Storage Disaster Recovery Service

Product Introduction

Issue	01
Date	2024-11-05

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 SDRS Infographics

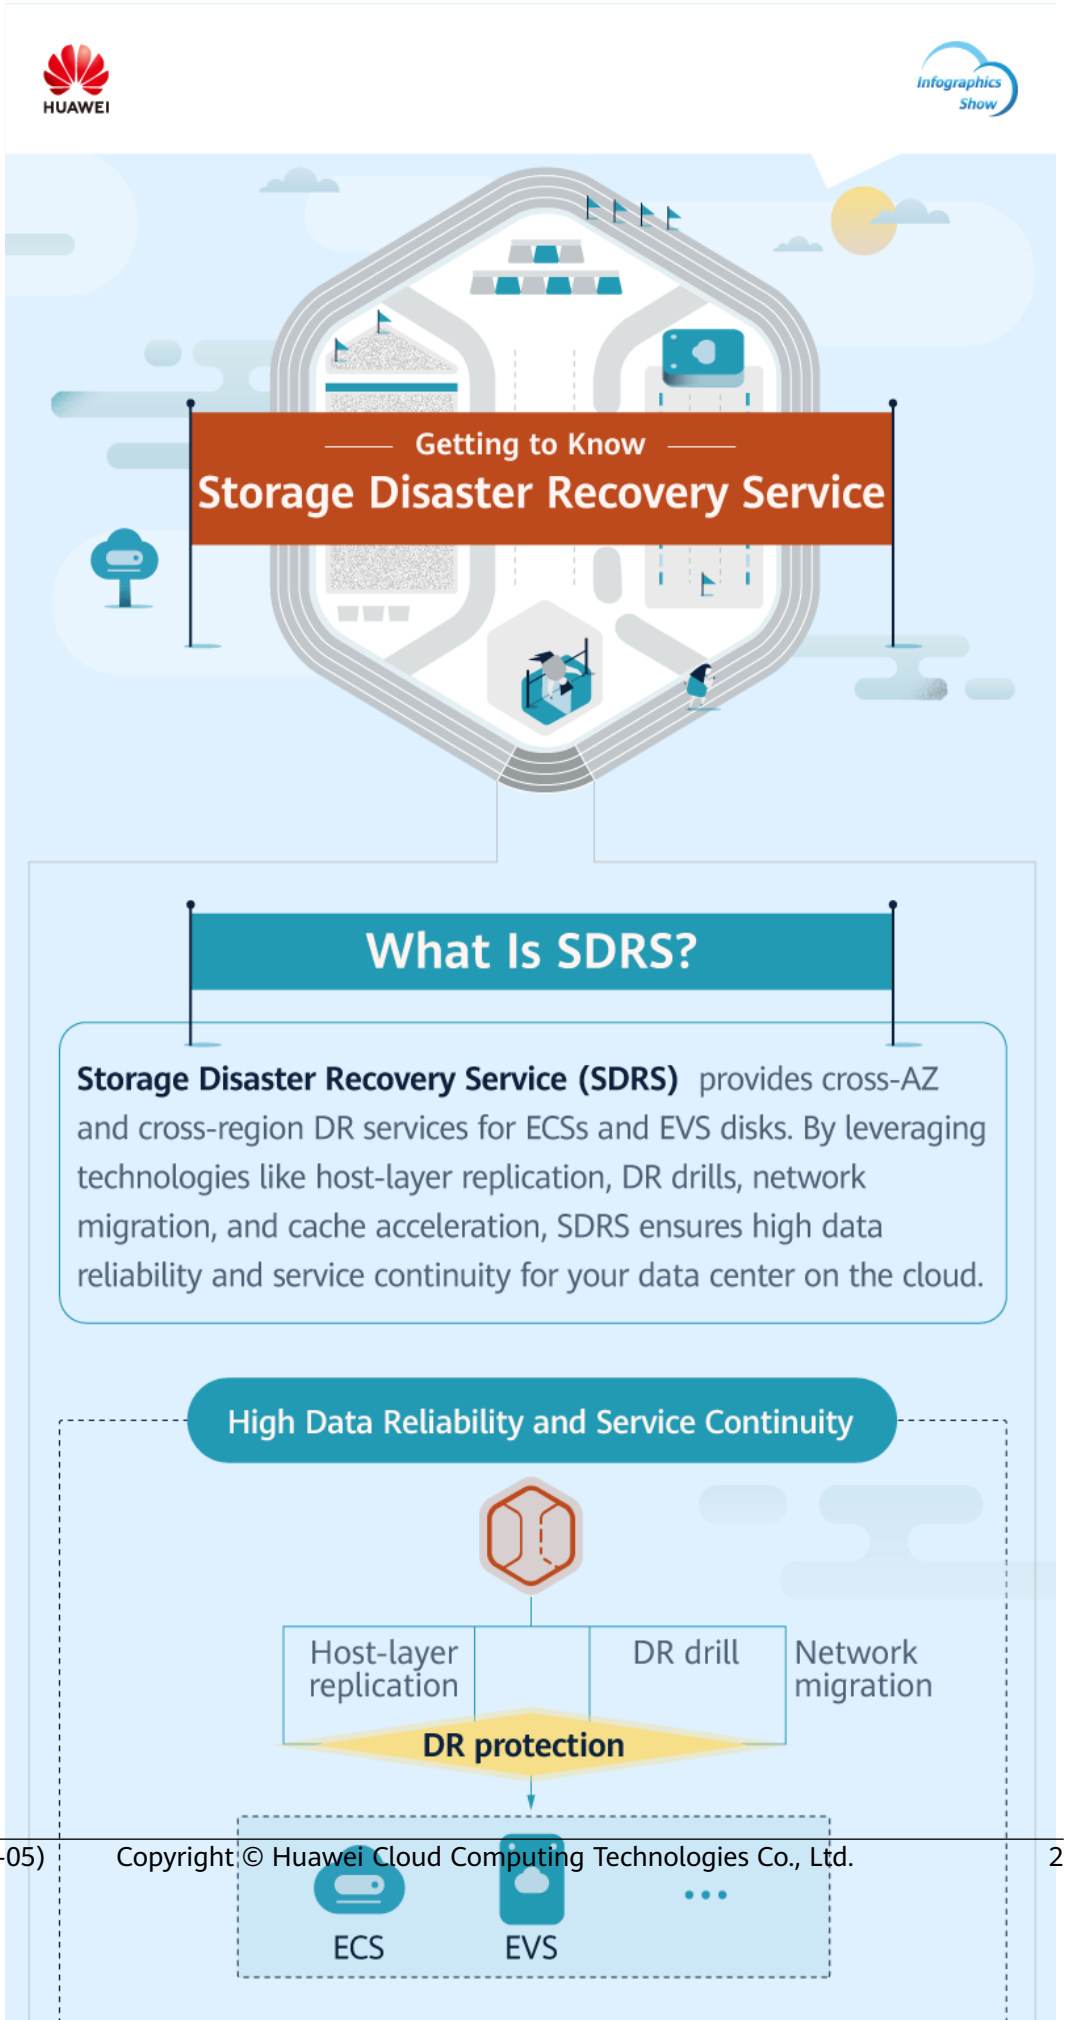

2 What Is SDRS?

Overview

Storage Disaster Recovery Service (SDRS) provides disaster recovery services for cloud services like Elastic Cloud Server (ECS) and Elastic Volume Service (EVS). By leveraging technologies, such as host-layer replication, data redundancy, and cache acceleration, SDRS can provide you with high data reliability and service continuity.

SDRS protects your applications by replicating the ECS data and configurations to a disaster recovery site. It allows applications to start and run at the disaster recovery site if any production site server stops.

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.

3 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 servers. 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 (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.

4 SDRS Application Scenarios

Cross-region/Cross-AZ Disaster Recovery

SDRS provides host-layer disaster recovery protection for Huawei Cloud servers with an RPO of just seconds. By leveraging host-layer asynchronous replication, it offers cross-region and 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 simple configurations.

