

Cloud Log Service

Product Introduction

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

Product Introduction

Overview

Features

Available Regions

Concepts

Log

Log Topic and Logset

Machine Group

Segment and Index

Topic Partition

Data Processing

Limits

Basic Resources

Log Collection

LogListener Limits

Metric Collection

Search and Analysis

Monitoring Alarm

Data Processing

Shipping and Consumption

Product Introduction

Overview

Last updated : 2024-01-20 16:29:01

Cloud Log Service (CLS) provides a one-stop log data solution. You can quickly and conveniently connect to it in five minutes to enjoy a full range of stable and reliable services from log collection, storage, and processing to search, analysis, consumption, shipping, dashboard generation, and alarming, with no need to care about resource issues such as scaling. It helps you improve the problem locating and metric monitoring efficiency in an all-around manner, making log Ops much easier.

Overview

CLS has the following features.

Log collection: CLS easily collects logs from different regions, channels, platforms, and data sources (e.g., various Tencent Cloud products) in real time.

Log storage: CLS offers two storage types: real-time storage and STANDARD_IA storage.

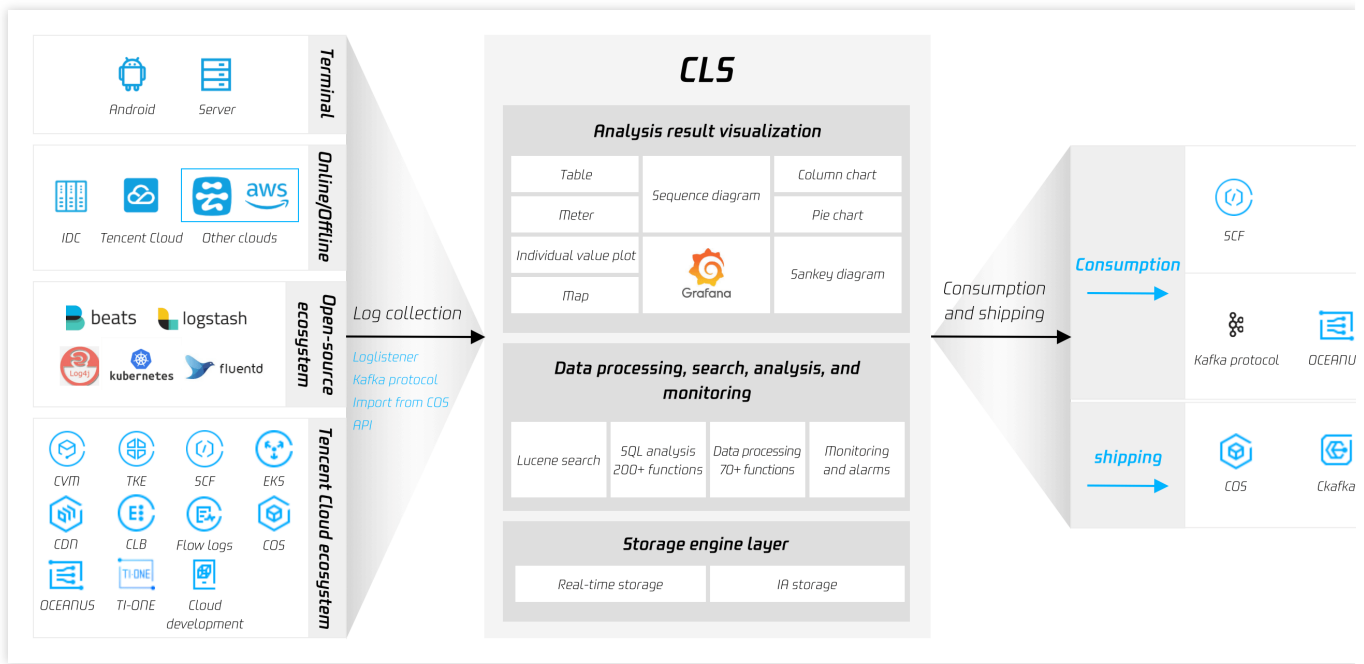
Log search and analysis: You can search for logs by keyword to quickly locate exception logs and use SQL statements to collect and analyze log statistics. This helps you get statistical metrics such as log quantity change trend over time and proportion of error logs.

Log data processing: CLS can filter, cleanse, mask, enrich, distribute, and structure logs.

Log shipping and consumption: CLS can ship logs to Tencent Cloud storage and middleware services and consume logs to stream computing services.

Dashboard: CLS can quickly generate custom dashboards for search and analysis results.

Alarming: CLS can trigger alarms for exception logs within seconds and notify you through phone, SMS, email, and custom API callback.



Log collection

CLS currently supports multiple methods for data collection, such as LogListener, API, Kafka, and COS import.

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: Call the API to upload logs, without the need to install LogListener. Multiple programming languages are supported.

Collection over Kafka: Upload logs to CLS by using the Kafka Producer SDK.

COS import: Import COS data to CLS.

Log storage

According to users' different requirements for log search latency and log processing capabilities, CLS provides two storage classes:

Real-time storage: It is suitable for users who require statistical analysis and provides log search within seconds, real-time statistical analysis, real-time monitoring, streaming consumption, and other application capabilities.

STANDARD_IA storage: It is suitable for infrequently accessed logs that do not require statistical analysis, such as archived audit logs. It provides the full-text log search capability, meeting users' requirements for backtracking and archiving historical logs. The overall usage costs of STANDARD_IA storage are 80% lower than those of real-time storage in the long run.

Log search and analysis

CLS provides real-time log search and analysis to help you quickly locate exception logs and collect system and business metric statistics.

Log search: Use keywords to search for logs by full text or field.

Statistical analysis: Use SQL statements to flexibly collect the system and business metric statistics in logs and display them in charts. This feature is compatible with the SQL-92 standard and supports over 200 SQL functions. It can collect the number of business requests by province and query the request error rate change trend by time.

Superior performance: Query results can be returned within seconds, and the search and analysis of hundreds of millions of logs are supported.

Log data processing

Real-time processing and distribution of log streams: Log streams are processed, and the processing result is generated in real time and can be distributed to multiple topics in different scenarios.

Log filtering and cleansing: Dirty data is cleansed, logs are filtered by condition, and index is enabled, effectively reducing costs.

Log structuring: Text logs are processed into structured data for future search, analysis, dashboard generation, and alarming.

Log masking: Sensitive information in the log text is masked, such as mobile number and ID number.

Log shipping and consumption

You can ship specified logs to other Tencent Cloud services, such as COS and CKafka. You can also consume CLS logs to Oceanus.

Log shipping to COS: CLS allows you to ship logs to COS buckets under your account. COS storage is economical and recommended.

Log shipping to CKafka: It is suitable for scenarios where CKafka is used as the source for data analysis and storage.

Real-time log consumption: Logs are consumed directly to big data components such as Flink, Oceanus, and Flume.

Dashboard

A dashboard is a global view of data analysis, where you can view multiple statistical charts of query and analysis results.

Dashboards: Dashboards save multiple statistical charts of query and analysis results to form a comprehensive scenario-specific view.

Preset dashboards: CLS provides preset dashboards for Tencent Cloud services such as CLB, TKE, and COS to make common monitoring capabilities out-of-the-box.

Template variables: Template variables can be data source variables and quickly filtered to switch between statistical analysis objects and dimensions in the charts of dashboards. They can be added to support more complex business scenarios.

Alarming

Alarms will be triggered within seconds in case of too many error logs or system and business metrics exceeding the threshold, to identify system and business exceptions.

Notification channels: Notifications can be sent via phone, SMS, email, Weixin, WeCom, and custom API callback (which can be connected to DingTalk and Lark).

Multi-dimensional analysis: When an alarm is triggered, raw logs can be further searched for and analyzed, and the result can be added to the alarm notification to facilitate root cause discovery.

