Storage Disaster Recovery Service

Quick Start

Issue 01

Date 2025-02-10

HUAWEI TECHNOLOGIES CO., LTD.

Copyright © Huawei Technologies Co., Ltd. 2025. All rights reserved.

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

Trademarks and Permissions

HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

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

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

Contents

1 Asynchronous Replication (Restricted Use)

1.1 Overview

Asynchronous replication provides server-level protection if production site applications cannot be recovered within a short period of time due to force majeure (fire and earthquake) or device faults (faulty software and hardware). You can quickly recover services at the disaster recovery site with simple configurations.

NOTE

Asynchronous Replication is currently in **restricted use**. To use it, **submit a service ticket** to contact technical support.

Figure 1-1 shows the configuration process of asynchronous replication.

Figure 1-1 Configuration process

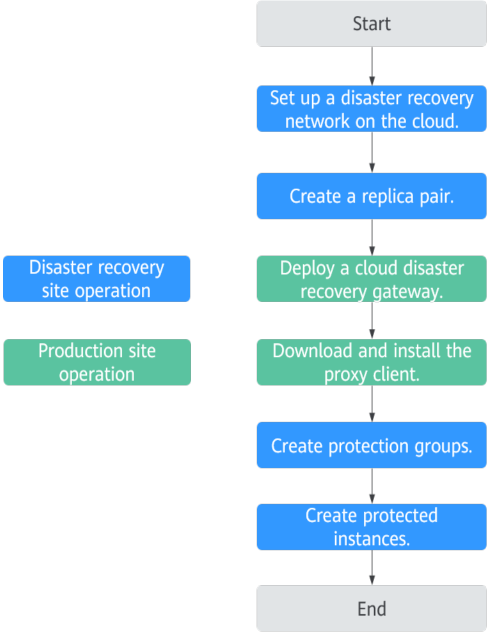
Start
Set up a disaster recovery network on the cloud.
Create a replica pair.
Disaster recovery site operation
Deploy a cloud disaster recovery gateway.
Production site operation
Download and install the proxy client.
Create protection groups.
Create protected instances.
End

1. Set up a disaster recovery network on the cloud.

   Create VPCs and subnets on the cloud based on your disaster recovery plan.

2. Create a replica pair to establish the replication relationship between the production site and the disaster recovery site.
3. Deploy a cloud disaster recovery gateway.

   The cloud disaster recovery gateway aggregates, compresses, and encrypts the data on all replicated production site servers, and synchronizes the data to the disaster recovery site.

4. Download and install the proxy client.

   A proxy client transmits the data on the server to the cloud disaster recovery gateway.

5. Create protected instances.

   A protected instance consists of a server and its replicated server. The system automatically identifies the production site servers with proxy client installed. Create protected instances for the servers that require disaster recovery. After a protected instance is created, protection is automatically enabled for data synchronization.

1.2 Preparation: Set Up a Disaster Recovery Network on the Cloud

Scenarios

A VPC provides an isolated virtual network for your disaster recovery site servers. You can configure and manage the network as required.

In asynchronous replication, data on the to-be-protected servers at the on-premises data center is continuously replicated to the cloud disaster recovery site through a network. When an outage occurs at your local data center, you can switch services to the disaster recovery site servers on the cloud to ensure service continuity.

Factors to consider when creating a network on the cloud:

- Scope of disaster recovery

  Select a region considering the following factors: physical distance between two sites, network performance, and costs. For example, keeping at least 100-km physical distance between the production site and disaster recovery site, less than 100-ms network latency, and cost-effective network (Direct Connect not used due to a tight budget).

- Network between the on-premises data center and the VPC on Huawei Cloud
  - Public network: suitable for scenarios that the data volume is stable and access to the cloud resources from the on-premises data center is infrequent.
  - VPN: suitable for scenarios that the data volume is stable and access to the Huawei Cloud resources from the on-premises data center is frequent. If some of your services are deployed on Huawei Cloud, and the on-premises data center services interact with the cloud services through a VPN, you can use this VPN for asynchronous replication.
  - Direct Connect: suitable for complex scenarios with a large volume of data. Make the plan based on the data change volume of your services.
- VPC CIDR block

  Provide IP addresses for the servers created during failovers or disaster recovery drills. To keep the IP addresses unchanged, set the VPC CIDR block to the same as the network segments of production site servers in the local data center. In this case, server IP addresses will remain the same during failovers or drills, without any additional configurations.

Procedure

Create a disaster recovery network on the cloud according to your overall network plan. For details, see **Creating a VPC**.

1.3 Step 1: Create a Replica Pair

Procedure

1. Log in to the SDRS console.
   a. Log in to the management console.
   b. In the upper left corner of the page, select the region where the service is located from the drop-down list.
   c. Click the service list icon in the upper left corner and choose **Storage** > **Storage Disaster Recovery Service**.
2. Go to **Create Replica Pair** page in either of the following ways:
   a. In the upper right corner of the SDRS console, click **Create Replica Pair**.
   b. Locate **Create Replica Pair** in the process flow and click **Create Now**.