Figure 4-1 Storage disaster recovery

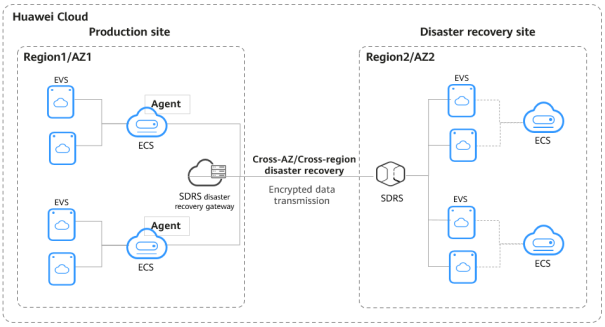

IDC Disaster Recovery

IDC disaster recovery (DR) is a DR solution that involves both public and private clouds. It allows you to set up disaster recovery for a VMware environment from a local data center, private cloud, or HCS Online environment to the public cloud. It also allows you to recover services on the cloud to ensure your data security and service continuity.

Figure 4-2 VMware disaster recovery

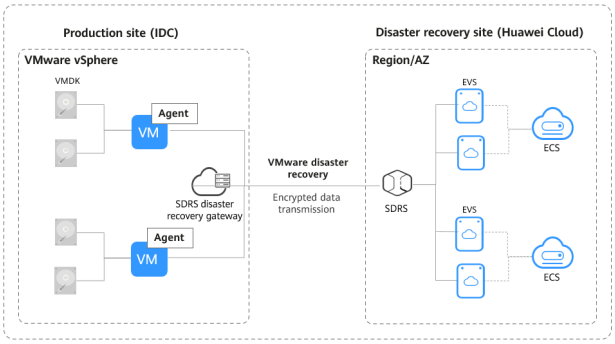

Figure 4-3 HCS Online disaster recovery

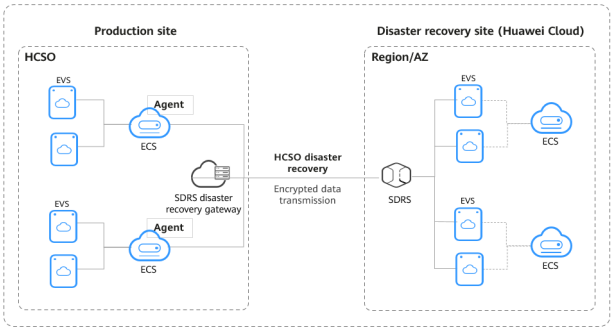

Disaster Recovery Drill

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.

5 SDRS Usage Restrictions

Before using SDRS, learn about the constraints listed in the following tables.

Table 5-1 Asynchronous replication constraints

Constraint	Description
Specifications	Constraints on specifications: • Minimum specifications recommended for the cloud disaster recovery gateway, production site servers, and disaster recovery site servers are 8 vCPUs and 16 GB memory. **NOTE** For disaster recovery site servers, you are advised to reserve 2 GB memory for reverse reprotection. • One disaster recovery gateway can protect a maximum of 20 protected instances and 58 disks.
Server	Constraints on servers: • Protected instances can be created for ECSs. • Only KVM ECSs are supported. • Both x86 and Arm ECSs are supported. • Protected instances cannot be created for ECSs of the following series: – General Computing-plus C7 – Memory-optimized M7 – Disk-intensive – Ultra-high I/O – GPU-accelerated – FPGA-accelerated • The OS must meet compatibility requirements. • If you delete a server or disks that have been used to create a protected instance, the protected instance will become invalid.

Constraint	Description
Disk	• Shared disks and encrypted disks are not supported. • Servers' local disks are not supported. • General Purpose SSD V2 disks cannot be used as disaster recovery disks. • Disks using LVM are not supported.
Network	• Round-trip time (RTT) ≤ 100 ms **NOTE** You can run the **ping** command and set the packet size to 64,000 bytes to test the latency between the production site and the OBS domain name at the disaster recovery site. **ping** *OBS domain name at the disaster recovery site* **-s 64000** • Bandwidth ≥ Changed data volume per minute during peak hours/60 seconds. Minimum bandwidth: 10 Mbit/s **NOTE** Obtain the total amount of changed data during the period with the maximum service load and calculate the peak write bandwidth of the service disk using the changed data volume and time period. To meet the RPO requirements, use the obtained peak bandwidth as the minimum bandwidth. In cross-region replication scenarios, if the network bandwidth is shared, obtain the peak bandwidth of each protected instance, and use the maximum peak bandwidth as the minimum bandwidth. Note: 1. Service load statistics are collected by server. 2. The baseline bandwidth of the cloud disaster recovery gateway must meet the bandwidth requirements. The recommended minimum baseline bandwidth is 2 Gbit/s. • Packet loss rate < 0.1%
Function	Constraints on functions: • All disks on an ECS are protected. After protected instances are created, new disks cannot be added to and existing disks cannot be removed from the protected instances, and existing disks cannot have their capacities expanded. • Data consistency between protected instances is not supported.

