Storage Disaster Recovery Service

FAQs

Issue 02

Date 2019-06-25

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

Trademarks and Permissions


and other Huawei trademarks are the property 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 Cloud 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.

Huawei Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road
Qianzhong Avenue
Gui'an New District
Gui Zhou 550029
People's Republic of China

Website: https://www.huaweicloud.com/intl/en-us/

Contents

1 Common Problems

1.1 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.

1.2 What Are the SDRS Advantages?

SDRS has the following advantages:

- Convenient recovery solution

  Using the SDRS console, you can configure and manage server replication and perform failovers and drills.
- Site server replication

  You can set up disaster recovery for site servers from the production site to the disaster recovery site.
- Replication on demand

  You can replicate servers from one AZ to another as required, reducing the costs and complexity for you to maintain another data center.

- Zero impact on applications

  You can replicate all applications on the servers. The replication has no impact on the applications.

- RPO target

  SDRS provides asynchronous replication for ECSs, and the recovery point object (RPO) is in seconds.

- RTO target

  Normally, the recovery time objective (RTO) is within 30 minutes, which does not include the time spent on DNS configuration, security group configuration, or customer script execution.

- Crash consistency

  Host-layer asynchronous replication ensures crash consistency between your production site and disaster recovery site. (SDRS only ensures crash consistency, not application consistency.)

- Disaster recovery drill

  By running disaster recovery drills, you can simulate recovery fault scenarios and formulate recovery plans. When a fault occurs, you can use the plans to recover services as quickly as possible.

- Flexible failover

  If the production site fails, you can fail over to the disaster recovery site in just a few clicks (purchasing, deploying, and starting disaster recovery servers and attaching disks with the most current data). Services can be recovered with only a few configurations.

- Cost-effective: When production site services are running properly, servers at the disaster recovery site are not created. You pay only for the disaster recovery site EVS disks and the Object Storage Service (OBS) buckets used.
- Simple deployment: Agent can be installed online without interrupting production services. The deployment is simple and fast.

1.3 What Are RPO and RTO?

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

- SDRS asynchronous replication is based on the continuous asynchronous replication on the host side. Normally, the RPO is not zero (RPO < 1 minute).

  This RPO can be achieved when the following conditions are met:

  a. Actual network bandwidth > Changed data volume per minute in peak hours/1 minute. Minimum bandwidth: 10 Mbit/s
  b. Network latency ≤ 100 ms; packet loss rate < 0.1%
  c. If the initial synchronization or differential data synchronization is not complete for a protected instance, this RPO is not guaranteed.
  d. During a planned switchover, ensure that the services and OS data in the cache has been flushed to disks.

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.

- With SDRS replication, RTO refers to the period from the time when you perform a planned or an unplanned failover to the time when ECSs are started on Huawei Cloud. Note that this period does not include the time that you spent on manual configuration, operations, network configuration, and script execution to try to start services. Normally, the RTO is not zero (RTO < 30 minutes).

  This RTO can be achieved when the following conditions are met:

  a. A failover is performed after the initial synchronization is complete.
  b. Your account has sufficient resource quotas, including ECS, EVS, and VPC quotas.
  c. The RTO of SDRS asynchronous replication does not include the time that you spent on manual configuration and service configuration.

