ALIBABA CLOUD

Alibaba Cloud

Log Service

Product Introduction

Document Version: 20220614

Alibaba Cloud

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.

2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.

3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.

4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).

5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.

6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

Style	Description	Example
Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	**Danger:** Resetting will result in the loss of user configuration data.
Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	**Warning:** Restarting will cause business interruption. About 10 minutes are required to restart an instance.
Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	**Notice:** If the weight is set to 0, the server no longer receives new requests.
Note	A note indicates supplemental instructions, best practices, tips, and other content.	**Note:** You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click **Settings**> **Network**> **Set network type**.
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click **OK**.
`Courier font`	Courier font is used for commands	Run the `cd /d C:/window` command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	`bae log list --instanceid` *Instance_ID*
[] or [a|b]	This format is used for an optional value, where only one item can be selected.	`ipconfig [-all|-t]`
{} or {a|b}	This format is used for a required value, where only one item can be selected.	`switch {active|stand}`

Table of Contents

1.What is Log Service?

Log Service is a cloud-native observation and analysis platform that provides large-scale, low-cost, and real-time services to process multiple types of data such as logs, metrics, and traces. Log Service allows you to collect, transform, query, analyze, visualize, ship, and consume data. You can also configure alerts in the console. Log Service helps enterprises improve their digital capabilities in terms of R&D, O&M, and data security.

Learning path

Log Service learning path provides common guides to use Log Service. For more information, see Log Service Learning Path. Log Service learning path also provides videos that match the user guides to help you understand Log Service.

Terms

Before you use Log Service, familiarize yourself with the following terms.

Term	Description
project	A project in Log Service is used to isolate resources of different users and control access to specific resources. For more information, see Project.
Logstore	A Logstore in Log Service is used to collect, store, and query log data. For more information, see Logstore.
Metricstore	A Metricstore in Log Service is used to collect, store, and query time series data. For more information, see Metricstore.
log	Logs are records of changes that occur in a system during the runtime of the system. The records contain information about the operations that are performed on specific objects and the results of the operations. The records are ordered by time. For more information, see Log.
log group	A log group is a collection of logs. A log group is the basic unit that is used to write and read logs. Logs in a log group contain the same metadata, such as the IP address and log source. For more information, see Log group.
metric	Metric data is stored as time series. For more information, see Metric.
trace	Trace data indicates the execution process of an event or a procedure in a distributed system. For more information, see Trace.
shard	A shard is used to control the read and write capacity of a Logstore. In Log Service, data is stored in shards. Each shard has an MD5 hash range, and each range is a left-closed and right-open interval. Each range does not overlap with the ranges of other shards. Each range must be within the entire MD5 hash range [00000000000000000000000000000000,ffffffffffffffffffffffffffffffff). For more information, see Shard.

Term	Description
topic	A topic is a basic management unit in Log Service. You can specify topics when you collect logs. This way, Log Service classifies logs by topic. For more information, see Topic.
endpoint	An endpoint of Log Service is a URL that is used to access a project and the data of the project. To access the projects in different regions, you must use different endpoints. To access the projects in the same region over an internal network or the Internet, you must also use different endpoints. For more information, see Endpoints.
AccessKey pair	An AccessKey pair is an identity credential that consists of an AccessKey ID and an AccessKey secret. The AccessKey ID and AccessKey secret are used for symmetric encryption and identity authentication. The AccessKey ID is used to identify a user. The AccessKey secret is used to encrypt and verify a signature string. The AccessKey secret must be kept confidential. For more information, see AccessKey pair.
region	A region is the physical location where a data center of Log Service is deployed. You can specify a region when you create a project. After the project is created, you cannot change the region. For more information, see Supported regions.

Features

Log Service provides the following features to meet the requirements of cloud-native observation and analysis in multiple business scenarios.

Usage

You can use Log Service by using the following methods.

Method	Description
Console	Log Service provides a web console to manage your Log Service resources. For more information, see Log Service console.
SDK	Log Service provides SDKs for various programming languages to facilitate secondary development. For more information, see SDK overview.
API	Log Service provides the API to manage your Log Service resources. This method requires you to sign API requests. For more information, see API overview. **Note** We recommend that you use SDKs to avoid signature verification.
CLI	Log Service provides a command-line interface (CLI) to manage your Log Service resources. For more information, see Command-line interface.

Billing

Log Service supports the pay-as-you-go billing method. You are charged based on your actual usage. Compared with self-managed ELK, Log Service allows you to reduce the total cost by 50%. For more information about metering items and billing items, see Billable items.

Activate Log Service

Click Activate Log Service to go to the buy page of Log Service.

2.Features

This topic describes the features of Log Service.

Data collection

Log Service can collect the following types of data by using more than 50 methods:

- Logs, time series data, and trace data from servers and applications
- Logs from IoT devices
- Logs from Alibaba Cloud services
- Data from mobile devices
- Data from open source software such as Logstash, Flume, Beats, FluentD, and Telegraph
- Data transferred over protocols such as HTTP, HTTPS, Syslog, Kafka, and Prometheus

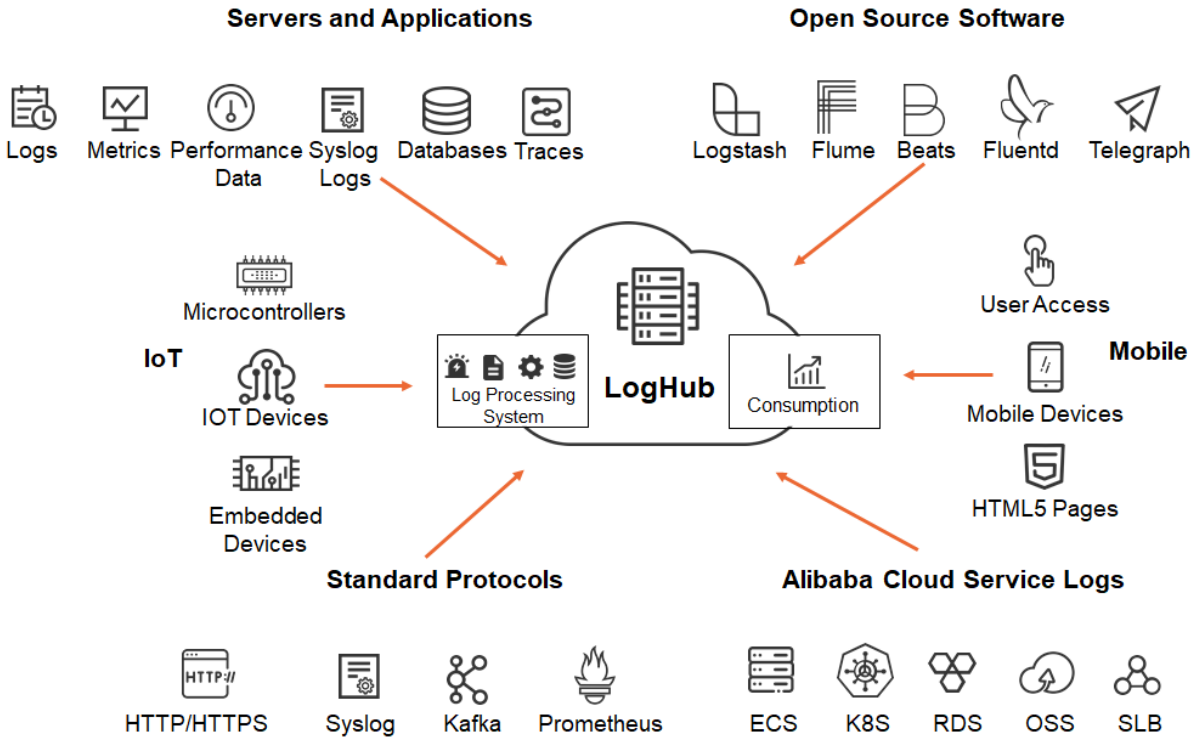

For more information, see Data collection overview.

Query and analysis

Log Service supports data query and analysis in real time.

- Log Service supports exact search, fuzzy search, full-text search, and field search.
- Log Service supports features such as contextual query, LogReduce, LiveTail, and reindexing.
- Log Service supports the SQL-92 syntax.
- Log Service provides the Dedicated SQL feature.

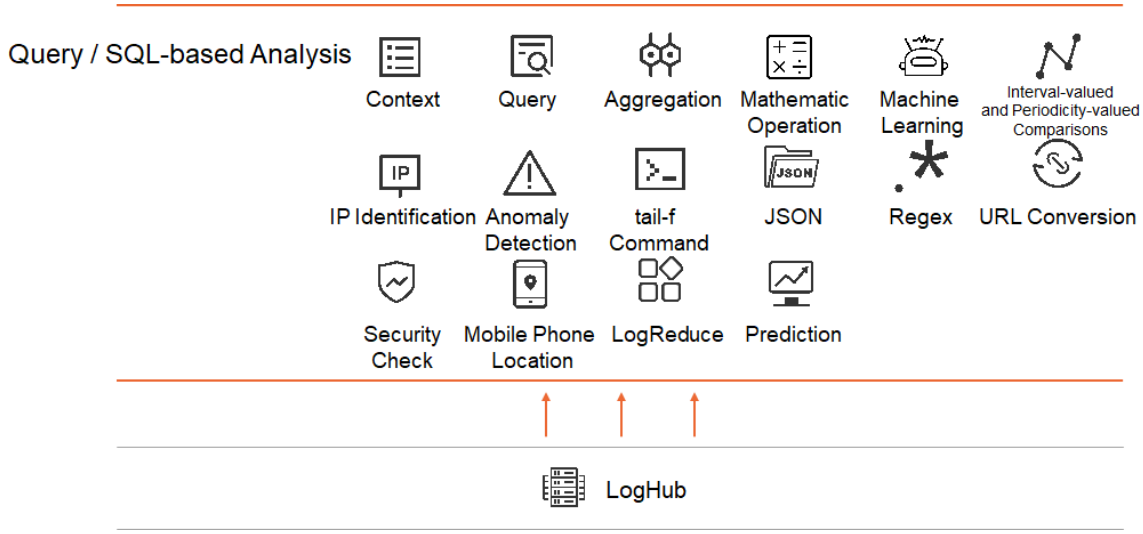

For more information, see Log search overview and Log analysis overview.

Data transformation

You can use the data transformation feature to standardize, enrich, transfer, mask, and filter data.

- Data standardization: Log Service can extract fields from logs in different formats and convert the log formats to obtain structured data for stream processing and computing in data warehouses.
- Data enrichment: Log Service can join the fields of logs and dimension tables to link logs with dimension information, which facilitates data analysis. For example, Log Service can join the fields of order logs and a user information table.
- Data transfer: Log Service can transfer logs from regions outside China to the region of the central project by using the global acceleration feature. This way, global logs can be managed in a centralized manner.
- Data masking: Log Service can mask sensitive information that is contained in data. The sensitive information includes passwords, mobile phone numbers, and addresses.
- Data filtering: Log Service can filter logs for those of key services. This helps further analysis.

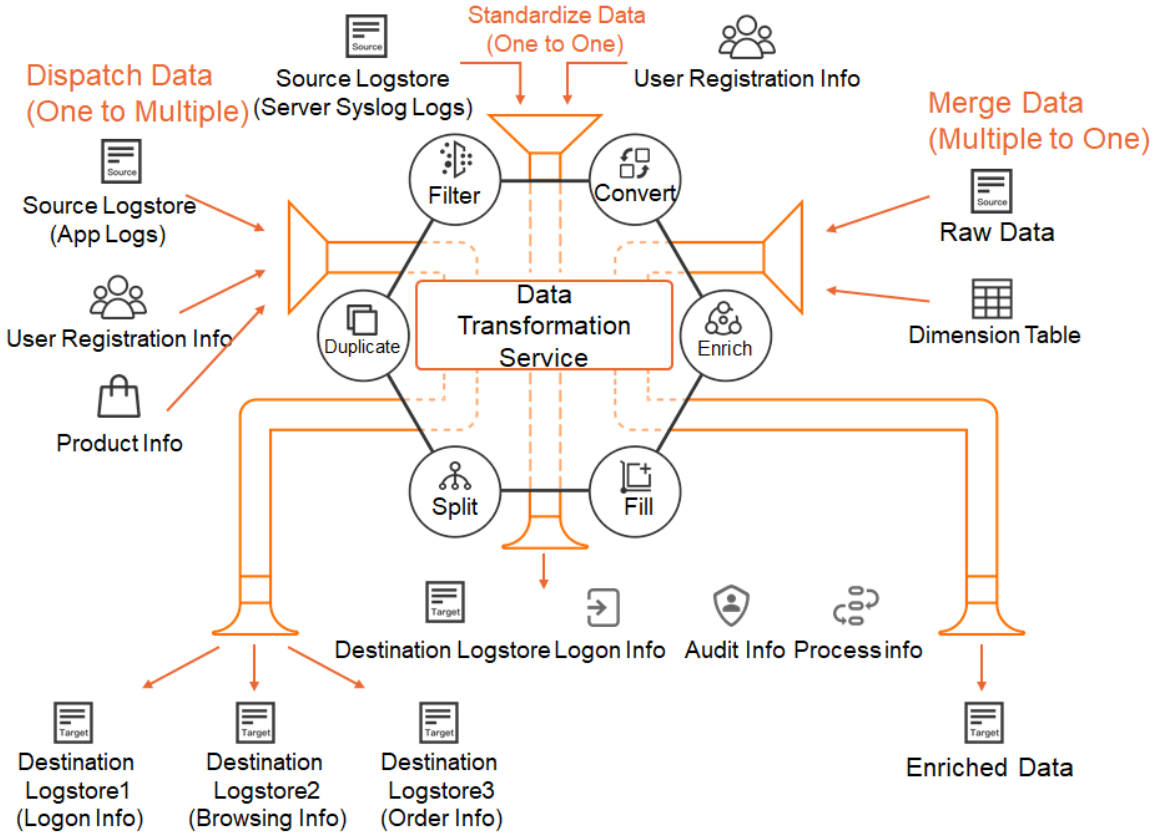

For more information, see Data transformation overview.

Consumption and shipping

You can use the consumption and shipping feature to consume data in real time by using SDKs or API operations of Log Service. You can also ship logs to other Alibaba Cloud services, such as Object Storage Service (OSS) and MaxCompute, in real time in the Log Service console.

- You can consume data by using third-party software, such as Splunk, QRadar, Logstash, and Flume.
- You can consume data by using different programming languages, such as Java, Python, and GO.
- You can consume data by using Alibaba Cloud services, such as Function Compute and Realtime Compute for Apache Flink.
- You can consume data by using different stream processing platforms, such as Apache Flink, Apache Spark, and Apache Storm.
- You can ship data to Alibaba Cloud services, such as OSS and MaxCompute.

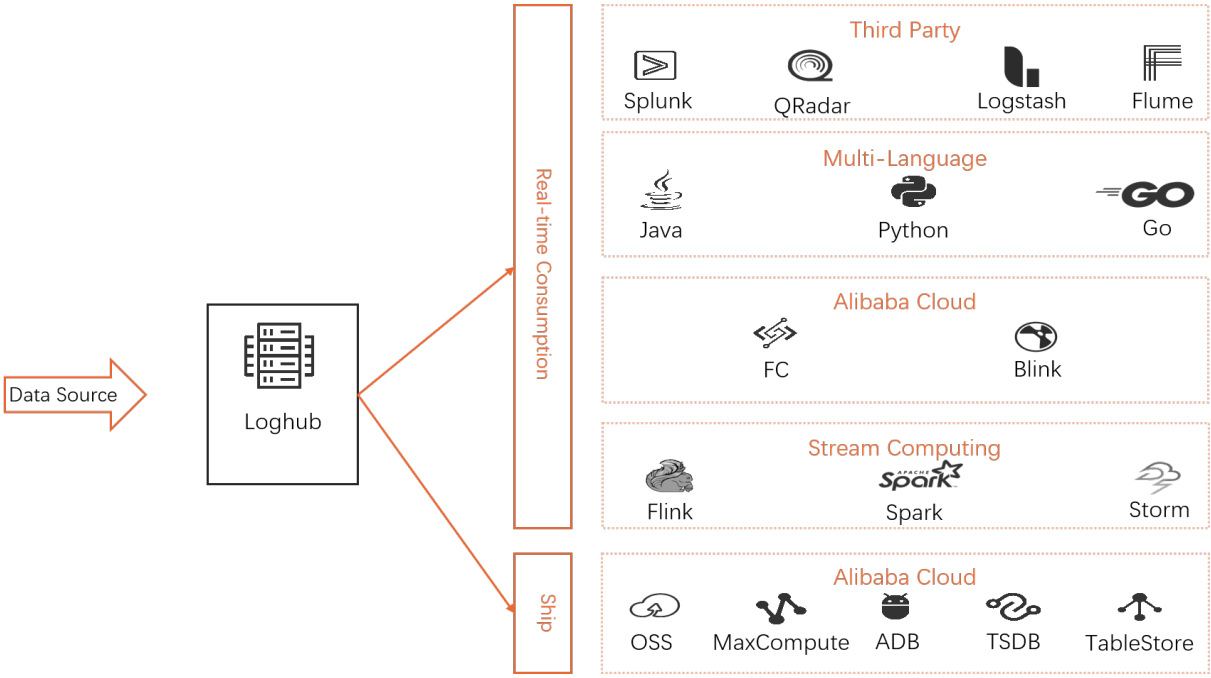

For more information, see Overview of real-time consumption and 数据投递概述.

Visualization

Log Service supports the visualization of query and analysis results.

- Built-in charts on dashboards: Log Service provides various statistical charts, such as tables, line charts, and column charts. You can select chart types to visualize query and analysis results on a dashboard and save the results to the dashboard.
- Third-party visualization tools: Log Service is compatible with third-party visualization tools, such as Grafana and DataV.

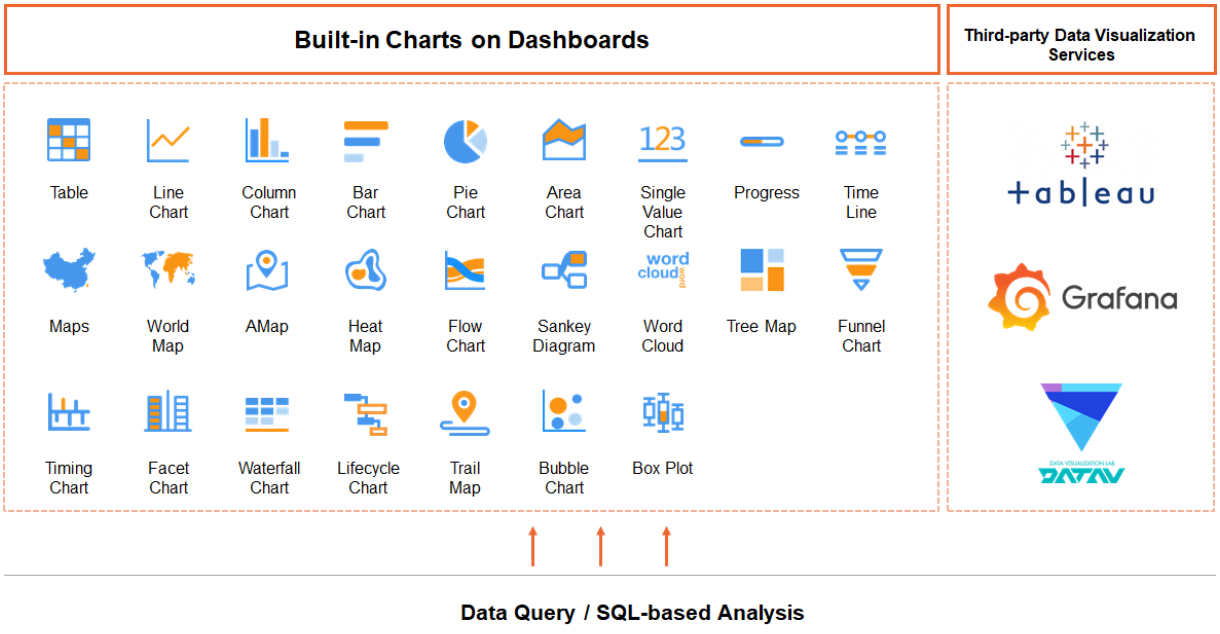

For more information, see Visualization overview.

Alerting

You can use the alerting feature of Log Service to configure alert monitoring, denoise alerts, manage alert incidents, and configure notification methods.

- Alert monitoring: The alert monitoring system can regularly check and evaluate query and analysis results based on alert monitoring rules, trigger or clear alerts, and send alert or recovery notifications to the alert management system.
- Alert management: The alert management system can process alerts based on alert policies. For example, the system can dispatch, suppress, deduplicate, silence, or merge alerts. After the alerts are processed, they are sent to the notification management system.
- Notification management: The notification management system can send alert notifications to specified recipients by using specified notification methods based on action policies. Recipients can be users, user groups, or on-duty groups.
- Alert ingestion: The alert ingestion system can ingest alerts from external monitoring systems such as Grafana and Prometheus by using webhooks. After the alerts are ingested, the alert ingestion system can manage the alerts and send alert notifications.

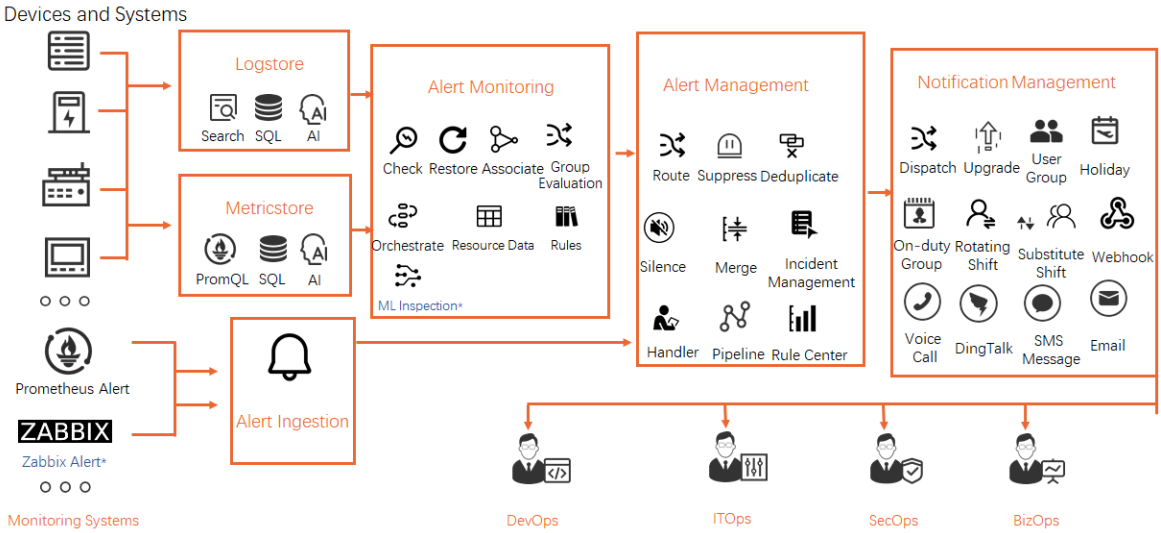

For more information, see Alert overview.

Log audit

Log Audit Service provides all features of Log Service and supports automated collection and third-party cloud services.

- Log Audit Service supports automated collection for cloud service logs across Alibaba Cloud accounts. The logs can be stored in a central project. Log Audit Service also allows you to audit the logs in a centralized manner.
- You can use Log Audit Service to audit the logs that are collected from the following Alibaba Cloud services: ActionTrail, Container Service for Kubernetes (ACK), OSS, Apsara File Storage NAS, Server Load Balancer (SLB), API Gateway, ApsaraDB RDS, Distributed Relational Database Service (DRDS), PolarDB for MySQL, Web Application Firewall (WAF), Anti-DDoS, Cloud Firewall, and Security Center.
- You can also use Log Audit Service to audit the logs that are collected from third-party cloud services or self-managed security operations centers (SOCs).
- Log Audit Service provides hundreds of built-in alert rules. You can enable the alert rules with only a few clicks. The alert rules help you monitor the compliance of hosts, databases, networks, and logs in

account security and permission management.

