Storage Disaster Recovery Service

# Product Introduction

**Issue** 10

**Date** 2021-09-25

HUAWEI TECHNOLOGIES CO., LTD.

## Trademarks and Permissions


and other Huawei trademarks are trademarks 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 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.

# Security Declaration

## Vulnerability

Huawei's regulations on product vulnerability management are subject to “Vul. Response Process”. For details about the policy, see the following website:https://www.huawei.com/en/psirt/vul-response-process
For enterprise customers who need to obtain vulnerability information, visit:https://securitybulletin.huawei.com/enterprise/en/security-advisory

# Contents



  

# 1 What Is SDRS?

## Overview

Storage Disaster Recovery Service (SDRS) provides cloud disaster recovery (DR) for your data centers. If your on-premises or cloud data center fails, you can fail over services to the DR center on Huawei Cloud, and then fail back the services after the production center recovers. This helps you improve service continuity and ensure the security and reliability of critical data.

## 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.

  

# 2 SDRS Advantages

SDRS has the following advantages:

- Convenient recovery solution

  Using the SDRS console, you can configure and manage server replication and perform switchovers or failovers.

- 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.

- Crash consistency

  Storage-based, real-time data synchronization keeps crash consistency for your data across two AZs. Specifically, application data might not be consistent during a failover, but disk data is always consistent.

- Disaster recovery drill

  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.

- Synchronous replication

  Replicate the servers to the disaster recovery site in real time to ensure zero recovery point objectives (RPO).

  Efficient network switchover: Synchronous replication streamlines resource management during failovers, for example, to reserve IP addresses and MAC addresses.

  Cost-effective: When production site services are running properly, servers at the disaster recovery site are stopped. This greatly reduces the disaster recovery TCO.

  Easy deployment: Disaster recovery agent is not required during server deployment. This makes the deployment simple and quick.

  

- Asynchronous replication (under restricted OBT)

  Continuously replicate the servers in the user's data center and ensures RPO within seconds.

  Efficient network switchover: Asynchronous replication streamlines resource management during failovers, for example, to reserve IP addresses and MAC addresses.

  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 OBS buckets used.

  

# 3 SDRS Application Scenarios

## Asynchronous Replication (Under Restricted OBT)

Asynchronous replication continuously replicates servers at local IDCs to Huawei Cloud with RPO in a few seconds. By leveraging asynchronous replication techniques at the host layer, it allows for cross-AZ disaster recovery and keeps crash consistency for your data. If production site services fail to recover within a short period of time due to force majeure (fire and earthquake) or device faults (software and hardware damage), you can quickly recover services at the disaster recovery site with some simple configurations.

**Figure 3-1** Asynchronous replication

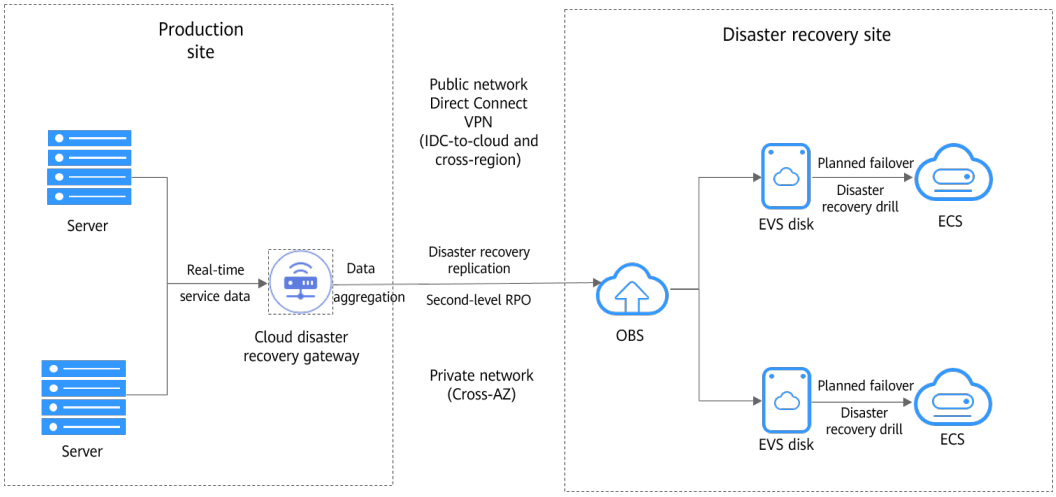


## Synchronous Replication

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

  

SDRS is suitable for stateful applications, such as Microsoft Office 365, that requires the storage of user data on the particular server running this application.

