Cloud Backup and Recovery

User Guide (Kuala Lumpur Region)

Issue	01
Date	2022-08-16

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions

Notice

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

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

Contents

1 Service Overview.. 1
1.1 What Is CBR?.. 1
1.2 Application Scenarios.. 5
1.3 Functions.. 5
1.4 Permissions Management.. 8
1.5 Constraints.. 10
1.6 CBR and Other Services.. 11
1.7 Basic Concepts.. 12
1.7.1 CBR Basic Concepts.. 12
1.7.2 Region and AZ.. 13
2 Getting Started.. 15
2.1 Step 1: Create a Vault.. 15
2.1.1 Creating a Server Backup Vault.. 15
2.1.2 Creating a Disk Backup Vault.. 17
2.1.3 Creating an SFS Turbo Backup Vault.. 19
2.2 Step 2: Associate a Resource with the Vault.. 21
2.3 Step 3: Create a Backup.. 23
2.3.1 Creating a Cloud Server Backup.. 23
2.3.2 Creating a Cloud Disk Backup.. 25
2.3.3 Creating an SFS Turbo Backup.. 26
3 Permissions Management.. 28
3.1 Creating a User and Granting CBR Permissions.. 28
3.2 Creating a Custom Policy.. 29
4 Vault Management.. 32
4.1 Querying a Vault.. 32
4.2 Deleting a Vault.. 34
4.3 Dissociating a Resource.. 35
4.4 Expanding Vault Capacity.. 35
4.5 Managing Vault Tags.. 36
4.6 Managing the Enterprise Projects of Vaults.. 37
5 Backup Management.. 39
5.1 Viewing a Backup.. 39

5.2 Sharing a Backup 41
5.3 Deleting a Backup 42
5.4 Using a Backup to Create an Image 43
5.5 Using a Backup to Create a Disk 44
5.6 Using a Backup to Create a File System 44

6 Policy Management 46
6.1 Creating a Backup Policy 46
6.2 Modifying a Policy 51
6.3 Deleting a Policy 51
6.4 Applying a Policy to a Vault 52
6.5 Removing a Policy from a Vault 52

7 Restoring Data 54
7.1 Restoring Data Using a Cloud Server Backup 54
7.2 Restoring Data Using a Cloud Disk Backup 56

8 Application-Consistent Backup 57
8.1 What Is Application-Consistent Backup? 57
8.2 Changing a Security Group 61
8.3 Installing the Agent 62
8.4 Creating an Application-Consistent Backup 66
8.5 Uninstalling the Agent 67

9 Managing Tasks 69

10 Monitoring 70
10.1 CBR Metrics 70

11 Quotas 72

12 FAQs 73
12.1 Concepts 73
12.1.1 What Are Full Backup and Incremental Backup? 73
12.1.2 What Are the Differences Between Backup and Disaster Recovery? 74
12.1.3 What Are the Differences Between Backups and Snapshots? 75
12.1.4 Why Is My Backup Size Larger Than My Disk Size? 76
12.1.5 What Are the Differences Between Cloud Server Backup and Cloud Disk Backup? 77
12.2 Backup 77
12.2.1 Do I Need to Stop the Server Before Performing a Backup? 77
12.2.2 Can I Back Up a Server Deployed with Databases? 78
12.2.3 How Can I Distinguish Automatic Backups From Manual Backups? 78
12.2.4 Can I Choose to Back Up Only Some Partitions of a Disk? 78
12.2.5 Does CBR Support Cross-Region Backup? 78
12.2.6 Will the Server Performance Be Affected If I Delete Its Backups? 79
12.2.7 Can I Use Its Backup for Restoration After a Resource Is Deleted? 79
12.2.8 How Many Backups Can I Create for a Resource? 79

12.2.9 Can I Stop an Ongoing Backup Task?....79

12.3 Restoration....79

12.3.1 Do I Need to Stop the Server Before Restoring Data Using Backups?....79

12.3.2 Can I Use a System Disk Backup to Recover an ECS?....79

12.3.3 Do I Need to Stop the Server Before Restoring Data Using Disk Backups?....79

12.3.4 Can a Server Be Restored Using Its Backups After It Is Changed?....80

12.3.5 Can a Disk Be Restored Using Its Backups After Its Capacity Is Expanded?....80

12.3.6 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?....80

12.3.7 What Changes Will Be Made to the Original Backup When I Use the Backup to Restore a Server?....80

12.3.8 Can I Stop an Ongoing Restoration Task?....80

12.4 Policies....81

12.4.1 How Do I Configure Automatic Backup for a Server or Disk?....81

12.4.2 Why Does the Retention Rule Not Take Effect After Being Changed?....81

12.4.3 How Do I Back Up Multiple Resources at a Time?....82

12.5 Others....83

12.5.1 Is There a Quota for CBR Vaults?....83

12.5.2 Can I Merge My Vaults?....83

13 Troubleshooting Cases....84

13.1 Failed to Attach Disks....84

13.2 Data Disks Are Not Displayed After a Windows Server Is Restored....85

13.3 Failed to Download or Install the Agent Required by Application-Consistent 87

13.4 A Server Created Using an Image Enters Maintenance Mode After Login....90

A Appendix....94

A.1 Agent Security Maintenance....94

A.1.1 Changing the Password of User rdadmin....94

A.1.2 Changing the Password of the Account for Reporting Alarms (SNMP v3)....95

A.1.3 Replacing the Server Certificate....97

A.1.4 Replacing CA Certificates....99

A.2 Change History....101

1 Service Overview

1.1 What Is CBR?

1.2 Application Scenarios

1.3 Functions

1.4 Permissions Management

1.5 Constraints

1.6 CBR and Other Services

1.7 Basic Concepts

1.1 What Is CBR?

Overview

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

CBR protects your services by ensuring the security and consistency of your data.

Product Architecture

CBR consists of backups, vaults, and policies.

Backup

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

- Cloud disk backup. This type of backup provides snapshot-based data protection for EVS disks.

- Cloud server backup. This type of backup uses the consistency snapshot technology for disks to protect data of ECSs and BMSs. The backups of servers without deployed databases are common server backups, and those of servers with deployed databases are application-consistent backups.
- SFS Turbo backup. This type of backup protects data of SFS Turbo file systems.

Vault

CBR uses vaults to store backups. Before creating a backup, you need to create at least one vault and associate the resource you want to back up with the vault. Then the generated resource backups are stored in the associated vault.

The backups of different types of resources must be stored in different types of vaults.

Policy

- Backup policies: To perform automatic backups, configure a backup policy by setting the execution times of backup tasks, the backup frequency, and retention rules, and then apply the policy to a vault.

Figure 1-1 Architecture of CBR

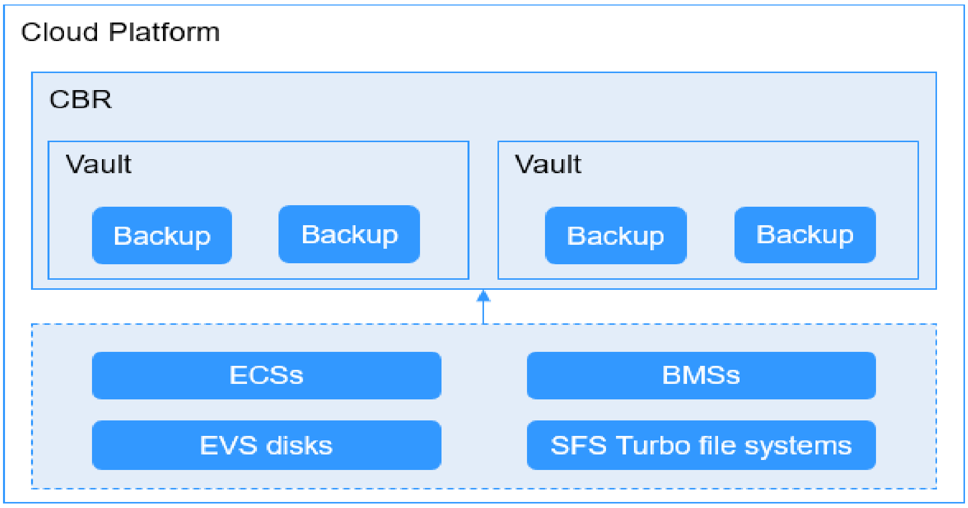

Differences Among the Backup Types

Table 1-1 Differences among the backup types

Item	Cloud Server Backup	Cloud Disk Backup	SFS Turbo Backup
Backup and restore object	All disks (system and data disks) on a server	One or more specific disks (system or data disks)	SFS Turbo file systems
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.	Data in the SFS Turbo file systems needs to be protected.
Advantages	All disks on a server are backed up at the same time, ensuring data consistency.	Backup cost is reduced without compromising data security.	Backup data and original file systems are stored separately. You can use the backup data to create a new file system.

Backup Mechanism

CBR in-cloud backups offer block-level backup. A full backup is performed only for the first backup and backs up all used data blocks. For example, if the size of a disk is 100 GB and the used space is 40 GB, the 40 GB of data is backed up. An incremental backup backs up only the data changed since the last backup, which is storage- and time-efficient. When a backup is deleted, only the data blocks that are not depended on by other backups are deleted, so that other backups can still be used for restoration. Both a full backup and an incremental backup can restore data to the state at a given backup point in time.

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

CBR stores backup data in OBS to enhance backup data security.

Backup Options

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

Table 1-2 describes 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 driven by a 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	Full backup for the first time and incremental backup subsequently, by default	Full backup for the first time and incremental backup subsequently, by default
Application scenario	Executed before patching or upgrading the OS or upgrading an application on a resource. A one-off backup can be used to restore the resource to the original state if the patching or upgrading fails.	Executed for routine maintenance of a resource. The latest backup can be used for restoration if an unexpected failure or data loss occurs.

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

Figure 1-2 Intermixed use of the two backup options

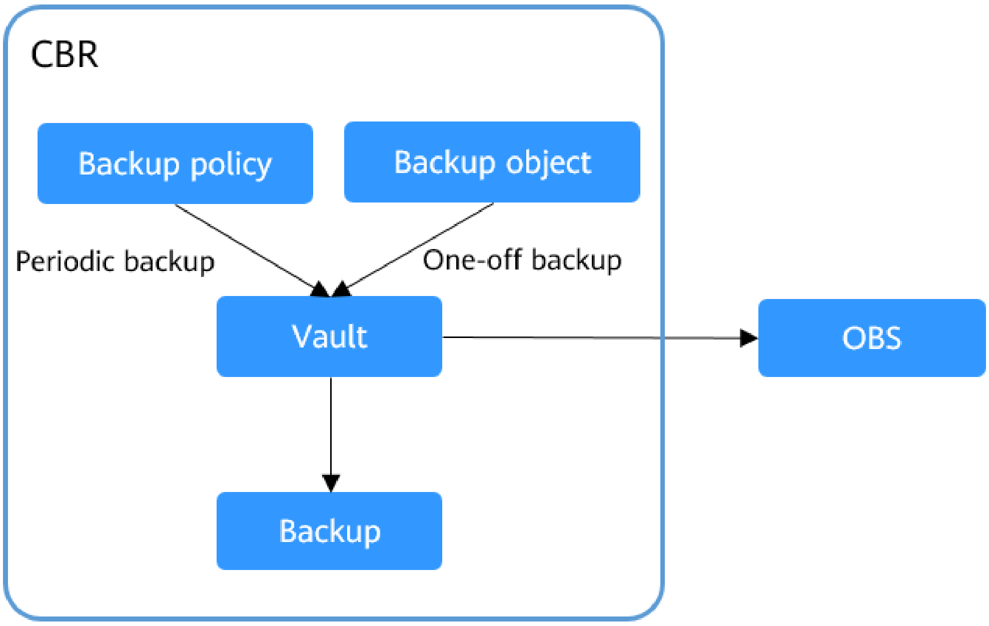

Method of Access

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

- Console

  Use the console if you prefer a web-based UI to perform operations. Log in to the console and choose **Cloud Backup and Recovery**.
- APIs

  Use APIs if you need to integrate CBR into a third-party system for secondary development. For details, see the *Cloud Backup and Recovery API Reference*.

1.2 Application Scenarios

CBR backs up resources to maximize user data security and consistency and ensure service continuity. CBR is suitable for data backup and restoration.

Data Backup and Restoration

CBR can be used to quickly restore data in the following scenarios:

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

  For any of the incidents above, you can use CBR to restore data to the latest backup point prior to the incident.

1.3 Functions

Table 1-3 lists the basic functions of CBR.

Before using this service, it is recommended that you go to **basic concepts** to learn more about CBR, such as about vault and backup policy.

Table 1-3 CBR basic functions

Category	Function	Description
Cloud disk backup	Backing up disks	Cloud disk backup provides snapshot-based data protection for EVS disks. You can use CBR to back up a single disk on a server to protect the data on that disk.

Category	Function	Description
	Policy-driven data backup	A backup policy allows a vault to automatically execute backup tasks at specified times or intervals. Periodic backups can be used to restore data quickly against data corruption or loss.
	Backup management	When a backup task is running or completed, you can set search criteria to filter backups from the backup list to manage them and view their details.
	Restoring disk data using backups	When a disk is faulty, or disk data is lost due to misoperations, you can use a backup to restore the disk.
	Creating disks using backups	You can use a disk backup to create a disk. After the disk is created, data on the new disk is the same as that in the disk backup.
	Sharing backups	You can share a server or disk backup with other accounts. Shared backups can be used to create disks or servers.
Cloud server backup	Backing up servers	Cloud server backup uses the consistency snapshot technology for disks to protect data of ECSs. You can use CBR to back up an entire server to protect the data on the server.
	Policy-driven data backup	A backup policy allows a vault to automatically execute backup tasks at specified times or intervals. Periodic backups can be used to restore data quickly against data corruption or loss.
	Backup management	When a backup task is running or completed, you can set search criteria to filter backups from the backup list to manage them and view their details.
	Restoring server data using backups	When a server is faulty, or server data is lost due to misoperations, you can use a backup to restore the server.

Category	Function	Description
	Sharing backups	You can share a server or disk backup with other accounts. Shared backups can be used to create disks or servers.
	Creating images using backups	Cloud server backup allows you to create images using ECS backups. You can use the images to provision ECSs to rapidly restore service running environments.
	Backing up database servers	Cloud server backup supports application-consistent backup in addition to crash-consistent backup. Application-consistent backup ensures the consistency of application data by backing up files and disks at the exact same time. It is suitable for backing up ECSs as well as the MySQL or SAP HANA databases running on them.
SFS Turbo backup	Backing up SFS Turbo file systems	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-driven data backup	A backup policy allows a vault to automatically execute backup tasks at specified times or intervals. Periodic backups can be used to restore data quickly against data corruption or loss.
	Backup management	When a backup task is running or completed, you can set search criteria to filter backups from the backup list to manage them and view their details.
	Restoring file system data using backups	When a file system is faulty or the data is lost due to misoperations, you can use a backup to restore the file system.
	Creating file systems using backups	You can use an SFS Turbo file system backup to create a new file system. After it is created, data on the new file system is the same as that in the backup.

1.4 Permissions Management

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

With IAM, you can use your cloud account to create IAM users for your employees, and assign permissions to the users to control their access to specific resource types. For example, some software developers in your enterprise need to use CBR resources but should not be allowed to delete the resources or perform any other high-risk operations. In this scenario, you can create IAM users for the software developers and grant them only the permissions required for using CBR resources.

If your cloud account does not require individual IAM users for permissions management, skip this section.

IAM can be used free of charge. You pay only for the resources in your account. For more information about IAM, see *Identity and Access Management User Guide*.

CBR Permissions

By default, new IAM users do not have permissions assigned. You need to add a user to one or more groups, and attach permissions policies or roles to these groups. Users inherit permissions from the groups to which they are added and can perform specified operations on cloud services based on the permissions.

CBR is a project-level service deployed and accessed in specific physical regions. To assign CBR permissions to a user group, specify the scope as region-specific projects and select projects for the permissions to take effect. If **All projects** is selected, the permissions will take effect for the user group in all region-specific projects. When accessing CBR, the users need to switch to a region where they have been authorized to use this service.

