

Storage Disaster Recovery Service

Best Practices

Issue 03
Date 2024-12-19



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

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

Trademarks and Permissions



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

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

Notice

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

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

Security Declaration

Vulnerability

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

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

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

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

Contents

1 Two-Site Three-Center Disaster Recovery (SDRS+CBR)	1
1.1 Solution Overview.....	1
1.2 Solution Principle.....	3
1.3 Resource and Cost Planning.....	6
1.4 Cross-AZ DR Implementation.....	7
1.4.1 Configuring Cross-AZ DR.....	8
1.4.2 Disaster Recovery Drill (Synchronous Replication).....	13
1.4.3 Restoring Services If the Production Site Fails.....	15
1.5 Cross-Region Backup Implementation.....	16
1.5.1 Configuring Cross-Region Backup.....	16
1.5.2 Verifying the Configuration.....	26
1.5.3 Restoring Services When Both the Production Site and Its Cross-AZ DR Site Become Faulty.....	26

1 Two-Site Three-Center Disaster Recovery (SDRS+CBR)

1.1 Solution Overview

What Is Two-Site Three-Center?

Two-site three-center is a disaster recovery (DR) solution used for service continuity. With three data center coexisted, the continuity of mission-critical services can be guaranteed even if any two centers fail. This significantly improves the availability of the solution.

The two-site three-center DR solution includes a production center, an intra-city DR center, and a remote DR center.

- Production center: It provides services for external applications.
- Intra-city DR center: It is usually dozens of kilometers away from the production center. Applications can be switched to the intra-city DR center without data loss. This is the level-1 DR protection of this solution.
- Remote DR center: It is usually built hundreds or thousands of kilometers away from the production center to protect against regional disasters. Data is periodically, asynchronously replicated to the remote DR center. This is the level-2 DR protection of this solution.

You can use Huawei Cloud SDRS and CBR to deploy two-site three-center DR. This solution responds quickly to small-scale failures and large-scale natural disasters, protecting service data and guaranteeing service continuity.

What Is SDRS?

Storage Disaster Recovery Service (SDRS) provides DR services for many cloud services, including Elastic Cloud Server (ECS), Elastic Volume Service (EVS), and Dedicated Distributed Storage Service (DSS). By leveraging technologies, such as storage replication, data redundancy, and cache acceleration, SDRS offers high data reliability and service continuity for users.

What Is CBR?

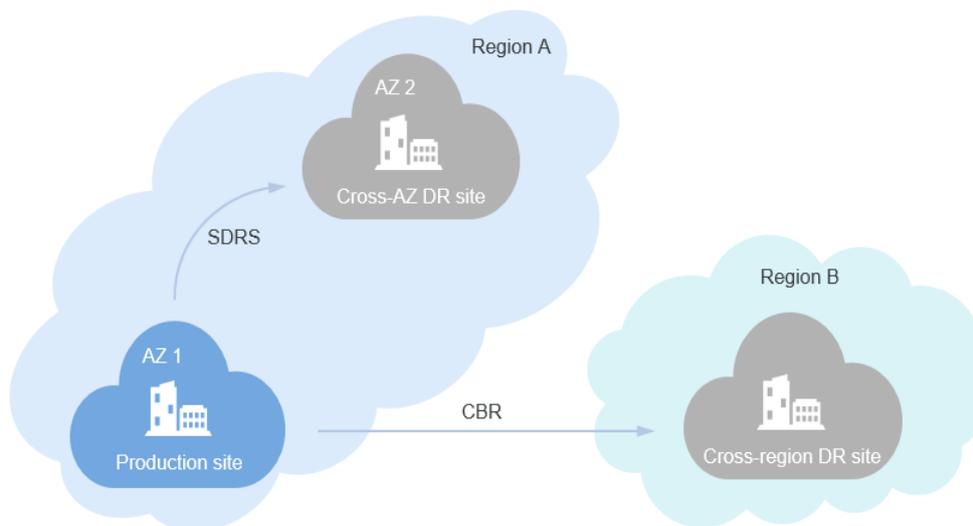
Cloud Backup and Recovery (**CBR**) enables you to back up ECSs, Hyper Elastic Cloud Servers (HECSs), Bare Metal Servers (BMSs), EVS disks, SFS Turbo file systems, a single or multiple files and databases on your local hosts, and on-premises VMware virtual environments with ease. In case of a virus attack, accidental deletion, or software or hardware fault, you can restore data to any point in the past when the data was backed up.

What Is the Two-Site Three-Center DR Solution (SDRS+CBR)?

In this solution, SDRS synchronizes data from the production site to a DR site in a different AZ of the same region for cross-AZ DR. CBR periodically backs up the production data to the other DR site in a different region for cross-region DR. The production site, cross-AZ DR site, and cross-region DR site all together form the geo-redundant architecture to meet the two-site three-center DR requirements.

The cross-AZ DR site corresponds to the intra-city DR center, and the cross-region DR site corresponds to the remote DR center.

In this document, the two-site three-center DR solution (SDRS+CBR) is called the SDRS+CBR DR solution for short.



In the SDRS+CBR DR solution, the solution is deployed in two regions and involves three sites: a production site, a cross-AZ DR site in the same region (cross-AZ DR center for short in the following document), and a cross-region DR site.

In this solution, SDRS synchronizes data from the production site to the cross-AZ DR site. CBR periodically backs up the production data to the cross-region DR site.

Application Scenarios

The SDRS+CBR DR solution is suitable if you use ECSs and EVS disks to run workloads.

Solution Advantages

Compared with two-site solutions (the cross-AZ DR solution and cross-region DR solution), the SDRS+CBR DR solution combines the advantages of both two-site solutions and can protect against both AZ-level failures and regional disasters.

- Disk data consistency is ensured during DR switchover if an AZ-level failure occurs.
- If a regional disaster occurs, this solution can restore the service data to the latest backup time point to ensure minimal data loss.

Table 1-1 compares the three types of DR solutions.

