Cloud Backup and Recovery

User Guide

Issue	01
Date	2025-11-24

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

Trademarks and Permissions


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 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 can be used to restore the original data in the event of data loss.

There are the following types of backups:

- Cloud server backup: uses the consistency snapshot technology to protect data for ECSs. Backups of non-database servers are 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.

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.

Figure 1-1 CBR architecture

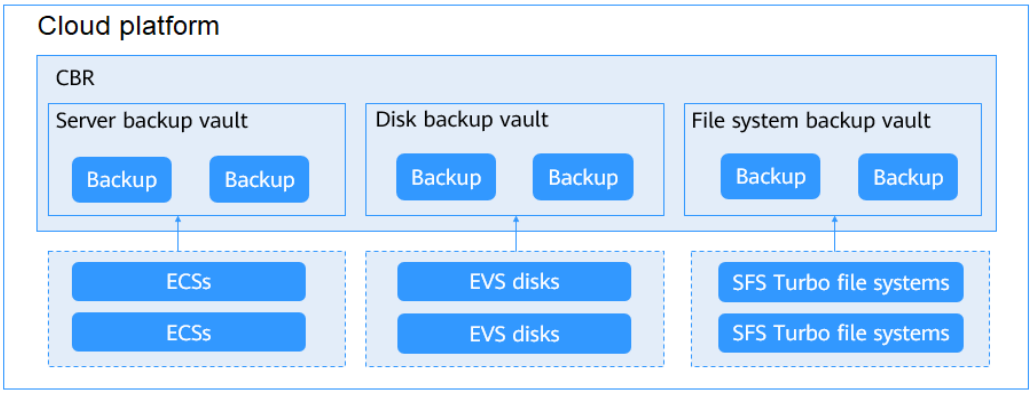

Differences Among the Backup Types

Table 1-1 Differences among the backup types

Item	Cloud Server Backup	Cloud Disk Backup	SFS Turbo Backup
What to back up	All disks (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
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.
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.

Backup Mechanism

CBR in-cloud backup offers block-level backup. The first backup is a full backup of all used data blocks. For example, if a disk size is 100 GB and 40 GB has been used, only the 40 GB 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 that are referenced by other backups will not be deleted, ensuring that these 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 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 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 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 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.

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 1-2 Use of the two backup options

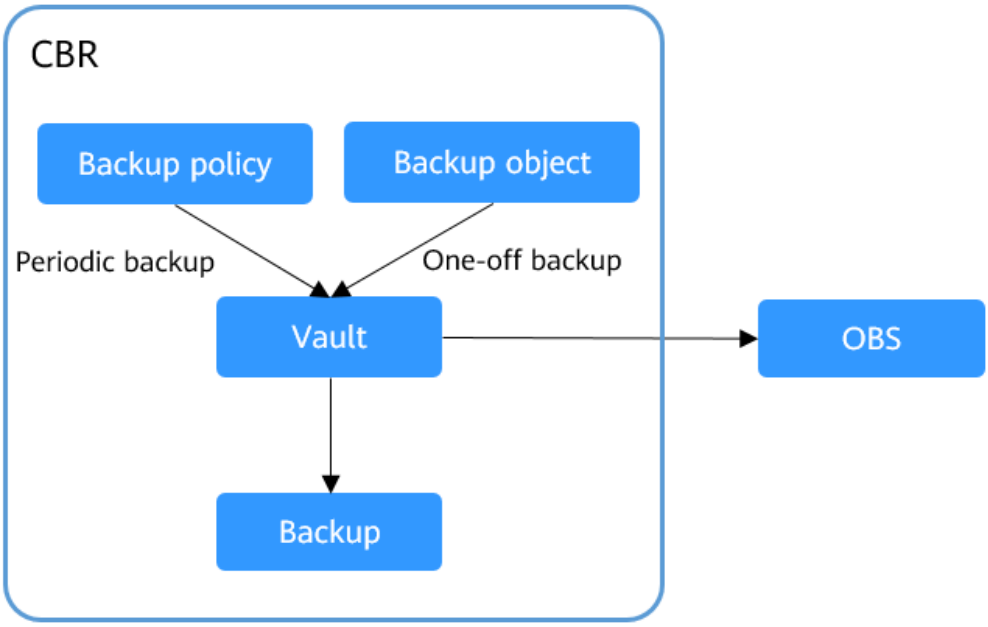

Access to CBR

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*.

1.2 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.

1.3 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

1.4 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 their 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 without the need to install the Agent on servers. It allows you to back up the entire servers.

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 their 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.

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 to control access to your cloud resources.

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

- User management permissions refer to the permissions to manage users, user groups, and user group permissions.
- Resource management permissions refer to the permissions to control operations on cloud service resources.

For details, see **System-defined Permissions**.

1.6 Constraints

General

- A vault can be associated with only one backup policy.
- A vault can be associated with only one replication policy.
- A vault can be associated with a maximum of 256 resources.
- A maximum of 32 backup policies and 32 replication 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.
- When Storage Disaster Recovery Service (SDRS) is used to set up disaster recovery for cloud servers, restorations can be performed at the disaster recovery site only after DR protection is disabled.
- 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 disk backups cannot be replicated to other regions.

Cloud Server Backup

- A maximum of 10 shared disks with a cloud server can be backed up.
- Only cloud server backups in the **Available** or **Locked** vaults can be used to create images and be replicated.
- Multiple EVS disks on a cloud server can be backed up using either crash-consistent or application-consistent backup method.
- Images cannot be created from backups if the number 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.
- Backups can be replicated only to supported regions.
  - A backup can be replicated only when it meets all of the following conditions:
    i. It is an ECS backup.
    ii. It contains system disk data.
    iii. It is in the **Available** state.
  - Only backups can be replicated. Backup replicas cannot be replicated again but can be used to create images.
  - A backup can be replicated to multiple regions but can have only one replica in each destination region. 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.

  - Only replication-supported regions can be selected as destination regions.

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.
- Backups can be replicated only to supported regions.
  - A backup can be replicated only when it meets all of the following conditions:
    i. It is generated from an SFS Turbo file system.
    ii. It is in the **Available** state.
  - Only backups can be replicated. Backup replicas cannot be replicated again but can be used to create SFS Turbo file systems.
  - A backup can be replicated to multiple regions but can have only one replica in each destination region. 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.
  - Only replication-supported regions can be selected as destination regions.

RDS Backup

Table 1-3 OSs supporting the Agent

Database	OS	Supported Versions
SQLServer 2008	Windows	Windows Server 2012, 2012 R2, 2019 for x86_64
SQLServer 2012	Windows	Windows Server 2012, 2012 R2, 2019 for x86_64
SQLServer 2019	Windows	Windows Server 2019 for x86_64
SQLServer 2014/2016/EE	Windows	Windows Server 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 for x86_64
	CentOS	CentOS 6 and 7 for x86_64
	EulerOS	Euler OS 2.2, 2.3 for x86_64
HANA 1.0/2.0	SUSE	SUSE Linux Enterprise Server 12 for x86_64

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 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 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
IAM is a self-service system for enterprise administrators to manage cloud resources. It provides user identity management and access control functions.	IAM	**User Permissions**

1.8 Basic Concepts

1.8.1 CBR Concepts

Vault

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

- Backup vaults store backups of a variety of resources, including servers and disks, and are classified into the following types:
  - **Server backup vaults**: store backups of non-database servers or database servers. You can associate servers with a server backup vault and apply a backup or replication policy to schedule automatic backups or replications.

  - **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.
- 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.

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.

- A one-off backup is named **manualbk**_*xxxx* and can be user- or system-defined.
- A periodic backup is automatically 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.

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.

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.

1.8.2 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

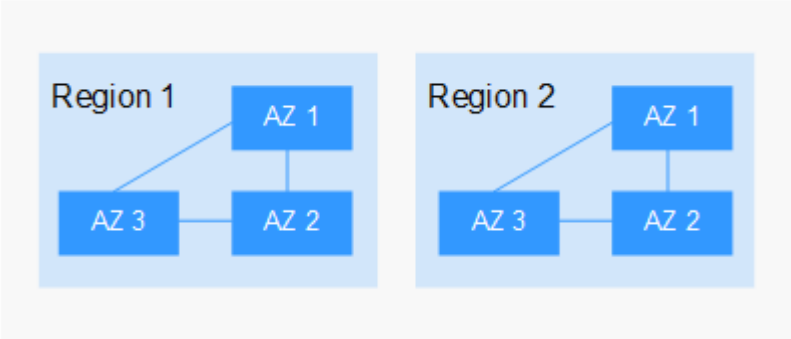

Selecting a Region

You are advised to select a region close to you or your target users. This helps ensure low access latency.

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 Before You Start

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

You can directly associate resources when purchasing a vault. You can configure a backup policy when purchasing a vault to perform automatic backup. Or you can back up manually after associating resources with your vault.

NOTE

When purchasing a vault, you can associate resources with the vault and apply a backup policy to perform automatic backup. You can also manually back up data later.

Figure 2-1 Backup process

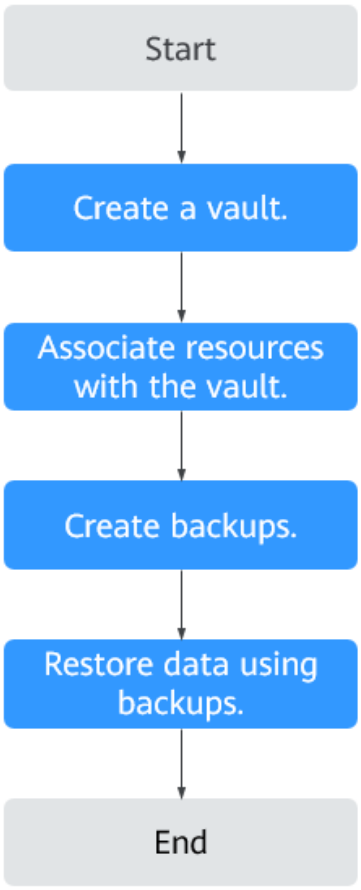

Operation Procedure

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 CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 In the navigation pane, choose **Cloud Server Backups**.

Step 3 In the upper right corner of the page, click **Create Server Backup Vault**.

Step 4 Select a protection type.

- **Backup**: A server backup vault stores server backups.
- **Replication**: A server replication vault stores replicas of server backups. If you select **Replication**, you do not need to select a server.

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

Step 5 Determine whether to enable application-consistent backup.

- If enabled, the vault can be used to store database server backups. For example, you can back up ECSs running MySQL or SAP HANA databases, because application-consistent backup ensures that the backed-up data is transactionally consistent. If an application-consistent backup task fails, CBR automatically performs a non-database server backup task instead. This non-database server backup will be stored in the database server backup vault.
- If disabled, only non-database server backup is performed on associated servers, which is usually used for ECSs not running databases.

Step 6 (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 or disks. You can also select specific disks on a server and associate them with the vault.

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 7 Specify the 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.

You can expand the vault capacity if it becomes insufficient.

Step 8 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 9 Specify a name for the vault.

The name can contain 1 to 64 characters. Only letters, digits, underscores (_), and hyphens (-) are allowed. 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 **Viewing 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 CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 In the navigation pane, choose **Cloud Disk Backups**.

Step 3 In the upper right corner of the page, click **Create Disk Backup Vault**.

Step 4 (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 5 Specify the vault capacity. This capacity is the total space that is required by the disks you want to associate with this vault. Plan the vault capacity and ensure that it is at least the same as the size of the disks you want to back up. If a backup policy is applied to the vault, plan more capacity as required. The capacity ranges from 10 GB to 10,485,760 GB.

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 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 7 Specify a name for the vault.

The name can contain 1 to 64 characters. Only letters, digits, underscores (_), or hyphens (-) are allowed. Example: **vault-612c**.

NOTE

You can also use the default name **vault**_*xxxx*.

Step 8 Complete the creation as prompted.

Step 9 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. In the upper left corner, click  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.
- **Replication**: An SFS Turbo replication vault stores replicas of SFS Turbo backups. If you select **Replication**, you do not need to select any SFS Turbo file system.

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 the vault capacity. This capacity is the total size of the file systems that you want to associate with this vault. Plan the vault capacity and ensure that it is at least the same as the size of the file systems you want to back up. If a backup policy is applied to the vault, plan more capacity as required. The capacity ranges from 10 GB to 10,485,760 GB.

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 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 Specify a name for the vault.

The name can contain 1 to 64 characters. Only letters, digits, underscores (_), or hyphens (-) are allowed. Example: **vault-612c**

NOTE

You can also use the default name **vault**_*xxxx*.

Step 8 Complete the creation as prompted.

Step 9 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 2-1 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 CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 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**. You can view the number of associated resources in the **Associated Servers**, **Associated Disks**, or **Associated File Systems** column.

NOTE

After a new disk is added to the associated server, the system automatically identifies the new disk and backs up the new disk.

----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.

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

Backing up a server does not impact the server performance.

The backup service experiences peak usage between 22:00 and 08:00, during which delays may occur. To ensure optimal performance, it is recommended that you evaluate your service types and stagger backups across 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 CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

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:

- Click **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, locate the target server and click **Perform Backup** in the **Operation** column.

Step 4 Set **Name** and **Description** for the backup.

Table 2-2 Parameter description

Parameter	Description	Remarks
Name	Name of the backup you are creating. A name can contain 1 to 64 characters. Only letters, digits, underscores (_), and hyphens (-) are allowed. **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. A full backup requires a larger capacity than an incremental backup.

Figure 2-2 Full Backup

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

On the **Backups** tab, 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.

The backup service experiences peak usage between 22:00 and 08:00, during which delays may occur. To ensure optimal performance, it is recommended that you evaluate your service types and stagger backups across 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 CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

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, 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 2-3 Parameter description

Parameter	Description	Remarks
Name	Name of the backup you are creating. A name can contain 1 to 64 characters. Only letters, digits, underscores (_), and hyphens (-) are allowed. **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

Parameter	Description	Remarks
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. A full backup requires a larger capacity than an incremental backup.

Figure 2-3 Full Backup

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

On the **Backups** tab, 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.

The backup service experiences peak usage between 22:00 and 08:00, during which delays may occur. To ensure optimal performance, it is recommended that you evaluate your service types and stagger backups across 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. In the upper left corner, click  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:

- Click **Perform Backup** in the **Operation** column. In the file system list, select the file system you want to back 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, locate the target file system and click **Perform Backup** in the **Operation** column.

Step 4 Set **Name** and **Description** for the backup.

Table 2-4 Parameter description

Parameter	Description	Remarks
Name	Name of the backup you are creating. **NOTE** You can also use the default name **manualbk**_*xxxx*. 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**.	manualbk_d819
Description	Description of the backup. It cannot exceed 255 characters.	--

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

On the **Backups** tab, 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 **Creating a File System from an SFS Turbo Backup**.

----End

3 Vault Management

3.1 Viewing 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 CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

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, which can be backup vault or replication vault. • A backup vault stores backups of servers and disks. • A replication vault stores replicas of backups.
Status	Vault status. Table 3-2 describes the vault statuses.

