Cloud Backup and Recovery

FAQs

Issue	02
Date	2020-04-08

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions

Notice

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

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

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process*. For details about this process, visit the following web page:
https://www.huawei.com/en/psirt/vul-response-process
For vulnerability information, enterprise customers can visit the following web page:
https://securitybulletin.huawei.com/enterprise/en/security-advisory

Contents

1 Concepts

1.1 What Are Full Backup and Incremental Backup?

Description

A full backup backs up all data at a certain time point.

An incremental backup backs up the changed data since the last full or incremental backup.

CBR uses the permanent incremental backup technology. A full backup is performed for a resource in the initial backup and incremental backups in all subsequent backups. If a full backup expires and is deleted, its next incremental backup will be regarded as the resource's full backup.

Suppose that server **X** has backups **A**, **B**, and **C** in time sequence. Backup A is a full backup, and backups B and C are incremental backups. Only changed data blocks are backed up in incremental backups and unchanged data blocks are indexed using pointers, so each incremental backup can be regarded as a virtual full backup.

If backup A is deleted, data blocks in backup A indexed by subsequent backups will not be deleted. Only data blocks that are exclusive to backup A are deleted. So, backups B and C can still be used to restore data. Or if backups A and B are deleted, backup C can also be used to restore data independently. There is no obvious difference between their restoration speeds.

Differences

- Backup duration: A full backup backs up the entire resource data which is usually larger than an incremental backup, so full backup take a longer time.
- Restoration duration: Both full backups and incremental backups can be restored. There is no obvious difference between their restoration speeds.
- Reliability: The latest incremental backup depends on the last full backup and intermediate incremental backups. If any backup data block is damaged, subsequent backups may be affected, which will reduce the backup reliability. All full backup data is independent and does not depend on previous backups. So full backups are more reliable.

You are advised to configure periodic full backup (for example, once every 30 days) and daily incremental backup to reduce the interval of full backup on which incremental backup depends and improve the reliability of backups.

NOTE

In extreme cases, the size of a backup is the same as the disk size. The used capacity in a full backup and the changed capacity in an incremental backup are calculated based on the data block change in a disk rather than the file change in the operating system. The size of a full backup cannot be evaluated based on the file capacity in the operating system, and the size of an incremental backup cannot be evaluated based on the file size change.

1.2 What Are the Differences Between Backup and Disaster Recovery?

The following table lists the main differences between backup and disaster recovery (DR).

Table 1-1 Differences between backup and DR

Item	Backup	DR
Purpose	To prevent data loss. It adopts the snapshot or backup techniques to generate data backups that can be used to restore data when data loss or corruption occurs.	To ensure service continuity. It takes the replication techniques (such as application-layer replication, host-based replication at the I/O layer, and storage-layer replication) to construct standby service hosts and data in a remote center, so that the remote center can take over services whenever the primary center is faulty.
Scenario	It offers protection against virus attacks, accidental deletions, software and hardware faults.	It enables failover upon software and hardware faults, as well as natural disasters, such as tsunami, fires, and earthquakes, to fast recover services. When the source AZ recovers, you can easily fail back to the source AZ.
Cost	The cost is 1 to 2% of the production system's cost.	The cost is 20 to 100% of the production system's, varying with the RPO/RTO requirements. For active-active DR, the service system deployed in the standby center is required to be the same as that in the active system. In this case, the cost on infrastructure doubles.

NOTE

Recovery Point Objective (RPO) specifies the maximum acceptable period in which data can be lost.

Recovery Time Objective (RTO) specifies the maximum acceptable amount of time for restoring the entire system after a disaster occurs.

1.3 What Are the Differences Between Backups and Snapshots?

Both backups and snapshots provide data redundancy for disks to improve data reliability. **Table 1-2** lists the differences between them.

Table 1-2 Differences between backups and snapshots

Item	Storage Solution	Data Synchronization	Service Recovery
Backup	Backup data is stored in OBS, instead of disks. This ensures data restoration upon disk data loss or corruption.	A backup is the data copy of a disk at a given point in time. CBR supports automatic backup by configuring backup policies. Deleting a disk will not clear its backups.	You can restore backups to their original disks or create new disks from the backups.
Snapshot	Snapshot data is stored with disk data. **NOTE** Creating a backup requires a certain amount of time because data needs to be transferred. Therefore, creating or rolling back a snapshot consumes less time than creating a backup.	A snapshot is the state of a disk at a specific point in time. If a disk is deleted, all the snapshots created for this disk will also be deleted. If you have reinstalled or changed the server OS, snapshots of the system disk are automatically deleted. Snapshots of the data disks can be used as usual.	You can use a snapshot to roll back its original disk or create a disk for data restoration and service recovery.

1.4 What Are the Differences Between Backups and Images?

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

Differences Between Backups and Images

Table 1-3 lists the differences between them.

Table 1-3 Differences between backups and images

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

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

NOTE

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

Relationship Between Backups and Images

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

1.5 What Are the Differences Between Cloud Server Backup and Cloud Disk Backup?

Table 1-4 describes the differences between cloud server backup and cloud disk backup.

Table 1-4 Differences between cloud server backup and cloud disk backup

Item	Cloud Server Backup	Cloud Disk Backup
Resources to be backed up or restored	All disks (system and data disks) on a server	One or more specified disks (system or data disks)
Recommended scenario	An entire cloud server needs to be protected.	Only data disks need to be backed up, because the system disk does not contain users' application data.
Advantages	All disks on a server are backed up at the same time, ensuring data consistency.	Backup cost is reduced without compromising data security.

2 Billing

2.1 How Is CBR Billed?

Billing Items

You are billed for the storage space and the data traffic generated if backup replication is used. Pricing of the storage space varies with vault types. See details in the following table.

Category	Billing Item	Description	Billing Mode
Storage space	Disk backup vault	If cloud disks need to be backed up, buy disk backup vaults to store the backups.	Pay-per-use Yearly/Monthly
	Server backup vault	If cloud servers (without applications) need to be backed up, buy server backup vaults to store the backups.	Pay-per-use Yearly/Monthly
	SFS Turbo backup vault	If SFS Turbo file systems need to be backed up, buy SFS Turbo backup vaults to store the backups.	Pay-per-use Yearly/Monthly

Category	Billing Item	Description	Billing Mode
	Database server backup vault	If cloud servers (with applications) need to be backed up, buy database server backup vaults to store the backups. You need to enable **Application-Consistent Backup** on the **Buy Server Backup Vault** page before using database server backup vaults. For more information, see **Application-Consistent Backup Overview**.	Pay-per-use Yearly/ Monthly
	Hybrid cloud backup vault	If backups of on-premises VMware VMs and OceanStor Dorado arrays need to be stored, buy hybrid cloud backup vaults.	Pay-per-use Yearly/ Monthly
	Replication vault	If you need to replicate backups to another region, buy replication vaults in the destination region.	Pay-per-use Yearly/ Monthly
Data traffic	Outbound traffic over the Internet	If hybrid cloud backups on the cloud are used to restore data to on-premises IDCs, outbound traffic is charged.	Free for a limited time
	Cross-region replication traffic	If backups or vaults are replicated to another region, traffic for cross-region replication is charged for the source region.	Pay-per-use

Billing Examples

Example 1

Purchase a pay-per-use vault for cloud servers without databases deployed:

If a user purchases a 400-GB server backup vault for their 100-GB cloud server in the LA-Mexico City1 region, the user is billed for the 400-GB server backup vault in CBR.

Example 2

Purchase a pay-per-use vault for cloud servers with databases deployed:

If a user purchases an 800-GB database server backup vault for their 100-GB database server in the LA-Mexico City1 region, the user is billed for the 800-GB database server backup vault in CBR.

Example 3

Replicate a backup to another region:

A user purchases a 100 GB server backup vault A in the LA-Mexico City1 region, and the backup data uses 40 GB of the storage space. This user also purchases a 200 GB replication vault B in the AP-Bangkok region and replicates data from vault A to vault B, without using the acceleration service. In this case, the user is billed for the 100 GB backup vault and the 200 GB replication vault, as well as the 40 GB cross-region replication data traffic.

2.2 How Large of a Vault Do I Need?

Manual Backup Scenario

If only manual backup is required, you are advised to set the vault capacity to at least twice the total capacity of the resources you want to back up.

Automatic Backup Scenario

If automatic backup is required, you are advised to set the vault capacity to twice the total capacity of the resources you want to back up. Alternatively, you can use the following formula to estimate the capacity of the vault to be created:

1. **Prepare the following data:**

Disk capacity (GB): a

Backup retention period (days): b

Daily changed data volume (GB): c

2. **Calculate the vault capacity using the following formula:**

Vault capacity (GB) = (a + b x c) x 120%

NOTE

- If you configure to keep backups by quantity, you can convert the number of retained backups to the backup retention period and then use the preceding formula. For example, if a company backs up data once a day and configures to retain seven backups, the retention period can be considered as seven days.
- Deleted, added, and changed data is all included when the daily changed data volume is calculated.

You can adjust the vault capacity based on your service needs after calculation.

If you only need to back up some of your files, use **file backup** to reduce the backup space and lower the backup costs.

Example

A financial company has an 800 GB cloud server, has used 200 GB of it, and its daily data changes are about 10 GB. As scheduled, the company's data is backed up twice at 02:00 and 20:00 every day, and backups are retained for a month. You can calculate the capacity of the server backup vault as follows: (Note that the vault capacity, not the used capacity, is used in the calculation.)

Vault capacity = (800 + 30 x 10) x 120% = 1,320 GB

2.3 What Is the Billing Cycle of Pay-per-Use Vaults?

Pay-per-use vaults are billed by the hour, and payments are made once a day after use. For example, if you purchased a pay-per-use vault at sometime between 18:00 to 19:00, the usage period of an hour would be billed (18:00 to 19:00).

2.4 How Do I Disable CBR?

If you have enabled CBR when purchasing an ECS but want to disable it afterward, go to the CBR console and then delete all vaults on the cloud server backup, cloud disk backup, SFS Turbo backup, and hybrid cloud backup pages. See **Figure 2-1**.

- If a message is displayed indicating that the ECS backup cannot be deleted, check whether the backup has been used to create an image and whether the image has been deleted.
- If you have not migrated resources to CBR, switch back to the CSBS or VBS console page to delete the backups.
- To delete CSBS backups displayed on the VBS console page, go to the **Backups** tab page on CSBS Console.

Figure 2-1 Deleting vaults

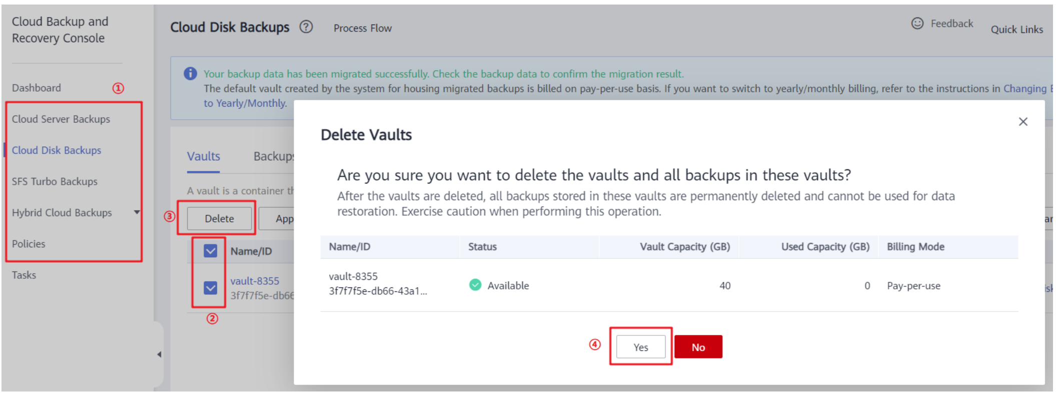

Your backup data has been migrated successfully. Check the backup data to confirm the migration result.
The default vault created by the system for housing migrated backups is billed on pay-per-use basis. If you want to switch to yearly/monthly billing, refer to the instructions in Changing B to Yearly/Monthly.

Are you sure you want to delete the vaults and all backups in these vaults?
After the vaults are deleted, all backups stored in these vaults are permanently deleted and cannot be used for data restoration. Exercise caution when performing this operation.

Name/ID	Status	Vault Capacity (GB)	Used Capacity (GB)	Billing Mode
vault-8355 3f7f7f5e-db66-43a1...	Available	40	0	Pay-per-use

2.5 How Do I Unsubscribe from CSBS or VBS Resource Packages?

If you want to migrate resources from CSBS and VBS to CBR, the original yearly/monthly CSBS and VBS packages cannot be migrated. In this case, you can unsubscribe from the packages yourself or **submit a service ticket**. For details about the unsubscription rules, see **Unsubscription Rules**.

Unsubscribing from a Resource Package By Self Service

Step 1 Log in to CSBS Console.

Step 2 Click **Billing** in the upper right corner of the page to go to the Billing Center.

Step 3 In the left navigation pane, choose **Orders** > **Unsubscriptions** to view the resources you purchased.

Step 4 Select an unwanted resource package and click **Unsubscribe**. See **Figure 2-2**.