Table 1-1 Solution comparison

DR Solution	Handle AZ-Level Failures	Perform AZ-Level Failover to Restore System to the State Before It Broke Down	Handle Regional Disasters
Cross-AZ DR (SDRS)	√	√	×
Cross-region DR (CBR)	√	×	√
Two-site three-center DR (SDRS+CBR)	√	√	√

Notes and Constraints

- Cloud-Init must be installed on Linux ECSs, and Cloudbase-Init on Windows ECSs. After you perform a failover for the first time and start the servers, the system will run Cloud-Init/Cloudbase-Init to inject the initial data. The passwords or key pairs for logging in to the production and DR servers will change.
- If you use SDRS to set up DR, DR servers or disks can be restored only when protection is disabled for the servers.

For details about SDRS usage restrictions, see [SDRS Usage Restrictions](#).

For details about CBR usage restrictions, see [CBR Usage Restrictions](#).

1.2 Solution Principle

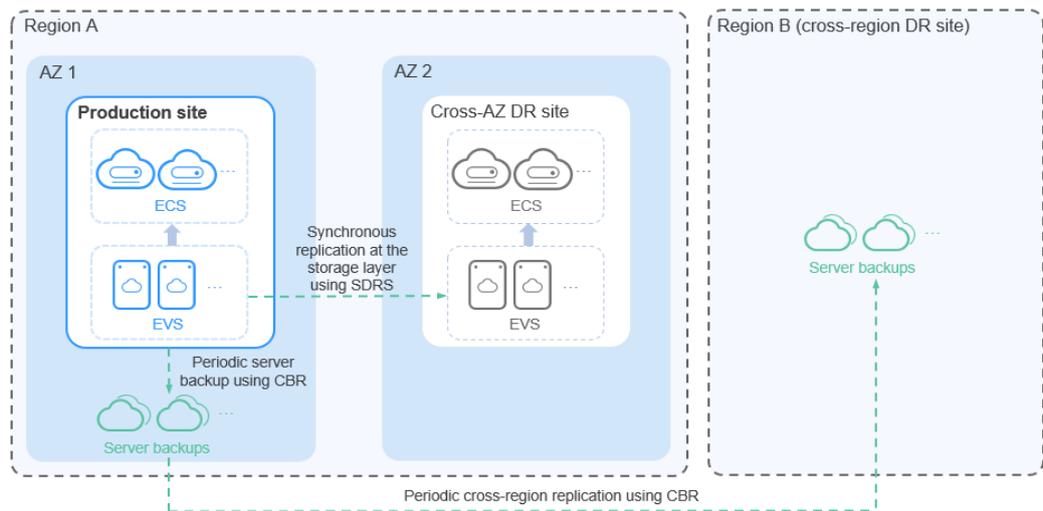
This section describes how the solution works in three fault scenarios: when the production site is functional; when the production site becomes faulty; when both the production site and cross-AZ DR site become faulty.

When the Production Site Is Functional

Figure 1-1 shows the work mechanism when the production site is functional.

- SDRS synchronizes the server data and configuration of the production site in AZ 1 in region A to the cross-AZ DR site in AZ 2 in region A. You can perform routine DR drills to periodically simulate fault recovery scenarios and formulate emergency recovery plans.
- CBR periodically backs up the entire server at the production site in region A and replicate the backup to the cross-region DR site in region B.

Figure 1-1 When the production site is functional



When the Production Site Becomes Faulty

If the production site becomes unavailable due to a small-scale fault such as a device fault, applications can be switched to the cross-AZ DR site without data loss.

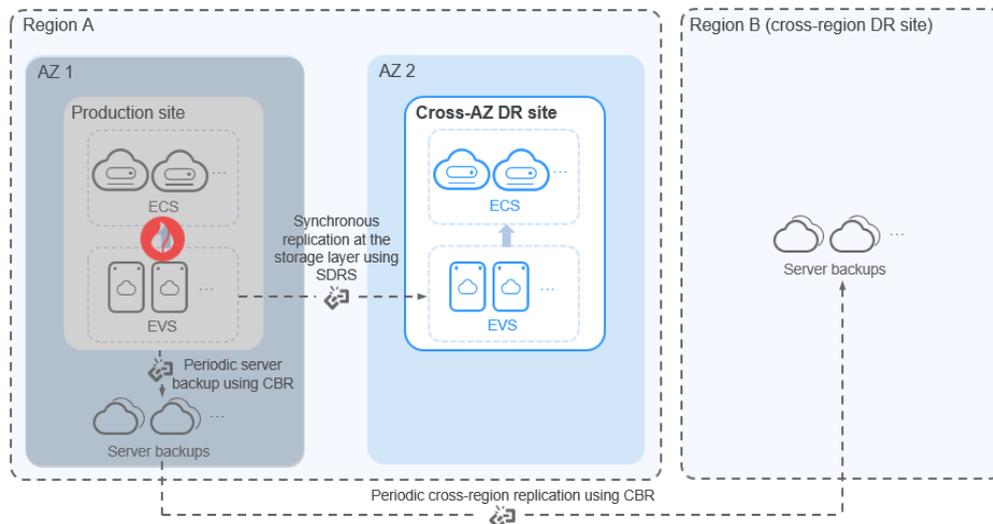
In this phase, the Recovery Point Objective (RPO) is 0, and the Recovery Time Objective (RTO) is within 30 minutes.

NOTE

RPO specifies the maximum acceptable period in which data might be lost.

RTO specifies the maximum acceptable amount of time for restoring the entire system after a disaster happens.

