

Cloud Backup and Recovery

Service Overview

Issue 05
Date 2026-03-26



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2026. 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.

Huawei Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road
Qianzhong Avenue
Gui'an New District
Gui Zhou 550029
People's Republic of China

Website: <https://www.huaweicloud.com/intl/en-us/>

Contents

1 CBR Infographic	1
2 What Is CBR?	3
3 Advantages	8
4 Application Scenarios	9
5 Functions	11
6 Security	14
6.1 Identity Authentication and Access Control	14
6.2 Data Protection	15
6.3 Auditing and Logging	16
6.4 Resilience	16
6.5 Risk Monitoring	16
6.6 Fault Recovery	16
6.7 Certificates	17
6.8 Trusted Services	18
7 Permissions Management	19
8 Notes and Constraints	27
9 CBR and Other Services	34
10 Basic Concepts	36
10.1 CBR Concepts	36
10.2 Project and Enterprise Project	38
10.3 Region and AZ	39

1 CBR Infographic



Next-Gen HUAWEI CLOUD CBR

Multi-level protection for your data



Sophie, good news! We have migrated our services to the cloud, and the efficiency is great, but what about data loss. Any ideas?

Well, you need backups. Security first, always! Use HUAWEI CLOUD Cloud



your data.

2 What Is CBR?

Overview

Cloud Backup and Recovery (CBR) enables you to easily back up Elastic Cloud Servers (ECSs), Bare Metal Servers (BMSs), Elastic Volume Service (EVS) disks, SFS Turbo file systems, local files and directories, and on-premises VMware virtual environments. In case of a virus attack, accidental deletion, or software or hardware fault, you can use CBR to restore data to any point in time 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 can be used to restore the original data in the event of data loss.

The following types of backups are available:

- Cloud server backup: uses the consistency snapshot technology to protect data for ECSs and BMSs. Backups of non-database servers are common server backups, and those of database servers are database server backups.
- Cloud disk backup: provides snapshot-based data protection for EVS disks.
- SFS Turbo backup: protects data for SFS Turbo file systems.
- Desktop backup: protects data for Workspace desktops.

Vault

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

There are backup policies and replication policies.

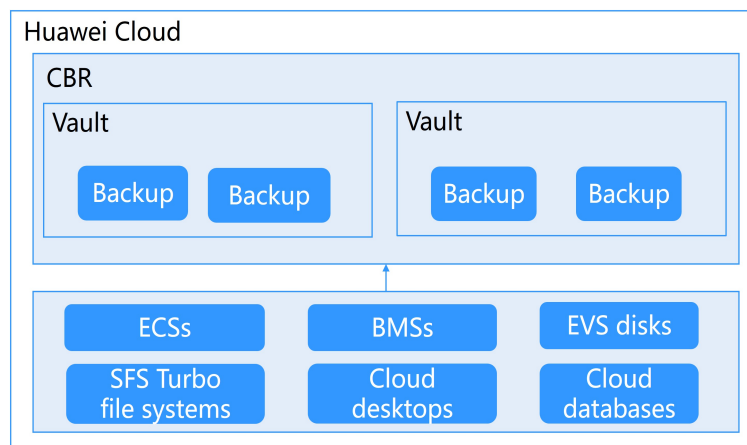
- A backup policy defines the timing, frequency, and retention of backups. Once applied to a vault, CBR will automatically back up data as specified.
- A replication policy determines the schedule and frequency for replicating data from one vault to another, as well as the retention period for each replica. Once applied, CBR automatically performs replication as specified. Backup replicas are stored in replication vaults.

Organizational policies

The organization administrator or an administrator delegated by CBR can centrally create and configure organizational backup policies and replication policies for member accounts in the organization.

- Organizational backup policies: An enterprise can use the organization's management account to create and configure organizational backup policies for member accounts.
- Organizational replication policies: An enterprise can use the organization's management account to create and configure organizational replication policies for member accounts.

Figure 2-1 CBR architecture



Differences Among the Backup Types

Table 2-1 Differences among the backup types

Item	Cloud Server Backup	Cloud Disk Backup	SFS Turbo Backup	Cloud Desktop Backup
What to back up	All disks (the system disk and data disks) on a server or certain disks and cloud servers running applications such as databases	One or more specific disks (the system disk or data disks)	SFS Turbo file systems	Entire Workspace desktop systems, including all disks

Item	Cloud Server Backup	Cloud Disk Backup	SFS Turbo Backup	Cloud Desktop Backup
When to use	You want to back up entire cloud servers.	You want to back up only data disks, as the system disk contains no user data.	You want to back up only SFS Turbo file systems.	You want to back up only Workspace desktops.
Advantages	All disks on a server are backed up at the same time to ensure data consistency.	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 restore file systems.	Desktop data and their backups are stored separately, and the backups can be used to restore desktops.