Figure 1-2 Service Overview

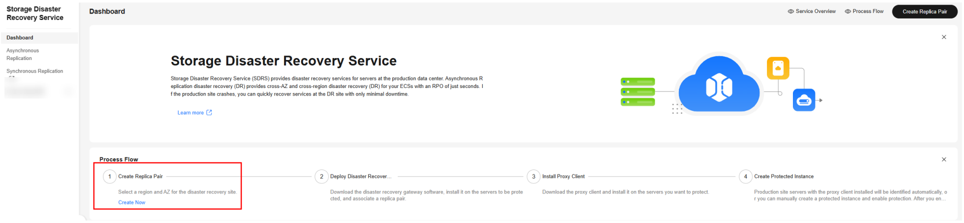

Storage Disaster Recovery Service

Dashboard

Asynchronous Replication

Synchronous Replication

Dashboard

Service Overview | Process Flow | Create Replica Pair

Storage Disaster Recovery Service

Storage Disaster Recovery Service (SDRS) provides disaster recovery services for servers at the production data center. Asynchronous Replication disaster recovery (DR) provides cross-AZ and cross-region disaster recovery (DR) for your ECSs with an RPO of just seconds. If the production site crashes, you can quickly recover services at the DR site with only minimal downtime.

Learn more

Process Flow

1 Create Replica Pair
Select a region and AZ for the disaster recovery site.
Create Now

2 Deploy Disaster Recover...
Download the disaster recovery gateway software, install it on the servers to be protected, and associate a replica pair.

3 Install Proxy Client
Download the proxy client and install it on the servers you want to protect.

4 Create Protected Instance
Production site servers with the proxy client installed will be identified automatically, or you can manually create a protected instance and enable protection. After you en...

3. Select the type of the replica pair you want to create and configure parameters by referring to the table below.

   a. **Cross-AZ**: The production site and disaster recovery site are located in different AZs of the same region.

   Figure 1-3 Creating a cross-AZ replica pair

   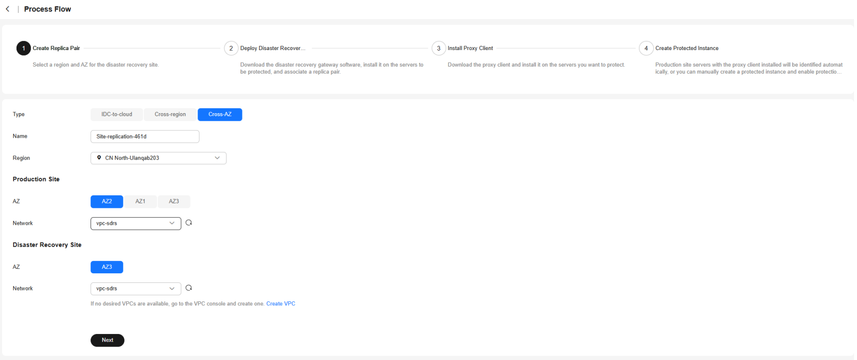

   Process Flow

   1 Create Replica Pair
   Select a region and AZ for the disaster recovery site.

   2 Deploy Disaster Recover...
   Download the disaster recovery gateway software, install it on the servers to be protected, and associate a replica pair.

   3 Install Proxy Client
   Download the proxy client and install it on the servers you want to protect.

   4 Create Protected Instance
   Production site servers with the proxy client installed will be identified automatically, or you can manually create a protected instance and enable protectio...

   Type: IDC-to-cloud | Cross-region | [✓ Cross-AZ]

   Name: Site-replication-461d

   Region: CN North-Ulanqab203

   Production Site

   AZ: [✓ AZ2] | AZ1 | AZ3

   Network: vpc-sdrs

   Disaster Recovery Site

   AZ: [✓ AZ3]

   Network: vpc-sdrs

   If no desired VPCs are available, go to the VPC console and create one. Create VPC

   Next

   b. **Cross-region**: The production site and disaster recovery site are located in different regions.

   Figure 1-4 Creating a cross-region replica pair

   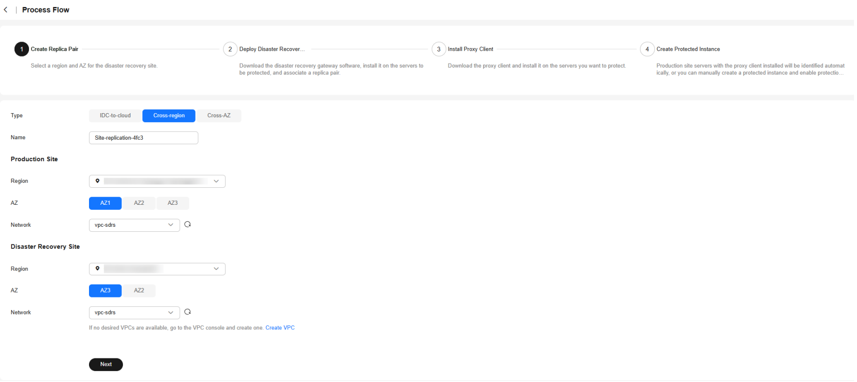

   Process Flow

   1 Create Replica Pair
   Select a region and AZ for the disaster recovery site.

   2 Deploy Disaster Recover...
   Download the disaster recovery gateway software, install it on the servers to be protected, and associate a replica pair.