Table 5-2 Synchronous replication constraints

Constraint	Description
Compute	Constraints on server types: • Kunpeng ECSs are not supported. • x86-based, GPU-accelerated and FPGA-accelerated ECSs are not supported.

Constraint	Description
Replication	Constraints on servers: • The servers must be from two AZs of the same region. • BMSs are not supported. • The following types of servers cannot be used to create protected instances: – Large Memory (Xen): This type of servers is bound to InfiniBand NICs. – Disk Intensive: This type of servers has local disks. – Ultra-high I/O: This type of servers has local disks.
	Constraints on EVS disks: • Disks used to create replication pairs cannot be deleted, and the disk snapshot cannot be used to roll back data.
Storage	Only servers using EVS or DSS disks are supported.
Application	Storage-based synchronous replication ensures disk data consistency but does not guarantee application data consistency. If your applications support crash consistency, you can run and replicate applications.
Deployment model	VPC migration: Servers at the production site and those at the disaster recovery site are in the same VPC. NIC migration and multiple NICs are supported for each server.
Service interconnection	Tag Management Service (TMS) is supported only via API.
Backup and restore	Only servers at the production site can be backed up and restored. Servers at the disaster recovery site can only be backed up.

NOTE

If the production site AZ becomes faulty, you can run a disaster recovery drill to recovery the services on servers.

Helpful Links

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?

6 SDRS Supported OSs

The following tables list the OS versions that have been verified in the lab. Site servers running the following OSs supports SDRS.

Windows OSs are in the PoC state.

Asynchronous Replication

Table 6-1 OSs supported by asynchronous replication

Type	Version	Kernel Version	Bit
CentOS	7.2	3.10.0-327.el7.x86_64 3.10.0-1127.19.1.el7.x86_64 3.10.0-1160.6.1.el7.x86_64	64
	7.3	3.10.0-514.el7.x86 64	64
	7.4	3.10.0-693.el7.x86 64	64
	7.5	3.10.0-862.el7.x86 64	64
	7.6	3.10.0-957.el7.x86_64 3.10.0-1160.6.1.el7.x86_64	64
	7.7	3.10.0-1062.el7.x86_64	64
	7.9	3.10.0-1160.el7.x86_64	64
Red Hat	7.4	3.10.0-693.el7.x86_64	64
	7.9	3.10.0-1160.el7.x86_64	64
	8.6	4.18.0-372.9.1.el8.x86_64	64
	8.8	4.18.0-477.10.1.el8_8.x86_64	64
	8.9	4.18.0-513.5.1.el8_9.x86_64	64
Windows	Windows Server 2016	-	64

Type	Version	Kernel Version	Bit
	Windows Server 2019	-	64
	Windows Server 2022	-	64
UnionTech OS V20	1050e	4.19.90-2201.4.0.0135.up1.uel20.aarch64	64
Kylin V10	sp2	4.19.90-24.4.v2101.ky10.aarch64	64

Notes and Constraints on VMware Disaster Recovery

- VMware versions that support disaster recovery to the cloud include VMware vSphere 6.7 and VMware vSphere 7.0.
- Disaster recovery can only be set up to the cloud for VMware VMs deployed using ESXi, and failback is not supported.
- For better performance and operation experience, disaster recovery is only supported for the OSs listed in **Table 6-2**. These OSs have passed the compatibility test.
- Only VMs whose system disks are configured with LVM on a single disk can be stored on the cloud. Those with LVM configured on multiple disks cannot be restored on the cloud.
- The bandwidth used for disaster recovery should be at least 100 Mbit/s. The size of each disk on a VMware VM should be an integral multiple of 10, in GB.
- If a VM runs Linux, you must configure the **disk.EnableUUID TRUE** VM parameter.
- Only SCSI disks are supported for VMware VMs.

