

Cloud Backup and Recovery

Service Overview

Issue 01
Date 2025-02-05



Copyright © Huawei Technologies Co., Ltd. 2025. 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 Technologies Co., Ltd.

Trademarks and Permissions



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

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

Notice

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

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

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process*. For details about this process, visit the following web page:

<https://www.huawei.com/en/psirt/vul-response-process>

For vulnerability information, enterprise customers can visit the following web page:

<https://securitybulletin.huawei.com/enterprise/en/security-advisory>

Contents

1 CBR Infographics.....	1
2 What Is CBR?.....	3
3 Advantages.....	7
4 Application Scenarios.....	8
5 Functions.....	9
6 Billing.....	12
7 Permissions Management.....	15
8 Constraints.....	18
9 CBR and Other Services.....	21
10 Basic Concepts.....	23
10.1 CBR Concepts.....	23
10.2 Project and Enterprise Project.....	25
10.3 Region and AZ.....	25

1 CBR Infographics




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 Backup and Recovery (CBR) to protect 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, and on-premises VMware virtual environments. 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 and BMSs. 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

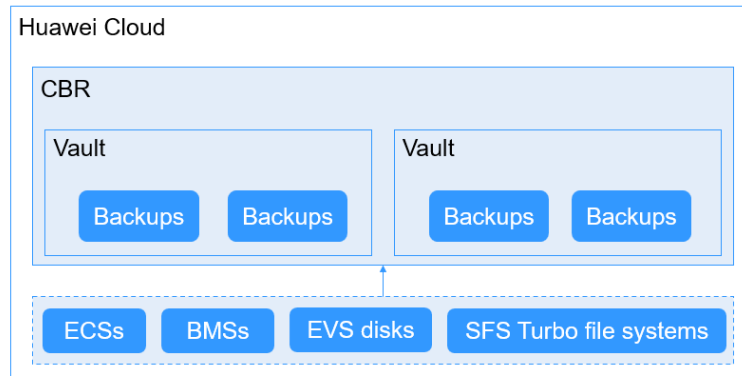
CBR stores backups in vaults. Before creating a backup, you need to create at least one vault and associate the resources you want to back up with the vaults. Then the resources can be backed up to the associated vaults.

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 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
What to back up	All disks (system and data disks) on a server or part of disks and cloud servers (with applications)	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 2-2 compares the two backup options.

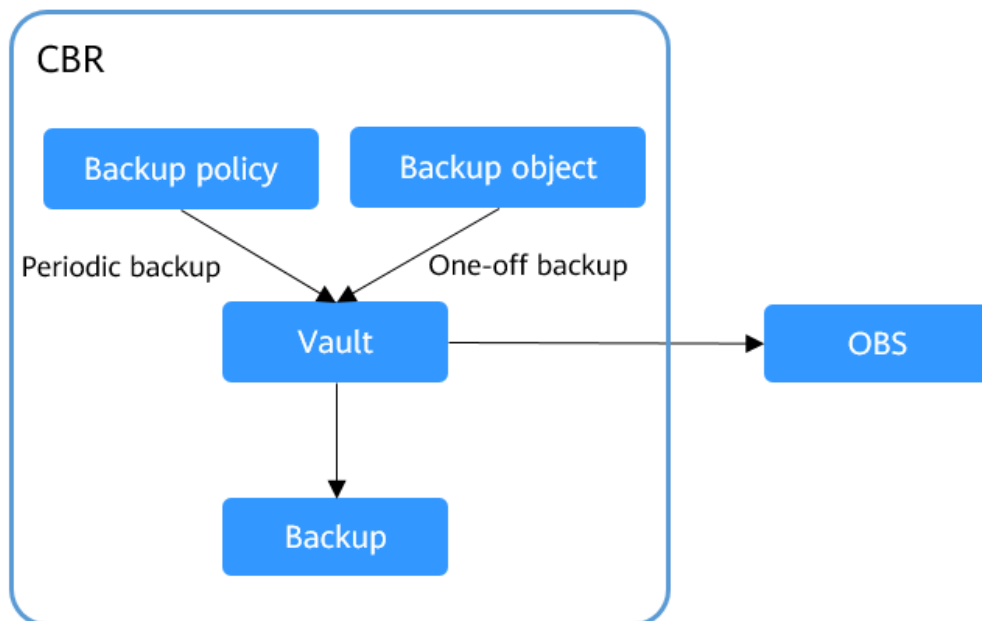
Table 2-2 One-off backup and periodic backup

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 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 2-2** shows the use of the two backup options.

Theoretically, you can create as many backups for a resource as needed. This number is not limited.

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-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 [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) 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.

4 Application Scenarios

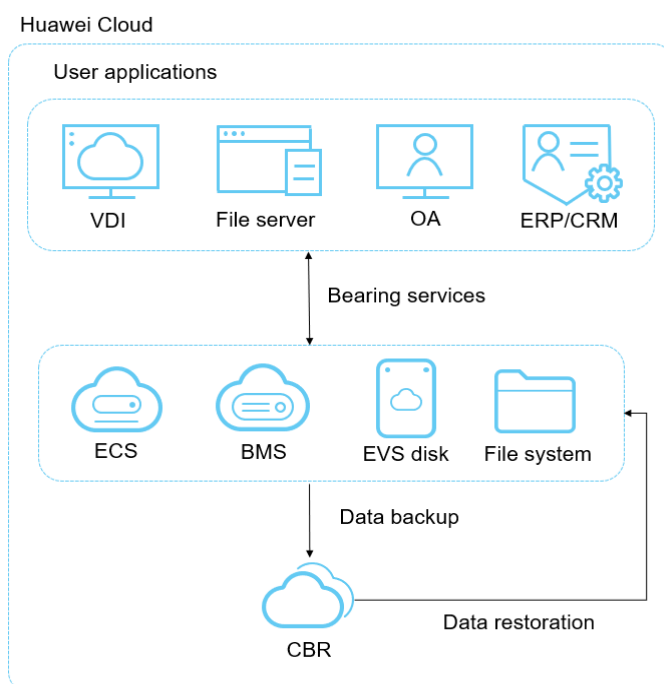
CBR is ideal for data backup and restoration. The backups can maximize your data security and consistency.

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

Figure 4-1 Data backup and restoration



5 Functions

Table 5-1 lists the functions of CBR.

Before using CBR functions, it is recommended that you learn about **basic CBR concepts**.

Table 5-1 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.
Cloud disk backup	Policy-based backup	You can create, modify, or delete a backup policy. A backup policy defines the schedule and retention for automatic backups.
Cloud disk backup	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.
Cloud disk backup	Disk restoration using backups	When a disk is faulty, or their data is lost, you can use a backup to quickly restore the data.
Cloud disk backup	Disk creation using backups	You can use a disk backup to create a disk that contains the same data as the backup.
Cloud disk backup	Backup sharing	You can share a disk backup with other accounts to allow them to use the backup to create disks.

Category	Function	Description
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. You can use CBR to back up an entire server to protect their data, especially when high data consistency is required, such as in RAID clusters.
Cloud server backup	Backup of specific disks on a server	You can create a single backup for multiple disks on a server to save the vault space.
Cloud server backup	Policy-based backup	You can create, modify, or delete a backup policy. A backup policy defines the schedule and retention for automatic backups.
Cloud server backup	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.
Cloud server backup	Server restoration using backups	When a server is faulty, or their data is lost, you can use a backup to quickly restore the data.
Cloud server backup	Backup sharing	You can share a server backup with other accounts to allow them to use the backup to create servers.
Cloud server backup	Image creation using server backups	You can create images from ECS backups and then use the images to quickly provision ECSs to restore service.
Cloud server backup	Database server backup	Cloud server backup supports application-consistent backup in addition to crash-consistent backup. You can use cloud server backup to back up ECSs running MySQL or SAP HANA databases, because application-consistent backup ensures that the backed-up data is transactionally consistent.

Category	Function	Description
SFS Turbo backup	Manual SFS Turbo backup	You can back up SFS Turbo file systems and use the backups to create a new SFS Turbo file system.
SFS Turbo backup	Policy-based backup	You can create, modify, or delete a backup policy. A backup policy defines the schedule and retention for automatic backups.
SFS Turbo backup	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.
SFS Turbo backup	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.

6 Billing

Billing Items

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

Category	Billing Item	Description	Billing Mode
Storage space	Disk backup vault	If cloud disks need to be backed up, buy disk backup vaults to store the backups.	Pay-per-use Yearly/ Monthly
	Server backup vault	If cloud servers (without applications) need to be backed up, buy server backup vaults to store the backups.	Pay-per-use Yearly/ Monthly
	SFS Turbo backup vault	If SFS Turbo file systems need to be backed up, buy SFS Turbo backup vaults to store the backups.	Pay-per-use Yearly/ Monthly
	Database server backup vault	If cloud servers (with applications) need to be backed up, buy database server backup vaults to store the backups. You need to enable Application-Consistent Backup on the Buy Server Backup Vault page before using database server backup vaults. For more information, see Application-Consistent Backup Overview .	Pay-per-use Yearly/ Monthly
	Hybrid cloud backup vault	If backups of on-premises VMware VMs need to be stored, buy hybrid cloud backup vaults.	Pay-per-use Yearly/ Monthly

Category	Billing Item	Description	Billing Mode
	Replication vault	If you need to replicate backups to another region, buy replication vaults in the destination region.	Pay-per-use Yearly/ Monthly
Data traffic	Outbound traffic over the Internet	If hybrid cloud backups on the cloud are used to restore data to on-premises IDCs, outbound traffic is charged.	Limited-time free trial

Billing Modes

Two billing modes are available: pay-per-use and yearly/monthly. Select a billing mode that best suits your business needs.

- **Pay-per-use**

You pay for the duration you use the resources. Prices are calculated by the hour, and no minimum fee is required.

- **Yearly/Monthly**

Yearly/monthly is a prepaid billing mode. You are billed based on the subscription duration you specify. This mode is ideal when the resource use duration is predictable. A longer subscription often means a lower cost.

CBR also provides replication traffic packages for cross-region replication. If you do not have such a package, you will be billed per use.

Billing Examples

Example 1

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

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

Changing Billing Mode

- Yearly/monthly is a prepaid billing mode. You are billed based on the subscription duration you specify. This mode is ideal when the resource use duration is predictable. A longer subscription often means a lower cost.
- Pay-per-use is a postpaid billing mode. You are billed based on your resource usage and you can increase or delete resources anytime.

Expiration

For details, see [Service Suspension and Resource Release](#).

Renewal

Choose **More > Renew** in the **Operation** column of the yearly/monthly vault to renew your subscription. For more information about renewal, including auto-renewal, exporting the renewal history, and changing subscriptions, see [Renewal Management](#).

7 Permissions Management

If you need to assign different permissions to personnel in your enterprise to access your CBR resources, 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 to securely access your Huawei Cloud resources.

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

If your Huawei Cloud account does not require individual IAM users for permissions management, you can skip this section.

IAM is a free service. You only pay for the resources in your account. For more information about IAM, see [IAM Service Overview](#).

CBR Permissions

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 projects in the specified regions, the users only have permissions for CBR resources in the selected projects. If you set **Scope** to **All resources**, the users have permissions for CBR resources in all region-specific projects. When accessing CBR resources, the users need to switch to the authorized region.

You can grant permissions by using roles and policies.

- **Roles:** A coarse-grained authorization strategy provided by IAM to assign permissions based on users' job responsibilities. Only a limited number of service-level roles are available for authorization. Huawei Cloud services depend on each other. When you grant permissions using roles, you also need to attach any existing role dependencies. Roles are not ideal for fine-grained authorization and least privilege access.

- **Policies:** A fine-grained authorization strategy that defines permissions required to perform operations on specific cloud resources under certain conditions. This type of authorization is more flexible and is ideal for least privilege access. For example, you can grant users only permission to manage ECSs of a certain type. A majority of fine-grained policies contain permissions for specific APIs, and permissions are defined using API actions. For the API actions supported by CBR, see [Permissions Policies and Supported Actions](#).

Table 7-1 lists all the system-defined permissions for CBR.

Table 7-1 System-defined permissions for CBR

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 BackupsAndVaults-FullAccess	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 7-2 lists the common operations supported by system-defined permissions of CBR.

Table 7-2 Common operations supported by system-defined permissions of CBR

Operation	CBR FullAccess	CBR BackupsAndVaultsFullAccess	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 resources	Supported	Supported	Not supported
Dissociating resources	Supported	Supported	Not supported
Creating policies	Supported	Not supported	Not supported

Operation	CBR FullAccess	CBR BackupsAndVaultsFullAccess	CBR ReadOnlyAccess
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
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
Associating vaults	Supported	Supported	Not supported
Batch adding or deleting vault tags	Supported	Supported	Not supported
Adding vault tags	Supported	Supported	Not supported
Editing tags	Supported	Supported	Not supported

8 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.
- Concurrent data restoration is not supported.
- Auto capacity expansion does not take effect if it is enabled after the vault is full.

Cloud Disk Backup

- Only disks in the **Available** or **In-use** state can be backed up.
- Frozen disks in the retention period cannot be backed up.
- A new disk must be at least as large as the backup's source disk.
- Backup and restoration of local disks are not supported.

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.
- Frozen servers in the retention period cannot be backed up.
- Cloud servers support crash-consistent backup, whereas database servers support application-consistent backup in addition to crash-consistent backup.
- 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.
- 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.

Hybrid Cloud Backup - VMware Backup

- VM backups from the following VMware vSphere versions can be synchronized to the cloud: 5.1, 5.5, 6.0, 6.5, 6.7.
- To achieve better performance and operation experience, only the OSs listed in [Table 8-1](#) allows for cloud recovery. These OSs have passed the the compatibility test. If you only need to restore data to on-premises VMware VMs, there is no restriction on the OS version.
- The VDDK version of VMware 6.7 or earlier VMs must be 6.0.3.
- Backups synchronized to the cloud cannot be used to create cloud servers.
- Backups synchronized to the cloud can only be restored to existing cloud servers running the same type of OS. Recovery of partial disks on a VM is not supported.
- Servers whose system disks are configured with LVM cannot be restored on cloud.
- Before the restoration, configure security groups according to the procedure. Otherwise, the restoration may fail.
- The backup bandwidth should be at least 100 Mbit/s. The sizes of disks on the VMware VMs to be backed up to the cloud should be integers in GB.
- When recovering data on the cloud from a Windows VMware backup, the boot partition number of the source VM system disk must be **2**.
- When recovering data on the cloud from a Linux VMware backup, the system disk of the source VM can have no more than two partitions.

Table 8-1 OSs that support restoration to the cloud

OS	Supported Version
Windows	Windows 7 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016 Windows Server 2019

OS	Supported Version
CentOS	CentOS 6.4 CentOS 6.5 CentOS 7.2 CentOS 7.3 CentOS 7.4 CentOS 7.5 CentOS 7.6 CentOS 7.7
Red Hat	Red Hat 6.4 Red Hat 6.5 Red Hat 7.2

Application-Consistent Backup

Table 8-2 OSs that support installation of the Agent

Database	OS	Version
SQL Server 2008/2012/2019	Windows	Windows Server 2008, 2008 R2, 2012, 2012 R2, and 2019 for x86_64
SQL Server 2014/2016/Enterprise Edition	Windows	Windows Server 2014, 2014 R2, and 2016 Datacenter for x86_64
MySQL 5.5/5.6/5.7	Red Hat	Red Hat Enterprise Linux 6 and 7 for x86_64
	SUSE	SUSE Linux Enterprise Server 11, 12, 15 SP1, 15 SP2 for x86_64
	CentOS	CentOS 6 and 7 for x86_64
	EulerOS	EulerOS 2.2 and 2.3 for x86_64
HANA 1.0/2.0	SUSE	SUSE Linux Enterprise Server 12 for x86_64

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 SFS Turbo file systems and uses the backup to create new file systems to restore lost or corrupted data.	Scalable File Service Turbo (SFS Turbo)	Creating an SFS Turbo Backup
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	Creating a Cloud Disk Backup
Cloud Trace Service (CTS) records operations on CBR resources, facilitating future queries, audits, and backtracking.	CTS	Auditing

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	7 Permissions Management
Tag Management Service (TMS) enables you to add preset tags to CBR vaults to facilitate vault management.	TMS	Managing Vault Tags
Cloud Eye allows you to check the usage and other performance metrics of CBR vaults. This can be done without requiring additional plug-ins.	Cloud Eye	CBR Metrics Creating an Alarm Rule

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

- **Server backup vaults:** store backups of non-database servers or database servers.
- **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 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.

Application-Consistent 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.

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.

Selecting a Region

If your target users are in Europe, select the **EU-Dublin** region.

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.