You can grant users permissions by using roles and policies.

- Roles: A type of coarse-grained authorization mechanism that defines permissions related to users responsibilities. Only a limited number of service-level roles for authorization are available. When using roles to grant permissions, you need to also assign other roles on which the permissions depend to take effect. However, roles are not an ideal choice for fine-grained authorization and secure access control.
- Policies: A type of fine-grained authorization mechanism that defines permissions required to perform operations on specific cloud resources under certain conditions. This mechanism allows for more flexible policy-based authorization, meeting requirements for secure access control. For example, you can grant ECS users only the permissions for managing a certain type of ECSs. Most policies define permissions based on APIs. For the API actions supported by CBR, see "Permissions Policies and Supported Actions".

Table 1-4 lists all the system-defined roles and policies supported by CBR.

Table 1-4 System-defined policies supported by CBR

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

Table 1-5 lists the common operations supported by each system-defined policy or role of CBR. Select the policies or roles as required.

Table 1-5 Common operations supported by each system-defined policy or role of CBR

Operation	CBR FullAccess	CBR BackupsAndVaultsFullAccess	CBR ReadOnlyAccess
Querying vaults	√	√	√
Creating vaults	√	√	×
Listing vaults	√	√	√
Updating vaults	√	√	×
Deleting vaults	√	√	×
Associating resources	√	√	×
Dissociating resources	√	√	×
Creating policies	√	×	×
Updating policies	√	×	×
Applying policies to a vault	√	√	×
Removing policies from a vault	√	√	×
Deleting policies	√	×	×

Operation	CBR FullAccess	CBR BackupsAndVaultsFullAccess	CBR ReadOnlyAccess
Performing backups	√	√	×
Updating subscriptions	√	√	×
Querying the Agent status	√	√	×
Deleting backups	√	√	×
Restoring data using backups	√	√	×
Associating vaults	√	√	×
Batch adding or deleting vault tags	√	√	×
Adding vault tags	√	√	×
Editing tags	√	√	×

1.5 Constraints

General

- A vault can be associated with only one backup policy.
- A vault can be associated with a maximum of 256 resources.
- A maximum of 32 backup policies can be created.
- Only the backups in a vault whose status is **Available** or **Locked** can be used for data restoration.
- Backups in a vault whose status is **Deleting** cannot be deleted.
- Backups cannot be downloaded to a local PC or uploaded to OBS.
- A vault and its associated servers or disks must be in the same region.

Cloud Disk Backup

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

Cloud Server Backup

- Shared disks on a server can be backed up, but there can be no more than 10 shared disks.
- Only backups in a vault whose status is **Available** or **Locked** can be used to create images.

- You can choose to back up only specified disks on a server, but such a backup of disks must be restored as a whole. File- or directory-level restoration is not supported.
- Images cannot be created using backups if the amount of resources associated with a server backup vault exceeds the quota.
- Only ECS backups can be used to create images.
- You are advised not to back up a server whose disk size exceeds 4 TB.

SFS Turbo Backup

- Only file systems in the **Available** state can be backed up.

1.6 CBR and Other Services

CBR-related Services

Table 1-6 CBR-related services

Function	Related Service	Reference
CBR backs up data of disks on an ECS, and restores backup data to disks of an ECS to restore lost or corrupted data. Generated backups can be used to create images to rapidly restore service running environments.	ECS	Creating a Cloud Server Backup Creating a Cloud Disk Backup
CBR backs up data of disks on a BMS, and restores backup data to disks of a BMS to restore lost or corrupted data. The backup and management processes for BMSs and ECSs are the same.	BMS	**1.1 What Is CBR?** Creating a Cloud Server Backup
CBR backs up data of SFS Turbo file systems. You can use backup data to create new file systems to restore lost or corrupted data.	SFS	Creating a File System Backup
CBR stores backup data in OBS to enhance backup data security.	OBS	**1.1 What Is CBR?**
CBR backs up data on disks. You can use backup data to create new disks.	EVS	
IAM is a self-service system for enterprises to manage cloud resources. It provides user identity management and access control functions.	IAM	**1.4 Permissions Management**

Function	Related Service	Reference
Tag Management Service (TMS) enables you to add preset tags to CBR vaults to facilitate filtering and management.	TMS	**4.5 Managing Vault Tags**

1.7 Basic Concepts

1.7.1 CBR Basic Concepts

Vault

CBR uses vaults to store backups of resources. Backup vaults are classified into the following types:

- **Server backup vaults**: They include those that only store backups of common servers and those that store backups of database servers. You can associate servers with a server backup vault and apply a backup policy to the vault. You can also replicate backups from a vault in one region to a replication vault in another region. Server backups can be used to restore server data.
- **Disk backup vaults**: store only disk backups. You can associate disks with a disk backup vault and apply a backup policy to the vault.
- **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 the vault.

Backup

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

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

- The name of a one-off backup is **manualbk_***xxxx*. It can be user- or system-defined.
- The name of a periodic backup is **autobk_***xxxx*, which is assigned automatically by the system.

Backup Policy

A backup policy is a set of rules for backing up data, including the policy name, policy status, execution time of backup tasks, backup frequency, and retention rule. A retention rule specifies for how long backups are retained or the number of backups that are retained. Automatic backups can be performed by applying a backup policy to a backup vault.

Application-Consistent Backup

There are three backup types in terms of backup consistency:

- Inconsistent backup: Files in an inconsistent backup contain data taken from different points in time. This typically occurs if changes are made to your files or the data on your disks while backup is running.
- Crash-consistent backup: A crash-consistent backup captures data that exists on disks as of the backup time, without backing up memory data or quiescing application systems. Backup consistency of application systems is not ensured. To complete this, disks are checked upon operating system restart to restore damaged data, for example, by using **chkdsk**, and log rollback is performed on databases to keep data consistent.
- Application-consistent backup: An application-consistent backup is a backup of application data that allows applications to achieve a quiescent and consistent state. This type of backup captures the contents of the memory and any pending writes that occurred during the backup process.

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.

1.7.2 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

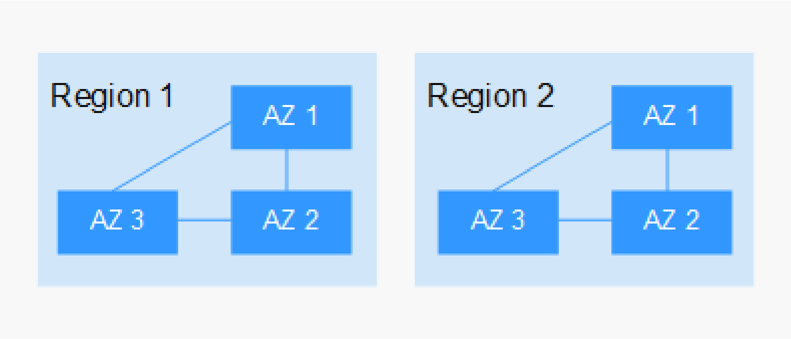

Selecting a Region

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

Selecting an AZ

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

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

Regions and Endpoints

Before you use an API to call resources, specify its region and endpoint. For more details, see **Regions and Endpoints**.

2 Getting Started

2.1 Step 1: Create a Vault

2.2 Step 2: Associate a Resource with the Vault

2.3 Step 3: Create a Backup

2.1 Step 1: Create a Vault

2.1.1 Creating a Server Backup Vault

This section describes how to create a server backup vault.

Procedure

Step 1 In the upper right corner of the page, click **Create Server Backup Vault**.

Step 2 Select a protection type.

- **Backup**: The created vault is a server backup vault, which stores cloud server backups.

Step 3 Enable or disable application-consistent backup.

- If application-consistent backup is enabled, the vault can be used to store database backups. Backing up memory data through application-consistent backup ensures application system consistency, which is suitable for ECSs containing MySQL or SAP HANA databases. If an application-consistent backup task fails, the system automatically performs a common server backup task instead. The common server backup will be stored in the application-consistent backup vault.
- If application-consistent backup is disabled, only common server backup is performed on associated servers, which is usually used for ECSs that do not run databases.

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

NOTE

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

Step 5 Specify the vault capacity, which ranges from 10 GB to 10,485,760 GB. You need to 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 automatic association is enabled and a backup policy is applied to the vault, more capacity is required.

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

Step 6 Determine whether to 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, the system 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 it.

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

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

Step 8 (Optional) Configure automatic resource association.

- If you select **Configure**, in the next backup period, unprotected resources will be automatically scanned and associated with the vault, and backups will be automatically executed.
- If you select **Skip**, resources will not be automatically associated with the vault you are creating.

You can filter unprotected resources by tag. If a tag is selected, only unprotected resources having the specified tag will be associated with the vault. Or, all unprotected resources will be associated.

Only existing tags can be selected. If no tag is available, create tags on the corresponding resource page. You can select a maximum of 5 tags to search for vaults. If you select more than one tag, the vaults containing any of the specified tags will be returned.

Step 9 (Optional) Add tags to the vault.

A tag is represented in the form of a key-value pair. Tags are used to identify, classify, and search for vaults. Vault tags are used to filter and manage vaults only. A vault can have a maximum of 10 tags.

Table 2-1 describes the parameters of a tag.

Table 2-1 Tag parameter description

Parameter	Description	Example Value
Key	Each tag has a unique key. You can customize the key or select the key of an existing tag created in TMS. A tag key: ● Can contain 1 to 36 Unicode characters. ● Can contain only letters, digits, hyphens (-), and underscores (_).	Key_0001
Value	A tag value can be repetitive or left blank. A tag value: ● Can contain 0 to 43 Unicode characters. ● Can contain only letters, digits, hyphens (-), and underscores (_).	Value_0001

Step 10 Specify a name for the vault.

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

NOTE

You can use the default name, which is in the format of **vault**_*xxxx*.

Step 11 Complete the creation as prompted.

Step 12 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 **4.1 Querying a Vault**.

----End

2.1.2 Creating a Disk Backup Vault

This section describes how to create a disk backup vault.

Procedure

Step 1 In the upper right corner of the page, click **Create Disk Backup Vault**.

Step 2 (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 a vault and must be in the **Available** or **In-use** state.
- You can associate disks with the vault you are creating if you skip this step.

Step 3 Specify the vault capacity. This capacity indicates the total size of the disks that you want to associate with this vault. You need to properly plan the vault capacity, which must be at least the same as the size of the disks you want to back up. The capacity ranges from 10 GB to 10485760 GB.

Step 4 Determine whether to 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, the system 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 it.

Step 5 (Optional) Configure automatic resource association.

- If you select **Configure**, in the next backup period, unprotected resources will be automatically scanned and associated with the vault, and backups will be automatically executed.
- If you select **Skip**, resources will not be automatically associated with the vault you are creating.

You can filter unprotected resources by tag. If a tag is selected, only unprotected resources having the specified tag will be associated with the vault. Or, all unprotected resources will be associated.

Only existing tags can be selected. If no tag is available, create tags on the corresponding resource page. You can select a maximum of 5 tags to search for vaults. If you select more than one tag, the vaults containing any of the specified tags will be returned.

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

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

Step 7 (Optional) Add tags to the vault.

A tag is represented in the form of a key-value pair. Tags are used to identify, classify, and search for vaults. Vault tags are used to filter and manage vaults only. A vault can have a maximum of 10 tags.

Table 2-2 describes the parameters of a tag.

Table 2-2 Tag parameter description

Parameter	Description	Example Value
Key	Each tag has a unique key. You can customize the key or select the key of an existing tag created in TMS. A tag key: ● Can contain 1 to 36 Unicode characters. ● Can contain only letters, digits, hyphens (-), and underscores (_).	Key_0001
Value	A tag value can be repetitive or left blank. A tag value: ● Can contain 0 to 43 Unicode characters. ● Can contain only letters, digits, hyphens (-), and underscores (_).	Value_0001

Step 8 Specify a name for the vault.

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

NOTE

You can use the default name, which is in the format of **vault**_*xxxx*.

Step 9 Complete the creation as prompted.

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

You can associate disks to the new vault or perform backup for the disks. For details, see **Vault Management**.

----End

2.1.3 Creating an SFS Turbo Backup Vault

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

Procedure

Step 1 Log in to CBR Console.

1. Log in to the management console.
2. Click 📍 in the upper left corner and select your region and project.
3. 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**: The created vault is an SFS Turbo backup vault, which stores backups of SFS Turbo file systems.

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 a vault and must be in the **Available** state.
- You can associate file systems with the vault you are creating if you skip this step.

Step 5 Specify the vault capacity. This capacity indicates the total size of the file systems that you want to associate with this vault. You need to properly plan the vault capacity, which must be at least the same as the size of the file systems you want to back up. The capacity ranges from 10 GB to 10485760 GB.

Step 6 Determine whether to 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, the system 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 it.

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

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

Step 8 (Optional) Add tags to the vault.

A tag is represented in the form of a key-value pair. Tags are used to identify, classify, and search for vaults. Vault tags are used to filter and manage vaults only. A vault can have a maximum of 10 tags.

Table 2-3 describes the parameters of a tag.

Table 2-3 Tag parameter description

Parameter	Description	Example Value
Key	Each tag has a unique key. You can customize the key or select the key of an existing tag created in TMS. A tag key: • Can contain 1 to 36 Unicode characters. • Can contain only letters, digits, hyphens (-), and underscores (_).	Key_0001

Parameter	Description	Example Value
Value	A tag value can be repetitive or left blank. A tag value: ● Can contain 0 to 43 Unicode characters. ● Can contain only letters, digits, hyphens (-), and underscores (_).	Value_0001

Step 9 Specify a name for the vault.

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

NOTE

You can use the default name, which is in the format of **vault**_*xxxx*.

Step 10 Complete the creation as prompted.

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

You can associate file systems to the new vault or perform backup for the file systems. For details, see **Vault Management**.

----End

2.2 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

- The servers you plan to associate with a vault must be in the **Running** or **Stopped** state.
- The disks you plan to associate with a vault must be in the **Available** or **In-use** state.
- The SFS Turbo file systems you plan to associate with a vault must be in the **Available** state.
- The servers you plan to associate with a vault must have at least one disk attached.
- The vault and the resources to be associated must be in the same region.
- The total capacity of the resources to be associated cannot be greater than the capacity of the vault.

Procedure

Step 1 On a backup page, locate the target vault and click **Associate Server**, **Associate File System**, or **Associate Disk**.

Step 2 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 3 Click **OK**. Then in the **Associated Servers** column in the vault list, you can view the number of resources that have been successfully associated.

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

----End

Automatic Association

Backup vaults support automatic association with resources that are not backed up. After successful association, resources will be backed up according to the backup policy applied to the vault.

- You can enable automatic association only when the vault's remaining capacity is greater than 40 GB. Remaining capacity of a vault = Total capacity of the vault - Capacity of resources associated with the vault. You can obtain the vault capacity and associated capacity in the **Basic Information** area on the details page of the vault. Specifically, if the capacity of a server backup vault is 800 GB and it has been associated with two 100 GB servers, the remaining capacity is 600 GB (800 GB – 200 GB). In this case, you can enable automatic association.
- If multiple vaults are enabled with automatic association, the system scans their backup policies and associates resources with the vault whose next scheduled backup time is the earliest.
- When the capacity of the vault selected by the system is used up, resources will be associated with the vault whose next scheduled backup time is the second earliest.
- If a backup policy with the earliest scheduled backup time is applied to more than one vault, the system randomly associates the resources with one of these vaults.
- If a vault is enabled with automatic association but no backup policy has been applied to it, no resources will be automatically associated with this vault. You can manually associate resources that have not been backed up with the vault.
- After the automatic association function is disabled for a vault, the vault stops automatically scanning for resources that have not been backed up. The associated resources are not affected.

Step 1 Log in to CBR Console.

1. Log in to the management console.
2. Click 📍 in the upper left corner and select your region and project.
3. Choose **Storage** > **Cloud Backup and Recovery**.

Step 2 On any backup page, locate the target vault.

Step 3 Choose **More** > **Enable Automatic Association** in the **Operation** column of the vault.