Table 6-2 Supported OSs of VMware production VMs

OS Type	Supported Version
Windows	Windows Server 2016 Windows Server 2019
CentOS	CentOS 7.2 CentOS 7.6

Synchronous Replication

Table 6-3 OSs supported by synchronous replication

Type	Version	Bit
Windows	Windows Server 2008 R2	64
	Windows Server 2012 R2	64
	Windows Server 2016	64
Red Hat	Red Hat Enterprise Linux 6.10	64
	Red Hat Enterprise Linux 7.5	64
CentOS	CentOS 6.5	64
	CentOS 6.8	64
	CentOS 6.9	64
	CentOS 7.2	64
	CentOS 7.3	64
	CentOS 7.4	64
SUSE	SUSE Enterprise 12 SP2	64
Ubuntu	Ubuntu 16.04 server	64

NOTE

- OS images are from the cloud platform public images.
- OSs supported by SDRS will be updated on an ongoing basis.

7 SDRS and Other Services

Table 7-1 Asynchronous replication and other services

Interaction	Related Service	Reference
Use ECSs to create protected instances in asynchronous replication.	ECS	• **Creating Protected Instances**
Use EVS disks to create replication pairs in asynchronous replication.	EVS	• **Creating Protected Instances**
Create a disaster recovery network on the cloud for asynchronous replication.	VPC	• **Creating a Replica Pair**
Cache the production site and disaster recovery site data.	OBS	-

Table 7-2 Synchronous replication and other services

Interaction	Related Service	Reference
Use ECSs to create protected instances in synchronous replication.	ECS	• **Creating Protected Instances**
Use EVS disks to create replication pairs in synchronous replication.	EVS	• **Creating a Replication Pair**
Use DSS disks to create replication pairs in synchronous replication.	DSS	• **Creating a Replication Pair**
Select a VPC for the protection group in synchronous replication.	VPC	• **Creating a Protection Group**

8 SDRS Billing

Billing Modes

SDRS supports pay-per-use billing (postpayment).

- With regard to pay-per-use, prepayment is not required. After you top up your Huawei Cloud account, the system calculates the resource usage and settles the bill every hour. Fees are then deducted from your account balance.

For detailed prices, see **Product Pricing Details**.

NOTE

Spot instances cannot be used as production site servers.

Pay per Use

- Billing item: protected instance usage duration
- Billing mode: billed and settled by hour. A minimum fee is not included.

Helpful Links

- **How Am I Billed for SDRS?**

9 Permissions Management

If you need to assign different permissions to personnel in your enterprise to access your SDRS resources, Identity and Access Management (IAM) is a good choice for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you to securely access your Huawei Cloud resources.

With IAM, you can use your HUAWEI ID to create IAM users, and assign permissions to the users to control their access to specific resources.

If your HUAWEI ID does not require individual IAM users for permissions management, skip this section.

IAM is a free service. You only pay for the resources in your account. For more information about IAM, see **IAM Service Overview**.

SDRS Permissions

By default, new IAM users do not have permissions. To assign permissions to new users, add them to one or more groups, and grant permissions to these groups. The users then inherit permissions from the groups to which the users belong, and can perform specific operations on cloud services.

SDRS is a project-level service deployed and accessed in specific physical regions. To assign SDRS permissions to a user group, specify the scope as region-specific projects and select projects (such as **ap-southeast-2**) for the permissions to take effect. If **All projects** is selected, the permissions will take effect for the user group in all region-specific projects. When accessing SDRS, the users need to switch to a region where they have been authorized to use SDRS.

Table 9-1 lists all the system-defined roles supported by SDRS. This role is dependent on other roles. When assigning SDRS roles to a user, you need to also assign the dependent roles to the user for the SDRS permissions to take effect.

Table 9-1 System-defined roles supported by SDRS

Role Name	Description	Dependencies
SDRS Administrator	Administrator permissions for SDRS	Dependent on the **Tenant Guest**, **Server Administrator**, and **VPC Administrator** policies. • **Tenant Guest**: A global policy, which must be assigned in the Global project. • **Server Administrator**: A project-level policy, which must be assigned in the same project as the **SDRS Administrator** policy. • **VPC Administrator**: A project-level policy, which must be assigned in the same project as the **SDRS Administrator** policy.