Figure 2-2 Unsubscribing from a resource package

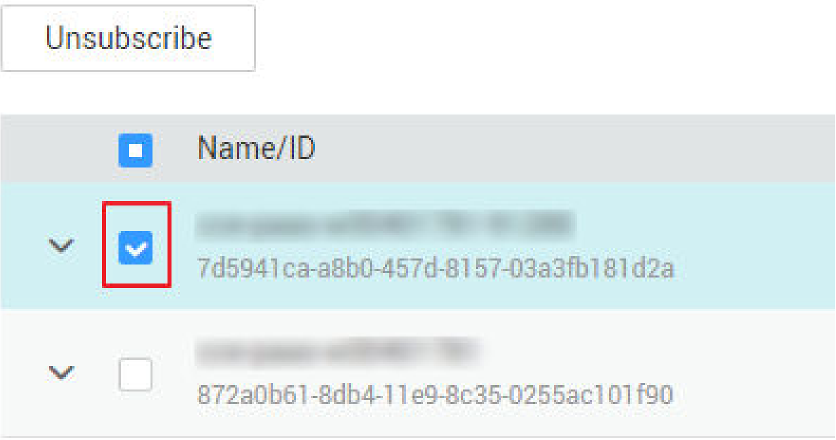
Unsubscribe
■ Name/ID
☑ [redacted]
7d5941ca-a8b0-457d-8157-03a3fb181d2a
□ [redacted]
872a0b61-8db4-11e9-8c35-0255ac101f90

----End

Unsubscribing from a Resource Package By Submitting a Service Ticket

Step 1 Log in to the management console.

Step 2 In the upper right corner of the page, choose **Service Tickets** > **Create Service Ticket**. The **Create Service Ticket** page is displayed. Click **Subscriptions** on the **Select Ticket Type** tab page.

Figure 2-3 Going to the Subscriptions page

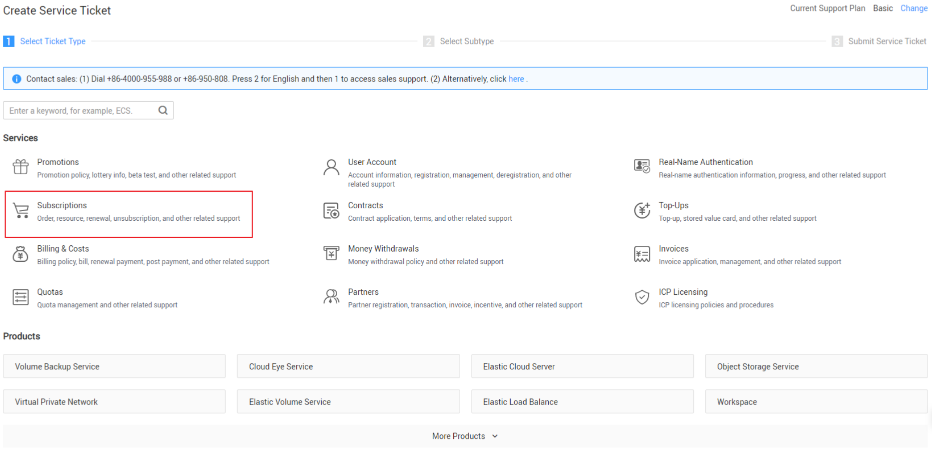

Contact sales: (1) Dial +86-4000-955-988 or +86-950-808. Press 2 for English and then 1 to access sales support. (2) Alternatively, click here .

Services

Promotions
Promotion policy, lottery info, beta test, and other related support

User Account
Account information, registration, management, deregistration, and other related support

Real-Name Authentication
Real-name authentication information, progress, and other related support

Subscriptions
Order, resource, renewal, unsubscription, and other related support

Contracts
Contract application, terms, and other related support

Top-Ups
Top-up, stored value card, and other related support

Billing & Costs
Billing policy, bill, renewal payment, post payment, and other related support

Money Withdrawals
Money withdrawal policy and other related support

Invoices
Invoice application, management, and other related support

Quotas
Quota management and other related support

Partners
Partner registration, transaction, invoice, incentive, and other related support

ICP Licensing
ICP licensing policies and procedures

Products

Step 3 On the **Select Subtype** tab page, choose **UnsubscriptionsUnsubscription - Unsubscription Problem** > **Create Service Ticket**.

Figure 2-4 Creating a service ticket

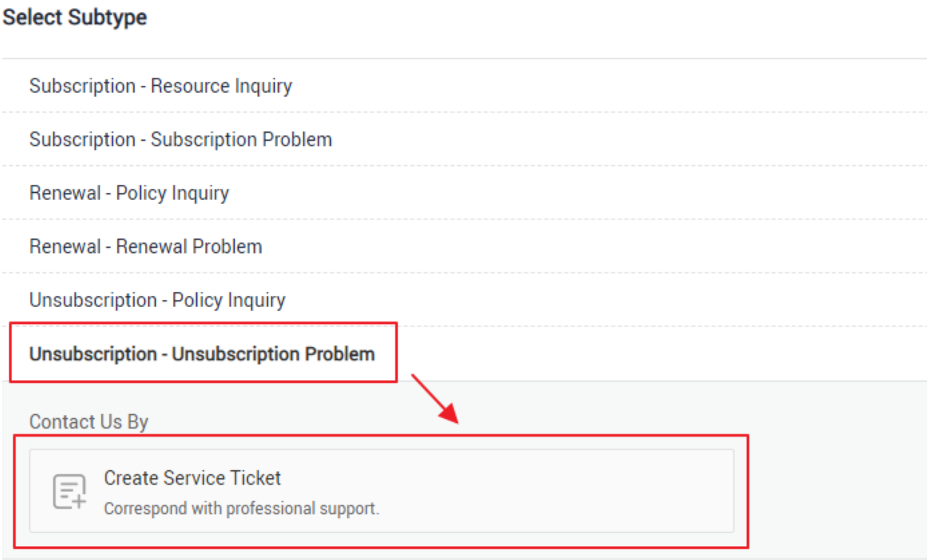

Select Subtype

Subscription - Resource Inquiry

Subscription - Subscription Problem

Renewal - Policy Inquiry

Renewal - Renewal Problem

Unsubscription - Policy Inquiry

Unsubscription - Unsubscription Problem

Contact Us By

Create Service Ticket
Correspond with professional support.

Step 4 Enter a description, for example, "unsubscribe from CSBS/VBS resource packages", in the text box next to **Problem Description** and set other parameters as required. Then click **Submit**.

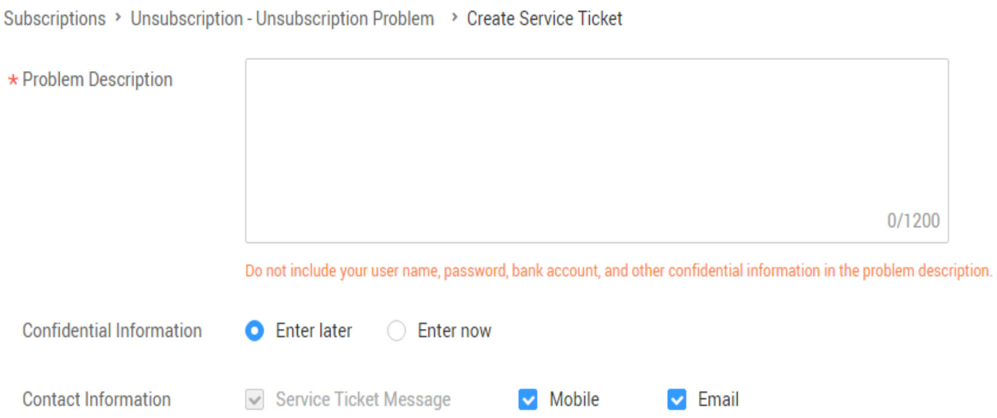

Subscriptions › Unsubscription - Unsubscription Problem › Create Service Ticket

* Problem Description 0/1200

Do not include your user name, password, bank account, and other confidential information in the problem description.

Confidential Information ● Enter later ○ Enter now

Contact Information ☑ Service Ticket Message ☑ Mobile ☑ Email

----End

2.6 Why Is a Message Displayed Indicating Insufficient User Rights When I Create a Policy?

If your account is in arrears or has no balance, you cannot create policies or add tags.

2.7 What Can I do If a Yearly/Monthly-Billed Vault Is About to Expire?

After a yearly/monthly-billed vault expires, the system will not automatically change you to the pay-per-use mode. For details about the resource handling during the retention period, see **Service Suspension and Resource Release**. If the resource package is not renewed before the retention period expires, the resource will be deleted.

- If you want to continue to use the vault, choose **More** > **Renew** in the **Operation** column of the vault to renew your subscription.
- If you do not need the vault anymore, choose **More** > **Delete** in the **Operation** column of the vault, or you can wait for the system to automatically delete it when the subscription expires.

2.8 How Do I Unsubscribe from a Vault?

If you no longer need a vault billed in yearly/monthly mode to store backups, you can unsubscribe from the vault in either of the following methods. For details about the unsubscription rules, see **Unsubscription Rules**.

Method 1

Step 1 Log in to the CBR console.

Step 2 Click the **Vaults** tab and locate the target vault. Click **More** > **Unsubscribe** in the **Operation** column.

Step 3 Complete the unsubscription operations as prompted.

----End

Method 2

Step 1 Log in to the CBR console.

Step 2 Click **Billing** in the upper right corner of the page to go to the Billing Center.

Step 3 In the left navigation pane, choose **Orders** > **Unsubscriptions** to view the yearly/monthly vaults you purchased.

Step 4 Select an unwanted vault and click **Unsubscribe**. See **Figure 2-5**.

Figure 2-5 Unsubscribing from a resource package

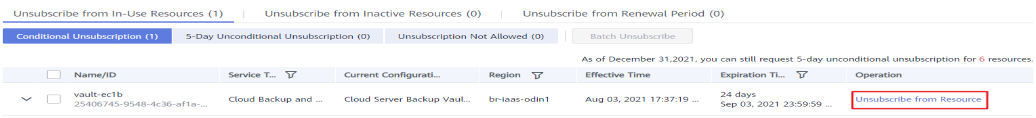
Unsubscribe from In-Use Resources (1)
Unsubscribe from Inactive Resources (0)
Unsubscribe from Renewal Period (0)
Conditional Unsubscription (1)
5-Day Unconditional Unsubscription (0)
Unsubscription Not Allowed (0)
Batch Unsubscribe

As of December 31,2021, you can still request 5-day unconditional unsubscription for 6 resources.

	Name/ID	Service T...	Current Configurati...	Region	Effective Time	Expiration Ti...	Operation
□	vault-ec1b 25406745-9548-4c36-af1a-...	Cloud Backup and ...	Cloud Server Backup Vaul...	br-iaas-odin1	Aug 03, 2021 17:37:19 ...	24 days Sep 03, 2021 23:59:59 ...	Unsubscribe from Resource

----End

2.9 What Should I Do If the Fee Is Too High When I Expand the Capacity of a Vault?

Symptom

A user purchases a yearly-billed vault (1 to 3 years). After the vault has been used for a period of time, the user wants to expand the capacity. However, the expansion cost is unexpectedly high.

Possible Cause

Due to system restrictions, when expanding a vault in the current yearly package, the fee cannot be calculated on the common yearly/monthly basis.

Solution

1. Plan the vault capacity properly before you purchase a yearly/monthly vault.
2. If you still want to expand the vault capacity, renew the vault by making up the used duration.

   For example, if you have purchased a vault for two years and need to expand its capacity after using it for six months, you need to renew the vault for another six months to add up to a two-year package. Then you can expand the vault on a relatively preferential price.

2.10 Can I Change a Server Backup Vault to a Disk Backup Vault or the Other Way Around?

No. You can unsubscribe from or delete the unwanted vault and then purchase a new one. For details about the unsubscription rules, see **Unsubscription Rules**.

2.11 Why Are CBR Backups Displayed on the VBS Console?

If you have migrated data from CSBS and VBS to CBR, and created a backup on the CBR console, the same backup record will be generated on the VBS console page. This is due to the underlying mechanism. The VBS console displays all backups generated by CBR, CSBS, and VBS. However, these backups will not be billed repeatedly.

To delete a CBR or CSBS backup, go to the corresponding console to delete it, and then it will no longer be displayed on the VBS console.

2.12 What Charges Will I Incur When Using a Backup to Create an Image?

You will be charged only for the backup vault capacity when you use a cloud server backup to create an image.

For detailed pricing, see **CBR pricing details**.

2.13 How Do I Purchase a CBR Replication Traffic Package?

CBR provides replication traffic packages. If you buy a package, the quota in the package is used for billing first. Any usage exceeding the package is billed based on the pay-per-use basis.

A package is bound to a specific region. The purchased package can be used only in the bound region. This region must be the source region of the replication.

Package resetting rules: During billing, the quota in the package is deducted first. Remaining quotas are cleared upon the monthly reset, and excessive usage is billed on the pay-per-use basis. The quota of a package defines the available traffic within one month from the date you bought the package.

Procedure

Step 1 Log in to the CBR console.

1. Log in to the management console.
2. Click 📍 in the upper left corner and select a region.
3. Choose **Storage** > **Cloud Backup and Recovery** > **Cloud Server Backup**.

Step 2 Click **Buy CBR Replication Traffic Package** in the upper right corner.

Step 3 Select a package type and usage duration.

Step 4 Confirm the information, click **Next**, and complete the payment.

----End

2.14 Can I Change the Protection Type of an Existing Vault?

No. The protection type of a purchased vault cannot be changed. You need to delete or unsubscribe from the vault and then create a new one.

