Multi-Site High Availability Service

Service Overview

Issue 01

Date 2024-11-29

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 What Is MAS?

Enterprises are posed many challenges during application development, for example:

- Unreliable single available zone (AZ) or region
- Small system capacity and high access latency
- Technology incompetence and high price due to the cloud vendor lock-in
- Technical bottlenecks brought by rapid service growth

Multi-Site High Availability Service (MAS) is part of Huawei's consumer solution for high availability of multi-active applications. It provides E2E service failover and disaster recovery (DR) drill capabilities from the traffic input and data to the application layer, for faster service recovery and better continuity.

MAS = Multi-active access + Application layer software development kits (SDKs) + Data synchronization channels + Unified management and control center. Implementable solution = MAS + Consulting + Ecosystem partners + DR specifications.

MAS has the following advantages:

1. Service-level high availability
2. E2E availability of traffic, services, and data
3. Recovery Time Objectives (RTO) and Recovery Point Objective (RPO) in seconds, ensuring service continuity
4. DR drills with higher efficiency and lower cost

MAS provides the following core capabilities:

- E2E management-traffic-application-data arbitration and multi-active DR
- Secure and reliable data synchronization channels
- Multi-active DR standards and specifications
- Consulting and implementation

Figure 1-1 shows the MAS service architecture.

Figure 1-1 MAS service architecture

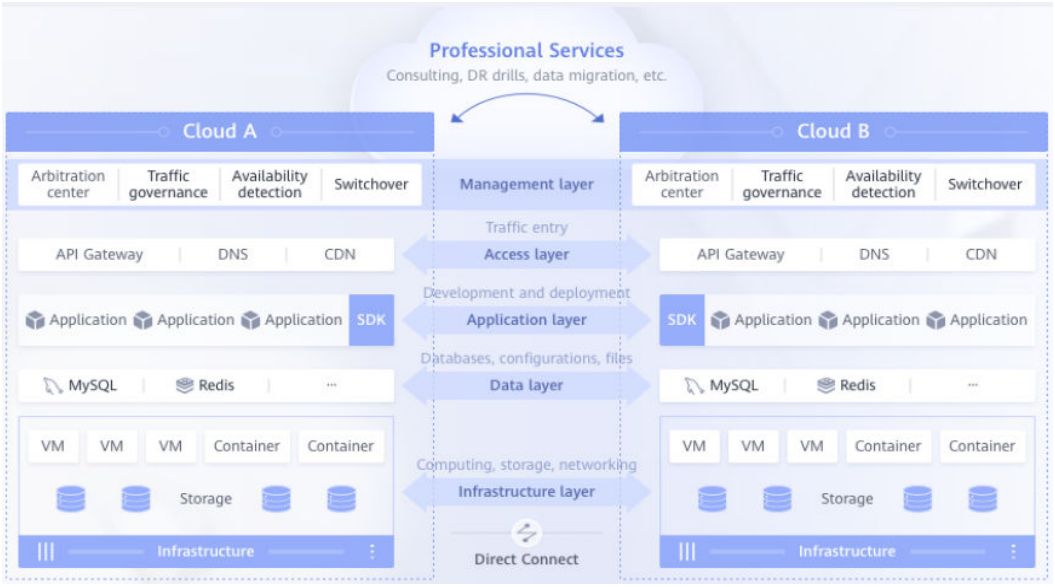

NOTE

- 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 as universal regions and dedicated regions. A universal region provides universal cloud services for common tenants. A dedicated region provides services of the same type only or 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.
- SDK is a type of software package for assisting in software development.
- RPO indicates the maximum data loss amount tolerated by the system.
- RTO indicates the maximum interruption duration tolerated by the system when a disaster occurs.

2 Application Scenarios

Service architecture is challenged by the system capacity/reliability and other business factors during development. MAS is applicable to the following scenarios:

- **Scenario 1: Intra-city cross-AZ**

  DR for intra-city cross-AZ applications and data, reaching 99.9% availability.
- **Scenario 2: Intra-city active-active**
  - DR for traffic-sharing applications and data that are intra-city, cross-AZ, and active-active, reaching 99.95% availability.
- **Scenario 3: Three AZs in two regions**
  - Synchronous DR across AZs and regions, providing the highest service continuity and 99.99% availability.
- **Scenario 4: Remote multi-active (cell-based)**
  - Suitable for latency-sensitive services that require 99.99%+ DR and partitioned data
  - Data partitioning for high availability, scalability, and reliability

Figure 2-1 MAS application scenarios

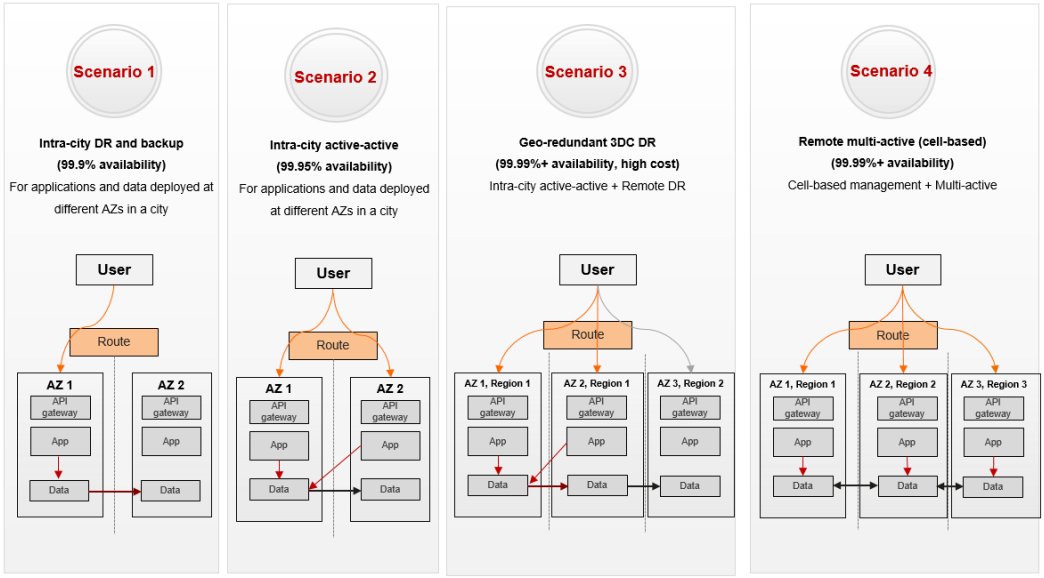

3 Edition Differences

Function Modules

You can choose the **Intra-city multi-active** or **Remote multi-active** edition.

The **Intra-city multi-active** edition provides the following **Features**:

- MySQL
- Oracle
- PostgreSQL
- Redis
- MongoDB
- Elasticsearch
- openGauss
- API

The **Remote multi-active** edition provides the following **Features**:

- MySQL
- Oracle
- PostgreSQL
- Redis
- MongoDB
- Elasticsearch
- openGauss
- API

Table 3-1 lists the details.

Table 3-1 Function modules

Edition	Description
Intra-city multi-active	Deploy your subsystems in the same region and less than 50 km from each other. • Service continuity: real-time traffic balancing and fast failover • Horizontal scaling: rapid service development • Traffic isolation: flexible traffic scheduling • Cost effectiveness: higher resource utilization with less redundancy costs
Remote multi-active	Deploy at least one subsystem in a different region and more than 300 km from any other subsystem. • Service continuity: cross-region DR and multi-active with doubled system capacity and improved SLA • Unified management: same architecture, operations, and multi-active management for the two DCs • Switchable workflows: automatic failover and DR with default/custom workflows • Enhanced architecture: zone-based management for smaller fault radius

Multi-Active Instances

You can create Platinum MAS instances. They can be deployed in intra-city multi-active or remote multi-active mode on Huawei Cloud.

4 Quotas

Table 4-1 lists the resource quotas available to you in MAS.

Table 4-1 MAS quotas

Resource	Quota (Count)	Remarks
Namespace	10	Number of namespaces supported by the system
MAS instance	10	Number of instances supported by the system
Application	200	Number of applications per instance
Monitor	100	Number of monitors per instance
Data source	200	Number of data sources per namespace
Data synchronization link	100	Number of links per namespace

5 Billing

Billing Items

MAS bills you for the enabled function modules, multi-active instances, and associated nodes. For details, see **Price Calculator**.

Table 5-1 MAS billing items

Billing Item	Description
Function module	Edition, features, and duration
MAS instance	Specifications, duration, and instance quantity
Node	Node quantity and usage duration **NOTE** This calculates the number of application processes monitored by a multi-active instance.

Billing Mode

There are two billing modes.

- **Yearly/Monthly**: billed by the purchase period. This is available only for MAS instances.
- **Pay-per-use**:
  - Function module: billed by the MAS instance edition and features on a daily basis
  - MAS instance: billed by the specifications on an hourly basis
  - Node: billed on a daily basis
- You can change your billing mode as required.

To change the billing mode, see **From Pay-per-Use to Yearly/Monthly** or **From Yearly/Monthly to Pay-per-Use**.

Renewal

Currently, MAS supports yearly/monthly and pay-per-use billing.

- Pay-per-use: You can use MAS as long as you need with sufficient balance.
- Yearly/monthly: An upfront payment is required when you purchase resources. No additional fees are incurred as you use the resources.

Renew the subscription on the management console. For details, see **Renewal Rules**.

Expiration and Overdue Payment

- Expiration

  If your subscription is not renewed before it expires, a retention period will apply. The retention period length depends on your account. After this period ends, your resources will be automatically deleted and cannot be restored, and the service cannot be renewed. For details about the retention period rule, see **Resource Suspension and Release**.

- Overdue payment

  If your account is in arrears, you can view the arrears details in the Billing Center. To prevent related resources from being stopped or released, top up your account within the specified period. For details, see **Topping Up an Account (Prepaid Direct Customers)**.

6 Permissions Management

If you need to assign different permissions to employees in your enterprise to access your MAS 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 secure access to your resources.

You can use your Huawei Cloud account to create IAM users, and assign permissions to the users to control their access to specific resources. For example, some software developers in your enterprise need to use MAS resources but should not be allowed to delete MAS instances or perform any other high-risk operations. In this scenario, you can create IAM users for the software developers and grant them only the permissions required for using MAS resources.

You can skip this section if you do not need fine-grained permissions management.

IAM is free of charge. You pay only for the resources in your account. For details, see **IAM Service Overview**.

MAS 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.

MAS is a project-level service deployed and accessed in specific physical regions. To assign MAS permissions to a user group, you need to specify region-specific projects for which the permissions will take effect. If you select **All projects**, the permissions will be granted for both the global service project and all region-specific projects. When accessing MAS, users need to switch to a region where they have been authorized to use MAS.

You can grant permissions by using roles and policies.

- Roles: a type of coarse-grained authorization mechanism that defines service-level permissions based on user responsibilities. This mechanism provides only a limited number of service-level roles for authorization. When using roles to grant permissions, you need to also assign other dependent roles for permissions to take effect. However, roles are not an ideal choice for fine-grained authorization and secure access control.

- Policies: a type of fine-grained authorization mechanism that defines permissions required to perform operations on specific cloud resources under certain conditions. This mechanism allows for more flexible policy-based authorization and secure access control. For example, you can grant MAS users only the permissions for performing specific operations on MAS instances.

  Table 6-1 lists all the system-defined roles and policies supported by MAS.

Table 6-1 System-defined roles and policies supported by MAS

Role/Policy Name	Description	Type	Dependency
MAS FullAccess	All permissions for MAS. Users granted these permissions can manage all MAS instances.	System-defined policy	None
MAS CommonOperations	Basic operation permissions for MAS, including the permissions to operate application, monitors, and components, but excluding the permissions to create or delete instances.	System-defined policy	None
MAS ReadOnlyAccess	Read-only permissions for MAS. Users granted these permissions can only view the MAS instance data.	System-defined policy	None

Table 6-2 lists the common operations supported by each system policy of MAS. Select the policies as required.

