

TDMQ for CKafka

Operation Guide

Product Documentation



Copyright Notice

©2013-2024 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by Tencent Cloud Computing (Beijing) Company Limited and its affiliated companies. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

Contents

Operation Guide

Instance Management

Creating Instance

Naming with Consecutive Numeric Suffixes or Designated Pattern String

Viewing Instance

Upgrading Instance

Downgrading Instance Configuration

Terminating/Returning Instances

Change from Pay-as-You-Go to Monthly Subscription

Upgrading Instance Version

Adding Routing Policy

Public Network Bandwidth Management

Connecting to Prometheus

AZ Migration

Setting Maintenance Time

Setting Message Size

Topic Management

Instance Topic

Creating Topic

Viewing Topic

Deleting Topic

Sending Message

Setting Topic Traffic Throttling Rule

Elastic Topic

Topic Management

Subscription Management

Consumer Group

Consumer Group Overview

Creating Consumer Group

Querying Consumer Group

Deleting Consumer Group

Setting Offset

Automatically Resetting Offset

Monitoring and Alarms

Viewing Monitoring Data

- Querying Advanced Monitoring (Pro Edition)
- Configuring Alarms
- Suggestions for Alarm Configuration
- Viewing Advanced Ops Features (Pro Edition)

Smart Ops

- Elastic Scaling
 - Disk Utilization Policy
 - Dynamic Partition Processing
 - Elastic Bandwidth

Permission Management

- Configuring ACL Policy
- CAM

Tag Management

- Tag Overview
- Using Tags
- Editing Tags

Querying Message

Event Center

Migration to Cloud

- Migration Scheme Overview
- Step 1. Purchasing CKafka Instance
- Step 2. Migrating Topic to Cloud
- Step 3. Migrating Data to Cloud
 - Scheme 1: Single-Producer Dual-Consumer Migration
 - Scheme 2: Single-Producer Single-Consumer Migration
 - Scheme 3: Migration with MirrorMaker
 - Scheme 4: Migrating Unconsumed Data

Data Compression

Operation Guide

Instance Management

Creating Instance

Last updated : 2024-01-09 14:47:31

Overview

This document describes how to create instances and topics in the CKafka console.

Directions

1. Log in to the [CKafka console](#).

2. Click **Instance List** on the left sidebar and click **Create** to enter the purchase page.

3. On the instance purchase page, set the configuration information for purchase based on your actual needs.

Billing Mode: Pro Edition instances support both **monthly subscription** and **pay-as-you-go**, while Standard Edition instances support **monthly subscription** only.

Specs Type: Select the Standard or Pro Edition based on your business needs.

Kafka Version: Select a Kafka version based on your business needs. For more information, see [Suggestions for CKafka Version Selection](#).

Region: Select a region close to the resource for client deployment.

AZ: Select an AZ based on your actual needs.

Standard Edition: This edition does not support multi-AZ deployment.

Pro Edition: If the current region supports multi-AZ deployment, you can select up to three AZs for deployment. For more information on multi-AZ deployment, see [Multi-AZ Deployment](#).

Product Specification: Select a model based on the peak bandwidth and disk capacity.

Message Retention: Select a value between 24 and 2,160 hours.

When the disk capacity is insufficient (that is, the disk utilization reaches 90%), old messages will be deleted in advance to ensure the service availability.

VPC: If you need to access other VPCs, you can modify the routing access rules as instructed in [Adding Routing Policy](#).

Public Network Bandwidth: CKafka provides 3 Mbps public network bandwidth free of charge by default. CKafka Pro Edition allows you to increase the bandwidth of the public network. You can pay for more public network bandwidth if you require it.

Tag: It is optional. For more information, see [Tag Overview](#).

Instance Name: When purchasing multiple instances, you can batch create instances by its numeric suffix (which is numbered in an ascending order) or its designated pattern string. For detailed directions, see [Naming with Consecutive Numeric Suffixes or Designated Pattern String](#).

Auto-Renewal: The instance and public network bandwidth will be automatically renewed monthly if the account balance is sufficient.

4. You can view detailed billable items by hovering over the configuration fees in the bottom-right corner.

5. Click **Buy Now** to complete the instance creation process.

<input type="checkbox"/>	ID/Name	Monitor...	Status	AZ	Instance Type	Configuration	Network Type	Instance Billing M...	Tag
<input type="checkbox"/>	test		Healthy	Guangzhou Zone 6	Standard Edition-Advanced Version: 1.1.1 Disk Type: Premium Cloud Storage	Topic Limit: 50 Partition Limit: 150 Peak Bandwidth: 150 MB/s Disk Capacity: 2500GB	VPC	Pay as you go	

Naming with Consecutive Numeric Suffixes or Designated Pattern String

Last updated : 2024-01-09 14:47:32

Overview

To allow you to name batch created CKafka instances according to a rule during creation, the features of automatically incrementing suffixed numbers and specifying pattern strings are provided.

Automatically incrementing suffixed numbers: If you purchase multiple CKafka instances, numbers will be suffixed to their names by default, such as ckafka1, ckafka2, ckafka3... For more information, see [Automatically incrementing suffixed numbers](#).

Specifying the pattern string:

Specifying one pattern string: This is suitable if you need to create n instances and specify their names with incrementing serial numbers starting from x (such as ckafka3, ckafka4, ckafka5...). For more information, see [Specifying one pattern string](#).

Specifying multiple pattern strings: This is suitable if you need to create n instances named with multiple prefixes and specified serial numbers for each prefix (such as ckafka3-big10-test, ckafka4-big11-test, ckafka5-big12-test). For more information, see [Specifying multiple pattern strings](#).

Directions

Automatically incrementing suffixed numbers

This feature allows you to name batch purchased instances with the same prefix and automatically incrementing suffixed numbers.


Note:

The created instances are suffixed with numbers starting from 1 by default. You cannot specify the starting number. The following example assumes that you want to purchase 3 CKafka instances and name them in the format of "ckafka+serial number" (i.e., ckafka1, ckafka2, and ckafka3).

Operations on the purchase page

Operations through API

1. Purchase 3 instances as instructed in [Creating Instance](#). Enter the instance name in the format of **prefix+serial number**. In this case, enter `ckafka` as the instance name:



Instance Name ⓘ

Support naming with consecutive numeric suffixes (for batch instance creation) or designated pattern string. You can enter 58 more characters.

2. Follow the prompts on the page and make the payment.

In the [ModifyInstanceAttributes](#) API, set the relevant fields:

Instance name: Set `InstanceName` to `ckafka` .

Specifying the pattern string

This feature allows you to name batch purchased instances in a complex format with specified serial numbers. You can use one or more pattern strings in instance names as needed. The instance name with a specified pattern string is in the format of **{R:x}**, where `x` indicates the starting number in the generated instance names and can be positive integers only.


Specifying one pattern string

The following example assumes that you want to create 3 instances and name them with incrementing numbers starting from 3.

Operations on the purchase page

Operations through API

1. Purchase instances as instructed in [Creating Instance](#). Enter the instance name in the format of **prefix+specified pattern string {R:x}**. In this case, enter `ckafka{R:3}` as the instance name:



Instance Name ⓘ

Support naming with consecutive numeric suffixes (for batch instance creation) or designated pattern string. You can enter 53 more characters.

2. Follow the prompts on the page and make the payment.

In the [ModifyInstanceAttributes](#) API, set the relevant fields:

Instance name: Set `InstanceName` to `ckafka{R:3}` .

Specifying multiple pattern strings

The following example assumes that you want to create 3 instances and name them with the prefixes of `ckafka` , `big` , and `test` and serial numbers after `ckafka` and `big` incrementing from 13 and 2 respectively (i.e., `ckafka13-big2-test`, `ckafka14-big3-test`, and `ckafka15-big4-test`).

Operations on the purchase page

Operations through API

1. Purchase 3 instances as instructed in [Creating Instance](#). Enter the instance name in the format of **"prefix+specified pattern string {R:x}-prefix+specified pattern string {R:x}-prefix"**. In this case, enter `ckafka{R:13}-big{R:2}-test` as the instance name:

Instance Name ⓘ

Support naming with consecutive numeric suffixes (for batch instance creation) or designated pattern string. You can enter 38 more cha

2. Follow the prompts on the page and make the payment.

In the [ModifyInstanceAttributes](#) API, set the relevant fields:

Instance name: Set `InstanceName` to `ckafka{R:13}-big{R:2}-test` .

Feature Verification

After you batch create instances through [automatically incrementing suffixed numbers](#) or [specifying the pattern string](#), you can verify the set instance names as follows:

Log in to the [CKafka console](#) and view the newly created instances. You can see that the batch purchased instances are named according to the rule you set:

ID/Name	Monitor...	Status	AZ	Instance T...	Configura...	Network T...	Instance B...	Tag	Operation
<input type="checkbox"/> ckafka- Renew ckafka14-big3-test		Creating ⓘ	Guangzhou Zone 7	Standard Edition- Small Version: 1.1.1 Disk Type: Premium Cloud Storage	Peak Bandwidth: 40 MB/s Disk Capacity: 300GB	VPC clue-test test	Pay as you go		Configure Ala Upgrade Terminate
<input type="checkbox"/> ckafka- Renew ckafka13-big2-test		Healthy	Guangzhou Zone 7	Standard Edition- Small Version: 1.1.1 Disk Type: Premium Cloud Storage	Peak Bandwidth: 40 MB/s Disk Capacity: 300GB	VPC clue-test test	Pay as you go		Configure Ala Upgrade Terminate

Viewing Instance

Last updated : 2024-01-09 14:47:31

Overview

This document describes how to view the configuration information and status of an instance in the CKafka console.

Directions

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and click the **ID/Name** of the target instance to enter its **Basic Info** page, where you can view various information of the instance such as the status, configuration information, access mode, message retention period, and automatic topic creation.

<input type="checkbox"/>	ID/Name	Monitor...	Status	AZ	Instance Type	Configuration	Network Type	Instance
<input type="checkbox"/>	test		Healthy	Guangzhou Zone 6	Standard Edition-Advanced Version: 1.1.1 Disk Type: Premium Cloud Storage	Topic Limit: 50 Partition Limit: 150 Peak Bandwidth: 150 MB/s Disk Capacity: 2500GB	VPC	Pay as you go

Note:

The private IP and port (such as `10.6.206.110:9092`) in **Configuration Information** represent the communication address used to get the backend service. There may be multiple ports in a real access address. If access control is configured on your server, open ports 9092–9192 to the internet on it. As brokers may be automatically scaled out, the number of ports that need to be opened will increase after the scale-out, and you should reserve a sufficient number of ports.

Once you enable the automatic topic creation feature for the server, when you use or access metadata of a topic that does not exist, it will be automatically created with the configured number of replicas and partitions.

The total number of topics that can be automatically created varies by instance specification. For more information, see [Billing Overview](#).

Health status description

CKafka runs an inspection program for each instance, which checks the instance's number of connections, disk utilization, production peak bandwidth, and consumption peak bandwidth. Different health status are displayed when these metrics exceed certain thresholds as detailed below:

Metric	Threshold (N)	Status Description
Number of connections (up to 5,000 by default)	$N \leq 80\%$	Healthy
	$80\% < N \leq 95\%$	Alarmed
	$N > 95\%$	Abnormal
Disk utilization	$N \leq 80\%$	Healthy
	$80\% < N \leq 95\%$	Alarmed
	$N > 95\%$	Abnormal
Production peak bandwidth (excluding the bandwidth of replicas)	$N \leq 80\%$	Healthy
	$80\% < N \leq 100\%$	Alarmed
	$N > 100\%$	Abnormal
Consumption peak bandwidth	$N \leq 80\%$	Healthy
	$80\% < N \leq 100\%$	Alarmed
	$N > 100\%$	Abnormal

Note:

The maximum number of connections is 5,000 by default, and you should determine the threshold based on the percentage of the maximum value. After the number of instance connections exceeds this maximum value, the client cannot create more connections. If this maximum value is unreasonable in your actual business, you can [submit a ticket](#) to increase it.

Upgrading Instance

Last updated : 2024-01-09 14:47:32

Overview

If the current instance specifications cannot meet your business needs, you can upgrade them in the console.

Note:

You can upgrade the peak bandwidth, disk capacity, partition count, and other instance specifications in the console.

The specification upgrade is a smooth upgrade, so it won't interrupt your business.

New ports may be added during the upgrade. You can view them by clicking **View All IPs and Ports** in the **Access Mode** section on the instance details page later. **When configuring the security group, you need to open all these ports.**

Prerequisites

You need to check the following before the upgrade:

1. Check whether the instance has unavailable public network routes, supporting networks, and VPC networks. For more information, see [Adding Routing Policy](#).
2. Check whether the instance has unsynced replicas. For more information, see [Viewing Topic](#).
3. Check whether the instance has unfinished tasks (data migration), abnormally created topics, or abnormally deleted topic data.

Note:

If there are unfinished tasks, we recommend that you wait for all the tasks to be completed before the upgrade. If task execution is abnormal, [submit a ticket](#) for assistance.

Directions

Scenario 1: Upgrading Standard Edition instance specification

Scenario 2: Upgrading Standard Edition to Pro Edition

Scenario 3: Upgrading Pro Edition instance specification

1. Log in to the [CKafka console](#).
2. On the instance list page, select **More > Upgrade** in the **Operation** column to enter the upgrade page.
3. Select the target specification for the instance.

Upgrade Configuration ✕

Original Configuration

Model	Peak bandwidth...	Disk capacity (GB)	Topic Limit	Partition Limit
Advanced	150	2500	50	150

Target Specification Type:

Target Model:

Peak Bandwidth: 150MB/s

Disk capacity: 2500 4000 2500 GB

Topic Limit: 50

Partition Limit: 150

New ports may be added during the upgrade. You can view them by clicking "View All IPs and Ports" in the "Access Mode" section on the instance details page later.

Original Price: **0.23USD/hr**

New Price: **0.23USD/hr**

4. Click **Submit** and complete the upgrade as prompted

1. Log in to the [CKafka console](#).

2. On the instance list page, select **More** > **Upgrade** in the **Operation** column to enter the upgrade page.

3. Select **Pro Edition** as the target specification type.

Upgrade Configuration ✕

Original Configuration

Model	Peak bandwidth...	Disk capacity (...)	Topic Limit	Partition Limit
Advanced	150	2500	50	150

Target Specification Type: Standard Edition Pro Edition

Target Instance Specification: 40-1,200 MB/sec Over 1,200 MB/sec

Peak Bandwidth: - 160 + MB/s
20 1200

Disk capacity: - 2500 + GB
1000 500000

Topic Limit: Up to 600 topic(s) available.

Partition Limit: - 1200 +

Rebalance Time: Rebalancing will not occur during the configuration change

Configuration Change Mode: The configuration change process takes about one minute. You don't need to select a configuration change mode as the process doesn't involve data migration between nodes.

The upgrade is riskless as data migration is not required.

New ports may be added during the upgrade. You can view them by clicking "View All IPs and Ports" in the "Access Mode" section on the instance details page later.

Original Price: [REDACTED]

New Price: [REDACTED]

Submit
Close

Product Specs: Select a model based on the peak bandwidth and disk capacity.

Instance Price: You need to pay for the upgrade based on the number of days.

Rebalancing Time:

If the backend detects that data migration is required during the upgrade, you can choose to upgrade right away or at the specified time (upgrade at night is recommended to reduce the impact on your business). The estimated upgrade duration is calculated by backend APIs.

If the backend detects that data migration is not required during the upgrade, you will see the message "Rebalancing will not occur during the upgrade".

Upgrade Mode: If the backend detects that instance migration is required during the upgrade, you can select an upgrade mode as needed. Otherwise, you do not need to select an upgrade mode.

Stable: CKafka will limit the speed of data migration during the upgrade to ensure that instances have optimal bandwidth performance. This option is recommended if you do not want to affect business operations.

High-speed: CKafka will not limit the speed of data migration during the upgrade, thus affecting the bandwidth available for data production and consumption of instances. This option is recommended for off-peak hours or when service suspension is allowed.

4. Click **Submit** to complete the instance upgrade. You can view the upgrade progress in the **Status** column.

5. If you choose to upgrade at the specified time, you can modify the time you specified in the **Status** column.

1. Log in to the [CKafka console](#).

2. On the instance list page, select **More > Upgrade** in the **Operation** column to enter the upgrade page.

3. Select the target specification for the instance.

Upgrade Configuration ✕

Original Configuration

Model	Peak bandwidth (MB...	Disk capacity (GB)	Partition Limit
40-1,200 MB/sec	40	500	800

Target Instance Specification

40-1,200 MB/sec

Over 1,200 MB/sec

Peak Bandwidth

− 80 + MB/s

20 1200

Disk capacity

− 500 + GB

500 500000

Topic Limit

Up to 450 topic(s) available.

Partition Limit

− 900 +

Rebalance Time

Rebalancing will not occur during the configuration change

Configuration Change Mode

The configuration change process takes about one minute. You don't need to select a configuration change mode as the process doesn't involve data migration between nodes.

The upgrade is riskless as data migration is not required.

New ports may be added during the upgrade. You can view them by clicking "View All IPs and Ports" in the "Access Mode" section on the instance details page later.

Original Price

New Price

Submit

Close

Product Specs: Select a model based on the peak bandwidth and disk capacity.

Instance Price: You need to pay for the upgrade based on the number of days.

Rebalancing Time:

If the backend detects that data migration is required during the upgrade, you can choose to upgrade right away or at the specified time (upgrade at night is recommended to reduce the impact on your business). The estimated upgrade duration is calculated by backend APIs.

If the backend detects that data migration is not required during the upgrade, you will see the message "Rebalancing will not occur during the upgrade".

Upgrade Mode: If the backend detects that instance migration is required during the upgrade, you can select an upgrade mode as needed. Otherwise, you do not need to select an upgrade mode.

Stable: CKafka will limit the speed of data migration during the upgrade to ensure that instances have optimal bandwidth performance. This option is recommended if you do not want to affect business operations.

High-speed: CKafka will not limit the speed of data migration during the upgrade, thus affecting the bandwidth available for data production and consumption of instances. This option is recommended for off-peak hours or when service suspension is allowed.

4. Click **Submit** to complete the instance upgrade. You can view the upgrade progress in the **Status** column.

5. If you choose to upgrade at the specified time, you can modify the time you specified in the **Status** column.

Possible Causes of Upgrade Failures

1. The disk resources in the current AZ cannot meet the requirements of this upgrade. We recommend that you contact the Tencent Cloud customer service to check whether there are sufficient resources.
2. If the high-speed mode is selected during the instance upgrade, and there are bandwidth-intensive production tasks in the cluster, the data migration delay will increase. You can view monitoring data to see whether there are excessive peaks of production and consumption traffic during the upgrade. For more information, see [Viewing Monitoring Data](#).
3. The upgrade process takes too long. The maximum message size configured for the broker to be migrated is 8 MB, while the size configured for the target broker is 1 MB. In this case, the target cluster cannot receive the oversized messages, causing the data migration failure and prolonging the upgrade duration. You can [submit a ticket](#) for assistance.
4. During the cluster upgrade or migration process, an exception may occur while updating the broker IP, causing the failure to pull the broker IP data of the new cluster. You can observe that no monitoring data is displayed for a period of time. You can [submit a ticket](#) for assistance. For more information, see [Viewing Monitoring Data](#).

Downgrading Instance Configuration

Last updated : 2024-01-09 14:47:31

Operation Scenarios

If there are too many redundant resources in your current instance specifications, you can consider degrade your **Pro Edition** instance specification via the console.

Note :

To maintain the cluster stability, each instance can only be downgraded 3 times a week. If this restriction needs to be lifted, you can [submit a ticket](#) for application.

You can use the console to directly downgrade your instance specifications, including peak bandwidth, disk capacity, and the number of partitions. The downgrade operation is a smooth change and will not interrupt your services.

Standard Edition instance specifications cannot be downgraded.

Prerequisites

You need to check the following before the downgrade:

1. Check whether the instance has unavailable public network routes, supporting networks, VPC networks and so on.

For more information, see [Adding Routing Policy](#).

2. Check whether the instance has unsynced replicas. For further details, see [Viewing Topic Details](#).

3. Check whether the instance has unfinished tasks (data migration), abnormally created topics, or abnormally deleted topic data.

4. Check the resource usage of the instance at least 7 days in the past.

Note :

If there are unfinished tasks, we recommend that you wait for all the tasks to be completed before the downgrade. If task execution is abnormal, [submit a ticket](#) for assistance.

Directions

1. Log in to the [CKafka console](#).

2. On the instance list page, select **More** > **Downgrade** in the **Operation** column to enter the upgrade page.

3. On the instance downgrade page, select the target specification for the downgrade.

Downgrade Configuration ✕

Original Configuration	<table style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="border-bottom: 1px solid #ccc;">Model</th> <th style="border-bottom: 1px solid #ccc;">Peak bandwidth (...)</th> <th style="border-bottom: 1px solid #ccc;">Disk Capacity (GB)</th> <th style="border-bottom: 1px solid #ccc;">Partition Limit</th> </tr> </thead> <tbody> <tr> <td style="padding: 2px 5px;">20-1,200 MB/s</td> <td style="padding: 2px 5px;">20</td> <td style="padding: 2px 5px;">200</td> <td style="padding: 2px 5px;">400</td> </tr> </tbody> </table>	Model	Peak bandwidth (...)	Disk Capacity (GB)	Partition Limit	20-1,200 MB/s	20	200	400		
Model	Peak bandwidth (...)	Disk Capacity (GB)	Partition Limit								
20-1,200 MB/s	20	200	400								
Target Instance Specification	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="border: 1px solid #ccc; padding: 2px 5px;">20-1,200 MB/s</td> <td style="padding: 2px 5px;">1,600 MB/s or greater</td> </tr> </table>	20-1,200 MB/s	1,600 MB/s or greater								
20-1,200 MB/s	1,600 MB/s or greater										
Usage (Last 7 Days)	<table style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="border-bottom: 1px solid #ccc;">Average Prod...</th> <th style="border-bottom: 1px solid #ccc;">Average Cons...</th> <th style="border-bottom: 1px solid #ccc;">Average Disk ...</th> <th style="border-bottom: 1px solid #ccc;">Topics</th> <th style="border-bottom: 1px solid #ccc;">Partitions</th> </tr> </thead> <tbody> <tr> <td style="padding: 2px 5px;">0.00/160</td> <td style="padding: 2px 5px;">0.00/160</td> <td style="padding: 2px 5px;">0.15/200</td> <td style="padding: 2px 5px;">1</td> <td style="padding: 2px 5px;">24</td> </tr> </tbody> </table>	Average Prod...	Average Cons...	Average Disk ...	Topics	Partitions	0.00/160	0.00/160	0.15/200	1	24
Average Prod...	Average Cons...	Average Disk ...	Topics	Partitions							
0.00/160	0.00/160	0.15/200	1	24							
Production Traffic Peak in Last 7 Days	0.00MB/s										
Consumption Traffic Peak in Last 7 Days	0.00MB/s										
Peak Storage (Last 7 Days)	0.75GB										
	If the peak bandwidth or storage is greater than the target specification in the last 7 days, you cannot perform specification downgrade.										
Peak Bandwidth	The lower limit of peak bandwidth has been reached.										
Disk Capacity	The lower limit of disk capacity has been reached.										
Topics	Up to 200 topics are available.										
Partitions	<p>The lower limit of partitions has been reached.</p> <p>You've already used 24 partitions. To ensure the normal operation of your business, you cannot reduce the current partition count.</p>										
Configuration Change Mode	<p>The configuration change process takes about one minute. You don't need to select a configuration change mode as the process doesn't involve data migration between nodes.</p> <p style="color: #e67e22;">The configuration change is risk-free as data migration is not required.</p>										
	<div style="display: flex; justify-content: center; gap: 10px;"> <div style="border: 1px solid #00aaff; padding: 5px 15px; color: white; cursor: pointer;">Submit</div> <div style="border: 1px solid #ccc; padding: 5px 15px; color: #00aaff; cursor: pointer;">Close</div> </div>										

Note :

If the peak bandwidth in the last 7 days or the peak storage in the last 7 days is greater than the target specification, downgrade is not advised.

4. Click **Submit** and complete the downgrade as prompted.

Terminating/Returning Instances

Last updated : 2024-01-09 14:47:32

Overview

If you no longer need a CKafka instance, you can terminate and release it.

The lifecycle of a CKafka instance refers to all statuses from the start of the instance to its release. Properly managing an instance in different statuses can ensure that applications running on it can provide services economically and efficiently. The statuses of an instance are as follows:

Status	Status Attribute	Description
Creating	Intermediate status	The instance has been created but is not running yet.
Running	Stable status	The instance is running normally, and the disk space, traffic, and the number of connections are all within the normal range.
Deleting	Intermediate status	The instance is being deleted through the console or API.
Isolated	Intermediate status	The instance is isolated through the console or API, or it is in a seven-day isolation status due to overdue payments.
Failed to create	Intermediate status	You have purchased an instance through the console or API but are not assigned the instance. In this case, submit a ticket .
Failed to delete	Stable status	CKafka fails to release resources after the instance is manually deleted or is not renewed 14 days after expiration (including the 14th day).

Directions

Manual termination

You can manually terminate a monthly-subscribed instance before it expires by following the steps below:

1. Log in to the [CKafka console](#).
2. On the instance list page, click **More > Terminate/Return** in the **Operation** column.
3. In the pop-up window, click **Submit** to terminate the instance.

Terminate Instance ×

You have selected **1 instance**, [View Details](#)^

No.	ID	Instance Name	Operation
1	ckafka-7!...za		Can be terminated

Are you sure you want to terminate selected CKafka instance?

i • Once terminated, all data are cleared and cannot be recovered. Please back up your data in advance.

• After the resource is terminated, the first resource will be unconditionally refunded to your Tencent Cloud account in 5 days. The normal refund amount will be returned to your Tencent Cloud account by the proportion of the cash and voucher amount paid for the purchase.

• If you purchase at a discount price or with a voucher, then the discount and voucher **cannot be refunded**.

Submit Close

Note:

After a monthly subscribed instance is terminated, it will be in “Isolated” status for seven days in the CKafka console. Isolated instances cannot produce or consume data. Data and configurations saved in CKafka will not be terminated, but **expired messages will still be automatically deleted according to the open-source Kafka mechanism**. For instances in the seven-day isolation status, you can go to the CKafka console to renew them by clicking **Renew** in the **Operation** column on the instance list page. Successfully renewed instances can go back to the running status and be normally used.

For more information, see [Refunds](#).

Automatic termination upon expiration/overdue payment

Monthly-subscribed instances can be retained in the CKafka console for up to seven natural days after they expire or have overdue payments. You can continue to use them if you renew them within 7 days after expiration. For details, see [Payment Overdue](#).

If the CKafka instance is not renewed within 7 (inclusive) days after expiration, its resources will be released within 24 hours after expiration, and data will be cleared and cannot be restored.

Deleting instances

A monthly-subscribed instance will be retained for seven days in the CKafka console after manual termination or automatic termination upon expiration/overdue payment, during which it will be in “Isolated” status. After 7 (inclusive) days, it will be completely released, and you can delete it completely.

Note:

After deletion, all data will be cleared and cannot be recovered. Please back up the data in advance.


1. Log in to the [CKafka console](#).
2. On the instance list page, click **More > Terminate/Return** in the **Operation** column.
3. In the pop-up window, click **Submit** to delete the instance.