Parameter	Description
Specifications	Vault specifications, which can be server backup or application-consistent backup. • A server backup vault stores backups of non-database servers. • A database server backup vault stores backups of database servers.
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. If an associated resource is deleted, the resource is still counted in the number of associated resources.

Step 3 On the **Vaults** tab, 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**.

Status	Attribute	Description
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 to view the error cause. If the error persists, contact technical support.

- Search a vault by its name or ID.

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 CBR console.
2. In the upper left corner, click and select a region.
3. Click and choose **Storage** > **Cloud Backup and Recovery**.

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 .

After deletion, the system displays the result in the upper right corner. You can return to the vault list. If the vault does not exist anymore, the deletion is successful.

----End

3.3 Dissociating Resources from a Vault

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

After a resource is dissociated from a vault, the vault's backup or replication 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 CBR console.
2. In the upper left corner, click and select a region.
3. Click and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 Locate the target vault and click its name.

Step 3 In this example, cloud servers will be used as an example 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 **OK**.

After the disassociation, the result is displayed in the upper right corner. If the target server is not displayed in the **Associated Servers** list, the disassociation is successful.

----End

3.4 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 CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 Locate 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** GB.

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

3.5 Changing Vault Specifications

Server backup vaults and replication vaults both have two specifications: those for common server backups and those for database backups.

- Common server backups are backups of non-database servers.
- Database server backups are backups of database servers.

If you need to back up database servers, change the specifications of the target vault from server backup to application-consistent backup.

You can only change the specifications of a vault from server backup to application-consistent backup, but not the other way around.

Procedure

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 On the **Cloud Server Backups** page, find the target vault. Choose **More** > **Change Specifications** in the **Operation** column of the vault.

Step 3 **Application-Consistent Backup** is preset for **Backup Type**. Click **Next**.

Step 4 Click **Submit** and complete the payment. The system automatically changes the vault specifications.

After the change, the vault specifications in the vault list are changed from **Server backup** to **Application-consistent backup**.

----End

3.6 Replicating a Vault Across Regions

Scenarios

Cloud server backup vaults allow you to **replicate all backups in vaults to replication vaults of the same account in another region**. Replicas of server backups in the destination region can be used to create images and provision servers. to quickly deploy services across regions.

There are two methods available for replicating a vault.

- Manual replication: Select a backup vault and manually replicate it.
- Policy-based replication: Configure a replication policy to periodically replicate backups that have not been replicated or failed to be replicated to the destination region.

Constraints

- Disk backup vaults cannot be replicated to other regions.
- Backups can be replicated to vaults in different regions. Any traffic costs associated with this operation will be attributed to the source region. Backup replicas take up the replication vault space.
- A server backup vault can be replicated only when it contains at least one backup that meets all the following conditions:
  a. It is an ECS backup.
  b. It contains system disk data.
  c. It is in the **Available** state.
- Only backup vaults of the same type can be replicated. Replicated vaults cannot be replicated again but their replicas can be used to create images.
- A backup vault can be replicated to different destination regions. For manual and policy-based vault 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.
- Only replication-supported regions can be selected as destination regions.

Procedure

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 On the **Vaults** tab, find the target backup vault.

Step 3 Choose **More** > **Create Replica** in the **Operation** column of the vault.

Step 4 In the displayed dialog box, configure the parameters as described in **Table 3-3**.

Table 3-3 Parameter description

Parameter	Description
Destination Region	Region to which the vault is replicated Only the regions that support replication will be displayed. • If the selected region contains only one project, you can directly select the region name. • If the selected region has multiple projects, the default project of the region is preselected. You can select another project if needed.
Destination Vault	A replication vault in the destination region.

Step 5 Click **OK**.

Step 6 After the replication is complete, you can switch to the destination region to view generated replicas. For details, see **Viewing a Vault**. You can then use replicas to create images.

----End

4 Backup Management

4.1 Viewing a Backup

Scenario

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.

Procedure

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

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 above 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** above the backup list.

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

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

----End

4.2 Sharing a Backup

Scenarios

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

Context

For sharers:

- Backups can only be shared among projects of 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.

For recipients:

- 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 sharing request. After accepting the backup, the recipient can use the backup to create 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.

Initiating Backup Sharing

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

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.

Step 4 Click the **Share Backup** tab.

Step 5 Enter the account name of the recipient.

Step 6 Click **Add**. The account and project to be added is displayed in the list. You can continue to add accounts. A backup can be shared to a maximum of **10** projects.

Step 7 Click **OK**.

Go back to the backup list, click the backup name to go to the backup details page, and click the **Share List** tab to view the shared backup.

----End

Accepting the Shared Backup

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click and select a region.
3. Click and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 Click the **Backups** tab on the cloud server or disk backup vault page 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: Determines whether to enable automatic association for the vault. If you select **Configure**, the vault automatically scans for and associates servers that have not been backed up and performs backup in the next backup period.

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

----End

Canceling Backup Sharing

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click and select a region.
3. Click and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 On the cloud server backup page or cloud disk backup page, 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.

Step 4 Click the **Cancel Sharing** tab, select the projects you want to cancel sharing, and click **OK**.

Return to the backup list, click the backup name to go to the backup details page, and click the **Share List** tab to view the unshared backups.

----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**.

NOTE

- Backups are not stored on a server. Deleting backups has no impact on the server performance.
- If a backup has been created and the next backup task is in progress, CBR will not allow you to delete the most recent backup. You can delete the backup only after the backup task is complete.
- CBR automatically creates snapshots during backup and retains the latest snapshot for each disk. If a disk already has a backup, after another backup, the old snapshot will be deleted and the latest one will be retained.

Prerequisites

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

Procedure

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 Click the **Backups** tab. 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 . Return to the backup list. If the target backup does not exist, the backup is deleted successfully.

----End

4.4 Replicating Backups Across Regions

Scenarios

Cross-region replication of server backup vaults allows you to **replicate backups from one region to another in the same account**.

Replicas of server backups can be used to create images and provision servers.

With cross-region replication, you can quickly deploy services in a different region. The new resources created from replicated backups are in the same state as the original resources when you took the backup.

You can replicate backups in either of the following methods on the CBR console:

- Select a backup from the backup list and manually replicate it.
- Select a backup vault and manually replicate it. Alternatively, you can configure a replication policy to periodically replicate backups that have not been replicated or failed to be replicated to the destination region.

This section uses the first method to describe how to replicate a backup. For details about the second method, see **Replicating a Vault Across Regions**.

Constraints

- A server backup can be replicated only when it meets all the following conditions:
  a. It is an ECS backup.
  b. It contains system disk data.
  c. It is in the **Available** state.
- Only backups or backup vaults can be replicated. Replicated backups and vaults cannot be replicated again but can be used to create images.
- A backup can be replicated to multiple regions but can have only one replica in each destination region. Replication rules vary with the replication method:

  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.

  For manual and policy-based vault 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.
- Only replication-supported regions can be selected as destination regions.

Create a Cross-region Backup Replication

Step 1 Click in the upper left corner of the management console and select the region where the backups to be replicated are located.

Step 2 Click the **Backups** tab. Locate the desired backup. For details, see **Viewing a Backup**.

Step 3 Choose **More** > **Create Replica** in the **Operation** column of the backup.

Step 4 In the displayed dialog box, configure the parameters as described in **Table 4-2**.

Table 4-2 Parameter description

Parameter	Description
Name	Replica name. You can enter a custom name or use the default name **resource name_xxxx**. A name can contain 1 to 64 characters. Only letters, digits, underscores (_), and hyphens (-) are allowed.
Description	Replica description. This parameter is optional. It cannot exceed 255 characters.
Destination Region	Region that the backup is replicated to. Only the regions that support replication will be displayed. • If the selected region contains only one project, you can directly select the region name. • If the selected region has multiple projects, the default project of the region is preselected. You can select another project if needed.
Destination Vault	A replication vault in the destination region You can replicate backups to vaults in multiple destination regions. Creating replica will replicate all backups in the source vault to the destination vault.

NOTE

The traffic for cross-region replication is the size of the replicated backup.

Step 5 Click **OK**.

Step 6 After the replication is complete, you can switch to the destination region to view generated replicas. For details, see **Viewing a Backup**. You can then use the backups to create images.

----End

5 Policy Management

5.1 Viewing the Policy of a Vault

After a policy is created, you can view the policy in the policy list or on the vault details page.

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. Click in the upper left corner and select a region.
3. Click and choose **Storage** > **Cloud Backup and Recovery** > **Policies**.

Step 2 View the created policy in the policy list. In the search box above the list, you can select an attribute or enter a keyword to search for policies.

Step 3 After creating a policy and applying it to a vault, you can view the policy on the vault details page. Locate the desired vault and click its name to go to the details page.

Step 4 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.

You can use the default backup policy or create a custom one to enable periodic data backups for resources.

The backup service experiences peak usage between 22:00 and 08:00, during which delays may occur. To ensure optimal performance, it is recommended that you evaluate your service types and stagger backups across discrete time periods.

Constraints

- 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.
- A user can create a maximum of 32 backup policies.
- 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.
- 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 subsequent backups.
- The minimum interval between two full backups is one day.

Procedure

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 Choose **Policies** in the 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 example, select the backup policy.	Backup policy

Parameter	Description	Example Value
Policy Name	Backup policy name You can enter a custom name or use the default name **policy_xxxx**. A name can contain 1 to 64 characters. Only letters, digits, underscores (_), and hyphens (-) are allowed.	backup_policy
Status	Whether to enable the backup policy. By default, this function is enabled. CBR backs up resources to vaults and deletes expired backups only after a backup policy is applied to the vaults.	Enabled
Execution Time	Time when the backup is executed. The default backup time is 22:00. Backups can be scheduled at the beginning of each hour, and you can select multiple hours. It is recommended that backups be performed during off-peak hours or when no services are running. The backup service experiences peak usage between 22:00 and 08:00, during which delays may occur. To ensure optimal performance, it is recommended that you evaluate your service types and stagger backups across discrete time periods. **NOTICE** ● 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. 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. ● The execution times refer to the local times of clients, not the time zone and times of the region.	00:00, 02:00

Parameter	Description	Example Value
Backup Cycle	Select a backup cycle. • Weekly Specifies on which days of each week the backup task will be executed. You can select multiple days. • Day-based Specifies the interval (every 1 to 30 days) for executing the backup task.	Every day 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. It is recommended that backups be performed during off-peak hours or when no services are running.

Parameter	Description	Example Value
Retention Rule	Rule that specifies how backups will be retained By default, backups are retained for one month. • **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	6 months

Parameter	Description	Example Value
	NOTE – 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 during 08:00 to 20:00 in batches**. For example, if a backup expired at 20:00 on November 23, 2024, it was deleted during 08:00 to 20:00 on November 24, 2024. In this way, backup data can be deleted during off-peak hours. – 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. Instead, it will be forcibly retained. If you delete the image created from the backup, the backup will be retained based on the original retention rule applied. (If the backup expires or is not within the most recent backups specified in the retention rule, CBR will automatically delete the backup.) – A maximum of 10 failed periodic backups are retained. They are retained for one month and can be deleted manually. – If a backup has been created and the next backup task is in progress, CBR will not allow you to delete the most recent backup. You can delete the backup only after the ongoing backup task is complete.	

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 Create Now. After the backup policy is created, you can view it in the backup policy list.

Step 5 Locate the desired vault and choose **More** > **Apply Backup Policy** in the **Operation** column 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 Creating a Replication Policy

A replication policy allows CBR to periodically replicate backups that have not been replicated or failed to be replicated to the destination region.

When both a backup time and a replication time are configured, ensure that replication starts after backup is complete. Or, replication may fail.

Constraints

- You can set replication policies for server backup vaults.
- A user can create a maximum of 32 replication policies.

Procedure

Step 1 Log in to the management console. Click  in the upper left corner and select a region.

Step 2 Click  and choose **Storage** > **Cloud Backup and Recovery**.

Step 3 Choose **Policies** in the navigation pane and click the **Replication Policies** tab. In the upper right corner, click **Create Policy**. Select **Replication Policy** for **Type**.

Step 4 Set the replication policy parameters. **Table 5-2** describes the parameters.

Table 5-2 Replication policy parameters

Parameter	Description	Example Value
Type	Select a policy type. In this section, we select the replication policy.	Replication policy

Parameter	Description	Example Value
Name	Replication policy name You can enter a custom name or use the default name **policy_xxxx**. A name can contain 1 to 64 characters. Only letters, digits, underscores (_), and hyphens (-) are allowed.	replication_policy
Status	Whether to enable the replication policy By default, this function is enabled. CBR replicates backups and deletes expired replicas only after a replication policy is enabled and applied to vaults.	Enabled
Replication Frequency	Select a replication frequency. By default, the replication task is executed automatically every Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, and Sunday. • **Weekly** Specifies on which days of each week the replication task will be executed. You can select multiple days. • **Day-based** Specifies the interval (every 1 to 30 days) for executing the replication task. If you select day-based replication, the first replication is supposed to be executed on the day when the replication policy is created. If the replication policy is created after the replication time, the first replication will be executed in the next replication period.	Every day
Execution Time	Execution times of the replication policy in a day The default replication time is 22:00. Replication tasks can be scheduled at the beginning of each hour, and you can select multiple hours. It is recommended that replication be performed during off-peak hours or when no services are running.	00:00, 02:00

Parameter	Description	Example Value
Destination Region	Select the region to which the backup data is to be replicated. Only the regions that support replication will be displayed. • If the selected region contains only one project, you can directly select the region name. • If the selected region has multiple projects, the default project of the region is preselected. You can select another project if needed.	-
Retention Rule	Rule that specifies how backup replicas will be retained in the destination region By default, backup replicas are retained for one month. • **Backup quantity** You can set the maximum number of backup replicas to retain for each resource. The value ranges from 2 to 99999. • **Time period** You can choose to retain backup replicas for one month, three months, six months, one year, or for any desired number (2 to 99999) of days. • **Permanent** **NOTE** – The system automatically deletes the earliest and expired backup replicas every other day to avoid exceeding the maximum number of backup replicas to retain or retaining any backup replica longer than the maximum retention period. – There will be delays for CBR to delete expired backup replicas, but normally these delays will not be over 24 hours. – The retention rules apply only to auto-generated backup replicas, but not manual ones. Manual backup replicas need to be deleted manually. – After a backup replica is used to create an image, the replica will not be deleted by the retention rule.	6 months

Step 5 Click **Create Now**. After the replication policy is created, you can view it in the replication policy list.

Step 6 Locate the desired vault and choose **More** > **Apply Replication Policy** to apply the replication policy to the vault. Then you can view the applied policy on the vault details page.

After the policy is applied, backups will be periodically replicated to the destination vault based on the policy.

