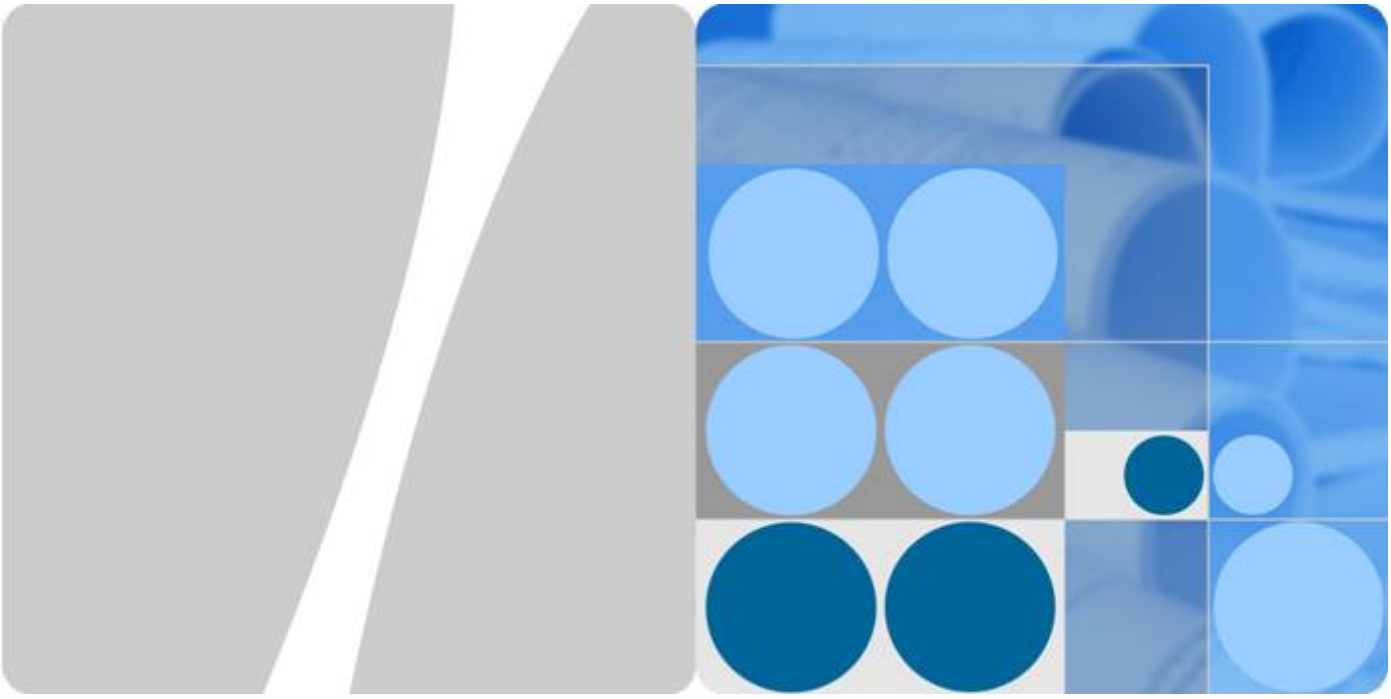

Storage Disaster Recovery Service

FAQs

Date **2019-08-28**

Contents

1 What Is DR?..... 1
2 What Are the Highlights of SDRS?..... 2
3 What Are RPO and RTO?..... 4
4 What Are the Differences Between DR and Backup?..... 5
5 How Am I Billed for SDRS?..... 6
6 How Do I Use a Resource Package?..... 7
7 Do I Need to Manually Create DR Resources?..... 8
8 What Can I Do When a Disk Detached from a DR Site Server Cannot Be Attached to Another Server?..... 9
9 What Can I Do When the EIP Cannot Be Pinged After I Perform a Planned Failover for a Protection Group Containing a SUSE Server?..... 12
10 What Can I Do If the NIC Names of the DR Drill Server and Production Site Server Are Different?..... 14
11 What Can I Do If hostname of the Production Site Server and DR Site Server Are Different After a Planned Failover or Failover?..... 16
12 Why NICs of DR Site Servers Are Not Displayed After I Perform a Failover?..... 18
13 What Are the Precautions If the Production Site Server Uses the Key Login Mode?..... 19
14 Change History..... 20