Delete Instance ×

You have selected **1 instance**, [View Details](#)^

No.	ID	Instance Name	Operation
1	ckafka-2v4rknej	未命名	Deletable

Are you sure you want to delete selected CKafka instance?

 • Once deleted, all data will be cleared and cannot be recovered. Please back up the data in advance.

Submit Close

Change from Pay-as-You-Go to Monthly Subscription

Last updated : 2024-01-09 14:47:32

Overview

For your convenience, CKafka allows you to change the billing mode from pay-as-you-go to monthly subscription, so that you can convert temporarily used pay-as-you-go instances into monthly subscribed instances for long-term and stable use. You can perform the change operation in the console.

Change Rules

You can follow the rules below to use the billing mode change feature in the CKafka console:

When a pay-as-you-go instance is changed to monthly subscribed instance, a renewal order will be generated. You must pay the order for the change to take effect. If you don't make the payment or your payment is unsuccessful, you can check and handle the order on the [Billing Center > Order Management](#) page.

For instances with billing mode changed from **Pay-as-You-Go** to **Monthly Subscription**, the five-day free returns are not supported.

Upon the successful completion of the payment and billing mode change, the instance will be billed on a monthly subscription basis, with the start time being the moment the change takes effect.

You can't perform repeated billing mode change on the instance before making the payment.

You won't be able to pay an unpaid order if the order amount has changed because of your instance configuration adjustment (such as adjusting configurations/bandwidth/disk and reinstalling system). You need to cancel the unpaid order in [Billing Center > Order Management](#) before performing a new change operation.

The change from pay-as-you-go to monthly subscription involves a change in the billing mode for instance and disk, **but not for public network bandwidth**.

The change from pay-as-you-go instances to monthly subscribed instances is irreversible.

The change from pay-as-you-go instances to monthly subscribed instances can't be performed in batches.

Directions

1. Log in to the [CKafka console](#).
2. On the instance list page, select the target pay-as-you-go instance, and click **More > Change from Pay-as-You-Go to Monthly Subscription** in the **Operation** column.

3. In the **Change from Pay-as-You-Go to Monthly Subscription** pop-up window, set the renewal duration and whether to automatically renew as needed.
4. Click **OK** to complete the change operation.

Upgrading Instance Version

Last updated : 2024-01-09 14:47:32

CKafka **Pro Edition** supports upgrading the version of open-source instances to the latest version for more features. It also supports upgrading the kernel minor version to use new features, improve performance, and fix bugs.

Scenarios

Automatic upgrade (**kernel minor version**)

Scenario 1: When a severe bug or security vulnerability occurs in CKafka, the system will perform kernel minor version upgrade during the maintenance time and send upgrade notifications through the Message Center and SMS.

Scenario 2: When a CKafka cluster is migrated due to the cluster configuration upgrade, storage capacity expansion/reduction, or CKafka version upgrade, the system will automatically upgrade the cluster's kernel to the latest minor version.

Manual upgrade (**kernel minor version + open-source Kafka version**)

You can also manually upgrade the kernel minor version and open-source Kafka version in the console.

Note

This feature is only supported for Pro Edition.

Cluster switch will be required after the version upgrade is completed, which will cause your instance to disconnect for seconds. We recommend that you conduct the switch during the instance maintenance time. Make sure that your business has a reconnection mechanism.

Instances can be upgraded from an earlier version to a later one but cannot be downgraded.

The instance's minor version is automatically detected and upgraded by the system, and you cannot select the target version.

If an instance is upgraded to a compatible version, no billing changes will be caused.

Directions

1. Log in to the [CKafka console](#).
2. On the left sidebar, select **Instance List** and click the ID/name of the target instance to enter the instance details page.
3. In the **Basic Info** module, click **Upgrade Version** on the right of the instance version, and set the target version and upgrade time.

Upgrade kernel minor version:

Current Version: The current kernel version.

Target Version: The target version is the latest broker version by default. The system will automatically detect the minor version. If the **Upgrade kernel minor version** option is grayed out, it indicates that the instance is already on the latest minor version. For differences between kernel minor versions, see [Broker Release Notes](#).

Execution Time: You can select **Immediate execution** or **Custom time** (any time within the next 24 hours). We recommend that you select a time during off-peak hours.

Upgrade Kafka version:

Current Version: The current Kafka version.

Target Version: The target Kafka version. For differences between Kafka versions, see [DOWNLOAD](#).

Execution Time: You can select **Immediate execution** or **Custom time** (any time within the next 24 hours). We recommend that you select a time during off-peak hours.

Note:

Currently, you cannot upgrade from Kafka 1.1.1 to v2.4.1.

4. Click **OK** to submit the upgrade task.
5. Select the **Event Center** tab at the top of the page, and you can see a record of instance version upgrade.
6. Click **View Details** in the **Operation** column of the record to view the detailed upgrade task progress.
7. Return to the instance list page. You can see that the status of the instance has changed to **Changing configurations...**, and you can also view the upgrade progress.

Adding Routing Policy

Last updated : 2024-01-09 14:47:32

Overview

This document describes how to configure a routing rule in the CKafka console to enhance network access control in public/private network transfers. For more information on public network access, see [Configuring ACL Policy](#).

Route Type	VPC	Public Domain Name Access
Access mode	PLAINTEXT SASL_PLAINTEXT SASL_SSL (only supported by Pro Edition instances) SASL_SCRAM (only supported by instances on v2.4.1; for existing instances, submit a ticket for application)	SASL_PLAINTEXT SASL_SSL (only supported by Pro Edition instances)

Directions

Note:

Up to five routes can be created for an instance and there can be only one public network route.

VPC

Public domain name access

Operation scenario: When purchasing an instance, if you select VPC and choose a corresponding VPC environment (such as VPC A), then CKafka services (such as data production and consumption) can only be accessed from VPC A. If you subsequently find that you need to access the CKafka services in VPC A from other VPCs (such as VPC B), you can select an appropriate routing policy for VPC by configuring the access mode.

Suggestion: To ensure security, this access mode supports user management and ACL policy configuration to manage user access permission. Configure as appropriate.

Directions:

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and then the **ID/Name** of your instance to enter the instance details page.
3. On the page you enter, click **Add a routing policy** in the "Access Mode" module.
4. In the pop-up window, select **VPC Network** as the route type and select the access mode and network.

Add a routing policy ×

Route Type

Access Mode

Network ↻

To change network settings, go to the console to [create VPC](#) or [create Subnet](#)

IP

If no IP is specified, one will be assigned automatically.

Note:

If you select VPC access, you can specify the IP to keep it unchanged when changing the access mode.

5. Click **Submit** to add the VPC network.

6. Click **View All IPs and Ports** in the **Operation** column to view the list of IPs and ports that should be opened.

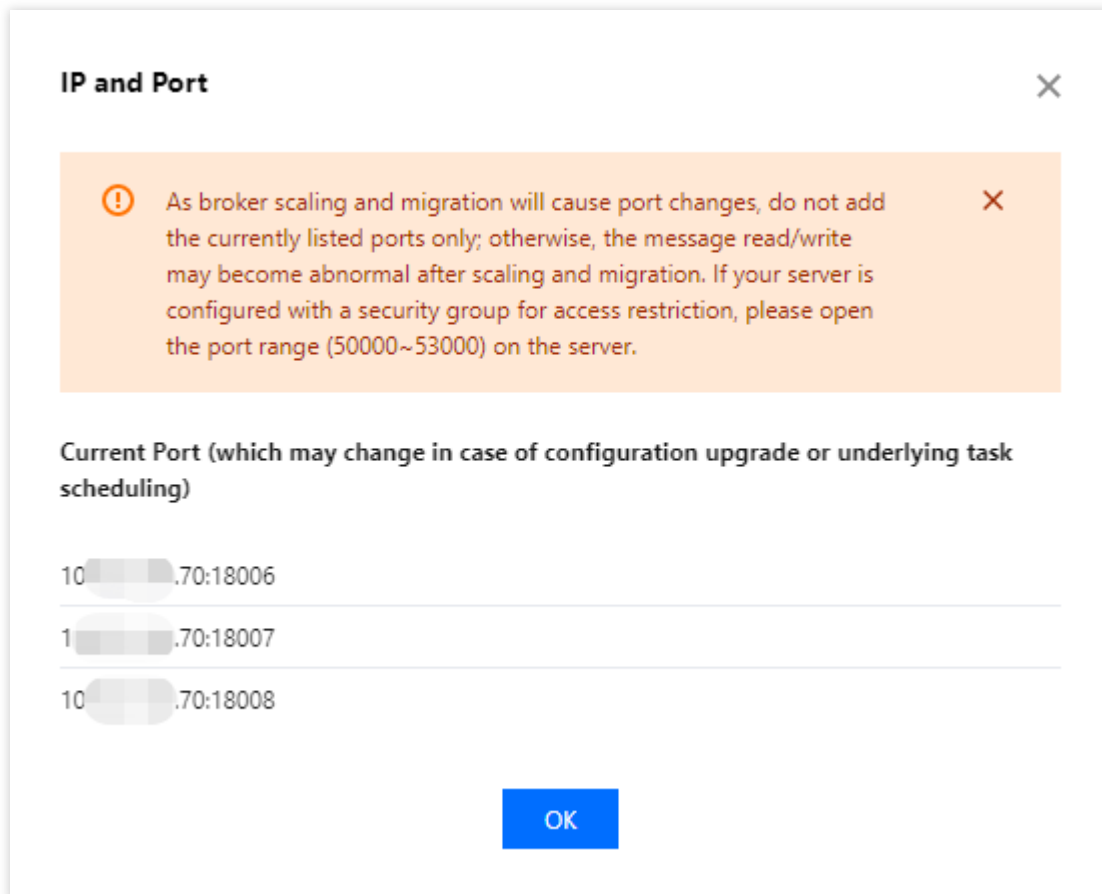
Note:

As broker scaling and migration will cause port changes, do not add only the current listed ports; otherwise, message read/write may become abnormal after scaling and migration. If your server is configured with security groups for access restrictions, you should open the following port ranges on the server:

Port range that should be opened for the VPC route: 9092–60000

Port range that should be opened for the public network route: 50000–53000

Port range that should be opened for the supportive route: 6000–12000



Operation scenario: If your consumer or producer is located in a self-built data center or another cloud, you can produce and consume data in CKafka through public network access.

Suggestion: To ensure security, Kafka offers various security authentication mechanisms, which mainly fall into two categories: SSL and SASL2. SASL/PLAIN is a more commonly used authentication method based on account and password. CKafka supports SASL_PLAINTEXT and SASL_SSL authentication. We recommend that you configure the authentication method as appropriate when selecting public domain name access.

Directions:

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and then the **ID/Name** of your instance to enter the instance details page.
3. On the page you enter, click **Add a routing policy** in the "Access Mode" module.
4. In the pop-up window, select **Public domain name access** as the route type and select the access mode and network.

Add a routing policy ×

i A single broker on a standard edition instance supports up to 1 Mbps public network bandwidth.

Route Type

Access Mode

This access mode provides user management and ACL policy configuration to manage user access permission.

5. Click **Submit** to add the public network routing policy.

6. Click **View All IPs and Ports** in the **Operation** column to view the list of IPs and ports that should be opened.

Note:

As broker scaling and migration will cause port changes, do not add only the current listed ports; otherwise, message read/write may become abnormal after scaling and migration. If your server is configured with security groups for access restrictions, you should open the following port ranges on the server:

Port range that should be opened for the VPC route: 9092–60000

Port range that should be opened for the public network route: 50000–53000

Port range that should be opened for the supportive route: 6000–12000

CKafka provides 3 Mbps public network bandwidth free of charge by default, which can be increased for Pro Edition instances. For detailed directions, see [Public Network Bandwidth Management](#).

IP and Port ✕

! As broker scaling and migration will cause port changes, do not add the currently listed ports only; otherwise, the message read/write may become abnormal after scaling and migration. If your server is configured with a security group for access restriction, please open the port range (50000~53000) on the server. ✕

Current Port (which may change in case of configuration upgrade or underlying task scheduling)

10

1

10

OK

Public Network Bandwidth Management

Last updated : 2024-01-09 14:47:32

Overview

CKafka transfers data over the private network by default. To access over the public network, you need to activate a public network route separately. For detailed directions, see [Adding Routing Policy](#). Currently, public network bandwidth of 3 Mbps is provided free of charge by default.

CKafka Pro Edition allows you to increase the bandwidth of the public network. You can pay for more public network bandwidth if you require it. For the specific prices, see [Billing Overview](#).

This document describes how to upgrade, adjust, and delete the public network bandwidth in the CKafka console.

Directions

CKafka adjusted the purchase entry and mode of public network bandwidth on January 7, 2022. After then, new users can purchase a public network bandwidth during instance purchase, and existing instances continue to be billed by the hourly bandwidth, which can be upgraded.

You can click the following tags to view the purchase methods and relevant operations for old and new instances.

[Public network bandwidth management \(new\)](#)

[Public network bandwidth management \(old\)](#)

Purchasing the public network bandwidth

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and click **Create** to enter the purchase page.
3. On the purchase page, set the instance information and select the public network bandwidth based on your actual needs.

Billing Mode: **Monthly bandwidth** and **hourly bandwidth** are supported. Currently, you cannot switch the billing mode of the public network bandwidth. For billing details, see [Billing Overview](#).

Bandwidth Size: Select the desired public network bandwidth size. CKafka provides 3 Mbps public network bandwidth free of charge by default. You can purchase more as needed.

4. Click **Buy Now**.

Upgrading the public network bandwidth

Note:

Currently, the public network bandwidth cannot be downgraded in the console. Therefore, make your purchase with caution. If you need to downgrade it, [submit a ticket](#) for assistance.

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and then the **ID/Name** of your instance to enter the **Basic Info** page.
3. On the instance's basic information page, click **Upgrade** next to the public network bandwidth in the **Configuration Information** module.
4. Modify the public network bandwidth in the pop-up window and click **Submit**.

Unsubscribing from the public network bandwidth

Note:

Currently, only the **bill-by-hour** public network bandwidth can be unsubscribed from, while the **bill-by-month** one cannot, which can be terminated together with the instance though.

You can unsubscribe from the public network bandwidth only when there is no public network route in your instance; therefore, you need to delete public network routes first before proceeding.

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and then the **ID/Name** of your instance to enter the **Basic Info** page.
3. On the instance's basic information page, click **Unsubscribe** next to the public network bandwidth in the **Configuration Information** module.
4. Click **Submit** in the pop-up window. After unsubscription, the bandwidth will stop being billed.

Deleting a public network route

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and then the **ID/Name** of your instance to enter the **Basic Info** page.
3. In the **Access Mode** module, click **Delete** in the **Operation** column for the public network route you want to delete and select the execution time in the pop-up window.

Right now: The public network route will be deleted immediately.

Custom time: You can select any time within the next 24 hours for scheduled deletion. The public network route will enter the state of waiting for deletion, and the scheduled deletion time can be modified in the **Operation** column.

Access Mode ?			
Access Type	Access Mode	Network	Operation
VPC Network	PLAINTEXT	vpc-fs6qq7yn subnet-rhayp718 10.0.11.14:9092	Delete View All IP
Public domain name access	SASL_PLAINTEXT	ckafka- 193 001	Delete View All IP

Note:

Deleting the public network route will result in rebalancing. Do so with caution.

Upgrading the public network bandwidth

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and then the **ID/Name** of your instance to enter the **Basic Info** page.
3. On the page you enter, click **Add a routing policy** in the **Access Mode** module.
4. In the pop-up window, select **Public domain name access** as the route type, select the public network bandwidth you need, and click **Submit**.

Adjusting the public network bandwidth

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and then the **ID/Name** of your instance to enter the **Basic Info** page.
3. On the page you enter, click the edit icon next to the public network bandwidth in the **Configuration Information** module.
4. Modify the public network bandwidth in the pop-up window and click **Submit**.

Note:

Public network bandwidth is billed by hour. If it is changed multiple times within one hour, fees will be charged based on the highest bandwidth value.

Deleting a public network route

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and then the **ID/Name** of your instance to enter the **Basic Info** page.
3. In the **Access Mode** module, click **Delete** in the **Operation** column for the public network route you want to delete and select the execution time in the pop-up window.

Right now: The public network route will be deleted immediately.

Custom time: You can select any time within the next 24 hours for scheduled deletion. The public network route will enter the status of waiting for deletion, and the scheduled deletion time can be modified in the **Operation** column.

Access Mode ?			
Access Type	Access Mode	Network	Operation
VPC Network	PLAINTEXT	vpc-fs6qq7yn subnet-rhayp718 10.0.11.14:9092	Delete View All IPs and
Public domain name access	SASL_PLAINTEXT	ckafka- 193 001	Delete View All IPs and

Note:

Deleting the public network route will result in rebalancing. Do so with caution.

Connecting to Prometheus

Last updated : 2024-01-09 14:47:31

Overview


By default, a CKafka Pro Edition instance allows connection with external monitoring services for instance monitoring through the provided access point, including monitoring metrics of open-source Kafka such as **unsynced replicas and incoming message rate**.

CKafka Pro Edition instances currently provide Prometheus to scrape the metric information of broker nodes, including basic monitoring metrics such as CPU, memory usage, and system load, as well as the metrics exposed by the broker's JMX.

Directions

1. Log in to the [CKafka console](#).
2. Select **Instance List** on the left sidebar and click the **ID** of the target instance to enter its basic information page.
3. Click **Get Monitoring Target** in the top-right corner of the **Prometheus Monitoring** module and select the VPC and subnet.







Get Monitoring Target

Network 

To change network settings, go to the console to [create VPC](#) or [create Subnet](#)


4. Click **Submit** to get the set of monitoring targets.

Prometheus Monitoring Get Monitoring Target

jmx exporter	node exporter	Operation
10.0.11.1:60001 	10.0.11.2:60002 	
10.0.11.1:60003 	10.0.11.2:60004 	Delete
10.0.11.2:60005 	10.0.11.2:60006 	


5. Download [Prometheus](#) and configure the monitoring scrape address.

5.1 Enter the directory of the Prometheus package and run the following command to decompress it.



```
tar -vxf prometheus-2.30.3.linux-amd64.tar.gz
```

5.2 Modify the `prometheus.yml` configuration file by adding the `jmx_exporter` and `node_exporter` scrape tasks.




```
scrape_configs:
  # The job name is added as a label `job=<job_name>` to any timeseries scraped from
  - job_name: "prometheus"
    # metrics_path defaults to '/metrics'
    # scheme defaults to 'http'.
    static_configs:
      - targets: ["localhost:9090"]

  - job_name: "broker-jmx-exporter"
    scrape_interval: 5s
    metrics_path: '/metrics'
```

```
static_configs:
  - targets: ['10.x.x.0:60001', '10.x.x.0:60003', '10.x.x.0:60005']
    labels:
      application: 'broker-jmx'
- job_name: "broker-node-exporter"
  scrape_interval: 10s
  metrics_path: '/metrics'
  static_configs:
    - targets: ['10.x.x.0:60002', '10.x.x.0:60004', '10.x.x.0:60006']
      labels:
        application: 'broker-node'
```

Here, `broker-jmx-exporter` is the tag configured for the `jmx` metric of the broker scraped by Prometheus, `Targets` contains the information of the mapped port, `broker-node-exporter` is the tag configured for the basic metrics of the node of the scraped broker, and `scrape_interval` is the frequency of scraping metric data.

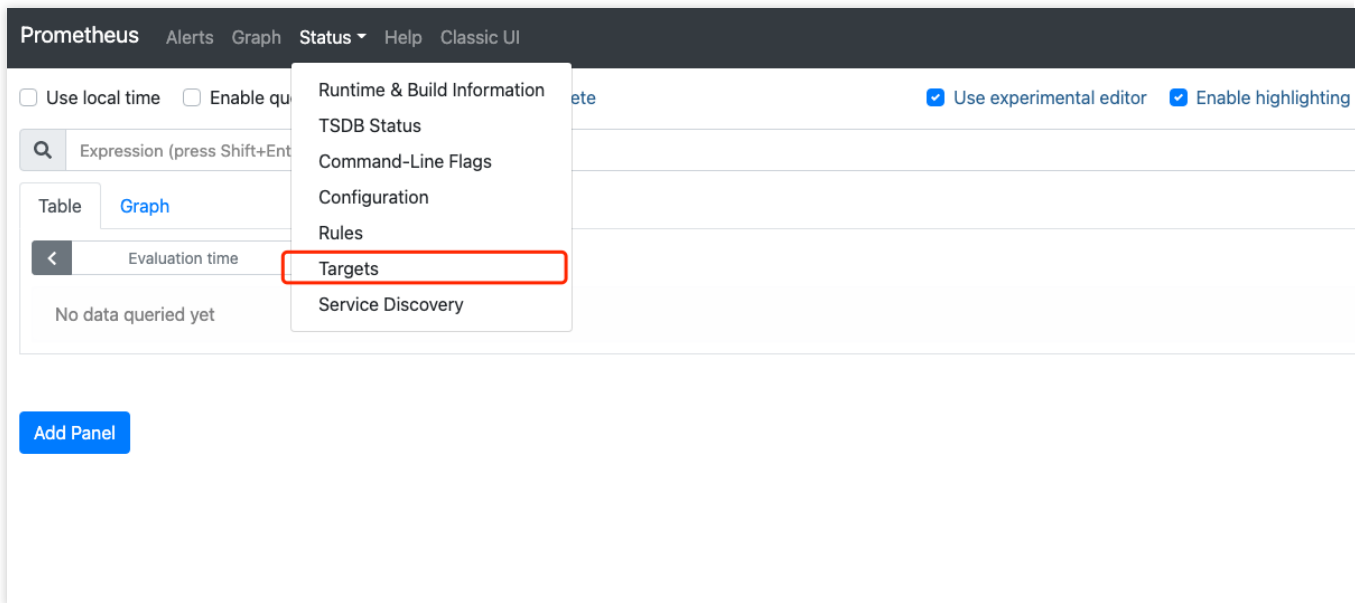
5.3 Start Prometheus.



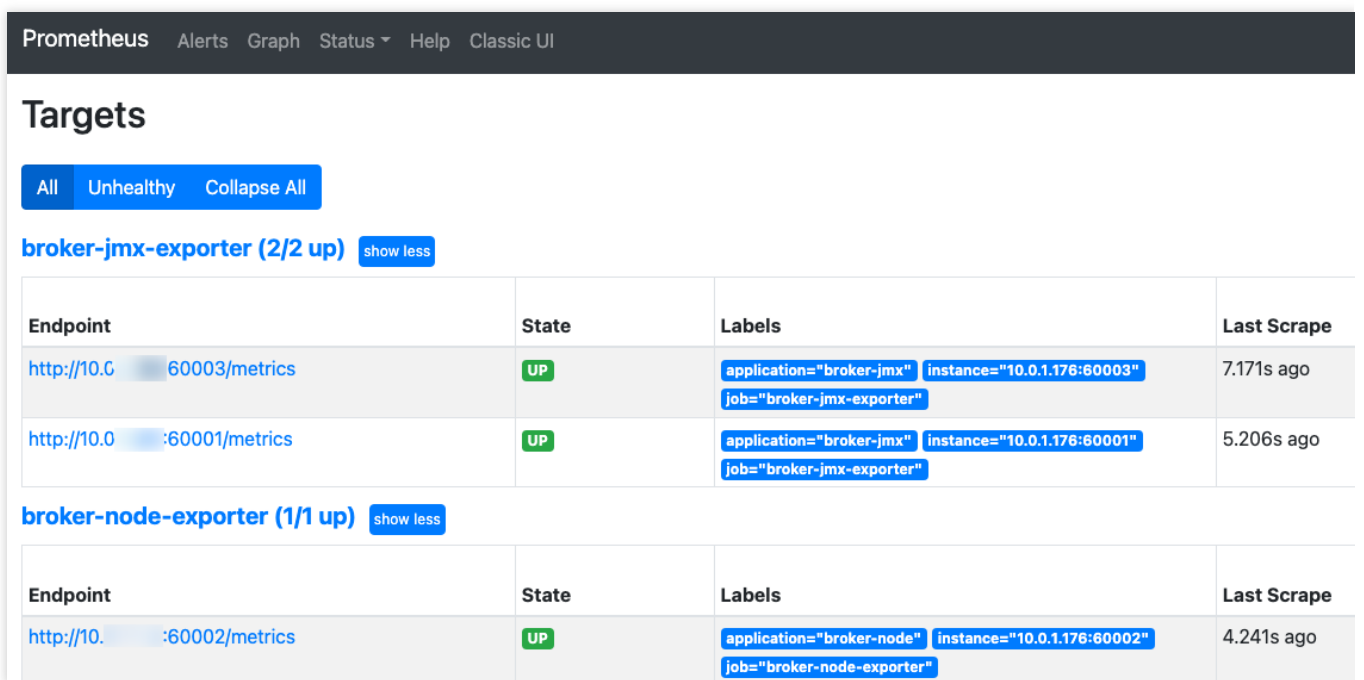
```
./prometheus --config.file=prometheus.yml --web.enable-lifecycle
```

5.4 Open the UI provided by Prometheus to check whether the connected targets are normal by entering

`http://localhost:9090` in the browser for example.



5.5 Confirm that all targets are in `UP` status.



If the targets are in `DOWN` status, check whether the network is accessible or check the `Error` option at the end of the status bar for the cause.

6. Query the monitoring metric data.

Click **Graph**, enter a metric name such as `node_memory_MemAvailable_bytes`, and click **Execute** to view the monitoring data.

Prometheus Alerts Graph Status Help Classic UI

Use local time Enable query history Enable autocomplete Use experimental editor Enable highlight

Q node_memory_MemAvailable_bytes

Table **Graph** Load time: 48ms Resol

< Evaluation time >

node_memory_MemAvailable_bytes(application="broker-node", instance="00", job="broker-node-exporter") 1171

Add Panel

AZ Migration

Last updated : 2024-01-09 14:47:32

Overview

You can migrate a **CKafka Pro Edition** instance to another AZ in the same region. All its attributes, configurations, and connection address will stay unchanged after the migration. The time it takes to migrate the instance is subject to the instance's data volume.

For example, you can migrate to a new AZ in the following scenarios:

If you want to modify an instance's type, but the current AZ doesn't support the new instance type, you can migrate the instance to an AZ supporting the new type.

If the current AZ has no remaining resources for scaling, you can also migrate the instance to another AZ in the same region with sufficient resources to meet your business needs.

Prerequisites

The instance is running.

The region where the instance is located has multiple AZs to support cross-AZ migration.

Billing Description

This feature is free of charge. There is no charge even for migrating an instance from a single AZ to multiple AZs.

Description

If the original instance is deployed in a single AZ, you can switch the AZ or upgrade to [multi-AZ deployment](#).

If the original instance is deployed in multiple AZs, you can switch AZs but cannot downgrade to single-AZ deployment.


Migration Type

Migration Type	Applicable Scenario

Migration from one AZ to another AZ	The AZ where the instance is located is under full load or other conditions that affect the instance performance.
Migration from one AZ to multiple AZs	You can improve the disaster recovery capability of the instance and implement cross-data center disaster recovery. Source and replica instances are located in different AZs. Multi-AZ instances can withstand higher levels of disasters than single-AZ instances. For example, the latter can tolerate server- and rack-level failures, while the former can tolerate data center-level failures.

Directions

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and click the **ID/Name** of the target instance to enter the basic information page.
3. In the **Basic Info** module, click **Edit** on the right of the current AZ and select the AZ you want to switch to.