Backup Mechanism

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

When a backup is deleted, data blocks that are referenced by other backups will not be deleted, ensuring that these 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 for a disk, CBR also creates a snapshot for it. CBR keeps only the latest snapshot. Every time it creates a new snapshot, 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 2-2](#) compares the two backup options.


Table 2-2 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 subsequent backups are incremental.	The first backup is a full backup and subsequent backups are incremental.
Application scenario	A one-off backup is usually performed before an OS or application is patched or upgraded. The backup can be used for restoration if the patching or upgrade fails.	Periodic backups are performed as part of routine maintenance. The latest backup can be used to restore data in the event of an unexpected failure or data loss.

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 irregularly perform a one-off backup for the most important resources. CBR can store backups in OBS to ensure backup data security. [Figure 2-2](#) shows the use of the two backup options.

Theoretically, you can create as many backups for a resource as needed. There is no limit to the number of backups you can create for a resource.

Figure 2-2 Use of the two backup options



Access to CBR

You can access the CBR service through the console or by calling HTTPS 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 [Cloud Backup and Recovery API Reference](#).

3 Advantages

Reliable

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

Efficient

Incremental 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) defines how much data your business can afford to lose in the event of a disruption.

Recovery Time Objective (RTO) defines how quickly you need to restore systems and resume operations after a disruption.

Easy to Use

CBR is easier to use than conventional backup systems. You can complete a 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.

You can replicate backups across regions and restore them in remote regions for remote backup and disaster recovery.

4 Application Scenarios


CBR is ideal for data backup and restoration. It 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

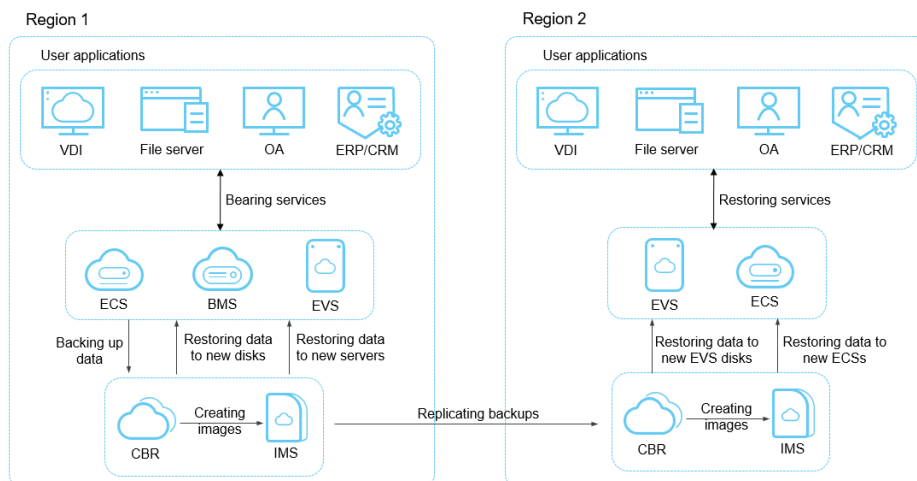
Figure 4-1 Data backup and restoration



Rapid Migration & Deployment

You can use cloud server backups to create images and then use such images to quickly provision new cloud servers with the same configuration as existing ones.

Figure 4-2 Rapid migration and deployment



5 Functions

This section describes main functions of CBR. You can check if a certain function is available in a region on the console.

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

Cloud Disk Backup

Manual disk backup

A cloud disk backup is a snapshot-based backup of EVS disks. You can back up a single disk or all disks to protect data on them.

Policy-based backup

With a backup policy, you can schedule regular backups of all disks to a vault, enabling fast restoration in case of data loss or corruption.

Backup management

You can set search criteria to quickly find the backup tasks you want to manage. Then you can view their details, share, restore, or delete them if needed.

Disk restoration from backups

When a disk is faulty, or its data is lost, you can use a backup to quickly restore the data.

Disk creation from backups

You can use a disk backup to create a disk that contains the same data as the backup.

Backup sharing

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

Cloud Server Backup

Manual server backup

Cloud server backup uses the consistency snapshot technology to protect data for ECSs and BMSs without the need to install the Agent on servers. It allows you to back up the entire servers. You are advised to use cloud server backup in scenarios that require high data consistency, such as RAID clusters.

Backup of specific disks on a server

You can create a single backup for multiple disks on a server to save the vault space.

Policy-based backup

With a backup policy, you can schedule regular backups of servers, enabling fast restoration in case of data loss or corruption.

Backup management

You can set search criteria to quickly find the backup tasks you want to manage. Then you can view their details, share, restore, or delete them if needed.

Server restoration from backups

When a server is faulty, or its data is lost, you can use a backup to quickly restore the data.

Backup sharing

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

Image creation from server backups

You can create images from ECS backups and then use the images to quickly provision ECSs to restore services.

With cross-region replication, you can replicate backups to destination regions and then create images and use the images to provision ECSs there.

