Cloud Log Service

FAQs

Product Documentation

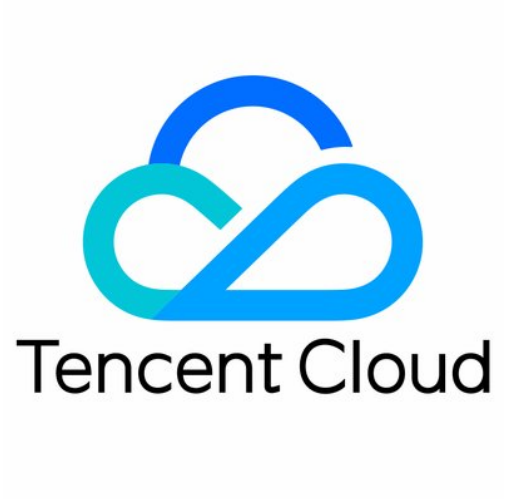

Contents

FAQs
Health Check
Index Configuration

Last updated：2024-01-20 17:11:57

How do I enable indexing configuration?

Index configuration is a necessary condition for using CLS for log search and analysis; that is, log search and analysis can be performed only after indexing is enabled. It is strongly recommended that you enable indexing.

Directions

1. Log in to the CLS console.
2. Click **Log Topic** on the left sidebar to enter the log topic list page.
3. Click the target log topic ID/name to enter the log topic management page.
4. Click the **Index Configuration** tab and click **Edit** to enter the index configuration page.
5. Toggle on **Index Status** to enable indexing.
6. You can further enable **Full-Text Index** or **Key-Value Index**.

Full-Text Index: Allows you to search for logs by keywords.

Key-Value Index: Allows you to search for logs by fields. Before enabling **Key-Value Index**, ensure that the log extraction mode in the log collection configuration is the structured mode (that is, logs are parsed into key-value pairs).

How do I enable key-value indexing configuration?

In key-value indexing, a raw log is split into multiple segments based on a field (key:value), and indexes are created based on the segments. You can search for logs based on key-value (key-value search). We strongly recommend you enable key-value indexing to maximize log search efficiency.

Note:

Before enabling key-value indexing, ensure that the log extraction mode in the log collection configuration is the structured mode (that is, logs are parsed into key-value pairs).

Directions

1. Log in to the CLS console.
2. Click **Log Topic** on the left sidebar to enter the log topic list page.
3. Click the target log topic ID/name to enter the log topic management page.
4. Click the **Index Configuration** tab and click **Edit** to enter the index configuration page.
5. Toggle on **Key-Value Index** to enable key-value indexing.
6. After enabling key-value indexing, you can also click **Auto Configure** to enable the system to automatically get the latest log collected as a sample and parse the fields in it into key-value indexes. You can perform fine tuning on the basis of automatic configuration to quickly obtain the final index configuration information.

Log Upload

Last updated：2024-01-20 17:11:57

What should I do if a message indicating a parameter error is displayed?

This error indicates that Request Parameters were specified incorrectly for log uploading via the API or SDK. Ensure that the parameters are specified correctly.

What are the possible causes for an authentication failure?

Authentication will fail when there is an authentication error during log uploading. The possible causes are as follows:

Cause	Solution
The key does not exist	Check whether the key has been deleted or disabled in the console. If its status is normal, check whether the key is entered correctly. Note that there must be no leading or trailing spaces. You can click here to view the key.
The signature is incorrect	Check the signature calculation process against the signature document under "Making API Requests".
The signature expire	Re-calculate the signature against the signature document under "Making API Requests".
The request is not authorized	There is no log upload permission. Go to the CAM console to add the CLS load upload permission for your account.
The key is invalid	The key format is incorrect. You can click here to view the key.
Others	If the error persists after you have checked the above causes, contact us.

What should I do if the logs to upload exceed the size limit?

Adjust the log size according to the following limits:

Limit	Description
Log upload	The number of logGroups in a .pb package cannot exceed 10.
	A logGroup in a .pb package can contain 1 to 10,000 log entries.
	A single value in log cannot exceed 1 MB.
	All values in the logGroup of a .pb package cannot exceed 5 MB.
	The package for a single upload request cannot exceed 6 MB before compression.

Why is throttling or frequency control triggered?

The following limits may be exceeded:

Limit	Description
Frequency limit on writes	For a single log topic partition: 500 QPS.
Bandwidth limit on writes	For a single log topic partition: 5 MB/s.

You are advised to reduce the frequency and traffic of log uploads. If you do not want to do so, enable auto-split for the log topic.

How do I handle an upload request error?

Contact us.

Collection
Machine Group Exception

Last updated：2024-01-20 17:11:57

Error Description

When the machine group is configured, LogListener may have an exception, such as disconnection from the CLS server and log upload failure. In this case, the machine group is exceptional, as shown below:

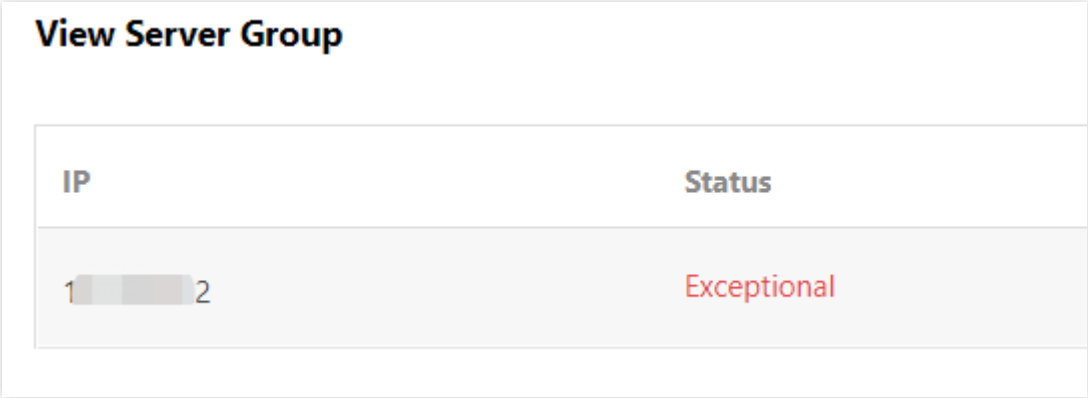

Troubleshooting Directions

Note:
These troubleshooting steps only apply to LogListener 2.2.4 or later. If you're using an earlier version, see Troubleshooting Earlier LogListener Versions.

1. Use the LogListener diagnostic tool

This tool helps you quickly check the LogListener operation, heartbeat and configuration.
Run the following CLI commands.

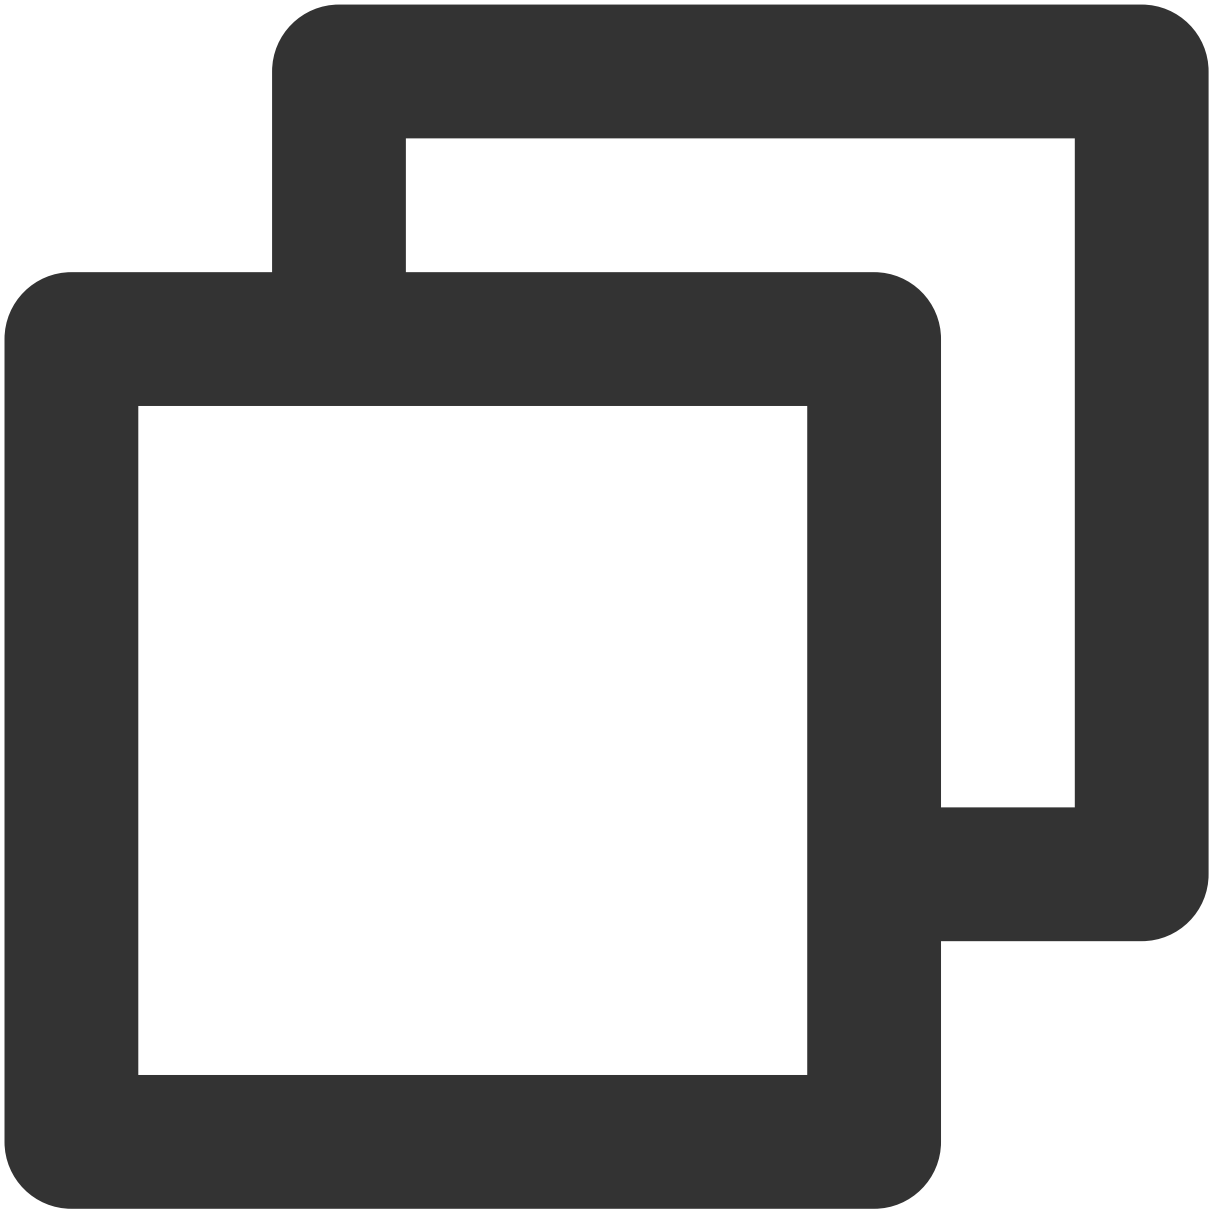

```
/etc/init.d/loglistenerd check
```

The following output indicates that LogListener is running properly.

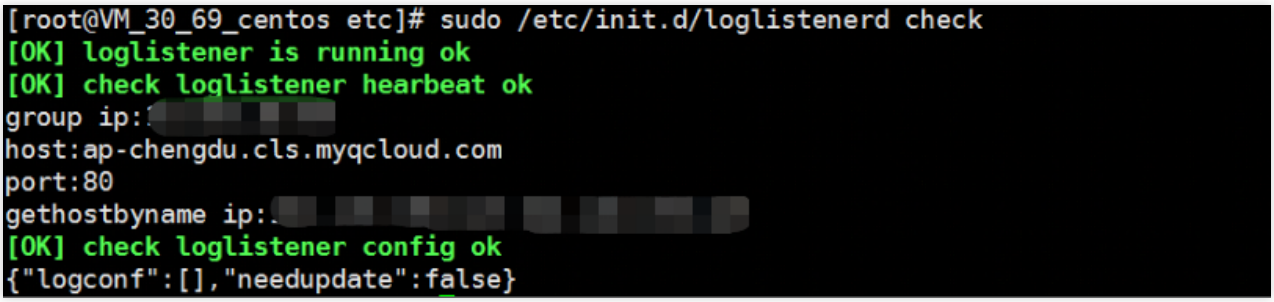

LogListener process exception

If the result returns “[ERROR] loglistener is not running” as shown in the following figure, it indicates that LogListener is not started. Run the `/etc/init.d/loglistenerd start` command to start it. For more information about the operation commands, see Using LogListener.

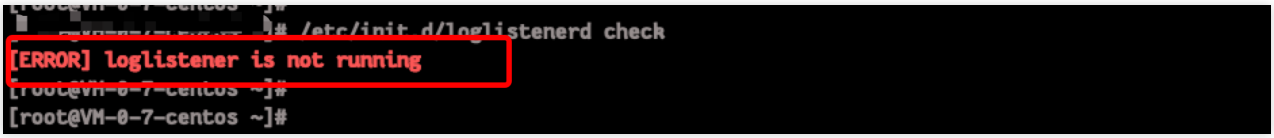

LogListener heartbeat exception

If the result returns “[ERROR] check loglistener heartbeat fail”, it indicates that LogListener has a heartbeat exception.

Many causes can lead to a LogListener heartbeat exception. Possible causes include:

Network error

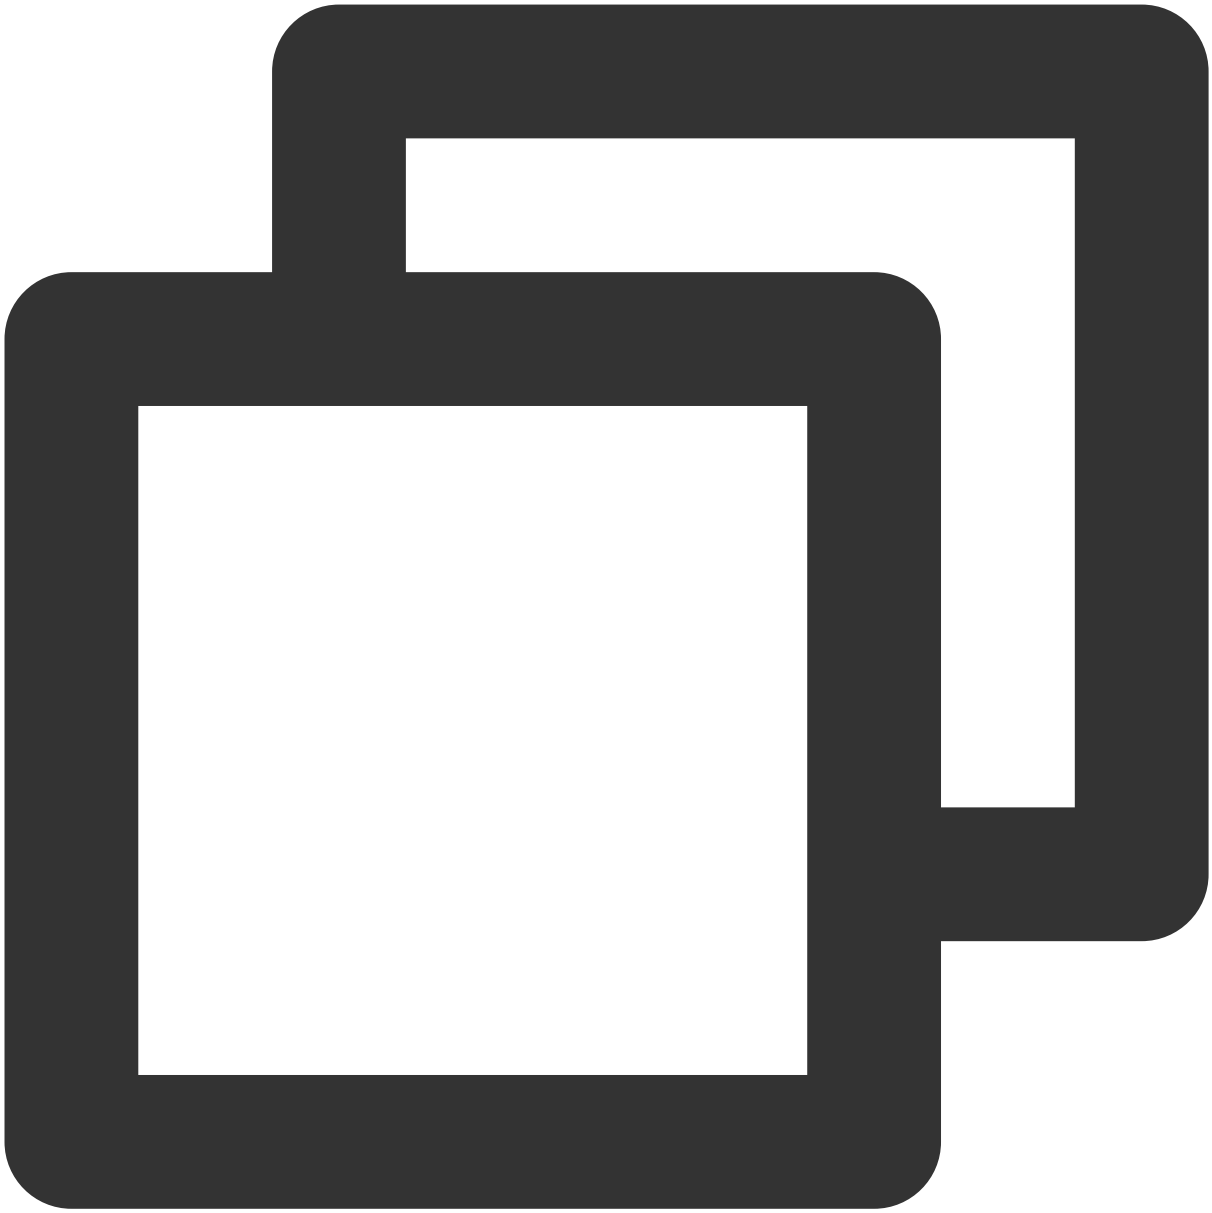

```
telnet <cls domain name> 80
```

Check the network connectivity. For more information about the CLS domain name, see Available Regions.

Incorrect key

To check the LogListener key, access the LogListener installation directory and run the following command.

