receiving QoS for a NIC.</p> <p>Example: Enabling the packet receiving QoS for eth0 and setting eth0 as an online NIC</p> <pre>tbwmcli -i eth0 online enable eth0 ingress success, dev is online</pre> <p>Example: Enabling the packet receiving QoS for eth0 and setting eth0 as an offline NIC</p> <pre>tbwmcli -i eth0 offline enable eth0 ingress success, dev is offline</pre> <p><b>NOTE</b><br/>Packet receiving QoS cannot be set for multiple offline NICs at the same time. Packet receiving QoS can only be set for one offline NIC but multiple online NICs.</p> <p>Example: Disabling the packet receiving QoS for eth0</p> <pre>tbwmcli -d eth0 ingress disable eth0 ingress success</pre> |
| <p>tbwmcli -d ethx</p>                                                    | <p>Forcibly disables QoS for a NIC and disables ifb.</p> <p>Example: Forcibly disabling QoS for eth0 and disable ifb</p> <pre>tbwmcli -d eth0 disable eth0 success</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <p>tbwmcli -p istats/estats</p>                                           | <p>Displays the internal statistics of the packet sender and receiver.</p> <p>Example: Displaying the internal statistics of the packet receiver</p> <pre>tbwmcli -p istats offline_target_bandwidth: 94371840online_pkts: 3626190offline_pkts: 265807online_rate: 0offline_rate: 13580offline_prio: 0</pre> <p>Example: Displaying the internal statistics of the packet sender</p> <pre>tbwmcli -p estats offline_target_bandwidth: 94371840online_pkts: 4805452offline_pkts: 373961online_rate: 0offline_rate: 19307offline_prio: 1</pre>                                                                                                                                                                                               |
| <p>tbwmcli -s path &lt;prio&gt;</p> <p>tbwmcli -p path</p>                | <p>Sets or queries the QoS priority of a cgroup.</p> <p>Currently, only two QoS priorities can be set.</p> <ul style="list-style-type: none"> <li>● <b>0</b>: Sets the cgroup online.</li> <li>● <b>-1</b>: Sets the cgroup offline.</li> </ul> <p>Example: Setting the priority of the test_online cgroup to <b>0</b></p> <pre>tbwmcli -s /sys/fs/cgroup/test_online 0 set prio success</pre> <p>Querying the priority of the test_online cgroup</p> <pre>tbwmcli -p /sys/fs/cgroup/test_online prio is 0</pre>                                                                                                                                                                                                                           |

| Command                                                                    | Description                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>tbwmcli -s bandwidth<br/>&lt;low, high&gt;<br/>tbwmcli -p bandwidth</p> | <p>Sets or queries the offline bandwidth range.</p> <p>Example: Setting the offline bandwidth range to 30 Mbit/s to 100 Mbit/s<br/> <b>tbwmcli -s bandwidth 30mb,100mb</b><br/> set bandwidth success</p> <p>Example: Querying the offline bandwidth range<br/> <b>tbwmcli -p bandwidth</b><br/> bandwidth is 31457280(B),104857600(B)</p>  |
| <p>tbwmcli -s waterline<br/>&lt;val&gt;<br/>tbwmcli -p waterlin</p>        | <p>Sets or queries the online network bandwidth threshold.</p> <p>Example: Setting the online network bandwidth threshold to 20 Mbit/s<br/> <b>tbwmcli -s waterline 20mb</b><br/> set waterline success</p> <p>Example: Querying the online network bandwidth threshold<br/> <b>tbwmcli -p waterline</b><br/> waterline is 20971520 (B)</p> |
| <p>tbwmcli -p devs</p>                                                     | <p>Checks the enabling status of all NICs in the system.</p> <pre> <b>tbwmcli -p devs</b> lo    Egress : disabled lo    Ingress : disabled eth0  Egress : disabled eth0  Ingress : enabled, it's offline ifb0  Egress : enabled </pre>                                                                                                      |
| <p>tbwmcli -c</p>                                                          | <p>Forcibly deletes the QoS settings of all NICs.</p>                                                                                                                                                                                                                                                                                       |
| <p>modprobe ifb<br/>numifbs=1</p>                                          | <p>Loads ifb.</p>                                                                                                                                                                                                                                                                                                                           |
| <p>rmmmod ifb</p>                                                          | <p>Uninstalls ifb.</p>                                                                                                                                                                                                                                                                                                                      |

# 10 Kernel Functions and Interfaces

---

## 10.1 OOM Process Control Policy

### Background

Both offline and online services can be configured in an OS. When Out Of Memory (OOM) occurs, the system preferentially ends the process that consumes the most memory in the offline service control group to reclaim the memory. However, some core services are often running offline. If the memory consumed by such services is reclaimed, the OS will be greatly affected.

To solve this problem, Huawei Cloud EulerOS adjusts the memory reclamation policy during OOM and adds the function of configuring cgroups priority. When the memory is insufficient, the kernel traverses cgroups, ends the processes for cgroups with low priorities, and reclaims the memory so that important offline services can keep running.

### Prerequisites

`vm.panic_on_oom` is enabled by default. When the system is out of memory, a kernel panic occurs. Before you use `memcg` OOM for priority configuration, you must ensure that `vm.panic_on_oom` is disabled by running `sysctl -w vm.panic_on_oom=0`.




## Interface Description

| Interface        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Example Value                                                                                           |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| memcg_qos_enable | <p>Specifies whether to enable memcg OOM priority configuration.</p> <ul style="list-style-type: none"><li>• <b>0</b>: Priority configuration is disabled. When OOM occurs, the process that consumes the most memory is ended based on the original OOM operation, and the memory is reclaimed.</li><li>• <b>1</b>: Priority configuration is enabled, and priorities are configured by cgroup. When OOM occurs, all processes in the cgroup with a lower priority are ended and the memory is reclaimed.</li><li>• <b>2</b>: Priority configuration is enabled, and priorities are configured by process. When OOM occurs, the largest process in the cgroup with a lower priority is ended and the memory is reclaimed.</li></ul>                                                                                                                                                                                                                                                                                                                                                                         | <p>The value is an integer ranging from <b>0</b> to <b>2</b>. The default value is <b>0</b>.</p>        |
| memory.qos_level | <p>Specifies how to configure the priorities of cgroups. A smaller value indicates a lower priority.</p> <ul style="list-style-type: none"><li>• When OOM occurs, the current cgroup is used as the parent cgroup, the process with the highest memory usage in the child cgroup with the lowest priority is ended, and the memory is reclaimed.</li><li>• When OOM occurs, cgroups with the same priority will be sorted based on their memory usage, and the cgroup with the largest memory usage is ended.</li></ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>• Before using <b>memory.qos_level</b>, ensure that <b>memcg_qos_enable</b> is set to <b>1</b> or <b>2</b>.</li><li>• By default, the value of <b>memory.qos_level</b> of a newly created cgroup inherits the value of <b>memory.qos_level</b> of the parent cgroup. The priority of a child cgroup is not restricted by the parent cgroup.</li><li>• If the priority of the parent cgroup is changed, the priorities of the child cgroups are automatically changed to be the same as that of the parent cgroup.</li></ul> | <p>The value is an integer ranging from <b>-1024</b> to <b>1023</b>. The default value is <b>0</b>.</p> |

## Interface Configuration Example

Create six cgroups A, B, C, D, E and F, configure the **memcg\_qos\_enable** interface, and set the memcg OOM priorities by specifying **memory.qos\_level**.



**Table 10-1** Data planning

| cgroup | Value of <b>memory.qos_level</b> | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| A      | -8                               | When the OOM operation is performed in the root cgroup, the kernel traverses all cgroups in the root cgroup and finally selects cgroups A and E, both with the lowest priority. Because A and E have the same priority, the kernel continues to compare the memory used by A and E. <ul style="list-style-type: none"> <li>• If you set <b>memcg_qos_enable</b> to <b>1</b>, the system preferentially ends all processes in the cgroup with a large memory usage and reclaims the memory.</li> <li>• If you set <b>memcg_qos_enable</b> to <b>2</b>, the system ends the process with the largest memory usage in two cgroups and reclaims the memory.</li> </ul> |
| B      | 10                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| C      | 1                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| D      | 2                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| E      | -8                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| F      | 3                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

1. Disable **vm.panic\_on\_oom**.  

```
sysctl -w vm.panic_on_oom=0
```
2. Enable memcg OOM priority configuration.  

```
echo 1 > /proc/sys/vm/memcg_qos_enable
```
3. Create cgroups A and B and set their memcg OOM priorities to **-8** and **10**.  

```
mkdir /sys/fs/cgroup/memory/A
mkdir /sys/fs/cgroup/memory/B
cd /sys/fs/cgroup/memory/A
echo -8 > memory.qos_level
cd /sys/fs/cgroup/memory/B
echo 10 > memory.qos_level
```
4. Create child cgroups C and D under cgroup A and child cgroups E and F under cgroup B, and set the memcg OOM priorities of cgroups C, D, E, and F to **1**, **2**, **-8**, and **3**.

```
mkdir /sys/fs/cgroup/memory/A/C
mkdir /sys/fs/cgroup/memory/A/D
mkdir /sys/fs/cgroup/memory/B/E
mkdir /sys/fs/cgroup/memory/B/F
cd /sys/fs/cgroup/memory/A/C
echo 1 > memory.qos_level
cd /sys/fs/cgroup/memory/A/D
echo 2 > memory.qos_level
cd /sys/fs/cgroup/memory/B/E
echo -8 > memory.qos_level
cd /sys/fs/cgroup/memory/B/F
echo 3 > memory.qos_level
```

## 10.2 Multi-level Memory Reclamation Policy

### Background

In high-density hybrid container deployments, offline services with a large number of I/O reads and writes consume a large amount of page cache. As a result, the idle memory of the system decreases, and memory reclamation is triggered when the idle memory watermark is reached. Online tasks enter the slow path for memory reclamation when applying for memory, causing latency and jitter.