Step 4 After the function is enabled, you can see **Automatic association** in the **Associated Servers** column of the vault list.

Step 5 (Optional) If the automatic association function is not required, choose **More** > **Disable Automatic Association** in the **Operation** column of the vault.

----End

2.3 Step 3: Create a Backup

2.3.1 Creating a Cloud Server Backup

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

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

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

During the cloud server backup, the performance of the server is not affected. To ensure data integrity, back up the server during off-peak hours when no write operation is performed on the disks.

Prerequisites

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

Procedure

Step 1 On the **Cloud Server Backups** page, click the **Vaults** tab and find the vault to which the server is associated.

Step 2 Choose **More** > **Perform Backup** in the **Operation** column. In the server list, select the server you want to back up. After a server is selected, it is added to the list of selected servers.

Step 3 Set the **Name** and **Description** for the backup. **Table 2-4** describes the parameters.

Table 2-4 Parameter description

Parameter	Description	Remarks
Name	Name of the backup you are creating. A name must contain 1 to 64 characters including digits, letters, underscores (_), or hyphens (-). **NOTE** You can use the default name, which is in the format of **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 4 Choose whether to enable full backup. If full backup is enabled, the system performs a full backup on every associated server, which requires a larger capacity compared to an incremental backup. See **Figure 2-1**.

Figure 2-1 Selecting full backup

Full Backup ? □ Enable

Step 5 (Optional) You can also click the vault name to open the detailed page of the vault. On the **Associated Servers** tab page, locate the target server. Click **Perform Backup** in the **Operation** column of the server.

Step 6 Click **Yes**. The system automatically creates a backup for the server.

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

NOTE

You can restart a server if necessary after the backup progress exceeds 10%. However, to ensure data integrity, you are advised to 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 **7.1 Restoring Data Using a Cloud Server Backup**.

----End

2.3.2 Creating a Cloud Disk Backup

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

During the cloud disk backup, the performance of the disk is not affected. To ensure data integrity, back up the disk during off-peak hours when no write operation is performed on the disk.

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

Step 2 Choose **More** > **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.

NOTE

The system will identify whether the selected disk is encrypted. If it is encrypted, its backup data will be stored in encrypted mode.

Step 3 Set the **Name** and **Description** for the backup. Table 2-5 describes the parameters.

Table 2-5 Parameter description

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

Step 4 Choose whether to enable full backup. If full backup is enabled, the system performs a full backup on every associated disk, which requires a larger capacity compared to an incremental backup. See **Figure 2-2**.

Figure 2-2 Selecting full backup

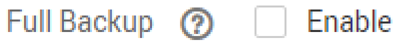

Full Backup ? □ Enable

Step 5 (Optional) You can also click the vault name to open the detailed page of the vault. On the **Associated Disks** tab page, locate the target disk. Click **Perform Backup** in the **Operation** column of the disk.

Step 6 Click **Yes**. The system automatically creates a backup for the disk.

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

NOTE

If you delete files from the disk during the backup, backup of the deleted files may fail. To ensure data integrity, you are advised to wait until the backup task is complete and then delete data and perform a backup again.

After the backup is complete, you can use the backup to restore disk data. For details, see **7.2 Restoring Data Using a Cloud Disk Backup**.

----End

2.3.3 Creating an SFS Turbo Backup

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

During the SFS Turbo file system backup, the performance of the file system is not affected. To ensure data integrity, back up the file system during off-peak hours when no write operation is performed on the file system.

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 CBR Console.

1. Log in to the management console.
2. Click 📍 in the upper left corner and select your region and project.
3. 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 Choose **More** > **Perform Backup** in the **Operation** column. In the file system list, select the file system to be backed up. After a file system is selected, it is added to the list of selected file systems.

Step 4 Set the **Name** and **Description** for the backup. **Table 2-6** describes the parameters.

Table 2-6 Parameter description

Parameter	Description	Remarks
Name	Name of the backup you are creating. A name must contain 1 to 64 characters including digits, letters, underscores (_), or hyphens (-). NOTE You can use the default name, which is in the format of **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 (Optional) You can also click the vault name to open the detailed page of the vault. On the **Associated File Systems** tab page, locate the target file system. Click **Perform Backup** in the **Operation** column of the file system.

Step 6 Click **Yes**. The system automatically creates a backup for the file system.

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

NOTE

If you delete files from the file system during the backup, backup of the deleted files may fail. To ensure data integrity, you are advised to wait until the backup task is complete and then delete data and perform a backup again.

After the backup is complete, you can create a new SFS Turbo file system using the backup. For details, see **5.6 Using a Backup to Create a File System**.

----End

3 Permissions Management

3.1 Creating a User and Granting CBR Permissions

3.2 Creating a Custom Policy

3.1 Creating a User and Granting CBR Permissions

This section describes how to use IAM to implement fine-grained permissions control for your CBR resources. With IAM, you can:

- Create IAM users for employees based on your enterprise's organizational structure. Each IAM user will have their own security credentials for accessing CBR resources.
- Grant only the permissions required for users to perform a specific task.
- Entrust a cloud account or cloud service to perform efficient O&M on your CBR resources.

If your cloud account does not require individual IAM users, skip this section.

This section describes the procedure for granting permissions (see **Figure 3-1**).

Prerequisites

Learn about the permissions (see **1.4 Permissions Management**) supported by CBR and choose policies or roles according to your requirements. For the system policies of other services, see "Permissions".

Process Flow

Figure 3-1 Process for granting CBR permissions

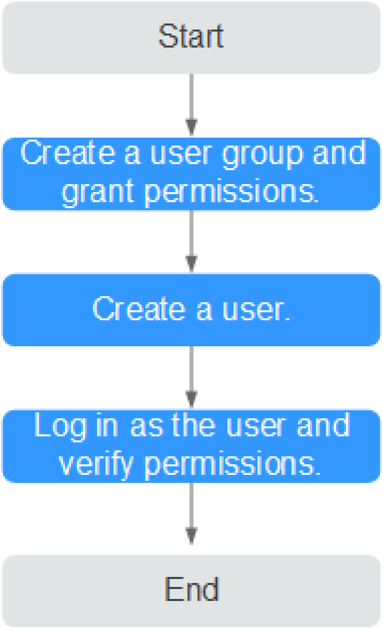
Start
Create a user group and grant permissions.
Create a user.
Log in as the user and verify permissions.
End

1. Create a user group and assign permissions to it.

   Create a user group on the IAM console, and assign the **CBR ReadOnlyAccess** policy to the group.

2. Create an IAM user and add it to the user group.

   Create a user on the IAM console and add the user to the group created in **1**.

3. Log in and verify permissions.

   Log in to CBR Console using the created user, and verify that the user has read-only permissions for CBR.

   – Choose **Service List** > **Cloud Backup and Recovery**. Then click **Buy Server Backup Vault** on CBR Console. If a message appears indicating that you have insufficient permissions to perform the operation, the **CBR ReadOnlyAccess** policy has already taken effect.

   – Choose any other service in **Service List**. If a message appears indicating that you have insufficient permissions to access the service, the **CBR ReadOnlyAccess** policy has already taken effect.

3.2 Creating a Custom Policy

Custom policies can be created to supplement the system-defined policies of CBR. For the actions supported for custom policies, see section "Permissions Policies and Supported Actions" in *Cloud Backup and Recovery API Reference*.

You can create custom policies in either of the following ways:

- Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy syntax.

- JSON: Edit JSON policies from scratch or based on an existing policy.

This section provides examples of common user-defined CBR policies.

Example Custom Policies

- Example 1: Allowing users to create, modify, and delete vaults

```
{
    "Version": "1.1",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "cbr:*:get*",
                "cbr:*:list*",
                "cbr:vaults:update",
                "cbr:vaults:delete",
                "cbr:vaults:create"
            ]
        }
    ]
}
```

- Example 2: Denying users to delete vaults and backups

  A policy with only "Deny" permissions must be used in conjunction with other policies to take effect. If the permissions assigned to a user contain both "Allow" and "Deny", the "Deny" permissions take precedence over the "Allow" permissions.

  The following method can be used if you need to assign permissions of the **CBR FullAccess** policy to a user but you want to prevent the user from deleting vaults and backups. Create a custom policy for denying vault and backup deletion, and attach both policies to the group to which the user belongs. Then, the user can perform all operations on CBR except deleting vaults or backups. The following is an example of a deny policy:

```
{
   "Version": "1.1",
   "Statement": [
      {
         "Effect": "Deny",
         "Action": [
            "cbr:backups:delete",
            "cbr:vaults:delete"
         ]
      }
   ]
}
```

- Example 3: Defining permissions for multiple services in a policy

  A custom policy can contain the actions of multiple services that are of the global or project-level type. The following is an example policy containing actions of multiple services:

```
{
   "Version": "1.1",
   "Statement": [
      {
         "Effect": "Allow",
         "Action": [
            "cbr:vaults:create",
            "cbr:vaults:update",
            "cbr:vaults:delete"
         ]
      },
      {
            "Effect": "Allow",
            "Action": [
                "sfs:shares:createShare"
            ]
        }
    ]
}
```

4 Vault Management

4.1 Querying a Vault

4.2 Deleting a Vault

4.3 Dissociating a Resource

4.4 Expanding Vault Capacity

4.5 Managing Vault Tags

4.6 Managing the Enterprise Projects of Vaults

4.1 Querying a Vault

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

Prerequisites

A vault has been created.

Viewing Vault Details

Step 1 View the basic information about vaults. Related parameters are described in the following table.

Table 4-1 Basic information parameters

Parameter	Description
Name/ID	Name and ID of the vault. Click the vault name to view details about the vault.
Type	Vault type
Status	Vault status. **Table 4-2** describes the vault statuses.

Parameter	Description
Specifications	Vault specifications, which can be server backup and application-consistent backup ● A server backup vault stores backups of common servers. ● An application-consistent backup vault stores backups of database servers.
Vault Capacity (GB)	Capacity used by backups in the vault. It shows the space used by backups and the vault capacity. For example: If **20/100** is displayed, 20 GB has been used out of the 100 GB vault capacity.
Associated Servers/ File Systems/Disks	Number of servers, file systems, and disks associated with the vault. You can click the number to view details of associated resources. The associated capacity shown on the details page is the total capacity of all the resources that have been associated with this vault.

Step 2 On any backup page, click the **Vaults** tab and set filter criteria to view the vaults.

- Select a value from the status drop-down list to query vaults by status. **Table 4-2** describes the vault statuses.

Table 4-2 Vault statuses

Status	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 when a capacity expansion is in progress. In this state, you cannot expand the vault capacity. However, you can perform other operations, such as applying a policy and associating servers, file systems, or disks. After the capacity expansion is complete, the vault status becomes **Available**.

Status	Status Attribute	Description
Deleting	An intermediate state	An intermediate state when a vault is being deleted. In this state, a progress bar is displayed indicating the deletion progress. If the progress bar remains unchanged for an extended time, an exception has occurred. Contact technical support.
Error	A stable state	A vault enters the **Error** state when an exception occurs during task execution. You can click **Tasks** in the navigation pane on the left to view the error cause. If the error persists, contact technical support.

- Search the vault by its name or ID.
- Click **Search by Tag** in the upper right corner to search for vaults by tag.
  - On the **Search by Tag** tab page that is displayed, enter an existing tag key and value and click +. The added tag search criteria are displayed under the text boxes. Click **Search** in the lower right corner.
  - You can use more than one tag for a combination search. Each time after a key and a value are entered, click +. The added tag search criteria are displayed under the text boxes. When more than one tag is added, the tags will be applied together for a combination search. A maximum of 10 tags can be added at a time.
  - You can click **Reset** in the lower right corner to reset the search criteria.

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

NOTE

For the values of used capacity and backup space, only the integer part is maintained, and the decimal part is rounded off. For example, the used backup space is displayed as 0 GB, but the backup space that has actually been used might be 0.2 GB.

----End

4.2 Deleting a Vault

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

All backups stored in the vault will be deleted once you delete a vault.

Prerequisites

- At least one vault exists.
- The vault is in the **Available** or **Error** state.

Procedure

Step 1 On any backup page, locate the vault to be deleted and choose **More** > **Delete** in the **Operation** column. All backups stored in the vault will be deleted once you delete a vault. Exercise caution when performing this operation.

Step 2 Click **Yes**.

----End

4.3 Dissociating a Resource

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

After a resource is dissociated, the backup policy of the vault no longer has any effect on the resource. In addition, all manual and automatic backups of this resource will be deleted. Deleted backups cannot be used for data restoration. Exercise caution when performing this operation.

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

Procedure

Step 1 On any backup page, locate the target vault and click the vault name.

Step 2 In this example, we will be using the **Cloud Server Backups** page to illustrate the process. Click the **Associated Servers** tab. Find the target server and click **Dissociate** in the **Operation** column.

After a resource is dissociated, the backup policy of the vault no longer has any effect on the resource. In addition, all manual and automatic backups of this resource will be deleted. Deleted backups cannot be used for data restoration. Exercise caution when performing this operation.

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

----End

4.4 Expanding Vault Capacity

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

Procedure

Step 1 On any backup page, locate the target vault and choose **More** > **Expand Capacity** in the **Operation** column.

Step 2 Enter the capacity to be added. The minimum value is **1**.

Step 3 Click **Next**. Confirm the settings and click **Submit**.

Step 4 Return to the vault list and check that the capacity of the vault has been expanded.

----End

4.5 Managing Vault Tags

You can add tags to a vault as well as edit and delete these tags. Vault tags are used to filter and manage vaults only.

Procedure

Step 1 Click the name of a vault and select the **Tag** tab in the displayed vault information page.

- Adding a tag
  a. Click **Add Tag** in the upper left corner.
  b. In the dialog box that is displayed, set the key and value of the new tag.

     A tag is represented in the form of a key-value pair. Tags are used to identify, classify, and search for cloud resources. Vault tags are used to filter and manage vaults only. A vault can have a maximum of 10 tags.

     Table 4-3 describes the parameters of a tag.

Table 4-3 Tag parameter description

Parameter	Description	Example Value
Key	Tag key. Each tag of a vault has a unique key. You can customize the key or select the key of an existing tag created in TMS. The naming rules for a tag key are as follows: ▪ It contains 1 to 36 Unicode characters. ▪ It can contain only letters, digits, hyphens (-), and underscores (_).	Key_0001

Parameter	Description	Example Value
Value	A tag value can be repetitive or left blank. The naming rules for a tag value are as follows: ■ It contains 0 to 43 Unicode characters. ■ It can contain only letters, digits, hyphens (-), and underscores (_).	Value_0001

c. Click **OK**.

- Editing a tag
  a. In the **Operation** column of the tag that you want to edit, click **Edit**.
  b. In the **Edit Tag** dialog box that is displayed, modify the tag value. **Table 4-3** describes the parameters.
  c. Click **OK**.
- Deleting a tag
  a. In the **Operation** column of the tag that you want to delete, click **Delete**.
  b. In the dialog box that is displayed, confirm the deletion information.
  c. Click **OK**.

----End

4.6 Managing the Enterprise Projects of Vaults

If you need to modify the enterprise project of a vault, go to the **Enterprise Management** page to move the vault from the original enterprise project to a new one.

Procedure

Step 1 Click **Enterprise** on the upper right of console page. By default, the **Overview** page of Enterprise Management is displayed.

Step 2 In the navigation pane of the **Enterprise Management** page, choose **Enterprise Project Management**.

Step 3 Locate the enterprise project from which the vault will be removed. Click **View Resources** in the **Operation** column. The **Resources** tab page is displayed. You can view resources in the current enterprise project.

Step 4 Select **Single Resource** for the removal mode.

Step 5 Select the destination enterprise project to which the vault is to be added and click **OK**.

After the vault is removed from the enterprise project, you can view it in the resource list of the destination enterprise project.

----End

5 Backup Management

5.1 Viewing a Backup

5.2 Sharing a Backup

5.3 Deleting a Backup

5.4 Using a Backup to Create an Image

5.5 Using a Backup to Create a Disk

5.6 Using a Backup to Create a File System

5.1 Viewing a Backup

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

Prerequisites

A backup task has been created.