Database server backup

Cloud server backup supports both crash-consistent and application-consistent backups. You can use it to back up ECSs running MySQL or SAP HANA databases, as application-consistent backups ensure transactional consistency by capturing in-memory data and pending I/O operations.

Cross-region replication

Cloud server backup enables you to replicate generated backups from one region to another. You can use the replicas in the destination region to create images and provision servers.

SFS Turbo Backup

Manual SFS Turbo backup

SFS Turbo backup allows you to back up SFS Turbo file systems. An SFS Turbo file system backup can be used to create a new SFS Turbo file system, preventing the loss of important data.

Policy-based backup

With a backup policy, you can schedule regular backups of SFS Turbo file systems, enabling fast restoration in case of data loss or corruption.

Backup management

You can set search criteria to quickly find the backup tasks you want to manage. Then you can view their details, share, restore, or delete them if needed.

File system creation from backups

You can use an SFS Turbo file system backup to create a file system that contains the same data as the backup.

Cross-region replication

SFS Turbo backup enables you to replicate SFS Turbo file system backups from one region to another. You can then use the replicated backup to create a file system in the destination region.

6 Security

- [6.1 Identity Authentication and Access Control](#)
- [6.2 Data Protection](#)
- [6.3 Auditing and Logging](#)
- [6.4 Resilience](#)
- [6.5 Risk Monitoring](#)
- [6.6 Fault Recovery](#)
- [6.7 Certificates](#)
- [6.8 Trusted Services](#)

6.1 Identity Authentication and Access Control

You can access CBR through the CBR console, APIs, or SDKs. No matter which method you choose, you actually use REST APIs.

These APIs require authentication, so you must first obtain the necessary credentials from Huawei Cloud Identity and Access Management (IAM) before using CBR. For details about IAM authentication, see [Authentication](#).

Access Control

You can use IAM to control access to your CBR resources.

Table 6-1 CBR access control

Method		Description	Reference
Permissions management	IAM permissions	IAM permissions define which actions are allowed or denied on your cloud resources. After creating an IAM user, the administrator needs to add it to a user group and grant the permissions required by CBR to the user group. Then, all users in this group automatically inherit those permissions.	Permissions management

6.2 Data Protection

CBR takes many measures to keep data secure and reliable.

Table 6-2 CBR data protection

Measure	Description
Transmission encryption (HTTPS)	To ensure transmission security, backups are stored in OBS buckets via HTTPS.
Storage data redundancy	<p>CBR allows you to create multi-AZ backup vaults so that your backups can be stored in multiple AZs of a region. If one AZ becomes unavailable, backups remain accessible from other AZs. This multi-AZ redundancy ensures high data reliability and availability.</p> <p>NOTE CBR relies on OBS's redundancy features to ensure backup redundancy. For details, see What Storage Redundancy Techniques Does OBS Use?</p>
Backup data encryption	If a disk you want to back up is encrypted, the backups generated for this disk will also be encrypted. When such a backup is used to restore data, the encrypted backup will first be decrypted and then restored to the target disk.
Cross-region replication	Cross-region replication allows you to automatically and asynchronously replicate backups from one region to a replication vault in a different region based on a replication policy. The cross-region disaster recovery capabilities it offers can meet your needs for remote backup.
Backup locking	To prevent backups from being deleted by mistake or maliciously, you can enable backup locking for vaults to improve data security. Once enabled, all backups in the vault enter the write once, read many (WORM) state. No one can delete the backups in their retention periods.

6.3 Auditing and Logging

Auditing

Cloud Trace Service (CTS) records operations on the Huawei Cloud resources in your account. You can use the logs generated by CTS to perform security analysis, track resource changes, audit compliance, and locate faults.

After you enable CTS and configure a tracker, CTS can record management and data traces of CBR for auditing.

For details about how to enable and configure CTS, see [CTS Getting Started](#).

For the CBR management and data traces supported by CTS, see [Auditing](#).

Logging

CBR shows critical asynchronous tasks on the web page. You can log in the CBR console, choose **Tasks** from the navigation page, and view the task list in the right pane. Alternatively, you can [query the task list](#) via the API.

6.4 Resilience

CBR uses a multi-level reliability architecture and technical measures such as cross-region replication, intra-region cross-AZ DR, and intra-AZ redundancy to guarantee data durability and reliability.

CBR backups are stored in OBS and enjoy a durability of 99.9999999999%.

For details, see [How Durable and Available Is OBS?](#)

6.5 Risk Monitoring

Cloud Eye is a comprehensive monitoring platform that provides real-time insights into resource usage and service health.

With Cloud Eye, you can monitor vaults and backups in real time and receive alarms and notifications for events such as backup creation or deletion failures.

For details on how to define alarm rules for CBR metrics, see [Monitoring](#).

6.6 Fault Recovery

CBR allows you to back up and restore cloud resources, such as ECSs, EVS disks, SFS Turbo file systems, and Workspace desktops.

If a resource fails, you can use backups to quickly restore data and services.