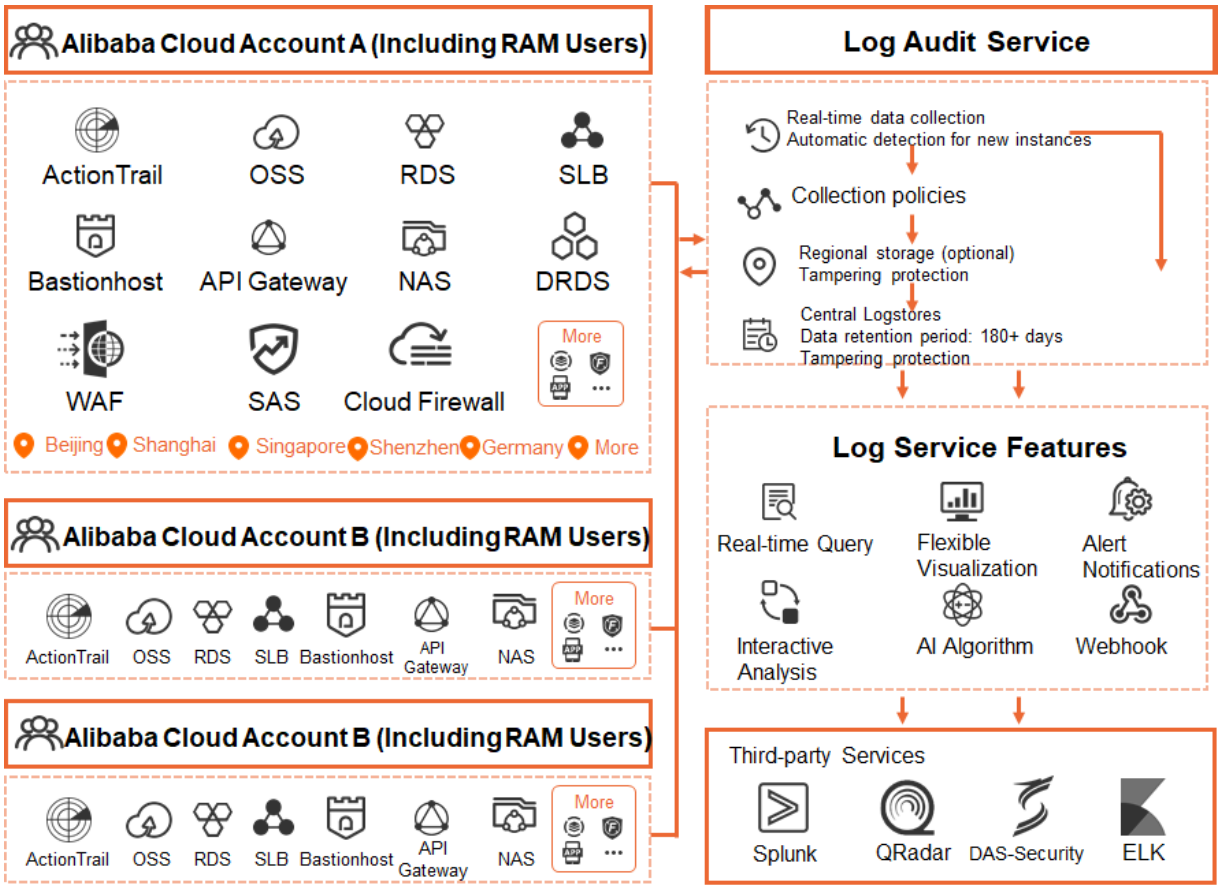

For more information, see Overview of Log Audit Service.

3.Architecture

This topic describes the architecture of Log Service.

The following figure shows the architecture of Log Service.

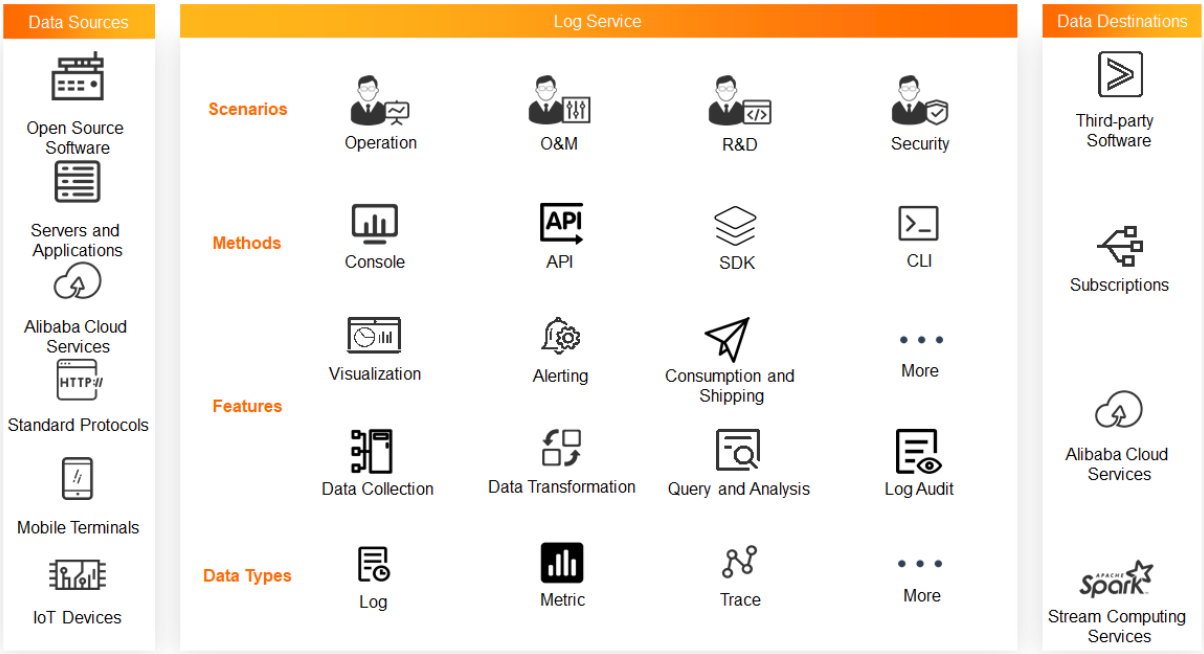

- Data sources

  Log Service allows you to collect data from open source software, servers and applications, Alibaba Cloud services, mobile terminals, and IoT devices. You can also collect data that is transferred over multiple protocols.

- Log Service
  - Data types

    Log Service provides a platform on which you can use various features to process large amounts of data at low cost and in real time. The types of data that you can process include logs, metrics, and traces. For more information, see Log, Metric, and Trace.

- Features
  - Data collection: Log Service allows you to collect data by using Logtail, SDKs, and protocols. For more information, see Data collection overview.
  - Data transformation: Data transformation is a fully managed feature that provides high availability and scalability in Log Service. You can use the data transformation feature to standardize, enrich, dispatch, mask, and filter data. For more information, see Data transformation overview.
  - Data query and analysis: Log Service allows you to query and analyze petabytes of data in real time. This feature supports more than 10 operators, more than 10 machine learning functions, and more than 100 SQL functions. This features also supports scheduled SQL jobs and dedicated SQL instances. For more information, see Log search overview and Log analysis overview.
  - Visualization: Log Service allows you to visualize query and analysis results. You can customize dashboards based on charts. For more information, see Visualization overview.
  - Alerting: Log Service provides the integrated alerting feature. This feature includes the alert monitoring, alert management, and notification management systems. The alerting feature is applicable to multiple scenarios such as DevOps, ITOps, AIOps, SecOps, and BizOps. For more information, see The alerting feature of Log Service.
  - Data consumption and shipping: Log Service allows you to consume data in real time by using Storm, Flume, or Flink. You can also ship data to Object Storage Service (OSS), MaxCompute, or Time Series Database (TSDB) in real time. For more information, see Data shipping overview and Data consumption overview.
  - Log audit: Log Service provides an automated and centralized method to collect and audit the logs of cloud services across Alibaba Cloud accounts in real time. For more information, see Log Audit Service.
- Methods

  You can manage your resources in Log Service by using the Log Service console, API, SDKs, or CLI.
- Scenarios

  Log Service is applicable to multiple scenarios such as operation, O&M, R&D, and security. For more information, see Scenarios.

- Data destinations

  Log Service allows you to export data to cloud services or third-party software by using the data consumption or data shipping feature.

4.Benefits

This topic describes the benefits of Log Service.

Unified import method

Log Service allows you to import data from multiple types of data sources.

Intelligence

Log Service provides AIOps capabilities to intelligently detect exceptions and analyze root causes.

Efficiency

Log Service can collect, query, and analyze tens of billions of log data rows in real time.

One-stop service

Log Service allows you to collect, transform, query, analyze, and visualize data. You can also configure alerts for the data.

Scalability

Log Service provides auto scaling capabilities for petabytes of data.

Cost-effectiveness

Log Service supports the pay-as-you-go billing method. You are charged only for the actual usage. The total cost of ownership (TCO) is reduced by more than 50%.

5.Scenarios

This topic describes typical scenarios in which you can use Log Service for your business. The scenarios include data collection and consumption, data extract, transform, and load (ETL) and stream processing, integration with data warehouses, and real-time query and analysis.

Data collection and consumption

You can use the LogHub module of Log Service to collect large amounts of log data in real time. The log data can be metrics, events, binary logs, text logs, and clickstream data.

Benefits:

- Ease of use: Log Service provides more than 50 data collection methods to help you build platforms. Log Service also delivers powerful configuration and management capabilities to reduce your O&M workloads.
- Elastic scalability: Log Service can handle traffic spikes and business growth.

Data collection and consumption

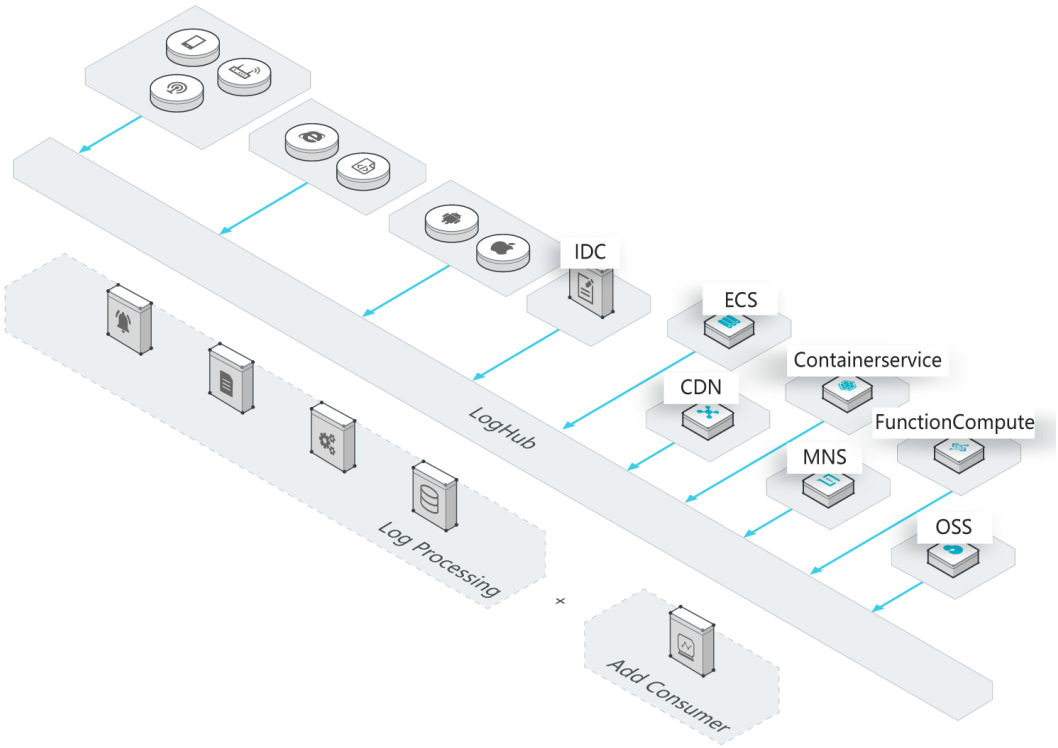

ETL and stream processing

LogHub can connect to multiple stream processing engines and services. LogHub can also monitor the processing progress and generate alerts. You can also SDKs or API operations to consume data based on your business requirements.

- Ease of use: Log Service provides comprehensive SDKs and programming frameworks for seamless integration with multiple stream processing engines.
- Monitoring and alerting: Log Service provides comprehensive metrics and an alerting mechanism upon latency.

- Elastic scalability: Log Service supports auto scaling to process petabytes of data without latency.

  ETL and stream processing

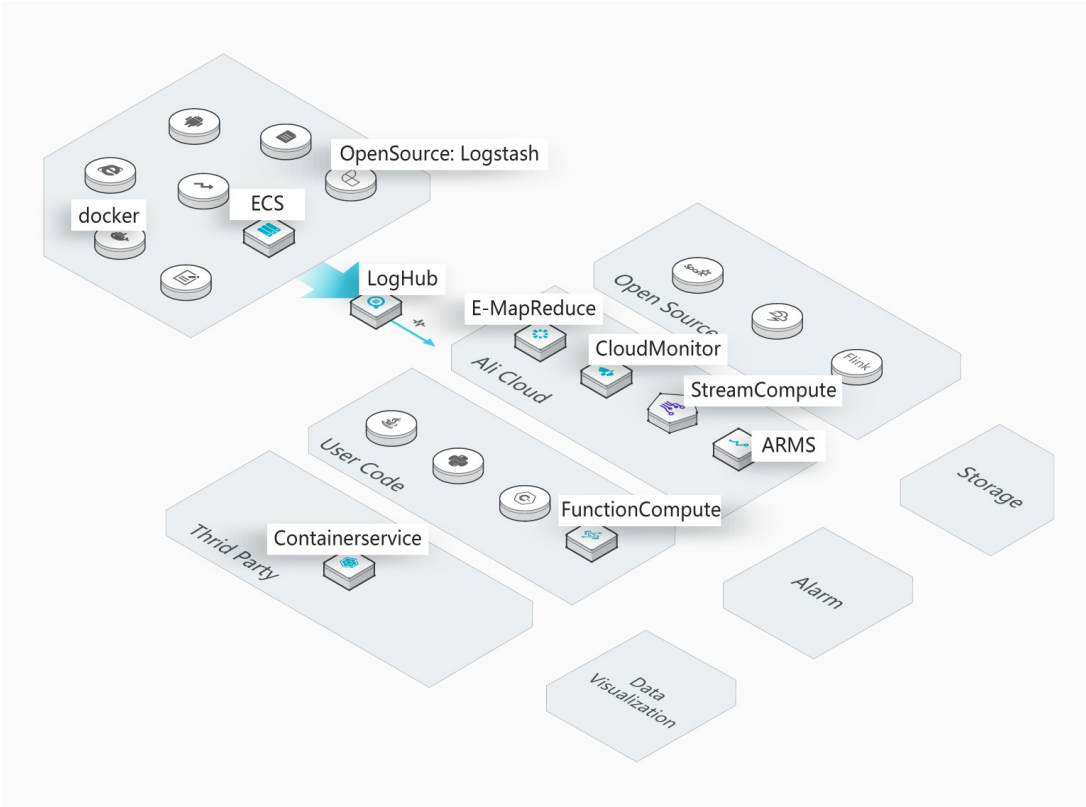

Integration with data warehouses

The LogShipper module of Log Service can ship LogHub data to storage services. During the shipping, you can compress the data, define custom partition formats, and specify row or column store.

- Large data capacity: An unlimited amount of data can be shipped to storage services.
- Multiple formats: Various storage formats such as row store, column store, and text files are supported.
- Flexible configurations: Different configurations are supported, which allows you to define custom partition formats.

  Integration with data warehouses

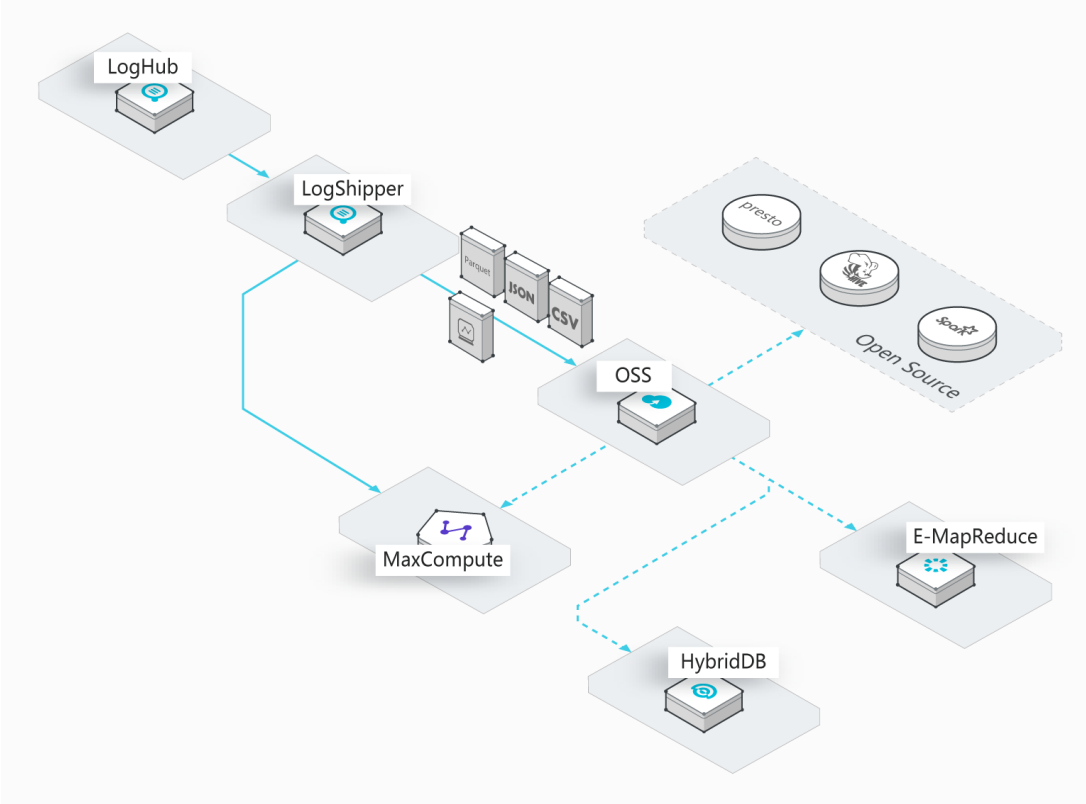

Real-time query and analysis

The LogAnalytics module allows you to index LogHub data in real time and query data by using keywords, fuzzy match, contextual query, or SQL aggregate functions. You can also query data within a specified range.

- Timeliness: You can perform real-time query after data is written to LogHub.
- High efficiency at low costs: You can index petabytes of data per day. Costs are 85% lower compared with self-managed systems.
- Strong analysis: Multiple query methods and SQL aggregate functions are supported. Log Service can also generate visualized reports and alerts.

Real-time query and analysis

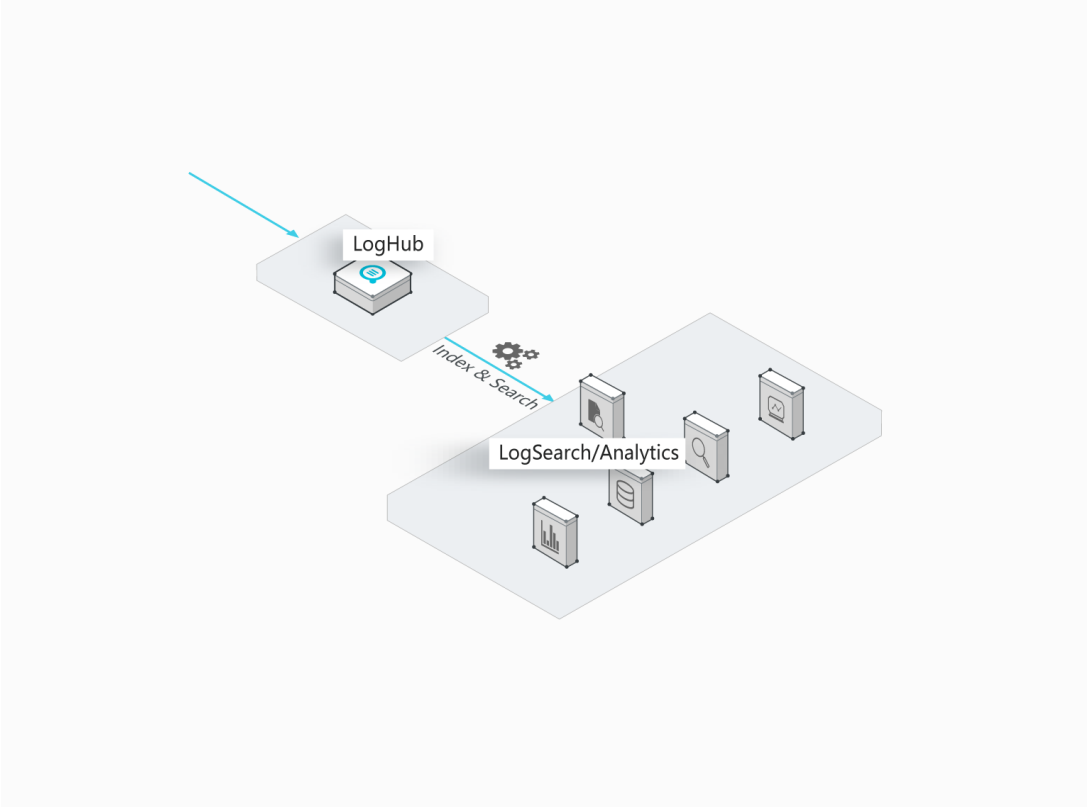
LogHub
Index & Search
LogSearch/Analytics

6.Terms

6.1. Terms

This topic introduces the terms that are used in Log Service.

Basic resources

Term	Description
project	A project in Log Service is used to isolate resources of different users and control access to specific resources. For more information, see Project.
Logstore	A Logstore in Log Service is used to collect, store, and query log data. For more information, see Logstore.
Metricstore	A Metricstore in Log Service is used to collect, store, and query time series data. For more information, see Metricstore.
log	Logs are records of changes that occur in a system during the runtime of the system. The records contain information about the operations that are performed on specific objects and the results of the operations. The records are ordered by time. For more information, see Log.
log group	A log group is a collection of logs. A log group is the basic unit that is used to write and read logs. Logs in a log group contain the same metadata, such as the IP address and log source. For more information, see Log group.
metric	Metric data is stored as time series. For more information, see Metric.
trace	Trace data indicates the execution process of an event or a procedure in a distributed system. For more information, see Trace.
shard	A shard is used to control the read and write capacity of a Logstore. In Log Service, data is stored in shards. Each shard has an MD5 hash range, and each range is a left-closed and right-open interval. Each range does not overlap with the ranges of other shards. Each range must be within the entire MD5 hash range [00000000000000000000000000000000,ffffffffffffffffffffffffffffffff). For more information, see Shard.
topic	A topic is a basic management unit in Log Service. You can specify topics when you collect logs. This way, Log Service classifies logs by topic. For more information, see Topic.
endpoint	An endpoint of Log Service is a URL that is used to access a project and the data of the project. To access the projects in different regions, you must use different endpoints. To access the projects in the same region over an internal network or the Internet, you must also use different endpoints. For more information, see Endpoints.

Term	Description
AccessKey pair	An AccessKey pair is an identity credential that consists of an AccessKey ID and an AccessKey secret. The AccessKey ID and AccessKey secret are used for symmetric encryption and identity authentication. The AccessKey ID is used to identify a user. The AccessKey secret is used to encrypt and verify a signature string. The AccessKey secret must be kept confidential. For more information, see AccessKey pair.
region	A region is the physical location where a data center of Log Service is deployed. You can specify a region when you create a project. After the project is created, you cannot change the region. For more information, see Supported regions.

Data collection

Term	Description
Logtail	Logtail is used by Log Service to collect logs. For more information, see Use Logtail to collect data.
Logtail configuration	A Logtail configuration is a set of policies that are used by Logtail to collect logs. The configuration includes the log source and collection method. For more information, see Logtail configurations.
machine group	A machine group is a virtual group that contains multiple servers. Log Service uses machine groups to manage the servers from which you want to collect logs by using Logtail. For more information, see Introduction.

Data query and analysis

Term	Description
query	You can configure filter conditions in search statements to obtain specific logs. For more information, see Log search overview.
analysis	You can invoke SQL functions on query results to perform statistical and analytical operations. You can also obtain analysis results. • Log Service supports the SQL-92 syntax for you to analyze log data. For more information, see Log analysis overview. • Log Service supports the SQL-92 syntax and the PromQL syntax for you to analyze time series data. For more information, see Overview of query and analysis on time series data.
query statement	A query statement is in the `Search statement | Analytic statement` format. A search statement can be separately executed. However, an analytic statement must be executed together with a search statement. The log analysis feature is used to analyze search results or all data in the Logstore. For more information, see Query and analysis.

