

Log Service

Pricing

Pricing

Billing method

Log Service is billed by resource usage on a tiered basis each month. **LogShipper** is free of charge. You can ship log data to MaxCompute and Object Storage Service (OSS) for storage and analysis at no cost. This document introduces 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 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 USD 0.01 is excluded from the bill.**

The resource billing items and unit prices are as follows.

Resource 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 API mode. Logs are generally	USD 0.045/GB	500 MB	If the size of the raw logs is 10 GB and the size of logs after compression is 1.5 GB, the logs are charged by 1.5 GB.

	compressed by 5 to 10 times.			
Storage space	The storage space is the sum of data size after compression and the indexed data size.	USD 0.002875/GB per day	500 MB	The data size is 1 GB per day. After compression, the data size is 200 MB. 10% of the data is used for index (the size is 100 MB). After a storage period of 30 days, the accumulative maximum storage size is $30 \times (100 + 200) = 9 \text{ GB}$. The maximum fee for storage is $0.002875 \times 9 \approx \text{USD } 0.026$ per day.
Indexing traffic	<ul style="list-style-type: none"> - The indexing traffic is calculated by actual index fields. Storage fee is collected in full during writing. - Traffic of fields having both FullTex 	USD 0.0875/GB	500 MB	For example, if 10% fields in 10 GB logs need to be queried, only the traffic used to query the fields is charged. The indexing traffic fee is USD 0.0875.

	<p>t indexes and KeyVal ue indexes is calcula ted only once. For differe nces betwe en differe nt indexes s, see Index setting s. - Indexe s occupy the storag e space after being create d. Theref ore, the storag e space fee is collect</p>			
--	---	--	--	--

	ed.			
Other billing items (to restrict the abuse of resources, the price of the following billing items is low by default)				
Active shard rent	Only shards currently in readwrite status are counted. Rent of merged/split shards is not collected.	USD 0.01/day	31/day	For example, one of the three shards is in readwrite status. The other two shards are merged and in readonly status. The rent of only one shard (USD 0.01/day) is collected.
Read/write count	The count of writing logs to Log Service is subject to the log generation speed. The background realization mechanism attempts to minimize the read/write count.	USD 0.03/1 million times	106 times	You can use Logtail for automatic batch sending and a total of USD 0.03 is collected for one million write operations.
Internet read traffic	The data traffic generated when Internet programs read log data collected by Log Service.	USD 0.2/GB	None	The fee of Internet read traffic is USD 0.4 when 2 GB data is shipped by Log Service to non-Alibaba Cloud products.

Deduction and outstanding payment

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

If the overdue bill is not paid off within 24 hours, your service will be stopped automatically. However, you will be still charged for the storage space you are using, so the overdue

amount will increase. We recommend that you pay off the overdue bill within 24 hours to avoid any business loss caused by service stop.

If the overdue bill is paid off within six days after the service is stopped, 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 will be cleared. The cleared data is not 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. Query the logs in real time and create a dashboard for online Operation & Maintenance (O&M).
2. Use a Java program to subscribe to log processing in real time.
3. Ship logs to OSS.

Billing details:

- Resource: One Logstore (one shard) is created per day. 31 Logstores (one shard) are created in a month in total, not exceeding the quota.
- Read and write traffic: The read and write traffic per day is $15/5$ (data after compression) * 2 (read + write) = 6 MB (after compression). The accumulated read and write traffic in a month 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 a 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 API requests per day. Each request generates a 200-byte log, and the size of logs generated per day is 20 GB. 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 lifecycle is two days) and imported to OSS for offline computing (Hive/Spark).

Billing details:

- Active shard rent: One shard is reserved. The fee is USD 0.01/day.

- Read/write count: Use Logtail for automatic batch sending and a total of USD 0.03 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 $2 * 0.045 = \text{USD } 0.09$.
 - The read traffic is the same as the write traffic in real-time computing, which is USD 0.09.
- Storage space: The storage size is $2 \text{ GB} * 2 \text{ (days)} = 4 \text{ GB}$. The storage fee per day is $4 * 0.002875 = \text{USD } 0.0115$.
- Importing logs to OSS is free of charge.

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

Case 3: Online log query and analysis

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

Billing details:

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

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

Cost advantages

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 less than 30% of the Kafka cost 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 [Collect Internet data](#).

- Operation & Maintenance (O&M) -free and auto scaling anytime and anywhere. See the contents about auto scaling in [Processing](#) - use ConsumerLib to process logs in order, at least once, and exactly once.

- LogShipper:

- No code/machine resources required. Flexible configuration and rich monitoring data.
- Linear scalability (PB grade/day), available for free currently.

- LogSearch/Analytics:

- At less than 15% 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 above-mentioned log management softwares 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 the billing model, for your reference only.

LogHub (LogHub vs Kafka)

	Focus	LogHub	Self-built middleware (such as Kafka)
Use	Add	Imperceptible	O&M required
	Expand	Imperceptible	O&M required
	Increase backups	Imperceptible	O&M required
	Multitenancy	Isolated	Might affect each other
Cost	Internet collection (10 GB/day)	USD 2/day	USD 16.1/day
	Internet collection (1 TB/day)	USD 162/day	USD 800/day
	Intranet collection (small data size)	\$	\$\$\$
	Intranet collection (moderate data size)	\$\$	\$\$\$
	Intranet collection (large data size)	\$\$\$	\$\$\$

For detailed cost comparison, see [LogHub vs Kafka in Internet data collection](#).

Log storage and query engine

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

Note: The price comparisons here are calculated basically based on the fact that softwares are deployed on Elastic Compute Service (ECS) and three copies have been configured.

For more information, see [Comparisons of log query solutions](#).

Cost optimization

Cost is related to two factors:

- Data volume. Data volume is determined by your business needs and cannot be optimized.
- Configurations. Configurations can be optimized. With the configurations matching with the amount of data, you can select the optimal solution to minimize the cost.

Optimizing configurations

Configurations can be optimized in the following two aspects:

Number of shards

The price for one shard is USD 0.04 per day, with a maximum data processing capability of 5 MB/s. Only shards in readwrite status are charged. Adjust the number of shards so that each shard can process data exactly at 5 MB/s. To reduce the number of shards, merge the shards.

Storage cycle of indexes

We recommend that you optimize the storage time of indexes based on your requirements for log query and storage.

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

If you want to query the logs within the past 90 days, and barely query the logs earlier than that, we recommend that you change the storage time of indexes to 90 days, and import the data to MaxCompute. To query data within 90 days, use Log Service. To query data earlier than that, use MaxCompute.

If you want to store and back up logs for a long time, we recommend that you configure the Object Storage Service (OSS) Shipper, and import the logs to OSS.

Other optimization recommendations

- Use Logtail: With the functions of batch processing and breakpoint transmission, data is transmitted with optimal algorithm while guaranteeing the real-timeliness. Logtail consumes 3/4 less resources than that of the open-source software (such as Logstash and Fluentd), thus reducing CPU consumption.
- Try to use large packages (64 KB–1 MB) to write logs by using API, thus reducing the number of requests.
- Only index key fields (for example, UserID and Action), without configuring indexes for useless fields.