Cloud Backup and Recovery

FAQs

Issue 03

Date 2020-04-08

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions


and other Huawei trademarks are trademarks 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 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 Concepts

1.1 What Are Full Backup and Incremental Backup?

After an initial full backup, a server continues to be backed up incrementally by default.

- The initial full backup covers only the used capacity of a disk. If a 100 GB disk contains 40 GB data, the initial backup consumes 40 GB backup space.
- Subsequent incremental backup backs up data changed since the last backup. If 5 GB data changed since the last backup, only the 5 GB changed data will be backed up.

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.

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, not by calculating 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	DR Range	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.	A backup and its source disk reside in the same AZ. Cloud server backups can be replicated across regions.	Data can be recovered and services can be restored by restoring the backup data to original disks or creating new disks from backups, ensuring excellent data reliability.
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.	A snapshot and its source disk reside in the same AZ.	You can use a snapshot to roll back its original disk or create a disk for data restoration and service recovery.

1.4 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:

- CBR in-cloud backups are 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.

1.5 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.
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	CBR applies to the following scenarios: • Data backup and restoration • Rapid service deployment and migration	IMS applies to the following scenarios: • Server migration to the cloud or between clouds • Deploying a specific software environment • Deploying software environments in batches • Backing up server operating environments

Item	CBR	IMS
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 CBR's billing standard, and image storage fees are charged according to OBS's billing standard.

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. Therefore, the size of the generated image is smaller than that of the backup.

1.6 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.

Item	Cloud Server Backup	Cloud Disk Backup
Advantages	All disks on a server are backed up at the same time, ensuring data consistency.	Backup cost is reduced without compromising data security.

1.7 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 1-1 Backup mechanism

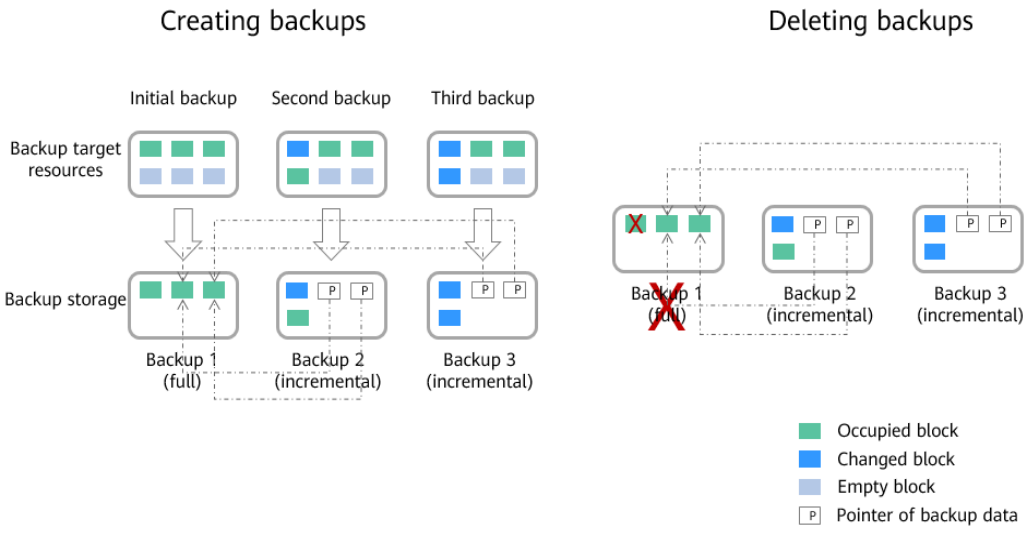

2 Billing

2.1 How Is CBR Billed?

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 generated.	Pay-per-use Yearly/ Monthly
	Server backup vault	If cloud servers (applications not included) need to be backed up, buy server backup vaults to store the backups generated.	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 generated.	Pay-per-use Yearly/ Monthly
	Database server backup vault	If cloud servers (applications included) need to be backed up, buy database server backup vaults to store the backups generated. How to purchase: Enable **Application-Consistent Backup** on the **Buy Server Backup Vault** page. For more information, see **Application-Consistent Backup Overview**.	Pay-per-use Yearly/ Monthly

Category	Billing Item	Description	Billing Mode
	Hybrid cloud backup vault	If on-premises VMware VMs and OceanStor Dorado arrays need to be backed up, buy hybrid cloud backup vaults to store the backups generated.	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 in-cloud, hybrid cloud backups are used to restore the data in on-premises IDCs, outbound traffic over the Internet is generated.	Free for a limited time
	Cross-region replication traffic	If backups or vaults are replicated to another region, cross-region replication traffic is generated in the source region.	Pay-per-use Yearly/ Monthly

Billing Examples

Case 1:

Pay-per-use vault for cloud servers without databases deployed:

A user has a 100 GB cloud server and a 400 GB server backup vault purchased in the CN North-Beijing4 region, and the cloud server is associated with the vault. The user is billed for the 400 GB server backup vault in CBR.

Case 2:

Pay-per-use vault for cloud servers with databases deployed:

A user has a 100 GB cloud server running databases and an 800 GB database server backup vault purchased in the CN North-Beijing4 region, and the cloud server is associated with the vault. The user is billed for the 800 GB database server backup vault in CBR.

Case 3:

Replicating a backup to another region, with pay-per-use billing:

A user purchases a 100 GB server backup vault A in the CN North-Beijing4 region, and the backup data uses 40 GB of the storage space. This user also purchases a 200 GB replication vault B in the CN South-Guangzhou 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 and 200 GB vaults, 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, use the following formula to estimate the required capacity of the vault to be created:

Vault capacity (GB) = [Disk capacity (GB) + Backup retention period (days)/ Backup cycle (days) x Daily data volume changes (GB)] x 120%

NOTE

If you set that backups are retained based on backup quantity, you can calculate the vault capacity you need by converting the number of retained backups to backup retention period. For example, if a company backs up data once a day and the number of retained backups is set to 7, then the retention period can be deemed as 7 days.