The image shows a 'Change AZ' dialog box with a close button (X) in the top right corner. It contains two sections: 'Deployment Mode' and 'Single-AZ'. Under 'Deployment Mode', there are two buttons: 'Single-AZ' (highlighted with a blue border) and 'Multi-AZ'. Under 'Single-AZ', there are five radio button options: 'Shanghai Zone 3', 'Shanghai Zone 4' (selected with a blue dot), 'Guangzhou Zone 1', 'Guangzhou Zone 3', and 'Guangzhou Zone 4'. At the bottom, there are two buttons: 'OK' (highlighted in blue) and 'Cancel'.

4. Click **OK**. It will take 5–10 minutes to complete the configuration adjustment. You can view the progress in the **Status** column in the instance list.

Setting Maintenance Time

Last updated : 2024-01-09 14:47:32

Overview

The maintenance time is critical for CKafka. The backend system performs maintenance operations on your CKafka instance from time to time to ensure its stability. To minimize the potential impact on your business, we recommend that you set an acceptable maintenance period for your business instance, usually during off-peak hours.


Note:

The default system maintenance time starts at 11:00 PM and lasts one hour. You can modify the time for CKafka Pro Edition instances but not for Standard Edition instances.

Before maintenance is carried out for CKafka, notifications will be sent to the recipients configured in your Tencent Cloud account by SMS and email.


Setting the maintenance time

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and then the **ID/Name** of your instance to enter the **Basic Info** page.
3. In the **Basic Info** module, click **Edit** next to the system maintenance time.
4. In the pop-up window, select **Maintenance Window** and **Maintenance Time**.



Modify Maintenance Window and Time

Maintenance Window Mon Tue Wed Thu Fri
 Sat Sun

Maintenance Time Start Time 23:30 

Duration 1 hr

OK Cancel

5. Click **OK**.

Setting Message Size

Last updated : 2024-01-09 14:47:32

Operation Scenarios

CKafka supports instance-level configuration of the default message size, which serves as the default initial value when a new topic is created. It can later be modified for each topic.

Configuring Default Message Size

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left navigation bar, and click **ID/Name** of the target instance to enter the instance details page.
3. In the instance details page, click **Modify Configuration** in the message configuration module.
4. After setting the message retention period and max message size, click Submit.

Message Configuration ✕

Message Retention Period day(s) ▼

The message retention period range is 1 min to 90 day(s)

Max Message Size MB ▼

Range: 1 KB to 12 MB
Default value for a new topic, which can be modified for the topic later

Setting Message Size for a Topic

1. Log in to the [CKafka console](#).
2. Choose **Instance List** from the left sidebar, and click **ID/Name** of the target instance to enter the instance details page.
3. On the instance details page, click on the **Topic Management** tab and then click **Edit** in the operation bar.
4. On the Edit Topic page, click Show Advanced Configuration, and configure **max.message.bytes** to modify the maximum message size for the topic.

max.message.bytes

MB



When the client sends data, the data sent to the same partition will be aggregated and sent uniformly. The server will compare the size and range (1 KB to 12 MB) of each batch of messages.

Topic Management

Instance Topic

Creating Topic

Last updated : 2024-01-09 14:47:31

Overview

Topic is the name of a category where messages can be stored and published. CKafka uses the concept of topic externally, that is, a producer writes messages to a topic, while a consumer read messages from it. To implement horizontal scaling, a topic actually consists of multiple [partitions](#). When a performance bottleneck occurs, you can scale out the topic by adding more partitions.

This document describes how to create a topic in the CKafka console.

Prerequisites

You have created an instance as instructed in [Creating Instance](#).

Directions

Step 1. Create a topic

1. Log in to the [CKafka console](#).
2. In the instance list, click the **ID/Name** of your instance to enter the instance details page.
3. On the instance details page, click **Topic Management** at the top and click **Create**.
4. In the **Create Topic** window, set the **Partition Count**, **Replica Count**, and other parameters.

Create Topic ✕

Name

Remarks

Partition Count ⓘ - 1 +
3000
 Max number of partitions per topic: 3000
[Suggestions](#) [🔗](#) about the partition count

Replica Count ⓘ 1 2 replicas 3

If you select n replicas, up to (n-1) replica(s) are allowed to be down. Supported partition count * replica count. Replica quota is 150, with 2 used in the quota. You can also create up to 74 partitions with 2 replicas now. For more partitions, you can upgrade instances. For rule details, see [Documentation](#) [🔗](#).

Tag + Add

Tags are used to categorize and manage resources from different dimensions. If the existing tags do not meet your requirements, please go to the [Tag](#) [🔗](#) console to manage tags.

Preset ACL Policy

[Show advanced configuration](#)

Submit Close

Name: The topic name. It cannot be changed once entered and can contain only letters, digits, underscores, or symbols ("- and ".").

Partition Count: It is a concept in physical partition, where one topic can contain one or more partitions. CKafka uses partition as an allocation unit. For more information on how to configure the partition count, see [Configuration Guide for Common Parameters in CKafka](#).

Replica Count: The number of partition replicas that are used to ensure the high availability of the partition. To ensure data availability, two replicas are enabled by default. Replicas are also counted into the number of partitions. For example, if you create 1 topic with 6 partitions and 2 replicas for each partition, then you have a total of 12 partitions (1 x 6 x 2).

Note:

A single-replica topic cannot guarantee high availability. Select the partition count with caution.

Tag: Set a resource tag. For more information, see [Tag Overview](#).

Preset ACL Policy: Select the preset ACL policy. For more information on ACL policy, see [Configuring ACL Policy](#).

5. Click **Submit**.

Step 2. Configure advanced topic parameters

1. On the instance list page, click the **ID/Name** of the target instance to enter the instance details page.
2. On the instance details page, select **Topic Management**.
3. In the **Operation** column, click **Edit > Show advanced configuration** and set the following parameters:

cleanup.policy delete ▼

Log can be deleted by retention time, or can be compacted by key (the compact mode is required for kafka connect)

min.insync.replicas 1

The minimum number of replicas specified by "min.insync.replicas" when the producer sets "request.required.acks" to -1

unclean.leader.election.enable

segment.ms ms ▼

Segment scrolling duration, with a range of 1 to 90 day(s)

retention.ms 2 day(s) ▼

Message retention period in topic dimension. Range: 1 min to 90 day(s)

retention.bytes B ▼

Message retention size in partition dimension. Range: 1 to 1024 GB
 Message retention size for a topic = partition count * retention.bytes. If both the message retention period and message retention size are set for a topic, the actual message retention will be determined by which threshold is reached first.

max.message.bytes B ▼

When the client sends data, the data sent to the same partition will be aggregated and sent uniformly. The server will compare the size and range (1 KB to 12 MB) of each batch of messages.

Submit
Close

The parameters are as detailed below:

Parameter	Default Value	Valid Values/Value Range	Description
cleanup.policy	delete	delete/compact	Log can be deleted by retention time, or can be compacted by key (the compact mode is required for kafka connect).
min.insync.replicas	1	-	When "producer" sets

			"request.required.acks" to 1, "min.insync.replicas" will specify the minimum number of replicas.
unclean.leader.election.enable	true	true/false	This parameter specifies whether a replica not in ISR can be set as a leader.
segment.ms	-	5 minutes–90 days	The segment shard rolling duration in ms. Minimum value: 86,400,000 ms.
retention.ms	The message retention period of the instance.	60,000 ms–90 days	The message retention period at the topic level.
retention.bytes	The message retention size of the instance.	1–1,024 GB	Message retention size for a topic = partition count * retention.bytes. If both the message retention period and message retention size are set for a topic, the actual message retention will be determined by which threshold is reached first.
max.message.bytes	-	1 KB–12 MB	The maximum message size at the topic level. If this parameter is left empty, it will be 1 MB by default.

Viewing Topic

Last updated : 2024-01-09 14:47:31

Overview

This document describes how to view the topic details and producer connection after you create a topic in the CKafka console.

Directions

Step 1. View topic details

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and then the **ID/Name** of your instance to enter the instance details page.
3. On the instance details page, click **Topic Management** to view the topic information and enter the topic list page.
4. On the topic list page, click the right triangle icon on the left of the topic name to view the topic details.

ID/Name	Monitoring	Partition Count	Replica Count	Tag	Remarks	Creation Time	Message Reten...	Status
topic-5thxt31w cccc		1	2			2022-11-22 11:09:35	1 day(s)	Normal

Partition Name	Leader	Replica	ISR	Start Offset	End Offset	Messages
partition-0	110204	110204,110206	110204,110206	1	1	0

Item	Description
Partition Name	Partition name.
leader	The leader processes all read/write requests in the partition, and the follower passively and periodically replicates the data from the leader.
Replica	Replica list.
ISR	Replicas with synced messages.
Start Offset	The last position of message consumption.
End Offset	The last position of message write.

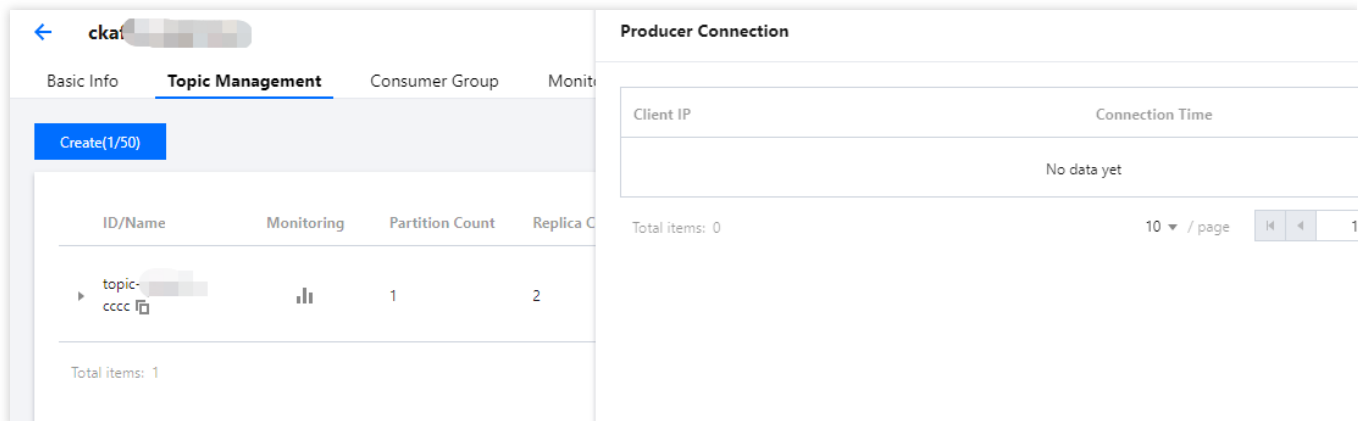
Messages	The number of stored messages.
Unsynced Replicas	The number of unsynced replicas. You can filter partitions with unsynced replicas.

Step 2. View the producer connection

Note:

Currently, you can view the producer connections only of Pro Edition instances on v2.4 or later.

1. On the instance list page, click the **ID/Name** of the target instance to enter the instance details page.
2. On the instance details page, select **Topic Management** and click **Producer Connection** in the **Operation** column to view the list of producers connected to the topic.



Deleting Topic

Last updated : 2024-01-09 14:47:32

Overview

This document describes how to delete a topic in the CKafka console.

Directions

1. Log in to the [CKafka console](#).
2. Select **Instance List** on the left sidebar and click the **ID/Name** of your instance to enter the instance details page.
3. On the instance details page, select **Topic Management** and click **Delete** in the **Operation** column.
4. In the window that pops up, click **Delete** to delete the topic.

Note:

Deleting a topic will also delete the messages stored in the topic. Do so with caution.

Deleting a topic is an async operation. After the configuration is deleted successfully, the ZooKeeper configuration will take effect in 1 minute. If a topic with the same name is created during this period, the system will prompt an error code [4000]10011, and you can try again later.

Sending Message


Last updated : 2024-01-09 14:47:32

Overview

This document describes how to send a message and view its record in message query after creating a topic in the CKafka console.

Directions

1. Log in to the [CKafka console](#).
2. Select **Instance List** on the left sidebar and click the **ID/Name** of your instance to enter the instance details page.
3. On the instance details page, select **Topic Management** and click **Send Message** in the **Operation** column.



Send Message ✕

Message Content
The message is only used for debugging and can contain up to 64 characters.

Message Key

Send to Specified Partition

OK

Message Content: Enter the content of the message to be sent, which is required.

Message Key: Enter the sending key, which is optional.

Send to Specified Partition: This parameter supports sending messages to the specified partition, which is disabled by default.

4. Click **OK** to send the message. In the **Sent the message successfully** pop-up window, click **Message Query** to view the message just sent.

Setting Topic Traffic Throttling Rule

Last updated : 2024-01-09 14:47:31

Overview

You can set topic traffic throttling rules to prevent the excessive traffic of one topic from affecting other topics.

Note:

Setting topic traffic throttling rules are only available to broker on versions 1.1.1, 2.4.1, and 2.8.1.

Directions

1. Log in to the [CKafka console](#).
2. Select **Instance List** on the left sidebar and click the **ID/Name** of your instance to enter the instance details page.
3. On the instance details page, select **Topic Management**.
4. Click **More > Traffic Throttling** in the **Operation** column to set the traffic throttling threshold.

Maximum Topic Production Traffic: This value excludes replica traffic and ranges from 1 MB/s to the maximum bandwidth purchased for the instance / number of replicas of the topic.

Maximum Topic Consumption Traffic: This value ranges from 1 MB/s to the maximum bandwidth purchased for the instance.

Note:

The underlying layer throttles the traffic for brokers, and the actual traffic throttling value (equal to an integer multiple of the number of brokers) may be slightly different from the set value.

For soft traffic throttling mechanism, see [Traffic Throttling](#).

Elastic Topic

Topic Management

Last updated : 2024-01-09 14:47:31

Overview

Topic is a category name where messages are stored and published. Producers write messages to topics, and consumers read messages from topics. To enable horizontal scaling, a topic is divided into multiple partitions. This allows you to horizontally scale your resources by adding more partitions in case of performance bottlenecks. A topic is the smallest unit of message subscription and publishing in Connector. You can use a topic to denote a type of streaming data.

CKafka supports the **elastic topic** type. You can directly create an elastic topic in the console with no need to purchase a CKafka instance first. You can use the SDK to send and receive messages through the elastic topic. You can also use the elastic topic as the data source or target of a data task.

Note:

Currently, billing has not been started for elastic topics. We will inform you of the billing start date in advance.

Use Limits

A single elastic topic can have up to 500 partitions, and the elastic topic bandwidth has no default size. It is recommended that the bandwidth not exceed 500 Mbps.

The maximum message size supported by an elastic topic is 12 MB.

Directions

Creating a topic


1. Log in to the [CKafka console](#).
2. Click **Elastic Topic** on the left sidebar, select the region, and click **Create Topic**.
3. Enter the topic information and click **OK**.

Topic Name: The topic name is in the format of "appid-[the name you enter]". It can only contain letters, digits, underscores, or symbols ("- and ".").

Remarks: Enter the optional remarks of the topic.

Partition Count: It is a concept in physical partition, where one topic can contain one or more partitions. CKafka uses partition as an allocation unit. A single topic can have up to 500 partitions.

Message Retention Period: This value can range from 1 minute to 90 days. When the disk capacity is insufficient (that is, the disk utilization reaches 90%), old messages will be deleted in advance to ensure the service availability.



The screenshot shows a dialog box for editing a topic. It contains the following fields and controls:

- Topic Name:** A text input field with the value "1305469081-" and a placeholder "Please enter a topic name".
- Remarks:** A text input field with the placeholder "Optional. Please enter remarks."
- Partition Count:** A slider control ranging from 1 to 500. The current value is 1. Below the slider, it says "Max number of partitions in a single topic: 500".
- Message Retention Period:** A text input field with the value "1" and a dropdown menu set to "day(s)". Below it, it says "Range: 1 min to 90 day(s)".
- Buttons:** "OK" (blue) and "Cancel" (white) buttons at the bottom left.

Querying a topic

You can search for topics by topic name in the search box in the top-right corner of the [Elastic Topic](#) list. Connector will perform a fuzzy match and display the search results.


Editing a topic

1. In the [elastic topic](#) list, click **Edit** in the **Operation** column of the target topic.
2. In the pop-up window, you can edit the topic remarks, partition count, and message retention period.
3. Click **OK** to complete your edits.

Sending a message

You can manually send messages to the specified topic in the Ckafka console.

1. In the [Elastic Topic](#) list, click **Send Message** in the **Operation** column of the target topic.
2. Enter the message content and message key. You can also specify to send the message to the specified partition.



The screenshot shows a dialog box titled "Send Message". It contains the following fields and controls:


- Message Content:** A text input field with the value "Hello World".
- Message Key:** A text input field with the value "test".
- Send to Specified Partition:** A toggle switch that is currently turned on (blue).
- Partition ID:** A dropdown menu with the value "0".
- Buttons:** "OK" (blue) and "Close" (white) buttons at the bottom.

3. Click **OK** to send the message. After the message is sent, it can be consumed by any subscribers to the topic.

Adding a subscription

You can manually create a subscription in the console.

1. In the [Elastic Topic](#), click the **ID** of the target topic to enter its basic information page.
2. Select the **Subscription** tab at the top, click **Create Subscription**, and enter the consumer name in the pop-up window.




3. Click **Submit**, and you can see the created subscription in the list.

Deleting a topic

Note:

After a topic is deleted, all unconsumed messages retained in it will be cleared; therefore, proceed with caution.

1. In the [Elastic Topic](#) list, click **Delete** in the **Operation** column of the target topic.
2. In the pop-up window, click **Delete**.



Subscription Management

Last updated : 2024-01-09 14:47:31

Overview

A subscription represents a specific consumer and its subscription to a topic. A consumer can consume all messages in a topic after subscribing to it. CKafka's subscription service provides the feature of saving the consumption offset on the server. This document describes how to add a subscription and reset the consumption offset in the CKafka console.

Prerequisites

You have created a topic as instructed in [Creating Topic](#).

Directions

Adding a subscription

Add a subscription as instructed in [Instance Topic > Creating Topic](#).

Setting an offset

In scenarios such as offline data processing, sometimes it is necessary to reset the offset to consume message history.

1. In the [Topic List](#), click the target topic ID to enter the topic's **Basic Info** page.
2. Select the **Subscription** tab at the top and click **Offset Settings** in the **Operation** column of the target subscription.

3. In the **Offset Settings** window, select **Topic** or **Partition** as the dimension of settings and click **Next**.

Set based on Topic: It is suitable for scenarios such as batch resetting of consumption offsets and setting consumption offsets according to businesses.

Set based on Partition: It is suitable for setting consumption offset for each partition after knowing the consumption situation of each Partition.

4. Select the topic or partition for which to reset the offset (if no topic is selected, the offset will be reset for all topics by default) and click **Next**.
5. Specify the offset.

Set Offset

1 **Select object** > 2 **Offset Settings**

Offset Settings Move the offset to the specified position

Please enter the offset

The adjustment range of offset is between its min and max values.
If the input position is less/greater than the offset range, it will be reset to the min/max offset.

Move the offset forward or backward several rows

Consume from the latest/start offset

Reset consumption position by time-in-point

Previous Submit

Note :

The offset value should be between the minimum offset and the maximum offset. If it is configured to be smaller than the minimum offset, the consumption will start from the minimum offset; if it is larger than the maximum offset, the consumption will start from the maximum offset.

Make sure that there are no consumers in a consumer group before resetting the group.

Viewing consumer details

On the subscription list page, click **View Consumer Details** in the **Operation** column to view the information of consumers in the consumer group and the relationship between a consumer and the topic subscribed to by the consumer.

On the subscription list page, click the small triangle on the left of a consumer group name to display the information of the topic subscribed to by the consumer group, including topic name, number of partitions, submitted offset position, maximum offset position, and number of unconsumed messages.

Click **View Details** in the **Operation** column to view offset consumption at the partition level.

Note :

As the offset information is maintained on the consumer side, the offset position is subject to the way the consumer submits the offset. It is displayed asynchronously and does not necessarily represent real-time consumption conditions.

Deleting a subscription

In the subscription list, click **Delete** in the **Operation** column of the target consumer and click **OK** to delete the consumer.

Note :

After you delete the consumer group, when the consumers in it re-establish connections for consumption, the offset will be reset, and consumption will start from the beginning.


Consumer Group

Consumer Group Overview

Last updated : 2024-01-09 14:47:32

Consumer Group Status Description

A consumer group can be in Dead, Empty, Preparing Rebalance, AwaitingSync, or Stable state. Most common states are Empty, Stable, and Dead. The transformation of the consumer group's state machine is as shown below:



Dead: consumer group does not have any members and metadata has been removed.

Empty: consumer group currently does not have any members. If all offsets in the group have expired, the group's state will be Dead. Generally, a newly created group is in Empty state by default.

As required by the open-source Kafka v0.10.x, **a consumer group will be automatically deleted if it does not have any members for more than 7 days.**

Stable: every consumer in the consumer group has joined and the state is stable.

Rebalance State Details

Cause of rebalance

According to the consumer group's state machine, the group may perform rebalance when the consumer group is in Empty, AwaitingSync, or Stable state. Rebalance may occur in the following scenarios:

A consumer subscribes to a topic.

A consumer is closed.

A consumer is considered to be in Dead state by the group coordinator.

If a consumer does not send a heartbeat to the group coordinator within the `session.timeout.ms` time period, the consumer will be considered to be in Dead state and rebalance will be initiated. For more information, please see [CKafka Common Parameter Configuration Guide](#).

The number of partitions has increased.

A topic that does not exist is subscribed to.

If you subscribe to a topic that has not been created yet, rebalance will occur when the topic is created. Similarly, if a topic that has already been subscribed to is deleted, rebalance will also occur.

Application crashes.

Rebalance process analysis

Taking the mechanism of Kafka v0.10 as an example, the rebalance process is analyzed as follows:

1. A consumer that wants to join a consumer group will send a JoinGroup request to the group coordinator. The first consumer to join a group will become the group leader.
2. The leader will receive a list of all consumers in the group from the coordinator and will be responsible for assigning partitions to consumers in the group. Partition assignment can be implemented through the PartitionAssignor API.
3. When the assignment is done, the leader will send the assignment result to the coordinator. The coordinator will then send the result to all consumers.

Each consumer can only view the partition assigned to it. The leader is the only consumer that can get the information of consumers in the consumer group and their partitions.

The above process will be executed once whenever rebalance occurs.

Creating Consumer Group

Last updated : 2024-01-09 14:47:32

Overview

This document describes how to create a consumer group directly in the CKafka console.


Note:

We recommend that you limit the number of consumer groups in each instance to 200 or less, since excess consumer groups will result in certain constraints.

Directions

Creating a consumer group

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and then the **ID/Name** of your instance to enter the instance details page.
3. On the instance details page, select the **Consumer Group** tab and click **Create Consumer Group**.



Create Consumer Group ✕

Consumer Group Name

Select Topic

4. In the pop-up window, enter the consumer group name and select the topic to which you want to subscribe. Note that the consumer group name cannot contain quotation marks ("").


Note:

You can select multiple topics.

5. Click **Submit**, and you can see the created consumer group in the consumer group list.

Disabling automatic consumer group creation

CKafka allows you to disable automatic consumer group creation in the console. After it is disabled, only existing consumer groups in the console can be used for consumption, while new data sync tasks cannot be created normally.

**Note:**

Only Pro Edition 2.4.1 and later are supported.

Querying Consumer Group

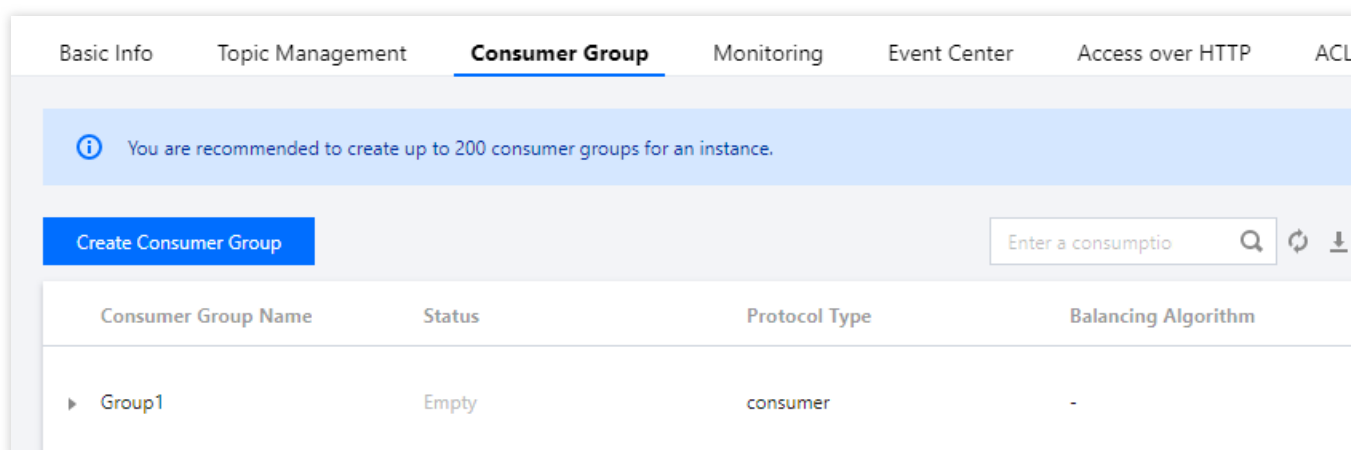
Last updated : 2024-01-09 14:47:32

Overview

This document describes how to view the consumer group information of an instance in the CKafka console.

Directions

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and then the **ID/Name** of your instance to enter the instance details page.
3. On the instance details page, click the **Consumer Group** tab to view the consumer group information of the current CKafka instance.



On the **Consumer Group** list page, click **View Consumer Details** in the **Operation** column to view the information of consumers in the consumer group and the relationship between a consumer and the topic subscribed to by the consumer.

On the **Consumer Group** list page, click the small triangle on the left of a consumer group name to display the information of the topic subscribed to by the consumer group, including topic name, number of partitions, submitted offset position, maximum offset position, and number of unconsumed messages. You can click **View Details** in the **Operation** column to view offset consumption at the partition level.

Note:

As the offset information is maintained on the consumer side, the offset position is subject to the way the consumer submits the offset. It is displayed asynchronously and does not necessarily represent real-time consumption conditions.

Deleting Consumer Group

Last updated : 2024-01-09 14:47:31

Overview

In some cases, a consumer group may not consume for a long time before consuming again. You can delete it, so that when the consumers in it establish a connection again, the offset will be reset, and consumption will start from the beginning.

Deleting a consumer group

Note:

Only empty consumer groups with a broker version not below v1.1.1 can be deleted.

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and then the **ID/Name** of your instance to enter the instance details page.
3. On the instance details page, click **Delete** in the **Operation** column of the consumer group to be deleted.

Consumer Group Name	Status	Protocol Type	Balancing Algorithm	Operation
▶ group2	Empty	-	-	Offset Settings Delete

Deleting a subscription to an associated topic

If the status of a consumer group is **Empty**, the subscriptions of a topic associated with it can be deleted.

On the **Consumer Group** list page, click the



icon next to the subscription to be deleted and click **OK** to delete it.

Setting Offset

Last updated : 2024-01-09 14:47:32

Overview

In scenarios such as offline data processing, sometimes it is necessary to reset the offset to consume earlier messages.

