

Cloud Backup and Recovery

User Guide

Issue 01
Date 2024-02-01



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

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

Trademarks and Permissions



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

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

Notice

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

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

Contents

1 Service Overview.....	1
1.1 What Is CBR?.....	1
1.2 Advantages.....	4
1.3 Application Scenarios.....	5
1.4 Functions.....	5
1.5 User Permissions.....	7
1.6 Constraints.....	7
1.7 CBR and Other Services.....	8
1.8 Basic Concepts.....	9
1.8.1 CBR Concepts.....	9
1.8.2 Project and Enterprise Project.....	10
1.8.3 Region and AZ.....	10
2 Getting Started.....	12
2.1 Overview.....	12
2.2 Step 1: Create a Vault.....	13
2.2.1 Creating a Server Backup Vault.....	13
2.2.2 Creating a Disk Backup Vault.....	15
2.2.3 Creating an SFS Turbo Backup Vault.....	17
2.3 Step 2: Associate a Resource with the Vault.....	19
2.4 Step 3: Create a Backup.....	20
2.4.1 Creating a Cloud Server Backup.....	20
2.4.2 Creating a Cloud Disk Backup.....	22
2.4.3 Creating an SFS Turbo Backup.....	23
3 Vault Management.....	26
3.1 Querying a Vault.....	26
3.2 Deleting a Vault.....	28
3.3 Dissociating a Resource.....	29
3.4 Migrating a Resource.....	29
3.5 Expanding Vault Capacity.....	30
4 Backup Management.....	32
4.1 Viewing a Backup.....	32
4.2 Sharing a Backup.....	33

4.3 Deleting a Backup.....	35
5 Policy Management.....	37
5.1 Viewing the Policy of a Vault.....	37
5.2 Creating a Backup Policy.....	37
5.3 Modifying a Policy.....	41
5.4 Deleting a Policy.....	41
5.5 Applying a Policy to a Vault.....	42
5.6 Removing a Policy from a Vault.....	42
6 Restoring Data.....	44
6.1 Restoring from a Cloud Server Backup.....	44
6.2 Creating an Image from a Cloud Server Backup.....	46
6.3 Restoring from a Cloud Disk Backup.....	47
6.4 Using a Backup to Create a Disk.....	48
6.5 Using a Backup to Create a File System.....	49
7 Managing Tasks.....	50
8 Monitoring.....	51
8.1 CBR Metrics.....	51
9 Quotas.....	53
10 FAQs.....	54
10.1 Concepts.....	54
10.1.1 What Are Full Backup and Incremental Backup?.....	54
10.1.2 What Are the Differences Between Backup and Disaster Recovery?.....	55
10.1.3 What Are the Differences Between Backups and Snapshots?.....	55
10.1.4 What Are the Differences Between Backups and Images?.....	56
10.1.5 What Are the Differences Between Cloud Server Backup and Cloud Disk Backup?.....	58
10.2 Billing.....	59
10.2.1 What Is the Billing Cycle of Pay-per-Use Vaults?.....	59
10.2.2 What Charges Will I Incur When Using a Backup to Create an Image?.....	59
10.3 Backup.....	59
10.3.1 Do I Need to Stop the Server Before Performing a Backup?.....	59
10.3.2 Can I Back Up a Server Deployed with Databases?.....	59
10.3.3 How Can I Distinguish Automatic Backups From Manual Backups?.....	59
10.3.4 Can I Choose to Back Up Only Some Partitions of a Disk?.....	60
10.3.5 Does CBR Support Cross-Region Backup?.....	60
10.3.6 Can I Back Up Two Disks to One Target Disk?.....	60
10.3.7 How Do I Replicate a Disk to the Same AZ in a Region as the Source Disk?.....	60
10.3.8 Will the Server Performance Be Affected If I Delete Its Backups?.....	60
10.3.9 Can I Use Its Backup for Restoration After a Resource Is Deleted?.....	60
10.3.10 How Many Backups Can I Create for a Resource?.....	60
10.3.11 Can I Use an Incremental Backup to Restore Data After a Full Backup Is Deleted?.....	60

10.3.12 Can I Stop an Ongoing Backup Task?.....	60
10.3.13 How Do I Reduce the Vault Space Occupied by Backups?.....	61
10.3.14 How Do I View the Size of Each Backup?.....	61
10.3.15 How Do I View My Backup Data?.....	61
10.3.16 How Long Will My Backups Be Kept?.....	62
10.4 Capacity.....	62
10.4.1 How Large of a Vault Do I Need?.....	62
10.4.2 Why Is My Backup Size Larger Than My Disk Size?.....	63
10.4.3 What Can I Do If the Vault Capacity Is Not Enough?.....	64
10.4.4 Why Does the Used Capacity of a Vault Change Only Slightly After I Deleted Unwanted Backups?	64
10.4.5 Will Backup Continue If the Usage of a Vault Reaches the Upper Limit?.....	65
10.5 Restoration.....	65
10.5.1 Do I Need to Stop the Server Before Restoring Data Using Backups?.....	65
10.5.2 Can I Use a System Disk Backup to Recover an ECS?.....	65
10.5.3 Do I Need to Stop the Server Before Restoring Data Using Disk Backups?.....	66
10.5.4 Can a Server Be Restored Using Its Backups After It Is Changed?.....	66
10.5.5 Can a Disk Be Restored Using Its Backups After Its Capacity Is Expanded?.....	66
10.5.6 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?.....	66
10.5.7 What Changes Will Be Made to the Original Backup When I Use the Backup to Restore a Server?	66
10.5.8 How Do I Restore Data to a New Server?.....	67
10.5.9 Can I Stop an Ongoing Restoration Task?.....	67
10.6 Policies.....	67
10.6.1 How Do I Configure Automatic Backup for a Server or Disk?.....	67
10.6.2 Why the New Retention Rule I Changed Is Not Applied?.....	68
10.6.3 How Do I Back Up Multiple Resources at a Time?.....	69
10.6.4 How Do I Retain My Backups Permanently?.....	69
10.6.5 How Can I Cancel Auto Backup?.....	70
10.6.6 How Can I Have the System Automatically Delete Backups That I No Longer Need?.....	70
10.6.7 Why Aren't My Backups Deleted Based on the Retention Rule?.....	70
10.7 Optimization.....	70
10.7.1 What Are Common Problems During Cloud-Init Installation?.....	70
10.7.2 What Can I Do If Injecting the Key or Password Using Cloud-Init Fails After NetworkManager Is Installed?.....	75
10.7.3 What Can Cloud-Init Do?.....	75
10.8 Others.....	75
10.8.1 Is There a Quota for CBR Vaults?.....	75
10.8.2 Can I Merge My Vaults?.....	75
10.8.3 How Do I Delete a Backup That Has Been Used to Create an Image While Retaining the Image?	76
10.8.4 Can I Export Disk Backup Data to Another Server?.....	76
10.8.5 Why Do I Need a Vault to Accept the Image Shared to Me?.....	76

10.8.6 Can I Download Backup Data to a Local PC?.....	76
10.8.7 How Do I Copy Disk Data to Another Account?.....	76
11 Troubleshooting Cases.....	77
11.1 Failed to Execute a Backup Task.....	77
11.2 Failed to Delete a Backup.....	78
11.3 Failed to Attach Disks.....	79
11.4 Data Disks Are Not Displayed After a Windows Server Is Restored.....	80
A Appendix.....	83
A.1 Change History.....	83

1 Service Overview

1.1 What Is CBR?

Overview

Cloud Backup and Recovery (CBR) enables you to easily back up Elastic Cloud Servers (ECSs), Elastic Volume Service (EVS) disks, and SFS Turbo file systems. In case of a virus attack, accidental deletion, or software or hardware fault, you can use the backup to restore data to any point when the data was backed up.

CBR Architecture

CBR involves backups, vaults, and policies.

Backup

A backup is a copy of a particular chunk of data and is usually stored elsewhere so that it may be used to restore the original data in the event of data loss. There are the following types of backups:

- Cloud disk backup: provides snapshot-based backups for EVS disks.
- Cloud server backup: uses the consistency snapshot technology to protect data for ECSs. Backups of non-database servers are non-database server backups, and those of database servers are application-consistent backups.
- SFS Turbo backup: backs up data of SFS Turbo file systems.

Vault

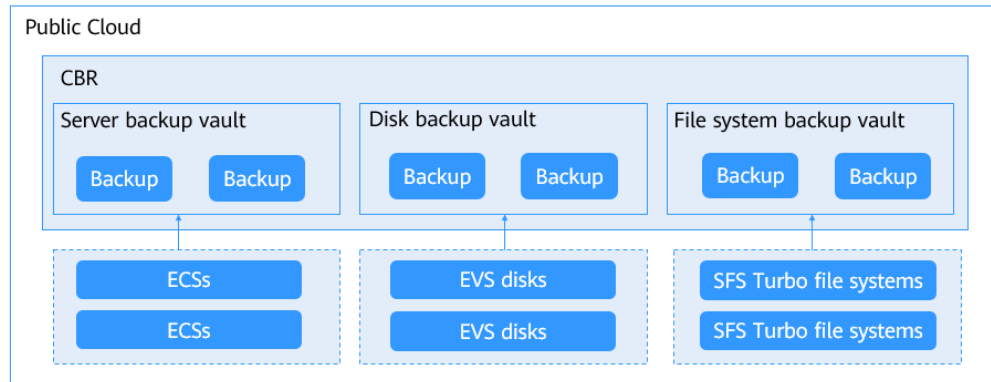
CBR stores backups in vaults. Before creating a backup, you need to create at least one vault and associate the resources you want to back up with the vaults. Then the resources can be backed up to the associated vaults.

Different types of resources must be backed up to different types of vaults. For example, cloud servers must be backed up to server backup vaults, not disk backup vaults or any other types of vaults.

Policy

- A backup policy defines when you want to take a backup and for how long you would retain each backup.

