

Log Service

Product Introduction

Product Introduction

The Log Service (or Log for short) is an all-in-one service for log-type data that allows you to quickly complete log data collection, consumption, shipping, query, and analysis without the need for development. Log Service can help you increase O&M efficiency, and build processing capabilities, that help you easily handle massive log volumes.


Core functions include:

Real-time log collection and consumption (LogHub)

Functions:

- The LogHub function of Log Service enables access to real-time log data (such as Metric, Event, BinLog, TextLog, and Click data) through ECS, containers, mobile terminals, open source software, and JS.
- A real-time consumption interface is provided for interconnection with real-time computing and service systems.

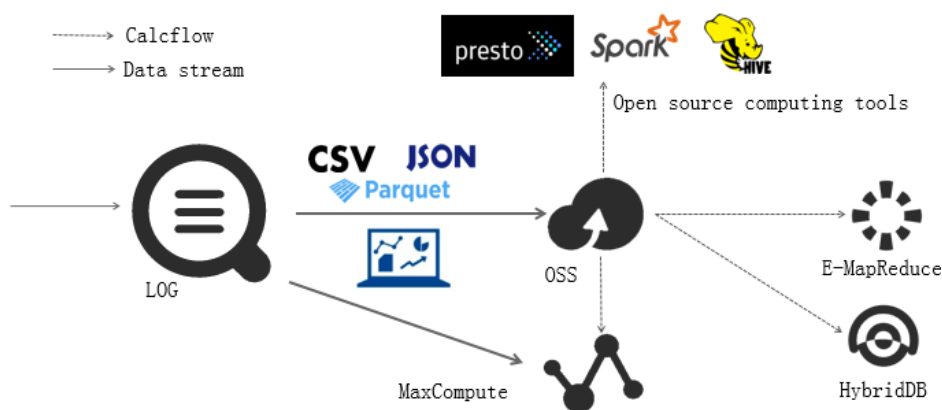
Purposes: ETL, stream computing, monitoring and alarm, machine learning, and iterative computing



LogShipper

Stable and reliable log shipping. Ships LogHub data to storage services for storage and big data analysis. Supports compression, user-defined partitions, and various storage methods such as row and column storage.

Purposes: data warehouse + data analysis, audit, recommendation system, and user profiling

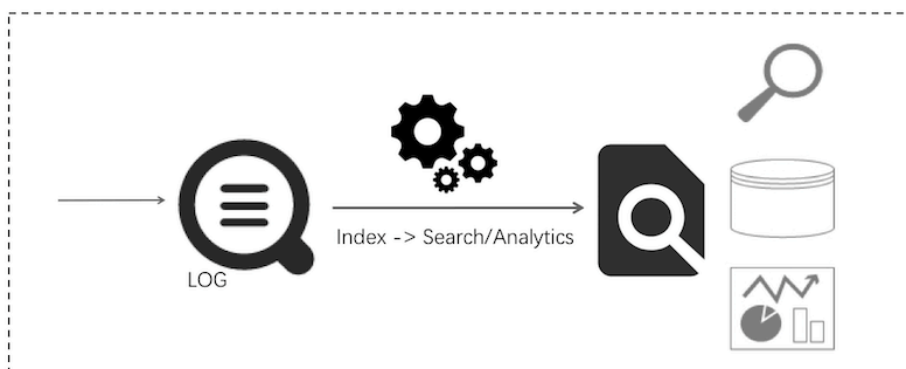


Query and real-time analysis (Search/Analytics)

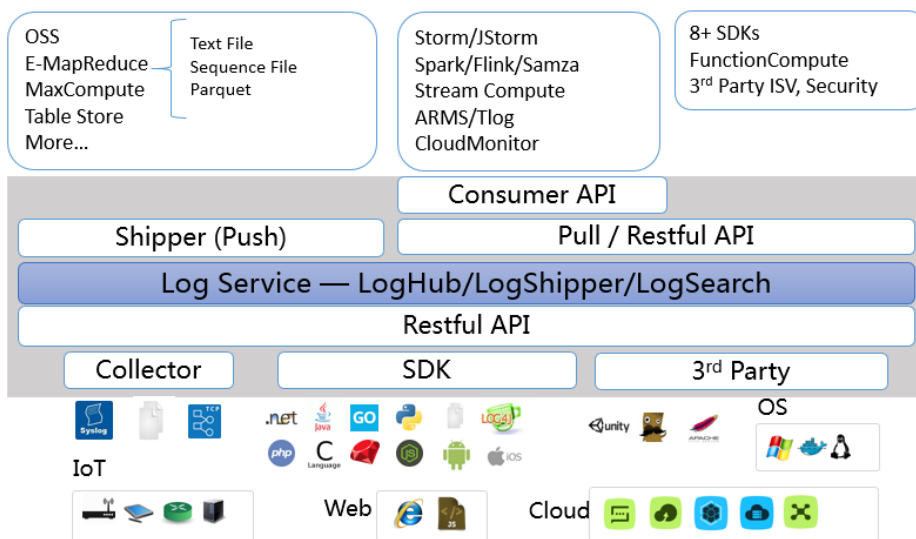
Real-time analysis data indexing and querying.