Viewing Backup Details

Step 1 On any backup page, click the **Backups** tab and set filter criteria to view the backups.

- You can search for backups by selecting a status from the **All statuses** drop-down list in the upper right corner of the backup list. **Table 5-1** describes the backup statuses.

Table 5-1 Backup statuses

Status	Status Attribute	Description
All statuses	--	All backups are displayed if this value is selected.
Available	A stable state	A stable state of a backup after the backup is created, indicating that the backup is available and 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 .
Error	A stable state	A backup enters the **Error** state when an exception occurs. A backup in this state cannot be used for restoration, and must be deleted manually. If manual deletion fails, contact technical support.

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

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

Step 2 Click the backup name to view details about the backup.

----End

5.2 Sharing a Backup

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

Context

Sharer

- Backups can only be shared among accounts in the same region.
- Encrypted backups cannot be shared. Backups cannot be shared across regions. Account to which a backup is shared must be in the same region as the backup.
- Accepted shared backups will be deleted once the sharer deletes the original backup. If a shared backup has been used to create new disks or servers, the created resources will not be deleted.

Recipient

- A recipient must have at least one backup vault to store the accepted shared backup, and the vault remaining space must be greater than the size of the backup to be accepted.
- A recipient can choose whether to accept a backup. After accepting a backup, the recipient can use the backup to create new servers or disks.
- Accepted shared backups will be deleted once the sharer deletes the original backup. If a shared backup has been used to create new disks or servers, the created resources will not be deleted.

Procedure for the Sharer

Step 1 On the cloud server or cloud disk backup page, click the **Backups** tab and set filter criteria to view the backups.

Step 2 Locate the target backup and choose **More** > **Share Backup** in the **Operation** column.

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

- Sharing a backup

1. Click the **Share Backup** tab.
2. Enter the account name of the tenant to whom the backup is to be shared.
3. Click **Add**. The account name and project to be added are displayed in the list. You can continue to add account names. A backup can be shared to a maximum of ten projects.
4. Click **OK**.

- Canceling sharing

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

2. On the displayed page, click the **Cancel Sharing** tab and select the backup that no longer needs to be shared. Then, click **OK**.

----End

Procedure for the Recipient

Step 1 On the cloud server or cloud disk backup page, click the **Backups** tab and then click **Backups Shared with Me**.

Step 2 Ensure that the recipient has at least one backup vault before accepting the shared backup. For how to purchase a backup vault, see **2.1 Step 1: Create a Vault**.

Step 3 Click **Accept**. On the displayed page, select the vault used to store the shared backup. Ensure that the vault remaining capacity is greater than the backup size.

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

Step 4 After a shared backup is accepted, it will be displayed in the backup list.

----End

5.3 Deleting a Backup

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

Context

CBR supports manual deletion of backups and automatic deletion of expired backups. The latter is implemented based on the backup retention rule in the backup policy. For details, see **6.1 Creating a Backup Policy**.

Prerequisites

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

Procedure

Step 1 On any backup page, click the **Backups** tab and locate the desired backup. For details, see **5.1 Viewing a Backup**.

Step 2 In the row of the backup, choose **More** > **Delete**. Alternatively, select the backups you want to delete and click **Delete** in the upper left corner to delete them in a batch.

Step 3 Click **Yes**.

----End

5.4 Using a Backup to Create an Image

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

Prerequisites

- Before backing up an ECS, ensure that the ECS has been optimized, and the Cloud-Init (for Linux) or Cloudbase-Init (for Windows) tool has been installed.
- A backup can be used to create an image in either of the following scenarios: 1. The backup is in the **Available** state. 2. The backup is in the **Creating** state which is marked with **Image can be created**.

  NOTE

  Once a backup creation starts, the backup enters the **Creating** state. After a period of time, a message stating "Image can be created" is displayed under **Creating**. In this case, the backup can be used for creating an image, even though it is still being created and cannot be used for restoration.
- The backup you want to use to create an image contains the system disk data.
- Only ECS backups can be used to create images.

Function Description

- Images created using a backup are the same, so CBR allows you to use a backup to create only one full-ECS image that contains the whole data of the system disk and data disks of an ECS, in order to save the image quota. After an image is created, you can use the image to provision multiple ECSs in a batch.
- A backup with an image created cannot be directly deleted. If you want 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. Therefore, the size of the generated image is smaller than that of the backup.

Procedure

Step 1 Click the **Backups** tab. Locate the desired backup. For details, see **5.1 Viewing a Backup**.

Step 2 In the row of the backup, choose **More** > **Create Image**.

Step 3 Create an image by referring to section "Creating a Full-ECS Image Using a Cloud Server Backup" in the *Image Management Service User Guide*.

Step 4 If you want to use an image to provision ECSs, see section "Creating ECSs Using an Image" in the *Image Management Service User Guide*.

----End

5.5 Using a Backup to Create a Disk

You can use a disk backup to create a disk. After the disk is created, data on the new disk is the same as that in the disk backup.

After a new disk is created using the backup data of a system disk, the new disk can only be mounted to the cloud server as a data disk and cannot be mounted as a system disk.

Procedure

Step 1 Click the **Backups** tab. Locate the desired backup. For details, see **5.1 Viewing a Backup**.

Step 2 If the status of the target backup is **Available**, click **Create Disk** in the **Operation** column of the backup.

Step 3 Set the disk parameters.

NOTE

For details about these parameters, see the parameter description table in section "Creating an EVS Disk" of the *Elastic Volume Service User Guide*.

Note the following items when setting disk parameters:

- You can choose the AZ to which the backup source disk belongs, or you can choose 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 source disk, initialize the disk 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 disk type.

Step 4 Click **Next**.

Step 5 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**. 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

5.6 Using a Backup to Create a File System

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

Procedure

Step 1 Click the **Backups** tab. Locate the desired backup. For details, see **5.1 Viewing a Backup**.

Step 2 If the status of the target backup is **Available**, click **Create File System** in the **Operation** column of the backup.

Step 3 Set the file system parameters.

NOTE

- For detailed parameter descriptions, see table "Parameter description" under "Creating an SFS Turbo File System" in the *Scalable File Service User Guide*.

Step 4 Click **Next**.

Step 5 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**. After the file system status has changed from **Creating** to **Available**, the file system is successfully created. After the status has changed from **Restoring** to **Available**, backup data has been successfully restored to the created file system.

----End

6 Policy Management

6.1 Creating a Backup Policy

6.2 Modifying a Policy

6.3 Deleting a Policy

6.4 Applying a Policy to a Vault

6.5 Removing a Policy from a Vault

6.1 Creating a Backup Policy

A backup policy allows a vault to automatically execute backup tasks at specified times or intervals. Periodic backups can be used to restore data quickly against data corruption or loss.

To implement periodic backup, you need to create a backup policy first. CBR will then periodically perform backups according to the execution time specified in the backup policy. You can choose to use the default backup policy provided by CBR or create one as needed.

You can apply backup policies to server backup vaults, SFS Turbo backup vaults, and disk backup vaults.

Context

- After a backup policy is enabled, CBR automatically backs up resources associated with the vaults that have been associated with the policy and periodically deletes expired backups.
- Each account can create a maximum of 32 backup policies.
- When expired backups are cleared based on a policy's retention rules, only automatic backups will be deleted. Manual backup will not be deleted.
- Only servers in the **Running** or **Stopped** state can be backed up.
- Only disks in the **Available** or **In-use** state can be backed up.

Procedure

Step 1 Choose **Policies** and click the **Backup Policies** tab. In the upper right corner, click **Create Policy**.

Step 2 Set the backup policy parameters. **Table 6-1** describes the parameters.

Table 6-1 Backup policy parameter description

Parameter	Description	Example Value
Type	Select a policy type. This section uses creating a backup policy as an example.	Backup policy
Name	Backup policy name A name must contain 1 to 64 characters including digits, letters, underscores (_), or hyphens (-).	backup_policy
Status	Whether to enable the backup policy.	Only after a backup policy is enabled will CBR automatically backs up servers and disks associated with the vaults applied with the policy and deletes expired backups.

Parameter	Description	Example Value
Execution Time	Execution time Backups can be scheduled at the beginning of each hour. Multiple selections are supported. **NOTICE** ● There may be a time difference between the scheduled backup time and the actual backup time. ● To back up a large amount of data, you are advised to set a less frequent backup schedule. If a backup task takes longer than the backup interval, the system will skip the next backup execution time. For example, as scheduled in a backup policy, a disk needs to be backed up at 00:00, 01:00, and 02:00. At 00:00, the disk starts being backed up. Because the high-volume 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. In this case, the system performs the next backup at 02:00. Therefore, only two backups will be generated in total, one at 00:00, and the other at 02:00.	00:00, 02:00 It is recommended that backups be performed during off-peak hours or when there are no services running.
Backup Cycle	Dates for performing backups ● **Week-based cycle** Specifies on which days of each week the backup task will be executed. You can select multiple days. ● **Custom cycle** Specifies the interval (every 1 to 30 days) for executing the backup task.	Every day If you select day-based backup, the first backup time is supposed to be on the day the backup policy is created. If the creation time of the backup policy is later than the latest execution time, the initial backup will be performed in the next backup cycle. It is recommended that backups be performed during off-peak hours or when there are no services running.

Parameter	Description	Example Value
Retention Rule	Rule that specifies how backups will be retained ● **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 cloud server backups to retain for one cloud server. The value ranges from 2 to 99999. ● You can also set long-term retention rules with advanced options. Long-term retention rules and quantity-based retention rules do not conflict. They will both be applied. – **Day-based**: The value ranges from **0** to **100**. – **Weekly**: The value ranges from **0** to **100**. – **Monthly**: The value ranges from **0** to **100**. – **Yearly**: The value ranges from **0** to **100**. For example, day-based advanced option retains the most recent backup by day. If a disk is backed up for multiple times in a day, only the most recent backup of that day is retained. If you set the value to 5, the system keeps the most recent backup from each of the last five days that have backups generated. If there are more than five backup files, the system automatically deletes the earliest backups. If the day-based, weekly, monthly, and yearly advanced options are all configured, the union backups are selected for retention. For example, if the number of retained day-based backups is set to 5 and the number of retained weekly backups is set to 1, five backups will be retained. The long-term retention rule and the quantity-based retention rule can be effective at the same time.	6 months

Parameter	Description	Example Value
	● **Permanent** **NOTE** – When the number of retained backups exceeds the preset value, the system automatically deletes the earliest backups. When the retention periods of retained backups exceed the preset value, the system automatically deletes all expired backups. By default, the system automatically clears data every other day. The deleted backup does not affect other backups for restoration. – Expired backups are not deleted right after they are expired. They will be deleted from 12:00 to 0:00 in batches. – This parameter applies only to backups generated based on a scheduled backup policy. Manual backups are not affected by this parameter and will not be automatically deleted. You can manually delete them from the backup list. – After a backup is used to create an image, the backup will not be counted as a retained backup and will not be deleted automatically. – A maximum of 10 backups are retained for failed periodic backup tasks. They are retained for one month and can be manually deleted.	

NOTE

More frequent backup intervals 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 cycle as needed.

Step 3 Click **OK**.

Step 4 Locate the desired vault and choose **More** > **Apply Backup Policy** to apply the created policy to the vault. You can view the applied policy on the vault details page.

After the policy is applied, data is periodically backed up to the vault based on the policy.

----End

Example

A user has a vault associated with one disk. At 10:00 a.m. on Monday, the user sets a backup policy for the vault, that is, executing a backup task at 02:00 a.m. every day and retaining a maximum of three backups. At 11:00 a.m. on Saturday, three backups are retained, which are generated on Wednesday, Thursday, and Friday. The backup generated at 2:00 a.m. on Tuesday has been automatically deleted.

6.2 Modifying a Policy

This section describes how to modify a policy.

Prerequisites

You have created at least one policy.

Procedure

Step 1 On any backup page, find the target vault and click the vault name to view the vault details.

Step 2 In the **Policies** area, click **Edit** in the row of a policy to open the policy editing page.

Related parameters are described in **Table 6-1**.

Step 3 Click **OK**.

If the policy retention rule is modified, the new rule takes effect depending on how you have changed it. For details, see **12.4.2 Why Does the Retention Rule Not Take Effect After Being Changed?**

Step 4 Alternatively, you can select **Policies** from the navigation tree on the left and edit the desired policy.

----End

6.3 Deleting a Policy

You can delete backup and replication policies if needed.

Prerequisites

You have created at least one policy.

Procedure

Step 1 Click the **Backup Policies** tab, locate the row that contains the policy you want to delete, and click **Delete**.

NOTE

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

Step 2 Confirm the information and click **Yes**.

----End

6.4 Applying a Policy to a Vault

A backup policy allows a vault to automatically execute backup tasks at specified times or intervals. Periodic backups can be used to restore data quickly against data corruption or loss.

Procedure

Step 1 On any backup page, find the target vault and choose **More** > **Apply Backup Policy**.

Step 2 You can select an existing backup policy from the drop-down list or create a new one. For how to create a policy, see **6.1 Creating a Backup Policy**.

Step 3 After the policy is successfully applied, you can view the details in the **Policies** area on the vault details page.

----End

6.5 Removing a Policy from a Vault

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

Prerequisites

A policy has been applied to the vault.

Procedure

Step 1 On any backup page, find the target vault and click the vault name to view the vault details.

Step 2 In the **Policies** area, click **Remove Policy**.

📖 NOTE

- If a backup task is being executed for a resource in the vault, the policy can be removed normally. However, the backup task will continue and backups will be generated.
- After a policy is removed, backups retained by time will expire based on the retention rule, but backups retained by quantity will not be automatically deleted. You can manually delete unwanted backups.

Step 3 Click **Yes**. The vault will no longer execute tasks as specified in this policy.

----End

7 Restoring Data

7.1 Restoring Data Using a Cloud Server Backup

7.2 Restoring Data Using a Cloud Disk Backup

7.1 Restoring Data Using a Cloud Server Backup

When disks on a server are faulty, or server data is lost due to misoperations, you can use a backup to restore the server.

Context

- Data on data disks cannot be restored to system disks.
- Data cannot be restored to servers in the **Faulty** state.

Prerequisites

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

Procedure

Step 1 Click the **Backups** tab. Locate the desired backup. For details, see **5.1 Viewing a Backup**.

Step 2 In the row of the backup, click **Restore Server**.

NOTICE

The historical data at the backup point in time will overwrite the current server data. The restoration cannot be undone.

Step 3 (Optional) Deselect **Start the server immediately after restoration**.

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

NOTICE

Servers are shut down during restoration. It is therefore recommended that you perform restoration operations during off-peak hours.

Step 4 In the **Specified Disk** drop-down list, select the target disk to which the backup will be restored.

NOTE

- If the server has only one disk, the backup is restored to the disk by default.
- If the server has multiple disks, the backup is respectively restored to the original disks by default. You can also restore the backup to another disk on the backup server by selecting the disk from the drop-down list. However, the specified destination disk must be at least as large as the backup source disk.
- Data on data disks cannot be restored to system disks.

NOTICE

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

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

Step 5 Click **Yes** and confirm the restoration is successful.

In the backup list, view the restoration status. When the backup enters the **Available** state and no new failed restoration tasks exist in **Tasks**, the restoration is successful. The resource is restored to the state at the time you took that backup.

For details about how to view failed restoration tasks, see **9 Managing Tasks**.

NOTICE

If a Windows server is restored, data disks may fail to be displayed due to Windows limitations.

After you use a cloud server backup to restore a logical volume group, the logical volume group needs to be attached again.

You need to manually online these data disks. For details, see **13.2 Data Disks Are Not Displayed After a Windows Server Is Restored**.

----End

7.2 Restoring Data Using a Cloud Disk Backup

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

Prerequisites

- The status of the disk to be restored must be **Available**.
- Before restoring the disk data, stop the server to which the disk is attached and detach the disk from the server. After the disk data is restored, attach the disk to the server and start the server.