Figure 1-1 CBR architecture



Differences Among the Backup Types

Table 1-1 Differences among the backup types

Item	Cloud Server Backup	Cloud Disk Backup	SFS Turbo Backup
What to back up	All disks (system and data disks) on a server	One or more specific disks (system or data disks)	SFS Turbo file systems
When to use	You want to back up entire cloud servers.	You want to back up only data disks.	You want to back up entire SFS Turbo file systems.
Advantages	All disks on a server are backed up at a time.	Only data of specific disks is backed up, which costs less than backing up an entire server.	File system data and their backups are stored separately, and the backups can be used to create new file systems.

Backup Mechanism

CBR in-cloud backup offers block-level backup. The first backup is a full backup and backs up all used data blocks. For example, if a disk size is 100 GB and 40 GB has been used, only the 40 GB of data is backed up. An incremental backup backs up only the data changed since the last backup to save the storage space and backup time.

When a backup is deleted, data blocks will not be deleted if they are depended on by other backups, ensuring that other backups can still be used for restoration.

Both a full backup and an incremental backup can be used to restore data to a given backup point in time.

When creating a backup of a disk, CBR also creates a snapshot for it. CBR keeps only the latest snapshot. Every time it creates a new snapshot during backup, it deletes the old snapshot.

CBR stores backups in OBS to ensure data security.

Backup Options

CBR supports one-off backup and periodic backup. A one-off backup task is manually created and is executed only once. Periodic backup tasks are automatically executed based on a user-defined backup policy.


Table 1-2 compares the two backup options.

Table 1-2 One-off backup and periodic backup

Item	One-Off Backup	Periodic Backup
Backup policy	Not required	Required
Number of backup tasks	One manual backup task	Periodic tasks triggered by a preset backup policy
Backup name	User-defined backup name, which is manualbk_XXXX by default	System-assigned backup name, which is autobk_XXXX by default
Backup mode	The first backup is a full backup and the consecutive backups are incremental.	The first backup is a full backup and the consecutive backups are incremental.
Application scenario	Executed before patching or upgrading the OS or upgrading an application. A one-off backup can be used for restoration if the patching or upgrading fails.	Executed for routine maintenance. The latest backup can be used for restoration if an unexpected failure or data loss occurs.

You can also use the two backup options together if needed. For example, you can associate resources with a vault and apply a backup policy to the vault to execute periodic backup for all the resources in the vault. Additionally, you can perform a one-off backup for the most important resources to enhance data security. **Figure 1-2** shows the use of the two backup options.

Figure 1-2 Use of the two backup options



Access to CBR

You can access the CBR service through the console or by calling HTTPS-based APIs.

- Console
Use the console if you prefer a web-based UI. Log in to the console and choose **Cloud Backup and Recovery**.
- APIs
Use APIs if you need to integrate CBR into a third-party system for secondary development. For details, see the *Cloud Backup and Recovery API Reference*.

1.2 Advantages

Reliable

CBR offers crash-consistent backup for multiple disks on a server. The backups protect against human errors, virus attacks, and natural disasters, and ensure your data security and reliability.

Efficient

Incremental forever backups shorten the time required for backup by 95%. With Instant Restore, CBR offers an RPO of as low as 1 hour and an RTO of only several minutes.

 **NOTE**

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

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

Easy to Use

CBR is easier to use than conventional backup systems. You can complete backup in just three steps, and no professional backup skills are required.

Secure

If the disks are encrypted, their backups are also encrypted to ensure data security.

1.3 Application Scenarios

CBR is ideal for data backup and restoration. The backups can maximize your data security and consistency.

Data Backup and Restoration

You can use CBR to quickly restore data to the latest backup point if any of the following incidents occur:

- Hacker or virus attacks
- Accidental deletion
- Application update errors
- System breakdown

1.4 Functions

[Table 1-3](#) lists the functions of CBR.

Before using CBR functions, it is recommended that you learn about [basic CBR concepts](#).

Table 1-3 CBR functions

Category	Function	Description
Cloud disk backup	Manual disk backup	Cloud disk backup provides snapshot-based backup for EVS disks on servers. You can back up specific disks to protect data on them.

Category	Function	Description
	Policy-based backup	You can create, modify, or delete a backup policy. A backup policy defines the schedule and retention for automatic backups.
	Backup management	You can set search criteria to quickly find the backups you want to manage. Then you can view their details, share, restore, or delete them if needed.
	Disk restoration using backups	When a disk is faulty, or their data is lost, you can use a backup to quickly restore the data.
	Disk creation using backups	You can use a disk backup to create a disk that contains the same data as the backup.
	Sharing a Backup	You can share a disk backup with other accounts to allow them to use the backup to create disks.
Cloud server backup	Manual server backup	Cloud server backup uses the consistency snapshot technology to protect data for ECSs. You can use CBR to back up an entire server to protect their data.
	Policy-based backup	You can create, modify, or delete a backup policy. A backup policy defines the schedule and retention for automatic backups.
	Backup management	You can set search criteria to quickly find the backups you want to manage. Then you can view their details, share, restore, replicate, or delete them if needed.
	Server restoration using backups	When a server is faulty, or their data is lost, you can use a backup to quickly restore the data.
	Sharing a Backup	You can share a server backup with other accounts to allow them to use the backup to create servers.

Category	Function	Description
	Image creation using server backups	You can create images from ECS backups and then use the images to quickly provision ECSs to restore service.
SFS Turbo backup	Manual SFS Turbo backup	You can back up SFS Turbo file systems and use the backups to create new SFS Turbo file systems.
	Policy-based backup	You can create, modify, or delete a backup policy. A backup policy defines the schedule and retention for automatic backups.
	Backup management	You can set search criteria to quickly find the backups you want to manage. Then you can view their details, share, restore, replicate, or delete them if needed.
	File system creation using backups	You can use an SFS Turbo file system backup to create a file system that contains the same data as the backup.

1.5 User Permissions

If you need to assign different permissions to employees in your enterprise to access your CBR resources, IAM is a good choice for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you secure access to your cloud resources.

The system provides two types of user permissions by default: user management and resource management.

- User management refers to management of users, user groups, and user group permissions.
- Resource management refers to control of operations on cloud service resources.

For details, see [System Permissions](#).

1.6 Constraints

General

- A vault can be associated with only one backup policy.
- A vault can be associated with a maximum of 256 resources.

- A maximum of 32 backup policies can be created.
- Only backups in the **Available** or **Locked** vaults can be used to restore data.
- Backups in a **Deleting** vault cannot be deleted.
- Backups cannot be downloaded to a local PC or uploaded to OBS.
- A vault and its associated servers or disks must be in the same region.

Cloud Disk Backup

- Only disks in the **Available** or **In-use** state can be backed up.
- A new disk must be at least as large as the backup's source disk.

Cloud Server Backup

- A maximum of 10 shared disks can be backed up with a cloud server.
- Only backups in the **Available** or **Locked** vaults can be used to create images.
- Images cannot be created from backups if the amount of resources associated with a server backup vault exceeds the quota.
- You are advised not to back up a server whose disk size exceeds 4 TB.

SFS Turbo Backup

- Only file systems in the **Available** state can be backed up.
- An SFS Turbo file system backup cannot be used to restore data to the original file system.

1.7 CBR and Other Services

CBR-related Services

Table 1-4 CBR-related services

Function	Related Service	Reference
CBR backs up data of an ECS and uses the backup to restore data for the ECS. You can also create images from ECS backups and use the images to quickly provision ECSs to restore services.	ECS	Creating a Cloud Server Backup Creating a Cloud Disk Backup
CBR backs up data of a BMS and uses the backup to restore data for the BMS. The backup and management processes for BMSs and ECSs are the same.	BMS	What Is CBR? Creating a Cloud Server Backup
CBR backs up data of SFS Turbo file systems and uses the backup to create new file systems to restore lost or corrupted data.	SFS	Creating an SFS Turbo Backup

Function	Related Service	Reference
CBR stores backups securely in OBS.	OBS	What Is CBR?
CBR backs up data of EVS disks and uses the backup to create new disks.	EVS	

1.8 Basic Concepts

1.8.1 CBR Concepts

Vault

CBR stores backups of a variety of resources in vaults, which are classified into the following types:

- **Server backup vaults:** store only backups of non-database servers. You can associate servers with a server backup vault and apply a backup policy to schedule automatic backups.
- **Disk backup vaults:** store only disk backups. You can associate disks with a disk backup vault and apply a backup policy to schedule automatic backups.
- **SFS Turbo backup vaults:** store only backups of SFS Turbo file systems. You can associate file systems with an SFS Turbo backup vault and apply a backup policy to schedule automatic backups.

Backup

A backup is a copy of a particular chunk of data and is usually stored elsewhere so that it may be used to restore the original data in the event of data loss. It can be generated either manually by a one-off backup task or automatically by a periodic backup task.

A one-off backup task is manually created and is executed only once. Periodic backup tasks are automatically executed based on a user-defined backup policy.

- A one-off backup is named **manualbk_XXXX** and can be user- or system-defined.
- A periodic backup is named **autobk_XXXX** by CBR.

Backup Policy

A backup policy is a set of rules that define the schedule and retention of backups. After you apply a backup policy to a vault, CBR automatically backs up data and retains backups based on that backup policy.

Instant Restore

Instant Restore restores data and creates images from backups, much faster than a normal restore.

Instant Restore is an enhanced function of CBR and requires no additional configuration. After Instant Restore is provided, you take less time to restore server data or create images.

Enhanced Backup

Enhanced backups are backups generated after Instant Restore is provided. Enhanced backups make it faster to restore server data or create images.

Before providing Instant Restore, CBR generates common backups. After providing Instant Restore, CBR first performs a full backup for each associated resource and then generates enhanced backups. CBR only generates enhanced backups currently.

For the same resource, an enhanced backup and a common backup have the same backup content and size. They only differ in the restoration speed.

1.8.2 Project and Enterprise Project

Project

A project is used to group and isolate OpenStack resources, such as the compute, storage, and network resources. A project can be a department or a project team. Multiple projects can be created for one account.

Enterprise Project

An enterprise project manages multiple resource instances by category. Resources and projects in different cloud service regions can be classified into one enterprise project. An enterprise can classify resources based on department or project group and put relevant resources into one enterprise project for management. Resources can be migrated between enterprise projects.