Term	Description
index	Indexes are a structure for storage. Indexes are used to sort one or more columns of data. You can query data only after you create indexes for the data. Log Service provides the following two types of indexes: • full-text index: Log Service splits an entire log entry into multiple words based on specific delimiters and creates indexes. In a query, the field names and field values are both plaintext. • field index: After you create field indexes, you can query log entries by specifying field names and field values in the key:value format. For more information, see Configure indexes.
Standard SQL	The Standard SQL feature is free of charge. It allows you to analyze data by executing SQL statements. The Standard SQL feature provides less resources than the Dedicated SQL feature.

Data transformation

Term	Description
domain-specific language (DSL)	DSL is a Python-compatible scripting language. DSL is used for data transformation in Log Service. For more information, see Language introduction.
transformation rule	A transformation rule is a data transformation script that is orchestrated by using the DSL for Log Service. For more information, see Syntax introduction.

Consumption and shipping

Term	Description
consumer group	You can use consumer groups to consume data in Log Service. A consumer group consists of multiple consumers. Each consumer consumes different log entries that are stored in a Logstore. For more information, see Use consumer groups to consume log data.

Alerting

Term	Description
alert	An alert indicates an alert event. If an alert is triggered based on a specific alert monitoring rule, the event is transferred to the notification management system. Log Service also provides alert-related features, entities, modules, and subsystems, such as the alert monitoring system and alert monitoring rules. For more information, see The alerting feature of Log Service.

Term	Description
alert monitoring system	The alert monitoring system is a subsystem that triggers alerts. The alert monitoring system contains alert monitoring rules and resource data. An alert monitoring rule is used to periodically monitor and evaluate query and analysis results. If an alert is triggered or cleared, an alert notification or recovery notification is sent to the alert management system based on monitoring rule orchestration.
alert management system	The alert management system is a subsystem that denoises alerts and manages alert states. The alert management system contains alert policies, alert incidents, and alert dashboards. The alert management system processes alerts based on alert policies. For example, the system can dispatch, suppress, deduplicate, silence, or merge alerts. After the alerts are processed, they are sent to the notification management system. The alert management system also allows you to switch incident phases and set handlers for incidents.
notification management system	The notification management system is a subsystem that manages notification methods and recipients. The notification management system contains action policies, alert templates, calendars, users, user groups, on-duty groups, and notification method quotas. The notification management system sends alert notifications to specified recipients by using specified notification methods based on action policies. Recipients can be users, user groups, or on-duty groups. The notification management system also allows you to escalate alerts and customize alert templates.
alert ingestion system	The alert ingestion system is a subsystem that ingests external alerts. The alert ingestion system contains alert ingestion services and alert ingestion applications. Each alert ingestion application provides a webhook to ingest external alerts from external services, such as Zabbix and Prometheus. Recovery notifications can also be ingested, After an external alert is received, the alert is preprocessed and sent to the alert management system for further processing.

6.2. Log

Logs are records of changes that occur in a system during the running of the system. The records contain information about the operations that are performed on specific objects and the results of the operations. The records are ordered by time.

Format

Log data is stored in different forms, such as log files, log events, binary logs, and metric data. Log Service uses a semi-structured data model to define logs. A log entry consists of the following fields: topic, time, content, source, and tags. Log Service has different format requirements on different log fields. The following table describes the log fields and provides the format requirements.

Field	Description	Format
Topic	The custom field in a log entry. This field can be used to identify the log topic. For example, you can set different log topics (access_log and operation_log) for website logs based on log types. For more information, see Topic.	The field value can be a string of up to 128 bytes, including an empty string. If the field is an empty string, it indicates that the log topic is not configured.
Time	The time when the log entry is generated, or the system time of the host where Logtail resides when the log data is collected. This field is a reserved field.	The value is a UNIX timestamp. It is the number of seconds that have elapsed since 00:00:00 Thursday, 1 January 1970.
Content	The content of the log entry. The content consists of one or more items. Each item is a key-value pair.	A key-value pair must comply with the following requirements: • The key is a UTF-8 encoded string of up to 128 bytes. The key can contain letters, digits, and underscores (_). The string is 1 to 128 bytes. The following fields cannot be used: ◦ __time__ ◦ __source__ ◦ __topic__ ◦ __partition_time__ ◦ _extract_others_ ◦ __extract_others__ • The value can be a string of up to 1 MB.
Source	The source of the log entry. For example, the value of this field can be the IP address of the server where the log entry is generated.	The value of this field can be a string of up to 128 bytes.
Tags	The tags of the log entry. The following log tags can be used: • Custom tags: the tags that you add when you call the PutLogs operation to write logs to a specified Logstore. • System tags: the tags added by Log Service. The tags include __client_ip__ and __receive_time__.	The field value is in the dictionary format. The keys and the values are strings. The field name is prefixed by __tag__:.

Examples

The following examples show a website access log that is collected to Log Service in different modes.

- Raw log entry

```
127.0.0.1 - - [01/Mar/2021:12:36:49  0800] "GET /index.html HTTP/1.1" 200 612 "-" "Mozill
a/5.0 (Macintosh; Intel Mac OS X 10_13_6) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/6
8.0.3440.106 Safari/537.36
```

- Log entry collected to Log Service in simple mode

- Log entry collected to Log Service in full regex mode

6.3. Log group

A log group is a group of logs that serves as a basic unit for read/write operations. The logs in a log group have the same metadata, such as the IP address and log source.

When you write logs to or read logs from Log Service, multiple logs are encapsulated into a log group. Then, you can write and read logs by log group. This method can reduce the number of read and write operations and improve business efficiency. The maximum length of a log group is 5 MB.

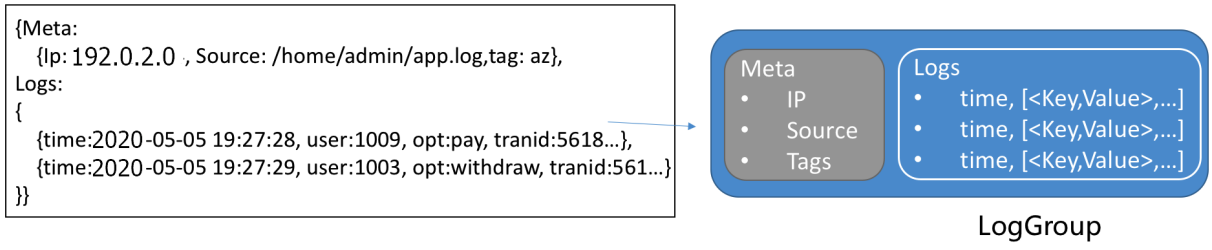

LogGroup

6.4. Project

A project in Log Service is used to separate different resources of multiple users and control access to specific resources.

A project contains resources such as Logstores, Metricstores, and machine groups, and provides an endpoint that you can use to access the resources of Log Service. We recommend that you use different projects to manage the data in different applications, services, or projects.

- A project organizes and manages Logstores or Metricstores. You can use Log Service to collect and

store the logs of different projects, services, or environments. You can specify different projects to facilitate log data consumption, export, and analysis.

- A project facilitates access control. You can grant a Resource Access Management (RAM) user the permissions to manage a specified project.
- A project provides an endpoint that you can use to access the resources in the project. Log Service allocates an exclusive endpoint to each project. You can use the endpoint of a project to read, write, and manage log data. For more information, see Endpoints.

6.5. Logstore

A Logstore in Log Service is a unit that is used to collect, store, and query logs.

Each Logstore belongs to a project. You can create multiple Logstores in a project. You can create multiple Logstores in a specific project based on your business requirements. In most cases, you can create a Logstore for each log type of an application. For example, if you want to collect the operation logs, application logs, and access logs of App A, you can create a project named app-a, and then create three Logstores named operation_log, application_log, and access_log in the project to store the logs.

You must specify a Logstore when you write, query, analyze, transform, consume, or ship logs.

- Log Service uses the Logstore as a collection unit to collect logs.
- Log Service uses the Logstore as a storage unit to store, transform, consume, and ship logs.
- Log Service creates indexes for the Logstore to query and analyze logs.

6.6. Metricstore

A Metricstore in Log Service is a unit that is used to collect, store, and query time series data.

Each Metricstore belongs to a project. You can create multiple Metricstore in a project. You can create multiple Metricstores in a project based on your business requirements. In most cases, you can create a Metricstore for each type of time series data. For example, if you want to collect the monitoring metrics of hosts, cloud services, and business applications, you can create a project named demo-monitor, and then create three Metricstores named host-metrics, cloud-service-metrics, and app-metrics in the project to store the metrics.

You must specify a Metricstore when you write, query, analyze, or consume time series data.

- Log Service uses Metricstore as a collection unit to collect time series data.
- Log Service uses Metricstore as a storage unit to store and consume time series data.
- Log Service allows you to query and analyze time series data by using the SQL-92 syntax and the PromQL syntax.

6.7. Metric

Log Service stores all the data in Metricstores as time series. Log Service uses the model of time series data defined by Prometheus. For more information about the model, see DATA MODEL in the Prometheus documentation. Each time series consists of samples with the same metric identifier.

Metric identifier

Each time series has a unique metric identifier that consists of a metric name and a label.

Metric names are strings and must match the **[a-zA-Z_:][a-zA-Z0-9_:]*** regular expression. In most cases, a metric name indicates a description of a time series. For example, http_request_total indicates that each sample of a time series represents the total number of received HTTP requests.

Labels are key-value pairs. Label keys must match the **[a-zA-Z_][a-zA-Z0-9_]*** regular expression. Label values can contain all characters except vertical bars (|). In most cases, a label indicates an attribute of a time series. For example, the value of the method key may be POST, and the value of the URL key may be /api/v1/get.

Samples

A sample indicates the value of a metric at a point in time. Each sample consists of a timestamp and a value. Timestamps are accurate to the nanosecond, and values are of the DOUBLE type.

Encoding format

When time series data is written to Log Service, the Protocol Buffer (Protobuf) format must be used. This format is also used to write log data. For more information, see Data encoding. The metric identifier and samples are contained in the content field. The following table describes the related subfields.

Key	Limit	Example
__name__	The metric name of the time series.	nginx_ingress_controller_response_size
__labels__	The labels of the time series. Format: {key}#$#{value}\|{key}#$#{value}|{key}#$#{value}. **Note** The labels must be sorted by key in alphabetical order.	app#$#ingress-nginx\|controller_class#$#nginx|controller_namespace#$#kube-system\|controller_pod#$#nginx-ingress-controller-589877c6b7-hw9cj
__time_nano__	The timestamp of a sample. It is accurate to the nanosecond.	1585727297293000000
__value__	The value of a sample.	36.0

6.8. Shard

Shards are used to control the read and write capacity of Logstores or Metricstores. In Log Service, data is stored in a shard.

MD5 value range

Each shard has an MD5 value range. and each range is a left-closed and right-open interval in the format of [BeginKey, EndKey). Each range does not overlap the ranges of other shards. The entire MD5 value range is within the following range: [00000000000000000000000000000000, ffffffffffffffffffffffffffffffff). You can specify the number of shards when you create a Logstore or Metricstore. Log Service evenly divides the entire MD5 value range for each shard.

- BeginKey: the start of a shard. The value is included in the MD5 value range of the shard.
- EndKey: the end of a shard. The value is excluded from the MD5 value range of the shard.

In this example, Logstore A has four shards. The following table shows the MD5 value range of each shard.

MD5 value range

Shard ID	Value range
Shard0	[00000000000000000000000000000000,40000000000000000000000000000000)
Shard1	[40000000000000000000000000000000,80000000000000000000000000000000)
Shard2	[80000000000000000000000000000000,c0000000000000000000000000000000)
Shard3	[c0000000000000000000000000000000,ffffffffffffffffffffffffffffffff)

To read data from a shard, you must specify the ID of the shard. To write data to a shard, you can use the load balancing method or specify a hash key.

- If you use the load balancing method, each data packet is randomly written to an available shard.
- If you specify a hash key, data is written to the shard whose MD5 value range includes the value of the specified hash key.

  For example, the shard range is shown in MD5 value range. If you specify 5F as a hash key to write data to a Logstore, the data is written to Shard1 because the MD5 value range of Shard1 contains the hash key 5F. If you specify 8C as a hash key, the data is written to Shard2 because the MD5 value range of Shard2 contains the hash key 8C.

Shard capacity

Each shard provides the following read capacity and write capacity:

- Write capacity: 5 MB/s or 500 times/s
- Read capacity: 10 MB/s or 100 times/s

We recommend that you adjust the number of shards based on the actual data traffic. If the data traffic exceeds the read or write capacity of a shard, you can split the shard into multiple shards to increase the capacity. If the data traffic is much lower than the read or write capacity of a shard, you can merge the shard with another shard to reduce the capacity and save costs.

For example, you have two shards that are in the readwrite state and the shards can provide a maximum write capacity of 10 MB/s. If you need to write data at a speed of 14 MB/s in real time, we recommend that you split one of your shards into two shards. This way, you can have three shards that are in the readwrite state. If you need to write data at a speed of 3 MB/s in real time, we recommend that you merge your shards.

> Notice
>
> - If the error code 403 or 500 is frequently returned when you write data by calling the Log Service API, you can go to the CloudMonitor console to check the traffic and status codes. Then, you can determine whether to increase the number of shards.
> - If the data traffic exceeds the capacity of your shards, Log Service attempts to provide the best possible services. However, Log Service cannot ensure the quality of the service.

Shard status

A shard can be in the readwrite (read and write) state or readonly (read-only) state.

When you create a shard, the shard is in the readwrite state. If you split or merge the shard, its status changes to readonly. The newly generated shard is in the readwrite state. The status of a shard does not affect the read capacity of the shard. Data can be written to the shards that are in the readwrite state, but cannot be written to the shards that are in the readonly state.

Splitting and merging

Log Service allows you to split and merge shards.

- After you split a shard, two more shards are added. The new shards are in the readwrite state and are listed under the original shard. The MD5 value range of the new shards incudes the MD5 value range of the original shard.

  Before you can split a shard, the shard must be in the readwrite state. After you split a shard, the status of the shard changes from readwrite to readonly. This specifies that data can still be read from the shard, but cannot be written to the shard.

- You can merge two shards into one shard. The new shard is in the readwrite state and is listed under the original shard. The MD5 value range of the new shard incudes the MD5 value range of the two original shards.

  When you merge shards, you must specify a shard that is in the readwrite state. This shard cannot be the last shard in the shard list. Log Service locates the shard whose MD5 value range is next to the specified shard and then merges the two shards. After you merge the shards, the status of the shards changes from readwrite to readonly. This specifies that data can still be read from the shards, but cannot be written to the shards.

6.9. Topic

A topic is a basic management unit in Log Service. When you collect logs, you can specify topics to identify the logs.

You can use topics to identify logs generated by different services, users, and instances. For example, if system A consists of an HTTP request processing module, a cache module, a logic processing module, and a storage module. You can set the log topic of the HTTP request processing module to http_module, the log topic of the cache module to cache_module, the log topic of the logic processing module to logic_module, and the log topic of the storage module to store_module. When the logs of the preceding modules are collected to the same Logstore, you can identify logs based on the topics.

If you do not need to identify logs in a Logstore, set the topic to *Null-Do not generate topic* when you collect logs. A topic can be an empty string, which indicates that the value of the topic is an empty string.

The following figure shows the relationships between Logstores, topics, and shards.

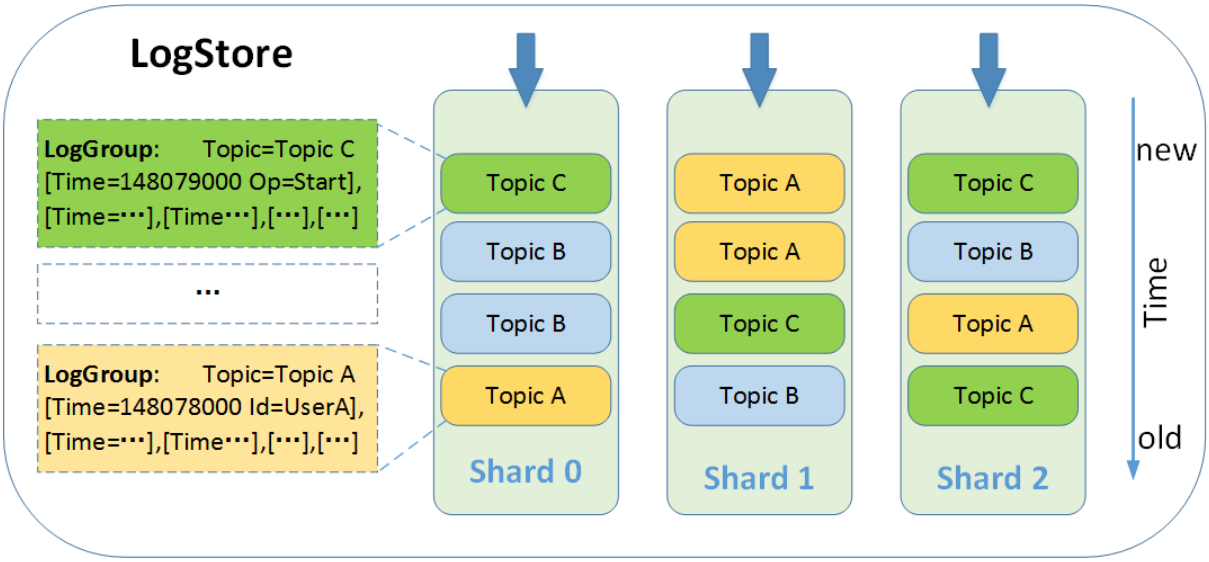

6.10. Trace

You can use a trace to record the processing information of a request, such as service calls and processing duration.

A trace corresponds to a call chain. Generally, a call chain represents the execution process of an incident or the process in a distributed system. In the OpenTracing standard, a call chain is a directed acyclic graph (DAG) that consists of multiple spans. Each span represents a named and timed segment that is continuously run in the trace.

The following figure shows an example of a distributed call. When the client initiates a request, the request is sent to the load balancer, processed by the authentication service and billing service, and then sent to the requested resources. Then, a result is returned to the client.

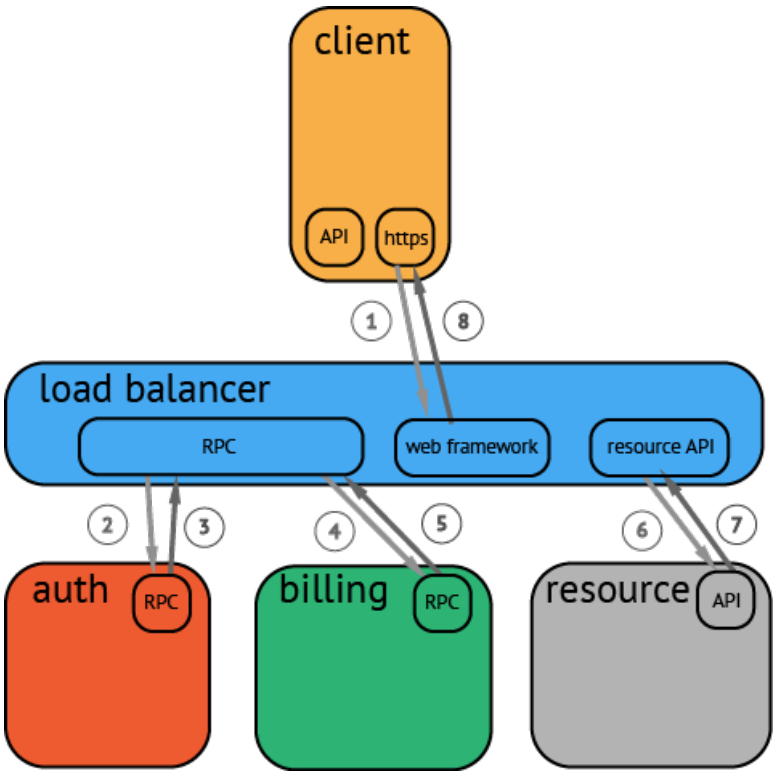

7.Limits

7.1. Basic resources

This topic describes the limits on the basic resources of Log Service.

Item	Limit	Adjustable
Project	You can create a maximum of 50 projects for an Alibaba Cloud account.	To increase the quota, submit a ticket.
Logstore	You can create a maximum of 200 Logstores for a project.	To increase the quota, submit a ticket.
Shard	• You can create a maximum of 400 shards for a project. • When you create a Logstore by using the console, you can create a maximum of 10 shards in the Logstore. When you create a Logstore by calling the API, you can create a maximum of 100 shards in the Logstore. You can split shards to increase the number of shards by using one of the preceding two methods.	To increase the quota, submit a ticket.
Logtail configuration	You can create a maximum of 200 Logtail configurations for a project.	To increase the quota, submit a ticket.
Log retention period	You can permanently retain logs. You can also specify a log retention period. Unit: days. Valid values: 1 to 3000.	N/A
Machine group	You can create a maximum of 200 machine groups for a project.	To increase the quota, submit a ticket.
Consumer group	You can create a maximum of 30 consumer groups for a Logstore.	You can delete consumer groups that are no longer used.
Saved search	You can create a maximum of 100 saved searches for a project.	N/A

Item	Limit	Adjustable
Dashboard	• You can create a maximum of 100 dashboards for a project. • You can add a maximum of 100 charts to a dashboard.	N/A
LogItem	• The maximum size of a log entry that is collected by calling the API is 1 MB. • The maximum size of a log entry that is collected by using Logtail is 512 KB.	N/A
Field name (key)	The maximum size of a field name (key) is 128 bytes.	N/A
Field value (value)	The maximum size of a field value (value) is 1 MB.	N/A
Log group	The maximum size of a log group is 5 MB.	N/A
Alert	You can create a maximum of 100 alerts for a project.	To increase the quota, submit a ticket.

7.2. Data read and write

This topic describes the limits of data read and write in Log Service.

Resource	Item	Limit	Description
Project	Write traffic	The maximum write traffic of raw data is 30 GB per minute.	If the limit is reached, the HTTP status code 403 and the "Inflow Quota Exceed" error message are returned. If you want to increase the limit, submit a ticket.
	Number of write operations	The maximum number of write operations is 600,000 per minute.	If the limit is reached, the HTTP status code 403 and the "Write QPS Exceed" error message are returned. If you want to increase the limit, submit a ticket.
	Number of read operations	The maximum number of read operations is 600,000 per minute.	If the limit is reached, the HTTP status code 403 and the "Read QPS Exceed" error message are returned. If you want to increase the limit, submit a ticket.

Resource	Item	Limit	Description
Shard	Write traffic	• If indexes are configured in a Logstore, the maximum write traffic of raw data is 5 MB per second. • If no indexes are configured in a Logstore, the maximum write traffic of raw data is 10 MB per second.	Not required. If the limit is reached, Log Service continues to read data. However, the quality of the service may be degraded.
	Number of write operations	The maximum number of write operations is 500 per second.	Not required. If the limit is reached, Log Service continues to read data. However, the quality of the service may be degraded.
	Read traffic	The maximum read traffic is 10 MB per second.	Not required. If the limit is reached, Log Service continues to read data. However, the quality of the service may be degraded.
	Number of read operations	The maximum number of read operations is 100 per second.	Not required. If the limit is reached, Log Service continues to read data. However, the quality of the service may be degraded.

7.3. Logtail

This topic describes the limits of Logtail.

Limits on log files

Item	Description
Log file encoding	Logtail supports log files that are encoded in UTF-8 and GBK. We recommend that you use UTF-8-encoded log files to improve processing performance. If log files are encoded in other formats, errors such as garbled characters and data loss may occur.
Log file size	The log file size is unlimited.
Log file rotation	Logtail supports log file rotation. Both `.log*` and `.log` files are supported in rotation.
Log collection behavior performed when log parsing is blocked	When log parsing is blocked, Logtail keeps the log file descriptor (FD) open. If log file rotation occurs multiple times during the blocking period, Logtail attempts to ensure that new log files are parsed in sequence. If the number of new files that are not parsed exceeds 20, Logtail does not process the excess log files.

Item	Description
Symbolic link	Monitored directories can be symbolic links.
Log size	The maximum size of a log is 512 KB. If a multi-line log is split by using a regular expression to match the first line, the maximum size of each log after splitting is still 512 KB. If the size of a log exceeds 512 KB, the log is forcibly split into multiple parts for collection. For example, if the size of a log is 1,025 KB, the log is split into three parts: 512 KB, 512 KB, and 1 KB. Then, the log parts are collected in sequence.
Regular expression	Logtail uses regular expressions that are compatible with Perl.
Multiple Logtail configurations for the same log file	You cannot use multiple Logtail configurations for the same log file. We recommend that you store data collected from a log file to the same Logstore. You can configure multiple tasks to subscribe to logs collected from different files. If you want to use multiple Logtail configurations for the same log file, configure symbolic links for log files to bypass this limit.
File opening behavior	When Logtail collects data from a log file, Logtail keeps the log file open. If the log file is not modified for more than 5 minutes and log rotation does not occur, Logtail closes the log file.
First log collection behavior	Logtail collects data only from incremental log files. If the size of a log file exceeds 1 MB the first time the modification to the log file is detected, Logtail collects data from the last 1 MB. If the log file size does not exceed 1 MB, Logtail collects data from the beginning of the log file. If the log file is not modified after the Logtail configuration is delivered, Logtail does not collect data from the log file.
Non-standard text logs	If a log contains \0, the log is truncated at the first occurrence of \0.