2.15 What Resources Do I Need to Implement Cross-Region Backup Replication?

To replicate backups to a different region, buy or create the following resources:

1. Cloud server backup vault in the source region. This vault is used to store cloud server backups generated in the production region.
2. Cloud server replication vault in the destination region. This vault is used to store cloud server backups replicated to the disaster recovery region.
3. (Optional) Cross-region replication traffic package. This package can be used to deduct the replication traffic generated during replication.

If a replication traffic package is not available, you will be billed for the replication traffic based on a pay-per-use basis. Buy a replication traffic package based on the size of the backups to be replicated. A traffic package can only be used to deduct traffic fees. Any capacity fee generated during replication cannot be deducted by the traffic package.

3 Backup

3.1 Do I Need to Stop the Server Before Performing a Backup?

No. You can back up servers that are in use. When a server is running, data is written into disks on the server, and some newly generated data is cached in the server memory. During a backup task, data in the memory will not be automatically written into disks, so the disk data and their backups may be inconsistent.

To ensure data integrity, you are advised to perform the backup during off-peak hours when no data is written to the disks. For applications that require strict consistency, such as databases and email systems, you are advised to enable application-consistent backup.

3.2 Can I Back Up a Server Deployed with Databases?

Yes. CBR provides application-consistent backup. For details about the function compatibility, see **Table 3-1**. For applications or databases with which the application-consistent function is incompatible, you are advised to suspend all data write operations before performing backup. If write operations cannot be suspended, you can stop the application systems or the server for offline backup. If you do not perform the preceding operations before backup, status of the server after restoration will be similar to restart upon an unexpected power failure. In this case, log rollback will be performed on databases to keep data consistent.

Table 3-1 OSs that support installation of the Agent

Database	OS	Version
SQL Server 2008/2012/2019	Windows	Windows Server 2008, 2008 R2, 2012, 2012 R2, and 2019 for x86_64

Database	OS	Version
SQL Server 2014/2016/ Enterprise Edition	Windows	Windows Server 2014, 2014 R2, and 2016 Datacenter for x86_64
MySQL 5.5/5.6/5.7	Red Hat	Red Hat Enterprise Linux 6 and 7 for x86_64
	SUSE	SUSE Linux Enterprise Server 11, 12, 15 SP1, 15 SP2 for x86_64
	CentOS	CentOS 6 and 7 for x86_64
	EulerOS	EulerOS 2.2 and 2.3 for x86_64
HANA 1.0/2.0	SUSE	SUSE Linux Enterprise Server 12 for x86_64

3.3 How Can I Distinguish Automatic Backups From Manual Backups?

They can be distinguished by name prefix:

- Automatic backups: **autobk_***xxxx*
- Manual backups: **manualbk_***xxxx* or custom names

3.4 Can I Choose to Back Up Only Some Partitions of a Disk?

No. The minimum backup granularity supported by CBR is disks.

3.5 Does CBR Support Cross-Region Backup?

You can replicate backups to a destination region and create images in the destination region using the generated replicas.

3.6 How Do I Migrate Server Data Across Regions Using Server Backups?

Context

A user has an ECS in the LA-Mexico City1 region, and the ECS has one system disk. To implement cross-region disaster recovery and fast service deployment in a different region, the user needs to create a same ECS with the same data in that region. In this case, CBR cloud server backup would be an ideal choice.

To replicate the ECS from LA-Mexico City1 to AP-Bangkok, the user needs to back up the ECS in LA-Mexico City1, replicates the backup to AP-Bangkok, uses the backup replica to create a full-ECS image, and then uses the image to create an ECS in AP-Bangkok. In this way, data on the ECS can be migrated to AP-Bangkok.

Procedure

Step 1 Log in to the ECS console, switch to the LA-Mexico City1 region, and locate the target ECS in the server list. See **Figure 3-1**.

Figure 3-1 Target ECS information

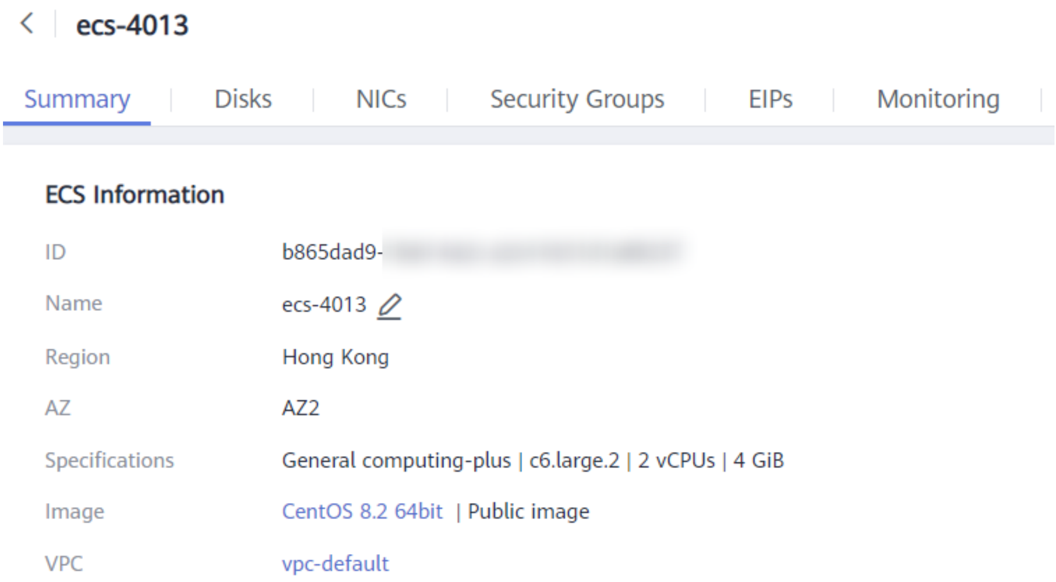

ecs-4013

ECS Information

ID	b865dad9-[redacted]
Name	ecs-4013
Region	Hong Kong
AZ	AZ2
Specifications	General computing-plus | c6.large.2 | 2 vCPUs | 4 GiB
Image	CentOS 8.2 64bit | Public image
VPC	vpc-default

Step 2 Choose **More** > **Create Backup** in the **Operation** column to go to the CBR console and purchase a vault. Make sure that the target ECS is selected when purchasing a vault.

Step 3 After the vault is created and associated with the ECS, back up the ECS. See **Figure 3-2**.

Figure 3-2 Perform Backup

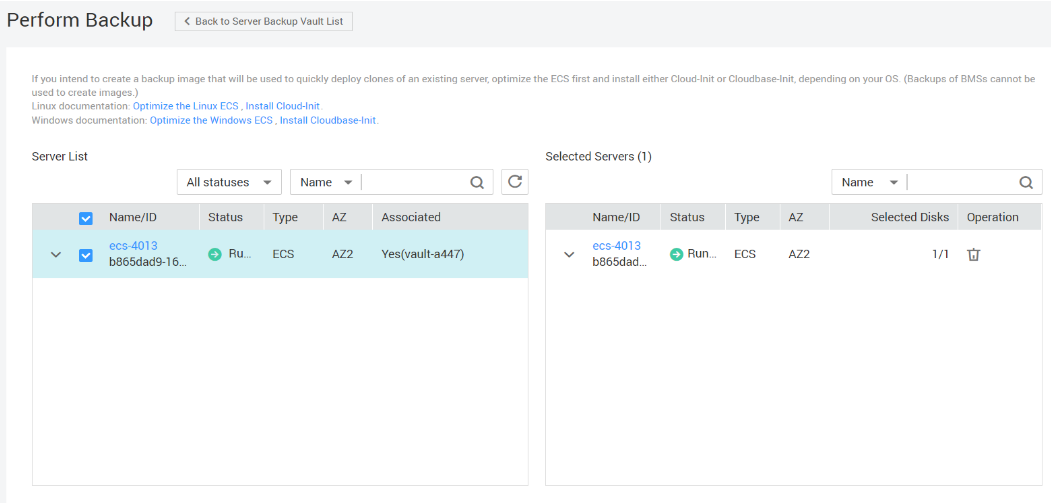

Perform Backup

If you intend to create a backup image that will be used to quickly deploy clones of an existing server, optimize the ECS first and install either Cloud-Init or Cloudbase-Init, depending on your OS. (Backups of BMSs cannot be used to create images.)
Linux documentation: Optimize the Linux ECS , Install Cloud-Init.
Windows documentation: Optimize the Windows ECS , Install Cloudbase-Init.

Server List

☑	Name/ID	Status	Type	AZ	Associated
☑	ecs-4013 b865dad9-16...	Ru...	ECS	AZ2	Yes(vault-a447)

Selected Servers (1)

Name/ID	Status	Type	AZ	Selected Disks	Operation
ecs-4013 b865dad...	Run...	ECS	AZ2	1/1	🗑

Step 4 Go to the cloud server backup page on the CBR console and confirm that the backup has been created in the vault. Then, choose **More** > **Create Replica** in the **Operation** column to replicate the backup to AP-Bangkok. See **Figure 3-3**.

If the destination vault does not exist, go to the destination region to create a replication vault first.

Figure 3-3 Creating a replica

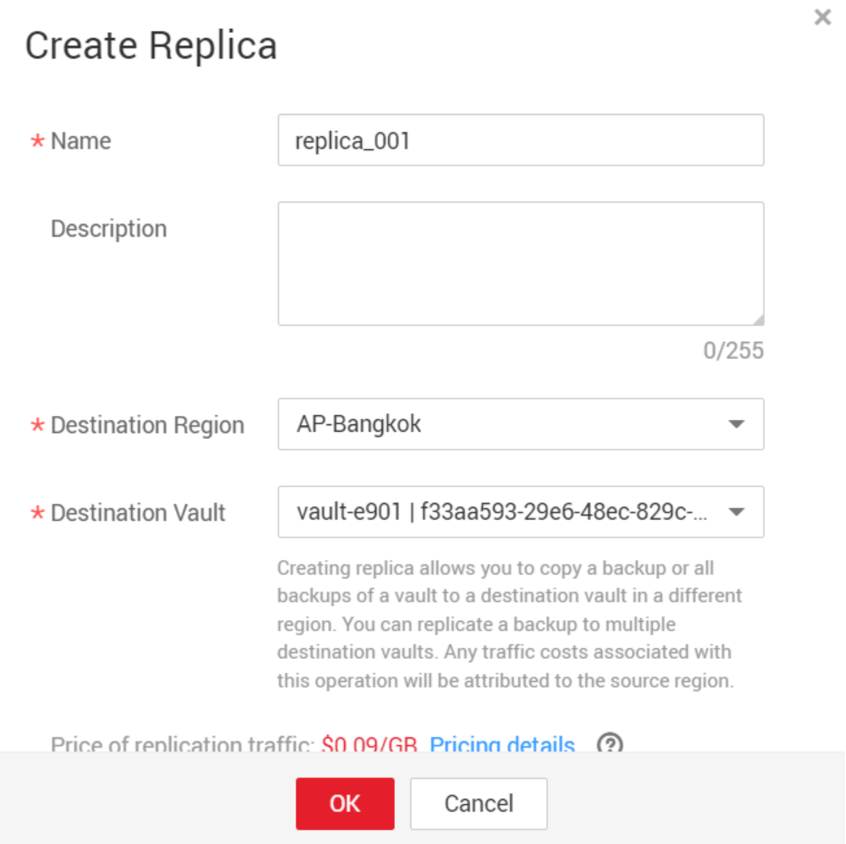

Create Replica

* Name: replica_001

Description:

0/255

* Destination Region: AP-Bangkok

* Destination Vault: vault-e901 | f33aa593-29e6-48ec-829c-...

Creating replica allows you to copy a backup or all backups of a vault to a destination vault in a different region. You can replicate a backup to multiple destination vaults. Any traffic costs associated with this operation will be attributed to the source region.

Price of replication traffic: $0.09/GB Pricing details

Step 5 Go to the AP-Bangkok region. Choose **Cloud Server Backups**, find the replicated backup **replica_001** in the backup list. An **R** icon is displayed next to the backup name. In the **Operation** column, click **Create Image**.

Create an image using the backup. See **Figure 3-4**. This backup can be deleted only after the image is deleted.