Features

Last updated : 2024-01-20 16:29:43

Rich Features

LogListener provides one-stop log service, including log collection, storage, search, transfer, and shipping features. LogListener parses log structures in many ways, including full text in a single line, full text in multi lines, separator, JSON, and regular expressions.

LogListener allows you to select from multiple data access methods based on your business needs. For more information, see [Collection Overview](#).

LogListener provides multiple search syntaxes and allows you to search logs by keyword, fuzzy match, or range.

Stable and Reliable

CLS has a highly scalable distributed storage architecture. It supports horizontal scaling and automatic service scaling, and easily stores and manages massive log data.

The CLS backend storage manages and stores logs in multiple copies, ensuring data reliability.

Simple and Efficient

LogListener supports simple, GUI-based configuration. You can quickly access CLS through LogListener.

Data can be consumed immediately after being written to CLS. Query results of hundreds of millions of data records can be returned in seconds.

The service is billed by actual usage. You do not need to build or maintain a log system, preventing resource idling and waste.

Ecological Expansion

The logs of certain cloud services have already been integrated into CLS. For more information, see [Accessing log sources](#).

Logs are shipped to Cloud Object Storage (COS), which supports long-term archive storage of logs.

Logs are shipped to CKafka, which supports real-time log consumption and facilitates further log processing and analysis.

Available Regions

Last updated : 2024-10-24 15:05:43

Overview

You can create logsets and log topics in different regions when using CLS. Regions are independent geographical areas where IDCs are located. Tencent Cloud regions are completely isolated. You can select the nearest region based on different business scenarios and the location of your targeted users to reduce log access latency and improve user experience.

Available regions

Region	Code
Beijing	ap-beijing
Guangzhou	ap-guangzhou
Shanghai	ap-shanghai
Chengdu	ap-chengdu
Nanjing	ap-nanjing
Chongqing	ap-chongqing
Hong Kong (China)	ap-hongkong
Silicon Valley	na-siliconvalley
Virginia	na-ashburn
Singapore	ap-singapore
Bangkok	ap-bangkok
Frankfurt	eu-frankfurt
Tokyo	ap-tokyo
Seoul	ap-seoul

Jakarta	ap-jakarta
São Paulo	sa-saopaulo

Note:

If CLS is integrated into other cloud products, you need to select a logset in the same region as the other cloud products. Cloud products in the same region access each other over a private network, which effectively reduces latency and improves access speed.

Domain name

CLS has different domain names for each of its modules, as described below:

LogListener

CLS API 3.0

API for log upload

Uploading Logs via Kafka

Kafka Consumption Logs

LogListener is a log collection client provided by CLS and can report local logs to CLS. It uses the following domain names:

Region	Code	Private Network Domain Name	Public Network 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

Virginia	na-ashburn	na-ashburn.cls.tencentyun.com	na-ashburn.cls.tencentcs.com
Singapore	ap-singapore	ap-singapore.cls.tencentyun.com	ap-singapore.cls.tencentcs.com
Bangkok	ap-bangkok	ap-bangkok.cls.tencentyun.com	ap-bangkok.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
Jakarta	ap-jakarta	ap-jakarta.cls.tencentyun.com	ap-jakarta.cls.tencentcs.com
São Paulo	sa-saopaulo	sa-saopaulo.cls.tencentyun.com	sa-saopaulo.cls.tencentcs.com

CLS API 3.0 is the latest version of CLS APIs that complies with the unified API specifications of Tencent Cloud. You can use the APIs to manage resources such as log topics and alarm policies. The APIs use the following domain names:

If you access CLS over the public network, you can also use the unified domain name

`cls.tencentcloudapi.com`. Your API request will be automatically resolved to a server nearest to the client.

For latency-sensitive businesses, we recommend you specify the domain name in a region.

Region	Code	Private Network Domain Name	Public Network Domain Name
Beijing	ap-beijing	cls.internal.tencentcloudapi.com	cls.ap-beijing.tencentcloudapi.com
Guangzhou	ap-guangzhou	cls.internal.tencentcloudapi.com	cls.ap-guangzhou.tencentcloudapi.com
Shanghai	ap-shanghai	cls.internal.tencentcloudapi.com	cls.ap-shanghai.tencentcloudapi.com
Chengdu	ap-chengdu	cls.internal.tencentcloudapi.com	cls.ap-chengdu.tencentcloudapi.com
Nanjing	ap-nanjing	cls.internal.tencentcloudapi.com	cls.ap-nanjing.tencentcloudapi.com
Chongqing	ap-chongqing	cls.internal.tencentcloudapi.com	cls.ap-chongqing.tencentcloudapi.com
Hong Kong (China)	ap-hongkong	cls.internal.tencentcloudapi.com	cls.ap-hongkong.tencentcloudapi.com

Silicon Valley	na-siliconvalley	cls.internal.tencentcloudapi.com	cls.na-siliconvalley.tencentcloudapi.com
Virginia	na-ashburn	cls.internal.tencentcloudapi.com	cls.na-ashburn.tencentcloudapi.com
Singapore	ap-singapore	cls.internal.tencentcloudapi.com	cls.ap-singapore.tencentcloudapi.com
Bangkok	ap-bangkok	cls.internal.tencentcloudapi.com	cls.ap-bangkok.tencentcloudapi.com
Frankfurt	eu-frankfurt	cls.internal.tencentcloudapi.com	cls.eu-frankfurt.tencentcloudapi.com
Tokyo	ap-tokyo	cls.internal.tencentcloudapi.com	cls.ap-tokyo.tencentcloudapi.com
Seoul	ap-seoul	cls.internal.tencentcloudapi.com	cls.ap-seoul.tencentcloudapi.com
Jakarta	ap-jakarta	cls.internal.tencentcloudapi.com	cls.ap-jakarta.tencentcloudapi.com
São Paulo	sa-saopaulo	cls.internal.tencentcloudapi.com	cls.sa-saopaulo.tencentcloudapi.com

The following domain names apply to APIs for log upload. Update all other APIs to API 3.0.

Region	Code	Private Network Domain Name	Public Network 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-	na-

		siliconvalley.cls.tencentyun.com	siliconvalley.cls.tencentcs.com
Virginia	na-ashburn	na-ashburn.cls.tencentyun.com	na-ashburn.cls.tencentcs.com
Singapore	ap-singapore	ap-singapore.cls.tencentyun.com	ap-singapore.cls.tencentcs.com
Bangkok	ap-bangkok	ap-bangkok.cls.tencentyun.com	ap-bangkok.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
Jakarta	ap-jakarta	ap-jakarta.cls.tencentyun.com	ap-jakarta.cls.tencentcs.com
São Paulo	sa-saopaulo	sa-saopaulo.cls.tencentyun.com	sa-saopaulo.cls.tencentcs.com

As described in [Uploading Logs via Kafka](#), you can use Kafka Producer SDKs or other Kafka agents to upload logs to CLS. This feature uses the following domain names:

