

Cloud Log Service

Product Introduction

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

Product Introduction

- Overview

- Available Regions

- Concept

 - Log Topic

 - Logset

 - Server Group

 - Topic Partition

Product Introduction

Overview

Last updated : 2020-08-10 16:12:44

Cloud Log Service (CLS) is a centralized logging solution. You can stop worrying about scaling or other resource issues and access CLS within just five minutes. CLS offers solutions for collecting, storing, searching, and analyzing logs, helping you identify business issues, monitor metrics, ensure security, and simplify operations.

Features

CLS has the following features.

- Log collection: collect logs to CLS from different log sources by using LogListener, API, etc.
- Log storage: store log data with CLS.
- Log index: enabling log index for log query can help you quickly identify log problems.
- Log shipping: you can ship specified logs to other cloud products to meet storage or other computing needs. For example, you can ship a log to a specified COS bucket to manage its lifecycle and meet the need for log auditing.

Log collection

CLS supports collecting logs by using LogListener and API, enabling you to easily collect log data from different regions, channels, platforms, and data sources in real time and collect logs from various Tencent Cloud products.

- LogListener real-time collection: use LogListener to collect logs. This method is easy to install, reliable, and secure. It supports most mainstream Linux operating systems, delivering high performance while occupying few resources.
- Collection via API: upload logs by calling an API, without the need to install LogListener. Multiple languages are supported.

Log index and query

- Real-time index: index collected log data in real time to enable data search.
- Excellent query performance: query results are returned within seconds. 100-million-level log data search and quick data locating are supported.
- Flexible query: features such as full-text search, multi-keyword search, and cross-topic query are supported.

Log shipping

- Shipping logs to COS: ship log data to COS buckets under your account to store and manage the log data.

Available Regions

Last updated : 2020-10-16 16:19:03

Overview

You can create logsets and log topics in different regions when using CLS. Regions refer to the geographical areas of physical IDCs. Networks are completely isolated between regions. You can select the nearest region based on your business scenario and target users' geographical location to reduce log access latency and improve access speed.

Available regions and their abbreviations

Region	Abbreviation	Private Domain Name	Public Domain Name
Beijing	ap-beijing	ap-beijing.cls.tencentyun.com	ap-beijing.cls.tencentcs.com
Guangzhou	ap-guangzhou	ap-guangzhou.cls.tencentyun.com	ap-guangzhou.cls.tencentcs.com
Shanghai	ap-shanghai	ap-shanghai.cls.tencentyun.com	ap-shanghai.cls.tencentcs.com
Chengdu	ap-chengdu	ap-chengdu.cls.tencentyun.com	ap-chengdu.cls.tencentcs.com
Nanjing	ap-nanjing	ap-nanjing.cls.tencentyun.com	ap-nanjing.cls.tencentcs.com
Chongqing	ap-chongqing	ap-chongqing.cls.tencentyun.com	ap-chongqing.cls.tencentcs.com
Hong Kong (China)	ap-hongkong	ap-hongkong.cls.tencentyun.com	ap-hongkong.cls.tencentcs.com
Silicon Valley	na-siliconvalley	na-siliconvalley.cls.tencentyun.com	na-siliconvalley.cls.tencentcs.com
Singapore	ap-singapore	ap-singapore.cls.tencentyun.com	ap-singapore.cls.tencentcs.com
Mumbai	ap-mumbai	ap-mumbai.cls.tencentyun.com	ap-mumbai.cls.tencentcs.com
Frankfurt	eu-frankfurt	eu-frankfurt.cls.tencentyun.com	eu-frankfurt.cls.tencentcs.com
Tokyo	ap-tokyo	ap-tokyo.cls.tencentyun.com	ap-tokyo.cls.tencentcs.com
Seoul	ap-seoul	ap-seoul.cls.tencentyun.com	ap-seoul.cls.tencentcs.com

Notes

1. If another Tencent Cloud product is integrated with CLS, try to select a logset in the same region as the product. Tencent Cloud products in the same region can access each other over the private network, which effectively reduces latency and improves access speed.
2. The classic network cannot access CLS over private network.