- Search: keyword, fuzzy match, context, and range
- Analytics: various means such as SQL aggregation
- Visualized: dashboard and report functions
- Interconnection: Grafana and JDBC/SQL92

Purposes: DevOps/online O&M, real-time log data analysis, security diagnosis and analysis, and operation and customer service systems



The Log Service system architecture is shown in the following diagram.



Logtail

Logtail helps you quickly collect logs through the following features:

- Non-invasive log collection based on log files
 - Only read files.
 - Unobtrusive during reading process.
- Secure and reliable
 - Supports file rotation, so data is not lost.
 - Supports local caching.
 - Provides network exception retry mechanism.
- Convenient management
 - Web client.
 - Visualization configuration.
- Comprehensive self-protection
 - Real-time monitoring of process CPU and memory.
 - Consumption and restrictions on CPU/memory usage.

Frontend servers

Frontend machines are built using LVS+Nginx. Its features are as follows:

- HTTP and REST protocols
- Horizontal scaling
 - Support horizontal scaling When traffic increases
 - Frontend machines can be quickly added to improve processing capabilities.
- High throughput, low latency
 - Pure asynchronous processing, a single request exception will not affect other requests.

- Lz4 compression is adopted to increase the processing capabilities of individual machines and reduce network bandwidth consumption.

Backend servers

The backend is a distributed process deployed on multiple machines. It provides real-time Logstore data persistence, indexing, query, and shipping to OSS. The features of the overall backend service are as follows:

- High data security
 - Each log you write is saved in triplicate.
 - Data is automatically recovered in case of any disk damage or machine downtime.
- Stable service
 - Logstores automatically migrate in case of a process crash or machine downtime.
 - Automatic server load balancing ensures that traffic is distributed evenly among different machines.
 - Strict quota restrictions that prevent abnormal behavior of a single user from affecting other users.
- Horizontal scaling
 - Horizontal scaling is performed using shards as the basic unit.
 - You can dynamically add shards as needed to increase throughput.

Managed security service

- Great accessibility allows you to set up and connect to the service in five minutes, and use Agents to collect data in any network environment.
- LogHub has all the functions of Kafka, and provides complete functional data, such as monitoring and alarms, and supports auto scaling (by PB/day). The use cost is less than 50% of the self-development cost.
- LogSearch/Analytics provide the query saving, dashboard, and alarm functions. The use cost is less than 20% of the self-development cost.
- **More than 30 Access Methods**, seamless interworking with cloud products (OSS, E-MapReduce, MaxCompute, Table Store, MNS and CDN) and open-source software (e.g. Storm and Spark).

Rich ecosystem

- LogHub supports over 30 collectors, including LogStash and Fluent, and can be easily connected using embedded devices, web pages, servers, and programs. It can also be interconnected with consumption systems such as Spark Streaming, Storm, and CloudMonitor.
- LogShipper supports a variety of data formats (including TextFile, SequenceFile, and Parquet)

and user-defined partitions. The data can be directly used by storage engines such as Presto, Hive, Spark, Hadoop, E-MapReduce, and HybridDB.

- LogSearch/Analytics have complete syntaxes and are compatible with SQL-92. The JDBC protocol for interconnection with Grafana will be provided.

Real-time processing

- LogHub: Data can be used as soon as being written. Logtail (data collection agent) can collect and transfer data to the server side within one second (in 99.9% cases).
- LogSearch/Analytics: Data can be searched and analyzed as soon as being written. When multiple search criteria are used, billions of data pieces can be searched within one second. When multiple aggregation conditions are used, hundreds of millions of data pieces can be analyzed within one second.

Complete API/SDK

- The service supports user-defined management and secondary development.
- All functions can be implemented using APIs/SDKs. SDKs for multiple languages are provided to facilitate service management.
- The search and analytic syntaxes are simple (compatible with SQL-92). The interfaces can be used to easily interconnected with the common open software environment (solution for interconnection with Grafana will be available soon).


