# Cloud Log Service

# Getting Started

# Product Documentation

# Contents

# Getting Started
# Getting Started in 5 Minute

Last updated : 2021-05-31 16:24:41

# Overview

CLS allows you to collect/store/search/ship logs, analyze charts, and monitor alarms. With CLS, you can collect logs for central management, search, and analysis. Also, you can set monitor alarms for log topics and ship the collected logs to COS or other products for further analysis.

To get you started, instructions to the following CLS features are described in this document:

- Collecting logs with LogListener
- Searching logs
- Shipping logs

## Directions

### 1. Activate CLS

First, you need to activate CLS on Tencent Cloud.

### 2. Download and install LogListener

LogListener is a client that collects log data and sends it to CLS in a fast and non-intrusive way. You can install it as follows:

#### 2.1 Check the network connection

To install LogListener, the source server and the CLS region must be able to connect to each other. Tencent's Cloud Virtual Machine (CVM) accesses CLS via a private network by default.
You can run the following command to check the network connectivity, where `<region>` is the abbreviation for the CLS region. For more information about regions, please see Available Regions.

```
shell
telnet <region abbreviation>.cls.tencentyun.com 80
```

#### 2.2 View/Create a key pair

Log in to the CAM console, view (or create) a key pair, and make sure that the key is enabled.

#### 2.3 Install LogListener

In this example, the CLS service runs in the CVM CentOS 7.2 (64-bit) environment. To download and install LogListener, see LogListener Installation Guide.

## 3. Create a log topic

CLS introduces various regions. To lower network latency, please create log resources in a region closest to your business region. For supported regions, please see Available Regions. Log topic management includes the management of logsets and log topics. A logset represents a project, while a log topic represents a type of service. A single logset may contain multiple log topics.

#### 3.1 Create a log topic

1. Log in to the CLS console.
2. In the left sidebar, click **Log Topic** to go to the management page.
3. Select a region and click **Create Log Topic**.
4. Configure as needed on the page that is displayed.
   - Log topic name: name of the log topic, such as topic_test
   - Logset: **Select an existing logset** is selected by default. Alternatively, you can select **Create Logset** and set the logset's name (e.g., cls_test) as needed.

     > Note :
     > A logset can be retained for 3−90 days. To retain it for a longer period, please submit a ticket.

5. Click **Confirm**.
   - A created log topic will be displayed in the log topic list.
   - You can click **Manage Logset** to view the created logset.

## 4. Create a server group

CLS uses a server group to manage a list of log source servers.

1. In the left sidebar, click **Server Group** to go to the management page.

2. Select the desired region and click **Create Machine Group**

3. Configure as needed on the page that is displayed.
   Multiple IPs can be input in a server group (one IP per line). For CVM instances, please input the private IP addresses directly. For more information, please see Server Group Management.

4. Click **OK**.
   After a server group is created, you can click **View** in the **Operation** column to check the connection between LogListener and servers.
   - Normal: The client's LogListener has successfully connected to CLS.
   - Abnormal: You can troubleshoot by referring to Server Group Exception.

## 5. Configure LogListener

1. In the left sidebar, click **Log Topic** to go to the management page.
2. Click the desired log topic ID/name to go to the log topic management page.
3. Select the **Collection Configuration** tab to specify the collection path, bound server group, and parsing mode.

> Note :
> The following describes how to collect logs using LogListener. For more information, please see Collection Methods.

### 5.1 Configure a collection path

The collection path needs to match the absolute path of the log file on the server. You are required to enter two parameters: the directory prefix and the file name in the format of **[directory prefix expression]//**//[file name expression]**. LogListener matches all paths with common prefixes that satisfy the **[directory prefix expression]**, and monitors all log files under these directories (including subdirectories) that satisfy the **[file name expression]**. The detailed parameter description is as follows:

| Field | Description |
| --- | --- |
| Directory prefix | The directory prefix for log files, which supports only the wildcard characters `*` and `?`. `*` indicates that one or more characters can be matched, while `?` indicates that any single character can be matched. |

| Field | Description |
| --- | --- |
| /**/ | Current directory and all its subdirectories |
| File name | A log file name supports only the wildcard characters `¥*` and `?`. `¥*` indicates that one or more characters can be matched, while `?` indicates that any single character can be matched. |

For example, if the absolute path of the file to be collected is `/cls/logs/access.log`, then the directory prefix entered for the collection path should be `/cls/logs`, and the file name `access.log`, as shown below:

### 5.2 Bind a server group

Select an existing server group, and associate it with the current log topic. Then, LogListener will monitor the log files in this server group according to the rules you set. You may bind a log topic to multiple server groups, but a log file will only be collected into one log topic.