To solve this problem, Huawei Cloud EulerOS 2.0 provides multi-level memory reclamation. You can set a memory warning value to trigger a memory reclamation task, which ensures available memory space. For memory reclamation, you can set multiple levels of memory protection watermarks to protect task availability.

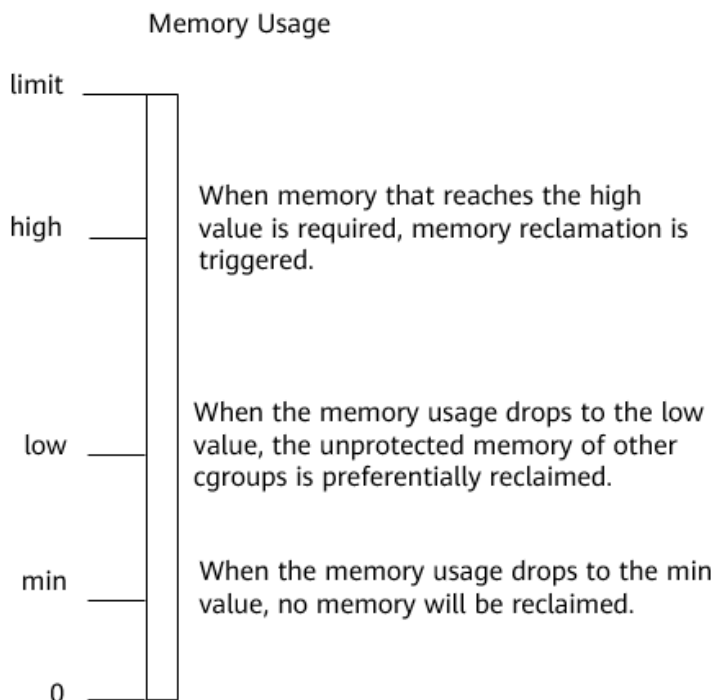
### Constraints

memory.min and memory.low take effect only on leaf cgroups. When a memory cgroup is created, memory.min and memory.low of the parent cgroup are cleared.

### Interface Description

The memory.min, memory.low, and memory.high interfaces exist in the non-root memory cgroup by default. You can write values to the files or read the current configuration. The proper value sequence is  $\text{memory.min} \leq \text{memory.low} < \text{memory.high}$ . The three values can be used independently or together.

The following figure shows the memory reclamation mechanism.




| Interface  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| memory.min | <p>Specifies the minimum amount of memory the cgroup must always retain. The default value is <b>0</b>. Even if there is no memory that can be reclaimed, the system will not reclaim the memory that is less than or equal to the value of this parameter. The read and write operations are described as follows:</p> <ul style="list-style-type: none"> <li>• Reading this interface can view the size (in byte) of the protection memory.</li> <li>• Writing to this interface can set the size of the protection memory. The unit is not limited.</li> <li>• The value ranges from <b>0</b> to <b>memory.limit_in_bytes</b>.</li> </ul>                                  |
| memory.low | <p>Specifies the best-effort memory protection. The default value is <b>0</b>. The system preferentially reclaims the memory of unprotected cgroups. If the memory is still insufficient, the system reclaims the memory between <b>memory.min</b> and <b>memory.low</b>.</p> <p>The read and write operations are described as follows:</p> <ul style="list-style-type: none"> <li>• Reading this interface can view the best-effort memory protection value, in bytes.</li> <li>• Writing to this interface can set the best-effort memory protection value. The unit is not limited.</li> <li>• The value ranges from <b>0</b> to <b>memory.limit_in_bytes</b>.</li> </ul> |

| Interface   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| memory.high | <p>Specifies the memory reclamation warning. The default value is <b>max</b>. When the memory usage of a cgroup reaches the high value, a synchronous memory reclamation task is triggered for the cgroup and its child cgroups. The memory is limited to a value lower than the high value to prevent OOM caused by the memory limit. The read and write operations are described as follows:</p> <ul style="list-style-type: none"> <li>• Reading this interface can view the Throttle limit, in bytes.</li> <li>• Writing to this interface can set the Throttle limit. The unit is not limited.</li> <li>• The value ranges from <b>0</b> to <b>memory.limit_in_bytes</b>.</li> </ul> |

## Interface Configuration Example

Create cgroups A, B, C, D, E and F and configure the memory.min interface.



**Table 10-2** Data planning

| cgroup | memory.limit_in_bytes | memory.min | memory.usage_in_bytes |
|--------|-----------------------|------------|-----------------------|
| A      | 200M                  | 0          | -                     |
| B      | -                     | 0          | -                     |
| C      | -                     | 75M        | 50M                   |
| D      | -                     | 25M        | 50M                   |
| E      | -                     | 0          | 50M                   |
| F      | -                     | 125M       | -                     |

1. Create cgroup A and set **memory.limit\_in\_bytes** to **200M**.

```

mkdir /sys/fs/cgroup/memory/A
echo 200M > /sys/fs/cgroup/memory/A/memory.limit_in_bytes

```

2. Create cgroup B.

```

mkdir /sys/fs/cgroup/memory/A/B

```

3. Create cgroup C, set **memory.min** to **75M**, and create a process that will use 50-MB cache in the cgroup.  

```
mkdir /sys/fs/cgroup/memory/A/B/C
echo 75M > /sys/fs/cgroup/memory/A/B/C/memory.min
```
4. Create cgroup D, set **memory.min** to **25M**, and create a process that will use 50-MB cache in the cgroup.  

```
mkdir /sys/fs/cgroup/memory/A/B/D
echo 25M > /sys/fs/cgroup/memory/A/B/D/memory.min
```
5. Create cgroup E, set **memory.min** to **0**, and create a process that will use 50-MB cache in the cgroup.  

```
mkdir /sys/fs/cgroup/memory/A/B/E
```
6. Create cgroup F, set **memory.min** to **125M**, and request 125-MB cache for memory protection.  

```
mkdir /sys/fs/cgroup/memory/A/F
echo 125M > /sys/fs/cgroup/memory/A/F/memory.min
```

Information similar to the following is displayed:

```
cgroup C: memory.min=75M, memory.usage_in_bytes=50M
```

```
cgroup D: memory.min=25M, memory.usage_in_bytes=25M
```

```
cgroup E: memory.min=0, memory.usage_in_bytes=0
```

```
cgroup B: memory.usage_in_bytes=75M
```

## 10.3 Multi-level Hybrid Scheduling of Kernel CPU cgroups

### Background

In hybrid deployments, the Linux kernel scheduler assigns more scheduling opportunities to high-priority tasks and minimizes the impact of low-priority tasks on kernel scheduling. However, the two-level scheduling of online and offline services cannot meet this requirement.

To solve the problem, Huawei Cloud EulerOS 2.0 allows multi-level scheduling of kernel CPU cgroups and provides the interface **/sys/fs/cgroup/cpu/cpu.qos\_level** to extend the scheduling levels from two to five, allowing users to set the priority for each cgroup separately.

### Constraints

Multi-level hybrid scheduling of the kernel CPU cgroups is developed based on the 5.10.0-60.18.0.50.r692\_16.hce2.x86\_64 kernel. **cpu.qos\_level** supports only cgroup-v1.

### Interface Description

Rules for **cpu.qos\_level** to take effect:

- Completely Fair Scheduler (CFS) selects **task\_group** level by level from top to bottom. **cpu.qos\_level** takes effect for child cgroups under the same parent cgroup.

- When a child cgroup is created, it inherits the `cpu.qos_level` value of the parent cgroup by default, but the `cpu.qos_level` value can be reconfigured.
- For QoS levels with the same priority, their resource competition complies with the policy of CFS.
- On the same CPU, the tasks whose **qos\_level** is less than **0** are always unconditionally preempted by tasks whose **qos\_level** is greater than or equal to is **0**, regardless of their levels.

