

Log Service

Pricing

Pricing

Log Service is billed by resource usage on a tiered basis each month. **LogShipper** is free of charge. You can use LogShipper to send your logs to MaxCompute and OSS for storage and analysis at no cost. This document describes the **billing method**, **deduction method** and **billing example** of Log Service in details.

Billing method

Description:

- **The billing cycle is one day:** The bill is sent to you every day and the service is billed by the resource usage in the day.
- **The FreeTier quota cycle is one month** and the remaining quota will be cleared at the end of the month. If your resource usage does not exceed the FreeTier quota, no charge is collected; otherwise, the part exceeding the quota is charged.
- **Any billing item of less than 0.01 USD is excluded from the bill.**

The resource billing items and unit prices are listed in the following table.

Billing items	Description	Price	FreeTier quota (per month)	Example
Read and write traffic	The read and write traffic is calculated by the traffic for transmitting compressed logs. Logs are automatically compressed in SDK/Logtail mode, but need to be manually compressed in APIs mode. Logs are generally compressed by 5 to 10 times.	0.045 USD/GB	500 MB	If the raw size of the logs is 10 GB per day and the size of the logs after compression is 1.5 GB, the logs are charged by 1.5 GB.
Storage space	The storage space is the sum of the size of the raw data	0.002875 USD/GBDay	500 MB	The size of the raw data is 1 GB per day. After

	after compression and the size of the index.			<p>compression, the data size is 200 MB, and 10% of the data is used for index (the size is 100 MB). After a storage period of 30 days, the maximum cumulative size of the stored data is $30 \times (100 + 200) = 9$ GB. The maximum fee is $0.002875 \times 9 \approx 0.026$ USD per day.</p>
Indexing traffic	<ul style="list-style-type: none"> - The indexing traffic is calculated by the actual index fields. Storage fee is collected in full during writing. - The traffic of fields having both FullText indexing 	0.0875 USD/GB	500 MB	<p>For example, if 10% fields in 10-GB logs need to be queried, only the traffic used to query the fields is charged. The index traffic fee is 0.0875 USD.</p>

	<p>s and KeyVal ue indexes is calculated only once. For differences between different indexes, see Index settings.</p> <p>- Indexes occupy the storage space and thus the storage space fee is collected.</p>			
<p><i>*Other billing items</i> (the following billing items are provided to restrict abuse of resources and the price is very low by default)</p>				
Active Shard rent	Only current read-write Shards are	0.01 USD/day	31 /day	For example, there are three Shards, one is

	counted. Rent of Merged/splitte d Shards is not collected.			in read-and-write state, and the other two are merged and in read-only sate. Only one shard rent (0.01 USD/day) is collected.
Read/write count	The write count of logs into the Log Service is subject to the log generation speed. The background realization mechanism assures the minimum read/write count as much as possible.	0.03 USD/1 million requests	1106	You can use Logtail for automatic batch sending and a total of 0.03 USD is collected for one million write operations.
Internet read traffic	It refers to the traffic consumed when Internet reads logs collected by Log Service.	0.2 USD/GB	None	The fee of Internet read traffic is 0.4 USD when 2 GB logs are shipped by Log Service to non-Alibaba Cloud products.

[Price Calculator \(Excel version\) downloading](#)

Deduction and arrears

A bill is generally provided within four hours after the current billing cycle. The system automatically deducts the bill amount from you amount balance. If the account balance is insufficient, your account is in arrears.

If the overdue bill is not paid off within 24 hours, your Log Service will stop automatically. However, you will be still charged for the storage space you are using, so the overdue amount will increase. You are recommended to renew the arrears within 24 hours to avoid any service loss caused by service stop.

If the arrears are paid off within six days, the service will be automatically restored. Otherwise, Alibaba Cloud will assume that you have chosen to stop using the service, the

project space will be reclaimed and data in it will be cleared. This data will not be recoverable.

Billing example

Case 1: FreeTier quota

You have three servers, of which one server generates 5 MB logs per day. You want to use a program to process the logs as follows:

1. Real-time query the logs and create a dashboard for online maintenance.
2. Use a Java program to subscribe for logs in real time.
3. Ship the logs to OSS.

Billing details:

- Resource: One Logstore (one Shard) is created per day. 31 Logstores (Shard) are created in the month in total, not exceeding the quota.
- Read and write traffic: The read and write traffic is $15/5$ (compressed data) * 2 (read + write) = 6 MB (after compression). The accumulated read and write traffic is $6 * 31$ (days) = 183 MB, not exceeding the quota.
- Indexing traffic: 15 (raw data) * 31 (days) = 465 MB, not exceeding the quota.
- Read/write count: The read/write count in the month is less than one million, not exceeding the quota.

You can use Log Service for online log analysis and processing for free each month.

Case 2: Real-time computing + Offline Computing (Lambda Architecture)

The website has 100 million access requests from APIs per day, each request generates a 200-byte log, and the size of logs generated is 20 GB per day. The peak traffic is 5 times of the average traffic, 1.16 MB/s (less than 5 MB/s). The logs are read once per day for real-time computing (the life cycle is two days) and imported to OSS for offline computing (Hive/Spark).

Billing details:

- Active Shard rent: One Shard is reversed. The fee is 0.01 USD/day.
- Read/write count: You use Logtail for automatic batch sending and a total of 0.03 USD is collected for one million write operations.
- Read and write traffic:
 - The write traffic is 20 GB. If the compression rate is 10%, the actual traffic is 2 GB and the fee is $2 * 0.045 = 0.09$ USD.
 - The read traffic is same to the write traffic in real-time computing, 0.09 USD.
- Storage space: The size of the stored data is $2 \text{ GB} * 2$ (Day)=4 GB. The storage fee per day is $4 * 0.002875 = 0.0115$ USD.