### 5.3 Configure a parsing mode

CLS supports various log parsing modes such as full text in a single line, separator, JSON, and full regex. The following log sample uses the separator mode (for more information, please see Separator Format).

```
sh
Tue Jan 22 14:49:45 2019;download;success;194;a31f28ad59434528660c9076517dc23b
```

- Selecting an extraction mode
  This sample uses separators to separator logs. Therefore, select **Separator** for the **Extraction Mode** field and **Semicolon** for **Separator**.
- Inputting a sample log and extracting key-value pairs
  Enter a complete log in the log sample field, and key-value pairs will be automatically extracted. Then, you can define a unique key for each key-value pair.
  In this example, the log is parsed into `Tue Jan 22 14:49:45 2019`, `download`, `success`, `194`, and `a31f28ad59434528660c9076517dc23b`. The keys defined for these 5 fields are `time`, `action`, `status`, `size`, and `hashcode` respectively. LogListener will then use this defined structure to collect data.

## 6. Configure indexes

CLS offers a log search and analysis feature based on segment indexing. We currently offer two index types, full-text index and key-value index. They can be managed on the index configuration tab in the log topic management page. Both indexes can be enabled at the same time.

| Index Type | Description |
| --- | --- |
| Full-Text | Breaks a full log into segments by delimiter, and executes keyword query based on the segments |
| Key-Value | Breaks a full log into key-value pairs according to the specifications, and executes field query based on the key-value pairs |

Here we use key-value index as an example to describe how to configure indexes. In the log topic management page, go to the **Index Configuration** tab, click **Edit**, and toggle the key to enable index status. Toggle on **Key-Value Index**. Click **Add** to add keys. Select a field type for each key, currently `long` , `double` , and `text` are supported. `text` type allows you to specify delimiters, which separate a character string into segments. Continuing the above example, enter `time` , `action` , `status` , `size` , and `hashcode` as key-value indexes, and set the field type of `size` to `long` .

Once the index rule is enabled, indexes will be created for new input data accordingly, and stored over a specified period of time depending on your configured storage cycle. Only logs for which indexes have been created can be queried for analysis. **Therefore, modifying an index rule or disabling an index only affects new input data. Unexpired legacy data will still be searchable.**

## 7. Ship logs

With CLS, you can ship data to COS or CKafka to store logs for a longer period at a lower cost. Moreover, you can analyze big data logs offline.

### 7.1 Ship logs to COS

To ship logs to COS, you can perform the following steps:

1. Create a COS bucket.
2. Log in to the CLS Console.
3. Select the desired log topic ID/ name to go to the management page.
4. Select the **Ship to COS** tab.
5. Click **Add Shipping Configuration** to create a shipping task.
   Currently, CLS supports shipping logs in CSV and JSON formats.

After a shipping task is created, CLS asynchronously ships data to the destination bucket. To view the shipping status, you can click the desired log topic and then choose the **Ship to COS** tab. Alternatively, you can click **Shipping task** in the left sidebar of the console.

### 7.2 Ship to CKafka

Only logs generated after the configuration can be shipped.
To ship logs to CKafka, you can perform the following steps:

1. Create a CKafka instance and topic.
2. Log in to the CLS Console.
3. Select the desired log topic ID/ name to go to the management page.
4. Select the **Ship to CKafka** tab.
5. Click **Edit**.
6. Select the desired CKafka instance and click **OK** to enable CKafka consumption.

Currently, CLS supports shipping original logs and JSON-formatted logs. To view the shipping status, you can click the consumed log topic and then select the **Ship to CKafka** tab.

# Collecting and Searching NGINX Access Logs

Last updated : 2021-05-19 16:15:16

## Overview

NGINX is a common reverse proxy server that handles a high volume of service requests in actual businesses. A high number of access logs are generated when the service is running, causing such problems as scattered logs and massive data within clusters. Therefore, effective log data collection and management is highly important for business OPS. Using NGINX access logs as an example, this document describes how to access NGINX logs through CLS.