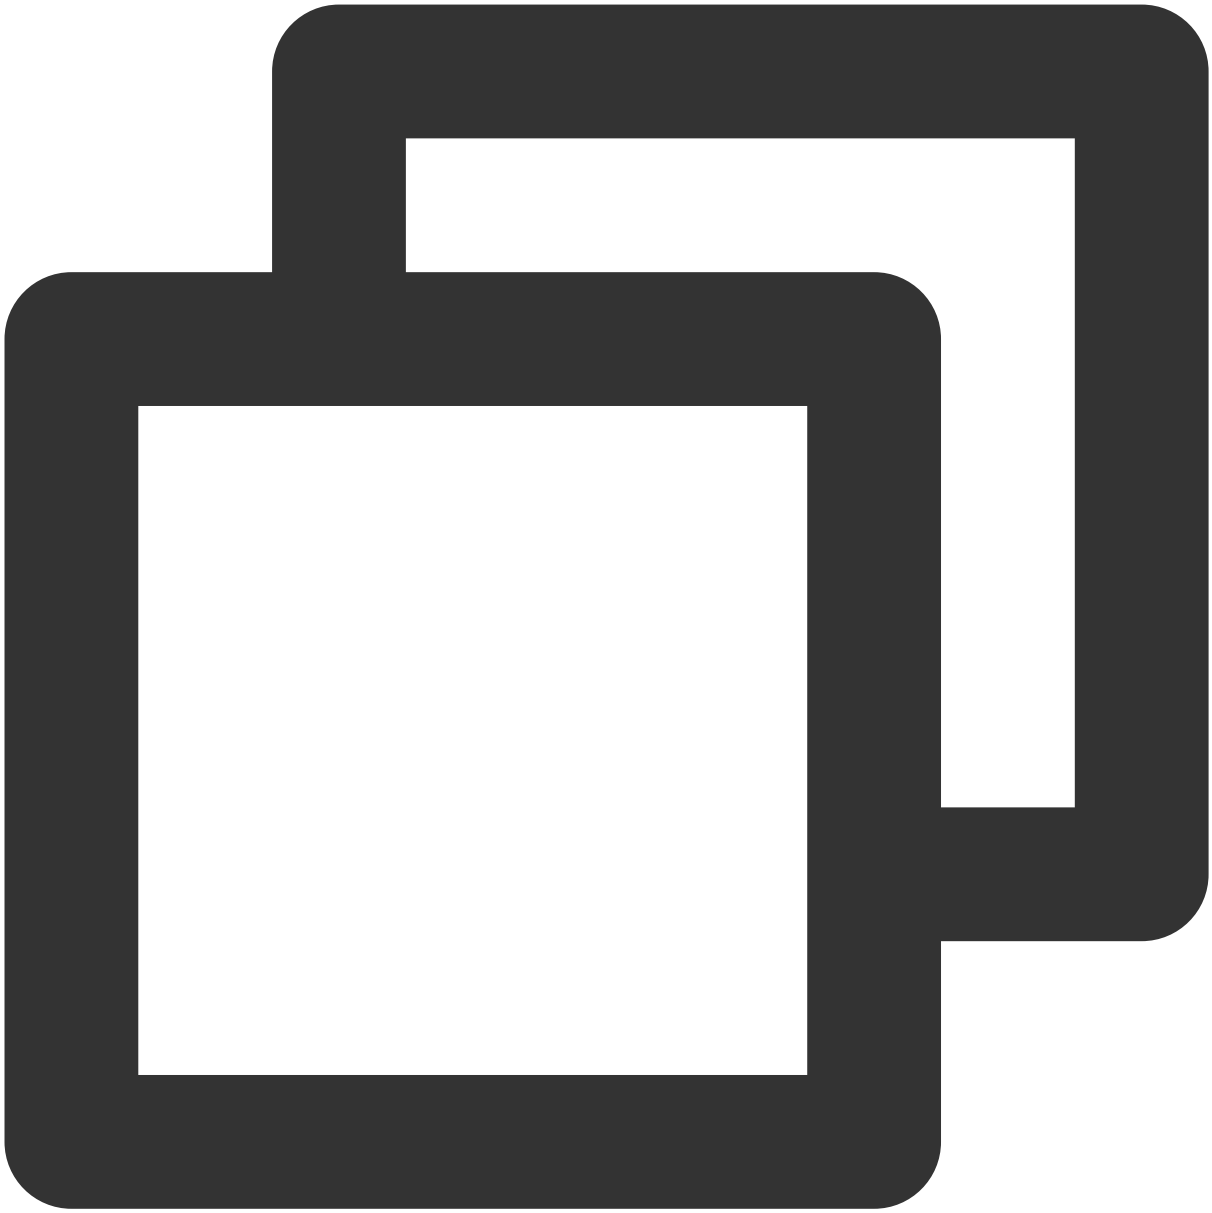

```
grep secret etc/loglistener.conf
```

2. Check for the IP address of the machine group

Check that the IP address added to the machine group is the one configured on LogListener during installation. Run the following command to check the IP address configured on LogListener.

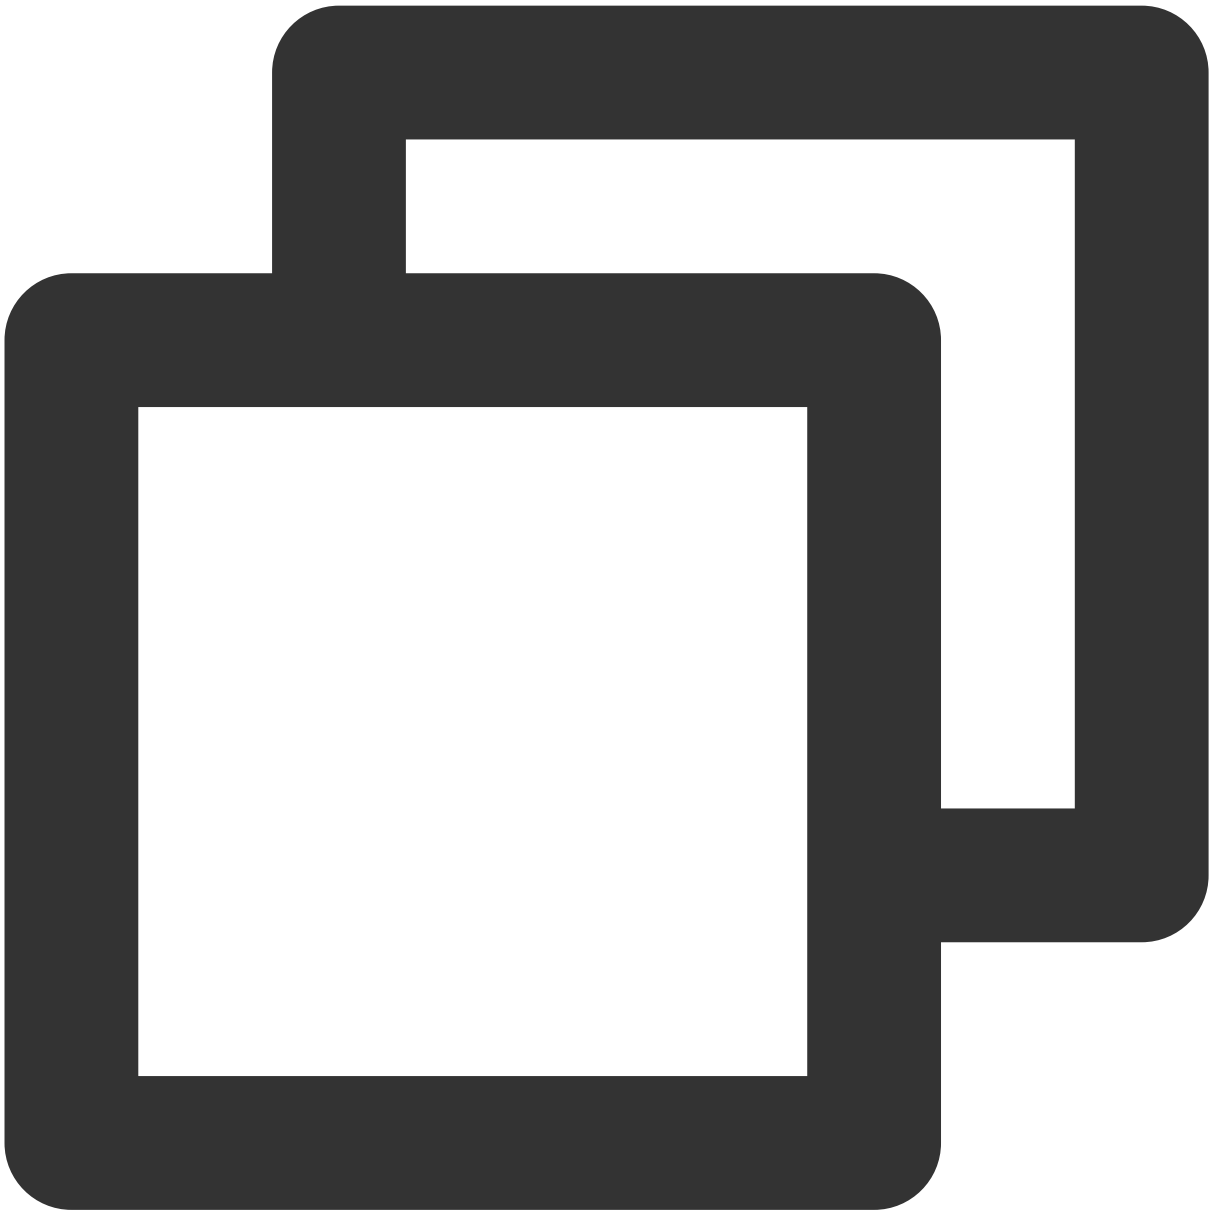

```
grep group_ip etc/loglistener.conf
```

Log in to the CLS console, and select **Server Group** in the left sidebar. On the **Server Group Management** page, view and verify that the IP address of the machine group is the same as that configured on LogListener.

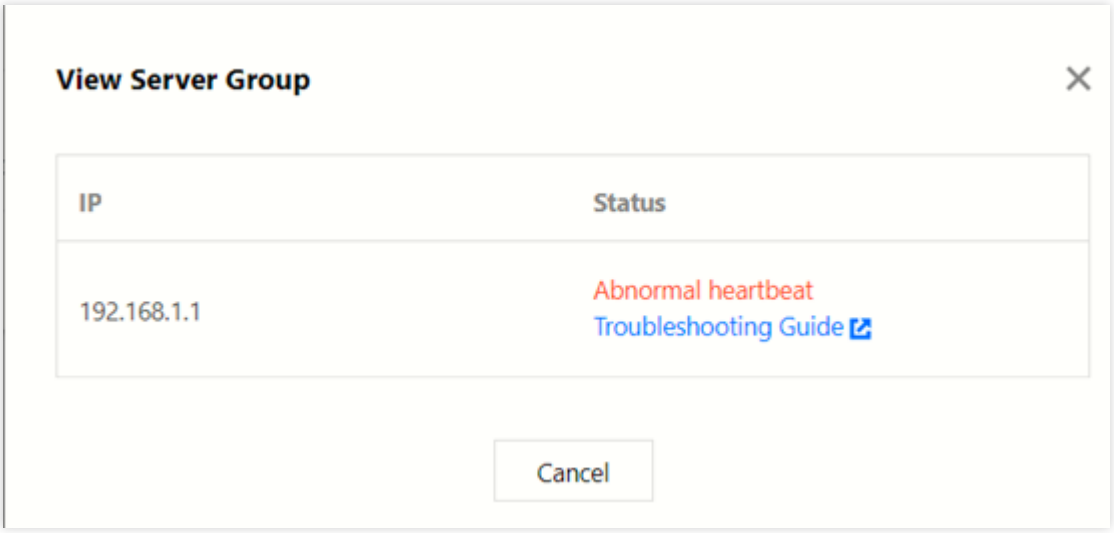

LogListener FAQs

Last updated：2024-01-20 17:11:57

Why cannot a single file be uploaded to multiple topics?

LogListener adopts the policy that a single file can be uploaded to only one topic.
For example, if you have two files `topicA` and `topicB` and you configure the two files as follows:
Set the collection path for `topicA` to `/data/log/\\*\\*/\\*.log`.
Set the collection path for `topicB` to `/data/log/test/\\*\\*/\\*.log`, `/data/log/\\*\\*/test\\*.log`, `/data/log/\\*\\*/\\*.log`, or other collection paths similar to that of `topicA`.
In this scenario, though a file matches two collection paths, the file will be uploaded to only one of the topics.
Therefore, it is recommended that different log topics be used to collect logs of different business types and that you configure collection paths as accurately as possible. If you need to upload a file to different topics, use soft links. You can create different soft links for the same file so that it can be collected to different topics via different soft links.

How do I configure collection paths?

Currently, the required collection path format is: Path prefix + "/**/" + Wildcard filename, for example, `/data/log + /\\*\\*/ + \\*.log ==> /data/log/\\*\\*/\\*.log`.
When setting a wildcard collection path, you need to set the collection path (prefix) as accurately as possible so that LogListener can provide service more efficiently. If the prefix of the collection path is incorrectly set, a large number of paths may be matched by the collection path. As a result, LogListener enters the abnormal status and cannot work.
For example, if the collection path is set to `//\\*\\*/\\*.log`, where the prefix is "/", LogListener will scan the entire root directory and cannot work.

What is the recommended log rotation scheme?

For log rotation, it is recommended that the file name after rotation not be matched by the wildcard collection path.
For example, assume that the collection path is `/var/log/xxxx/\\*\\*/\\*.log` and the log file to be collected is `test.log`. When `test.log` is rotated to `test.2021-07-13.4.log`, LogListener can identify that `test.2021-07-13.4.log` is the rotated file of `test.log` and still label it as `test.log`. Therefore, the checkpoint files stored by LogListener do not include the collection record of the `test.2021-07-13.4.log` file.
However, when LogListener is restarted, LogListener will scan files according to `/var/log/xxxx/\\*\\*/\\*.log` and detect that the `test.2021-07-13.4.log` file matches the matching rule but has no collection record. Then LogListener considers the file a new file and collects it.
Therefore, it is recommended not to match rotated files when matching the wildcard collection path to avoid the case where LogListener collects the rotated files after LogListener restarts.

The recommended log rotation scheme is as follows: if you need to collect `test.log`, you are advised to name the file after rotation to `test.log.2021-07-13.xxx` so that it will not be matched by `\\*.log`.

What should be noted about the LogListener upgrade?

During the iteration of LogListener, the API parameters of the collection path are modified. The collection paths set in versions earlier than v2.2.8 are not supported in the latest version.
Therefore, if you are to upgrade a version earlier than v2.2.8, you need to configure the collection paths again as wildcard paths in the console after LogListener is upgraded.

How do I configure the regular expression collection mode during LogListener collection configuration?

When configuring collection in the console, if you select a collection mode related to regular expressions, although the console provides a tool for extracting regular expression key-value indexes, the tool does not provide automatic generation of regular expressions for Chinese content. If you need to extract regular expressions from Chinese text, you can write your own regular expressions and verify them in the console or using other third-party tools.

What can I do if no log is uploaded when I access CLS through the LogListener for the first time?

The LogListener configuration may be incorrect. The common cases are as follows:
The configured server domain name does not match. As a result, LogListener cannot obtain the collection configuration of the current region, and no collection service is running.
LogListener is added to the IP machine group, but LogListener is configured with label information. As a result, the collection configuration cannot be pulled from the current region, and no collection service is running.
The secret ID or key configured in LogListener is incorrect, or the permission is insufficient. As a result, logs cannot be uploaded.
Environment problem (for example, the public network access is not enabled in the VPC subnet). If cross-region upload is configured, the configuration does not take effect. In fact, LogListener still communicates with the local server.
Generally, in this case, you can log in to LogListener, go to the LogListener installation directory, and run the `./bin/check` command to check the following information:
Whether the domain name is correct.
Whether heartbeats are reported properly.
Whether collection configuration is pulled properly.

Log collection failed due to mixed usage of machine groups. What should I do?

At present, machine groups are divided into two categories, and their usage methods are independent of each other:
IP machine group: a machine IP must be manually added to the machine group in the console, and the `group_label` field in `loglistener.conf` on the corresponding machine must be empty.

Label machine group: the machine group label is set in the console, and the `group_label` field in `loglistener.conf` on the corresponding machine must be set to the same label.
The above two usage methods are incompatible. If they are used together, LogListener will not be able to pull the correct collection configuration, resulting in no collection.

In what situations will LogListener collect logs?

When a group of files queue for LogListener collection, the first file in the queue is collected first, and only when the end of the first file is read at a certain time, the position of the first file is given up. That is, not all files can enjoy the collection resources equally in a unit of time.
If the writing speed of a single file is always greater than the collection speed, and the latest position of the file cannot be consumed due to the slow collection speed, the file occupies collection resources for a long time, and as a result other files cannot be collected.

What should I do in the case of topic collection blocking?

In a certain period of time, if the generation speed of a single file is greater than the collection speed, LogListener will continue to collect this file, and the collection of other files will be blocked.

What are the rules for filters?

The filter rule is collection after matching, not discarding after matching. LogListener does not collect logs that do not match.

How do I use non-root permission to start LogListener?

You are advised to use the root permission to start LogListener. If you need to use LogListener with non-root permission, see "Configuring Non-Root Permission to start LogListener".

How do I pin the LogListener process to a CPU?

For the CPU pinning, use the taskset tool and run the `taskset -cp ${cpu number} ${pid>}` command.

How do I control the high memory and resource usage of LogListener?

We recommend that you upgrade LogListener to the latest version and set `memory_tight_mode = true`.
Use CGroup to control CPU and MEM usage.

Does LogListener support log collection via a soft link?

Yes. But LogListener earlier than version 2.3.0 does not collect those log files in soft links, or in shared file directories of NFS, CIFS, etc.

Can LogListener upload data to multiple log topics?

Yes, provided that these log topics are in the same region.

A log file will only be collected into one log topic.

Are machines automatically added to a machine group when LogListener is initialized?

Yes, provided that you configure the machine group by machine ID. For more information, please see Machine Group Management.

In what situations will LogListener upload logs?

More than 4 MB logs are cached.
More than 10,000 logs are cached.
LogListener finishes reading a file.

What does the maximum performance of LogListener mean?

Collecting logs with full text in a single line: 115 MB/sec.
Collecting logs with full text in multi lines: 40 MB/sec.
Collecting JSON logs: 25 MB/sec.
Collecting CSV logs: 50 MB/sec.
Collecting full RegEx logs: 18 Mb/sec, depending on the regex complexity.

How do I modify the LogListener configuration after the server IP address is changed?

If you configure the machine group by machine ID, you don't need to modify the LogListener configuration. This method is recommended when the server IP frequently changes. For more information, see Configuring the machine group by machine ID.
If you configure the machine group by IP address, modify the configuration as follows:
a. Add the new IP address to the `group_ip` field in the configuration file.

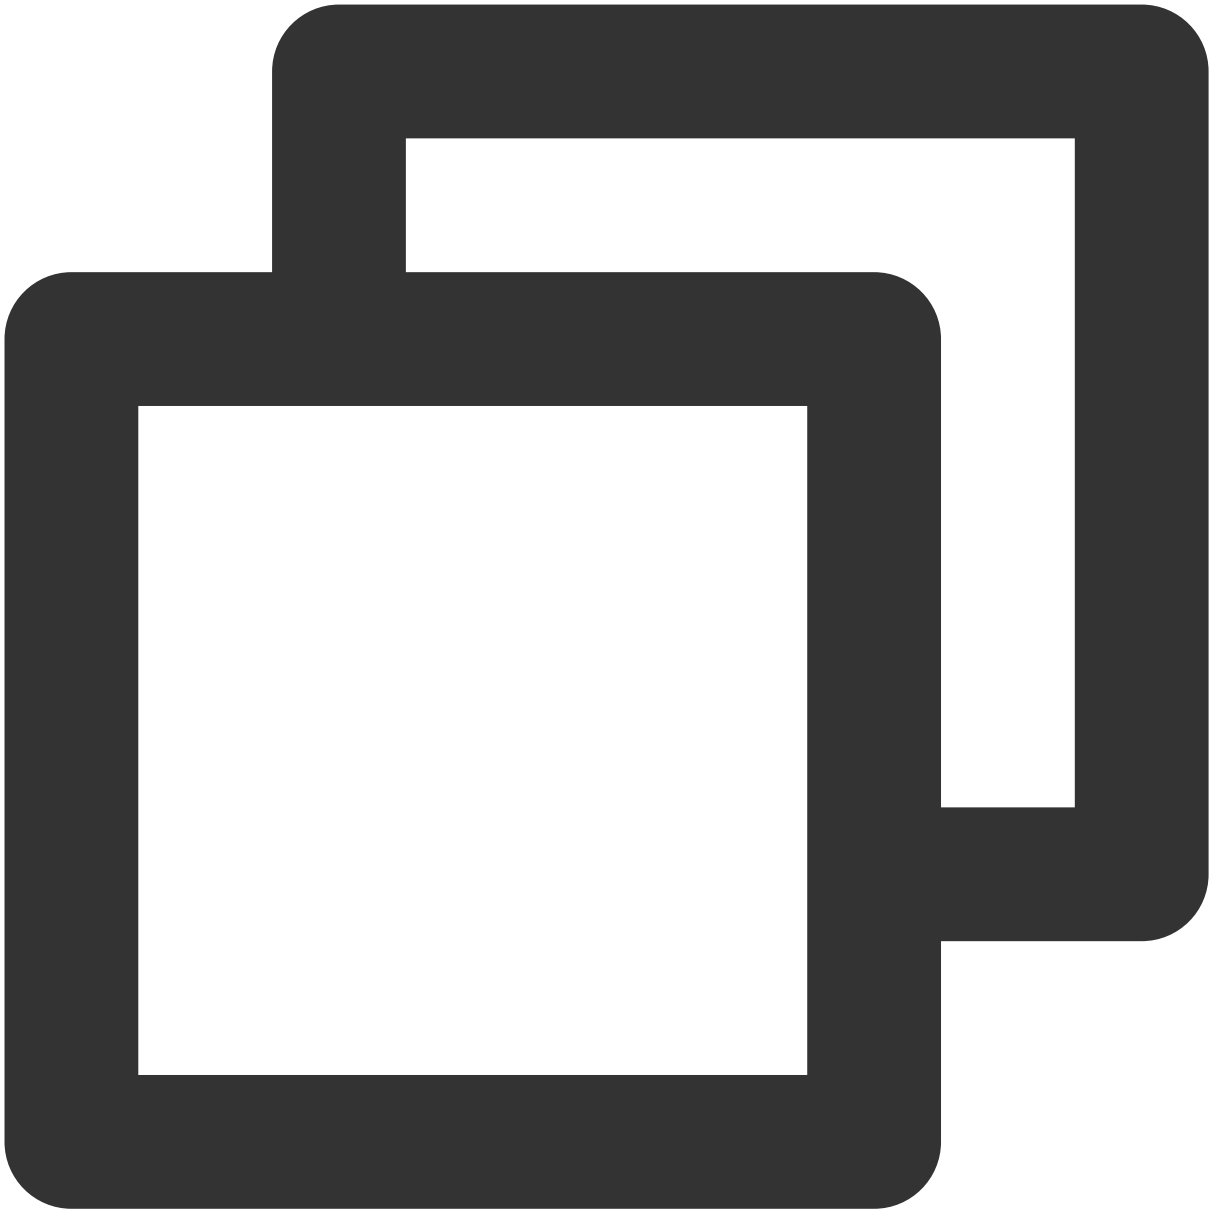

```
sed -i '' "s/group_ip *=.*/group_ip = ${group_ip}/" etc/loglistener.conf
```

b. Restart LogListener.