When a high-priority task is scheduled:


- Online tasks can unconditionally preempt the CPU resources of offline tasks. During multi-core scheduling, online tasks can preferentially preempt the CPU resources of offline tasks on other cores. In the hyper-thread scenario, online tasks with priority 2 can evict offline tasks on the SMT.
- When a task with a higher priority is woken up, the task is accelerated by time slicing and can immediately preempt the CPU resources of the task with a lower priority to obtain a response at a lower latency (the minimum running time slice of CFS is ignored).

**Table 10-3** Interface description

| Interface                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>cpu.qos_level</code> | <p>Specifies the CPU priorities of the cgroups. The value is an integer ranging from -2 to 2. The default value is 0.</p> <ul style="list-style-type: none"> <li>• <code>cpu.qos_level &gt;= 0</code><br/>Tasks in the cgroup are online tasks, which can unconditionally preempt offline tasks.<br/>A lower value indicates a lower priority (<math>0 &lt; 1 &lt; 2</math>). Online tasks with a higher priority can obtain more CPU resources than those with a lower priority.</li> <li>• <code>cpu.qos_level &lt; 0</code><br/>Tasks in the cgroup are offline tasks. The priority of -1 is higher than that of -2, meaning that tasks at level -1 have more CPU resources than tasks at level -2.<br/>If the parent cgroup is set offline, the child cgroup can inherit only the priority of the parent cgroup, and the priority cannot be changed.</li> </ul> |

## Interface Configuration Example

Create cgroups A, B, and C, and configure the `cpu.qos_level` interface.



**Table 10-4** Data planning

| cgroup | cpu.qos_level Value |
|--------|---------------------|
| A      | 1                   |
| B      | -2                  |
| C      | 2                   |

1. Create cgroup A and child cgroups B and C, and set their CPU scheduling priorities to **1**, **-2**, and **2**.

Tasks in cgroups A and C can unconditionally preempt CPU resources of the tasks in cgroup B. cgroup C preferentially preempts CPU resources because the priority of cgroup C is higher than that of cgroup A.

```
mkdir -p /sys/fs/cgroup/cpu/A
echo 1 > /sys/fs/cgroup/cpu/A/cpu.qos_level
mkdir -p /sys/fs/cgroup/cpu/A/B
echo -2 > /sys/fs/cgroup/cpu/A/B/cpu.qos_level
mkdir -p /sys/fs/cgroup/cpu/A/C
echo 2 > /sys/fs/cgroup/cpu/A/C/cpu.qos_level
```

2. Add the task1, task2, and task3 processes to cgroup B.

The CPU scheduling priority of the task1, task2, and task3 processes is **-2**.

```
echo $PID1 > /sys/fs/cgroup/cpu/A/B/tasks
echo $PID2 > /sys/fs/cgroup/cpu/A/B/tasks
echo $PID3 > /sys/fs/cgroup/cpu/A/B/tasks
```

3. Add the task4 and task5 processes to cgroup C.

The CPU scheduling priority of the task4 and task5 processes is **2**.

```
echo $PID4 > /sys/fs/cgroup/cpu/A/C/tasks
echo $PID5 > /sys/fs/cgroup/cpu/A/C/tasks
```

4. View the CPU scheduling priority and processes of cgroup B.

```
[root@localhost cpu_qos]# cat /sys/fs/cgroup/cpu/A/B/cpu.qos_level
-2
[root@localhost boot]# cat /sys/fs/cgroup/cpu/A/B/tasks
1879
1880
1881
```

5. View the CPU scheduling priority and processes of cgroup C.

```
[root@localhost cpu_qos]# cat /sys/fs/cgroup/cpu/A/C/cpu.qos_level
2
[root@localhost boot]# cat /sys/fs/cgroup/cpu/A/C/tasks
1882
1883
```

## 10.4 Kernel Exception Events Analysis

### Background

When Huawei Cloud EulerOS is running, there are some inevitable kernel events, such as **soft lockup**, **RCU (Read-Copy Update) stall**, **hung task**, **global OOM**, **cgroup OOM**, **page allocation failure**, **list corruption**, **bad mm\_struct**, **I/O error**, **EXT4-fs error**, **Machine Check Exception (MCE)**, **fatal signal**, **warning**, and **panic**. This section describes those events and how you can trigger them.



## Soft Lockup

A soft lockup is the symptom of a task or kernel thread not releasing a CPU for a period longer than allowed (20 seconds by default).

- Details

A soft lockup is triggered by the watchdog mechanism of the Linux kernel. The kernel starts a FIFO real-time kernel thread (watchdog) with the highest priority for each CPU. The thread name is watchdog/0, watchdog/1, and so on. The thread invokes the watchdog function every 4 seconds by default. Each time the function is invoked, an hrtimer will be reset to expire after a soft lockup threshold, which is 2 times the duration specified by **watchdog\_thresh** (a kernel parameter) and defaults to 20 seconds.

Within this duration, if watchdog is not scheduled and the hrtimer expires, the kernel prints a soft lockup exception similar to the following:

```
BUG: soft lockup - CPU#3 stuck for 23s! [kworker/3:0:32]
```

- Triggering method

Disable interrupts or preemption to result in an infinite loop.

## RCU Stall

An RCU stall is an exception that RCU kernel threads are not scheduled within the RCU grace period.

- Details

RCU readers are allowed to access any data, and RCU records information about these readers. When RCU writers are updating data, they copy a backup and modify the data on the backup. After all readers exit, writers replace the old data at a time.

Writers can only replace the old data after all readers stop referencing the old data. This period of time is a grace period.

If the readers do not exit even after the grace period expires and the writers wait for a period longer than the grace period, an RCU stall will be reported.

- Triggering method

Stimulate a scenario described in **Documentation/RCU/stallwarn.txt** to trigger RCU stalls. An example is that CPU keeps looping in the RCU read-side critical section and keeps looping when the interrupt or preemption function is disabled.

## Hung Task

When the kernel detects that a process is in the **D** state for a period longer than the specified time, a hung task exception is reported.

- Details

One status of a process is **TASK\_UNINTERRUPTIBLE**, which is also called the **D** state. A process in the **D** state can be woken up only by **wake\_up**. When the kernel introduces the **D** state, the process waits for the I/O to complete. When I/Os are normally, the process should not be in the **D** state for a long time.

The kernel creates a thread (khungtaskd) to periodically traverse all processes in the system and check whether there is a process that is in the **D** state for a

period longer than the preset duration (120 seconds by default). If there is such a process, related warnings and process stacks will be printed and reported. If **hung\_task\_panic** is configured (through proc or kernel startup parameters), a panic is initiated directly.

- Triggering method

Create a kernel thread, set it to the **D** state, and use the scheduler to release the time slice.

## Global OOM

The Linux OOM killer is a memory management mechanism. When there is less available memory, the kernel kills some processes to release some memory to ensure system continuity.

- Details

When the kernel allocates memory to a process but the system memory is insufficient, OOM will occur. The OOM killer traverses all processes, scores the processes based on their memory usage, selects a process with the highest score, and terminates this process to release memory.

The kernel source code is `linux/mm/oom_kill.c`, and the core function is `out_of_memory()`. The following describes the processing flow:

- a. The kernel instructs the modules that are registered with `oom_notify_list` in the system to release some memory. If these modules release some memory, it will take no more actions. If the memory fails to be reclaimed, it will go to the next step.
- b. Generally, the OOM killer is triggered when the kernel is allocating memory to a process. If the process has a pending SIGKILL or is exiting, the kernel will terminate this process to release memory. Otherwise, the kernel will go to the next step.
- c. The kernel checks the settings of the system administrator using `panic_on_oom` and determines whether to perform OOM killer or panic in case of OOM. If the kernel selects panic, the system will crash and restart. If the kernel selects OOM killer, it will go to the next step.
- d. The kernel enters the OOM killer and checks the system settings. The system administrator can terminate the process that attempts to request memory and causes OOM, or other processes. If the system administrator chooses to terminate the current process, the OOM killer stops. Otherwise, the kernel will go to the next step.
- e. The kernel invokes `select_bad_process` to select appropriate processes, and then invokes `oom_kill_process` to terminate the selected processes. If `select_bad_process` does not select any process, the kernel will enter the panic state.

- Triggering method

Execute the program that occupies large memory until the memory is insufficient.

## cgroup OOM

- Difference from global OOM

The memory of cgroup OOM is different from that of global OOM. When the memory usage of processes in the cgroup exceeds the upper limit, the cgroup kills the processes to release the memory.

- Triggering method  
Execute the program that occupies large memory until the memory is insufficient.

## Page Allocation Failure

A page allocation failure is an error reported by the system when a program fails to apply for an idle page. When a program applies for memory of an order, but there is no idle page whose order is higher than the required order in the system memory, the kernel reports an error.