Directions

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and click the **ID/Name** of the target instance to enter the instance details page.
3. On the instance details page, select **Consumer Group** and click **Offset Settings** in the **Operation** column.
4. In the **Offset Settings** window, select **Topic** or **Partition** as the dimension of settings and click **Next**.
5. Select the topic or partition for which to reset the offset (if no topic is selected, the offset will be reset for all topics by default) and click **Next**.
6. Specify the offset.

Offset Settings

✓ Select Topic > 2 Offset Settings

Offset Settings Move the offset to the specified position

The adjustment range of offset is between its min and max values.
If the input position is less/greater than the offset range, it will be reset to the min/max offset position.

Move the offset forward or backward several rows

Consume from the latest/starting position

Reset consumption position by time-in-point

Note:

The offset value should be between the minimum offset and the maximum offset. If it is configured to be smaller than the minimum offset, the consumption will start from the minimum offset; if it is larger than the maximum offset, the consumption will start from the maximum offset.

Make sure that there are no consumers in a consumer group before resetting the group.

Automatically Resetting Offset

Last updated : 2024-01-09 14:47:32

This document describes the concept and usage of the `auto.offset.reset` parameter.

What is `auto.offset.reset` ?

The `auto.offset.reset` parameter defines where to start consumption if the offset of the partition to be consumed cannot be obtained. For example, it specifies how the offset will be initialized if no offset is configured for the broker (such as upon initial consumption or when the offset expired for more than seven days) or how it will be reset if the error `OFFSET_OUT_OF_RANGE` occurs.

The `auto.offset.reset` parameter has the following valid values:

`earliest` : Reset to the minimum offset in the partition.

`latest` : Reset to the maximum offset in the partition. This is the default value.

`none` : Report an `OffsetOutOfRangeException` exception without resetting the offset.

When will `OFFSET_OUT_OF_RANGE` occur?

This error indicates that the offset committed by the client is out of the offset range allowed by the broker. For example, if `LogStartOffset` and `LogEndOffset` of partition 1 in `topicA` are 100 and 300 respectively, but the offset committed by the client is less than 100 or greater than 300, then the broker will return this error, and the offset will be reset.

This error may occur in the following cases:

If an offset is set on the client but not used for consumption during a period of time, it will be deleted on the broker due to log scrolling after the message retention period set for the topic has elapsed. In this case, if it is recommitted by the client after the deletion, this error will occur.

If the client commits an abnormal offset due to issues such as SDK bugs and network packet loss, this error will occur.

If there are unsynced replicas on the broker and the leader is switched, follower replicas will be truncated. In this case, if the offset committed by the client falls in the truncated range, this error will occur. For more information on leader switch, see [Client > What is leader switch?](#).

<code>auto.offset.reset=none</code>	Description
-------------------------------------	-------------

Background

You don't want the offset to be automatically reset, as your business doesn't allow such large-scale repeated consumption.

Note:

In this case, the consumer group will report an error for failing to find the offset in its first consumption. Therefore, you need to manually set the offset in `catch` .

Note

After `auto.offset.reset` is set to `none` , automatic offset reset can be avoided; however, as the automatic reset mechanism is disabled, when a new partition is added, the client doesn't know where to start consuming the new partition, and an exception will occur. In this case, you need to manually set a consumer group offset and start consuming.

Instructions

During consumption, if you set `auto.offset.reset` to `none` for the consumer, you need to capture the exception `NoOffsetForPartitionException` and set the offset in `catch` on your own. You can select one of the following methods based on your actual business needs:


Specify the offset. You need to maintain the offset, which is convenient for retries.

Specify to start consumption from the beginning.

Use the nearest available offset.

Obtain and set the offset based on the timestamp.

Below is the sample code:



```
package com.tencent.tcb.operation.ckafka.plain;

import com.google.common.collect.Lists;
import com.tencent.tcb.operation.ckafka.JavaKafkaConfigurer;
import java.time.Instant;
import java.time.temporal.ChronoUnit;
import java.util.ArrayList;
import java.util.Collection;
import java.util.HashMap;
import java.util.List;
import java.util.Map;
```

```
import java.util.Map.Entry;
import java.util.Properties;
import org.apache.kafka.clients.CommonClientConfigs;
import org.apache.kafka.clients.consumer.ConsumerConfig;
import org.apache.kafka.clients.consumer.ConsumerRecord;
import org.apache.kafka.clients.consumer.ConsumerRecords;
import org.apache.kafka.clients.consumer.KafkaConsumer;
import org.apache.kafka.clients.consumer.NoOffsetForPartitionException;
import org.apache.kafka.clients.consumer.OffsetAndTimestamp;
import org.apache.kafka.clients.producer.ProducerConfig;
import org.apache.kafka.common.PartitionInfo;
import org.apache.kafka.common.TopicPartition;
import org.apache.kafka.common.config.SaslConfigs;

public class KafkaPlainConsumerDemo {

    public static void main(String args[]) {
        // Set the path of the JAAS configuration file.
        JavaKafkaConfigurer.configureSaslPlain();

        // Load `kafka.properties`.
        Properties kafkaProperties = JavaKafkaConfigurer.getKafkaProperties();

        Properties props = new Properties();
        // Set the access point. Obtain the access point of the corresponding topic
        props.put(ProducerConfig.BOOTSTRAP_SERVERS_CONFIG, kafkaProperties.getPrope

        // Set the access protocol.
        props.put(CommonClientConfigs.SECURITY_PROTOCOL_CONFIG, "SASL_PLAINTEXT");
        // Set the PLAIN mechanism.
        props.put(SaslConfigs.SASL_MECHANISM, "PLAIN");
        // Set the maximum interval between two polls.
        // If the consumer does not return a heartbeat message within the interval,
        props.put(ConsumerConfig.SESSION_TIMEOUT_MS_CONFIG, 30000);
        // Set the maximum number of messages that can be polled at a time.
        // Do not set this parameter to an excessively large value. If polled messa
        props.put(ConsumerConfig.MAX_POLL_RECORDS_CONFIG, 30);
        // Set the method for deserializing messages.
        props.put(ConsumerConfig.KEY_DESERIALIZER_CLASS_CONFIG,
            "org.apache.kafka.common.serialization.StringDeserializer");
        props.put(ConsumerConfig.VALUE_DESERIALIZER_CLASS_CONFIG,
            "org.apache.kafka.common.serialization.StringDeserializer");
        // Set the consumer group for the current consumer instance. You need to ap
        // The instances in the same consumer group consume messages in load balanc
        props.put(ConsumerConfig.GROUP_ID_CONFIG, kafkaProperties.getProperty("grou

        // Consumption offset Note If `auto.offset.reset` is set to `none`, the con
```

```
props.put(ConsumerConfig.AUTO_OFFSET_RESET_CONFIG, "none");
// Construct a consumer object. This generates a consumer instance.
KafkaConsumer<String, String> consumer = new KafkaConsumer<String, String>(
// Set one or more topics to which the consumer group subscribes.
// We recommend that you configure consumer instances with the same `GROUP_
List<String> subscribedTopics = new ArrayList<String>();
// If you want to subscribe to multiple topics, add the topics here.
// You must create the topics in the console in advance.
String topicStr = kafkaProperties.getProperty("topic");
String[] topics = topicStr.split(",");
for (String topic : topics) {
    subscribedTopics.add(topic.trim());
}
consumer.subscribe(subscribedTopics);
// Consume messages in loop.
while (true){
    try {
        ConsumerRecords<String, String> records = consumer.poll(1000);
        // All messages must be consumed before the next poll, and the total
        for (ConsumerRecord<String, String> record : records) {
            System.out.println(
                String.format("Consume partition:%d offset:%d", record.
        }
    } catch (NoOffsetForPartitionException e) {
        System.out.println(e.getMessage());

        // If you set `auto.offset.reset` to `none`, you need to capture the
        // Sample 1. Specify the offset. You need to maintain the offset, with
        Map<Integer, Long> partitionBeginOffsetMap = getPartitionOffset(consumer,
        Map<Integer, Long> partitionEndOffsetMap = getPartitionOffset(consumer,
        consumer.seek(new TopicPartition(topicStr, 0), 0);

        // Sample 2. Specify to start consumption from the beginning
        consumer.seekToBeginning(Lists.newArrayList(new TopicPartition(topicStr, 0),

        // Sample 3. Use the nearest available offset
        consumer.seekToEnd(Lists.newArrayList(new TopicPartition(topicStr, 0),

        // Sample 4. Obtain and set the offset based on the timestamp. For
        Map<TopicPartition, Long> timestampsToSearch = new HashMap<>();
        Long value = Instant.now().minus(300, ChronoUnit.SECONDS).toEpochMilli();
        timestampsToSearch.put(new TopicPartition(topicStr, 0), value);
        Map<TopicPartition, OffsetAndTimestamp> topicPartitionOffsetAndTimestamp
            .offsetsForTimes(timestampsToSearch);
        for (Entry<TopicPartition, OffsetAndTimestamp> entry : topicPartitionOffsetAndTimestamp
            .entrySet()) {
            TopicPartition topicPartition = entry.getKey();
```

```
        OffsetAndTimestamp entryValue = entry.getValue();
        consumer.seek(topicPartition, entryValue.offset()); // Specify
    }

    }
}

/**
 * Get the earliest and nearest offsets of the topic
 * @param consumer
 * @param topicStr
 * @param beginOrEnd true begin; false end
 * @return
 */
private static Map<Integer, Long> getPartitionOffset(KafkaConsumer<String, Stri
    boolean beginOrEnd) {
    Collection<PartitionInfo> partitionInfos = consumer.partitionsFor(topicStr)
    List<TopicPartition> tp = new ArrayList<>();
    Map<Integer, Long> map = new HashMap<>();
    partitionInfos.forEach(str -> tp.add(new TopicPartition(topicStr, str.parti
    Map<TopicPartition, Long> topicPartitionLongMap;
    if (beginOrEnd) {
        topicPartitionLongMap = consumer.beginningOffsets(tp);
    } else {
        topicPartitionLongMap = consumer.endOffsets(tp);
    }
    topicPartitionLongMap.forEach((key, beginOffset) -> {
        int partition = key.partition();
        map.put(partition, beginOffset);
    });
    return map;
}
}
```

Monitoring and Alarms

Viewing Monitoring Data

Last updated : 2024-01-09 14:47:32

Overview

CKafka allows you to monitor the resources created under your account, such as instances, topics, and consumer groups, so that you can keep track of the status of your resources in real time and troubleshoot possible issues to ensure stable business operations.

This document describes how to view monitoring metrics and their descriptions in the CKafka console.

Directions

1. Log in to the [CKafka console](#).
2. In the instance list, click the **ID/Name** of the target instance to enter the instance details page.
3. At the top of the instance details page, Click **Monitor**, select resource tags, and set the time range to view the monitoring data about the instance.

Monitoring information display

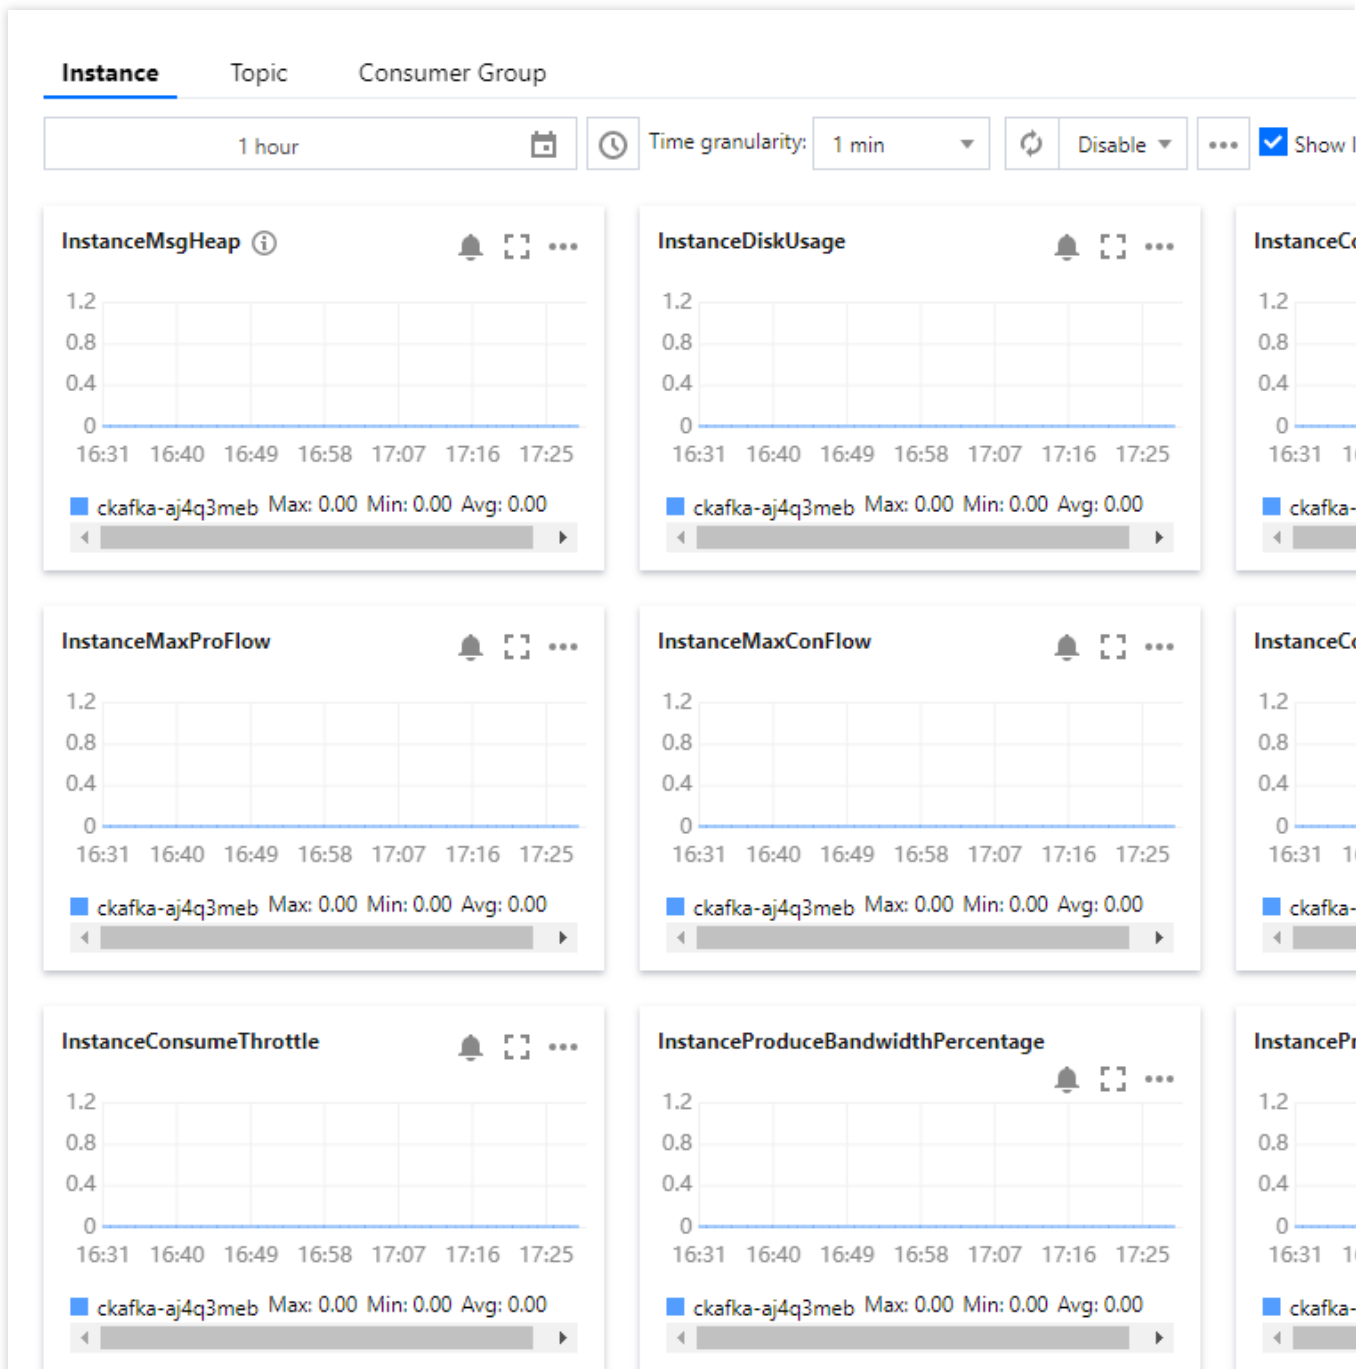
Note:

You can click the following tabs to view the monitoring data of instances, topic, and consumer groups.

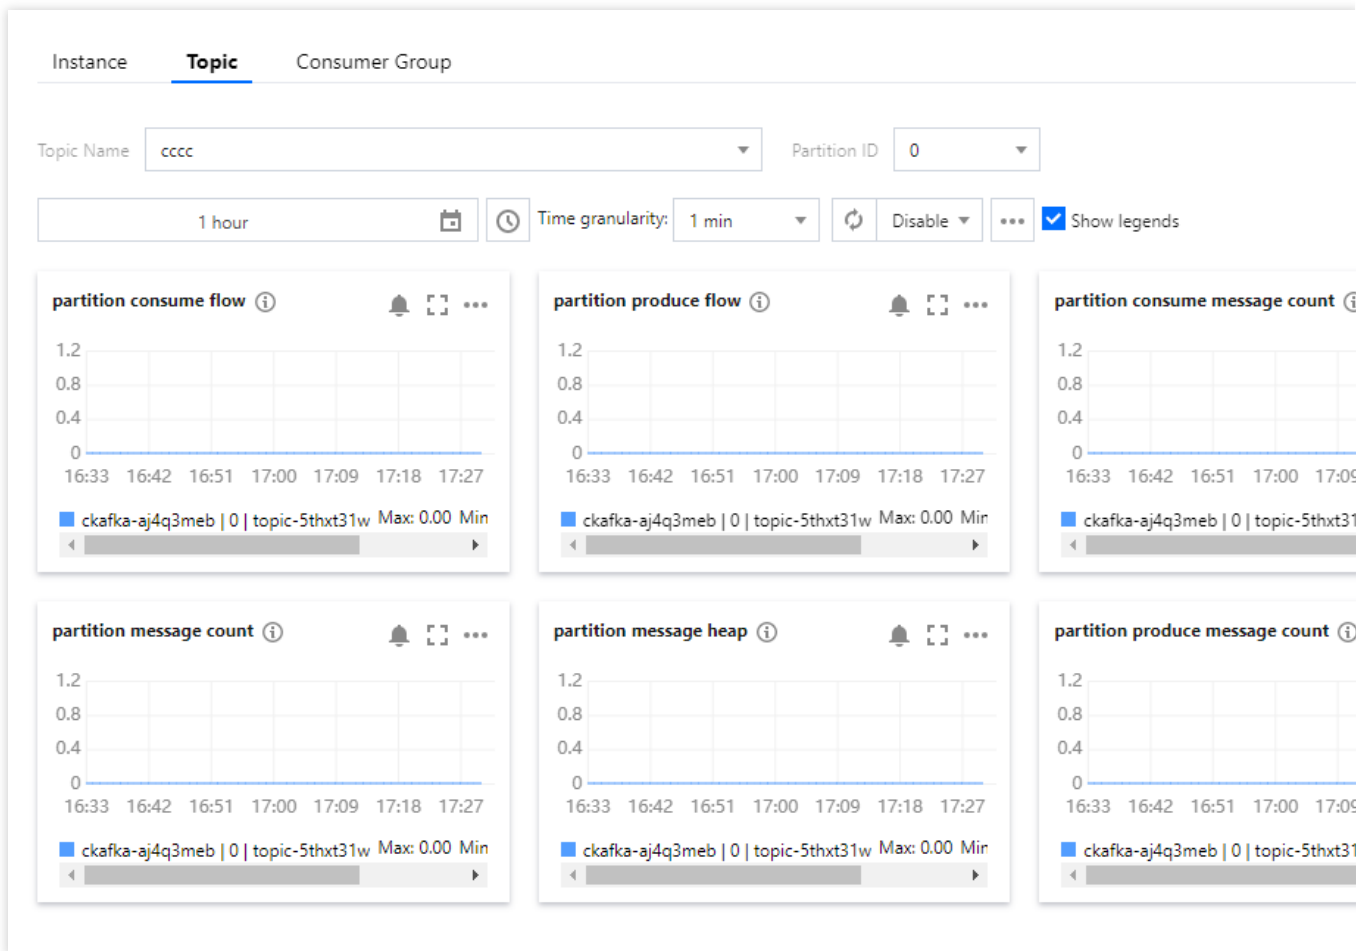
Instance

Topic

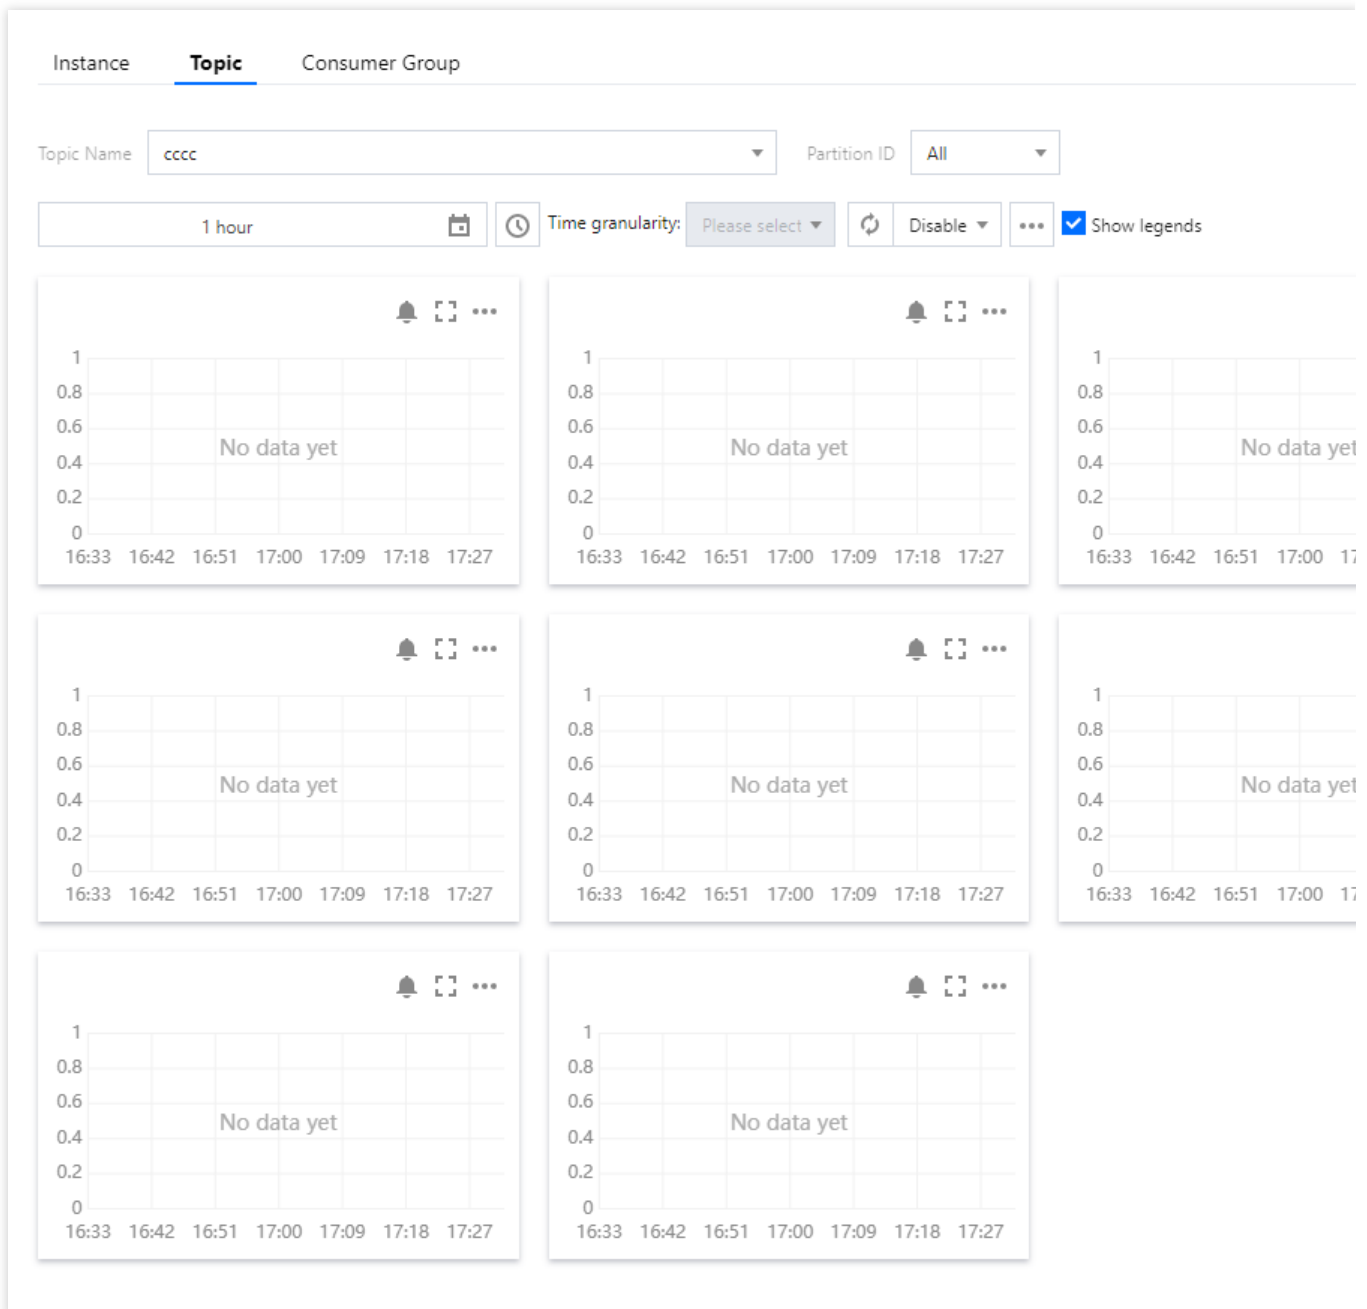
Consumer Group



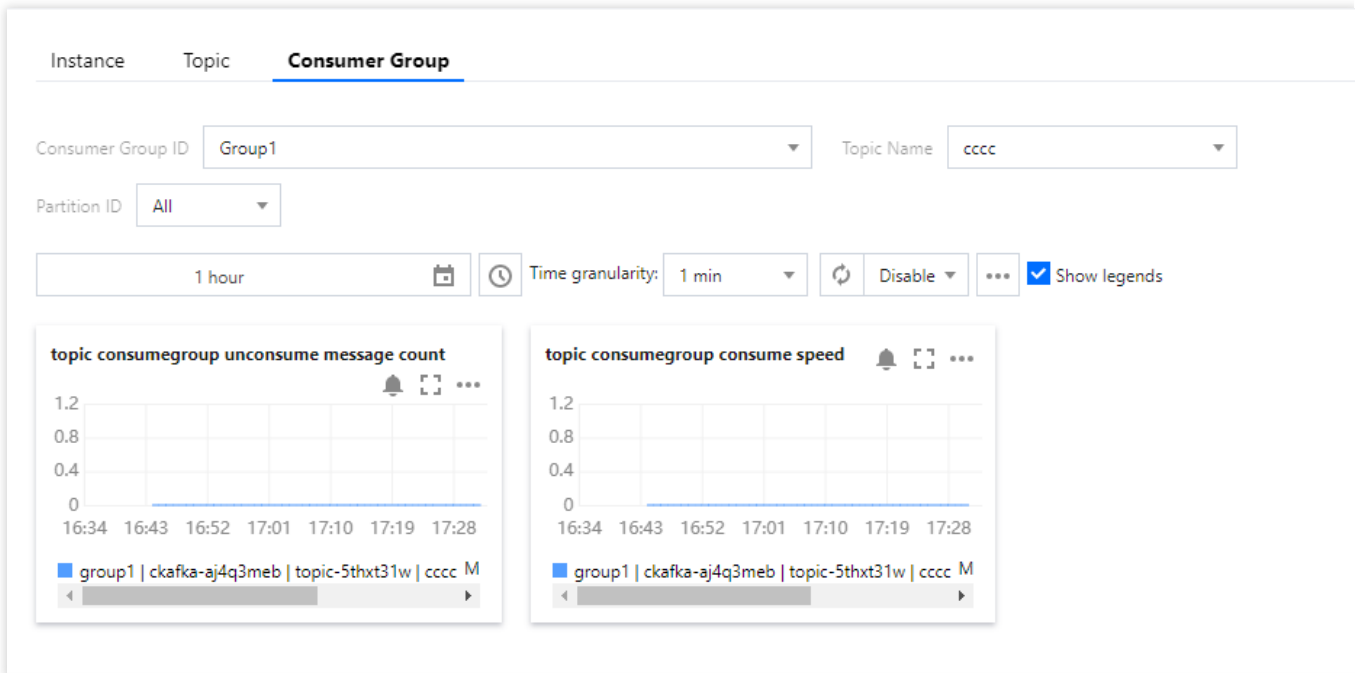
After selecting a partition, you can view its monitoring data.