1 What Is DR?

Disaster recovery (DR) is a broad concept. In a broad sense, disaster recovery is a systematic project that includes all contents related to service continuity. Regarding IT, DR provides a computing system that protects user service systems from various disasters. DR is the ability to provide continuous services after an accident, such as a fire, earthquake, or power outage. This is achieved by setting up two or more IT systems with the same functions in dispersed areas. If one system stops, the services can quickly recover on the other system and continue to run properly.

The purpose of DR is to ensure maximum service continuity in the event of a natural or human-caused disaster on production systems.

2 What Are the Highlights of SDRS?

SDRS provides the following features:

- Convenient service recovery

  SDRS provides a management console. You can configure and manage server replication and perform a planned failover or failover.

- server replication

  You can establish a replication relationship between the production site and the DR site.

- Replication on demand

  You can replicate servers in an AZ to another AZ as required, thereby reducing the costs and complexity of maintaining another data center.

- Zero impact on applications

  Applications running on servers can be replicated, and the replication will not have any impact on the applications.

- High RTO and RPO standard

  For SDRS, recovery time objective (RTO) refers to the period from the time when users perform a planned failover or failover at the production site to the time when the servers at the DR site start to run. This period does not include any time for DNS configuration, security group configuration, or customer script execution, and is within 30 minutes.

  SDRS provides continuous and synchronous replication for servers to ensure zero recovery point objective (RPO).

- Crash consistency

  Real-time data synchronization based on storage ensures that data synchronized across two AZs maintains crash consistency. Specifically, the application data across AZs may not be consistent but the disk data across the two AZs remains consistent.

- Seamless DR drills

  You can easily perform DR drills without affecting ongoing replication.

- Flexible failover

  You can perform a planned failover for an expected service interruption to prevent data loss, or perform a failover for unexpected failures to restore services quickly.

- Efficient network switchover

  SDRS simplifies program resource management during failovers, including reserving IP addresses and MAC addresses, all of which facilitates efficient network switchovers.

- Excellent performance for the price

  When services are running properly, servers at the DR site are stopped and thereby will not be billed. This greatly reduces the DR TCO.

- Simple deployment

  servers do not need the DR agent plug-in. The service deployment is simple and quick.

3 What Are RPO and RTO?

Recovery Point Objective (RPO): the maximum data loss amount tolerated by the system.

Recovery Time Objective (RTO): the maximum service interruption duration tolerated by the system. It refers to the requirement for the recovery duration of an information system failure or service function failure caused by a disaster.

Figure 3-1 RPO and RTO

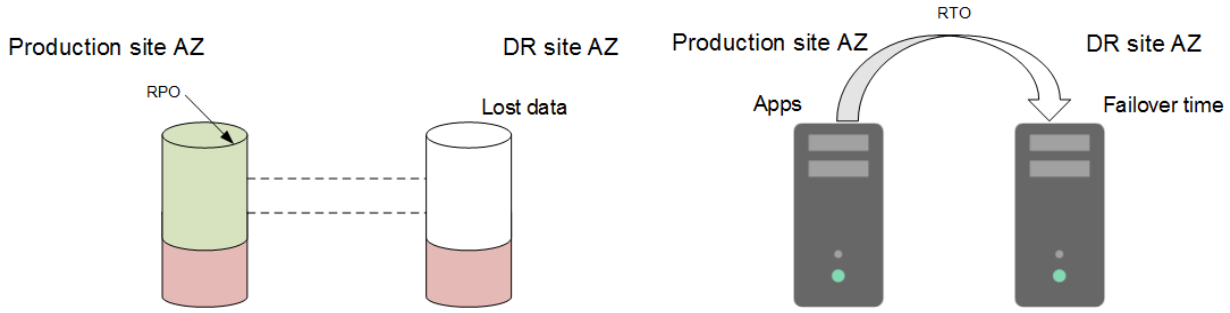

4 What Are the Differences Between DR and Backup?

Differences between disaster recovery (DR) and backup are as follows:

- DR is used to prevent impacts on the systems caused by natural disasters, such as fires and earthquakes. A production site and its DR site must be located with a certain secure distance. Backup is to prevent impacts on the systems caused by inappropriate manual operations, virus infection, and logic errors. It is used to restore the service system data. Usually, a system and its backup are deployed in the same data center.
- A DR system protects data but more focuses on protecting service continuity. A data backup system only ensures that data generated at different time points can be restored. Generally, the system performs the full backup for the first time, which takes a long period of time. The subsequent backup is incremental and can be completed within a short period of time.
- The highest DR standard is to implement zero RPO. You can set a maximum of 24 automatic backup policies at different time points in one day to restore data to different backup points.
- If a disaster occurs, such as earthquakes or fires, a DR system takes only several minutes to perform a failover, but a backup system takes several hours or even dozens of hours to restore the data.

5 How Am I Billed for SDRS?

Billing Mode

- The bill for SDRS is for the service only and does not include that for the protected resources, such as the servers and disks.
- SDRS is billed by number of protected instances and usage duration (hour). SDRS supports two billing modes: pay-per-use and resource package. For details, see **Product Pricing Details**.
  - Billing in pay-per-use mode: Price of a single protected instance = Usage duration (hour) x Unit price
  - Billing by resource package: It is similar to the traffic package for mobile phone users. Each resource package contains 744-hour availability duration. For details, see **How Do I Use a Resource Package?**.
- For resources at the production site created by users, the billing mode uses the mode selected by the users. For resources on the DR site created by SDRS, the servers and disks are billed in pay-per-use mode. Servers created at the DR site are stopped by default. For details about the billing information for the resources, see the price details of each resource.

Billing After a Planned Failover or Failover

After users perform a planned failover or failover, the billing mode of the protected resources will not change, even if the production site and DR site have changed.

Billing Mode for DR Drills

During a DR drill, SDRS automatically create DR drill resources. The created resources will be billed according to the target service price standard. No extra cost is generated for the DR drill.

6 How Do I Use a Resource Package?

- A resource package is similar to a mobile phone traffic package. A one-month resource package contains 744 hours' validity period. A resource package can be used by multiple protection instances. For example, if you purchase a one-month resource package and create two protected instances, the resource package can provide 372 hours' validity period for each protected instance.
- After you purchase a resource package, the system will deduct the total usage duration of the protected instances from the resource package by default based on the number of protected instances and the usage duration of the protected instances. When the validity period in the resource package is used up, the system will bill the used resources in pay-per-use mode.
- The validity period of a yearly resource package is reset by month. If the usage duration of the protected instances in one month is shorter than the validity period of the resource package, the unused duration will expire in next month. If you exceed the usage duration included in the package, you will be billed on a pay-per-use basis for the excess usage duration. After a month ends, the resource package resets.
- You are recommended to purchase resource packages based on the number of protected instances.

7 Do I Need to Manually Create DR Resources?

When you use SDRS to create a protected instance, the resources at the production site are the existing resources that you have manually created. SDRS will automatically create the resources required by the DR site. For resources on the DR site created by SDRS, the servers and disks are billed in pay-per-use mode. Servers created at the DR site are stopped by default. For details about the billing information for the resources, see the price details of each resource.

8 What Can I Do When a Disk Detached from a DR Site Server Cannot Be Attached to Another Server?

Symptom

A user has created a protection group and protected instance in the East China region. **Table 8-1** lists the details.

Table 8-1 Protection group details

	AZ	Specifications	Data Disk Name
Production site	AZ2	c3.large.2	volume-0001
DR site	AZ1	c3.large.2	volume-0001

The user deletes this protected instance but does not delete the DR site server in the protected instance. Then, the user switches to the server page and detach the disks from the DR site server, and attempts to attach the disks to another server in AZ1. The user finds that some disks can be attached successfully, but some failed.

Restrictions and Limitations

This section applies to only the East China and North China regions.