----End

Example

A user applies a replication policy to a vault in a given region at 11:00 a.m. on Thursday. According to this policy, backups will be replicated to the destination region at 02:00 a.m. every day, and only two backup replicas will be retained. According to this vault's backup policy, two backups are automatically generated at 00:00 every day. At 12:00 p.m. on Saturday, the replication vault will contain two backup replicas, which are replicated on Saturday. Backup replicas generated at 02:00 a.m. on Friday have been automatically deleted according to the replication policy.

5.4 Modifying a Policy

You can modify existing policies if needed.

Prerequisites

At least one policy has been created.

Procedure

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

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

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

Related parameters are described in **Table 5-1** and **Table 5-2**.

Step 4 Click **OK**. You can view the rules column of the new policy in the policy list.

The new retention rule may not apply to existing backups. For details, see **Why Isn't the New Retention Rule Being Applied?**

Alternatively, select **Policies** from the navigation pane and edit the desired policy.

----End

5.5 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 CBR console.
2. In the upper left corner, click and select a region.
3. Click and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 In the navigation pane, choose **Policies**.

Step 3 Click the **Backup Policies** or **Replication 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 or replicas generated based on the policy. You can manually delete unwanted backups or replicas.

Step 4 Confirm the information and click . After the policy is deleted, the system displays a message in the upper right corner. If the policy does not exist in the policy list anymore, it is deleted successfully.

----End

5.6 Applying a Policy to a Vault

You can apply a backup or replication policy to a vault to execute backup or replication tasks at specified times or intervals. The backups or replicas can be used to restore data quickly in the event of data corruption or loss.

Constraints

You can create multiple policies, but can apply only one backup or replication policy to a vault.

Procedure

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 Locate the target vault and choose **More** > **Apply Backup Policy** or **More** > **Apply Replication 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** and **Creating a Replication Policy**.

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

----End

5.7 Removing a Policy from a Vault

Scenarios

If you need to cancel auto backup or replication of a vault, disassociate the policy from the vault, or disable the policy. This section describes how to remove a policy from a vault.

Prerequisites

A policy has been applied to the vault.

Procedure

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

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

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

NOTE

- You can remove a policy from a vault when the vault resources are being backed up. In this case, backup tasks 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 **OK**. After the disassociation, you can go to the **Policy** area on the vault details page to view the results. If the policy does not exist, the disassociation is successful. The vault will no longer execute tasks as specified in this policy.

----End

6 Database Server Backup

6.1 Database Server Backup

Overview

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. CBR cloud server backup uses the consistency snapshot technology for disks to protect data of ECSs. If you back up multiple EVS disks separately, the backup time points of the EVS disks are different. As a result, the backup data of the EVS disks is inconsistent.
- 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.

Figure 6-1 compares these backup types in detail.

CBR supports both crash-consistent backup (also called cloud server backup) and application-consistent backup (also called database server backup).

Crash-consistent backup does not back up data in memory or pending I/O operations and cannot be used to restore applications. If your server is running a MySQL or SAP HANA database, you can use application-consistent backup. An application-consistent backup captures application information both in memory and in pending I/O operations and can be used to quickly restore applications.

Figure 6-1 Backup consistency

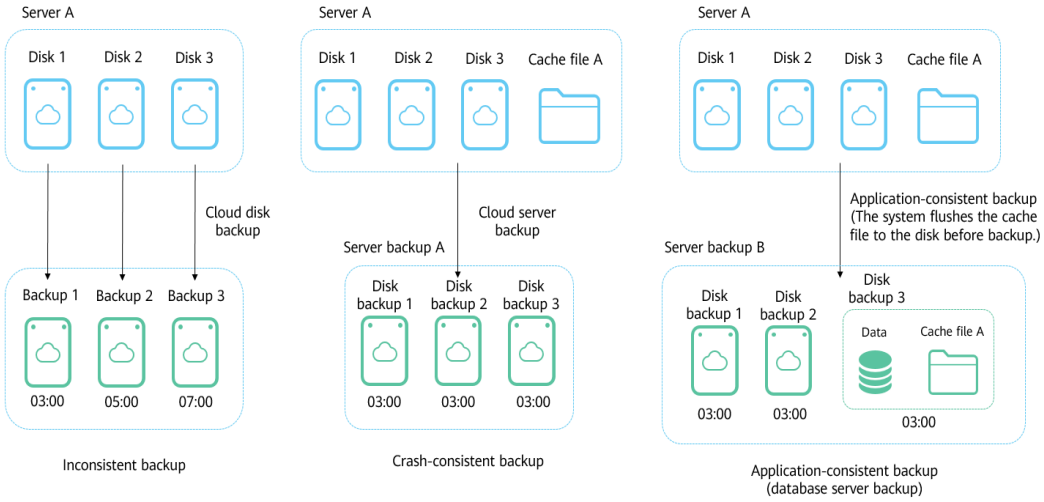

Differences Between Database Server Backup and Cloud Server Backup

Table 6-1 Differences between database server backup and cloud server backup

Item	Database Server Backup	Cloud Server Backup
Object	Cloud servers with MySQL or SAP HANA database deployed	Cloud servers without databases
Granularity	Cloud server	Cloud server
Vault	Server backup vault	Server backup vault
Recommended scenario	Data of cloud servers and their databases such as MySQL or SAP HANA database needs to be backed up. All data and application configurations need to be restored in case of an error.	Only data of cloud servers needs to be backed up. Such data needs to be restored in case of an error. If you use cloud server backup to back up database servers such as MySQL or SAP HANA database servers, certain database configurations may not be fully restored from the backups. This can result in startup issues when the database is restarted.

NOTICE

There are two types of vaults to store server backups. Those store backups of non-database servers are common server backup vaults, and those store backups of database servers are database server backup vaults.

OSs Supporting the Agent

Table 6-2 lists the OSs that can run the Agent.

Table 6-2 OSs supporting the Agent

Database	OS	Supported Versions
SQLServer 2008	Windows	Windows Server 2012, 2012 R2, 2019 for x86_64
SQLServer 2012	Windows	Windows Server 2012, 2012 R2, 2019 for x86_64
SQLServer 2019	Windows	Windows Server 2019 for x86_64
SQLServer 2014/2016/EE	Windows	Windows Server 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 for x86_64
	CentOS	CentOS 6 and 7 for x86_64
	EulerOS	Euler OS 2.2, 2.3 for x86_64
HANA 1.0/2.0	SUSE	SUSE Linux Enterprise Server 12 for x86_64

NOTE

For Windows servers, ensure that Volume Shadow Copy is running properly. Otherwise, the backup will fail.

For databases that are not in the compatibility list, you can create a custom script to back up the database server by referring to Using a Custom Script to Implement Application-Consistent Backup in the *Cloud Backup and Recovery Best Practices*.

Process

Figure 6-2 shows the database server backup process.

Figure 6-2 Database server backup process

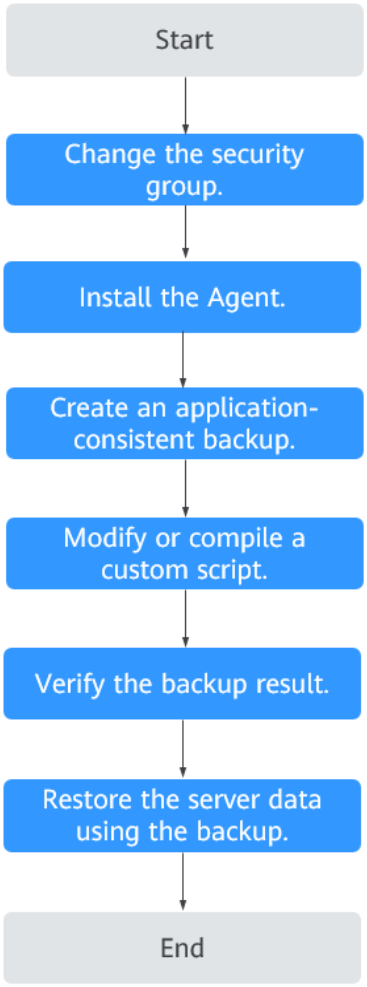

Step 1 Change the security group: Before performing a database server backup task, change the security group of the server you want to back up.

Step 2 Install the Agent: Change the security group and install the Agent in any sequence. Ensure that the two operations are completed before backing up the desired server.

Step 3 Create a database server backup: After creating a server backup vault for storing database server backups, associate it with the desired database server and then create a database server backup.

Step 4 Modify or compile a custom script: After backing up a database server on the CBR console, modify or compile a custom script on the database of the server. For details, see the *Cloud Backup and Recovery Best Practices*.

Step 5 Verify the backup result: After the backup is performed, verify that the backup succeeds. For details, see the *Cloud Backup and Recovery Best Practices*.

Step 6 Use the database server backup to restore server data. After restoration, the database applications and data will return to the exact state they were in at the

time the backup was created. For details, see **Restoring from a Cloud Server Backup**.

----End

6.2 Changing Security Group Rules

Context

A security group is a collection of inbound and outbound rules for ECSs that have the same security protection requirements and are mutually trusted in a VPC. You can create different inbound and outbound rules for the security group to protect the ECSs that are added to this security group. The system creates a security group for each cloud account by default. The default outbound rule allows all outgoing data packets. ECSs in a security group can access each other without the need to add rules. You can also create custom security groups by yourself.

When creating a security group, you must add the inbound and outbound rules and enable the ports required for database server backup to prevent backup failures.

Operation Instructions

Before using database server backup, you need to change security group rules. To ensure network security, CBR has not set any inbound rule, so you need to manually set it.

In the inbound direction, allow traffic from 100.125.0.0/16 over any port from 59526 to 59528. In the outbound direction, allow traffic destined for 100.125.0.0/16 over any port from 1 to 65535. The default outbound rule allows all data packets destined for 0.0.0.0/0 (any IP address). Therefore, you can also use the default outbound rule.

Procedure

Step 1 Log in to the ECS console.

1. Log in to the ECS console.
2. Click  in the upper left corner and select a region.
3. Under **Compute**, click **Elastic Cloud Server**.

Step 2 In the navigation pane, choose **Elastic Cloud Server** or **Bare Metal Server**. On the page displayed, select the target server. Go to the server details page.

Step 3 On the **Security Groups** tab, locate the target security group, and click **Manage Rule**. On the **Security Groups** tab, locate the target security group, and click **Manage Rule**.

Step 4 On the **Inbound Rules** tab, click **Add Rule**. The **Add Inbound Rule** dialog box is displayed. Select **TCP (Custom ports)** for **Protocol & Port**, enter **59526-59528**, select **IP address** for **Source** and enter **100.125.0.0/16**. After supplementing the description, click **OK** to finish setting the inbound rule. You can view the added inbound rule on the **Inbound Rules** tab of the security group.

Step 5 On the **Outbound Rules** tab, click **Add Rule**. The **Add Outbound Rule** dialog box is displayed. Select **TCP (Custom ports)** for **Protocol & Port**, enter **1-65535**, select **IP address** for **Destination** and enter **100.125.0.0/16**. After supplementing the description, click **OK** to finish setting the outbound rule. You can view the added outbound rule on the **Outbound Rules** tab of the security group.

----End

6.3 Installing the Agent

Before enabling database server backup, you also need to install the Agent on your ECSs.

If Agent is not installed on servers, database server backup will fail, but a common server backup can be performed instead. To ensure a successful database server backup, download and install the Agent first.

Operation Instructions

- During the Agent installation, the system requires the **rdadmin** user's permissions to run the installation program. To improve O&M security, change the user **rdadmin**'s password of the Agent OS regularly and disable this user's remote login permission. For details, see **Changing the Password of User rdadmin**.
- **Table 6-3** lists the OSs that can have the Agent installed.

Table 6-3 OSs supporting the Agent

Database	OS	Supported Versions
SQLServer 2008	Windows	Windows Server 2012, 2012 R2, 2019 for x86_64
SQLServer 2012	Windows	Windows Server 2012, 2012 R2, 2019 for x86_64
SQLServer 2019	Windows	Windows Server 2019 for x86_64
SQLServer 2014/2016/ EE	Windows	Windows Server 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 for x86_64
	CentOS	CentOS 6 and 7 for x86_64
	EulerOS	Euler OS 2.2, 2.3 for x86_64

Database	OS	Supported Versions
HANA 1.0/2.0	SUSE	SUSE Linux Enterprise Server 12 for x86_64

NOTE

For Windows servers, ensure that Volume Shadow Copy is running properly. Otherwise, the backup will fail.

- Table 6-4 lists the supported SHA256 values.

Table 6-4 SHA256 values

Package Name	SHA256 Value
Cloud Server Backup Agent-CentOS6-x86_64.tar.gz	f0c59ccb4443bcb6e874bf6e3c574914f9f8b27f3f7379e2d81956a9972802f3
Cloud Server Backup Agent-CentOS7-x86_64.tar.gz	2d3028cb794e1699bae9f65746a60ae99be17d5c4c5e7ebe6b45ff261db9c3c7
Cloud Server Backup Agent-EulerOS2-x86_64.tar.gz	4fb4cf9cb6f5b0e6c13d8ad8bf928754cb95332ee645a97fd0bb3fcbeb53d003
Cloud Server Backup Agent-RedHat6-x86_64.tar.gz	6ae3838fb5644f0f47282c211fe20c6b57a7c5c1d683cd5a1f55860d259b2054
Cloud Server Backup Agent-RedHat7-x86_64.tar.gz	40fa68a808d9da04672678b2773e3345ea6c9dee3c17d598acb66a023cc5cacc
Cloud Server Backup Agent-SuSE11-x86_64.tar.gz	346cc9f1fc0a41a817abb2db61e657a4d615449e13bc46f1c1cfbadc0b281f47
Cloud Server Backup Agent-SuSE12-x86_64.tar.gz	625279b9c9d17ddcc4210b78242efebacdad73f808b86754659d243ece85a400
Cloud Server Backup Agent-WIN64.zip	b7b2067ac89f1fec635d82e3fe2ea794ce6482f9880838f34924b383be44ac4e

NOTICE

During Agent installation, the system automatically opens an available port between 59,526 and 59,528 on the ECS to allow communication. If port 59,526 is occupied, it will try 59,527, then 59,528.

Prerequisites

- The username and password for logging in to the console have been obtained.
- The security group has been configured.
- The **Agent Status** of the ECS is **Not installed**.
- If you use Internet Explorer, you need to add the websites you will use to trusted sites.

Installing the Agent for a Linux OS (Method 1)

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click and select a region.
3. Click and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 In the navigation pane, choose **Cloud Server Backups**.

Step 3 Click the **Agent Installation** tab.

Step 4 In method 1, select the corresponding Agent version as required, and copy the installation command in step 2.

Step 5 On the ECS page, select the target server and click **Remote Login** in the **Operation** column to log in to the ECS.

NOTE

Ensure that the package's SHA256 value is the same as that listed in **Table 6-4**.

For how to obtain the software package, go to method 2. Specifically, click **Download**, and then on the displayed page, select a version based on the target ECS OS and click **OK**.

Step 6 Paste the installation command in step 2 to the ECS and run the command as user **root**.

The system displays a message indicating that the agent is installed successfully. See **Figure 6-3**.