Figure 1-2 When the production site becomes faulty



When Both the Production Site and Cross-AZ DR Site Become Faulty

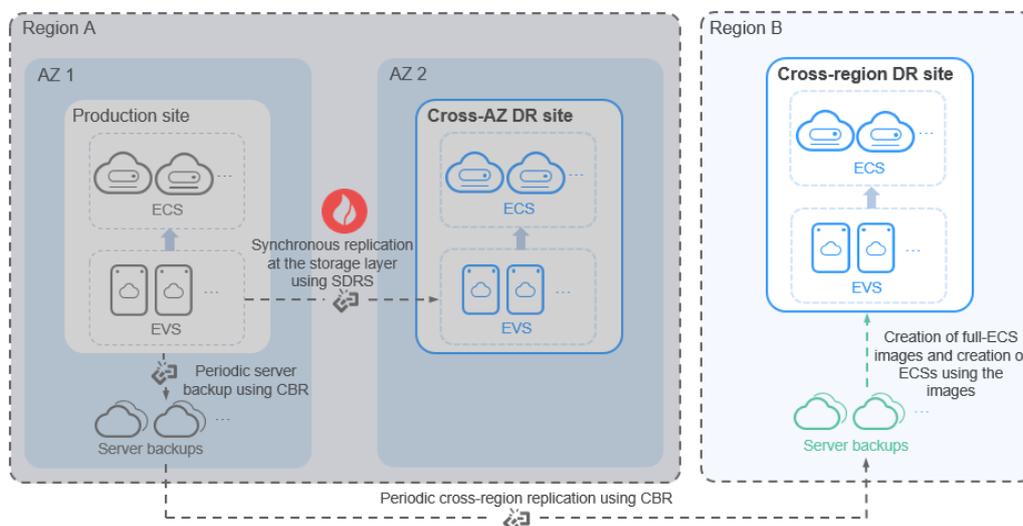
If the production site and cross-AZ DR site become unavailable due to a large-scale disaster such as a natural disaster, applications can be switched to the cross-region DR site. You can create full-ECS images using the server backups periodically replicated to region B, use the full-ECS images to create ECSs, and restore applications at the cross-region DR site to ensure service continuity.

In this phase, the RPO ranges from 0 to the backup interval. The minimum backup interval is 1 hour, and the RTO is within 30 minutes.

NOTE

In the cross-region DR phase, the RPO equals the difference between the time when a disaster occurs and the time when the latest backup file is generated.

Figure 1-3 When both the production site and cross-AZ DR site become faulty



1.3 Resource and Cost Planning

This section describes the resource planning in this best practice.

Cross-AZ DR Deployment

Table 1-2 Resource planning for cross-AZ DR deployment

Region	Resource	Resource Name	Resource Description	Quantity	Monthly Price (USD)
AP-Bangkok (production site)	ECS	ecs-test	Flavor: s6.small.1 AZ: az-01 You can select a maximum of five production site servers at a time.	1	43.40
	VPC	vpc1	CIDR block: 10.0.0.0/16	1	00.00
AP-Bangkok (DR site)	ECS	ecs-test2	Flavor: s6.small.1 AZ: az-02	1	43.40
	VPC	vpc1	CIDR block: 10.0.0.0/16 (Optional) Extension NIC: subnet-01 (192.168.0.0/24)	1	00.00

Cross-Region DR Deployment

Table 1-3 Resource planning for cross-region DR deployment

Region	Resource	Resource Name	Resource Description	Quantity	Monthly Price (USD)
AP-Bangkok (production site)	ECS	ecs-test	Flavor: s6.small.1 AZ: az-01 You can select a maximum of five production site servers at a time.	1	43.40
	VPC	vpc-demo	CIDR block: 10.0.0.0/16	1	00.00

Region	Resource	Resource Name	Resource Description	Quantity	Monthly Price (USD)
	CBR	backup_policy	Backup policy: Automatically performs weekly backup at 00:00 on the following days: Monday, Wednesday, Friday, Sunday	1	00.00
		replication_0001	Replication policy: Automatically replicate backups weekly at 03:00 on the following days: Monday, Wednesday, Friday, Sunday. The destination region is AP-Singapore , and the destination vault is vault_replication .	1	00.00
		vault-f61e	Backup vault: The vault capacity is 500 GB. The ecs-test ECS is associated with the vault, and the backup_policy backup policy and replication_0001 replication policy are applied to the vault.	1	18.00
AP-Singapore (DR site)	VPC	vpc-demo	CIDR block: 10.0.0.0/16	1	00.00
	CBR	vault_replication	Replication vault: 500 GB	1	14.00

NOTICE

The price provided is for reference only. See the actual price on the Huawei Cloud console. For more information, see [Pricing Details](#).

1.4 Cross-AZ DR Implementation

1.4.1 Configuring Cross-AZ DR

Scenarios

In the event of a production site failure caused by force majeure (such as fire) or device faults (software or hardware faults), if the faults cannot be rectified within a short period of time, you can use SDRS to set up cross-AZ DR protection for your ECSs with zero RPO.

SDRS uses synchronous replication at the storage layer to provide cross-AZ DR protection, and thereby meeting the data consistency requirements. If the production site fails, you can quickly restore services at the cross-AZ DR site with a few clicks.

The procedure of setting up cross-AZ DR with SDRS is as follows:

1. [Creating a Protection Group](#)
2. [Creating a Protected Instance](#)
3. [Enabling Protection](#)

Notes and Constraints

- If a production site server has been added to an ECS group, you cannot specify a DeH to create the DR site server.
- After a protected instance is created, the DR site server and production site server have the same name by default, but their IDs are different.
- To change a server name, click the name of the target server on the protected instance details page to go to the server details page and then change the name.
- After a protected instance is created, and protection is enabled for production site servers, any modifications made to **Hostname, Name, Agency, ECS Group, Auto Recovery, Security Group, and Tags** of production site servers will not be synchronized to DR site servers. Log in to the console and manually make the same modifications on DR site servers.
- If protection is enabled for servers that are added by scaling actions, when the auto scaling group scales in, such servers cannot be deleted.
- If the production site servers run Windows and use key pairs for login, ensure that the key pairs exist when you create the protected instance. Or, DR site servers may fail to create, resulting in the protected instance creation failure.

NOTE

- If a key pair used by a production site server has been deleted, re-create the key pair and use the same name.
- If the production site servers run Linux and use key pairs for login, the key pair information will not be displayed on the details pages of the created DR site servers, but login using the production server key pairs is not affected.
- If the production site servers of a protection group are added to an enterprise project, the DR site servers created will not be automatically added to the enterprise project. Manually add them to the project as needed.

Creating a Protection Group

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select the region and project of the production site.

Step 3 Click **Service List** and choose **Storage > Storage Disaster Recovery Service**.

The **Storage Disaster Recovery Service** page is displayed.

Step 4 Click **Create Protection Group**.

The **Create Protection Group** page is displayed.

Step 5 Set the protection group parameters by referring to [Table 1-4](#).

 **NOTE**

All the parameters in [Table 1-4](#) are mandatory.

