

Distributed Message Service for Kafka

Getting Started

Issue 01
Date 2024-07-17



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2024. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Cloud Computing Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI 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.

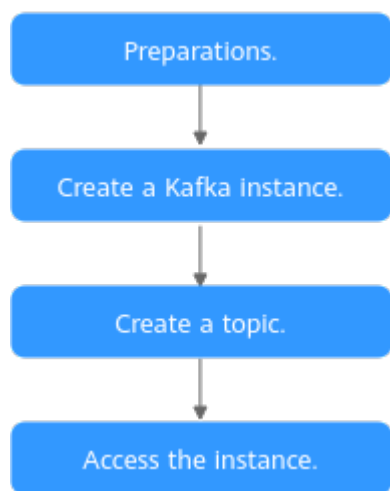
Contents

1 Getting Started with Kafka for Message Production and Consumption.....	1
2 Common Practices.....	14

1 Getting Started with Kafka for Message Production and Consumption

This section takes the example of creating a Kafka instance (ciphertext access and SASL_SSL) and accessing it on the client (private network, within a virtual private cloud (VPC)) for message production and consumption to get you quickly started with Distributed Message Service (DMS) for Kafka.

Figure 1-1 Procedure for using DMS for Kafka



1. **Step 1: Preparations**

A Kafka instance runs in a VPC. Before creating a Kafka instance, ensure that a VPC is available.

After a Kafka instance is created, download and install the Kafka open-source client on your ECS before producing and consuming messages.

2. **Step 2: Create a Kafka Instance**

You can select the specification and quantity, and enable ciphertext access and SASL_SSL when creating a Kafka instance.

When connecting to a Kafka instance with SASL_SSL enabled, SASL is used for authentication. Data is encrypted with SSL certificates for high-security transmission.

3. **Step 3: Create a Topic**
Topics store messages created by producers and subscribed by consumers. This section uses the example of creating a topic on the console.
4. **Step 4: Connect to a Kafka Instance to Produce and Consume Messages**
Before connecting to a Kafka instance with SASL_SSL enabled, download the certificate and configure the connection in the client configuration file.

Step 1: Preparations

Step 1 Register with Huawei Cloud and complete real-name authentication.

For details, see [Signing up for a HUAWEI ID and Enabling Huawei Cloud Services](#) and [Real-Name Authentication](#).

If you already have a Huawei account and have completed real-name authentication, skip this step.

Step 2 Top up the account.

Ensure that your account has sufficient balance before creating a Kafka instance. For details about how to top up an account, see [Topping Up an Account](#).

Step 3 Grant Kafka instance permissions.

To achieve fine-grained management of your cloud resources, create Identity and Access Management (IAM) user groups and users and grant specified permissions to the users. For more information, see [Creating a User and Granting DMS for Kafka Permissions](#).

Step 4 Create a VPC and subnet.

Before creating a Kafka instance, ensure that a VPC and a subnet are available. For details about how to create a VPC and a subnet, see [Creating a VPC and Subnet](#).

NOTICE

The VPC must be created in the same region as the Kafka instance.

Step 5 Create a security group and add security group rules.

Before creating a Kafka instance, ensure that a security group is available. For details about how to create a security group, see [Creating a Security Group](#).

To connect to Kafka instances, add the security group rules described in [Table 1-1](#). Other rules can be added based on site requirements.

Table 1-1 Security group rules


Direction	Protocol	Port	Source address	Description
Inbound	TCP	9093	0.0.0.0/0	Accessing a Kafka instance over a private network within a VPC (in ciphertext)

NOTE

After a security group is created, it has a default inbound rule that allows communication among ECSs within the security group and a default outbound rule that allows all outbound traffic. If you access your Kafka instance using the private network within a VPC, you do not need to add the rules described in [Table 1-1](#).

Step 6 Construct a client for message production and consumption.

This section uses a Linux elastic cloud server (ECS) as the client. Before creating a Kafka instance, create an ECS with an EIP, install the JDK, configure environment variables, and download an open-source Kafka client.

1. Log in to the console, click  in the upper left corner, click **Elastic Cloud Server** under **Computing**, and then create an ECS.
For details about how to create an ECS, see [Purchasing an ECS](#). If you already have an available ECS, skip this step.
2. Log in to an ECS as user **root**.
3. Install Java JDK and configure the environment variables **JAVA_HOME** and **PATH**.
 - a. Download a JDK.