You are advised to adjust the value based on your actual needs after the calculation.

Example

Financial company A has an 800 GB cloud server and has used 200 GB of it. Use the disk capacity (800 GB) rather that the used capacity for calculation. The daily data change volume is 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 one month. The capacity of a server backup vault required by the company can be estimated using the following formula:

Required vault capacity = [(800+30/(1/2) x 10] x 120% = 1,680 GB

2.3 What Is the Billing Cycle and Pay Time of the Pay-per-Use Mode?

The vaults are billed hourly and are paid after usage.

2.4 How Do I Disable CBR?

If you have enabled CBR when purchasing an ECS but want to disable it afterward, go to 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

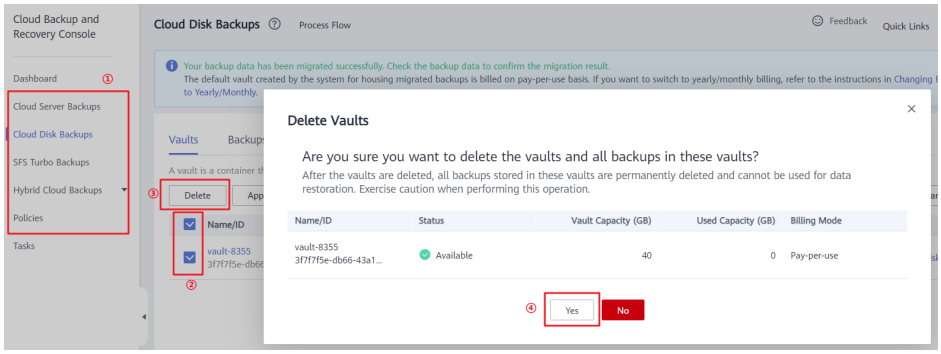

2.5 How Do I Unsubscribe from CSBS or VBS Resource Packages?

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 **Unsubscriptions and Changes** > **Unsubscriptions**. The resources you purchased are displayed.

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

Figure 2-2 Unsubscribing from a resource package

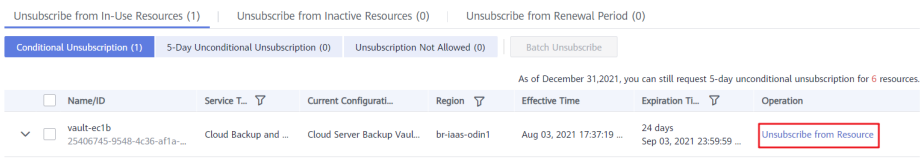

----End

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 **Billing Center** on the **Select Ticket Type** tab page.

Figure 2-3 Going to the Subscriptions page

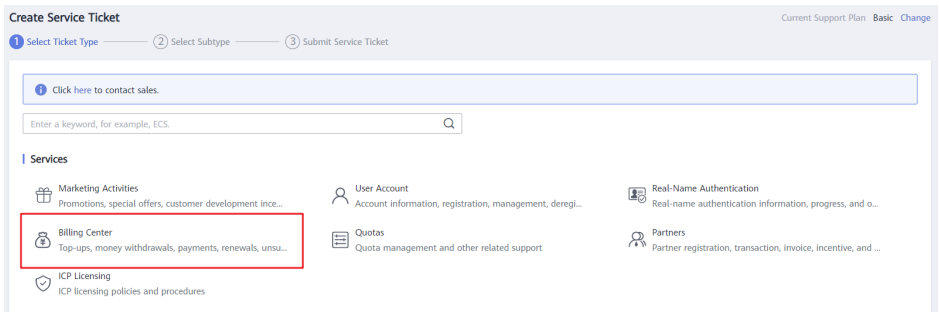

Step 3 On the **Select Subtype** tab page, choose **Unsubscriptions** > **Create Service Ticket**.

Figure 2-4 Creating a service ticket

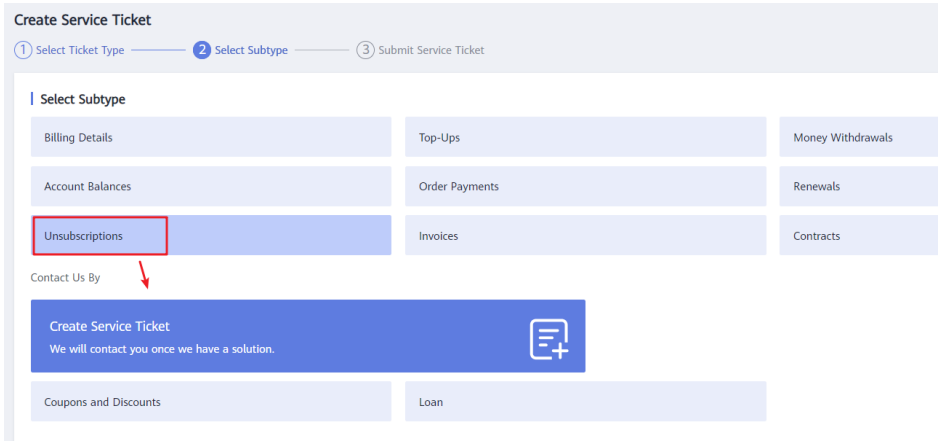

Step 4 Enter a description, for example, "unsubscribe from CSBS/VBS resource packages", in the text box next to **Problem Description** and set other parameters as required. Then click **Submit**.