For details, see [Functions](#).

6.7 Certificates

Compliance Certificates

Huawei Cloud services and platforms have obtained various security and compliance certifications from authoritative organizations, such as International Organization for Standardization (ISO). You can [download](#) them from the console.


Figure 6-1 Downloading compliance certificates



Resource Center

Huawei Cloud also provides the following resources to help users meet compliance requirements. For details, see [Resource Center](#).

Figure 6-2 Resource center



6.8 Trusted Services

A trusted service is a Huawei Cloud service that is authorized by Organizations to access and manage organizational resources.

When a management account creates organizational policies, CBR is automatically enabled as a trusted service.

As a trusted service, CBR can access information about organizational units (OUs) and member accounts, and manage resources across the entire organization.

7 Permissions Management

If you need to assign different permissions to personnel in your enterprise, Identity and Access Management (IAM) is a good choice for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you control access to CBR resources. If your Huawei Cloud account does not require IAM for permissions management, you can skip this section.

IAM is a free service. You only pay for the resources used.

With IAM, you can control access to specific Huawei Cloud resources. For example, if you want some software developers in your enterprise to use CBR resources but do not want them to delete CBR resources or perform any other high-risk operations, you can grant permission to use CBR resources but not permission to delete them.

IAM supports role/policy-based authorization and identity policy-based authorization.

The following table describes the differences between the two authorization models.

Table 7-1 Differences between role/policy-based authorization and identity policy-based authorization

Authorization Model	Authorization Using	Permissions	Authorization Method	Scenario
Role/Policy-based authorization	User-permissions-authorization scope	<ul style="list-style-type: none"> • System-defined roles • System-defined policies • Custom policies 	Granting roles or policies to principals	To grant permissions to a user, you must first add the user to a user group and then define the scope of authorization. It is hard to provide fine-grained permissions control using user group-based authorization and a limited number of condition keys. This method is suitable for small- and medium-sized enterprises.
Identity policy-based authorization	Policies	<ul style="list-style-type: none"> • System-defined policies • Custom identity policies 	<ul style="list-style-type: none"> • Assigning identity policies to principals • Attaching identity policies to principals 	You can authorize a user by attaching an identity policy to it. User-specific authorization and a variety of key conditions allow for more fine-grained permissions control. However, this model can be hard to set up. It requires a certain amount of expertise and is suitable for medium- and large-sized enterprises.

Assume that you want to grant IAM users permission to create ECSs in CN North-Beijing4 and OBS buckets in CN South-Guangzhou. With role/policy-based authorization, the administrator needs to create two custom policies and attach both to the IAM users. With identity policy-based authorization, the administrator only needs to create one custom policy with the condition key **g:RequestedRegion**, and then apply the policy to the users. It is more flexible than role/policy-based authorization.

Policies and actions in the two authorization models are not interoperable. Identity policy-based authorization is recommended. For details about system-defined permissions, see [Role/Policy-based Authorization](#) and [Identity Policy-based Authorization](#).

Role/Policy-based Authorization

CBR supports authorization with roles and policies. New IAM users do not have any permissions assigned by default. You need to first add them to one or more groups and attach policies or roles to these groups. The users then inherit

permissions from the groups and can perform specified operations on cloud services based on the permissions they have been assigned.

CBR is a project-level service deployed for specific regions. When you set **Scope** to **Region-specific projects** and select the specified project (for example, **ap-southeast-2**) in the specific region (for example, **AP-Bangkok**), users only have permissions for CBR resources in the selected project. If you set **Scope** to **All resources**, users have permissions for CBR resources in all region-specific projects. When accessing CBR resources, users need to switch to the authorized region.

Table 7-2 lists all system-defined permissions for CBR. System-defined policies in role/policy-based authorization are not interoperable with those in identity policy-based authorization.

Table 7-2 System-defined permissions for CBR

Role/Policy Name	Description	Type
CBR FullAccess	Administrator permissions for CBR. Users with these permissions can operate and use all vaults, backups, and policies.	System-defined policy
CBR BackupsAndVaultsFullAccess	Common user permissions for CBR. Users with these permissions can create, view, and delete vaults and backups, but cannot create, update, or delete policies.	System-defined policy
CBR ReadOnlyAccess	Read-only permissions for CBR. Users with these permissions can only view CBR data.	System-defined policy

Table 7-3 lists the common operations supported by CBR's system-defined permissions.