Figure 6-3 Agent installed successfully on the Linux ECS

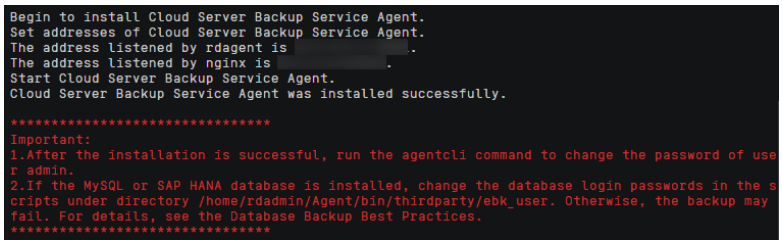

If the execution fails, run the **yum install -y bind-utils** command to install the dig module. If the installation still fails, use method 2 to install the Agent for a Linux OS.

Step 7 After the installation is complete, Agent is running properly. To implement application-consistent backup for MySQL, SAP HANA, or other types of databases, modify or compile a custom script by referring to the *Cloud Backup and Recovery Best Practices*.

----End

Installing the Agent for a Linux OS (Method 2)

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click and select a region.
3. Click and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 Click the **Agent Installation** tab.

Step 3 In method 2, click **Download**. On the displayed download page, select the version to be downloaded based on the OS of the target ECS, and click **OK**.

Step 4 After downloading the Agent to a local directory, check that the package's SHA256 value is the same as that listed in **Table 6-4**.

Step 5 Use a file transfer tool, such as Xftp, SecureFX, or WinSCP, to upload the Agent installation package to your ECS.

Step 6 After the upload, go to the ECS page. Select the target server and click **Remote Login** in the **Operation** column to log in to the ECS.

Step 7 Run the **tar -zxvf** command to decompress the Agent installation package to any directory and run the following command to go to the **bin** directory:

cd bin

Step 8 Run the following command to run the installation script:

sh agent_install_ebk.sh

Step 9 The system displays a message indicating that the client is installed successfully. See **Figure 6-4**.

Figure 6-4 Successful client installation for Linux

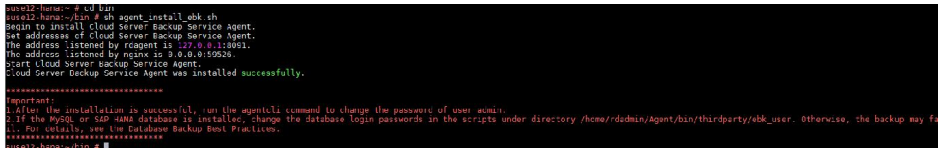

Step 10 If the MySQL or SAP HANA database has been installed on the ECS, run the following command to encrypt the password for logging in to the MySQL or SAP HANA database:

/home/rdadmin/Agent/bin/agentcli encpwd

Step 11 Use the encrypted password in **previous step** to replace the database login password in the script in **/home/rdadmin/Agent/bin/thirdparty/ebk_user/**.

Step 12 After the installation is complete, Agent is running properly. To implement application-consistent backup for MySQL, SAP HANA, or other types of databases, modify or compile a custom script by referring to the *Cloud Backup and Recovery Best Practices*.

----End

Installing the Agent for a Windows OS (Method 1)

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click and select a region.
3. Click and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 Click the **Agent Installation** tab.

Step 3 In method 1, click **Download**. Save the downloaded installation package to a local directory.

Step 4 After downloading the Agent to a local directory, check that the package's SHA256 value is the same as that listed in **Table 6-4**.

Step 5 Use a file transfer tool, such as Xftp, SecureFX, or WinSCP, to upload the Agent installation package to your ECS.

Step 6 Log in to the ECS console and access the ECS as the administrator.

Step 7 Decompress the installation package to any directory and go to the *Installation path***bin** directory.

Step 8 Double-click the **agent_install_ebk.bat** script to start the installation.

Step 9 The system displays a message indicating that the client is installed successfully. See **Figure 6-5**.

Figure 6-5 Successful client installation for Windows

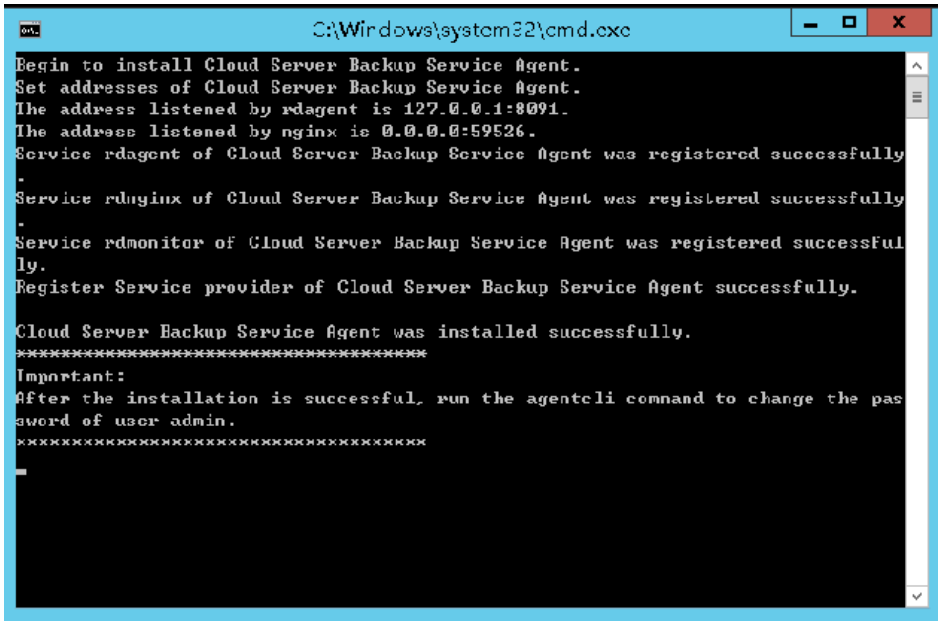

----End

Installing the Agent for a Windows OS (Method 2)

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click and select a region.
3. Click and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 Click the **Agent Installation** tab.

Step 3 On the ECS page, select the target server and click **Remote Login** in the **Operation** column to log in to the ECS as the administrator.

Step 4 Copy the installation commands in step 2 of method 2 to the server and run the command in the Command Prompt.

Step 5 Copy the installation command in step 3 of method 2 to the browser. The following uses ***region1*** as the example region. Then press **Enter** to download the installation package.

https://csbs-agent-*region1*.obs.*region1*.xxxxx/Cloud Server Backup Agent-WIN64.zip

Step 6 After downloading the Agent to a local directory, check that the package's SHA256 value is the same as that listed in **Table 6-4**.

Step 7 Decompress the file to obtain the installation file. Decompress the installation package to any directory and go to the *Installation path***bin** directory.

Step 8 Double-click the **agent_install_ebk.bat** script to start the installation.

Step 9 The system displays a message indicating that the client is installed successfully. See **Figure 6-6**.

Figure 6-6 Successful client installation for Windows

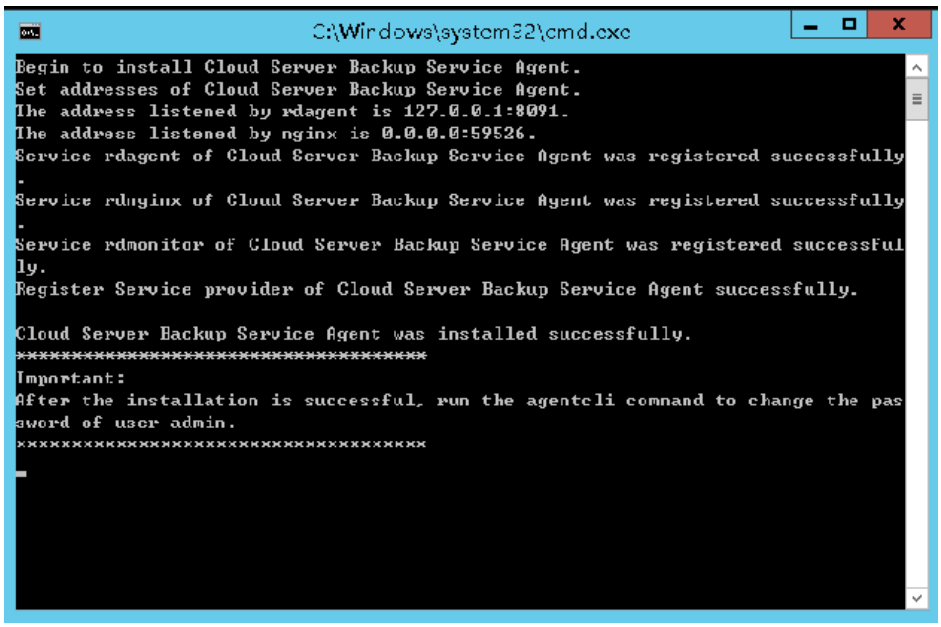

----End

6.4 Creating a 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.

Constraints

- Application-consistent backup is currently not supported for cluster applications, such as, MySQL Cluster. It is supported only for applications on standalone servers.
- To ensure services run normally, perform an application-consistent backup during off-peak hours (for example, at night or on weekends) to minimize the impact on services. If an application-consistent backup is performed during peak hours, the backup may fail or database services may be blocked.

Procedure

Step 1 Log in to the CBR console.

1. Log in to the CBR console.

2. In the upper left corner, click and select a region.
3. Click and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 Create a vault for application-consistent backups by referring to **Creating a Server Backup Vault**. Select **Enable** for **Application-Consistent Backup**.

Step 3 Associate the cloud servers with the created vault. Ensure that the Agent has been installed on the servers.

Step 4 Create a cloud server backup by referring to **Creating a Cloud Server Backup**.

- If an application-consistent backup is created successfully, a blue letter "A" is displayed next to the backup name in the backup list.
- If an application-consistent backup fails to be created, CBR automatically creates a cloud server backup instead and stores the backup in the vault, and a gray letter "A" is displayed next to the backup name in the backup list. You can view the failure cause in the **Management Information** area on the backup details page.

Step 5 Return to the cloud server backup page as prompted. If the backup execution fails, rectify the fault based on the failure details shown on the page.

----End

Follow-up Procedure

If data is lost due to virus attacks or database faults, you can restore the data by following instructions in **Restoring from a Cloud Server Backup** and **Creating an Image from a Cloud Server Backup**.

6.5 Uninstalling the Agent

Scenarios

This section describes how to uninstall the Agent when application-consistent backup is no longer needed.

Prerequisites

The username and password for logging in to an ECS have been obtained.

Uninstalling the Agent for Linux

Step 1 Log in to the ECS and run the **su -root** command to switch to user **root**.

Step 2 In the **home/rdadmin/Agent/bin** directory, run the following command to uninstall the Agent. **Figure 6-7** displays an example. If the word **successfully** in green is displayed, the Agent is uninstalled successfully.

sh agent_uninstall_ebk.sh

Figure 6-7 Agent uninstalled successfully from Linux

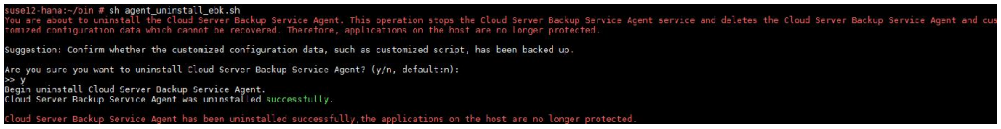

----End

Uninstalling the Agent for Windows

Step 1 Log in to the ECS.

Step 2 In the *Installation path***bin** directory, double-click **agent_uninstall_ebk.bat**. The window for uninstalling the Agent is displayed.

After the uninstallation is complete and successful, the window will be automatically closed.

Figure 6-8 Agent uninstalled successfully from Windows

----End

7 Data Restoration

7.1 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.

NOTE

The server is stopped before a data restoration, and automatically starts up 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.

Constraints

- A data disk backup cannot be restored to the system disk.
- Data cannot be restored to servers in the **Faulty** state.
- A backup replica cannot be used for restoration.
- An ongoing restoration task cannot be terminated.
- Data cannot be restored for a cloud server that is being backed up.

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 CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 Click the **Backups** tab. Locate the desired backup. For details, see **Viewing a Backup**.

Step 3 Click **Restore Server** in the **Operation** column.

WARNING

- The current server data will be overwritten by the data captured at the time of backup. **The restoration cannot be undone.**
- Servers will be shut down during restoration. It is recommended that you perform restoration during off-peak hours.

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

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

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

WARNING

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 the disks are restored, data may become inconsistent and the application may fail to start.

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.
- A data disk backup cannot be restored to the system disk.

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 data is restored to the state when that backup was created.

For details about how to view failed restoration tasks, see **Task Management**.

NOTE

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

----End

Helpful Links

- Due to Windows 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**.
- **Can I Use a System Disk Backup to Recover an ECS?**
- **Can a Server Be Restored Using Its Backups After It Is Changed?**
- **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?**
- **What Changes Will Be Made to the Original Backup When I Use the Backup to Restore a Server?**

7.2 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.

CBR cross-region replication allows you to replicate backups to destination regions and then create images. You can use the images to provision ECSs.

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.

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 the 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.

Constraints

- NOTE
  - Once 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 must contain the system disk data.

Procedure

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

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

7.3 Restoring from a Cloud Disk Backup

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

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.

Prerequisites

- The disk to be restored is **Available**.
- The server is stopped and the disk is detached 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 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.

Procedure

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click  and select a region.

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

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**.

WARNING

- The EVS disk data will be overwritten by the data captured at the time of backup. **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 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 a Non-Shared Disk" in the *Elastic Volume Service User Guide*.

----End

7.4 Creating a Disk from a Cloud Disk Backup

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

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 CBR console.
2. In the upper left corner, click 📍 and select a region.
3. Click ☰ and choose **Storage** > **Cloud Backup and Recovery**.

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**. The **Create Disk** page is displayed.

Step 4 Configure the disk parameters.

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

7.5 Creating a File System from an SFS Turbo Backup

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 CBR console.
2. In the upper left corner, click  and select a region.
3. Click  and choose **Storage** > **Cloud Backup and Recovery**.

Step 2 Click the **Backups** tab and locate the desired backup. For details, see **Viewing a Backup**.

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

Step 4 Configure the file system parameters.

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 Turbo 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 very fast. After the file system status changes 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

8 (Optional) Resource Migration from CSBS/VBS

Context

The cloud platform has launched the next-generation backup service, CBR. If you have backups in CSBS or VBS but want to switch to CBR to manage these historical backups, you can migrate them to CBR in a few clicks.

If you have never used CSBS or VBS, or do not need the historical backups anymore, skip this section.

Migration Rules

During migration, CBR will automatically create vaults based on the types of your historical resources.