**Figure 3-2** Synchronous replication

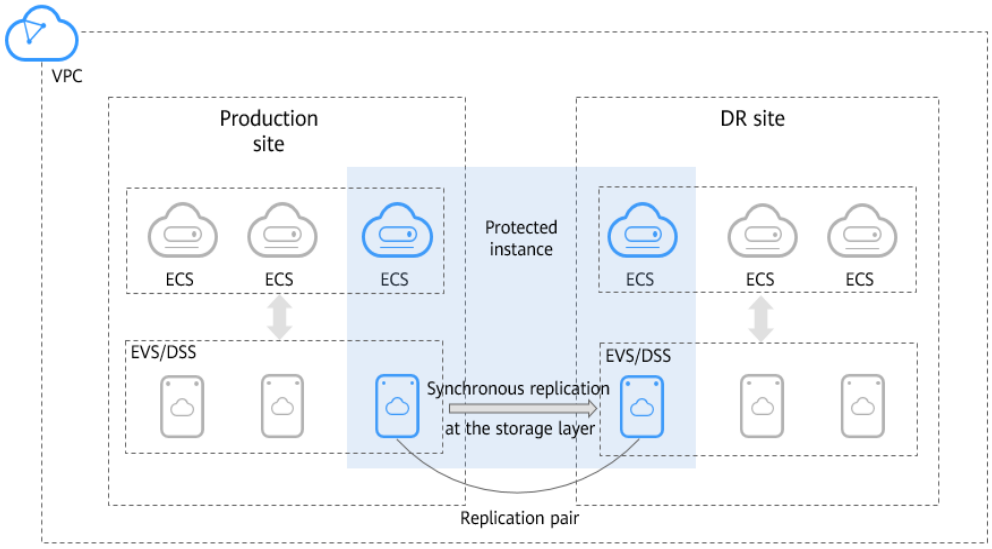


## Disaster Recovery Drill

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.

## Server Migration

With server replication capabilities, SDRS allows you to migrate on-premises servers to the cloud, or cloud servers from one AZ or region to another.

  

# 4 SDRS Usage Restrictions

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

**Table 4-1** Asynchronous replication constraints

| Constraint | Description |
| --- | --- |
| Architecture | Constraints on architecture types and versions:<br>• For Huawei Cloud Stack, the version must be 6.5 or later.<br>• For Huawei Cloud Stack Online, all versions are supported. |
| Specifications | Constraints on specifications:<br>• Memory on each cloud disaster recovery gateway, production site server, and disaster recovery site server must be greater than 2 GB.<br>• One disaster recovery gateway can protect a maximum of 20 protected instances. |
| Server | Constraints on servers:<br>Only KVM servers can be used for disaster recovery. |
| Function | Constraints on functions:<br>• Shared EVS disks are not supported.<br>• Failback is not supported if the production site servers are billed on a yearly/monthly basis.<br>• Bare Metal Servers (BMSs) are not supported.<br>• EVS disk creation, deletion, and capacity expansion are not allowed on protected instances.<br>• The system disk of a disaster recovery gateway server cannot use LVM. |

  

**Table 4-2** Synchronous replication constraints

| Constraint | Description |
| --- | --- |
| Compute | Constraints on server types:<br>● Kunpeng ECSs are not supported.<br>● x86-based, GPU-accelerated and FPGA-accelerated ECSs are not supported. |
| Replication | Constraints on servers:<br>● The servers must be from two AZs of the same region.<br>● BMSs are not supported.<br>● The following types of servers cannot be used to create protected instances:<br>– Large Memory (Xen): This type of servers is bound to InfiniBand NICs.<br>– Disk Intensive I (Xen): This type of servers has local disks.<br>– Disk Intensive II (KVM): This type of servers has local disks. |
| | Constraints on EVS disks:<br>● 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?**

  

