Storage Disaster Recovery Service

Quick Start

Issue 05

Date 2023-04-07

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 Asynchronous Replication

1.1 Configuration Process

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 some simple configurations.

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

Figure 1-1 Configuration process

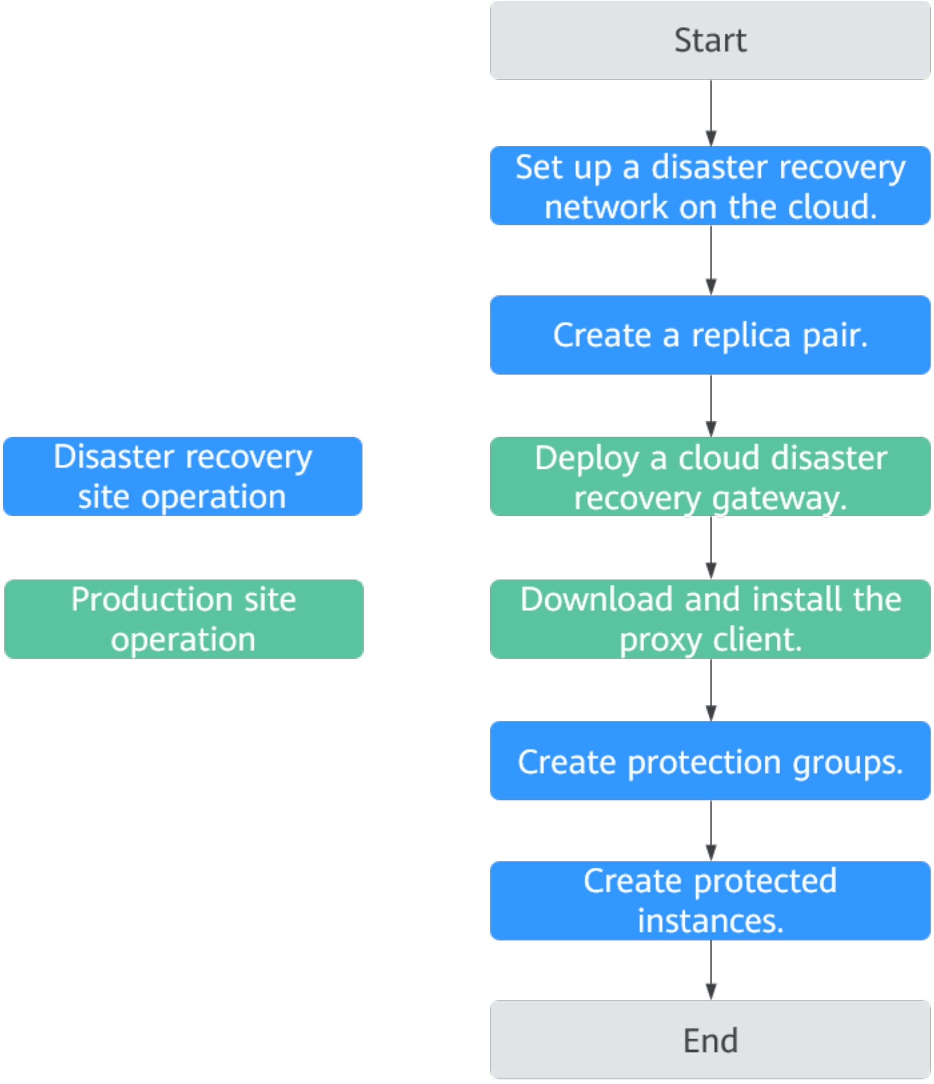
Start
Set up a disaster recovery network on the cloud.
Create a replica pair.
Disaster recovery site operation
Deploy a cloud disaster recovery gateway.
Production site operation
Download and install the proxy client.
Create protection groups.
Create protected instances.
End

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.
   - **IDC-to-cloud**: Select a region, an AZ, and a VPC for the disaster recovery site to connect the on-premises data center with the created cloud disaster recovery network. In this way, replication is established between the two sites.