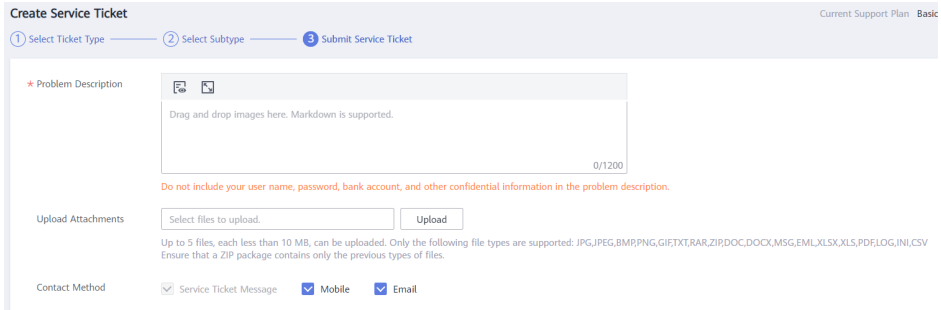

----End

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 **Grace Period and Retention Period**. 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 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 CBR Console.

Step 2 Click **Billing** in the top navigation bar to go to the **Billing Center** page.

Step 3 In the left navigation pane, choose **Unsubscriptions and Changes** > **Unsubscriptions**. The resources you purchased are displayed.

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

Figure 2-5 Unsubscribing from a resource package

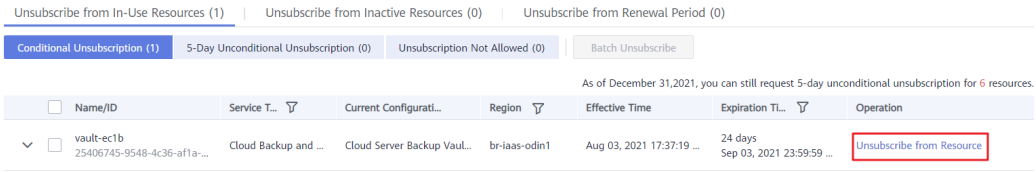

----End

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 unwanted vaults 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 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 prices, see the **CBR pricing details**.

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 onto disks on the server, and some newly generated data is stored in the server memory as cached data. During a backup task, the data in the memory will not be automatically written onto disks, resulting in data inconsistency between disks and their backups.

To ensure data integrity, back up the server during off-peak hours when no write operation is performed on 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	Windows	Windows Server 2008, 2008 R2, 2012, and 2012 R2 for x86_64

Database	OS	Version
SQL Server 2014/2016/EE	Windows	Windows Server 2012, 2012 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 and 12 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 customized 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?

Background Information

A user has an ECS with only one system disk in CN North-Beijing4. 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 CN North-Beijing4 to CN South-Guangzhou, the user needs to back up the ECS in CN North-Beijing4, replicates the backup to CN

South-Guangzhou, uses the backup replica to create a full-ECS image, and then uses the image to create an ECS in CN South-Guangzhou. This way, data on the ECS has been migrated to CN South-Guangzhou.

Procedure

Step 1 Log in to the ECS console, switch to the CN North-Beijing4 region, and locate the target ECS in the server list. See **Figure 3-1**.

Figure 3-1 Target ECS information

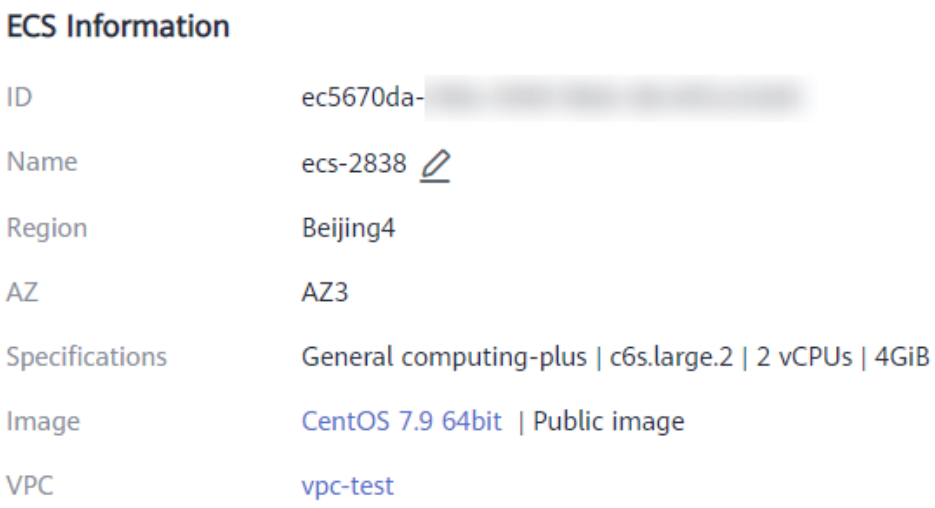

Step 2 Choose **More** > **Create Backup** in the **Operation** column. Switch to CBR Console to purchase a vault. Make sure that the target ECS is selected when purchasing a vault.