   3 Install Proxy Client
   Download the proxy client and install it on the servers you want to protect.

   4 Create Protected Instance
   Production site servers with the proxy client installed will be identified automatically, or you can manually create a protected instance and enable protectio...

   Type: IDC-to-cloud | [✓ Cross-region] | Cross-AZ

   Name: Site-replication-4fc3

   Production Site

   Region:

   AZ: [✓ AZ1] | AZ2 | AZ3

   Network: vpc-sdrs

   Disaster Recovery Site

   Region:

   AZ: [✓ AZ3] | AZ2

   Network: vpc-sdrs

   If no desired VPCs are available, go to the VPC console and create one. Create VPC

   Next

   c. **IDC-to-cloud**: The production site is deployed in a local data center.

Figure 1-5 Creating an IDC-to-cloud replica pair

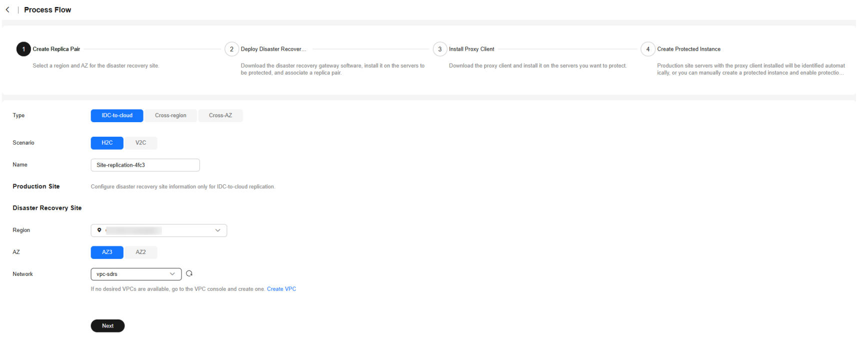

< | **Process Flow**

1 Create Replica Pair — Select a region and AZ for the disaster recovery site.

2 Deploy Disaster Recover... — Download the disaster recovery gateway software, install it on the servers to be protected, and associate a replica pair.

3 Install Proxy Client — Download the proxy client and install it on the servers you want to protect.

4 Create Protected Instance — Production site servers with the proxy client installed will be identified automatically, or you can manually create a protected instance and enable protectio...

Type: IDC-to-cloud (selected) | Cross-region | Cross-AZ

Scenario: H2C (selected) | V2C

Name: Site-replication-4fc3

Production Site Configure disaster recovery site information only for IDC-to-cloud replication.

Disaster Recovery Site

Region: [redacted]

AZ: AZ3 (selected) | AZ2

Network: vpc-sdrs

If no desired VPCs are available, go to the VPC console and create one. Create VPC

Table 1-1 Parameter description

Parameter		Description	Example Value
Type		Type of the replica pair The supported types are **IDC-to-cloud**, **Cross-region** and **Cross-AZ**.	Cross-AZ
Scenario		Select the replication scenario you want to set up. Supported scenarios are: **H2C** (HCS Online DR to the public cloud) and **V2C** (VMware DR to the public cloud). **NOTE** This field shows up only when you are creating an IDC-to-cloud replica pair.	H2C
Name		Name of the replica pair The name can contain letters, digits, underscores (_), hyphens (-), or periods (.), can be no more than 64 characters long, and cannot contain spaces.	Site-replication-001
Production Site **NOTE** You only need to configure the production site when creating a cross-region or cross-AZ replica pair.	Region	Region where the production site resides **NOTE** You only need to select a region when creating a cross-region replica pair.	-
	AZ	AZ where the production site servers reside **NOTE** You only need to select an AZ when creating a cross-region or cross-AZ replica pair.	AZ1
	Network	VPC where the production site servers reside	VPC01

Parameter		Description	Example Value
Disaster Recovery Site	Region	Region where the disaster recovery site resides Select the region you selected when you set up the disaster recovery network. For details, see **Preparation: Set Up a Disaster Recovery Network on the Cloud**. **NOTE** You only need to select a region when creating an IDC-to-cloud or a cross-region replica pair.	-
	AZ	AZ where the disaster recovery site servers reside	AZ2
	Network	VPC where the disaster recovery site servers reside	VPC02

4. Click **Next** to go to he **Deploy Disaster Recovery Gateway** page.

1.4 Step 2: Deploy the Cloud Disaster Recovery Gateway

Procedure

1. Create an ECS to deploy the cloud disaster recovery gateway.

   NOTE

   - The disaster recovery gateway must be **deployed separately**. Do not deploy the gateway and proxy client on the same server.
   - The region, AZ, and VPC configurations of the ECS must be the same as those of the production site servers.
   - The minimum specifications supported by a cloud diaster recovery gateway are 2 vCPUs and 4 GB memory. You are advised to select specifications with 8 vCPUs and 16 GB memory or higher.