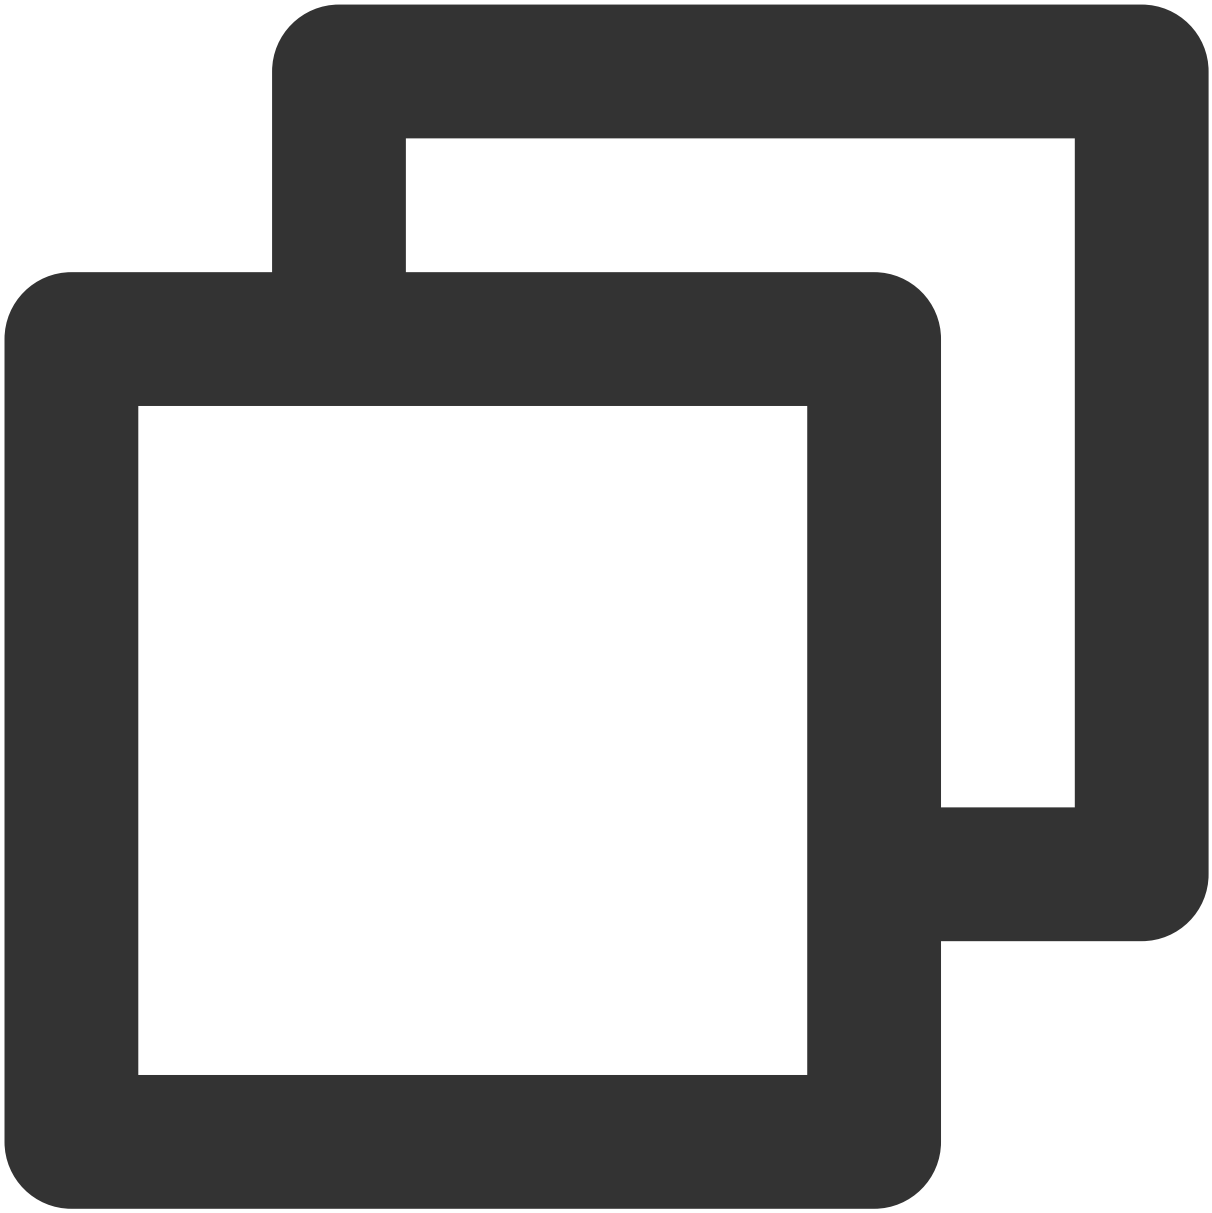

```
/etc/init.d/loglistenerd restart
```

c. Log in to the CLS console and select **Machine Group Management** on the left sidebar. Locate the machine group to which the server binds and click **Edit**. In the pop-up window, replace the old IP address with the new one, and click **OK**.

LogListener Installation Exception

Last updated：2024-01-20 17:11:57

For details about how to install and use LogListener, see LogListener Installation Guide.

Possible causes

Loglistener may not be installed correctly for the following reasons:

1. The kernel version only supports 64-bit.
2. The installation method is incorrect.
3. The latest features rely on a later version of LogListener.

Directions

1. Check the kernel version.

The executable file in the bin directory under the LogListener installation directory only supports Linux 64-bit kernel. Execute the command **uname -a** to check whether the kernel version is x86_64.

2. Check the installation command.

Be sure to perform operations according to the LogListener Installation Guide.

3. Check the LogListener version.

Some of new CLS features may be available only for the latest version of Loglistener. In this case, please download and install the latest version. For step-by-step directions, see LogListener Installation Guide.

4. Verify the LogListener installation.

Check for process and heartbeat of LogListener and check whether it can properly obtain collection configuration of users. To do this, please see LogListener Diagnostic Tool.

Container Log Collection

Last updated：2024-01-20 17:11:57

Installation and Upgrade

How do I deploy the log collection component in a TKE cluster?

1. Log in to the TKE console.
2. On the left sidebar, click **Ops Feature Management** to enter the feature management page.
3. Find the target cluster and click **Settings**.
4. In the pop-up window, click **Edit** in the **Log Collection** column.
5. Select **Enable Log Collection** and click **OK**.
6. Click **Close**.

How do I upgrade the log collection component in a TKE cluster?

1. Log in to the TKE console.
2. On the left sidebar, click **Ops Feature Management** to enter the feature management page.
3. Find the target cluster and click **Settings**.
4. In the pop-up window, click **Edit** in the **Log Collection** column.
5. Click **Upgrade Component**.

Network and Permission

What should I do if the TencentCloud API domain name is inaccessible?

cls-provisioner, the component for communication between the TKE log collection component and CLS, uses a TencentCloud API domain, which must be kept accessible. If the component deployment fails, or you find the following error in logs, the domain name is inaccessible.

As shown above, a cls-provisioner start exception occurred, and you can find in logs that the `cls.internal.tencentcloudapi.com` domain name is inaccessible.

The private and public domain names are accessible by default from servers in Tencent Cloud. A common cause of this problem is that the DNS configuration on the TKE node is modified. You can fix it in the following two ways:

Add the default Tencent Cloud DNS configuration to the DNS configuration on the TKE node server.
If the DNS of the host is the server CoreDNS, add Tencent Cloud DNS configure to CoreDNS.

Note:

We recommend you check the DNS configuration on the relevant TKE nodes first when the domain name is inaccessible in the TKE cluster.

What should I do if the CLS log upload domain name is inaccessible?

Log upload domain names are different from TencentCloud API domain names and are in the format of `<region>.cls.tencentcs.com` (public network domain) or `<region>.cls.tencentcs.com` (private network domain). For more information, see Available Regions.

Solution:

Make the domain name accessible on the corresponding cluster node server.

What should I do if a no permission error was reported during communication between cls-provisioner and CLS?

Sometimes an error similar to the following may be reported during communication between cls-provisioner and CLS:

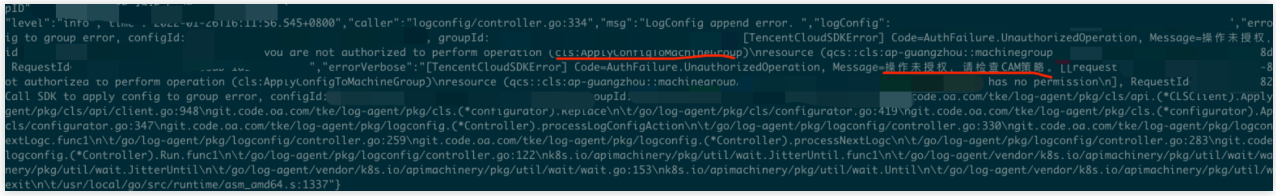

Solution:

Associate the `QcloudAccessForTKERoleInOpsManagement` policy in the TKE_QCSRole role under the account that created the TKE cluster.

Collection

What should I do if logs collected to CLS are truncated?

In some cases, the output type of user logs is standard output, but logs collected to CLS are truncated. This happens as json-tool, the default log collection tool of Docker, limits the size of single-line logs. Therefore, logs exceeding 16 KB in size will be truncated.

Solution:

Modify the log output configuration to make the size of printed single-line logs below 16 KB.

What should I do if the same logs are collected repeatedly?

If you find that some logs are collected repeatedly in the CLS console, you can check the log output path first to see whether logs are output to the persistent storage created by PV/PVC.
If logs are output to the persistent storage, when a business Pod is recreated, logs will be collected again. You can run the following command to view the YAML definition of the Pod:

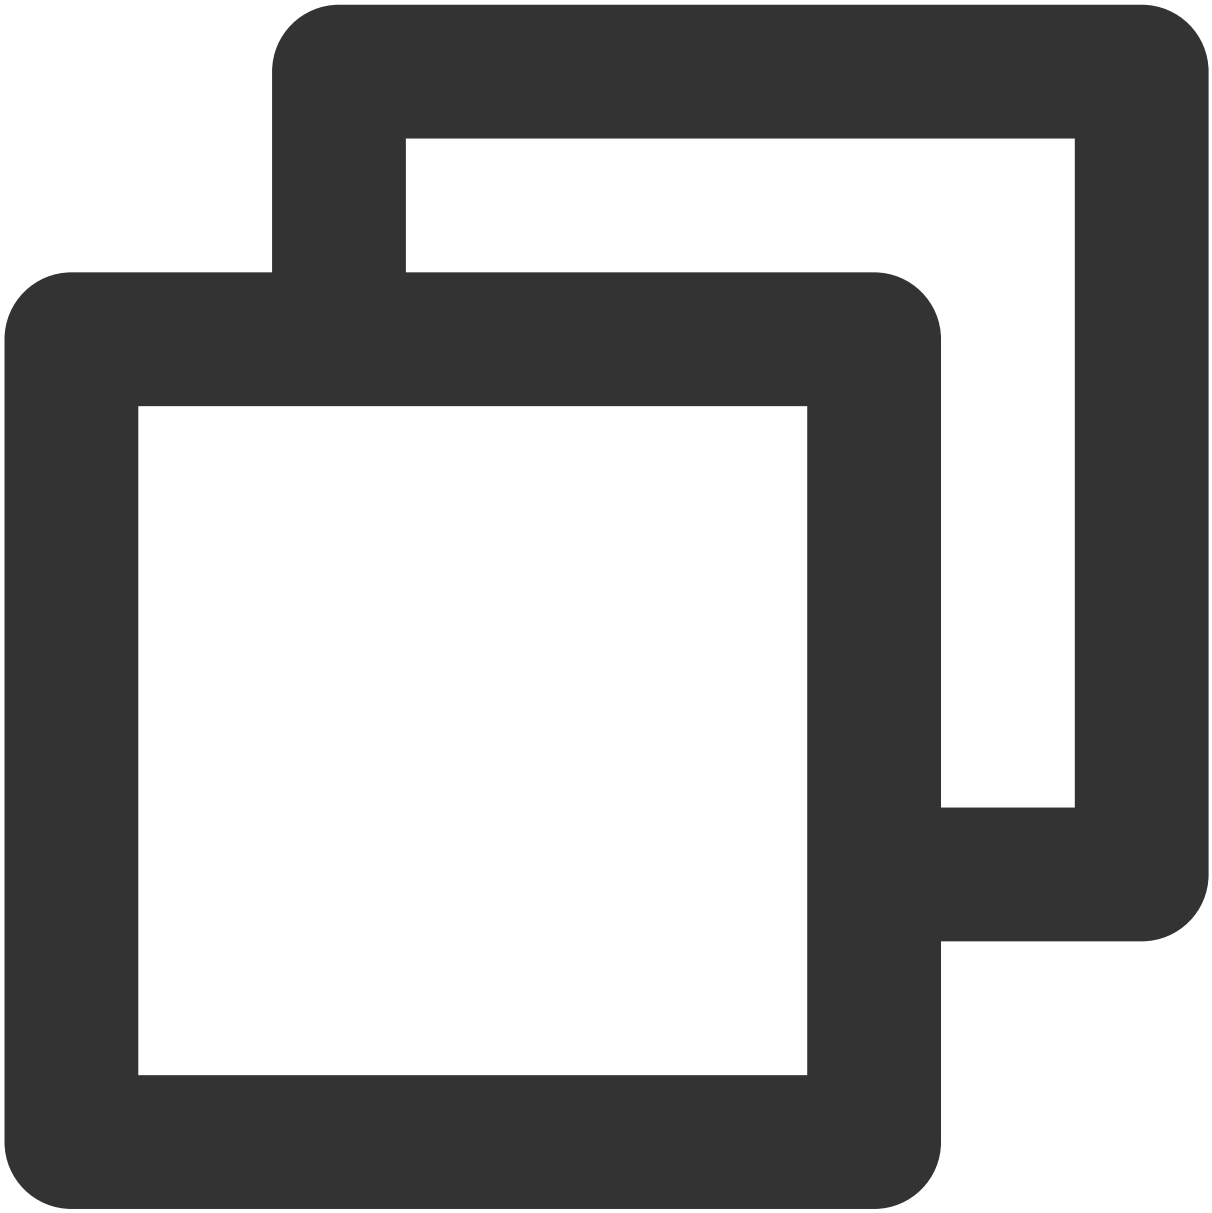

```
kubectl get pods <pod_name> -n <namespace> -o yaml | less
```

If information similar to the following is returned, logs are output to the persistent storage.
The business uses CFS, and CFS is mounted to the container.

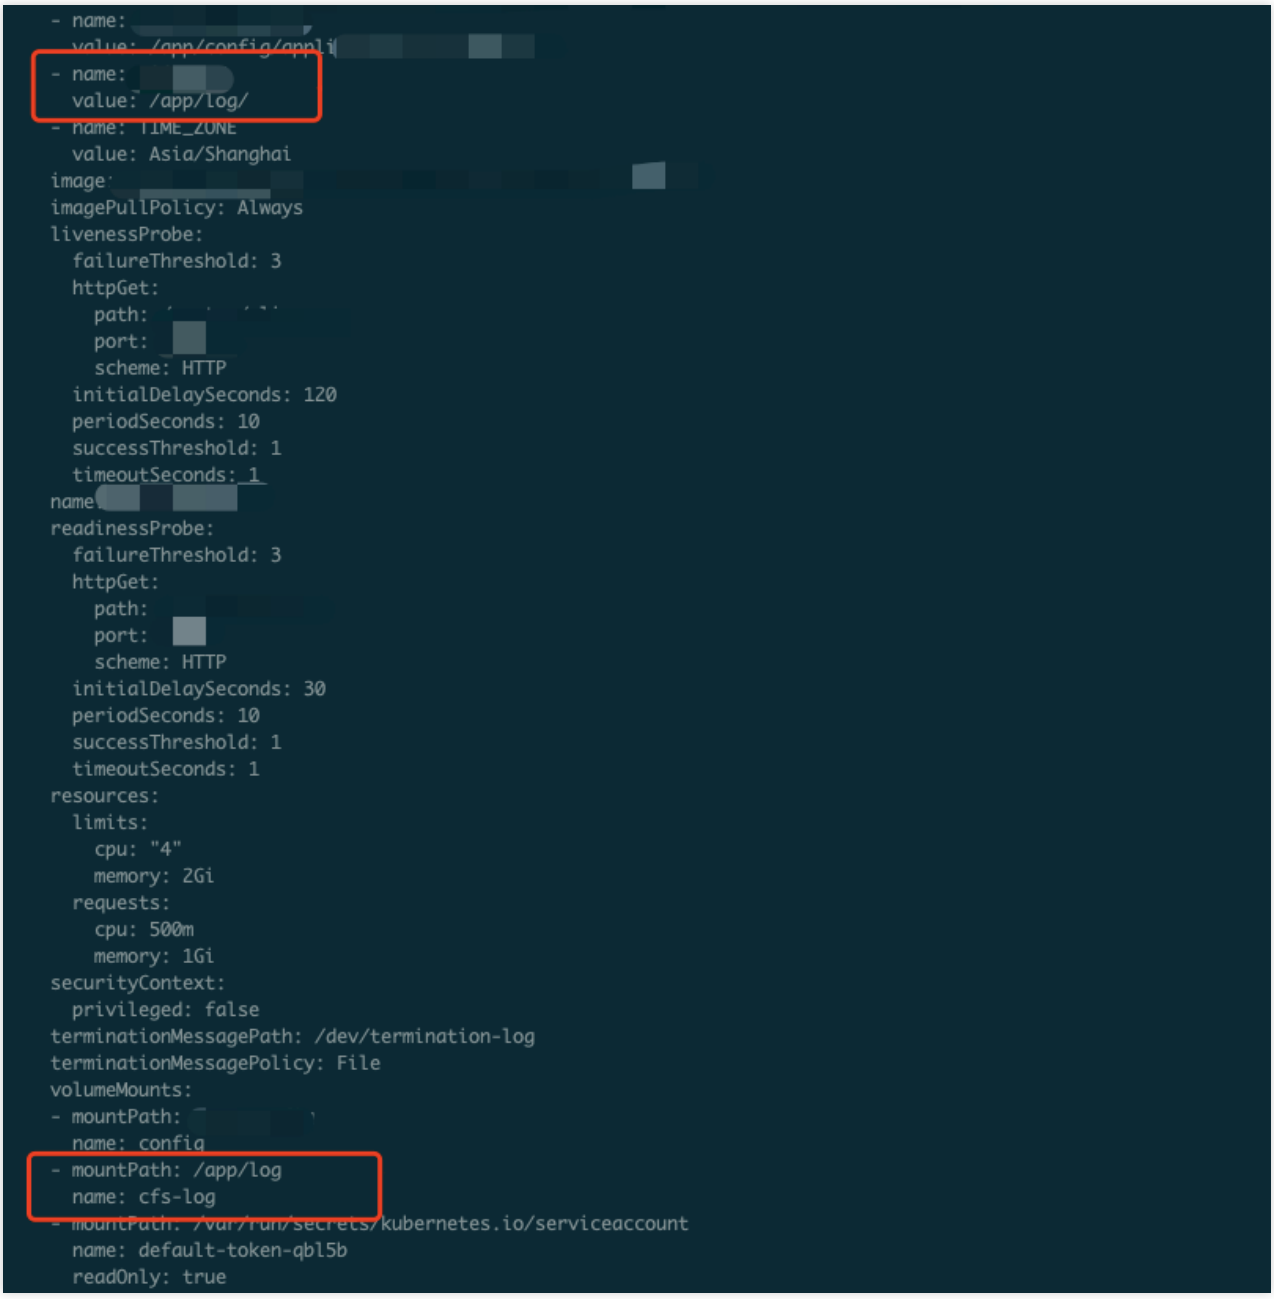

CFS is used.