# 5 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.

**Table 5-1** OSs supported by asynchronous replication

| Type | Version | Kernel Version | Bit |
| --- | --- | --- | --- |
| CentOS | 7.2.1511 | 3.10.0-327.el7.x86_64<br>3.10.0-1160.6.1.el7.x86_64<br>3.10.0-1127.19.1.el7.x86_64 | 64 |
| CentOS | 7.6.1810 | 3.10.0-1160.6.1.el7.x86_64<br>3.10.0-957.el7.x86_64 | 64 |
| Red Hat | 8.6 | 4.18.0-372.9.1.el8.x86_64 | 64 |
| Windows | Windows Server 2016 | - | 64 |
| | Windows Server 2019 | - | 64 |

**Table 5-2** 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 |

  

| Type | Version | Bit |
| --- | --- | --- |
| 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.

  

# 6 SDRS and Other Services

**Table 6-1** SDRS and other services

| Interaction | Related Service | Reference |
| --- | --- | --- |
| Use ECSs to create protected instances in synchronous replication. | ECS | • Creating a Protected Instance |
| 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 |
| Create a cloud disaster recovery network in asynchronous replication or select a VPC for a protection group in synchronous replication. | VPC | • Creating a Protection Group |
| Use CTS to record SDRS operations for querying, auditing, or backtracking later. | Cloud Trace Service (CTS) | • Interconnecting with CTS |

  

# 7 SDRS Billing

## Billing Modes

SDRS provides two billing modes. You can prepay for yearly/monthly packages or pay per use.

- With regard to prepayment, you need to buy a resource package before use. Resources used are then offset by the package quota. Any resource usage exceeds the quota is billed on a pay-per-use basis.
- 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.

## Yearly/Monthly

- Billing item: resource package

  Only resource packages containing protected instances are available.

  NOTE

  A resource package cannot be used across regions.

- Billing mode: yearly/monthly package, which offers more preferences than pay-per-use billing
  - Resource packages are charged as a one-time payment. You can choose whether a resource package takes effect immediately upon purchase or at a specified time.
  - Within the validity period of a resource package, the package quota resets at the beginning of each subscription month. If your usage exceeds the package quota, you will be billed on a pay-per-use basis for the subsequent usage.

  

NOTE

For example, you purchase a resource package containing 5 protected instances at 14:00 March 3, 2020, and the resource package takes effect immediately with a validity of one year. Then, there will be 5 protected instances available in each subscription month (from the third day of a month to the third day of the next month, in this example, from 14:00 March 3 to 14:00 April 3). If the quota in a subscription month is not used up, the remaining quota will be cleared at the beginning of the next subscription month.

- Renewal: If your purchased resource package is about to expire, you can renew the subscription to extend the validity period. The renewal fee varies with the renewal duration.
- Package unsubscription is not supported. After a resource package expires, you can still use SDRS, but you will be billed on a pay-per-use basis for the resources used. Ensure that your account has a sufficient balance.

## Pay per Use

- Billing item: protected instance usage duration
- Billing mode: billed and settled by hour. A minimum fee is not included.
- Billing mode change: By default, SDRS uses pay-per-use billing (protected instance usage duration is rounded up to the closest hours). You can change your billing mode if needed. Alternatively, you can purchase resource packages for a yearly or monthly period based on your service needs, to get cheaper prices.

## Helpful Links

- **How Am I Billed for SDRS?**
- **How Do I Use a Resource Package?**

  

# 8 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 8-1** 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 8-1** 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.<br>• **Tenant Guest**: A global policy, which must be assigned in the Global project.<br>• **Server Administrator**: A project-level policy, which must be assigned in the same project as the **SDRS Administrator** policy.<br>• **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

# 9 Product Concepts

## 9.1 SDRS Concepts

**Table 9-1** 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. In synchronous replication, the production site refers to the AZ where your servers reside. It is specified when you create a protection group. |
| 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.<br>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. |
| 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. |

  

| Concept | Description |
| --- | --- |
| 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. For 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.<br>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 9-2** 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 continuously synchronize the data to the disaster recovery site. |
| Proxy client | Continuously transmit the data on the proxy client-installed 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. |
| 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**. |

  

| Concept | Description |
| --- | --- |
| 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 9-3** 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.<br>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.<br>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.<br>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.<br>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. |

  

# 9.2 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. AZs within a region are interconnected using high-speed optical fibers, to support cross-AZ high-availability systems.