- Details  
Linux uses the buddy system to efficiently allocate and manage memory. All idle page tables (with a size of 4 KB per page table) are linked to an array containing 11 elements. Each element in the array forms a linked list with consecutive page tables of the same size. The number of page tables is 1, 2, 4, 8, 16, 32, and 64, or 128, 256, 512, and 1,024. The maximum continuous memory that can be allocated at a time is 4 MB, the memory of 1,024 continuous 4-KB page tables.  
Assume that you apply for memory that contains 256 page tables and whose order is 6. The system searches for the ninth, tenth, and eleventh linked lists in the array in sequence. If the previous linked list is empty, there is no free memory of this order. The system searches for the next linked list until the last linked list.  
If all linked lists are empty and the application fails, the kernel will report a page allocation failure and display the following error message to indicate that the memory page whose order is 6 fails to be requested.  

```
page allocation failure:order:6
```
- Triggering method  
Use `alloc_pages` to continuously apply for high-order memory pages (for example, `order=10`) and do not release the memory pages until the application fails.

## List Corruption

A list corruption error is reported when the kernel fails to check the validity of a linked list. There are two error types: `list_add` corruption and `list_del` corruption.

- Details  
The kernel provides `list_add` and `list_del` to check the validity of the linked list and to add or delete an entry from the linked list if it is valid. If the linked list is invalid, a list corruption error is reported. The kernel source code is `lib/list_debug.c`.

```
bool __list_add_valid(struct list_head *new, struct list_head *prev,
 struct list_head *next)
{
 if (CHECK_DATA_CORRUPTION(next->prev != prev,
 "list_add corruption. next->prev should be prev (%px), but was %px. (next=%px).\n",
 prev, next->prev, next) ||
 CHECK_DATA_CORRUPTION(prev->next != next,
 "list_add corruption. prev->next should be next (%px), but was %px. (prev=%px).\n",
 next, prev->next, prev) ||
 CHECK_DATA_CORRUPTION(new == prev || new == next,
 "list_add double add: new=%px, prev=%px, next=%px.\n",
 new, prev, next))
 return false;
 return true;
}
EXPORT_SYMBOL(__list_add_valid);

bool __list_del_entry_valid(struct list_head *entry)
{
 struct list_head *prev, *next;

 prev = entry->prev;
 next = entry->next;

 if (CHECK_DATA_CORRUPTION(next == LIST_POISON1,
 "list_del corruption. %px->next is LIST_POISON1 (%px).\n",
 entry, LIST_POISON1) ||
 CHECK_DATA_CORRUPTION(prev == LIST_POISON2,
 "list_del corruption. %px->prev is LIST_POISON2 (%px).\n",
 entry, LIST_POISON2) ||
 CHECK_DATA_CORRUPTION(prev->next != entry,
 "list_del corruption. prev->next should be %px, but was %px.\n",
 entry, prev->next) ||
 CHECK_DATA_CORRUPTION(next->prev != entry,
 "list_del corruption. next->prev should be %px, but was %px.\n",
 entry, next->prev))
 return false;
 return true;
}
EXPORT_SYMBOL(__list_del_entry_valid);
```

This error is typically caused by abnormal memory operations, such as memory corruption and memory damage.

- Triggering method  
Use the standard kernel interface of list.h to create a linked list, illegally modify the previous or next pointer of a linked list entry, and then call the kernel list\_add or list\_del interface.

## Bad mm\_struct

A bad mm\_struct error is reported when one or more mm\_struct data structures in the kernel are corrupted or damaged.

- Details  
mm\_struct is an important data structure in the Linux kernel. It is used to trace the virtual memory area of a process. If the data structure is damaged, the process or system may break down. This error is usually caused by memory exceptions. For example, the memory in mm\_struct is corrupted or memory overwriting occurs.
- Triggering method  
Bad mm\_struct is triggered when there is a hardware error or Linux kernel code error.

## I/O Error

An I/O error is reported when an input/output operation fails. This error may be printed when the driver of the I/O device such as the NIC or disk is abnormal or the file system is abnormal.

- Details  
The condition under which the code fails to be executed is the cause of this error. Common causes are hardware faults, disk damage, file system errors,

driver problems, and permission problems. For example, if an error occurs when the system attempts to read data from or write data to a disk, an I/O error is reported.

- Triggering method

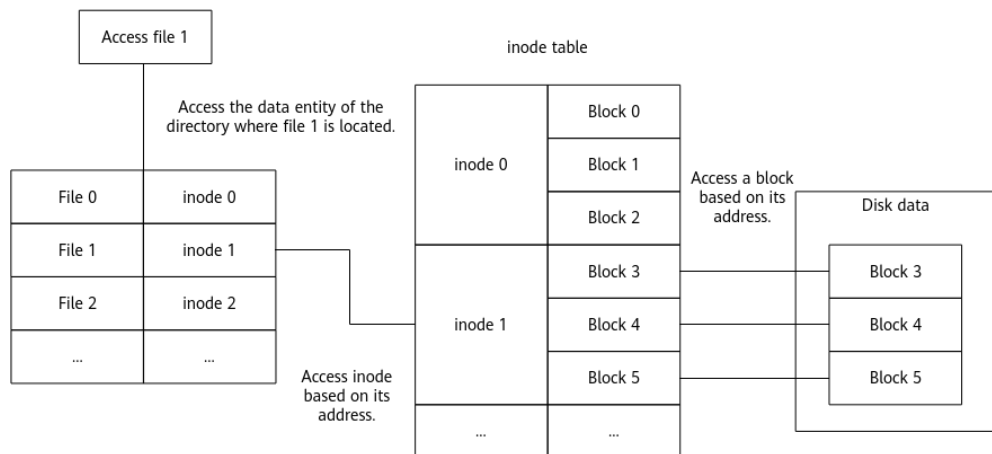
When the system is reading data from or writing data to the disk, remove the disk to damage the disk data.

## EXT4-fs Error

EXT4-fs errors typically indicate problems with the ext4 file system.

- Details

A sector is the minimum file storage unit on a storage device. Multiple consecutive sectors form a block. inode stores the metadata of a file, including the creator, creation date, file size, attributes, and the number of blocks. If the inode information in EXT4 format fails to be verified, an EXT4-fs error will be reported.



The kernel ext4 verification uses checksum to verify inode information. When there is a partition table error or the disk is damaged, the kernel returns the EIO (Input/Output Error) error code and the system reports "EXT4-fs error checksum invalid".

- Triggering method

Forcibly remove the disk and add it back to read the data.

## MCE

An MCE is a type of hardware error that occurs when a CPU detects a hardware problem. The interrupt number is 18, and the exception type is abort.

- Details

MCEs are caused by bus faults, memory ECC errors, cache errors, TLB errors, or internal clock errors. In addition to hardware faults, inappropriate BIOS configurations, firmware bugs, and software bugs may also cause MCEs.

When an MCE is reported, the OS checks a group of registers called Machine-Check MSR and executes the corresponding function based on the error codes of the registers. (The function varies depending on the chip architecture.)

- Triggering method

An MCE is reported when there is a bus fault, memory ECC error, cache error, TLB error, or internal clock error.

## Fatal Signal

If a signal cannot be ignored or handled through user-defined processing functions, it is a fatal signal, such as SIGKILL, SIGSTOP, and SIGILL.

- Details

The signal mechanism is an asynchronous notification mechanism for communication between processes in the system. When a signal is sent to a process but the OS interrupts the process, all non-atomic operations are interrupted.

If a signal is a SIGKILL, SIGSTOP, or SIGILL signal, it is a fatal signal.

```
#define sig_fatal(t, signr) \
 (!siginmask(signr, SIG_KERNEL_IGNORE_MASK|SIG_KERNEL_STOP_MASK) && \
```

- Triggering method

Use a user-mode program to execute invalid instructions or run **kill -9** to kill the process.

## Warning

Warning is an action taken to report a kernel issue that requires immediate attention when the OS is running. Warning prints the call stack information when the issue occurs. The OS continues to run after a warning.

- Details

Warning is triggered when macros such as WARN, WARN\_ON, and WARN\_ON\_ONCE are invoked.

There are several causes of invoking a warning macro. You need to trace the call stack to locate the cause. A warning macro does not change the system status and does not provide guidance for handling the warning.

- Triggering method

Construct a warning when the system is invoking a macro.

## Panic

A kernel panic refers to the action taken by the OS when it detects a fatal internal error and cannot securely handle the error. When an exception occurs during kernel running, the kernel uses the kernel\_panic function to print all the information obtained when the exception occurs.

- Details

There are various causes for the exception. Common causes include kernel stack overflow, division by zero, memory access out of bounds, and kernel deadlock. When this exception occurs, locate the cause of kernel\_panic based on the invoking information printed for the exception.