Table 1-4 Protection group parameters

Parameter	Description	Example Value
Region	Geographic area where servers reside If the region is not what you desired, click the drop-down list to change it.	AP-Bangkok
DR Direction	<ul style="list-style-type: none">• Production Site: Select the AZ where production site servers reside.• DR Site: Select the AZ where DR site servers reside.	Production Site: az-01 DR Site: az-02
Deployment Model	Only the VPC migration model is supported currently. All resources at the production and DR sites belong to the same VPC.	VPC migration
VPC	VPC where the protection group belongs	vpc-test
Protection Group Name	Enter the protection group name, which is used for group classification and search.	protection_group_001

Step 6 Click **Create Now**.

Step 7 Click **Back to Protection Group List** to go back to the SDRS homepage and view the protection group status.

If the protection group is displayed and its status is **Available** () , the protection group has been created successfully.

----End

Creating a Protected Instance

Step 1 In the pane of the protection group where protected instances are to be added, click **Protected Instances**.

The protection group details page is displayed.

Step 2 On the **Protected Instances** tab, click **Create Protected Instance**.

The **Create Protected Instance** page is displayed.

Step 3 Set the protected instance parameters by referring to [Table 1-5](#).

Table 1-5 Protected instance parameters

Parameter	Description	Example Value
Protection Group Name	Name of the protection group where you want to create protected instances. You do not need to configure it.	protection_group_001
Protection Group ID	ID of the protection group	2a663c5c-4774-4775-a321-562a1ea163e3
DR Direction	DR direction of the protection group. You do not need to configure it.	-
Production Site	AZ where production site servers reside. You do not need to configure it.	az-01
Deployment Model	Deployment model of the protection group where protected instances will be created. You do not need to configure it.	VPC migration
VPC	VPC of the protection group where protected instances will be created. You do not need to configure it.	vpc1
Production Site Server	<p>Mandatory</p> <p>In the list, select the production site servers you want to protect and specify specifications for corresponding DR site servers.</p> <ul style="list-style-type: none">You can select a maximum of five production site servers at a time.If you select a server that uses a shared disk, ensure that all the servers that this shared disk attached to are selected. <p>NOTE</p> <ul style="list-style-type: none">If Server Type of the protection group is ECS, select specifications for the DR site servers from the DR Site Server Specifications drop-down list. A production site server and its DR site server can have different specifications.	ecs-test > s3.small.1

Parameter	Description	Example Value
DR Site Server	<p>Mandatory when Server Type of the protection group is ECS</p> <p>You can choose to deploy DR site servers on ECSs or DeHs.</p> <p>DeHs are physical hosts dedicated for a specified user. You can create ECSs on a DeH to enhance isolation, security, and performance of your ECSs.</p> <p>NOTE If a production site server has been added to an ECS group, you cannot specify a DeH to create the DR site server.</p>	ECS
DeH	<p>Mandatory when DR Site Server is set to DeH.</p> <p>Select a DeH for deploying the DR site server. If multiple production site servers are selected, corresponding DR site servers will be created on the same DeH.</p>	deh-01
DR Site VPC	<p>VPC where the DR site servers reside</p> <p>DR site VPC is the same as VPC. You do not need to configure it.</p>	vpc1
DR Site Primary NIC	<p>Optional</p> <p>Primary NIC on the DR site server</p> <p>You can use the primary NIC automatically allocated by the system or specify a primary NIC based on your network plan.</p> <p>After you select a NIC, select an available subnet from the drop-down list and configure the private IP address.</p> <p>NOTE You cannot specify the IP addresses if you create multiple protected instances at a time.</p>	subnet-01 (192.168.0.0/24)
DR Site Disk	<p>Mandatory</p> <p>The following two options are available:</p> <ul style="list-style-type: none">• EVS• DSS If you select DSS for DR Site Disk, Storage Pool is mandatory.	EVS
Storage Pool	<ul style="list-style-type: none">• If you select EVS for DR Site Disk, Storage Pool is not required.• If you select DSS for DR Site Disk, Storage Pool is mandatory.	dss-01

Parameter	Description	Example Value
Protected Instance Name	Mandatory Enter a protected instance name, which is used for instance classification and search.	Protected-Instance-test

 NOTE

DR Site Disk and **Storage Pool** are available only when **DSS** is selected.

Step 4 Click **Create Now**.

Step 5 On the **Confirm** page, confirm the protected instance information.

- If you do not need to modify the information, click **Submit**.
- If you need to modify the configuration, click **Previous**.

Step 6 Click **Back to Protection Group Details Page** and view the protected instances of the protection group.

If the protected instance status changes to **Available** or **Protecting**, the protected instance has been created successfully.

 NOTE

After a protected instance is created, the system automatically creates replication pairs for all the disks of the protected instance and backs up the disks.

How to query replication pairs:

1. Go the protection group details page.
2. Click the **Replication Pairs** tab.

On this tab page, you can view the statuses of the replication pairs, corresponding protected instances, and production site disks.

----End

Enabling Protection

Step 1 In the pane of the desired protection group, click **Enable Protection**.

Step 2 In the displayed dialog box, confirm the protection group information and click **Yes**.

After protection is enabled, data synchronization starts.

 NOTE

The amount of time required to complete the synchronization depends on the disk capacity. Larger disk capacity requires longer time.

----End

1.4.2 Disaster Recovery Drill (Synchronous Replication)

Scenarios

Disaster recovery drills are used to simulate fault scenarios, formulate recovery plans, and verify whether the plans are applicable and effective. Services are not affected during disaster recovery drills. When a fault occurs, you can use the plans to quickly recover services, thus improving service continuity.

SDRS allows you to run disaster recovery drills in isolated VPCs (different from the disaster recovery site VPC). During a disaster recovery drill, drill servers can be quickly created based on the disk snapshot data. This way, drill servers will have the same server specifications and disk types as the production site servers.

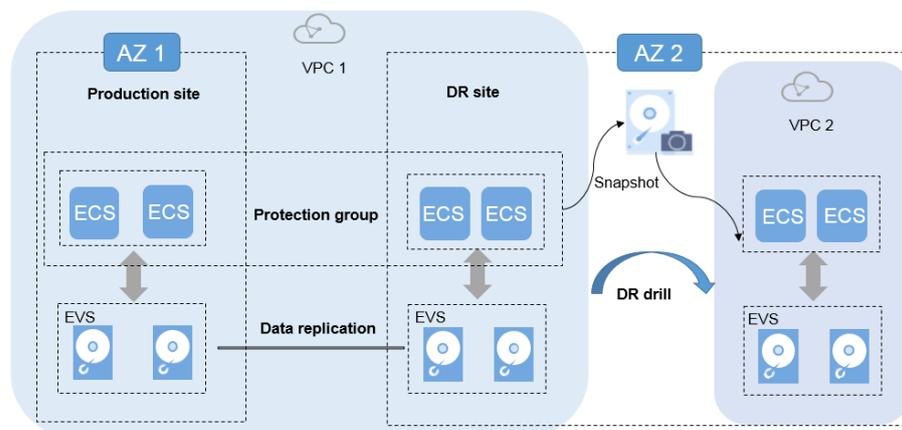
NOTE

After drill servers are created, production site servers and drill servers will independently run at the same time, and data will not be synchronized between these servers.

