

Message Queue CKafka

Glossary

Product Documentation



Copyright Notice

©2013-2022 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.

Glossary

Last updated : 2022-11-07 15:34:49

Broker

A broker is a proxy server and a service provider.

In CKafka, a broker is an individual CKafka server mainly used to receive messages sent by producers, assign offsets, and save messages to disks. It also receives requests from consumers and other brokers, processes them based on request types, and returns responses.

CKafka

For more information, please see [Cloud Kafka](#).

Partition

A partition is a physical concept for storing messages. It is the basis for CKafka's horizontal scalability. The parallel processing power of CKafka can be increased by adding servers and assigning partitions on them.

- Each topic can be divided into multiple partitions and has at least one partition.
- Different partitions under the same topic contain different messages.
- Different partitions under the same topic will be assigned to different brokers.

Offset

An offset is a unique identifier of a message in a partition.

Replica

A replica is a message backup. Each partition can have multiple replicas that have the same messages (this is subject to the sync mechanism as replicas may not be exactly the same at the same time).

Each partition in CKafka has at least two replicas to ensure high service availability.

Consumer Group

A consumer group is a set of consumers. In CKafka, multiple consumers can form a consumer group, and each consumer can only belong to one consumer group. Consumer group guarantees that each partition under the topic to which it subscribes to is assigned to exactly one consumer in the group.

You are recommended to specify a consumer group ID when messages are consumed. Otherwise, CKafka system will randomly generate a consumer group, which may reach the maximum number of consumer groups allowed. For more information, please see [CKafka Billing Overview](#).

Cloud Kafka

Cloud Kafka (CKafka) is a distributed, high-throughput, and highly scalable messaging system fully compatible with open-source Kafka API (v0.9 and 0.10). Based on the publish/subscribe model, it enables async interaction between producer and consumer by decoupling the messages and thereby eliminates the wait time on both sides. It supports data compression and offline/real-time data processing, making it ideal for scenarios such as collection of compressed logs and aggregation of monitoring data.

ZooKeeper

ZooKeeper is a software program that provides centralized services for distributed applications, such as configuration maintenance, domain name service, distributed sync, and group service. In CKafka, it is mainly used to store the metadata of the cluster, elect leaders, and tolerate faults.