**Figure 9-1** shows the relationship between regions and AZs.

**Figure 9-1** Regions and AZs

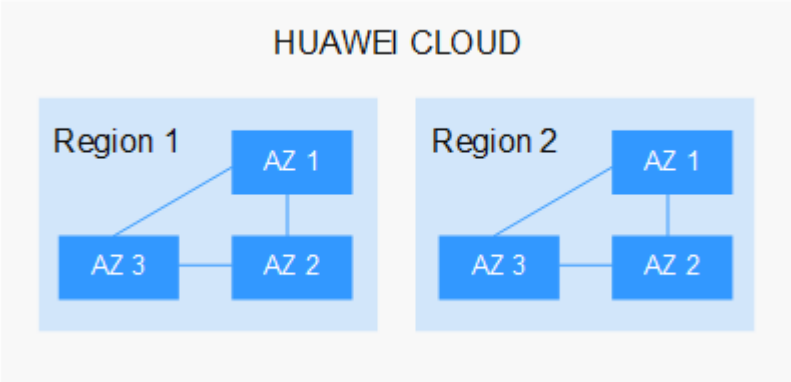


Huawei Cloud provides services in many regions around the world. You can select a region and an AZ based on requirements. For more information, see **Huawei Cloud Global Regions**.

## 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.

  - If your target users are in Asia Pacific (excluding the Chinese mainland), select the **CN-Hong Kong**, **AP-Bangkok**, or **AP-Singapore** region.
  - If your target users are in Africa, select the **AF-Johannesburg** region.

  - If your target users are in Latin America, select the **LA-Santiago** region.

    NOTE

    The **LA-Santiago** region is located in Chile.

- Resource price

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

## Selecting an AZ

When deploying resources, consider your applications' requirements on disaster recovery (DR) and network latency.

- For high DR capability, deploy resources in different AZs within the same region.
- For lower network latency, deploy resources in the same AZ.

## Regions and Endpoints

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

  

# 10 Change History

| Released On | Description |
| --- | --- |
| 2021-09-25 | This issue is the tenth official release.<br>Added the following content:<br>Added the description of asynchronous replication. |
| 2021-06-21 | This issue is the ninth official release.<br>Modified the following content:<br>Deleted the RTO target in **SDRS Advantages**.<br>Deleted the following content:<br>Deleted section "SDRS Compatible Applications and Versions." |
| 2021-04-20 | This issue is the eighth official release.<br>Modified the following content:<br>Deleted the description that general purpose SSD disks cannot be used to create replication pairs and that high I/O (performance-optimized I) and ultra-high I/O (latency-optimized) disks on SAP HANA ECSs, HPC ECSs, and HL1 ECSs cannot be used to create replication pairs in **SDRS Usage Restrictions**. |
| 2020-06-20 | This issue is the seventh official release.<br>Modified the following content:<br>Modified constraints in **SDRS Usage Restrictions**. Specifically, Kunpeng servers are not supported. |
| 2020-04-29 | This issue is the sixth official release.<br>Modified the following content:<br>Modified constraints in **SDRS Usage Restrictions**. Specifically, shared disks are supported. |

  

| Released On | Description |
| --- | --- |
| 2020-03-31 | This issue is the fifth official release.<br>Added the following content:<br>**SDRS Billing** |
| 2019-07-25 | This issue is the fourth official release.<br>Modified the following content:<br>Added constraints in **SDRS Usage Restrictions**. Specifically, added the constraint that servers at the production site can be backed up and restored. Servers at the disaster recovery site can be backed up only. |
| 2019-07-25 | This issue is the third official release.<br>Modified the following content:<br>Added constraints in **SDRS Usage Restrictions**. Specifically, added the constraint that servers at the production site can be backed up and restored. Servers at the disaster recovery site can be backed up only. |
| 2019-05-30 | This issue is the second official release.<br>Added the following content:<br>• **Region and AZ**<br>Modified the following content:<br>• Deleted constraints in **SDRS Usage Restrictions**. Specifically, deleted the constraint that the specifications of the disaster recovery site server must be consistent with those of the production site server. |
| 2019-05-24 | This issue is the first official release. |

  