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

Figure 3-2 Performing the backup

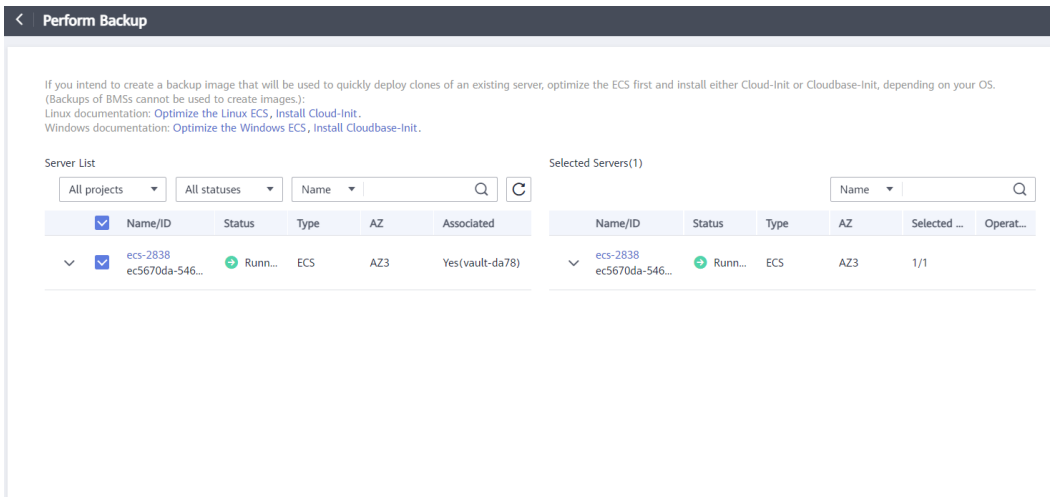

Step 4 Switch to the cloud server backup page on CBR Console, and confirm that the backup is successfully created in the destination vault. Then, choose **More** >

Create Replica in the **Operation** column to replicate the backup to CN South-Guangzhou. 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

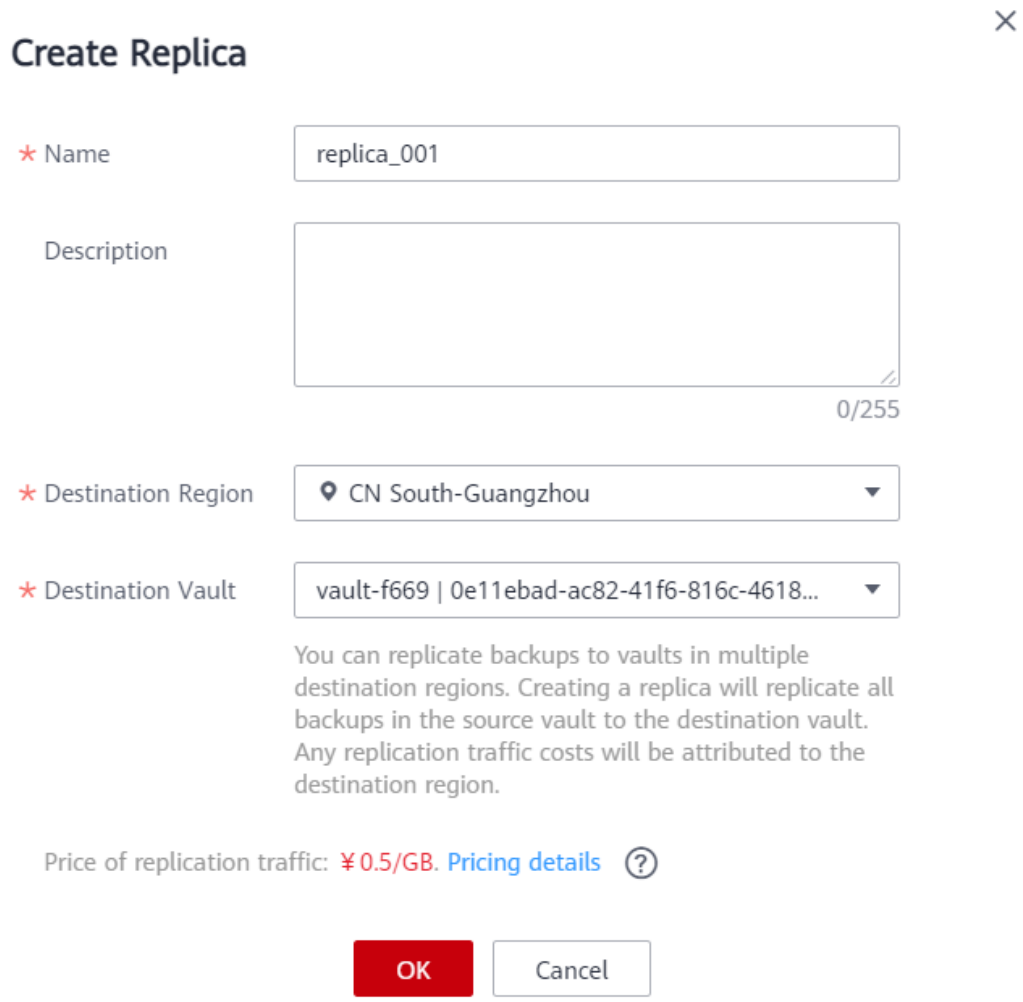

Step 5 Switch to the CN South-Guangzhou region. Choose **Cloud Server Backup**, find the replicated backup **replica_001** in the backup list. An **R** icon is displayed next to the backup name. In the **Operation** column, click **Create Image**.