Constraints

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

Procedure

Step 1 Click the **Backups** tab. Locate the desired backup. For details, see **5.1 Viewing a Backup**.

Step 2 In the row of the backup, click **Restore Disk**.

NOTICE

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

Step 3 Click **Yes**. You can check whether the data is successfully restored on the **Backups** tab page of **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 at the time you took that backup.

Step 4 After the restoration is complete, re-attach the disk to the server. For details, see section "Attaching an Existing Non-Shared Disk" in the *Elastic Volume Service User Guide*.

----End

8 Application-Consistent Backup

8.1 What Is Application-Consistent Backup?

8.2 Changing a Security Group

8.3 Installing the Agent

8.4 Creating an Application-Consistent Backup

8.5 Uninstalling the Agent

8.1 What Is Application-Consistent Backup?

Overview

There are three backup types in terms of backup consistency:

- Inconsistent backup: Files in an inconsistent backup contain data taken from different points in time. This typically occurs if changes are made to your files or the data on your disks while backup is running. 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 data that exists on disks as of the backup time, without backing up memory data or quiescing application systems. Backup consistency of application systems is not ensured. To complete this, disks are checked upon operating system restart to restore damaged data, for example, by using **chkdsk**, and log rollback is performed on databases to keep data consistent.
- Application-consistent backup: An application-consistent backup is a backup of application data that allows applications to achieve a quiescent and consistent state. This type of backup captures the contents of the memory and any pending writes that occurred during the backup process.

Figure 8-1 compares these backup types in detail.

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

If a MySQL or SAP HANA database is deployed on a server, you can use CBR application-consistent backup to back up the server data and application cache. Crash-consistent backup backs up only data and some application caches without interrupting services. If a system fails or data loss occurs, you can use an application-consistent backup to quickly restart services. A crash-consistent backup, however, may fail to restore some application configurations.

Figure 8-1 Backup consistency

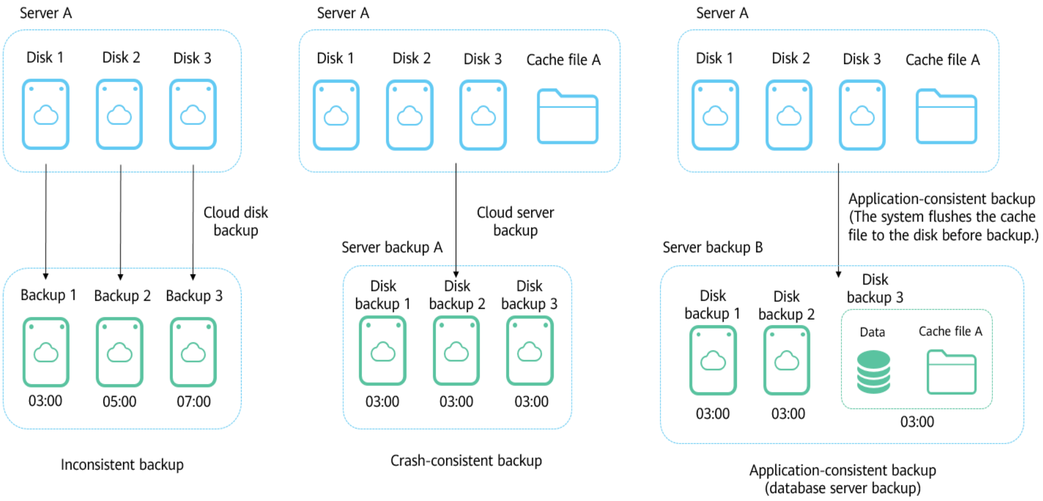

Differences Between Application-Consistent Backup Cloud Server Backup

Item	Application-Consistent Backup	Cloud Server Backup
Backup and restore object	Cloud servers with MySQL or SAP HANA database deployed	Cloud servers excluding database applications
Backup unit	Cloud server	Cloud server
Vault type	Server backup vault	Server backup vault
Recommended scenario	Data of cloud servers as well as their deployed 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. All data needs to be restored in case of an error. In the event that a cloud server backup is performed for a server deployed with a MySQL or SAP HANA database, if the backup is then used to restore data, some database configurations may fail to be restored, and issues may occur after the database is restarted.

Application Scope

Table 8-1 lists the OSs that support the installation of Agent.

Table 8-1 OSs that support installation of the Agent

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

For the databases not included in this list, you can customize a script to perform application-consistent backup by referring to section "Using a Custom Script to Implement Application-Consistent Backup" in the *Cloud Backup and Recovery Best Practices*.

Process

Figure 8-2 shows the application-consistent backup process.

Figure 8-2 Application-consistent backup process

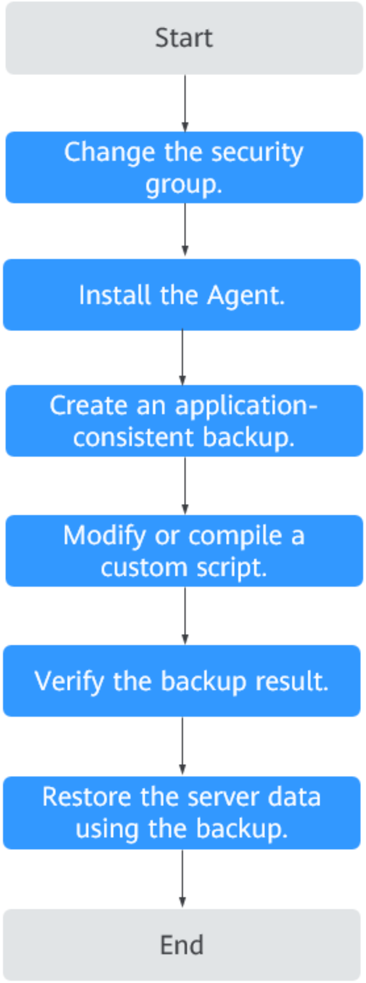
Start
Change the security group.
Install the Agent.
Create an application-consistent backup.
Modify or compile a custom script.
Verify the backup result.
Restore the server data using the backup.
End

Step 1 Change the security group: Before performing an application-consistent backup task, change the security group of the server you want to back up. For details, see **8.2 Changing a Security Group**.

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. For details, see **8.3 Installing the Agent**.

Step 3 Create an application-consistent backup: After creating a server backup vault for storing application-consistent backups, associate it with the desired database server and then create an application-consistent backup. For details, see **8.4 Creating an Application-Consistent Backup**.

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

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

Step 6 Use the backup to restore server data: Use the application-consistent backup to restore server data. The restored database applications and data are the same as those at the backup point in time. For details, see **7.1 Restoring Data Using a Cloud Server Backup**.

----End

8.2 Changing a Security Group

Context

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

When creating a security group, you must add the inbound and outbound access rules and enable the ports required for application-consistent backup to prevent application-consistent backup failures.

Operation Instructions

Before using the application-consistent backup function, you need to change the security group. To ensure network security, CBR has not set the inbound direction of a security group, so you need to manually configure it.

In the outbound direction of the security group, ports 1 to 65535 on the 100.125.0.0/16 network segment must be configured. In the inbound direction, ports 59526 to 59528 on the 100.125.0.0/16 network segment must be configured. The default outbound rule is 0.0.0.0/0, that is, all data packets are permitted. If the default rule in the outbound direction is not modified, you do not need to configure the outbound direction.

Procedure

Step 1 In the navigation pane on the left, choose **Elastic Cloud Server**. On the page displayed, select the target server. Go to the server details page.

Step 2 Click the **Security Groups** tab and select the target security group. On the right of the ECS page, click **Modify Security Group Rule** for an ECS.

Step 3 On the **Security Groups** page, click the **Inbound Rules** tab, and then click **Add Rule**. The **Add Inbound Rule** dialog box is displayed. Select **TCP** for **Protocol/Application**, enter **59526-59528** in **Port & Source**, select **IP address** for **Source** and enter **100.125.0.0/16**. After supplementing the description, click **OK** to complete the setting of the inbound rule.

Step 4 Click the **Outbound Rules** tab, and then click **Add Rule**. The **Add Outbound Rule** dialog box is displayed. Select **TCP** for **Protocol/Application**, enter **1-65535** in **Port & Source**, select **IP address** for **Destination** and enter **100.125.0.0/16**. After supplementing the description, click **OK** to complete the setting of the outbound rule.

----End

8.3 Installing the Agent

Before enabling application-consistent backup, change the security group and successfully install the Agent on your ECSs.

If application-consistent backup is enabled but Agent is not installed on servers, application-consistent backup will fail, and a common server backup will be performed instead. To ensure that application-consistent backup is properly executed, 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 **A.1.1 Changing the Password of User rdadmin**.
- **Table 8-2** lists OSs that support installation of the Agent.

Table 8-2 OSs that support installation of the Agent

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

NOTICE

To install the Agent, the system will open the firewall of a port from 59526 to 59528 of the ECS. When port 59526 is occupied, the firewall of port 59527 is enabled, and so on.

Prerequisites

- You have obtained a username and its password for logging in to the management console.
- 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 Click the **Agent Installation** tab.

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

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

Step 4 Paste the installation command in step 2 to the server and run the command as user **root**. 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 5 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 Click the **Agent Installation** tab.

Step 2 In method 2, click **Download**. On the displayed download client dialog box, select the version to be downloaded based on the operating system type of the target ECS, and click **OK**.

Step 3 After downloading the Agent, use a file transfer tool, such as Xftp, SecureFX, or WinSCP, to upload the Agent installation package to your ECS.

Step 4 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 5 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 6 Run the following command to run the installation script:

sh agent_install_ebk.sh

Step 7 The system displays a message indicating that the client is installed successfully. See **Figure 8-3**.

Figure 8-3 Successful client installation for Linux

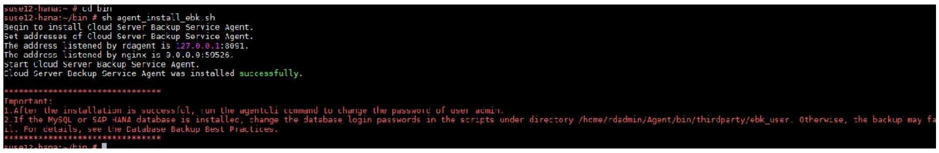

```
suse12-hana:~ # cd bin
suse12-hana:~/bin # sh agent_install_ebk.sh
Begin to install Cloud Server Backup Service Agent.
Set addresses of Cloud Server Backup Service Agent.
The address listened by rdagent is 127.0.0.1:8091.
The address listened by nginx is 0.0.0.0:59526.
Start Cloud Server Backup Service Agent.
Cloud Server Backup Service Agent was installed successfully.

*******************************
Important:
1.After the installation is successful, run the agentcli command to change the password of user admin.
2.If the MySQL or SAP HANA database is installed, change the database login passwords in the scripts under directory /home/rdadmin/Agent/bin/thirdparty/ebk_user. Otherwise, the backup may fa
il. For details, see the Database Backup Best Practices.
*******************************
suse12-hana:~/bin #
```

Step 8 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 9 Use the encrypted password in **Step 8** to replace the database login password in the script in **/home/rdadmin/Agent/bin/thirdparty/ebk_user/**.

Step 10 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 Click the **Agent Installation** tab.

Step 2 In method 1, click **Download**. Save the downloaded installation package to a local directory.

Step 3 After downloading the Agent, use a file transfer tool, such as Xftp, SecureFX, or WinSCP, to upload the Agent installation package to your ECS.

Step 4 Log in to the console and then log in to the ECS as the administrator.

Step 5 Decompress the installation package to any directory and go to the *Installation path***\bin** directory.

Step 6 Double-click the **agent_install_ebk.bat** script to start the installation.

Step 7 The system displays a message indicating that the client is installed successfully. See **Figure 8-4**.

Figure 8-4 Successful client installation for Windows

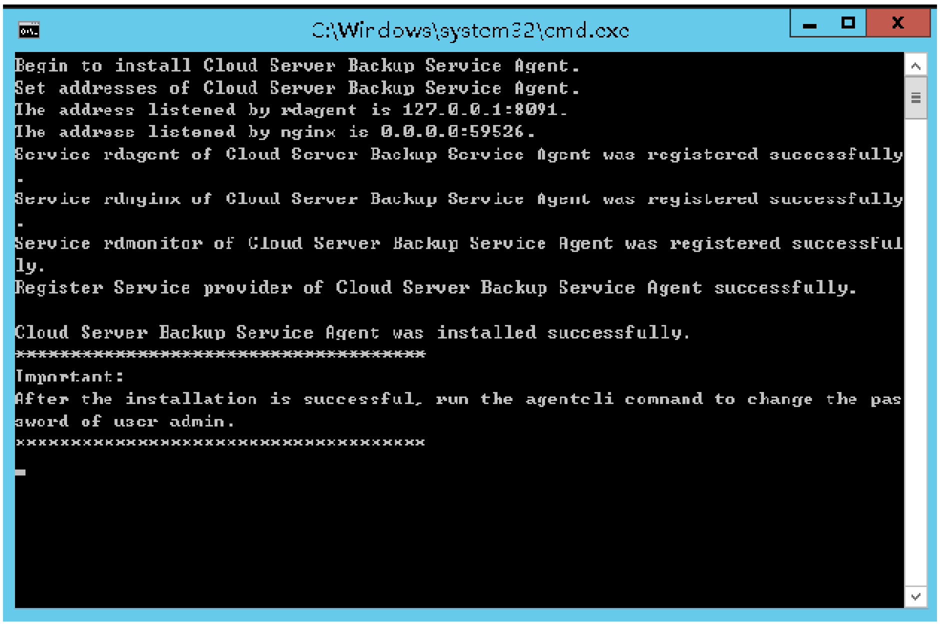

```
Begin to install Cloud Server Backup Service Agent.
Set addresses of Cloud Server Backup Service Agent.
The address listened by rdagent is 127.0.0.1:8091.
The address listened by nginx is 0.0.0.0:59526.
Service rdagent of Cloud Server Backup Service Agent was registered successfully
.
Service rdnginx of Cloud Server Backup Service Agent was registered successfully
.
Service rdmonitor of Cloud Server Backup Service Agent was registered successful
ly.
Register Service provider of Cloud Server Backup Service Agent successfully.

Cloud Server Backup Service Agent was installed successfully.
*************************************
Important:
After the installation is successful, run the agentcli comnand to change the pas
sword of user admin.
*************************************
```

----End

Installing the Agent for a Windows OS (Method 2)

Step 1 Click the **Agent Installation** tab.

Step 2 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 3 Copy the installation commands in step 2 of method 2 to the server and run the command in the Command Prompt.

Step 4 Copy any IP address in the response name, paste it in the address box of the browser, and replace **0.0.0.0** in the following address with the address. Replace ***region1*** with the actual region. The following command uses ***region1*** as an example. Then, press **Enter** in the browser to download the installation package.

http://*0.0.0.0*/csbs-agent-*region1*/Cloud Server Backup Agent-WIN64.zip

Step 5 Decompress the file to obtain the installation file. Decompress the installation package to any directory and go to the *Installation path***\bin** directory.

Step 6 Double-click the **agent_install_ebk.bat** script to start the installation.

Step 7 The system displays a message indicating that the client is installed successfully. See **Figure 8-5**.