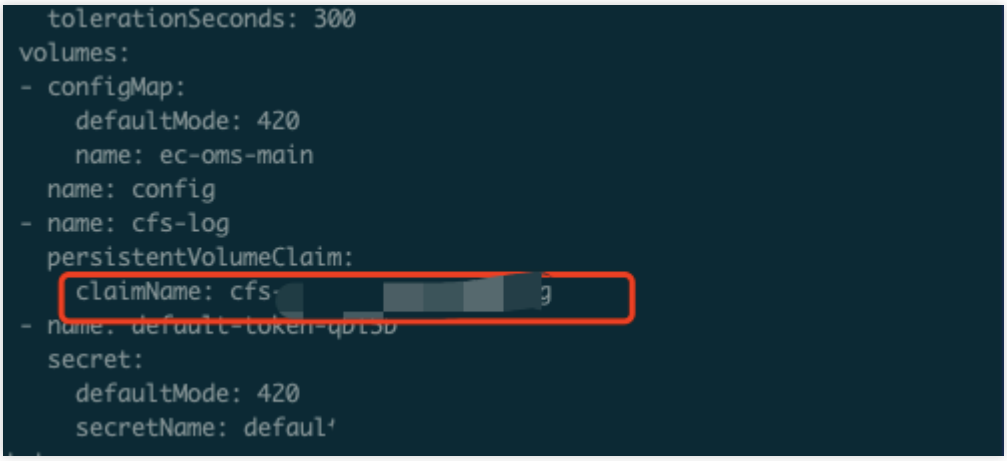

Solution:

If logs don't need to be persistently stored, you can enable log collection in the container cluster to collect logs to CLS.

If logs need to be persistently stored, you can modify the collection policy to **Incremental** collection when configuring the LogListener rule in the CLS console; however, incremental collection cannot ensure that all logs will be collected.

What should I do if some logs are not collected?

LogListener currently doesn't support logs stored in NFS. It subscribes to Linux kernel events to get the file update information instead of actively scanning target files.

As NSF file update information is generated on the NFS server, and file update events cannot be generated in the local kernel, such information cannot be perceived by LogListener. Therefore, NFS file logs cannot be collected in real time.

What should I do if the collection configuration doesn't match the Pod?

You can troubleshoot as follows:

1. Run the following command to check whether Pod labels match the collection configuration:

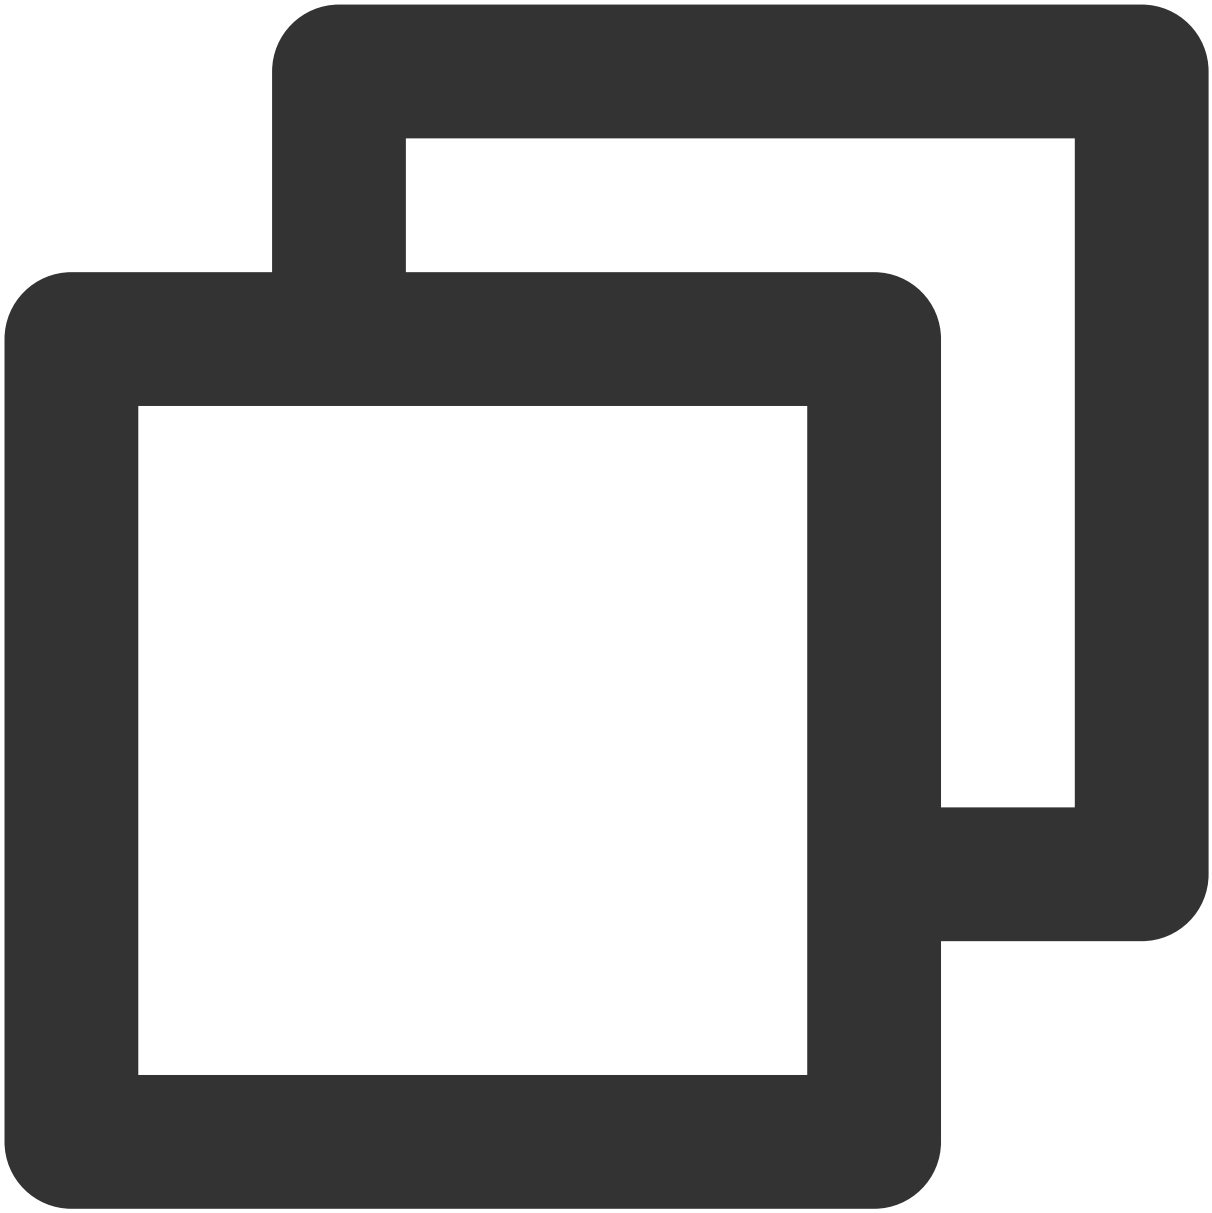

```
kubectl get pods <pod_name> -n <namespace> --show-labels
```

If so, proceed to the next step.

If not, modify the collection configuration according to the correct content.

2. Check whether the Pod workload (Deployment or StatefulSets) match the collection configuration.

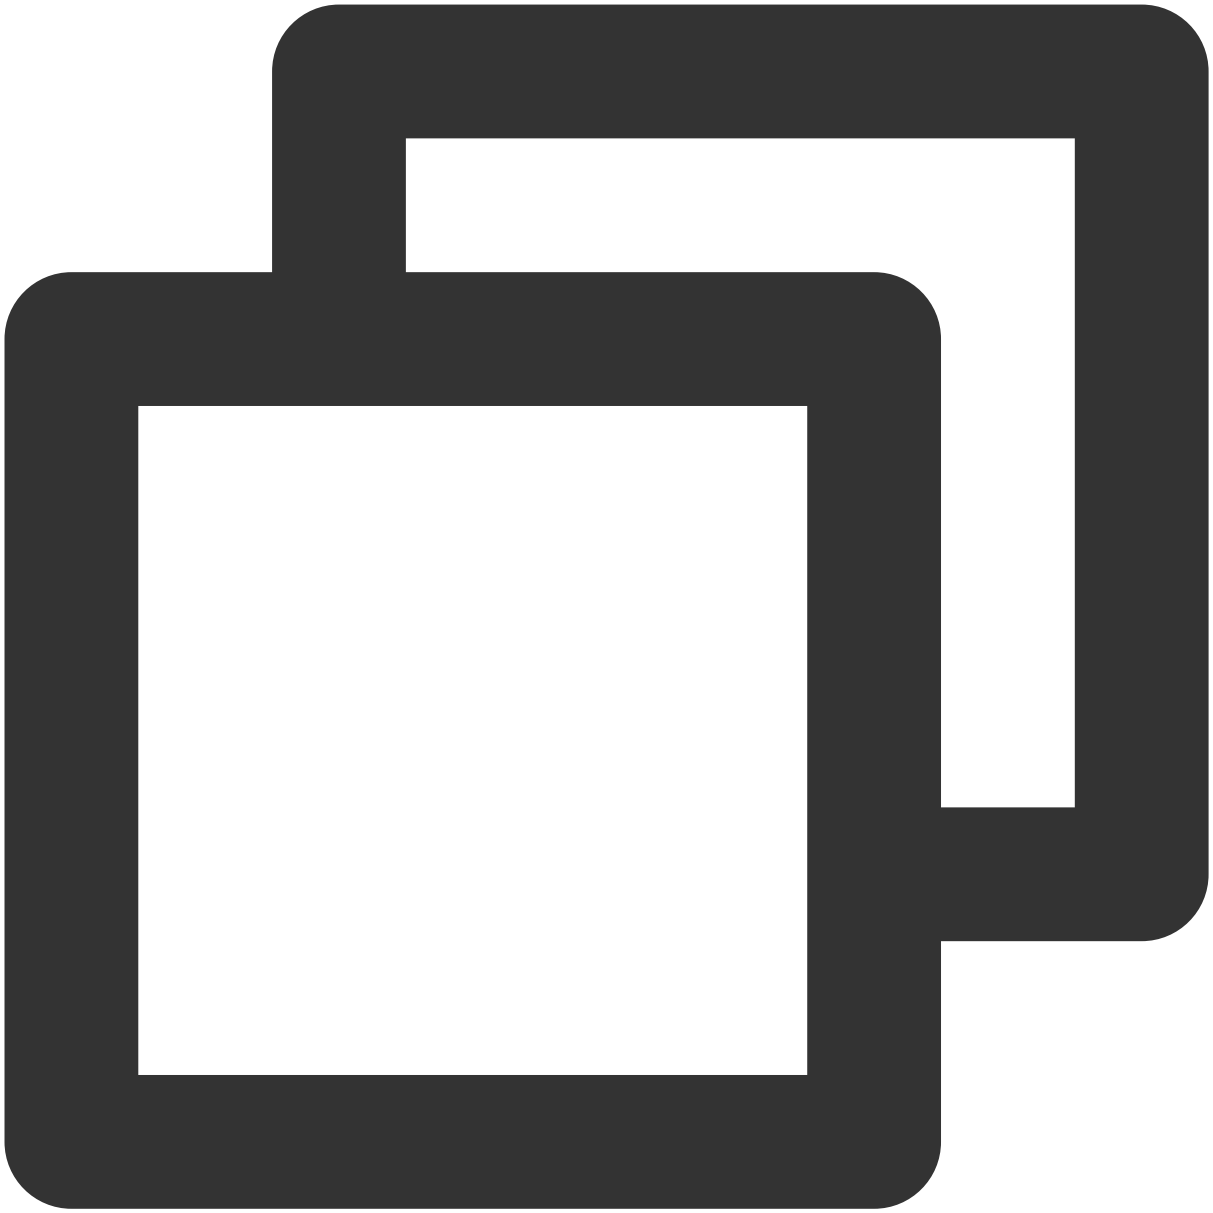

```
kubectl get pods -n <namespace> |grep testa
```

If so, proceed to the next step.

If not, modify the collection configuration according to the correct content.

3. Use the following command to view the YAML definition of the Pod and check whether the container name matches that specified in the collection configuration.

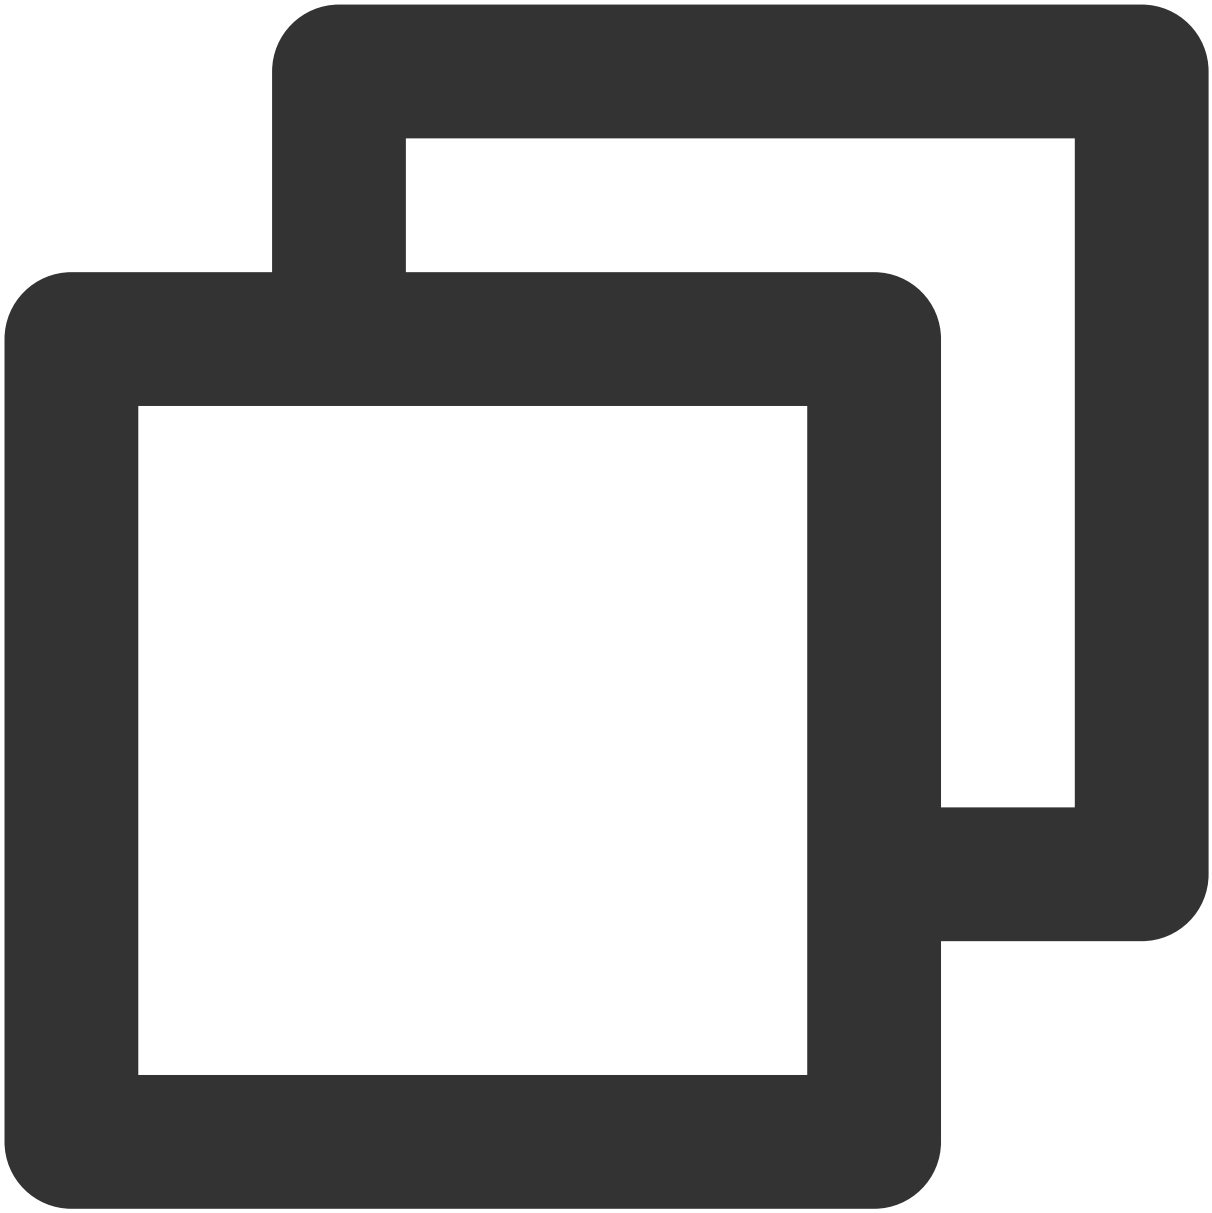

```
kubectl get pods <pod_name> -n <namespace> -o yaml
```

If so, the task is completed.

If not, modify the collection configuration according to the correct content.

What should I do if the collection path is incorrect?

When collecting container files or host files, check whether the collection directory path is correct and contains logs complying with the collection rule.

Can I use soft links for log files?

In container file collection scenarios, matched log files cannot have soft links.

In Kubernetes scenarios, **CLS collects logs by parsing the location of the container file on the host. As a container soft link points to a path within the container, if a matched file to be collected has a soft link, it cannot be reached correctly**.

Solution:

Modify the collection rule path and matched files to use the actual log file path and log files to be collected, so as to avoid matching soft links.

What are the limits for collection scale?

As resources are restricted when LogListener is used to collect container logs, the numbers of directories and files listened on are also limited as follows:

Directories listened on: 5,000

Files listened on: 10,000

You may encounter such problems when collecting container or host files. Generally, the following information will be displayed in LogListener logs if expired log files are not cleared:

Files and directories exceeding the limits will not be listened on by LogListener; therefore, some target log files may not be collected.

For more information, see LogListener Limits.

Solution:

In the log directory, run the `tree` command to check whether the current numbers of directories and files in the entire directory structure reach the LogListener limits.

If not, run `tree -L 5` under `/var/log/tke-log-agent` on the host of the business container to check whether the host limits are reached

As LogListener limits take effect for the entire host, if the number of files listened on in a container doesn't reach the threshold, the reason may be that the number of files in all containers on the host has reached the limit.

If so, archive expired logs in time to reduce the resources consumed by directories and files listened on by LogListener.

What should I do if a volume is defined in Dockerfile?

In Docker scenario, run the `docker history $image` command to view the image rebuild information.

In Containerd scenarios, run the `crictl inspecti $image` command to view the image rebuild information. The following information is returned. You can see that the `/logs/live-srv` volume is customized in the Dockerfile, which also happens to be the log directory. Such an operation prevents the log collection component from finding correct log files.