Possible Causes

The underlying resource pool of AZ2 in East China supports only KVM. The underlying resource pool of AZ1 in East China supports both Xen and KVM. This means that the underlying resource pool of the production site AZ supports only KVM, and that of the DR site AZ supports both Xen and KVM.

Therefore, after the disks are detached from the DR site server, the disks can be attached to only KVM servers (KVM instances).

Handling Method

Query the virtualization type of the server in the DR site AZ and attached the detached disk to a KVM instance. To do so, perform the following operations:

Step 1 Query the virtualization type of the target server.

1. Log in to the management console and click **Elastic Cloud Server** under **Computing**.
2. Use the search function above the server list to search for the target server.

   Figure 8-1 Searching for the target server

   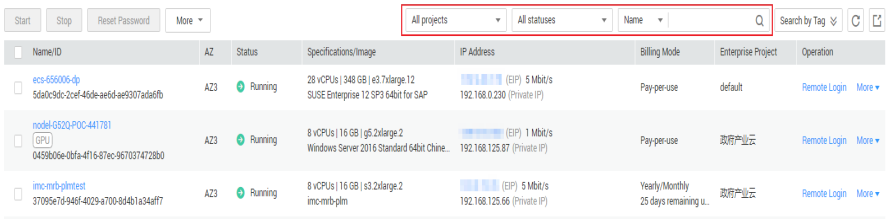

3. In the **Specifications/Image** column, query the specifications of the target server.

   Figure 8-2 Querying the specifications of the target server

   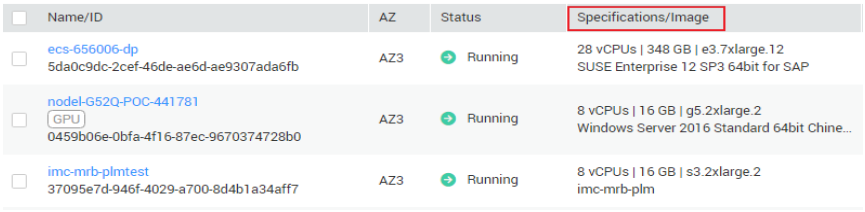

4. Check the virtualization type of the server based on the **specifications list**.

   For example, the virtualization type of s3.large.2 is KVM, then servers of the s3.large.2 specifications are KVM instances.

   NOTE

   The current DR site AZ of the protection group supports both KVM and Xen, but the disk attached to the DR site server supports only KVM. Therefore, this disk can be attached to only KVM instances.

Step 2 Attach the disk.

1. Log in to the management console and click **Elastic Volume Service** under **Storage**.

   The EVS page is displayed.
2. In the disk list, select the disk to be detached from the DR site server.

   The disk status is **Available**.
3. Click **Attach** in the **Operation** column.

Figure 8-3 Attaching the disk

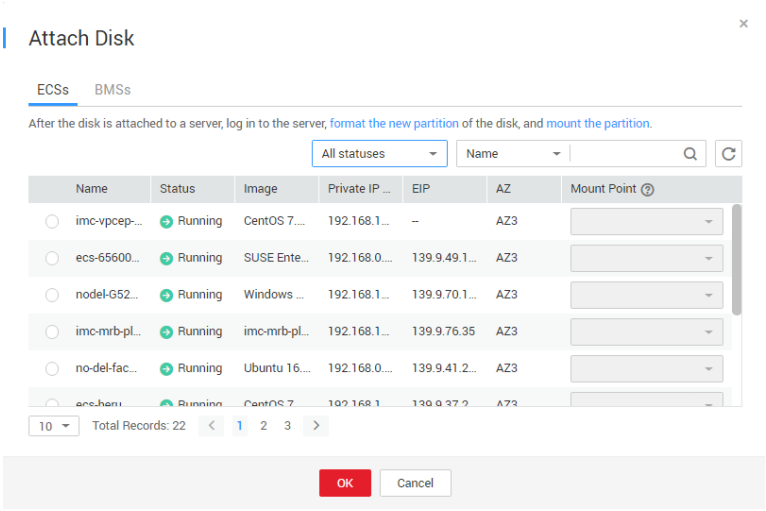

4. Select the KVM instance whose virtualization type has been verified in **Step 1**, and click **OK**.

----End

9 What Can I Do When the EIP Cannot Be Pinged After I Perform a Planned Failover for a Protection Group Containing a SUSE Server?