To guarantee that services can be switched to the disaster recovery site when an outage occurs, it is recommended that you run disaster recovery drills regularly to check that:

- Data between the production site and disaster recovery site is consistent at the moment you create a disaster recovery drill.
- Services run properly at the disaster recovery site after a planned failover.

Figure 1-4 Disaster recovery drill



Precautions

- If the disaster recovery site servers of a protection group are added to an enterprise project, the drill servers created will not be automatically added to the enterprise project. Manually add them to the project as needed.
- If an existing drill VPC is used for a new drill, the subnet ACL rule of the drill VPC will be different from that of the protection group VPC. Manually set them to be the same as needed.
- If a custom route table is configured and associated with a subnet in the protection group VPC, the corresponding route table will not be automatically created in the drill VPC. Manually create one as needed.

- If the disaster recovery site servers run Windows and use key pairs for login, ensure that the key pairs exist when you create the drill. Otherwise, drill servers may fail to create, resulting in the drill creation failure.

 **NOTE**

If a key pair has been deleted, recreate the key pair with the same name.

- If the disaster recovery site servers run Linux and use key pairs for login, the key pair information will not be displayed on the server details page, but login using the key pairs is not affected.
- After a disaster recovery drill is created and before it is executed, modifications made to **Hostname, Name, Agency, ECS Group, Security Group, Tags**, and **Auto Recovery** of disaster recovery site servers will not synchronize to drill servers. Log in to the console and manually make the modifications for the drill servers.
- If the synchronization progress of replication pairs in the protection group is not all 100%, the created drill servers may fail to start. It is recommended that you run disaster recovery drills after all replication pairs are synchronized.

Prerequisites

- The protection group is in the **Available, Protecting, Failover complete, Enabling protection failed, Disabling protection failed, Planned failover failed, Re-enabling protection failed, or Failover failed** state.
- Do not run disaster recovery drills before the first time data synchronization between the production site servers and disaster recovery site servers completes. Otherwise, drill servers may not start properly.

Procedure

Step 1 Log in to the management console.

Step 2 Click Service List and choose **Storage > Storage Disaster Recovery Service**.

The **Storage Disaster Recovery Service** page is displayed.

Step 3 In the pane of the protection group to which a DR drill is to be added, click **DR Drills**.

The protection group details page is displayed.

Step 4 On the **DR Drills** tab, click **Create DR Drill**.

The **Create DR Drill** dialog box is displayed.

Step 5 Configure **Name** and **Drill VPC**.

Table 1-6 Parameter description

Parameter	Description	Example Value
Name	DR drill name	DR drill servername

Parameter	Description	Example Value
Drill VPC	<p>VPC that used for a DR drill. It cannot be the same as the VPC of the DR site server. The value can be Automatically create or Use existing.</p> <ul style="list-style-type: none">• Automatically create: The system automatically creates a drill VPC and subnet for the protection group.• Use existing: The system uses an existing VPC as the drill VPC. If you select to use an existing VPC, the subnet CIDR block of the drill VPC must be consistent with that of the production group VPC. <p>NOTE The drill VPC cannot be the same as the VPC of the protection group.</p>	vpc-f9f7

Step 6 Click **OK**.

After the disaster recovery drill is created, you can log in to a drill server and check whether services are running properly.

----End

1.4.3 Restoring Services If the Production Site Fails

Scenarios

If the production site servers and disks fail due to physical faults, you can fail over services to the DR site servers and disks to guarantee service continuity.

DR site servers and disks can be used immediately after a failover. You can power on the servers, or use CBR to restore data to a specified data restore point.

 **NOTE**

- After a failover is started, data synchronization stops.
- After the failover is complete, the status of the protection group changes to **Failover complete**. Go to the protected instance details page and start the DR site server.

Constraints and Limitations

- For Linux servers with Cloud-Init installed, if you have changed **hostname** of the production site server before you perform a failover for the first time, this modification will not synchronize to the DR site server.
To resolve this problem, see [What Can I Do If hostname of the Production Site Server and DR Site Server Are Different After a Planned Failover or Failover?](#)
- The protection group must contain replication pairs.
- Protection is enabled for the protection group, and the protection group is in the **Protecting**, **Failover failed**, or **Planned failover failed** state.

Restoring Services

Step 1 Log in to the management console.

Step 2 Click **Service List** and choose **Storage > Storage Disaster Recovery Service**.

The **Storage Disaster Recovery Service** page is displayed.

Step 3 In the pane of the desired protection group, click **Protected Instances**.

The protection group details page is displayed.

Step 4 In the upper right corner of the page, click **More** and choose **Fail Over** from the drop-down list.

The **Fail Over** dialog box is displayed.

Step 5 Click **Fail Over**.

During the failover, do not start or stop the servers in the protection group. Or, the failover may fail.

----End

1.5 Cross-Region Backup Implementation

1.5.1 Configuring Cross-Region Backup

Scenarios

If both the production site and the cross-AZ DR site become unavailable due to a regional disaster such as an earthquake, you can use CBR for cross-region DR.

Specifically, CBR periodically backs up the servers and disks associated with a vault based on a backup policy, and then replicates the backups to the cross-region DR site based on a replication policy. If both the production site and the cross-AZ DR site fail, you can restore services at the cross-region DR site with only a few clicks.

The procedure of setting up cross-region DR with CBR is as follows:

1. [Creating a Backup Policy](#)
2. [Purchasing a Server Backup Vault and Applying the Backup Policy](#)

3. [Creating a Replication Policy](#)
4. [Purchasing a Server Replication Vault and Applying the Replication Policy](#)

Constraints and Limitations

- Each user can create a maximum of 32 backup policies and 32 replication policies.
- A maximum of 64 ECSs can be associated with a backup policy.
- Automatic backup jobs require enabling the backup policy. The system automatically backs up ECSs associated with the backup policy and deletes expired backups.