     NOTE

     If you select **IDC-to-cloud** for **Type**, you only need to specify a region and an AZ for the disaster recovery site.
   - **Cross-region**: Select a region and a VPC for the production site, and select a region, an AZ, and a VPC for the disaster recovery site to connect the two sites. In this way, replication is established between the two sites.

> NOTE
>
> If you select **Cross-region** for **Type**, the regions you specify for the production site and disaster recovery site must be different.

- **Cross-AZ**: Select an AZ and a VPC for both the production site and disaster recovery site individually to connect the two sites. In this way, replication is established between the two sites.

> NOTE
>
> If you select **Cross-AZ** for **Type**, you only need to specify the AZs for the two sites because they are already in the same region.

3. Deploy a cloud disaster recovery gateway.

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

   - **IDC-to-cloud**: The system generates the cloud disaster recovery gateway software based on the replica pair information you provided. You can download the software from the console and deploy it in your on-premises data center.
   - **Cross-region** and **Cross-AZ**: Run the command provided on the console to obtain the cloud disaster recovery gateway software and then deploy it at the production site.

> ⚠ CAUTION
>
> - Because communication between the disaster recovery gateway and proxy client is not encrypted, it is recommended that you deploy them in the same security group and only allow ECSs within the security group to communication with each other. For details, see **Security Group Configuration Examples**.
> - Regardless of whether the production server runs Linux or Windows, it is recommended that the disaster recovery gateway be deployed on a Linux server.

4. Download and install the proxy client.

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

   - **IDC-to-cloud**: Download the proxy client from the console based on the production site server OS and install the client on the server.
   - **Cross-region** and **Cross-AZ**: Select the OS and version of the production site server and run the command provided on the console to obtain the proxy client software.

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 Step 1: Set Up a Disaster Recovery Network on the Cloud

Scenarios

A VPC provides an isolated virtual network for your disaster recovery 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 the services to the disaster recovery servers on the cloud to guarantee the 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
  - 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 reserve the IP addresses, set the VPC CIDR block to the same as the network segments of the on-premises data center servers. 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 2: Establish the Disaster Recovery Relationship

Scenarios

On the SDRS console, establish a disaster recovery relationship by creating a replica pair, deploy a cloud disaster recovery gateway, download and install the proxy client, and create protected instances. After the disaster recovery relationship is established, the system automatically starts data synchronization.

Procedure

Step 1 Create a replica pair.

1. Log in to the management console.
2. Click **Service List** and choose **Storage** > **Storage Disaster Recovery Service**.

   The **Storage Disaster Recovery Service** page is displayed.
3. In the navigation pane, choose **Asynchronous Replication**. In the upper right corner of the displayed page, click **Create Replica Pair**.
4. On the displayed page, set the parameters as prompted.

Table 1-1 Parameter description

Parameter		Description	Example Value
Type		Type of the replica pair There are three types: **IDC-to-cloud**, **Cross-region**, and **Cross-AZ**. Only Huawei Cloud Stack is supported for IDC-to-cloud replication currently.	IDC-to-cloud
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** Set the production site only when you are creating a cross-region or cross-AZ replica pair.	Region	Region where the production site resides **NOTE** Set the production site region only when you are creating a cross-region replica pair.	-
	AZ	AZ where the production site servers reside **NOTE** Set the production site AZ only when you are creating a cross-AZ replica pair.	AZ1

Parameter		Description	Example Value
	Network	VPC where the production site servers reside	VPC01
Disaster Recovery Site	Region	Region where the disaster recovery site resides Select the region created in **Step 1: Set Up a Disaster Recovery Network on the Cloud**. **NOTE** Set the disaster recovery site region only when you are 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 Select the VPC created in **Step 1: Set Up a Disaster Recovery Network on the Cloud**.	VPC02

5. Click **Next**.

   The **Deploy Disaster Recovery Gateway** page is displayed.

Step 2 Deploy a cloud disaster recovery gateway.

1. Obtain the gateway package and upload it to a directory on the target server.
   - **IDC-to-cloud**: Upload the software package to the DR gateway node.
   - **Cross-region** and **Cross-AZ**: Copy the command provided on the console, go to the directory where you want to install the gateway, and paste and run the command to obtain the package.