- Triggering method

Read address 0 in kernel mode.

# 11 xGPU

---

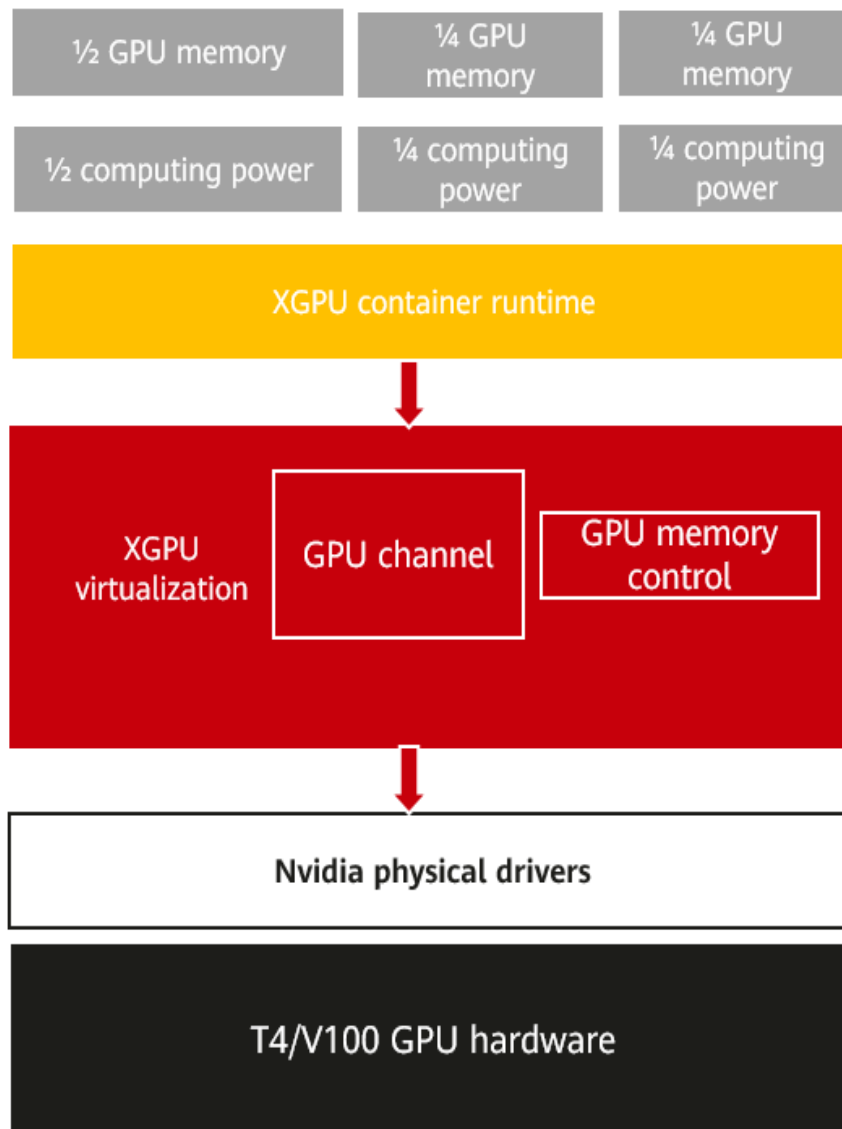
## 11.1 Overview

xGPU is a GPU virtualization and sharing service developed by Huawei Cloud. This service isolates GPU resources, ensuring service security and helping save costs by improving GPU resource utilization.

### Architecture

xGPU uses the in-house kernel drivers to provide vGPUs for containers. This service can isolate the GPU memory and compute while delivering high performance, ensuring that GPU hardware are fully used for training and inference. You can run commands to easily configure vGPUs in a container.


**Figure 11-1** xGPU architecture



## Why Using xGPU

- Lower cost  
As GPU technology develops, the price of a single GPU is rising though it can offer larger GPU compute. In some scenarios, an AI application does not require an entire GPU. xGPU enables multiple containers to share one GPU and isolates GPU resources to keep services securely isolated, improving GPU hardware utilization and reducing the resource cost.
- Flexible resource allocation  
xGPU allows you to flexibly allocate physical GPU resources based on your service requirements.
  - You can allocate resources by GPU memory or compute as required.





- You can isolate either of GPU memory or compute, and specify weights to allocate the GPU compute. The GPU compute is allocated at a granularity of 1%. It is recommended that the minimum GPU compute be greater than or equal to 4%.
- Robust compatibility  
xGPU is compatible with open source container technologies such as Docker, Containerd, and Kubernetes.
- Easy of use  
You do not need to recompile your AI applications and replace the Compute Unified Device Architecture (CUDA) library when your services are running.

## 11.2 Installing and Using xGPU

This section describes how to install and use xGPU.

### Constraints

- xGPU is supported only on NVIDIA Tesla T4 and V100.
- The Huawei Cloud EulerOS kernel version is 5.10 or later.
- The NVIDIA driver of version 535.54.03 has been installed on GPU-accelerated ECSs.
- Docker 18.09.0-300 or later has been installed on GPU-accelerated ECSs.
- GPU memory isolation provided by GPU virtualization cannot be used for unified virtual memory (UVM) through CUDA API `cudaMallocManaged()`. For more information, see [NVIDIA official document](#). Use other methods to apply for the GPU memory. For example, call `cudaMalloc()`.
- Limited by GPU virtualization, when applications in a container are initialized, the GPU compute monitored by NVIDIA System Management Interface (`nvidia-smi`) may exceed the upper limit of the GPU compute available for the container.
- When the first CUDA application is created, a percentage of the GPU memory (about 3 MiB on NVIDIA Tesla T4) is requested from each GPU. This GPU memory is shared by all CUDA applications. It is considered as management overhead and is not controlled by the xGPU service.
- During the initialization of the XGPU module, ensure that no other process on the node is using NVIDIA resources.

## Installing xGPU

To install xGPU, contact customer service.

## Using xGPU

The following table describes the environment variables of xGPU. When creating a container, you can specify the environment variables to configure the GPU compute and memory that a container engine can obtain using xGPU.

**Table 11-1** Environment variables that affect xGPU

| Environment Variable        | Value Type | Description                                                                                                                                                                                                                       | Example                                                                             |
|-----------------------------|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| GPU_IDX                     | Integer    | Specifies the GPU that can be used by a container.                                                                                                                                                                                | Assigning the first GPU to a container:<br>GPU_IDX=0                                |
| GPU_CONTAINER_MEM           | Integer    | Specifies the size of the GPU memory that will be used by a container, in MiB.                                                                                                                                                    | Allocating 5,120 MiB of the GPU memory to a container:<br>GPU_CONTAINER_MEM=5120    |
| GPU_CONTAINER_QUOTA_PERCENT | Integer    | Specifies the percentage of the GPU compute that will be allocated to a container from a GPU. The GPU compute is allocated at a granularity of 1%. It is recommended that the minimum GPU compute be greater than or equal to 4%. | Allocating 50% of the GPU compute to a container:<br>GPU_CONTAINER_QUOTA_PERCENT=50 |

| Environment Variable   | Value Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Example                                                                               |
|------------------------|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| GPU_POLICY             | Integer    | <p>Specifies the GPU compute isolation policy used by the GPU.</p> <ul style="list-style-type: none"> <li>● <b>0</b>: native scheduling. The GPU compute is not isolated.</li> <li>● <b>1</b>: fixed scheduling</li> <li>● <b>2</b>: average scheduling</li> <li>● <b>3</b>: preemptive scheduling</li> <li>● <b>4</b>: weight-based preemptive scheduling</li> <li>● <b>5</b>: hybrid scheduling</li> <li>● <b>6</b>: weighted scheduling</li> </ul> <p>For details, see <a href="#">GPU Compute Scheduling Examples</a>.</p> | <p>Setting the GPU compute isolation policy to fixed scheduling:<br/>GPU_POLICY=1</p> |
| GPU_CONTAINER_PRIORITY | Integer    | <p>Specifies the priority of a container.</p> <ul style="list-style-type: none"> <li>● <b>0</b>: low priority</li> <li>● <b>1</b>: high priority</li> </ul>                                                                                                                                                                                                                                                                                                                                                                    | <p>Creating a high-priority container:<br/>GPU_CONTAINER_PRIORITY=1</p>               |

The following example describes how you can use xGPU to allow two NVIDIA-Docker containers to share one GPU.