Prerequisites

- Confirm that the following operations have been performed before you back up an ECS:
 - The Linux ECS has been optimized (see [Optimizing a Linux Private Image](#)), and Cloud-Init has been installed (see [Installing Cloud-Init](#)).
 - The Windows ECS has been optimized (see [Optimizing a Windows Private Image](#)), and Cloudbase-Init has been installed (see [Installing Cloudbase-Init](#)).

Creating a Backup Policy

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select the region and project of the production site.

Step 3 Click **Service List** and choose **Storage > Cloud Backup and Recovery > Policies**.

Step 4 Click **Create Policy** in the upper right corner. On the displayed page, click **Backup policy** to create a custom backup policy. See [Figure 1-5](#).

Figure 1-5 Creating a backup policy

The screenshot shows a 'Create Policy' dialog box with the following configuration:

- Type:** Backup policy (selected), Replication policy
- Name:** backup_policy
- Status:** Enabled (toggle switch)
- Execution Time:** 02:00 (selected from a grid of hourly slots)
- Backup Cycle:** Custom cycle, Every 1 days
- Retention Rule:** Time period, 6 months

Step 5 Set the backup policy parameters [Table 1-7](#) describes the parameters.

Table 1-7 Backup policy parameter description

Parameter	Description	Example Value
Type	Type of the policy. In this section, we select the 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

Parameter	Description	Example Value
Status	Whether to enable the backup policy.	Only after a backup policy is enabled and applied will CBR automatically back up the vault resources and deletes expired backups.
Backup Frequency	<p>Select a backup frequency.</p> <ul style="list-style-type: none"> Weekly Defines on which days of each week the backup task will be executed. You can select multiple days. Day based Defines the interval (every 1 to 30 days) for executing the backup task. 	<p>Every day</p> <p>If you select day-based backup, the first backup is supposed to be executed on the day when the backup policy is created. If the execution time on the day you create the backup policy has passed, the first backup will be executed in the next backup cycle.</p> <p>It is recommended that backups be performed during off-peak hours or when no services are running.</p>
Execution Time	<p>Execution times of the backup policy in a day</p> <p>Backups can be scheduled at the beginning of each hour, and you can select multiple hours.</p> <p>NOTICE</p> <p>If a large amount of data needs to be backed up, you are advised to make backup less frequent to prevent the system from skipping any execution time.</p> <p>For example, a disk is scheduled to be backed up at 00:00, 01:00, and 02:00. A backup task starts at 00:00. Because a large amount of incremental data needs to be backed up or a heap of backup tasks are executed at the same time, this backup task takes 90 minutes and completes at 01:30. The system performs the next backup at 02:00. In this case, only two backups are generated in total, one at 00:00, and the other at 02:00.</p>	<p>00:00, 02:00</p> <p>It is recommended that backups be performed during off-peak hours or when no services are running.</p>

Parameter	Description	Example Value
Retention Rule	<p>Defines how backups will be retained</p> <ul style="list-style-type: none"> ● Time period You can choose to retain backups for one month, three months, six months, one year, or for any desired number (2 to 99999) of days. ● Backup quantity You can set the maximum number of 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 will be both applied. <ul style="list-style-type: none"> - Day-based: 0-100 - Weekly: 0-100 - Monthly: 0-100 - Yearly: 0-100 <p>A resource may be backed up multiple times in a day. If day-based backup is configured, only the most recent backup of that day is retained. If you set Day-based to 5, the most recent backup of each of the last five days that have backups generated will be retained and the earliest backups will be deleted automatically. If day-based, weekly, monthly, and yearly retention rules are all configured, all the rules will apply and the union set of backup replicas will be retained. For example, if Day-based is set to 5 and Weekly to 1, five backup replicas will be retained. The long-term retention rule and the quantity-based retention rule both apply.</p> <ul style="list-style-type: none"> ● Permanent 	6 months

Parameter	Description	Example Value
	<p>NOTE</p> <ul style="list-style-type: none"> - The system automatically deletes the earliest and expired backups every other day to avoid exceeding the maximum number of backups to retain or retaining any backup longer than the maximum retention period. - The retention rules apply only to auto-generated backups, but not manual backups. Manual backups need to be deleted manually. - After a backup is used to create an image, the backup will not be deleted by the retention rule. - A maximum of 10 backups are retained for failed periodic backup tasks. They are retained for one month and can be deleted manually. 	

 **NOTE**

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

Step 6 Click **OK**.

----End

Purchasing a Server Backup Vault and Applying the Backup Policy

Step 1 Log in to CBR Console. In the navigation pane on the left, choose **Cloud Server Backup**.

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

Step 3 Set **Protection Type** to **Backup**.

A server backup vault stores server backups.

 **NOTE**

If you need to purchase a replication vault to store the generated backup replicas, set **Protection Type** to **Replication (cross-region)**.

Step 4 Configure application-consistent backup.

- If application-consistent backup is enabled, the vault can be used to store database backups. For example, you can back up ECSs running MySQL or SAP HANA databases, because application-consistent backup ensures that the backed-up data is transactionally consistent. If an application-consistent backup task fails, CBR automatically performs a non-database server backup task instead. This non-database server backup will be stored in the database server backup vault.

- If application-consistent backup is not enabled, only non-database server backup is performed on associated servers, which is usually used for ECSs not running databases.

Step 5 Click **Configure** for **Associate Server**. In the displayed server list, select the servers or disks that you want to back up.

After the servers or disks are selected, they are added to the list of selected servers. For details, see [Figure 1-6](#). You can also select specific disks on a server and associate them with the vault.

NOTICE

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

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

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

Figure 1-6 Selecting servers

The screenshot shows two panels: 'Server List' and 'Selected Servers (1)'. The 'Server List' panel has filters for 'All projects', 'All statuses', and a search bar. It contains a table with columns: Name/ID, Status, Type, AZ, and Associated. One server is selected and highlighted in light blue. The 'Selected Servers (1)' panel has a search bar and a table with columns: Name/ID, Status, Type, AZ, Selected Disks, and Operation. It shows the same server selected, with '2/2' disks selected and a trash icon in the Operation column.

Name/ID	Status	Type	AZ	Associated
cce-iaas-w00... 2aacc7ed-6f0...	St...	ECS	AZ2	No

Name/ID	Status	Type	AZ	Selected Disks	Operation
cce-iaas-... 2aacc7e...	Sto...	ECS	AZ2	2/2	🗑️

NOTE

- The selected servers must be in the **Running** or **Stopped** state.

Step 6 Specify the vault capacity.

Its value ranges from the total size of the associated servers 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 7 Click **Configure** for **Auto Backup** and select a backup policy from the drop-down list.

The backup policy will be applied to this vault, and all servers associated with this vault will be automatically backed up based on the policy.

Step 8 If you have subscribed to the Enterprise Project Management Service (EPS), add the vault to an existing enterprise project.

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

 **NOTE**

If the **CBR FullAccess** permissions have been assigned to IAM users, enterprise projects will not be displayed for you to choose from when you create a vault. Go to the Enterprise Project Management console and assign the **CBR FullAccess** permissions to the target user group.

Step 9 (Optional) Add tags to the vault.

Tags are key-value pairs, which 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 1-8 describes the tag parameters.

Table 1-8 Tag parameter description

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

Step 10 Specify a name for the vault.

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

 **NOTE**

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

Step 11 Pay for the order as prompted.

Step 12 Go back to the **Cloud Server Backup** page. You can view the created vault in the vault list.

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

----End

Creating a Replication Policy

Step 1 Log in to CBR Console. In the navigation pane on the left, choose **Policies**.

Step 2 Click **Create Policy** in the upper right corner. On the displayed page, click **Replication policy** to create a custom replication policy. See [Figure 1-7](#).

Figure 1-7 Creating a replication policy

Create Policy [Back to CBR Policy List](#)

Type: Backup policy | **Replication policy**

* Name: backup_policy

Status:

* Execution Time:

<input checked="" type="checkbox"/> 00:00	<input type="checkbox"/> 01:00	<input checked="" type="checkbox"/> 02:00	<input type="checkbox"/> 03:00	<input type="checkbox"/> 04:00
<input type="checkbox"/> 05:00	<input type="checkbox"/> 06:00	<input type="checkbox"/> 07:00	<input type="checkbox"/> 08:00	<input type="checkbox"/> 09:00
<input type="checkbox"/> 10:00	<input type="checkbox"/> 11:00	<input type="checkbox"/> 12:00	<input type="checkbox"/> 13:00	<input type="checkbox"/> 14:00
<input type="checkbox"/> 15:00	<input type="checkbox"/> 16:00	<input type="checkbox"/> 17:00	<input type="checkbox"/> 18:00	<input type="checkbox"/> 19:00
<input type="checkbox"/> 20:00	<input type="checkbox"/> 21:00	<input type="checkbox"/> 22:00	<input type="checkbox"/> 23:00	

[Select All](#) [Invert Selection](#)

* Replication Cycle: Week-based cycle | Custom cycle Every days

* Retention Rule: Replica quantity | Time period | Permanent

[Learn more](#)

* Destination Region:

Price of replication traffic: ¥0.50/GB. [Pricing details](#) [?](#)

Step 3 Set the replication policy and **Create Now**.

----End

Purchasing a Server Replication Vault and Applying the Replication Policy

Step 1 Log in to CBR Console. In the navigation pane on the left, choose **Cloud Server Backup**.

Step 2 Click  in the upper left corner and select the region and project of the cross-region DR site.

Step 3 Click **Buy Server Backup Vault** in the upper right corner of the page. For details, see [Step 1](#) to [Step 12](#) in [Purchasing a Server Backup Vault and Applying the Backup Policy](#).

Step 4 Click  in the upper left corner and select the region and project of the production site.

Step 5 On the **Vaults** tab page, locate the row that contains the target vault, click **More** in the **Operation** column, and choose **Apply Replication Policy**.

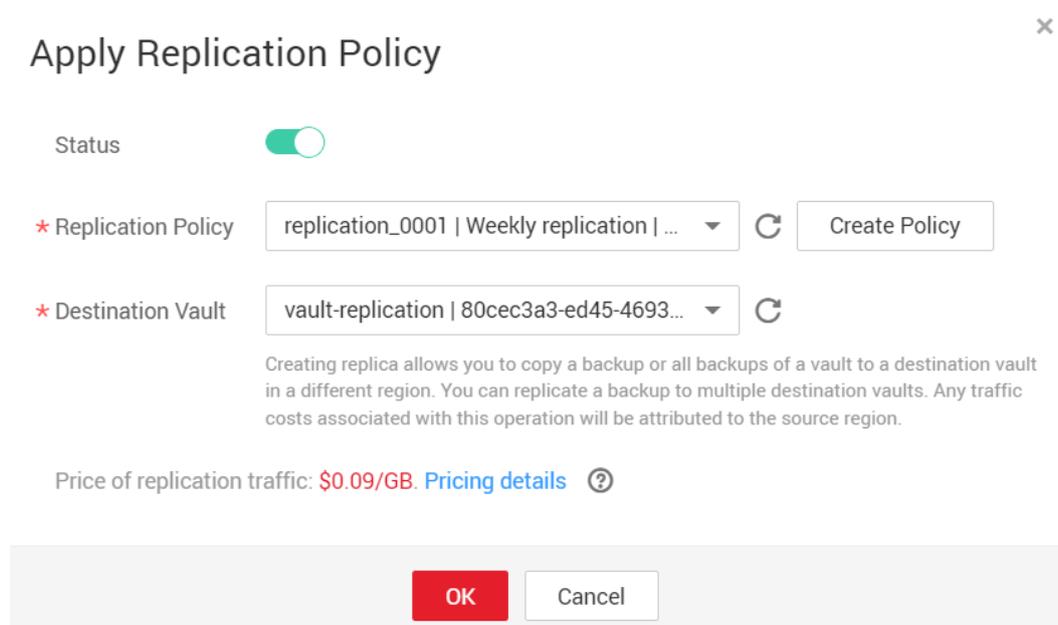
The **Apply Replication Policy** dialog box is displayed.

NOTE

The backup policy has been applied to the server backup vault when you purchase the vault. Now you only need to apply the replication policy.