2. Install and configure the cloud disaster recovery gateway.

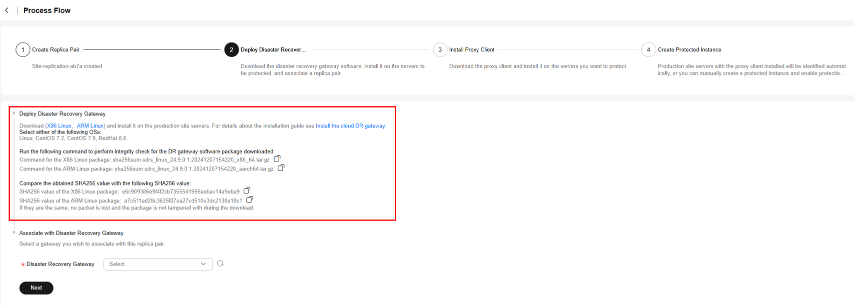
Process Flow
1 Create Replica Pair
2 Deploy Disaster Recover...
3 Install Proxy Client
4 Create Protected Instance
Disaster Recovery Gateway
Select
Next

Site-replication-ab7a created

Download the disaster recovery gateway software, install it on the servers to be protected, and associate a replica pair.

Download the proxy client and install it on the servers you want to protect.

Production site servers with the proxy client installed will be identified automatically, or you can manually create a protected instance and enable protectio...

Deploy Disaster Recovery Gateway

Download (X86 Linux、ARM Linux) and install it on the production site servers. For details about the installation guide see Install the cloud DR gateway.
Select either of the following OSs:
Linux: CentOS 7.2, CentOS 7.9, RedHat 8.6.

Run the following command to perform integrity check for the DR gateway software package downloaded:
Command for the X86 Linux package: sha256sum sdrs_linux_24.9.0.1.20241207154220_x86_64.tar.gz
Command for the ARM Linux package: sha256sum sdrs_linux_24.9.0.1.20241207154220_aarch64.tar.gz

Compare the obtained SHA256 value with the following SHA256 value:
SHA256 value of the X86 Linux package: e5c909386e9f4f2cb73555d1956aebac14a9eba9
SHA256 value of the ARM Linux package: a7c511ad28c3625f07ea27cd510e3dc2138e18c1
If they are the same, no packet is lost and the package is not tampered with during the download.

Associate with Disaster Recovery Gateway

Select a gateway you wish to associate with this replica pair.

For details, see **Installing a Disaster Recovery Gateway** and **Configuring a Disaster Recovery Gateway**.

3. Associate the replica pair with the disaster recovery gateway.

   Select the disaster recovery gateway you have deployed to associate with this replica pair.

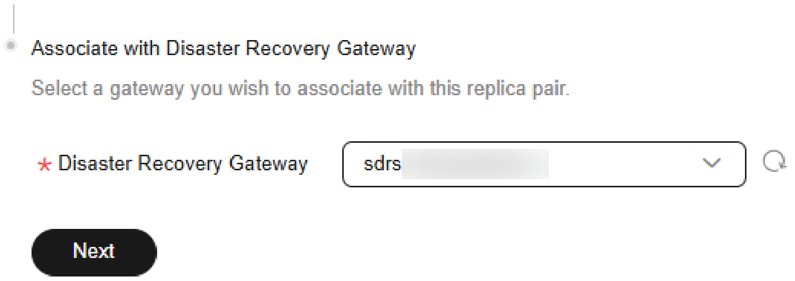

Associate with Disaster Recovery Gateway

Select a gateway you wish to associate with this replica pair.

* Disaster Recovery Gateway: sdrs[redacted]

   NOTE

   If the gateway cannot be found, see **Why Can't I Find the Disaster Recovery Gateway When Associating a Replica Pair with It?**

4. Click **Next**.

   The cloud disaster recovery gateway is deployed, and the **Install Proxy Client** page is displayed.