Helpful Links

- **IAM Service Overview**
- **Creating User Groups and Users and Granting SDRS Permissions**

10 Product Concepts

10.1 SDRS Concepts

Table 10-1 General concepts

Concept	Description
Production site	Data center that independently runs services in normal cases. In asynchronous replication, the production site refers to your on-premises data center or the location where your ECSs reside. In synchronous replication, the production site refers to the AZ where your servers reside. It is specified when you create protection groups.
Disaster recovery site	Data center that does not run services when the production site works properly. It is used to back up data in real time. When the production site fails (planned or unexpected), the disaster recovery site can take over services after a switchover or failover. It can reside in the same city as the service management center or in another city. The production site and disaster recovery site must be in two AZs of a same region.
Protection group	Manage the servers you want to replicate. One protection group manages servers in one VPC. If you have multiple VPCs, create multiple protection groups.
Protected instance	A protected instance consists of one server and its replicated server. One protected instance belongs to one protection group only. The AZs of instance servers are the same as those of the protection group's production site and disaster recovery site.

Concept	Description
VBD	Virtual Block Device (VBD) is the default device type of EVS disks. VBD EVS disks support only basic SCSI read/write commands. This disk type is suitable for enterprise office applications as well as development and testing.
SCSI	Small Computer System Interface (SCSI) is another EVS device type. SCSI EVS disks support transparent SCSI command transmission and allow the server OS to directly access the underlying storage media. In addition to basic SCSI read/write commands, SCSI EVS disks support advanced SCSI commands, such as persistent SCSI reservations, which are used for clustered applications to guarantee data security.
RPO	Recovery point objective. It is a service switchover policy with minimal data loss. Data recovery points are used as objectives to ensure that the data used for disaster recovery switchovers is the latest backup data.
RTO	Recovery time objective. It is the target time spent for critical services to recover to an acceptable level. RTO is set to minimize the impacts on the services. In SDRS, RTO refers to the period of time from when you perform a switchover or failover at the production site to the time when the servers at the disaster recovery site start to run. This period does not include the time spent on DNS configuration, security group configuration, or customer script execution, and is within 30 minutes.
Disaster recovery drill	Verify that disaster recovery site servers can take over services from production site servers after a failover. By running disaster recovery drills, you can simulate recovery scenarios and formulate recovery plans. When a fault occurs, you can use the plans to recover services as quickly as possible.

Table 10-2 Asynchronous replication concepts

Concept	Description
Replica pair	A replica pair consists of a production site and a disaster recovery site. The replication relationship is established between two sites.
Cloud disaster recovery gateway	Aggregate and compress data on all replicated production site servers, and synchronize the data to the disaster recovery site.
Proxy client	Transmit data on the server to the cloud disaster recovery gateway.
Enabling protection	If services are running at the production site and the data synchronization stops, you can enable protection to start data synchronization.

Concept	Description
Failover	A failover switches the services from the production site to the disaster recovery site. After a failover, data synchronization stops and the protected instance status changes to **Failover completed**.
Failback	After a failover, services are running at the disaster recovery site. You can fail back to your production site with a failback. After the failback, data synchronization stops.
Reverse reprotection	After a failover, data is not automatically synchronized from the disaster recovery site to the production site, and protection is disabled for protected instances. To start data synchronization from the disaster recovery site to the production site, perform a reverse reprotection.
Reprotection	After a failback, data is not automatically synchronized from the production site to the disaster recovery site, and protection is disabled for protected instances. To start data synchronization from the production site to the disaster recovery site, perform a reprotection.
Disabling protection	If services are running at the production site, and data synchronization is in progress or completed, you can disable protection to stop data synchronization.

Table 10-3 Synchronous replication concepts

Concept	Description
Replication pair	A replication pair consists of one EVS disk and its replicated disk. One replication pair belongs to one protection group and can be attached to a protected instance in this group.
Switchover	Temporarily stop servers at the production site and switch over services to the disaster recovery site for planned outages. After a switchover, the disaster recovery direction is from the disaster recovery site to the production site. Servers and EVS disks at the disaster recovery site are ready to start.
Failover	A failover forcibly stops the servers and disks at the production site and sets the servers and disks at the disaster recovery site to ready-to-start state. This operation affects all the protected instances in the protection group. After a failover, you need to manually start the servers at the disaster recovery site. In addition, the protection group status changes to **Failover complete**, and data synchronization of the protection group stops. You need to enable reprotection to recover data synchronization.