Concept

Log Topic

Last updated : 2019-11-18 18:50:20

Log topic is a basic management unit of the Cloud Log Service (CLS). A [logset](#) can contain multiple log topics. A log topic corresponds to a project or application. It is recommended to collect logs of the same type from different servers under the same log topic. For example, a business project may have 3 types of logs, operation logs, application logs, and access logs. You should create a log topic for each log type.

The CLS system manages different log data based on each log topic. Each log topic can be configured with different data sources, index rules, and shipping rules. Therefore, log topic is the basic unit for configuring and managing log data on the CLS system. You need to configure corresponding rules after creating a log topic to perform log collection, search, analysis, and shipping.

In terms of functions, the log topic service can:

- Collect logs to log topics.
- Store and manage logs based on log topics.
- Search and analyze logs by log topics.
- Ship logs to other platforms based on log topics.
- Download and consume logs from log topics.

Logset

Last updated : 2020-06-30 15:39:47

A logset is a project management unit within the Cloud Log Service, used to distinguish between the logs of different projects. A logset corresponds to a project or application. It is recommended to manage the logs of different projects and products using different logsets. For example, your company runs two types of services, e-commerce services and payment services. You should create a logset for each service type.

A logset can contain multiple [log topics](#) and each log topic has the following basic attributes:

- Logset name: the name of a logset.
- Logset ID: the unique ID of a logset.
- Region: the [regions](#) to which the logset belongs.
- Storage period: the storage period for the data in the current logset. The period can be 1-90 days.

Server Group

Last updated : 2019-11-18 18:45:41

A server group is a list of servers for which logs need to be collected. In server groups, Cloud Log Service manages all servers for which logs are collected using [LogListener](#).

A server group can contain multiple servers. Generally, service applications of the same type are deployed on multiple servers, and the log collection configuration of these servers are similar. Therefore, you can assign these servers to the same server group, and then associate the server group with the corresponding log topic. You can associate a log topic with multiple server groups, and a server group can be associated with multiple log topics.

To define a server group, add the IP address of a server to the server group. Then, the server group will be able to identify the server using the IP address.

Topic Partition

Last updated : 2020-08-04 10:45:10

Overview

Partition is the smallest read/write unit of Cloud Log Service (CLS). A log topic can be divided into several partitions and must have at least one partition. CLS uses the value range of MD5 as the valid range and controls the overall throughput performance by merging or splitting partitions. A log topic supports up to 50 partitions. We recommend that you use and work with topic partitions rationally to prevent waste of resources.

Basic attributes of a partition:

- **Partition number:** every partition has a unique number under the same log topic, which is assigned by the system after partition creation or operation.
- **Partition range:** every partition has a left-closed and right-open interval.
- **Partition status:**
 - Read-write: the current partition allows read and write.
 - Read-only: the current partition is read-only and no data can be written to it.

Partition Range

With a partition range, logs can be written in the HashKey mode. The valid range of a log topic is the value range of MD5, which is `[00000000000000000000000000000000, ffffffffffffffffffffffffffffffff)`. All read-write partitions segment the entire value range of a log topic and each occupies a left-closed and right-open interval to ensure that every log record collected is written to the corresponding partition.

CLS provides two write modes: load-balancing mode and HashKey mode.

- **Load-balancing mode:** every data packet is written to a random log topic partition.
- **HashKey mode:** every data packet is written to the topic partition that contains the current Key value.

For example, a log topic has three read-write partitions and their ranges are as follows:

Partition No.	Status	Partition Range

Partition No.	Status	Partition Range
1	Read-write	[00000000000000000000000000000000, 7fffffffffffffffffffffffffffffffff)
2	Read-write	[7fffffffffffffffffffffffffffffffff, a0000000000000000000000000000000)
3	Read-write	[a0000000000000000000000000000000, ffffffffffffffffffffffffffffffffff)

If the write mode is HashKey, log data with the Key value of `2fffffffffffffffffffffffffffffffff` is written to partition 1 and log data with the Key value of `9f000000000000000000000000000000` is written to partition 2.

Read and Write Capability of Partitions

Each partition has a certain level of read and write capability. We recommend that you plan partition count based on the actual log traffic of your business. Split partitions when the traffic is higher than the read and write capability of the log topic, and merge partitions when the traffic is lower than the read and write capability of the log topic to save resources.