Table 8-1 Migration rules

Before Migration	After Migration
Servers or disks are associated with a backup policy.	If backups have been generated, CBR will create a vault with the same name (up to 64 characters) as the policy name (regardless of whether the policy is enabled) and apply the policy to the vault after the vault is created. If no backup is generated, CBR will create a vault only when the policy is enabled. The policy applying rule and vault naming rule are the same as above.
Servers or disks are associated with a backup or replication policy.	If no backup is generated and the policy is disabled, only the policy will be migrated.
Backup or replication policies are not associated with any resource.	The policies will be migrated.

Before Migration	After Migration
Application-consistent backup is enabled.	CBR will create a database server backup vault and name the vault with the policy name.
Backup replicas are generated.	CBR will create a replication vault named **default** to store generated backup replicas.
An image is created using a backup and a tag is added to the image.	The backup will fail to be migrated. Go to the IMS console, delete the tag and then migrate the backup again. After the backup is migrated, add the tag if needed.

Other backups, including manual backups, will be stored in a server backup vault named **default**. Different vaults will be created based on different types of resources. For example, CBR will create a disk backup vault to store the migrated disk backups.

After the migration, backups created using CBR will also be displayed on the VBS console, but you will be billed only once.

NOTE

To delete backups from the VBS console, find these backups in CBR and delete them. Then, the backups will also be deleted from the VBS console.

Based on the preceding rules, the capacity of each vault created by the system is predefined as 1.2 times the total backup size.

For example, you have a 100-GB ECS and a 50-GB ECS. The used storage capacity of the two ECSs is 20 GB and 10 GB, respectively. You have manually backed up the two ECSs using cloud server backup. During migration, the capacity of the vault automatically created will be 1.2 times the total backup size. In this example, the total backup size multiplied by 1.2 is 36 GB. So CBR will automatically create a 36-GB vault.

Constraints

- The vaults you have created cannot be used for migration. Resources will be automatically migrated to system-created vaults.
- Backup resources of one account only need to be migrated once.
- During the migration, do not delete the CBR backup vault that is automatically created. Use it after the migration is complete.
- After resources are migrated, disk backups and server backups will be automatically stored in CBR vaults. No further operations are required.

Procedure

Step 1 Log in to the CBR console.

1. Log in to the CBR console.
2. In the upper left corner, click  and select a region.

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

Step 2 Click **Migrate to CBR** in the upper right corner. Read the content in the displayed dialog box and click **OK**.

Step 3 The system will automatically migrate resources. After the migration, a vault named **default** will be created and a message will be displayed in the upper part of the page indicating that the migration is successful.

----End

FAQ

1. Why Are CBR Backups Displayed on the VBS Console?

   If you have migrated data from CSBS or VBS to CBR and created a backup on the CBR console, the same backup record will be generated on the VBS console. This is due to an underlying mechanism. The VBS console displays all backups generated by CBR, CSBS, and VBS.

2. How Do I Delete Backups from the VBS Console?

   After you migrated data from CSBS or VBS to CBR, backups displayed in the VBS console cannot be deleted alone. Find these backups in CBR and delete them. Then, the backups will also be deleted from the VBS console.

3. What Are the Differences Between CBR, CSBS, and VBS?

   CBR integrates CSBS and VBS. In addition, CBR supports SFS Turbo backup.

4. What Can I Do If a Resource Has Been Associated with CSBS or VBS?

   Choose **Cloud Server Backup Service** or **Volume Backup Service** from the service list. On the corresponding service console, check whether there are resources associated with policies on the **Policies** tab. If so, disassociate the resources from the policy and go to the CBR console to associate the resources with a vault.

9 Task Management

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

Prerequisites

There is at least one task.

Viewing a Task

Step 1 Log in to the CBR console.

1. Log in to the CBR 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, or 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

10 Cloud Eye Monitoring

10.1 Viewing CBR Monitoring Data

Scenarios

This section describes the namespaces, metrics, and dimensions of CBR metrics reported to Cloud Eye. You can view the metrics and alarms generated by CBR on the Cloud Eye console or by calling APIs.

Namespace

SYS.CBR

Metrics

Table 10-1 CBR metrics

Metric ID	Metric Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
used_vault_size	Used Vault Size	Used capacity of the vault	≥ 0	GB	1024 (IEC)	Vault	15 min
vault_util	Vault Usage	Capacity usage of the vault	0~100	%	N/A	Vault	15 min

Dimensions

Key	Value
instance_id	Vault name or ID. You can obtain the vault name and ID by *Querying Vaults*.

Viewing Monitoring Statistics

Step 1 Log in to the management console.

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

- 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.
- 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

11 Recording CBR Operations Using CTS

You can use Cloud Trace Service (CTS) to trace operations in CBR.

Prerequisites

CTS has been enabled.

Key Operations Recorded by CTS

Table 11-1 CBR operations that can be recorded by CTS

Operation	Resource Type	Trace Name
Creating a policy	policy	createPolicy
Updating a policy	policy	updatePolicy
Deleting a policy	policy	deletePolicy
Associating a policy with a vault	vault	associatePolicy
Disassociating a policy from a vault	vault	dissociatePolicy
Creating a vault	vault	createVault
Modifying a vault	vault	updateVault
Deleting a vault	vault	deleteVault
Removing resources	vault	removeResources
Adding resources	vault	addResources
Performing a backup	vault	createVaultBackup
Creating a backup	backup	createBackup
Deleting a backup	backup	deleteBackup
Restoring a backup	backup	restoreBackup

Viewing Audit Logs

For how to view audit logs, see section "Querying Real-Time Traces" in the *Cloud Trace Service User Guide*.

Disabling or Enabling a Tracker

The following procedure illustrates how to disable an existing tracker on the CTS console. After the tracker is disabled, CTS will stop recording operations, but you can still view existing operation records.

Step 1 Log in to the CTS console.

Step 2 In the upper left corner, click  and select a region.

Step 3 Click **Service List** and choose **Management & Deployment** > **Cloud Trace Service**.

Step 4 Choose **Tracker List** in the navigation pane.

Step 5 In the tracker list, click **Disable** in the **Operation** column.

Step 6 Click **OK**.

Step 7 After the tracker is disabled, the available operation changes from **Disable** to **Enable**. To enable the tracker again, click **Enable** and then click **OK**. CTS will start recording operations again.

----End

12 Quotas

What Is Quota?

Quotas can limit the number or amount of resources available to users, such as the maximum number of 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. In the upper right corner of the page, click  .

   The **Quotas** page is displayed.
3. 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?

The system does not support online quota adjustment.

If you need to adjust a quota, contact the operations administrator.

13 FAQs

13.1 Concepts

13.1.1 What Are Full Backup and Incremental Backup?

Definitions

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

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

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

Suppose that server **X** has backups **A**, **B**, and **C**. Backup A is a full backup, and backups B and C are incremental backups. Only changed data blocks are backed up in incremental backups and unchanged data blocks are referenced.

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

Figure 13-1 Backup mechanism

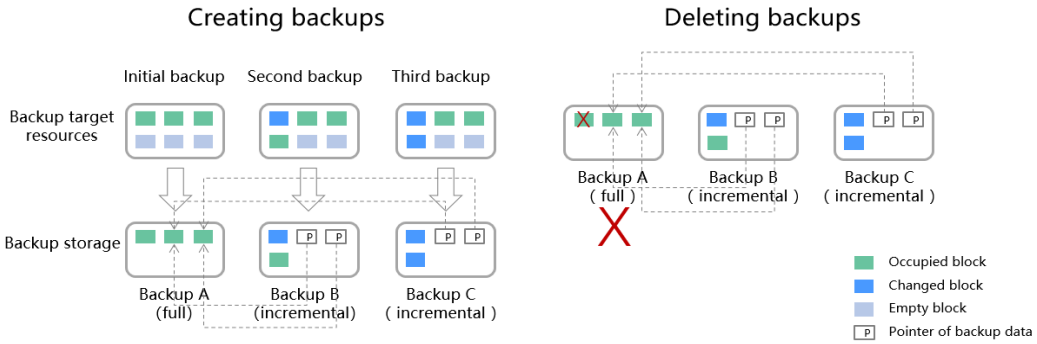

Differences

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

  You are advised to configure periodic full backup (for example, once every 30 days) and daily incremental backup. This approach shortens the backup interval that incremental backups rely on, enhancing backup reliability.

NOTE

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

13.1.2 What Are the Differences Between Backup and Disaster Recovery?

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

Table 13-1 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.

13.1.3 What Are the Differences Between Backups and Snapshots?

Both backups and snapshots provide data redundancy for disks to improve data reliability. **Table 13-2** lists the differences between them.

Table 13-2 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	Snapshots are stored on the physical disks that provide storage resources for EVS disks. Therefore, snapshots do not use the EVS disk space. **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.

13.1.4 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 13-3 lists the differences between them.

Table 13-3 Differences between backups and images

Item	CBR	IMS
Concept	A backup contains the status, configuration, and data of a cloud server or disk stored at a specific time point for recovery in case of a fault. It is used to ensure data security and improve availability.	An image provides all information required for starting a cloud server. It is used to create a cloud server and deploy software environments in batches. A system disk image contains an OS and pre-installed application software for running services. A data disk image contains service data. A full-ECS image contains data of the system disk and data disks.
Usage method	• Data storage location: Unlike server or disk data, backups are stored in OBS. Deleting a disk will not clear its backups. • Operation object: A server or disk can be backed up at a given point in time. CBR supports automatic backup and automatic deletion by configuring backup policies. • Usage: Backups can be used to restore data to the original server or disk, or to create a new disk or full-ECS image. • Support exporting to a local PC: No	• Data storage location: Unlike server or disk data, backups are stored in OBS. If a server or disk that is created using an image is deleted, the image will not be cleared. • Operation object: The system disk and data disks of a server can be used to create private images. You can also create private images using external image files. • Usage: System disk images or full-ECS images can be used to create new servers, and data disk images can be used to create new disks for service migration. • Support exporting to a local PC: Yes However, full-ECS images cannot be exported to a local PC.
Application scenarios	• Data backup and restoration • Rapid service deployment and migration	• Server migration to the cloud or between clouds • Deploying a specific software environment • Deploying software environments in batches • Backing up server operating environments

Item	CBR	IMS
Advantages	Supports automatic backup. Data on a server or disk at a certain time point can be retained periodically or quantitatively.	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.

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

Table 13-4 describes the differences between cloud server backup and cloud disk backup.

Table 13-4 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 disks and data disks) or some 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.

13.2 Backup

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

No. You can back up servers that are in use.

When a server is running, data is written into disks on the server, and some newly generated data is cached in the server memory. During a backup task, data in the memory will not be automatically written into disks, so the disk data and their backups may be inconsistent.

To ensure data integrity, you are advised to perform the backup during off-peak hours when no data is written to the disks.

For applications that require strict consistency, such as databases and email systems, you are advised to enable application-consistent backup.

13.2.2 Can I Back Up a Server Deployed with Databases?

Yes. CBR provides application-consistent backup. For details about the function compatibility, see **Table 13-5**. For applications or databases with which the application-consistent function is incompatible, you are advised to suspend all data write operations before performing backup. If write operations cannot be suspended, you can stop the application systems or the server for offline backup. If you do not perform the preceding operations before backup, status of the server after restoration will be similar to restart upon an unexpected power failure. In this case, log rollback will be performed on databases to keep data consistent.

Table 13-5 OSs supporting the Agent

Database	OS	Supported Versions
SQLServer 2008	Windows	Windows Server 2012, 2012 R2, 2019 for x86_64
SQLServer 2012	Windows	Windows Server 2012, 2012 R2, 2019 for x86_64
SQLServer 2019	Windows	Windows Server 2019 for x86_64
SQLServer 2014/2016/EE	Windows	Windows Server 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 for x86_64
	CentOS	CentOS 6 and 7 for x86_64
	EulerOS	Euler OS 2.2, 2.3 for x86_64
HANA 1.0/2.0	SUSE	SUSE Linux Enterprise Server 12 for x86_64

13.2.3 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

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

No.

The minimum backup granularity that CBR supports is disk.

CBR backs up the status, configuration, and data of a disk at a certain time point for restoration in case of a fault.

13.2.5 Does CBR Support Cross-Region Backup?

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

You can replicate backups in either of the following methods on the CBR console:

- Select a backup from the backup list and manually replicate it.
- Select a backup vault and manually replicate it. Alternatively, you can configure a replication policy to periodically replicate backups that have not been replicated or failed to be replicated to the destination region.

13.2.6 Can I Back Up Data of Two Disks to One Backup?

No.

CBR generates backups for each individual disk.

Data of two disks cannot be backed up to one backup.

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

1. Back up the desired disk.

2. After the backup is successful, locate the backup generated.

3. Create a disk using the backup in the same AZ.

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

Yes.

Resources and backups are not stored together. If resources are deleted, backups will not be deleted at the same time.

So, you can still use the backups to restore resources to a backup point in time.

13.2.9 How Many Backups Can I Create for a Resource?

This number is not limited.

Theoretically, you can create as many backups for a resource as needed.

By default, full backup at the first time and incremental backup subsequently.

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

Yes.

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

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

13.2.11 Can I Stop an Ongoing Backup Task?

No.

An ongoing backup task cannot be stopped.

You need to wait until the backup is complete and then perform operations on the backup.

13.2.12 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 large files are deleted from a disk, the data remains in the disk. 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 its security by yourself.

13.2.13 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, click the name of the target backup to view its details.

13.2.14 How Do I View My Backup Data?

You can check your backup data in the following ways:

NOTE

You cannot check the data that is being backed up on the CBR console.

Cloud 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. Use the disk backup to create a disk. For details, see **Creating a Disk from a Cloud Disk Backup**.
2. Attach the created disk to the new server. For details, see Section "Attaching a Non-shared Disk" in the *Elastic Volume Service User Guide* or Section "Attaching a Shared Disk" in the *Elastic Volume Service User Guide*.
3. Log in to the server to view the data.

SFS Turbo Backups

1. Use the SFS Turbo backup to create a file system. For details, see **Creating a File System from an SFS Turbo Backup**.
2. Mount the file system to a server.

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

13.2.15 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 usually 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 Isn't the New Retention Rule Being Applied?**

13.3 Capacity

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

Symptoms

- There were files on a server, and the server was backed up. After deleting some files, the server was backed up again. By viewing the used vault capacity, we came to a conclusion that the backup size was not changed or even became bigger.

  For example, a user backed up a server that contains 100 files. The used vault capacity was A. Then the user deleted 10 files and backed up the server again. The used vault capacity changed to B. B may be the same as A or even larger.
- 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. If there are data changes, CBR further checks whether data blocks in the sector are all zero blocks. If so, the data blocks will not be backed up and will not be counted in the new backup capacity. If there are non-zero blocks, they will be backed up and will be counted in the new backup capacity. 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.

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

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

You can expand the vault capacity, reduce the number of retained backups, or migrate vault resources.

- 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. If a backup policy has been applied to the vault, disable the backup policy or remove the policy from the vault. Then, automatic backup will stop, and the vault's used space will not change.