1.8.3 Region and AZ


Concept

A region and availability zone (AZ) identify the location of a data center. You can create resources in a specific region and AZ.

- A region is a physical data center, which is completely isolated to improve fault tolerance and stability. The region that is selected during resource creation cannot be changed after the resource is created.
- An AZ is a physical location where resources use independent power supplies and networks. A region contains one or more AZs that are physically isolated but interconnected through internal networks. Because AZs are isolated from each other, any fault that occurs in one AZ will not affect others.

Figure 1-3 shows the relationship between regions and AZs.

Figure 1-3 Regions and AZs



Selecting a Region

Select a region closest to your target users for lower network latency and quick access.

Selecting an AZ

When deploying resources, consider your applications' requirements on disaster recovery (DR) and network latency.

- For high DR capability, deploy resources in different AZs within the same region.
- For lower network latency, deploy resources in the same AZ.

Regions and Endpoints


Before you use an API to call resources, specify its region and endpoint. For more details, see [Regions and Endpoints](#).

2 Getting Started

2.1 Overview

This section describes how to use CBR to back up cloud servers, cloud disks, and file systems. The following figure illustrates the process.

Figure 2-1 Backup process



1. Create a backup vault of the right type based on the resources you want to protect. See the following sections for more information:
 - [Creating a Server Backup Vault](#)
 - [Creating a Disk Backup Vault](#)
 - [Creating an SFS Turbo Backup Vault](#)
2. Associate resources with the vault if you have not done so during vault creation. For details, see [Step 2: Associate a Resource with the Vault](#).
3. Create backups for the associated resources. Backups are stored in vaults. See the following sections for more information:
 - [Creating a Cloud Server Backup](#)
 - [Creating a Cloud Disk Backup](#)
 - [Creating an SFS Turbo Backup](#)
4. Use backups to restore the resources from virus attacks or accidental deletion. See the following sections for more information:
 - [Restoring from a Cloud Server Backup](#)
 - [Restoring from a Cloud Disk Backup](#)



2.2 Step 1: Create a Vault

2.2.1 Creating a Server Backup Vault

This section describes how to create a server backup vault.

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 In the upper right corner of the page, click **Create Server Backup Vault**.

Step 3 Select a protection type.

- **Backup:** A server backup vault stores server backups.

Step 4 (Optional) In the server list, select the servers or disks you want to back up. After the servers or disks are selected, they are added to the list of selected servers. You can also select specific disks on a server and associate them with the vault.

NOTICE

To avoid data inconsistency after restoration, you are advised to back up the entire server.

If you want to back up only some of the disks to reduce costs, ensure that data on the backed up disks does not depend on the disks that are not backed up. Or, data inconsistency may occur.

For example, the data of an Oracle database is scattered across different disks. If only some of the disks are backed up, restoration restores only the data of the disks that have been backed up, with data on the rest of the disks unchanged. As a result, the data may be inconsistent and the Oracle database may fail to start.

 **NOTE**

- The selected servers must have not been associated with any vault and must be in the **Running** or **Stopped** state.
- You can also associate servers with the vault you are creating later if you skip this step.

Step 5 Specify a vault capacity ranging from 10 GB to 10,485,760 GB. Properly plan the vault capacity, which must be at least the same as the size of the servers you want to back up. Also, if a backup policy is applied to the vault, more capacity is required.

As the vault's used space grows, you can expand the vault capacity if it becomes insufficient.

Step 6 Configure auto backup.

- If you select **Configure**, you must then select an existing backup policy or create a new policy. After the vault is created, CBR will apply the policy to this vault, and all servers associated with this vault will be automatically backed up based on this policy.
- If you select **Skip**, servers associated with this vault will not be automatically backed up until you apply a backup policy to the vault.

Step 7 If you have subscribed to the Enterprise Project Management Service (EPS), add the vault to an existing enterprise project.

EPS provides a unified method to manage cloud resources by project, allowing you to manage resources, users, and user groups in your projects. The default enterprise project is **default**.

Step 8 (Optional) Add tags to the vault.

Tags are key-value pairs, which are used to identify, classify, and search for vaults. You can add a maximum of 10 tags for a vault, and vault tags are only used for vault search and management.

Table 2-1 describes the parameters of a tag.

Table 2-1 Tag parameter description

Parameter	Description	Example Value
Key	<p>Each tag has a unique key. You can customize a key or select the key of an existing tag created in TMS.</p> <p>A tag key:</p> <ul style="list-style-type: none"> • Can contain 1 to 36 Unicode characters. • Can contain only letters, digits, hyphens (-), and underscores (_). 	Key_0001
Value	<p>A tag value can be repetitive or left blank.</p> <p>A tag value:</p> <ul style="list-style-type: none"> • Can contain 0 to 43 Unicode characters. • Can contain only letters, digits, hyphens (-), and underscores (_). 	Value_0001

Step 9 Specify a name for the vault.

The name must contain 1 to 64 characters including digits, letters, underscores (_), or hyphens (-), for example, **vault-f61e**.

 **NOTE**

You can also use the default name **vault_XXXX**.

Step 10 Complete the creation as prompted.

Step 11 Go back to the **Cloud Server Backups** page. You can see the created vault in the vault list.

You can associate servers with the vault and perform backup for the servers. For details, see [Querying a Vault](#).



----End

2.2.2 Creating a Disk Backup Vault

This section describes how to create a disk backup vault.

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 In the upper right corner of the page, click **Create Disk Backup Vault**.

Step 3 (Optional) In the disk list, select the disks you want to back up. After disks are selected, they are added to the list of selected disks.

 **NOTE**

- The selected disks must have not been associated with any vault and must be in the **Available** or **In-use** state.
- You can also associate disks with the vault you are creating later if you skip this step.

Step 4 Specify a vault capacity ranging from 10 GB to 10,485,760 GB. Properly plan the vault capacity, which must be at least the same as the size of the disks you want to back up.

Step 5 Configure auto backup.

- If you select **Configure**, you must then select an existing backup policy or create a new policy. After the vault is created, CBR will apply the policy to this vault, and all disks associated with this vault will be automatically backed up based on this policy.
- If you select **Skip**, disks associated with this vault will not be automatically backed up until you apply a backup policy to the vault.

Step 6 If you have subscribed to the EPS service, add the vault to an existing enterprise project.

EPS provides a unified method to manage cloud resources by project, allowing you to manage resources, users, and user groups in your projects. The default enterprise project is **default**.

Step 7 (Optional) Add tags to the vault.

Tags are key-value pairs, which are used to identify, classify, and search for vaults. You can add a maximum of 10 tags for a vault, and vault tags are only used for vault search and management.

[Table 2-2](#) describes the parameters of a tag.

Table 2-2 Tag parameter description

Parameter	Description	Example Value
Key	Each tag has a unique key. You can customize a key or select the key of an existing tag created in TMS. A tag key: <ul style="list-style-type: none"> • Can contain 1 to 36 Unicode characters. • Can contain only letters, digits, hyphens (-), and underscores (_). 	Key_0001

Parameter	Description	Example Value
Value	<p>A tag value can be repetitive or left blank.</p> <p>A tag value:</p> <ul style="list-style-type: none"> • Can contain 0 to 43 Unicode characters. • Can contain only letters, digits, hyphens (-), and underscores (_). 	Value_0001

Step 8 Specify a name for the vault.

The name must contain 1 to 64 characters including digits, letters, underscores (_), or hyphens (-), for example, **vault-612c**.

 **NOTE**

You can also use the default name **vault_XXXX**.

Step 9 Complete the as prompted.

Step 10 Go back to the **Cloud Disk Backups** page. You can see the created vault in the vault list.

You can associate disks to the new vault or perform backup for the disks. For details, see [Vault Management](#).



----End

2.2.3 Creating an SFS Turbo Backup Vault

This section describes how to create an SFS Turbo backup vault.

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. Click  and choose **Storage > Cloud Backup and Recovery > SFS Turbo Backups**.

Step 2 In the upper right corner of the page, click **Create SFS Turbo Backup Vault**.

Step 3 Select a protection type.

- **Backup:** An SFS Turbo backup vault stores SFS Turbo backups.

Step 4 (Optional) In the file system list, select the file systems to be backed up. After file systems are selected, they are added to the list of selected file systems.

 **NOTE**

- The selected file systems must have not been associated with any vault and must be in the **Available** state.
- You can also associate file systems with the vault you are creating later if you skip this step.

Step 5 Specify a vault capacity ranging from 10 GB to 10,485,760 GB. Properly plan the vault capacity, which must be at least the same as the size of the file systems you want to back up.

Step 6 Configure auto backup.

- If you select **Configure**, you must then select an existing backup policy or create a new policy. After the vault is created, CBR will apply the policy to this vault, and all file systems associated with this vault will be automatically backed up based on this policy.
- If you select **Skip**, file systems associated with this vault will not be automatically backed up until you apply a backup policy to the vault.

Step 7 If you have subscribed to the EPS service, add the vault to an existing enterprise project.

EPS provides a unified method to manage cloud resources by project, allowing you to manage resources, users, and user groups in your projects. The default enterprise project is **default**.

Step 8 (Optional) Add tags to the vault.

Tags are key-value pairs, which are used to identify, classify, and search for vaults. You can add a maximum of 10 tags for a vault, and vault tags are only used for vault search and management.

Table 2-3 describes the parameters of a tag.

Table 2-3 Tag parameter description

Parameter	Description	Example Value
Key	Each tag has a unique key. You can customize a key or select the key of an existing tag created in TMS. A tag key: <ul style="list-style-type: none"> • Can contain 1 to 36 Unicode characters. • Can contain only letters, digits, hyphens (-), and underscores (_). 	Key_0001
Value	A tag value can be repetitive or left blank. A tag value: <ul style="list-style-type: none"> • Can contain 0 to 43 Unicode characters. • Can contain only letters, digits, hyphens (-), and underscores (_). 	Value_0001

Step 9 Specify a name for the vault.

The name must contain 1 to 64 characters including digits, letters, underscores (_), or hyphens (-), for example, **vault-612c**.