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

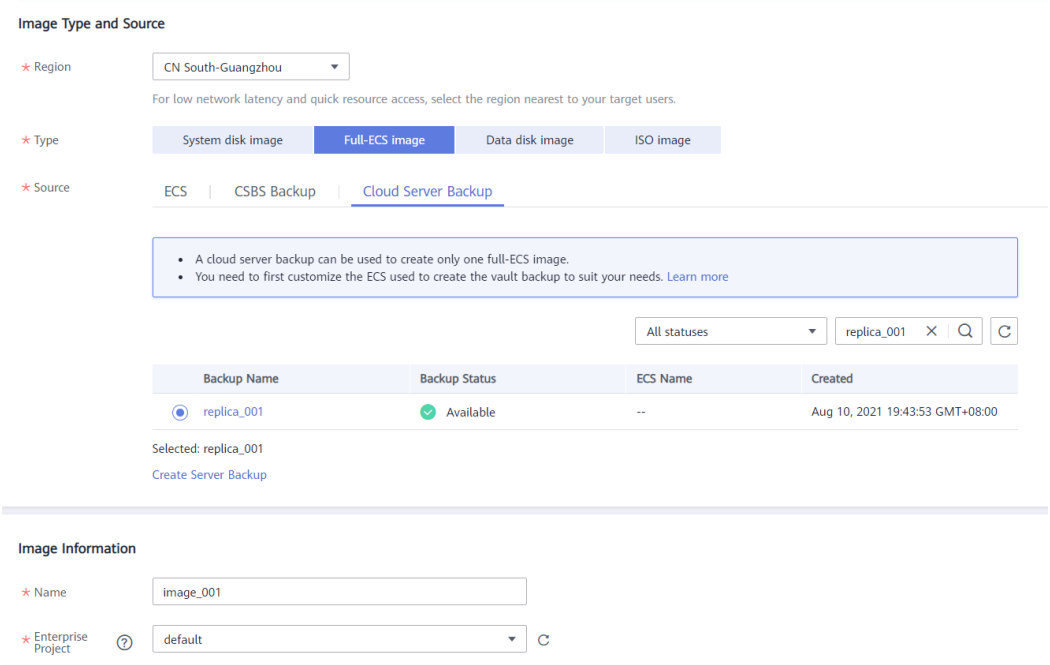

Step 6 On the IMS console in CN South-Guangzhou, find the new image created and click **Apply for Server** in the **Operation** column to create a server. See **Figure 3-5**. Set other parameters based on service requirements.

Figure 3-5 Creating a server

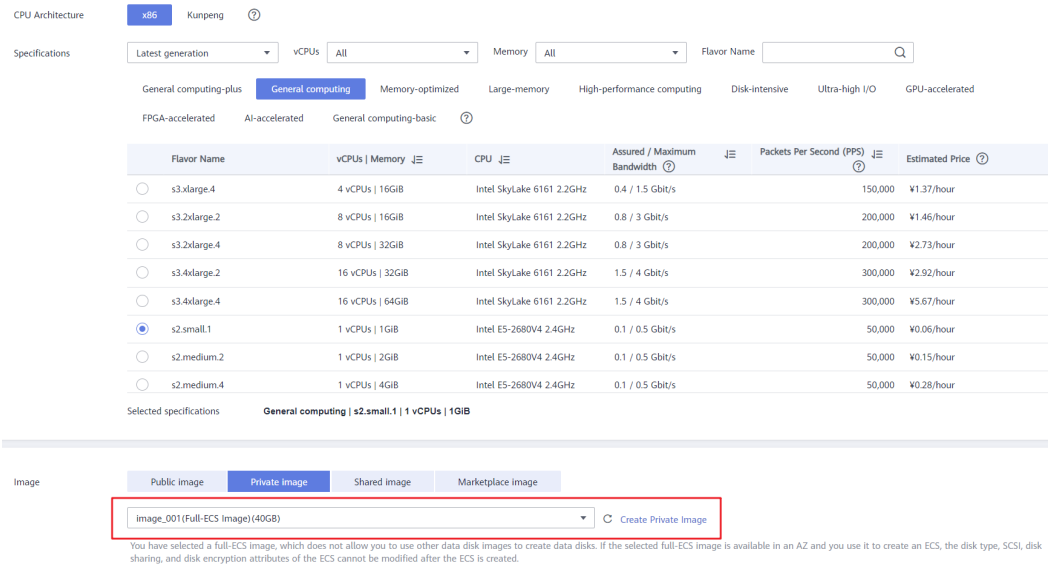

Step 7 View the server details page. It can be seen that the ECS data has been migrated from CN North-Beijing4 to CN South-Guangzhou. See **Figure 3-6**.

Figure 3-6 Successful cross-region ECS data migration

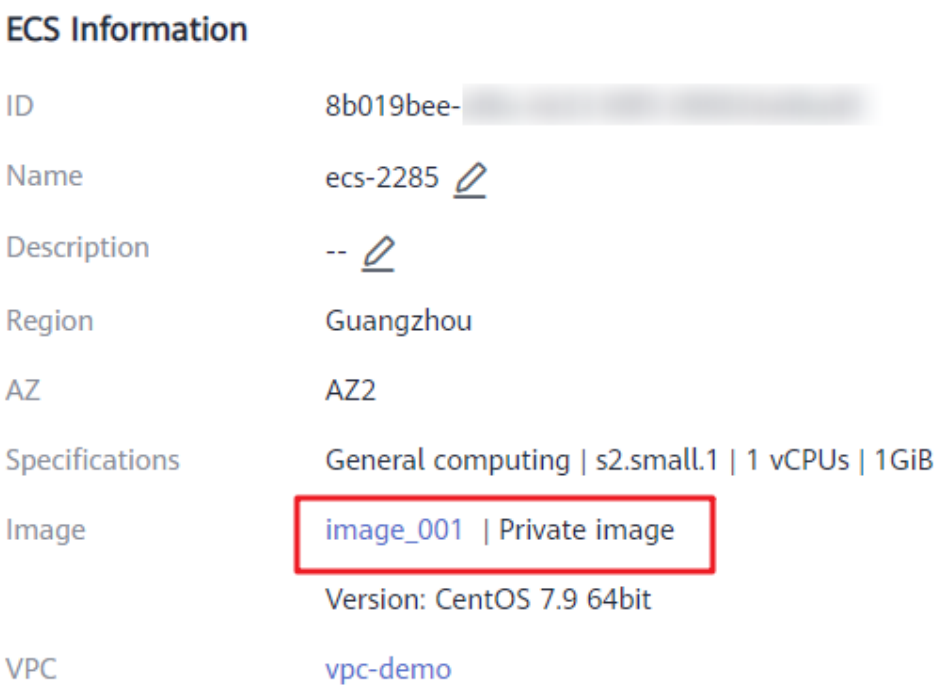

----End

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?