Figure 3-4 Creating an image

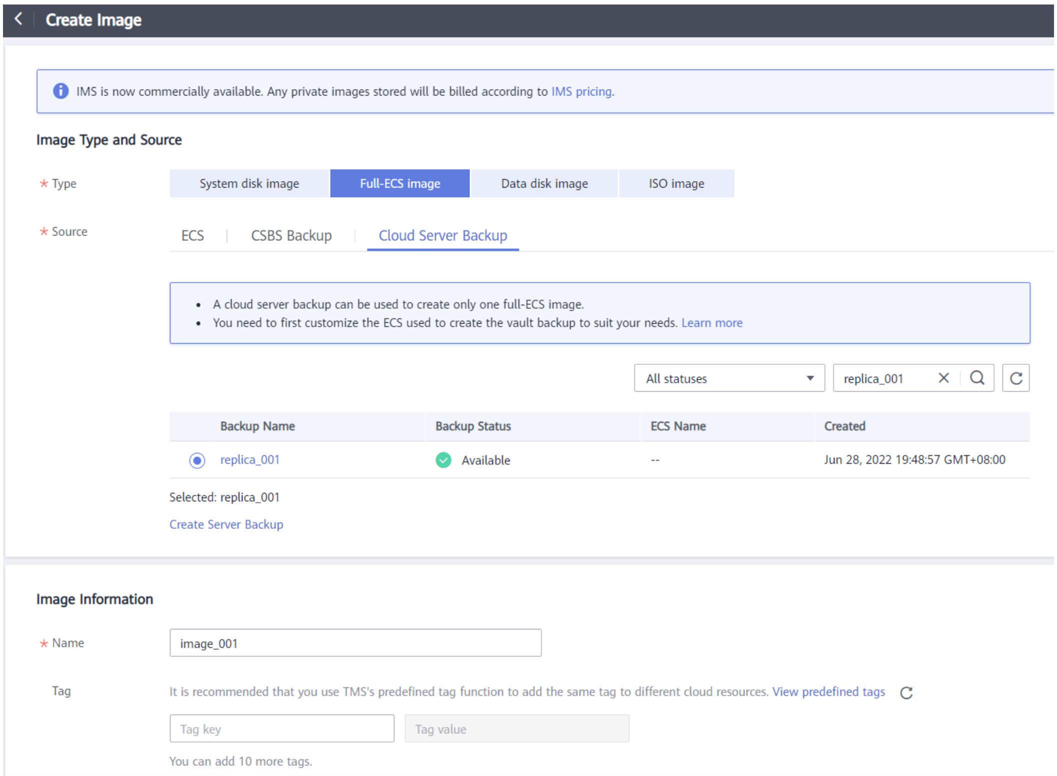

Create Image

IMS is now commercially available. Any private images stored will be billed according to IMS pricing.

Image Type and Source

* Type: System disk image | Full-ECS image | Data disk image | ISO image

* Source: ECS | CSBS Backup | Cloud Server Backup

- A cloud server backup can be used to create only one full-ECS image.
- You need to first customize the ECS used to create the vault backup to suit your needs. Learn more

All statuses | replica_001

	Backup Name	Backup Status	ECS Name	Created
●	replica_001	✓ Available	--	Jun 28, 2022 19:48:57 GMT+08:00

Selected: replica_001

Create Server Backup

Image Information

* Name: image_001

Tag: It is recommended that you use TMS's predefined tag function to add the same tag to different cloud resources. View predefined tags

Tag key | Tag value

You can add 10 more tags.

Step 6 On the IMS console in AP-Bangkok, find the created image and click **Apply for Server** in the **Operation** column to create a server. See **Figure 3-5**. Configure other parameters based on service requirements.

Figure 3-5 Creating a server

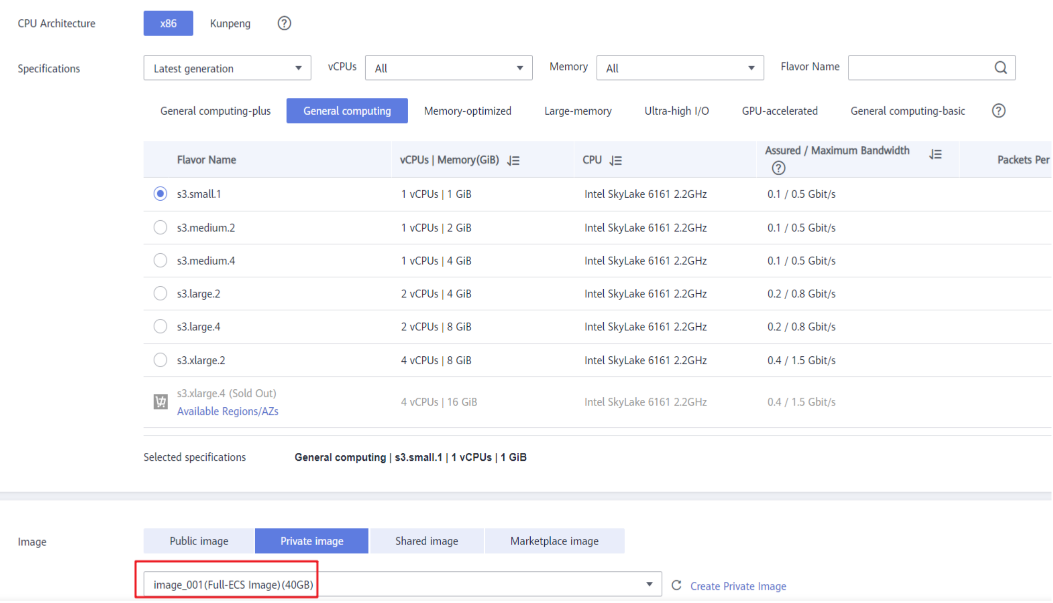

CPU Architecture: x86 | Kunpeng

Specifications: Latest generation | vCPUs: All | Memory: All | Flavor Name:

General computing-plus | General computing | Memory-optimized | Large-memory | Ultra-high I/O | GPU-accelerated | General computing-basic

	Flavor Name	vCPUs | Memory(GiB)	CPU	Assured / Maximum Bandwidth	Packets Per
●	s3.small.1	1 vCPUs | 1 GiB	Intel SkyLake 6161 2.2GHz	0.1 / 0.5 Gbit/s	
○	s3.medium.2	1 vCPUs | 2 GiB	Intel SkyLake 6161 2.2GHz	0.1 / 0.5 Gbit/s	
○	s3.medium.4	1 vCPUs | 4 GiB	Intel SkyLake 6161 2.2GHz	0.1 / 0.5 Gbit/s	
○	s3.large.2	2 vCPUs | 4 GiB	Intel SkyLake 6161 2.2GHz	0.2 / 0.8 Gbit/s	
○	s3.large.4	2 vCPUs | 8 GiB	Intel SkyLake 6161 2.2GHz	0.2 / 0.8 Gbit/s	
○	s3.xlarge.2	4 vCPUs | 8 GiB	Intel SkyLake 6161 2.2GHz	0.4 / 1.5 Gbit/s	
	s3.xlarge.4 (Sold Out) Available Regions/AZs	4 vCPUs | 16 GiB	Intel SkyLake 6161 2.2GHz	0.4 / 1.5 Gbit/s	

Selected specifications: **General computing | s3.small.1 | 1 vCPUs | 1 GiB**

Image: Public image | Private image | Shared image | Marketplace image

image_001(Full-ECS Image)(40GB)

Step 7 View the server details page. You can see that the ECS data has been migrated from LA-Mexico City1 to AP-Bangkok. See **Figure 3-6**.

Figure 3-6 Successful cross-region ECS data migration

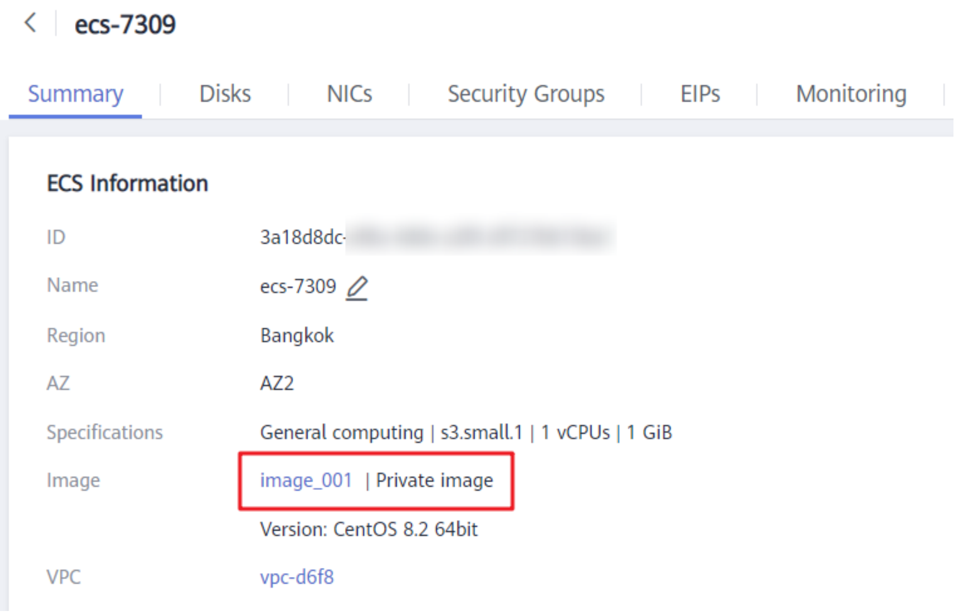

ECS Information

ID: 3a18d8dc-[redacted]

Name: ecs-7309

Region: Bangkok

AZ: AZ2

Specifications: General computing | s3.small.1 | 1 vCPUs | 1 GiB

Image: image_001 | Private image
Version: CentOS 8.2 64bit

VPC: vpc-d6f8

----End

3.7 Can I Back Up Two Disks to One Target Disk?

No. One target disk corresponds to one source disk. The data of two disks cannot be backed up to one target disk.

3.8 How Do I Replicate a Disk to the Same AZ in a Region as the Source Disk?

Back up the desired disk. Then use the disk backup to create a new disk, and select the same AZ as that of the source disk for the new one.

3.9 Can I Migrate Backups Between Vaults?

Backups can be migrated between vaults. For details, see **Migrating a Resource**.

3.10 Will the Server Performance Be Affected If I Delete Its Backups?

No. Backups are not stored on a server. Therefore, deleting its backups has no impact on the server performance.

3.11 Can I Use Its Backup for Restoration After a Resource Is Deleted?

Yes. Resources and backups are not stored together. If a resource is deleted, its backup still stays in your CBR vault. You can use the backup to restore the resource to a backup point in time.

3.12 How Many Backups Can I Create for a Resource?

You can create as many backups for a resource as needed.

3.13 Can I Use an Incremental Backup to Restore Data After a Full Backup Is Deleted?

Yes.

CBR allows you to use any backup, no matter it is a full or incremental one, to restore the full data of a resource. By virtue of this, manual or automatic deletion of a backup will not affect the restoration function.

Suppose server **X** has backups **A**, **B**, and **C** (in time sequence) and every backup involves data changes. If backup **B** is deleted, you can still use backup **A** or **C** to restore data.

3.14 Can I Stop an Ongoing Backup Task?

No. An ongoing backup task cannot be stopped.

3.15 How Do I Reduce the Vault Space Occupied by Backups?

Symptom

The size of a disk backup is much greater than the used space of the disk displayed on a server. Even if you delete large files from the disk and back up the disks again, the backup size does not reduce significantly.

Possible Cause

After files are deleted from a disk, the data remains though it is no longer available. When you use CBR to back up a disk, all disk data including the invisible data will be backed up. For the backup principles, see **Why Is My Backup Size Larger Than My Disk Size?**.

Solution

Currently, CBR cannot help reduce the backup size. You can use a third-party tool to do this but need to evaluate the security of the tool by yourself.

3.16 How Do I View the Size of Each Backup?

You cannot view the size of each backup.

However, you can view the size of all backups for each resource. On the **Backups** tab page, click the name of the target backup to view its details. See **Figure 3-7**.

Figure 3-7 Checking the size of all backups of a server

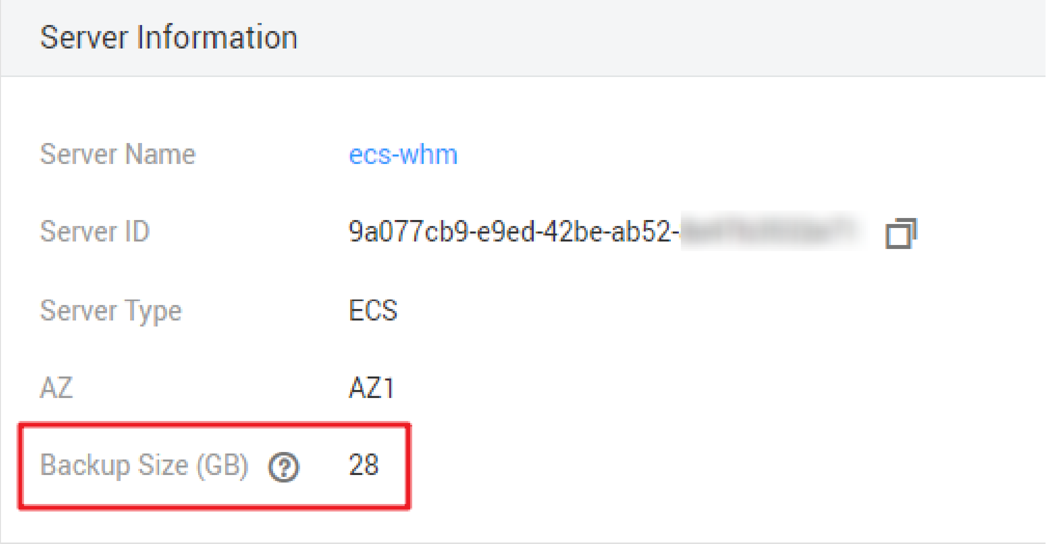

Server Information

Server Name ecs-whm

Server ID 9a077cb9-e9ed-42be-ab52-[redacted]

Server Type ECS

AZ AZ1

Backup Size (GB) 28

3.17 How Do I View My Backup Data?

You can check your backup data in the following ways:

NOTE

Backup data cannot be viewed on the CBR console.

Server Backups

1. Create an image from a server backup. For details, see **Using a Backup to Create an Image**.
2. Use the image to create a server. For details, see **Creating an ECS from an Image**.
3. Log in to the server to view the data.