Typical Log Service application scenarios include data collection, real-time computing, data warehousing and offline analysis, product operation and analysis, and O&M and management. The following lists the typical application scenarios. For more details about referrers and their configuration, see [Best Practices](#).

Data collection and integration

The LogHub function of Log Service enables access to massive real-time log data (including Metric, Event, BinLog, TextLog, and Click data) at low costs.

Advantages of the solution:

- Easy to use: Over 30 real-time data collection methods are provided for you to quickly set up your platform. The powerful configuration and management capabilities can ease O&M workload. Nodes are available across China and the rest of the world.
- Highly elastic: It helps easily address traffic peaks and service growth.



ETL/Stream Processing

LogHub can work with all types of real-time computing and services, provides complete progress monitoring and alerting functions, and supports SDK/API-based custom consumption.

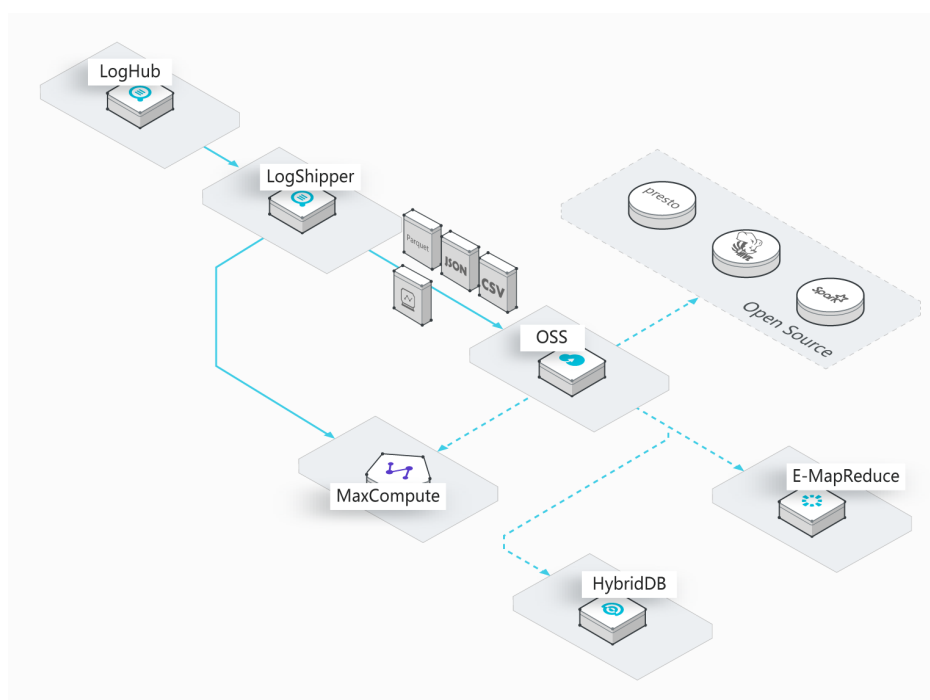
- Easy to operate: It provides various SDKs and programming frameworks and can be seamlessly interconnected with various stream computing engines.
- Rich functions: Various monitoring data and alarm postponing are provided.
- Elastic: PB-grade elasticity and zero latency.



Data Warehouse

LogShipper enables data in LogHub be shipped to storage services and stored in a range of formats such as compression, custom partition, and row/column.


- Massive data: The amount of data is not limited.
- Rich storage formats: Various storage formats are supported, such as row storage, column storage, and TextFile.
- Flexible configuration: Configurations such as user-defined partitions are supported.



Log search and analytics

Log search and analytics support real-time searching data in LogHub, and provide rich means of inquiry: keywords, fuzzy, context, scope, SQL aggregation.

- Strong real-time performance: Write after the query.
- Massive low cost: Support PB/Day indexing capacity, and the cost is 15% of self-built program.
- Strong analytical skills: Support a variety of query means and SQL for aggregation analysis, and provide visualization and alarm function.



Basic concepts

Log

Log is an abstraction of system changes during the running process. The log content is a time-ordered collection of some operations and the corresponding operation results of specified objects. LogFile, Event, BinLog, and Metric data are different carriers of logs. In LogFile, every log file is composed of one or more logs, and every log describes a single system event. The log is the minimum data unit processed in Log Service.

Log group

A log group is a collection of logs and is the basic unit for writing and reading.

Log topic

Logs in a Logstore can be classified by log topics. You can specify the topic when writing and querying logs.

Project

The project is the resource management unit in Log Service and is used to isolate and control

resources. You can manage all the logs and the related log sources of an application by using projects. Projects manage the information of all your Logstores and the log collection machine configuration, and serve as the portals where you can access the Log Service resources.