**Table 11-2** Data planning

| Parameter                   | Container 1 | Container 2 | Description                                                                                         |
|-----------------------------|-------------|-------------|-----------------------------------------------------------------------------------------------------|
| GPU_IDX                     | 0           | 0           | Two containers use the first GPU.                                                                   |
| GPU_CONTAINER_QUOTA_PERCENT | 50          | 30          | Allocate 50% of GPU compute to container 1 and 30% of GPU compute to container 2.                   |
| GPU_CONTAINER_MEM           | 5120        | 1024        | Allocate 5,120 MiB of the GPU memory to container 1 and 1,024 MiB of the GPU memory to container 2. |
| GPU_POLICY                  | 1           | 1           | Set the first GPU to use the fixed scheduling policy.                                               |

| Parameter              | Container 1 | Container 2 | Description                                                      |
|------------------------|-------------|-------------|------------------------------------------------------------------|
| GPU_CONTAINER_PRIORITY | 1           | 0           | Give container 1 a high priority and container 2 a low priority. |

Example:

```
docker run --rm -it --runtime=nvidia -e GPU_CONTAINER_QUOTA_PERCENT=50 -e GPU_CONTAINER_MEM=5120 -e GPU_IDX=0 -e GPU_POLICY=1 -e GPU_CONTAINER_PRIORITY=1 --shm-size 16g -v /mnt:/mnt nvcr.io/nvidia/tensorrt:19.07-py3 bash
docker run --rm -it --runtime=nvidia -e GPU_CONTAINER_QUOTA_PERCENT=30 -e GPU_CONTAINER_MEM=1024 -e GPU_IDX=0 -e GPU_POLICY=1 -e GPU_CONTAINER_PRIORITY=0 --shm-size 16g -v /mnt:/mnt nvcr.io/nvidia/tensorrt:19.07-py3 bash
```

## Viewing procfs Directory

At runtime, xGPU generates and manages multiple proc file systems (procfs) in the `/proc/cgpu_km` directory. The following are operations for you to view xGPU information and configure xGPU settings.

1. Run the following command to view the node information:

```
ls /proc/xgpu/
0 container version
```

The following table describes the information about the procfs nodes.

| Directory | Read/Write Type | Description                                                                                                                                                                                              |
|-----------|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 0         | Read and write  | xGPU generates a directory for each GPU on a GPU-accelerated ECS. Each directory uses a number as its name, for example, 0, 1, and 2. In this example, there is only one GPU, and the directory ID is 0. |
| container | Read and write  | xGPU generates a directory for each container running in the GPU-accelerated ECSs.                                                                                                                       |
| version   | Read-only       | xGPU version                                                                                                                                                                                             |

2. Run the following command to view the items in the directory of the first GPU:

```
ls /proc/xgpu/0/
max_inst meminfo policy quota utilization_line utilization_rate xgpu1 xgpu2
```

The following table describes the information about the GPU.

| Directory        | Read/Write Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| max_inst         | Read and write  | Specifies the maximum number of containers. The value ranges from <b>1</b> to <b>25</b> . This parameter can be modified only when no container is running.                                                                                                                                                                                                                                                                                                                                                                               |
| meminfo          | Read-only       | Specifies the total available GPU memory.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| policy           | Read and write  | Specifies the GPU compute isolation policy used by the GPU. The default value is <b>1</b> . <ul style="list-style-type: none"><li>● <b>0</b>: native scheduling. The GPU compute is not isolated.</li><li>● <b>1</b>: fixed scheduling</li><li>● <b>2</b>: average scheduling</li><li>● <b>3</b>: preemptive scheduling</li><li>● <b>4</b>: weight-based preemptive scheduling</li><li>● <b>5</b>: hybrid scheduling</li><li>● <b>6</b>: weighted scheduling</li></ul> For details, see <a href="#">GPU Compute Scheduling Examples</a> . |
| quota            | Read-only       | Specifies the total weight of the GPU compute.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| utilization_line | Read and write  | Specifies the GPU usage threshold for online containers to suppress offline containers. If the GPU usage exceeds the value of this parameter, online containers completely suppresses offline containers. If the GPU usage does not exceed the value of this parameter, online containers partially suppresses offline containers.                                                                                                                                                                                                        |
| utilization_rate | Read only       | Specifies the GPU usage.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| xgpuIndex        | Read and write  | Specifies the <b>xgpu</b> subdirectory of the GPU. In this example, <b>xgpu1</b> and <b>xgpu2</b> belong to the first GPU.                                                                                                                                                                                                                                                                                                                                                                                                                |

3. Run the following command to view the directory of a container:

```
ls /proc/xgpu/container/
9418 9582
```

The following table describes the content in the directory.

| Directory   | Read/Write Type | Description                                                                          |
|-------------|-----------------|--------------------------------------------------------------------------------------|
| containerID | Read and write  | Specifies the container ID. xGPU generates an ID for each container during creation. |

4. Run the following commands to view the directory of each container:

```
ls /proc/xgpu/container/9418/
xgpu1
ls /proc/xgpu/container/9582/
xgpu2
```

The following table describes the content in the directory.

| Directory | Read/Write Type | Description                                                                                                                                                                           |
|-----------|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| xgpuIndex | Read and write  | Specifies the <b>xgpu</b> subdirectory of the container. In this example, <b>xgpu1</b> is the subdirectory of container 9418, and <b>xgpu2</b> is the subdirectory of container 9582. |

5. Run the following command to view the xgpuIndex directory:

```
ls /proc/xgpu/container/9418/xgpu1/
meminfo priority quota
```

The following table describes the content in the directory.

| Directory | Read/Write Type | Description                                                                                                                                                                                                                                                                                                                                                   |
|-----------|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| meminfo   | Read-only       | Specifies the GPU memory allocated using xGPU and the remaining GPU memory. For example, <b>3308MiB/5120MiB, 64% free</b> indicates that 5,120 MiB of memory is allocated and 64% is available.                                                                                                                                                               |
| priority  | Read and write  | Specifies the priority of a container. The default value is <b>0</b> . <ul style="list-style-type: none"><li><b>0</b>: low priority</li><li><b>1</b>: high priority</li></ul> This parameter is used in hybrid scenarios where there are both online and offline containers. High-priority containers can preempt the GPU compute of low-priority containers. |

| Directory | Read/Write Type | Description                                                                                                                                                |
|-----------|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| quota     | Read-only       | Specifies the percentage of the GPU compute allocated using xGPU.<br>For example, <b>50</b> indicates that 50% of the GPU compute is allocated using xGPU. |

You can run the following commands to perform operations on GPU-accelerated ECSs. For example, you can change the scheduling policy and modify the weight.

| Command                                                              | Description                                                                   |
|----------------------------------------------------------------------|-------------------------------------------------------------------------------|
| echo 1 > /proc/xgpu/0/policy                                         | Changes the scheduling policy to weight-based preemptive scheduling.          |
| cat /proc/xgpu/container/\${containerID}/\${xgpuIndex}/meminfo       | Queries the memory allocated to a container.                                  |
| cat /proc/xgpu/container/\${containerID}/\${xgpuIndex}/quota         | Queries the GPU compute weight specified for a container.                     |
| cat /proc/xgpu/0/quota                                               | Queries the weight of the remaining GPU compute on the first GPU.             |
| echo 20 > /proc/xgpu/0/max_inst                                      | Sets the maximum number of containers allowed on the first GPU to <b>20</b> . |
| echo 1 > /proc/xgpu/container/\${containerID}/\${xgpuIndex}/priority | Sets the xGPU in a container to high priority.                                |
| echo 40 > /proc/xgpu/0/utilization_line                              | Sets the suppression threshold of the first GPU to 40%.                       |

## Upgrading xGPU

You can perform cold upgrades on xGPU.

- Run the following command to stop all running containers:  
**docker ps -q | xargs -I {} docker stop {}**
- Run the following command to upgrade the RPM package of xGPU:  
**rpm -U hce\_xgpu**

## Uninstalling xGPU

- Run the following command to stop all running containers:  
**docker ps -q | xargs -I {} docker stop {}**

2. Run the following command to uninstall the RPM package:  
`rpm -e hce_xgpu`

## Monitoring Containers Using xgpu-smi

You can use `xgpu-smi` to view xGPU information, including the container ID, GPU compute usage and allocation, and GPU memory usage and allocation.

The following shows the `xgpu-smi` monitoring information.

```
=====
| HUAWEI CLOUD XGPU-SMI | XGPU Version: 1.0 |
=====
| Container-Id | GPU | GPU-Util/Limit | GPU-Memory-Usage/Limit |
=====
| 800a09ae74bc | 1 | 0% / 30% | 0M / 2000M |
| b8402c9316e4 | 0 | 0% / 10% | 0M / 6000M |
=====
```

You can run the `xgpu-smi -h` command to view the help information about the `xgpu-smi` tool.

```
[root@localhost ~]#
[root@localhost ~]# xgpu-smi -h
xgpu-smi provides monitoring information about xgpu containers.
The data is presented in either a plain text or a json format, via stdout or a file.

Usage:
xgpu-smi [flags]
xgpu-smi [command]

Available Commands:
clean Release XGPU
list List XGPU containers and used GPUs
query Show XGPU containers' information and GPUs' available resource information in JSON format

Flags:
-c, --container string filter the specific containers' information
-f, --filename string log to a specified file, rather than to stdout
-g, --gpu string filter the specific GPUs' information
-h, --help help for xgpu-smi
-r, --remote-runtime-endpoint string endpoint of remote runtime endpoint (default "unix:///var/run/containerd/containerd.sock")

Use "xgpu-smi [command] --help" for more information about a command.
[root@localhost ~]#
[root@localhost ~]#
```

## 11.3 GPU Compute Scheduling Examples

When you want to create an xGPU device, the xGPU service sets time slices ( $X$  ms) for each GPU based on the maximum number of containers (`max_inst`) to allocate the GPU compute to containers. The time slices are represented by processing unit 1, processing unit 2, ..., processing unit  $N$ . The following describes different scheduling policies, and the maximum number of containers is 20 (`max_inst=20`).