- Fee of logs imported into OSS is not charged.

The maximum fee per day is $0.01 + 0.03 + 0.09 * 2 + 0.115 = 0.335$ USD.

Case 2: Log online query and analysis

The service has one million API access requests from APIs per day, each request generates a 200-byte log, and the size of logs generated is 200 MB per day. Logs of the last 30 days are saved for query.

Billing details:

- Active Shard rent: One Shard is reversed. The fee is 0.01 USD/day.
- Read/write count: You use Logtail for automatic batch sending and a total of 0.03 USD is collected for one million write operations.
- Read/write traffic: The write traffic is 200 MB. The size of the compressed data is 0.05 GB. The read traffic fee per day is $0.05 * 0.045 = 0.00225$ USD.
- Indexing traffic: The indexing traffic per day is $0.2 * 0.0875 = 0.0175$ USD.
- Storage space: 200 MB + 50 MB (compressed raw data) = 250 MB. The storage peak size is $250 * 30 = 7.5$ GB. The storage fee per day is $7.5 * 0.002875 = 0.022$ USD.

The maximum fee per day is $0.01 + 0.03 + 0.00225 + 0.0175 + 0.22 = 0.08175$ USD.

Log Service has the following cost advantages in three log processing scenarios:

- LogHub:
 - A more cost-effective choice for users in 98% scenarios compared to building Kafka with purchased cloud hosts + cloud disks. At 30% or less of the cost of Kafka for small websites.
 - Provides RESTful APIs and supports data collection on mobile devices, saving you the cost of the gateway servers for log collection. See [Collection-Public Network Data](#).
 - Maintenance-free and auto scaling anytime and anywhere. See the auto scaling part in [Processing-Ensure Logs are Processed in Order, at Least Once, and Exactly Once Using ConsumerLib](#).
- LogShipper:
 - No code/machine resources required, flexible configuration, and rich monitoring data.
 - Linear scalability (PBs per day), available for free currently.
- LogSearch/Analytics:
 - At 15% or less of the cost of purchasing cloud hosts + self-building ELK, and offers dramatic enhancement in query capability and data processing scale. See [Comparison Report](#).
 - A better choice than the abovementioned log management software solutions for its ability to seamlessly integrate with various popular stream computing + offline computing frameworks to allow for unobstructed flow of logs.

Cost comparison

The following is the comparison of Log Service and self-built solutions in billing model, for your reference only.

LogHub (LogHub vs Kafka)

	Focus	LogHub	Self-built middleware (such as Kafka)
Use	Add	Imperceptible	Maintenance required
	Scale up	Imperceptible	Maintenance required
	Increase backups	Imperceptible	Maintenance required
	Multitenancy	Isolated	May interfere with each other
Cost	Internet collection (10 GB/day)	CNY2/day	CNY16.1/day
	Internet collection (1 TB/day)	CNY162/day	CNY800/day
	Intranet collection (small data volume)	\$	\$\$\$
	Intranet collection (moderate data volume)	\$\$	\$\$\$
	Intranet collection (large data volume)	\$\$\$	\$\$\$

For detailed cost comparison, see LogHub vs Kafka in Public Network Data Collection

Log Storage and Query Engine

	Focus	LogSearch	ES (Lucene Based)	NoSQL	Hive
Scale	Scale	PB	TB	PB	PB
Cost	Store (CNY/GB *day)	0.0115	3.6	0.02	0.035
	Write (CNY/GB)	0.35	5	0.4	0
	Query	0	0	0.2	0.3

	(CNY/GB)				
	Speed-query	Millisecond level-second level	Millisecond level-second level	Millisecond level	Minute level
	Speed-Statistics	Weak+	Strong-	Weak	Strong
Latency	Write->Queryable	Real time	Minute level	Real time	Ten minutes level

Note: The price comparisons here are calculated based on the fact that the software solutions are deployed on ECS and three copies have been configured.

For more information, see [Comparisons of Log Query Solutions](#).

Cost is related to two factors:

- Data volume. The amount of user data depends on the volume of business and cannot be optimized.
- Configuration. Configuration can be optimized. Only with an appropriate configuration and the best solution can the cost be minimized.

Optimizing the configuration

Configuration can be optimized in the following two ways:

Number of shards

The price for one shard is CNY 0.04 per day, with a maximum data processing capability of 5 Mbit/s. Only shards with read/write accesses are charged. You can adjust the number of shards, so that each shard can process the data exactly at 5 Mbit/s. To reduce the number of shards, you can perform the “合并分区” action.

– Storage cycle of indexes

We recommend that you optimize the storage time of indexes based on your own log query and storage needs. For more information on how to change the storage cycle of indexes, see the [document](#).

If you collect the logs for StreamCompute, we recommend that you only use LogHub without creating indexes.

If you have the need to query the logs within the past 90 days, and barely query the logs earlier than that, you are advised to change the storage cycle of indexes to 90 days and

import the data into MaxCompute. You can query the data within the past 90 days by using the Log Service, and query the data earlier than 90 days ago by using MaxCompute.

If you have the need for long-time storage and backup of the logs, we recommend that you configure the OSS Shipper, and import the logs into OSS.

Other optimization recommendations

- Use Logtail: With capabilities of batch processing and resumable uploads and downloads, the data can be transferred with optimal algorithm while delivering good real-time performance. Logtail consumes 3/4 less resources than that of the open source software (such as Logstash and FluentD), thus reducing CPU usage.
- Try to use large packages (64 KB - 1 MB) to write into logs using API, thus reducing the number of requests.
- Only perform indexes for key fields (for example, UserID and Action) instead of the useless fields.