Table 6-2 Common operations supported by each system policy

Operation	MAS FullAccess	MAS CommonOperations	MAS ReadOnlyAccess
Enabling function modules	√	×	×
Modifying function modules	√	×	×
Deleting function modules	√	×	×

Operation	MAS FullAccess	MAS CommonOperations	MAS ReadOnlyAccess
Creating namespaces	√	×	×
Modifying namespaces	√	√	×
Deleting namespaces	√	×	×
Viewing namespace details	√	√	√
Creating applications	√	√	×
Modifying applications	√	√	×
Deleting applications	√	√	×
Buying multi-active instances	√	×	×
Modifying multi-active instances	√	√	×
Deleting multi-active instances	√	×	×
Viewing multi-active instances	√	√	√
Creating data sources	√	×	×
Modifying data sources	√	√	×
Deleting data sources	√	×	×
Viewing data source details	√	√	√
Creating sync links	√	×	×

Operation	MAS FullAccess	MAS CommonOperations	MAS ReadOnlyAccess
Editing synchronization links	√	√	×
Deleting sync links	√	×	×
Viewing sync link details	√	√	√
Switching over monitors	√	√	×
Creating monitors	√	√	×
Configuring monitors	√	√	×
Modifying monitors	√	√	×
Switching over data centers	√	√	×
Deleting monitors	√	√	×
Configuring connection pools	√	√	×
Configuring SDK accesses	√	√	×
Adding secret keys	√	√	×
Modifying secret keys	√	√	×
Deleting secret keys	√	√	×
Creating notifications	√	√	×
Modifying notifications	√	√	×
Deleting notifications	√	√	×

Operation	MAS FullAccess	MAS CommonOperations	MAS ReadOnlyAccess
Configuring DC-level automatic switchover	√	√	×
Creating credentials	√	×	×
Deleting credentials	√	×	×
Using credentials	√	√	×
Checking credentials	√	√	√

7 Basic Concepts

Namespace

A namespace is a collection of all resources (traffic input, multi-active areas, data synchronization records, and monitoring information) for a MAS instance. You can create a namespace for each system (such as OA and payment systems) to isolate resources for security and efficiency.

Instance

An instance is the minimum resource unit provided by MAS.

A MAS instance is an independent resource space where all operations are performed. Resources of different instances are isolated from each other. You can use one or more MAS instances as required.

AZ

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.

For maximum availability, deploy the MAS instance in multiple AZs. For minimum latency, deploy the MAS instance in the same AZ. The selected AZs must match your deployment architecture.

Enterprise Project

An enterprise project makes it easy to manage projects and group cloud resources and users.

Application

An application is a service system with functions and consists of one or more application components.

Monitor

A monitor detects and displays the real-time status of your database, and automatically triggers traffic switching when the database is abnormal.

8 Related Services

Virtual Private Cloud

A VPC is an isolated virtual network environment on the cloud. You can configure IP address ranges, subnets, and security groups in a VPC.

MAS runs in a VPC. The VPC service manages IP addresses and bandwidth and provides security groups. You can configure access rules for security groups to secure the access to MAS.

Relational Database Service

RDS is a cloud-based web service that is reliable, scalable, easy to manage, and immediately ready for use.

You can monitor RDS MySQL instances using MAS and switch data centers in the event of a failure.

Distributed Cache Service

DCS is an online, distributed, in-memory cache service compatible with Redis and Memcached. It is reliable, scalable, usable out of the box, and easy to manage. DCS supports single-node, master/standby, hot standby, cluster, and read/write splitting instances. It can meet your requirements for high read/write performance, fast data access, and auto scaling.

You can monitor DCS Redis instances using MAS and switch data centers in the event of a failure.

Data Replication Service

DRS is an easy-to-use, stable, and efficient cloud service for online database migration and real-time database synchronization.

MAS has weak dependency on DRS. You can use the DRS plug-in in MAS to implement online database migration and real-time database synchronization.

Log Tank Service

LTS can collect, analyze, and store log data. You can use LTS for efficient device O&M, service trend analysis, security audits, and monitoring.

MAS strongly depends on LTS and uses LTS to store workflow execution logs.

IAM

IAM provides identity authentication, permissions management, and access control.

With IAM, you can control access to MAS.

Cloud Eye

Cloud Eye is a secure, scalable, and integrated monitoring service. With Cloud Eye, you can monitor your MAS service and configure alarm rules and notifications.

Cloud Trace Service (CTS)

CTS allows you to collect, store, and query cloud resource operation records for security analysis, compliance auditing, resource tracking, and fault locating.

After you enable CTS, the system starts to record MAS operations. CTS stores operation records from the last seven days.