Concept	Description
Enabling protection	Protection can be enabled after a protection group is created or data synchronization stops. Once protection is enabled, data synchronization starts, and you can view the synchronization progress on the console. This operation affects all the protected instances and replication pairs in the protection group. After you click **Enable Protection**, the status of the protection group changes to **Synchronizing**, and **Disable Protection** is not available.
Reprotection	Reprotection can be enabled after a failover. Once reprotection is enabled, data synchronization starts, and you can view the synchronization progress on the console. This operation affects all the protected instances and replication pairs in the protection group. After you click **Reprotect**, the status of the protection group changes to **Reprotecting**, and **Disable Protection** becomes unavailable.
Disabling protection	Protection can be disabled after data synchronization of a protection group is complete. After disabling protection, the status of the protection group changes to **Available**.
Attaching a replication pair	Attach the two disks in a replication pair to the servers in a protected instance.
Detaching replication pair	Detach the two disks in a replication pair from the servers in a protected instance.
Disaster recovery direction	Data replication direction. After you create a protection group, data is replicated from the production site to the disaster recovery site. A switchover or failover changes the disaster recovery direction of a protection group.
Protection group status	Status of a protection group, after you create, delete, switch over, fail over, enable protection for, or disable protection for a protection group. For details, see **Protection Group Status**.
Synchronization status	Data replication status between the production and disaster recovery sites.
VPC	VPC of the protection group. A VPC facilitates internal network management and configuration, allowing secure and quick modifications to networks. By defaults, servers in the same VPC can communicate with each other, but those in different VPCs cannot.

10.2 Region and AZ

Concept

A region and availability zone (AZ) identify the location of a data center. You can create resources in a specific region and AZ.

- Regions are divided based on geographical location and network latency. Public services, such as Elastic Cloud Server (ECS), Elastic Volume Service (EVS), Object Storage Service (OBS), Virtual Private Cloud (VPC), Elastic IP (EIP), and Image Management Service (IMS), are shared within the same region. Regions are classified into universal regions and dedicated regions. A universal region provides universal cloud services for common tenants. A dedicated region provides specific services for specific tenants.
- An AZ contains one or more physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proof, and electricity facilities. Within an AZ, computing, network, storage, and other resources are logically divided into multiple clusters.

Figure 10-1 shows the relationship between regions and AZs.

Figure 10-1 Regions and AZs

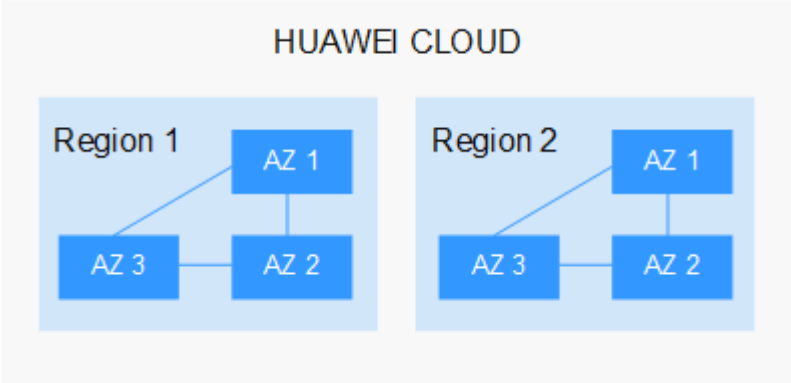

Huawei Cloud provides services in many regions around the world. You can select a region and an AZ based on requirements. For more information, see **Huawei Cloud Global Regions**.

Selecting a Region

When selecting a region, consider the following factors:

- Location

  It is recommended that you select the closest region for lower network latency and quick access.
- Resource price

  Resource prices may vary in different regions. For details, see **Product Pricing Details**.

Regions and Endpoints

Before you use an API to call resources, specify its region and endpoint. For more details, see **Regions and Endpoints**.