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

Symptoms

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

Possible Causes

The backup mechanism of CBR:

- By default, CBR performs a full backup for a resource for the first time and backs up all used data blocks. 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 13-2 Backup mechanism

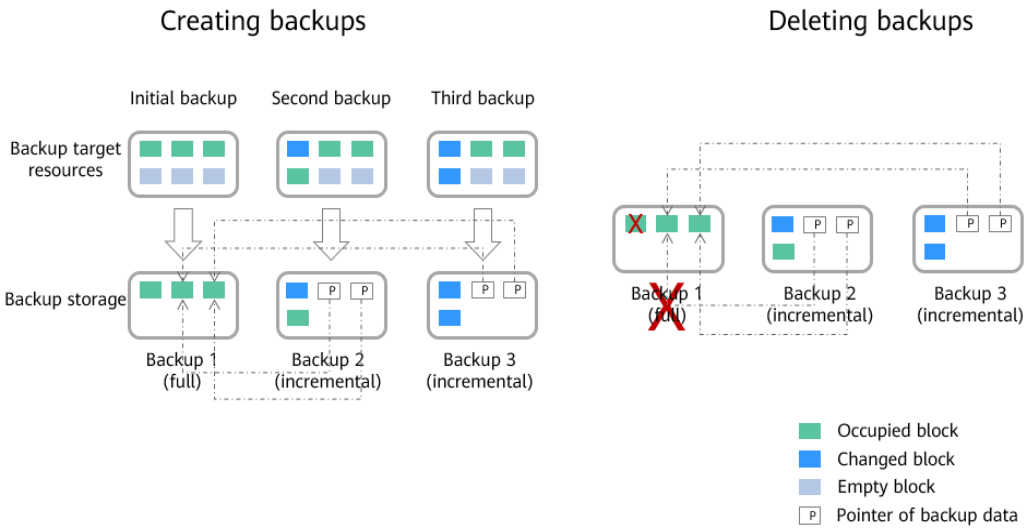

13.3.4 Will Backup Continue If the Usage of a Vault Reaches the Upper Limit?

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

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

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

13.4 Restoration

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

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

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

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.

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

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

You can also use a system disk backup to create an EVS disk.

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

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

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

After the disk data is restored, attach the disk to the server and start the server.

13.4.4 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.

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

Yes. After restoration, the capacity of the expanded disk goes back to the original capacity before expansion.

If you want to use the capacity added to the disk, you need to attach the restored disk to a server, log in to the server, and then manually modify the file system configuration.

For details, see section "Disk Capacity Expansion" in the *Elastic Volume Service User Guide*.

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

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

If you have installed the password reset plug-in, you can reset the password on the management console.

You can reset the password online or offline.

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

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

13.4.8 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, take the following steps:

  NOTE

  Data consistency is not guaranteed using method 2.

  a. Back up the disks on the original server.

     Back up each disk of the original server. For details, see **Creating a Cloud Disk Backup**.

  b. Create new disks from the backups.

     Use the backup of each disk to create a disk. For details, see **Creating a Disk from a Cloud Disk Backup**.

  c. Attach the created disk to the new server. For details, see Section "Attaching a Non-shared Disk" in the *Elastic Volume Service User Guide* or Section "Attaching a Shared Disk" in the *Elastic Volume Service User Guide*.

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

Use a backup to create an EVS disk and associate the EVS disk with a server.

For details, see section "Attaching a Non-Shared Disk" in the *Elastic Volume Service User Guide* or section "Attaching a Shared Disk" in the *Elastic Volume Service User Guide*.

Then copy data in the data disk to the system disk.

13.4.10 Can I Stop an Ongoing Restoration Task?

No. An ongoing restoration task cannot be stopped.

The backup data will overwrite the current data, and the restoration cannot be undone.

13.5 Policies

13.5.1 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.

13.5.2 Why Isn't the New Retention Rule Being 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 performs manual backups or waits 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 performs a manual backup or waits for the system to automatically create one at 2:00 a.m. on Friday, the system will apply the latest retention policy after the backup task completes. 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.

13.5.3 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 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.

13.5.4 How Do I Retain My Backups Permanently?

You can retain your backups permanently through manual or automatic backup (backup policies).

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.

13.5.5 How Can I Cancel Auto Backup or Auto Replication?

To cancel automatic backup or replication, remove the policy from the vault. For details, see section "Removing a Policy from a Vault" in the *Cloud Backup and Recovery User Guide*.

You can also disable the policy.

If you need to enable the policy again, apply the policy again or enable the policy.

13.5.6 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.

13.5.7 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 Isn't the New Retention Rule Being Applied?**
3. The backups are created manually. The policy's retention rule does not apply to manual backups. They can only be deleted manually.

13.6 Optimization

13.6.1 What Are Common Problems During Cloud-Init Installation?

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

For details about how to install and configure Cloud-Init, see 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 13-3 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 13-4 Checking Cloud-Init status

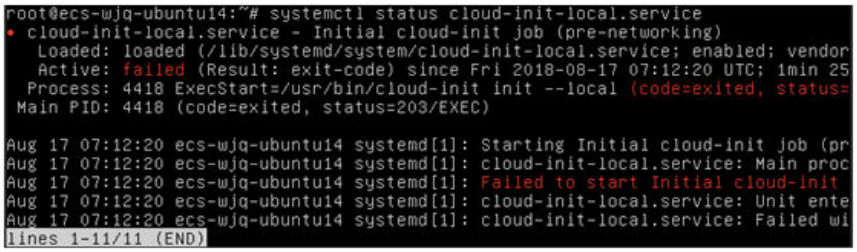

Figure 13-5 Checking Cloud-Init status

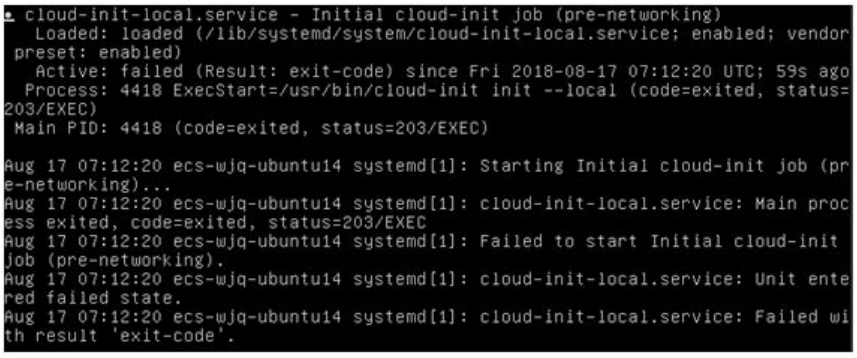

d. Run the **cp /usr/local/cloud-init /usr/bin/** command to copy the **cloud-init** file to the **usr/bin** directory, and then run the following command to restart Cloud-Init:

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

Figure 13-6 Restarting Cloud-Init

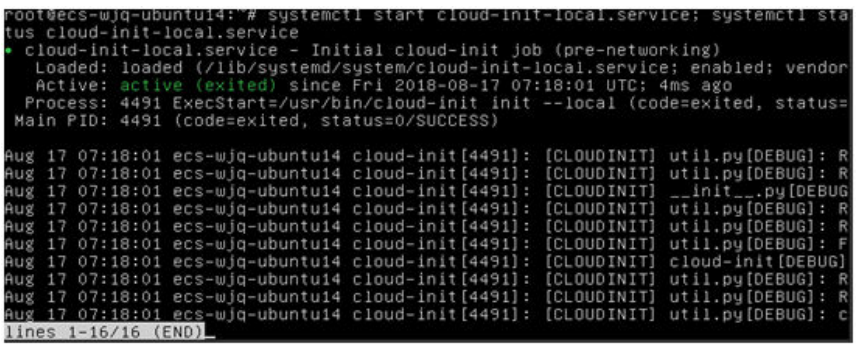

e. Run the following commands to check the Cloud-Init status:

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

Ubuntu14.04: chkconfig and systemctl Not Installed

- Symptom

  chkconfig is not installed.

- Solution

  Run the following commands to install chkconfig on Ubuntu14.04:

  # **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 13-7 Information returned when Cloud-Init automatic start is successfully set

   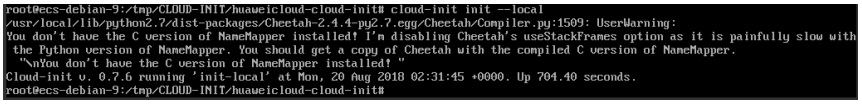

   Cause analysis: The compilation fails because the GNU compiler collection (GCC) is not installed.

   Solution

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

   yum -y install gcc

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

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

Information similar to the following is displayed:

Figure 13-8 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 13-9 Checking the Cloud-Init status

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

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

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

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

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

CentOS 7/Fedora 28: Required C Compiler Not Installed

- Symptom

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

  cloud-init init --local

  The following information is displayed:

```
/usr/lib/python2.5/site-packages/Cheetah/Compiler.py:1532: UserWarning:
You don't have the C version of NameMapper installed! I'm disabling Cheetah's useStackFrames
```



```
option as it is painfully slow with the Python version of NameMapper. You should get a copy of
Cheetah with the compiled C version of NameMapper.
 "\nYou don't have the C version of NameMapper installed!
```

- Possible Cause

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

- Solution

  Run the following command to install GCC:

  yum -y install gcc

  Reinstall Cloud-Init.

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

- Symptom

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

  Figure 13-10 NIC not started

  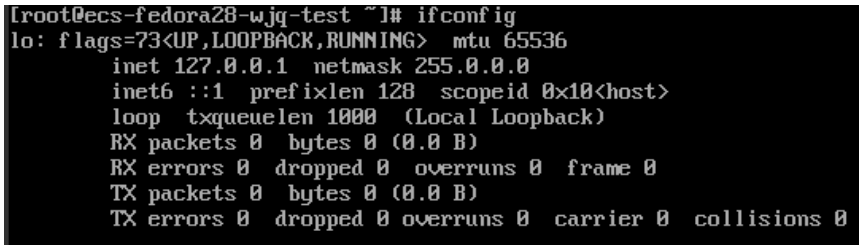

- Solution

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

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

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

Solution

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

13.6.3 What Can Cloud-Init Do?

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

Installation Methods

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.

13.7 Others

13.7.1 Is There a Quota for CBR Vaults?

There are no quota limits for other vaults.

You can create as many vaults as needed.

There is no limit on the number of backups that can be created for a single resource. You can create multiple backups for a resource.

13.7.2 Can I Merge My Vaults?

No. Vaults cannot be merged.

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

You can create an ECS using an image.

Then create a new image from the ECS. Delete the original image and then delete the backup.

For details, see section "Creating an Image from a Cloud Server Backup" in the *Cloud Backup and Recovery User Guide*.

13.7.4 Can I Export Disk Backup Data to Another Server?

Yes.

Use a backup to create an EVS disk, associate the EVS disk with a server, and export data to that server.

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*.

13.7.5 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.

13.7.6 Can I Download Backup Data to a Local PC?

No.

CBR backup data cannot be downloaded to a local PC.

You can view backup data and use backups to restore data on the cloud at any time.

13.7.7 How Do I Copy Disk Data to Another Account?

If the two accounts are in the same region, you can use CBR backup sharing to copy disk data between accounts.

If two accounts are not in the same region, data sharing between accounts is not allowed.

14 Troubleshooting Cases

14.1 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 14-1 Troubleshooting

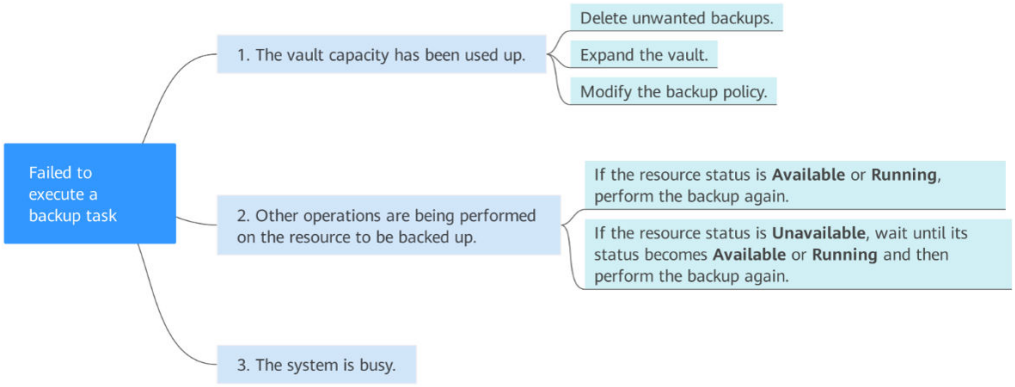

Table 14-1 Troubleshooting

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 during 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 will stop, and the vault's used capacity 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**.

14.2 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 14-2 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.

14.3 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 **xfs_admin -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

14.4 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 14-2**.

Figure 14-2 Data disks in the offline state

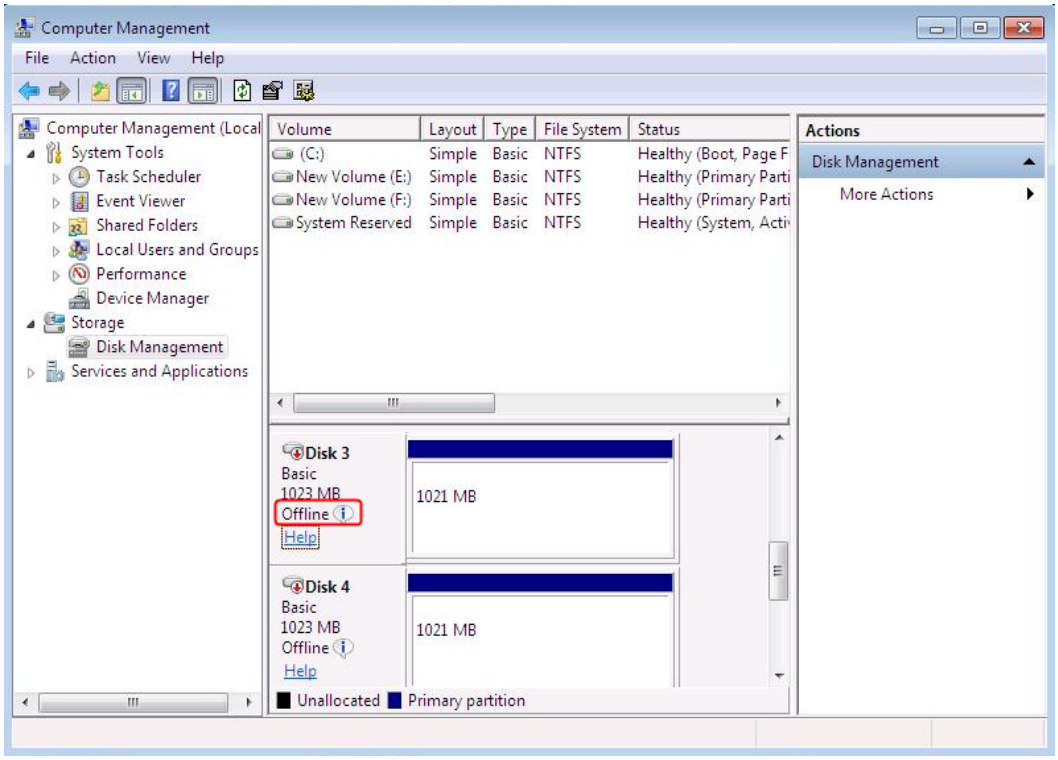

Step 4 Right-click a data disk in the offline state and choose **Online**, as shown in **Figure 14-3**.