1.5 Step 3: Download and Install the Proxy Client

Procedure

1. Select the OS and version used on the production site servers.

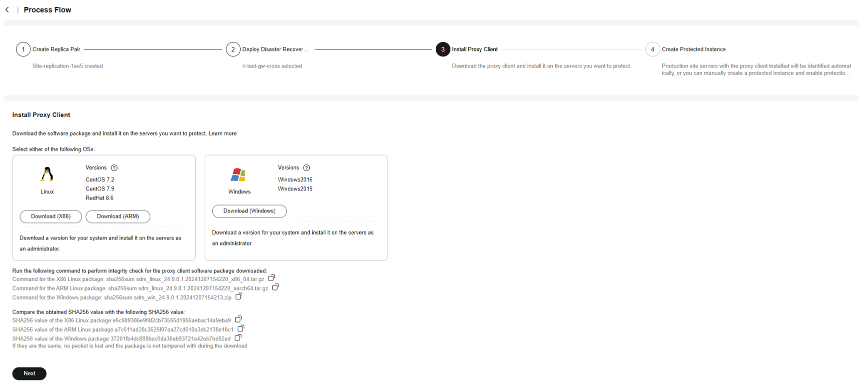

Process Flow

1 Create Replica Pair: Site-replication-1ee5 created

2 Deploy Disaster Recover...: rl-test-gw-cross selected

3 Install Proxy Client: Download the proxy client and install it on the servers you want to protect.

4 Create Protected Instance: Production site servers with the proxy client installed will be identified automatically, or you can manually create a protected instance and enable protectio...

Install Proxy Client

Download the software package and install it on the servers you want to protect. Learn more

Select either of the following OSs:

Linux
Versions: CentOS 7.2, CentOS 7.9, RedHat 8.6

Download a version for your system and install it on the servers as an administrator.

Windows
Versions: Windows2016, Windows2019

Download a version for your system and install it on the servers as an administrator.

Run the following command to perform integrity check for the proxy client software package downloaded:

Command for the X86 Linux package: sha256sum sdrs_linux_24.9.0.1.20241207154220_x86_64.tar.gz

Command for the ARM Linux package: sha256sum sdrs_linux_24.9.0.1.20241207154220_aarch64.tar.gz

Command for the Windows package: sha256sum sdrs_win_24.9.0.1.20241207154213.zip

Compare the obtained SHA256 value with the following SHA256 value:

SHA256 value of the X86 Linux package:e5c909386e9f4f2cb73555d1956aebac14a9eba9

SHA256 value of the ARM Linux package:a7c511ad28c3625f07ea27cd510e3dc2138e18c1

SHA256 value of the Windows package:37201fb4dc888bec0da36eb93721e42eb7bd82ed

If they are the same, no packet is lost and the package is not tampered with during the download.

2. Install the proxy client.

   For details, see **Installing a Proxy Client**.

3. Click **Next**.

   The proxy client is installed, and the **Create Protected Instance** page is displayed.

1.6 Step 4: Create a Protection Group and Protected Instances

Procedure

1. Configure the protected instance information. **Table 1-2** describes the parameters.
2. Click **Next**.

   The page for you to confirm disaster recovery information is displayed.

   Confirm the configuration and click **Submit**. The configuration is complete, and the **Asynchronous Replication** page is displayed.

   Figure 1-6 Create Protected Instance

   NOTE

   If no production site servers are displayed, see **Why Is No Production Site Server Displayed When I Create Protected Instances?**

3. If no protection group is created, click **Create Protection Group** to create one.

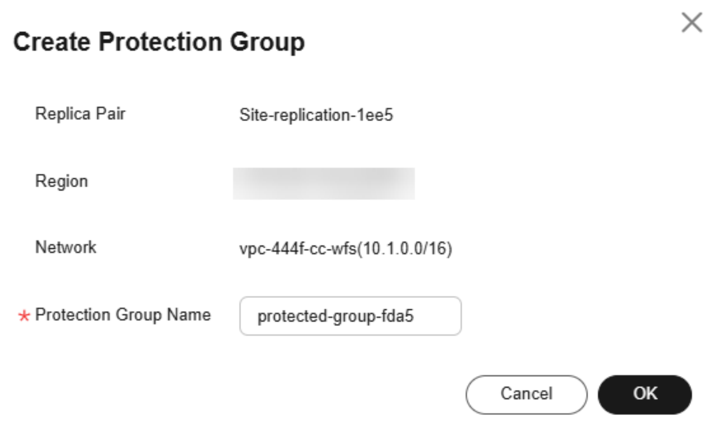

Create Protection Group

Replica Pair Site-replication-1ee5

Region [redacted]

Network vpc-444f-cc-wfs(10.1.0.0/16)

* Protection Group Name protected-group-fda5

Table 1-2 Parameter description

Parameter	Description	Example Value
Production Site Server	- Select production site servers you want to protect. If no server is available, it may be because the agent deployment failed. - Select the disk type for each disaster recovery site disk. - Enter a name for each protected instance. The name can contain letters, digits, underscores (_), hyphens (-), or periods (.), can be no more than 64 characters long, and cannot contain spaces.	-
Protection Group	Select a protection group for the protected instances. If you create protected instances first time ever or the current protection group does not meet your requirements, click **Create Protection Group** to create a new one. It is recommended that you add servers of a specific business to the same protection group. In this case, you can run DR drills, start protection and perform failovers for the entire group.	protected-group-01

2 Synchronous Replication (for Installed Base Operations)

2.1 Configuration Process

SDRS provides server-level protection (RPO = 0) if production site applications cannot be recovered within a short period of time due to force majeure (fire and earthquake) or device faults (faulty software and hardware). Storage-layer synchronous replication provides cross-AZ DR protection to meet data consistency requirements. If the production site fails, you can quickly restore services at the cross-AZ DR site with a few clicks.