   Software name: **sdrs_linux_amd64_*xxx*_with_certs.tar.gz**, in which *xxx* indicates the software version

2. Install and configure the cloud disaster recovery gateway.

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

   ⚠ CAUTION

   To ensure that servers can run properly, make sure that the ports described in **Port Description (Asynchronous Replication)** are not used.

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

   Select the disaster recovery gateway you have deployed to associate with this replica pair. If your desired gateway is not available, it may be because the gateway deployment failed.

4. Click **Next**.

   The **Install Proxy Client** page is displayed.

Step 3 Download and install the proxy client.

1. If you need to install the proxy client on other nodes, use the **sdrs_linux_amd64_xxx_with_certs.tar.gz** package for installation. For details, see **Installing and Upgrading the Proxy Client**.
2. Click **Next**.

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

Step 4 Create protected instances.

1. Set the parameters as prompted.

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 one protection group. In this case, you can run DR drills, start protection and perform failovers for the entire group.	protected-group-01

2. Click **Next**.

   The **Details** page is displayed.
3. Confirm the configuration and click **Submit**.

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

----End

2 Synchronous Replication

2.1 Configuration Process

Synchronous replication replicates servers from one AZ to another in real time with zero RPO. By leveraging synchronous replication techniques at the storage layer, it allows for cross-AZ disaster recovery and keeps crash consistency for your data. If production site services fail to recover within a short period of time due to force majeure (fire and earthquake) or device faults (software and hardware damage), you can quickly recover services at the disaster recovery site with some simple configurations.

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

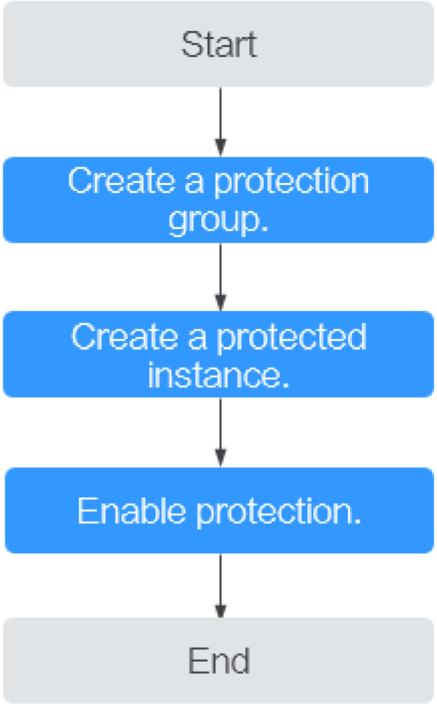

2.2 Create a Protection Group

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

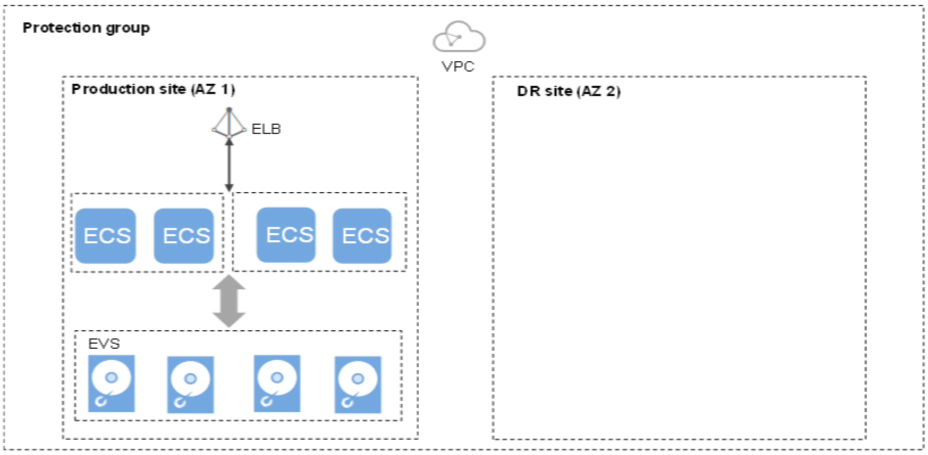

Procedure

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 Click **Create Protection Group**.

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

Step 4 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
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 5 Click **Create Now**.