Figure 14-3 Setting a data disk to be online

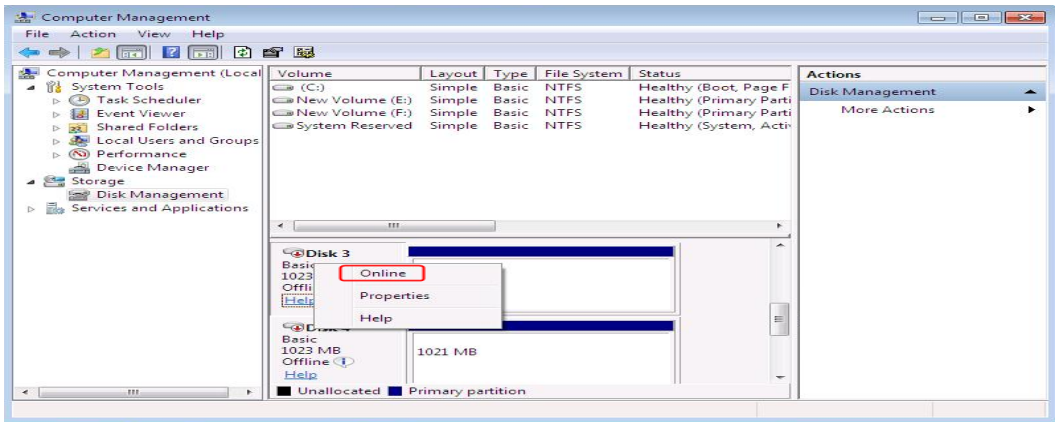

After the data disk status changes to **Online**, the data disk will be displayed in the disk list, as shown in **Figure 14-4**.

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

Figure 14-4 Viewing online data disks

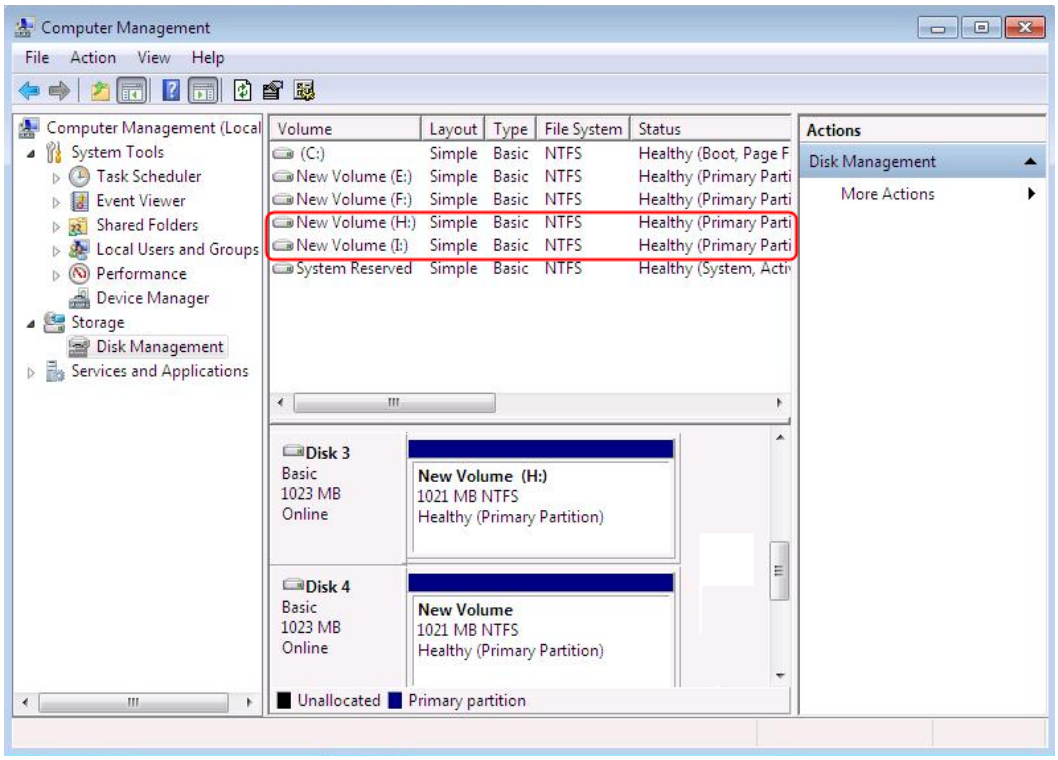

----End

14.5 Failed to Download or Install the Agent Required by Application-Consistent

Symptom

The system displays a message indicating that the script cannot be downloaded or the Agent fails to be installed in Linux mode 2.

Possible Causes

- The DNS cannot resolve the OBS domain name.
- The OpenSSL version installed on the target server is outdated.

Solution (When the DNS Cannot Resolve the OBS Domain Name)

Cause: The DNS cannot resolve the domain name.

You need to manually change the DNS server address. Obtain the IP address from technical support. If the problem persists, try later or use the Linux mode 1 to install it.

Procedure (Linux)

Step 1 Log in to the server as the **root** user.

Step 2 Run the **vi /etc/resolv.conf** command to edit the **/etc/resolv.conf** file. Add the DNS server IP address above the existing name server information, as shown in Figure 14-5.

Figure 14-5 Configuring DNS

The format is as follows:

```
nameserver DNS server IP address
```

Step 3 Press **Esc**, input **:wq**, and press **Enter** to save the changes and exit the vi editor.

Step 4 Run the following command to check whether the IP address is added. If yes, the operation is complete.

cat /etc/resolv.conf

----End

Procedure (Windows)

Step 1 Go to the ECS console and log in to the ECS running Windows Server 2012.

Step 2 Click **This PC** in the lower left corner.

Step 3 On the page that is displayed, right-click **Network** and choose **Properties** from the drop-down list. The **Network and Sharing Center** page is displayed, as shown in Figure 14-6. Click **Local Area Connection**.

Figure 14-6 Page for network and sharing center

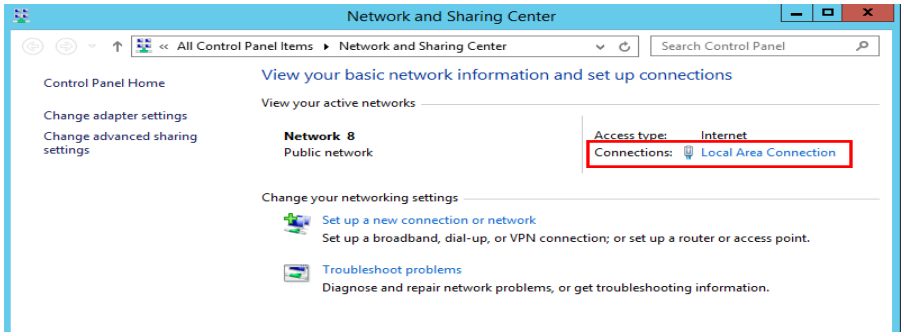

Step 4 In the **Activity** area, select **Properties**. See Figure 14-7.

Figure 14-7 Local area connection

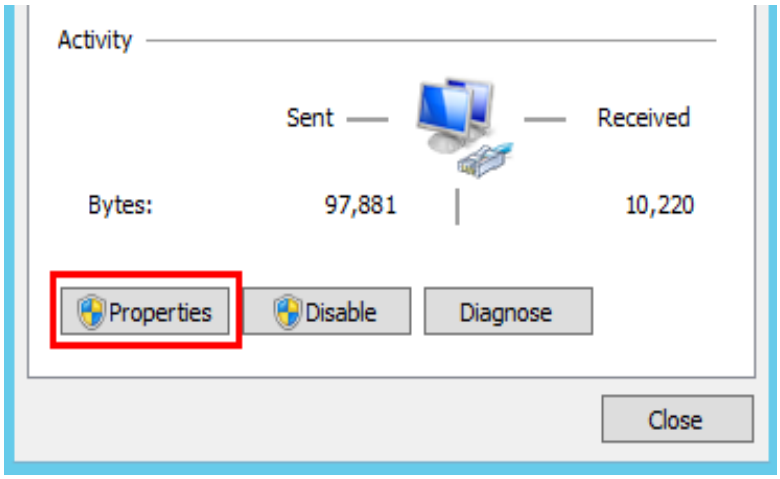

Step 5 In the **Local Area Connection Properties** dialog box that is displayed, select **Internet Protocol Version 4 (TCP/IPv4)** and click **Properties**. See **Figure 14-8**.

Figure 14-8 Local area connection properties

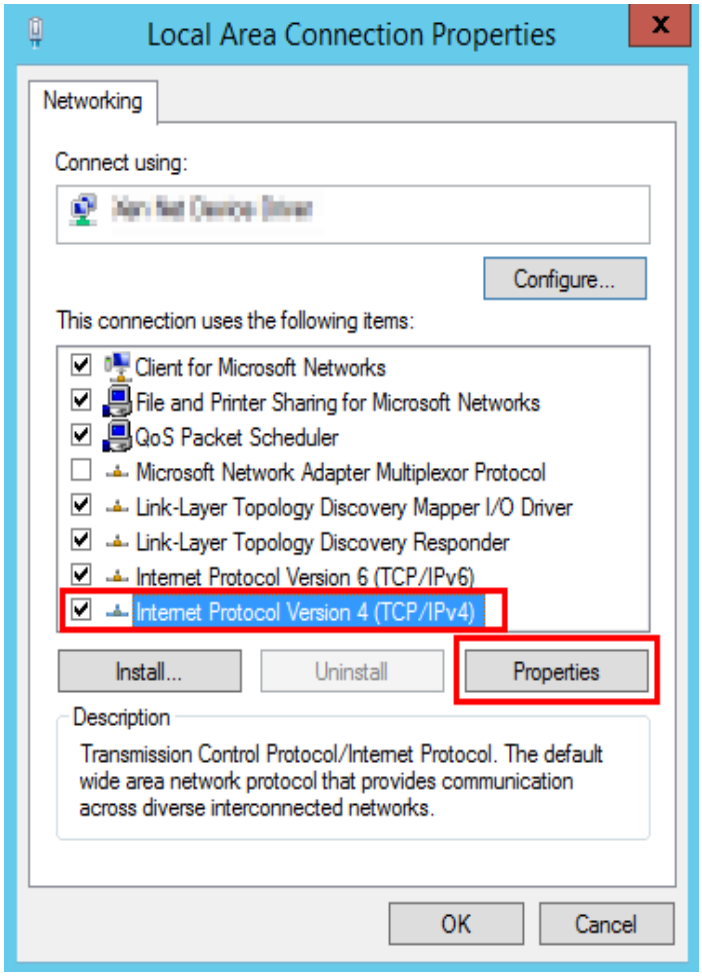

Step 6 In the dialog box that is displayed, select **Use the following DNS server addresses:** and configure DNS, as shown in **Figure 14-9**. You need to manually

change the DNS server address. Obtain the IP address from technical support. Then click **OK**.

Figure 14-9 Configuring DNS

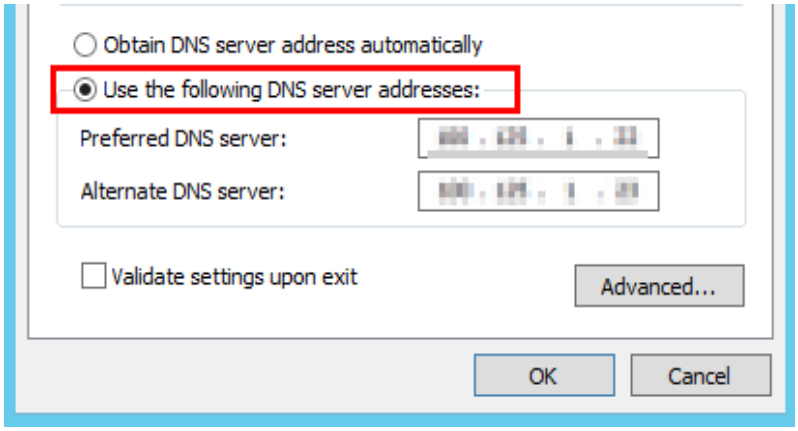

----End

Solution for Cause 2

Cause 2: The OpenSSL version of the target server is too early.

Step 1 Use a remote management tool (such as PuTTY or Xshell) to connect to your ECS through the elastic IP address.

Step 2 Select the Agent version based on your needs, copy the command of installation mode 2 to the server, and change **https** to **http** in wget. Run the command as the **root** user.

----End

14.6 A Server Created Using an Image Enters Maintenance Mode After Login

Symptom

A server is created using the image of a cloud server backup. However, upon login to the server, the server enters maintenance mode and cannot be used.

Possible Cause

After the server is created, the **/etc/fstab** file on the system disk of the new server retains configuration parameters from the source server. As a result, the UUIDs referenced in the file do not match the new data disks, causing the ECS instance to encounter an error when loading **/etc/fstab** and enter maintenance mode.

Solution

The following uses CentOS as an example.

Step 1 After creating an ECS using an image, log in to the ECS console, click **Remote Login** in the row of the ECS.

Step 2 On the maintenance mode page that is displayed, access the system as prompted.

Figure 14-10 Maintenance mode of the system

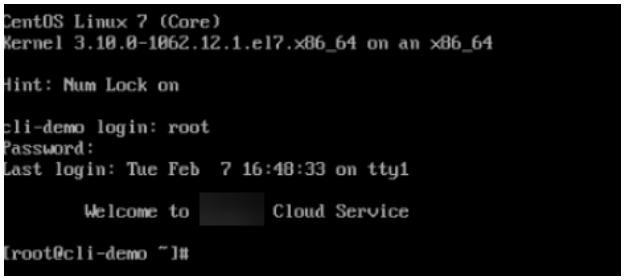

Step 3 Run the **cat /etc/fstab** command to check the disk attachment information.

Figure 14-11 Data disk UUIDs

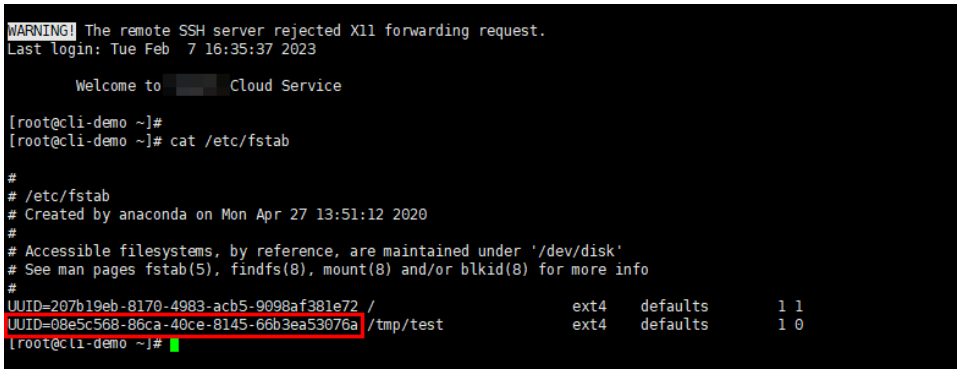

Step 4 Run the **vi /etc/fstab** command to open the file, press **i** to enter the editing mode, and delete the attachment information of all data disks. Then, press **Esc** to exit the editing mode and run **:wq!** to save the change and exit.