Figure 8-5 Successful client installation for Windows

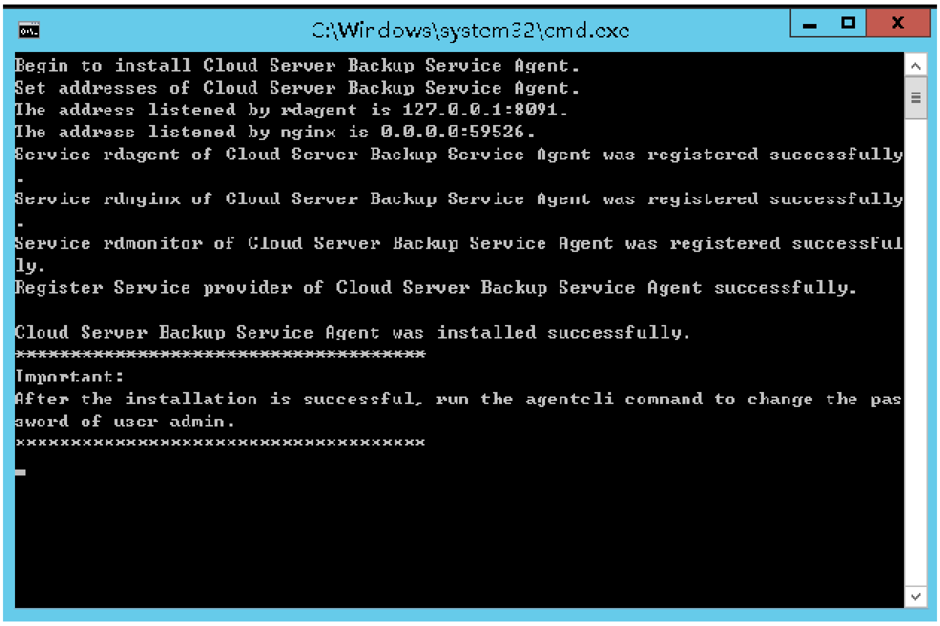

```
Begin to install Cloud Server Backup Service Agent.
Set addresses of Cloud Server Backup Service Agent.
The address listened by rdagent is 127.0.0.1:8091.
The address listened by nginx is 0.0.0.0:59526.
Service rdagent of Cloud Server Backup Service Agent was registered successfully
.
Service rdnginx of Cloud Server Backup Service Agent was registered successfully
.
Service rdmonitor of Cloud Server Backup Service Agent was registered successful
ly.
Register Service provider of Cloud Server Backup Service Agent successfully.

Cloud Server Backup Service Agent was installed successfully.
*************************************
Important:
After the installation is successful, run the agentcli comnand to change the pas
sword of user admin.
*************************************
```

----End

8.4 Creating an Application-Consistent Backup

CBR supports application-consistent backup in addition to crash-consistent backup. Application-consistent backup ensures the consistency of application data by backing up files and disks at the exact same time. It is suitable for backing up ECSs as well as the MySQL or SAP HANA databases running on them.

Constraints

- Application-consistent backup is currently not supported for cluster applications, such as, MySQL Cluster. It is supported only for applications on standalone servers.
- You are advised to perform application-consistent backup in off-peak hours.

Procedure

Step 1 Create a vault for application-consistent backups by referring to **2.1.1 Creating a Server Backup Vault**. Select **Enable** for **Application-Consistent Backup**.

Step 2 Associate the cloud servers with the created vault. Ensure that the Agent has been installed on the servers.

Step 3 Create a cloud server backup by referring to **2.3.1 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, the system 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 4 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 **7.1 Restoring Data Using a Cloud Server Backup** and **5.4 Using a Backup to Create an Image**.

8.5 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 8-6** displays an example. If the word **successfully** in green is displayed, the Agent is uninstalled successfully.

sh agent_uninstall_ebk.sh

Figure 8-6 Agent uninstalled successfully from Linux

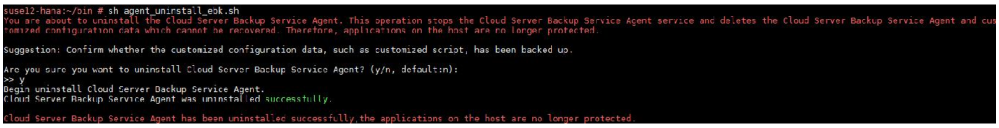

```
suse12-hana:~/bin # sh agent_uninstall_ebk.sh
You are about to uninstall the Cloud Server Backup Service Agent. This operation stops the Cloud Server Backup Service Agent service and deletes the Cloud Server Backup Service Agent and cus
tomized configuration data which cannot be recovered. Therefore, applications on the host are no longer protected.

Suggestion: Confirm whether the customized configuration data, such as customized script, has been backed up.

Are you sure you want to uninstall Cloud Server Backup Service Agent? (y/n, default:n):
>> y
Begin uninstall Cloud Server Backup Service Agent.
Cloud Server Backup Service Agent was uninstalled successfully.

Cloud Server Backup Service Agent has been uninstalled successfully,the applications on the host are no longer protected.
```

----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. See **Figure 8-7**.

Figure 8-7 Agent uninstalled successfully from Windows

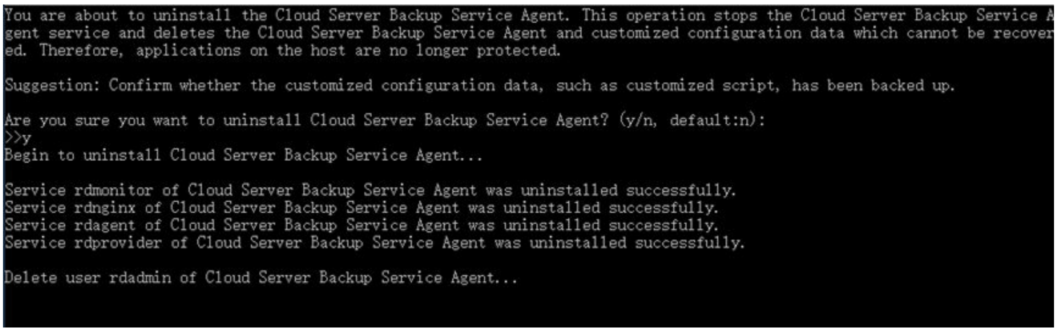

```
You are about to uninstall the Cloud Server Backup Service Agent. This operation stops the Cloud Server Backup Service A
gent service and deletes the Cloud Server Backup Service Agent and customized configuration data which cannot be recover
ed. Therefore, applications on the host are no longer protected.

Suggestion: Confirm whether the customized configuration data, such as customized script, has been backed up.

Are you sure you want to uninstall Cloud Server Backup Service Agent? (y/n, default:n):
>>y
Begin to uninstall Cloud Server Backup Service Agent...

Service rdmonitor of Cloud Server Backup Service Agent was uninstalled successfully.
Service rdnginx of Cloud Server Backup Service Agent was uninstalled successfully.
Service rdagent of Cloud Server Backup Service Agent was uninstalled successfully.
Service rdprovider of Cloud Server Backup Service Agent was uninstalled successfully.

Delete user rdadmin of Cloud Server Backup Service Agent...
```

----End

9 Managing Tasks

This section describes how to view tasks. The tasks list shows policy-driven backup tasks that have been executed over the past 30 days.

Prerequisites

At least one failed task exists.

Procedure

Step 1 Log in to CBR Console.

1. Log in to the management console.
2. Click 📍 in the upper left corner and select your region and project.
3. Choose **Storage** > **Cloud Backup and Recovery** > **Tasks**.

Step 2 You can filter tasks by project, task type, task status, task ID, resource ID, resource name, vault ID, vault name, and time.

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

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

----End

10 Monitoring

10.1 CBR Metrics

10.1 CBR Metrics

Scenarios

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

Namespace

SYS.CBR

Metrics

Table 10-1 CBR metrics

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

Dimensions

Key	Value
instance_id	Vault name/ID

Viewing Monitoring Statistics

Step 1 Log in to the management console.

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

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

Step 3 You can view the vault monitoring data by metric or monitored duration.

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

----End

11 Quotas

What Is Quota?

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

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

How Do I View My Quotas?

1. Log in to the management console.
2. Click 📍 in the upper left corner and select the desired region and project.
3. In the upper right corner of the page, click 📊 .

   The **Service Quota** page is displayed.
4. View the used and total quota of each type of resources on the displayed page.

   If a quota cannot meet service requirements, apply for a higher quota.

How Do I Apply for a Higher Quota?

1. Log in to the management console.
2. In the upper right corner of the page, choose **Resources** > **My Quotas**.

   The **Service Quota** page is displayed.
3. Click **Increase Quota**.
4. On the **Create Service Ticket** page, configure parameters as required.

   In **Problem Description** area, fill in the content and reason for adjustment.
5. After all necessary parameters are configured, select **I have read and agree to the Tenant Authorization Letter and Privacy Statement** and click **Submit**.

12 FAQs

12.1 Concepts

12.2 Backup

12.3 Restoration

12.4 Policies

12.5 Others

12.1 Concepts

12.1.1 What Are Full Backup and Incremental Backup?

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

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

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

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

NOTE

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

12.1.2 What Are the Differences Between Backup and Disaster Recovery?

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

Table 12-1 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.

12.1.3 What Are the Differences Between Backups and Snapshots?

Both backups and snapshots provide data redundancy for disks to improve data reliability. Table 12-2 lists the differences between them.

Table 12-2 Differences between backups and snapshots

Item	Storage Solution	Data Synchronization	DR Range	Service Recovery
Backup	Backup data is stored in OBS, instead of disks. This ensures data restoration upon disk data loss or corruption.	A backup is the data copy of a disk at a given point in time. CBR supports automatic backup by configuring backup policies. Deleting a disk will not clear its backups.	A backup and its source disk reside in the same AZ.	Data can be recovered and services can be restored by restoring the backup data to original disks or creating new disks from backups, ensuring excellent data reliability.

Item	Storage Solution	Data Synchronization	DR Range	Service Recovery
Snapshot	Snapshot data is stored with disk data. **NOTE** Creating a backup requires a certain amount of time because data needs to be transferred. Therefore, creating or rolling back a snapshot consumes less time than creating a backup.	A snapshot is the state of a disk at a specific point in time. If a disk is deleted, all the snapshots created for this disk will also be deleted. If you have reinstalled or changed the server OS, snapshots of the system disk are automatically deleted. Snapshots of the data disks can be used as usual.	A snapshot and its source disk reside in the same AZ.	You can use a snapshot to roll back its original disk or create a disk for data restoration and service recovery.

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

Symptoms

- There is no difference or an increase in size between the original backup and a backup generated after a file is deleted.
- The ECS backup size is larger than the used disk space obtained from the file system.

Possible Causes

Possible causes are as follows:

- The backup mechanism itself causes this problem. The cloud server backups, cloud disk backups, and SFS Turbo backups created using CBR are all block-level backups. Different from file-level backups, block-level backups are performed by sector (512 bytes) each time.
- The metadata of the file systems on the disk occupies disk space.
- To reduce performance overhead, the file system adds a delete marker for the deleted file, but does not erase the data that has been written to the sector,

and the metadata on the sector still exists. Block-level backups cannot detect whether data on a sector is deleted or not, but only determine whether a backup needs to be performed by checking whether all data blocks are zero blocks.

- CBR determines whether data in each sector changes by comparing two snapshots. Data changes include data addition, modification, and deletion. Backup is not performed if there are no data changes. If there are data changes, CBR further checks whether data blocks in the sector are all zero blocks. If so, backup is also not performed. Backups are performed only when there are non-zero blocks. If the data is deleted but metadata in the sector is not, the data block is also recognized as a non-zero block, and backups will be performed.

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

Table 12-3 describes the differences between cloud server backup and cloud disk backup.

Table 12-3 Differences between cloud server backup and cloud disk backup

Item	Cloud Server Backup	Cloud Disk Backup
Resources to be backed up or restored	All disks (system and data disks) on a server	One or more specified disks (system or data disks)
Recommended scenario	An entire cloud server needs to be protected.	Only data disks need to be backed up, because the system disk does not contain users' application data.
Advantages	All disks on a server are backed up at the same time, ensuring data consistency.	Backup cost is reduced without compromising data security.

12.2 Backup

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

To ensure data integrity, back up the server during off-peak hours when no write operation is performed on the disks. For applications that require strict consistency, such as databases and email systems, you are advised to enable application-consistent backup.

12.2.2 Can I Back Up a Server Deployed with Databases?

Yes. CBR provides application-consistent backup. For details about the function compatibility, see **Table 12-4**. 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 12-4 OSs that support installation of the Agent

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

12.2.3 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

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

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

12.2.5 Does CBR Support Cross-Region Backup?

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

12.2.6 Will the Server Performance Be Affected If I Delete Its Backups?

No. Backups are not stored on a server. Therefore, deleting its backups has no impact on the server performance.

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

Yes. Resources and backups are not stored together. If a resource is deleted, its backup still stays in your CBR vault. You can use the backup to restore the resource to a backup point in time.

12.2.8 How Many Backups Can I Create for a Resource?

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

12.2.9 Can I Stop an Ongoing Backup Task?

No. An ongoing backup task cannot be stopped.

12.3 Restoration

12.3.1 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.

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

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

You can also use a backup of the system disk to create new disks. However, newly created disks cannot be used as system disks.

12.3.3 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.

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

Yes. If a server has been backed up and then changed such as adding, deleting, or expanding disks, its backups can still be used to restore data. You are advised to back up data again after the change.

If you have added a disk after backup and then use backups to restore data, data on the newly added disk will not change.

If you have deleted a disk after backup and then use backups to restore data, data on the deleted disk will not be restored.

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

Yes. After restoration, the capacity of the expanded disk goes back to the original capacity before expansion. If you want to use the capacity added to the disk, you need to attach the restored disk to a server, log in to the server, and then manually modify the file system configuration. For detailed operations, see sections about post-expansion operations on disks in the *Elastic Volume Service User Guide*.

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

12.3.7 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.

12.3.8 Can I Stop an Ongoing Restoration Task?

No. An ongoing restoration task cannot be stopped.

12.4 Policies

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

1. Go to the Cloud Backup and Recovery console and purchase a backup vault. You are advised to set the vault capacity to at least twice the total capacity of the resources you want to back up.
2. Associate resources with the vault during or after the purchase.
3. After the resources are associated, go to the **Policies** page to configure a backup policy. You are advised to back up data during off-peak hours, for example, early in the morning. Set the backup retention rule as needed. If your vault capacity is small, set a small value for the number of backups to be kept or the days that backups will be retained. Retention rule does not apply to manual backups.
4. Apply the policy you defined to the vault. The system then will back up the resources that are associated with the vault at the specified time and retains the backups based on the retention rule.

12.4.2 Why Does the Retention Rule Not Take Effect After Being Changed?

There are the following scenarios for a retention rule change:

Rule Type Unchanged, Only with a New Backup Quantity Configured

The new policy will take effect for the backups generated based on the old policy. After a backup is generated, regardless of an automatic or a manual backup, 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 takes effect immediately. If the user then perform manual backups or wait until the system automatically create a backup at 02:00 a.m. on Friday, the system will verify and use the latest retention rule after the backup task is complete. In this case, only one most recent backup will be kept. Manual backups are not affected by policies, so they will not be deleted.

Rule Type Changed from Backup Quantity to Time Period/Permanent

The new policy will take effect only for the new backups generated. 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 takes effect immediately. If the user then perform manual backups or wait until the system automatically create a backup at 02:00 a.m. on Friday, the system will verify and use the latest retention rule after the backup task is complete. The three backups generated based on the old policy still exist (the number of backups does not exceed the quantity set in the old retention rule). They will not be automatically deleted and need to be manually deleted 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 take effect only for the new backups generated. 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 take effect for the backups generated based on the old policy.

New policy taking effect for 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 taking effect for 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.

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

1. Log in to CBR Console and click **Cloud Server Backups** or **Cloud Disk Backups** on the left navigation pane. On the displayed page, create a backup

vault. It is recommended that the capacity of the vault be at least twice the total size of resources to be backed up.

2. Associate resources with the vault during or after the creation.
3. After the resources are associated, choose **More** > **Perform Backup** in the **Operation** column of the target vault. You can manually back up two or more resources at a time.

   Alternatively, you can set a backup policy for the vault. In this way, the system will automatically back up the associated resources at the scheduled time.

12.5 Others

12.5.1 Is There a Quota for CBR Vaults?

No. You can create as many vaults as needed.

12.5.2 Can I Merge My Vaults?

No. Vaults cannot be merged.

You can expand the capacity of one vault and migrate the resources from another vault to it.

13 Troubleshooting Cases

13.1 Failed to Attach Disks

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

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

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

13.1 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

13.2 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 13-1**.