Step 6 In the displayed dialog box, select the created replication policy and select the replication vault at the cross-region DR site as the destination vault.

Figure 1-8 Applying a replication policy



Step 7 Confirm the information and click **OK**.

Step 8 Go to the **Policies** page and view the policy status.

----End

1.5.2 Verifying the Configuration

Scenarios

After using CBR to configure the cross-region DR, you can verify whether the periodic backup and replication are configured successfully for the cross-region DR site.

Prerequisites

- You have configured the cross-region DR.
- You have configured the time for automatic backup and time for copy replication.

If the backup time specified in the backup policy is 00:00 and the replication time specified in the replication policy is 01:00, you can verify the backup and replication at 01:15 because the backup and replication will take certain time.

Procedure

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select the region and project of the production site.

Step 3 Click **Service List** and choose **Storage > Cloud Backup and Recovery > Cloud Server Backup Service**.

Step 4 Click the **Backups** tab.

The backups periodically created based on the backup policy are displayed in the list.

Step 5 Click  in the upper left corner and select the region and project of the cross-region DR site.

Step 6 On the cloud server backup service page, click the **Backups** tab.

Backups generated during the replication are displayed in the list. If letter **R** is displayed next to a backup name, the backup is successfully replicated to the region of the cross-region DR site.

----End

1.5.3 Restoring Services When Both the Production Site and Its Cross-AZ DR Site Become Faulty

Scenarios

This section describes how to restore services when the production site and its cross-AZ DR site become unavailable due to a major disaster. You can use the

latest cloud server backup on the cross-region DR site to restore services at the production site to the backup point of the backup source.

Prerequisites

- Confirm that the following operations have been performed before you back up an ECS:
 - The Linux ECS has been optimized (see [Optimizing a Linux Private Image](#)), and Cloud-Init has been installed on it (see [Installing Cloud-Init](#)).
 - The Windows ECS has been optimized (see [Optimizing a Windows Private Image](#)), and Cloudbase-Init has been installed on it (see [Installing and Configuring Cloudbase-Init](#)).
- The backup used to create images is in the **Available** state or 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 used to create images contains the system disk data.
- The backup is an ECS backup.

Procedure

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select the region and project of the cross-region DR site.

Step 3 Click **Service List** and choose **Storage > Cloud Backup and Recovery > Cloud Server Backup**.

Step 4 On the **Cloud Server Backup** page, click the **Backups** tab.

Backups generated during the replication are displayed in the list, and a letter **R** is displayed next to a backup name.

Step 5 Click **Create Image** in the **Operation** column.

The **Image Type and Source** page is displayed.

Step 6 Set the parameters as prompted. See [Figure 1-9](#).

Set **Region** to the region where the cross-region DR site is located, **Image Type** to **Full-ECS image**, and **Source** to **Cloud Server Backup**.

Figure 1-9 Creating an image

* Region: China North-Beijing

Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.

* Image Type: System disk image, **Full-ECS image**, Data disk image

* Source: ECS, **Cloud Server Backup**

- A cloud server backup can be used to create only one full-ECS image.
- You need to first customize the ECS used to create the vault backup to suit your needs. [Learn more](#)

All statuses | autobk_694a | X | Q | C

Backup Name	Backup Status	ECS Name	Created
au...	Available		Jan 10, 2025 09:09:00 AM CST

Selected: autobk_694a
[Create Server Backup](#)

Image Information

* Name: image-0001

* Enterprise Project: default | C | ?

Step 7 Click **Next**.

The **Details** page is displayed.

Step 8 Confirm the configuration and click **Submit**.**Step 9** Click **Back to Image List**.

The IMS console is displayed.

Step 10 In the image list, locate the newly created image and click **Apply for Server** in the **Operation** column.

The ECS configuration page is displayed.

Step 11 Configure the parameters as required. Set **Region** to the region where the cross-region DR site is located. Select **Private image** for **Image** and then select the created image from the drop-down list. See [Figure 1-10](#).**Figure 1-10** Configuring the image

Image: Public image, **Private image**, Shared image, Marketplace image

▼ [Image ID] | Create Private Image

It is recommended that you [view known OS issues](#) and then use the private image to create ECSs.

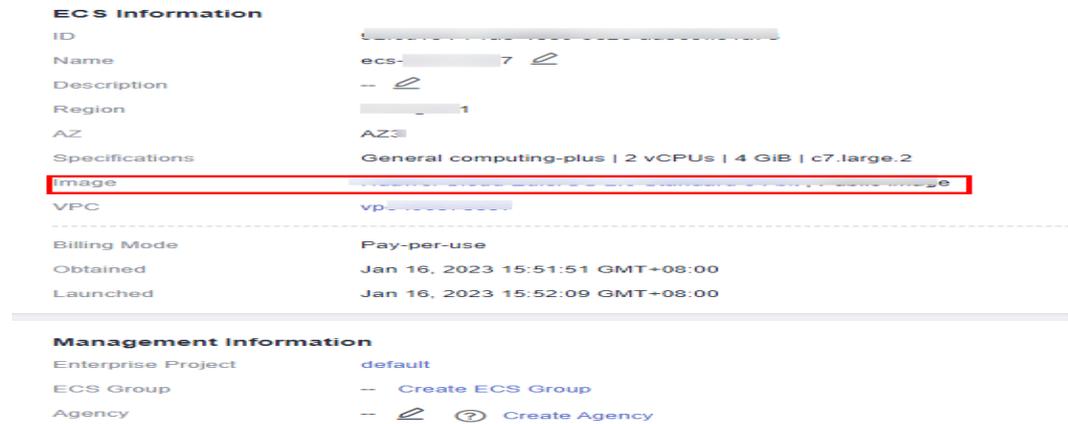
Step 12 Configure all required parameters and click **Next**.**Step 13** Go back to the ECS list.

The ECS console is displayed.

Step 14 Click the name of the created ECS to go to the ECS details page.

You can see that the ECS image is the newly created image. See [Figure 1-11](#).

Figure 1-11 ECS details page



Step 15 Log in to the ECS. You can see that the service has been restored to the backup time point.

NOTE

If you use the latest cloud server backup A' on the cross-region DR site to create the image in [Step 5](#), the services will be restored to the backup point of the backup A'.

----End

Related Operations

A backup used to create an image can be deleted only after the image is deleted.