Disk Backups

1. Create a new disk from a disk backup. For details, see **Using a Backup to Create a Disk**.
2. Attach the created disk to a server. For details, see **Attaching a Non-Shared Disk** or **Attaching a Shared Disk**.
3. Log in to the server to view the data.

SFS Turbo Backups

1. Create a new SFS Turbo file system from an SFS turbo backup. For details, see **Using a Backup to Create a File System**.
2. Mount the file system to a server.
   - To mount the file system to a Linux server, see **Mounting an NFS File System to ECSs (Linux)**.
   - To mount the file system to a Windows server, see **Mounting an NFS File System to ECSs (Windows)**.
3. Log in to the server to view the data.

3.18 How Long Will My Backups Be Kept?

Manual backup: The name of a manual backup is usually in the format of **manualbk_***xxxx* or is customized. If you do not delete manual backups and your account balance is sufficient, manual backups will always be kept. If your account balance is not enough or not topped up timely, manual backups will be automatically released and cannot be recovered.

Automatic backup: The name of an automatic backup is in the format of **autobk_***xxxx*. If a retention rule has been set in the policy, automatic backups will be kept and deleted based on the retention rule. If the policy's retention rule has been changed during the backup execution, some automatic backups may not be deleted. For details, see **Why the New Retention Rule I Changed Is Not Applied?**

3.19 How Do I Implement Periodic Full Backups for My Resources?

Context

Huawei Cloud CBR by default performs a full backup for a resource in the initial backup and incremental backups in all subsequent backups.

CBR now allows for periodic full backups in addition to the initial backup. You can configure a policy to perform a full backup after every N incremental backups. This further improves backup data security and meets periodic full backup needs.

Figure 3-8 Periodic full backup process

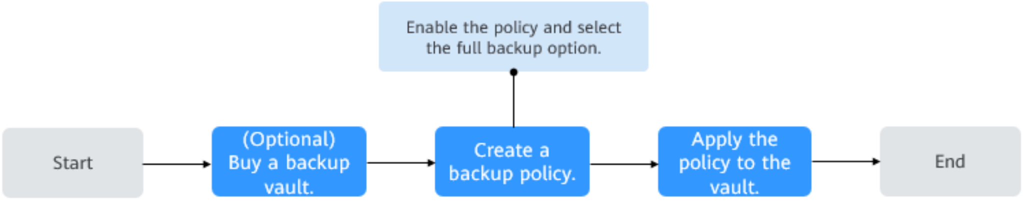
Enable the policy and select the full backup option.
Start
(Optional) Buy a backup vault.
Create a backup policy.
Apply the policy to the vault.
End

Procedure

Step 1 Log in to the CBR console.

Step 2 Choose **Policies** and click the **Backup Policies** tab. In the upper right corner, click **Create Policy**. See **Figure 3-9**.

Figure 3-9 Create Policy

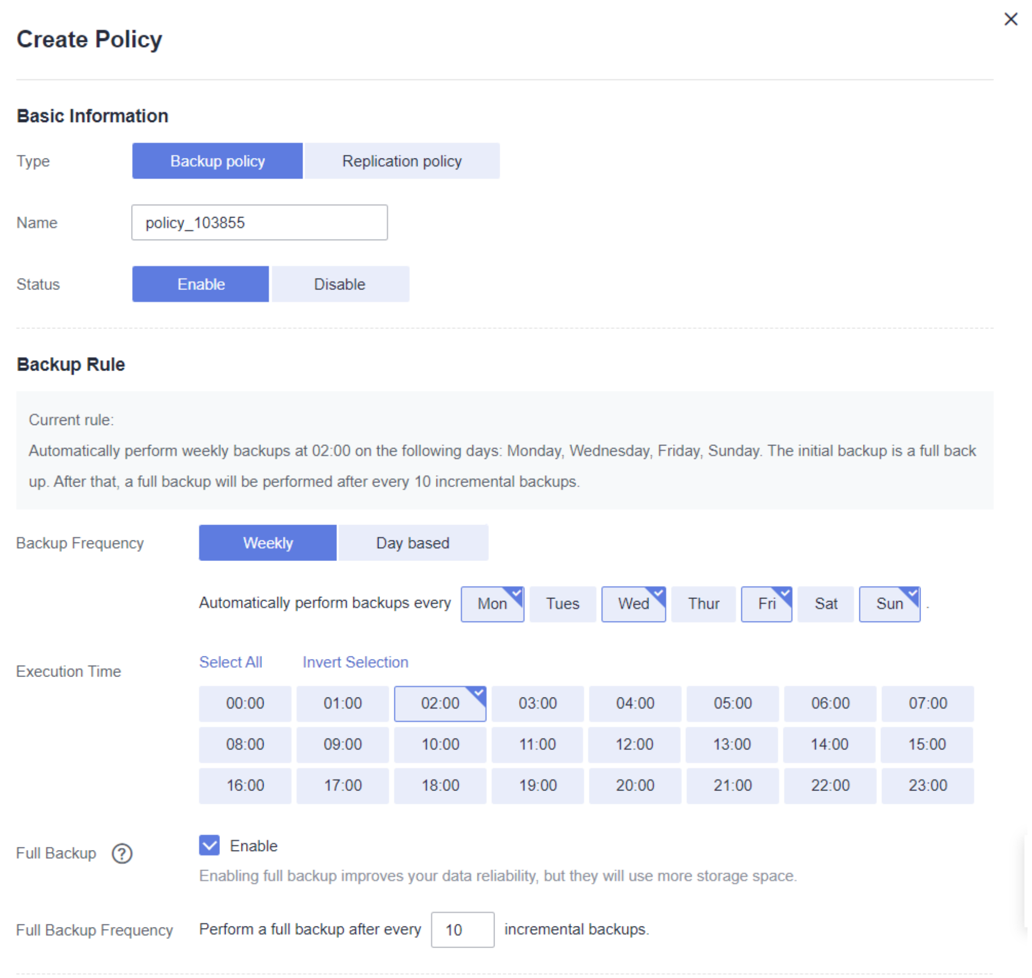

Create Policy

Basic Information

Type: Backup policy | Replication policy

Name: policy_103855

Status: Enable | Disable

Backup Rule

Current rule:
Automatically perform weekly backups at 02:00 on the following days: Monday, Wednesday, Friday, Sunday. The initial backup is a full back up. After that, a full backup will be performed after every 10 incremental backups.

Backup Frequency: Weekly | Day based

Automatically perform backups every Mon Tues Wed Thur Fri Sat Sun .

Execution Time: Select All Invert Selection

00:00	01:00	02:00	03:00	04:00	05:00	06:00	07:00
08:00	09:00	10:00	11:00	12:00	13:00	14:00	15:00
16:00	17:00	18:00	19:00	20:00	21:00	22:00	23:00

Full Backup ? ☑ Enable
Enabling full backup improves your data reliability, but they will use more storage space.

Full Backup Frequency: Perform a full backup after every 10 incremental backups.

Step 3 Set backup policy parameters according to **Backup policy parameter description**.

Step 4 Select **Enable** for **Full Backup** and set the full backup frequency. Then the system automatically checks whether the configured frequency is appropriate. If message shows that full backups cannot be executed, change the frequency. See **Figure 3-10**.

Figure 3-10 Configuring periodic full backup

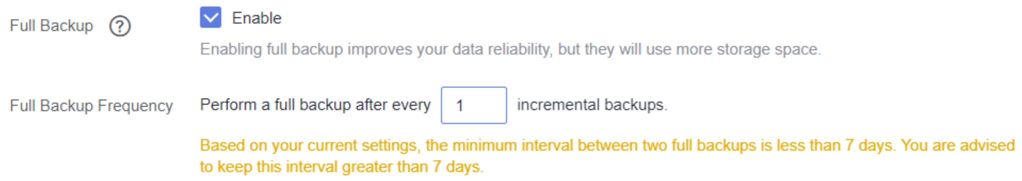

Full Backup ? ☑ Enable
Enabling full backup improves your data reliability, but they will use more storage space.

Full Backup Frequency: Perform a full backup after every 1 incremental backups.
Based on your current settings, the minimum interval between two full backups is less than 7 days. You are advised to keep this interval greater than 7 days.

Step 5 Click **OK**.

Step 6 Locate the desired vault and choose **More** > **Apply Backup Policy** to apply the created policy to the vault. You can view the applied policy on the vault details page.

After the policy is applied, periodic full backups will be automatically performed based on the policy.

----End

4 Capacity

4.1 Why Is My Backup Size Larger Than My Disk Size?

Symptoms

- There is no difference or an increase in size between the original backup and a backup generated after a file is deleted.
- The ECS backup size is larger than the used disk space obtained from the file system.

Possible Causes

Possible causes are as follows:

- The backup mechanism itself causes this problem. The cloud server backups, SFS Turbo backups, and cloud disk backups created using CBR are all block-level backups. Different from file-level backups, block-level backups are performed by sector (512 bytes) each time.
- The metadata of the file systems on the disk occupies disk space.
- To reduce performance overhead, the file system adds a delete marker for the deleted file, but does not erase the data that has been written to the sector, and the metadata on the sector still exists. Block-level backups cannot detect whether data on a sector is deleted or not, but only determine whether a backup needs to be performed by checking whether all data blocks are zero blocks.
- CBR determines whether data in each sector changes by comparing two snapshots. Data changes include data addition, modification, and deletion. Backup is not performed if there are no data changes. If there are data changes, CBR further checks whether data blocks in the sector are all zero blocks. If so, backup is also not performed. Backups are performed only when there are non-zero blocks. If the data is deleted but metadata in the sector is not, the data block is also recognized as a non-zero block, and backups will be performed.

Solution

If you only need to back up some of your files, use **file backup** to reduce the backup space and lower the backup costs.

4.2 What Can I Do If the Vault Capacity Is Not Enough?

If your vault capacity is used up, CBR will not continue to back up your resources. New backups will never overwrite previous backups.

You can handle the storage insufficiency by either expanding the vault capacity or reducing the number of retained backups.

- Expanding vault capacity

  If you want to retain the generated backups, expand the vault capacity. For details, see **Expanding Vault Capacity**.
- Reducing the number of retained backups
  a. Locate the target vault and delete unwanted backups by referring to **Deleting a Backup**.
  b. If the vault is applied with a backup policy, you can:
     i. Decrease the backup frequency, shorten the retention period (automatically deleting expired backups), or reduce the number of servers associated with the vault.
     ii. Disable the backup policy or remove the policy from the vault. To disable the policy, see **Modifying a Policy**. To remove the policy, see **Removing a Policy from a Vault**. Then, automatic backup will stop, and the vault used space will not change.

4.3 Why Does the Used Capacity of a Vault Change Only Slightly After I Deleted Unwanted Backups?

Symptoms

After unwanted backups are deleted from the vault, the used capacity of the vault decreases by only 1 GB to 2 GB.

Possible Causes

The backup mechanism of CBR:

- By default, CBR performs a full backup for a resource for the first time and backs up all used data blocks. All subsequent backups are incremental. An incremental backup backs up only the data blocks changed since the last backup.
- Each incremental backup is a virtual full backup. Correlated data blocks are indexed by using pointers.
- When you delete a backup, no matter manually or automatically, only data blocks that are not referenced by other backups will be deleted.

Figure 4-1 Backup mechanism

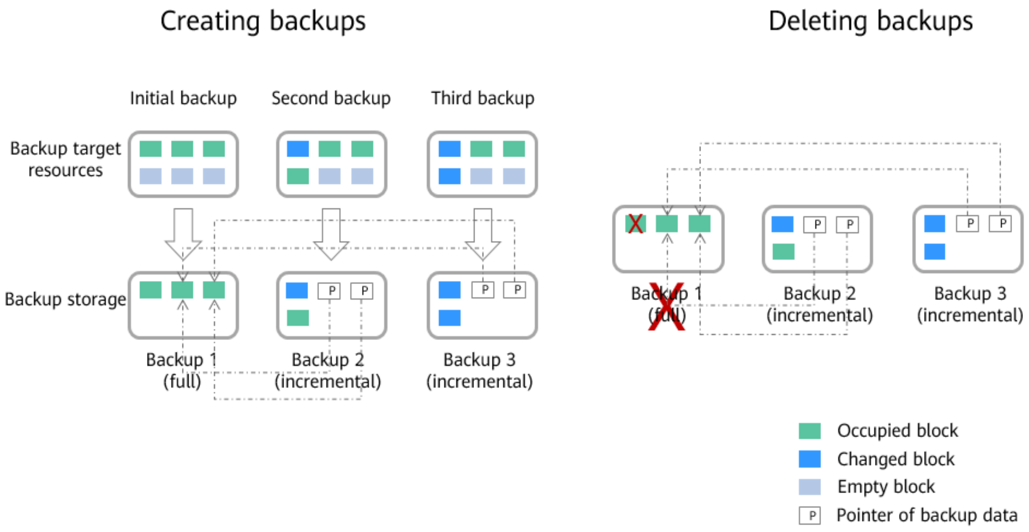

4.4 Will Backup Continue If the Usage of a Vault Reaches the Upper Limit?

If the vault is just used up or its remaining space is not enough for the next backup, the next backup can be executed successfully.

However, backup stops once the usage of the vault exceeds the upper limit.

4.5 How Do I Adjust the Vault Capacity Alarm Threshold?

Symptom

If you set a vault capacity threshold, you will receive an alarm when your vault capacity exceeds the configured threshold.

Solution

1. If the configured threshold is inappropriate, go to Cloud Eye to modify the alarm threshold. For details, see **Creating an Alarm Rule**.