 **NOTE**

You can also use the default name **vault_XXXX**.

Step 10 Complete the creation as prompted.

Step 11 Go back to the **SFS Turbo Backups** page. You can see the created vault in the vault list.

You can associate file systems to the new vault or perform backup for the file systems. For details, see [Vault Management](#).

----End

2.3 Step 2: Associate a Resource with the Vault

If you have already associated servers, file systems, or disks when creating a vault, skip this step.

After a server backup vault, SFS Turbo backup vault, or disk backup vault is created, you can associate servers, file systems, or disks with the vault to back up these resources.

Prerequisites



- A vault can be associated with a maximum of 256 resources.
- The servers you plan to associate with a vault must have at least one disk attached.
- The vault and the resources you plan to associate with it must be in the same region.
- The total size of the resources to be associated cannot be greater than the vault capacity.
- Resources can be associated only when they are in the statuses in the table below.

Table 2-4 Resource statuses available for association

Resource Type	Status
Cloud server	Running or Stopped
Cloud disk	Available or In-use
SFS Turbo file system	Available

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 On a backup page, locate the target vault and click **Associate Server**, **Associate File System**, or **Associate Disk**.

Step 3 In the resource list, select the resources you want to associate with the vault. After resources are selected, they are added to the list of selected resources.

Step 4 Click **OK**. Then on the **Associated Servers** tab page, you can view the number of resources that have been associated.

 **NOTE**

If a new disk is attached to an associated server, CBR automatically identifies the new disk and includes the new disk in subsequent backup tasks.

----End

2.4 Step 3: Create a Backup

2.4.1 Creating a Cloud Server Backup

This section describes how to quickly create a cloud server backup.

The backup process for BMSs is the same as that for ECSs.

If you do not need an ECS for the moment, you can back up the ECS and then delete it. When you want an ECS later, you can create an image from the ECS backup and use the image to create ECSs.

Backing up a server does not impact the server performance.

Peak hours of the backup service are from 22:00 to 08:00, during which there may be delays. So you are advised to evaluate your service types and schedule backups in discrete time periods.



Prerequisites

- Only servers in the **Running** or **Stopped** state can be backed up.
- At least one server backup vault is available.

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 On the **Cloud Server Backups** page, click the **Vaults** tab and find the vault to which the server is associated.

Step 3 Perform backup in either of the following ways:

- Choose **More > Perform Backup** in the **Operation** column. In the server list, select the server you want to back up. After a server is selected, it is added to the list of selected servers.

- Click the vault name to go to the vault details page. On the **Associated Servers** tab page, locate the target server and click **Perform Backup** in the **Operation** column.

Step 4 Set **Name** and **Description** for the backup. [Table 2-5](#) describes the parameters.

Table 2-5 Parameter description

Parameter	Description	Remarks
Name	Name of the backup you are creating. A name must contain 1 to 64 characters including digits, letters, underscores (_), or hyphens (-). NOTE You can also use the default name manualbk_XXXX . If multiple servers are to be backed up, the system automatically adds suffixes to their backup names, for example, backup-0001 and backup-0002 .	manualbk_d819
Description	Description of the backup. It cannot exceed 255 characters.	--

Step 5 Choose whether to enable full backup. If full backup is enabled, CBR performs a full backup on every associated server, which requires a larger capacity compared to an incremental backup. See [Figure 2-2](#).

Figure 2-2 Full Backup

Full Backup  Enable

Step 6 Click **OK**. CBR automatically creates a backup for the server.

On the **Backups** tab page, if the status of the backup is **Available**, the backup task is successful.

 **NOTE**

A server can be restarted if the backup progress exceeds 10%. However, to ensure data integrity, restart it after the backup is complete.

After the backup is complete, you can use the backup to restore server data or create an image. For details, see [Restoring from a Cloud Server Backup](#) and [Creating an Image from a Cloud Server Backup](#).

----End

2.4.2 Creating a Cloud Disk Backup

This section describes how to quickly create a cloud disk backup.

Backing up a server does not impact the disk performance.



Peak hours of the backup service are from 22:00 to 08:00, during which there may be delays. So you are advised to evaluate your service types and schedule backups in discrete time periods.

Prerequisites

A disk can be backed up only when its status is **Available** or **In-use**. If you have performed operations such as expanding, attaching, detaching, or deleting a disk, refresh the page first to ensure that the operation is complete and then determine whether to back up the disk.

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 On the **Cloud Disk Backups** page, click the **Vaults** tab and find the vault to which the disk is associated.

Step 3 Perform backup in either of the following ways:

- Click **Perform Backup** in the **Operation** column. In the disk list, select the disk you want to back up. After a disk is selected, it is added to the list of selected disks.
- Click the vault name to go to the vault details page. On the **Associated Disks** tab page, locate the target disk and click **Perform Backup** in the **Operation** column.

 **NOTE**

CBR will identify whether the selected disk is encrypted. If it is encrypted, the backups will be automatically encrypted.

Step 4 Set **Name** and **Description** for the backup. [Table 2-6](#) describes the parameters.

Table 2-6 Parameter description

Parameter	Description	Remarks
Name	Name of the backup you are creating. A name must contain 1 to 64 characters including digits, letters, underscores (_), or hyphens (-). NOTE You can also use the default name manualbk_XXXX . If multiple disks are to be backed up, the system automatically adds suffixes to their backup names, for example, backup-0001 and backup-0002 .	manualbk_d819
Description	Description of the backup. It cannot exceed 255 characters.	--

Step 5 Choose whether to enable full backup. If full backup is enabled, CBR performs a full backup on every associated disk, which requires a larger capacity compared to an incremental backup. See [Figure 2-3](#).

Figure 2-3 Full Backup

Full Backup  Enable

Step 6 Click **OK**. CBR automatically creates a backup for the disk.

On the **Backups** tab page, if the status of the backup is **Available**, the backup task is successful.

 **NOTE**

If you delete data from the disk during the backup, the deleted data may fail to be backed up. Therefore, to ensure data integrity, delete the target data after the backup is complete, and then perform the backup.

After the backup is complete, you can use the backup to restore disk data. For details, see [Restoring from a Cloud Disk Backup](#).

----End

2.4.3 Creating an SFS Turbo Backup

This section describes how to quickly create an SFS Turbo file system backup.

To ensure data integrity, you are advised to back up the file system during off-peak hours when no data is written to the file system.



Peak hours of the backup service are from 22:00 to 08:00, during which there may be delays. So you are advised to evaluate your service types and schedule backups in discrete time periods.

Prerequisites

A file system can be backed up only when its status is **Available** or **In-use**. If you have performed operations such as expanding, mounting, unmounting, or deleting a file system, refresh the page first to ensure that the operation is complete and then determine whether to back up the file system.

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. Click  and choose **Storage > Cloud Backup and Recovery > SFS Turbo Backups**.

Step 2 On the **SFS Turbo Backups** page, click the **Vaults** tab and find the vault to which the file system is associated.

Step 3 Perform backup in either of the following ways:

- Choose **More > Perform Backup** in the **Operation** column. In the file system list, select the file system to be backed up. After a file system is selected, it is added to the list of selected file systems.
- Click the vault name to go to the vault details page. On the **Associated File Systems** tab page, locate the target file system and click **Perform Backup** in the **Operation** column.

Step 4 Set **Name** and **Description** for the backup. [Table 2-7](#) describes the parameters.

Table 2-7 Parameter description

Parameter	Description	Remarks
Name	<p>Name of the backup you are creating.</p> <p>A name must contain 1 to 64 characters including digits, letters, underscores (_), or hyphens (-).</p> <p>NOTE You can also use the default name manualbk_XXXX.</p> <p>If multiple file systems are to be backed up, the system automatically adds suffixes to their backup names, for example, backup-0001 and backup-0002.</p>	manualbk_d819
Description	<p>Description of the backup.</p> <p>It cannot exceed 255 characters.</p>	--

Step 5 Click **OK**. CBR automatically creates a backup for the file system.

On the **Backups** tab page, if the status of the backup is **Available**, the backup task is successful.

 **NOTE**

If you delete data from the file system during the backup, the deleted data may fail to be backed up. Therefore, to ensure data integrity, delete the target data after the backup is complete, and then perform the backup.

After the backup is complete, you can create a new SFS Turbo file system using the backup. For details, see [Using a Backup to Create a File System](#).

----End

3 Vault Management

3.1 Querying a Vault



You can set search criteria for querying desired vaults in the vault list.

Prerequisites

A vault has been created.

Viewing Vault Details

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 On the **Vaults** tab, view basic information about all vaults. Related parameters are described in the following table.

Table 3-1 Basic information parameters

Parameter	Description
Name/ID	Name and ID of the vault. Click the vault name to view details about the vault.
Type	Vault type
Status	Vault status. Table 3-2 describes the vault statuses.
Specifications	Vault specifications <ul style="list-style-type: none"> • A server backup vault stores backups of non-database servers.



Parameter	Description
Vault Capacity (GB)	Capacity used by the backups in the vault. It shows the space used by backups and the total vault capacity. For example: If 20/100 is displayed, 20 GB has been used out of the 100 GB vault capacity.
Associated Servers/ File Systems/Disks	Number of servers, file systems, and disks associated with the vault. You can click the number to view details of associated resources. The associated capacity shown on the details page is the total capacity of all the resources that have been associated with this vault.

Step 3 On the **Vaults** tab page, set filter criteria to view specific vaults.

- Select a value from the status drop-down list to query vaults by status. [Table 3-2](#) describes the vault statuses.

Table 3-2 Vault statuses

Status	Attribute	Description
All statuses	--	All vaults are displayed if this value is selected.
Available	A stable state	A stable state after a vault task is complete. This state allows most of the operations.
Locked	An intermediate state	An intermediate state displayed when a capacity expansion is in progress. If a vault is in this state, you can perform operations, such as applying a policy and associating servers, file systems, or disks. However, the following operations are not allowed on such a vault: expanding the vault capacity and changing the vault specifications. Once those operations are complete, the vault status will become Available .
Deleting	An intermediate state	An intermediate state displayed when a vault is being deleted. In this state, a progress bar is displayed indicating the deletion progress. If the progress bar remains unchanged for an extended time, an exception has occurred. Contact technical support.
Error	A stable state	A vault enters the Error state when an exception occurs during task execution. You can click Tasks in the navigation pane on the left to view the error cause. If the error persists, contact technical support.