Symptom

The production site server in a protection group runs the SUSE OS. After users enable protection and perform a planned failover for the protection group, the EIP of the server cannot be pinged.

Root Cause

After the planned failover, the server NIC name may already change. If the NIC has an EIP bound, the EIP cannot be pinged.

Handling Method

After the planned failover, delete the **/etc/udev/rules.d/70-persistent-net.rules** file on the DR site server and then restart it. The procedure is as follows.

Step 1 Log in to the DR site server.

1. Log in to the management console and click **Elastic Cloud Server** under **Computing**.
2. In the server list, select the DR site server.
3. Locate the row containing the server and click **Remote Login** in the **Operation** column.

   Log in to the server as prompted.

Step 2 Run the following command to delete the file:

rm /etc/udev/rules.d/70-persistent-net.rules

Step 3 Run the following command to restart the DR site server:

reboot

----End

10 What Can I Do If the NIC Names of the DR Drill Server and Production Site Server Are Different?

Symptom

The production site server runs the SUSE OS. After users create a DR drill using this server, the NIC names of the DR drill server are different from those of the production site server.

The following is an example:

A production site server running Novell SUSE Linux Enterprise Server 12 SP3 64-bit has five NICs attached. Log in to the production site server and query the NIC names (eth0 to eth4).

Figure 10-1 Production site server NIC names

```
wyhtest-0001:~ # ifconfig -a|grep eth
eth0      Link encap:Ethernet  HWaddr FA:16:3E:00:F1:5F
eth1      Link encap:Ethernet  HWaddr FA:16:3E:D4:9C:AD
eth2      Link encap:Ethernet  HWaddr FA:16:3E:1E:3A:29
eth3      Link encap:Ethernet  HWaddr FA:16:3E:83:B0:9E
eth4      Link encap:Ethernet  HWaddr FA:16:3E:D7:B5:45
wyhtest-0001:~ #
```

Log in to the DR drill server and query the NIC names (eth5 to eth9).

Figure 10-2 DR drill server NIC names

```
wyhtest-0001:~ #
wyhtest-0001:~ # ifconfig -a|grep eth
eth5      Link encap:Ethernet  HWaddr FA:16:3E:D4:9C:AD
eth6      Link encap:Ethernet  HWaddr FA:16:3E:1E:3A:29
eth7      Link encap:Ethernet  HWaddr FA:16:3E:D7:B5:45
eth8      Link encap:Ethernet  HWaddr FA:16:3E:83:B0:9E
eth9      Link encap:Ethernet  HWaddr FA:16:3E:00:F1:5F
```

The NIC names of the DR drill server are different from those of the production site server.

Root Cause

The NIC names may change when users create a DR drill.

Handling Method

After the DR drill, delete the **/etc/udev/rules.d/70-persistent-net.rules** file on the DR drill server and then restart it. The procedure is as follows.

Step 1 Log in to the DR drill server.

1. Log in to the management console and click **Elastic Cloud Server** under **Computing**.
2. In the server list, select the DR drill server.
3. Locate the row containing the server and click **Remote Login** in the **Operation** column.

   Log in to the server as prompted.

Step 2 Run the following command to delete the file:

rm /etc/udev/rules.d/70-persistent-net.rules

Step 3 Run the following command to restart the DR drill server:

reboot

----End

11 What Can I Do If hostname of the Production Site Server and DR Site Server Are Different After a Planned Failover or Failover?

Symptom

Users have changed **hostname** of the production site server before they perform a planned failover or failover for the first time. After the planned failover or failover, users start the DR site server and find that **hostname** of the DR site server is not updated accordingly.

Possible Causes

For Linux servers, if you have changed **hostname** of the production site server before you perform a planned failover or failover for the first time, this modification will not synchronize to the DR site server.

Prerequisites

- The production site server is a Linux server with Cloud-Init installed.
- **hostname** of the production site server has been changed.