### NGINX log format

You can run the log_format command to define the format of NGINX logs (access.log). The definition of each field and how to configure the index in default format are as follows.

```
log_format main '$remote_addr - $remote_user [$time_local] "$request"'
'$status $body_bytes_sent "$http_referer"'
'"$http_user_agent" "$http_x_forwarded_for"';
```

The fields are defined as follows:

| Field Name | Description |
| --- | --- |
| remote_addr | Client IP address |
| remote_user | Client name |
| time_local | Local server time |
| method | HTTP request method |
| url | URL |
| protocol | Protocol type |
| status | HTTP request status code |
| body_bytes_sent | Number of bytes sent to client |
| http_referer | Access source page URL |

| Field Name | Description |
| --- | --- |
| http_user_agent | Client browser information |
| http_x_forwarded_for | Tracking and recording of actual client IP address when the frontend has a proxy server |

# Directions

**1. Creating a logset and a log topic**

(1) Log in to the CLS Console. Click **Logset Management** in the left sidebar to enter the logset management page.
(2) On the top of the page, select a region, then click **Create Logset** to create a logset. For example, create a logset named nginx_project.
(3) Click **Logset Name** to enter the log topic management page.
(4) Click **Add Log Topic** and create a log topic. For example, create a log topic named nginx_access. The created log topic is displayed in the log topic list.

**2. Creating a server group**

> We recommend using the CLS collector to collect logs from the NGINX cluster. For more information about how to download and install the collector, see Installation Guide.

(1) In the left sidebar in the console, click **Server Group Management** to enter the server group management page.
(2) On the top of the page, select a region, then click **Create Server Group** to create a server group named nginx_group. You can enter multiple server IP addresses for one server group, one IP address per row. For Cloud Virtual Machine (CVM), you can enter private IP addresses. For more information, see Server Group Management.

**3. Configuring a collection rule**

(1) In the left sidebar in the console, click **Logset Management**, then enter the management page for the created logset and log topic respectively.
(2) On the log topic management page, click **Collection Configuration**, specify a collection path and parsing mode for the log topic, and bind the log topic to a server group.

- Set the log path and bind the log topic to a server group
  For example, set the target collection path to the local log path

`/usr/local/webserver/nginx/logs/access.log` , and bind the log topic to the server group nginx_group. The settings are shown in the following figure.

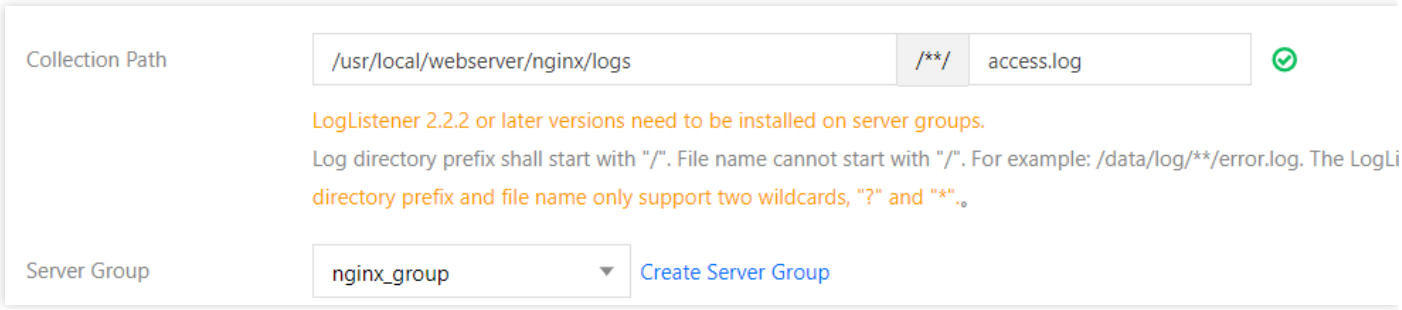


- Extract the key-value
  Set the key-value extraction mode to Full RegEx, enter a log sample, and verify the regular expression for extraction rule.
  For example, log_format of the NGINX access log is defined as follows:

```
log_format main '$remote_addr - $remote_user [$time_local] "$request"'
'$status $body_bytes_sent "$http_referer"'
'"$http_user_agent" "$http_x_forwarded_for"';
```

A complete sample of the NGINX access log is as follows:

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

The corresponding regular expression for key-value extraction is as follows:

```
(¥S+)¥s¥S+¥s(¥S+)¥s[([^]]+)]¥s&quot;([^&quot;]+)¥s(¥S+)¥s([^&quot;]+)&quot;¥s(¥d+)¥s(¥d+)¥s&quot;([^&quot;]+)&quot;¥s&quot;([^&quot;]+)&quot;¥s&quot;([^&quot;]+)&quot;$
```

After the regular expression for extraction passes verification, name a key-value for each field:

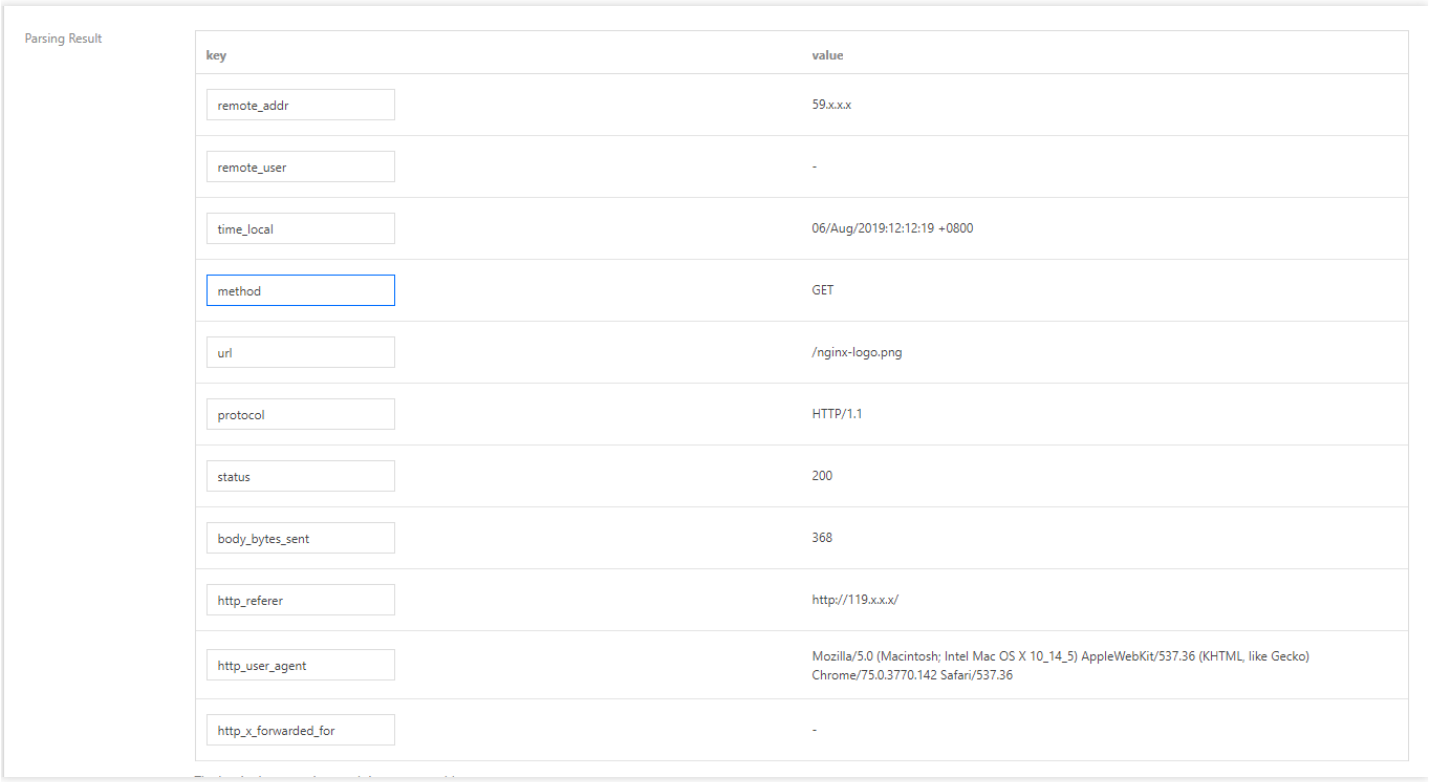


### 4. Configuring an index rule

(1) In the left sidebar in the console, click **Logset Management**, then enter the management page for the created logset and log topic respectively.

(2) On the log topic management page, click **Index Configuration**. Configure the index field on the

index configuration management page based on the definition of the NGINX log format field.