- Search a vault by its name or ID.
- Click **Search by Tag** in the upper right corner to search for vaults by tag.
 - On the displayed **Search by Tag** page, enter an existing tag key and value and click . The added tag search criteria are displayed under the text boxes. Click **Search** in the lower right corner.
 - You can add a maximum of 10 tags by clicking . They will be applied together for a combination search.
 - You can click **Reset** in the lower right corner to reset the search criteria.

Step 4 Click the name of a specific vault to view vault details.

 **NOTE**

The values of used capacity and backup space are rounded off to integers. CBR will display 0 GB for any backup space less than 1 GB. For example, there may be 200 MB backup space used, but it will be displayed as 0 GB on the console.

----End

3.2 Deleting a Vault

You can delete unwanted vaults to reduce storage space usage and costs.



Once you delete a vault, all backups stored in the vault will be deleted.

Prerequisites

- There is at least one vault.
- The vault is in the **Available** or **Error** state.

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 Find the target vault and choose **More > Delete** in the **Operation** column. All backups stored in the vault will be deleted once you delete a vault.

Step 3 Click **Yes**.

----End

3.3 Dissociating a Resource



If you no longer need to back up an associated resource, dissociate it from your vault.

After a resource is dissociated, the vault's backup policy no longer applies to the resource. In addition, all manual and automatic backups of this resource will be deleted. Deleted backups cannot be used to restore data.

Dissociating a resource from a vault does not affect the performance of services on the resource.

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 Find the target vault and click its name.

Step 3 In this example, we will be using the **Cloud Server Backups** page to illustrate the process. Click the **Associated Servers** tab. Find the target server and click **Dissociate** in the **Operation** column.

Step 4 Confirm the information and click **Yes**.

----End

3.4 Migrating a Resource

Migrating a resource means that you dissociate a resource from a vault and then associate it to another vault. All backups of the resource will be migrated to the destination vault.



Constraints

- Resources can be migrated only when the source and destination vaults are in the **Available** or **Locked** state.
- The source and destination vaults for resource migration must be of the same specifications.
- The remaining capacity of the destination vault must be greater than the size of resource backups to be migrated.
- Cross-account resource migration is currently not supported.

- The source and destination vaults must be in the same region.

Procedure

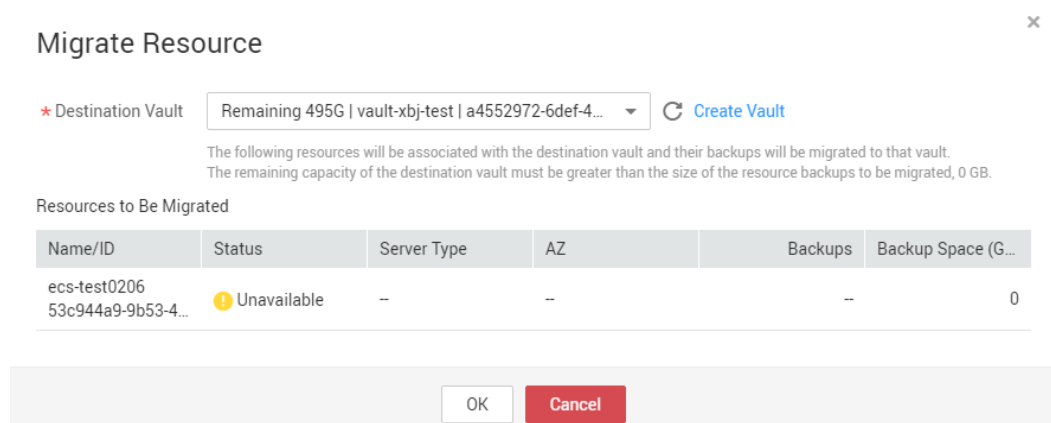
Step 1 Log in to the CBR console.

1. Log in to the management console.
2. Click  in the upper left corner and select a region.
3. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 Find the target vault and click its name. In this example, we will be using the **Cloud Server Backups** page to illustrate the process.

Step 3 Click the **Associated Servers** tab. Find the target server and click **Migrate** in the **Operation** column. See [Figure 3-1](#).

Figure 3-1 Migrating a resource



Step 4 Select the destination vault and click **OK**.

Step 5 View the migration progress on the **Tasks** page. If **Status** changes to **Successful**, the resource has been migrated.

Step 6 Go to the destination vault to confirm that the resource has been associated and all its backups have been migrated.



----End

3.5 Expanding Vault Capacity

You can expand the size of a vault if its total capacity is insufficient.

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 Find the target vault and choose **More > Expand Capacity** in the **Operation** column.

Step 3 Enter the capacity to be added. The minimum value is **1**.

Step 4 Click **Next**. Confirm the settings and click **Submit**.

Step 5 Return to the vault list and check that the capacity of the vault has been expanded.

----End

4 Backup Management

4.1 Viewing a Backup



In the backup list, you can set search criteria to filter backups and view their details. The results contain backup tasks that are running or have completed.

Prerequisites

At least one backup task has been created.

Viewing Backup Details

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 Click the **Backups** tab and set filter criteria to view the backups.

- You can search for backups by selecting a status from the **All statuses** drop-down list in the upper right corner of the backup list. [Table 4-1](#) describes the backup statuses.

Table 4-1 Backup statuses

Status	Status Attribute	Description
All statuses	--	All backups are displayed if this value is selected.

Status	Status Attribute	Description
Available	A stable state	A stable state of a backup after the backup is created, indicating that the backup is currently not being used. This state allows most of the operations.
Creating	An intermediate state	An intermediate state of a backup from the start of a backup job to the completion of this job. In the Tasks list, a progress bar is displayed for a backup task in this state. If the progress bar remains unchanged for an extended time, an exception has occurred. Contact technical support.
Restoring	An intermediate state	An intermediate state when using the backup to restore data. In the Tasks list, a progress bar is displayed for a restoration task in this state. If the progress bar remains unchanged for an extended time, an exception has occurred. Contact technical support.
Deleting	An intermediate state	An intermediate state from the start of deleting the backup to the completion of deleting the backup. In the Tasks list, a progress bar is displayed for a deletion task in this state. If the progress bar remains unchanged for an extended time, an exception has occurred. Contact technical support.
Error	A stable state	A backup enters the Error state when an exception occurs. A backup in this state cannot be used for restoration, and must be deleted manually. If manual deletion fails, contact technical support.

- You can search for backups by clicking **Advanced Search** in the upper right corner of the backup list.

You can search by specifying a backup status, backup name, backup ID, vault ID, server name, server ID, server type, or the creation date.

Step 3 Click the backup name to view details about the backup.

----End

4.2 Sharing a Backup

You can share a server or disk backup with other accounts. Shared backups can be used to create servers or disks.

Context

Sharer



- Backups can only be shared among accounts in the same region. They cannot be shared across regions.
- Encrypted backups cannot be shared.
- When a sharer deletes a shared backup, the backup will also be deleted from the recipient's account, but the disks or servers previously created using the backup will be retained.

Recipient

- A recipient must have at least one backup vault to store the accepted shared backup, and the vault's remaining space must be greater than the size of the backup to be accepted.
- A recipient can choose to accept or reject a backup. After accepting a backup, the recipient can use the backup to create new servers or disks.
- When a sharer deletes a shared backup, the backup will also be deleted from the recipient's account, but the disks or servers previously created using the backup will be retained.

Procedure for the Sharer

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 Click the **Backups** tab and set filter criteria to view the backups.

Step 3 Locate the target backup and choose **More > Share Backup** in the **Operation** column.



The backup name, server or disk name, backup ID, and backup type are displayed.

- Sharing a backup
 1. Click the **Share Backup** tab.
 2. Enter the account name of the recipient.
 3. Click **Add**. The account and project to be added will be displayed in the list. You can add multiple account names. A backup can be shared to a maximum of ten projects.
 4. Click **OK**.
- Canceling sharing
 1. Click the **Cancel Sharing** tab, select the projects you want to cancel sharing, and click **OK**.

----End

Procedure for the Recipient

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 Click the **Backups** tab and then click **Backups Shared with Me**.

Step 3 Ensure that the recipient has at least one backup vault before accepting the shared backup. For how to purchase a backup vault, see [Step 1: Create a Vault](#).

Step 4 Click **Accept**. On the displayed page, select the vault used to store the shared backup. Ensure that the vault's remaining capacity is greater than the backup size.

Automatic Association: Determine whether to enable automatic association for the vault. If you select **Configure**, the vault automatically scans and associates in the next backup period servers that have not been backed up and performs backup.

Step 5 View the shared backup you accepted in the backup list.

----End

4.3 Deleting a Backup

You can delete unwanted backups to reduce space usage and costs.



CBR supports manual deletion of backups and automatic deletion of expired backups. The latter is executed based on the backup retention rule in the backup policy. For details, see [Creating a Backup Policy](#).

Prerequisites

- There is at least one backup.
- The backup to be deleted is in the **Available** or **Error** state.

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

- Step 2** Click the **Backups** tab and locate the desired backup. For details, see [Viewing a Backup](#).
- Step 3** Choose **More > Delete** from the **Operation** column. Alternatively, select the backups you want to delete in a batch and click **Delete** in the upper left corner to delete them.
- Step 4** Click **Yes**.
- End

5 Policy Management

5.1 Viewing the Policy of a Vault

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 Find the target vault and click the vault name to view its details.

Step 3 View the policy applied to the vault.

----End

5.2 Creating a Backup Policy

A backup policy allows CBR to automatically back up vaults at specified times or intervals. Periodic backups can be used to restore data quickly against data corruption or loss.

To implement periodic backups, you need a backup policy first. You can use the default backup policy or create one as needed.