Table 7-3 Common operations supported by system-defined permissions

Operation	CBR FullAccess	CBR BackupsAndVaults-FullAccess	CBR ReadOnlyAccess
Querying vaults	Supported	Supported	Supported
Creating vaults	Supported	Supported	Not supported
Listing vaults	Supported	Supported	Supported
Updating vaults	Supported	Supported	Not supported
Deleting vaults	Supported	Supported	Not supported
Associating with resources	Supported	Supported	Not supported

Operation	CBR FullAccess	CBR BackupsAndVaults-FullAccess	CBR ReadOnlyAccess
Dissociating from resources	Supported	Supported	Not supported
Creating policies	Supported	Not supported	Not supported
Updating policies	Supported	Not supported	Not supported
Applying policies to vaults	Supported	Supported	Not supported
Removing policies from vaults	Supported	Supported	Not supported
Deleting policies	Supported	Not supported	Not supported
Synchronizing backups	Supported	Supported	Not supported
Replicating vaults	Supported	Supported	Not supported
Performing backups	Supported	Supported	Not supported
Updating subscriptions	Supported	Supported	Not supported
Querying the Agent status	Supported	Supported	Not supported
Deleting backups	Supported	Supported	Not supported
Restoring data from backups	Supported	Supported	Not supported
Replicating backups	Supported	Supported	Not supported
Batch adding or removing tags from vaults	Supported	Supported	Not supported
Adding tags to vaults	Supported	Supported	Not supported
Editing tags	Supported	Supported	Not supported

Identity Policy-based Authorization

CBR supports authorization with identity policies. [Table 7-4](#) lists all the system-defined identity policies for CBR. System-defined policies in identity policy-based authorization are not interoperable with those in role/policy-based authorization.

Table 7-4 System-defined identity policies for CBR

Identity Policy Name	Description	Type
CBRFullAccessPolicy	Administrator permissions for CBR. Users with these permissions can operate all vaults, backups, and policies.	System-defined identity policy
CBRBackupsAndVaults-FullAccessPolicy	Common user permissions for CBR. Users with these permissions can create, view, and delete vaults and backups, but cannot create, update, or delete policies.	System-defined identity policy
CBRReadOnlyAccessPolicy	Read-only permissions for CBR. Users with these permissions can only view CBR data.	System-defined identity policy

[Table 7-5](#) lists the common operations supported by CBR's system-defined identity policies.

Table 7-5 Common operations supported by system-defined identity policies

Operation	CBRFullAccessPolicy	CBRBackupsAndVaultsFullAccessPolicy	CBRReadOnlyAccessPolicy
Querying vaults	Supported	Supported	Supported
Creating vaults	Supported	Supported	Not supported
Listing vaults	Supported	Supported	Supported
Updating vaults	Supported	Supported	Not supported
Deleting vaults	Supported	Supported	Not supported
Associating with resources	Supported	Supported	Not supported
Dissociating from resources	Supported	Supported	Not supported
Creating policies	Supported	Not supported	Not supported

Operation	CBRFullAccessPolicy	CBRBackupsAndVaultsFullAccessPolicy	CBRReadOnlyAccessPolicy
Updating policies	Supported	Not supported	Not supported
Applying policies to vaults	Supported	Supported	Not supported
Removing policies from vaults	Supported	Supported	Not supported
Deleting policies	Supported	Not supported	Not supported
Synchronizing backups	Supported	Supported	Not supported
Replicating vaults	Supported	Supported	Not supported
Performing backups	Supported	Supported	Not supported
Updating subscriptions	Supported	Supported	Not supported
Querying the Agent status	Supported	Supported	Not supported
Deleting backups	Supported	Supported	Not supported
Restoring data from backups	Supported	Supported	Not supported
Replicating backups	Supported	Supported	Not supported
Batch adding or removing tags from vaults	Supported	Supported	Not supported
Adding tags to vaults	Supported	Supported	Not supported
Editing tags	Supported	Supported	Not supported

Roles or Policies that the CBR Console Depends on

Table 7-6 Roles or policies that the CBR console depends on

Console Function	Dependent Services	Roles or Policies Required
Associating ECSs with a vault	ECS	<p>When an IAM user associates ECSs with a vault on the CBR console, the permissions of querying the ECS list and details are required. The user can either use the CBRFullAccessPolicy policy or add the required actions to a custom policy.</p> <p>Required actions:</p> <ul style="list-style-type: none"> ecs:cloudServers:listServerVolumeAttachments ecs:cloudServers:list ecs:cloudServers:showServer
Associating EVS disks with a vault	EVS	<p>When an IAM user associates EVS disks with a vault on the CBR console, the permissions of querying the EVS disk list and details are required. The user can either use the CBRFullAccessPolicy policy or add the required actions to a custom policy.</p> <p>Required actions:</p> <ul style="list-style-type: none"> evs:volumes:list
Associating SFS Turbo file systems with a vault	SFS Turbo	<p>When an IAM user associates SFS Turbo file systems with a vault on the CBR console, the permissions of querying the SFS Turbo file system list and details are required. The user can either use the CBRFullAccessPolicy policy or add the required actions to a custom policy.</p> <p>Required actions:</p> <ul style="list-style-type: none"> sfsturbo:shares:getAllShares
Associating Workspace desktops with a vault	WorkSpace	<p>When an IAM user associates Workspace desktops with a vault on the CBR console, the permissions of querying the Workspace desktop list and details are required.</p> <p>The user can either use the CBR FullAccess policy or add the following actions to a custom policy:</p> <ul style="list-style-type: none"> workspace:desktops:listDetail vpc:securityGroups:get vpc:publicIps:list vpc:ports:get

Console Function	Dependent Services	Roles or Policies Required
Querying a backup and registering an image	IMS	<p>When an IAM user uses a cloud server backup to create a private image on the CBR console, the permission of querying the IMS image list is required.</p> <p>The user can either use the CBRFullAccessPolicy policy or add the required actions to a custom policy.</p> <p>Required actions: ims:images:list</p>

Helpful Links

- [What Is IAM?](#)

8 Notes and Constraints

This section describes the constraints on using CBR.