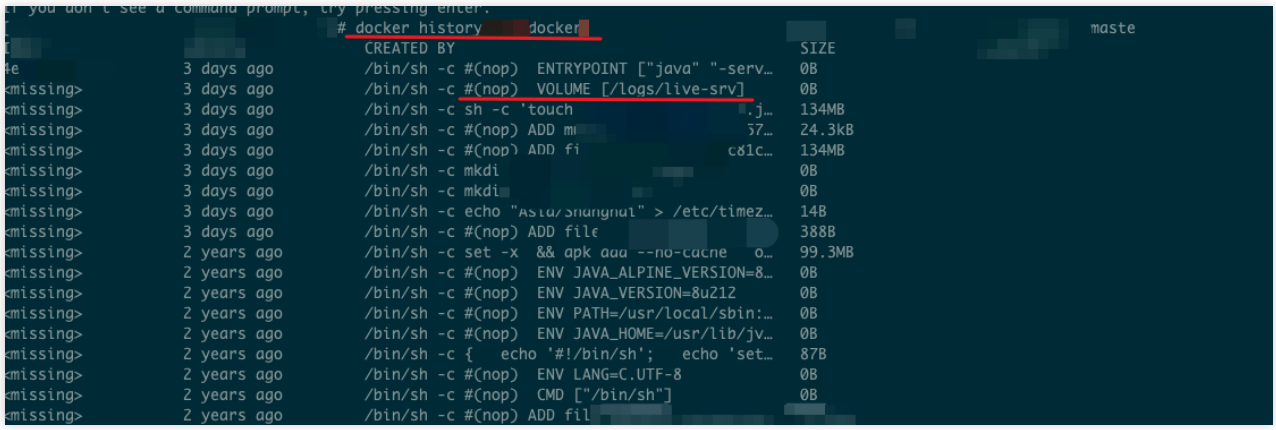

Solution:

Modify the Dockerfile to remove the volume, rebuild the image, and redeploy the service.

Modify the directory to which service logs are written, so that logs are not written to the volume path defined in the Dockerfile.

Others

What should I do if the identified container engine type is incorrect?

In some Docker scenarios, bugs on earlier versions may be triggered, causing a log collection component start failure and generating panic logs.

This happens mainly because the Docker configuration of TKE cluster nodes is customized, leading to the error as shown below:

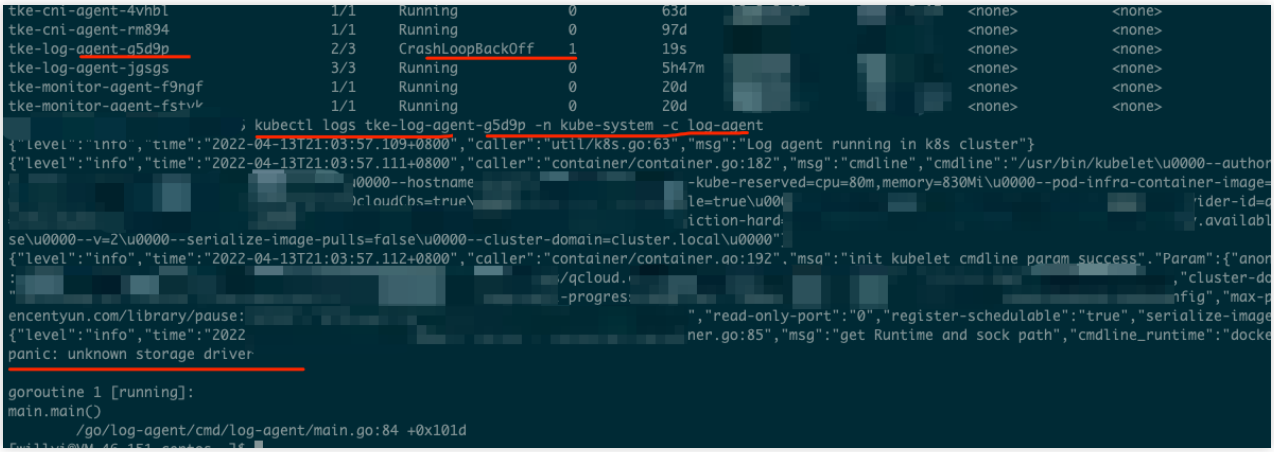

Solution:

Add `"storage-driver": "overlay2"` to the `/etc/docker/daemon.json` configuration file as shown below:

```
[root@VM-0-5-centos ~]# cat /etc/docker/daemon.json
{
  "bridge": "none",
  "debug": false,
  "default-runtime": "runc",
  "exec-opts": [],
  "exec-root": "",
  "graph": "/var/lib/docker",
  "group": "",
  "insecure-registries": [],
  "ip-forward": true,
  "ip-masq": false,
  "iptables": false,
  "ipv6": false,
  "labels": [],
  "live-restore": true,
  "log-driver": "json-file",
  "log-level": "warn",
  "log-opts": {
    "max-file": "10",
    "max-size": "100m"
  },
  "max-concurrent-downloads": 10,
  "registry-mirrors": [
    "https://mirror.ccs.tencentyun.com"
  ],
  "runtimes": {},
  "selinux-enabled": false,
  "storage-driver": "overlay2",
  "storage-opts": [
    "overlay2.override_kernel_check=true"
  ]
}
[root@VM-0-5-centos ~]#
```

Upgrade the log collection component version in the TKE console. As this problem has been fixed on new versions of the component, you don't need to modify the Docker configuration.

What should I do if a subdirectory is set in `filePattern`?

As shown below, a subdirectory is set in the `filePattern` parameter, making it unable to collect logs.

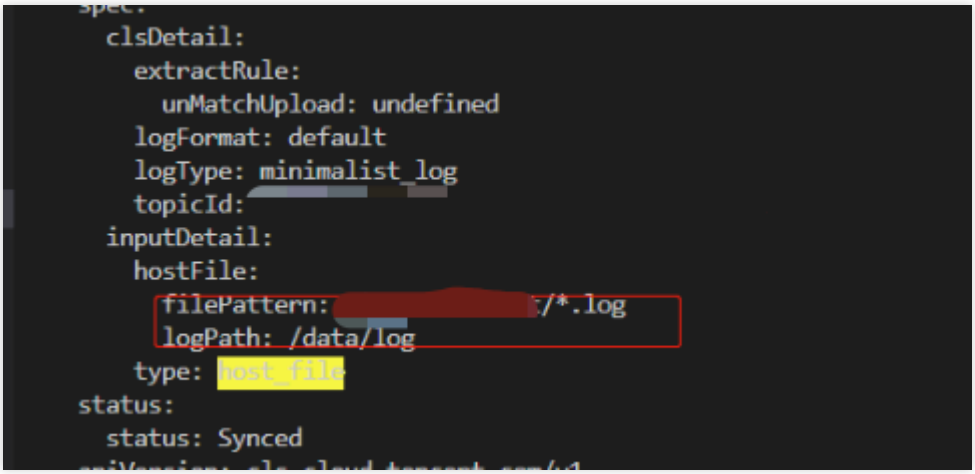

Solution:

Set the log file directory in the `logPath` parameter, and set only the file type parameter in `filePattern` .

Self-Built Kubernetes Log Collection Troubleshooting Guide

Last updated：2024-01-20 17:11:57

After installing and deploying LogListener in the self-built Kubernetes cluster, you can configure collection by creating a LogConfig object or in the console to start log collection.
In case of any log collection exceptions, perform troubleshooting as follows.

1. Check the LogConfig status

View all collection configurations of the cluster: kubectl get logconfig
View a specific collection configuration: kubectl get logconfig xxx -o yaml
View the LogConfig sync status. If the status is not `Synced` , an exception occurred, and you can view the exception message in `reason` . In normal cases, the status is `success` .
If the LogConfig status is `Synced` as shown below, the collection exception is due to another issue:

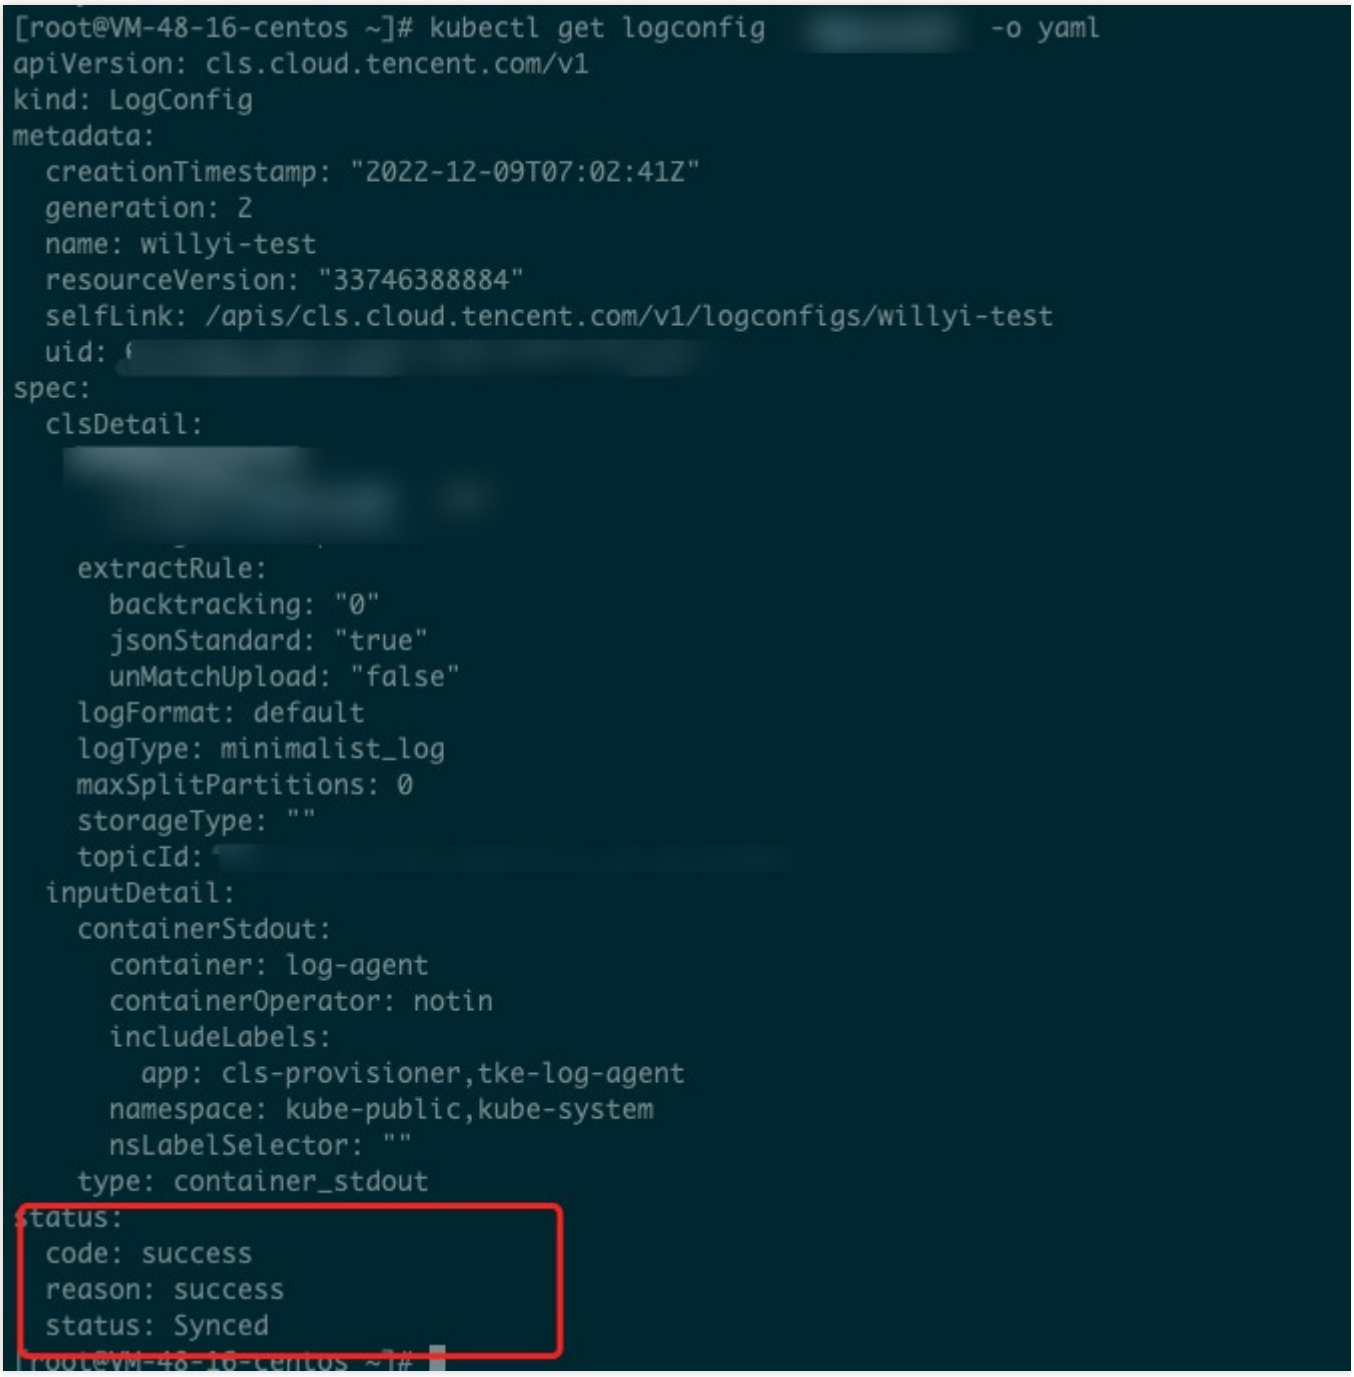

To further identify the sync error cause, check the `cls-provisioner` log.

2. View the cls-provisioner log

Identify the Pod of `cls-provisioner` : kubectl get pods -n kube-system -o wide |grep cls-provisioner

View the log: kubectl logs cls-provisioner-xxx -n kube-system

As shown below:

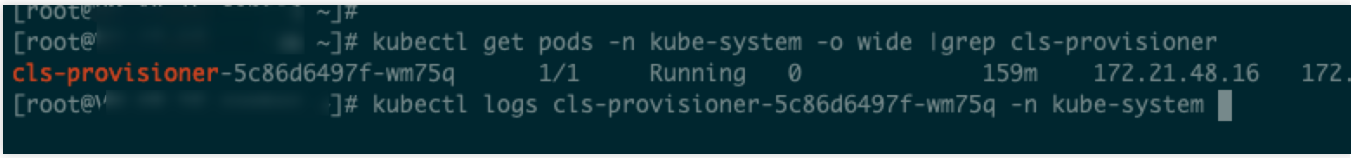

View the `cls-provisioner` log to check the sync error cause.

Note:

The `cls-provisioner` component communicates with the CLS server. Specifically, it converts and syncs the LogConfig collection configuration to the CLS server. In this way, the collector can get the collection configuration from the server for normal log collection.

3. View the collector log

If the log collection exception persists when the collection configuration is synced properly, check the collector log.

Check whether the soft link is created successfully.

Taking the collection of standard output as an example:

The soft link of the standard output log of the Pod to be collected will be created in `/var/log/tke-log-agent/<Collection configuration name (LogConfig name)>/stdout-docker-json`, after which logs can be collected properly.

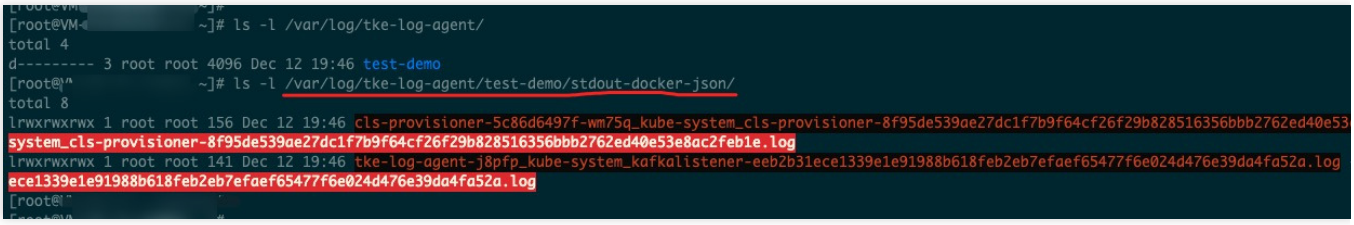

In the case of Docker, if the runtime is containerd, the path will be `/var/log/tke-log-agent/<Collection configuration name (LogConfig name)>/stdout-containerd`.

The soft link of the containerd file to be collected is created as follows:

/var/log/tke-log-agent/<Collection configuration name (LogConfig name)>/

As shown below:

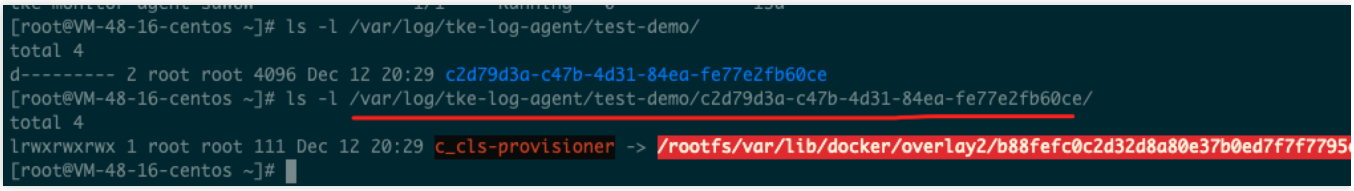

Check whether the soft link created as instructed above is OK. If no soft link is created, an exception occurred. If the creation is successful, further check the log of LogListener.

View the log of LogListener.

kubectl get pods -n kube-system -o wide |grep tke-log-agent

First, find the Pod of `tke-log-agent` on the host of the Pod of the log collection exception and view the LogListener log.

kubectl logs tke-log-agent-xxx -n kube-system -c loglistener

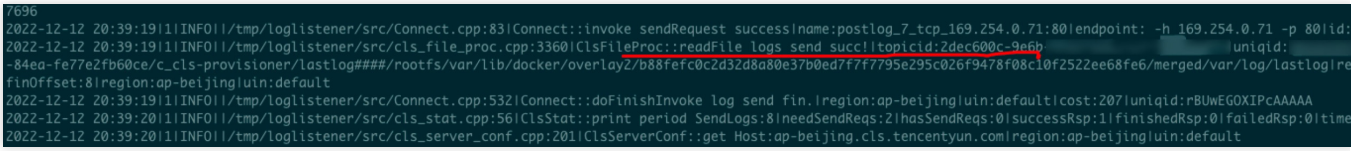

Check whether **readFile logs send succ!|topicid** or a similar description exists as shown above, and if so, the log is successfully collected to the target topic; if not, the collection is abnormal, and you can contact us for assistance.

If the log has been collected to the topic but cannot be found, check whether full-text index is enabled for the topic.

Log Search
Log Search Failure

Last updated：2024-01-20 17:11:57

The log search may fail sometimes. In case of a search failure, use the following methods for troubleshooting.

Checking Search Criteria

A log search failure is often caused by an incorrect time range or search statement. To address this issue, first select a larger time range (such as `last 30 minutes`), leave the search bar empty, and search for logs.
If logs are found, it indicates that the log search is available. We recommend that you check the search syntax and rules or modify the time range.

Checking Index Configuration

The index configuration is required for CLS log search. On the top right of the **Search Analysis** page, click **Index Configuration** to enable both full-text index and key-value index. For more information, see Enabling Index.
Note:
The index configuration takes effect in about 1 minute. The new configuration is only effective for log data written subsequently.

Checking Log Collection

Log collection from Tencent Cloud services

To collect logs from other Tencent Cloud services including TKE and CLB, see Collection for Tencent Cloud Services to verify the configuration. If you have any question, please contact smart customer service.

Log collection by LogListener client

If you're using CLS's LogListener client to collect logs, perform the following steps for troubleshooting:
1. Check the machine group.
On the top right of the **Search Analysis** page, click **LogListener Collection Configuration** to check the machine group from which you want to collect logs.
Note:
If the server is exceptional, see Server Group Exception.