2. If the vault space is about to used up, expand the vault capacity or delete the vault. For details, see **What Can I Do If the Vault Capacity Is Not Enough?**

5 Restoration

5.1 Do I Need to Stop the Server Before Restoring Data Using Backups?

The system shuts down the server before restoring server data, and automatically starts up the server after the restoration is complete.

If you deselect **Start the server immediately after restoration**, you need to manually start the server after the restoration is complete.

5.2 Can I Use a System Disk Backup to Recover an ECS?

Yes. However, before the recovery, you need to detach the system disk to be recovered from the ECS.

You can also use a backup of the system disk to create new disks. However, newly created disks cannot be used as system disks.

5.3 Do I Need to Stop the Server Before Restoring Data Using Disk Backups?

Yes. Before restoring the disk data using a disk backup, you must stop the server to which the disk is attached, and detach the disk from the server. After the disk data is restored, attach the disk to the server and start the server.

5.4 Can a Server Be Restored Using Its Backups After It Is Changed?

Yes. If a server has been backed up and then changed (adding, deleting, or expanding disks), its backups can still be used to restore data. You are advised to back up data again after the change.

If you have added a disk after a backup and then use the backup to restore data, data on the new disk will not change.

If you have deleted a disk after a backup and then use the backup to restore data, data on the deleted disk cannot be restored.

If you have changed the server OS after a backup and then use a system disk backup to restore the original system disk, the system disk data cannot be restored because the disk UUID has changed. If the entire server is backed up and you changed the server OS, the backup can be used to restore data only after you have changed back to the original OS.

5.5 Can a Disk Be Restored Using Its Backups After Its Capacity Is Expanded?

Yes. After restoration, the capacity of the expanded disk goes back to the original capacity before expansion. If you want to use the capacity added to the disk, you need to attach the restored disk to a server, log in to the server, and then manually modify the file system configuration. For detailed operations, see sections about post-expansion operations on disks in the *Elastic Volume Service User Guide*.

5.6 What Can I Do if the Password Becomes a Random One After I Use a Backup to Restore a Server or Use an Image to Create a Server?

For details about how to reset the password, see **Passwords** in the *Elastic Cloud Server User Guide*.

5.7 What Changes Will Be Made to the Original Backup When I Use the Backup to Restore a Server?

- For Linux:
  - Check whether drivers related to the PV driver exist. If yes, delete them.
  - Modify the **grub** and **syslinux** configuration files to add the OS kernel boot parameters and change the disk partition name to **UUID=***UUID of the disk partition*.
  - Change the names of the disk partitions in the **/etc/fstab** file to **UUID=***UUID of the disk partition*.
  - Delete services of VMware tools.
  - Linux OSs automatically copy the built-in VirtIO driver to initrd or initramfs.
- For Windows:
  - Inject the VirtIO driver offline to solve the problem that the system cannot start when UVP VMTools is not installed.

5.8 How Do I Restore Data to a New Server?

You can restore data on your original server to a new server in either of the following ways:

- Method 1:

  Create an image using the backup of the original server and then use the image to create a new server. For details, see **Using a Backup to Create an Image**.

- Method 2:

  If a new server has already been created, perform the following steps:

  NOTE

  Data consistency is not guaranteed using method 2.

  a. Back up the disks on the original server.

     Ensure that all disks on the server are backed up. For how to back up server disks, see **Creating a Cloud Disk Backup**.

  b. Create new disks from the backups.

     Create new disks using their backups one by one. For details, see **Using a Backup to Create a Disk**.

  c. Attach the new disks to the new server. For details, see **Attaching a Non-Shared Disk** or **Attaching a Shared Disk**.

5.9 How Do I Restore a Data Disk Backup to a System Disk?

You can **use a disk backup to create a new disk** and **attach the new disk to a server**. Then copy data in the data disk to the system disk.

5.10 Can I Use CBR to Restore Data to Any Point When the Data Was Backed Up?

Yes. You can do as follows to verify this.

Procedure

Step 1 Log in to a server and create a file named **test1**.

Figure 5-1 Viewing the file

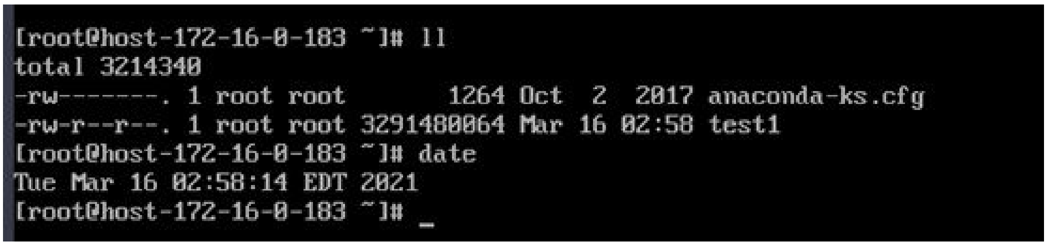

```
[root@host-172-16-0-183 ~]# ll
total 3214340
-rw-------. 1 root root       1264 Oct  2  2017 anaconda-ks.cfg
-rw-r--r--. 1 root root 3291480064 Mar 16 02:58 test1
[root@host-172-16-0-183 ~]# date
Tue Mar 16 02:58:14 EDT 2021
[root@host-172-16-0-183 ~]# _
```

Step 2 Log in to the CBR console and create a backup for the server.

Figure 5-2 Creating a backup for the server

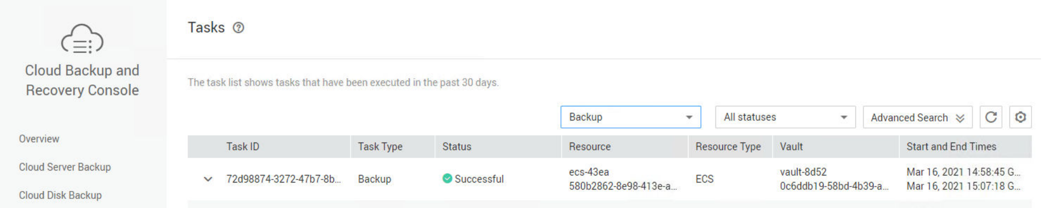

Tasks

The task list shows tasks that have been executed in the past 30 days.

Backup | All statuses | Advanced Search

Task ID	Task Type	Status	Resource	Resource Type	Vault	Start and End Times
72d98874-3272-47b7-8b...	Backup	Successful	ecs-43ea 580b2862-8e98-413e-a...	ECS	vault-8d52 0c6ddb19-58bd-4b39-a...	Mar 16, 2021 14:58:45 G... Mar 16, 2021 15:07:18 G...

Step 3 Log in to the server again and delete the **test1** file.

Figure 5-3 Deleting the file

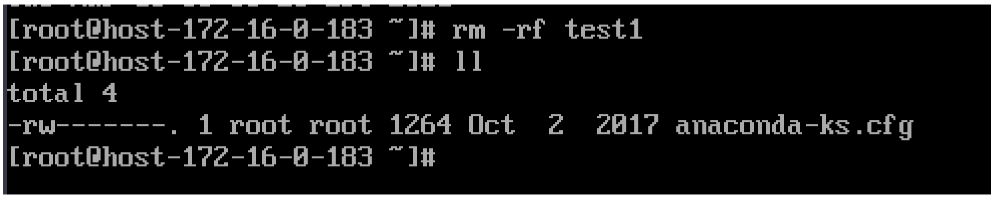

```
[root@host-172-16-0-183 ~]# rm -rf test1
[root@host-172-16-0-183 ~]# ll
total 4
-rw-------. 1 root root 1264 Oct  2  2017 anaconda-ks.cfg
[root@host-172-16-0-183 ~]#
```

Step 4 On the CBR console, use the server backup you created to restore data.

Figure 5-4 Restoring data

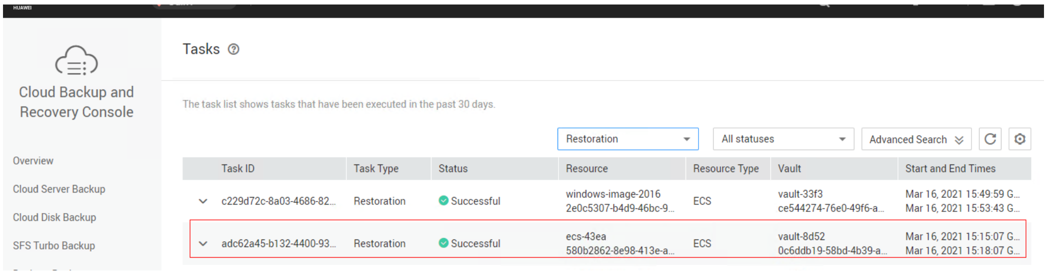

Tasks

The task list shows tasks that have been executed in the past 30 days.

Restoration | All statuses | Advanced Search

Task ID	Task Type	Status	Resource	Resource Type	Vault	Start and End Times
c229d72c-8a03-4686-82...	Restoration	Successful	windows-image-2016 2e0c5307-b4d9-46bc-9...	ECS	vault-33f3 ce544274-76e0-49f6-a...	Mar 16, 2021 15:49:59 G... Mar 16, 2021 15:53:43 G...
adc62a45-b132-4400-93...	Restoration	Successful	ecs-43ea 580b2862-8e98-413e-a...	ECS	vault-8d52 0c6ddb19-58bd-4b39-a...	Mar 16, 2021 15:15:07 G... Mar 16, 2021 15:18:07 G...

Step 5 Log in to the server and confirm that the data has been restored to the state when the backup was created.

Figure 5-5 Confirming the restoration result

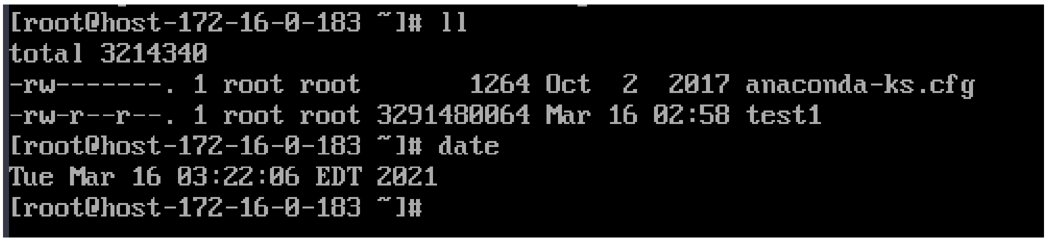

```
[root@host-172-16-0-183 ~]# ll
total 3214340
-rw-------. 1 root root       1264 Oct  2  2017 anaconda-ks.cfg
-rw-r--r--. 1 root root 3291480064 Mar 16 02:58 test1
[root@host-172-16-0-183 ~]# date
Tue Mar 16 03:22:06 EDT 2021
[root@host-172-16-0-183 ~]#
```

----End

5.11 Can I Stop an Ongoing Restoration Task?

No. An ongoing restoration task cannot be stopped.

6 Policies

6.1 How Do I Configure Automatic Backup for a Server or Disk?

1. Go to the Cloud Backup and Recovery console and purchase a backup vault. You are advised to set the vault capacity to at least twice the total capacity of the resources you want to back up.
2. Associate resources with the vault during or after the purchase.
3. Go to the **Policies** page to configure a backup policy. You are advised to set the backup execution time at off-peak hours, for example, early in the morning. Set the backup retention rule as needed. If your vault capacity is small, set a small value for the number of backups to be kept or the days that backups will be retained. Retention rule does not apply to manual backups.
4. Apply the policy you defined to the vault. The system then will back up the resources that are associated with the vault at the specified time and retains the backups based on the retention rule.

6.2 Why the New Retention Rule I Changed Is Not Applied?

The scenarios of a retention rule change are as follows:

Rule Type Unchanged, with Only a New Backup Quantity Configured

The new rule will be applied to the backups generated based on the old policy. After a backup is generated, regardless of an automatic or a manual one, the system verifies and uses the latest retention rule.

Example: A user has a vault associated with a disk. At 10:00 a.m. on Monday, the user applies a backup policy to the vault, based on which a backup task will be executed at 02:00 a.m. every day and three most recent backups will be kept. At 10:00 a.m. on Thursday, three backups are kept. Then the user changes the number of backups kept from three to one, and the new policy will be applied immediately. If the user then perform manual backups or wait until the system automatically create a backup at 02:00 a.m. on Friday, the system will verify and use the latest retention rule after the backup task is complete. In this case, only one most recent backup will be kept. Manual backups are not affected by policies, so they will not be deleted.

Rule Type Changed from Backup Quantity to Time Period/Permanent

The new rule will be applied only to the new backups. Backups generated based on the old policy will not be automatically deleted.

Example: A user has a vault associated with a disk. At 10:00 a.m. on Monday, the user applies a backup policy to the vault, based on which a backup task will be executed at 02:00 a.m. every day and three most recent backups will be kept. At 10:00 a.m. on Thursday, three backups are kept. Then the user changes the retention rule type from backup quantity to time period and sets to retain the backups from the last one month. The new policy will be applied immediately. If the user then perform manual backups or wait until the system automatically create a backup at 02:00 a.m. on Friday, the system will verify and use the latest retention rule after the backup task is complete. The three backups generated based on the old policy will still be kept (the number of backups does not exceed the quantity set in the old retention rule). They will not be automatically deleted and you need manually delete them if needed. Backups generated based on the new policy will be kept based on the new retention rule.