Figure 14-12 /etc/fstab after being updated

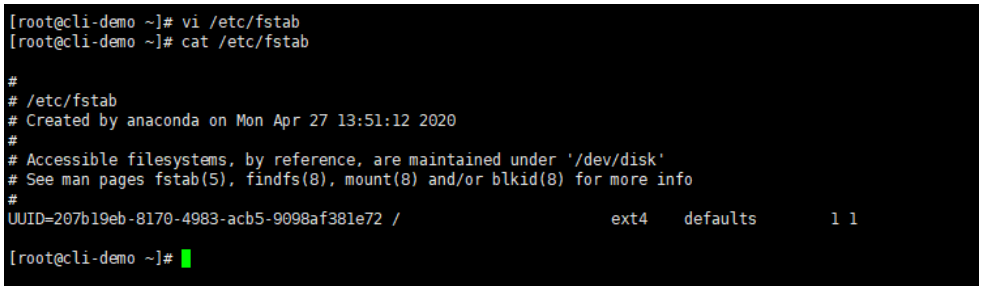

Step 5 Run the **reboot** command to restart the system.

Figure 14-13 Normal bootup page

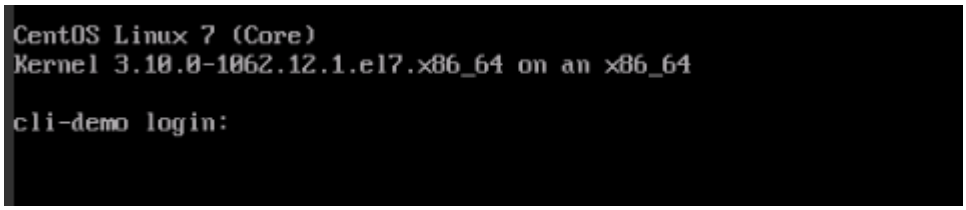

Step 6 After entering the system, attach the data disks manually.

Figure 14-14 Attaching the data disks manually

Step 7 Run the **blkid** command to obtain the UUID information of the data disks.

Figure 14-15 Obtaining UUIDs of data disks

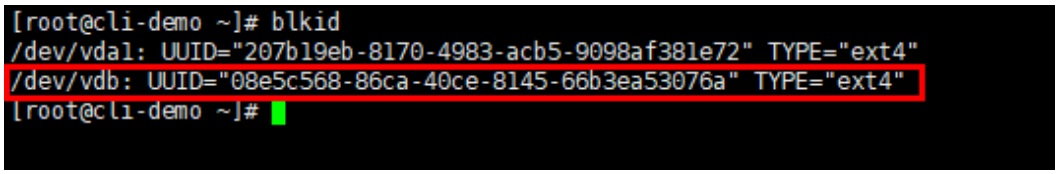

Step 8 Run the **vi /etc/fstab** command to open the file, press **i** to enter the editing mode, and add the attachment information of all data disks. Then, press **Esc** to exit the editing mode and run **:wq!** to save the change and exit.

Figure 14-16 Adding attachment information of data disks

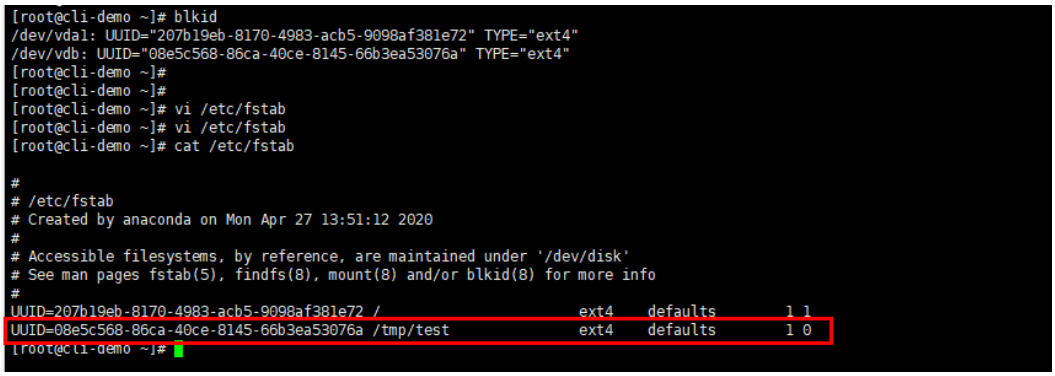

After the information is added, the system will automatically attach the data disks on restart.

----End

A Appendix

A.1 Agent Security Maintenance

A.1.1 Changing the Password of User rdadmin

Scenarios

- To improve O&M security, you are advised to change the user **rdadmin**'s password of the client OS regularly and disable this user's remote login permission.
- In Linux, user **rdadmin** does not have a password.
- This section describes how to change the password of user **rdadmin** in Windows Server 2012. Change the password according to actual situation in other versions.

Prerequisites

- The username and password for logging in to the console have been obtained.
- The username and password for logging in to a Windows ECS have been obtained.

Procedure

Step 1 Go to the ECS console and log in to the Windows ECS.

Step 2 Choose **Start** > **Control Panel**. In the **Control Panel** window, click **User Accounts**.

Step 3 On the displayed **User Account Control** dialog box, select **rdadmin** and click **Reset Password**.

Step 4 Enter the new password and click **OK**.

Step 5 In **Task Manager**, click the **Services** tab and then click **Open Service**.

Step 6 Select RdMonitor and RdNginx respectively. In the displayed dialog box, select **Login**, change the password to the one entered in **Step 4**, and click **OK**.

----End

A.1.2 Changing the Password of the Account for Reporting Alarms (SNMP v3)

To enhance the system O&M security, you are advised to change the password of the account for reporting alarms.

Prerequisites

- The username and password for logging in to the console have been obtained.
- The username and password for logging in to a server have been obtained.

Context

This section introduces the procedures in Windows and Linux.

NOTICE

There may be security risks if you use the same password for SNMP v3 authentication and data encryption. To ensure system security, you are advised to set different passwords for SNMP v3 authentication and data encryption.

Obtain the initial authentication password from technical support.

NOTE

The password must meet the following complexity requirements:

- Contains 8 to 16 characters.
- Contains at least one of the following special characters: `~!@#$%^&*()-_=+\|[{}];:'",<.>/?
- Contains at least two of the following types of characters:
  - Uppercase letters
  - Lowercase letters
  - Numeric characters
- Cannot be the same as the username or the username in reverse order.
- Cannot be the same as the old passwords.
- Cannot contain spaces.

Changing the Password of the Account for Reporting Alarms (SNMP v3) in Windows

Step 1 Log in to the server where the Agent is installed.

Step 2 Open the CLI and go to the *Installation path***bin** directory.

Step 3 Run the **agentcli.exe chgsnmp** command, enter the server login password, and press **Enter**.

```
Please choose operation:
1: Change authentication password
2: Change private password
3: Change authentication protocol
4: Change private protocol
5: Change security name
6: Change security Level
7: Change security model
8: Change context engine ID
9: Change context name
Other: Quit
Please choose:
```

NOTE

admin is the username configured during the Agent installation.

Step 4 Select the SN of the authentication password or data encryption password that you want to change and press **Enter**.

Step 5 Type the old password and press **Enter**.

Step 6 Type a new password and press **Enter**.

Step 7 Type the new password again and press **Enter**.

The password is changed.

----End

Changing the Password of the Account for Reporting Alarms (SNMP v3) in Linux

Step 1 Log in to the Linux server using the server password.

Step 2 Run the **TMOUT=0** command to prevent PuTTY from exiting due to session timeout.

NOTE

After you run the preceding command, the system continues to run even when no operation is performed, which brings security risks. To ensure system security, run the **exit** command to exit the system after you finish the operations.

Step 3 Run the **su - rdadmin** command to switch to user **rdadmin**.

Step 4 Run the **/home/rdadmin/Agent/bin/agentcli chgsnmp** command, enter the server login password, and press **Enter**.

NOTE

The installation path of the Agent is **/home/rdadmin/Agent**.

```
Please choose operation:
1: Change authentication password
2: Change private password
3: Change authentication protocol
4: Change private protocol
5: Change security name
6: Change security Level
7: Change security model
8: Change context engine ID
```



```
9: Change context name
Other: Quit
Please choose:
```

Step 5 Select the SN of the authentication password or data encryption password that you want to change and press **Enter**.

Step 6 Type the old password and press **Enter**.

Step 7 Type a new password and press **Enter**.

Step 8 Type the new password again and press **Enter**.

The password is changed.

----End

A.1.3 Replacing the Server Certificate

For security purposes, you may want to use a Secure Socket Layer (SSL) certificate issued by a third-party certification authority. The Agent allows you to replace authentication certificates and private key files as long as you provide the authentication certificates and private-public key pairs. The update to the certificate can take effect only after the Agent is restarted, hence you are advised to update the certificate during off-peak hours.

Prerequisites

- The username and password for logging in to the console have been obtained.
- The username and password for logging in to a server have been obtained.
- New certificates in the X.509v3 format have been obtained.

Context

- The Agent is pre-deployed with the Agent CA certificate **bcmagentca**, private key file of the CA certificate **server.key** (), and authentication certificate **server.crt**. All these files are saved in **/home/rdadmin/Agent/bin/nginx/conf** (if you use Linux) or **\bin\nginx\conf** (if you use Windows).
- You need to restart the Agent after replacing a certificate to make the certificate effective.

Replacing the Server Certificate in Linux

Step 1 Log in the Linux server with the Agent installed.

Step 2 Run the **TMOUT=0** command to prevent PuTTY from exiting due to session timeout.

NOTE

After you run the preceding command, the system continues to run even when no operation is performed, which brings security risks. To ensure system security, run the **exit** command to exit the system after you finish the operations.

Step 3 Run the **su - rdadmin** command to switch to user **rdadmin**.

Step 4 Run the **cd /home/rdadmin/Agent/bin** command to go to the script path.

NOTE

The installation path of the Agent is **/home/rdadmin/Agent**.

Step 5 Run the **sh agent_stop.sh** command to stop the Agent running.

Step 6 Place the new certificates and private key files in the specified directory.

NOTE

Place new certificates in the **/home/rdadmin/Agent/bin/nginx/conf** directory.

Step 7 Run the **/home/rdadmin/Agent/bin/agentcli chgkey** command.

The following information is displayed:

```
Enter password of admin:
```

NOTE

admin is the username configured during the Agent installation.

Step 8 Type the login password of the Agent and press **Enter**.

The following information is displayed:

```
Change certificate file name:
```

Step 9 Enter a name for the new certificate and press **Enter**.

NOTE

If the private key and the certificate are the same file, names of the private key and the certificate are identical.

The following information is displayed:

```
Change certificate key file name:
```

Step 10 Enter a name for the new private key file and press **Enter**.

The following information is displayed:

```
Enter new password:
Enter the new password again:
```

Step 11 Enter the protection password of the private key file twice. The certificate is then successfully replaced.

Step 12 Run the **sh agent_start.sh** command to start the Agent.

----End

Replacing the Server Certificate in Windows

Step 1 Log in to the Windows server with the Agent installed.

Step 2 Open the CLI and go to the *Installation path***bin** directory.

Step 3 Run the **agent_stop.bat** command to stop the Agent running.

Step 4 Place the new certificates and private key files in the specified directory.

> **NOTE**
>
> Place new certificates in the *installation path***bin\nginx\conf** directory.

Step 5 Run the **agentcli.exe chgkey** command.

The following information is displayed:

```
Enter password of admin:
```

> **NOTE**
>
> **admin** is the username configured during the Agent installation.

Step 6 Enter a name for the new certificate and press **Enter**.

> **NOTE**
>
> If the private key and the certificate are the same file, names of the private key and the certificate are identical.

The following information is displayed:

```
Change certificate key file name:
```

Step 7 Enter a name for the new private key file and press **Enter**.

The following information is displayed:

```
Enter new password:
Enter the new password again:
```

Step 8 Enter the protection password of the private key file twice. The certificate is then successfully replaced.

Step 9 Run the **agent_start.bat** command to start the Agent.

----End

A.1.4 Replacing CA Certificates

Scenarios

A CA certificate is a digital file signed and issued by an authentication authority. It contains the public key, information about the owner of the public key, information about the issuer, validity period, and certain extension information. It is used to set up a secure information transfer channel between the Agent and the server.

If the CA certificate does not comply with the security requirements or has expired, replace it for security purposes.

Prerequisites

- The username and password for logging in to an ECS have been obtained.
- A new CA certificate is ready.

Replacing CA Certificates in Linux

Step 1 Log in the Linux server with the Agent installed.

Step 2 Run the following command to prevent logout due to system timeout:

TMOUT=0

Step 3 Run the following command to switch to user **rdadmin**:

su - rdadmin

Step 4 Run the following command to go to the path to the Agent start/stop script:

cd /home/rdadmin/Agent/bin

Step 5 Run the following command to stop the Agent running:

sh agent_stop.sh

Step 6 Run the following command to go to the path to the CA certificate:

cd /home/rdadmin/Agent/bin/nginx/conf

Step 7 Run the following command to delete the existing CA certificate:

rm bcmagentca.crt

Step 8 Copy the new CA certificate file into the **/home/rdadmin/Agent/bin/nginx/conf** directory and rename the file **bcmagentca.crt**.

Step 9 Run the following command to change the owner of the CA certificate:

chown rdadmin:rdadmin bcmagentca.crt

Step 10 Run the following command to modify the permissions on the CA certificate:

chmod 400 bcmagentca.crt

Step 11 Run the following command to go to the path to the Agent start/stop script:

cd /home/rdadmin/Agent/bin

Step 12 Run the following command to start the Agent:

sh agent_start.sh

----End

Replacing CA Certificates in Windows

Step 1 Log in to the ECS with the Agent installed.

Step 2 Go to the *Installation path***bin** directory.

Step 3 Run the **agent_stop.bat** script to stop the Agent.

Step 4 Go to the *Installation path***nginx\conf** directory.

Step 5 Delete the **bcmagentca.crt** certificate file.

Step 6 Copy the new CA certificate file into the *Installation path***nginx\conf** directory and rename the file **bcmagentca.crt**.

Step 7 Go to the *Installation path***bin** directory.

Step 8 Run the **agent_start.bat** script to start the Agent.

----End

A.2 Change History

Released On	Description
2021-07-22	This issue is the first official release.

