**Business Recovery Service**

# FAQs

**Issue**      02
**Date**     2025-09-10

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

# 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 wide concept.

In a broad sense, DR is a systematic project covering all the content related to business continuity. In terms of IT, DR provides a computer 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 Advantages of BRS?

BRS has the following advantages:

- Convenient recovery solution

  Using the BRS 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

BRS provides asynchronous replication for servers. The recovery point object (RPO) is in seconds.

- RTO target

  Normally, BRS offers a recovery time objective (RTO) of no more than 30 minutes. This 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. (BRS 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 (creating, 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.

- BRS 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.
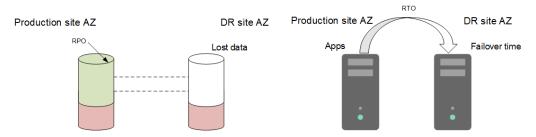
- In BRS replication, RTO refers to the period from the time when you perform a planned or an unplanned failover to the time when ECSs are started on Huawei Cloud. Note that this period does not include the time you spent on

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 BRS asynchronous replication does not include the time 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 data 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 BRS?

## Billing Modes

● Billing for BRS 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 BRS 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

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

# 1.6 What Are the Differences Between Synchronous Replication and Asynchronous Replication?

Asynchronous replication:

The DR solution is based on the **host layer** asynchronous data replication. The RPO is within seconds. Asynchronous replication supports the following DR scenarios: cross-region DR, cross-AZ DR, HCS Online DR to Huawei Cloud.

Synchronous replication:

The DR solution is based on the **storage layer** synchronous data replication. Synchronous replication supports only cross-AZ DR and has high requirements on the network. The sale of synchronous replication is **discontinued**. For details, **Discontinuation Notice**.

# 2 Synchronous Replication (for Installed Base Operations)

## 2.1 Do I Need to Manually Create DR Resources?

When you use BRS to create a protected instance, the resources at the production site are the existing resources that you have manually created. BRS will automatically create the resources required by the DR site.

For resources on the DR site created by BRS, 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.

### Possible Causes

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

### Solution

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.

## Possible Causes

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

## Solution

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 be automatically synchronized 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?

BRS allows you to perform switchovers or failovers by protection group. When the production site is working normally, you can perform a switchover according to plans. When a fault occurs at the production site, you 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.

### Possible Causes

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
```

# 3.3 What Can I Do If the Name of a Production Site Server or the Host Name Reported by the Gateway Is Incorrect and Always Displayed as "localhost"?

**Symptom**

The name of a production site server or the host name reported by the gateway was incorrect and always displayed as "localhost".

## Possible Causes

The host uses minimal installation, or the network or domain name service is not installed.

## Solution

- Do not use minimal installation when installing production site servers and gateways. Or, BRS may become abnormal due to lack of certain functions.
- If minimal installation needs to be used, manually add a line of the hostname and host IP address in the **/etc/hosts** file.

  Example: "192.168.0.1 sdrs-hostname"

# 3.4 What Can I Do If the Disaster Recovery Site VM Is Not Started After a Switchover?

## Possible Cause

BRS asynchronous replication supports crash consistency. During a switchover, data in the memory may be lost, so the disaster recovery site VM may need to use the crash recovery capability of the application or file system to restore data. In most cases, the restored data is consistent and can be normally used.

## Solution

Example: After a switchover, the disaster recovery site VM failed to be started. An XFS exception occurred (LSN verification failed), and the OS entered the emergency mode.



Solution: Run **journalctl** to check that the fault is caused by an XFS file system error. Enter the emergency mode of the OS and run **xfs_repair** to rectify the fault.

An operation example is provided as follows:

**Step 1** Check the root file system, for example **rhel-root**, in the **/dev/mapper** directory.

```
xfs_repair /dev/mapper/rhel-root
```

If an error is reported, check the error cause or run the following command:

```
xfs_repair –L /dev/mapper/rhel-root
```

Information similar to the following is displayed.