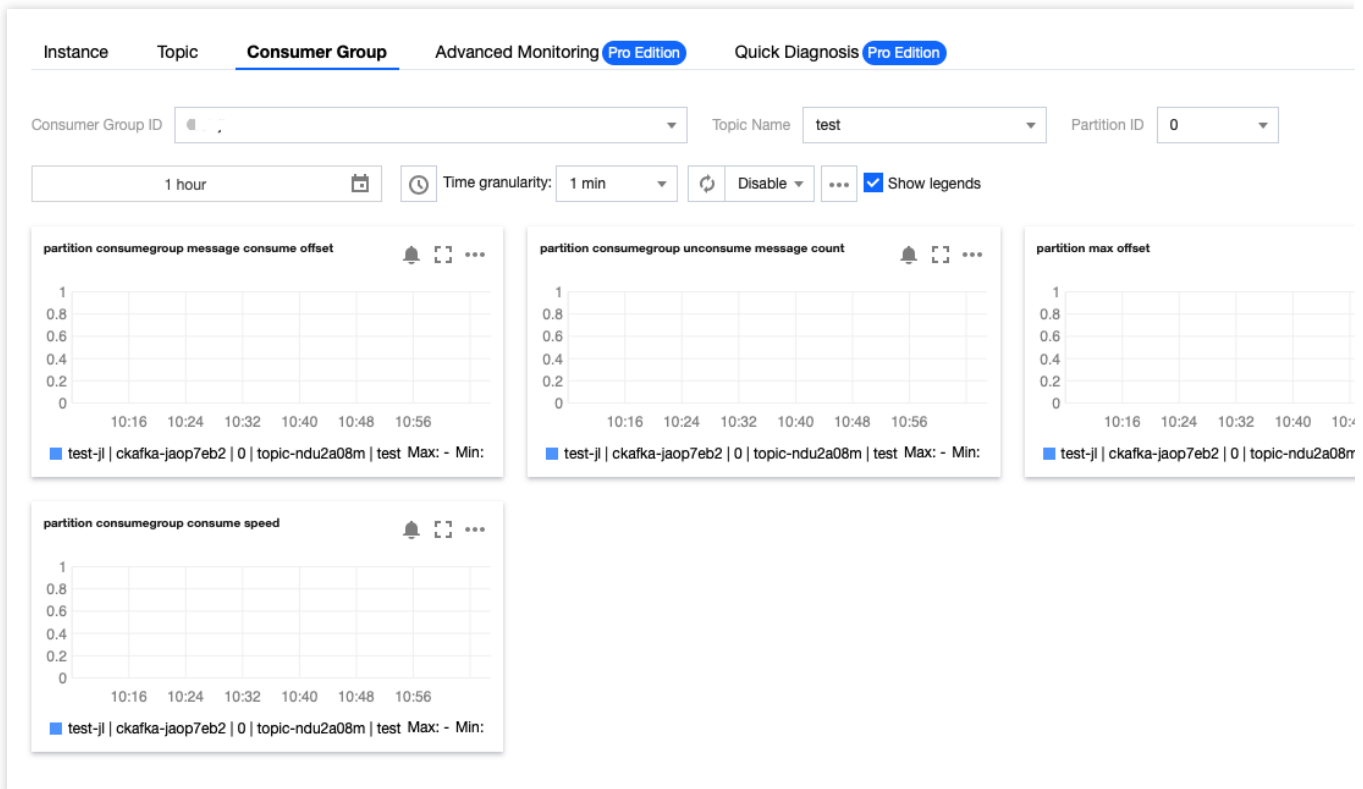
If you don't select a partition, the current monitoring data at the topic level will be displayed.



Consumer group - topic monitoring



Consumer group - partition monitoring:



Descriptions of metrics

The following lists the CKafka monitoring metrics and their description. Based on user feedback and statistics, we recommend that you configure monitoring and alarming policies for metrics marked with the word **recommended**. Configure policies properly as needed.

Instance
Topic

Consumer Group

Monitoring Metric	Description
Instance Disk Usage (MB)	Instance disk usage (including replicas), which is the latest value in the selected time period.
Disk Utilization (%) (Recommended)	Ratio of the currently used disk capacity to the total disk capacity of the instance in percentages.
Instance Connections (Count) (Recommended)	Number of connections between clients and servers in the cluster.
Max Production Traffic in Instance (MB) (Recommended)	Instance message production bandwidth peak (excluding traffic generated by replicas). This is the basis for calculating the production bandwidth utilization in the instance.
Max Consumption Traffic in Instance (MB) (Recommended)	Instance message consumption bandwidth peak (there is no concept of replica for consumption). This is the basis for calculating the consumption bandwidth utilization in the instance.
Instance Consumption Bandwidth Percentage (%)	Ratio of the instance consumption bandwidth to the quota.
Instance Consumption Throttling Occurrences (Count)	Number of instance consumption throttling occurrences.
Instance Production Bandwidth Percentage (%)	Ratio of the instance production bandwidth to the quota.
Instance Production Throttling Occurrences (Count)	Number of instance production throttling occurrences.
Messages Stored in Instance (Count)	Total number of stored messages in the instance (excluding those produced by replicas), which is the latest value in the selected time period.
Instance Consumption Traffic (MB)	Consumption traffic of the instance (excluding traffic generated by replicas), which is the total traffic in the selected time period.
Messages Produced in Instance (Count)	Number of messages produced in the instance, which is the total number of messages in the selected time period.
Instance Production Traffic (MB)	Production traffic of the instance (excluding traffic generated by replicas), which is the total traffic in the selected time period.
Messages Consumed in Instance (Count)	Number of messages consumed in the instance, which is the total number of messages in the selected time period.

Topic monitoring

Monitoring Metric	Description
Max Production Traffic (MB/sec)	Maximum production traffic in topic (excluding replica traffic).
Max Consumption Traffic (MB/sec)	Maximum consumption traffic in topic.
Consumed Messages (Count)	Number of messages consumed in the topic, which is the total number of messages in the selected time period.
Consumption Traffic (MB)	Consumption traffic in the topic (excluding traffic generated by replicas), which is the total traffic in the selected time period.
Used Disk Capacity (MB) (Recommended)	Total size of messages in the topic (excluding those produced by replicas) that actually use disk capacity, which is the latest value in the selected time period.
Produced Messages (Count)	Number of messages produced in the topic, which is the total number of messages in the selected time granularity.
Production Traffic (MB)	Production traffic in the topic (excluding traffic generated by replicas), which is the total traffic in the selected time granularity.
Persisted Messages (Count)	Total number of stored messages in the topic (excluding those produced by replicas), which is the latest value in the selected time period.
Topic Consumption Throttling Occurrences (Count/sec)	Number of topic consumption throttling occurrences, which is the total number of occurrences in the selected time granularity.
Topic Production Throttling Occurrences (Count/sec)	Number of topic production throttling occurrences, which is the total number of occurrences in the selected time granularity.

Topic - partition monitoring

Monitoring Metric	Description
Consumed Messages (Count)	Number of messages consumed in the partition, which is the total number of messages in the selected time period.
Consumption Traffic (MB)	Consumption traffic in the partition (excluding traffic generated by replicas), which is the total traffic in the selected time period.

Persisted Messages (Count)	Total number of stored messages in the partition (excluding those produced by replicas), which is the latest value in the selected time period.
Used Disk Capacity (MB)	Total size of messages in the partition (excluding those produced by replicas) that actually use disk capacity, which is the latest value in the selected time period.
Produced Messages (Count)	Number of messages produced in the partition, which is the total number of messages in the selected time granularity.
Production Traffic (MB)	Production traffic in the partition (excluding traffic generated by replicas), which is the total traffic in the selected time granularity.

Consumer group - topic monitoring

Monitoring Metric	Description
Max Offset in Topic	Maximum offset in all partitions of the topic.
Consumption Offset in Topic	Maximum consumption offset of consumer groups in all partitions of the topic.
Unconsumed Messages in Topic (Count)	Total number of messages not consumed by the consumer group in all partitions of the topic.
Consumption Speed in Topic (Count/min)	Total number of messages consumed by consumer groups from all partitions in the topic per second.

Consumer group - partition monitoring:

Monitoring Metric	Description
Consumption Speed in Partition (Count/min)	Number of messages consumed by consumer groups in the partition per second.
Current Consumption Offset	Current consumption offset of consumer groups in the partition.
Max Offset for Current Partition	Maximum offset for the current partition.
Unconsumed Messages (Count) (Recommended)	Number of messages not consumed by consumer groups in the partition.

For more information on the metric names and descriptions, see [Topic Monitoring Metrics](#).

CKafka monitoring API documentation

For more information, see the following documents:

[Topic Monitoring Metrics](#)

[Instance Monitoring Metrics](#)

[Consumer Group Monitoring Metrics](#)

Querying Advanced Monitoring (Pro Edition)

Last updated : 2024-01-09 14:47:32

Overview

CKafka Pro Edition supports advanced monitoring. You can view metrics such as core services, production, consumption, and broker GC in the console, making it easier for you to troubleshoot CKafka issues.

This document describes how to view advanced monitoring metrics in the console and explains their meanings.

Directions

1. Log in to the [CKafka console](#).
2. In the instance list, click the **ID/Name** of the target instance to enter the instance details page.
3. At the top of the instance details page, click **Monitoring > Advanced Monitoring**, select the metric to be viewed, and set the time range to view the monitoring data.

Monitoring information display

Note:

You can click the following tabs to view the monitoring information of the core service, production, consumption, instance resource, and broker GC.

Core Services

Production

Consumption

Instance Resources


Broker GC

Instance Topic Consumer Group **Advanced Monitoring** Pro Edition Quick Diagnosis Pro Edition

▲ **Broker Status**

Live Broker Nodes

0%



Advanced Metric

Core Services Production Consumption Instance Resources Broker GC

Real-Time Last 24 hours Last 7 days 2022-05-16 16:31:34 ~ 2022-05-16 17:31:34 Time Granularity 1 minute

Network Status

Current Status

Busy



Request Queue Depth

Current Status

0



Unsynced Replicas [View details of unsynced replicas](#)

Current Status

0



ZK Reconnections

Current Time Range

0 times



ISR Expansions

Current Time Range

0 times



ISR Shrinks

Current Time Range

0 times




Advanced Metric


Core Services **Production** Consumption Instance Resources Broker GC

Real-Time Last 24 hours Last 7 days 2022-05-16 16:33:15 ~ 2022-05-16 17:33:15 Time Granularity 1 minute Percentile p95

Total Production Duration



Production Duration Distribution




Click the bar chart on the left to view details

Advanced Metric


Core Services Production **Consumption** Instance Resources Broker GC

Real-Time Last 24 hours Last 7 days 2022-05-16 16:33:35 ~ 2022-05-16 17:33:35 Time Granularity 1 minute Percentile p95

Total Consumption Duration



Consumption Duration Distribution



Click the bar chart on the left to view details

Advanced Metric

Core Services Production Consumption **Instance Resources** Broker GC

1 hour Time granularity: 1 min Disable Show legends


CpuUsage



DiskUsage

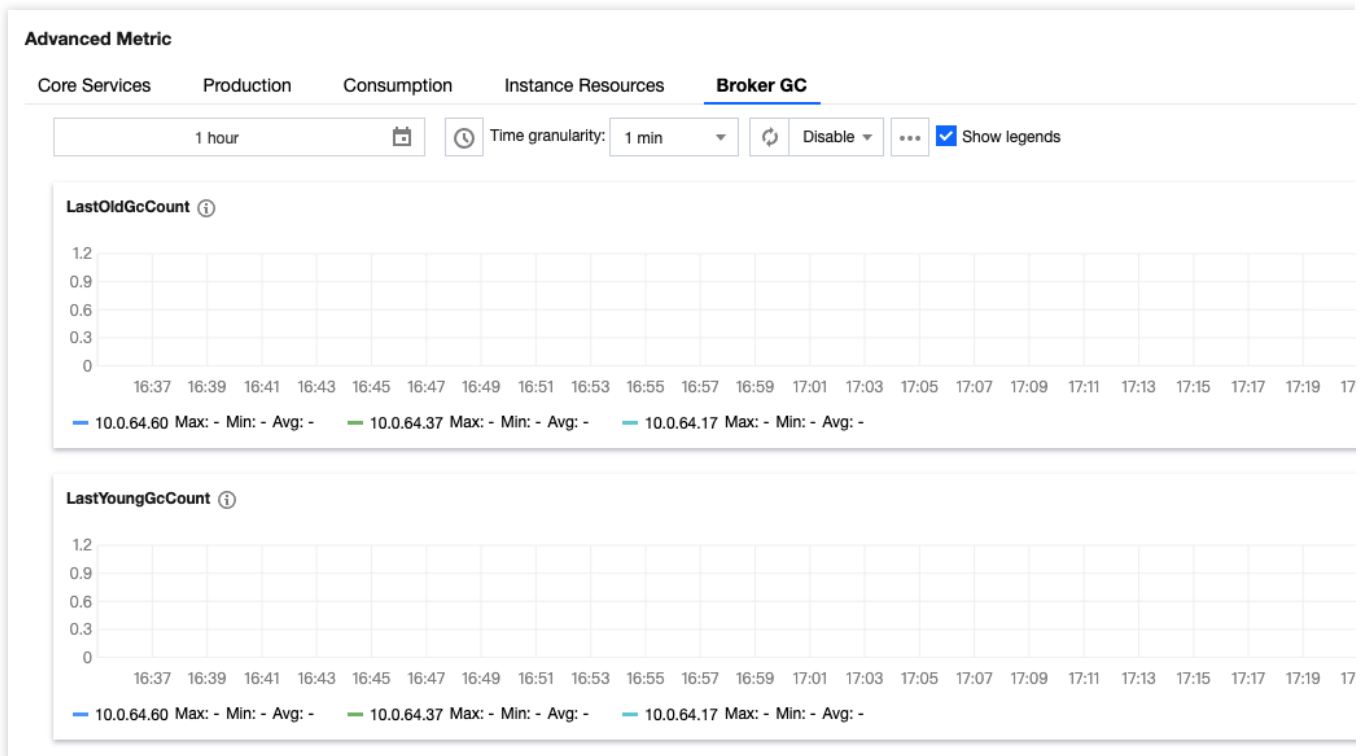


LanIntraffic



LanOuttraffic





Monitoring metric description

Note:

You can click the following tabs to view the detailed descriptions of monitoring metrics of the core service, production, consumption, instance resource, and broker GC.

Core Service Monitoring

Production

Consumption

Instance Resources

Broker GC

Monitoring Metric	Description	Normal Range
Network Status	This value is used to measure the current remaining I/O resources for concurrent request processing. The closer to 1 it is, the idler the instance is.	This value generally fluctuates between 0.5 and 1. If it falls below 0.3, the load is considered high.
Request Queue Depth	This value indicates the number of production requests that have not been processed. If this value is too large, it may be because that the number of concurrent requests is high,	If this value stays at 2000, the cluster load is considered high. If it is below 2000, it can be ignored.

	the CPU load is high, or the disk I/O hits a bottleneck.	
Unsynced Replicas	This value indicates the number of unsynced replicas in the cluster. When there are unsynced replicas of an instance, there may be a health problem with the cluster.	If this value stays above 5 (this is because that some built-in topic partitions of Tencent Cloud may be offline and has nothing to do with the business), the cluster needs to be fixed. In case that the broker occasionally fluctuates, this value may surge and then become stable, which is normal.
ZK Reconnections	This value indicates the number of reconnections of the persistent connection between the broker and ZooKeeper. Network fluctuations and high cluster loads may cause disconnections and reconnections, thus leading to a leader switch. For more information, see FAQs > Client >What is leader switch? .	There is no normal range for this metric. The number of ZK reconnections is cumulative, so a large number does not necessarily mean that there is a problem with the cluster. This metric is for reference only.
ISR Expansions	This value is the number of Kafka ISR expansions. It will increase by 1 when an unsynced replica catches up with the data from the leader and rejoins the ISR.	There is no normal range for this metric. Expansions occur when the cluster fluctuates. No action is required unless this value stays above 0.
ISR Shrinks	This value is the number of Kafka ISR shrinks. It is counted when the broker is down and ZooKeeper reconnects.	There is no normal range for this metric. Shrinks occur when the cluster fluctuates. No action is required unless this value stays above 0.

Monitoring Metric	Description	Normal Range
Total Production Duration	This value indicates the total duration of a production request, which is based on metrics such as the request queue duration, local processing duration, and delayed response duration. At each point in time, the total duration is not equal to the sum of the following five metrics, because each metric is averaged.	This value generally ranges between 0 and 100 ms. It is normal to have a value falling in the range of up to 1000 ms when the data volume is high. No action is required unless this value stays above 1000 ms.
Request Queue	This value indicates the amount of time a	This value generally ranges

Duration	production request waits in the queue of requests to be received. It means that the request packet waits for subsequent processing.	between 0 and 50 ms. It is normal to have a value falling in the range of up to 200 ms when the data volume is high. No action is required unless this value stays above 200 ms.
Local Processing Duration	This value indicates the amount of time a production request is processed by the leader broker, i.e., the duration between the request packet is obtained from the request queue and it is written to the local page cache.	This value generally ranges between 0 and 50 ms. It is normal to have a value falling in the range of up to 200 ms when the data volume is high. No action is required unless this value stays above 200 ms.
ACK Wait Time	This value indicates the amount of time a production request waits for data to be synced. It is greater than 0 only when the client ack is -1; in other words, it is 0 as long as ack is 1 or 0.	This value generally ranges between 0 and 200 ms. It is normal to have a value falling in the range of up to 500 ms when the data volume is high. No action is required unless this value stays above 500 ms. This value for a multi-AZ instance is greater than that for a single-AZ instance when ack is -1. For more information, see Multi-AZ Deployment .
Response Delay	This value indicates the amount of time it takes the system to delay returning a packet to a production request. This value will always be 0 as long as the traffic of the instance does not exceed the purchased traffic, and it will be greater than 0 if the traffic is throttled.	This value will be 0 as long as the instance does not exceed the limit. If the limit is exceeded, there will be a delay of 0–5 minutes proportional to the excess; in other words, the maximum value is 5 minutes.
Response Queue Duration	This value indicates the amount of time a production request waits in the response queue. It means that the request packet waits to be sent to the client.	This value generally ranges between 0 and 50 ms. It is normal to have a value falling in the range of up to 200 ms when the data volume is high. No action is required unless this value stays above 200 ms.

Monitoring Metric	Description	Normal Range
-------------------	-------------	--------------

<p>Total Consumption Duration</p>	<p>This value indicates the total duration of a consumption request, which is based on metrics such as the request queue duration and local processing duration. At each point in time, the total duration is not equal to the sum of the following five metrics, because each metric is averaged.</p>	<p>This value generally ranges between 500 and 1000 ms (the default <code>fetch.max.wait.ms</code> on the client is 500 ms). It is normal to have a value falling in the range of up to 5000 ms when the data volume is high.</p>
<p>Request Queue Duration</p>	<p>This value indicates the amount of time a consumption request waits in the request queue. It means that the request packet waits for subsequent processing.</p>	<p>This value generally ranges between 0 and 50 ms. It is normal to have a value falling in the range of up to 200 ms when the data volume is high. No action is required unless this value stays above 200 ms.</p>
<p>Local Processing Duration</p>	<p>This value indicates the amount of time it takes a consumption request to pull data from the leader broker, i.e., reading data from the local disk.</p>	<p>This value generally ranges between 0 and 500 ms. It is normal to have a value falling in the range of up to 1000 ms when the data volume is high. No action is required unless the value stays above 1000 ms, because the consumer sometimes may read cold data, which will consume a lot of time.</p>
<p>Consumption Wait Duration</p>	<p>The default <code>fetch.max.wait.ms</code> on the client is 500 ms. This value indicates the amount of time the client allows the server to wait before returning any packet to the client when the client cannot read any data.</p>	<p>This value is generally around 500 ms (the default <code>fetch.max.wait.ms</code> on the client is 500 ms), subject to the client's parameter settings.</p>
<p>Response Delay</p>	<p>This value indicates the amount of time it takes the system to delay returning a packet to a consumption request. This value will always be 0 as long as the traffic of the instance does not exceed the purchased traffic, and it will be greater than 0 if the traffic is throttled.</p>	<p>This value will be 0 as long as the instance does not exceed the limit. If the limit is exceeded, there will be a delay of 0–5 minutes proportional to the excess; in other words, the maximum value is 5 minutes.</p>
<p>Response Queue Duration</p>	<p>This value indicates the amount of time a consumption request waits in the response queue. It means that the request packet waits to be sent to the client.</p>	<p>This value generally ranges between 0 and 50 ms. It is normal to have a value falling in the range</p>

		of up to 200 ms when the data volume is high. No action is required unless this value stays above 200 ms.
--	--	--

Monitoring Metric	Description	Normal Range
CPU Utilization (%)	This is the percentage of CPU time used by a process in a period of time to the total CPU time.	This value generally falls between 1 and 100. If it is above 90 in more than 5 consecutive statistical periods, the system load is very high and needs to be handled.
Disk Utilization (%)	This is the usage of a disk mounted to a CVM instance.	This value generally falls between 0 and 100. If it exceeds 80, capacity expansion is required.
Private Network Inbound Bandwidth (MB)	This is the bandwidth that a CVM instance can reach for communication in the cluster. It limits the private network bandwidth and packet receiving capabilities according to different specifications.	This value is generally greater than 0 (CVM monitoring in the cluster will generate data). If there is no inbound bandwidth, the CVM service is abnormal or the network is unreachable.
Private Network Outbound Bandwidth (MB)	This is the bandwidth that a CVM instance can reach for communication in the cluster. It limits the private network bandwidth and packet sending capabilities according to different specifications.	This value is generally greater than 0 (CVM monitoring in the cluster will generate data). If there is no outbound bandwidth, the CVM service is abnormal or the network is unreachable.
Memory Utilization (%)	This is the percentage of the total memory space minus the used memory space to the total memory space.	This value generally falls between 1 and 100. If it is above 90, the program uses too much memory and some processes need to be handled.
Public Network Inbound Bandwidth (MB)	This is the bandwidth that CVM can reach in public network communication. The maximum	It will be greater than 0 if there is public network inbound traffic; otherwise, it will be 0.

	public network bandwidth and packet receiving capabilities vary by specification.	
Public Network Outbound Bandwidth (MB)	This is the bandwidth that CVM can reach in public network communication. The maximum public network bandwidth and packet sending capabilities vary by specification.	It will be greater than 0 if there is public network outbound traffic; otherwise, it will be 0.

Monitoring Metric	Description	Normal Range
Young Generation Collections	Young GC count of the broker.	This value generally ranges between 0 and 300. If it stays above 300, the GC parameters need to be adjusted.
Old Generation Collections	Full GC count of the broker.	This value is generally 0. Actions are required if it is greater than 0.

Causes of monitoring metric exceptions

The following describes causes of certain monitoring metric exceptions.

Metric	Exception Cause
CPU Utilization (%)	When you find that it is above 90% in more than 5 consecutive statistical periods, you can first check whether there are message compression and message format conversion. If the client machine has sufficient CPU resources, we recommend that you enable Snappy compression. You can observe the request queue depth at the same time. If this value is too large, it may be because that the request volume is too high, which may also cause a high CPU load.
Unsynced Replicas (Count)	When this value is above 0, there are unsynced replicas in the cluster. This is usually due to broker node exceptions or network issues. You can troubleshoot through the broker logs.
Full GC Count (Count)	An occasional old GC occurrence may be caused by disk I/O or CVM issues. You can check to see whether brokers with the same IP have the same issue. If so, submit a ticket for assistance.

Request Queue
Depth (Count)

If the client's production and consumption time out but the CVM load remains normal, it means the request queue length of the CVM instance has reached the upper limit, which is 500 by default. You can [submit a ticket](#) to adjust it appropriately according to the purchased resource configuration.

Configuring Alarms

Last updated : 2023-08-24 11:03:05

Overview

Tencent Cloud provides the monitoring service for all users by default; therefore, you do not need to manually activate it. Tencent Cloud Observability Platform (TCOP) will start collecting monitoring data only after a Tencent Cloud product is used.

CKafka allows you to monitor the resources created under your account, including instances, topics, and consumer groups, so that you can keep track of the status of your resources in real time. You can configure alarm rules for monitoring metrics. When a monitoring metric reaches the set alarm threshold, TCOP will notify you of exceptions in time via email, SMS, WeChat, phone call, etc.

Directions

Configuring an alarm policy

The created alarm can determine whether an alarm notification should be sent based on the comparison between the monitoring metric and the given threshold in the selected time period. You can take appropriate precautionary or remedial measures in a timely manner when the alarm is triggered because the status of CKafka changes. Properly creating alarm policies can help improve the robustness and reliability of your applications.

Note

Be sure to configure alarms for your instance to prevent exceptions caused by traffic spikes or specification limits.

1. Log in to the [CKafka console](#).
2. In the instance list, click **Configure Alarm Policy** in the **Operation** column to enter the alarm configuration page.
3. On the alarm configuration page, select a policy type and instance, and set the alarm rule and notification template.

Monitoring Type: Select **Tencent Cloud services**.

Policy Type: Select **CKafka**.

Alarm Object: Select the CKafka resource for which to configure the alarm policy.

Trigger Condition: You can select **Select template** or **Configure manually**. The latter is selected by default. For more information on manual configuration, see the description below. For more information on how to create a template, see [Creating a trigger condition template](#).

Note

Metric: For example, if you select 1 minute as the statistical period for the "Disk Utilization" metric, then if the disk utilization exceeds the threshold for N consecutive data points, an alarm will be triggered.

Alarm Frequency: For example, "Alarm once every 30 minutes" means that there will be only one alarm triggered every 30 minutes if a metric exceeds the threshold in several consecutive statistical periods. Another alarm will be triggered only if the metric exceeds the threshold again in the next 30 minutes. For the metrics for which we recommend that you configure an alarm policy, see [Monitoring and Alarm Policies Recommended for CKafka](#).

Notification Template: You can select an existing notification template or create one to set the alarm recipient objects and receiving channels.

4. Click **Complete**.

Note

For more information on alarms, see [Creating Alarm Policy](#).

Creating a trigger condition template

1. On the **Configure Alarm Policy** page, select **Select Template** for **Trigger Condition** and click **Create Trigger Condition Template**.

2. On the template creation page, configure the policy type.

Policy Type: Select **CKafka**.

Apply preset trigger conditions: Select this option and the system recommended alarm policy will be displayed.

3. After confirming that everything is correct, click **Save**.

Create

Template Name

Remarks

Policy Type Apply preset trigger conditions ⓘ

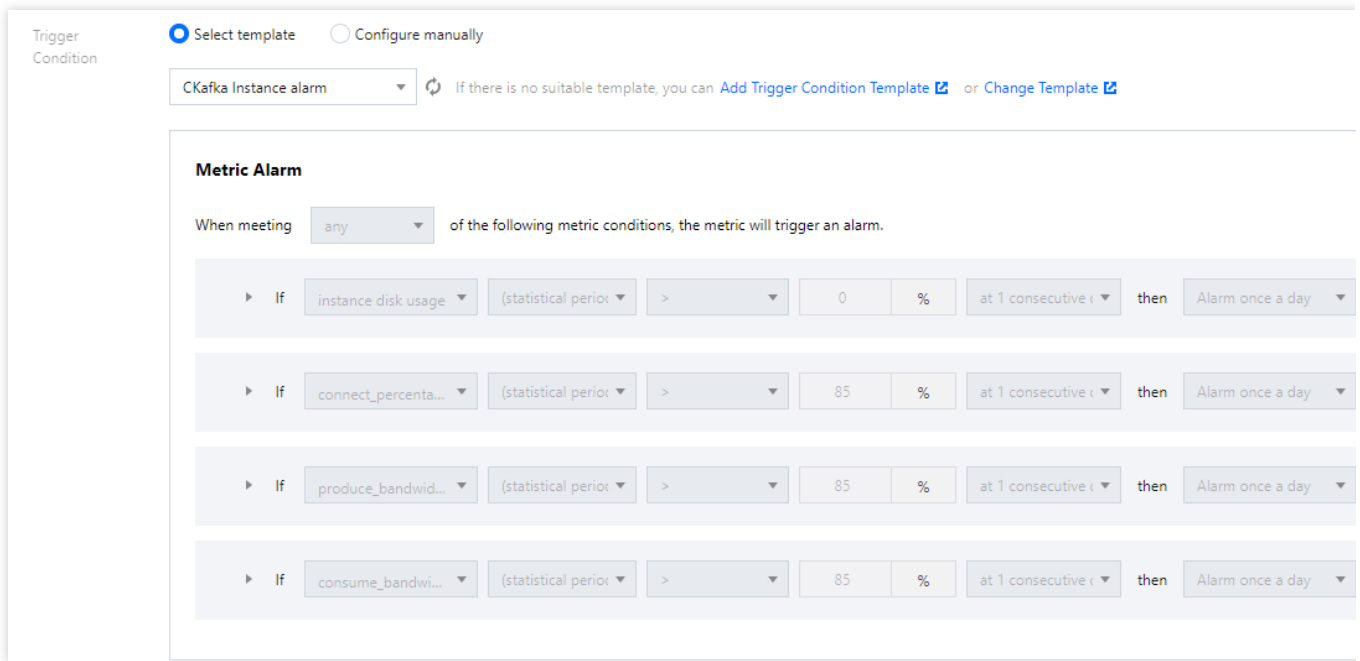
Trigger Condition Metric alarm

Condition met conditions are met.

if	<input type="text" value="instance disk usage"/>	<input type="text" value="Statistical Period"/>	<input type="text" value=">"/>	<input type="text" value="0"/>	<input style="background-color: #ccc;" type="text" value="%"/>	<input type="text" value="Last for 1 per"/>	then	<input type="text" value="Alarm once ever"/>
	ⓘ							
if	<input type="text" value="connect_percentage"/>	<input type="text" value="Statistical Period"/>	<input type="text" value=">"/>	<input type="text" value="85"/>	<input style="background-color: #ccc;" type="text" value="%"/>	<input type="text" value="Last for 1 per"/>	then	<input type="text" value="Alarm once ever"/>
	ⓘ							
if	<input type="text" value="produce_bandwidth_"/>	<input type="text" value="Statistical Period"/>	<input type="text" value=">"/>	<input type="text" value="85"/>	<input style="background-color: #ccc;" type="text" value="%"/>	<input type="text" value="Last for 1 per"/>	then	<input type="text" value="Alarm once ever"/>
	ⓘ							
if	<input type="text" value="consume_bandwidth_"/>	<input type="text" value="Statistical Period"/>	<input type="text" value=">"/>	<input type="text" value="85"/>	<input style="background-color: #ccc;" type="text" value="%"/>	<input type="text" value="Last for 1 per"/>	then	<input type="text" value="Alarm once ever"/>
	ⓘ							

[Add](#)

4. Return to the alarm policy creation page and click **Refresh**. The alarm policy template just configured will be displayed.



Monitoring and Alarm Policies Recommended for CKafka

For more information on metrics that may affect the business data stability, see [CKafka Data Reliability Description](#). Based on user feedback, we recommend that you configure alarm policies in the following 3 dimensions (6 metrics in total) for CKafka. You need to configure them reasonably based on your actual business conditions.

Instance monitoring:

Monitoring Metric	Description
Production Peak Bandwidth (MB/sec)	Peak traffic generated when the instance produces messages (excluding the traffic generated by replicas).
Consumption Peak Bandwidth (MB/sec)	Peak traffic generated when the instance consumes messages (there is no replica concept in consumption).
Disk Utilization (%)	Ratio of the currently used disk capacity to the total disk capacity of the instance in percentages.
Instance Connections (Count)	Number of connections between the client and server.

Topic monitoring:

Monitoring Metric	Description
Used Disk Capacity (MB)	Total size of messages in the topic (excluding those produced by replicas) that actually use disk capacity, which is the latest value in the selected time period.

Consumer group:

Monitoring Metric	Description
Unconsumed Messages	Number of unconsumed messages in the consumer group.

Suggestions for Alarm Configuration

Last updated : 2024-01-09 14:47:32

CKafka not only provides a number of monitoring metrics for running CKafka clusters to monitor their health, but also allows you to configure alarms for key metrics, so that you can identify cluster problems and address them in a timely manner. For more information, see [Viewing Monitoring Information](#) and [Configuring Alarm](#).

This document describes some metrics that require special attention during your use of CKafka, as well as recommended alarm configurations:

Metric	Recommended Alarm Configuration	Description
Disk utilization (%)	The statistical period is 1 minute, and if this value is above 80% in 5 consecutive periods, an alarm will be triggered once every 30 minutes.	The average disk utilization refers to the average of the disk utilization values of all nodes in the cluster. If the disk utilization is too high, the node will not have enough disk space to sustain messages assigned to it, making it unable to store messages. We recommend you clear the data or scale the cluster in time when this value exceeds 75%.
Unconsumed messages (count)	The statistical period is 5 minutes, and if this value is above 8,000 in 10 consecutive periods, an alarm will be triggered once every 30 minutes.	Too many retained messages will cause the broker node disk utilization to soar up, so no more messages can be received, and the service will stop. In this case, you should perform scaling.
Production peak bandwidth (MB/s)	The statistical period is 1 minute, and if this value is above the purchased bandwidth specification of the instance in 5 consecutive periods, an alarm will be triggered once every 10 minutes.	This refers to the maximum value of traffic per second within one minute. You can determine whether the purchased traffic cap is exceeded and upgrade the specification accordingly.

Viewing Advanced Ops Features (Pro Edition)

Last updated : 2024-01-09 14:47:31

Overview

CKafka Pro Edition supports advanced Ops features. You can view the number of TCP connections, unsynced replica details, and topic/consumer group rankings in the console, making it easier for you to troubleshoot CKafka issues.

Directions

1. Log in to the [CKafka console](#).
2. In the instance list, select a region and click the **ID/Name** of the target instance to enter the instance details page.
3. At the top of the instance details page, click **Monitoring > Dashboard** and set the time range to view related ranking information.

TCP Connections: It displays the total number of TCP connections to the current broker. It allows you to view the connections to each server more easily when the number of instance connections is about to be used up.

TCP Connections <small>Total number of TCP connections in the broker: 3</small>	
Source Broker IP	
10.0.0.15	

Total items: 1 20 ▼ / page

Details of Unsynced Replicas: It displays the details of unsynced replicas.

Details of Unsynced Replicas			
Topic ID	Topic Name	Partition Name	Unsynced
No data yet			

Total items: 0 20 ▼ / page

Ranking:

Topic: It displays the top 10 topics in terms of production and consumption traffic as well as disk usage.

Ranking

Topic Consumer Group

[Real-Time](#)
[Last 24 hours](#)
[Last 7 days](#)
[Last 30 days](#)
2022-05-16 16:29:35 ~ 2022-05-16 17:29:35

Topic Traffic (Top 10)
Consumption

Topic ID	Topic Name	Partition C...	Replica Co...	Topic Traffic (MB/min)	Topic I
No data yet					

Top 10 Topics by Disk Usage

Topic ID	Topic Name	Partition C...	Replica Co...	Topic Traffic (MB/min)	Used I
No data yet					

Consumer Group: It displays the top 10 consumer groups in terms of consumption speed.

Ranking

Topic Consumer Group

[Real-Time](#)
[Last 24 hours](#)
[Last 7 days](#)
[Last 30 days](#)
2022-05-16 16:29:35 ~ 2022-05-16 17:29:35

Consumer Group Consumption Speed (Top 10)

Consumer Group Name	Consumpti
No data yet	

Smart Ops

Elastic Scaling

Disk Utilization Policy

Last updated : 2023-08-24 10:29:19

Overview

CKafka supports automatic adjustment of the disk utilization. After the disk utilization reaches the threshold, you can set the **Dynamic Message Retention Policy** to reduce the message retention time or set the **Automatic Disk Capacity Expansion** to adjust the disk space.

Dynamic Message Retention Policy: After you set the message retention time, expired messages will be deleted. In case of message surges, normal production and consumption cannot be performed after the disk is full. If you set a dynamic retention policy, when the disk utilization reaches a certain percentage, a certain proportion of data will be automatically expired to avoid the above issue.

Automatic Disk Capacity Expansion: When the disk load gets heavy, messages cannot be produced and consumed normally. After the automatic disk capacity expansion policy is set, when the disk load reaches the trigger threshold, the disk capacity will be automatically adjusted according to the policy, thus avoiding this problem.

Policy Type	Supported by CKafka Standard Edition	Supported by CKafka Pro Edition
Dynamic message retention	Yes	Yes
Automatic disk capacity expansion	No	Yes

Note

You can enable either dynamic message retention or automatic disk capacity expansion at any time.

Feature Entry

Standard Edition

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and then the **ID/Name** of your instance to enter the instance details page.

3. In the **Disk Utilization Policy** module on the instance's basic information page, you can configure a policy to adjust the disk space.

Pro Edition

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and then the **ID/Name** of your instance to enter the instance details page.
3. Select the **Smart Ops** tab at the top of the page, select the **Auto Scaling** tab, and configure the disk space adjustment policy in the **Disk Utilization Policy** module.

Configuring a Disk Utilization Policy

Dynamic Message Retention Policy

Automatic Disk Capacity Expansion Policy

1. In the **Disk Utilization Policy** module, enable **Dynamic Retention Policy**.

Note

The default dynamic policy reduces the message retention time by 10% when the disk utilization reaches 75%.

Policy Type	Status	Policy Description	Adjustment Record	Oper...
Dynamic Message Retention Policy	<input checked="" type="checkbox"/>	When the disk utilization reaches 90%, the message retention period will be reduced by 10%. Minimum retention period: 6 hr.	View	Configurat
Automatic Disk Capacity Expansion	<input type="checkbox"/>	When the disk utilization reaches 90%, the disk capacity will be automatically expanded by 10%. Max disk capacity: 5000 GB.	View	Configurat

2. Click **Configure** in the **Operation** column of the dynamic retention policy to configure **Dynamic Policy** and **Minimum Retention Period**.

Dynamic Adjustment Record			
Topic	Retention Period (min) Before Adjustment	Retention Period (min) After Adjustment	Adjusted At
test1	1440	-	-
test2	1440	-	-

Total items: 2 10 / page 1 / 1 page

OK Cancel

Dynamic Policy: After message retention time adjustment is triggered, the broker will delete the oldest historical data based on the new retention time. This feature has a certain delay.

Minimum Retention Period: It can be 1 minute to 30 hours. If the dynamic retention period is lower than this parameter, no dynamic adjustment will be triggered.

3. Click **View** in the **Adjustment Record** column to enter the event center and view the adjustment records of automatic disk capacity expansion.

1. In the **Disk Utilization Processing** module, enable **Automatic Disk Capacity Expansion**.

Note

When the disk utilization reaches 75%, the disk capacity will be automatically expanded by 10% by default. The maximum disk capacity is as displayed in the console, which is the upper limit of the disk corresponding to the current instance specification.

Dynamic Adjustment Record			
Topic	Retention Period (min) Before Adjustment	Retention Period (min) After Adjustment	Adjusted At
test1	1440	-	-
test2	1440	-	-

Total items: 2 10 / page 1 / 1 page

OK Cancel

2. Click **Configure** in the **Operation** column of the automatic disk capacity expansion policy to configure **Dynamic Policy** and **Max Disk Capacity**.

Automatic Disk Capacity Expansion Policy ✕

Dynamic Policy When the disk utilization reaches % ,

Automatic Disk Capacity Expansion %

The disk capacity is expanded in increments of 100 GB.

Max Disk Capacity GB

Automatic disk capacity expansion will no longer be triggered after the disk capacity has been expanded to the specified value.

Dynamic Policy: After the disk load reaches the trigger threshold, the disk capacity will be automatically adjusted according to the capacity expansion policy, with a certain delay.

Note

The disk capacity will be expanded in increments of 100 GB.

Max Disk Capacity: Automatic disk capacity expansion will no longer be triggered after the disk capacity has been expanded to the specified value.

3. Click **View** in the **Adjustment Record** column to enter the event center and view the adjustment records of automatic disk capacity expansion.

Dynamic Partition Processing

Last updated : 2023-04-27 17:28:38

Overview

CKafka supports dynamic partition processing. After you enable **Automatic Partition Balancing**, CKafka will automatically check the partition distribution of topics on your specified time and then select a time during off-peak hours to initiate partition balancing.

Note

This feature is only available for CKafka Pro Edition instances.

Directions

1. Log in to the [CKafka console](#).
2. Click **Instance List** on the left sidebar and then the **ID/Name** of your instance to enter the instance details page.
3. Select the **Smart Ops** tab at the top and click **Elastic Scaling**.
4. In the **Dynamic Partition Processing** module, enable **Automatic Partition Balancing**.
5. Click **Configure** in the **Operation** column to set the automatic partition balancing policy.

Custom time: You can customize the time to initiate partition balancing. We recommend that you select a time during off-peak hours to avoid affecting your business.

Auto-selected time: CKafka will select the time to initiate partition balancing during off-peak hours calculated by automatic analysis.

6. Click **View** in the **Adjustment Record** column to enter the event center and view the adjustment records of automatic partition balancing.

Elastic Bandwidth

Last updated : 2024-05-27 15:58:29

Background

TDMQ for Ckafka offers fixed bandwidth specifications, and supports the upgrade and downgrade feature. This facilitates users to make adjustments when there are changes in business. However, in the actual business operation scenarios, we often find unpredictable changes in business traffic. Also, due to potential data migration, the expansion time cannot be controlled, thus causing unexpected impacts on the business.

To address such issues, the Pro Edition of TDMQ for Ckafka introduces the capability of elastic bandwidth. That is, when the business exceeds the cluster's fixed specifications, a certain degree of elastic space will still be provided to ensure the normal operation of the business.

Feature Overview

The Pro Edition of TDMQ for Ckafka supports elastic bandwidth configuration. This is suitable for scenarios where there is occasional sudden traffic on the business. The part exceeding the instance specifications is charged on a pay-as-you-go basis, and the incremental part that can be added has an upper limit. If the business traffic exceeds the instance specifications for most of the time, it is recommended to quickly scale up or down using the existing upgrade and downgrade capability.

Billing Description

The pay-as-you-go method is adopted for elastic bandwidth. Once the billing mode is enabled, no fees will be incurred if elastic bandwidth is not used. The Pro Edition's console offers relevant monitoring metrics to help with observation.

The billing items for elastic bandwidth consist of two parts: elastic production traffic and elastic consumption traffic.

The related billing methods are outlined in the following table:

Billing Items	Calculation Formula	Settlement Method
Elastic Production Fees	Total Elastic Production Traffic (GB) (excluding traffic generated by replicas) x Elastic Pricing (CNY/GB) x Number of Replicas	The time is settled hourly. If the time is less than one hour, it is counted as one hour. Billing is pushed daily
Elastic Consumption	Total Elastic Consumption Traffic (GB) (excluding traffic generated by replicas) x Elastic Pricing	The time is settled hourly. If the time is less than one hour, it is

Fees	(CNY/GB)	counted as one hour. Billing is pushed daily
------	----------	---

Elastic Bandwidth Space

Available space of elastic bandwidth. It is related to Pro Edition cluster specifications, and below is the corresponding elastic interval.

Bandwidth Range (MB/s)	Upward Adjustment Margin (MB/s)
20~100	30
120~300	60
320~600	100
620~900	200
920~1200	600
1,600 and above	700

Use Limits

1. Current elastic bandwidth configuration does not support deactivation once enabled.
2. This feature is in allowlist grayscale testing. To enable, you may [Contact Us](#).

Permission Management

Configuring ACL Policy

Last updated : 2024-01-09 14:49:39

Overview

This document describes how to configure SASL authentication and ACL (access control list) rules in the CKafka console to enhance access control in public/private network transfers and permission control in production and consumption of resources such as topics.

Note:

Kafka offers various security authentication mechanisms, which mainly include SSL and SASL. SASL/PLAIN is a more commonly used authentication method that is based on account and password. CKafka supports SASL_PLAINTEXT authentication. For more information, see [Adding Routing Policy](#).

An ACL helps you define a set of permission rules to allow/deny users to read/write topic resources through IPs.

Directions

Creating a user

1. Log in to the [CKafka console](#).
2. On the topbar, select a region and click the **ID/Name** of the target instance.
3. On the instance details page, click the **ACL Policy Management** tab at the top.
4. On the **User Management** tab, click **Create** and enter the username and password to create a user.

Configuring an ACL policy

1. On the **ACL Policy Management** page, select the **Policy List** tab.
2. Click **Batch Set** to grant permissions to the user.

Note:

If only allow rules are configured, any IPs not in those rules cannot connect to the instance.

If only deny rules are configured, any IPs not in those rules can connect to the instance.

If both allow and deny rules are configured, only IPs in allow rules can connect to the instance.

Instances on v2.4.1 or later

Instances on other versions

You can grant permissions to the user through the **Topics**, **Topic name prefix**, or **Preset Rule** option.

Note:

You can enter multiple IPs or IP ranges and separate them by ; when configuring the ACL policy. If the IP is empty, the permission will be added for **all IPs** by default.

Topics: Select multiple topics that need to be configured with the same ACL policy.

Topic name prefix: Fuzzy match topics that need to be configured with the same ACL policy by topic name prefix.

You need to specify the fuzzy matching rule name. After this is set, when a new topic with a name containing the specified prefix is added, the system will automatically configure the specified ACL policy for it.

Note:

Up to five fuzzy match rules can be set.

Preset Rule: A set of rules can be preset and automatically applied during subsequent topic creation.

Note:

Up to five preset rules can be set.

You can grant permissions to the user through the **Topics** or **Preset Rule** option.

Note:

You can enter multiple IPs or IP ranges and separate them by ; . If the IP is empty, the permission will be added for **all IPs** by default.

Topics: Select multiple topics that need to be configured with the same ACL policy.

Preset Rule: A set of rules can be preset and automatically applied during subsequent topic creation.

Note:

Up to five preset rules can be set.

Add ACL Policy
✕

ℹ️ ACL policy: allow/deny users to read/write to topic resource "" by IP.

Match By Topics Preset Rule

Select Topic **Please select**

ID/Name

topic-5thxt31w
cccc

Selected (0)

ID/Name

↔

ACL Policy

Permissi...	User	IP/IP Range	Policy
Allow ▼	Please select an item (selec ▼	Enter IPs/IP ranges (all by default) and se	write ▼
Add Rule			

Submit
Close

Next steps: After the authorization is completed, the user can access CKafka through the SASL access point and consume messages by using the PLAIN mechanism. For more information, see the [SDK documentation](#).

Use limits

1. Enabling routing only affects the authentication method during access, while the set ACL policy takes effect globally.

2. If you use the PLAINTEXT method to access CKafka while enabling public network access routing, the ACL previously set for the topics will still take effect. If you want PLAINTEXT access to be unaffected, add the read/write permissions of all users for the topics that PLAINTEXT needs to access.

Note:

When adding an ACL policy, you don't need to select any user, and read/write permissions are added to **all users** by default.

The effect after addition is as follows:

Permission Type	User	IP	Policy	Resource Name	Op
Allow	All	All	write	cccc	Del

3. If a topic is already being used by another Tencent Cloud service (e.g., log shipping in CLS, message dump in SCF, and component consumption in EMR), enabling ACL policy is equivalent to imposing restrictions on the permissions of these linked capabilities, and they will directly become unavailable. Therefore, be sure to do so with caution. In such cases, we recommend that you produce the same data to another topic for separate processing instead of configuring a unified ACL policy on the same topic.

Viewing a preset rule

1. On the ACL policy management page, select **Preset Rule**.
2. In the preset rule list, click **Details** in the **Operation** column to view the details of a rule.

Deleting a preset rule

1. On the ACL policy management page, select **Preset Rule**.
2. In the preset rule list, click **Delete** in the **Operation** column to delete a rule.

The impact of deleting the preset rule varies by the type of rule match:

If the rule is a fuzzy match rule, it will no longer be automatically applied to new topics or take effect for topics to which it is already applied.

If the rule is not a fuzzy match rule, it will no longer be automatically applied to new topics but will still take effect for topics to which it is already applied.

CAM

Last updated : 2024-01-09 14:49:39

Basic CAM Concepts

The root account authorizes sub-accounts by associating policies. The policy setting can be specific to the level of **[API, Resource, User/User Group, Allow/Deny, and Condition]**.

Account

Root account: It owns all Tencent Cloud resources and can access any of its resources.

Sub-account: It includes sub-users and collaborators.

Sub-user: It is created and fully owned by a root account.

Collaborator: It has the identity of a root account. After it is added as a collaborator of the current root account, it becomes one of the sub-accounts of the current root account and can switch back to its root account identity.

Identity credential: It includes login credentials and access certificates. **Login credential** refers to a user's login name and password. **Access certificate** refers to Tencent Cloud API keys (`SecretId` and `SecretKey`).

Resource and permission

Resource: An object that is operated in Tencent Cloud services, such as a CVM instance, a COS bucket, or a VPC instance.

Permission: It is an authorization that allows or forbids users to perform certain operations. By default, **a root account has full access to all resources under it**, while **a sub-account does not have access to any resources under its root account**.

Policy: It is a syntax rule that defines and describes one or more permissions. The **root account** performs authorization by **associating policies** with users/user groups.

[View CAM documentation >>](#)


Relevant Documents

Content	Document
Understand the relationship between policies and users	Concepts
Understand the basic structure of policies	Element Reference
Check CAM-enabled products	Overview

Sample CAM Policy

Full access policy for CKafka

Grant a sub-user full access (including resource creation and management) to the CKafka service.



```
{  
  "version": "2.0",  
  "statement": [  
    {  
      "action": [  

```


```
        "name/ckafka:*",
        "name/monitor:GetMonitorData"
    ],
    "resource": "*",
    "effect": "allow"
}
]
```

You can also configure the system's [full read/write policy](#) to support this permission.

1. Log in to the [CAM console](#).
2. Click [Policies](#) on the left sidebar.
3. In the policy list, click **Create Custom Policy**.
4. In the **Select Policy Creation Method** pop-up window, select **Create by Policy Syntax**.
5. In **Template Type**, search for **CKafka**, select **QcloudCKafkaFullAccess** (full access to CKafka), and click **Next**.
6. Click **Complete**.

Read-only policy for a CKafka instance

1. Create a policy with the Policy Generator and grant permission for listing operations and product monitoring.



```
{
  "version": "2.0",
  "statement": [
    {
      "effect": "allow",
      "action": [
        "name/ckafka:ListInstance",
        "name/monitor:GetMonitorData"
      ],
      "resource": [
        "*"
      ]
    }
  ]
}
```




```
    ]  
  }  
]  
}
```

2. Grant read-only access to the specified instance.

Note:

List* APIs don't support authentication at the resource level.



```
{  
  "version": "2.0",
```

```
"statement": [  
  {  
    "effect": "allow",  
    "action": [  
      "name/monitor:GetMonitorData",  
      "name/ckafka:Get*"  
    ],  
    "resource": [  
      "qcs::ckafka:gz::ckafkaId/uin/$createUin/$instanceId"  
    ]  
  }  
]  
}
```

You can also configure the system's [read-only policy](#) to support this permission.

1. Log in to the [CAM console](#).
2. Click [Policies](#) on the left sidebar.
3. In the policy list, click **Create Custom Policy**.
4. In the **Select Policy Creation Method** pop-up window, select **Create by Policy Syntax**.
5. In **Template Type**, search for **CKafka**, select **QcloudCkafkaReadOnlyAccess** (read-only access to CKafka), and click **Next**.
6. Click **Complete**.

Tag Management

Tag Overview

Last updated : 2024-01-09 14:49:39

Overview

Tag is a key-value pair provided by Tencent Cloud which is used to identify resources on the cloud.

You can use tags to classify CKafka resources based on various factors such as service, usage, and person in charge. With tags, you can quickly filter out the corresponding resources. The key-value pair of the tag does not mean anything to Tencent Cloud semantically and will be parsed and matched strictly according to the string.

Use Limits

Quantity

Each cloud resource can have up to 50 tags.

Tag key restrictions

Tag keys cannot start with `qcloud`, `tencent`, or `project` because they are reserved by the system.

Tag keys can only start with letters, figures, spaces, or Chinese characters, supporting `+`, `-`, `=`, `.`, `_`, `:`, `/`, and `@`.

The maximum length of a tag key is 255 characters.

Tag value restrictions

Tag values can only start with letters, figures, spaces, or Chinese characters, supporting `+`, `-`, `=` value, `.`, `_`, `:`, `/`, and `@`.

The maximum length of a tag value is 127 characters.

Examples

Case background

A company has 10 CKafka instances on Tencent Cloud, which are owned by three departments: Ecommerce, Game, and Entertainment. The instances are used for services such as marketing, game A, game B, and post-production.

The OPS owners of the three departments are John, Jane, and Harry, respectively.

Setting tags

To facilitate management, the company categorizes its CKafka resources with tags and defines the following tag key-value pairs.

Tag Key	Tag Value
Department	Ecommerce, Game, and Entertainment
Business	Marketing, Game A, Game B, and Post-production
OPS Owner	John, Jane, and Harry

Bind the tag key-value pairs to CKafka. The relation between resources and the tag key-value pairs is shown in the table below.

ID	Department	Service	OPS Owner
ckafka-1jqwv1	Ecommerce	Marketing	Harry
ckafka-1jqwv12	Ecommerce	Marketing	Harry
ckafka-1jqwv13	Game	Game A	John
ckafka-1jqwv13	Game	Game B	John
ckafka-1jqwv14	Game	Game B	John
ckafka-1jqwv15	Game	Game B	Jane
ckafka-1jqwv16	Game	Game B	Jane
ckafka-1jqwv17	Game	Game B	Jane
ckafka-1jqwv18	Entertainment	Post-production	Harry
ckafka-1jqwv19	Entertainment	Post-production	Harry
ckafka-1jqwv110	Entertainment	Post-production	Harry

Using tags

Filter out the CKafka instances that Harry is in charge of:

Filter out the CKafka resources OPS owner Harry is in charge of according to the filtering rules. For specific filtering steps, please see [Using Tags](#).

Filter out the CKafka instances Jane of the game department is responsible for:

Filter out the CKafka instances where the department is "Game" and the OPS owner is "Jane". For detailed directions, please see [Using Tags](#).

Using Tags

Last updated : 2024-01-09 14:49:39

Overview

This document describes how to filter out the corresponding instance resources by tag in the CKafka Console.

Directions

1. Log in to the [CKafka Console](#).
2. Select the region at the top of the instance list page.
3. In the search box in the top-right corner of the instance list, click the blank space and select **Tag**.
4. Enter the tag key in **Tag**.

For example, to filter out instances bound to a specified tag key (such as `name`), enter `name` (case sensitive).

5. Press **Enter** or click **Search** icon to search.

Editing Tags

Last updated : 2024-01-09 14:49:39

Overview

This document describes how to edit the tags of target instances in the CKafka Console.

Directions

Note:

You can edit tags for single or multiple instances, with up to 20 instances for tag batch editing.

1. Log in to the [CKafka Console] (<https://console.tencentcloud.com/ckafka>).
2. On the CKafka instance list page, select the instance whose tag needs to be edited, and click **Edit Tag** at the top of the list.
3. Add, modify, or delete tags in the tag editing operation box.
4. Click **OK** to complete editing.

Querying Message

Last updated : 2024-01-09 14:49:39

Overview

When encountering a consumption exception, you can troubleshoot the problem by querying the message in the CKafka console. This document describes how to query messages in the CKafka console.

The CKafka console supports query by offset and time, which are used in different scenarios:

Query by offset: You know the ID of a topic's partition to which the message is sent and the offset of the message.

Query by time: You don't know the offset of the message but know when the message is sent.

Note:

The console does not list all the messages after the specified time or offset. You can query at most 20 messages at a time.

Refrain from frequently querying as message querying also consumes bandwidth.

Directions

Query by offset

1. Log in to the [CKafka console](#).
2. Select **Instance List** on the left sidebar, select a region, and click the **ID/Name** of the target instance to enter the instance details page.
3. On the instance details page, select **Topic Management** and click **More> Message Query** in the **Operation** column on the right of the target topic.
4. Select **Query by offset**, enter the partition ID and starting offset, and click ****Query*** to view the message details.

i Message query consumes the bandwidth resources of CKafka instances. Please narrow down the query range and do not query frequently.
The query results display up to 20 data entries starting from the specified offset or time point

Instance

Topic

Query Type Query by offset Query by start time

Partition ID

Start Offset

Query

Partition ID	Offset	Timestamp	Operation
! Not found message(ckafka[#FailedOperation]) Retry			

Partition ID: Topic partition to which the message is sent.

Offset: Consumer offset of the message.

Timestamp: `timestamp` in `ProducerRecord` .

Operation:

Click **Download Message** to download the message in the format of "header, key, value".

Click **Message Details** to view the key and value of the message.

Message Details

i The currently queried message has been force converted to String type. If garbled characters appear, please analyze the serialization format and encoding format of your message.

Key	No data yet
Value	hello world

OK

Query by time

1. Log in to the [CKafka console](#).
 2. Select **Instance List** on the left sidebar, select a region, and click the **ID/Name** of the target instance to enter the instance details page.
 3. On the instance details page, select **Topic Management** and click **More> Message Query** in the **Operation** column on the right of the target topic.
 4. Select **Query by start time**, enter the partition ID and time, and click **Query** to view the message details.
- If you select **All partitions**, messages of all partitions will be displayed; if you specify a partition, messages of the specified partition will be displayed.

Instance

Topic

Query Type

Partition ID

Time

Partition ID	Offset	Timestamp	Operation
0	0	2021-03-02 11:32:15	View Message D
0	1	2021-03-02 11:33:13	View Message D

Partition ID: Topic partition to which the message is sent.

Offset: Consumer offset of the message.


Timestamp: `timestamp` in `ProducerRecord`.

Operation:

Click **Download Message** to download the message, including its headers, key, and value.

Click **View Details** to view the details of the message.

Message Details

 The currently queried message has been force converted to String type. If garbled characters appear, please analyze the serialization format and encoding format of your message.

Key

No data yet

Value

hello world

OK

Event Center

Last updated : 2024-01-09 14:49:39

The event center manages, stores, analyzes, and displays the event data generated by CKafka in a unified manner.

You can easily view the event data details there.

The events currently supported by CKafka include:

Instance events

Automatic disk capacity expansion

Dynamic disk message retention policy

Partition balancing

Kafka version upgrade

Kernel minor version upgrade

AZ change

Diagnosis events

Scheduled instance diagnosis

Instant instance diagnosis

Broker changes

Consumer group member heartbeat timeout

Consumer group member update

Consumer group rebalancing

Leader switch

Cluster node launch

Cluster node removal

This document describes how to view event details and configure event notification rules in the CKafka console.

Entering the event center

1. Log in to the [CKafka console](#).
2. Click **Instance Management** on the left sidebar and click the **ID** of the target instance to enter the instance details page.
3. Select the **Event Center** tab at the top of the page, set the time range (last 6 hours, last 7 days, last 30 days, or a custom time range), and select the target event type.

Basic Info Topic Management Consumer Group Monitoring **Event Center** Access over HTTP ACL Policy Management

Last 7 days Last 30 days 2023-04-07 17:18:31 ~ 2023-04-14 17:18:31 Instance Event / Dynamic disk message retention policy

Event Name	Status	Event Type	Affected Object	Last Updated	Operation
No data yet					

4. On the **Event Details** page, you can perform the following operations:

Click **View Details** in the **Operation** column to view the event details on the right.

Click **Configure Alarm Policy** in the **Operation** column to configure a notification template for the event and send event details to the target object.

Migration to Cloud

Migration Scheme Overview

Last updated : 2024-01-09 14:49:39

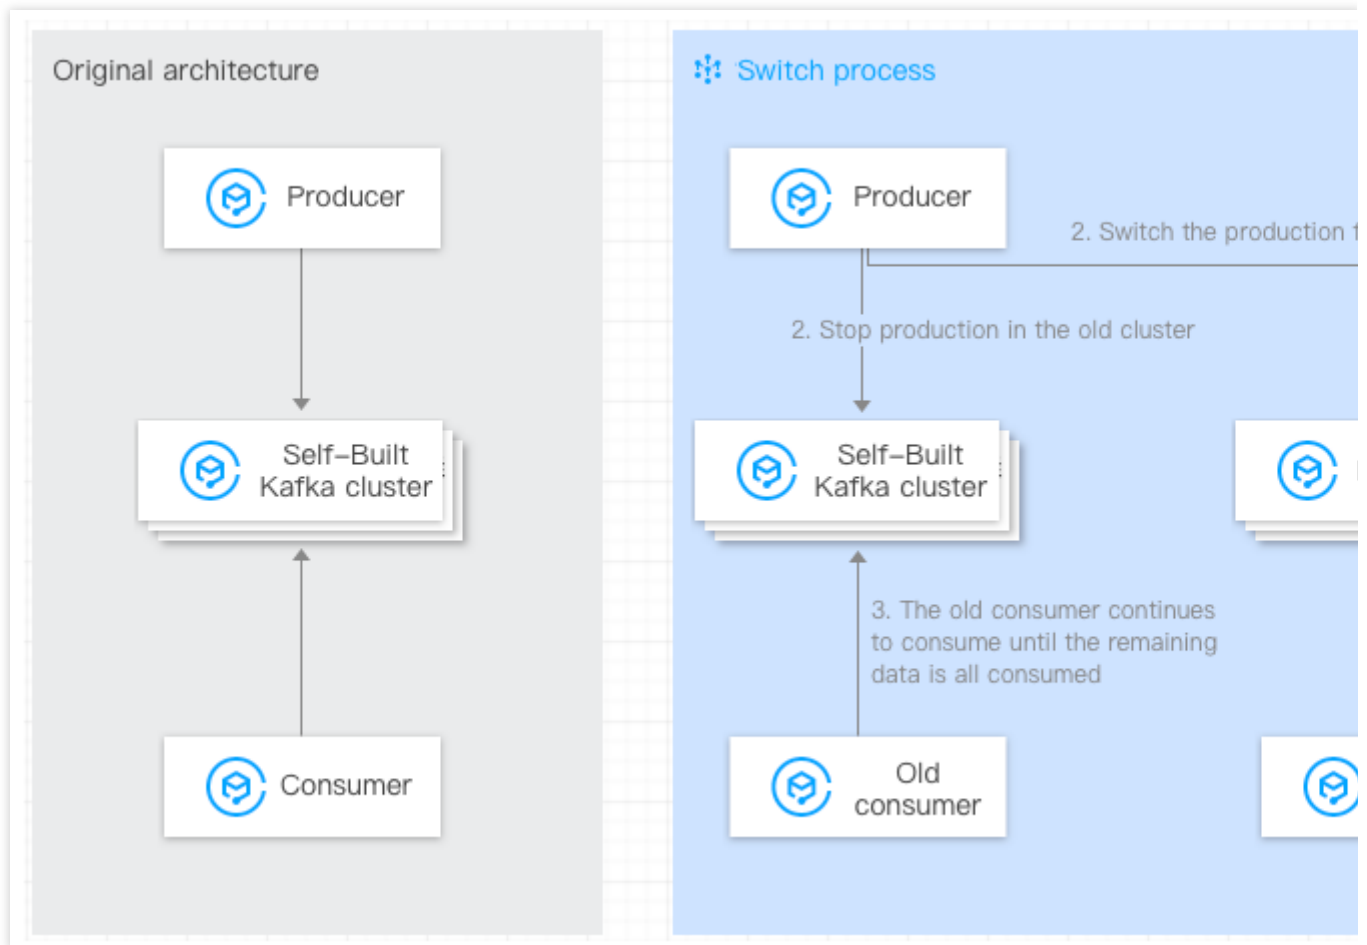
Overview

This document describes the schemes for migrating a self-built Kafka cluster to a CKafka cluster. You can choose an appropriate scheme according to your own business scenarios.

Scheme Description

Scheme 1. Single-Producer dual-consumer

This scheme is simple, clear, and easy to implement, with no data heap for smooth transition.



Scheme ideas:

1. [Migrate the topic metadata.](#)
2. Keep the original consumer in the self-built Kafka cluster intact.
3. Start a new consumer in CKafka and configure the bootstrap-server of the new CKafka cluster for consumption.
4. Wait for all consumers to listen on the new CKafka cluster.
5. Switch the production of the self-built cluster to the new CKafka cluster (by configuring the bootstrap-server of the new CKafka cluster).
6. Deactivate the original consumer in the self-built Kafka cluster after it finishes consuming the remaining data in the cluster.


Pros and cons of the scheme:

Pros: the overall migration process is simple, clear, and easy to implement, with no data heap for smooth transition.

Cons: an additional consumer is required.

Scheme 2. Single-Producer single-consumer

This scheme is simple, clear, and easy to implement.



Scheme ideas:

1. [Migrate the topic metadata.](#)
2. Switch the production of the self-built Kafka cluster to the new CKafka cluster (by configuring the bootstrap-server of the new CKafka cluster).
3. Wait for the consumer in the self-built cluster to consume the remaining data.
4. Switch the old consumer to the new CKafka cluster for consumption (by configuring the bootstrap-server of the new CKafka cluster).

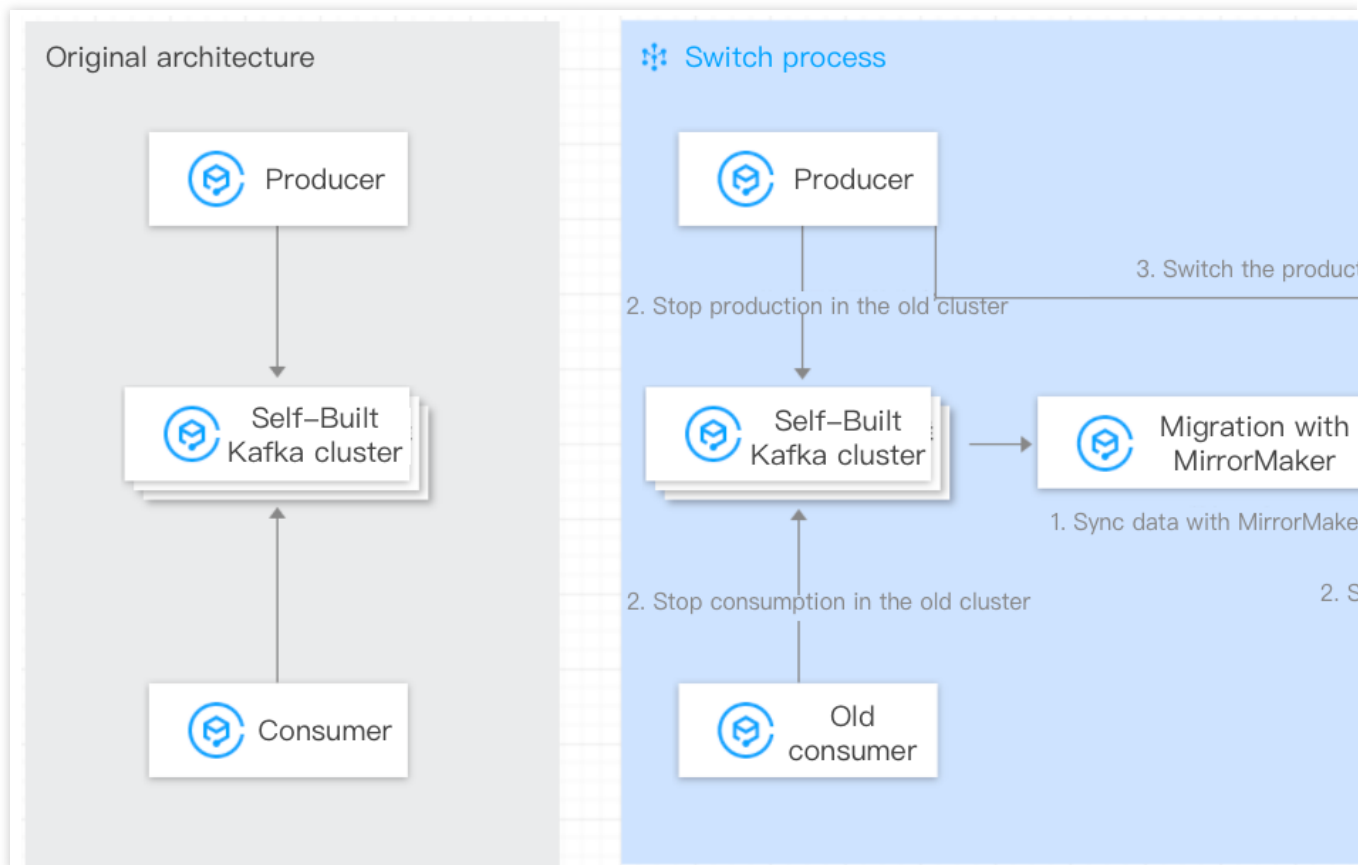
Pros and cons of the scheme:

Pros: the overall migration process is simple, clear, and easy to implement for smooth transition.

Cons: after the production is switched to the CKafka cluster, before the old consumer is switched to the CKafka cluster, there will be a certain amount of data heap in the CKafka cluster.

Scheme 3. Migration with MirrorMaker

This scheme migrates the existing data in the self-built Kafka cluster to CKafka.



Scheme ideas:

1. [Migrate the topic metadata.](#)
2. Keep the original consumer in the self-built Kafka cluster intact.
3. Start the data sync feature of [MirrorMaker](#).
4. Wait for the data sync to complete, modify the consumer configuration, and switch the consumer.
5. Wait for the data sync to complete, modify the producer configuration, and switch the producer.
6. The migration is completed.

Pros and cons of the scheme:

Pros: the overall migration process is simple, clear, and easy to implement, where historical data can be synced to the CKafka cluster.

Cons: the consumer needs to consume data from the beginning after switch to the destination cluster, so consumption idempotency must be ensured.

Step 1. Purchasing CKafka Instance

Last updated : 2024-01-09 14:49:39

Overview

This document describes how to use the **Specification Calculator** in the CKafka console to find the appropriate CKafka instance specification for your self-built Kafka cluster that needs to be migrated to the cloud.

Directions

1. Log in to the [CKafka console](#).
2. On the left sidebar, select **Migration to Cloud**, select the region you want to migrate to, and click **Specification Calculator**.
3. Enter the specification of your self-built Kafka cluster on the specification calculator page.

Specification Calculator

- 1 **On-Premises Kafka Specification**
- >
- 2 **Recommended Specification**

Kafka Version

0.10.x

[Suggestions for CKafka Version Selection](#)

Select 0.10.x if the CKafka version is v0.10 or earlier

Peak Application Bandwidth

40

MB/s

Peak application bandwidth = max(peak production bandwidth * the number of replicas, peak consumption bandwidth)

Disk

300

GB

Estimate this value based on your disk's current message heap peak

Total Partitions

60

Count

This value indicates the total number of partitions to migrate and is proportional to the number of replicas. For example, a single-replica topic has five partitions while a dual-replica topic has ten. CKafka does not support single-replica topics.

Multi-AZ Deployment

Yes No

Data Compression

Disabled

CKafka does not support gzip format. For details, see [Data Compression](#)

Next

Parameter	Description
Kafka Version	Select the version of your self-built Kafka cluster. For more information on CKafka version selection, please see Suggestions for CKafka Version Selection .
Peak Business Bandwidth	Peak business bandwidth = max (production peak bandwidth * number of replicas, consumption peak bandwidth).
Disk	Estimate according to the current peak value of actual disk usage.
Total Partitions	Total number of partitions of the topic to be migrated. The number of replicas should be considered. For example, if one topic replica has 5 partitions, then two replicas have 10

	partitions in total. CKafka does not support single-replica topics.
Multi-AZ Deployment	Select whether to deploy in multiple AZs according to your business needs. For more information, see Multi-AZ Deployment .
Data Compression	CKafka does not support the Gzip compression format. For more information, see Data Compression .

4. Click **Next** to get the recommended CKafka instance specification.
5. Click **Buy Now** to redirect to the instance purchase page.
6. Confirm the purchase information, click **Buy Now**, wait 5–10 minutes, and you can see that the instance is created on the instance list page.

Create Edit Tag Delete Terminate Please enter key								
ID/Name	Monitor	Status	AZ	Instance Type	Configuration	Network Type	Billing Mode	
<input type="checkbox"/> ckafka-l-5		Healthy	Changsha Zone 1	Pro Edition-Basic Type Version: 2.8.1 Disk Type: SSD	Topic Limit: 400 Partition Limit: 800 Peak Bandwidth: 40 MB/s Disk Capacity: 500GB	VPC abc abc	Monthly subsc Expire on 202	

Step 2. Migrating Topic to Cloud

Last updated : 2024-01-09 14:49:40

Overview

This document describes how to use the migration tool provided by CKafka to migrate topics from a self-built Kafka cluster to a CKafka instance.

Prerequisites

You have [purchased a CKafka instance](#).

You have [downloaded Python 2.x](#).


Directions

1. Download the [migration tool](#) and decompress it on a server that can connect to the broker and ZooKeeper of the self-built instance.
2. Enter the configuration parameters in the `ckafka-migrate.py` file.

Note:

Make sure that the networks of the server where the subsequent migration operation is performed, the CKafka instance, and the self-built Kafka cluster are interconnected.

The user corresponding to the key of TencentCloud API needs to have the write permission for CKafka. We recommend that you use the key pair of the root account.



```
# your local broker ip:port list
# Broker list of the self-built instance ["broker1:port1","ip2:port2"]
bootstrapServers = ["$ip:$port"]

# your local zk ip:port list
# ZooKeeper list of the self-built instance ["zk1:port1","zk2:port2"]. This is an o
sourceZk = ["$ip:$port"]
# This parameter is required if ZooKeeper has enabled authentication. Format: [("di
zkAuthData = []
# This parameter is required if the ZooKeeper path has a prefix. Sample value: "/cl
zkPathPrefix = ""
```

```

# your cloud instanceId
# CKafka instance ID "ckafka-xxx"
instanceId = "$yourinstanceId"
# topic regex, just migrate match topics
# Regular expression of the topic name. If it is not empty, only matched topics will
topicRegex = ""

# your secretId and secretKey
# Account key pair
secretId = "$yoursecretId"
secretKey = "$yoursecretKey"

# your cloud instance region
# CKafka instance region and code:
# Guangzhou: ap-guangzhou; Shanghai: ap-shanghai; Nanjing: ap-nanjing; Beijing: ap-
# Hong Kong (China): ap-hongkong; Singapore: ap-singapore; Mumbai: ap-mumbai; Tokyo
# Virginia: na-ashburn; Toronto: na-toronto; Taipei (China): ap-taipei; Tianjin: ap-
# Shenzhen Finance: ap-shenzhen-fsi; Shenzhen: ap-shenzhen; Frankfurt: eu-frankfurt
# Beijing Finance: ap-beijing-fsi; Bangkok: ap-bangkok; Changsha: ap-changsha-ec; J
# Hangzhou ec: ap-hangzhou-ec; São Paulo: sa-saopaulo; Shanghai Auto-Driving Cloud:
region = "ap-tokyo"

# If you make sure the migrate topic List, please modify checkFlag = 1
# Check flag. If it is 0, only the list of topics to be migrated will be displayed,
# 0: List the topics to be migrated and then stop the script
# 1: List the topics to be migrated and then start migration
checkFlag = 0

# force transfer your cloud-topic config to migrate
# If the value is 0, a local topic will not be migrated to CKafka if its attributes
# For example, if the CKafka topic supports only 1, 2 or 3 replicas, but a local to
# 0: Skip the unmatched topic or topic attributes if the numbers of replicas or att
# 1: Forcibly migrate all topics to CKafka if the numbers of replicas or attributes
force = 0

```

Parameter	Description
bootstrapServers	Broker list of the self-built instance ["ip1:port1","ip2:port2"]
sourceZk	ZooKeeper list of the self-built instance ["zk1:port1","zk2:port2"]. This is an optional configuration. If it is not configured, the source cluster information will be obtained via <code>bootstrapServers</code> .
zkAuthData	This parameter is required if ZooKeeper has enabled authentication. Format: [("digest", "\$user:\$password")]

zkPathPrefix	This parameter is required if the ZooKeeper path has a prefix. Sample value: "/cluster1".
instanceId	ID of the CKafka instance you purchased in Step 1. Purchasing CKafka Instance . You can copy the ID on the Instance List page in the CKafka console.
secretId	Account key pair – ID
secretKey	Account key pair – password
region	The deployment region you selected in Step 1. Purchasing CKafka Instance . The region codes are available in the script annotation.
checkFlag	Check flag. If it is 0, only the list of topics to be migrated will be displayed and migration will not start; otherwise, the topics will start to be migrated.
topicRegex	Regular expression of the topic name. If it is empty, all topics will be migrated; otherwise, only matched topics will be migrated.
force	It specifies whether to forcibly migrate topics. If the value is 0, a local topic will not be migrated to CKafka if its attributes do not match those of CKafka. If the value is 1, the topic attribute values out of the value range of CKafka will be forcibly converted to the closest values within the value range.

3. Set the `checkFlag` parameter in `ckafka-migrate.py` to 0, run `python ckafka-migrate.py` to execute the script, and check the list of topics that need to be migrated based on the output result.

Note:

If some self-built topics are missing, this may be because their names are invalid, or the numbers of replicas or attribute values of the self-built and CKafka topics do not match.

```
[root@VM-0-21-centos ~/migrateCloud/migrateToCkafkaTool]# python ckafka-migrate.py
check topic List to migrate: ["test1", "test0", "test108", "test109", "test5", "test4", "test7", "test6", "test102", "test103", "test100",
"test101", "test106", "test107", "test104", "test105", "test182", "test183", "test180", "test181", "test186", "test187", "test184", "test
185", "test188", "test189", "test46", "test47", "test44", "test45", "test42", "test43", "test40", "test41", "test48", "test49", "test8", "
test179", "test178", "test177", "test176", "test175", "test174", "test173", "test172", "test171", "test170", "test55", "test54", "test57",
"test56", "test51", "test50", "test53", "test52", "test59", "test58", "test168", "test169", "test164", "test165", "test166", "test167", "
test160", "test161", "test162", "test163", "test60", "test61", "test62", "test63", "test64", "test65", "test66", "test67", "test68", "tes
t69", "test2", "test159", "test158", "test151", "test150", "test153", "test152", "test155", "test154", "test157", "test156", "test79", "tes
t78", "test77", "test76", "test75", "test74", "test73", "test72", "test71", "test70", "test146", "test147", "test144", "test145", "test142
", "test143", "test140", "test141", "test148", "test149", "test82", "test83", "test80", "test81", "test86", "test87", "test84", "test85",
"test88", "test89", "test3", "test133", "test132", "test131", "test130", "test137", "test136", "test135", "test134", "test139", "test138",
"test91", "test90", "test93", "test92", "test95", "test94", "test97", "test96", "test99", "test98", "test35", "test19", "test18", "test11
", "test10", "test13", "test12", "test15", "test14", "test17", "test16", "test120", "test121", "test122", "test123", "test124", "test125",
"test126", "test127", "test128", "test129", "test24", "test25", "test26", "test27", "test20", "test21", "test22", "test23", "test28", "te
st29", "test", "test119", "test118", "test115", "test114", "test117", "test116", "test111", "test110", "test113", "test112", "test195", "te
st194", "test197", "test196", "test191", "test190", "test193", "test192", "test199", "test198", "test33", "test32", "test31", "test30", "
test37", "test36", "test9", "test34", "test39", "test38"]
```

4. Set the `checkFlag` parameter in `ckafka-migrate.py` to 1 and run `python ckafka-migrate.py` to execute the script to start migrating the topics.

```
[root@VM-0-21-centos ~/migrateCloud/migrateToCkafkaTool]# python ckafka-migrate.py
preper to migrate these topics ["test1", "test0", "test108", "test109", "test5", "test4", "test7", "test6", "test102", "test103", "test100", "test101", "test106", "test107", "test104", "test105", "test182", "test183", "test180", "test181", "test186", "test187", "test184", "test185", "test188", "test189", "test46", "test47", "test44", "test45", "test42", "test43", "test40", "test41", "test48", "test49", "test8", "test179", "test178", "test177", "test176", "test175", "test174", "test173", "test172", "test171", "test170", "test55", "test54", "test57", "test56", "test51", "test50", "test53", "test52", "test59", "test58", "test168", "test169", "test164", "test165", "test166", "test167", "test160", "test161", "test162", "test163", "test60", "test61", "test62", "test63", "test64", "test65", "test66", "test67", "test68", "test69", "test2", "test159", "test158", "test151", "test150", "test153", "test152", "test155", "test154", "test157", "test156", "test79", "test78", "test77", "test76", "test75", "test74", "test73", "test72", "test71", "test70", "test146", "test147", "test144", "test145", "test142", "test143", "test140", "test141", "test148", "test149", "test82", "test83", "test80", "test81", "test86", "test87", "test84", "test85", "test88", "test89", "test3", "test133", "test132", "test131", "test130", "test137", "test136", "test135", "test134", "test139", "test138", "test91", "test90", "test93", "test92", "test95", "test94", "test97", "test96", "test99", "test98", "test35", "test19", "test18", "test11", "test10", "test13", "test12", "test15", "test14", "test17", "test16", "test120", "test121", "test122", "test123", "test124", "test125", "test126", "test127", "test128", "test129", "test24", "test25", "test26", "test27", "test20", "test21", "test22", "test23", "test28", "test29", "test", "test119", "test118", "test115", "test114", "test117", "test116", "test111", "test110", "test113", "test112", "test195", "test194", "test197", "test196", "test191", "test190", "test193", "test192", "test199", "test198", "test33", "test32", "test31", "test30", "test37", "test36", "test9", "test34", "test39", "test38"] to cloud instance ckafka-lxmap3v4 .....

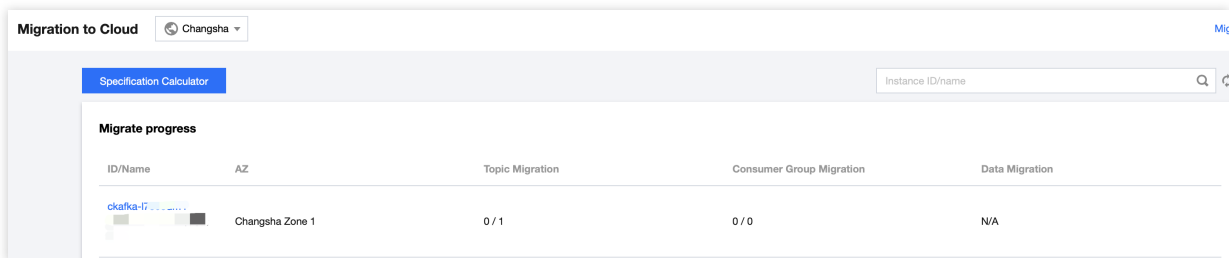
create Migreate to Cloud Task Success!

1/201 migrate topic test1 to ckafka-lxmap3v4 success,topicId:topic-9c9q3i2i

2/201 migrate topic test0 to ckafka-lxmap3v4 success,topicId:topic-o57vfbw2
```

5. Log in to the [CKafka console](#), view the task list on the **Migration to Cloud** page, and wait for the topics to be completely migrated.

The task list is as shown below:



The result of successful migration is as shown below:

```
199/201 migrate topic test34 to ckafka-lxmap3v4 success,topicId:topic-a4bo0ske
200/201 migrate topic test39 to ckafka-lxmap3v4 success,topicId:topic-bv7pezvi
201/201 migrate topic test38 to ckafka-lxmap3v4 success,topicId:topic-m96t1b0e
migrate topics to cloud finished!
```

Step 3. Migrating Data to Cloud

Scheme 1: Single-Producer Dual-Consumer Migration

Last updated : 2024-01-09 14:49:39

Overview

This document describes how to use the single-producer dual-consumer scheme to migrate data from a self-built Kafka cluster to a CKafka cluster.

Prerequisites

You have already [purchased a CKafka instance](#).


You have already [migrated the topic to CKafka](#).

Directions

If your requirement for message ordering is not high, you can migrate the data while it is consumed by multiple consumers in parallel.


The single-producer dual-consumer scheme is simple, clear, and easy to implement, with no data heap for smooth transition; however, it requires adding a new consumer.

The migration steps are as follows:



1. Keep the old consumer intact, start a new consumer in CKafka, and configure the bootstrap-server of the new CKafka cluster for consumption.


You need to configure the accessed network of the CKafka instance as the IP in `--bootstrap-server` by copying the information in the **Network** column in the **Access Mode** section on the **Instance Details** page in the console.



```
./kafka-console-consumer.sh --bootstrap-server xxx.xxx.xxx.xxx:9092 --from-beginning
```

2. Switch the production flow so that the producer produces data to the CKafka instance.

Change the IP in the `broker-list` to the accessed network of the CKafka instance and `topicName` to the topic name in the CKafka instance:



```
./kafka-console-producer.sh --broker-list xxx.xxx.xxx.xxx:9092 --topic topicName
```

3. The original consumer does not need to be configured and can continue to consume the data in your self-built Kafka cluster. After such data is all consumed, the migration is completed.

Note:

The above commands are test commands. In actual business operations, just modify the broker address configured for the corresponding application and then restart the application.

Scheme 2: Single-Producer Single-Consumer Migration

Last updated : 2024-01-09 14:49:39

Overview

This document describes how to use the single-producer single-consumer scheme to migrate data from a self-built Kafka cluster to a CKafka cluster.

Prerequisites

You have already [purchased a CKafka instance](#).

You have already [migrated the topic to CKafka](#).

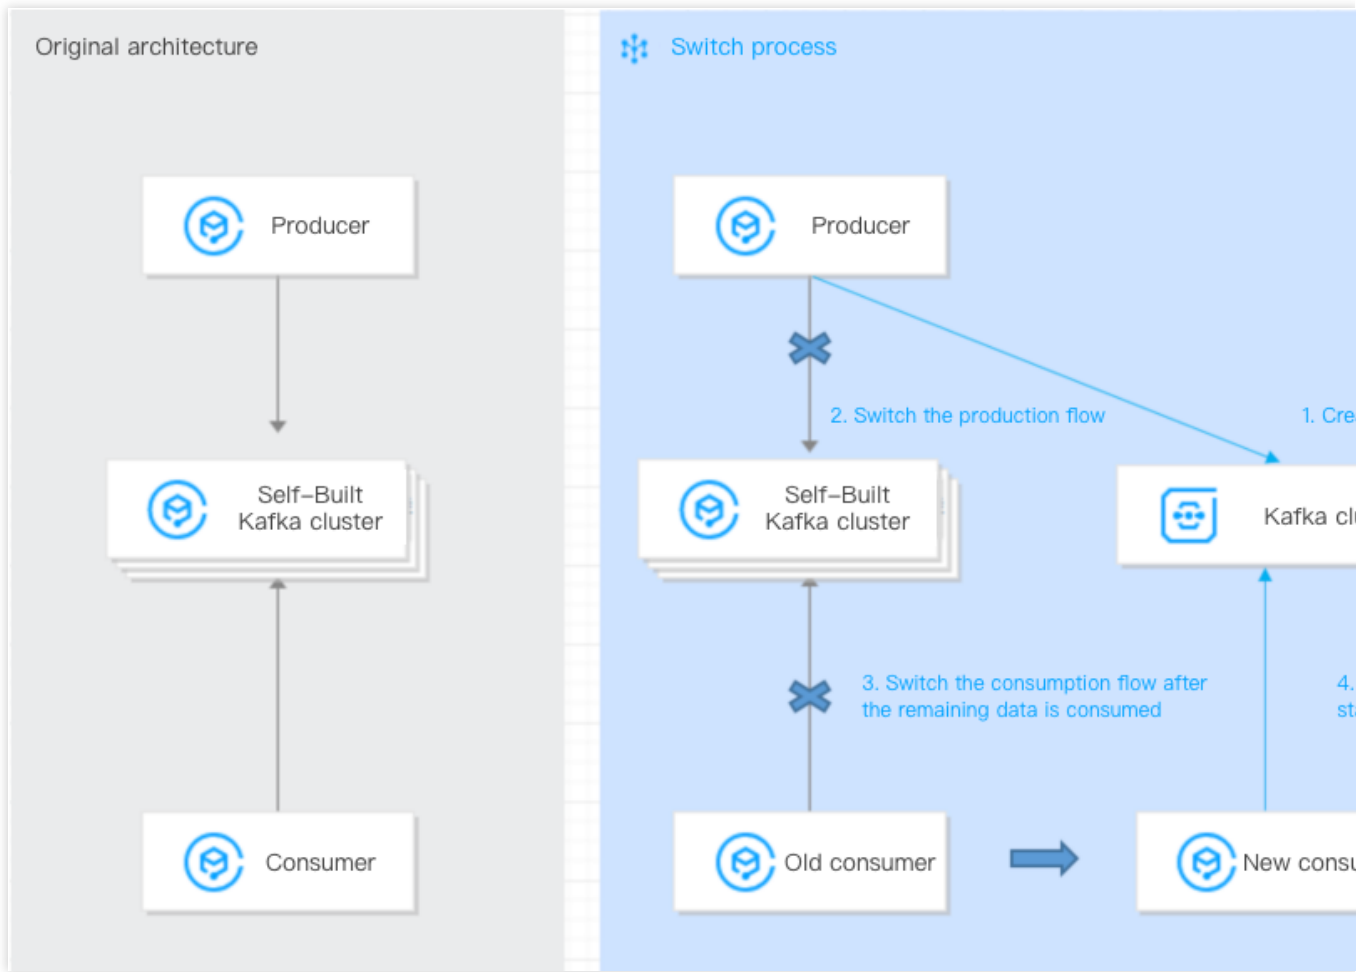
Directions

The prerequisite for guaranteeing message ordering is to strictly limit data consumption to only one consumer.

Therefore, timing of the migration is vital.


The single-producer single-consumer scheme is simple, clear, and easy to implement; however, after the production is switched to the new cluster, before the old consumer is switched to the new cluster, there will be a certain amount of data heap in the new cluster.

The migration steps are as follows:



1. Switch the production flow so that the producer produces data to the CKafka instance.

Configure the accessed network of the CKafka instance as the IP in `broker-list` by copying the information in the **Network** column in the **Access Mode** section on the **Instance Details** page in the console, and change the `topicName` to the topic name in the CKafka instance.




```
./kafka-console-producer.sh --broker-list xxx.xxx.xxx.xxx:9092 --topic topicName
```

2. The original consumer does not need to be configured and can continue to consume the data in your self-built Kafka cluster until all data is consumed.
3. When the consumption by the original consumer is completed, switch to the new CKafka cluster for consumption through the following configuration (let only one consumer consume the data to guarantee message ordering). If you add a new consumer, you need to configure the accessed network of the CKafka instance as the IP in `--bootstrap-server` .

Note:

If the original consumer is a CVM instance, it can continue to consume the data.



```
./kafka-console-consumer.sh --bootstrap-server xxx.xxx.xxx.xxx:9092 --from-beginning
```

4. The consumer continues to consume data in the CKafka cluster after switch, and the migration is completed (if the original consumer is a CVM instance, it can continue to consume the data).

Note:

The above commands are test commands. In actual business operations, just modify the broker address configured for the corresponding application and then restart the application.

Scheme 3: Migration with MirrorMaker

Last updated : 2024-01-09 14:49:40

Overview

This document describes how to use MirrorMaker to migrate data from a self-built Kafka cluster to a CKafka cluster. Kafka MirrorMaker can back up data in a self-built Kafka cluster to a CKafka cluster in the following way: MirrorMaker can use a consumer to consume messages in the self-built Kafka cluster and then use a producer to transfer the message data to the CKafka cluster. After this, you can transfer the production and consumption configurations of the client to the accessed network of the CKafka cluster so as to complete data migration.

Prerequisites

You have already [purchased a CKafka instance](#).

You have already [migrated the topic to CKafka](#).


Directions

1. [Download MirrorMaker](#) and decompress it on the local server.

Note:

This document takes [kafka_2.11-1.1.1.tgz](#) as an example:

2. Configure the `consumer.properties` file.




```
# list of brokers used for bootstrapping knowledge about the rest of the cluster
# format: host1:port1,host2:port2 ...
bootstrap.servers=localhost:9092

# consumer group id
group.id=test-consumer-group

partition.assignment.strategy=org.apache.kafka.clients.consumer.RoundRobinAssignor
# What to do when there is no initial offset in Kafka or if the current
# offset does not exist any more on the server: latest, earliest, none
#auto.offset.reset=
```

Parameter	Description
<code>bootstrap.servers</code>	List of broker access points of the self-built cluster.
<code>group.id</code>	ID of the consumer group used during data migration. It must be different from the r groups in the self-built cluster.
<code>partition.assignment.strategy</code>	Partition assignment policy, such as <code>partition.assignment.strategy=org.apache.kafka.clients.con</code>

3. Configure the `producer.properties` file.




```
# list of brokers used for bootstrapping knowledge about the rest of the cluster
# format: host1:port1,host2:port2 ...
bootstrap.servers=localhost:9092

# specify the compression codec for all data generated: none, gzip, snappy, lz4
compression.type=none
```

Parameter	Description
bootstrap.servers	Accessed network of the CKafka instance, which can be copied from the Network column in t

	<p>Access Mode section on the Instance Details page in the console.</p> <div data-bbox="459 309 1481 479" style="border: 1px solid #ccc; padding: 10px; margin: 10px auto; width: fit-content;"> <p>Access Mode ?</p> <p>VPC Network PLAINTEXT 10.1.1.1 3092 Delete</p> </div>
<p>compression.type</p>	<p>Data compression type. CKafka does not support the Gzip compression format.</p>

4. Start MirrorMaker in the `.bin` directory to start migration.




```
sh bin/kafka-mirror-maker.sh --consumer.config config/consumer.properties --produce
```

Note:

`whitelist` is a regular expression, and topics whose names hit it will be migrated.

5. Run `kafka-consumer-groups.sh` in the `bin` directory to view the consumption progress of the self-built cluster.



```
bin/kafka-consumer-groups.sh --new-consumer --describe --bootstrap-server self-buil
```

Note:

`group` is the ID of the consumer group used during data migration.

```
[root@MM-0-36-centos ~]# /usr/local/services/kafka_2.11-1.1.1/bin/kafka-consumer-groups.sh --bootstrap-s
group mm1-migrate-ckafka
Note: This will not show information about old Zookeeper-based consumers.
```

TOPIC	PARTITION	CURRENT-OFFSET	LOG-END-OFFSET	LAG	CONSUMER-ID
test26	0	-	461	-	mm1-migrate-ckafka-0-598a37fa
.0.36	mm1-migrate-ckafka-0				
test193	0	-	460	-	mm1-migrate-ckafka-0-598a37fa
.0.36	mm1-migrate-ckafka-0				
test89	0	-	461	-	mm1-migrate-ckafka-0-598a37fa
.0.36	mm1-migrate-ckafka-0				
test114	0	-	461	-	mm1-migrate-ckafka-0-598a37fa
.0.36	mm1-migrate-ckafka-0				
test1	3	-	2251	-	mm1-migrate-ckafka-0-598a37fa
.0.36	mm1-migrate-ckafka-0				

Subsequent Operations

After data migration, transfer the production and consumption configurations of the client to the access point of the CKafka cluster.

Scheme 4: Migrating Unconsumed Data

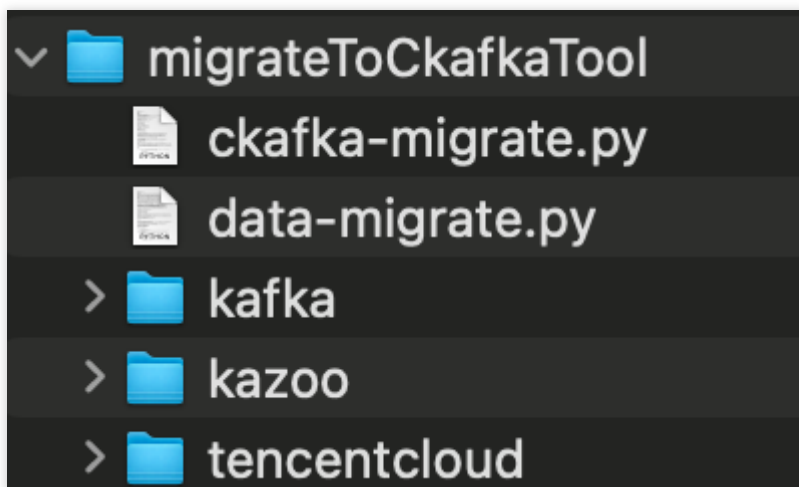
Last updated : 2024-01-09 14:49:39

Overview

This document describes how to migrate unconsumed messages in a source cluster to the topics of a target cluster during the migration to the cloud.

Prerequisites

1. The source cluster has stopped message production and consumption.
2. The source cluster retains the unconsumed messages long enough to prevent them from being automatically deleted when they expire during the migration.
3. The migration script must be written in Python 2.x series on a version later than v2.7.1, preferably v2.7.5.
4. You have downloaded the migration tool [migrateToCkafkaTool](#). Then go to the `migrateToCkafkaTool` directory, modify the configurations of the `data-migrate.py` file, and run `python data-migrate.py`.



How It Works

The script first scans the list of all consumer groups in the source cluster to get the list of topics (with unconsumed messages) to which the consumer has subscribed. Then the script will get the **submitted offset** and **end offset** of these topics (if a topic is subscribed to by multiple consumer groups, the script will get the smallest submitted offset). After consuming messages between the two offsets, the script will produce these messages to the corresponding topic partitions of the target cluster.

Directions

1. Create a topic in the target cluster

In the following example, the source cluster is `ckafka-47bd7goz` and the target cluster is `ckafka-kzamzogr`. We have created four topics with the same partition count as the source cluster: `test1`, `test2`, `test3`, and `test4` in the target cluster.

In the source cluster `ckafka-47bd7goz`, there are two consumer groups named `test123-group` and `test34-group`. The former subscribes to topics `test1`, `test2`, and `test3`, and the latter, `test3` and `test4`.

2. Download the migration toolkit

After downloading the migration tool, open the script and enter the address configurations of the source and target clusters. Set the value of `checkFlag` to `0` and run the script to check the topics to be migrated and their offsets.

```
# your local broker ip:port list
# Broker list of the self-built instance ["broker1:port1","ip2:port2"]
bootstrapServers = ["██████████"] → Source instance

# your cloud kafka instance address
# Cloud instance address "x.x.x.x:9092"
instanceAddress = "██████████" → Target instance

# A dedicated group for migration data, default migrateLocalToCkafkaGroup, can be customized but s
# the same name as the existing consumer group in the cluster, do not delete it after the self-bu
# generated, otherwise the migration progress will be lost
# A dedicated consumer group for data migration, which is 'migrateLocalToCkafkaGroup' by default and can be customized. The name of the custom consumer gro
group. Once you create a custom consumer group in the self-built cluster, don't delete it; otherwise, migration progress data will get lost. migrateGroupId = "mig

# topic regex,just migrate match topics
# Topic name regex, which is left empty by default, indicating that all topics will be migrated; otherwise, only topics that match the regex will be migra
topicRegex = ""

# The mapping between the local topic and the cloud topic, if not set, the default local and cloud
# The mapping between the local and cloud topics. If it is not set, the local and cloud topics will have the same name.
# example: localTopicAndCloudTopicMap = {"topic1-local": "topic1-cloud", "topic2-local": "topic2-c
localTopicAndCloudTopicMap = {}

# if you make sure the migrate topic List,please modify checkFlag = 1
# Check flag. If it is set to 0, the list of topics to be migrated will only be displayed but not migrated. You can set it to 0 first to check the topics to be migrated, confirm that everything
# 0: List the partition offset data of the topics to be migrated and stop the script.
# 1: List the partition offset data of the topics to be migrated and start migration.
checkFlag = 0 → Set the value to 0 first. If the script runs normally, set the value to 1 and run

logFileName = "data-migrate"
```

After the script is run, some information will be output, and a text log will be written to the current directory.

```
ckafka-migrate.py data-migrate_220914_12:09:56.log ← data-migrate.py kafk
```

3. Check the output information

Check the **Prepare to migrate** information, which is the offset information of the topics to be migrated, based on the output information displayed on the screen or the text log file.

```
-----Topic:[test3] Partition:(0) migrate-----
Prepare to migrate local-topic:[test3] partition:0 startOffset:187800 ~ endOffset:264322 data to

-----Topic:[test2] Partition:(0) migrate-----
Prepare to migrate local-topic:[test2] partition:0 startOffset:1873 ~ endOffset:2665 data to

-----Topic:[test2] Partition:(1) migrate-----
Prepare to migrate local-topic:[test2] partition:1 startOffset:1860 ~ endOffset:2621 data to

-----Topic:[test2] Partition:(2) migrate-----
Prepare to migrate local-topic:[test2] partition:2 startOffset:1830 ~ endOffset:2591 data to
```

For example, you can check the subscription information of the topic `test3` in the source cluster. It is subscribed to by the consumer groups `test123-group` and `test34-group` at the same time.

Theoretically, if a topic is subscribed to by multiple consumer groups, it will be synced from the smallest submitted offset, which is 187800 in this example as expected.

```
-----Topic:[test3] Partition:(0) migrate-----
Prepare to migrate local-topic:[test3] partition:0 startOffset:187800 ~ endOffset:264322 data to
```

If any messages in topic `test1` of the source cluster have expired and the consumer group has a submitted offset falling between expired messages, the topic `test1` will be synced from the smallest offset of non-expired messages.

Let's say you want to migrate the data from partition 0 of `test1`. The script will prompt you that messages have expired when the smallest offset (5226) of non-expired messages exceeds the offset (3713) submitted by the consumer group. Therefore, the offset 5226 becomes the start offset of the migration task and the largest offset of partition 0, making it impossible to migrate any messages from this partition. In this case, you will see a prompt saying `skip migrate...`, indicating that the migration of data in this partition will be skipped.

```
-----Topic:[test1] Partition:(0) migrate-----
the local-topic:[test1] partition:0 topic_begionOffset:5226 large than group_startOffset:3713,and reset startOffset to 5226
skip migrate local-topic:[test1] partition:0 startOffset:5226 ~ endOffset:5226 data to cloud-topic:[test1],beacuse startOffset:5226
-----Topic:[test1] Partition:(1) migrate-----
the local-topic:[test1] partition:1 topic_begionOffset:5340 large than group_startOffset:3784,and reset startOffset to 5340
skip migrate local-topic:[test1] partition:1 startOffset:5340 ~ endOffset:5340 data to cloud-topic:[test1],beacuse startOffset:5340
```

4. Start the migration

After checking that the output information is correct in the previous step, set the value of `checkFlag` to 1 and start the migration.

```
-----Topic: [test3] Partition:(0) migrate-----
Prepare to migrate local-topic:[test3] partition:0 startOffset:187800 ~ endOffset:264322 da
begin migrate local-topic:[test3] partition:0 startOffset:187800 ~ endOffset:264322 da
finish migrate local-topic:[test3] partition:0 startOffset:187800 ~ endOffset:264322 da

-----Topic: [test2] Partition:(0) migrate-----
Prepare to migrate local-topic:[test2] partition:0 startOffset:1873 ~ endOffset:2665 da
begin migrate local-topic:[test2] partition:0 startOffset:1873 ~ endOffset:2665 data t
finish migrate local-topic:[test2] partition:0 startOffset:1873 ~ endOffset:2665 data t

-----Topic: [test2] Partition:(1) migrate-----
Prepare to migrate local-topic:[test2] partition:1 startOffset:1860 ~ endOffset:2621 da
begin migrate local-topic:[test2] partition:1 startOffset:1860 ~ endOffset:2621 data t
finish migrate local-topic:[test2] partition:1 startOffset:1860 ~ endOffset:2621 data t

-----Topic: [test2] Partition:(2) migrate-----
Prepare to migrate local-topic:[test2] partition:2 startOffset:1830 ~ endOffset:2591 da
begin migrate local-topic:[test2] partition:2 startOffset:1830 ~ endOffset:2591 data t
finish migrate local-topic:[test2] partition:2 startOffset:1830 ~ endOffset:2591 data t
```

5. Check whether the number of messages is the same before and after the migration

Taking `test3` for example, if you want to migrate 76,522 unconsumed messages from `test123-group`, and all these messages have actually been written to the `test3` topic of the target instance, data migration is considered successful.

Data Compression

Last updated : 2024-01-09 14:49:40

Overview

Data compression can reduce network I/O transmission traffic and disk usage. This document describes the message formats supported for data compression and how to configure data compression based on your needs.

Message Format

Currently, CKafka supports two versions of message formats: v1 and v2 (imported in Kafka 0.11.0.0). CKafka is compatible with Kafka 0.9, 0.10, 1.1, 2.4, and 2.8.

Different configurations apply to different versions, which are described as below:

Message format conversion is mainly for compatibility with consumer programs on legacy versions. A CKafka cluster usually has message formats on multiple versions (v1 and v2).

The broker will decompress and recompress messages on a new version to convert them to the legacy format.

Message format conversion affects performance greatly as it requires extra compression and decompression operations. It also makes Ckafka's excellent zero-copy feature useless. Therefore, **you must use the same message format**.

Zero-copy: This feature can prevent costly data copy in kernel state when data is transferred in disks or over networks to implement fast data transfer.

Compression Algorithm Comparison

Snappy is the officially recommended compression algorithm. Its analysis process is as follows:

The performance of a compression algorithm is evaluated mainly based on two metrics: compression ratio and compression/decompression throughput.

Versions earlier than CKafka 2.1.0 support three compression algorithms: Gzip, Snappy, and LZ4.

The comparison of **performance metric** between the three algorithms in the actual use of CKafka is as shown below:

Compression ratio: LZ4 > Gzip > Snappy

Throughput: LZ4 > Snappy > Gzip

Comparison of **physical resource** usage is as shown below:

Bandwidth: As Snappy has the lowest compression ratio, its network bandwidth usage is the highest.

CPU: CPU usage is similar for each compression algorithm. Snappy uses more CPU resources during compression, while Gzip uses more CPU resources during decompression.


Therefore, the recommended order of the three compression algorithms under normal circumstances is LZ4 > Gzip > Snappy.

This recommended order has been well tested in most cases in the production environment. However, in extreme cases, LZ4 will increase the CPU load.

The analysis shows that LZ4 performs differently depending on the source data. Therefore, we recommend that you use the more stable Snappy compression algorithm if you are more concerned about the CPU usage.

Configuring Data Compression

A producer can use the following method to configure data compression:



```
Properties props = new Properties();
props.put("bootstrap.servers", "localhost:9092");
props.put("acks", "all");
props.put("key.serializer", "org.apache.kafka.common.serialization.StringSerializer");
props.put("value.serializer", "org.apache.kafka.common.serialization.StringSerializer");

// After the producer is started, all its produced message sets will be compressed,
// Note that different versions have different configurations. Currently, versions
// 0.11.0 and later support lz4 compression.
props.put("compression.type", "lz4");
```

```
Producer<String, String> producer = new KafkaProducer<>(props);
```

In most cases, after receiving a message from the producer, the broker will retain

Note

When data is sent to CKafka, `compression.codec` cannot be set.

Gzip compression is not supported by default. To use it, [submit a ticket](#).

As Gzip compression causes high CPU consumption, if it is used, all messages will become `InValid` .

The program cannot run properly when the LZ4 compression is used. Possible causes include:

The message format is incorrect. The default message version of CKafka is v0.10.2. You need to use the message format v1.

The setting method for SDK varies by CKafka client. You can query the setting method in the open-source community (such as the description for [C/C++ Client](#)) to set the version of the message format.