Limits on checkpoints

Item	Description
Checkpoint timeout period	If a log file is not modified for more than 30 days, the checkpoint of the log file is deleted.
Checkpoint storage policy	Checkpoints are stored at intervals of 15 minutes and at the point in time when Logtail exits. For more information about how to change the values of the related parameters, see 设置Logtail启动参数.
Checkpoint storage path	By default, checkpoints are stored in the `/tmp/logtail_checkpoint` directory. For more information about how to change the values of the related parameters, see 设置Logtail启动参数.

Limits on Logtail configurations

Item	Description

Item	Description
Configuration update	A configuration update requires approximately 30 seconds to take effect.
Dynamic loading of Logtail configurations	Logtail configurations can be dynamically updated. An update of a Logtail configuration does not affect other Logtail configurations.
Number of Logtail configurations	The number of Logtail configurations is unlimited. However, we recommend that you create no more than 100 Logtail configurations for a server.
Multi-tenant isolation	Logtail configurations are isolated.

Limits on resources and performance metrics

Item	Description
Throughput for log processing	The default transmission speed of raw logs is limited to 20 MB/s. Log data is uploaded after it is encoded and compressed. The compression ratio ranges from 5:1 to 10:1. If the speed exceeds the limit, log data may be lost. For more information about how to change the values of the related parameters, see 设置Logtail启动参数.
Maximum processing speed for logs	Single-core processing speed: The maximum processing speed is 100 MB/s in simple mode, 40 MB/s in delimiter mode, and 30 MB/s in JSON mode. By default, the maximum processing speed is 20 MB/s in full regex mode. The maximum processing speed in full regex mode varies based on the complexity of regular expressions. If multiple processing threads are started, the performance can be improved by 150% to 300%.
Number of monitored directories	Logtail limits the depth of monitored directories to reduce the consumption of user resources. If the upper limit is reached, Logtail stops monitoring additional directories or log files. Logtail can monitor a maximum of 3,000 directories, including subdirectories.
Number of monitored files	By default, a Logtail configuration on each server can be used to monitor a maximum of 10,000 files. The Logtail on each server can monitor a maximum of 100,000 files. Excessive files are not monitored. If the upper limit is reached, you can perform the following operations: • Use more exact names to specify the monitored directories in each Logtail configuration. • Increase the value of the mem_usage_limit parameter to raise the threshold of available Logtail memory resources. For more information, see 设置Logtail启动参数. You can raise the threshold to no more than 2 GB. This way, the maximum number of files that can be monitored by using each Logtail configuration is increased to 100,000, and the maximum number of files that the Logtail on each server can monitor is increased to 1,000,000.

Item	Description
Default resources	By default, Logtail can occupy up to 40% of the CPU and 256 MB of memory. If logs are generated at a high speed, you can change the values of the related parameters. For more information, see 设置Logtail启动参数.
Policy used to process excessive resource consumption	If the amount of resources occupied by Logtail remains higher than the upper limit for more than 5 minutes, Logtail is forcibly restarted. The restart may cause data loss or duplication.

Limits on troubleshooting

Item	Description
Network error handling	If a network error occurs, Logtail automatically retries the data collection task and adjusts the retry interval.
Processing of threshold-crossing events	If a data transmission speed exceeds the upper limit of a Logstore, Logtail blocks log collection and automatically retries the data collection task.
Maximum retry period before timeout	If data fails to be transmitted and the issue lasts for more than 6 hours, Logtail discards the data.
Status self-check	If an exception occurs, Logtail restarts. For example, if an application unexpectedly exits or the resource usage exceeds the specified upper limit, Logtail restarts.

Other limits

Item	Description
Log collection latency	In most cases, a latency of less than 1 second exists between the point in time at which a log is written to disk and the point in time at which Logtail collects the log. However, if the log collection is blocked, the latency increases.
Log upload policy	Before Logtail uploads logs, Logtail aggregates the logs in the same file. Logtail starts to upload logs when the number of logs exceeds 2,000, the total size of logs exceeds 2 MB, or the log collection duration exceeds 3 seconds.

7.4. Query and analysis

This topic describes the limits of query and analysis in Log Service.

Query

Item	Description	Remarks

Item	Description	Remarks
Number of keywords	The number of keywords that are used as search conditions. The number of logical operators is not included. You can specify up to 30 keywords in a search statement.	None
Size of a field value	The maximum size of a field value is 10 KB. The excess part is not involved in searching.	If the size of a field value is greater than 10 KB, logs may fail to be obtained by using keywords, but the logs are actually stored in the Logstore.
Maximum number of concurrent search statements	Each project supports up to 100 concurrent search statements.	For example, 100 users can concurrently execute search statements in all Logstores of a project.
Returned result	The returned logs are displayed on multiple pages. Each page displays up to 100 logs.	None
Maximum size of a log	Log Service performs the Document Object Model (DOM) operation only on the first 10,000 characters of a log due to browser performance limits.	If a log contains more than 10,000 characters, the following message appears in the Log Service console: The log contains log data of more than 10,000 characters, and some display will be downgraded.
Fuzzy search	In a fuzzy search, Log Service matches up to 100 words that meet the specified conditions and returns the logs that meet the search conditions and contain one or more of these words. For more information, see Fuzzy search.	None
Data sorting in search results	By default, search results are displayed in descending order of time, which is accurate to minutes.	None

Analytics

Item	Standard SQL	Dedicated SQL
Number of concurrent analytic statements	Each project supports a maximum of 15 concurrent analytic statements. For example, 15 users can concurrently execute analytic statements in all Logstores of a project.	Each project supports a maximum of 100 concurrent analytic statements. For example, 100 users can concurrently execute analytic statements in all Logstores of a project.

Item	Standard SQL	Dedicated SQL
Data volume	Each shard supports only 1 GB of data for a single analytic statement.	An analytic statement can scan a maximum of 200 billion rows of data at the same time.
Method to enable	By default, Standard SQL is enabled.	A switch is provided for you to manually enable Dedicated SQL. For more information, see Enable Dedicated SQL.
Resource usage fee	Free of charge.	You are charged based on the actual CPU time. For more information, see Billable items.
Applicable scope	You can analyze only the data that is written to Log Service after the log analysis feature is enabled. If you want to analyze historical data, you must reindex the historical data. For more information, see Reindex logs for a Logstore.	You can analyze only the data that is written to Log Service after the log analysis feature is enabled. If you want to analyze historical data, you must reindex the historical data. For more information, see Reindex logs for a Logstore.
Returned result	By default, an analytic statement returns a maximum of 100 rows of data. If you want to view more data, use a LIMIT clause. For more information, see LIMIT clause.	By default, an analytic statement returns a maximum of 100 rows of data. If you want to view more data, use a LIMIT clause. For more information, see LIMIT clause.
Size of a field value	The log analysis feature can analyze a maximum of 16,384 bytes (16 KB) of data in the value of each field. If the size of a field value exceeds 16 KB, the excess content is not analyzed. You can change the maximum size for each field value when you configure indexes. Valid values: 64 to 16384. Unit: bytes. For more information, see Configure indexes.	The log analysis feature can analyze a maximum of 16,384 bytes (16 KB) of data in the value of each field. If the size of a field value exceeds 16 KB, the excess content is not analyzed. You can change the maximum size for each field value when you configure indexes. Valid values: 64 to 16384. Unit: bytes. For more information, see Configure indexes.
Timeout period	The maximum timeout period for a single analytic statement is 55 seconds.	The maximum timeout period for a single analytic statement is 55 seconds.
Number of decimal places in the value of a double-type field	The value of a double-type field can contain a maximum of 52 decimal places. If the number of decimal places is greater than 52, the accuracy of the field value is compromised.	The value of a double-type field can contain a maximum of 52 decimal places. If the number of decimal places is greater than 52, the accuracy of the field value is compromised.

7.5. Alerting

This topic describes the limits of the alerting feature in Log Service.

Alerting (New)

Category	Item	Limit
Monitoring and alerting	Number of alert monitoring rules	You can create up to 100 alert monitoring rules in each project. If you want to increase the quota, you can submit a ticket. The quota can be increased to 200 in each project.
	Query statements	For more information about the limits on query statements, see Query and analysis.
	Results of a single query statement	If more than 1,000 rows are returned for a single query statement, only the first 1,000 rows are used for a set operation. If you specify three query statements but do not select **No Merge**, only the first 100 rows that are returned for each query are used for set operations.
	Number of query statements for set operations	You can specify one to three query statements for set operations.
	Field value length	If the length of a field value exceeds 1,024 characters, the field value is truncated, and only the first 1,024 characters are used for analysis.
	Time range for a query statement	The time range for a query statement cannot exceed 24 hours.
Alert management	Regions in which incident management is supported	Incident management is supported in the following regions: • Chinese mainland: China (Chengdu), China (Heyuan), China (Hohhot), China (Guangzhou), China (Qingdao), China (Ulanqab), China (Hong Kong), China (Hangzhou), China (Shenzhen), China East 2 Finance, and China South 1 Finance • Outside the Chinese mainland: Singapore (Singapore), UAE (Dubai), Malaysia (Kuala Lumpur), India (Mumbai), Japan (Tokyo), Germany (Frankfurt), UK (London), Australia (Sydney), US (Silicon Valley), and US (Virginia)

Category	Item	Limit
	Number of incidents	Up to 1,000 incidents can be retained in 30 days. If the number of incidents exceeds 1,000, the earliest incident is automatically overwritten by the latest incident.
	Number of incident comments	Up to 10 comments can be added for each incident. If the number of comments for an incident exceeds 10, the earliest comment is automatically overwritten by the latest comment.
	Notification methods	• Voice Call This notification method supports only mobile phone numbers in the Chinese mainland, which are prefixed with 86. **Note** ◦ If a voice call is not answered, Log Service sends a text message. ◦ You are charged for a voice call regardless of whether the call is answered. You are not charged for the text message that is sent upon a non-answered voice call. • DingTalk ◦ Only custom keywords are supported in the security settings of a DingTalk chatbot. ◦ Each DingTalk chatbot can send up to 20 alert notifications every minute. • Enterprise WeChat ◦ Only custom keywords are supported in the security settings of an Enterprise WeChat robot. ◦ Each Enterprise WeChat robot can send up to 20 alert notifications every minute. • Lark ◦ Only custom keywords are supported in the security settings of a Lark robot. ◦ Each Lark robot can send up to 20 alert notifications every minute. ◦ You can set **Notified Contacts** to **No Reminder** or **All**. You cannot set the parameter to **Specified Members**. • Webhook-Custom ◦ The webhook URL must be accessible over the Internet.

Category	Item	Limit
Notification management		◦ If a webhook is successfully invoked, the HTTP status code 200 is returned. If the HTTP status code 200 is not returned, Log Service considers the request as failed. • Function Compute Only the functions that start with `sls-ops-` are supported. For more information, see Notification methods.
	Notification content	• SMS Message A notification can contain up to 256 characters. • Voice Call A notification can contain up to 256 characters. • Email A notification cannot exceed 8 KB in size. • DingTalk A notification cannot exceed 8 KB in size. • Enterprise WeChat ◦ A notification cannot exceed 4 KB in size. ◦ If you set **Notified Contacts** to **All** or **Specified Members**, Enterprise WeChat supports only plain text in alert notifications, instead of the Markdown syntax. • Lark A notification cannot exceed 8 KB in size. • Slack A notification cannot exceed 8 KB in size. • Webhook-Custom A notification cannot exceed 16 KB in size. • Notifications A notification cannot exceed 8 KB in size. • Function Compute A notification cannot exceed 16 KB in size. • EventBridge A notification cannot exceed 16 KB in size.

Category	Item	Limit
	Alert template variables	The value of a variable cannot exceed 2 KB in size. If a value exceeds 2 KB in size, the value is truncated.
	Notification quota	A recipient can receive up to 9,999 emails, text messages, or phone calls per day. For more information, see Configure notification quotas.

Alerting (Old)

Item	Description
Associated query statements	You can associate an alert rule with a maximum of three query statements.
Field value size	If a field value exceeds 1,024 characters in length, Log Service extracts only the first 1,024 characters for data processing.
Trigger condition	The trigger condition has the following limits: • Each trigger condition must be 1 to 128 characters in length. • If a query result includes more than 100 rows, Log Service only checks whether the first 100 rows meet the trigger condition. • Log Service checks whether a trigger condition is met for a maximum of 1,000 times for the specified query statements.
Query time range	The maximum time range that you can specify for each query is 24 hours.
Voice calls	If a voice call is not answered, Log Service sends an SMS notification. You are charged only once for the voice call regardless of whether the call is answered. The SMS notification does not incur fees.

7.6. Log applications

This topic describes the limits of log applications.

Log Audit Service

8.Security and compliance

8.1. Overview

Log Service provides security protection features such as server-side encryption, fine-grained access control, and service logs. Log Service also provides Log Audit Service that allows you to collect and audit the logs of Alibaba Cloud services in real time across multiple accounts in an automatic and centralized manner. Log Service provides data insight capabilities and can meet your business requirements for data security and compliance.

Compliance certifications

Log Service is certified with the following standards:

- ISO 9001, ISO 20000, ISO 27001, ISO 27017, ISO 27018, ISO 22301, ISO 27701, and ISO29151
- BS10012
- CSA STAR
- MLPS level 3
- SOC
- C5
- Hong Kong Standard on Auditing (HKSA)
- Philippine Standards on Auditing (PSA)
- MTCS
- OSPAR
- PCI DSS

Security capabilities

The following table describes the security features that are provided by Log Service.

Feature	Description
Access control	Log Service provides features that can be used together with Resource Access Management (RAM) policies and Security Token Service (STS) temporary credentials. You can use these features to manage resource access in Log Service.
Data encryption	Log Service supports server-side encryption and encrypted transmission based on the SSL or TLS protocol to protect data from potential security risks on the cloud.
Data reliability	The underlying storage system of Log Service uses a three-replica mechanism to ensure high reliability.
Log auditing and monitoring	Log Service provides the service log feature and integrates with CloudMonitor to help you monitor and audit your service data.

Feature	Description
Log auditing for cloud services	Log Service provides Log Audit Service that allows you to collect and audit the logs of Alibaba Cloud services in real time across multiple accounts in an automatic and centralized manner.

8.2. Access control

Log Service provides features that can be used together with Resource Access Management (RAM) policies and Security Token Service (STS) temporary credentials. You can use these features to manage resource access in Log Service.

User-based RAM policies

RAM is a resource access control service provided by Alibaba Cloud. You can configure RAM policies based on users. You can manage user permissions when you configure RAM policies. You can create RAM users for employees, systems, and applications. You can grant users permissions to access the resources of your Alibaba Cloud account. You can also manage the permissions that are granted to specific users on specific resources. For example, you can create a RAM policy to grant users read-only permissions on specific resources in a project or a Logstore.

A RAM policy is in the JSON format. You can write a RAM policy when you specify the Action, Effect, Resource, and Condition elements in the Statement field. You can add multiple statements to a policy to help you manage authorization in a more efficient manner.

For more information, see RAM overview.

Temporary access authorization based on STS

RAM policies allow you to access resources for a long period of time. If you want users to access resources only for a short period of time, you can use STS to create temporary credentials. You can call STS API operations to obtain temporary AccessKey pairs and tokens. Then, you can send the AccessKey pairs and tokens to temporary users to access Log Service. The permissions that are obtained by using STS are restricted and have time limits. The risk of temporary credentials being leaked does have the same level of risks as other credentials.

You can use STS to grant temporary access to Log Service. You can use STS to grant a third-party application or a RAM user that you manage an access credential that has a custom validity period and custom permissions.

For more information, see Use STS to enable cross-account access to Log Service resources.

8.3. Data encryption

Log Service provides server-side encryption and encrypted transmission based on the SSL or TLS protocol to protect data from potential security risks on the cloud.

Server-side encryption

Log Service allows you to use Key Management Service (KMS) to encrypt data for secure storage. KMS is a secure and easy-to-use management service that is provided by Alibaba Cloud. You can use KMS to ensure the privacy, integrity, and availability of your keys at low cost. You can use the keys in a secure and convenient manner. You can also develop encryption and decryption solutions based on your business requirements. You can view and manage the keys in the KMS console. For more information, see Overview.

KMS stores and manages customer master keys (CMKs) that are used to encrypt data keys. KMS also generates data keys that can be used to encrypt and decrypt large amounts of data. Envelope encryption provided by KMS can protect your data and data keys from unauthorized access. You can use the default CMK stored in KMS or generate a CMK by using your BYOK materials or the BYOK materials that are provided by Alibaba Cloud.

Encrypted transmission based on the SSL or TLS protocol

Log Service can be accessed by using HTTP or HTTPS. SSL or TLS is a cryptographic protocol that provides secure communication and ensures data integrity between a web server and a client.

8.4. Data reliability

Log Service uses a three-replica mechanism to ensure high data reliability.

The underlying storage system of Log Service uses a three-replica mechanism to ensure high data reliability. Three replicas are provided for all data. Then, the replicas are stored in different servers of a cluster by using a specific distributed storage algorithm. This way, the storage system ensures that three data replicas are distributed to different physical disks on different servers. If a single hardware device fails, it does not cause data loss. The data in the three replicas remains consistent.

8.5. Log auditing and monitoring

Log Service provides the service log feature and integrates with CloudMonitor to help you monitor and audit your service data.

Service logs

The service log feature of Log Service allows you to record operational events in a project into logs. This feature also provides dashboards that allow you to analyze data from multiple dimensions. You can use this feature to view the service status of Log Service in real time and improve O&M efficiency.

When the service log feature is enabled for a project, the generated log data is classified and stored in one of the dedicated Logstores. By default, Log Service automatically creates the following two dedicated Logstores:

For more information, see Service logs.

CloudMonitor

You can use CloudMonitor to monitor the metrics of Log Service. The metrics include write traffic, overall QPS, and service status. You can configure alert rules to monitor log collection and shard usage. You can also detect related exceptions.

For more information, see Use CloudMonitor.

8.6. Log auditing for cloud services

Log Service is an observability platform that you can use to audit the logs of multiple Alibaba Cloud services.

However, you can read data only from the Logstores in Log Audit Service. In addition, the log data of other cloud services cannot be written to the dedicated Logstore of a cloud service. You cannot write data to the Logstores or change the data in the Logstores. Log Audit Service allows you to collect and audit the logs of Alibaba Cloud services in real time across multiple accounts in an automatic and centralized manner. Log Audit Service provides the following benefits:

- Centralized log collection
  - Log collection across accounts: You can collect logs from multiple Alibaba Cloud accounts to a project within one Alibaba Cloud account. You can configure multi-account collection in custom authentication mode or resource directory mode. The resource directory mode is recommended. For more information, see Configure multi-account collection.
  - Ease of use: You need to only configure collection policies once. Then, Log Audit Service collects logs in real time from Alibaba Cloud resources that belong to different accounts when new resources are detected. The new resources include newly created ApsaraDB RDS instances, SLB instances, and OSS buckets.
  - Centralized storage: Logs are collected and stored in the central project of a region. This way, you can query, analyze, and visualize the collected logs in a more efficient manner. You can also configure alerts for the logs and perform secondary development.
- Comprehensive audit
  - Log Audit Service supports all features of Log Service. For example, you can query, analyze, transform, visualize, and export logs, and configure alerts for logs. Log Audit Service also allows you to audit logs in a centralized manner.
  - You can use Log Audit Service together with Alibaba Cloud services, open source software, and third-party SOCs to create more value from data.

For more information, see Log Audit Service.

9.Supported regions

A region is a geographical area where data centers reside. After a project is created, you cannot change the region of the project. This topic describes the regions that are supported by Log Service.

The following table describes the mappings among regions, cities, and region IDs.

Region	City	Region ID
China (Qingdao)	Qingdao	cn-qingdao
China (Beijing)	Beijing	cn-beijing
China (Zhangjiakou)	Zhangjiakou	cn-zhangjiakou
China (Hohhot)	Hohhot	cn-huhehaote
China (Ulanqab)	Ulanqab	cn-wulanchabu
China (Hangzhou)	Hangzhou	cn-hangzhou
China (Shanghai)	Shanghai	cn-shanghai
China (Shenzhen)	Shenzhen	cn-shenzhen
China (Heyuan)	Heyuan	cn-heyuan
China (Guangzhou)	Guangzhou	guangzhou
China (Chengdu)	Chengdu	cn-chengdu
China (Hong Kong)	Hong Kong	cn-hongkong
Singapore (Singapore)	Singapore	ap-southeast-1
Australia (Sydney)	Sydney	ap-southeast-2
Malaysia (Kuala Lumpur)	Kuala Lumpur	ap-southeast-3
Indonesia (Jakarta)	Jakarta	ap-southeast-5
Philippines (Manila)	Manila	ap-southeast-6
Thailand (Bangkok)	Bangkok	ap-southeast-7
India (Mumbai)	Mumbai	ap-south-1
Japan (Tokyo)	Tokyo	ap-northeast-1
South Korea (Seoul)	Seoul	ap-northeast-2
US (Silicon Valley)	Silicon Valley	us-west-1

Region	City	Region ID
US (Virginia)	Virginia	us-east-1
Germany (Frankfurt)	Frankfurt	eu-central-1
UK (London)	London	eu-west-1
UAE (Dubai)	Dubai	me-east-1
Russia (Moscow)	Moscow	rus-west-1

10.FAQ

This topic provides answers to some frequently asked questions that may help you prepare to use Log Service.

What is Log Service?

Log Service is a cloud-native observation and analysis platform that provides large-scale, low-cost, and real-time services to process multiple types of data such as logs, metrics, and traces. Log Service allows you to collect, transform, query, analyze, visualize, ship, and consume data. You can also configure alerts in the console. Log Service helps enterprises improve their digital capabilities in terms of R&D, O&M, and data security.

How can Log Service contribute to the success of business?

- Log storage and analysis: You can use Log Service to store and analyze logs that are collected from infrastructures, applications, containers, and cloud services. You can also use Log Service to store and analyze website access logs, metrics, and traces.
- Intelligence and operations efficiency: Log Service supports intelligent clustering, anomaly detection, and root cause analysis.
- Alert monitoring: You can use Log Service to monitor business applications and user experience data, visualize data, and configure alerts.
- Troubleshooting: You can use Log Service to monitor infrastructures, identify issues, and improve troubleshooting efficiency.
- Log audit: You can use Log Service to manage and audit logs for multiple accounts.
- Unified data portal: You can use Log Service to store heterogeneous data. Log Service provides a stable monitoring backend platform and can further connect to a real-time computing platform.
- Data processing: You can use the data transformation feature of Log Service to cleanse, mask, and enrich data. You can also use the data shipping feature to ship data to data warehouses or data lakes.

What are the benefits of Log Service?

- Unified access: Log Service provides a unified portal for logs, metrics, and traces. Log Service can collect data from various data sources, such as IoT, mobile apps, and servers. You can also use Log

Service to collect logs from cloud services, open source systems, multi-cloud environments, and on-premises servers.

- End-to-end monitoring platform: Log Service allows you to collect, transform, analyze, visualize, and ship data. You can also configure alerts for the data and view the alerts in the console.
- Intelligence and efficiency: Log Service supports intelligent anomaly detection and root cause analysis. You can analyze tens of billions of data records in seconds.
- Scalability and low costs: Log Service provides PB-level scalability per day and is fully managed and maintenance-free. The pay-as-you-go billing method is supported. You are charged only for the amount of consumed resources. The total cost of ownership (TCO) can be reduced by more than 50%.

Which types of data can be collected by Log Service?

Log Service can collect the following types of data:

- Log
- Metric
- Trace

Which methods can be used by Log Service to collect data?

Log Service supports more than 50 data collection methods. Log Service allows you to collect text logs and mobile app data by using Logtail. You can upload logs by using the Kafka or Syslog protocol. You can also collect logs from Alibaba Cloud services. For more information, see Collect logs, Collect metrics, and Collect traces.

Does Alibaba Cloud use the data that I store in Log Service?

Alibaba Cloud does not use or disclose your data without authorization. Alibaba Cloud processes user data only based on your service requirements or the requirements of laws and regulations. For more information, see Service terms.