Step 6 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 Create a Protected Instance

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

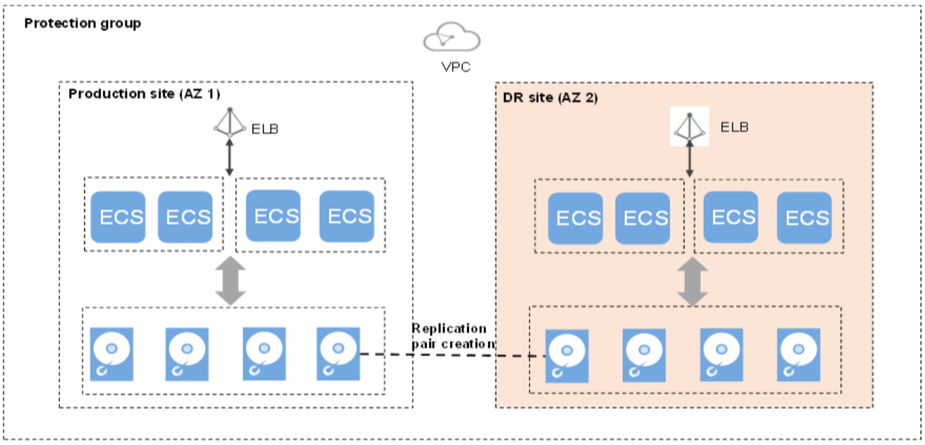

Notes

- 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.

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

The protection group details page is displayed.

Step 4 On the **Protected Instances** tab, click **Create Protected Instance**.

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

Step 5 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.	-
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

Parameter	Description	Example Value
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
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

Parameter	Description	Example Value
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
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)

Parameter	Description	Example Value
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
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 6 Click **Create Now**.

Step 7 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 8 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

If you want to enable protection for all resources in a specified protection group, you can perform steps provided in this section.

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

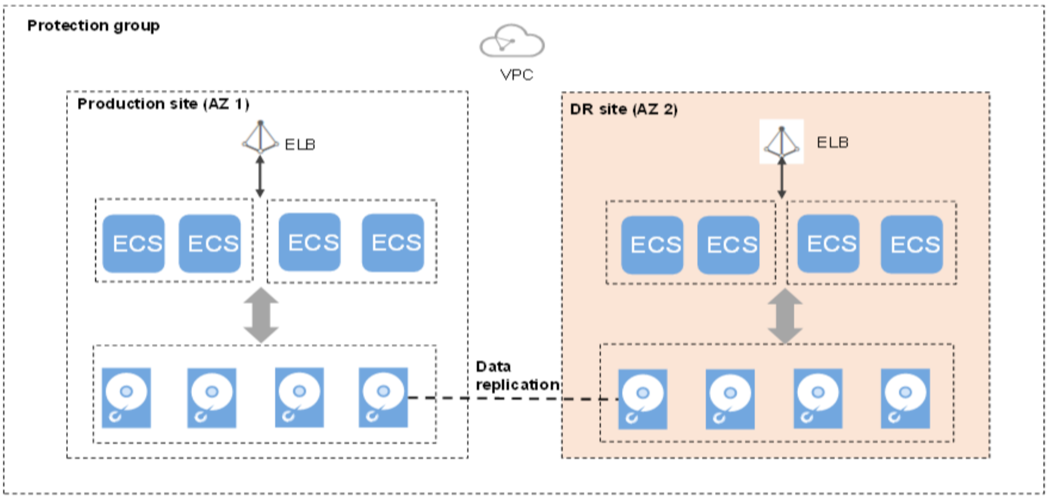

Prerequisites

- 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 pane of the desired protection group, click **Enable Protection**.

Step 4 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

3 Appendix

3.1 Installing and Upgrading a Disaster Recovery Gateway

Scenarios

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 some simple configurations.

Prerequisites

- LVM is not used on the system disk of the server used to deploy the disaster recovery gateway.
- Because communication between the disaster recovery gateway and proxy client is not encrypted, it is recommended that you deploy them in the same security group and only allow ECSs within the security group to communication with each other. For details, see **Security Group Configuration Examples**.