Constraints

- Backup policies can be applied to the following types of vaults: server backup vaults, disk backup vaults, SFS Turbo backup vaults
- A backup policy must be enabled before it can be used for periodic backups.
- A maximum of 32 backup policies can be created in each account.
- When expired backups are deleted, automatic backups will be deleted, but manual backups will not.

- Only servers in the **Running** or **Stopped** state and disks in the **Available** or **In-use** state can be backed up.
- CBR by default performs a full backup for a resource in the initial backup and incremental backups in all subsequent backups.
- The minimum interval between two full backups is 1 day.

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. Click  and choose **Storage > Cloud Backup and Recovery**.

Step 2 Choose **Policies** in the left navigation pane and click the **Backup Policies** tab. In the upper right corner, click **Create Policy**.

Step 3 Set the backup policy parameters. [Table 5-1](#) describes the parameters.

Table 5-1 Backup policy parameters

Parameter	Description	Example Value
Type	Select a policy type. In this section, we select the backup policy.	Backup policy
Name	Backup policy name A name must contain 1 to 64 characters including digits, letters, underscores (_), or hyphens (-).	backup_policy
Status	Whether to enable the backup policy.	Only after a backup policy is enabled and applied will CBR automatically backs up the vault resources and deletes expired backups.

Parameter	Description	Example Value
Execution Time	<p>Execution time</p> <p>Backups can be scheduled at the beginning of each hour, and you can select multiple hours.</p> <p>NOTICE</p> <ul style="list-style-type: none"> There may be a time difference between the scheduled backup time and the actual backup time. If a large amount of data needs to be backed up, you are advised to make backup less frequent to prevent the system from skipping any execution time. <p>For example, a disk is scheduled to be backed up at 00:00, 01:00, and 02:00. A backup task starts at 00:00. Because a large amount of incremental data needs to be backed up or a heap of backup tasks are executed at the same time, this backup task takes 90 minutes and completes at 01:30. CBR performs the next backup at 02:00. In this case, only two backups are generated in total, one at 00:00, and the other at 02:00.</p> <ul style="list-style-type: none"> The execution times refer to the local times of clients, not the time zone and times of the region. 	<p>00:00, 02:00</p> <ul style="list-style-type: none"> It is recommended that backups be performed during off-peak hours or when no services are running. Peak hours of the backup service are from 00:00 to 06:00, during which there may be scheduling delays. So you are advised to evaluate your service types and schedule backups outside of the backup peak hours.
Backup Cycle	<p>Select a backup cycle.</p> <ul style="list-style-type: none"> Week-based cycle Specifies on which days of each week the backup task will be executed. You can select multiple days. Custom cycle Specifies the interval (every 1 to 30 days) for executing the backup task. 	<p>Every day</p> <p>If you select day-based backup, the first backup is supposed to be executed on the day when the backup policy is created. If the execution time on the day you create the backup policy has passed, the first backup will be executed in the next backup cycle.</p> <p>It is recommended that backups be performed during off-peak hours or when no services are running.</p>

Parameter	Description	Example Value
Retention Rule	<p>Rule that specifies how backups will be retained</p> <ul style="list-style-type: none"> ● Time period You can choose to retain backups for one month, three months, six months, one year, or for any desired number (2 to 99999) of days. ● Backup quantity You can set the maximum number of backups to retain for each resource. The value ranges from 2 to 99999. ● Permanent <p>NOTE</p> <ul style="list-style-type: none"> - The system automatically deletes the earliest and expired backups every other day to avoid exceeding the maximum number of backups to retain or retaining any backup longer than the maximum retention period. - Expired backups are not deleted right after they are expired. They will be deleted from 12:00 to 00:00 in batches. - The retention rules apply only to auto-generated backups, but not manual backups. Manual backups need to be deleted manually. - If a backup is used to create an image, the backup will not be deleted by the retention rule. - A maximum of 10 backups are retained for failed periodic backup tasks. They are retained for one month and can be deleted manually. 	6 months

 **NOTE**

More frequent backups create more backups or retain backups for a longer time, protecting data to a greater extent but occupying more storage space. Set an appropriate backup frequency as needed.

Step 4 Click **OK**.

 **NOTE**

You can locate the desired vault and choose **More > Apply Backup Policy** to apply the policy to the vault. Then you can view the applied policy on the vault details page. After the policy is applied, data will be periodically backed up to the vault based on the policy.

----End

Example

At 10:00 a.m. on Monday, a user sets a backup policy for their vault to instruct CBR to execute a backup task at 02:00 a.m. every day and retain a maximum of three backups. As of 11:00 a.m. on Saturday, three backups will be retained, which are generated on Thursday, Friday, and Saturday. The backups generated at 02:00 a.m. on Tuesday and Wednesday have been automatically deleted.

5.3 Modifying a Policy



You can modify a policy to better suit your services.

Prerequisites

At least one policy has been created.

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 Find the target vault and click the vault name to view its details.

Step 3 In the **Policies** area, click **Edit** in the row of the policy to be edited.

Related parameters are described in [Table 5-1](#).

Step 4 Click **OK**.

If the retention rule is modified, the new rule does not necessarily apply to existing backups. For details, see [Why the New Retention Rule I Changed Is Not Applied?](#)

Step 5 Alternatively, select **Policies** from the navigation pane on the left and edit the desired policy.

----End

5.4 Deleting a Policy



You can delete policies if they are no longer needed.

Prerequisites

At least one policy has been created.

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. Click  and choose **Storage > Cloud Backup and Recovery**.

Step 2 Click the **Backup Policies** tab, locate the row that contains the policy you want to delete, and click **Delete**.

NOTE

Deleting a policy will not delete the backups generated based on the policy. You can manually delete unwanted backups.

Step 3 Confirm the information and click **Yes**.



----End

5.5 Applying a Policy to a Vault

You can apply a backup policy to a vault to execute backup tasks at specified times or intervals.

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 Find the target vault and choose **More > Apply Backup Policy**.

Step 3 Select an existing backup policy from the drop-down list or create a new one. For how to create a policy, see [Creating a Backup Policy](#).

Step 4 After the policy is successfully applied, view details in the **Policies** area of the vault details page.

----End

5.6 Removing a Policy from a Vault



If you no longer need automatic backup for a vault, remove the policy from the vault.

Prerequisites

A policy has been applied to the vault.

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 Find the target vault and click the vault name to view its details.

Step 3 In the **Policies** area, click **Remove Policy**.

NOTE

- If a policy is removed when a backup task is being executed for a resource in the vault, the backup task will continue and backups will be generated.
- After a policy is removed, backups retained by **Time period** will expire based on the retention rule, but backups retained by **Backup quantity** will not. You need manually delete unwanted backups.

Step 4 Click **Yes**.

Tasks will no longer be executed based on this policy for the vault.

----End

6 Restoring Data

6.1 Restoring from a Cloud Server Backup

When disks on a server are faulty or their data is lost, you can use a backup to restore the server to its state when the backup was created.

Constraints



- When restoring from a cloud server backup, backup of a data disk cannot be restored to the system disk.
- Data cannot be restored to servers in the **Faulty** state.

Prerequisites

- Disks are running properly on the server whose data needs to be restored.
- The server has at least one **Available** backup.

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 Click the **Backups** tab. Locate the desired backup. For details, see [Viewing a Backup](#).

Step 3 In the row of the backup, click **Restore Server**.

NOTICE

The current server data will be overwritten by the data captured at the time of backup. The restoration cannot be undone.

Step 4 (Optional) Deselect **Start the server immediately after restoration**.

If you do so, manually start the server after the restoration is complete.

NOTICE

Servers will be shut down during restoration, so you are advised to perform a restoration during off-peak hours.

Step 5 In the **Destination Disk** drop-down list, select the target disk to which the backup will be restored.

 **NOTE**

- If the server has only one disk, the backup is restored to that disk by default.
- If the server has multiple disks, the backup is restored to the original disks by default. You can also restore the backup to a different disk of at least the same size as the original disk.
- When restoring from a cloud server backup, backup of a data disk cannot be restored to the system disk.

NOTICE

If the number of disks to be restored is greater than the number of disks that were backed up, restoration may cause data inconsistency.

For example, if the Oracle data is scattered across multiple disks and only some of them are restored, data inconsistency may occur and the application may fail to start.

Step 6 Click **Yes** and confirm that the restoration is successful.

You can view the restoration status in the backup list. When the backup enters the **Available** state and no new restoration tasks failed, the restoration is successful. The resource is restored to the state when that backup was created.

For details about how to view failed restoration tasks, see [Managing Tasks](#).

NOTICE

If you use a cloud server backup to restore a logical volume group, you need to attach the logical volume group again.

Due to Window limitations, data disks may fail to be displayed after a Windows server is restored. If this happens, manually bring these data disks online. For details, see [Data Disks Are Not Displayed After a Windows Server Is Restored](#).

----End

6.2 Creating an Image from a Cloud Server Backup

CBR allows you to create images using ECS backups. You can use the images to provision ECSs to rapidly restore service running environments.

You can also use server backups to create images and then provision servers to restore data if your servers were accidentally deleted.

Prerequisites

- The ECS has been optimized before being backed up, and the Cloud-Init (for Linux) or Cloudbase-Init (for Windows) tool has been installed.
- The backup is in the **Available** state, or the backup is in the **Creating** state that is marked with "Image can be created."

NOTE

Once a backup creation starts, the backup enters the **Creating** state. After a period of time, a message stating "Image can be created" is displayed under **Creating**. In this case, the backup can be used for creating an image, even though it is still being created and cannot be used for restoration.



- The backup contains the system disk data.
- Only ECS backups can be used to create images.

Notes

- Images created using a backup are the same, so CBR allows you to use a backup to create only one full-ECS image that contains the whole data of the system disk and data disks of an ECS, in order to save the image quota. After an image is created, you can use the image to provision multiple ECSs in a batch.
- A backup with an image created cannot be deleted directly. To delete such a backup, delete its image first. If a backup is automatically generated based on a backup policy and the backup has been used to create an image, the backup will not be counted as a retained backup and will not be deleted automatically.
- 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.

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 Click the **Backups** tab. Locate the desired backup. For details, see [Viewing a Backup](#).

Step 3 In the row of the backup, choose **More > Create Image**.

Step 4 Create an image by referring to section "Creating a Full-ECS Image from a CBR Backup" in the *Image Management Service User Guide*.

Step 5 Use the image to provision ECSs when needed. For details, see section "Creating an ECS from an Image" in the *Image Management Service User Guide*.

----End

6.3 Restoring from a Cloud Disk Backup

You can use a disk backup to restore the disk to its state when the backup was created.

Prerequisites



- The disk to be restored is **Available**.
- Before restoring the disk data, 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.

Constraints

- Backups can only be restored to original disks. If you want to restore a backup to a different disk, use the backup to create a new disk.

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 Click the **Backups** tab. Locate the desired backup. For details, see [Viewing a Backup](#).

Step 3 In the row of the backup, click **Restore Disk**.

NOTICE

- The backup data will overwrite the current disk data, and the restoration cannot be undone.
 - If the restore button is grayed out, stop the server, detach the disk, and then try again. After the disk data is restored, attach the disk to the server and start the server.
-

Step 4 Click **Yes**. You can check whether data is successfully restored on the **Backups** tab page of **Cloud Disk Backups** or on the EVS console.

When the status of the backup changes to **Available**, the restoration is successful. The resource is restored to the state when that backup was created.

Step 5 After the restoration is complete, re-attach the disk to the server. For details, see section "Attaching an Existing Non-Shared Disk" in the *Elastic Volume Service User Guide*.

----End



6.4 Using a Backup to Create a Disk

You can create new disks from backups. Once created, the new disks will contain the backup data.

The new disks created using system disk backups can only be used as data disks on servers. They cannot be used as system disks.

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 Click the **Backups** tab. Locate the desired backup. For details, see [Viewing a Backup](#).

Step 3 Click **Create Disk** in the **Operation** column of the backup. The button is available only when the backup status is **Available**.

Step 4 Configure the disk parameters.

NOTE

See the parameter description table in section "Create an EVS Disk" of the *Elastic Volume Service User Guide* for more information.

Pay attention to the following:

- You can choose the AZ to which the backup source disk belongs, or a different AZ.
- The new disk must be at least as large as the backup's source disk.

If the capacity of the new disk is greater than that of the backup's source disk, format the additional space by following the steps provided in section "Extending Disk Partitions and File Systems" of the *Elastic Volume Service User Guide*.

- You can create a disk of any type regardless of the backup's source disk type.

Step 5 Click **Next**.

Step 6 Go back to the disk list. Check whether the disk is successfully created.

You will see the disk status change as follows: **Creating, Available, Restoring, Available**. You may not notice the **Restoring** status because Instant Restore is supported and the restoration speed is very fast. After the disk status has changed from **Creating** to **Available**, the disk is successfully created. After the status has changed from **Restoring** to **Available**, backup data has been successfully restored to the created disk.



----End

6.5 Using a Backup to Create a File System

In case of a virus attack, accidental deletion, or software or hardware fault, you can use an SFS Turbo file system backup to create a new file system. Once created, data on the new file system is the same as that in the backup.

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. Click  and choose **Storage > Cloud Backup and Recovery**. Select a backup type from the left navigation pane.

Step 2 Click the **Backups** tab and locate the desired backup. For details, see [Viewing a Backup](#).

Step 3 Click **Create File System** in the **Operation** column of the backup. The button is available only when the backup status is **Available**.

Step 4 Configure the file system parameters.

NOTE

- You can learn about the parameter descriptions in table "Parameter description" under "Creating an SFS Turbo File System" in "Create a File System" of the *Scalable File Service User Guide*.

Step 5 Click **Create Now**.

Step 6 Go back to the file system list and check whether the file system is successfully created.

You will see the file system status change as follows: **Creating, Available, Restoring, Available**. You may not notice the **Restoring** status because Instant Restore is supported and the restoration speed is very fast. After the file system status has changed from **Creating** to **Available**, the file system is successfully created. After the status has changed from **Restoring** to **Available**, backup data has been successfully restored to the created file system.

----End

7 Managing Tasks


You can view tasks in the task list, which shows policy-driven tasks that have been executed over the past 30 days.

Prerequisites


At least one task exists.

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 > Tasks**.

Step 2 Filter tasks by task type, task status, task ID, resource ID, resource name, vault ID, vault name, and time.

Step 3 Click  in front of the task to view the task details.

If a task fails, you can view the failure cause in the task details.

----End

8 Monitoring

8.1 CBR Metrics

Scenarios

This section describes metrics reported by CBR as well as their namespaces and dimensions. You can use the console or APIs provided by Cloud Eye to query the metrics generated for CBR.

Namespace

SYS.CBR

Metrics

Table 8-1 CBR metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
used_vault_size	Used Vault Size	Used capacity of the vault Unit: GB	≥ 0	Vault	15 min
vault_util	Vault Usage	Capacity usage of the vault	0~100 %	Vault	15 min

Dimensions

Key	Value
instance_id	Vault name/ID

Viewing Monitoring Statistics

Step 1 Log in to the management console.

Step 2 View the monitoring graphs using either of the following methods.

- Method 1: Choose **Storage > Cloud Backup and Recovery**. In the vault list, locate the vault whose monitoring data you want to view and choose **More > View Monitoring Data** in the **Operation** column.
- Method 2: Choose **Management & Deployment > Cloud Eye > Cloud Service Monitoring > Cloud Backup and Recovery**. In the vault list, click **View Metric** in the **Operation** column of the vault whose monitoring data you want to view.

Step 3 View the vault monitoring data by metric or monitored duration.

For more information, see the *Cloud Eye User Guide*.

----End



9 Quotas

What Is Quota?

Quotas can limit the number or amount of resources available to users, such as the maximum number of ECSs or EVS disks that can be created.

If the existing resource quota cannot meet your service requirements, you can apply for a higher quota.

How Do I View My Quotas?

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. In the upper right corner of the page, click  .
The **Service Quota** page is displayed.
4. View the used and total quota of each type of resources on the displayed page.
If a quota cannot meet service requirements, apply for a higher quota.

How Do I Apply for a Higher Quota?

1. Log in to the management console.
2. In the upper right corner of the page, choose **Resources > My Quotas**.
The **Service Quota** page is displayed.
3. Click **Increase Quota** in the upper right corner of the page.
4. On the **Create Service Ticket** page, configure parameters as required.
In the **Problem Description** area, fill in the content and reason for adjustment.
5. After all necessary parameters are configured, select **I have read and agree to the Ticket Service Protocol and Privacy Statement** and click **Submit**.

10 FAQs

10.1 Concepts

10.1.1 What Are Full Backup and Incremental Backup?

CBR by default performs a full backup for a resource in the initial backup and incremental backups in all subsequent backups. If a resource has been backed up for many times, and then all of its generated backups are deleted, and the resource is backed up again, the system will also perform a full backup for the resource.

- 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 use backup **A** or **C** to restore data. If backup **A** and backup **B** are both deleted, you can still use backup **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.

10.1.2 What Are the Differences Between Backup and Disaster Recovery?

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

Table 10-1 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.

10.1.3 What Are the Differences Between Backups and Snapshots?

Both backups and snapshots provide data redundancy for disks to improve data reliability. [Table 10-2](#) lists the differences between them.

Table 10-2 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.

10.1.4 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 10-3](#) lists the differences between them.

Table 10-3 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	<ul style="list-style-type: none"> ● 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 	<ul style="list-style-type: none"> ● 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	<ul style="list-style-type: none"> ● Data backup and restoration ● Rapid service deployment and migration 	<ul style="list-style-type: none"> ● 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.

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.

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

Table 10-4 describes the differences between cloud server backup and cloud disk backup.

Table 10-4 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.

10.2 Billing

10.2.1 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).

10.2.2 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.

10.3 Backup

10.3.1 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.

10.3.2 Can I Back Up a Server Deployed with Databases?

Yes. To back up applications requiring strict consistency, such as databases and email systems, you are advised to suspend all write operations and then perform backup. If write operations cannot be suspended, you can stop the application systems or the server for offline backup. Without doing these, status of the server after restoration is similar to restart upon an unexpected power failure and log rollback will be performed on databases to keep data consistent.

10.3.3 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

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

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

10.3.5 Does CBR Support Cross-Region Backup?

No. CBR supports only backup and restoration within a region but not across regions.

10.3.6 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.

10.3.7 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.

10.3.8 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.

10.3.9 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.

10.3.10 How Many Backups Can I Create for a Resource?

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

10.3.11 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.

10.3.12 Can I Stop an Ongoing Backup Task?

No. An ongoing backup task cannot be stopped.

10.3.13 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.

10.3.14 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.

10.3.15 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 [Creating an Image from a Cloud Server Backup](#).
2. Use the image to create a server. For details, see section "Creating an ECS from an Image" in the *Image Management Service User Guide*.
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 section "Attaching a Non-Shared Disk" or "Attaching a Shared Disk" in the *Elastic Volume Service Getting Started*.

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 section "Mounting an NFS File System to ECSs (Linux)" in the *Scalable File Service Getting Started*.
 - To mount the file system to a Windows server, see section "Mounting an NFS File System to ECSs (Windows)" in the *Scalable File Service Getting Started*.
3. Log in to the server to view the data.

10.3.16 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, manual backups will always be kept.

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?](#)

10.4 Capacity

10.4.1 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 vault capacity:

1. **Prepare the following data:**

Disk capacity (GB): a

Backup retention period (days): b

Backup frequency (days): c

Daily changed data volume (GB): d

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

Vault capacity: (GB) = $[a + (b/c) \times d] \times 120\%$

 **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.

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:

$$\text{Vault capacity} = [800 + (30/1) \times 10] \times 120\% = 1,320 \text{ GB}$$

10.4.2 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.

10.4.3 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..
- Reducing the number of retained backups
 - a. Locate the target vault and delete unwanted backups by referring to .
 - 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 remove the policy, see . Then, automatic backup will stop, and the vault used space will not change.