2. Check if LogListener obtains the collection configuration from the CLS server.

Run the following CLI commands:

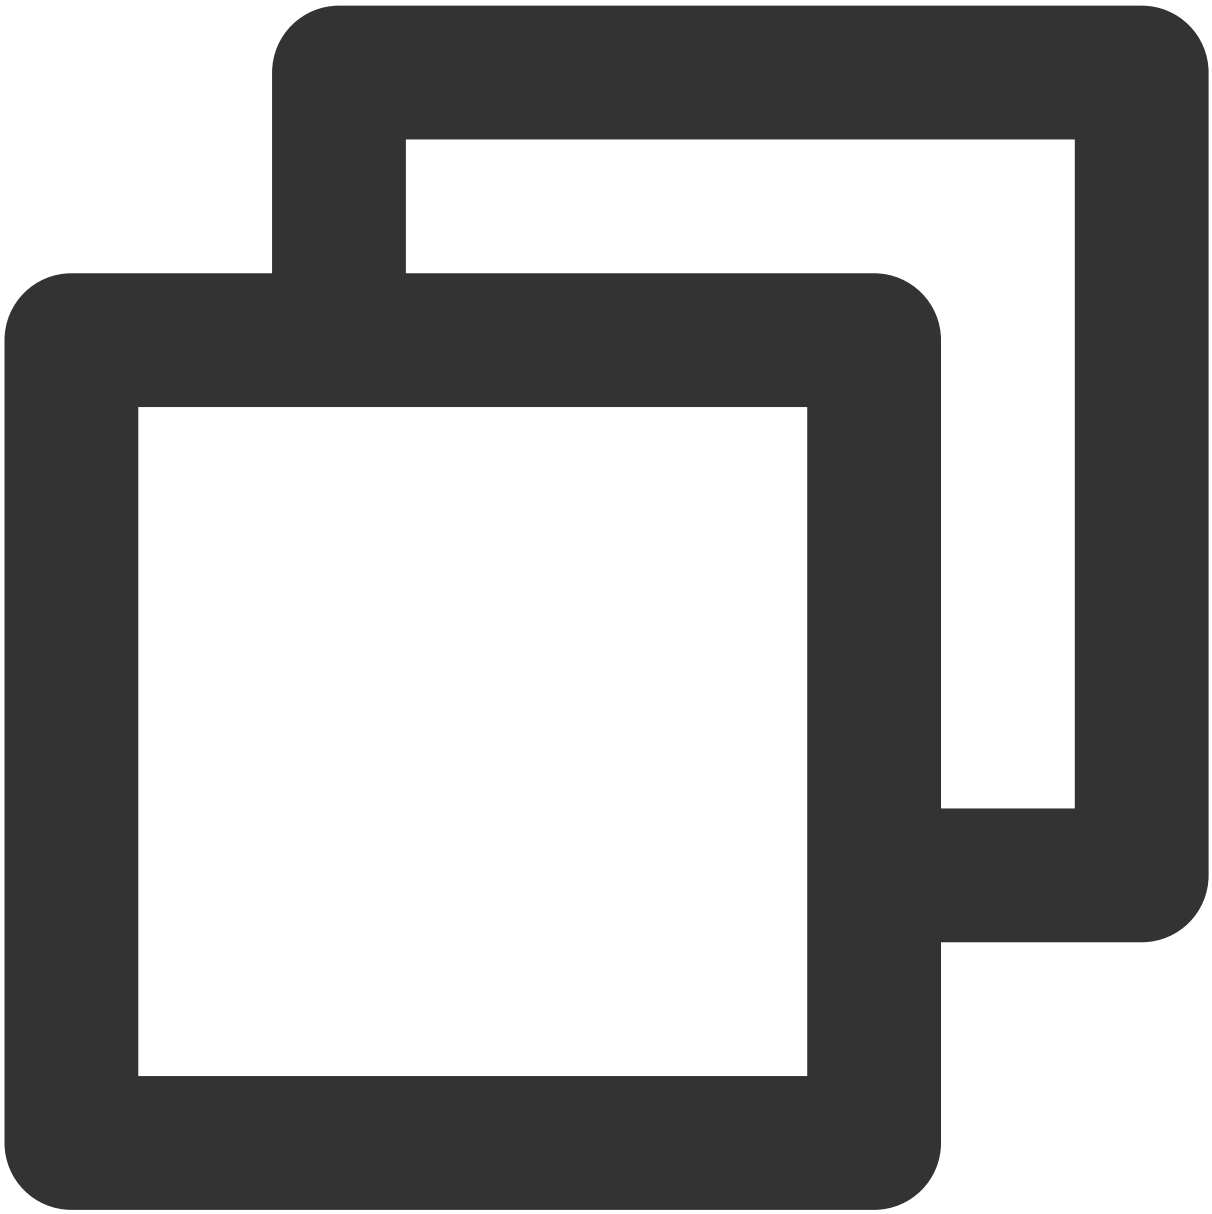

```
/etc/init.d/loglistenerd check
```

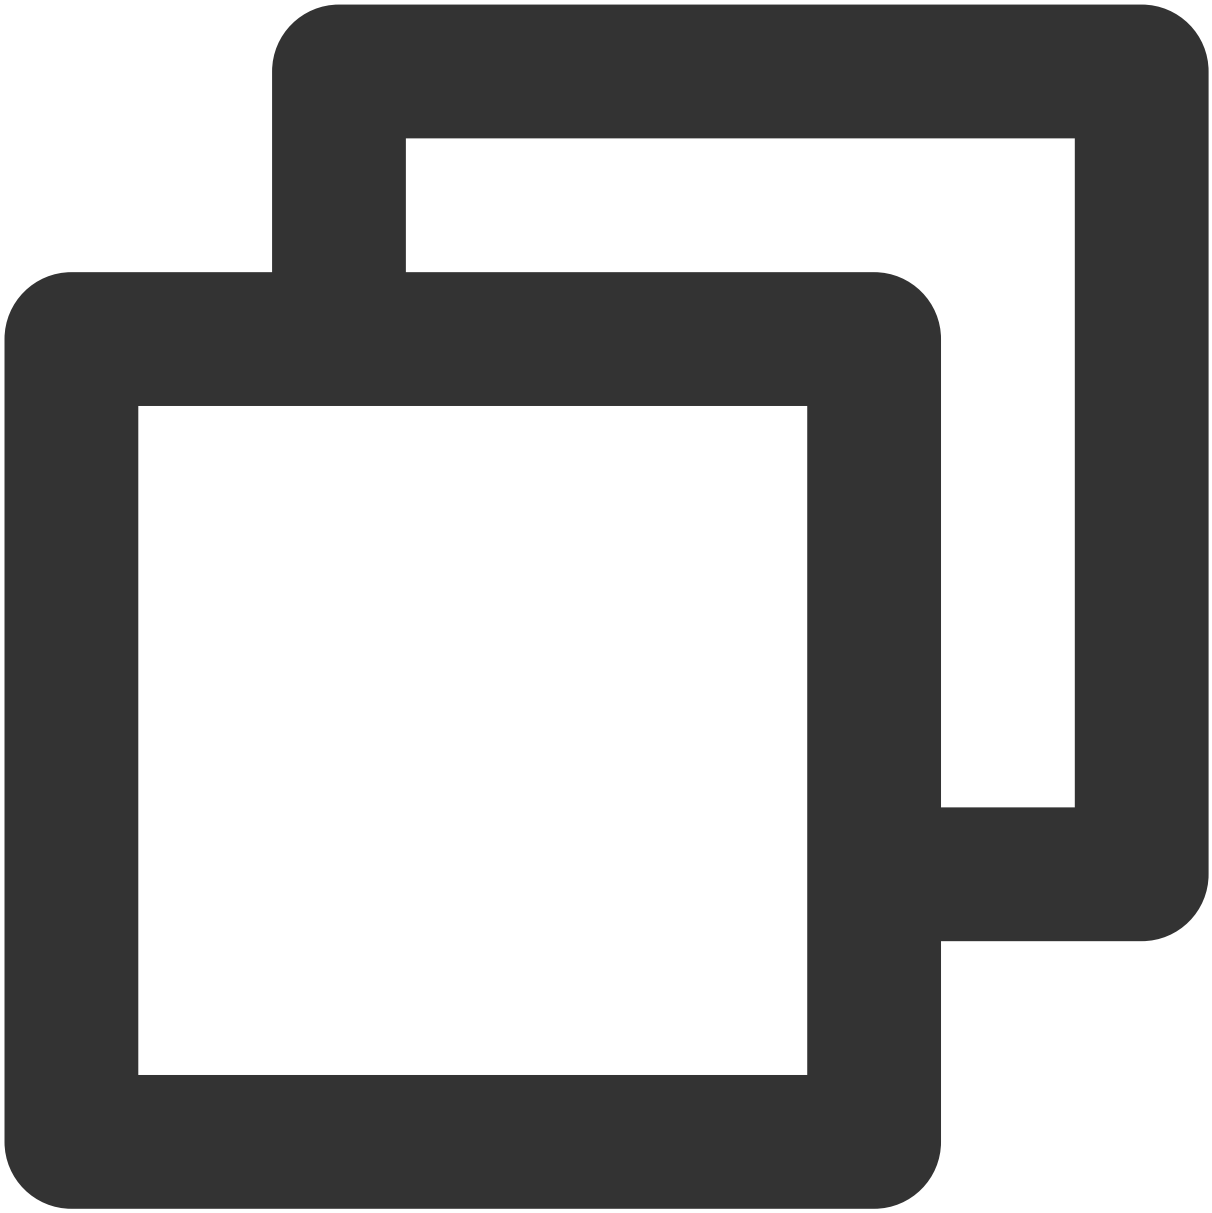

```
If the result returns “[OK] check loglistener config ok” as shown in the following
![](https://main.qcloudimg.com/raw/95022fc7832b36e2e8d51b6fe8ed3ab7.jpg)
The `logconf` field in the result refers to the collection configuration. If this
```

3. Use the latest version of LogListener.

Run the following command to check the version number. See LogListener Installation Guide to install the latest version of LogListener.

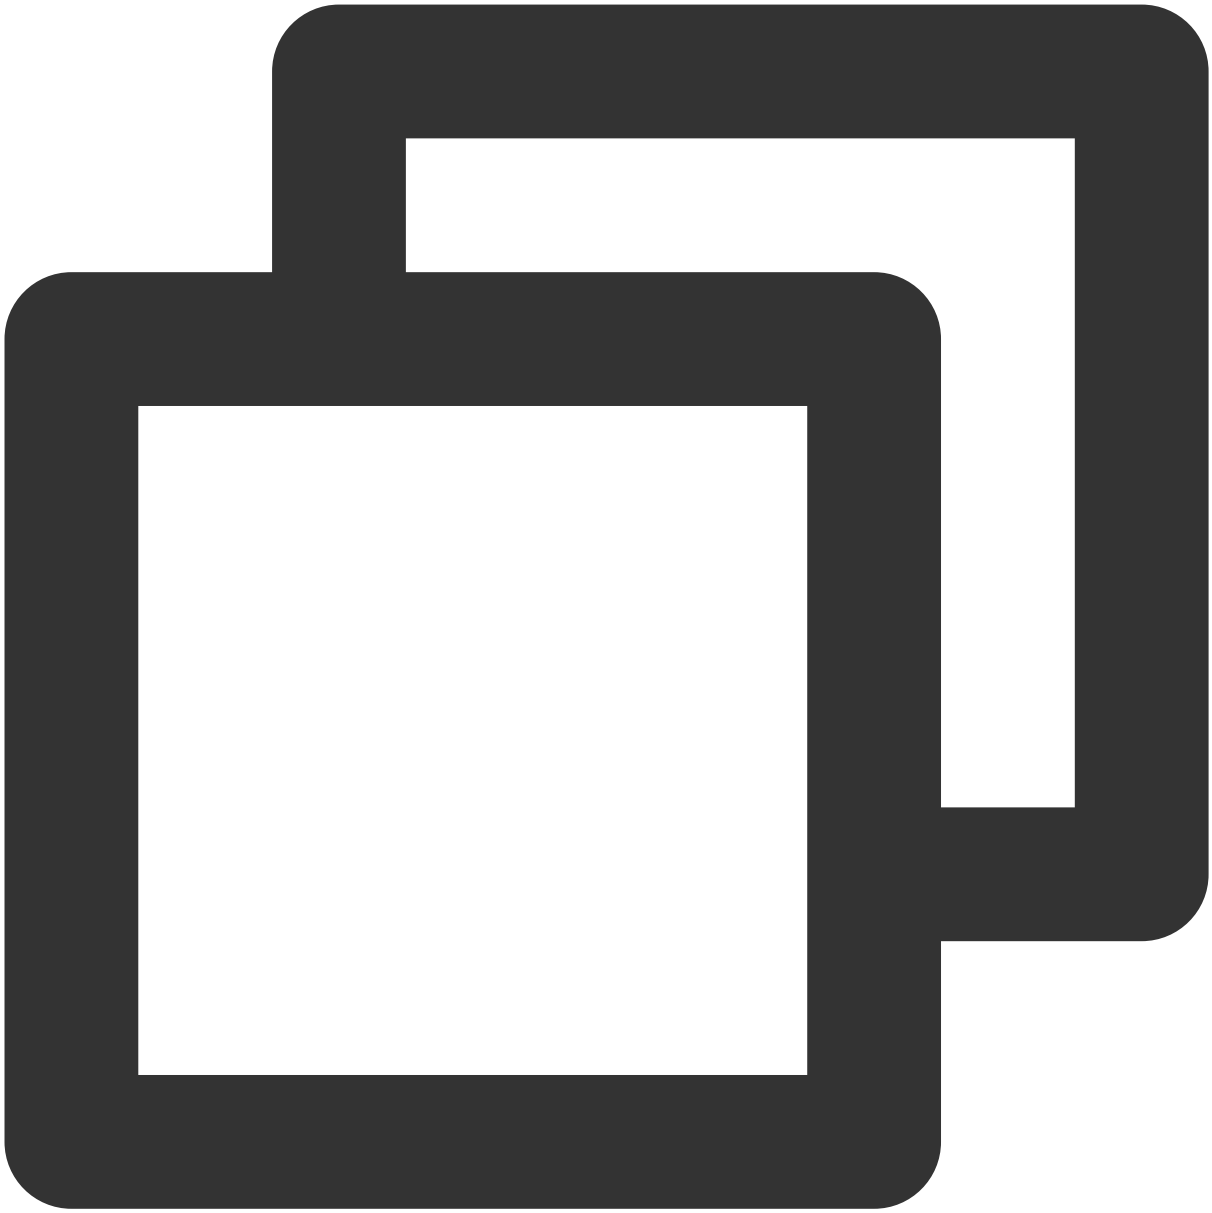

```
/etc/init.d/loglistenerd -v
```

Note:

LogListener earlier than 2.3.0 cannot collect log files in soft links.

4. Check that logs are successfully reported.

4.1 Open the LogListener Debug log and access the LogListener installation directory. Set **level** to `DEBUG` in the `etc/loglistener.conf` configuration file and restart LogListener.

```
<log>
    level   = DEBUG
    path    = log/
    name    = loglistener
    size    = 10000000
    num     = 10
</log>
```

4.2 Run the following command to restart LogListener.

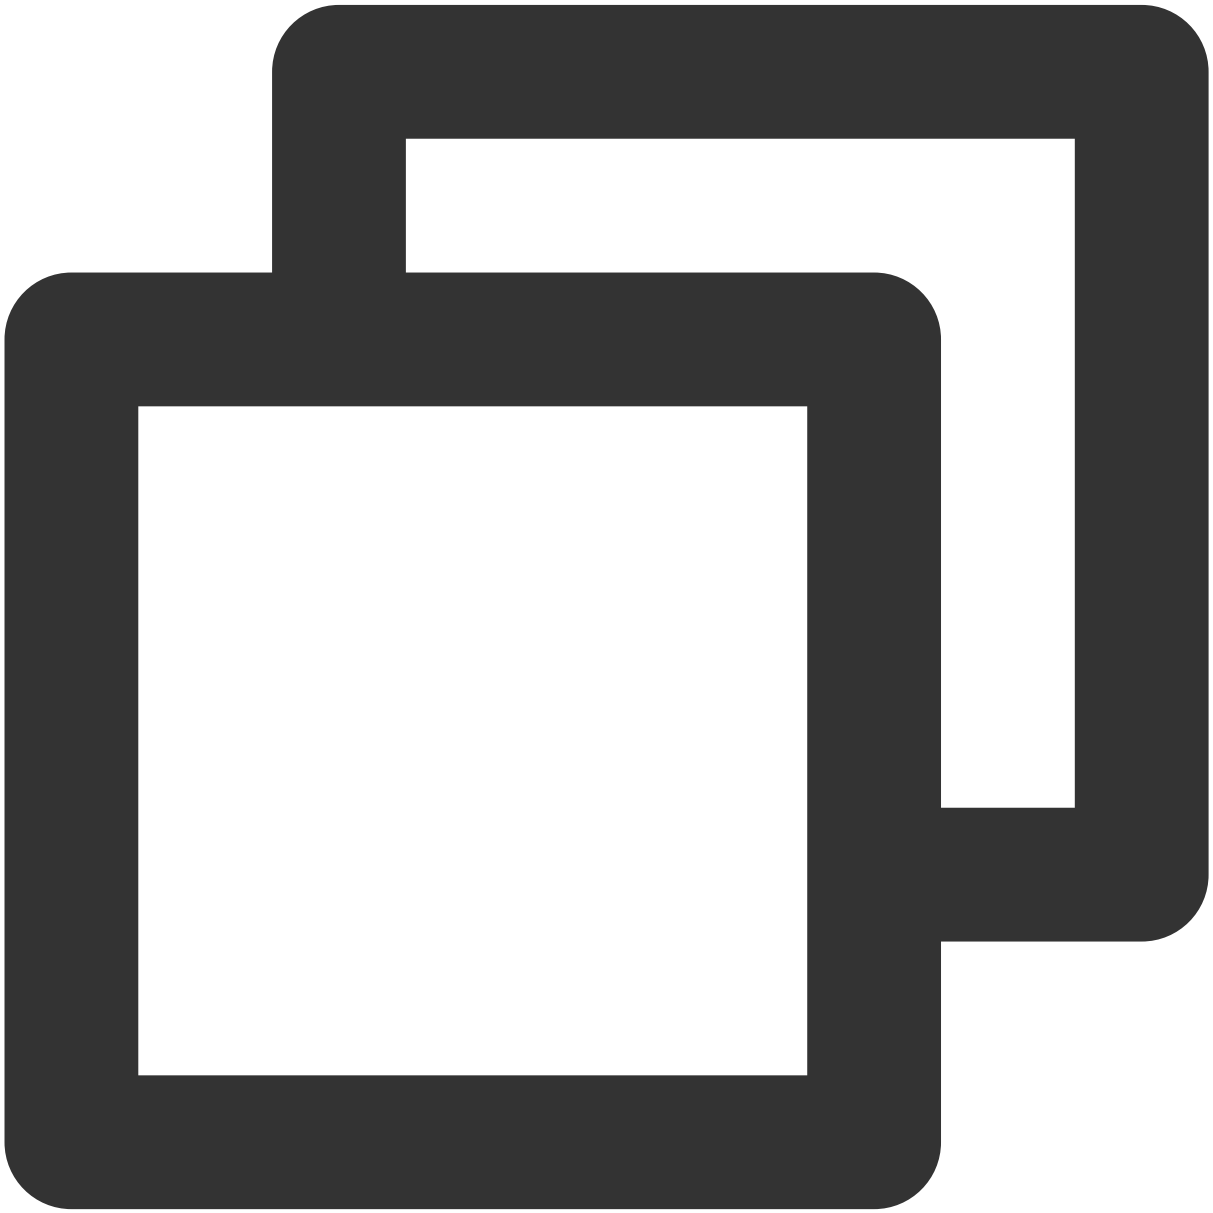

```
/etc/init.d/loglistenerd restart
```