Installation Procedure

In the following example, **sdrs_linux_amd64_23.12.0.20240103211150.tar.gz** is the package (23.12.0) used for gateway installation.

Step 1 Obtain the disaster recovery gateway package and upload it to a directory on the target server.

- **IDC-to-cloud**: Manually upload the gateway package to the target server.
- **Cross-region** and **Cross-AZ**: Copy the command provided on the console, go to the directory where you want to install the gateway, and paste and run the command to obtain the package.

Step 2 Log in to the server, go to the directory containing the gateway package, and run the following command as user **root** to decompress the package:

tar -zxvf sdrs_linux_amd64_23.12.0.20240103211150.tar.gz

Step 3 Go to the directory containing the installation script.

cd sdrs_linux_amd64_23.12.0.20240103211150

Step 4 Install the gateway.

sh install.sh --drm-ip=*drm ip* **--dra-ip=***dra ip* **--role=gateway**

In the command, *drm ip* and *dra ip* indicate the primary NIC IP addresses of server where the disaster recovery gateway is deployed.

If the command output contains the following information, the gateway has been installed:

```
...
Installed DRM successfully.
Installed SDRS successfully.
...
```

Step 5 Check whether the gateway is enabled.

ps -ef | grep java | grep drm

Information similar to the following is displayed:

```
service 3806 1 7 Aug31 ? 01:13:29 /opt/cloud/sdrs/drm/tools/jre/bin/java -Djava.security.egd=file:/dev/./
urandom -jar /opt/cloud/sdrs/drm/drm-20.8.0.jar --service.kernel.security.wcc.config_path=file:/opt/cloud/
sdrs/drm/classes/ --spring.config.location=/opt/cloud/sdrs/drm/classes/application.properties
```

If the command output contains the **drm** process, the gateway has been enabled.

Step 6 Run the following command in the **/opt/cloud/sdrs** directory to install the disaster recovery gateway:

sh register_gateway.sh

Figure 3-1 Executing the script

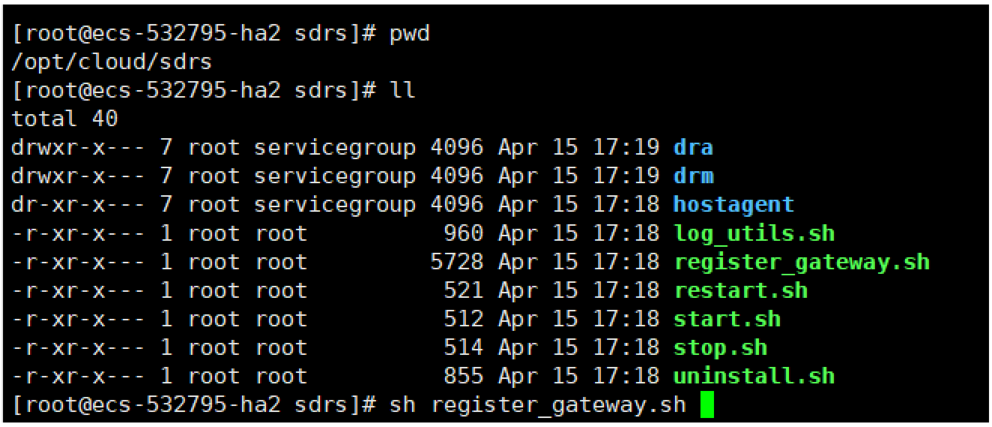

```
[root@ecs-532795-ha2 sdrs]# pwd
/opt/cloud/sdrs
[root@ecs-532795-ha2 sdrs]# ll
total 40
drwxr-x--- 7 root servicegroup 4096 Apr 15 17:19 dra
drwxr-x--- 7 root servicegroup 4096 Apr 15 17:19 drm
dr-xr-x--- 7 root servicegroup 4096 Apr 15 17:18 hostagent
-r-xr-x--- 1 root root          960 Apr 15 17:18 log_utils.sh
-r-xr-x--- 1 root root         5728 Apr 15 17:18 register_gateway.sh
-r-xr-x--- 1 root root          521 Apr 15 17:18 restart.sh
-r-xr-x--- 1 root root          512 Apr 15 17:18 start.sh
-r-xr-x--- 1 root root          514 Apr 15 17:18 stop.sh
-r-xr-x--- 1 root root          855 Apr 15 17:18 uninstall.sh
[root@ecs-532795-ha2 sdrs]# sh register_gateway.sh
```