Logstore

The Logstore is a unit in Log Service for the collection, storage, and query of log data. Each Logstore belongs to a project, and each project can create multiple Logstores.

Shard

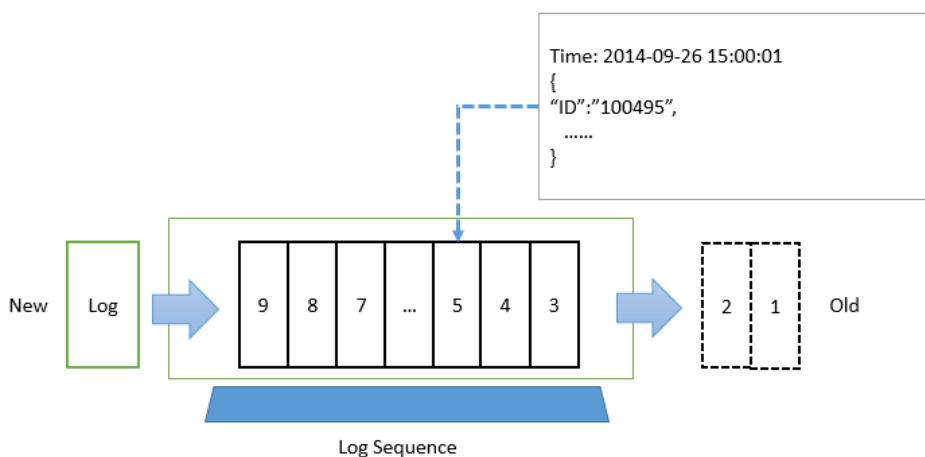
Each Logstore is divided into several shards and each shard is composed of MD5 left-closed and right-open intervals. Each interval range does not overlap with others and the total range of all the intervals is the entire MD5 value range.

Half a century ago, the term “log” was associated with a thick notebook written by a ship captain or operator. Nowadays, with the advent of computers, logs are generated and used everywhere. Servers, routers, sensors, GPS devices, orders, and various IoT devices describe the world we live from different angles by generating and using logs. With the computing power, we continuously update our recognition to the whole world and system by collecting, processing, and using logs.

What is a log?

Consider the example of a ship captain’s log. In addition to a recorded timestamp, a log can contain almost all sorts of information, including a text record, an image, weather conditions, and the sailing course. After centuries passed, now the “ship captain’s log” has been expanded to various areas such as orders, payment records, user accesses, and database operations.

The reason why logs are widely used and enduring is that logs are the simplest storage abstraction. Logs are a collection of chronological records that can only be added. The following figure is what logs (time-series data) look like.



We can add a record to the end of a log and read the log records from left to right. Each record has a unique log record number with a sequence.

The log sequence is determined by "time". From the preceding figure, we can see that the log time sequence is from right to left. The new event is recorded, and the old event is gradually out of sight. But a log records when and what happened. This is the foundation of recognition and reasoning, no matter to computers, humans, or the whole world.

Logs in Log Service

Log is an abstraction of system changes during the running process. The log content is a time-ordered collection of some operations and the corresponding operation results of specified objects. LogFile, Event, BinLog, and Metric data are different carriers of logs. In LogFile, every log file is composed of one or more logs, and every log describes a single system event. The log is the minimum data unit processed in Log Service.

Log Service defines a log by using the semi-structured data mode. This mode is composed of four data fields: Topic, Time, Content, and Source.

Meanwhile, Log Service has different format requirements for different log fields. See the following table for details.

Data field	Meaning	Format
Topic	A custom field used to mark a batch of logs. For example, access logs can be marked according to sites.	Any string up to 128 bytes, including a null string. By default, this field is a null string.
Time	A reserved field in the log used to indicate the log generation time. Generally this field is generated directly based on the time in the log.	An integer in the standard UNIX time format. The unit is in seconds. This field indicates the number of seconds since 1970-1-1 00:00:00 UTC.
Content	A field used to record the specific log content. The log content is composed of one or more content items, and each content item is a Key-Value pair.	The key is a UTF-8 encoded string up to 128 bytes, and can contain letters, underscores, and numbers. It cannot start with a number or use any of the following keywords: <code>_time_</code> , <code>_source_</code> , <code>_topic_</code> , <code>_partition_time_</code> , <code>_extract_others_</code> , and <code>_extract_others_</code> . The value can be any string up to 1024*1024 bytes.
Source	A field used to indicate the source of the log. For example, the IP address of the machine where the log is	Any string up to 128 bytes. By default, this field is null.

	generated.	
--	------------	--

Various log formats are used in actual usage scenarios. For ease of understanding, the following example describes how to map an original Nginx access log to the Log Service log data model. Assume that the IP address of your Nginx server is 10.249.201.117. The following is an original log of this server.

```
10.1.168.193 - - [01/Mar/2012:16:12:07 +0800] "GET /Send?AccessKeyId=8225105404 HTTP/1.1" 200 5 "-"
"Mozilla/5.0 (X11; Linux i686 on x86_64; rv:10.0.2) Gecko/20100101 Firefox/10.0.2"
```

Map the original log to the Log Service log data model as follows.