Figure 13-1 Data disks in the offline state

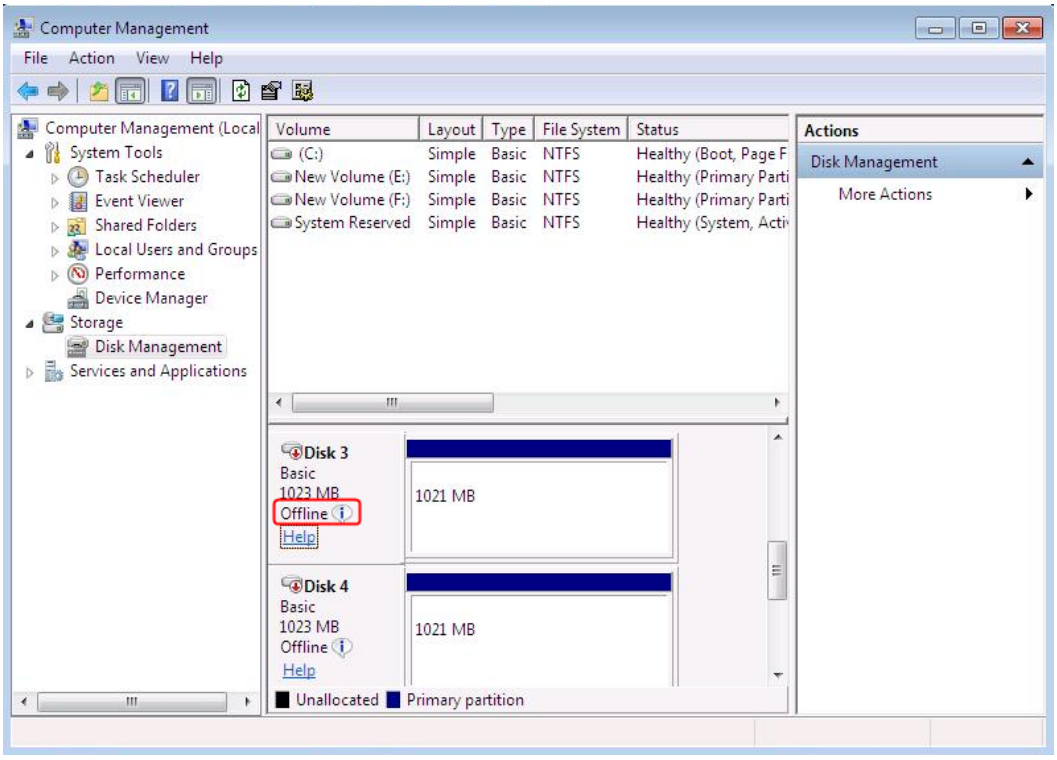

Volume	Layout	Type	File System	Status
(C:)	Simple	Basic	NTFS	Healthy (Boot, Page F
New Volume (E:)	Simple	Basic	NTFS	Healthy (Primary Parti
New Volume (F:)	Simple	Basic	NTFS	Healthy (Primary Parti
System Reserved	Simple	Basic	NTFS	Healthy (System, Acti

Disk 3
Basic
1023 MB
Offline
Help
1021 MB

Disk 4
Basic
1023 MB
Offline
Help
1021 MB

Step 4 Right-click a data disk in the offline state and choose **Online**, as shown in **Figure 13-2**.

Figure 13-2 Setting a data disk to be online

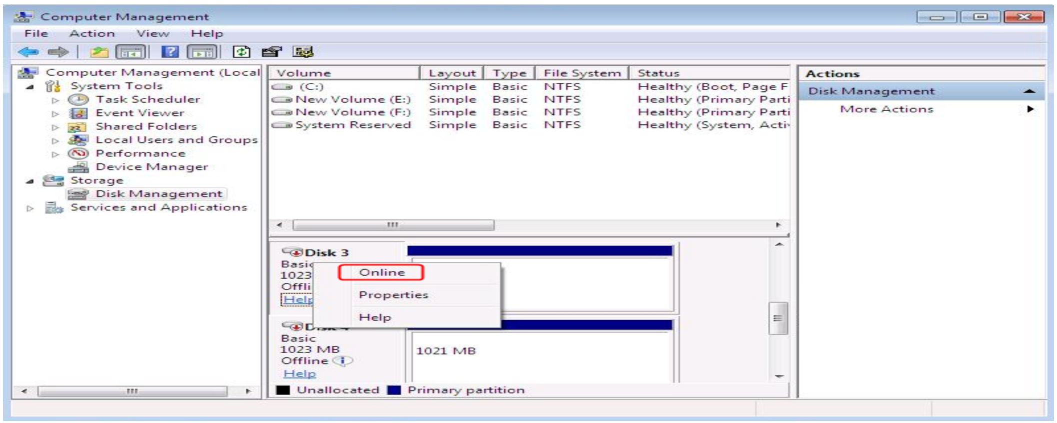

Volume	Layout	Type	File System	Status
(C:)	Simple	Basic	NTFS	Healthy (Boot, Page F
New Volume (E:)	Simple	Basic	NTFS	Healthy (Primary Parti
New Volume (F:)	Simple	Basic	NTFS	Healthy (Primary Parti
System Reserved	Simple	Basic	NTFS	Healthy (System, Acti

Disk 3
Basic
1023
Offli
Help

Basic
1023 MB
Offline
Help
1021 MB

After the data disk status changes to **Online**, the data disk will be displayed in the disk list, as shown in **Figure 13-3**.

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

Figure 13-3 Viewing online data disks

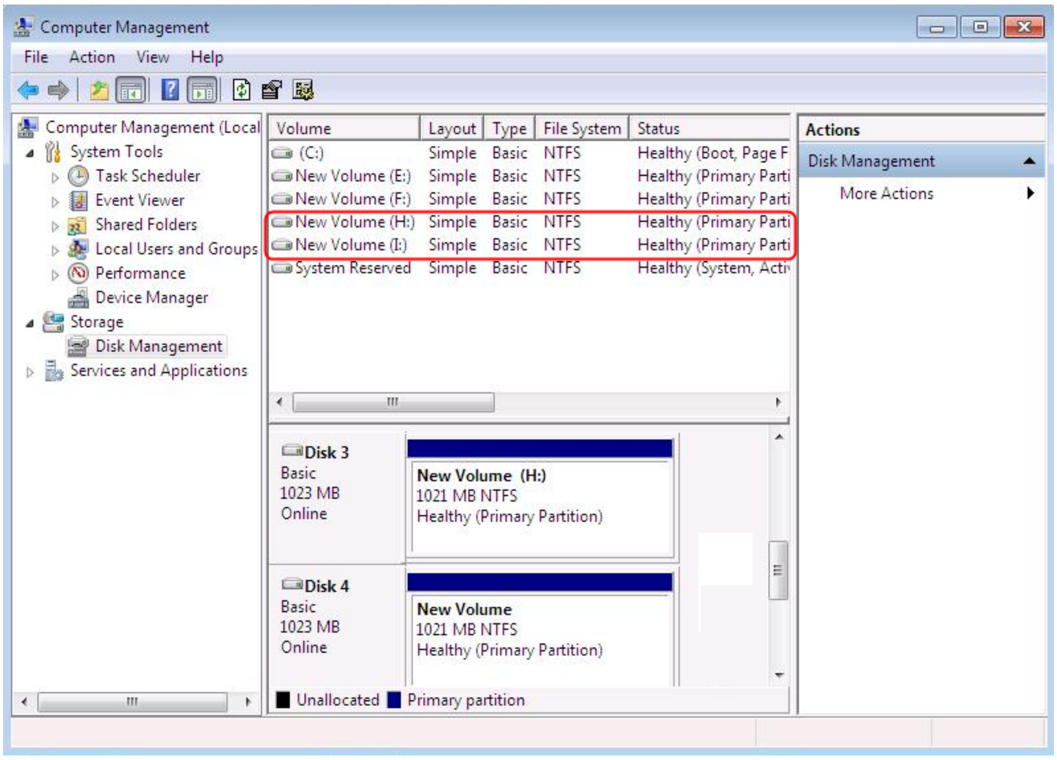

Volume	Layout	Type	File System	Status
(C:)	Simple	Basic	NTFS	Healthy (Boot, Page F
New Volume (E:)	Simple	Basic	NTFS	Healthy (Primary Parti
New Volume (F:)	Simple	Basic	NTFS	Healthy (Primary Parti
New Volume (H:)	Simple	Basic	NTFS	Healthy (Primary Parti
New Volume (I:)	Simple	Basic	NTFS	Healthy (Primary Parti
System Reserved	Simple	Basic	NTFS	Healthy (System, Acti

Disk 3
Basic
1023 MB
Online

New Volume (H:)
1021 MB NTFS
Healthy (Primary Partition)

Disk 4
Basic
1023 MB
Online

New Volume
1021 MB NTFS
Healthy (Primary Partition)

■ Unallocated ■ Primary partition

----End

13.3 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

- Cause 1: The DNS cannot resolve the OBS domain name.
- Cause 2: The OpenSSL version of the target server is too early.

Solution for Cause 1

Cause 1: 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 13-4**.

Figure 13-4 Configuring DNS

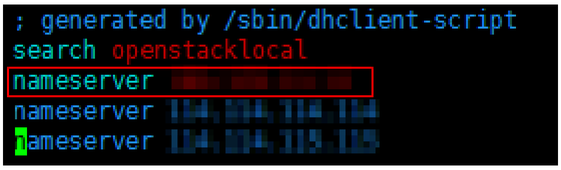

```
; generated by /sbin/dhclient-script
search openstacklocal
nameserver [redacted]
nameserver [redacted]
nameserver [redacted]
```

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 13-5**. Click **Local Area Connection**.

Figure 13-5 Page for network and sharing center

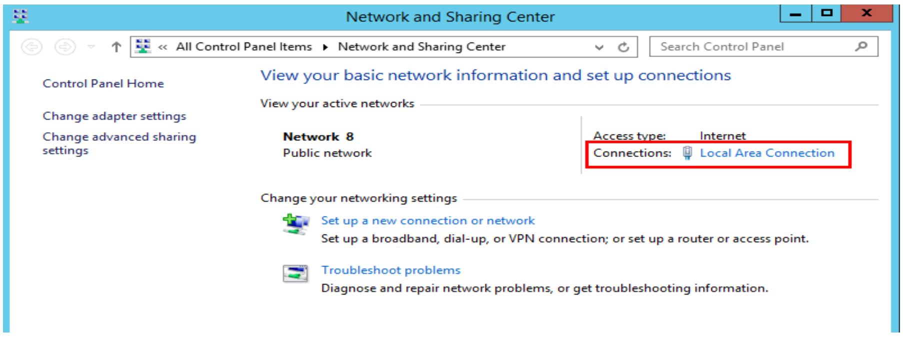

Network and Sharing Center

« All Control Panel Items ▸ Network and Sharing Center

Search Control Panel

Control Panel Home

Change adapter settings

Change advanced sharing settings

View your basic network information and set up connections

View your active networks

Network 8
Public network

Access type: Internet
Connections: Local Area Connection

Change your networking settings

Set up a new connection or network
Set up a broadband, dial-up, or VPN connection; or set up a router or access point.

Troubleshoot problems
Diagnose and repair network problems, or get troubleshooting information.

Step 4 In the **Activity** area, select **Properties**. See **Figure 13-6**.

Figure 13-6 Local area connection

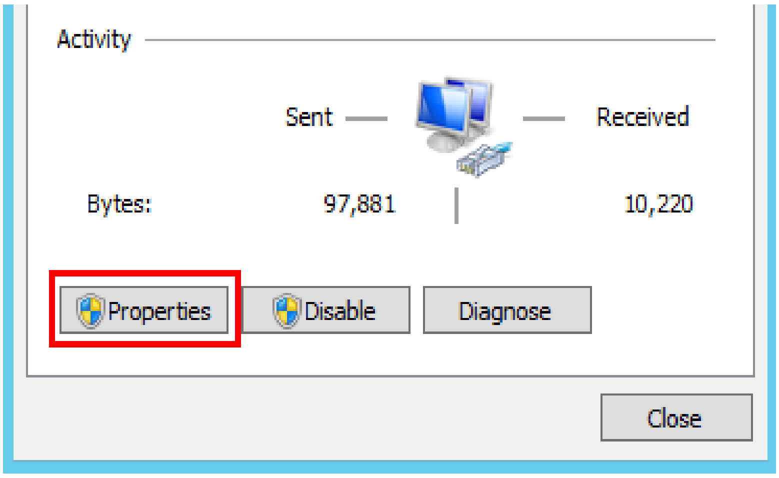

Activity

Sent — Received

Bytes: 97,881 10,220

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 13-7**.

Figure 13-7 Local area connection properties

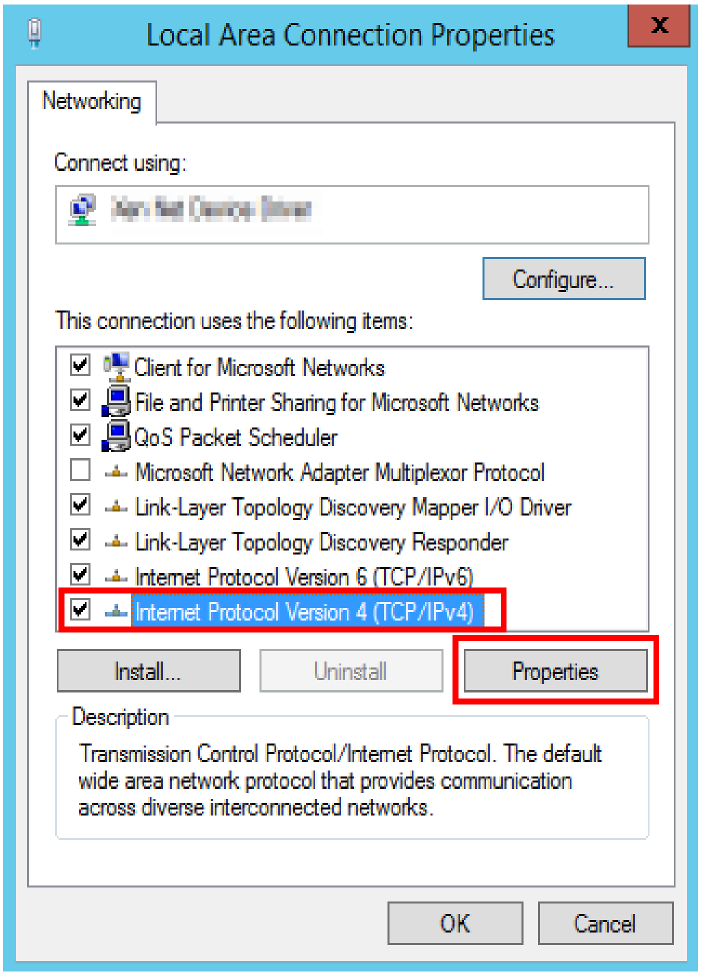

Local Area Connection Properties

Networking

Connect using:

[redacted]

This connection uses the following items:

- ☑ Client for Microsoft Networks
- ☑ File and Printer Sharing for Microsoft Networks
- ☑ QoS Packet Scheduler
- ☐ Microsoft Network Adapter Multiplexor Protocol
- ☑ Link-Layer Topology Discovery Mapper I/O Driver
- ☑ Link-Layer Topology Discovery Responder
- ☑ Internet Protocol Version 6 (TCP/IPv6)
- ☑ Internet Protocol Version 4 (TCP/IPv4)

Description

Transmission Control Protocol/Internet Protocol. The default wide area network protocol that provides communication across diverse interconnected networks.

Step 6 In the dialog box that is displayed, select **Use the following DNS server addresses:** and configure DNS, as shown in **Figure 13-8**. You need to manually change the DNS server address. Obtain the IP address from the administrator. Then click **OK**.