4.3 Run the following commands to check whether logs are successfully reported.

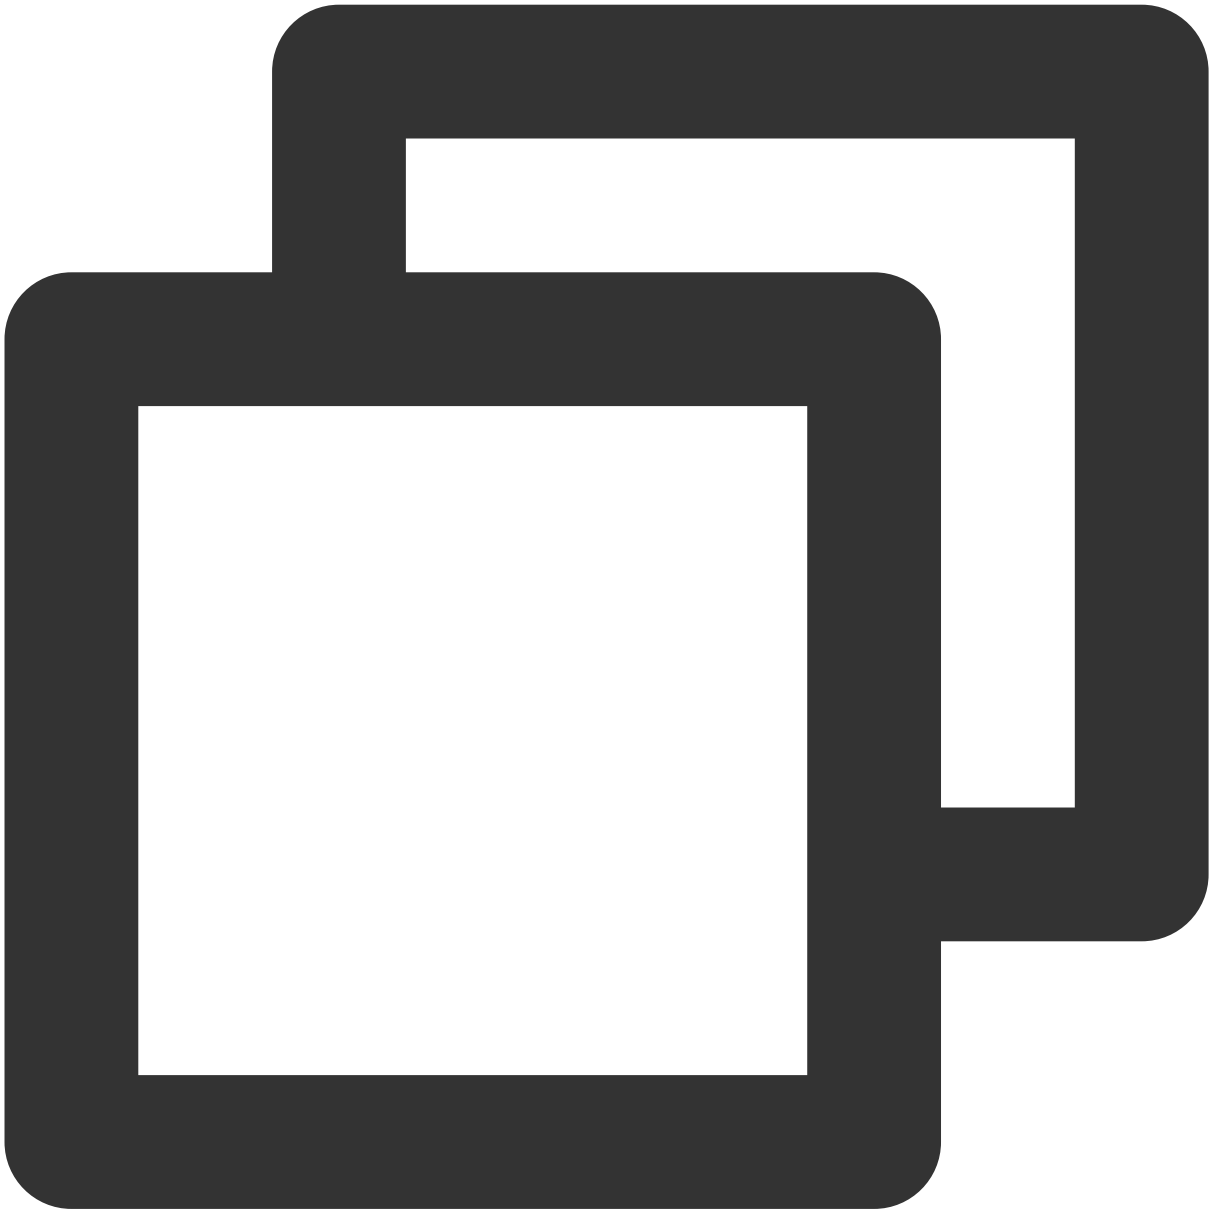

```
tail -f log/loglistener.log | grep "ClsFileProc::readFile" | grep send
```

If log information similar to that shown in the following figure is displayed, logs are successfully reported to the CLS server.

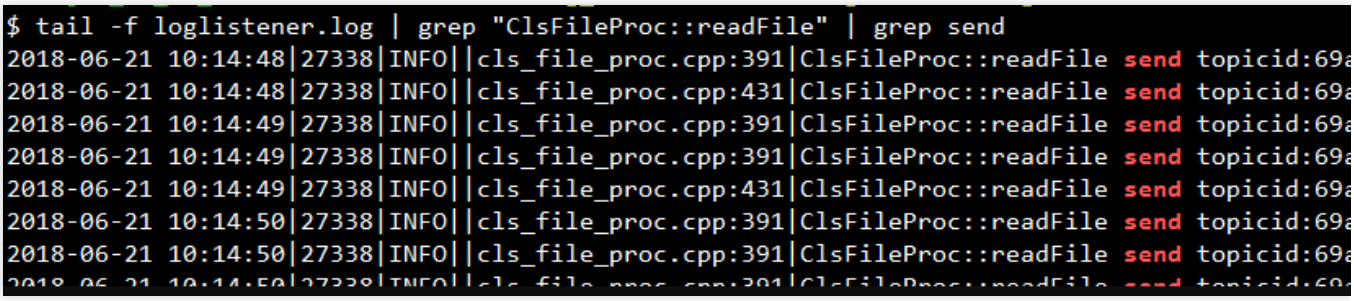

Note:

If logs are reported through HTTP, you can capture packets from port 80 to verify whether logs are successfully reported.

If logs are not reported, perform the following steps for troubleshooting:

a. Run the following commands in the installation directory to check whether the LogListener collection configuration is correct.

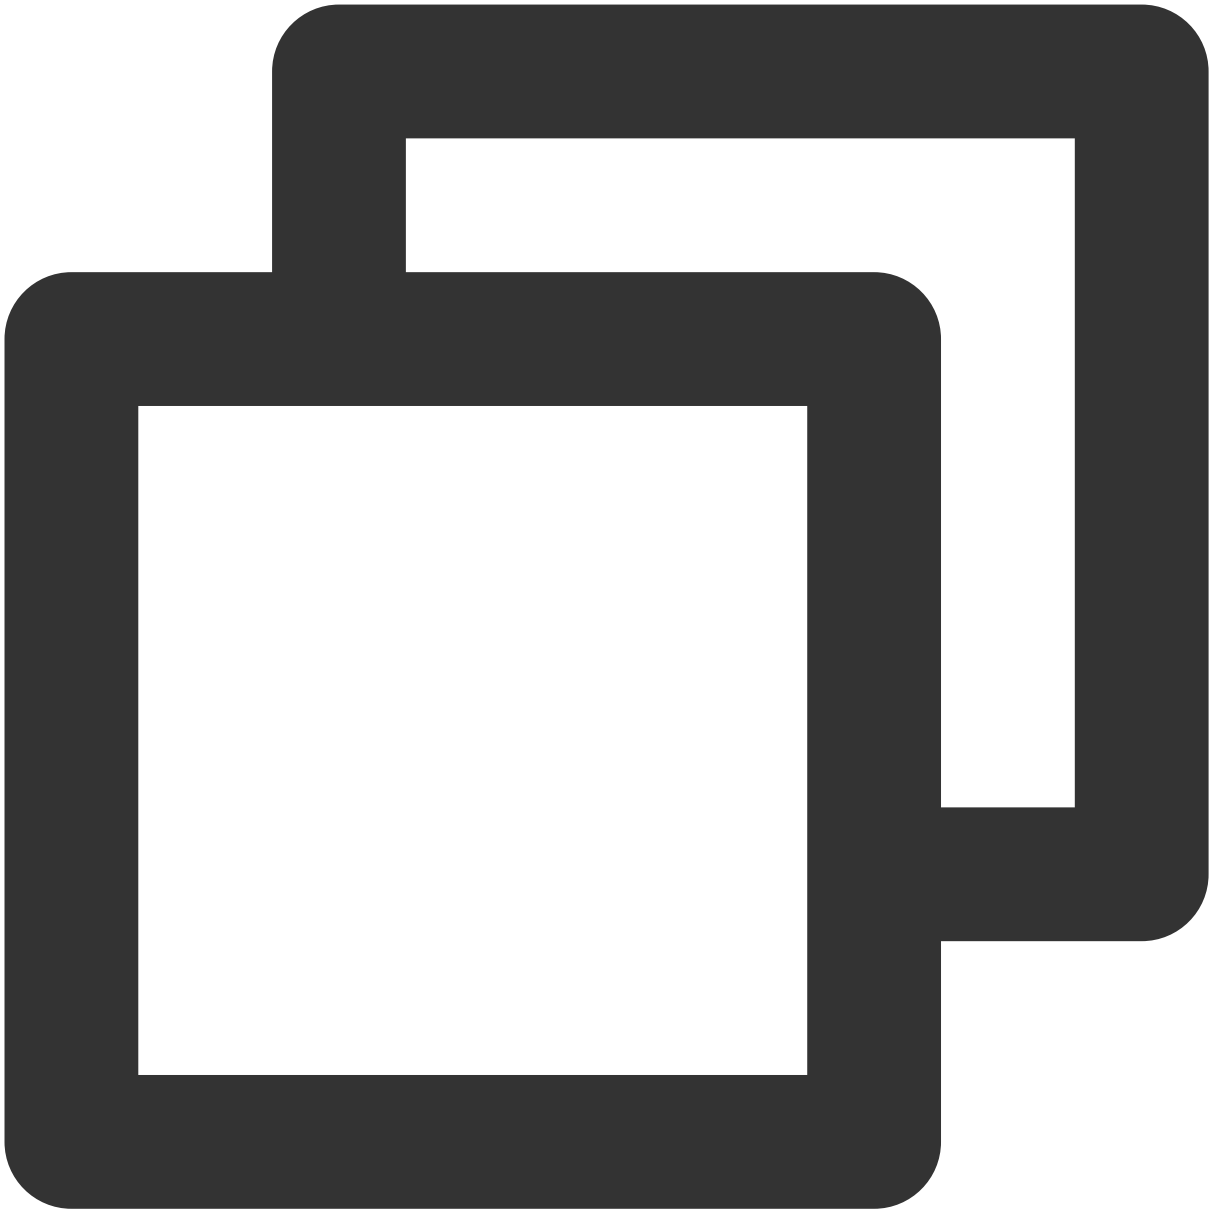

```
tail -f log/loglistener.log | grep "ClsServerConf::load"
```

If the configuration has been delivered to LogListener, log information is as follows:

```
$ tail -f log/loglistener.log | grep "ClsServerConf::load"
2018-06-21 10:01:49|20706|DEBUG||cls_server_conf.cpp:24|ClsServerConf::loa
"path":"/log","topicid":"56ed3e87-c895-49ba-a1cc-2f2c30e57a35"},{"extract_
a0207f-f3ec-4beb-a50f-9572546c1e8c"}],"needupdate":false}
```

In the delivered configuration, check whether the information of `log_type` and `path` is correct:

`log_type` indicates the log parsing type. Valid values: `minimalist_log` (full text in a single line), `delimiter_log` (separator), `json_log` (JSON logs), and `regex_log` (full text in multi lines).
`path` indicates the log collection directory.

b. Run the following command in the installation directory to check whether files are correctly listened to:

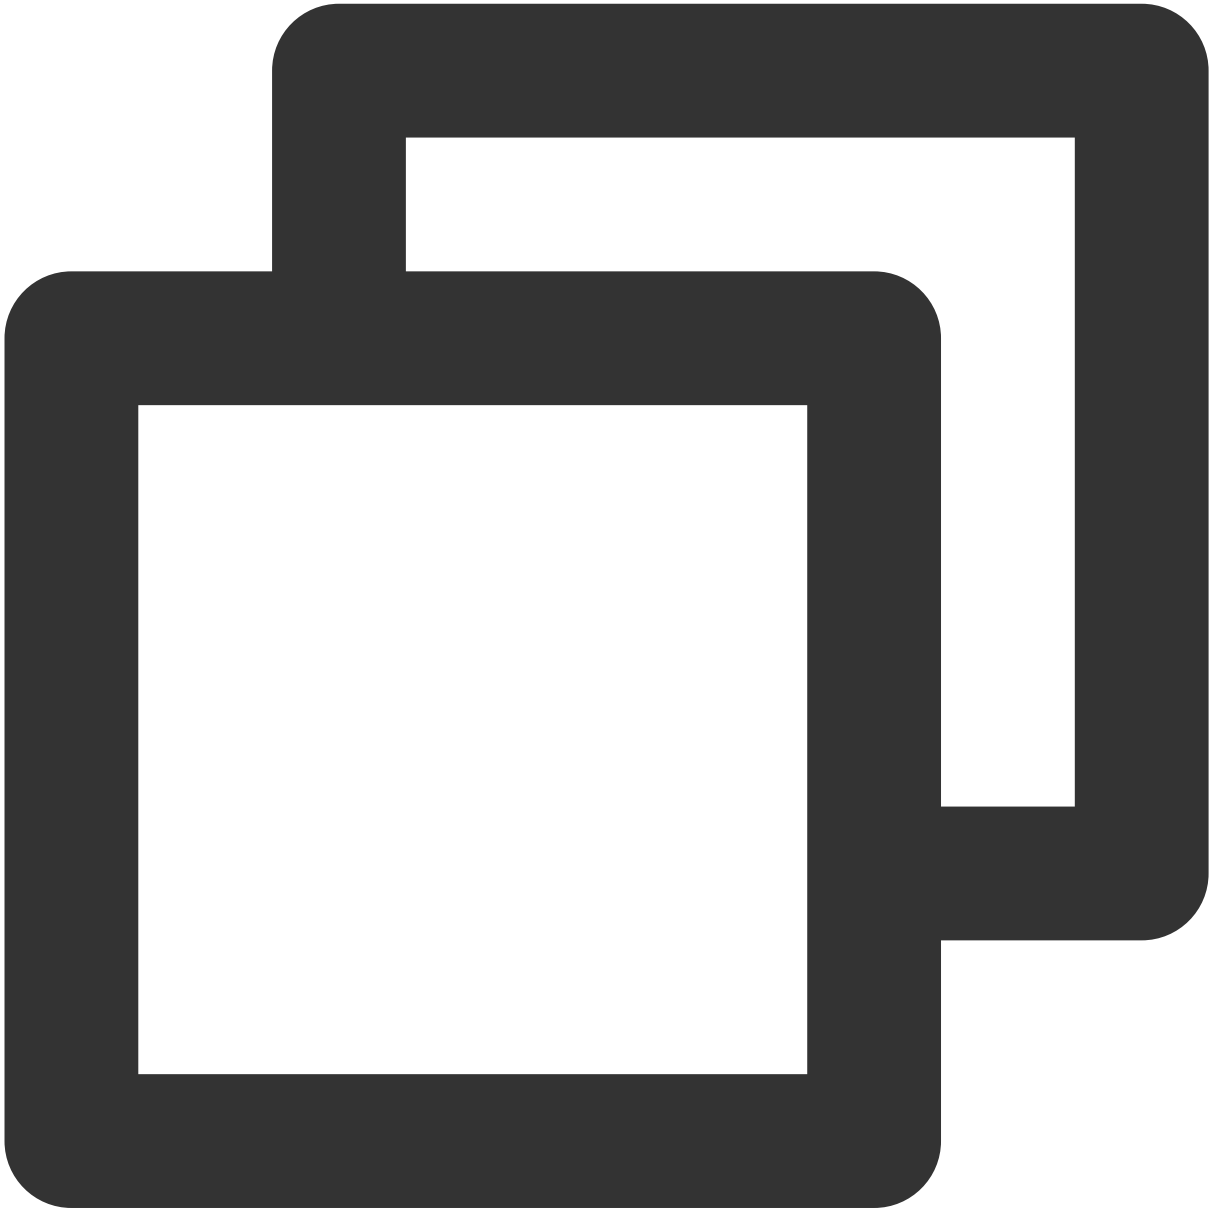

```
grep [Name of the reported log file] log/loglistener.log
```

If no log information is displayed, run the `grep regex_match log/loglistener.log` command to search for log information and check whether the regular expression is correctly configured in the console. If the content shown in

the following figure is displayed, the file name match based on the regular expression fails. In this case, please log in to the console and change the regular expression.

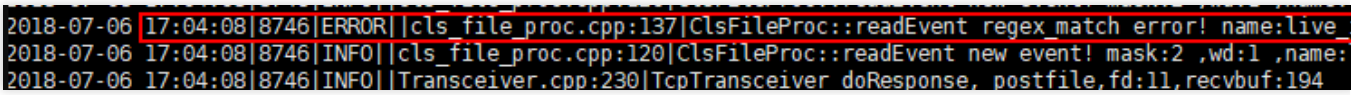

c. Check whether the log regular expression parse is correct.

For the extraction modes of full regular expression and full text in multi lines, regular expressions need to be specified. For full text in multi lines, the first line regular expression must match the entire content of the first line, instead of the beginning part of the first line.

Use the log content shown in the following figure as an example. Lines beginning with `INFO` , `ERROR` , and `WARN` are the first lines of logs. In addition to `(INFO|ERROR|WARN)` , the characters following `INFO` , `ERROR` , and `WARN` also need to be matched.

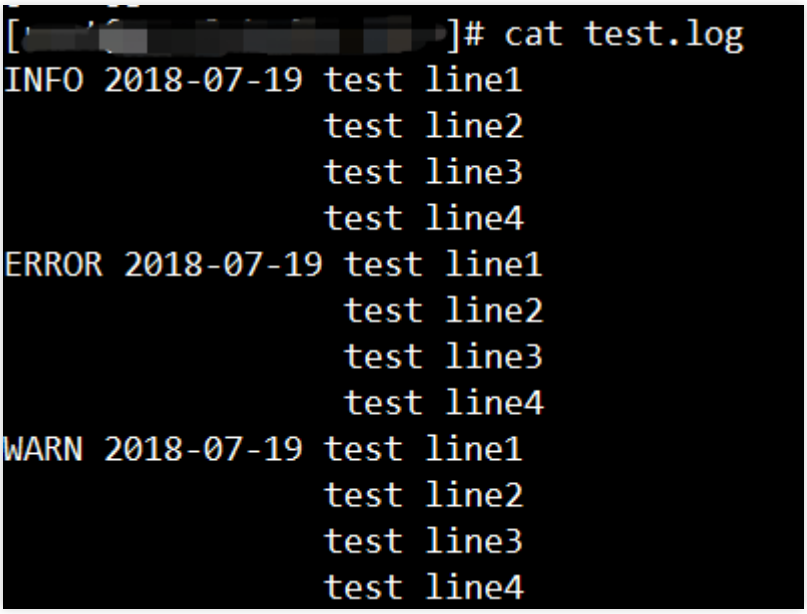

Incorrect configuration: `^(INFO|ERROR|WARN)`

Correct configuration: `^(INFO|ERROR|WARN).*`

5. A file can only be collected to one log topic and a single log line cannot exceed 1 MB.

Meet these requirements to ensure the complete log collection.