Feature	Item	Description
Frequency control	Write requests	Every partition supports a maximum of 500 QPS of write operations. It will reject requests and return the status code 429 with an error message of "SpeedQuotaExceed" when the limit is exceeded.
	Read requests	Every partition supports a maximum of 200 QPS of read operations. It will reject requests and return the status code 429 with an error message of "SpeedQuotaExceed" when the limit is exceeded.
Traffic throttle	Write traffic	Every partition supports the write traffic up to 5 MB/sec. It will truncate data and return the status code 429 with an error message of "SpeedQuotaExceed" when the limit is exceeded.

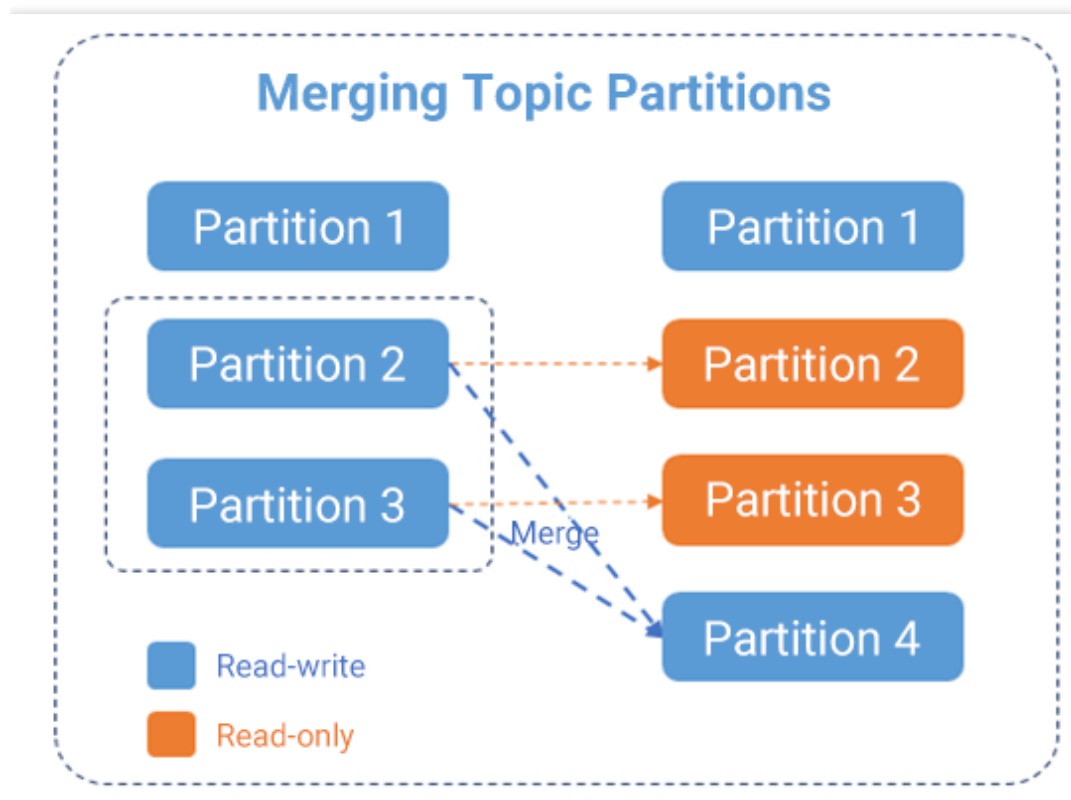
Partition Status

A partition can be in **read-write** or **read-only** mode. Only read-write partitions support data writing. Read-only partitions does not allow data writing but can be consumed within their lifetime. All

partitions are readable and writable when they are created, but merging and splitting operations will change the status to read-only.

Merging partitions

Two adjacent read-write partitions can be merged into one partition. After the merge, the two original partitions become read-only, which only allow data consumption, but not data writing. The new partition is readable and writable and covers the range of the two original partitions.



Splitting a partition

A read-write partition can be split into two partitions with smaller ranges. When splitting a partition, you need to specify the MD5 value of a split point, which must be larger than the value of the start point and smaller than the value of the end point. After the split, the original partition becomes read-only, which only allows data consumption, but not data writing. The new partitions are readable and writable and cover the range of the original partition.

