Distributed Message Service for RocketMQ

Service Overview

Issue 01

Date 2024-03-29

HUAWEI

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.

Contents

1 Infographics

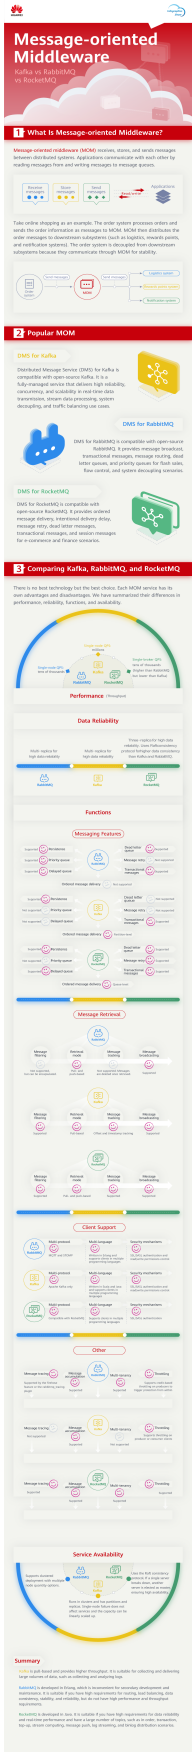
Message-oriented Middleware
Kafka vs RabbitMQ vs RocketMQ
1 What Is Message-oriented Middleware?
2 Popular MOM
DMS for Kafka
DMS for RabbitMQ
DMS for RocketMQ
3 Comparing Kafka, RabbitMQ, and RocketMQ
Performance
Data Reliability
Functions
Messaging Features
Message Retrieval
Client Support
Other
Service Availability
Summary

2 What Is DMS for RocketMQ?

Distributed Message Service (DMS) for RocketMQ is message-oriented middleware that delivers low latency, high flexibility, high throughput, dynamic expansion, easy management, and abundant messaging functions.

DMS for RocketMQ has the following features:

- Compatibility with open-source RocketMQ clients.
- Abundant messaging functions, including ordered message delivery, delayed messages, scheduled messages, message retry, dead letter messages, and transactional messages, which meet diverse needs in e-commerce and finance scenarios.
- Monitoring and analysis functions, including message tracing, message tracking, trace analysis, dead letter message export, monitoring and alarms, which allow you to monitor your services and keep them up and running.

Product Architecture

Figure 2-1 Product architecture

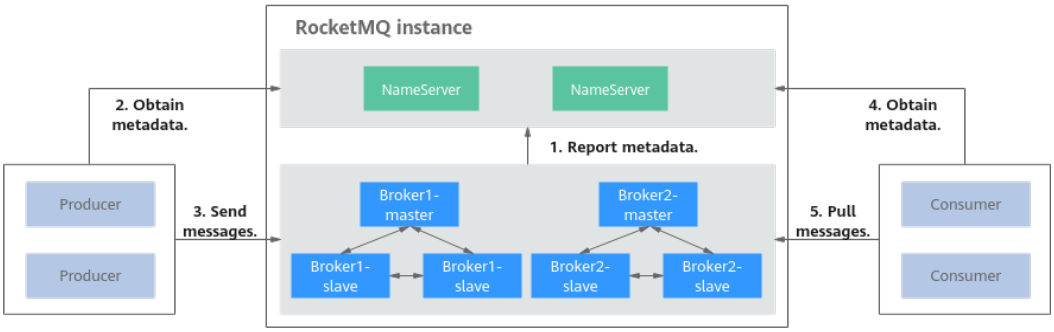

Architecture description:

- Brokers receive and store messages sent by producers or forward messages to consumers. A broker consists of one master and two slave nodes.
- NameServers receive and store broker metadata.
- Producers obtain metadata from NameServers and send messages to brokers.

- Consumers obtain metadata from NameServers and then pull messages from brokers.

Message Types

DMS for RocketMQ supports four messages types.

- Normal messages: Messages that do no have any features of delayed messages, ordered messages, or transactional messages.
- Delayed/Scheduled messages: Messages that are delivered to consumers after a specific period after being sent from producers to DMS for RocketMQ.
- Ordered messages: Messages that are retrieved in the exact order that they are created.
- Transactional message: Messages that achieve eventual consistency, delivering distributed transaction processing similar to X/Open XA.

Advanced Features

DMS for RocketMQ provides four advanced features.

- Message filtering: Consumers can tag messages in the subscribed topics to retrieve only the required messages.
- Message retry: DMS for RocketMQ retries sending messages that fail to be retrieved. If the message retrieval still fails after the maximum number of retries is reached, the messages are sent to the dead letter queue.

  Table 2-1 describes the retry mechanism in DMS for RocketMQ.