Figure 3-2 Script execution in progress

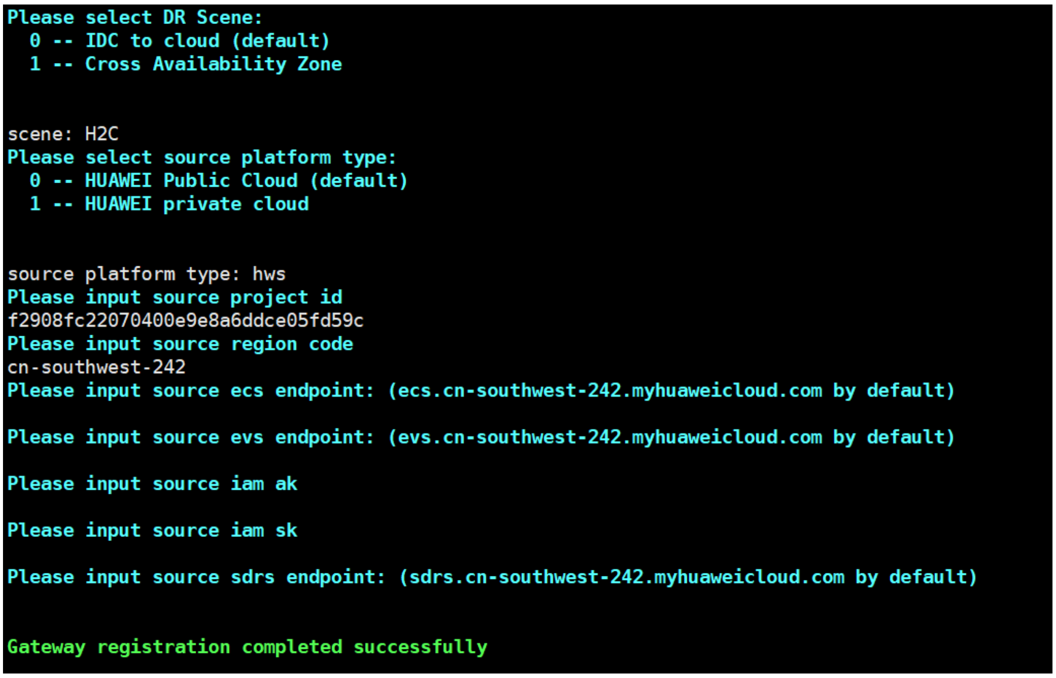

```
Please select DR Scene:
  0 -- IDC to cloud (default)
  1 -- Cross Availability Zone


scene: H2C
Please select source platform type:
  0 -- HUAWEI Public Cloud (default)
  1 -- HUAWEI private cloud


source platform type: hws
Please input source project id
f2908fc22070400e9e8a6ddce05fd59c
Please input source region code
cn-southwest-242
Please input source ecs endpoint: (ecs.cn-southwest-242.myhuaweicloud.com by default)

Please input source evs endpoint: (evs.cn-southwest-242.myhuaweicloud.com by default)

Please input source iam ak

Please input source iam sk

Please input source sdrs endpoint: (sdrs.cn-southwest-242.myhuaweicloud.com by default)


Gateway registration completed successfully
```

Table 3-1 describes the variables used during script execution.