Figure 2-1 shows the cross-AZ DR configuration process.

📖 **NOTE**

When you create a protected instance, the system creates a replication pair for the disks of the servers at the production and DR site by default.

Figure 2-1 Cross-AZ DR configuration process

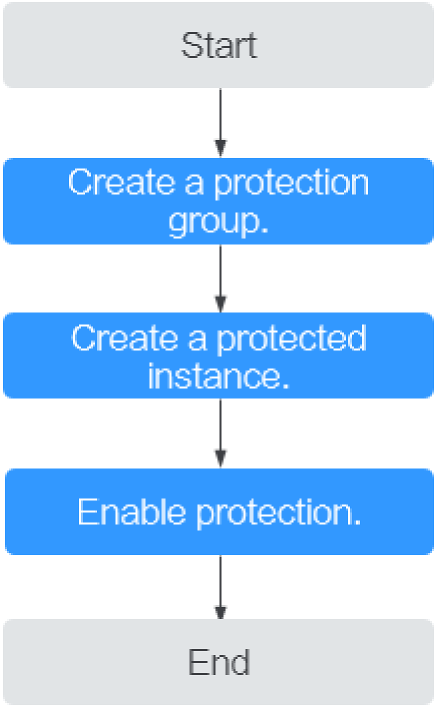

2.2 Step 1: Create a Protection Group

Scenarios

You can specify two AZs as the source and target AZs, and create a protection group. Then, you can create protected instances and replication pairs in this protection group.

Verify the servers at the production and DR sites before you create a protection group. In this version, only the VPC migration deployment model is supported. Specifically, the servers at the production and DR sites must be in different AZs but in the same VPC.

Figure 2-2 Creating a protection group

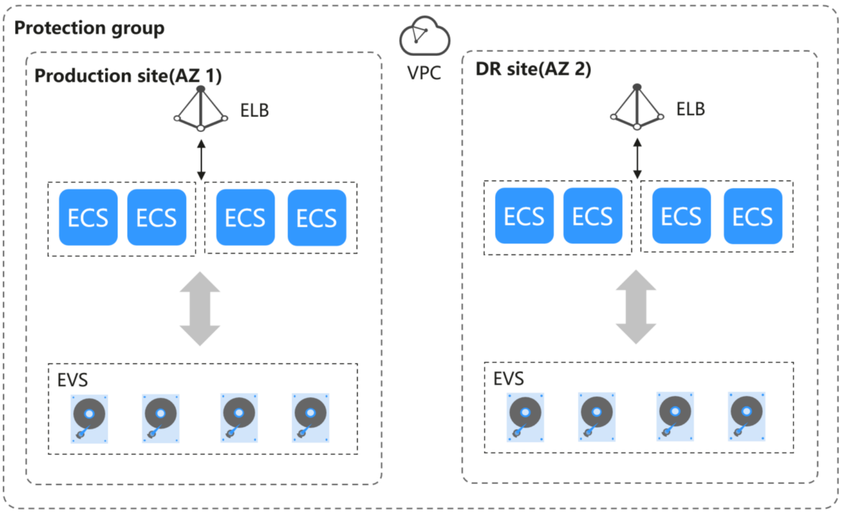

Procedure

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 navigation pane on the left, choose **Synchronous Replication**.

Step 4 Click **Create Protection Group** in the upper right corner.

Step 5 Configure the basic information about the protection group listed in **Table 2-1**.

NOTE

Parameters listed in **Table 2-1** are mandatory.

Table 2-1 Parameter description

Parameter	Description	Example Value
Region	A region is a geographic area where resources used by servers are located. If the region is incorrect, click the drop-down list for correction.	AP-Bangkok
DR Direction	• Production site: Select the AZ of the production site server. • DR site: Select the AZ of the DR site server.	Production site: az-01 DR site: az-02

Parameter	Description	Example Value
Deployment Model	Currently, only the VPC migration model is supported. All resources at the production and DR sites belong to the same VPC.	VPC migration
VPC	Specifies the VPC where the protection group is located.	vpc-test
Protection Group Name	Enter the protection group name. It is used for group classification and search.	protection_group_001

Step 6 Click **Create Now**.

Step 7 Click **Back to Protection Group List** to return to the SDRS homepage and query the protection group status.

If the protection group is displayed in the **Storage Disaster Recovery Service** page and its status is **Available** (✓), the protection group has been created successfully.

----End

2.3 Step 2: Create Protected Instances

Scenarios

You can create protected instances using the servers that you want to perform DR protection. If the current production site encounters an unexpected large-scale server failure, you can call the related protection group API to perform a failover, ensuring that services running on protected instances are not affected.

Select a protection group for each server to be replicated and create a protected instance. When you create a protected instance, the server and disk will be created at the DR site for the production site server and disk. The server specifications can be configured as required. Specifically, the specifications of the DR site server can be different from those of the production site server. The disks of the production site and DR site are of the same specifications and can automatically form a replication pair.