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 data, including the invisible data, on the disk 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

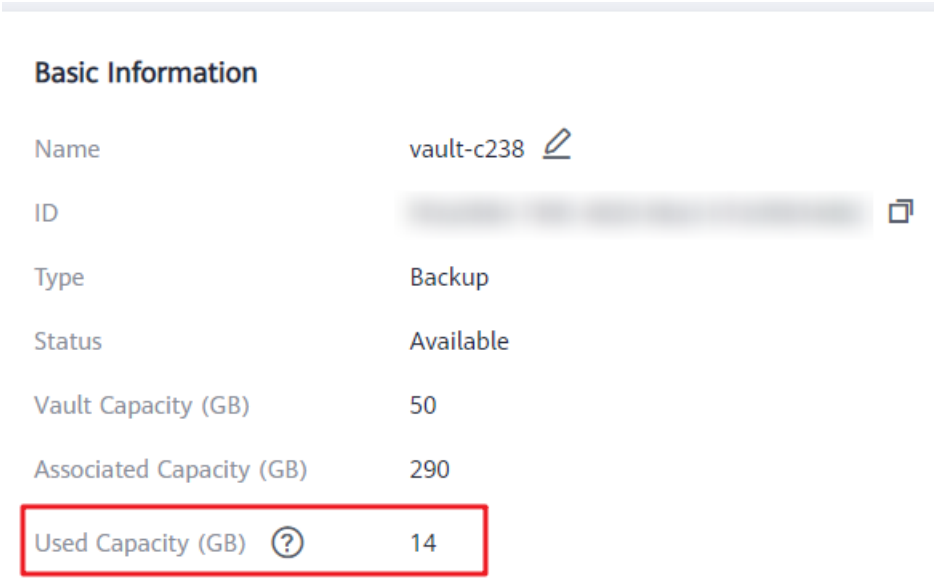

3.17 How Do I View My Backup Data?

Backup data cannot be viewed on CBR Console. You can check your backup data in the following ways:

Server Backups

To view server backup data, create an image from a server backup, use the image to create a server, and then log in to the server to view your data.

Disk Backups

To view disk backup data, create a new disk from a disk backup, attach the disk to a server, and then log in to the server to view your data.

SFS Turbo Backups

To view SFS Turbo backup data, create a new SFS Turbo file system from an SFS turbo backup, mount the file system to a server, and then log in to the server to view your data.

4 Restoration

4.1 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.

4.2 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.

4.3 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.

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

Yes. If a server has been backed up and changed such as 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 backup, using the backup to restore data will not change the data on the newly added disk.

If you have deleted a disk after backup, data restoration using the backup will not include the deleted disk.

4.5 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*.

4.6 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 **Resetting the Password for Logging In to an ECS** in the *Elastic Cloud Server User Guide*.

4.7 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.

4.8 How Do I Restore Data on the Original Server to a New Server?

To restore backup data to a new server, you can create an image using a server backup and then use the image to create a new server. For details, see **Using a Backup to Create an Image**.

If a new server has already been created, you can back up the disks of the original server, create new disks using the backups, and then attach the new disks to the new server. In this case, however, data consistency may not be ensured.

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

You can use a disk backup to create a new disk following instructions in **Using a Backup to Create a Disk** and attach the new disk to a server following instructions in **Attaching a Non-Shared Disk**. Then copy data in the data disk to the system disk.

4.10 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 4-1 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 CBR Console and create a backup for the server.

Figure 4-2 Creating a backup for the server

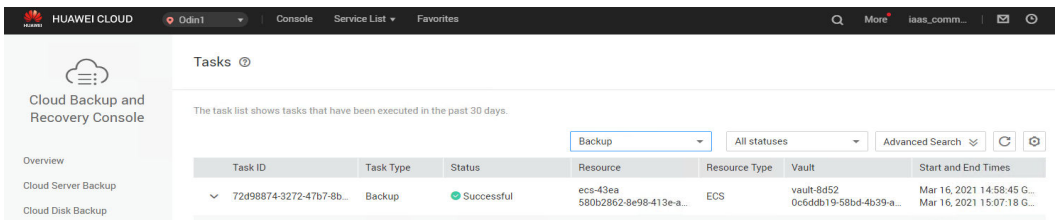

Step 3 Log in to the server again and delete the **test1** file.

Figure 4-3 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 CBR Console, use the server backup you created to restore data.

Figure 4-4 Restoring data

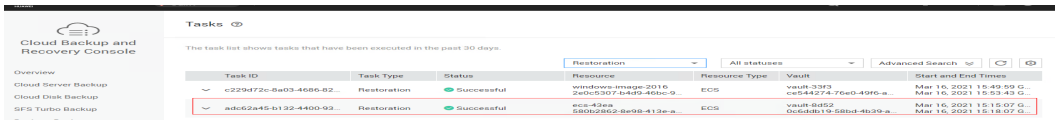

Step 5 Log in to the server and confirm that the data has been restored to the state when the backup was created.