Specifications

Table 8-1 Specifications

Item	Limitation	Description
Number of backups	Unlimited	-
Backup capacity (GB)	Unlimited	You are advised not to back up a server whose disk size exceeds 4 TB.

Naming

Table 8-2 Naming

Item	Description
Vault name	A name can contain 1 to 64 characters. Only letters, digits, underscores (_), and hyphens (-) are allowed.
Backup name	A name can contain 1 to 64 characters. Only letters, digits, underscores (_), and hyphens (-) are allowed.
Tag key	<ul style="list-style-type: none"> The tag key can contain 1 to 36 Unicode characters. It cannot be left blank, cannot start or end with spaces or contain non-printable ASCII (0-31) characters or any of the following special characters: =*<>\\, /

Item	Description
Tag value	<ul style="list-style-type: none"> • The tag value can contain a maximum of 43 Unicode characters. • It can be an empty string, but cannot start or end with spaces or contain non-printable ASCII (0-31) characters or the following special characters: =*<>\\, /

Operations

Table 8-3 Operations

Operations	Constraints
Restoring data	<ul style="list-style-type: none"> • Only backups in the Available or Locked vaults can be used to restore data. • Concurrent data restoration is not supported. • An SFS Turbo file system backup cannot be used to restore data to the original file system. • Restoring from a cloud server backup: <ol style="list-style-type: none"> 1. A data disk backup cannot be restored to the system disk. 2. Data cannot be restored to servers in the Faulty state. • Restoring from a cloud disk backup: <ol style="list-style-type: none"> 1. Backup and restoration of local disks are not supported. 2. You can back up specific disks on a server, but such a backup must be restored as a whole. File- or directory-level restoration is not supported. 3. If the server OS is changed after the system disk is backed up, the system disk backup cannot be restored to the original system disk due to reasons such as disk UUID changes. You can use the system disk backup to create a new disk and copy data to the original system disk. 4. Backups can only be restored to original disks and cannot be restored to other disks. If you want to restore a backup to a different disk, use the backup to create a new disk. • Restoring from a file backup: <ol style="list-style-type: none"> 1. The Agent on the server must be Normal. 2. You are advised not to restore file backups when applications are running. Stop the applications and then restore files.

Operations	Constraints
Creating a backup policy	<ul style="list-style-type: none"> • A user can create a maximum of 32 backup policies in each region. A vault can be associated with only one backup policy. • You can apply backup policies to server backup vaults, SFS Turbo backup vaults, and disk backup vaults. • A backup policy must be enabled before it can be used for periodic backups. Expired backups are also periodically deleted. • When both a backup time and a replication time are configured, ensure that replication starts after backup is complete. Or, replication may fail. • When expired backups are deleted, automatic backups will be deleted, but manual backups will not. • The initial backup is a full backup, and subsequent backups are incremental ones.
Creating a replication policy	<ul style="list-style-type: none"> • A user can create a maximum of 32 replication policies in each region. A vault can be associated with only one replication policy. • You can only apply replication policies to server backup vaults, SFS Turbo backup vaults, and hybrid cloud backup vaults.
Associating with resources	<ul style="list-style-type: none"> • A vault can be associated with a maximum of 256 resources.

Operations	Constraints
Creating a backup	<ul style="list-style-type: none"> ● Only servers in the Running or Stopped state can be backed up. ● Only disks in the Available or In-use state can be backed up. ● Only file systems in the Available state can be backed up. ● Only desktops in the Available or In-use state can be backed up. ● Desktops in the Frozen state can be associated with a desktop vault, but cannot be backed up. ● Frozen disks and servers in the retention period cannot be backed up. ● If a resource is being backed up, other policy-based or manual backups will not be performed on this resource. ● Resources with more than 64 disks cannot be backed up. ● 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. ● The minimum interval between two full backups is one day. ● You are not advised to back up the same disk using both cloud disk backup and cloud server backup. If a disk has been backed up using both cloud disk backup and cloud server backup, the backup type of the disk may be displayed incorrectly (the first backup of the disk may be displayed as an incremental backup). However, the backup can still be used to restore resource data. ● You are not advised to associate a shared disk with multiple ECSs for cloud server backup. It is recommended that you associate disks except for the shared disk for cloud server backup and back up the shared disk separately. For details, see Associating Resources with a Vault. ● Shared encrypted disks associated with multiple cloud servers cannot be backed up at the same time. If you need to back up multiple cloud servers at the same time, back up one cloud server first. After the backup is complete, back up all the cloud servers at the same time.
Deleting a backup	<ul style="list-style-type: none"> ● Only backups in the Available or Error state can be deleted. ● Backups in a Deleting vault cannot be deleted.