The server at the DR site is in the Stopped state after the protected instance created. These automatically created resources, including the DR site servers and disks, cannot be used before a switchover or failover.

Figure 2-3 Creating a protected instance

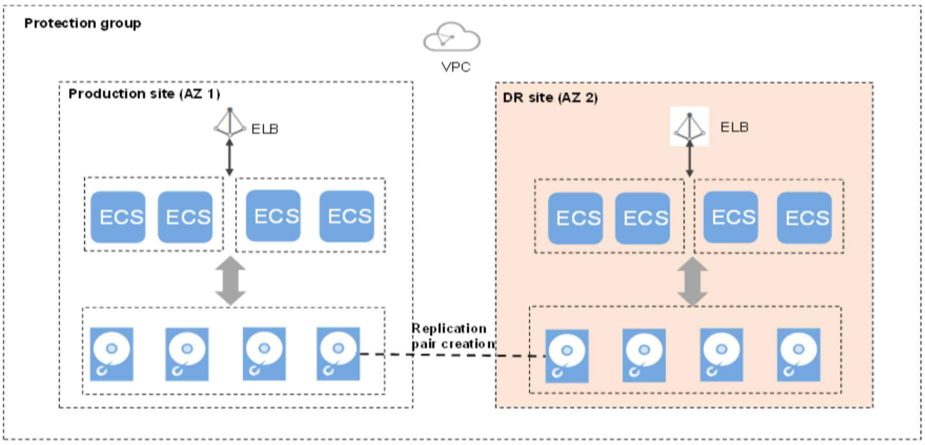

Notes

- If a production site server has been added to an ECS group, you are not allowed to specify a DeH to create the DR site server for the production site server.
- When a protected instance is created, the default name of the server at the DR site is the same as that of the server at the production site, but their IDs are different.
- To modify a server name, switch to the protected instance details page and click the server name to switch to the server details page.
- After you create a protected instance and enable protection for the server at the production site, modifications to the **Hostname**, **Name**, **Agency**, **ECS Group**, **Security Group**, **Tags**, and **Auto Recovery** configurations of the production site server will not synchronize to the DR site server. You can log in to the management console and manually add the configuration items to the servers at the DR site.
- If protection is enabled for servers created during capacity expansion of an Auto Scaling (AS) group, these servers cannot be deleted when the capacity of the AS group is reduced.
- If the server at the production site runs Windows and you choose the key login mode, ensure that the key pair of the server exists when you create a protected instance. Otherwise, the server at the DR site may fail to create, causing the protected instance creation failure.

  NOTE

  If the key pair of the server at the production site has been deleted, create a key pair with the same name.
- When you create a protected instance, if the production site server runs Linux and uses the key login mode, the key pair information will not be displayed on the details page of the DR site server after the DR site server is created. You can use the key pair of the production site server to log in to the DR site server.
- If the production site server is added to Enterprise Project, the created DR site server will not be automatically added to Enterprise Project. You need to manually add the server to Enterprise Project if needed.

- Spot instances cannot be used as production site servers.

Prerequisites

- The protection group is in the **Available** or **Protecting** state.
- No protected instances have been created for the production site server.
- Resources of the target specifications for the server to be protected are not sold out at the DR site.
- The server that you use to create a protected instance and the protection group are in the same VPC.

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 navigation pane on the left, choose **Synchronous Replication**.

Step 4 In the pane of the protection group for which protected instances are to be added, click **Protected Instances**.

The protection group details page is displayed.

Step 5 On the **Protected Instances** tab, click **Create Protected Instance**.

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

Step 6 Configure the basic information about the protected instance, as described in **Table 2-2**.

Table 2-2 Parameter description

Parameter	Description	Example Value
Protection Group Name	Indicates the name of the protection group to which the protected instance to be created belongs. You do not need to configure it.	protection_group_001
Protection Group ID	Indicates the ID of the protection group to which the protected instance to be created belongs.	2a663c5c-4774-4775-a321-562a1ea163e3
DR Direction	Indicates the replication direction of the protection group to which the protected instance to be created belongs. You do not need to configure it.	-

Parameter	Description	Example Value
Production Site	Indicates the AZ of the production site server. You do not need to configure it.	az-01
Deployment Model	Indicates the deployment model of the protection group to which the protected instance to be created belongs. You do not need to configure it.	VPC migration
VPC	Indicates the VPC of the protection group to which the protected instance to be created belongs. You do not need to configure it.	vpc1
Production Site Server	This parameter is mandatory. In the server list, select the server and specifications to be used to create the protected instance. ● You can select a maximum of five production site servers at a time. ● If a shared disk is attached to a server, you need to select all servers to which the shared disk is attached. **NOTE** ● If **Server Type** of the protection group is **ECS**, select the DR site server specifications. The specifications of the production site server and DR site server can be different. Select the specifications from the **DR Site Server Specifications** drop-down list.	ecs-test > s3.small.1