Figure 4-5 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

4.11 Can I Stop an Ongoing Restoration Task?

No. An ongoing restoration task cannot be stopped.

5 Policies

5.1 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. After the resources are associated, go to the **Policies** page to configure a backup policy. You are advised to back up data during 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.

5.2 Why Does the Retention Rule Not Take Effect After Being Modified?

A retention rule does not take effect immediately after being modified in the following scenarios:

The Retention Rule Remains to Be by Backup Quantity, With Only the Configured Quantity Changed

The new retention rule does not take effect on the backups generated by the original backup policy, but only on the new backups. The backups generated by the original backup policy still follow the original retention rule.

For example: At 10:00 a.m. on Monday, a user sets the backup execution time of policy A to 02:00 a.m. every day, sets the rule to retaining three backups, and

associates a disk. At 10:00 a.m. on Thursday, three backups are retained. Then the user modifies the retention rule to retaining one backup. However, this new retention rule does not take effect immediately. Instead, it retains only the generated backup at 02:00 a.m. on Friday. At then, the three backups generated by the original policy have been deleted.

The Retention Rule Is Changed From by Backup Quantity to by Retention Period or Permanent

The new retention rule does not take effect on the backups generated by the original backup policy. The backups generated by the original backup policy will not be automatically deleted.

For example: At 10:00 a.m. on Monday, a user sets the backup execution time of policy A to 02:00 a.m. every day, sets the rule to retaining three backups, and associates a disk. At 10:00 a.m. on Thursday, three backups are retained. Then the user modifies the retention rule to deleting the backup generated one month ago. At 02:00 a.m. on Friday, the backup generated by the new backup policy is retained and the three backups generated by the original policy still exist.

The Retention Rule Is Changed From by Retention Period to By a New Retention Period or Permanent

The new retention rule does not take effect on the backups generated by the original backup policy, but only on the new backups. The backups generated by the original backup policy still follow the original retention rule.

For example: At 10:00 a.m. on August 5, a user sets the backup execution time of policy A to 02:00 a.m. every day, sets the retention rule to deleting the backup generated one month ago, and associates a disk. At 10:00 a.m. on August 8, three backups are retained. Then the user modifies the retention rule to deleting the backup generated three months ago. At 02:00 on September 6, the backup generated by the original backup policy on August 6 is deleted, and the backup generated on August 9 will be deleted two months later.

The Retention Rule Is Changed From by Retention Period to by Quantity

The new retention rule does not take effect on the backups generated by the original backup policy, but only on the new backups. The backups generated by the original backup policy still follow the original retention rule.

For example: At 10:00 a.m. on August 5, a user sets the backup execution time of policy A to 02:00 a.m. every day, sets the retention rule to deleting the backup generated one month ago, and associates a disk. At 10:00 a.m. on August 8, three backups are retained. Then the user modifies the retention rule to retaining one backup. At 10:00 on August 10, the backup generated on August 9 has been deleted and the backup generated on August 10 is retained. However, the backups generated on August 6, 7, and 8 still exist. At 02:00 on September 6, the backup generated on August 6 is deleted.

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

1. Log in to CBR Console and click **Cloud Server Backup** or **Cloud Disk Backup** 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.

5.4 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 Optimization

6.1 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.

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

For details about how 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 6-1 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 6-2 Checking Cloud-Init status

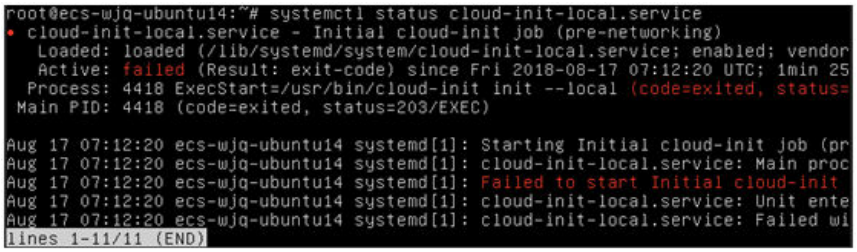

Figure 6-3 Checking Cloud-Init status

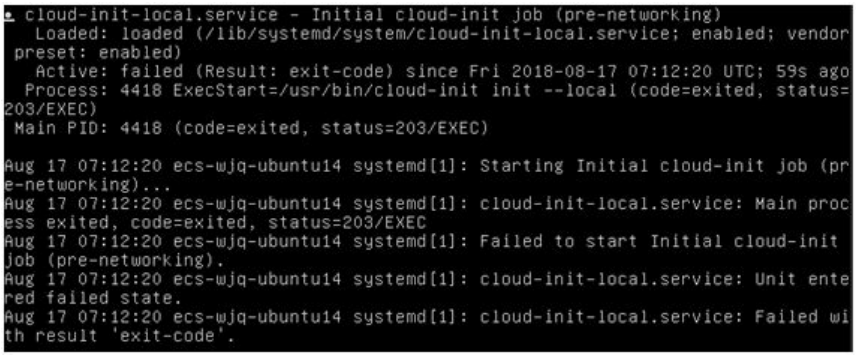

d. Run the **cp /usr/local/cloud-init /usr/bin/** command to copy the **cloud-init** file to the **usr/bin** directory, and then run the following command to restart Cloud-Init:

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

Figure 6-4 Restarting Cloud-Init

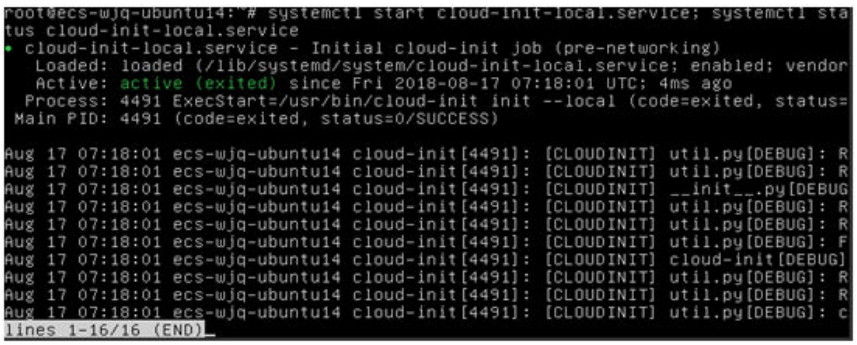

e. Run the following commands to check the Cloud-Init status:

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 6-5 Information returned when Cloud-Init automatic start is successfully set

   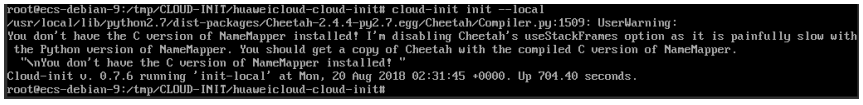

   Cause analysis: The compilation fails because the GNU compiler collection (GCC) is not installed.

   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 6-6 Failed to set Cloud-Init automatic start

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 6-7 Checking the Cloud-Init status

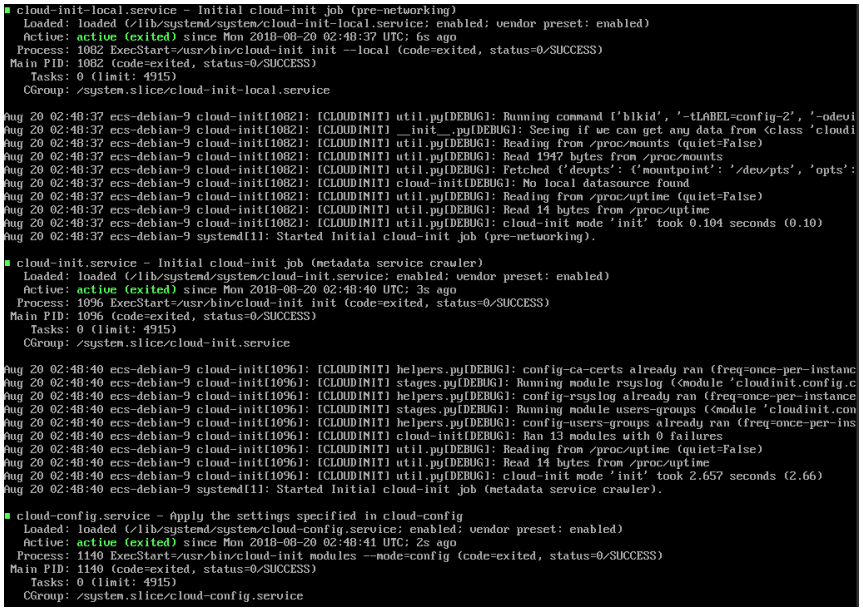

CentOS 7/Fedora 28: Required C Compiler Not Installed

- 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!
  ```

- Cause analysis

  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 6-8 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**.

6.2 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**.

6.3 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

You are advised to install Cloud-Init or Cloudbase-Init on the server restored by using backups.

- For Windows OSs, download and install Cloudbase-Init.
  For details about how to install Cloudbase-Init, see **Installing Cloudbase-Init**.
- For Linux OSs, download and install Cloud-Init.
  For details about how to install Cloud-Init, see **Installing Cloud-Init**.
  For details about how to configure Cloud-Init, see **Configuring Cloud-Init**.

7 Others

7.1 Is There a Quota for CBR Vaults?

No. You can create as many vaults as needed.

7.2 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.

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

If the storage capacity of a vault is used up, the system will not continue to back up your resources. New backups will never overwrite previous backups. Take the following measures when the storage capacity of your vault is not enough:

1. Locate the target vault and delete unwanted backups by following instructions in **Deleting a Backup**.
2. If you want to retain the generated backups, expand the vault. For details, see **Expanding a Vault**.
3. If a backup policy has been applied to the vault, disable the backup policy or remove the policy from the vault. For details, see **Disabling a Backup Policy** and **Removing a Policy from a Vault**. Then, automatic backup is disabled, and the storage capacity of the vault will not change. Alternatively, you can prolong the automatic backup interval or reduce the number of backups to be retained by editing the backup policy, or reduce the number of servers associated with the vault.

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

If the usage of a vault just reached the upper limit, or has not reached yet but its remaining capacity is insufficient for the next backup, backup can still be executed for once.

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

7.5 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.

7.6 Can I Download Backup Data to a Local PC?

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

7.7 How Do I Copy Disk Data to Another Account?

If two accounts are in the same region, you can use backup sharing of CBR to copy disk data to another account. For details, see **Sharing a Backup**. If two accounts are not in the same region, data sharing between accounts is not allowed.

7.8 What Should I Do If I Receive an Email or SMS Message Indicating a Backup Failure?

If you receive such an email or SMS message, a backup task has failed in one or more vaults in a region.

There is a delay in sending emails or SMS messages. You will receive the notification the next day after the backup fails. If you have dealt with the capacity issue (expanding capacity or taking other measures) and ensure that the vault capacity is normal, ignore the email and SMS message.

Log in to CBR Console, go to the region prompted in the notification, and check whether there are vaults whose progress bars in the **Vault Capacity** column are displayed in red. If yes, take measures mentioned in **What Can I Do If the Vault Capacity Is Insufficient?**