Search Analysis Error

Last updated：2024-01-20 17:11:57

Common error messages, causes and solutions are as follows:

Error Message	Cause	Solution
QueryError [illegal_argument_exception.Cannot search on field [xxx] since it is not indexed.]	Key-value index is not enabled for the query field `xxx`	Enable key-value index for this field. For details, please see Key-Value Index.
QueryError [illegal_argument_exception.Cannot search on Full-Text since it is not indexed.]	Full-text index is not enabled	Enable full-text index for this field. For details, please see Full-Text Index.
QueryError [illegal_argument_exception.syntax error on field [and|or|not], or full text search is closed]	The search condition does not support lowercase logical operators, which will be regarded as normal fields for full-text search	Use the uppercase logical operators AND|OR|NOT. If you do not need to use logical operators but to search for and/or/not, please enable full-text index.
QueryError [number_format_exception.For input string: ">"]	Syntax error of numerical comparison statement	Check whether there are special symbols such as spaces around the numerical comparison symbols. An example of the correct format: status:>400
QueryError [circuit_breaking_exception. Analysis data is too large,please reduce the scope of data query]	The query data volume is too large	Reduce the query time range appropriately and specify more precise query conditions. If the error persists, contact technical support.
QueryError [parse_exception.parse_exception: Cannot parse 'xxx': '*' or '?' not allowed as first character in WildcardQuery	Fuzzy query by prefix is not allowed, e.g. content:*examplecontent:*example	We recommend using separators to split a field into multiple ones. For details, please see Configuring Index
QueryError [sql_illegal_argument_exception.cannot	`cast` cannot convert dates in 13/Jul/2021:17:04:34 format. Only	Modify the format of the time field or use the

cast [13/Jul/2021:17:04:34] to [datetime]: failed to parse date field [13/Jul/2021:17:04:34] with format [date_optional_time]]	ISO standard format and millisecond-level UNIX timestamp are supported, e.g. yyyy-MM-dd'T'HH:mm:ss.SSSZ or yyyy-MM-dd.	__TIMESTAMP__ built-in field
QueryError [verification_exception.Cannot order by non-grouped column [xxx], expected [xxx] or an aggregate function	The statistics feature is not enabled for the field `xxx` and thus it cannot be used for sorting	Enable statistics for this field. For details, please see Log Analysis Overview
QueryError [verification_exception.Cannot use non-grouped column [xxx], expected [xxx]]	The statistics feature is not enabled for the query field `xxx`	Enable statistics for this field. For details, please see Log Analysis Overview
QueryError [verification_exception.Field [xxx] of data type [text] cannot be used for grouping]	The statistics feature is not enabled for the field `xxx` and thus it cannot be used for grouping	Enable statistics for this field. For details, please see Log Analysis Overview
QueryError [verification_exception.Unknown column [xxx]]	The query field `xxx` does not exist	Check whether the field name is correct
QueryError [verification_exception.Unknown function [xxxxxx]]	The function xxxxxx does not exist.	Check whether the function name is correct. In addition, this error also occurs if some functions are used together with Histogram functions. In that case, use a Time Completion Function to replace the Histogram function.
QueryError [verification_exception.argument of [FUNCNAME(xxx)] must be [numeric], found value [xxx] type [text]]	The type of the input parameter of the `FUNCNAME` function is incorrect. For example, if the `level` field of the `SUM(level)` function is of the text type, an error will be reported	Check whether the field type meets the function requirements
QueryError [parse_exception.Failed to parse query [xxx]]	Syntax error of query statement	Check the error position specified in the error information
QueryError [line X:X: XXX]	Syntax error of query statement	Check the error position and cause specified in the error information

Internal error. Please try again later RequestId:[7be994d4-xxxx-xxxxx-xxxx-9c38xxxx65de]	CLS internal error	Contact technical support and provide the `RequestId` in the error information.
SyntaxError[xxx]	There is a syntax error in part of the SQL statement	Please see the detailed tips in the error message to fix the syntax error, where line x,column x does not contain the search condition part (i.e. "|" and the part before it)
SearchTimeout	The query timed out	Reduce the scope of data query and SQL complexity as appropriate, or try again later.
LimitExceeded.LogSearch	The search concurrency exceeds the limit	Reduce the query frequency (including API call frequency) and try again later. If the current query frequency is not high, and the error persists, contact technical support.

.

Others

Last updated：2024-01-20 17:11:57

What is CLS?

Cloud Log Service (CLS) provides a one-stop log data solution. You can quickly and conveniently connect to it in five minutes to enjoy a full range of stable and reliable services from log collection, storage, and processing to search, analysis, consumption, shipping, dashboard generation, and alarming, with no need to care about resource issues such as scaling. It helps you improve the problem locating and metric monitoring efficiency in an all-around manner, making log Ops much easier.
CLS has the following features.
Log collection: CLS easily collects logs from different regions, channels, platforms, and data sources (e.g., various Tencent Cloud products) in real time.
Log storage: CLS offers two storage types: real-time storage and IA storage.
Log search and analysis: You can search for logs by keyword to quickly locate exception logs and use SQL statements to collect and analyze log statistics. This helps you get statistical metrics such as log quantity change trend over time and proportion of error logs.
Log data processing: CLS can filter, cleanse, mask, enrich, distribute, and structure logs.
Log shipping and consumption: CLS can ship logs to Tencent Cloud storage and middleware services and consume logs to stream computing services.
Dashboard: CLS can quickly generate custom dashboards for search and analysis results.
Alarming: CLS can trigger alarms for exception logs within seconds and notify you through phone, SMS, email, and custom API callback.

How is a log defined in CLS?

Logs are record data generated during the running of an application system, such as user operation logs, API access logs, and system error logs. Logs are usually stored in text format on the host where the application system resides. A log corresponding to a system running record may contain one line of text (single-line log) or multiple lines of text (multi-line log).
For more information and examples, see Log and Log Group.

How long can a log be retained?

CLS provides the log lifecycle management feature. You can set the log validity period to 1–3,600 days or permanent when creating a log topic. Once expired, the data will be cleared and no longer incur storage fees.

What are the differences between a logset and a log topic?

A log topic is a basic unit for log data collection, storage, search and analysis on the CLS platform. The massive amounts of logs collected are managed by log topic. For example, you can configure log collection rules and storage

time, search for and analyze logs, and download, consume, and ship logs by log topic.
A logset is a class of log topics and can contain multiple log topics. A logset itself does not store any log data, but just makes it easier for users to manage log topics.
For more information and examples, see Log Topic and Logset.

How many logs can a single log topic collect?

To collect high numbers of logs, a single topic contains multiple partitions, each of which has up to 500 write QPS and 5 MB/s write traffic. If there are many logs to be collected, we recommend you enable the automatic partition splitting feature (which is enabled by default). A single log topic can contain up to 50 partitions and thus has up to 50 * 500 = 25000 write QPS and 50 * 5 = 250 MB/s write traffic.
Here, the write request quantity and write traffic are not simply equal to the number of logs and log volume respectively. As multiple logs will be packaged and compressed into a log group during log upload, the actually supported number of logs and log volume are far greater than the above values. Logs will be automatically packaged and compressed if you use LogListener, so you don't need to care about the specific packaging policy.

What are an index and a segment?

Index configuration is necessary for log search and analysis in CLS. Only after index is enabled can CLS search for and analyze logs. Index creation is to split a raw log into multiple segments with the specified symbol and add an inverted index to such segments.
For more information and examples, see Segment and Index.

What are the differences between full-text index and key-value index?

Full-text index: A raw log is split into multiple segments, and indexes are created based on the segments. You can query logs based on keywords (full-text search). For example, entering `error` means to search for logs that contain the keyword `error` .
Key-Value search: A raw log is split into multiple segments based on a field (key:value), and indexes are created based on the segments. You can query logs based on key-value (key-value search). For example, entering `level:error` means to search for logs with a `level` field whose value contains `error` .
For more information and examples, see Configuring Indexes.

What are the differences between search and analysis?

Search: You can search for matched raw logs by specified criteria. For example, you can enter `status:404` to search for application request logs whose response status code is 404.
Analysis: You can use SQL statements to collect and analyze the statistics of logs meeting specified search criteria. For example, you can enter `status:404 | select count(*) as logCounts` to get the number of application request logs whose response status code is 404.
For more information and examples, see Overview and Syntax Rules.

How high is the performance of log search and analysis?

Search performance: Results can be returned within seconds for tens of billions of logs.
Analysis performance: Results can be returned within seconds for hundreds of millions of logs and within one minute for tens of billions of logs. The performance is subject to the complexity of the SQL statement used for analysis. If the statement is very complex, the performance may be lower.

How long does it take for logs to become searchable after generation?

The delay is within one minute if LogListener is used for log collection. If an API or SDK is used to collect logs, it takes no more than one minute for logs to become searchable after API call.

Can I use CLS if my business is not in Tencent Cloud?

Yes. CLS has no restrictions on the log source. You can collect logs to CLS as long as the log source can be connected to CLS over the network. For specific regions supported by CLS and corresponding domain names, see Available Regions.

How do I configure LogListener after the server IP address is changed?

If the server is bound to a machine group by server ID, you don't need to modify the LogListener configuration. Therefore, if the server IP needs to be changed frequently, we recommend you configure a machine group by server ID. For more information, see Machine Group Management.
If you configure the machine group by IP address, modify the configuration as follows:
1. Modify the `/etc/loglistener.conf` file under the LogListener installation directory. Here, the `/user/local` installation directory is used as an example:

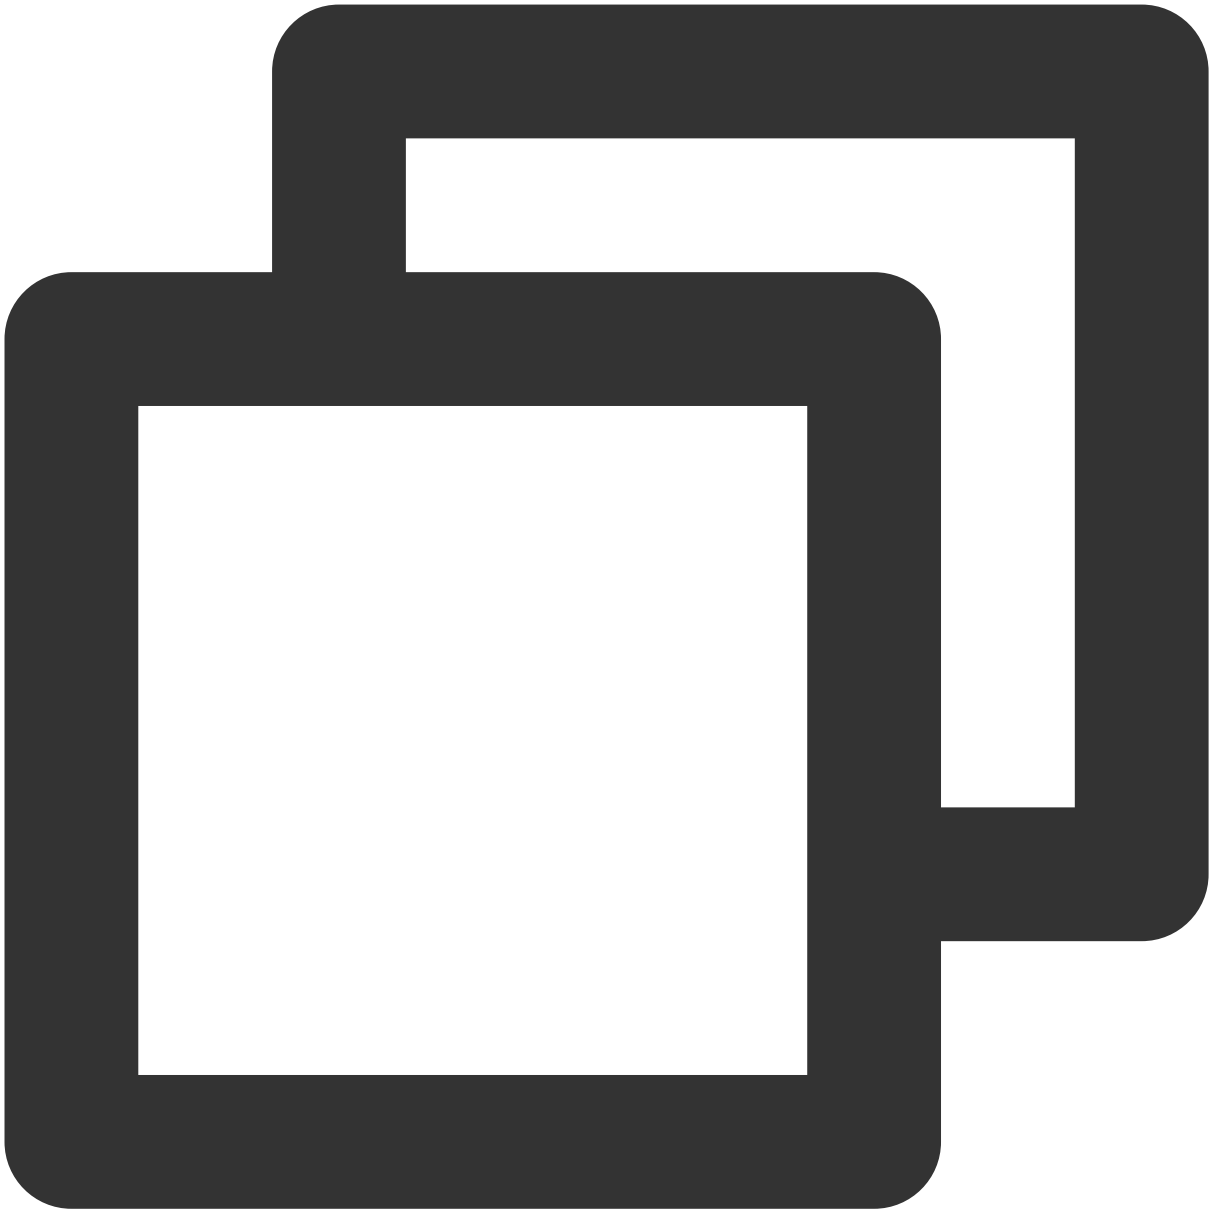

```
vi /usr/local/loglistener-2.3.0/etc/loglistener.conf
```

2. Press **i** to enter the edit mode.
3. Enter the changed IP address in `group_ip` in the configuration file.
4. Press **Esc**, enter **:wq**, and press **Enter** to save the configuration and exit the editor.
5. Run the following command to restart LogListener.

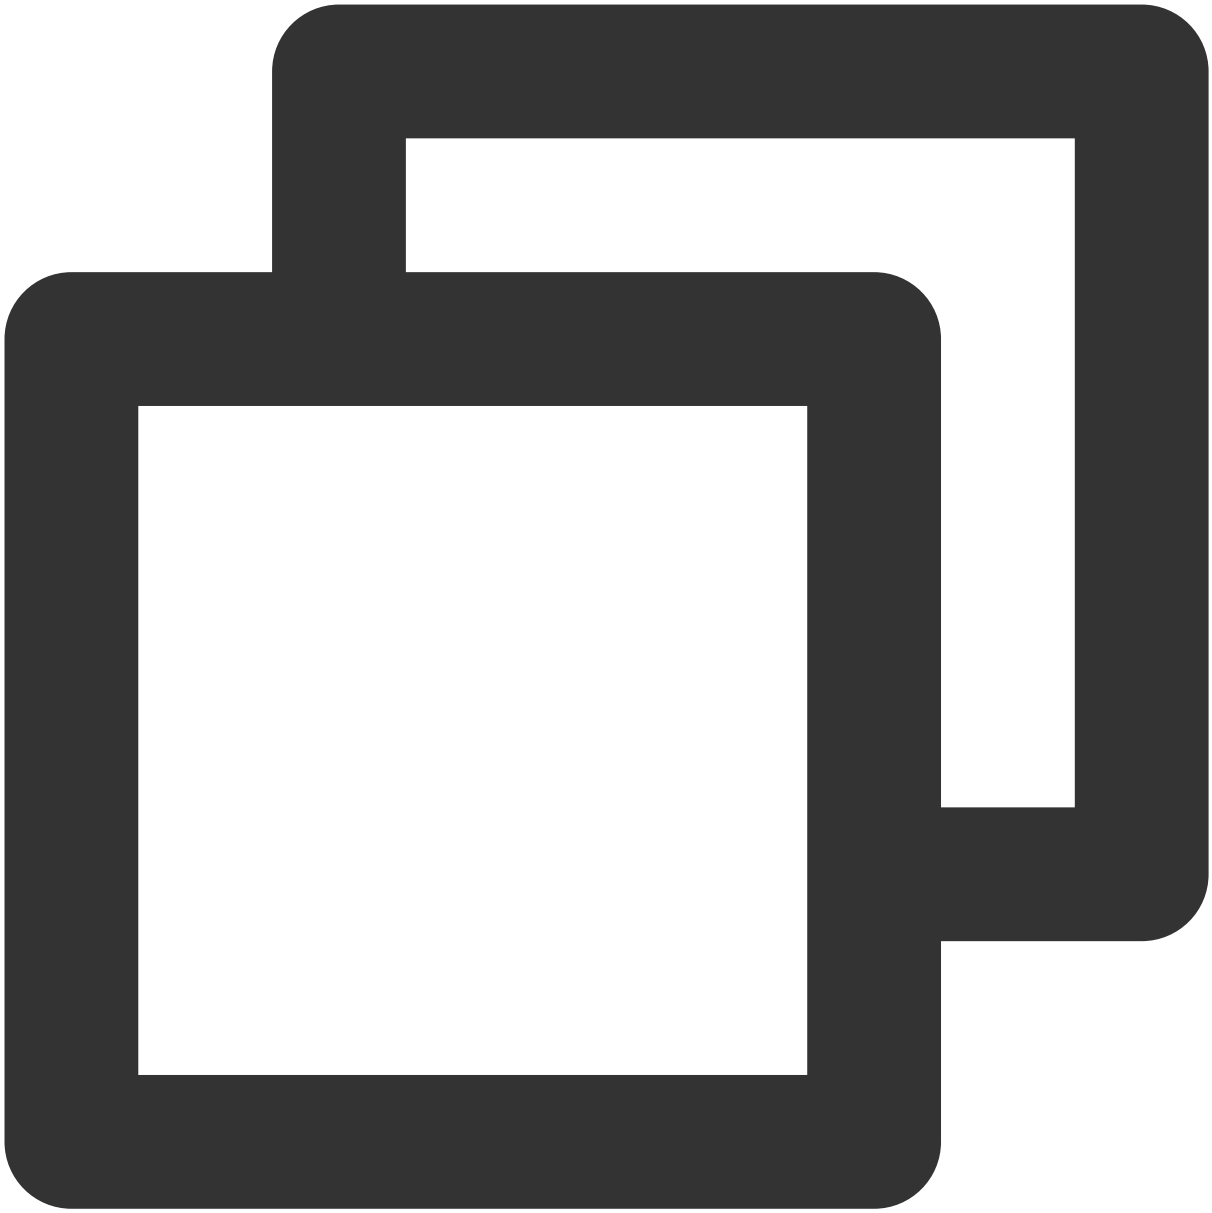

```
/etc/init.d/loglistenerd restart
```

6. Log in to the CLS console and select **Machine Group Management** on the left sidebar. Locate the machine group to which the server is bound and click **Edit**. In the pop-up window, replace the old IP address with the new one and click **OK**.

How do I troubleshoot when the testing alarm notification channel reported an error or failed to receive the testing message?

Case 1: The page displays "Message sending failed".

Hover over the "Message sending failed" message to view the error code and detailed failure cause. Common error codes are as listed below:

Error Code	Description	Troubleshooting Method
-1004	Message sending via this notification channel failed.	This is generally because no mobile numbers or email addresses have been configured or verified for all recipients or recipient groups. You can view and configure them in the user list.
-1005	Some messages failed to be sent via notification channels; for example, messages failed to be sent to some users or via some notification channels.	This is generally because no mobile numbers or email addresses have been configured or verified for some recipients or recipient groups. You can view and configure them in the user list.
-1006	The custom API callback reported an error.	Troubleshoot based on the specific failure cause. Common errors include: invalid URI for request: The URL is Invalid. i/o timeout: API access timed out. Check the API address and whether it can be directly accessed over the public network. callback custom error with status:xxx: An API response error occurred. Check whether the API address and backend service are normal. ssrf attack: The callback API address must be directly accessible over the public network. This error may occur if the API address is a Tencent Cloud private network address.

Case 2: The page displays "Sent", but the testing message is not received.

Common reasons of different receipt channels are as listed below:

Receipt Channel	Reason
Email, SMS, and phone	To avoid disturbing users with repeated notifications, only one testing message can be sent to the same user via each channel per day.
Custom API callback (DingTalk and Lark bot addresses)	The testing message didn't meet DingTalk or Lark's API requirements and was ignored. In this case, the feature of the testing notification channel is meaningless. You can directly configure the appropriate request headers and content in the alarm policy according to DingTalk and Lark's API requirements to send alarm messages.
Custom API callback (other addresses)	CLS determines whether a message was sent successfully based on the HTTP response status code. Check whether the custom API has other business logic limits while the HTTP response status code is normal.