Data field	Content	Description
Topic	""	Use the default value (null string).
Time	1330589527	The precise log generation time, indicating the number of seconds since 1970-1-1 00:00:00 UTC. The time is converted from the timestamp of the original log.
Content	Key-Value pair	Specific log content.
Source	"10.249.201.117"	Use the IP address of the server as the log source.

You can decide how to extract the original content of the log and combine them into Key-Value pairs. The following table is shown as an example.

Key	Value
ip	"10.1.168.193"
method	"GET"
status	"200"
length	"5"
ref_url	"_ "
browser	"Mozilla/5.0 (X11; Linux i686 on x86_64; rv:10.0.2) Gecko/20100101 Firefox/10.0.2"

A log group is a collection of logs and is the basic unit for writing and reading.

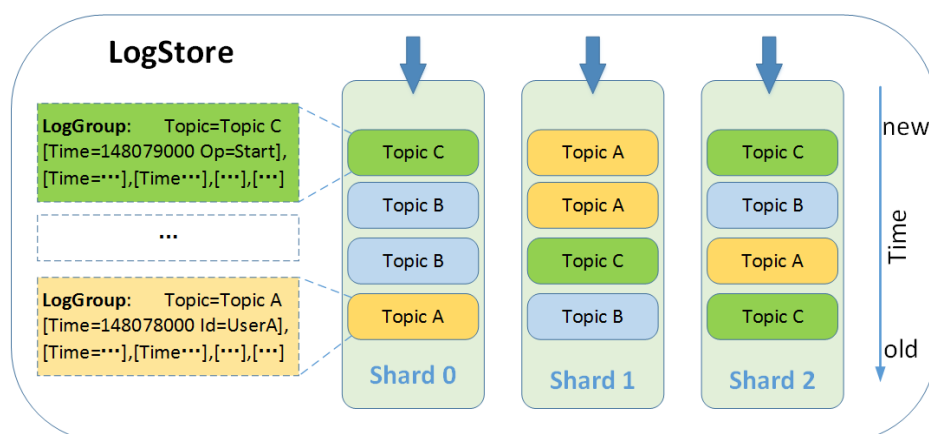
The maximum capacity of a log group is up to 4096 logs or 10 MB.



Logs in a Logstore can be classified by log topics. You can specify the topic when writing and querying logs. For example, as a platform user, you can use your user ID as the log topic when writing logs. In this way, you can select to only view your own logs based on the log topics when querying logs. If you do not need to classify the logs in a Logstore, use the same topic for all of the logs.

Note: A null string is a valid log topic and is the default log topic when writing and querying logs. So if you do not need to use the log topic, the easiest way is to use the default log topic, the null string, when writing and querying logs.

The following image describes the relationship among Logstore, log topic, and log.



The project is the resource management unit in Log Service and is used to isolate and control resources. You can manage all the logs and the related log sources of an application by using projects. Projects manage the information of all your Logstores and the log collection machine configuration, and serve as the portals where you can access the Log Service resources.

Specifically, projects provide the following functions:

Projects help you organize and manage different Logstores. In actual use, you may use Log Service to centrally collect and store the logs of different projects, products, or environments. You can classify different logs for management in different projects to facilitate subsequent usage, export, or index of logs. In addition, projects are the carriers of the log access permission management.

Projects serve as the portals where you can access the Log Service resources. Log Service

allocates a unique access portal for each created project. The access portal supports writing, reading, and managing logs through network.

You can perform the following project operations by using the Log Service management console:

- Create a project.
- View project list.
- Manage projects.
- Delete a project.

The Logstore is a unit in Log Service for the collection, storage, and query of log data. Each Logstore belongs to a project, and each project can create multiple Logstores. You can create multiple Logstores for a project according to your actual needs. Typically, an independent Logstore is created for each type of logs in an application. For example, you have a game application “big-game” , and three types of logs are on the server: operation_log, application_log, and access_log. You can first create a project named “big-game” , and then create three Logstores under this project for these three types of logs to collect, store, and query logs respectively.

You need to specify the Logstore for writing and querying logs.

Specifically, Logstores provide the following functions: