

Message Queue CKafka

FAQ

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.

Contents

FAQ

FAQs for Features

About Configurations

FAQ on CKafka

FAQ

FAQs for Features

Last updated : 2020-09-04 16:28:57

What is the difference between CKafka and CMQ?

CMQ provides finance-grade message transmission with high reliability and high data persistence while ensuring strong data consistency.

CKafka is suitable for scenarios that require higher throughput and lower reliability (such as log aggregation). In addition, CKafka is compatible with existing Kafka users, enabling migration at zero cost with full instance exclusivity.

Which version of open-source Kafka is compatible with CKafka?

Currently, CKafka is fully compatible with open-source Kafka API 0.9, 0.10, 1.1, and 2.4, allowing users to migrate data to the cloud at zero cost.

Which version of open-source Kafka is the current CKafka based on?

The current CKafka is based on Apache Kafka 0.10, 1.1, and 2.4. We recommend using an SDK for production and consumption according to the Apache Kafka's version.

Does CKafka expose ZooKeeper?

CKafka does not expose ZooKeeper or its address.

Does CKafka support public network access?

Currently, CKafka transfers data over the private network by default. As public network access runs the risk of issues such as delay and network environment security, we do not recommend long-term use of public network transfer.

If you have a temporary need for public network transfer, please contact your Tencent Cloud account manager for evaluation and assistance.

Does CKafka support message compression?

Currently, CKafka supports open-source Snappy and LZ4 message compression formats. Because gzip compression consumes more CPU resources, it is currently not supported.

We recommend disabling message compression when testing.

Configuration Method: in the configuration file of the producer, set the `compression.type` parameter to `snappy` or `lz4`. The default value is `none`, indicating that the feature is disabled.

Can a Kafka client connect directly to the CKafka service?

CKafka is compatible with open-source Kafka 0.9, 0.10, 1.1, and 2.4. You can connect to the message center via a Kafka client and deploy codes to Tencent Cloud services to produce or consume messages.

What are the restrictions on a CKafka instance?

Different instance specifications have different restrictions on peak throughput, disk capacity, number of topics at the instance level, and number of partitions at the instance level. For more information, see [Billing Overview](#).

Will CKafka lose messages?

- Open-source Apache Kafka does not guarantee no message loss. As CKafka is optimized for availability, Tencent Cloud promises a CKafka availability of over 99.95%.
- CKafka users can enable ACK during production to avoid message loss and improve message reliability.
- Cluster changes and upgrades complete in seconds and will not affect user experience.
- CKafka is mainly used in big data processing scenarios that require high throughput and high performance but relatively low data reliability. In extreme cases, a small number of messages may be lost. If you require zero message loss and have relatively performance requirements, we recommend CMQ.

How does CKafka guarantee security?

CKafka guarantees security by using following features:

- Tenant isolation: the network access of instances is isolated among different accounts by default.
- Permission control: CKafka has an additional allowlist-based authentication mechanism at the application layer and supports SASL authentication.
- Security protection: services such as multi-dimensional security protection and anti-DDoS attacks are provided.

About Configurations

Last updated : 2020-07-30 11:41:36

How do I choose an appropriate number of CKafka replicas?

To ensure data reliability, you are recommended to choose two or three replicas for data storage when creating a topic. Currently, CKafka has banned the creation of single-replica topics. If you have a single-replica topic in your account, please migrate it as follows:

1. Create a topic, select the same partition parameter, and select "dual-replica";
2. Produce messages in the new topic while the existing single-replica topic continues to be consumed;
3. Modify the consumer configuration after consumption is completed to subscribe to the new topic for consumption.

How do I fix production/consumption errors when a new client is connected to CKafka?

- Check whether telnet works. It might be a network issue. Check if 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. If yes, you need to configure the correct IP for access.

The CKafka message retention period is configured to 1 minute. Will the heaped messages be deleted immediately after 1 minute?

Not necessarily. Message deletion is related to not only the retention period configuration but also the data size of produced messages.

CKafka's smallest unit for heaped message deletion is partition-level file segment. The current file segment size is 1 GB. If the message heap does not reach the size of one file segment, the messages will not be deleted. If there are 10 partitions, and their total size does not reach 10 GB within 1 minute, then files will not be rolled, and messages will not be deleted.