Table 3-1 Parameter description

Site	Parameter	Description	How to Obtain	Example Value
Replication	DR Scene	Replication scenario	● **0**: IDC-to-cloud ● **1**: Cross-AZ	0 (default value)
Production site	source platform type	Type of the production site	● **0**: Huawei Cloud public cloud ● **1**: Huawei Cloud private cloud	0 (default value)
	source project id	Project ID	Log in to the console and choose **My Credentials** > **API Credentials** to view the project ID.	51af777371904892a49a0c3e3e53de44
	source region code	Destination region ID	Obtain the SDRS endpoint by referring to **SDRS Endpoints**.	sdrs.cn-east-2.myhuaweicloud.com
	source ecs endpoint	ECS endpoint	Obtain the ECS endpoint by referring to **ECS Endpoints**.	-
	source evs endpoint	EVS endpoint	Obtain the EVS endpoint by referring to **EVS Endpoints**.	-

Site	Parameter	Description	How to Obtain	Example Value
	source iam ak	Access key ID	Obtain AK/SK by referring to **How Do I Obtain an Access Key (AK/SK)?**	-
	source iam sk	Secret access key		-
Disaster recovery site	target platform type	Type of the disaster recovery site	● **0**: Huawei Cloud public cloud ● **1**: Huawei Cloud private cloud	0 (default value) -
	target project id	Project ID	Log in to the console and choose **My Credentials** > **API Credentials** to view the project ID.	0605767cb280d5762fd6c0133d6bea3f
	target sdrs endpoint	SDRS endpoint	Obtain the SDRS endpoint by referring to **SDRS Endpoints**.	sdrs.cn-east-2.myhuaweicloud.com
	target iam ak	Access key ID	Obtain AK/SK by referring to **How Do I Obtain an Access Key (AK/SK)?**	RZSAMHULWKKE71N0XHUT
	target iam sk	Secret access key		K7bXplAT0pEpy4SAiN2fHUwEtxvgmK3IqyhqnMTA

Step 7 After the installation is complete, delete the installation package and decompressed files.

----End

Upgrade Procedure

In the following example, **sdrs_linux_amd64_23.12.0.20240103211150.tar.gz** is the package (23.12.0) used for gateway upgrade.

Step 1 Obtain the disaster recovery gateway package and upload it to a directory on the target server.

- **IDC-to-cloud**: Manually upload the gateway package to the target server.
- **Cross-region** and **Cross-AZ**: Copy the command provided on the console, go to the directory where you want to install the gateway, and paste and run the command to obtain the package.

Step 2 Log in to the server, go to the directory containing the gateway package, and run the following command as user **root** to decompress the package:

```
tar -zxvf sdrs_linux_amd64_23.12.0.20240103211150.tar.gz
```

Step 3 Go to the directory containing the upgrade script.

cd sdrs_linux_amd64_23.12.0.20240103211150

Step 4 Upgrade the gateway.

sh upgrade.sh

If the command output contains the following information, the gateway has been upgraded:

```
...
Upgrade SDRS successfully.
```

----End

3.2 Installing and Upgrading the Proxy Client

Scenarios

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

Prerequisites

- If the firewall is enabled on the server you want to deploy the proxy client, enable port 59526 on the firewall.
- Because communication between the disaster recovery gateway and proxy client is not encrypted, it is recommended that you deploy them in the same security group and only allow ECSs within the security group to communication with each other. For details, see **Security Group Configuration Examples**.

Installation Procedure

In the following example, **sdrs_linux_amd64_23.12.0.20240103211150.tar.gz** is the proxy client installation package (23.12.0) used for CentOS.

Step 1 Obtain the proxy client package and upload it to a directory on the target server. Ensure the package integrity by comparing the sha256 value in advance.

- **IDC-to-cloud**: Manually upload the proxy client package to the target server.
- **Cross-region** and **Cross-AZ**: Copy the command provided on the console, go to the directory where you want to install the proxy client, and paste and run the command to obtain the package.

Step 2 Log in to the server, go to the directory containing the proxy client package, and run the following command as user **root** to decompress the package:

tar -zxvf sdrs_linux_amd64_23.12.0.20240103211150.tar.gz

Step 3 Go to the directory containing the installation script.

cd sdrs_linux_amd64_23.12.0.20240103211150

Step 4 Install the proxy client.

```
sh install.sh --hostagent-ip=hostagent ip --drm-ip=drm ip --role=all
```

In the preceding command, *hostagent ip* indicates the IP address of the proxy client. Set *hostagent ip* to the IP address of the primary NIC of the server you want to install the proxy client. *drm ip* indicates the IP address of the management gateway.

If the command output contains the following information, the proxy client has been installed:

```
...
Installed SDRS successfully.
...
```

Step 5 After the installation is complete, delete the installation package and decompressed files.