Operations	Constraints
File backup	<ul style="list-style-type: none"> ● Before backing up a file, make sure it is not being used or modified by any application. If the file is being used or modified but the backup client has only read permissions, the backup will be incomplete. ● Before backing up a file, make sure it is not being used by any process. If the file is being used but the backup client has only read permissions, the backup will be incomplete. ● One client with the Agent installed can have a maximum of 8 directories added. ● Each server can only have one Agent installed. ● There is no limitation on the number of resources where the Agent can be installed. ● A single directory can contain a maximum of 500,000 files, and you are advised to reserve at least 4 GB of memory on each backup client to perform file backups. ● A backup client can have a maximum of 100 directories added. ● A path must be an absolute path, for example, a path starting with /, C:\, or D:\. One path can contain a maximum of 200 characters. ● The maximum bandwidth allowed for file backup transmission is 16 Gbit/s. If the maximum bandwidth is reached, flow control will be triggered. ● File backup cannot back up the files stored in SFS file systems that are mounted to cloud servers. ● Backup may fail on directories with frequent file writes in Windows. ● At least 50 Mbit/s network bandwidth is required in cross-cloud or cross-region file backup scenarios. ● Root directories of Windows hosts or servers, such as C:\ and D:\ cannot be backed up.
Database server backup	<ul style="list-style-type: none"> ● Application-consistent backup is currently not supported for cluster applications, such as MySQL Cluster. It is supported only for applications on standalone servers. ● Application-consistent backup blocks database transactions. You are advised to back up databases during off-peak hours. ● If a database is under heavy load—such as during long-running transactions or when large amounts of data are cached—the freeze operation may time out, causing the application-consistent backup to fail. In severe cases, upper-layer service requests may fail. ● Point-in-Time Recovery (PITR) is not supported temporarily.

Operations	Constraints
Migrating a resource	<ul style="list-style-type: none"> • Resources can be migrated only when the source and destination vaults are in the Available or Locked state. • The source and destination vaults must be of the same types. For example, resources in a server backup vault can be migrated to another server backup vault, but cannot be migrated to another disk backup vault. • The remaining capacity of the destination vault must be greater than the size of the resource backup to be migrated. • Cross-account resource migration is currently not supported. • Backups cannot be migrated between single-AZ and multi-AZ backup vaults. • The source and destination vaults must be in the same region.
Auto capacity expansion	<ul style="list-style-type: none"> • Auto capacity expansion does not take effect if it is enabled after the vault is full.
Replicating a backup	<ul style="list-style-type: none"> • Cloud disk backups cannot be replicated to other regions. • Only server backups in the Available or Locked vaults can be replicated. • Only replication-supported regions can be selected as destination regions. • Only backups can be replicated. Backup replicas cannot be replicated again but can be used to create images or SFS Turbo file systems. • A backup vault can be replicated to different destination regions. For manual and policy-based replication, a vault can only be replicated to a destination region once. It cannot be replicated to that region again, even if its backups have been deleted. For manual replication, a backup can be manually replicated to the destination region as long as it has no replica in that region. A backup can be manually replicated again if its replica in the destination region has been deleted. <p>For more information, see Replicating a Backup Across Regions.</p>

Operations	Constraints
Backup locking	<ul style="list-style-type: none"> ● Backup locking cannot be disabled after it is enabled. ● After backup locking is enabled, associated resources cannot be dissociated or migrated. ● Backup locking does not affect normal backup, restoration, and replication operations. ● After backup locking is enabled, policy-based backups can only be automatically deleted upon expiration. ● Manual backups are not affected by backup locking and can be manually deleted. ● After backup locking is enabled, pay-per-use vaults cannot be deleted if they contain backups, but yearly/monthly vaults can be unsubscribed from. ● Backup locking is not supported for VMware backups.

9 CBR and Other Services

CBR-related Services

Table 9-1 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 Scalable File Service Turbo (SFS Turbo) file systems and uses the backup to create new file systems to restore lost or corrupted data.	SFS Turbo	Creating an SFS Turbo Backup
CBR backs up data of Workspace desktops and uses the backup to restore lost or corrupted data.	Workspace	Creating a Desktop Backup
CBR stores server backups securely in OBS.	OBS	What Is CBR?
CBR backs up data of EVS disks and uses the backup to create new disks.	EVS	Creating a Cloud Disk Backup
Cloud Trace Service (CTS) records operations on CBR resources, facilitating future queries, audits, and backtracking.	CTS	Recording CBR Operations Using CTS

Function	Related Service	Reference
IAM is a self-service system for enterprise administrators to manage cloud resources. It provides user identity management and access control functions. When multiple users within an enterprise need to use CBR, the enterprise administrator can use IAM to create IAM users and control these users' access to CBR resources.	IAM	Permissions management
Tag Management Service (TMS) enables you to add preset tags to CBR vaults to facilitate vault management.	TMS	Managing Vault Tags

10 Basic Concepts

[10.1 CBR Concepts](#)

[10.2 Project and Enterprise Project](#)

[10.3 Region and AZ](#)

10.1 CBR Concepts

Vault

CBR stores backups in vaults. Vaults can be either backup vaults or replication vaults.