Does CKafka support automatic topic creation (auto.create.topic)?

Currently, CKafka supports the open-source API for automatic topic creation. On the [instance list page](#), click an instance ID to enter the instance details page. On the **Basic Info** tab there, view or edit the configuration information, message retention period, and automatic topic creation of the instance.

What should I do if a lot of CKafka messages heap up?

CKafka uses exactly the same mechanism and principle as open-source Kafka. You can troubleshoot the problem in the following steps:

1. Determine how many consumers in your business are consuming messages.
2. If the consumers' consumption capabilities are poor, simply add more consumers.

How does CKafka throttle traffic?

To ensure stability of the service, CKafka implement network traffic control strategies on both inputting and outputting messages.

- Throttling occurs when the total traffic of the user's all replicas exceeds the purchased peak traffic.
- When the producer side is throttled, CKafka will extend the response time of a TCP connection. The delay period depends on how much the instantaneous traffic exceeds the limit. It is similar to the principle of road traffic control. The more traffic flow, the higher the delay value from the delay algorithm, up to 5 minutes.
- When the consumer side is throttled, CKafka will reduce the size of each `fetch.request.max.bytes` request to control the traffic.

How does throttling affect the production and consumption of messages?

Throttling only affects the message sending and receiving rate but not the content of messages.

How do I determine whether CKafka has been throttled?

1. In the instance list, you can see the health status of each cluster. If it's "Warning", you can hover your mouse over it to view the detailed data. The data displays your peak traffic and the throttling times, by which you can determine whether this instance has been throttled.
2. You can click the **Monitor** tab to view the max traffic value. If **the value of max traffic multiplied by replica quantity is greater than that of the purchased peak bandwidth**, you can determine that at least one throttling has occurred.

FAQ on CKafka

Last updated : 2020-07-30 11:22:55

Which version of open-source Kafka is compatible with Cloud Kafka?

CKafka service is compatible with the open-source Kafka.api of version 0.9 and later, achieving free migration to the cloud for users.

What is TOPIC?

Topic is the category to which each message published in the Cloud Kafka cluster belongs. That is to say, Cloud Kafka is topic-oriented. Users need to create a topic before read and write.

What is PARTITION?

Partition is a physical concept. Each topic can be divided into one or more partitions. Partition is used to scale out topic throughput. Published messages are written to different partitions and read by several consumers at the same time. As the assignment unit of Cloud Kafka is partition, the parallel throughput of topic is directly proportional to the number of partitions.

What is the difference between Cloud Kafka and CMQ?

CMQ provides financial-level message transmission with high reliability and high data persistence while ensuring strong data consistency.

Cloud Kafka is suitable for scenarios requiring higher throughput and relatively lower reliability, such as log aggregation. In addition, Cloud Kafka is compatible with the regular users of Kafka, with zero migration cost and complete exclusive instance.

Can the Kafka client be directly connected to Cloud Kafka?

Cloud Kafka is compatible with the open-source Kafka of version 0.9 and later. You can connect to the message center through the Kafka client, and deploy codes to Tencent Cloud services to produce or consume messages.

How does Cloud Kafka ensure security?

Cloud Kafka ensures security with the following security features:

Tenant Isolation: The network accesses of instances are naturally isolated between accounts.

Permission Control: Cloud Kafka provides an authentication mechanism for the source ip allowlist in the extra application layer.

Security Protection: Services including multi-dimensional security protection and anti-DDoS attacks are provided.

Will CKafka lose messages?

1. The open-source Apache Kafka does not guarantee no loss of messages. As CKafka is optimized in availability, Tencent Cloud promises that the availability of CKafka exceeds 99.95%.
2. CKafka customers can enable ACK during production to avoid message loss as much as possible and improve reliability.
3. Cluster change and upgrade are transparent to customers. Clusters can be changed instantly.
4. CKafka is mainly used in big data processing scenarios that require high throughput and high performance but average data reliability. A small amount of messages may be lost in extreme circumstances. For scenarios that require no message loss but average performance, it is recommended to use CMQ.

What are the restrictions on CKafka products?

Restrictions on product forms:

1. A maximum of 50 partitions and 20 groups can be created for each instance.
2. A maximum of 8 partitions and 3 copies can be created for each topic.
3. The idleness time of a consumer group is 1 month.