----End

Upgrade Procedure

NOTE

If the production services are running at the production site, upgrading the proxy client will resynchronize data.

In the following example, **sdrs_linux_amd64_23.12.0.20240103211150.tar.gz** is the proxy client upgrade package (23.12.0) used for CentOS.

Step 1 Obtain the proxy client package and upload it to a directory on the target server.

- **IDC-to-cloud**: Manually upload the proxy client package to the target server.
- **Cross-region** and **Cross-AZ**: Copy the command provided on the console, go to the directory where you want to install the proxy client, and paste and run the command to obtain the package.

Step 2 Log in to the server, go to the directory containing the proxy client package, and run the following command as user **root** to decompress the package:

```
tar -zxvf sdrs_linux_amd64_23.12.0.20240103211150.tar.gz
```

Step 3 Go to the directory containing the upgrade script.

```
cd sdrs_linux_amd64_23.12.0.20240103211150
```

Step 4 Upgrade the proxy client.

```
sh upgrade.sh
```

If the command output contains the following information, the proxy client has been upgraded:

```
...
Upgrade SDRS successfully.
```

----End

3.3 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 switchover.

Figure 3-3 Disaster recovery drill

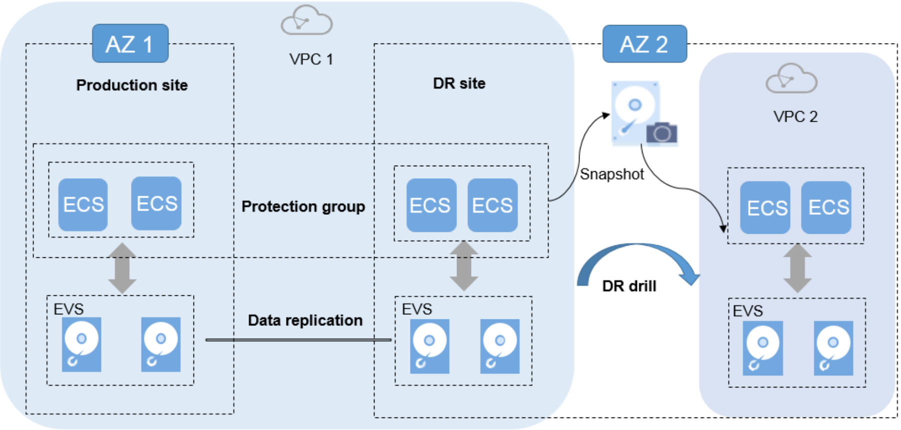

Precautions

- 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**, **Switchover 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 3-2 Parameter description

Parameter	Description	Example Value
Name	DR drill name	DR drill servername

Parameter	Description	Example Value
Drill VPC	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**. ● **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. **NOTE** The drill VPC cannot be the same as the VPC of the protection group.	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

3.4 Port Description (Asynchronous Replication)

Table 3-3 DR gateway port description

Port	Protocol	Description
29210	TCP	Used to communicate with proxy clients.
29211	TCP	Used to receive control commands.
7443	tcp	Used for API communication.

Table 3-4 Production and DR site server port description

Port	Protocol	Description
8091	TCP	Used to transfer messages between proxy clients.
59526	TCP	Used to communicate with the DR gateway.
29210	TCP	The local listening port used to communicate with proxy clients after a failover.
29211	tcp	The local listening port used to receive control commands after a failover.
7443	TCP	The local listening port used for API communication after a failover.

4 Change History

Released On	Description
2023-04-07	This issue is the fifth official release. Added the following content: Added a recommendation that the production site servers used to install the disaster recovery gateway and proxy client be deployed in the same security group in **Configuration Process**.
2021-09-25	This issue is the fourth official release. Added the following content: **Asynchronous Replication**.
2020-04-29	This issue is the third official release. Modified the following content: Modified restrictions in **Create a Protected Instance**. Specifically, shared disks are supported.
2019-09-30	This issue is the second official release. Modified the following content: Added descriptions about DeHs in **Create a Protected Instance**.
2019-05-24	This issue is the first official release.