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

Symptoms

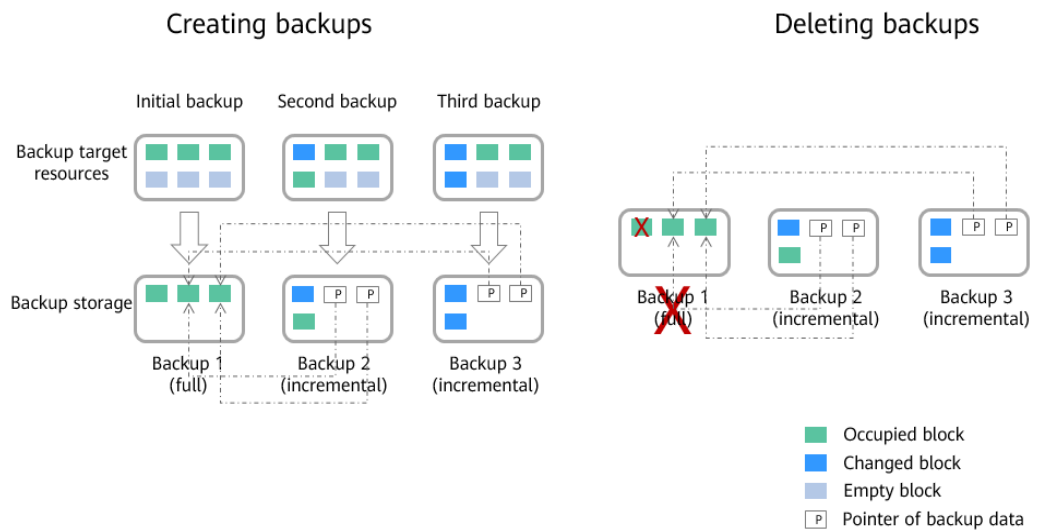
After unwanted backups are deleted from the vault, the used capacity of the vault decreases by only 1 GB to 2 GB.

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 10-1 Backup mechanism



10.4.5 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.

10.5 Restoration

10.5.1 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.

10.5.2 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.

10.5.3 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.

10.5.4 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.

10.5.5 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*.

10.5.6 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 section "Passwords" in the *Elastic Cloud Server User Guide*.

10.5.7 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.

10.5.8 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 [Creating an Image from a Cloud Server Backup](#).

- 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 section "Attaching a Non-Shared Disk" or "Attaching a Shared Disk" in the *Elastic Volume Service Getting Started*.

10.5.9 Can I Stop an Ongoing Restoration Task?

No. An ongoing restoration task cannot be stopped.

10.6 Policies

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

1. Go to the Cloud Backup and Recovery console and create 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 creation.
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.

10.6.2 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.

10.6.3 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, create 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 creation.
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.

10.6.4 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.

10.6.5 How Can I Cancel Auto Backup?

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

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

1. Log in to the CBR console and create a backup vault.
2. Associate resources with the vault during or after the creation.
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.

10.6.7 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.

10.7 Optimization

10.7.1 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 section "Installing Cloud-Init" in the *Image Management Service User Guide*.

To configure Cloud-Init, see section "Configuring Cloud-Init" in the *Image Management Service User Guide*.

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 10-2 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 10-3 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 10-4 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 10-5 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
 preset: enabled)
 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=
0/SUCCESS)
 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 10-6 Information returned when Cloud-Init automatic start is successfully set

```
root@ecs-debian-9:/tmp/CLDUP-INIT/haueicloud-c-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/CLDUP-INIT/haueicloud-c-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 10-7 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 10-8 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', '-odev
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 'depts': {'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; 2s 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-instanc
Aug 20 02:48:40 ecs-debian-9 cloud-init[1096]: [CLOUDINIT] stages.py[DEBUG]: Running module users-groups (module 'cloudinit.com
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 10-9 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-ethX`, and comment out `HWADDR`.

10.7.2 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.

10.7.3 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 Linux OSs, download and install Cloud-Init.

10.8 Others

10.8.1 Is There a Quota for CBR Vaults?

There are no quotas on CBR vaults. You can create as many vaults as needed.

10.8.2 Can I Merge My Vaults?

No. Vaults cannot be merged.

10.8.3 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.

10.8.4 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.

10.8.5 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.

10.8.6 Can I Download Backup Data to a Local PC?

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

10.8.7 How Do I Copy Disk Data to Another Account?

If two accounts are in the same region, you can use CBR backup sharing to copy disk data to another account. Cross-region backup sharing is currently not supported.

11 Troubleshooting Cases

11.1 Failed to Execute a Backup Task

Symptom

A manual or scheduled backup task fails.

Troubleshooting

Possible causes are listed here in order of their probability.

If the fault persists after you have ruled out one cause, move on to the next one.

Figure 11-1 Troubleshooting

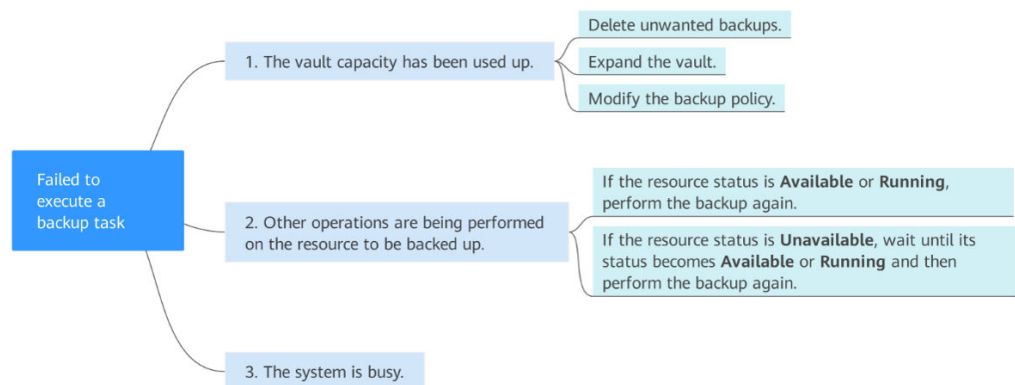


Table 11-1 Troubleshooting

Possible Cause	Solution
The vault capacity has been used up.	For detailed handling measures, see The Vault Capacity Has Been Used Up.

Possible Cause	Solution
Other operations are being performed on the resource to be backed up.	For detailed handling measures, see Other Operations Are Being Performed on the Resource to Be Backed Up .
The system is busy.	Back up the data in off-peak hours or contact technical support.

The Vault Capacity Has Been Used Up

Backup stops once the usage of the vault exceeds the upper limit. Take the following measures when the storage capacity of your vault is not enough:

- Log in to the CBR console, locate the target vault, and delete unwanted backups by following instructions in "Deleting a Backup".
- If you want to retain the generated backups, expand the vault capacity..
- If a backup policy has been applied to the vault, disable the backup policy or remove the policy from the vault.. Then, automatic backup is disabled, and the storage capacity of the vault will not change. You can also modify the backup policy or dissociate resources from the vault.

Other Operations Are Being Performed on the Resource to Be Backed Up

Automatic backup may fail when other operations, such as backup and restoration, are being performed on the target resource.

1. Log in to the CBR console and check whether the status of the target resource is **Available** or **Running**.
2. If yes, perform the backup again. If no, wait until the status of the target resource becomes **Available** or **Running**.

11.2 Failed to Delete a Backup

Symptom

The **Delete** button is unavailable, or a backup fails to be deleted.

Troubleshooting

Possible causes are listed here in order of their probability.

If the fault persists after you have ruled out one cause, move on to the next one.

Table 11-2 Troubleshooting

Possible Cause	Solution
The backup is being created or is being used for restoration.	Wait until the backup is created or the restoration completes and then delete it.
The backup is a cloud server backup and has been used to create an image, so the Delete button for the backup is grayed out.	Delete the image before deleting the backup.
A system exception occurs.	Delete the backup again. If the problem persists, contact technical support.
Insufficient permissions.	Check whether the account has the deletion permission on the IAM console.

11.3 Failed to Attach Disks

Symptom

Failed to attach disks despite following the procedure: Create EVS disks using the same disk backup (XFS file system backup) and attach them to the same server (to which multiple EVS disks with XFS file system backup have been attached). Running the **mount** command to attach disks fails.

Possible Cause

The superblock of an EVS disk (with XFS file systems) stores a universally unique identifier (UUID) about the file system. If a server has multiple disks (with XFS file systems), multiple UUIDs will exist on the server. Multiple disks may have the same UUID, which can cause the file system mounting to fail.

Troubleshooting Methods

When attaching an EVS disk, use parameters without UUID control or reallocate a new UUID to ensure that each UUID is unique.

Solution

Step 1 Log in to the server to which EVS disks failed to be attached.

Step 2 Resolve the problem in either of the following ways:

- Use a parameter without UUID when attaching an EVS disk: Run **mount -o nouuid /dev/Device name /Mount path**, for example:
mount -o nouuid /dev/sda6 /mnt/aa

- Reallocate a new UUID: Run `xfstool -U generate /dev/Device name`.

 NOTE

Because setting a parameter without UUID requires you to execute the command every time, you are advised to reallocate a new UUID.

----End

11.4 Data Disks Are Not Displayed After a Windows Server Is Restored

Symptom

When a Windows server is restored, the data disks are not displayed.


Possible Cause

Due to the limitations of Windows operating systems, data disks are in offline mode after a server is restored.

Solution


- Step 1** On the Windows desktop, right-click the **My Computer** icon.
- Step 2** Choose **Manage** from the shortcut menu. The **Computer Management** page is displayed.
- Step 3** In the navigation tree, choose **Storage > Disk Management**.
Data disks are in the offline state, as shown in [Figure 11-2](#).

Figure 11-2 Data disks in the offline state



Step 4 Right-click a data disk in the offline state and choose **Online**, as shown in Figure 11-3.


Figure 11-3 Setting a data disk to be online



After the data disk status changes to **Online**, the data disk will be displayed in the disk list, as shown in Figure 11-4.

In addition, the data disk will be properly displayed on the server.

Figure 11-4 Viewing online data disks



----End

A Appendix

A.1 Change History

Released On	Description
2023-02-28	This issue is the second official release, which incorporates the following change: Added support for resource migration.
2020-11-06	This issue is the first official release.