Rule Type Changed from Time Period to Time Period/Permanent

The new policy will only be applied to the new backups. Backups generated based on the old policy will be kept based on the old policy.

Example: A user has a vault associated with a disk. At 10:00 a.m. on August 5, the user applies a backup policy to the vault, based on which a backup task will be executed at 02:00 a.m. every day and the backups generated from the last one month will be kept. At 10:00 a.m. on August 8, three backups are kept. Then the user changes the backup retention time from the last one month to the last three months. At 02:00 a.m. on September 6, the backup generated on August 6 based on the old policy will be deleted. The backup generated on August 9 will be deleted two months later based on the new policy.

Rule Type Changed from Time Period to Backup Quantity

Both the old and new policies will be applied to the backups generated based on the old policy. The union set of the old and new rules will be applied.

New policy applied to old backups

Example: A user has a vault associated with a disk. At 10:00 a.m. on August 5, the user applies a backup policy to the vault, based on which a backup task will be executed at 02:00 a.m. every day and the backups generated from the last one month will be kept. At 10:00 a.m. on August 8, three backups are kept. Then the user changes the retention rule type from time period to backup quantity and sets to retain the most recent seven backups. At 10:00 a.m. on August 15, the backups generated on August 9, 10, 11, 12, 13, 14, and 15 will be kept. The backups generated on August 6, 7, and 8 have been deleted based on the new policy.

Old policy applied to old backups

Example: A user has a vault associated with a disk. At 10:00 a.m. on August 5, the user applies a backup policy to the vault, based on which a backup task will be executed at 02:00 a.m. every day and the backups generated from the last three days will be kept. At 10:00 a.m. on August 8, three backups are kept. Then the user changes the retention rule type from time period to backup quantity and sets to retain the most recent seven backups. At 10:00 a.m. on August 10, the backups generated on August 8, 9, and 10 will be kept. The backups generated on August 6 and 7 have been deleted based on the old policy.

6.3 How Do I Back Up Multiple Resources at a Time?

1. Log in to the CBR console and click **Cloud Server Backups** or **Cloud Disk Backups** on the left navigation pane. On the displayed page, purchase a backup vault. It is recommended that the capacity of the vault be at least twice the total size of resources to be backed up.
2. Associate resources with the vault during or after the purchase.
3. After the resources are associated, choose **More** > **Perform Backup** in the **Operation** column of the target vault. You can manually back up two or more resources at a time.

   Alternatively, you can set a backup policy for the vault. In this way, the system will automatically back up the associated resources at the scheduled time.

6.4 How Do I Retain My Backups Permanently?

Manual Backups

You can permanently keep backups that you manually created as long as you do not delete them and your account balance is sufficient.

Automatic Backups

To keep automatically generated backups permanently, set **Retention Rule** to **Permanent** or set the retention period to **99999** days and make sure your account balance is sufficient.

6.5 How Can I Cancel Auto Backup or Auto Replication?

To cancel auto backup or auto replication, remove the policy from the vault or disable the policy.

6.6 How Can I Have the System Automatically Delete Backups That I No Longer Need?

1. Log in to the CBR console and purchase a backup vault. Set an appropriate vault capacity by referring to **How Large of a Vault Do I Need?**.

2. Associate resources with the vault during or after the purchase.
3. Go to the **Policies** page to configure a backup policy. You are advised to set the backup execution time at off-peak hours, for example, early in the morning. Set the backup retention rule as needed. If your vault capacity is small, set a small value for the number of backups to be kept or the days that backups will be retained. Ensure that the vault has enough space to keep all backups automatically generated based on the policies before the retention rule takes effect. Or, auto backup will fail, and the quantity-based retention rule may not take effect. Retention rules are not applied to manual backups.
4. Apply the backup policy to your vault. The system will back up the resources associated with the vault at the specified time and keep backups based on the retention rule.

6.7 Why Aren't My Backups Deleted Based on the Retention Rule?

1. The policy applied to the vault is not enabled. Go to the **Policies** page to enable the policy.
2. The policy's retention rule was changed during the backup execution. For details, see **Why the New Retention Rule I Changed Is Not Applied?**
3. The backups are created manually. The policy's retention rule does not apply to manual backups. They can only be deleted manually.

7 Optimization

7.1 What Are Common Problems During Cloud-Init Installation?

You are advised to install Cloud-Init after the restoration to ensure the new server restored by using backups support custom configurations.

To install Cloud-Init, see **Installing Cloud-Init**.

To configure Cloud-Init, see **Configuring Cloud-Init**.

This section illustrates the FAQs encountered when installing Cloud-Init and their solutions.

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

- Symptom

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

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

  Information similar to the following is displayed:

  Figure 7-1 Failed to set Cloud-Init automatic start

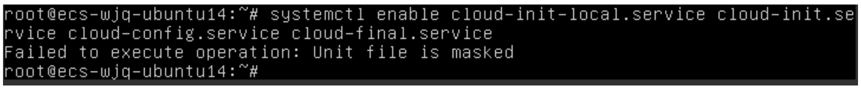

```
root@ecs-wjq-ubuntu14:~# systemctl enable cloud-init-local.service cloud-init.se
rvice cloud-config.service cloud-final.service
Failed to execute operation: Unit file is masked
root@ecs-wjq-ubuntu14:~#
```

- Solution

  a. Run the following command:

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

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

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

c. Run the following commands to check the Cloud-Init status:

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

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

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

Figure 7-2 Checking Cloud-Init status

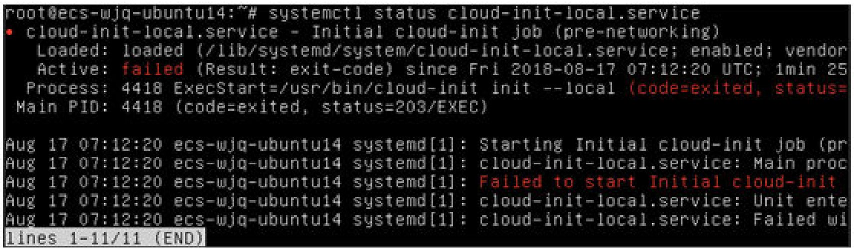

```
root@ecs-wjq-ubuntu14:~# systemctl status cloud-init-local.service
● cloud-init-local.service - Initial cloud-init job (pre-networking)
   Loaded: loaded (/lib/systemd/system/cloud-init-local.service; enabled; vendor
   Active: failed (Result: exit-code) since Fri 2018-08-17 07:12:20 UTC; 1min 25
  Process: 4418 ExecStart=/usr/bin/cloud-init init --local (code=exited, status=
 Main PID: 4418 (code=exited, status=203/EXEC)

Aug 17 07:12:20 ecs-wjq-ubuntu14 systemd[1]: Starting Initial cloud-init job (pr
Aug 17 07:12:20 ecs-wjq-ubuntu14 systemd[1]: cloud-init-local.service: Main proc
Aug 17 07:12:20 ecs-wjq-ubuntu14 systemd[1]: Failed to start Initial cloud-init
Aug 17 07:12:20 ecs-wjq-ubuntu14 systemd[1]: cloud-init-local.service: Unit ente
Aug 17 07:12:20 ecs-wjq-ubuntu14 systemd[1]: cloud-init-local.service: Failed wi
lines 1-11/11 (END)
```

Figure 7-3 Checking Cloud-Init status

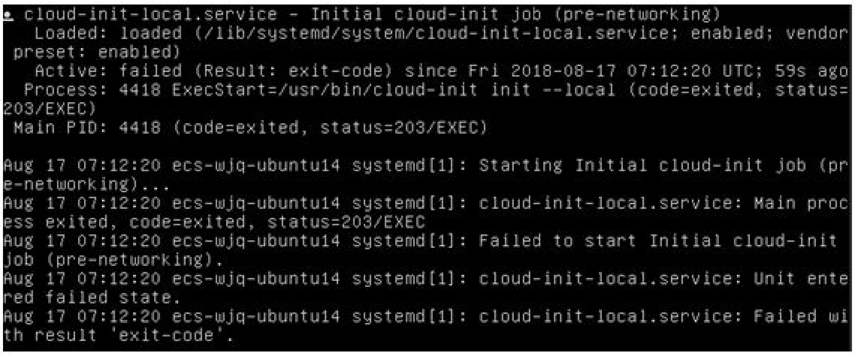

```
● cloud-init-local.service - Initial cloud-init job (pre-networking)
   Loaded: loaded (/lib/systemd/system/cloud-init-local.service; enabled; vendor
 preset: enabled)
   Active: failed (Result: exit-code) since Fri 2018-08-17 07:12:20 UTC; 59s ago
  Process: 4418 ExecStart=/usr/bin/cloud-init init --local (code=exited, status=
203/EXEC)
 Main PID: 4418 (code=exited, status=203/EXEC)

Aug 17 07:12:20 ecs-wjq-ubuntu14 systemd[1]: Starting Initial cloud-init job (pr
e-networking)...
Aug 17 07:12:20 ecs-wjq-ubuntu14 systemd[1]: cloud-init-local.service: Main proc
ess exited, code=exited, status=203/EXEC
Aug 17 07:12:20 ecs-wjq-ubuntu14 systemd[1]: Failed to start Initial cloud-init
job (pre-networking).
Aug 17 07:12:20 ecs-wjq-ubuntu14 systemd[1]: cloud-init-local.service: Unit ente
red failed state.
Aug 17 07:12:20 ecs-wjq-ubuntu14 systemd[1]: cloud-init-local.service: Failed wi
th result 'exit-code'.
```

d. Run the **cp /usr/local/cloud-init /usr/bin/** command to copy the **cloud-init** file to the **usr/bin** directory, and then run the following command to restart Cloud-Init:

   # **systemctl restart cloud-init-local.service cloud-init.service cloud-config.service cloud-final.service**

Figure 7-4 Restarting Cloud-Init

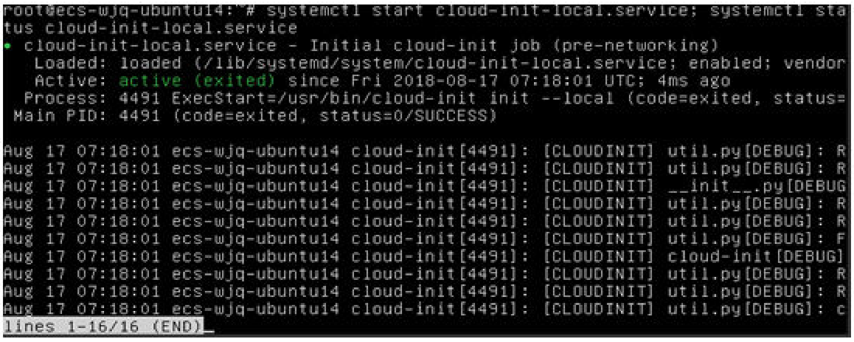

```
root@ecs-wjq-ubuntu14:~# systemctl start cloud-init-local.service; systemctl sta
tus cloud-init-local.service
● cloud-init-local.service - Initial cloud-init job (pre-networking)
   Loaded: loaded (/lib/systemd/system/cloud-init-local.service; enabled; vendor
   Active: active (exited) since Fri 2018-08-17 07:18:01 UTC; 4ms ago
  Process: 4491 ExecStart=/usr/bin/cloud-init init --local (code=exited, status=
 Main PID: 4491 (code=exited, status=0/SUCCESS)

Aug 17 07:18:01 ecs-wjq-ubuntu14 cloud-init[4491]: [CLOUDINIT] util.py[DEBUG]: R
Aug 17 07:18:01 ecs-wjq-ubuntu14 cloud-init[4491]: [CLOUDINIT] util.py[DEBUG]: R
Aug 17 07:18:01 ecs-wjq-ubuntu14 cloud-init[4491]: [CLOUDINIT] __init__.py[DEBUG
Aug 17 07:18:01 ecs-wjq-ubuntu14 cloud-init[4491]: [CLOUDINIT] util.py[DEBUG]: R
Aug 17 07:18:01 ecs-wjq-ubuntu14 cloud-init[4491]: [CLOUDINIT] util.py[DEBUG]: R
Aug 17 07:18:01 ecs-wjq-ubuntu14 cloud-init[4491]: [CLOUDINIT] util.py[DEBUG]: F
Aug 17 07:18:01 ecs-wjq-ubuntu14 cloud-init[4491]: [CLOUDINIT] cloud-init[DEBUG]
Aug 17 07:18:01 ecs-wjq-ubuntu14 cloud-init[4491]: [CLOUDINIT] util.py[DEBUG]: R
Aug 17 07:18:01 ecs-wjq-ubuntu14 cloud-init[4491]: [CLOUDINIT] util.py[DEBUG]: R
Aug 17 07:18:01 ecs-wjq-ubuntu14 cloud-init[4491]: [CLOUDINIT] util.py[DEBUG]: c
lines 1-16/16 (END)_
```

e. Run the following commands to check the Cloud-Init status:

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

Ubuntu14.04: chkconfig and systemctl Not Installed

- Symptom

  chkconfig is not installed.