Parameter	Description	Example Value
DR Site Server	This parameter is mandatory when **Server Type** of the protection group is **ECS**. You can select to use ECSs to create DR site servers or to deploy DR site servers on DeHs. DeHs are physical hosts dedicated for a specified user. You can create servers on a DeH to enhance isolation, security, and performance of your ECSs. **NOTE** If a production site server has been added to an ECS group, you are not allowed to specify a DeH to create the DR site server for the production site server.	ECS
DeH	This parameter is mandatory when **DR Site Server** is set to **DeH**. Select a DeH for deploying the DR site server. If multiple production site servers are selected, the DR site servers will be created on the same DeH.	deh-01
DR Site VPC	Indicates the VPC of the DR site server. Its value is the same as the **VPC** value and do not need to be configured.	vpc1

Parameter	Description	Example Value
DR Site Primary NIC	This parameter is optional. Indicates the primary NIC on the DR site server. You can use the primary NIC automatically allocated by the system or specify a primary NIC based on your network plan. After your select a NIC, select an available subnet from the drop-down list and configure the private IP address. **NOTE** ● IP addresses cannot be specified if you create multiple protected instances at a time. ● After a successful restoration, the IP addresses of DR site servers are the same as those of production site servers. They cannot be customized. In addition, because DR site servers and production site servers are in the same subnet, you do not need to bind EIPs to DR site servers.	subnet-01 (192.168.0.0/24)
DR Site Disk	This parameter is mandatory. The following two options are available: ● **EVS** ● **DSS** If you select **DSS** for **DR Site Disk**, **Storage Pool** is mandatory.	EVS
Storage Pool	● 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	This parameter is mandatory. Enter the protected instance name. It is used for protected instance classification and search.	Protected-Instance-test

NOTE

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

Step 7 Click **Create Now**.

Step 8 On the **Confirm** page, you can confirm the protected instance information.

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

Step 9 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 the disks of the protected instance and backs up all the disks.

Query the replication pairs.

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

   On this tab, you can query the statuses of the replication pairs, target protected instance, and production site disk.

----End

2.4 Enable Protection

Scenarios

You can enable protection for all resources in a protection group.

When data is written to the disks of the production site server, SDRS synchronizes the data to the disks of the DR site server in real time. Both the production site and DR site can use Cloud Server Backup Service (CSBS) and Volume Backup Service (VBS) to back up the servers and disks.

Figure 2-4 Enabling protection

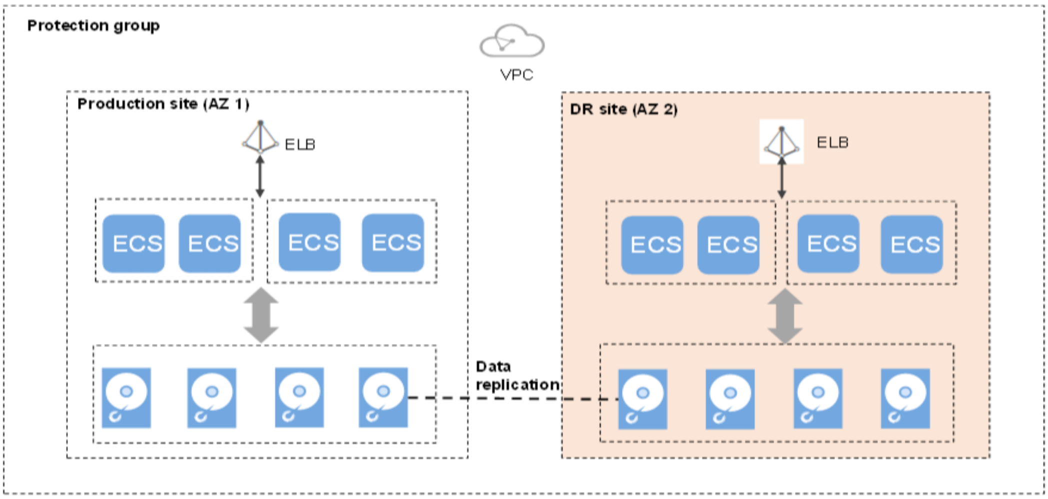

Prerequisites

- The protection group has replication pairs.
- The protection group is in the **Available** or **Enabling protection failed** state.
- After you create a protected instance and enable protection on servers at the production site, modifications to the **Hostname**, **Name**, **Security Group**, **Agency**, **ECS Group**, **Tags**, and **Auto Recovery** configurations of servers on the production site will not synchronize to the servers at the DR site. You can manually add the configuration items to the servers at the DR site on the management console.

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 navigation pane on the left, choose **Synchronous Replication**.

Step 4 In the pane of the desired protection group, click **Enable Protection**.

Step 5 In the displayed dialog box, click **Yes**.

Once protection is enabled, data synchronization starts.

NOTE

The synchronization time is in direct proportion to the disk capacity. Larger disk capacity requires longer synchronization time.

----End

Copyright © Huawei Technologies Co., Ltd.