Table 2-1 Message retry mechanism

Retrieval Type	Retry Interval	Maximum Retries
Ordered retrieval	The interval is set by the **suspendTimeMillis** parameter. The default interval is 1000 ms, that is, 1s.	The maximum number of retries is set by the **setMaxReconsume-Times** method on the consumer. If this parameter is not set, there will be unlimited retries.
Normal retrieval	The interval varies by the retry number, as shown in Table 2-2.	Set during consumer group creation. Value range: 1–16 times

Table 2-2 Intervals between retries in normal retrievals

Retry No.	Interval	Retry No.	Interval
1	10s	9	7 min
2	30s	10	8 min
3	1 min	11	9 min
4	2 min	12	10 min
5	3 min	13	20 min
6	4 min	14	30 min
7	5 min	15	1 h
8	6 min	16	2 h

- Delayed messages: After being sent from producers to DMS for RocketMQ, messages are delivered to consumers only after a **fixed period**. A producer can specify one of the 18 delay levels listed in **Table 2-3**.

Table 2-3 Delay levels

Delay Level	Delay	Delay Level	Delay
1	1s	10	6 min
2	5s	11	7 min
3	10s	12	8 min
4	30s	13	9 min
5	1 min	14	10 min
6	2 min	15	20 min
7	3 min	16	30 min
8	4 min	17	1 h
9	5 min	18	2 h

- Scheduled messages: After being sent from producers to DMS for RocketMQ, messages are delivered to consumers only after a specified time point. In DMS for RocketMQ, you can schedule messages to be delivered at **any time** within one year. You can also cancel scheduled messages.

  NOTE

  This function is supported only for instances created on or after March 30, 2022.

3 Product Advantages

DMS for RocketMQ has the following advantages:

- Instant availability: You can build your own message service in the cloud in just a few steps. After you create a RocketMQ instance, you can quickly access it at its connection addresses. DMS for RocketMQ is fully compatible with open-source RocketMQ, allowing you to migrate your applications to the cloud with no change to the application code.
- Fully hosted services: DMS for RocketMQ provides automatic deployment, after-sales services, and comprehensive O&M solutions including monitoring and alarms. You can focus on your service development without having to worry about deployment and O&M.
- Low latency: For the deployment on Huawei Cloud, intranet access latency is reduced to microseconds.
- Elasticity and high reliability: Raft-based clustering facilitates fault detection and failover, ensuring service continuity and reliability.
- Dynamic scaling: Clusters can be scaled out dynamically as service demand changes.
- Easy management: Monitoring and alarms, link diagnosis, and message tracing facilitate fault locating and routine maintenance.
- Abundant messaging functions: Ordered message delivery, intentional delivery delay, message retry, dead letter messages, message filtering, and transaction messages meet needs in diverse scenarios.

4 Application Scenarios

E-Commerce

E-commerce applications face the following challenges:

- E-commerce applications usually involve order placement, payment, and notification processing. Multiple service systems work together to complete a task, and upper-layer service systems depend on lower-layer service systems. If there are a large number of layers, the performance of user-facing services will deteriorate, affecting user experience.
- During e-commerce shopping events, service systems that provide subscription-based notifications are required.

DMS for RocketMQ can help you cope with e-commerce challenges.

- DMS for RocketMQ decouples service systems to improve their processing capabilities and response speeds.
- DMS (for RocketMQ) allows for scheduled and delayed message delivery to meet the requirements for subscription-based notifications.

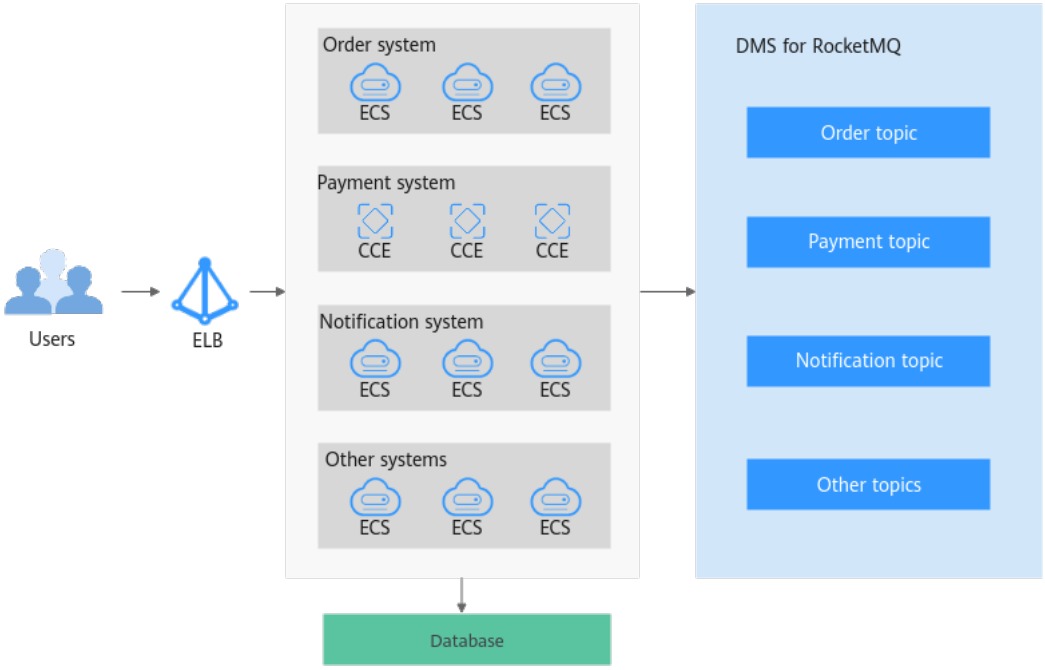

Finance

Unlike traditional finance, Internet finance requires fast responses to the rapidly changing Internet landscape. Common functions in traditional finance are extracted and then packaged into microservice modules to form a middle platform which provides basic capabilities for upper-layer service systems. With its decoupling and linking advantages, DMS for RocketMQ improves the processing capability and response speed of each microservice module, playing an indispensable role in the service middle platform.

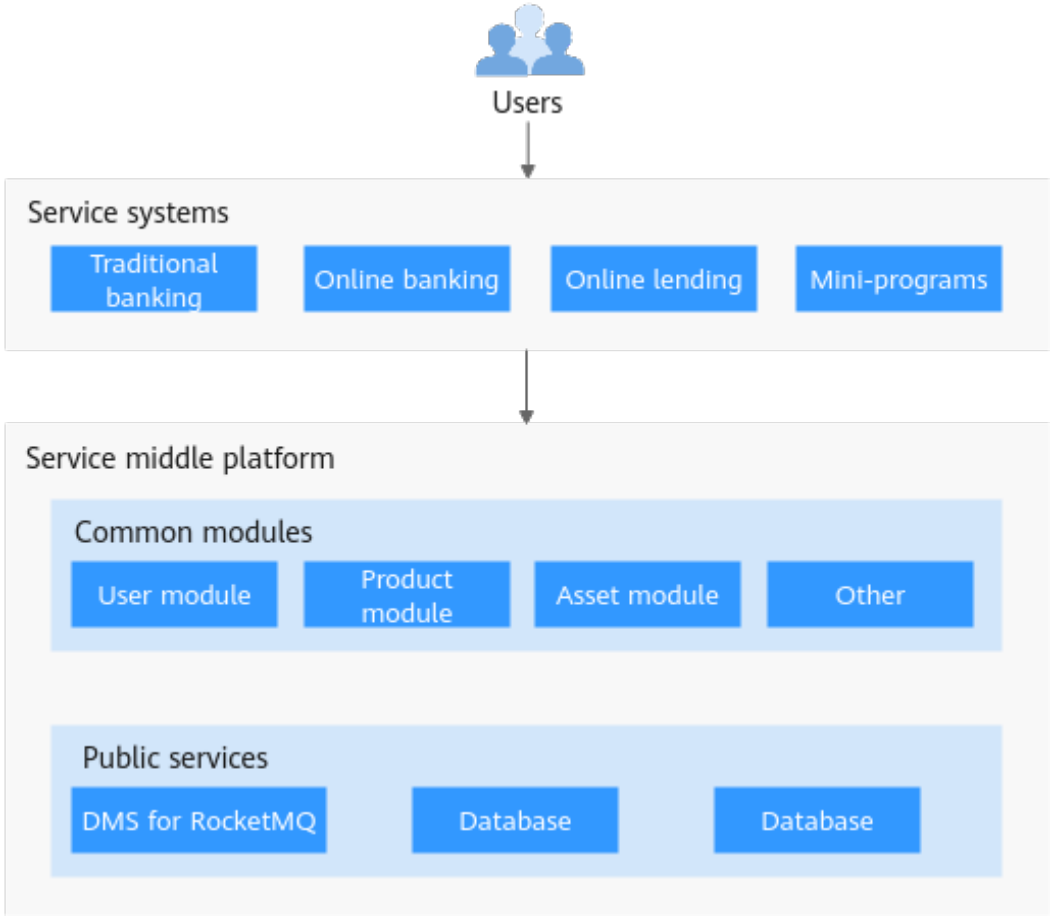

IoT

Massive device access is typical in IoT scenarios. A massive number of devices bring in massive amounts of data. However, each service component only desires specific data. Therefore, it is important to quickly identify data that services desire from a large amount of data. To address this challenge, DMS for RocketMQ provides message filtering. When sent from devices, messages are tagged to specify the message type. Services can specify only the desired types of messages to retrieve.

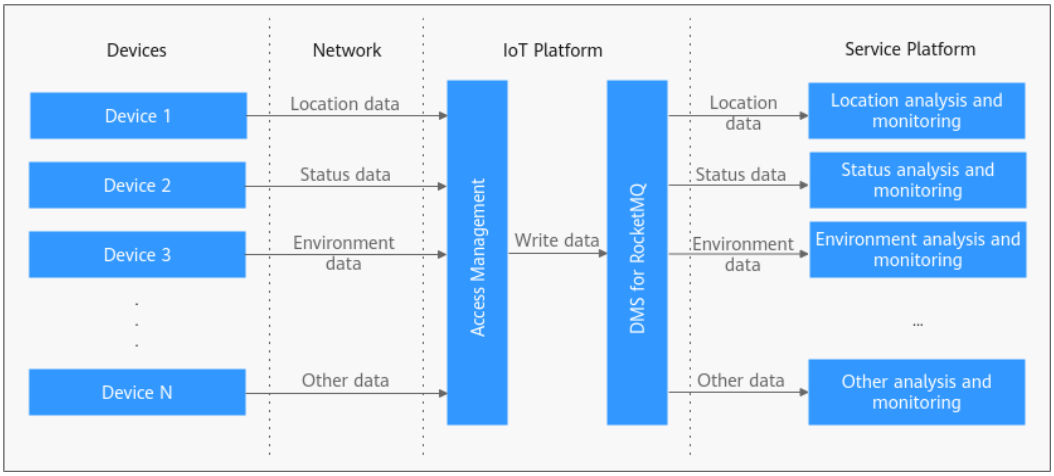

5 Specifications

DMS for RocketMQ is compatible with open-source RocketMQ 4.8.0. Their specifications are described as follows.

DMS for RocketMQ 4.8.0

The specifications of a RocketMQ 4.8.0 instance are defined on six dimensions:

- ECS flavor: flavor of the ECS used by the instance
- Number of brokers: scale of the instance
- Storage space: storage capacity available for each broker
- TPS per broker: TPS performance of each broker
- Max. topics per broker: number of topics allowed for each broker
- Max. consumer groups per broker: number of consumer groups allowed for each broker

Table 5-1 lists the flavors of DMS for RocketMQ 4.8.0.

NOTE

TPS (Transaction per second) indicates the total number of produced and consumed messages per second, that is, the total throughput of message production and consumption per second for a flavor.

Table 5-1 Instance flavors (DMS for RocketMQ 4.8.0)

Flavor	Brokers	Storage Space (GB per Broker)	TPS per Broker	Max. Topics per Broker	Max. Consumer Groups per Broker
rocketmq.4u8g.cluster.small	1–2	300–60,000	15,000	2000	2000
rocketmq.4u8g.cluster	1–10	300–600,000	20,000	4000	4000
rocketmq.8u16g.cluster	1–10	300–900,000	25,000	8000	8000

Flavor	Brokers	Storage Space (GB per Broker)	TPS per Broker	Max. Topics per Broker	Max. Consumer Groups per Broker
rocketmq.12u24g.cluster	1–10	300–900,000	28,000	12,000	12,000
rocketmq.16u32g.cluster	1–10	300–900,000	30,000	16,000	16,000

6 Comparing RocketMQ, Kafka, and RabbitMQ

Table 6-1 Functions

Feature	RocketMQ	Kafka	RabbitMQ
Priority queue	Not supported	Not supported	Supported. It is recommended that the priority be set to 0–10.
Delayed queue	Supported	Not supported	Supported
Dead letter queue	Supported	Not supported	Supported
Message retry	Supported	Not supported	Not supported
Retrieval mode	Pull-based and push-based	Pull-based	Pull-based and push-based
Message broadcasting	Supported	Supported	Supported
Message tracking	Supported	Supports offset and timestamp tracking.	Not supported. Once a message retrieval has been acknowledged, RabbitMQ will be notified that the message can be deleted.

Feature	RocketMQ	Kafka	RabbitMQ
Message accumulation	Supported	Supports higher accumulation performance than RabbitMQ thanks to high throughput.	Supported
Persistence	Supported	Supported	Supported
Message tracing	Supported	Not supported	Supported by the firehose feature.
Message filtering	Supported	Supported	Not supported, but can be encapsulated.
Multi-tenancy	Supported	Not supported	Supported
Multi-protocol	Compatible with RocketMQ.	Only supports Apache Kafka.	RabbitMQ is based on AMQP and supports MQTT and STOMP.
Multi-language	Supports clients in multiple programming languages.	Kafka is written in Scala and Java and supports clients in multiple programming languages.	RabbitMQ is written in Erlang and supports clients in multiple programming languages.
Throttling	Planned	Supports throttling on producer or consumer clients.	Supports credit-based throttling on producers, a mechanism that triggers protection from within.
Ordered message delivery	Message order is maintained within a queue.	Supports partition-level FIFO.	Not supported. Supports FIFO only for single-threaded message queuing without advanced features such as delayed queues or priority queues.
Security	Supports SSL authentication.	Supports SSL and SASL authentication and read/write permissions control.	Supports SSL.

Feature	RocketMQ	Kafka	RabbitMQ
Transactional messages	Supported	Supported	Supported

7 Comparing DMS for RocketMQ and Open-Source RocketMQ

DMS for RocketMQ maintains open-source compatibility while supporting or enhancing features in the open-source versions.

Table 7-1 Differences between DMS for RocketMQ and open-source RocketMQ

Feature	DMS for RocketMQ	Open-Source RocketMQ
Delayed/Scheduled messages	• Delayed message: There are 18 delay levels. The longest delay is two hours. • Scheduled messages: The message delivery delay can be customized. The longest delay is one year.	Only 18 delay levels are supported. The longest delay is two hours.
Ordered messages	Supported	Supported
Message retry	Supported	Supported
Dead letter messages	Supported	Supported
Cluster consumption	Supported	Supported
Message broadcasting	Supported	Supported
Dead letter queues	Supported	Supported
Resetting message retrieval start position	Supported	Supported
Message query	Supported	Supported
Encrypted transmission	Supported	Supported
Message tracing	Supported	Supported

Feature	DMS for RocketMQ	Open-Source RocketMQ
Transactional messages	Supported. The performance is ten times higher than that of open-source RocketMQ when a large number of transactional messages are stacked.	Supported. However, if a large number of transactional messages are stacked, the performance is poor.
Exporting dead letter messages	Supported	Not supported
Data dumping	Planned	Not supported
Instance diagnosis	One-click diagnosis of retrieval issues	Not supported
Instance monitoring	Historical monitoring data can be displayed in graphs. More than 18 metrics can be monitored.	Only current monitoring data can be viewed. Metrics such as CPU and memory cannot be monitored.
ACL	Flexible configuration with a few clicks	Complex configuration
O&M and capacity expansion	Fast expansion with a few clicks	Manual expansion, which is complex

8 Security

8.1 Shared Responsibilities

Huawei guarantees that its commitment to cyber security will never be outweighed by the consideration of commercial interests. To cope with emerging cloud security challenges and pervasive cloud security threats and attacks, Huawei Cloud builds a comprehensive cloud service security assurance system for different regions and industries based on Huawei's unique software and hardware advantages, laws, regulations, industry standards, and security ecosystem.

Figure 8-1 illustrates the responsibilities shared by Huawei Cloud and users.

- **Huawei Cloud**: Ensure the security of cloud services and provide secure clouds. Huawei Cloud's security responsibilities include ensuring the security of our IaaS, PaaS, and SaaS services, as well as the physical environments of the Huawei Cloud data centers where our IaaS, PaaS, and SaaS services operate. Huawei Cloud is responsible for not only the security functions and performance of our infrastructure, cloud services, and technologies, but also for the overall cloud O&M security and, in the broader sense, the security and compliance of our infrastructure and services.
- **Tenant**: Use the cloud securely. Tenants of Huawei Cloud are responsible for the secure and effective management of the tenant-customized configurations of cloud services including IaaS, PaaS, and SaaS. This includes but is not limited to virtual networks, the OS of virtual machine hosts and guests, virtual firewalls, API Gateway, advanced security services, all types of cloud services, tenant data, identity accounts, and key management.

Huawei Cloud Security White Paper elaborates on the ideas and measures for building Huawei Cloud security, including cloud security strategies, the shared responsibility model, compliance and privacy, security organizations and personnel, infrastructure security, tenant service and security, engineering security, O&M security, and ecosystem security.

Figure 8-1 Huawei Cloud shared security responsibility model

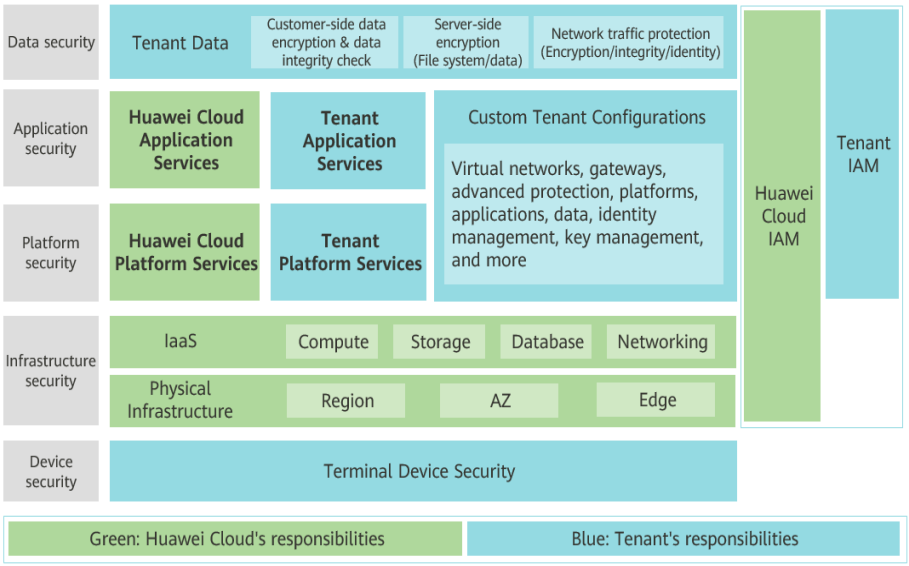

8.2 Identity Authentication and Access Control

Identity Authentication

Access requesters must present the identity credential for identity validity verification when accessing DMS for RocketMQ on the console or by calling APIs. DMS for RocketMQ uses Identity and Access Management (IAM) to provide three identity authentication modes: **passwords**, **access keys**, and **temporary access keys**. In addition, DMS for RocketMQ provides **login protection** and **login authentication policies** to harden identity authentication security.

Access Control

You can assign different permissions for DMS for RocketMQ to employees in your organization for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you secure access to your Huawei Cloud resources. For details about DMS for RocketMQ permissions, see **Permissions Management**.

8.3 Data Protection

DMS for RocketMQ takes different measures to keep data secure and reliable.

Table 8-1 DMS for RocketMQ data protection methods and features

Measure	Description	Reference
DR and multi-active	To meet reliability requirements of your data and services, you can deploy a RocketMQ instance in a single AZ (single equipment room) or across AZs (intra-city DR).	**Deploying an Instance Within an AZ or Across AZs**
Data replication	Each master has two slaves. To ensure data consistency, the slaves synchronize data from the master. When a network exception or node fault occurs, the master/slave roles are automatically switched by using the Raft protocol.	-
Data persistence	Exceptions may occur during daily running of the service system. Some service systems require high reliability, including high availability of instances, data security, and recoverability, so that backup data can be used to restore instances if an exception occurs, ensuring service running.	-

8.4 Audit and Logs

Cloud Trace Service (CTS) records operations on the cloud resources in your account. You can use the logs generated by CTS to perform security analysis, trace resource changes, audit compliance, and locate faults.

After you enable CTS and configure a tracker, CTS records management traces of DMS for RocketMQ for auditing.

For details about how to enable and configure CTS, see **Enabling CTS**.

For details about the DMS for RocketMQ management traces that can be recorded, see **Operations Logged by CTS**.

Figure 8-2 CTS

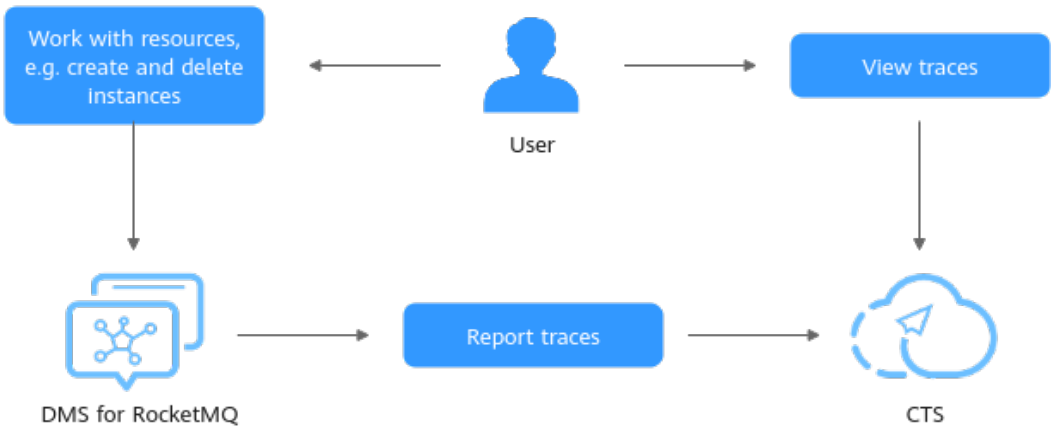

8.5 Resilience

DMS for RocketMQ provides a three-level reliability architecture and uses cross-AZ DR, intra-AZ instance DR, and instance data replication to ensure service durability and reliability.

Table 8-2 Reliability architecture of DMS for RocketMQ

Reliability Solution	Description
Cross-AZ DR	DMS for RocketMQ provides cross-AZ instances that support cross-AZ DR. When an AZ is abnormal, the instances can still provide services.
Intra-AZ instance DR	Each master has two slaves. To ensure strong data consistency, the slaves synchronize data from the master. When a node fault occurs, the master/slave roles are automatically switched by using the Raft protocol.
Data DR	Data DR is implemented through data replication.

8.6 Security Risks Monitoring

DMS for RocketMQ uses Cloud Eye to help you monitor your RocketMQ instances and receive alarms and notifications in real time. You can obtain key information about instances in real time, such as service requests, resource usage, traffic, number of connections, and number of accumulated messages.

For details about DMS for RocketMQ metrics and how to create alarm rules, see **RocketMQ Metrics**.

8.7 Certificates

Compliance Certificates

Huawei Cloud services and platforms have obtained various security and compliance certifications from authoritative organizations, such as International Organization for Standardization (ISO). You can **download** them from the console.

Figure 8-3 Downloading compliance certificates

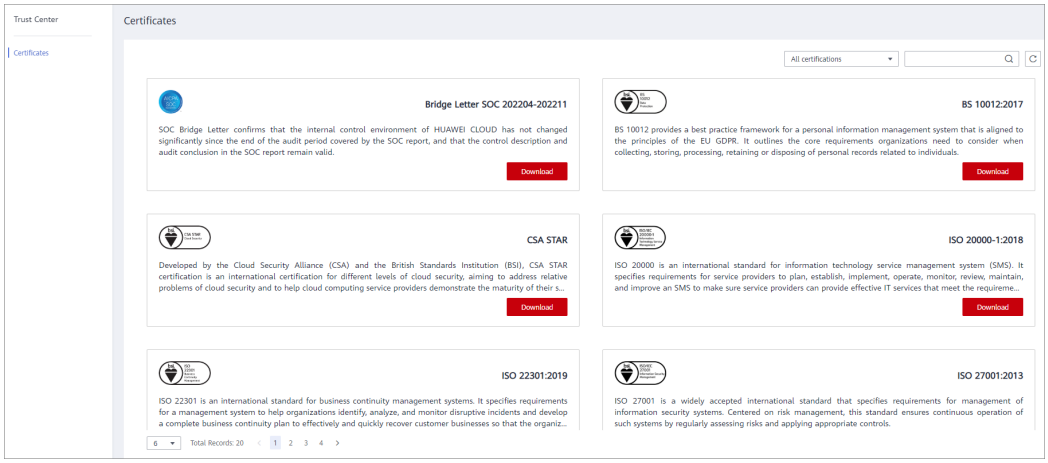

Resource Center

Huawei Cloud also provides the following resources to help users meet compliance requirements. For details, see **Resource Center**.

Figure 8-4 Resource center

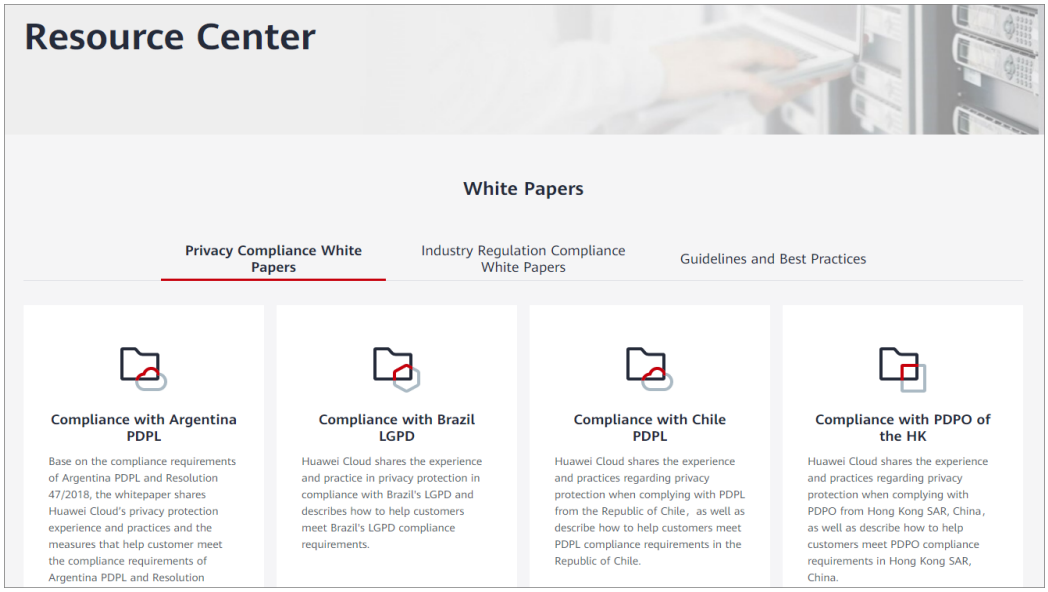

9 Restrictions

This section describes the restrictions on using DMS for RocketMQ.

Table 9-1 DMS for RocketMQ restrictions

Item	Restrictions
Maximum number of topics	Depends on the **instance specifications**. This limit cannot be changed. When this limit is reached, no more topics can be created.
Maximum number of consumer groups	Depends on the **instance specifications**. This limit cannot be changed. When this limit is reached, no more consumer groups can be created.
Message size	The maximum size of a message is 4 MB. The maximum size of a message attribute is 16 KB. The message size cannot be changed. When this limit is reached, messages will fail to be sent.
Message retention duration	The default retention period is 48 hours, which can be modified to up to 720 hours. Messages will be automatically deleted after the retention period expires.
Consumer offset reset	You can reset the retrieval start position to any time within two days.
Delay of scheduled messages	The maximum delay is one year. This limit cannot be changed. You can schedule messages to be delivered at any time within one year.
Request-Reply	Not supported.
Configuration parameters	Cannot be modified through open-source APIs.

Item	Restrictions
Broker faults	When some brokers of an instance are faulty, topics, consumer groups, and users cannot be created, modified, or deleted, but can be queried.

10 Related Services

- Virtual Private Cloud (VPC)

  RocketMQ instances run in VPCs and use the IP addresses and bandwidth of VPC. Security groups of VPCs enhance the security of network access to the instances.

- Cloud Eye

  Cloud Eye is an open platform that provides monitoring, alarm reporting, and alarm notification on your resources in near real time.

- Cloud Trace Service (CTS)

  Cloud Trace Service (CTS) generates traces to provide you with a history of operations performed on cloud service resources. The traces include operation requests sent using the Huawei Cloudmanagement console or open APIs, as well as the operation results. You can view all generated traces to query, audit, and backtrack performed operations.

- Elastic Cloud Server (ECS)

  An ECS is a basic computing unit that consists of vCPUs, memory, OS, and EVS disks. RocketMQ instances run on ECSs. A broker corresponds to three ECSs.

- Elastic Volume Service (EVS)

  EVS provides block storage services for ECSs. All RocketMQ data, such as messages, metadata, and logs, is stored in EVS disks.

- Elastic IP (EIP)

  The EIP service provides independent public IP addresses and bandwidth for Internet access. RocketMQ instances bound with EIPs can be accessed over public networks.

- Tag Management Service (TMS)

  TMS is a visualized service for fast and unified cross-region tagging and categorization of cloud services.

  Tags facilitate RocketMQ instance identification and management.

11 RocketMQ Concepts

Topic

A topic is a category of messages. It is the basic unit for creating and retrieving messages.

Queue

A topic consists of multiple queues. A larger number of queues indicates higher retrieval concurrency.

Producer

A producer sends messages to the server.

Producer Group

Producers in a group send the same type of messages with the same logic.

Consumer

A consumer obtains messages from the server.

Consumer Group

A consumer group contains consumers that have similar retrieval behavior.

Broker

Brokers are a cluster of ECSs that process services.

NameServer

A lightweight registry that stores metadata. Before producing or consuming messages, producers and consumers must obtain metadata from a NameServer.

12 Permissions Management

You can use Identity and Access Management (IAM) to manage DMS for RocketMQ permissions and control access to your resources. IAM provides identity authentication, permissions management, and access control, helping you secure access to your Huawei Cloud resources.

You can create IAM users for your employees, and assign permissions to these users on a principle of least privilege (PoLP) basis to control their access to specific resource types. For example, you can create IAM users for software developers and assign specific permissions to allow them to use DMS for RocketMQ resources but prevent them from being able to delete resources or perform any high-risk operations.

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

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

NOTE

DMS for RocketMQ permissions policies are based on DMS. Therefore, when assigning permissions, select DMS permissions policies.

DMS for RocketMQ Permissions

By default, new IAM users do not have any permissions assigned. To assign permissions to these new users, add them to one or more groups, and attach permissions policies or roles to these groups.

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

You can grant users permissions by using roles and policies.

- Roles: A type of coarse-grained authorization mechanism that provides only a limited number of service-level roles. When using roles to grant permissions,

you also need to assign dependency roles. However, roles are not an ideal choice for fine-grained authorization and secure access control.

- Policies: A fine-grained authorization strategy that defines permissions required to perform operations on specific cloud resources under certain conditions. This mechanism allows for more flexible policy-based authorization for securer access control. For example, you can grant DMS for RocketMQ users only the permissions for managing DMS for RocketMQ instances. Most policies define permissions based on APIs. For the API actions supported by DMS for RocketMQ, see **Permissions Policies and Supported Actions**.

Table 12-1 lists all the system-defined roles and policies supported by DMS for RocketMQ.

Table 12-1 System-defined roles and policies supported by DMS for RocketMQ

Role/Policy Name	Description	Type	Dependencies
DMS FullAccess	Administrator permissions for DMS. Users granted these permissions can perform all operations on DMS.	System-defined policy	None
DMS UserAccess	Common user permissions for DMS, excluding permissions for creating, modifying, deleting, and scaling up instances.	System-defined policy	None
DMS ReadOnlyAccess	Read-only permissions for DMS. Users granted these permissions can only view DMS data.	System-defined policy	None
DMS VPCAccess	VPC operation permissions to assign to DMS agencies.	System-defined policy	None
DMS KMSAccess	KMS operation permissions to assign to DMS agencies.	System-defined policy	None
DMS Administrator	Administrator permissions for DMS.	System-defined role	This role depends on the **Tenant Guest** and **VPC Administrator** roles.

Table 12-2 lists the common operations supported by each DMS for RocketMQ system policy or role. Select the policies or roles as required.

Table 12-2 Common operations supported by each system-defined policy or role of DMS for RocketMQ

Operation	DMS FullAccess	DMS UserAccess	DMS ReadOnlyAccess	DMS VPCAccess	DMS KMSAccess
Creating an instance	√	×	×	×	×
Modifying instances	√	×	×	×	×
Deleting instances	√	×	×	×	×
Modifying instance specifications	√	×	×	×	×
Querying instance information	√	√	√	×	×

Helpful Links

- What Is IAM?
- Creating a User and Granting DMS for RocketMQ Permissions
- Permissions Policies and Supported Actions