### Native Scheduling (policy=0)


Native scheduling indicates that the GPU compute scheduling method of NVIDIA GPUs. In the native scheduling policy, xGPU is used only for GPU memory isolation.






## Fixed Scheduling (policy=1)

Fixed scheduling indicates that the GPU compute is allocated to containers based on a fixed percentage. For example, 5% of the GPU compute is allocated to container 1, and 15% of the GPU compute is allocated to container 2, as shown in the figure.



## Average Scheduling (policy=2)


Average scheduling indicates that each container can have the same percentage of GPU compute ( $1/\text{max\_inst}$ ). For example, if **max\_inst** is set to 20, each container obtains 5% of the GPU compute, as shown in the following figure.



## Preemptive Scheduling (policy=3)

Preemptive scheduling indicates that each container obtains one time slice, and xGPU starts scheduling from processing unit 1. However, if a processing unit is not allocated to a container or no process in the container is using the GPU, the processing unit will be skipped, and scheduling will start from the next slice. Processing units in gray are skipped and do not participate in scheduling.


In this example, containers 1, 2, and 3 each occupy 33.33% of the GPU compute.



## Weight-based Preemptive Scheduling (policy=4)

Weight-based preemptive scheduling indicates that time slices are allocated to containers based on the GPU compute percentage of each container. xGPU starts scheduling from processing unit 1. However, if a processing unit is not allocated to a container, the processing unit will be skipped, and scheduling will start from the next slice. For example, if 5%, 5%, and 10% of the GPU compute is allocated to containers 1, 2, and 3, respectively, container 1 and container 2 each occupy 1 processing unit, and container 3 occupies 2 processing units. Processing units in gray are skipped and do not participate in scheduling.

In this example, containers 1, 2, and 3 occupy 25%, 25%, and 50% of the GPU compute, respectively.

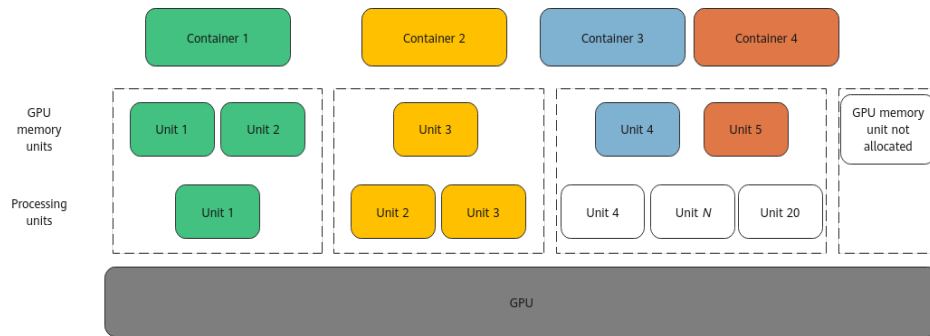


## Hybrid Scheduling (policy=5)

Hybrid scheduling indicates that a single GPU supports the isolation of the GPU memory and the isolation of both the GPU compute and GPU memory. The isolation of both the GPU compute and GPU memory is the same as that of the fixed scheduling (policy=1). The containers with only GPU memory isolated share the remaining GPU compute after the GPU compute is allocated to the containers with both GPU compute and memory isolated. If `max_inst` is set to `20`, containers 1 and 2 have both GPU compute and memory isolated, and 5% of the GPU compute is allocated to container 1 and 10% of the GPU compute is allocated to container 2. Containers 3 and 4 have only the GPU memory isolated. Because container 1 occupies one processing unit, container 2 occupies two processing units, containers 3 and 4 share the remaining 17 processing units. In addition, if no processes in container 2 are using the GPU, container 1 occupies one processing unit, container 2 occupies 0 processing units, and containers 3 and 4 share the remaining 19 processing units.

In hybrid scheduling, whether GPU compute isolation is enabled for containers is determined by whether `GPU_CONTAINER_QUOTA_PERCENT` is set to `0`. All

containers whose `GPU_CONTAINER_QUOTA_PERCENT` is 0 share the idle GPU compute.




**NOTE**

The hybrid scheduling policy is not available for high-priority containers.

### Weighted Scheduling (policy=6)

Weighted scheduling indicates that time slices are allocated to containers based on the percentage of the CPU compute of each container. Its compute isolation is not as good as that in weight-based preemptive scheduling. xGPU starts scheduling from processing unit 1. However, if a processing unit is not allocated to a container or no process in the container is using the GPU, the processing unit will be skipped, and scheduling will start from the next slice. For example, if 5%, 5%, and 10% of the GPU compute is allocated to containers 1, 2, and 3, respectively, container 1 and container 2 each occupy 1 processing unit, and container 3 occupies 2 processing units. In the following figure, processing units in white are idle and for container 3, and processing units in both white and gray are skipped and do not participate in scheduling.

In this example, containers 1, 2, and 3 occupy 50%, 50%, and 0% of the GPU compute, respectively.



**NOTE**

Weighted scheduling involves preemption of idle compute. When a container switches between idle and busy statuses, the compute of other containers is affected, and there will be compute fluctuation. When a container switches from the idle state to the busy state, the latency for the container to preempt the compute does not exceed 100 ms.

# 12 Configuring Repo Sources and Installing Software for Huawei Cloud EulerOS

Huawei Cloud EulerOS uses RPM packages to manage software and provides the official Repo source matching the OS to release software packages and update packages. You can run the **dnf/yum** commands for common software management functions, such as installation, upgrade, and uninstallation.

## Configuring the Official Repo Source

By default, the official Repo source is configured in the `/etc/yum.repos.d/hce.repo` file of the default Huawei Cloud EulerOS image purchased through the ECS. Take Huawei Cloud EulerOS 2.0 as an example. The file content is as follows:

```
[base]
name=HCE $releasever base
baseurl=https://repo.huaweicloud.com/hce/$releasever/os/$basearch/
enabled=1
gpgcheck=1
gpgkey=https://repo.huaweicloud.com/hce/$releasever/os/RPM-GPG-KEY-HCE-2

[updates]
name=HCE $releasever updates
baseurl=https://repo.huaweicloud.com/hce/$releasever/updates/$basearch/
enabled=1
gpgcheck=1
gpgkey=https://repo.huaweicloud.com/hce/$releasever/updates/RPM-GPG-KEY-HCE-2

[debuginfo]
name=HCE $releasever debuginfo
baseurl=https://repo.huaweicloud.com/hce/$releasever/ debuginfo/$basearch/
enabled=0
gpgcheck=1
gpgkey=https://repo.huaweicloud.com/hce/$releasever/ debuginfo/RPM-GPG-KEY-HCE-2
```

The fields are described as follows:

- **name**: name of the Repo source.
- **baseurl**: address of the server where the software repository is located. The value can be in the format of **http://**, **ftp://**, or **file://**.

- **enabled**: whether to enable the software repository. The value **1** indicates that the software repository is enabled, and the value **0** indicates that the software repository is disabled.
- **gpgcheck**: whether to perform GNU Privacy Guard (GPG) verification. The value **1** indicates that GPG verification is enabled, and the value **0** indicates that GPG verification is disabled.
- **gpgkey**: address for storing the public key that is used for GPG verification.

---

**⚠ CAUTION**

Modifying this file may affect the software installation and upgrade of the OS. You are advised not to modify this file.

---

## Configuring a Third-Party Repo Source

To add a third-party repo source, perform the following steps (the image source from the OpenEuler community is used as an example):

1. Add the **openEuler.repo** file to the **/etc/yum.repos.d/** directory. The file name can be customized, but the file name extension must be **.repo**. Run the **vim /etc/yum.repos.d/openEuler.repo** command to edit the file.
2. Set the repository name, for example, **[openEuler]**. The repository name must be unique and can be changed based on site requirements.
3. Configure the **name** field. For example, set it to **openEuler repository**, which is the detailed description of the repository. You can change the description based on site requirements.
4. Set **baseurl** to **https://repo.openeuler.org/openEuler-22.03-LTS/OS/x86\_64/**, which is the link where the software package is obtained. For details, see the product documentation of OpenEuler or the Repo provider.
5. Configure the **gpgcheck** field. The value **1** indicates that GPG verification is performed on the installed software package.
6. Configure the **enabled** field. The value **1** indicates that the Repo source is enabled.
7. Set **gpgkey** to **https://repo.openeuler.org/openEuler-22.03-LTS/OS/x86\_64/RPM-GPG-KEY-openEuler**, which is the link where the public key used for GPG verification is obtained.