NOTE

Use Oracle JDK instead of ECS's default JDK (for example, OpenJDK), because ECS's default JDK may not be suitable. Obtain Oracle JDK 1.8.111 or later from [Oracle's official website](#).

- b. Decompress the JDK.

```
tar -zxvf jdk-8u321-linux-x64.tar.gz
```


Change **jdk-8u321-linux-x64.tar.gz** to your JDK version.
- c. Open the **.bash_profile** file.

```
vim ~/.bash_profile
```
- d. Add the following content:

```
export JAVA_HOME=/root/jdk1.8.0_321
export PATH=$JAVA_HOME/bin:$PATH
```


Change **/root/jdk1.8.0_321** to the path where you install JDK.
- e. Press **Esc**. Enter the following line and press **Enter**. Save the **.bash_profile** file and exit.

```
:wq
```
- f. Run the following command to make the change take effect:

```
source .bash_profile
```
- g. Check whether the JDK is installed.

```
java -version
```

If the following message is returned, the JDK is installed.

```
java version "1.8.0_321"
```

4. Download an open-source Kafka client.

```
wget https://archive.apache.org/dist/kafka/2.7.2/kafka_2.12-2.7.2.tgz
```
5. Run the following command to decompress the package:

```
tar -zxf kafka_2.12-2.7.2.tgz
```

----End

Step 2: Create a Kafka Instance

Step 1 Go to the [Buy Instance page](#).

Step 2 Set basic instance information shown in [Figure 1-2](#). [Table 1-2](#) lists the configuration details.

Table 1-2 Basic instance settings

Parameter	Description
Billing Mode	Select Pay-per-use , which is a postpaid mode. You can pay after using the service, and will be billed for your usage duration. The fees are calculated in seconds and settled by hour.
Region	DMS for Kafka in different regions cannot communicate with each other over an intranet. Select a nearest location for low latency and fast access. Select AP-Singapore.
Project	Projects isolate compute, storage, and network resources across geographical regions. For each region, a preset project is available. Select AP-Singapore (default).
AZ	An AZ is a physical region where resources use independent power supply and networks. AZs are physically isolated but interconnected through an internal network. Select AZ1 , AZ2 , and AZ3 .
Instance Name	You can customize a name that complies with the rules: 4–64 characters; starts with a letter; can contain only letters, digits, hyphens (-), and underscores (_). Enter kafka-test .
Enterprise Project	This parameter is for enterprise users. An enterprise project manages project resources in groups. Enterprise projects are logically isolated. Select default .
Specifications	Select Cluster to create a cluster Kafka instance.

Parameter	Description
Version	Kafka version. Cannot be changed once the instance is created. Select 2.7 .
Broker Flavor	Select a broker flavor as required. Select kafka.2u4g.cluster .
Brokers	Specify the number of brokers as required. Enter 3 .
Storage Space per Broker	Select the disk type and specify the disk size as required. Total storage space = Storage space per broker × Broker quantity. After the instance is created, you cannot change the disk type. Select Ultra-high I/O and enter 100 .
Capacity Threshold Policy	Select Automatically delete : When the disk reaches the disk capacity threshold (95%), messages can still be produced and consumed, but the earliest 10% of messages will be deleted to ensure sufficient disk space. Use this policy for services intolerant of interruptions. However, data may be lost.

Figure 1-2 Configuring instance basic information

Billing Mode: Yearly/Monthly Pay-per-use

Region:

Project:

AZ: ⓘ
Select one AZ or at least three AZs. Do not select two AZs.
 The more AZs selected, the better the reliability and SLA coverage.

Instance Name:

Enterprise Project: ⓘ [View Enterprise Project](#)

Specifications: Cluster Custom Single-node ⓘ
Customize your Kafka instance.

Version:
Software version of the instance.
 Fixed once the instance is created. Use the same version as your client.