Figure 13-8 Configuring DNS

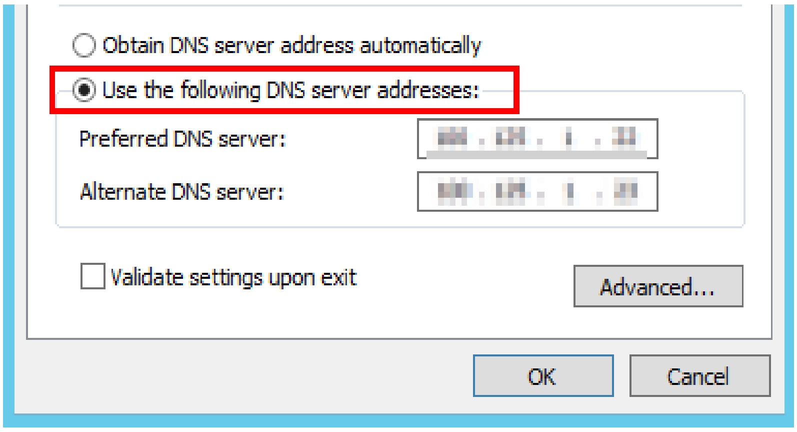

○ Obtain DNS server address automatically

● Use the following DNS server addresses:

Preferred DNS server: [redacted]

Alternate DNS server: [redacted]

□ Validate settings upon exit

----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

13.4 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 creation, the configuration parameters contained in the **/etc/fstab** file in the system disk of the new server are that of the backup source server, causing the UUID information to be inconsistent with the new data disks. As a result, the ECS encounters an error when uploading **/etc/fstab** during the bootup and enters 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 13-9 Maintenance mode of the system

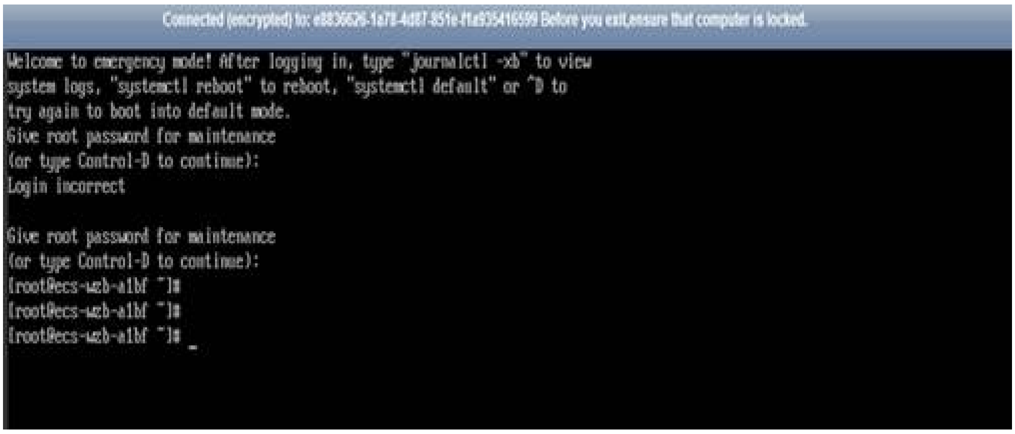

Connected (encrypted) to: e8836626-1a78-4d87-851e-f1a935416599 Before you exit,ensure that computer is locked.

```
Welcome to emergency mode! After logging in, type "journalctl -xb" to view
system logs, "systemctl reboot" to reboot, "systemctl default" or ^D to
try again to boot into default mode.
Give root password for maintenance
(or type Control-D to continue):
Login incorrect

Give root password for maintenance
(or type Control-D to continue):
[root@ecs-wzb-a1bf ~]#
[root@ecs-wzb-a1bf ~]#
[root@ecs-wzb-a1bf ~]# _
```

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

Figure 13-10 Data disk UUIDs

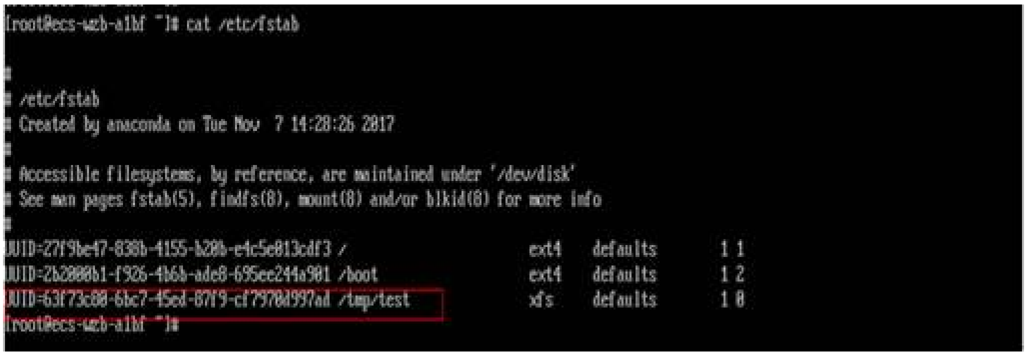

```
[root@ecs-wzb-a1bf ~]# cat /etc/fstab

#
# /etc/fstab
# Created by anaconda on Tue Nov  7 14:28:26 2017
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
UUID=27f9be47-838b-4155-b28b-e4c5e813cdf3 /                    ext4    defaults        1 1
UUID=2b2808b1-f926-4b6b-ade8-695ee244a981 /boot                ext4    defaults        1 2
UUID=63f73c08-6bc7-45ed-87f9-cf7970d997ad /tmp/test            xfs     defaults        1 0
[root@ecs-wzb-a1bf ~]#
```

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 13-11 /etc/fstab after being updated

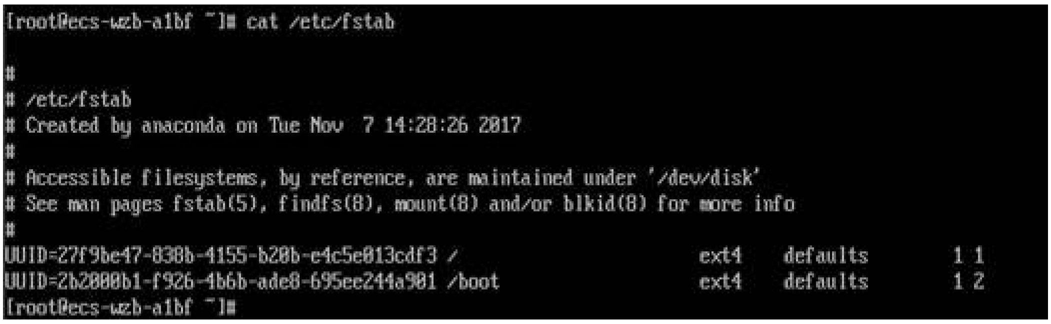

```
[root@ecs-wzb-a1bf ~]# cat /etc/fstab

#
# /etc/fstab
# Created by anaconda on Tue Nov  7 14:28:26 2017
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
UUID=27f9be47-838b-4155-b28b-e4c5e013cdf3 /                    ext4    defaults        1 1
UUID=2b2080b1-f926-4b6b-ade8-695ee244a901 /boot                ext4    defaults        1 2
[root@ecs-wzb-a1bf ~]#
```

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

Figure 13-12 Normal bootup page

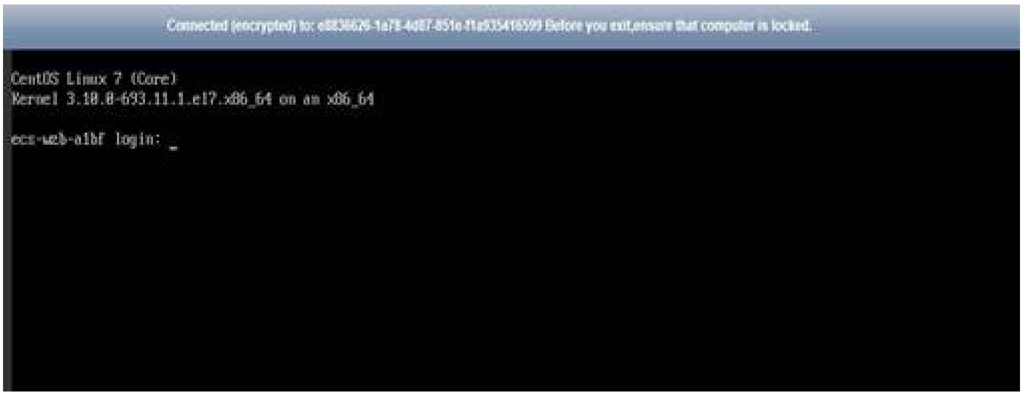

Connected (encrypted) to: e8836626-1a75-4d87-851e-f1e935416599 Before you exit,ensure that computer is locked.

```
CentOS Linux 7 (Core)
Kernel 3.10.0-693.11.1.el7.x86_64 on an x86_64

ecs-wzb-a1bf login: _
```

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

Figure 13-13 Attaching the data disks manually

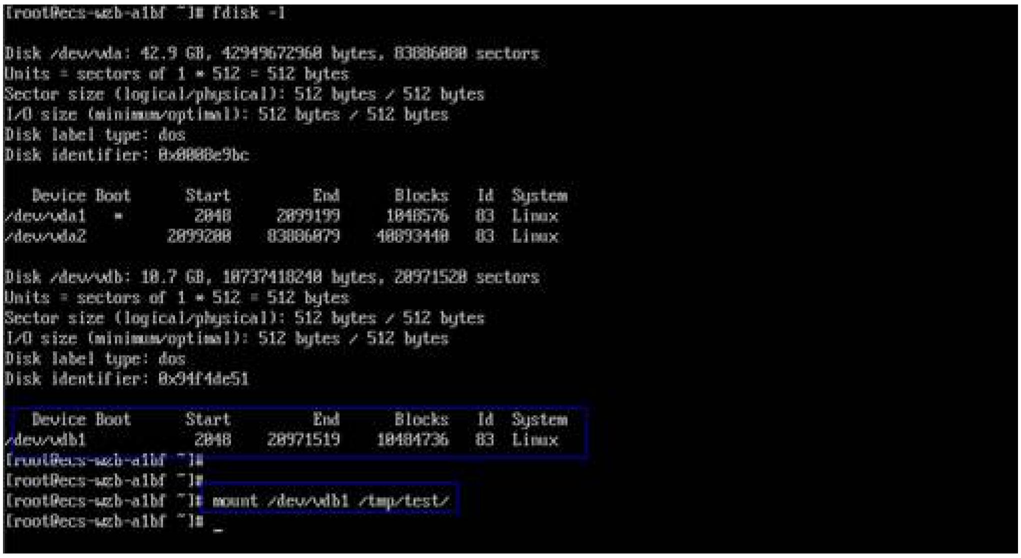

```
[root@ecs-wzb-a1bf ~]# fdisk -l

Disk /dev/vda: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x0008e9bc

   Device Boot      Start         End      Blocks   Id  System
/dev/vda1   *        2048     2099199     1048576   83  Linux
/dev/vda2         2099200    83886079    40893440   83  Linux

Disk /dev/vdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x94f4de51

   Device Boot      Start         End      Blocks   Id  System
/dev/vdb1            2048    20971519    10484736   83  Linux
[root@ecs-wzb-a1bf ~]#
[root@ecs-wzb-a1bf ~]#
[root@ecs-wzb-a1bf ~]# mount /dev/vdb1 /tmp/test/
[root@ecs-wzb-a1bf ~]# _
```

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

Figure 13-14 Obtaining UUIDs of data disks

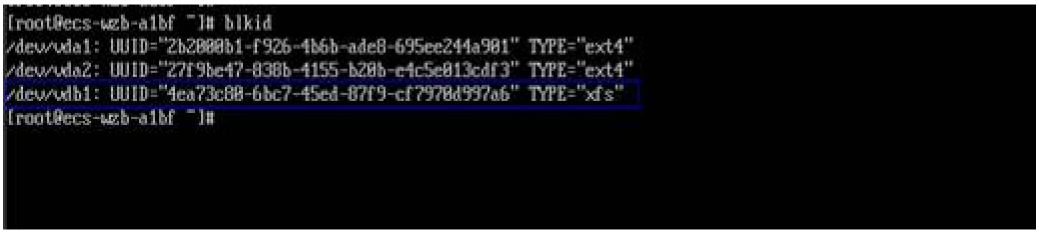

```
[root@ecs-wzb-a1bf ~]# blkid
/dev/vda1: UUID="2b2000b1-f926-4b6b-ade8-695ee244a901" TYPE="ext4"
/dev/vda2: UUID="27f9be47-830b-4155-b20b-e4c5e013cdf3" TYPE="ext4"
/dev/vdb1: UUID="4ea73c80-6bc7-45ed-87f9-cf7970d997a6" TYPE="xfs"
[root@ecs-wzb-a1bf ~]#
```

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 13-15 Adding attachment information of data disks

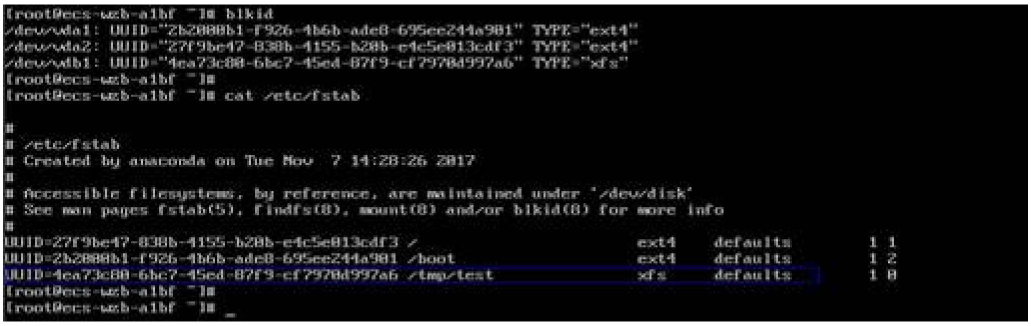

```
[root@ecs-wzb-a1bf ~]# blkid
/dev/vda1: UUID="2b2080b1-f926-4b6b-ade8-695ee244a981" TYPE="ext4"
/dev/vda2: UUID="27f9be47-838b-4155-b28b-e4c5e813cdf3" TYPE="ext4"
/dev/vdb1: UUID="4ea73c80-6bc7-45ed-87f9-cf7970d997a6" TYPE="xfs"
[root@ecs-wzb-a1bf ~]#
[root@ecs-wzb-a1bf ~]# cat /etc/fstab

#
# /etc/fstab
# Created by anaconda on Tue Nov  7 14:28:26 2017
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
UUID=27f9be47-838b-4155-b28b-e4c5e813cdf3 /                       ext4    defaults        1 1
UUID=2b2080b1-f926-4b6b-ade8-695ee244a981 /boot                   ext4    defaults        1 2
UUID=4ea73c80-6bc7-45ed-87f9-cf7970d997a6 /tmp/test               xfs     defaults        1 0
[root@ecs-wzb-a1bf ~]#
[root@ecs-wzb-a1bf ~]# _
```

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

- For O&M security purposes, you are advised to change the user **rdadmin**'s password of the Agent OS regularly and disable this user's remote login permission.
- In Linux, user **rdadmin** has no password.
- This section describes how to change the password of user **rdadmin** in Windows 2012. For other versions, change the password according to actual situation.

Prerequisites

- You have obtained a username and its password for logging in to the management console.
- 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 Click **User Accounts**. The **User Account Control** dialog box is displayed. 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

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

Context

This section introduces the procedures in Windows and Linux.

NOTICE

If the authentication password and data encryption password for SNMP v3 of the Agent are the same, security risks exist. To ensure system security, you are advised to set different passwords for authentication and data encryption.

Obtain the initial authentication password.

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.

Procedure (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. Type the login password of the Agent 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 authorization 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

Procedure (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 the preceding command is executed, the system remains running even when no operation is performed, which results in security risks. For security purposes, run the **exit** command to exit the system after you finish performing operations.

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

Step 4 Run the **/home/rdadmin/Agent/bin/agentcli chgsnmp** command. Type the login password of the Agent 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 authorization 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

- You have obtained a username and its password for logging in to the management console.
- 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.

Procedure (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 the preceding command is executed, the system remains running even when no operation is performed, which results in security risks. For security purposes, run the **exit** command to exit the system after you finish performing 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

Procedure (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.

Procedure (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

Procedure (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
2022-08-16	This issue is the first official release.