Figure 1-1 RPO and RTO

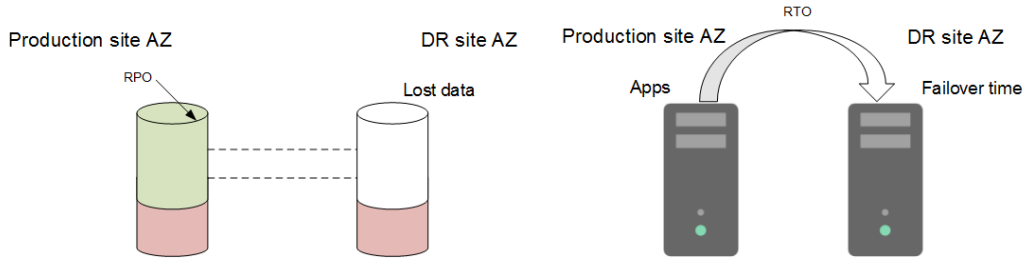

1.4 What Are the Differences Between DR and Backup?

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

- DR protects data centers against hardware faults or natural disasters, so it requires a safe distance (intra-city or remote) between the production site and disaster recovery site. Backups are used to restore data in the event of unintended actions, virus attacks, or logic errors. Backups are usually stored in the same data center as the service system.
- A DR system protects data but focuses more on protecting service continuity. A data backup system only ensures that data generated at different time points can be restored. Normally, a full backup is performed for the first time, which takes a long period of time. Subsequent backups are all incremental backups and can be done quicker.
- Disaster recovery can help you achieve an RPO of a few seconds. Backup allows you to set a backup policy to back up at up to 24 time points in one day, so you can restore data to different backup points.
- If a disaster occurs, such as earthquakes, fires, or dat center failure, a disaster recovery system takes only minutes to perform a failover, but a backup system takes hours or even dozens of hours to restore the data.

1.5 How Am I Billed for SDRS?

Billing Modes

- Billing for SDRS is for the service only and does not cover protected resources such as ECSs and EVS disks.

- Billing modes of production site resources purchased by you will not change. Billing modes of DR site resources created by SDRS will be billed based on a pay-per-use basis. See the ECS and EVS pricing details for more information.

Billing After Switchovers

During a switchover, disaster recovery site servers are created and billed based on the ECS billing standards.

Billing for Disaster Recovery Drills

SDRS automatically creates drill resources, which are billed based on a pay-per-use basis. No additional drill pricing applies.

2 Synchronous Replication (for Installed Base Operations)

2.1 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.

2.2 What Can I Do When the EIP Cannot Be Pinged After I Perform a Switchover 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 switchover for the protection group, the EIP of the server cannot be pinged.

Root Cause

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

Handling Method

After the switchover, 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

2.3 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 2-1 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 2-2 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

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

Symptom

Users have changed **hostname** of the production site server before they perform a switchover or failover for the first time. After the switchover 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 switchover 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 Switchover 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 switchover or failover.

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

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

If you have not modified configuration file **/etc/cloud/cloud.cfg** before the switchover 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.

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

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

- During a switchover, 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 switchover.
- 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.

2.6 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, 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 a key pair for login, you can create a protected instance or DR drill regardless of whether the key pair exists or not. 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.

2.7 What Should I Pay Attention to When Logging In to the Server After the First Time Ever I Executed a Switchover, Failover, or DR Drill?

After you have performed a switchover, failover, or DR drill for the first time:

- If your servers are installed with Cloud-Init/Cloudbase-Init, Cloud-Init/ Cloudbase-Init will start along with the server's first startup to inject the

initial data. In this case, the password or key pair used to log in to the production site server, disaster recovery site server, or drill server will change.

- If your servers are not installed with Cloud-Init/Cloudbase-Init, the password or key pair used to log in to the production site server, disaster recovery site server, or drill server will not change.

The following uses a switchover or failover as the example operation. For the login constraints on drill servers, see those for DR site servers.

In the following example, Server A and server B are deployed. **Table 2-1** shows the servers before and after the operation.

Table 2-1 Servers before and after a switchover or failover

-	Production Site Server	Disaster Recovery Site Server
Before	Server A	Server B
After	Server B	Server A

Detailed login constraints are described as follows:

Scenario 1: Server A runs Windows and does not have Cloudbase-Init installed. After the first time switchover or failover:

- If your servers use password for login, you can use the password of Server A to log in to the production site server (Server B) or disaster recovery site server (Server A).
- If your servers use key pair for login, you can use the obtained password of Server A to log in to the production site server (Server B) or disaster recovery site server (Server A).

NOTE

After the first time switchover or failover, the password or key pair remains the same for the subsequent switchovers or failovers. In this example:

You can use the password of Server A to log in to the production site server or disaster recovery site server.

