

Message Queue CKafka Troubleshooting Guide Product Documentation



Copyright Notice

©2013-2019 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.

Troubleshooting Guide

Last updated : 2021-03-19 16:16:12

What if a production/consumption error occurs after a new client is connected to the service?

- Check whether telnet works. It might be a network issue. Check whether Kafka and the producer are in the same network.
- Check whether the accessed vip - port is correctly configured.
- Check whether the topic allowlist is enabled, and if so, you need to configure the correct IP for access.

What if no messages can be seen when testing with the client in the Kafka console?

-If the "latest" option is used, a consumer can only get the last messages, and production needs to be ongoing in order to see the corresponding messages.

- Change to the "earliest" option for data consumption.

What if an error persists after a period of production?

- View whether traffic throttling is performed. Check whether there is a surge in traffic on the monitoring page and upgrade the instance specification for solution.
- View whether capacity throttling is performed. Check the instance capacity on the monitoring page and upgrade the instance specification or modify the data retention period for solution.

Traffic throttling description

- Traffic throttling is at the instance level and affects all topics in the instance.
- After the capacity is used up, consumption can continue but production cannot.
- Traffic = actual traffic * number of replicas
- Heap = actual heap (replica heap is also counted)

What if consumption is exceptional?

- View whether traffic throttling is performed. Check whether there is a surge in traffic in Barad and upgrade the instance specification for solution.
- Check whether the number of consumer groups exceeds the limit.
- If rebalance occurs frequently for network reasons, we recommend you adjust the client timeout period.
- Check whether expired offsets are pulled. Messages will be deleted after expiration, and if they are pulled with expired offsets, the pull will fail.

What if data heaps exceptionally?

Open-source Kafka supports setting a timestamp field and type in a message. Currently, two timestamp types are supported: `CreateTime` and `LogAppendTime`.

- `CreateTime` indicates the local time on the client. As client time may vary from server time, please check whether the entered time is correct. If it differs too much from the current Beijing time, the CKafka service will not be able to expire and delete data promptly based on the normal message retention period, which may cause exceptional message heap.
- `LogAppendTime` indicates the time when a message is produced to the CKafka service, which is the CKafka server time and thus recommended.

What if details are missing in the consumer group list?

Issue

The consumer group list in the CKafka console contains consumer group name, but no consumption details are displayed on the details page; for example, the consumer group CR has no details

displayed in the figure below:

Partition Name	Submitted offset position	Max Offset Position	Number of unconsumed messa...	Operation
No data yet				

Troubleshooting

1. View the consumer groups of the instance.

```
$ bin/kafka-consumer-groups.sh --bootstrap-server 127.0.0.1:9092 --list
CR
```

You should be able to see the names of all current consumer groups. If this is not the case, check whether the consumer clients are normal.

2. View the details of a specific consumer group of the instance.

```
$ bin/kafka-consumer-groups.sh --bootstrap-server 127.0.0.1:9092 --describe --group CR
```

Note: This will not show information **about** old Zookeeper-based consumers.

You will find that there are no details for this consumer group, which means that the consumer client did not use the `consumerGroup` mechanism to consume data, that is, the client did not submit consumption details to the server. As the server did not store the consumption data, no details will be displayed.

3. Check whether the problem is caused by the server.

Use the native consumer group command to specify the consumer group `test1` for consumption as shown below:

```
$ bin/kafka-console-consumer.sh --bootstrap-server 127.0.0.1:9092 --from-beginning --topic test --group test1
```

You should be able to see the normally displayed consumer group in the console, and you can run the `--describe` command to view the details as shown below:

```
$ bin/kafka-consumer-groups.sh --bootstrap-server 127.0.0.1:9092 --describe --group test
```

GROUP	TOPIC	PARTITION	CURRENT-OFFSET	LOG-END-OFFSET	LAG	CONSUMER-ID	HOST	CLIENT-ID
test	test	0	-	67	-	sarama/127.0.0.12020-06-11 23:33:03:110-a113e7d1-e36e-44f5-867c-f3e2fa2a57ab	/127.0.0.1	sarama

How it works

There are two data consumption modes in Kafka: consumer group mode and custom partition consumption mode.

- When consumption is performed in consumer group mode, the client will coordinate consumption through the consumption coordinator, and after data consumption is completed, it will send an offset storage request to the server, which will store information such as the consumed topic, partition progress, and client.
- When consumption is performed in custom partition consumption mode, the client will not automatically submit an offset storage request to the server. In this case, the server will not be able to see information related to consumption.

What if the linked capabilities of other Tencent Cloud services fail after an ACL policy is configured for a topic?

When you add an ACL policy for a topic, the policy will prevent all other ineligible requests from accessing the topic, including those initiated by other Tencent Cloud services (e.g., log shipping in CLS, message dump in SCF, and component consumption in EMR) connected to CKafka. Therefore, before adding an ACL policy for a topic, you must determine whether the topic is being used in other scenarios through the service information or the monitoring information in the console; otherwise, problems with other linked features may occur.

In such cases, if you have to use an ACL policy, we recommend you produce messages to a new topic for permission grant instead of reusing the original topic.