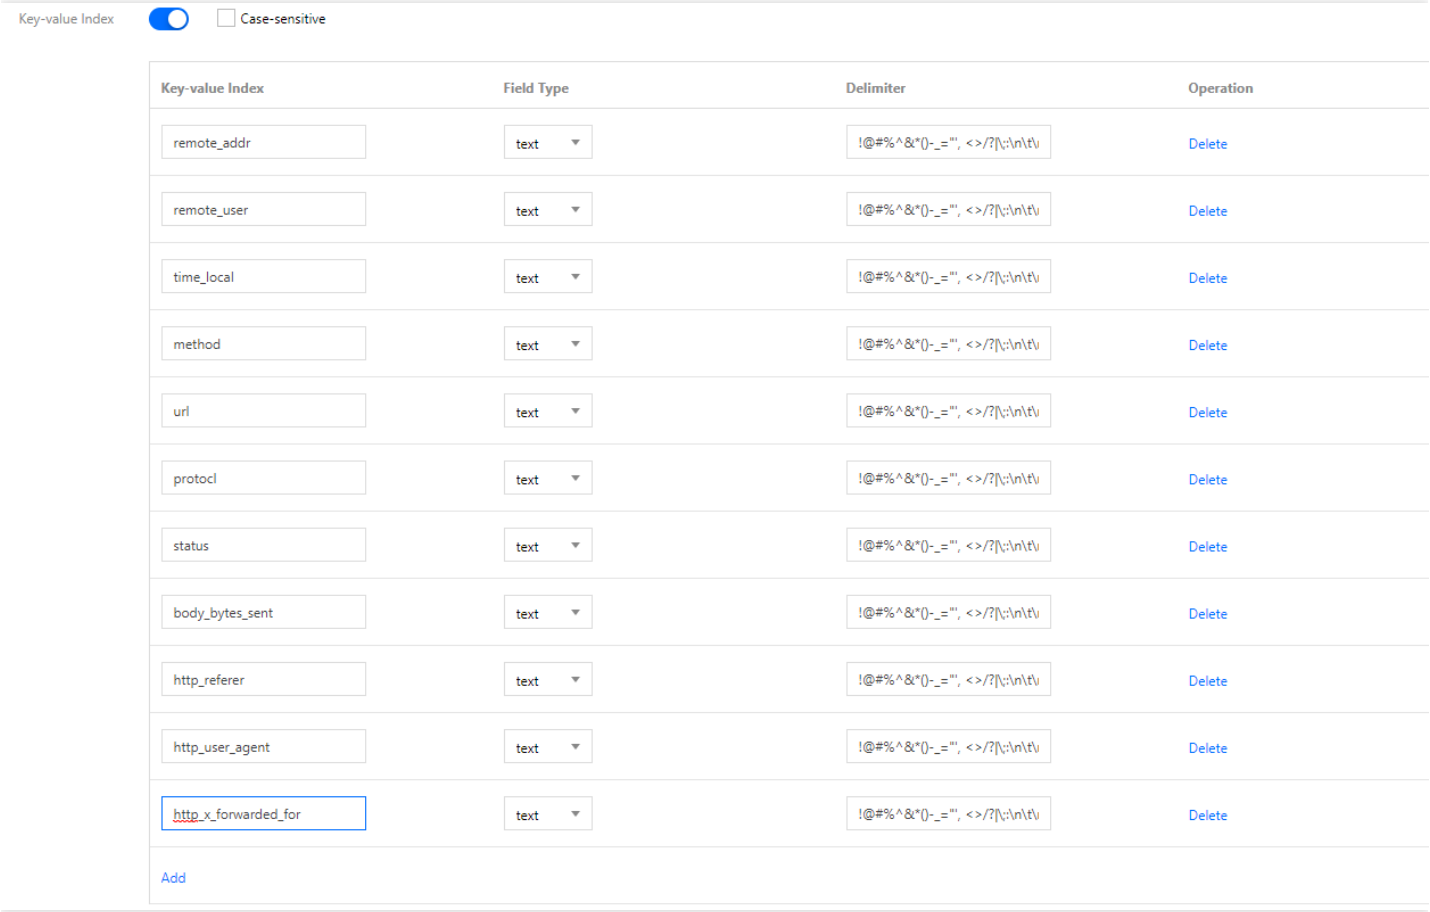


### 5. Querying a NGINX log

(1) In the left sidebar in the console, click **Log Search** to enter the log search page. Select the corresponding logset and log topic.

(2) Click **Search** to query the NGINX log.