Does Alibaba Cloud store their own data in Log Service?

Yes. Log Service is the internal logging and monitoring platform of Alibaba Group. Log Service has demonstrated its stability and reliability during Double 11 events over the years. Alibaba Cloud developers also use Log Service in many projects.

How does Log Service ensure its stability during traffic spikes?

Log Service provides a self-adaptive data architecture that supports auto scaling to process petabytes of data per day. This architecture helps Log Service handle traffic spikes and business peaks.

How do I store data in Log Service at low costs?

You can use the data shipping feature to ship your data from Log Service to Object Storage Service (OSS) for low-cost storage. For more information, see What is OSS? OSS provides the following storage classes to cover various data storage scenarios from hot data storage to cold data storage: Standard, Infrequent Access (IA), Archive, and Cold Archive.

How long can I store data in Log Service?

Log Service can permanently store your data. You can also modify the data retention period based on your business requirements.

How do I get started with Log Service?

Before you use Log Service, make sure that you have an Alibaba Cloud account and have completed real-name verification. If you do not have an Alibaba Cloud account, you are prompted to create an Alibaba Cloud account when you activate Log Service. For more information, visit the Create Your Alibaba Cloud Account page.

After you create an Alibaba Cloud account, open the Log Service product page. Then, click **Get it Free**.

The default billing method is pay-as-you-go. To reduce costs, we recommend that you purchase a subscription resource plan. For more information, see Log Service pricing.

How am I charged for Log Service?

You are separately charged for the billable items of Log Service. For example, if you collect logs, you are charged a write traffic fee. If you store logs, you are charged a log storage fee. For more information, see Billable items.

What do I do if I have specific requirements on Log Service?

If you have specific requirements on Log Service, submit a ticket to contact the technical support team of Log Service.

11.Case studies

11.1. Chanjet

The O&M and development team of Chanjet uses Log Servcie to prevent frequent false alerts, detect suspicious sites, and identify errors. Log Service helps Chanjet improve the O&M efficiency and reduce the costs of O&M and communication. Log Service supports the healthy and stable operation of all cloud services of Chanjet and builds a benchmark in the field of IT O&M.

Company profile

Chanjet Information Technology Co., Ltd. is a member enterprise of Yonyou. Chanjet aims to provide social, personalized, service-oriented, and small-scale business management support for micro and small enterprises. To help micro and small enterprises to fulfill digital transformation in terms of finance and management, Chanjet empowers these enterprises with advanced technologies. Chanjet helps them promote online business, change traditional business models, and maintain sustainable profit growth. To meet the various requirements of micro and small enterprises on cloud services, Chanjet utilizes the high-frequency interactions between Software as a Service (SaaS) business and customers to extract more value for the customers. Chanjet plans to expand its business from the SaaS market to the Backend as a Service (BaaS) market. Chanjet aims to become the largest one-stop service platform for micro and small enterprises in China. For more information, visit the official website of Chanjet.

Scenarios

The IT O&M and development team of Chanjet is responsible for developing cloud services, such as Haokuaiji, Haoshengyi, and Yidaizhang. This process includes the O&M, deployment, and release of these services. The team built a MIDAS intelligent O&M platform that provides data import, data processing, and scenario analysis features.

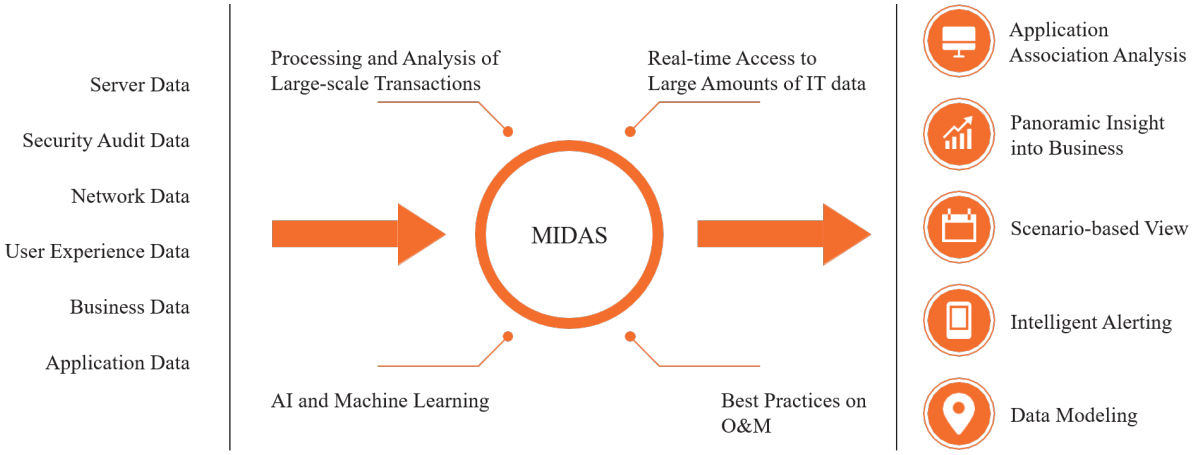

Challenges

In the early development stage of the MIDAS intelligent O&M platform, Chanjet used the self-managed Elasticsearch, Logstash, and Kibana (ELK) stack to perform data analysis for O&M. In the process of business growth, an increasing number of applications and systems are connected to the MIDAS intelligent O&M platform. As a result, Chanjet faces various challenges, such as platform issues and service stability issues.

- High concurrency

Tens of thousands of sites send data at the same time. As a result, terabytes of logs and messages are generated per day. The self-managed ELK stack delivers poor performance, and a large number of development resources are required to improve the performance.

- Heterogeneous data

  Different types of data such as access logs, system logs, application logs, notifications, and messages are collected in different formats. As a result, data cleansing is difficult.

- Multiple sources

  Logs are collected from different data sources, such as networks, servers, mobile apps, websites, and Docker containers. The APIs that are used to import data require high real-time performance. In this case, unified data management is required.

- Diversified requirements

  Each product department of Chanjet has specific requirements for the collected data. For example, a department needs to monitor alerts, troubleshoot issues, analyze or mine data, perform data mining, and subscribe to data analysis reports. In addition, the department needs to consume the data by using different methods.

Solution

To resolve the preceding issues, Chanjet used Log Service to build the MIDAS intelligent O&M platform.

- Efficient data collection and transmission

  Log Service has powerful data import capabilities. Chanjet uses Log Service to import different types of data that is in different formats from the networks, servers, mobile apps, and containers in its hybrid cloud architecture. The data includes access logs, system logs, application logs, and messages. Terabytes of data can be efficiently written to Log Service per day.

- Flexible data processing and storage

  Chanjet has deployed a configuration management database (CMDB) and defined the relationships between the resources and each department. Chanjet uses Log Service to segment and serialize raw logs, and then analyzes the logs based on specific scenarios. Chanjet ships data to external services based on specific shipping policies. Then, the external services use Ansible or forward messages to manage the shipped data in a centralized manner and support subsequent scenario analysis.

- Intelligent anomaly detection and location

  Chanjet built the MIDAS intelligent O&M platform by using the time series data analysis feature and supported functions of Log Service. For example, Chanjet uses the interval-valued comparison and periodicity-valued comparison functions to obtain the latest value of a metric in real time. After the specific values are obtained, the related alerts that are triggered based on the values have higher accuracy than the alerts that are triggered based on the original threshold. Chanjet uses the prediction and anomaly detection functions of Log Service to quickly locate exceptions from a large number of metrics, mark the exceptions, and identify system failures in a short time. To predict system failures, Chanjet uses Log Service to collect data from all modules of the MIDAS intelligent O&M platform, marks the collected data, and integrates application configurations with the data. Then, Chanjet performs root cause analysis on the system failures based on the time series.

The following figure shows the architecture of the MIDAS intelligent O&M platform that is built based on Log Service by Chanjet.

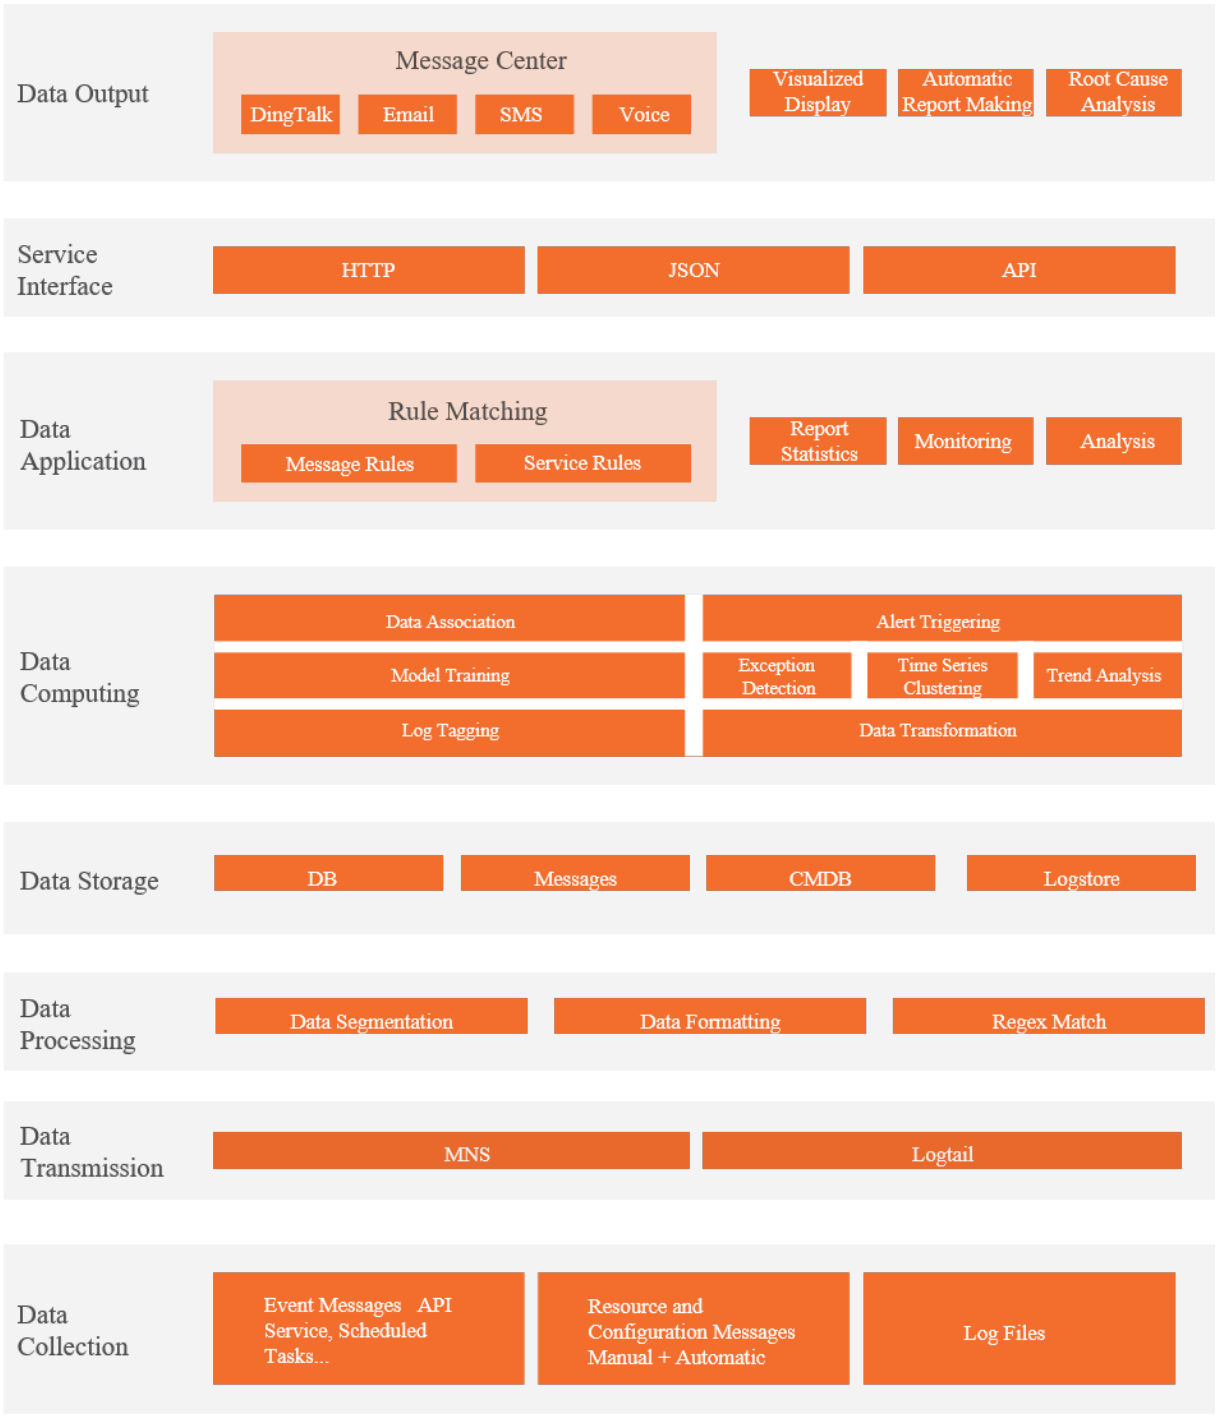

11.2. miHoYo

Log Service teams up with miHoYo to support the online game named Genshin Impact from the internal test and public preview to the commercial release. Log Service helps miHoYo manage logs from tens of millions of players. Log Service is highly recognized and praised by miHoYo for outstanding performance and stability.

Company profile

miHoYo, founded in 2012, focuses on mobile games and comics that are originated from Chinese animation. miHoyo is a video game development and animation studio that has successively released excellent and widely favored mobile games based on Chinese animations, such as Zombiegal Kawaii, Guns GirlZ, and Honkai Impact 3. In 2020, miHoYo released Genshin Impact, which is the first open-world mobile game of anime, comics, and games (ACG) subculture. Only one month after its release, Genshin Impact has became the most profitable mobile game around the world compared with its counterparts. For more information, visit the official website of miHoYo.

Scenarios

Genshin Impact is an open-world adventure game that is developed and released by miHoYo. Genshin Impact was available for public preview on September 28, 2020. After Genshin Impact was released, it turned to be a phenomenal success around the world. Genshin Impact has become a popular game with a large number of players even though the game is a new comer in the market. Therefore, miHoYo requires a logging platform with high stability, elastic scalability, and high performance to meet its growing requirements on data analysis. This way, miHoYo can fulfill intensive operations for Genshin Impact.

Challenges

The Genshin Impact team faces the following challenges after Genshin Impact was released:

- Data growth

  Only one month after the release of Genshin Impact, the number of active players per month exceeded 10 million. The data volume is also increasing at a high speed with the rapid growth of players and the optimization of the operation system. Therefore, the Genshin Impact team requires a solution to query and analyze large amounts of data with high performance.

- Centralized data collection

  Genshin Impact becomes popular worldwide and tops the App Store bestseller list in the United States, Germany, Canada, South Korea, and Singapore. However, the globalization of the game also increases the difficulty in centralized data collection and requires higher stability.

- Traffic surges

  Genshin Impact is an open-world adventure game that offers promotional activities based on different seasons and cultural backgrounds of each region in the game map. Holiday activities and version updates also result in traffic surges. Therefore, the Genshin Impact team requires a logging platform that has the elastic scaling feature to implement rapid scaling during peak hours.

Solution

To meet the requirements of the Genshin Impact team on the performance, data collection, and elastic scalability of its business monitoring platform, Alibaba Cloud provides Log Service as the solution to help the team fix the issues that occur during data collection, query, and monitoring.

- High performance: Log Service can process billions of rows of data within seconds. Log Service supports ten s of petabytes of data throughput with an SLA of 99.95%. Log Service provides multiple features that can be used to extract, aggregate, and visualize heterogeneous data. Log Service also provides the alerting and AIOps-based anomaly detection features. The team can use these features to collect and analyze various types of data based on specific business scenarios. The data includes the monitoring logs of business and services, the operation and audit logs of cloud services, and the metrics of game operations.

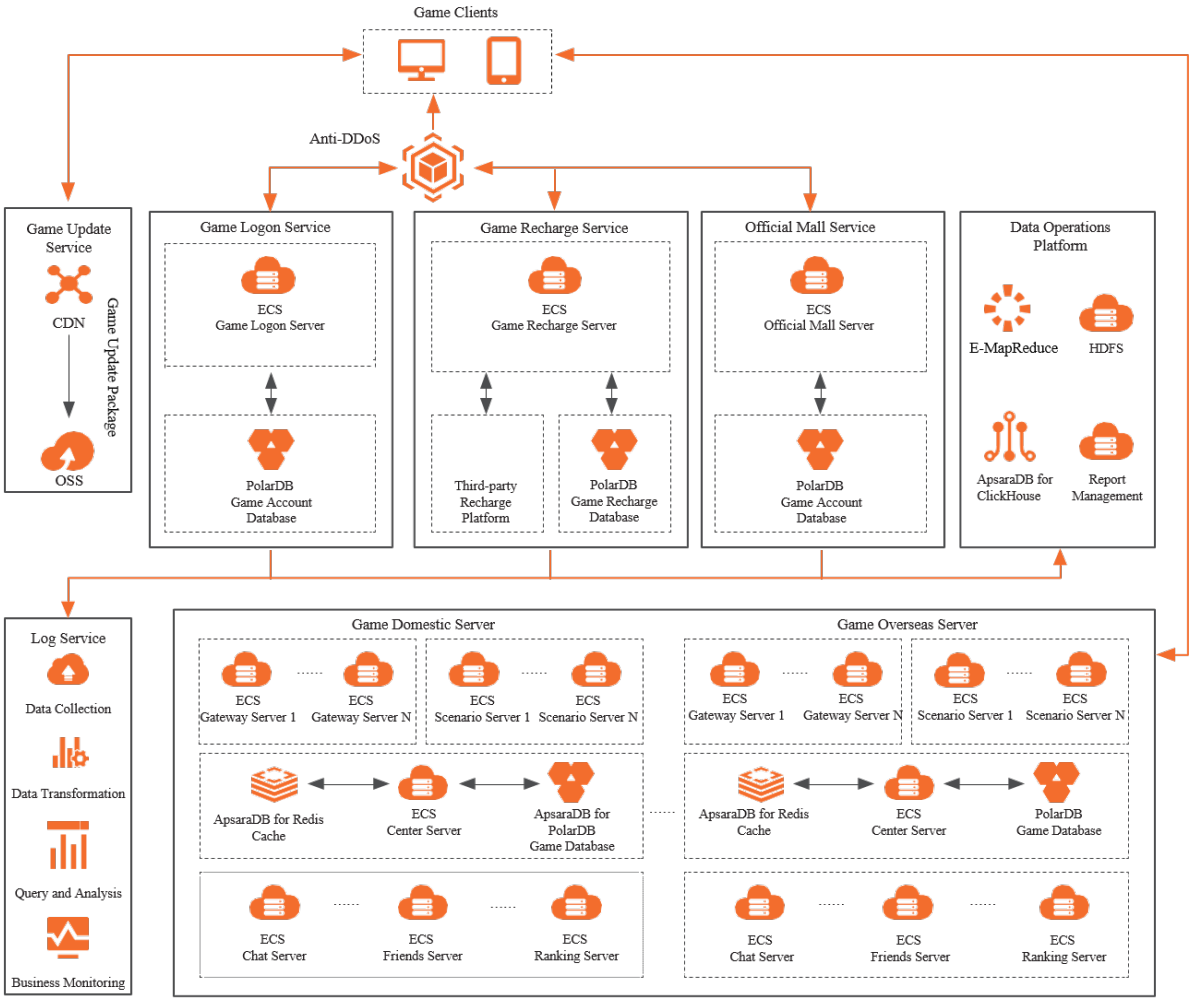

- Centralized data collection: Log Service supports more than 50 data sources. Data from clients, web pages, or protocols can be collected by using SDKs or APIs. Log Service provides the resumable upload feature to ensure the reliability of data collection. Log Service also supports multiple transmission methods in different scenarios by combining the internal network and a virtual private cloud (VPC), or by using the Internet based on global acceleration. Log Service helps the Genshin Impact team collect service logs from different regions with high efficiency and reliability and manage the logs in a unified manner.

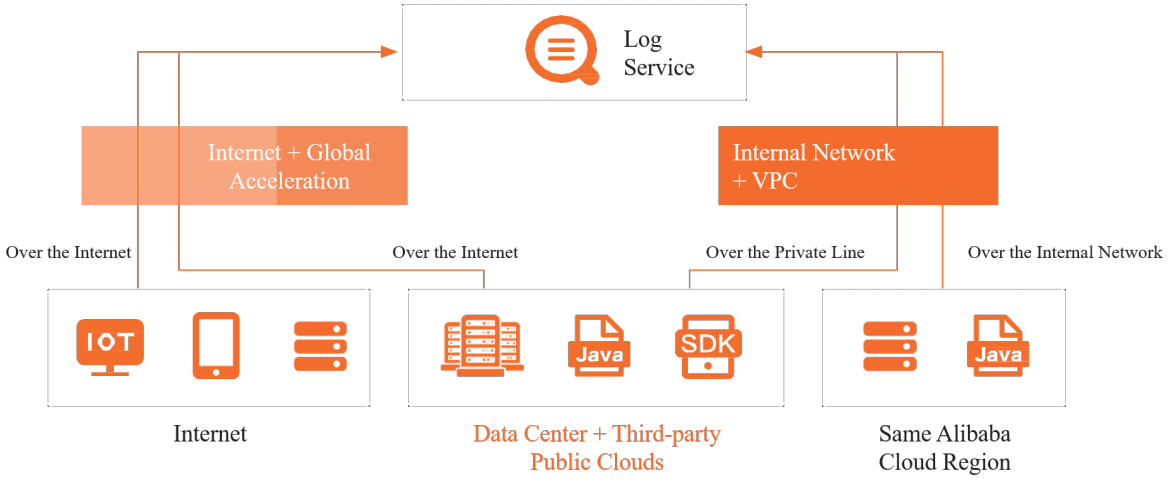

- Elastic scalability: Log Service supports linearly elastic scaling for petabytes of data per day based on the actual traffic. Log Service is adaptive to the traffic surges that are caused by promotional activities in Genshin Impact. The Genshin Impact team uses Log Service to overcome the instability caused by frequent traffic surges.

11.3. Sandbox Network

Log Service helps Sandbox Network collect and manage data in a centralized manner in global business scenarios in which services are deployed on multiple clouds. Sandbox Network has increased its troubleshooting efficiency by 30% to meet its rapid business growth.

Company profile

Sandbox Network is committed to building a global user-generated content (UGC) gaming platform that helps players convert their ideas into games. Blockman GO, a representative game of Sandbox Network, has been released online and well received by more than 50 million players around the world. For more information, visit the official website of Sandbox Network.

Scenarios

Sandbox Network is a well-known UGC gaming platform that has released multiple games worldwide and attracted more than 50 million players. As the number of self-developed and user-generated games increases, Sandbox Network expands their business to keep pace with the demands. The business development requires more fine-grained O&M and operations on the business platform of Sandbox Network. Upstream data and downstream data both need to be integrated so that multidimensional monitoring can be implemented. This helps improve platform stability and user experience. In addition, the analysis of user behavior helps iteratively improve services.

Challenges

Sandbox Network faces the following challenges:

- Centralized log collection in global business scenarios in which services are deployed on multiple clouds

  Sandbox Network has a global presence. The company has deployed its business clusters inside and outside mainland China, such as India and Hong Kong (China). In addition, Sandbox Network uses multi-cloud deployment solutions and a cross-border and cross-cloud architecture. Therefore, Sandbox Network requires a centralized method to collect data.

- Centralized management on heterogeneous data that is collected from various sources in different formats

  The business data of Sandbox Network includes the system data of each cloud platform, behavioral data of each game player, and logs of different clients such as iOS and Android clients. Heterogeneous data that is collected from various sources in different formats poses a big challenge on centralized data management.

- Rapid and elastic scalability that is required by the UGC gaming platform

  Traffic spikes occur when new games are released, promotional activities are launched for existing games, or hot-selling games come out. To ensure business stability, Sandbox Network requires a service that supports rapid and elastic scalability.

Solution

To help Sandbox Network address the preceding challenges, Alibaba Cloud provides Log Service to collect and manage data in a centralized manner.

- Centralized data collection: Log Service supports more than 50 data sources. For example, you can use Logtail to collect logs from Elastic Compute Service (ECS), data centers, and the servers of other cloud service providers. In multi-cloud deployment scenarios, Sandbox Network can use Logtail to collect data in a centralized manner and use Log Service as a centralized data analysis platform. Log Service also supports multiple transmission methods in different business scenarios. For example, you can use an internal network and a virtual private cloud (VPC) or the Internet and global acceleration to transmit data.
- Centralized management of heterogeneous data: Log Service is a fully managed and highly available service that provides data transformation and supports auto scaling. Log Service is suitable for various scenarios in which Sandbox Network needs to standardize, enrich, distribute, reindex, or aggregate data. Log Service helps Sandbox Network transform and standardize heterogeneous data for subsequent data analysis. To prevent the issues that may occur in multi-account and cross-cloud scenarios, Sandbox Network integrates the Log Service console into their business platform to manage data. This simplifies the authorization process. In addition, Log Service provides various dashboards to help different business teams of Sandbox Network visualize data as needed.
- Elastic scalability: Log Service supports auto scaling to process petabytes of data. Log Service helps Sandbox Network respond to traffic spikes and business growth.

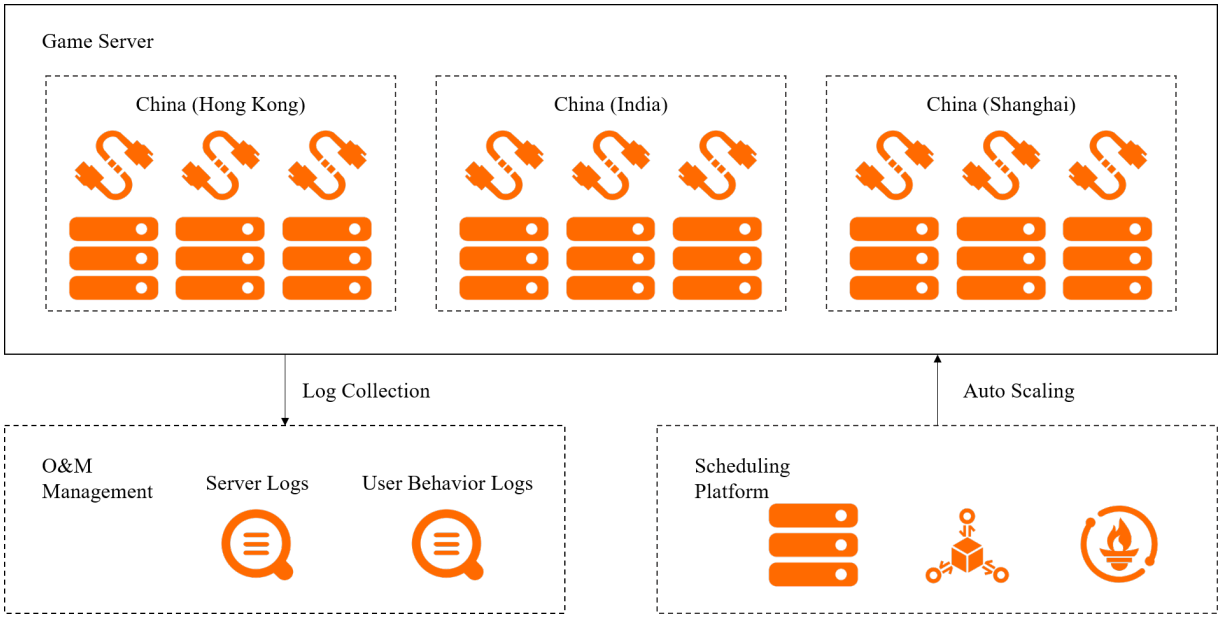

11.4. Milian Technology

Log Service helps Milian Technology aggregate disperse data and improve troubleshooting efficiency. Log Service also provides Milian Technology with multiple data analysis methods. Milian Technology has improved its capabilities in IT O&M, data operations, and risk management by using Log Service.

Company profile

Yidui is a brand of Beijing Milian Technology Co., Ltd. that was established in 2015. Milian Technology was awarded "National High Tech Enterprise" and "Zhongguancun New High-tech Enterprise" certifications. In 2020, Milian Technology completed Series B financing that is worth almost USD 100 million. Yidui is an online dating app that is committed to providing real-time Internet dating services. Yidui integrates real-time video calls, live streams, and matchmakers to provide single users with an innovative dating experience. Yidui has broken new ground for the video dating community and has become a highlight of the online matchmaking industry in 2019. For more information, visit the official website of Yidui.

Scenarios

Yidui is a real-time video dating platform that helps single users build relationships. Yidui provides various methods to help its users interact with each other by using text messages, voice calls, photos, and videos. Yidui can push the best matching live stream to its new users. The platform also provides the search feature to help users find love. In the early stage of its development, Yidui launched the search and recommender features based on Elasticsearch and MySQL. However, these features can no longer meet the increasing demands on data collection, processing, and analysis due to the growth of business and data volume and the upgrade of the service architecture. Therefore, Yidui requires a powerful platform to analyze large amounts of data and satisfy the requirements of increasing number of users.

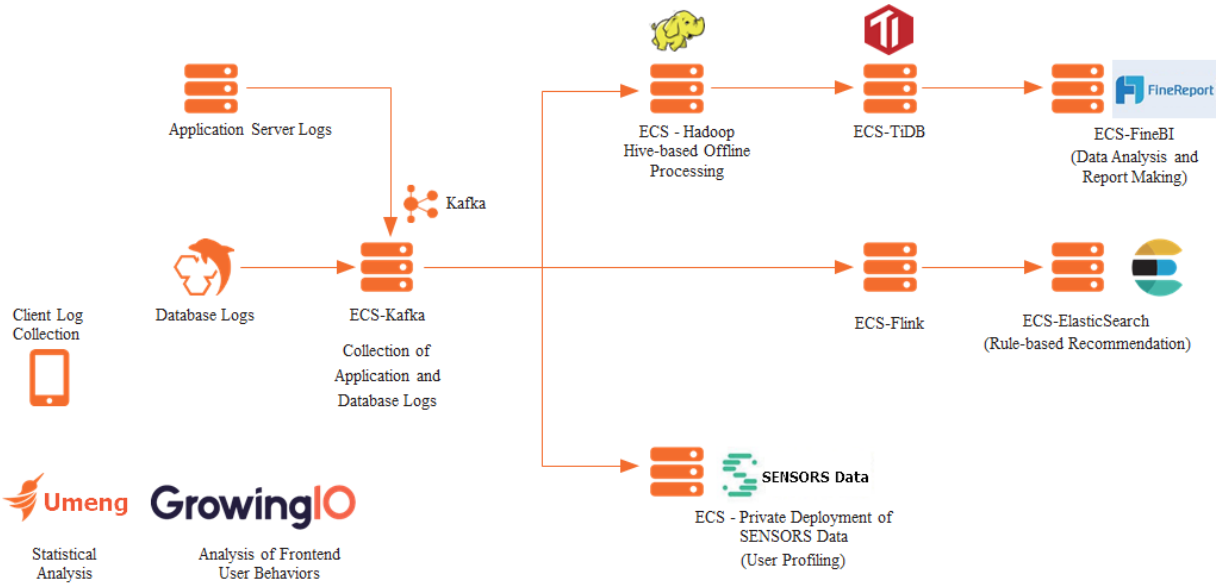

Challenges

Yidui faces the following challenges:

- Disperse data sources

  Yidui uses multiple compute and storage engines, such as databases, big data platforms, and third-party services. In this case, Yidui requires unified data planning and management to prevent data silos and improve the efficiency of development and management.

- Traffic surges

  The growth of business and the increasing number of users and livestream matchmaking activities increase system complexity and generate more logs. However, the data query feature of the self-managed Elasticsearch platform is unavailable during traffic peaks. In addition, system errors cannot be identified or fixed with high efficiency, resulting in poor user experience.

- Poor data analysis capabilities

  Yidui is dedicated to offering data-driven services, which cover statistical reports, algorithm-based recommendations, risk management, interactive operational queries, and behavior analysis. However, Yidui does not have powerful data analysis.

Solution

To enable Yidui to address the preceding challenges, Alibaba Cloud provides Log Service to help Yidui aggregate disperse data from multiple sources, query data during traffic peaks, and analyze large amounts of data.

- Unified log collection
  - Website logs can be collected and then sent to Log Service by using the web tracking feature.
  - App logs can be collected and then sent to Log Service by using Log Service SDK for iOS or Android.
  - Server logs can be collected and then sent to Log Service by using Logtail.
- Intelligent O&M platform

  Yidui has built an intelligent O&M platform based on Log Service. Yidui has built a platform to quickly analyze and fix exceptions to ensure the stability of its service system. Log Service provides unified log collection capabilities to help Yidui collect access logs from applications, systems, servers, databases, and network security services. Log Service helps Yidui query data within seconds. Log Service also provides the LogReduce feature and the AIOps-based anomaly detection feature, and supports multiple methods for alert notifications. The intelligent O&M platform is also used to improve user experience. Log Service provides the analysis feature for multidimensional metrics and enables Yidui to render analysis results into visualized charts.
- Associated services

  To provide a unified data analysis service, Log Service connects with other Alibaba Cloud services, such as MaxCompute, StreamCompute, and Machine Learning Platform for AI.
  - API Gateway: provides associated services with an egress gateway.
  - Quick BI: supports drag-and-drop operations and provides various visualized charts for interactive reports.
  - DataV: provides interactive data dashboards for data analysis.

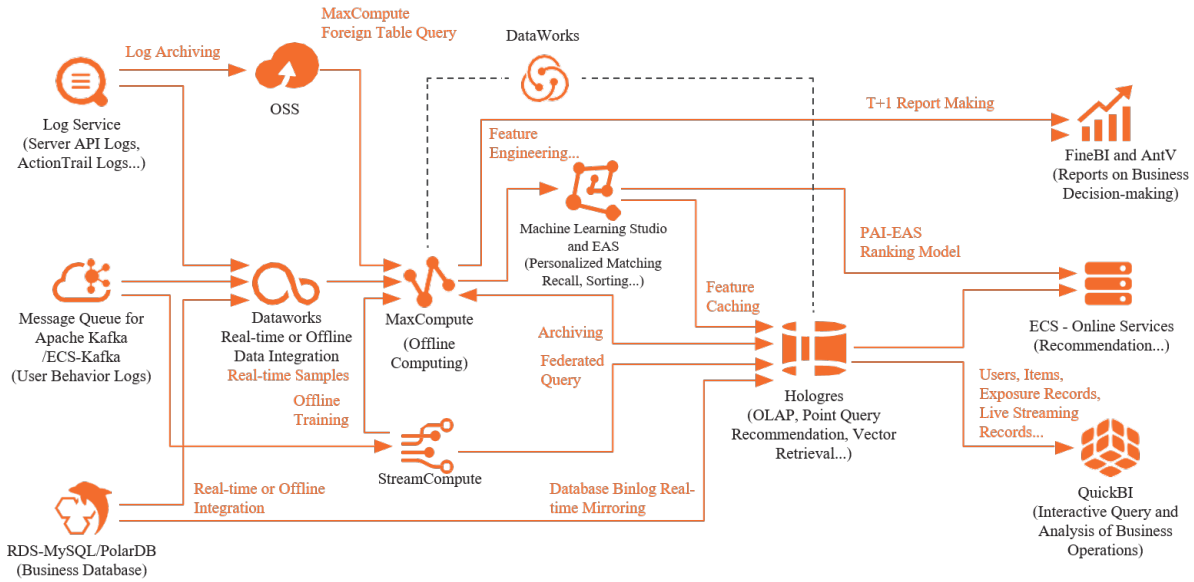

11.5. Hellobike

Hellobike saves 30% of costs by using Log Service instead of Kafka, Elasticsearch, and ClickHouse to collect and manage log data. Log Service meets the requirements of Hellobike on system stability, capacity scalability, and log query and analysis.

Company profile

Hellobike is a leading platform for local transportation and life services in China. Hellobike is committed to providing more convenient transportation methods and better inclusive life services based on digital dividends. For more information, visit the website of Hellobike.

Scenarios

Hellobike provides bicycle and moped rental services on a time-sharing basis. The company launched the bike-sharing service to solve the "last mile" issue. Hellobike is dedicated to building an efficient O&M system and a management system by applying the outcomes of technological innovation to intelligent terminals. To do this, Hellobike installs smart locks on its bicycles, equips the bicycles with positioning chips, establishes an intelligent planning and scheduling system at the backend, and promotes intensive operations based on intelligent O&M ports. Hellobike has expanded its business to more than 400 cities, with more than 400 million registered users whose accumulated riding distance has reached 23.7 billion kilometers. The company has implemented online real-time scheduling based on smart locks and seamlessly integrated bicycle data and app data. Therefore, Hellobike needs to collect, analyze, and store the data in real time.

Challenges

In the previous architecture of Hellobike, data is collected to Kafka, and then written to the Elasticsearch, Logstash, and Kibana (ELK) stack for log queries and written to ClickHouse for log analysis. Terabytes of incremental data is written to Elasticsearch per day, posing a big challenge on the stability of Elasticsearch. The write operations of Elasticsearch are delayed during data queries. The data query feature is unavailable in most cases because the system needs to write large amounts of data first. To address this bottleneck, the system creates only one replica for the data that is created on the current day, and creates more replicas on the next day. However, this solution greatly compromises data reliability. In addition, self-managed Kafka, Elasticsearch, and ClickHouse require high costs.

Solution

Log Service helps Hellobike collect and query terabytes of logs in real time and provides scalable storage.

- Real-time data collection

  Log Service provide native support for the Kafka protocol. Each client of Hellobike can seamlessly migrate data to Log Service by only setting the Kafka address to the Kafka protocol address of Log Service.

- Elastic scaling

  Log Service uses shards to read and write data. If traffic increases, you can manually split shards to expand the write throughput. You can also enable the automatic sharding feature. If the traffic reaches the upper limit, Log Service automatically increases the number of shards.

- Data query and analysis

  Log Service provides data query and analysis capabilities. You can query specific data by keyword or numeric range. You can also perform a recursive query for JSON fields or perform a union query. Log Service allows you to analyze logs by using the SQL-92 syntax. You can analyze tens of billions of log entries within seconds. The SQL-92 syntax supports more than 200 functions. You can use the SQL JOIN syntax to perform an associated analysis on the data in Log Service and the data from Object Storage Service (OSS) or MySQL databases.

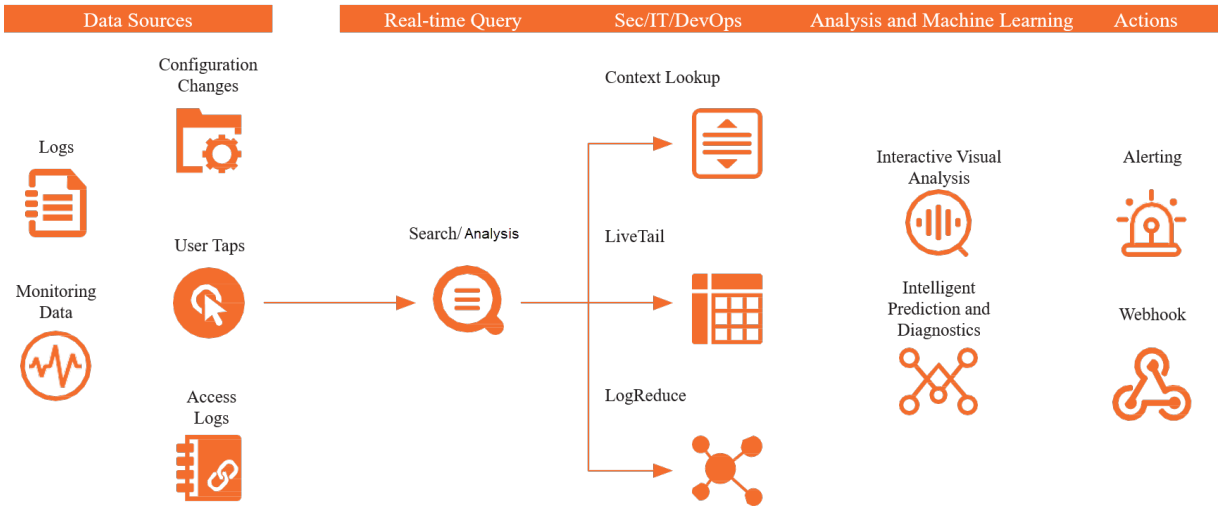

12.Competitive analysis

12.1. Comparison between Log Service and the ELK stack in log query and analysis scenarios

This topic describes the main features and benefits of Log Service. In this case, it compares Log Service with the Elasticsearch, Logstash, and Kibana (ELK) stack.

Background information

The ELK stack is a solution that is widely used to analyze logs in real time. For information about the case studies, visit the open source ELK community.

Log Service is a solution dedicated to log query and analysis scenarios. The service is developed based on the monitoring and diagnosis tool that is used to develop the Apsara distributed operating system. To meet the requirement of growth in both users and business, Log Service is improved to support log analysis in Ops scenarios, such as DevOps, Market Ops, and SecOps. The service has surpassed the challenges in scenarios such as Double 11, Ant Financial Double 12, Spring Festival red envelopes, and international business.

Overview

Apache Lucene is an open source search engine software library that is supported by the Apache Software Foundation. Apache Lucene provides full-text searching, full-text indexing, and text analysis capabilities. Elastic developed Elasticsearch in 2012 based on the Lucene library. It also launched the ELK stack in 2015 as an integrated solution to collect, store, and query logs. Lucene is designed to retrieve information based on documents. Its log processing capabilities are limited in many aspects, such as the data volume, query capability, intelligent grouping (LogReduce), and other custom features.

Log Service uses a self-developed log storage engine. In the past three years, Log Service has been applied to tens of thousands of applications. Log Service supports indexing for petabytes of data per day and serves tens of thousands of developers. It can be used to query and analyze data hundreds of millions of times per day. Log Service serves as the log analysis engine for multiple Alibaba Cloud services, such as SQL audit, EagleEye, Cloud Map, FLIGGY Trace, and Ditecting.

Log query is the most basic requirement of DevOps. In the industry research report 50 Most Frequently Used UNIX/Linux Commands, the tar and grep commands are the top two commands that are used by programmers to query logs.

The following list compares the ELK stack with Log Service in log query and analysis scenarios from five aspects:

- Ease of use: the convenience to get started and use the service.
- Features: the query and analysis features.
- Performance: the query capabilities, analysis capabilities, and latency.
- Capacity: the data processing capacity and the scalability.
- Cost: the cost of using features.

Ease of use

Log analysis is based on the following process:

- Collection: uses stable method to write data.
- Configuration: configures data sources.
- Capacity expansion: expands storage space and scales out servers.
- Usage: provides query and analysis features. These features are described in the "Features" section of this topic.
- Export: exports data to other systems for further processing, such as for stream computing and for data backup in Object Storage Service (OSS).
- Multi-tenancy: shares data with other users and ensures data security.

The following table includes a list that is used to compare the ease of use between Log Service and the ELK stack.

Item	Sub-item	Self-managed ELK stack	Log Service
Data collection	API	Restful API	• Restful API • JDBC
	Client	Various clients in the ecosystem, including Logstash, Beats, and Fluentd	• Logtail • Other clients such as Logstash
Configuration	Resource object	Provides indexes to classify logs.	• Project • Logstore Provides projects in which Logstores can be created to store logs.
	Method	API + Kibana	• API + SDK • Console
Capacity expansion	Storage	• Requires more servers. • Requires more disks.	Requires no added servers or disks.
	Computing	Requires more servers.	Requires no added servers.
	Configurations	• Configures servers by using the configuration management system. • The beta release of Logstash supports centralized configuration.	Provides the console and API for configurations, without the need of a configuration management system.

Item	Sub-item	Self-managed ELK stack	Log Service
	Collection point	Applies configurations and installs Logstash on server groups by using the configuration management system.	Provides the console and API for configurations, without the need of a configuration management system.
	Capacity	Does not support flexible capacity expansion.	Supports flexible and elastic capacity expansion.
Export	Method	• API • SDK	• API • SDK • Kafka-like consumer API • Consumer API for stream computing engines, such as Spark, Storm, and Flink • Consumer API for stream computing class libraries, such as Python and Java class libraries
Multi-tenancy	Security	Commercial versions	• https • Encrypted with signatures • Multi-tenant data isolation • Access control
	Traffic shaping	No traffic shaping	• Traffic shaping by project • Traffic shaping by shard
	Multi-tenant	Supported by Kibana	Supported by the provision of accounts and granted by the related permissions

Comparison results:

- The ELK stack ecosystem provides multiple tools such as write, installation, and configuration tools.
- Log Service is a managed service that is easy to access, configure, and implement. You can integrate Log Service with your services and use Log Service within 5 minutes.

- Log Service is a service that uses the software as a service (SaaS) model. This can be used to address challenges in terms of capacity or concurrency. Log Service supports elastic scaling and does not require O&M.

Query and analysis features

The query feature allows you to query log entries that meet specified search conditions. The analysis feature allows you to calculate and analyze data.

For example, you want to calculate the number and traffic of read requests whose status code is greater than 200 based on IP addresses. You can use one of the following methods to analyze log data:

- Query log entries and analyze the query results.
- Analyze all log entries in a Logstore.

```
1. Status in (200,500] and Method:Get* | select count(1) as c, sum(inflow) as sum_inflow, i
p group by Ip
2. * | select count(1) as c, sum(inflow) as sum_inflow, ip group by Ip
```

- **Basic query capabilities**

  The following table includes a list that is used to compare the ELK stack with Log Service based on Elasticsearch 6.5 Indices.

Data type	Feature	ELK	Log Service
Text	Query by index	Supported	Supported
	Delimiter	Supported	Supported
	Chinese delimiter	Supported	Supported
	Prefix	Supported	Supported
	Suffix	Supported	-
	Fuzzy query	Supported	Supported by using SQL statements
	Wildcard	Supported	Supported by using SQL statements
Numeric	long	Supported	Supported
	double	Supported	Supported
Nested	Json	Supported	-
Geo	Geo	Supported	Supported by using SQL statements
IP	Query by IP address	Supported	Supported by using SQL statements

Comparison results:

- The ELK stack supports more data types and provides stronger native query capabilities than Log Service.
- Log Service allows you to use SQL statements instead of fuzzy match or Geo functions to query data. However, the query performance is slightly compromised. The following examples provide further details about how to use SQL statements to query data:

```
To query data that matches the specified substring, execute the following query statement
:
* | select content where content like '%substring%' limit 100
To query data that matches the specified regular expression, execute the following query
statement:
* | select content where regexp_like(content, '\d+m') limit 100
To query parsed JSON-formatted data that matches the specified conditions, execute the fo
llowing query statement:
* | select content where json_extract(content, '$.store.book')='mybook' limit 100
To create an index for JSON-formatted data, execute the following query statement:
field.store.book='mybook'
```

- **Extended query capabilities**

  In log query scenarios, you may need to perform the following operations based on the query results:

  - If you find an error log entry, you can check the context to identify the parameters that cause the error.
  - If you find an error, you can run the tail -f command to display raw log entries and run the grep command to search for similar errors.
  - If you obtain millions of log entries from a query by using keywords, you can filter out 90% of log entries related to known issues and focus on unknown issues.

  To resolve the preceding issues, Log Service provides the following closed-loop solutions:

  - Contextual query: queries the context of a log entry in the raw log file and displays the results on multiple pages. You do not need to log on to the server to query the context.
  - LiveTail: uses the tail-f command to display raw log entries in real time.
  - LogReduce: dynamically groups logs based on different patterns to detect exceptions.

- LiveTail

  In the traditional O&M model, you must run the `tail -f` command on the server to monitor logs in real time. If you want a more specific result, run the `grep` or `grep -v` command to filter data by keyword. LiveTail provided in the Log Service console allows you to monitor and analyze online log data in real time. This reduces your O&M workloads.

  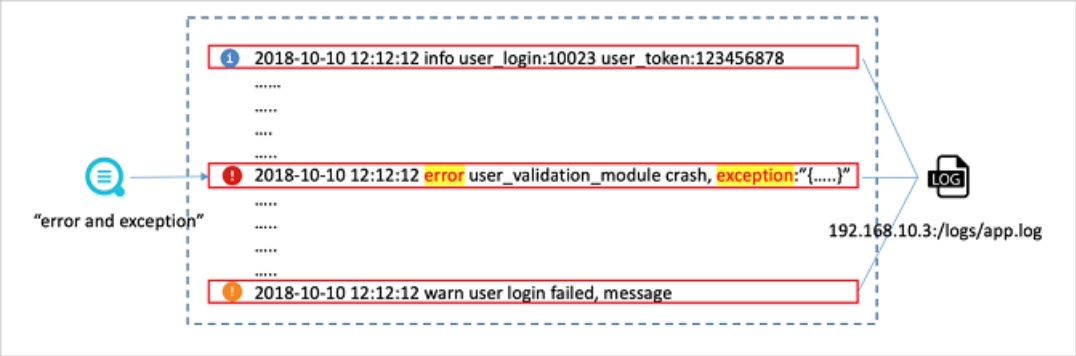

  LiveTail supports the following features:

  - Supports data collected from Docker and Kubernetes containers, servers, Log4j Appenders, and other data sources.
  - Monitors log data in real time, and allows you to filter data by keyword.
  - Delimits log fields to facilitate the queries for log entries that contain specific delimiters.

- LogReduce

  Large volumes of log data generated every day by rapid business development has created the following issues:

  - Potential system exceptions are difficult to identify.
  - Suspicious logons by intruders cannot be detected in real time.
  - System behavior changes caused by version updates cannot be detected due to large amounts of information. In addition, recorded logs have various formats and have no topics. Therefore, the logs are difficult to group. **LogReduce** provided in Log Service can group logs based on different patterns and provide an overview of the logs. LogReduce supports the following features:
    - Various formats of logs such as Log4j logs, JSON-formatted logs, and syslog logs are supported.
    - Before logs are grouped, these logs can be filtered based on specified conditions. Raw log entries can be retrieved based on the signature of log entries grouped in a specific pattern.
    - Compares the patterns of different time periods.
    - The precision of log grouping can be adjusted based on your business requirements.
    - Hundreds of millions of log entries can be grouped in seconds.

Analysis capabilities

Elasticsearch supports data aggregation based on the doc values. Elasticsearch 6.x supports data grouping and aggregation by using the SQL syntax. Log Service supports the RESTful API and JDBC API and is compatible with the SQL-92 standard. Log Service supports complete SQL statements, including basic aggregate functions. In addition, Log Service allows you to perform JOIN operations on internal and external data sources, and implement machine learning and pattern analysis.

> **Note** The following examples compare the analysis capabilities of the ELK and Log Service. For more information, see ES 6.5 Aggregation and Real-time log analysis.

In addition to SQL-92 standard functions, Log Service also provides the following features that are specific to log analysis scenarios:

- Interval-valued comparison and periodicity-valued comparison functions

  You can nest the interval-valued comparison and periodicity-valued comparison functions in SQL statements. You can use this method to calculate the changes of a single field value, multiple field values, and curves in different time windows.

  ```
  * | select compare( pv , 86400) from (select count(1) as pv from log)
  ```

  ```
  *|select t, diff[1] as current, diff[2] as yestoday, diff[3] as percentage from(select t, compare( pv , 86400) as diff from (select count(1) as pv, date_format(from_unixtime(__time__), '%H:%i') as t from log group by t) group by t order by t) s
  ```

- Join internal and external data sources for data query

  You can join the data in Log Service with external data sources for data query and analysis. The following list shows the supported data sources and JOIN operations:

  - You can perform JOIN operations on the data in Logstores, MySQL databases, and OSS buckets (CSV files).
  - You can perform LEFT OUTER JOIN, RIGHT OUTER JOIN, FULL OUTER JOIN, and INNER JOIN operations on data.
  - You can use SQL statements to query data in external tables and join data in Log Service with external tables.

  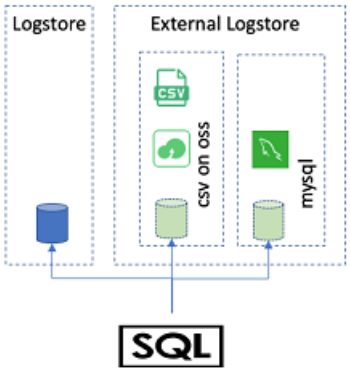

  The following example provides further details about how to join the data in Log Service with external tables:

```
sql
To create an external table, execute the following query statement:
* | create table user_meta ( userid bigint, nick varchar, gender varchar, province varcha
r, gender varchar,age bigint) with ( endpoint='oss-cn-hangzhou.aliyuncs.com',accessid='LT
A288',accesskey ='EjsowA',bucket='testossconnector',objects=ARRAY['user.csv'],type='oss')
To join the data in Log Service with the external table, execute the following query stat
ement:
* | select u.gender, count(1) from chiji_accesslog l join user_meta1 u on l.userid = u.us
erid group by u.gender
```

- Geolocation functions

  You can use the built-in geolocation functions to identify users based on IP addresses and mobile phone numbers. The following list shows the available geolocation functions:

  - IP functions: identify the country, province, city, longitude and latitude, and Internet service provider (ISP) of an IP address.
  - Phone number functions: identify the ISP, province, and city where a mobile phone number is registered.
  - Geohash functions: encode the longitude and latitude of a location.

  The following example provides further details about how to use geolocation functions in query statements:

  ```
  sql
  * | SELECT count(1) as pv, ip_to_province(ip) as province WHERE ip_to_domain(ip) != 'intr
  anet' GROUP BY province ORDER BY pv desc limit 10
  * | SELECT mobile_city(try_cast("mobile" as bigint)) as "city", mobile_province(try_cast(
  "mobile" as bigint)) as "province", count(1) as "number of requests" group by "province",
  "city" order by "number of requests" desc limit 100
  ```

  - GeoHash
  - IP functions

- Security check functions

  Log Service provides security check functions based on the globally shared asset library of WhiteHat Security. You can use security check functions to check whether an IP address, domain name, or the URL of a log is secure.

  - security_check_ip
  - security_check_domain
  - security_check_url

- Machine learning and time series functions

  Log Service provides machine learning and intelligent diagnostic functions. This can be used to perform the following operations:

  - Automatically learns historical data regularities and predicts the future trends.
  - Detects imperceptible exceptions in real time, and combines analysis functions to analyze error causes.
  - Intelligently detects exceptions and inspects the system based on the periodicity-valued comparison function and alerting feature. This feature provides an efficient method to analyze data in scenarios, such as intelligent O&M, security, and operations.

Machine learning and intelligent diagnostic functions provide the following features:

- Prediction: fits a baseline based on the historical data.
- Exception detection, change point detection, and inflection point detection: detect exceptions.
- Multi-period detection: detects the periodicity of time series data.
- Time series clustering: finds time series curves that have different curve shapes.

- Pattern analysis functions

  Pattern analysis functions provide a fast and efficient method to detect data patterns and identify issues. You can use pattern analysis functions to perform the following operations:

  - Identifies patterns that frequently occur. For example, you can use pattern analysis functions to identify the ID of the user who has sent 90% of invalid requests.
  - Identifies the factor that most affects two patterns.
    - In requests whose latency is greater than 10 seconds, the ratio of combined dimensions that contain an ID is significantly higher than the ratio of other combined dimensions.
    - The ratio of the ID in Pattern B is lower than the ratio in Pattern A.
    - Patterns A and B are significantly different from each other.

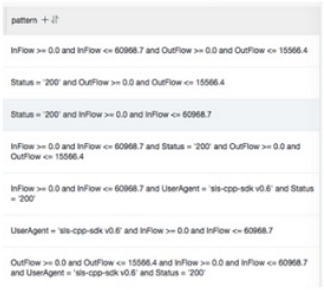

	In A	Ratio	In B	Ratio
ID=1002	95	68%	100	3%
ID=1002 Method=Get	80	74%	100	3%
ID=1002 Method=Get IP=120343	40	37%	250	7.5%
ID=1003	20	19%	500	15%

Performance

The following examples compare the data write, query, and aggregation performance of Log Service and the ELK stack by using the same dataset.

- Test environment
  - Test configurations

Item	Self-managed ELK stack	Log Service
Runtime environment	Four Elastic Compute Service (ECS) instances, each configured with four CPU cores, 16 GB of memory, and ultra disks or standard SSDs	-
Shard	10	10
Copies	2	3 (the default value that is invisible to users)

- Test data
  - Five columns of the double data type, five columns of the long data type, and five columns of the text data type are tested. The test data is displayed in 256, 512, 768, 1,024, and 1,280 dictionaries.
  - Fields are randomly sampled from the test data.
  - Raw data size: 50 GB.
  - Number of raw log entries: 162,640,232 (about 160 million).

  The following example provides further details of the sample test log entry:

```
timestamp:August 27th 2017, 21:50:19.000
long_1:756,444 double_1:0 text_1:value_136
long_2:-3,839,872,295 double_2:-11.13 text_2:value_475
long_3:-73,775,372,011,896 double_3:-70,220.163 text_3:value_3
long_4:173,468,492,344,196 double_4:35,123.978 text_4:value_124
long_5:389,467,512,234,496 double_5:-20,10.312 text_5:value_1125
```

- Test data write operations

  The Bulk API is called to write batch data to the ELK stack and the PostLogstoreLogs API is called to write batch data to Log Service. The following table includes a list that is used to compare the test results.

Item	Sub-item	Self-managed ELK stack	Log Service
Latency	Average write latency	40 ms	14 ms
Storage	Data volume of a copy	86G	58G
	Expansion rate: Data volume/Raw data size	172%	116%

  > **Note** The storage fee incurred in Log Service for 50 GB of data includes the fee that is incurred for writing 23 GB of compressed data. It also includes the fee that is incurred for writing 27 GB of indexes.

  Comparison results:

  - Log Service has a lower data write latency than the ELK stack.
  - The size of the raw data is 50 GB. The size of the stored data exceeds 50 GB because the test data is random. In most scenarios, the size of the stored data after compression is smaller than the size of the raw data. The data stored in the self-managed ELK stack expands to 86 GB. In this case, the expansion rate is 172%, which is 58% higher than that in Log Service. This expansion rate is close to the recommended value 220% when you write new data to the ELK stack.

- Test data read operations (query and analysis)

- Test scenarios

  Log query and aggregation are used as example scenarios. The average latency is calculated in the two scenarios when the number of concurrent read requests is 1, 5, and 10. The following examples provide further details of the two scenarios.

  - Use aggregate functions (AVG, MIN, MAX, SUM, and COUNT) on the five columns of the long data type and group the values of five numeric columns. Then, sort the calculated values by the COUNT value to obtain the first 1,000 rows of data. You can execute the following query statement:

    ```
    select count(long_1) as pv,sum(long_2),min(long_3),max(long_4),sum(long_5)
      group by text_1 order by pv desc limit 1000
    ```

  - Use a keyword, for example, value_126, to obtain the number of log entries that contain the keyword, and display the first 100 rows of data. You can execute the following query statement:

    ```
    value_126
    ```

- Test results

Scenario	Number of concurrent read requests	Latency of the self-managed ELK stack (Unit: seconds)	Latency of Log Service (Unit: seconds)
Log analysis	1	3.76	3.4
	5	3.9	4.7
	10	6.6	7.2
Log query	1	0.097	0.086
	5	0.171	0.083
	10	0.2	0.082

- Comparison results:
  - Both Log Service and the ELK stack can query and analyze 160 million log entries within seconds.
  - In the log analysis scenario, the latency is similar between Log Service and the ELK stack. The ELK stack uses SSDs and delivers better I/O performance than Log Service when a large amount of data is read.
  - In the log query scenario, Log Service has a much shorter latency than the ELK stack. When the number of concurrent requests increases, the latency of the ELK stack increases. However, the latency of Log Service remains stable and slightly decreases.

Capacity

- Log Service allows you to index petabytes of data per day and query dozens of terabytes of data within seconds at a time. Log Service supports elastic scaling and scale-out for the processing scale.
- The ELK stack is applicable to scenarios where data is written in units of GB to TB per day and stored in units of TB. The processing capacity is limited by the following factors:

- The size of a cluster: A cluster that consists of about 20 nodes has optimal performance. A large cluster in the industry can contain 100 nodes and is often split into multiple clusters for data processing.
- Write capacity: The number of shards cannot be modified after the shards are created. Therefore, the maximum number of available nodes cannot be increased if more write capacity is required due to the increasing throughput.
- Storage capacity: If the data size in the primary shard reaches the maximum disk capacity, you must migrate the shard to a larger disk, or allocate more shards to the current disk. In this case, you can create an index, specify more shards, and rebuild existing data.

Use case

Customer A is one of the major news websites in China. This customer has thousands of servers and hundreds of developers. The O&M team used an ELK cluster to process NGINX logs. However, the cluster frequently fails to process large amounts of data.

- When a user queries a large amount of data, the cluster is unavailable to other users.
- The cluster is fully occupied during peak hours to collect and process data. This compromises data integrity and the accuracy of query results.
- When the business grows, out of memory (OOM) errors frequently occur due to invalid memory settings and heartbeat synchronization failures. As a result, the query results are unavailable and inaccurate, and the cluster is inoperable for developers.

In June 2018, the O&M team started to use Log Service to fix the preceding issues.

1. The team used Logtail to collect online logs, and called the API to integrate log collection and server management configurations into the O&M system.
2. The team embedded the query page of Log Service into the unified logon and O&M platform to isolate business permissions from account permissions.
3. The team embedded the console page of Log Service into the platform of customer A. This way, the development team can use an efficient method to query logs. The team also configured Grafana plug-ins to monitor business and configured DataV to create dashboards in Log Service.

> **Note** For more information, see the following topics:
> - Embed console pages and share log data
> - Connect Log Service to Grafana
> - Connect Log Service with DataV
> - Connect Log Service with Jaeger

After two months, the team improved O&M in the following aspects:

- The number of queries per day was significantly increased. Developers used the O&M platform to query and analyze logs. This improved the efficiency of development. The O&M team also revoked the permissions of online logon.
- In addition to NGINX logs, the O&M platform also imported application logs, mobile device logs, and container logs into Log Service. The amount of processed data increased by 9 times.
- More applications were developed. For example, Jaeger plug-ins were integrated with Log Service to build a log tracing system. Alerts and charts were configured to detect online errors in real time.
- Various platforms are connected with the unified O&M platform that provides a centralized method to collect data. This prevents repeated data collections. In addition, the Spark and Flink

platforms of the big data team can consume log data in real time.

Conclusion

Elasticsearch applies to scenarios where you want to update, query, and delete data. Therefore, Elasticsearch is widely used in fields such as data query, data analysis, and application development. The ELK stack applies to log analysis scenarios because it can maximize the flexibility and performance of Elasticsearch. Log Service is designed for log data query and analysis scenarios by providing multiple unique features. The ELK stack covers a wider range of scenarios. However, Log Service provides deeper analysis features for specific scenarios. You can select one of the two services based on your business requirements.

12.2. Comparison of monitoring and analysis platforms

This topic compares multiple monitoring and analysis platforms and provides O&M and site reliability engineering (SRE) teams with multiple solutions to build monitoring and analysis platforms based on their business requirements.

Background information

O&M and SRE teams play key roles and undertake multiple complex tasks, such as application deployment, performance and availability monitoring, alert monitoring, shift scheduling, capacity planning, and business support. The rapid development of cloud native services, containerized services, and microservices has accelerated the iteration of related features and posed the following challenges to O&M and SRE teams:

- Widespread business lines
  - Business lines are widely spread to clients, frontend web applications, and backend applications.
  - Up to dozens of business lines need to be supported at the same time.
- Manpower shortage

  The O&M and SRE teams in some companies consist of only 1% of employees or less as compared with their R&D teams.
- High pressure on online service stability
  - O&M and SRE teams must frequently respond to emergencies and fix urgent issues.
  - Complex business processes and the tremendous number of system components have put great pressure on service troubleshooting and recovery.
- Disperse and inefficient monitoring and analysis platforms
  - Various types of data are monitored in different dimensions. This results in excessive scripts, monitoring tools, and data silos.
  - Various types of data are stored in different monitoring systems that do not support associated analysis. As a result, root causes cannot be located in an efficient manner.
  - Each system may have thousands of alert rules. This increases management costs and results in excessive alerts. In addition, the specified conditions in an alert rule may be incorrectly or incompletely evaluated.

To address the preceding challenges, O&M and SRE teams require an easy-to-use and powerful monitoring and analysis platform. This platform can be used to analyze data with high efficiency, improve the work efficiency, locate root causes in a timely and accurate manner, and ensure service continuity.

Data-related issues to be fixed by a monitoring and analysis system

To ensure service stability and support business development, O&M and SRE teams need to collect and analyze large amounts of data. The data includes hardware and network metrics, business data, and user behavior data. After data is collected, O&M and SRE teams also need a suitable system to convert, store, process, and analyze the data based on business requirements. O&M and SRE teams face the following data-related issues:

- Diversified data types
  - System data: hardware metrics related to CPU, memory, network, and disk, and system logs
  - Golden signals: latency, traffic, errors, and saturation
  - Access logs
  - Application logs: Java application logs and error logs
  - User behavior data: website clicks
  - Tracking points in apps: statistics of tracking points in Android and iOS apps
  - Framework data: Kubernetes framework data
  - Call chains: trace data
- Diversified requirements

  O&M and SRE teams use the preceding types of data to ensure service stability. The teams also need to support other business teams based on the following data management requirements:
  - Monitoring and alerting: Small amounts of stream data can be processed within seconds or minutes.
  - Customer service and troubleshooting: Data can be queried by keyword within seconds.
  - Risk management: Traffic can be processed within seconds.
  - Operations and data analysis: Large-scale data can be analyzed within seconds or hours. For example, data can be analyzed in online analytical processing (OLAP) scenarios.
- Difficult resource estimation

  The data amount of fast-growing business is difficult to estimate at the early stage due to the following reasons:
  - No reference can be used to estimate the data volume of newly connected data sources.
  - The rapid business growth results in data surges.
  - The previous data storage methods and data retention periods no longer meet the changing data management requirements.

Solutions to build a monitoring and analysis platform

O&M and SRE teams need a monitoring and analysis platform to process data that is collected from various sources in different formats. To meet diversified business requirements, the O&M and SRE teams may need to use and maintain multiple systems based on the following open source services:

- Telegraf+Influxdb+Grafana

Telegraf is a lightweight and plugin-driven server agent that can be used to collect metric data from operating systems, databases, and middleware. The collected time series data can be written to and read from InfluxDB that is used to store and analyze the data. Then, the related analysis results can be rendered into visualized charts and interactively queried by using Grafana.

- Prometheus

  Prometheus is a basic tool that is used to process time series data in the cloud native ecosystem of Kubernetes. Prometheus can ingest metric data that is flexibly collected by using exporters. Prometheus can also connect with Grafana for data visualization.

- ELK

  The Elasticsearch, Logstash, and Kibana (ELK) stack is an open source component that is most commonly used to perform multidimensional log query and analysis. The ELK stack provides fast, flexible, and powerful query capabilities to meet most query requirements of R&D, O&M, and customer service teams.

- Tracing systems

  In a microservice and distributed system, call chains are complex. If no suitable tracing system is used, root causes of errors cannot be located in an efficient manner. Tracing systems such as Zipkin, Jeager, OpenTelemetry, and SkyWalking are applicable solutions. OpenTelemetry is the industry standard and SkyWalking is a Chinese tracing system. However, these tracing systems do not provide data storage components and must be used together with Elasticsearch or Cassandra to store trace data.

- Kafka+Flink

  To meet the requirements for data cleansing and risk management, a platform is required to support real-time stream processing. In most cases, the platform is built by combining Kafka and Flink.

- ClickHouse, Presto, and Druid

  In scenarios where operational analytics and data analysis reports are required, data is often imported to an OLAP engine for higher capabilities in real-time data analysis. This way, large amounts of data can be analyzed and various ad hoc queries can be performed within seconds to minutes.

Different system components are used to process different types of data. The data is distributed to these components or a data copy is stored in multiple systems, increasing system complexity and usage costs.

When the amount of data increases, O&M and SRE teams face big challenges in terms of component stability, system performance, cost control, and support for a large number of business lines.

Challenges for monitoring and analysis platforms

To maintain multiple systems and support various business lines, O&M and SRE teams must address challenges in the following aspects:

- System stability
  - System dependencies: Data is distributed to multiple systems that depend on each other. If a system has an issue, other associated systems are affected. For example, if the write speed of the downstream system Elasticsearch becomes slow, the storage usage of Kafka clusters that are used to cache data becomes high. In this case, the Kafka clusters may have no more space for written data.
  - Traffic bursts: Traffic bursts frequently occur in the Internet. Traffic bursts also occur when large amounts of data are written to a system. O&M and SRE teams must ensure that the system can run as expected and the read and write operations are not affected by traffic bursts.

  - Resource isolation: Data entries can be separately stored based on their respective priorities in different physical storage resources. On one hand, if data entries are isolated by using only this method, a large number of cluster resources are wasted and the O&M costs are greatly increased. On the other hand, if data is shared by cluster resources, the interference between these resources must be minimized. For example, a query of a large amount of data in a system may break down all clusters of the system.
  - Technical issues: O&M and SRE teams need to optimize a large number of parameters for multiple systems and spend long periods of time in comparing and adjusting optimization solutions to meet various requirements in different scenarios.

- Predictable performance
  - Data size: The size of data in a system has a significant impact on the system performance. For example, O&M and SRE teams need to predict whether tens of millions to hundreds of millions of rows of time series data can be read from and written to a system, and whether 1 billion to 10 billion rows of data can be queried in Elasticsearch.
  - Quality of service (QoS) control: Each system consists of limited resources, which must be allocated and managed based on the QPS and concurrent requests for different data types. In some cases, the hardware resources need to be degraded to prevent other performance metrics from declining. However, most open source components are developed without considering QoS control.

- Cost control
  - Resource costs: The deployment of each system component consumes hardware resources. If a data copy exists in multiple systems at the same time, large amounts of hardware resources are consumed. In addition, the amount of business data is difficult to estimate. In most cases, the amount of business data is overestimated, resulting in wastes of resources.
  - Data access costs: O&M and SRE teams require a tool or platform to support automatic access to business data. The tool or platform can help O&M and SRE teams adapt data formats, manage environments, maintain configurations, and estimate resources. This way, O&M and SRE teams can focus on more important business.
  - Support costs: Technical support is required to resolve issues that may occur when multiple systems are used. These issues include unsuitable usage modes and invalid parameter settings. Bugs in open source software may result in additional costs.
  - O&M costs: O&M and SRE teams need to spend long periods of time in fixing hardware and software issues of each system. They also need to replace hardware, scale up or down the system capacity, and upgrade software.
  - Cost sharing: To effectively utilize resources and make strategic budgets and plans, O&M and SRE teams must estimate resource consumption based on actual business lines as accurately as possible. This requires a monitoring and analysis platform to support separate billing based on effective data metering.

Example scenario

Background information

- In this example, a company has developed more than 100 applications, and each application generates NGINX access logs and Java application logs.
- The log size of each application is greatly changing by 1 GB to 1 TB per day. A total of 10 TB of logs are generated by all the applications per day. The data retention period of each application ranges from 7 to 90 days and the average data retention period is 15 days.

- Log data is used for business monitoring and alerting, online troubleshooting, and real-time risk management.

Business architecture

- Beats: collects data and sends the data to Kafka in real time.
- Kafka: temporarily stores data and sends data to Flink for real-time consumption and sends data to Elasticsearch.
- Flink: monitors data, analyzes data, and manages risks in real time.
- Elasticsearch: queries logs, analyzes logs, and troubleshoots errors.

The preceding architecture describes how data is processed in a monitoring and analysis platform. The following workloads must be taken into consideration when the O&M and SRE team of the company creates solutions based on Elasticsearch.

- Capacity planning: Disk capacity = Size of raw data × Expansion coefficient × (1 + Number of replicas) × (1 + Percentage of reserved space). The expansion coefficient ranges from 1.1 to 1.3. The number of replicas is 1. The percentage of reserved space for temporary files is 25%. In this case, the actual disk capacity is 2.75 to 3.5 times the size of raw data. If the _all parameter is enabled, more reserved space is required for data expansion.
- Cold and hot data isolation: If all data is stored in an SSD, the storage cost is high. Some indexed data can be stored in an HDD based on the priority and timestamp of the indexed data. The indexed data can also be migrated by using the rollover feature of Elasticsearch.
- Index setting: The NGINX access logs and Java application logs of each application are periodically indexed in chronological order. The number of shards is specified based on the actual scenario and the storage size of each shard can be 30 GB to 50 GB. However, the amount of log data of each application is difficult to accurately estimate. If issues occur when data is written and queried, the number of shards may need to be adjusted and the reindexing process consumes more resources.
- Kafka consumer setting: Logstash is used to consume data from Kafka and the data is then written to Elasticsearch. During this process, the number of partitions of Kafka topics must match the value of the logconsumer_threads parameter. Otherwise, data cannot be evenly consumed from each partition.
- Parameter optimization for Elasticsearch: The O&M and SRE team must first evaluate the performance of Elasticsearch in write throughput, latency visibility, data security, and data query. Then, the team can make full use of Elasticsearch by optimizing the related parameters based on the CPU and memory of clusters. Common parameters include the number of threads, memory control, translog settings, queue length, the intervals between various operations, and merge scheduling.
- Memory: The memory size of a JVM heap is less than 32 GB and the remaining space is used as the

cache memory of the operating system. Frequent garbage collection (GC) operations affect application performance and may result in the unavailability of a Java application.
    - The memory usage of a master node is related to the number of shards in a cluster. In most cases, the number of shards in the cluster is less than 100,000. In the default settings of Elasticsearch, the maximum number of shards in a master node is 1,000. The numbers of indexes and shards need to be specified based on the actual scenario.
    - The memory of a data node is based on the size of index data. If the finite-state transducer (FST) of Elasticsearch resides in the memory for a long period of time, the system recovers the cache memory. In this case, the data query performance is reduced even though the mmap method is provided in Elasticsearch version 7.3 and later. If a previous version is used, the data size of a node must be controlled.
- Query and analysis: The O&M and SRE team needs to spend long periods of time in continuously testing the data query and analysis performance of Elasticsearch by using the trial and error method.
    - The mappings must be configured based on the actual scenario. For example, nested mappings need to be minimized when the text and keyword fields are specified.
    - Queries of large amounts of data or complex query statements (for example, deeply nested GROUP BY clauses) must be avoided to prevent sharp consumption of system resources. To prevent the out of memory (OOM) issue that may occur when an aggregate query or fuzzy search is performed on a large data set, the size of the data set must be limited.
    - The number of segments needs to be controlled. The force merge API can be called as needed. In addition, the amount of disk I/O and resource consumption after a force merge needs to be evaluated.
    - A filter context and a query context must be selected based on the actual scenario. In scenarios where analytics are not required, the query cache feature of Elasticsearch can be used with higher efficiency in a filter context than a query context.
    - If scripts are used to query and analyze data, system instability or other performance issues may occur.
    - If the routing parameter is used to route search requests to a specific shard, the performance can be improved.
- Data corruption: If a crash occurs, files may be corrupted. If a segment or translog file is corrupted, the data in the related shard cannot be loaded, and damaged data must be cleaned up manually or by using tools. However, data loss occurs during this process.

The preceding issues may occur when the O&M and SRE team uses and maintains an Elasticsearch cluster. If the data size increases to hundreds of terabytes and a large number of data access requests occur, the stability of the cluster is difficult to maintain. The same issues also exist in other systems. Therefore, the O&M and SRE team needs to spend more time in fixing these issues.

Solution for an integrated cloud service

To meet the requirements of O&M and SRE teams for a monitoring and analysis platform, the Alibaba Cloud Log Service team provides a high-performance and cost-effective solution that is easy-to-use, stable, and reliable. This solution helps O&M and SRE teams fix the issues that may occur when the teams build the platform and improve work efficiency. In the early stage of its development, Log Service was the internal logging system that supports only Alibaba Group and Ant Group. After years of evolution, Log Service has become a cloud native observability and analysis platform that supports petabytes of data of multiple types, such as logs, metrics, and traces.

- Simple data import methods

  - Logtail: Logtail is proven easy-to-use, reliable, and powerful after being tested on millions of servers. Logtail can be managed in the Log Service console.
  - SDKs or Producers: Data from mobile terminals can be collected by using SDK for Java, C, Go, iOS, or Android or by using the web tracking feature.
  - Cloud native data collection: Custom Resource Definitions (CRDs) can be created to collect data from Container Service for Kubernetes (ACK).
- Real-time data consumption and ecosystem connection
  - Log Service supports elastic scalability that can be implemented within seconds. Petabytes of data can be written to and consumed from Log Service in real time.
  - Log Service is natively compatible with Flink, Storm, and Spark Streaming.
- Large-scale data query and analysis
  - Tens of billions of rows of data can be queried within seconds.
  - Log Service supports the SQL-92 syntax, interactive queries, machine learning algorithms, and security check functions.
- Data transformation
  - Compared with the traditional extract, transform, and load (ETL) process, the data transformation feature of Log Service can reduce up to 90% of development costs.
  - Log Service is a fully-managed service that provides high availability and elastic scalability.
- Metric data

  Cloud native metric data can be imported to Log Service. A hundred million rows of metric data from Prometheus can be collected and stored.
- Unified trace data import methods

  The OpenTelemetry protocol can be used to collect traces. The OpenTracing protocol can be used to import traces from Jaeger or Zipkin to Log Service. Traces from OpenCensus and SkyWalking can also be imported to Log Service.
- Data monitoring and alerting

  Log Service provides an integrated alerting feature that can be used to monitor data, trigger alerts, denoise alerts, manage alert incidents, and dispatch alert notifications.
- AIOps-based anomaly detection

  Log Service offers unsupervised process monitoring and fault diagnosis and supports manual data labeling. These features greatly improve monitoring efficiency and accuracy.

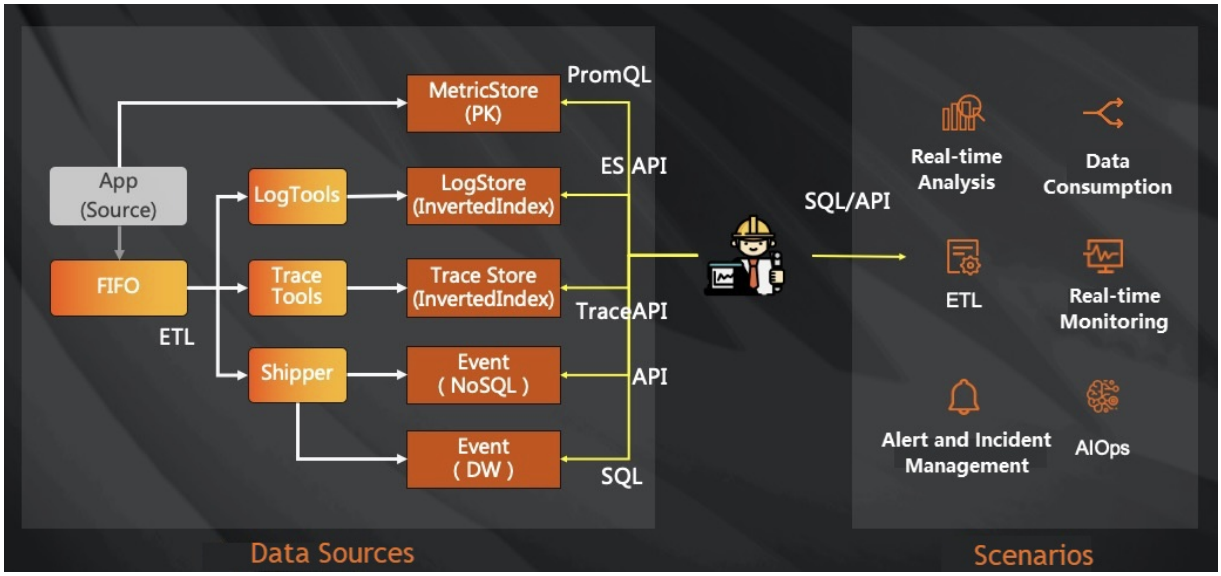

Compared with the solutions that combine multiple open source systems, Log Service is a one-stop logging service. Log Service provides only one system to meet all data monitoring and analysis requirements of O&M and SRE teams. O&M and SRE teams can use Log Service instead of multiple combined systems, such as Kafka, Elasticsearch, Prometheus, and OLAP. Log Service has the following advantages:

- Lower O&M complexity
  - Log Service is an out-of-the-box and maintenance-free cloud service.
  - Log Service supports visualization management. Data can be accessed within 5 minutes. The cost of business support can be greatly reduced by using Log Service.
- Lower cost
  - Data is retained as only one copy, and the data copy does not need to be transferred among multiple systems.
  - Resources are scaled in or out based on the actual scenario. No reserved resources are required.
  - Comprehensive technical support is provided to reduce labor costs.
- Better resource permission management
  - Log Service provides complete consumption data for internal separate billing and cost optimization.
  - Log Service supports full permission control and resource isolation to prevent information leakage.

To help O&M and SRE teams support business with higher efficiency, Log Service is committed to providing a large-scale, low-cost, and real-time monitoring and analysis platform for logs, metrics, and traces.

12.3. Alerting

The new alerting feature of Log Service is an inclusive intelligent O&M system that allows you to monitor data, denoise alerts, manage alert incidents, and distribute alert notifications. This topic provides a comparison between the new alerting feature of Log Service and various open source alerting systems.

New alerting feature of Log Service

The new alerting feature of Log Service is used to monitor various data such as log data and time series data, receive third-party alerts, denoise alerts, manage alert incidents, and distribute alert notifications. The new alerting feature of Log Service supports 40-plus more use scenarios than the original alerting feature of Log Service and can meet various monitoring, alerting, and O&M requirements of R&D, O&M, security, and operations engineers. For more information, see The alerting feature of Log Service.

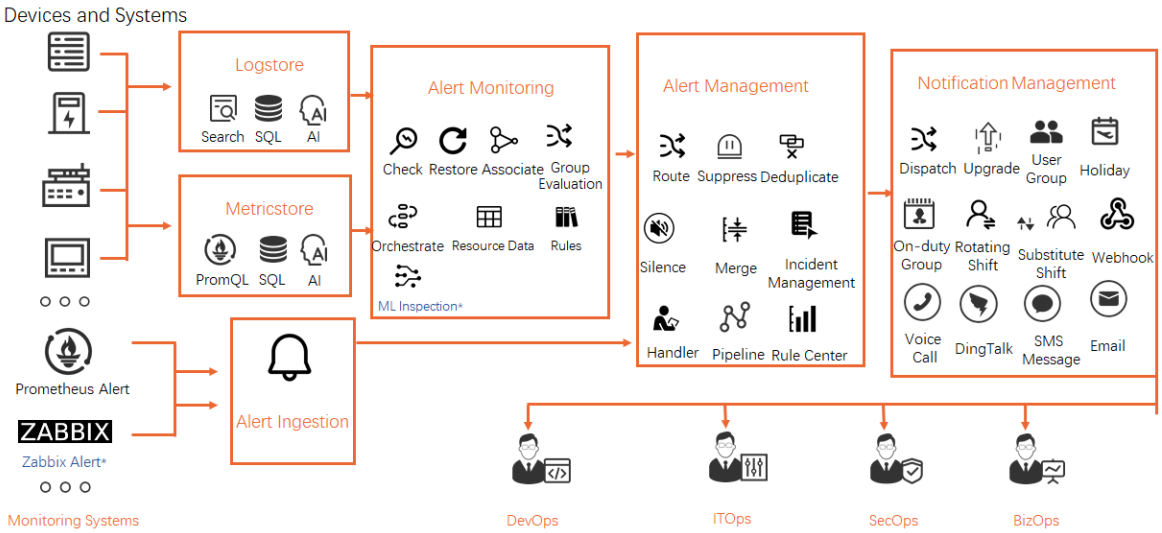

The following figure shows the benefits of the new alerting feature provided by Log Service.

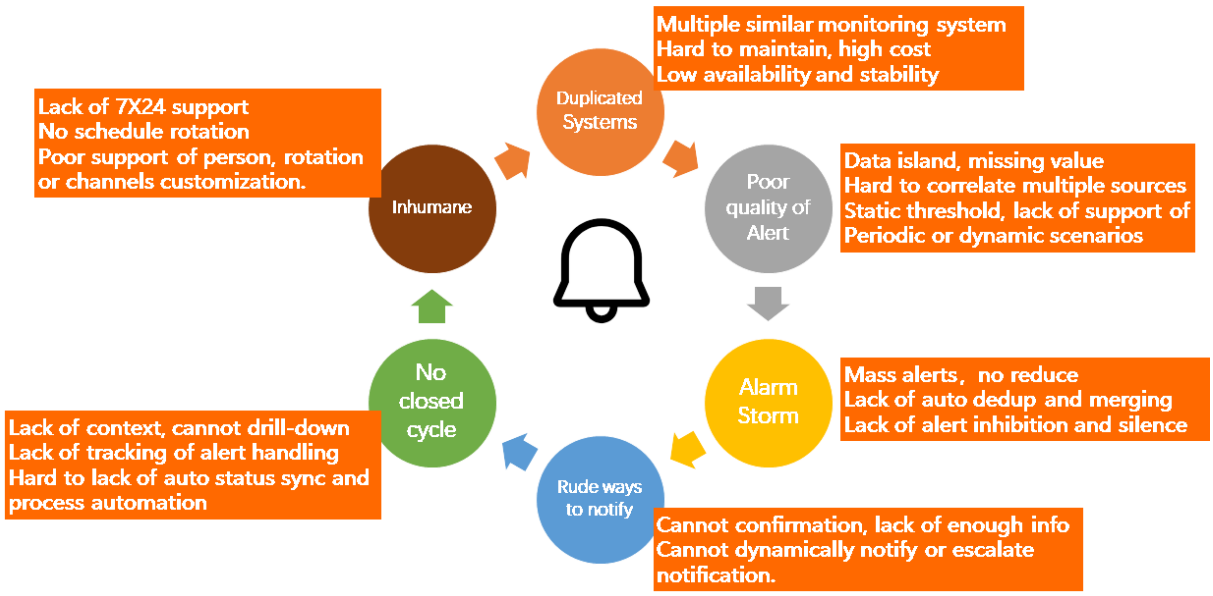

Comparison between Log Service and the ELK Stack

The ELK Stack is a combination of the following three open source projects: Elasticsearch, Logstash, and Kibana. This combination does not provide an alerting feature. If you want to configure alerts for your own ELK Stack, you must purchase the X-Pack extension. The X-Pack extension provides the following two alerting features: Elasticsearch Watcher and Kibana 7.x+Alert. These two alerting features are independent of each other and cannot be coordinated or associated.

Category	Item	Log Service	ELK Stack
Durability	Alerting service availability	Log Service supports an alerting service availability that is greater than 99.9% and an alerting data storage durability that is greater than 99.99999999%.	The ELK Stack is deployed in a distributed architecture. You must manually configure storage options.
Cost-effectiveness	Fee	You are not charged for subscriptions, monitoring, or alert management. No O&M labor is required. You need only to pay a small amount of fee for the text messages and voice messages that are sent to notify you of the generated alerts.	You must pay for subscriptions, O&M labor, purchased machines, and third-party text messages and voice messages.
Monitoring and alerting	Maximum amount of log data and time series data	Log Service can monitor petabytes of data.	The ELK Stack can monitor terabytes of data.
	Syntax	Log Service supports SQL92 and related extensions, PromQL, and alerting syntax extensions.	• Elasticsearch Watcher supports Elasticsearch DSL. • Kibana 7.x+Alert supports a limited number of filter operations and aggregate operations.
	Machine learning	Log Service supports more than a dozen AI algorithms that are used for prediction, exception detection, and root cause analysis.	The ELK Stack supports the machine learning algorithms of the X-Pack extension.
	Data collaboration	Log Service allows you to monitor data across multiple data stores, projects, regions, and accounts in a collaborative manner.	The ELK Stack allows you to merge and analyze data that has homogeneous indexes in the same cluster.
	Alerting in the event of no data	Log Service can trigger alerts if no data is detected.	The ELK Stack cannot trigger alerts if no data is detected.
	Alert clearance	Log Service can clear alerts.	The ELK Stack cannot clear alerts.
	Tag and label	Log Service supports tags and labels.	The ELK Stack supports custom tags.
	Dynamic adjustment of exception severity	Log Service can dynamically adjust the severity of exceptions.	The ELK Stack cannot dynamically adjust the severity of exceptions.

Category	Item	Log Service	ELK Stack
	Evaluation by group	Log Service allows you to create custom groups and can evaluate data by group.	• Elasticsearch Watcher cannot evaluate data by group. • Kibana 7.x+Alert automatically groups data and can evaluate data by group.
	Monitoring control	• Log Service can trigger consecutive alerts for a metric based on the threshold that you specify. • Log Service can pause and automatically resume monitoring activities based on the time that you specify.	Elasticsearch Watcher can pause and automatically resume monitoring activities based on ACK messages.
Alert management	Alert management	• Log Service allows you to deduplicate, merge, denoise, and silence alerts. • Log Service allows you to manage alert incidents and specify owners.	The ELK Stack does not allow you to manage alerts.
Notification management	Notification distribution	Log Service can dynamically distribute alert notifications. Log Service also allows you to escalate alert levels, manage contact groups, configure calendars, configure shift schedules, and control the quota for the notification channel that you specify.	The ELK Stack does not allow you to manage the distribution of alert notifications.
	Notification channel	Log Service supports notification channels such as text messages, voice messages, DingTalk, emails, webhooks, and Alibaba Cloud Message Center. Log Service also supports webhook-based notification channels such as Enterprise WeChat, Lark, and Slack.	The ELK Stack supports notification channels such as emails and webhooks. The ELK Stack does not support text messages or voice messages. • Elasticsearch Watcher supports PagerDuty, JIRA, and Slack. • Kibana 7.x+Alert supports IBM Resilient, Microsoft Teams, and ServiceNow.

Comparison between Log Service and Prometheus+Loki 2.0

Prometheus+Loki 2.0 is a combination of the following three open source projects: Prometheus, Loki, and Alertmanager. Prometheus monitors time series data, and Loki monitors log data. Prometheus and Loki send alerts to Alertmanager. Alertmanager centrally manages the alerts.

Category	Item	Log Service	Prometheus+Loki 2.0
Durability	Alerting service availability	Log Service supports an alerting service availability that is greater than 99.9% and an alerting data storage durability that is greater than 99.99999999%.	In Prometheus+Loki 2.0, some alerting services are deployed in distributed architectures, and some are deployed in standalone architectures. The storage layer of Prometheus+Loki 2.0 is deployed in a standalone architecture.
Cost-effectiveness	Fee	You are not charged for subscriptions, monitoring, or alert management. No O&M labor is required. You need only to pay a small amount of fee for the text messages and voice messages that are sent to notify you of the generated alerts.	You must pay for O&M labor, purchased machines, and third-party text messages and voice messages.
Monitoring and alerting	Maximum amount of log data and time series data	Log Service can monitor petabytes of data.	• Prometheus+Loki 2.0 can monitor hundreds of GB of log data. • Prometheus+Loki 2.0 can monitor terabytes of time series data.
	Syntax	Log Service supports SQL92 and related extensions, PromQL, and the alerting syntax.	• Prometheus+Loki 2.0 uses LogQL to process log data. • Prometheus+Loki 2.0 uses PromQL to process time series data.
	Machine learning	Log Service supports more than a dozen AI algorithms that are used for prediction, exception detection, and root cause analysis.	Prometheus+Loki 2.0 does not support machine learning algorithms.
	Data collaboration	Log Service allows you to monitor data across multiple data stores, projects, regions, and accounts in a collaborative manner.	Prometheus+Loki 2.0 allows you to run joins cross metrics in the same cluster by using PromQL.
	Alerting in the event of no data	Log Service can trigger alerts if no data is detected.	Prometheus+Loki 2.0 cannot trigger alerts if no data is detected.
	Alert clearance	Log Service can clear alerts.	Prometheus+Loki 2.0 can clear alerts.

Category	Item	Log Service	Prometheus+Loki 2.0
	Tag and label	Log Service supports tags and labels.	Prometheus+Loki 2.0 supports tags and labels.
	Dynamic adjustment of exception severity	Log Service can dynamically adjust the severity of exceptions.	Prometheus+Loki 2.0 cannot dynamically adjust the severity of exceptions.
	Evaluation by group	Log Service allows you to create custom groups and can evaluate data by group.	Prometheus+Loki 2.0 provides determinate groups and can evaluate data by group.
	Monitoring control	• Log Service can trigger consecutive alerts for a metric based on the threshold that you specify. • Log Service can pause and automatically resume monitoring activities based on the time that you specify.	Prometheus+Loki 2.0 can trigger consecutive alerts for a metric based on the threshold that you specify. However, Prometheus+Loki 2.0 cannot pause or automatically resume monitoring activities based on the time that you specify.
Alert management	Alert management	• Log Service allows you to deduplicate, merge, denoise, and silence alerts. • Log Service allows you to manage alert incidents and specify owners.	Prometheus+Loki 2.0 allows you to deduplicate, merge, denoise, and silence alerts. However, Prometheus+Loki 2.0 does not allow you to manage alert incidents or specify owners.
Notification management	Notification distribution	Log Service can dynamically distribute alert notifications. Log Service also allows you to escalate alert levels, manage contact groups, configure calendars, configure shift schedules, and control the quota for the notification channel that you specify.	Prometheus+Loki 2.0 can only dynamically distribute alert notifications.

Category	Item	Log Service	Prometheus+Loki 2.0
	Notification channel	Log Service supports notification channels such as text messages, voice messages, DingTalk, emails, webhooks, and Alibaba Cloud Message Center. Log Service also supports webhook-based notification channels such as Enterprise WeChat, Lark, and Slack.	Prometheus+Loki 2.0 supports notification channels such as emails, Enterprise WeChat, webhooks, PagerDuty, Pushover, Slack, Opsgenie, and VictorOps. You cannot define the body of a webhook URL. Prometheus+Loki 2.0 does not support text messages or voice messages. Prometheus+Loki 2.0 also supports DingTalk, Lark, and Slack that are integrated with third-party plug-ins.

Comparison between Log Service and InfluxDB 2.0

InfluxDB uses open source OSS 2.0 and Kapacitor to build a monitoring and alerting system. If you want to deploy your InfluxDB system as a cluster, you must purchase the InfluxDB Enterprise edition. The monitoring and alerting system of InfluxDB can monitor only time series data.

Category	Item	Log Service	InfluxDB 2.0
Durability	Alerting service availability	Log Service supports an alerting service availability that is greater than 99.9% and an alerting data storage durability that is greater than 99.99999999%.	The InfluxDB Enterprise edition is deployed in a distributed architecture and allows you to configure storage options. Open source InfluxDB is deployed in a standalone architecture.
Cost-effectiveness	Fee	You are not charged for subscriptions, monitoring, or alert management. No O&M labor is required. You need only to pay a small amount of fee for the text messages and voice messages that are sent to notify you of the generated alerts.	You must pay for subscriptions, O&M labor, purchased machines, and third-party text messages and voice messages.
	Maximum amount of log data and time series data	Log Service can monitor petabytes of data.	• InfluxDB cannot monitor log data. • InfluxDB can monitor terabytes of time series data.
	Syntax	Log Service supports SQL92 and related extensions, PromQL, and alerting syntax extensions.	InfluxDB supports Flux.

Category	Item	Log Service	InfluxDB 2.0
Monitoring and alerting	Machine learning	Log Service supports more than a dozen AI algorithms that are used for prediction, exception detection, and root cause analysis.	InfluxDB supports Loud ML.
	Data collaboration	Log Service allows you to monitor data across multiple data stores, projects, regions, and accounts in a collaborative manner.	InfluxDB allows you to run joins across multiple buckets in the same cluster by using Flux.
	Alerting in the event of no data	Log Service can trigger alerts if no data is detected.	InfluxDB cannot trigger alerts if no data is detected.
	Alert clearance	Log Service can clear alerts.	InfluxDB cannot clear alerts.
	Tag and label	Log Service supports tags and labels.	InfluxDB supports only simple tags.
	Dynamic adjustment of exception severity	Log Service can dynamically adjust the severity of exceptions.	InfluxDB can dynamically adjust the severity of exceptions.
	Evaluation by group	Log Service allows you to create custom groups and can evaluate data by group.	InfluxDB does cannot evaluate data by group.
	Monitoring control	• Log Service can trigger consecutive alerts for a metric based on the threshold that you specify. • Log Service can pause and automatically resume monitoring activities based on the time that you specify.	InfluxDB does not support monitoring control.
Alert management	Alert management	• Log Service allows you to deduplicate, merge, denoise, and silence alerts. • Log Service allows you to manage alert incidents and specify owners.	InfluxDB allows you only to denoise alerts.

Category	Item	Log Service	InfluxDB 2.0
Notification management	Notification distribution	Log Service can dynamically distribute alert notifications. Log Service also allows you to escalate alert levels, manage contact groups, configure calendars, configure shift schedules, and control the quota for the notification channel that you specify.	InfluxDB can only dynamically distribute alert notifications.
	Notification channel	Log Service supports notification channels such as text messages, voice messages, DingTalk, emails, webhooks, and Alibaba Cloud Message Center. Log Service also supports webhook-based notification channels such as Enterprise WeChat, Lark, and Slack.	InfluxDB supports notification channels such as emails, webhooks, exec, PagerDuty, Pushover, Slack, Opsgenie, VictorOps, and HipChat. You cannot define the body of a webhook URL. InfluxDB does not support text messages or voice messages.