Flavor Name	TPS Limit per Broker	Maximum Partitions per Bro...	Recommended Consumer G...	Traffic per Broker (MB/s)
<input type="radio"/> kafka.2u4g.cluster.small	20,000	100	15	40
<input checked="" type="radio"/> kafka.2u4g.cluster	30,000	250	20	100
<input type="radio"/> kafka.4u8g.cluster	100,000	500	100	200
<input type="radio"/> kafka.8u16g.cluster	150,000	1,000	150	375
<input type="radio"/> kafka.12u24g.cluster	200,000	1,500	200	625
<input type="radio"/> kafka.16u32g.cluster	250,000	2,000	200	750

Currently Selected: **kafka.2u4g.cluster** | TPS Limit per Broker 30,000 | Maximum Partitions per Broker 250 | Recommended Consumer Groups per Broker 20 | Traffic per Broker (MB/s) 100

Brokers: ⓘ
Cannot be decreased once the instance is created.

Total traffic: 300 MB/s (Total traffic = Traffic per broker x Broker quantity)
 Service traffic:
 * Recommended read traffic: 50 MB/s
 * Recommended write traffic: 50 MB/s
[Learn more](#) ⓘ about total traffic and service traffic.
 Tip: Reserve 30% of traffic to accommodate traffic fluctuations.

Storage space per broker: GB ⓘ
Total storage space 300 GB
 After the instance is created, you cannot change the disk type or reduce the storage space. Select a disk type for your required service I/O. [Learn more](#) ⓘ
 Total disk bandwidth: 510 MB/s = Disk bandwidth per broker x Broker quantity
 Disk bandwidth per broker: 170 MB/s


Capacity Threshold Policy: Automatically delete Stop production ⓘ

Step 3 Configure the instance network as shown in [Figure 1-3](#). For details, see [Table 1-3](#).

Table 1-3 Configuring instance network

Parameter	Description
VPC	The VPC and subnet cannot be changed once the instance is created. Select the VPC and subnet prepared in Step 4 .
Security Group	Select the security group prepared in Step 5 .

Figure 1-3 Configuring instance network



Step 4 Set the instance access mode, as shown in [Figure 1-4](#). For details, see [Table 1-4](#).

Table 1-4 Setting the instance access mode

Parameter	Sub-Parameter	Description
Private Network Access	Plaintext Access	Skip it.
	Ciphertext Access	<p>When this parameter is enabled, SASL authentication is required when a client connects to the Kafka instance.</p> <ol style="list-style-type: none"> Ciphertext Access is enabled. SASL_SSL is selected. Username and Password can be set. Enter "test" for the username. The username cannot be changed once ciphertext access is enabled. A username should contain 4 to 64 characters, start with a letter, and contain only letters, digits, hyphens (-), and underscores (_).A password must meet the following requirements: <ul style="list-style-type: none"> Contains 8 to 32 characters. Contains at least three types of the following characters: uppercase letters, lowercase letters, digits, and special characters `~!@#\$%^&*()-_+=\ []{};:'",<.>?` and spaces, and cannot start with a hyphen (-). Cannot be the username spelled forwards or backwards. SASL/PLAIN is enabled.
	Private IP Addresses	Select Auto : The system automatically assigns IP addresses from the subnet.
Public Network Access	-	Skip it.

Figure 1-4 Configuring the instance access mode

The screenshot displays the configuration interface for instance access mode, divided into Private Network Access and Public Network Access sections.

Private Network Access:

- Plaintext Access: Disabled (toggle off).
- Ciphertext Access: Enabled (toggle on).
- Security Protocol: SASL_SSL (selected), SASL_PLAINTEXT.
- Security Note: Data is encrypted for transmission for high security.
- Cross-VPC Access Protocol: SASL_SSL. Fixed once the instance is created.
- Private IP Addresses: Auto (selected), Manual.

Public Network Access:

- Public Network Access: Disabled (toggle off).
- Username: test.
- Username Note: The username will be used for accessing the instance with SSL authentication and cannot be modified once the instance is created.
- Password: [Redacted]
- Confirm Password: [Redacted]
- SASL/PLAIN: Enabled (toggle on).
- SASL/PLAIN Note: Enable to support both SCRAM-SHA-512 (enabled by default) and PLAIN.

Step 5 Skip **Advanced Settings**.

Step 6 Click **Buy**.

Step 7 Confirm the instance settings.

Step 8 Return to the **DMS for Kafka** page and check whether the instance has been created.

It takes 3 to 15 minutes to create an instance. During this period, the instance status is **Creating**.

- If the instance is created successfully, its status changes to **Running**.
- If the instance is in the **Creation failed** state, delete it, and create a new one. If the instance creation fails again, contact customer service.

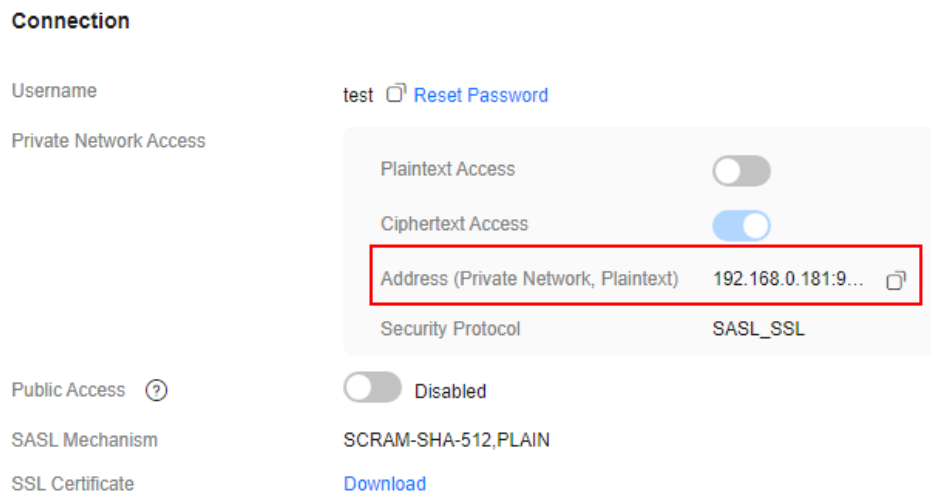
NOTE

Instances that fail to be created do not occupy other resources.

Step 9 After the instance is created, click its name to go to the instance details page.

Step 10 In the **Connection** area, view and record the connection address.

Figure 1-5 Kafka instance addresses (private network) for intra-VPC access



----End

Step 3: Create a Topic

- Step 1** On the **DMS for Kafka** page, click a Kafka instance.
- Step 2** In the navigation pane, choose **Topics**.
- Step 3** Click **Create Topic**.
- Step 4** Enter the topic name, specify other parameters by referring to **Table 1-5**, and click **OK**.

Table 1-5 Topic parameters

Parameter	Description
Topic Name	<p>Customize a name that contains 3 to 200 characters, starts with a letter or underscore (_), and contains only letters, digits, periods (.), hyphens (-), and underscores (_).</p> <p>The name must be different from preset topics:</p> <ul style="list-style-type: none"> • __consumer_offsets • __transaction_state • __trace • __connect-status • __connect-configs • __connect-offsets <p>Cannot be changed once the topic is created.</p> <p>Enter topic-01.</p>
Partitions	<p>If the number of partitions is the same as that of consumers, the larger the partitions, the higher the consumption concurrency.</p> <p>Enter 3.</p>

Parameter	Description
Replicas	Data is automatically backed up to each replica. When one Kafka broker becomes faulty, data is still available. A higher number of replicas delivers higher reliability. Enter 3 .
Aging Time (h)	How long messages will be preserved in the topic. Messages older than this period cannot be consumed. They will be deleted, and can no longer be consumed. Enter 72 .
Synchronous Replication	Skip it. When this option is disabled, leader replicas are independent from follower replica synchronization. They receive messages and write them to local logs, then immediately send the successfully written ones to the client.
Synchronous Flushing	Skip it. When this option is disabled, messages are produced and stored in memory instead of written to the disk immediately.
Message Timestamp	Select CreateTime : time when the producer created the message.
Max. Message Size (bytes)	Maximum batch processing size allowed by Kafka. If message compression is enabled in the client configuration file or code of producers, this parameter indicates the size after compression. Enter 10,485,760 .
Description	Skip it.

Figure 1-6 Creating a topic

Create Topic ✕

Topic Name

Partitions ? Value range: 1 to 200
Cannot be decreased once the topic is created.

Replicas ? Value range: 1 to 3
queue_term_fenBenCount_suggest_label
Cannot be greater than the broker quantity.

Aging Time (h) ? Value range: 1 to 720
How long messages will be preserved in the topic. Messages older than this period will be deleted and cannot be consumed.

Synchronous Replication ?

Synchronous Flushing ?

Message Timestamp ?

Max.Message Size (bytes) ?

Description
0/200 ↗

----End

Step 4: Connect to a Kafka Instance to Produce and Consume Messages

Step 1 Prepare the file for production and consumption configuration.

1. Log in to a Linux ECS.
2. Download the **client.jks** certificate and upload it to the **/root** directory on the ECS.

To obtain the certificate: On the Kafka console, click the Kafka instance to go to the **Basic Information** page. Click **Download** next to **SSL Certificate** in the **Connection** area. Decompress the package to obtain the client certificate file **client.jks**.

 NOTE

`/root` is the path for storing the certificate. Change it to the actual path if needed.

3. Go to the `/config` directory on the Kafka client.

```
cd kafka_2.12-2.7.2/config
```

4. Add the following commands in both the `consumer.properties` and `producer.properties` files (PLAIN is used as an example).

```
sasl.jaas.config=org.apache.kafka.common.security.plain.PlainLoginModule required \  
username="*****" \  
password="*****";  
sasl.mechanism=PLAIN
```

```
security.protocol=SASL_SSL  
ssl.truststore.location={ssl_truststore_path}  
ssl.truststore.password=dms@kafka  
ssl.endpoint.identification.algorithm=
```

Description:

- `username` and `password` are specified when enabling ciphertext access during instance creation.
- `ssl.truststore.location` is the path for storing the certificate obtained in [Step 1.2](#).
- `ssl.truststore.password` is certified by the server, which must be set to `dms@kafka` and cannot be changed.
- `ssl.endpoint.identification.algorithm` decides whether to verify the certificate domain name. In this example, **leave this parameter blank, which indicates disabling domain name verification**.

- Step 2 Go to the `/bin` directory on the Kafka client.

```
cd ../bin
```

- Step 3 Produce messages.

```
./kafka-console-producer.sh --broker-list ${connection address} --topic ${topic name} --producer.config ../  
config/producer.properties
```

Description:

- `{connection address}`: the connection address obtained in [Step 10](#)
- `{topic name}`: the topic name obtained in [Step 4](#)

For example, `192.xxx.xxx.xxx:9093`, `192.xxx.xxx.xxx:9093`, `192.xxx.xxx.xxx:9093` are the connection addresses of the Kafka instance.

After running this command, you can send messages to the Kafka instance by entering the information as prompted and pressing **Enter**. Each line of content will be sent as a message.

```
[root@ecs-kafka bin]#./kafka-console-producer.sh --broker-list  
192.xxx.xxx.xxx:9093,192.xxx.xxx.xxx:9093,192.xxx.xxx.xxx:9093 --topic topic-01 --producer.config ../config/  
producer.properties  
>Hello  
>DMS  
>Kafka!  
>^C[root@ecs-kafka bin]#
```

Press **Ctrl+C** to cancel.

- Step 4 Consume messages.

```
./kafka-console-consumer.sh --bootstrap-server ${connection address} --topic ${topic name} --from-  
beginning --consumer.config ../config/consumer.properties
```

Description:

- *{connection address}*: the connection address obtained in [Step 10](#)
- *{topic name}*: the topic name obtained in [Step 4](#)

Sample:

```
[root@ecs-kafka bin]# ./kafka-console-consumer.sh --bootstrap-server
192.xxx.xxx.xxx:9093,192.xxx.xxx.xxx:9093,192.xxx.xxx.xxx:9093 --topic topic-01 --from-beginning --
consumer.config ../config/consumer.properties
Hello
Kafka!
DMS
^CProcessed a total of 3 messages
[root@ecs-kafka bin]#
```

Press **Ctrl+C** to cancel.

----End

Related Information

- Learn more about the [basic concepts of Kafka](#).
- Learn more about [Distributed Message Service for Kafka Pricing](#).
- To view messages on the console, see [Viewing Kafka Messages](#).
- To view consumption offsets, see [Viewing and Resetting Kafka Consumption Offsets](#).
- To view monitoring metrics of a Kafka instance, see [Viewing Kafka Monitoring Metrics](#).

2 Common Practices

You can use the common practices provided by DMS for Kafka to meet your service requirements.

Table 2-1 Common practices

Practice	Description
Using MirrorMaker to Synchronize Data Across Clusters	MirrorMaker can synchronize data from the source cluster to the target cluster.
Configuring an Alarm Rule for Accumulated Messages	Configure an alarm rule so that the SMN service will notify you by SMS or email when the number of accumulated messages in a consumer group exceeds the threshold.