```
$ sudo xfs_repair /dev/mmcblk0
Phase 1 - find and verify superblock...
Phase 2 - using internal log
        - zero log...
        - scan filesystem freespace and inode maps...
        - found root inode chunk
Phase 3 - for each AG...
        - scan and clear agi unlinked lists...
        - process known inodes and perform inode discovery...
        - agno = 0
        - agno = 1
        - agno = 2
        - agno = 3
        - process newly discovered inodes...
Phase 4 - check for duplicate blocks...
        - setting up duplicate extent list...
        - check for inodes claiming duplicate blocks...
        - agno = 1
        - agno = 2
        - agno = 0
        - agno = 3
Phase 5 - rebuild AG headers and trees...
        - reset superblock...
Phase 6 - check inode connectivity...
        - resetting contents of realtime bitmap and summary inodes
        - traversing filesystem ...
        - traversal finished ...
        - moving disconnected inodes to lost+found ...
Phase 7 - verify and correct link counts...
done
```

**Step 2** Check whether the repair is successful. Run the following command:

mount /dev/mapper/rhel-root /

If the mount is successful, the file system has been repaired.

**Step 3** Restart the OS.

Run **reboot**. After the reboot, the VM OS can be accessed.

**----End**

# 3.5 How Do I Obtain the Installation Package on a Production Site Server from the Gateway?

You can use remote transfer tools, such as SFTP or SCP, to obtain the installation package from the gateway.

Example:

Log in to a Windows VM, open PowerShell, and run the following commands:

sftp root@***.***.*.*
get sdrs_win_******.zip C:\Users\Administrator\Desktop\sdrs_win_******.zip

# 3.6 How Do I Enable or Disable an ECS Firewall and Add a Port Exception to the Firewall?

You can enable or disable an ECS firewall and add port exceptions to the firewall.
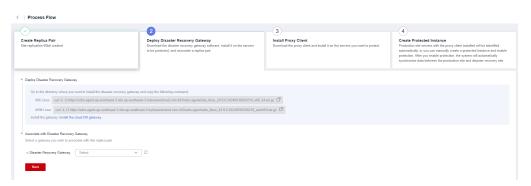
For the operation procedure on a Windows ECS, see **How Do I Disable a Windows ECS Firewall and Add a Port Exception on a Windows ECS Firewall?**

For the operation procedure on a Linux ECS, see **How Do I Disable a Linux ECS Firewall and Add a Port Exception on a Linux ECS Firewall?**

# 3.7 Why Can't I Find the Disaster Recovery Gateway When Associating a Replica Pair with It?

## Symptom

After a disaster recovery gateway is installed and configured, the gateway cannot be found on the console.
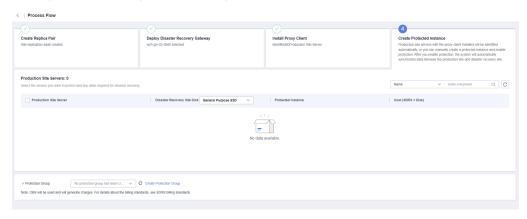


## Possible Causes

Possible causes are as follows:

- **Cause 1**: In a cross-AZ scenario, the gateway is not in the AZ where you create the replica pair.

  Solution 1: Create a replica pair in the AZ where the gateway resides.

  Solution 2: Create a server and deploy the gateway in the AZ where the replica pair is created.

- **Cause 2**: In a cross-region or cross-AZ scenario, the gateway is not in the VPC you selected when creating the replica pair.

  Solution: Create a new replica pair and select the VPC where the gateway belongs.

# 3.8 Why Is No Production Site Server Displayed When I Create Protected Instances?

## Symptom

After the proxy client is installed, corresponding production site servers are not displayed on the page for creating protected instances.



## Possible Causes

Possible causes are as follows:

- **Cause 1**: The proxy client installation fails. Check whether the OSs of the production site servers are included in **Supported OSs**.

  Solution: Reinstall the proxy client on servers running the supported OSs.

- **Cause 2**: The proxy client and disaster recovery gateway are deployed on the same server.

  Solution: Deploy the disaster recovery gateway on an independent server and reinstall the proxy client.