Region	Code	Private Network Domain Name	Public Network Domain Name
Beijing	ap-beijing	bj-producer.cls.tencentyun.com	bj-producer.cls.tencentcs.com
Guangzhou	ap-guangzhou	gz-producer.cls.tencentyun.com	gz-producer.cls.tencentcs.com
Shanghai	ap-shanghai	sh-producer.cls.tencentyun.com	sh-producer.cls.tencentcs.com
Chengdu	ap-chengdu	cd-producer.cls.tencentyun.com	cd-producer.cls.tencentcs.com
Nanjing	ap-nanjing	nj-producer.cls.tencentyun.com	nj-producer.cls.tencentcs.com
Chongqing	ap-chongqing	cq-producer.cls.tencentyun.com	cq-producer.cls.tencentcs.com
Hong Kong (China)	ap-hongkong	hk-producer.cls.tencentyun.com	hk-producer.cls.tencentcs.com
Silicon Valley	na-siliconvalley	usw-	usw-

		producer.cls.tencentyun.com	producer.cls.tencentcs.com
Virginia	na-ashburn	use-producer.cls.tencentyun.com	use-producer.cls.tencentcs.com
Singapore	ap-singapore	sg-producer.cls.tencentyun.com	sg-producer.cls.tencentcs.com
Bangkok	ap-bangkok	th-producer.cls.tencentyun.com	th-producer.cls.tencentcs.com
Frankfurt	eu-frankfurt	de-producer.cls.tencentyun.com	de-producer.cls.tencentcs.com
Tokyo	ap-tokyo	jp-producer.cls.tencentyun.com	jp-producer.cls.tencentcs.com
Seoul	ap-seoul	kr-producer.cls.tencentyun.com	kr-producer.cls.tencentcs.com
Jakarta	ap-jakarta	jkt-producer.cls.tencentyun.com	jkt-producer.cls.tencentcs.com

As described in [Consumption over Kafka](#), you can use Kafka Consumer SDKs or other big data components to consume the data to data warehouses. This feature uses the following domain names:

Region	Code	Private Network Domain Name	Public Network Domain Name
Beijing	ap-beijing	kafkaconsumer-ap-beijing.cls.tencentyun.com	kafkaconsumer-ap-beijing.cls.tencentcs.com
Guangzhou	ap-guangzhou	kafkaconsumer-ap-guangzhou.cls.tencentyun.com	kafkaconsumer-ap-guangzhou.cls.tencentcs.com
Shanghai	ap-shanghai	kafkaconsumer-ap-shanghai.cls.tencentyun.com	kafkaconsumer-ap-shanghai.cls.tencentcs.com
Chengdu	ap-chengdu	kafkaconsumer-ap-chengdu.cls.tencentyun.com	kafkaconsumer-ap-chengdu.cls.tencentcs.com
Nanjing	ap-nanjing	kafkaconsumer-ap-nanjing.cls.tencentyun.com	kafkaconsumer-ap-nanjing.cls.tencentcs.com
Chongqing	ap-chongqing	kafkaconsumer-ap-chongqing.cls.tencentyun.com	kafkaconsumer-ap-chongqing.cls.tencentcs.com
Hong Kong (China)	ap-hongkong	kafkaconsumer-ap-hongkong.cls.tencentyun.com	kafkaconsumer-ap-hongkong.cls.tencentcs.com

Silicon Valley	na-siliconvalley	kafkaconsumer-na-siliconvalley.cls.tencentyun.com	kafkaconsumer-na-siliconvalley.cls.tencentcs.com
Virginia	na-ashburn	kafkaconsumer-na-ashburn.cls.tencentyun.com	kafkaconsumer-na-ashburn.cls.tencentcs.com
Singapore	ap-singapore	kafkaconsumer-ap-singapore.cls.tencentyun.com	kafkaconsumer-ap-singapore.cls.tencentcs.com
Bangkok	ap-bangkok	kafkaconsumer-ap-bangkok.cls.tencentyun.com	kafkaconsumer-ap-bangkok.cls.tencentcs.com
Frankfurt	eu-frankfurt	kafkaconsumer-eu-frankfurt.cls.tencentyun.com	kafkaconsumer-eu-frankfurt.cls.tencentcs.com
Tokyo	ap-tokyo	kafkaconsumer-ap-tokyo.cls.tencentyun.com	kafkaconsumer-ap-tokyo.cls.tencentcs.com
Seoul	ap-seoul	kafkaconsumer-ap-seoul.cls.tencentyun.com	kafkaconsumer-ap-seoul.cls.tencentcs.com
Jakarta	ap-jakarta	kafkaconsumer-ap-jakarta.cls.tencentyun.com	kafkaconsumer-ap-jakarta.cls.tencentcs.com

Concepts

Log

Last updated : 2024-01-20 16:31:37

Log

Logs are record data generated during the running of an application system, such as user operation logs, API access logs, and system error logs. Logs are usually stored in text format on the host where the application system resides. A log corresponding to a system running record may contain one line of text (single-line log) or multiple lines of text (multi-line log).

Logs can be uploaded the CLS via [LogListener](#) or other methods such as APIs or SDKs.

Example of a single-line log:

```
59.x.x.x - - [06/Aug/2019:12:12:19 +0800] "GET /nginx-logo.png HTTP/1.1" 200
368 "http://119.x.x.x/" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_14_5)
AppleWebKit/537.36 (KHTML, like Gecko) Chrome/75.0.3770.142 Safari/537.36" "-"
```

Example of a multi-line log:

```
java.net.SocketTimeoutException:Receive timed out
    at j.n.PlainDatagramSocketImpl.receive0(Native Method) [na:1.8.0_151]
    at
j.n.AbstractPlainDatagramSocketImpl.receive(AbstractPlainDatagramSocketImpl.java:143) [^]
    at j.n.DatagramSocket.receive(DatagramSocket.java:812) [^]
    at o.s.n.SntpClient.requestTime(SntpClient.java:213) [classes/]
    at o.s.n.SntpClient$1.call(^:145) [^]
    at ^.call(^:134) [^]
    at o.s.f.SyncRetryExecutor.call(SyncRetryExecutor.java:124) [^]
    at o.s.f.RetryPolicy.call(RetryPolicy.java:105) [^]
    at o.s.f.SyncRetryExecutor.call(SyncRetryExecutor.java:59) [^]
    at o.s.n.SntpClient.requestTimeHA(SntpClient.java:134) [^]
    at ^.requestTimeHA(^:122) [^]
    at o.s.n.SntpClientTest.test2h(SntpClientTest.java:89) [test-classes/]
    at s.r.NativeMethodAccessorImpl.invoke0(Native Method) [na:1.8.0_151]
```

The main components of a log are as follows:

Component	Description	Example
__TIMESTAMP__	Log time, in the format of a UNIX timestamp in milliseconds	1640005601188
__FILENAME__	Log source file	/data/log/nginx/access.log

__SOURCE__	Log source IP	10.0.1.2
Log body	<p>Log body content, in <code>key:value</code> format (<code>key</code> is the field name, and <code>value</code> is the field value.)</p> <p>If the log extraction mode is full text in a single line or full text in multi lines (the entire raw log is reported without splitting), the entire log is stored in the <code>__CONTENT__</code> field.</p> <p>If the log extraction mode is any other mode (such as splitting with a separator), each part of the raw log corresponds to a <code>key:value</code> pair.</p>	<p>Full text in a single line or full text in multi lines:</p> <p>10.20.20.10;[2018-07-16 13:12:57];GET /online/sample HTTP/1.1;200</p> <p>Other modes:</p> <p>IP: 10.20.20.10 request: GET /online/sample HTTP/1.1 status: 200 time: [2018-07-16 13:12:57]</p>
Metadata	<p>Simple description or categorization of logs. For example, the cluster or container of a TKE log is stored in the <code>key:value</code> format, where <code>key</code> starts with <code>__TAG__</code>.</p>	<code>__TAG__.clusterId:1skzv59c</code>

Log group

A log group (LogGroup) is a collection of multiple logs. In the process of log upload, to increase data read/write efficiency, multiple logs can be packaged into one log group and then sent to CLS.

Logs in the same log group have the same basic information (`__TIMESTAMP__`, `__FILENAME__`, `__SOURCE__`, metadata, etc.).

Log Topic and Logset

Last updated : 2024-01-20 16:32:37

Log topic

A log topic is a basic unit for log data collection, storage, search and analysis on the CLS platform. The massive amounts of logs collected are managed by log topic. For example, you can configure log collection rules and storage time, search for and analyze logs, and download, consume, and ship logs by log topic.

A log topic usually corresponds to an application/service. It is recommended to collect logs of the same application/service from different servers to the same log topic. For example, a payment service (payService) is deployed on dozens of machines and has two types of logs: access logs (access_log) and error logs (error_log). You can create two log topics, payService_access_log_topic and payService_error_log_topic, to collect the two types of logs on the machines respectively. You can use these two log topics to centrally search for and analyze all logs on the machines.

Log topics and applications/services are not strictly in one-to-one mapping relationships. If the log structures of two services are similar, and the logs of the two services need to be analyzed in a centralized manner, you can report the logs of the two services to the same log topic.

Logset

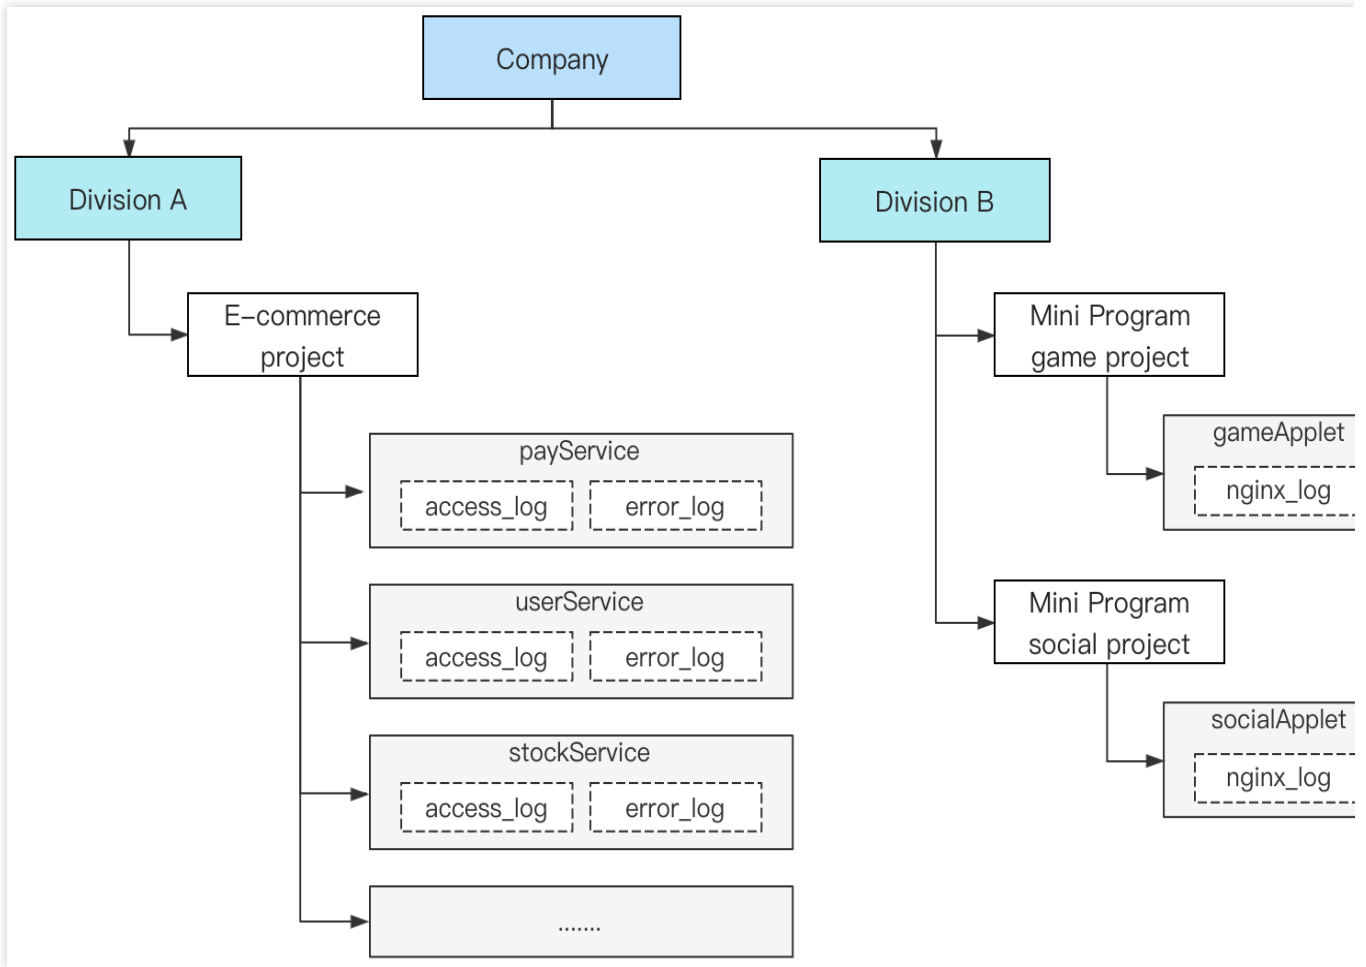
A logset is a class of log topics and can contain multiple log topics. A logset itself does not store any log data, but just makes it easier for users to manage log topics.

A logset usually corresponds to a project/business in a company. You are advised to add the log topics of multiple applications/businesses of a project/business to the same logset. For example, if a company's e-commerce project contains multiple services (payService, userService, stockService, etc.), you can create a logset named e_commerce_logset and add the log topics of these services to this logset. In this way, when the company has multiple projects, the project personnel only need to view the log topics in the logset corresponding to the project, without interference from the log topics of other projects.

Note:

The logset to which a log topic belongs can be specified when the log topic is created and cannot be changed once specified.

Example



As shown in the figure above, the company has two departments:

Department A has an e-commerce project that adopts the microservice architecture, and each service has two types of logs: access logs (access_log) and error logs (error_log).

Department B has two projects: Mini Program game project and Mini Program social project. The two projects adopt a simple technical architecture and each has one type of logs: nginx access logs (nginx_log).

To use CLS to monitor the logs of the applications of the company, you can create the following logsets and log topics:

Logset	Log Topic	Label
e_commerce_logset	payService_access_log_topic	Department: Department A
e_commerce_logset	payService_error_log_topic	Department: Department A
e_commerce_logset	userService_access_log_topic	Department: Department A
e_commerce_logset	userService_error_log_topic	Department: Department A
e_commerce_logset	stockService_access_log_topic	Department: Department A

e_commerce_logset	stockService_error_log_topic	Department: Department A
e_commerce_logset	Department: Department A
gameApplet_logset	gameApplet_nginx_log_topic	Department: Department B
socialApplet_logset	socialApplet_nginx_log_topic	Department: Department B

The labels are used to identify the departments to which the logsets and log topics belong. Combined with permission policies, the personnel in each department can view the data of their own department only.

For department A, the e_commerce_logset logset covers the log topics of all e-commerce services. If other projects are added later, you only need to create a logset.

For department B, although the current technical architecture is relatively simple, with only two types of logs in total, two logsets are created, each with only one log topic, for the following purposes:

Support for subsequent architecture expansion: If the service scale increases and the technical architecture changes to the microservice architecture, you can continue to use the current logset and add log topics to the current logset. In this way, projects do not affect each other.

Flexible response to project adjustment: If the entire department is used as a logset, and a project needs to be changed to an independent department or moved to another department, log adjustment will be difficult because the logset of a log topic cannot be changed directly, and you may need to collect logs again. If each project corresponds to a logset, this situation does not exist, and you only need to adjust the labels corresponding to the logset and log topic for department adjustment.

Machine Group

Last updated : 2024-01-20 16:33:07

Machine group

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

A machine group can contain multiple machines. When an application/service is deployed on multiple machines with the same log file path, the machines can be grouped into a machine group. In this way, you only need to configure the log data collection rule once in the console, and the rule takes effect on all the machines in the machine group in batches.

A machine group can be associated with a log topic. That is, all the logs in the machine group are reported to the same log topic. You can also associate a machine group with multiple log topics. That is, logs of different paths in a machine group are reported to different log topics.

Machine groups can be defined in two ways:

IP address: add an IP address list to a machine group. Then the machines corresponding to the IPs will be automatically added to the machine group.

Label: add labels to machines when installing [LogListener](#). Then the machines with the labels will be automatically added to the machine group.

Examples

An e-commerce project has 6 machines and 3 services (payService, userService, and stockService). The deployment method is as follows:

payService is deployed on 2 machines, and there are 2 log file paths. userService and stockService are deployed on 4 machines, and there are 4 log file paths. Each of the log types (access_log and error_log) of each of the services needs to be reported to an independent log topic.

Machine	Service Deployed	Log File Path
192.168.1.1	payService	/data/log/payService/access_log.log /data/log/payService/error_log.log
192.168.1.2	payService	/data/log/payService/access_log.log /data/log/payService/error_log.log
192.168.1.3	userService,stockService	/data/log/userService/access_log.log /data/log/userService/error_log.log /data/log/stockService/access_log.log /data/log/stockService/error_log.log
192.168.1.4	userService,stockService	/data/log/userService/access_log.log

		/data/log/userService/error_log.log /data/log/stockService/access_log.log /data/log/stockService/error_log.log
192.168.1.5	userService,stockService	/data/log/userService/access_log.log /data/log/userService/error_log.log /data/log/stockService/access_log.log /data/log/stockService/error_log.log
192.168.1.6	userService,stockService	/data/log/userService/access_log.log /data/log/userService/error_log.log /data/log/stockService/access_log.log /data/log/stockService/error_log.log

When deploying LogListener on servers, you can add a label for each machine according to the services running on the machine. See the table below.

Machine	Service Deployed	Label
192.168.1.1	payService	payService
192.168.1.2	payService	payService
192.168.1.3	userService,stockService	userService,stockService
192.168.1.4	userService,stockService	userService,stockService
192.168.1.5	userService,stockService	userService,stockService
192.168.1.6	userService,stockService	userService,stockService

Then create 3 machine groups in the console and define them with labels payService, userService, and stockService. Then, associate the machine groups with log topics and add the corresponding collection configuration. By now, log reporting configuration is completed.

Log Topic	Associated Machine Group	Collection Path
payService_access_log_topic	payService	/data/log/payService/access_log.log
payService_error_log_topic	payService	/data/log/payService/error_log.log
userService_access_log_topic	userService	/data/log/userService/access_log.log
userService_error_log_topic	userService	/data/log/userService/error_log.log
stockService_access_log_topic	stockService	/data/log/stockService/access_log.log

stockService_error_log_topic	stockService	/data/log/stockService/error_log.log
------------------------------	--------------	--------------------------------------

If a service needs to be expanded in the future, you only need to add the service label to the machines to be added so that the new machines can be automatically added to the corresponding machine group for log collection, significantly improving Ops and deployment efficiency.

Segment and Index

Last updated : 2024-01-20 16:33:38

Segmentation


Segmentation explanation

To search for a long log, you usually only need to search for part of the log content. Assume that you need to search for the following complete log that contains `sample` :

```
10.20.20.10;[2018-07-16 13:12:57];GET /online/sample HTTP/1.1;200
```

The full text of the log is long, and the log contains many other content in addition to `sample` . The complete log and the search condition are not identical. Therefore, `sample` cannot be used as the search condition to search for the log. To meet this search requirement, the full text of the log needs to be divided into multiple fragments, and each fragment is called a word, and this process is called "segmentation".

For example, the above sample log can be divided based on symbols `@&()=',';:<>[]{} / \n\t\r` , that is, wherever the symbols occur, segmentation is executed. The segmentation result is as follows:



If the search condition is `sample` , the preceding log after segmentation contains `sample` and therefore is considered to meet the search condition.

The search condition itself is also segmented, such as the 3 types of search conditions in the figure below:

Search condition	Segmented Logs
Search condition A ✓	
Search condition B ✗	
Search condition C ✓	

Search condition A: `\\/online\\/login`

`\\` is used to escape the `/` symbol (this symbol is a reserved symbol of the search syntax and therefore needs to be escaped).

The escaped `/` symbol is a delimiter, so the actual search condition is `online OR login`. A log containing `online or login` is considered to meet the search condition.

The example log provided above **meets** the search condition.

Search condition B: `"/online/login"`

Being enclosed by double quotation marks, the `/` symbol does not need to be escaped.

The content in the double quotation marks is also divided into two words. However, the double quotation marks indicate that only a log that **contains both the two words in the exact order** as that in the search condition is considered to meet the search condition.

The example log provided above does not contain `login` and therefore **does not meet** the search condition.

Search condition C: `"/online/sample"`

The example log provided above contains both `online` and `sample` in the exact order as that in the search condition and therefore is considered to **meet** the search condition.

In most scenarios, if the search condition contains delimiters, you are advised to use double quotation marks to avoid adding escape characters while accurately searching for desired logs. For more information about the search syntax, see [Syntax and Rules](#).

Segmentation configuration

The basis of log segmentation falls into two categories. For configuration details, please see [Configuring Indexes](#).

Delimiter: you can customize the symbols for log segmentation, including English symbols and `\\n\\t\\r`. For example, in the example above, `@&()=',';:<>[]{} / \\n\\t\\r` are delimiters.

Allow Chinese Characters: Chinese is special, and Chinese symbols cannot be used as delimiters, and segmentation according only to symbols often fails to achieve the desired effect. For example, if a log in Chinese is

用户登录失败，密码错误

and you want the log to be searched for by

"登录失败"

, segmentation according to symbols cannot be used to meet the search requirement. In this case, you can set the log to **Allow Chinese Characters** by referring to [Configuring Indexes](#) so that CLS will automatically segment every Chinese character into a separate word.

Index

To enable quick log searches, CLS performs a lot of preprocessing, including segmentation, on the logs uploaded to the platform, and this process is called "indexing". Indexes determine the conditions under which logs can be searched and analyzed. Therefore, before uploading log data, you need to set a reasonable index rule for the corresponding log topic to facilitate subsequent search and analysis. For more information, please see [Configuring Indexes](#).

Index rules take effect only for newly uploaded log data, not historical data. Index creation involves index traffic and index storage, resulting in certain usage costs.

Topic Partition

Last updated : 2024-01-20 16:34:10

Overview

Partition is the smallest read/write unit of 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.

Note:

Topic partitioning is a complex concept. You are advised to use the [Auto Split](#) function in actual practice. CLS automatically adjusts topic partitioning based on the volume of log data.

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, ff). 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
1	Read-write	[00000000000000000000000000000000, 7fffffffffffffffffffffffffffffffffffffff)

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 a partition is split, the original partition becomes read-only, which only allows data consumption, but not data writing. The new partition is readable and writable and covers the range of the original partition.

Data Processing

Last updated : 2024-12-20 16:02:12

Data Processing

Data processing refers to the process of filtering, cleaning, desensitizing, enriching, and distributing log data to the target log topic. It can be understood as log ETL (Extract-Transform-Load).

Source Log Topic: Input of data processing task.

Target Log Topic: Output of data processing task.

Target name: Custom target topic name, which improves the readability of the target topic (business attribute) and is used in the function `log_output("alias")` when outputting logs to the specified target topic. Data processing tasks must have an output target topic, otherwise, the task cannot be created.

DSL Processing Functions: DSL (Domain Specific Language) is a log data processing function developed by CLS for the requirements of log ETL. The functions are simple and easy to use, with high processing performance. The underlying layer is implemented based on Flink, which can process logs in real time.

Timed SQL Analysis

Timed SQL Analysis refers to the process of periodically querying log data (supporting retrieval and SQL) based on the specified time window and saving the query results to the target log topic.

Source Log Topic: Input of scheduled SQL task.

Target Log Topic: Output of timed SQL task.

Scheduling Range: Time range of logs that supports query, for example, the log data from 00:00:00 on January 1, 2023 to 00:00:00 on March 31, 2023.

Scheduling Period: Periodic query, with the value ranging from 1 to 1,440 minutes. If daily reports need to be generated, it can be configured as 1,440 minutes.

SQL Time Window: Time window for specifying query statements. When it is used with the scheduling cycle, you can get tumbling window and sliding window.

Rolling window: Non-overlapping query window. For example, if the scheduling period is 60 minutes, the SQL time window is 60 minutes. Typical scenario: hourly report.

Sliding Window: Query windows with overlapping. For example, the scheduling cycle is 1 minute and the SQL time window is 60 minutes. Typical scene: Create a sequence diagram for active users within 1 hour, with a time axis granularity of 1 minute.

Delayed Execution: For query of the delayed time in the Advanced Settings of the console, with the value ranging from 60 to 120 seconds. There is usually a delay when the log index is generated, and the log cannot be queried

before the index is generated. Therefore, set a delay of 60 seconds for query, and the index is already generated when the time elapses (99.9% of the index data will be generated within 5 seconds).

Limits

Basic Resources

Last updated : 2024-12-16 14:56:26

Quota Limits

Note:

If the current quotas cannot meet the needs, [submit a ticket](#) to apply for more quotas.

Related word notes are as follows:

Related Words	Description
Logset	Up to 100 logsets can be created for each account in each region.
	Up to 500 log topics and metric topics can be created for each logset.
	The name can contain 1 to 255 characters. Supported characters: uppercase letters, lowercase letters, digits, underscores (<code>_</code>), and hyphens (<code>-</code>).
Log Topic	The length of topic names ranges from 1 to 255 characters, allowing characters a-z, A-Z, 0-9, <code>_</code> , and <code>-</code> .
	Up to 50 partitions can be created under each log topic, and up to 10 partitions can be created when a new log topic is added.
	Each topic can be bound with up to 100 machine groups.
	Log topics of standard storage can be stored for 1-3600 days or permanently.
	Log topics of infrequent access storage can be stored for 7-3600 days or permanently.
	The log settlement function can be enabled only when the log topic of standard storage type is stored for more than 7 days.
Metric Topic	The length of topic names ranges from 1 to 255 characters, allowing characters a-z, A-Z, 0-9, <code>_</code> , and <code>-</code> .
Machine Group	Each account can create up to 200 machine groups in different regions.
	Each IP machine group can be bound with up to 200 machine IPs.
	Each label machine group can be bound with up to 50 Labels.

Dashboard

Each account can create up to 100 dashboard subscription tasks.

Log Collection

LogListener Limits

Last updated : 2024-01-20 16:34:39

This document describes LogListener's capabilities and limits of file collection, configuration, resources, performance, and error handling related to data collection.

File Collection Capabilities and Limits

Item	Capability and Limit
File encoding	Supports UTF8 and GBK codec. Note: only LogListener v2.6.2 and above support GBK codec.
Soft link	Supported.
Size of a single log	The size of a single-line log is limited to 512 KB. If the size of a single-line log exceeds 512 KB, the log will be truncated and only the first 512 KB is retained. For a multi-line log, after it is divided according to the first-line regular expression, the maximum limit for a single log is 1 MB.
Regular expression	Perl-compatible regular expressions are supported.
First log collection behavior	LogListener supports full/incremental collection policies: Full collection: after LogListener is started for the first time after installation, all logs that meet requirements are collected, including files that have not been written. Incremental collection: after LogListener is started for the first time after installation, existing files will be collected from the latest position. Note: Full/Incremental collection policies are available starting from LogListener v2.6.2.
Log file rotation	Supported. (It is recommended that the file name after rotation not be matched by the wildcard collection path.)
Collection behavior when log parsing fails	We recommend you enable the feature of uploading parsing-failure logs. If the feature is enabled, parsing-failure logs will be uploaded to the preset index in the format of full text in a single line. Otherwise, the logs will be discarded.
File opening	LogListener opens files when reading files for log collection and closes the files when the reading is completed.

Checkpoint Management Capabilities and Limits

Item	Capability and Limit
Checkpoint storage location	The default save path is the <code>data</code> directory in the LogListener installation directory. You can customize a save path in the <code>loglistener.conf</code> file.
Checkpoint storage policy	LogListener stores two copies of checkpoint metadata: One copy records the information only of checkpoints where upload is completed, and the information is persisted to a disk in real time. The other copy records the information of checkpoints where upload is completed and checkpoints where upload is not completed, and the information is periodically persisted to a disk. Persistence takes precedence when the program exits.

Resource and Performance Capabilities and Limits

Item	Capability and Limit
Default CPU resource limits	LogListener does not limit CPU resources by default. You can configure CPU resource limits in the configuration file. Without CPU resource limits, the current implementation architecture of LogListener can achieve a maximum CPU usage of 110% (up to 100% for a single business thread and about 10% for a control thread).
Default memory resource limits	LogListener's default memory threshold is 2 GB. You can change the threshold to 300 MB or higher as needed.
Default bandwidth resource limits	LogListener does not limit bandwidth resources by default. You can configure bandwidth resource limits in the <code>loglistener.conf</code> file.
Resource overrun handling policy	If LogListener overruns CPU and memory resources for more than 5 minutes, LogListener will be forced to restart automatically.
Log compression	Collected logs are shipped as compact logs by default. You can modify log compression configuration in the <code>loglistener.conf</code> file.
Number of monitored directories	The recommended maximum number of monitored directories is 5,000. If this limit is exceeded, log collection may fail.
Number of monitored files	The recommended maximum number of monitored files is 10,000. If this limit is exceeded, log collection may fail.

Error Handling

Item	Capability and Limit
Network error handling	Except exceptions requiring special handling (such as log topic deletion), all other errors (such as network exceptions, timeout, frequency control, and overdue payment) will be retried.
Maximum retry timeout duration	If data fails to be sent after retrying for more than 1 hour, the system discards the data. The default behavior is to retry at an interval, and the retry interval becomes longer and longer until the maximum retry timeout duration is exceeded.
Maximum number of retries	The maximum number of retries can be set in the <code>loglistener.conf</code> configuration file and is not set by default. By default, the system retries until the maximum retry duration is exceeded, and then discards the data. If the maximum number of retries is set, the system retries until the maximum number of retries is exceeded, and then discards the data.

File Collection Rules

Item	Capability and Limit
Log upload policy	LogListener automatically aggregates and uploads logs of the same file when any of the following aggregation conditions is met: 10,000 logs, the total logset size reaching 1 MB, or the log collection time exceeding 3 seconds.
File collection handling policy	A single target file (a file that can be matched by the collection path) can be uploaded to only one log topic. The same file cannot be matched by the collection paths of multiple topics. If you need to upload a file to multiple topics, use soft links. You can create different soft links for the same target file so that it can be collected to different topics via different soft links.
Log collection delay	In the case of real-time collection, data collection, transmission, and storage to disks will be completed within 1 minute, and the logs are retrievable in the console. If the log volume is large, or LogListener can use only a limited amount of resources, there will be some collection delay.

Machine Group Related Rules

Item	Capability and Limit
Machine group logic	At present, machine groups are divided into the following two categories, and their

	<p>usage methods are independent of and incompatible with each other. If the two usage methods are used together, LogListener will not be able to pull the correct collection configuration, resulting in no collection.</p> <p>IP machine group: a machine IP must be manually added to the machine group in the console, and the <code>group_label</code> field in <code>loglistener.conf</code> on the corresponding machine must be empty.</p> <p>Label machine group: the machine group label is set in the console, and the <code>group_label</code> field in <code>loglistener.conf</code> on the corresponding machine must be set to the same label.</p>
Relationships between machine groups and log topics	<p>A single log topic can be bound to multiple machine groups.</p> <p>A single machine group can be bound to multiple log topics.</p>
Relationships between machine groups and machines	<p>A single machine can be added to multiple machine groups.</p> <p>For an IP machine group, there is no limit to the number of machine groups that a machine can join.</p> <p>For a label machine group, the maximum number of machine groups that a machine can join is 20.</p>
Limits on the labels of label machine groups	<p>Currently, the label of a label machine group can contain up to 32 characters.</p> <p>Up to 20 labels can be configured for a single label machine group.</p>

Collection Path/Collection Blocklist Usage

Item	Capability and Limit
Collection blocklist	<p>This feature is used to specify the content to ignore in the collection path. Currently, collection blocklists support two modes:</p> <p>FILE mode: specify the files to ignore in the collection path.</p> <p>PATH mode: specify the subdirectories to ignore in the collection path. Wildcard filtering is supported.</p> <p>Notes:</p> <p>The FILE and PATH modes can be used together.</p> <p>A collection blocklist is to exclude content from the collection path, and for both the FILE and PATH modes, the specified path to exclude must be a subset of the collection path.</p>

Metric Collection

Last updated : 2024-12-16 14:58:40

The restrictions for metric collection by using metric topics are as follows:

Restriction Item	Description
Metric name	Supports English letters, numbers, underscores, and colons. It should conform to the regular expression <code>[a-zA-Z_:][a-zA-Z0-9_]*</code> .
Label name	Supports English letters, numbers, and underscores. It should conform to the regular expression <code>[a-zA-Z_][a-zA-Z0-9_]*</code> .
Label value	No special restrictions, supporting all types of Unicode characters.
Sample value	A float64 type value
Sample timestamp	Millisecond precision
Metric upload frequency control	Related to the number of partitions. The maximum number of partitions is 50, and the partition write request limit for individual topics is 500 QPS. It is recommended to enable partition auto-split during metric topic creation .
Metric upload flow control	Related to the number of partitions. The maximum number of partitions is 50, and the partition write traffic limit for individual topics is 5 MB/s. It is recommended to enable partition auto-split during metric topic creation .

Search and Analysis

Last updated : 2024-12-16 14:59:44

Log Topic

The restrictions for log topic search and analysis are shown in the table below:

Type	Restriction Item	Description
Index	Number of Fields	Up to 300 fields can be added for a single log topic key-value index.
	Field Name	Letters, digits, and special characters (except for <code>*\\",</code>) are supported. The name cannot start with <code>_</code> (except for the <code>__CONTENT__</code> field). It is not allowed to contain JSON parent-child fields at the same time, such as <code>a</code> and <code>a.b</code> .
	Field Hierarchy	When the key-value index for multi-level JSON is configured, there shall be no more than 10 levels for the Key, such as <code>a.b.c.d.e.f.g.h.j.k</code> .
	Delimiter	Only English symbols, <code>\n\t\r</code> and escape characters <code>\</code> are supported.
	Field Length	After the field is enabled for statistics, only the first 32,766 bytes are involved in SQL operation. The excess part cannot execute SQL, but the log will still be completely stored.
	Token Length	After the field is tokenized, only the first 10,000 characters of a single word are involved in search. The excess part cannot be searched, but the log will still be completely stored.
	Numeric Field Precision and Range	The data range supported by the long type field is $-1E15 \sim 1E15$. Data outside of this range may lose accuracy or do not support search. The data range supported by the double type field is $-1.79E+308 \sim +1.79E+308$. If the number of floating-point coding bits exceeds 64, precision will be lost. Suggestions on index configuration of ultra-long numerical field: If the field does not need to be searched by value range comparison, it can be stored as text type. If the field needs to be searched by value range comparison, it can be stored as double type and some precision may be lost.
	Activation Mechanism	Index configuration is effective only for newly collected data. The index rules after editing are effective only for newly written logs and existing data will not be updated. To update existing data, rebuild the index.

	Modifying index configuration	When the index configurations are created, modified and deleted, a single user has a maximum of 10 tasks being executed at the same time. If the number is exceeded, other tasks need to wait for the previous ones to completed and a single task usually takes no more than 1 minute.
	Reindexing	<p>A single log topic allows to run only one index rebuilding task at a time. A single log topic can have 10 index rebuilding task records at most. An index task can only be created after the task records that are no longer needed are deleted. For logs within the same time range, the index can only be rebuilt once. You need to delete previous task records before rebuilding.</p> <p>The write traffic of logs corresponding to the selected time range cannot exceed 5 TB.</p> <p>The index rebuilding time range is subject to the log time. When there is a deviation of more than 1 hour between the log uploading time and the index rebuilding time range (for example, a 02:00 log is uploaded to CLS at 16:00, and the log index from 00:00 to 12:00 is rebuilt), it will not be rebuilt and cannot be retrieved subsequently. When a new log is submitted to the time range for which the index has been rebuilt, the index will not be rebuilt and the log will not be searched later.</p>
Query	Sentence Length	Search and analysis statements support up to 12,000 characters.
	Query Concurrency	A single log topic supports 15 concurrent queries, including search and analysis queries.
	Fuzzy search	Fuzzy prefix search is not supported, e.g. search for error by <code>*error</code> .
	Phrase search	Wildcards in phrase search can only match with 128 words that meet the conditions, and all logs containing these 128 words will be returned. The more accurate the specified word is, the more accurate the query result will be.
	Logical group nesting depth	<p>When parentheses are used for logical grouping of search conditions under CQL syntax rules, up to 10 layers can be nested. This limitation is not applicable to Lucene syntax rules.</p> <p>For example, <code>(level:ERROR AND pid:1234) AND service:test</code> is 2-layer nesting and can perform search normally. The statement below is 11-layer nesting, and an error will be reported when the search is performed.</p> <pre>status:"499" AND ("0.000" AND (request_length:"528" AND ("https" AND (url:"/api" AND (version:"HTTP/1.1" AND ("2021" AND ("0" AND (upstream_addr:"169.254.128.14" AND (method:"GET" AND (remote_addr:"114.86.92.100")))))))))))</pre>
	Memory usage	The server memory occupied by statistical analysis each time cannot exceed 3 GB.

	(analysis)	Generally, this limitation may be triggered when group by, distinct() and count(distinct()) are used, because there are too many values in the statistical field after deduplication through group by or distinct(). It is recommended to optimize the query statement and use fields with fewer values to perform grouping statistics on data, or use approx_distinct() to replace count(distinct()).
	Query Result	If raw logs are queried, a maximum of 1,000 raw logs can be returned for each query.
		If statistical analysis results are queried, a maximum of 100 results can be returned for each query by default. If a LIMIT clause is used, a maximum of 1 million results can be returned for each query.
		The maximum data packet returned for the query cannot be greater than 49 MB. GZIP compression (Header Accept-Encoding:gzip) can be enabled if the corresponding API is used.
	Timeout Time	The timeout time for a single query is 55 seconds, including search and analysis.
	Query Latency	The delay from log reporting to being available for search and analysis is less than 1 minute.
Download	Log Count	The maximum number of logs downloaded at a time is 50 million.
	Task Quantity	For a single log topic, a maximum of 2 tasks can be in the File Generating status, and other unfinished tasks are in the Waiting status and queuing for execution. A single log topic can have up to 1,000 tasks at the same time, including completed tasks in the status of File Generated.
	File Retention Duration	The generated log files will be retained for only 3 days.
Related external data	Quantity Limit	A single log topic can be associated with up to 20 external data.
	Query Timeout	When an external database is queried, the timeout period is 50 seconds.
	MySQL Version	It is compatible with MySQL 5.7, 8.0 and later versions. MySQL 5.6 has not received a complete compatibility test. It is necessary to actually test whether the SQL statements are executed properly.
	CSV file size	The file shall not exceed 50 MB and compression is not supported.

Metric topic

The restrictions for metric topic search and analysis are shown in the table below:

Restriction Item	Description
Query Concurrency	A single metric topic supports up to 15 concurrent queries.
Query data volume	A single query can involve up to 200,000 time series, with a maximum of 11,000 data points per time series in the query results.
Timeout Time	The timeout time for a single query is 55 seconds, including search and analysis.

Monitoring Alarm

Last updated : 2024-12-16 15:01:53

The restrictions for monitoring and alarm are shown in the table below:

Type	Restriction Item	Description
Alarm Policy	Monitoring Object	A maximum of 20 log topics can be selected for a single alarm policy, and cross-regio
	Query Statement	The limitations are the same as Search and Analysis , including statement length and
	Number of Query Statement	When each execution statement uses the same log topic, up to 3 execution statemen When each execution statement selects log topics separately, up to 10 execution sta
	Execution statement query time range	The maximum time range is the last 24 hours.
	Execution Cycle	The minimum execution cycle is 1 minute and the maximum execution cycle is 1,440
	Trigger by group	The execution statement results can be divided into 1,000 groups at most. The 1,000 Note: Triggering by group is an optional function, and this limitation is not involved if it is no
	Trigger condition judgment times	When the result of the execution statement is multiple values (after grouping), the res
Alarm Type	Phone	The same phone number can receive up to 1 phone call alarm within 30 seconds, up The same phone number can receive a maximum of 3 phone call alarms from 22:00 1 The above limitations are not limited to phone call alarms generated by the Cloud Log To receive alarms through phone calls, it is recommended to add the following phone 02155931000, 02155931001, 02155931002, 02155931003, 02155931004
	WeCom, DingTalk, and Lark	Each Chatbot cannot send more than 20 messages per minute.

	Custom webhooks	Public network access should be supported. Private network addresses are not supported. When the status code returned by the interface is 200, it indicates that callback is successful.
Alarm Records	Storage Period	30 days
Alarm Notification	Alarm Notification Variables	<p>The following alarm notification variables are subject to data entry count or data size:</p> <ul style="list-style-type: none"> <code>Message</code>: The maximum size is 10 KB. <code>QueryResult</code>: Each execution statement can contain up to 1,000 data entries, and the total size is 10 KB. <code>TriggerResult</code>: Each execution statement can contain up to 1,000 data entries, and the total size is 10 KB. <code>QueryLog</code>: Each execution statement can contain up to 10 data entries, and the total size is 10 KB. <code>AnalysisResult</code>: The maximum data size of each multidimensional analysis result is 10 KB. <p>Related raw logs: Each statement can contain up to 10 data entries, depending on the number of fields and their proportions.</p> <p>Top 5 fields and their proportions: Each statement can contain up to 5 data entries.</p> <p>Custom search and analysis (SQL): Each statement can contain up to 1,000 data entries, and the total size is 10 KB.</p> <p>Custom search and analysis (search): Each statement can contain up to 10 data entries, and the total size is 10 KB.</p>

Data Processing

Last updated : 2024-12-17 17:55:46

The restrictions for data processing are shown in the table below:

Type	Restriction Item	Description
Data Processing	Number of Data Processing Tasks	One log topic can create 50 data processing tasks.
	Number of Target Topics	One data processing task can have up to 1000 target topics.
	The number of debugging data	By default, 100 pieces of logs are loaded from the source topic for you to debug the DSL function.
	Processing Capability	10 ~ 20 MB/s/partition (uncompressed data), with partition referring to the that of the source log topic.
	Processing Task Latency	Up to 99.9% of the data will be processed in the data processing module within 1 second.
	Impact of starting and stopping processing tasks on data	When the processing task is restarted within 12 hours after it is suspended, processing will be started from the data when it was suspended. When the processing task is restarted after 12 hours, processing will be started from the latest data.
	Cross-region/Cross-account	Not supported.
	Processing History Data	Not supported.
Scheduled SQL analysis	Number of scheduled SQL tasks	One log topic can be associated with 10 scheduled SQL tasks. (The log topic can be the source log topic of the task or the target log topic.)
	Query Concurrency	A single log topic supports 15 concurrent queries (Scheduled SQL, search and analysis, Alarms, Dashboards).
	Viewing the query result	<p>When the query results are statistical analysis results, 100 results will be returned at a time by default. When SQL LIMIT syntax is used, a maximum of 1 million results will be returned at a time.</p> <p>The maximum size of the returned data packet of query results is 49 MB, and gzip compression can be enabled when API is used.</p>

Query Timeout Duration	The timeout period for a single query is 55 seconds.
SQL time window	1 minute to 7 days.
Scheduling cycle	1 minute to 24 hours.
Custom timestamp	The minimum custom timestamp interval is 15 seconds. For example, 15:00:00, 15:00:15, and 15:00:30.
SQL execution delayed	The default value is 60 seconds and the maximum value is 120 seconds. It takes about 5 seconds to generate 99.9% of the log indexes, and 60 seconds in extreme cases. Therefore, set a delay for execution to ensure that the query can successfully obtain index data.
Cross-region/Cross-account	Not supported.

Shipping and Consumption

Last updated : 2024-12-17 17:57:09

The restrictions are shown in the table below:

Type	Restriction Item	Description
Ship to COS	Number of Tasks Submitted	For each log topic, a maximum of 10 shipping to COS tasks can be configured.
	Deliverable Data Scope	Data (including historical and real-time data) of log topics in the lifecycle.
	Task Submission Delay	The latency of shipping to COS is about 60 seconds. Note: If you set the shipping interval to be 5 minutes, logs from 12:00 to 12:05, meeting the interval of 5 minutes, will be shipped as a whole to COS in a batch. With a latency of 60 seconds, this batch of logs will arrive at COS around 12:06.
	Trigger Delivery	One shipping will be triggered when a partition meets any of the following conditions: Size of logs to be shipped (accumulated log data). Range: 5 to 256 MB (uncompressed). Shipping interval. Range: 5 to 15 minutes.
	Cross-account/Cross-region	Cross-account: supported Cross-region: not supported
Ship to CKafka	Number of Tasks Submitted	For each log topic, one shipping to CKafka task can be configured.
	Deliverable Data Scope	At present, only real-time log data can be shipped, and historical data is not supported.
	Task Submission Delay	The latency of shipping 99.9% of logs to the CKafka are within 5 seconds. (The target CKafka functions properly, has a throughput capacity sufficient to process messages from the CLS and the disk is not full or throttled). Note: If the timestamp (value of <code>_TIMESTAMP_</code>) of a log collected to CLS is 12:00:00, the log will arrive at CKafka before 12:00:05.

	Cross-account/Cross-region	Cross-account: supported Cross-region: not supported
Consumption over Kafka	Versions	Kafka protocol version supported: 1.1.1 and earlier.
	Number of consumer groups	Each log topic can be configured with 20 consumer groups.
	Authentication Method	Only SASL_PLAINTEXT is supported.
	Consumable data range	Once this feature is enabled, consumption will start from the latest data. Data will be retained for 2 hours. If data is not consumed within 2 hours, it will not be saved. Note: For example, if you enable consumption over Kafka protocol at 12:00, you will consume data from 12:00. Until 15:00, you can consume data from 13:00 to 15:00. Data before 13:00 cannot be consumed.
	Consumption latency	The consumption latency of 99.9% log data is within 5 seconds. Note: If the timestamp (value of <code>_TIMESTAMP_</code>) of a log collected to CLS is 12:00:00, the log will arrive at the consumer before 12:00:05.
Cross-account/Cross-region	Cross-account: Account A can be a collaborator of account B to consume the log topics of account B. Permissions need to be set for the collaborator. For details, see the permission configuration . Cross-region: Cross-region consumption can be implemented through public network.	