Solution 1 (If You Have Not Performed a Planned Failover or Failover)

Set **preserve_hostname: false** to **preserve_hostname: true** in the Cloud-Init configuration file **/etc/cloud/cloud.cfg** to ensure that **hostname** of the production site server and DR site server are the same after you enable protection.

The procedure is as follows:

1. Log in to the production site server.
2. Run the following command to edit the **/etc/cloud/cloud.cfg** configuration file:

   sudo vim /etc/cloud/cloud.cfg
3. Modify **preserve_hostname**.

- If **preserve_hostname: false** is already available in the **/etc/cloud/cloud.cfg** configuration file, change it to **preserve_hostname: true**.
- If **preserve_hostname: false** is unavailable in the **/etc/cloud/cloud.cfg** configuration file, add **preserve_hostname: true** before **cloud_init_modules**.

4. Perform a planned failover or failover.

   After the planned failover or failover, **hostname** of the DR site server and production site server are the same.

Solution 2 (If You Have Performed a Planned Failover or Failover)

If you have not modified configuration file **/etc/cloud/cloud.cfg** before the planned failover or failover, log in to the DR site server and manually change **hostname** of the DR site server to that of the production site server.

12 Why NICs of DR Site Servers Are Not Displayed After I Perform a Failover?

SDRS performs planned failovers or failovers for protection groups. After a planned failover or failover, the production site and DR site exchange. When the production site is functional, users can perform a planned failover according to plans. When a fault occurs at the production site, users need to perform a failover.

- During a planned failover, NICs of the production site servers and DR site servers will exchange. In this way, servers can be accessed using the same IP addresses and MAC addresses after the planned failover.
- During a failover, production site servers do not work properly, NICs of the production site servers will be migrated to the DR site servers. The primary NICs of the DR site servers will enter the to-be-used state. Therefore, after a failover, the original production site servers will not have NICs. When users perform reprotection for the protection group after the original production site servers restore, the primary NICs in the to-be-used state detached from the original DR site servers will be attached to the original production site servers. Then, the NICs of the production site servers and DR site servers are exchanged.

13 What Are the Precautions If the Production Site Server Uses the Key Login Mode?

- If the production site server runs Windows and uses the key login mode, you need to ensure that the key pair of the production site server exists when you create a protected instance or a DR drill. Otherwise, the DR site server or DR drill server may fail to create, causing the protected instance creation or DR drill creation failure.
- If the production site server runs Linux and uses the key login mode, you can create a protected instance and DR drill no matter whether the key pair of the production site server exists. However, the key pair information will not be displayed on the details page of the created DR site server or DR drill server. You can use the key pair of the production site server to log in to the created DR site server and DR drill server.

14 Change History

Released On	Change History
2019-08-28	This issue is the sixth official release. Added the following content: • **5 How Am I Billed for SDRS?** • **7 Do I Need to Manually Create DR Resources?** • **6 How Do I Use a Resource Package?** • **12 Why NICs of DR Site Servers Are Not Displayed After I Perform a Failover?**
2019-06-25	This issue is the fifth official release. Added the following content: • **11 What Can I Do If hostname of the Production Site Server and DR Site Server Are Different After a Planned Failover or Failover?**

Released On	Change History
2019-04-22	This issue is the fourth official release. Deleted the following content: • What Are the Standards of DR? • What Can I Do If an Error Occurred When I Expand the Replication Pair Capacity? • What Can I Do When My Protection Group Quotas Are Insufficient? Modified the following content: • **8 What Can I Do When a Disk Detached from a DR Site Server Cannot Be Attached to Another Server?** • **9 What Can I Do When the EIP Cannot Be Pinged After I Perform a Planned Failover for a Protection Group Containing a SUSE Server?** • **10 What Can I Do If the NIC Names of the DR Drill Server and Production Site Server Are Different?**
2019-03-30	This issue is the third official release. Added the following content: • **10 What Can I Do If the NIC Names of the DR Drill Server and Production Site Server Are Different?**
2019-01-30	This issue is the second official release. Modified the following content: • Deleted concepts of the source end, target end, preferred end, and added the concepts of the production site and DR site.
2018-05-30	This issue is the first official release.