- Solution

  Run the following commands to install chkconfig:

  # **apt-get update**

  # **apt-get install sysv-rc-conf**

  # **cp /usr/sbin/sysv-rc-conf /usr/sbin/chkconfig**

  After the installation completes, run the following command to query the Cloud-Init version:

  cloud-init -v

  Information similar to the following is displayed:

  ```
  -bash:/usr/bin/cloud-init:not found this command
  ```

  Solution: Run the following command to copy the **cloud-init** file to the **usr/bin** directory:

  # **cp /usr/local/bin/cloud-init /usr/bin**/

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

1. After Cloud-Init is installed, run the following command to query its version:

   cloud-init -v

   Information similar to the following is displayed:

   ```
   -bash:/usr/bin/cloud-init:not found this command
   ```

   Solution: Run the **# cp /usr/local/bin/cloud-init /usr/bin/** command to copy the **cloud-init** file to the **usr/bin** directory.

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

   Information similar to the following is displayed:

   Figure 7-5 Information returned when Cloud-Init automatic start is successfully set

   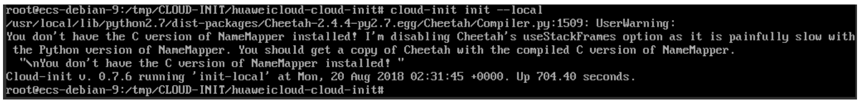

   ```
   root@ecs-debian-9:/tmp/CLOUD-INIT/huaweicloud-cloud-init# cloud-init init --local
   /usr/local/lib/python2.7/dist-packages/Cheetah-2.4.4-py2.7.egg/Cheetah/Compiler.py:1509: UserWarning:
   You don't have the C version of NameMapper installed! I'm disabling Cheetah's useStackFrames option as it is painfully slow with
    the Python version of NameMapper. You should get a copy of Cheetah with the compiled C version of NameMapper.
     "\nYou don't have the C version of NameMapper installed! "
   Cloud-init v. 0.7.6 running 'init-local' at Mon, 20 Aug 2018 02:31:45 +0000. Up 704.40 seconds.
   root@ecs-debian-9:/tmp/CLOUD-INIT/huaweicloud-cloud-init#
   ```

   Cause analysis: The compilation fails because the GNU compiler collection (GCC) is not installed.

   Solution

   After GCC is installed, run the following command to install Cloud-Init:

   yum -y install gcc

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

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

   Information similar to the following is displayed:

Figure 7-6 Failed to set Cloud-Init automatic start

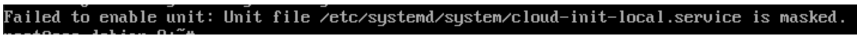

```
Failed to enable unit: Unit file /etc/systemd/system/cloud-init-local.service is masked.
```

Solution

a. Run the following command:

   # **systemctl unmask cloud-init-local.service cloud-init.service cloud-config.service cloud-final.service**

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

   # **systemctl enable cloud-init-local.service cloud-init.service cloud-config.service cloud-final.service**

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

   # **systemctl restart cloud-init-local.service cloud-init.service cloud-config.service cloud-final.service**

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

Figure 7-7 Checking the Cloud-Init status

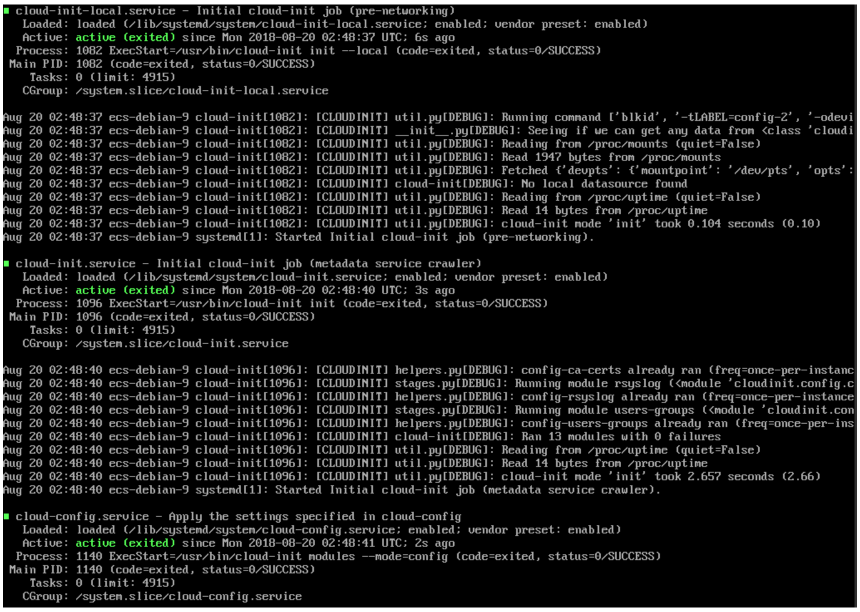

```
● cloud-init-local.service - Initial cloud-init job (pre-networking)
   Loaded: loaded (/lib/systemd/system/cloud-init-local.service; enabled; vendor preset: enabled)
   Active: active (exited) since Mon 2018-08-20 02:48:37 UTC; 6s ago
  Process: 1082 ExecStart=/usr/bin/cloud-init init --local (code=exited, status=0/SUCCESS)
 Main PID: 1082 (code=exited, status=0/SUCCESS)
    Tasks: 0 (limit: 4915)
   CGroup: /system.slice/cloud-init-local.service

Aug 20 02:48:37 ecs-debian-9 cloud-init[1082]: [CLOUDINIT] util.py[DEBUG]: Running command ['blkid', '-tLABEL=config-2', '-odevi
Aug 20 02:48:37 ecs-debian-9 cloud-init[1082]: [CLOUDINIT] __init__.py[DEBUG]: Seeing if we can get any data from <class 'cloudi
Aug 20 02:48:37 ecs-debian-9 cloud-init[1082]: [CLOUDINIT] util.py[DEBUG]: Reading from /proc/mounts (quiet=False)
Aug 20 02:48:37 ecs-debian-9 cloud-init[1082]: [CLOUDINIT] util.py[DEBUG]: Read 1947 bytes from /proc/mounts
Aug 20 02:48:37 ecs-debian-9 cloud-init[1082]: [CLOUDINIT] util.py[DEBUG]: Fetched {'devpts': {'mountpoint': '/dev/pts', 'opts':
Aug 20 02:48:37 ecs-debian-9 cloud-init[1082]: [CLOUDINIT] cloud-init[DEBUG]: No local datasource found
Aug 20 02:48:37 ecs-debian-9 cloud-init[1082]: [CLOUDINIT] util.py[DEBUG]: Reading from /proc/uptime (quiet=False)
Aug 20 02:48:37 ecs-debian-9 cloud-init[1082]: [CLOUDINIT] util.py[DEBUG]: Read 14 bytes from /proc/uptime
Aug 20 02:48:37 ecs-debian-9 cloud-init[1082]: [CLOUDINIT] util.py[DEBUG]: cloud-init mode 'init' took 0.104 seconds (0.10)
Aug 20 02:48:37 ecs-debian-9 systemd[1]: Started Initial cloud-init job (pre-networking).

● cloud-init.service - Initial cloud-init job (metadata service crawler)
   Loaded: loaded (/lib/systemd/system/cloud-init.service; enabled; vendor preset: enabled)
   Active: active (exited) since Mon 2018-08-20 02:48:40 UTC; 3s ago
  Process: 1096 ExecStart=/usr/bin/cloud-init init (code=exited, status=0/SUCCESS)
 Main PID: 1096 (code=exited, status=0/SUCCESS)
    Tasks: 0 (limit: 4915)
   CGroup: /system.slice/cloud-init.service

Aug 20 02:48:40 ecs-debian-9 cloud-init[1096]: [CLOUDINIT] helpers.py[DEBUG]: config-ca-certs already ran (freq=once-per-instanc
Aug 20 02:48:40 ecs-debian-9 cloud-init[1096]: [CLOUDINIT] stages.py[DEBUG]: Running module rsyslog (<module 'cloudinit.config.c
Aug 20 02:48:40 ecs-debian-9 cloud-init[1096]: [CLOUDINIT] helpers.py[DEBUG]: config-rsyslog already ran (freq=once-per-instance
Aug 20 02:48:40 ecs-debian-9 cloud-init[1096]: [CLOUDINIT] stages.py[DEBUG]: Running module users-groups (<module 'cloudinit.con
Aug 20 02:48:40 ecs-debian-9 cloud-init[1096]: [CLOUDINIT] helpers.py[DEBUG]: config-users-groups already ran (freq=once-per-ins
Aug 20 02:48:40 ecs-debian-9 cloud-init[1096]: [CLOUDINIT] cloud-init[DEBUG]: Ran 13 modules with 0 failures
Aug 20 02:48:40 ecs-debian-9 cloud-init[1096]: [CLOUDINIT] util.py[DEBUG]: Reading from /proc/uptime (quiet=False)
Aug 20 02:48:40 ecs-debian-9 cloud-init[1096]: [CLOUDINIT] util.py[DEBUG]: Read 14 bytes from /proc/uptime
Aug 20 02:48:40 ecs-debian-9 cloud-init[1096]: [CLOUDINIT] util.py[DEBUG]: cloud-init mode 'init' took 2.657 seconds (2.66)
Aug 20 02:48:40 ecs-debian-9 systemd[1]: Started Initial cloud-init job (metadata service crawler).

● cloud-config.service - Apply the settings specified in cloud-config
   Loaded: loaded (/lib/systemd/system/cloud-config.service; enabled; vendor preset: enabled)
   Active: active (exited) since Mon 2018-08-20 02:48:41 UTC; 2s ago
  Process: 1140 ExecStart=/usr/bin/cloud-init modules --mode=config (code=exited, status=0/SUCCESS)
 Main PID: 1140 (code=exited, status=0/SUCCESS)
    Tasks: 0 (limit: 4915)
   CGroup: /system.slice/cloud-config.service
```

CentOS 7/Fedora 28: Required C Compiler Not Installed

- Symptom

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

  cloud-init init --local

  The following information is displayed:

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

- Possible Cause

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

- Solution

  Run the following command to install GCC:

  yum -y install gcc

  Reinstall Cloud-Init.

CentOS 7/Fedora: Failed to Use the New Password to Log In to the Server Created from a Backup After Cloud-Init Is Successfully Installed

- Symptom

  After Cloud-Init is installed, the new password cannot be used to start the new server. After logging in to the server using the old password, you find the NIC is not started.

  Figure 7-8 NIC not started

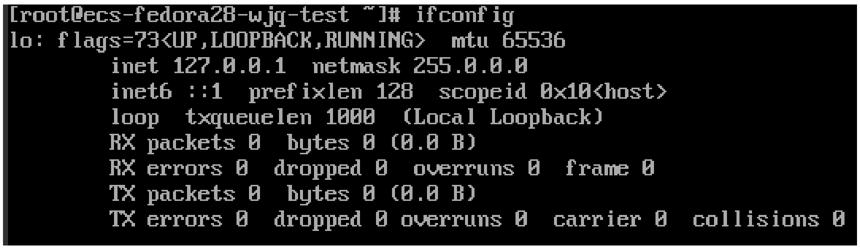

```
[root@ecs-fedora28-wjq-test ~]# ifconfig
lo: flags=73<UP,LOOPBACK,RUNNING>  mtu 65536
        inet 127.0.0.1  netmask 255.0.0.0
        inet6 ::1  prefixlen 128  scopeid 0x10<host>
        loop  txqueuelen 1000  (Local Loopback)
        RX packets 0  bytes 0 (0.0 B)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 0  bytes 0 (0.0 B)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

- Solution

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

7.2 What Can I Do If Injecting the Key or Password Using Cloud-Init Fails After NetworkManager Is Installed?

A major cause is that the version of Cloud-Init is incompatible with that of NetworkManager. In Debian 9.0 and later versions, NetworkManager is incompatible with Cloud-Init 0.7.9.

Solution

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

For details, see **Installing Cloud-Init**.

7.3 What Can Cloud-Init Do?

Cloud-Init initializes specified custom configurations, such as the host name, key, and user data, of a newly created server.

Installation Methods

If you have restored a server using a backup, it is recommended that you install Cloud-Init or Cloudbase-Init on the server.

- For Windows OSs, download and install Cloudbase-Init.

  For details, see **Installing and Configuring Cloudbase-Init**.

- For Linux OSs, download and install Cloud-Init.

  To install Cloud-init, see **Installing Cloud-Init**.

  To configure Cloud-Init, see **Configuring Cloud-Init**.

8 Others

8.1 Is There a Quota for CBR Vaults?

There are no quotas on CBR vaults, except that a maximum of 10 hybrid cloud backup vaults can be created. You can create as many vaults as needed.

8.2 Can I Merge My Vaults?

No. Vaults cannot be merged.

8.3 How Do I Delete a Backup That Has Been Used to Create an Image While Retaining the Image?

Use the image to create a server and the server to create another image. Delete the original image and then you can delete the backup.

8.4 Can I Export Disk Backup Data to Another Server?

You can export disk backup data by creating a new disk using a disk backup and then attaching the new disk to a server.

8.5 Why Do I Need a Vault to Accept the Image Shared to Me?

Before accepting a shared full-ECS image, you need a vault to store the image. Later, this vault is used to store the ECSs provisioned.

An accepted full-ECS image does not occupy the vault space. Do not delete this vault. Or, ECSs will fail to be provisioned using the accepted image.

8.6 Can I Download Backup Data to a Local PC?

No. CBR backup data cannot be downloaded to a local PC.

8.7 How Do I Copy Disk Data to Another Account?

If the two accounts are in the same region, you can use CBR backup sharing to copy disk data between accounts. For details, see **Sharing a Backup**. Cross-region backup sharing is currently not supported.