Scenario 2: Server A runs Windows and already has Cloudbase-Init installed. After the first time switchover or failover:

- When your servers use password for login,

  If Cloudbase-Init is not started (normally within 3 to 5 minutes after the production site server starts), you can use the password of Server B for login.

  After Cloudbase-Init is started, the login password of Server B becomes invalid. Reset the password and use the new password for login.
- When your servers use key pair for login,

  If Cloudbase-Init is not started (normally within 3 to 5 minutes after the production site server starts), you can use the obtained login password of Server B for login.

  After Cloudbase-Init is started, the obtained login password of Server B becomes invalid. Obtain the password again.

NOTE

After the first time switchover or failover, the password or key pair remains the same for the subsequent switchovers or failovers. In this example:

- Login using a password: Reset the password of Server B and use the new password for login.
- Login using a key pair: Obtain the password of Server B again and use the obtained password to log in to Server B.

Scenario 3: Server A runs Linux. After the first time switchover or failover:

- If your servers use password for login, you can use the password of Server A to log in to the production site server (Server B) or disaster recovery site server (Server A). Specifically:

  If the login password of Server A is not changed before the operation, use this password for login.

  If the login password of server A has been changed before the operation, use the new password for login.

  NOTE

  For ECS OSs other than CoreOS, the login password does not change after the first time switchover or failover.

  For ECSs running CoreOS, the login password of Server A will restore to the initial one after the first time switchover or failover. In this case, use the login password configured when Server A is created to log in to production site Server A or disaster recovery site Server B.

- If your server uses key pair for login, use the SSH key pair of Server A to log in to production site Server B or disaster recovery site Server A.

2.8 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.

3 Asynchronous Replication

3.1 How Do I Handle the drm Process Start Failure?

Symptom

The drm process failed to start after the disaster recovery gateway is deployed or proxy client is installed.

Root Cause

Possible causes are as follows:

- Cause 1: The **service** account does not have the write permission on the **/dev/null** directory.
- Cause 2: The server cannot resolve the **hostname** domain name.

Solution 1

Assign the read and write permissions on the **/dev/null** directory to the **service** account.

Step 1 Log in to the server where the disaster recovery gateway or proxy client resides.

Step 2 Run the following command as user **root** to change the permission on the **/dev/null** directory:

chmod 666 /dev/null

Step 3 Run the following command to check the permission on the **/dev/null** directory:

ll /dev/null

The permission is successfully changed if information similar to the following is displayed:


```
crw-rw-rw- 1 root root 1, 3 Apr  9 09:21 /dev/null
```

----End

Solution 2

Add the domain name resolution of **hostname**.

Step 1 Log in to the server where the disaster recovery gateway or proxy client resides.

Step 2 Run the following command as user **root** to add the domain name resolution of **hostname**:

echo "127.0.0.1 `hostname`" >> /etc/hosts

Step 3 Run the following command to check whether the domain name resolution of **hostname** takes effect:

ping `hostname`

The configuration is successful if information similar to the following is displayed:

```
PING test-dr (127.0.0.1) 56(84) bytes of data.64 bytes from localhost (127.0.0.1): icmp_seq=1 ttl=64
time=0.022 ms
```

----End

3.2 Failed to Install and Configure Disaster Recovery Gateway When Process drm Exists But Port 7443 Is Not Listened

Check the available system entropy on the VM.

```
cat /proc/sys/kernel/random/entropy_avail
```

If the returned value is less than 500, install haveged to ensure that the available entropy meets the encryption program's requirements. Insufficient system entropy will result in blocked threads, slow listening on port 7443, or port 7443 not listened even if the drm process is installed.

Solution

Download and install haveged to supplement entropy for **/dev/random**.

View the entropy value and compare it with the initial value.

```
cat /proc/sys/kernel/random/entropy_avail
```


4 Change History

Released On	Description
2021-09-25	This issue is the third official release. Added the following content: Added an FAQ in **3 Asynchronous Replication**.
2019-06-25	This issue is the second official release. Added the following content: • **2.4 What Can I Do If hostname of the Production Site Server and DR Site Server Are Different After a Switchover or Failover?**
2019-04-10	This issue is the first official release.