- Backup vaults store resource backups.
 - **Server backup vaults:** store backups of non-database servers or database servers. You can associate servers with vaults and apply an automatic backup policy to schedule automatic backups. The backups can then be used to restore server data.
 - **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.
 - **Desktop backup vaults:** store only backups of Workspace desktops. You can associate desktops with a desktop backup vault and apply a backup policy to schedule automatic backups.
- Replication vaults store only replicas of backups, and such replicas cannot be replicated again. Replication vaults that store replicas of server backups include those for non-database servers and those for database servers.

For details about how to quickly create a vault, see [Purchasing a Server Backup Vault](#).

Backup

A backup is a copy of a particular chunk of data and is usually stored elsewhere so that it can 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.

- One-time backup names can be user-defined or assigned automatically by the system. System-generated names follow the **manualbk_XXXX** format.
- A periodic backup is automatically named **autobk_XXXX** by CBR.

For details about how to quickly create a backup, see [Creating a Server Backup](#).

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.

For details about how to create a backup policy, see [Creating a Backup Policy](#).

Cross-region Replication

Replication is the process of replicating backups from one region to another. You can use the replicas in the destination region to create images and provision servers.

You can manually replicate a single cloud server backup. You can also configure replication rules in a policy to periodically replicate backups, including those that have not been replicated or failed to be replicated to the destination region.

For example, if you want to back up a server, select **Backup** for the vault protection type. If you want to replicate backups of this server to a different region, select **Replication** for the vault in this different region.

For details about how to replicate backups across regions, see [Replicating Backups Across Regions](#).

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. If **Instant Restore Support** is **Yes** in the backup details, the backup is an enhanced backup. Otherwise, the backup is a common backup.

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 for new resources 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.

Database Server Backup

There are three types of backups in terms of backup consistency:

- **Inconsistent backup:** An inconsistent backup contains data taken from different points in time. This typically occurs if changes are made to your files or disks during the backup.
- **Crash-consistent backup:** A crash-consistent backup captures all data on disks at the time of the backup and does not capture data in memory or any pending I/O operations. Although it cannot ensure application consistency, disks are checked by **chkdsk** upon operating system restart to restore damaged data and undo logs are used by databases to keep data consistent.
- **Application-consistent backup:** An application-consistent backup captures data in memory or any pending I/O operations and allows applications to achieve a quiescent and consistent state.

CBR cloud server backup supports both crash-consistent backup and application-consistent backup (also called database server backup). Install the Agent before enabling application-consistent backup to prevent the database server backup from failing.

For details about how to create a database server backup, see [Database Server Backup](#).

Periodic Full Backup

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

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

Periodic full backups occupy more storage space than incremental backups.

For details about how to create a periodic full backup, see the descriptions about the full-backup parameters in [Creating a Backup Policy](#).

10.2 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.

10.3 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.

- Regions are divided based on geographical location and network latency. Public services, such as Elastic Cloud Server (ECS), Elastic Volume Service (EVS), Object Storage Service (OBS), Virtual Private Cloud (VPC), Elastic IP (EIP), and Image Management Service (IMS), are shared within the same region. Regions are classified into universal regions and dedicated regions. A universal region provides universal cloud services for common tenants. A dedicated region provides specific services for specific tenants.
- An AZ contains one or more physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proof, and electricity facilities. Within an AZ, computing, network, storage, and other resources are logically divided into multiple clusters.

Figure 10-1 shows the relationship between regions and AZs.

Figure 10-1 Regions and AZs



Huawei Cloud provides services in many regions around the world. You can select a region and an AZ based on requirements. For more information, see [Huawei Cloud Global Products and Services](#).

Selecting a Region

When selecting a region, consider the following factors:

- Location

It is recommended that you select the closest region for lower network latency and quick access.

- If your target users are in Asia Pacific (excluding the Chinese mainland), select the **CN-Hong Kong**, **AP-Bangkok**, or **AP-Singapore** region.
- If your target users are in Africa, select the **AF-Johannesburg** region.
- If your target users are in Latin America, select the **LA-Santiago** region.

 **NOTE**

The **LA-Santiago** region is located in Chile.

- Resource price

Resource prices may vary in different regions. For details, see [Product Pricing Details](#).

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](#).