The final **openEuler.repo** file is as follows:

```
[openEuler]
name=openEuler repository
baseurl=https://repo.openeuler.org/openEuler-22.03-LTS/OS/x86_64/
gpgcheck=1
enabled=1
gpgkey=https://repo.openeuler.org/openEuler-22.03-LTS/OS/x86_64/RPM-GPG-KEY-openEuler
```

**NOTICE**

You can set the **priority** field to specify the priority of each Repo source. If the Huawei Cloud EulerOS official source is preferentially used, add **priority=1** to the **hce.repo** configuration (a smaller value indicates a higher priority), and then add **priority=2** to the third-party source configuration. You can adjust the value based on site requirements. For details about the complete openEuler repository configuration, see [Obtaining the openEuler Extended Software Packages](#).

 **NOTE**

To upgrade the software package, refer to [Upgrading Huawei Cloud EulerOS and RPM Packages](#).

## Common yum and dnf Commands

In Huawei Cloud EulerOS 1.1, you can only perform software management operations by running the **yum** commands. In Huawei Cloud EulerOS 2.0, both the **yum** and **dnf** commands are supported. The following are common commands for software management.

| Function                                     | yum Command                       | dnf Command                       | Example Command                                                  |
|----------------------------------------------|-----------------------------------|-----------------------------------|------------------------------------------------------------------|
| Installing software packages                 | yum install<br><Software package> | dnf install<br><Software package> | yum install gcc: installs GCC.                                   |
| Uninstalling software packages               | yum remove<br><Software package>  | dnf remove<br><Software package>  | yum remove gcc: uninstalls GCC.                                  |
| Listing installed software packages          | yum list<br>installed             | dnf list<br>installed             | yum list installed: lists all software packages in the system.   |
| Searching for software packages              | yum search<br><Software package>  | dnf search<br><Software package>  | yum search gcc: searches for the GCC package in the Repo source. |
| Querying information about software packages | yum info<br><Software package>    | dnf info<br><Software package>    | yum info gcc: queries GCC software package information.          |

# 13 Change History

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2024-03-30  | <p>This issue is the eighth official release.</p> <p>Added the following:</p> <ul style="list-style-type: none"><li><a href="#">Configuring Repo Sources and Installing Software for Huawei Cloud EulerOS</a></li></ul> <p>Modified the following:</p> <ul style="list-style-type: none"><li>Added backup description in <a href="#">Migrating an OS to Huawei Cloud EulerOS</a>.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                       |
| 2023-12-30  | <p>This issue is the seventh official release.</p> <p>Modified the following:</p> <ul style="list-style-type: none"><li>Added the content of weighted scheduling (policy=6) in <a href="#">GPU Compute Scheduling Examples</a>.</li><li>Added the commands used on CLI for dynamic acceleration in <a href="#">Dynamic Acceleration</a>.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                |
| 2023-11-17  | <p>This issue is the sixth official release.</p> <p>Modified the following:</p> <ul style="list-style-type: none"><li>Modified the constraints and optimized the procedure in <a href="#">Creating a Docker Image and Starting a Container</a>.</li><li>Added examples of software compatibility scanning on the Arm architecture in <a href="#">Evaluating Software Compatibility</a>.</li><li>Updated the version of the migration tool installation package and modified the OS migration description in <a href="#">Installing and Using xGPU</a>.</li><li>Modified the constraints in <a href="#">Constraints</a>.</li><li>Modified the constraints in <a href="#">Constraints</a>.</li><li>Modified the constraints in <a href="#">Installing and Using xGPU</a>.</li></ul> |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2023-09-27  | <p>This issue is the fifth official release.</p> <p>Added the following:</p> <ul style="list-style-type: none"> <li>● <a href="#">Kernel Exception Events Analysis</a></li> <li>● <a href="#">Dynamic Acceleration</a></li> </ul> <p>Modified the following:</p> <ul style="list-style-type: none"> <li>● Supported the .ko format.</li> <li>● Added support for scanning nested directories in occupiers.</li> <li>● Added compatibility evaluation for EulerOS 2.9/2.10 and Huawei Cloud EulerOS 2.0 of Arm servers in <a href="#">Overview of x2hce-ca</a>.</li> <li>● Updated the conflict package list in <a href="#">Conflicting Packages</a>.</li> <li>● Added constraints in <a href="#">Constraints</a>.</li> <li>● Added the introduction to dynamic acceleration and the advantages and disadvantages of static acceleration and dynamic acceleration in <a href="#">Overview</a>.</li> <li>● Added the description of the dynamic acceleration configuration file and parameters in <a href="#">Configuration File</a>.</li> <li>● Added descriptions about hybrid scheduling <a href="#">Installing and Using xGPU</a> and <a href="#">GPU Compute Scheduling Examples</a>.</li> </ul> |



| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2023-08-03  | <p>This issue is the fourth official release.</p> <p>Modified the following:</p> <ul style="list-style-type: none"><li>• Added the <b>motd_setup</b> parameter in <a href="#">Preparations</a>.</li><li>• Added the <b>--nocheck</b> parameter in <a href="#">OSMT Command Help Information</a>.</li><li>• Updated operations in <a href="#">Obtaining the openEuler Extended Software Packages</a>.</li><li>• Added the <b>hce-wae-auto</b> command for optimizing applications in <a href="#">Overview</a> and <a href="#">Static Acceleration</a>.</li><li>• Optimized the GPU compute allocation granularity from 5% to 1%, modified the corresponding description, and added the restriction that the xGPU function is supported only on the Nvidia Tesla T4, V100, A100, A30, A40 and A800 in <a href="#">Overview</a>.</li><li>• Added the GPU_POLICY, GPU_CONTAINER_PRIORITY, max_inst and priority parameters and related content, optimized the GPU compute isolation policy from 0-1 to 0-4; added the content of monitoring containers using the xgpu-smi tool in <a href="#">Installing and Using xGPU</a>.</li></ul> <p>Added the following:</p> <ul style="list-style-type: none"><li>• <a href="#">Creating a Docker Image and Starting a Container</a></li><li>• <a href="#">Configuration File</a></li><li>• <a href="#">GPU Compute Scheduling Examples</a></li></ul> |
| 2023-03-30  | <p>This issue is the third official release.</p> <p>Added the following:</p> <ul style="list-style-type: none"><li>• <a href="#">Pod Bandwidth Management Tool</a></li><li>• <a href="#">Kernel Functions and Interfaces</a></li><li>• <a href="#">OOM Process Control Policy</a></li><li>• <a href="#">Multi-level Memory Reclamation Policy</a></li><li>• <a href="#">Multi-level Hybrid Scheduling of Kernel CPU cgroups</a></li></ul> <p>Modified the following:</p> <ul style="list-style-type: none"><li>• Added the conflict package list of CentOS 8.4 and 8.5 in <a href="#">Conflicting Packages</a>.</li><li>• Updated operations in <a href="#">Static Acceleration</a>.</li></ul> <p>Deleted the following:</p> <ul style="list-style-type: none"><li>• <i>Common Issues</i> in the <i>Workload Accelerator</i></li><li>• <i>perf</i> in the <i>Workload Accelerator</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2023-01-17  | <p>This issue is the second official release.</p> <p>Added the following:</p> <ul style="list-style-type: none"> <li>● <a href="#">Constraints</a></li> <li>● <a href="#">Follow-up Operations</a></li> <li>● <a href="#">Tools</a></li> </ul> <p>Modified the following:</p> <ul style="list-style-type: none"> <li>● Added the CentOS 8.3 conflict package list in <a href="#">Conflicting Packages</a>.</li> <li>● Updated the dnf list updates command and added step 3 in <a href="#">Using dnf or yum to Upgrade</a>.</li> <li>● Added constraints in <a href="#">Constraints</a>.</li> <li>● Optimized operations in <a href="#">Version Upgrade and Rollback</a>.</li> <li>● Updated the osmt.conf configuration and added parameters and descriptions in <a href="#">Preparations</a>.</li> <li>● Optimized operations in <a href="#">Manual Update Using osmt update</a>.</li> <li>● Optimized operations in <a href="#">Automatic Update Using osmt-agent</a>.</li> <li>● Added the rollback description in <a href="#">Rolling Back RPM Packages</a>.</li> </ul> |
| 2022-11-30  | <p>This issue is the first official release.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |