Log Tank Service

FAQs

Issue	01
Date	2025-02-08

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

Trademarks and Permissions

HUAWEI and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei Cloud and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Contents

1 Overview

This document provides answers to frequently asked questions related to Log Tank Service (LTS).

Consultation

- How Do I Migrate Logs from a Third-Party Cloud to Huawei Cloud LTS?
- Can Logs Stored in LTS Be Used for Security Compliance Audit?
- What Are the Recommended Scenarios for Using LTS?
- What Are the Advantages of LTS Compared with Self-built ELK Stack?

Log Management

- How Do I Save Only One Copy of Logs After Multi-Account Log Aggregation Is Configured?

Host Management

- What Do I Do If ICAgent Installation Fails in Windows and the Message "SERVICE STOP" Is Displayed?
- What Do I Do If ICAgent Upgrade Fails on the LTS Console?
- What Do I Do If I Could Not Query New Logs on LTS?
- What Do I Do If ICAgent Restarts Repeatedly After Being Installed?
- What Do I Do If ICAgent Is Displayed as Offline on the LTS Console After Installation?
- What Do I Do If I Do Not See a Host with ICAgent Installed on the LTS Console?
- How Do I Create a VPC Endpoint on the VPCEP Console?
- How Do I Obtain an AK/SK Pair?
- How Do I Install ICAgent by Creating an Agency?

Log Ingestion

- What Do I Do If LTS Cannot Collect Logs After I Configure Host Log Ingestion?

- Will LTS Stop Collecting Logs After the Free Quota Is Used Up If I Disable "Continue to Collect Logs When the Free Quota Is Exceeded" in AOM?
- What Do I Do If the CPU Usage Is High When ICAgent Is Collecting Logs?
- What Kinds of Logs and Files Does LTS Collect?
- How Do I Disable the Function of Collecting CCE Standard Output Logs to AOM on the LTS Console?
- What Log Rotation Scheme Should I Use for ICAgent to Collect Logs?
- Does LTS Use the Log4j Plug-in to Report Logs?
- How Long Does It Take to Generate Logs After Configuring Log Ingestion?
- What Do I Do If LTS Cannot Collect Logs After I Configure Log Ingestion with ICAgent?

Log Search and Analysis

- How Often Is the Data Loaded in the Real-Time Log View in LTS?
- What Do I Do If I Cannot View Reported Logs in LTS?
- Can I Manually Delete Logs on the LTS Console?
- What Do I Do If I Could Not Search for Logs on LTS?

Log Transfer

- Does LTS Delete Logs That Have Been Transferred to OBS Buckets?
- What Are the Common Causes of LTS Log Transfer Abnormalities?
- How Do I Transfer CTS Logs to an OBS Bucket?
- What Do I Do If I Cannot View Historical Data in an OBS Bucket After Transferring Data from LTS to OBS?
- What Do I Do If I Cannot Find a New Partition in a DLI Table After Logs Are Transferred to DLI?

2 Consultation

2.1 How Do I Migrate Logs from a Third-Party Cloud to Huawei Cloud LTS?

If you use the log service of a third-party cloud vendor, a large amount of log data is on the vendor's object storage. You can perform the following operations to migrate the logs to Huawei Cloud:

- Hot logs (for search and analysis): typically stored for 7 to 14 days for application O&M. These logs do not need to be migrated to Huawei Cloud. Directly enable LTS on Huawei Cloud. After your services have run on Huawei Cloud for 14 days, the logs stored by the third-party cloud vendor will be aged. For details, see **Log Management**.
- Archived logs (for security compliance): stored for more than 180 days for compliance audit. Generally, these logs are transferred to object storage. Use a tool to migrate logs from third-party object storage services to Huawei Cloud OBS. For details, see **Migrating Data from Third-Party Cloud Service Vendors to OBS**.

2.2 Can Logs Stored in LTS Be Used for Security Compliance Audit?

Yes. LTS protects logs reported to it from being modified and tampered with. You can store logs in LTS for a long time and query them at any time for security compliance audit.

If the log retention duration of a log group is used, change the duration by referring to **Log Groups**. If the log retention duration of a log stream is used, change the duration by referring to **Log Streams**.

2.3 What Are the Recommended Scenarios for Using LTS?

Cloud Host Application Logs

Scenario description: The following suggestions are applicable when a user application system is deployed on cloud hosts and LTS is used to centrally collect and search for logs. Generally, a user application system consists of multiple components (microservices). Each component is deployed on at least two cloud hosts.

Suggestions:

- Log collection: The log collector ICAgent is recommended. Install ICAgent on the cloud hosts and configure the log collection path by referring to **Collecting Logs from ECS**. ICAgent is completely decoupled from application systems and does not require code modification. You are not advised to use SDKs or APIs to collect logs because this mode is complex and the application system stability may be affected due to improper code compilation.
- Log group planning: Place the logs of an application system in a log group. The name of the log group can be the same as that of the application system.
- Log stream planning:
  - If your logs are irregular, you can collect logs of similar components, for example, Java, PHP, and Python components, to the same log stream. This approach reduces the number of log streams for easier management. If your number of components is small (for example, less than 20), you can collect logs of each component to different log streams.
  - For logs that support structuring parsing, such as the Nginx gateway logs, it is recommended that logs of the same format be collected into the same log stream. A unified log format within a log stream allows for SQL-based analysis in visualized charts.
- Permission isolation: LTS log streams support isolation by enterprise project. By setting enterprise projects for log streams, you can set different log stream access permissions for different Identity and Access Management (IAM) users.

Containerized Application Logs

Scenario description: The following suggestions are applicable when a user application system is deployed on Kubernetes clusters and LTS is used to centrally collect and search for logs. A user application system consists of multiple workloads, each with at least two instances.

Suggestions:

- Log collection:
  - ICAgent is recommended. You can configure the log collection path by referring to **Collecting Logs from CCE**. ICAgent is completely decoupled from application systems and does not require code modification. You are not advised to use SDKs or APIs to collect logs because this mode is

complex and the application system stability may be affected due to improper code compilation.

  - Containerized application logs can be collected as container standard output, container files, node files, and Kubernetes events. Container files are recommended. In contrast to container standard output, container files can be mounted to hosts persistently and the output content can be controlled by users. In contrast to node files, container files collect metadata such as namespaces, workloads, and pods, facilitating log search.

- Log group planning: Place all logs of a CCE cluster in a log group. The log group alias (modifiable) can be the same as the CCE cluster name, and the original log group name (non-modifiable) is recommended to be **k8s-log-**{*cluster ID*}.
- Log stream planning:
  - If your logs are irregular, you can collect logs of similar components, for example, Java, PHP, and Python components, to the same log stream. This approach reduces the number of log streams for easier management. If your number of components is small (for example, less than 20), you can collect logs of each component to different log streams.
  - For logs that support structuring parsing, such as the Nginx gateway logs, it is recommended that logs of the same format be collected into the same log stream. A unified log format within a log stream allows for SQL-based analysis in visualized charts.
- Permission isolation: LTS log streams support isolation by enterprise project. By setting enterprise projects for log streams, you can set different stream access permissions for different IAM users.

Cloud Service Log Analysis

- Ingesting cloud service logs to LTS: LTS can **collect logs from cloud services**. You need to enable the log function on the corresponding cloud service console to collect logs to a specified log group or log stream.
- Optimal status: Many cloud service logs support structuring parsing. You can configure structuring parsing rules for them on the log structuring page. For details, see **Log Structuring**. After structuring parsing, you can use SQL statements to analyze the logs in a visualized manner.

Application Monitoring Alarms

Scenario description: The following suggestions are applicable when logs are used to monitor application systems in real time and detect system faults in advance.

Currently, SQL alarms are available to all users in regions CN South-Guangzhou, CN North-Beijing4, CN East-Shanghai1, CN-Hong Kong, CN Southwest-Guiyang1, AP-Singapore, and CN South-Shenzhen. They are also available to whitelisted users in regions AP-Bangkok, CN North-Beijing1, AP-Jakarta, and CN East-Shanghai2.

Suggestions:

- Alarm statistics mode: LTS supports **keyword alarms** and SQL alarms. For irregular logs such as run logs of Java programs, keyword alarms are applicable. For regular logs such as Nginx gateway logs, SQL alarms are applicable. You can use SQL statements to analyze structuring logs and obtain the required metrics to configure alarms.

- Alarm rule configuration: Generally, alarms need to be triggered as soon as possible. The recommended alarm rule statistics period is 1 minute. You can use the default message templates of LTS to send alarms. If you have personalized requirements, you can modify the default templates and save them as **message templates** for sending alarms.
- Configuring log alarms for key cloud services, such as Elastic Load Balance (ELB) and API Gateway (APIG): ELB is often used as the entry of application systems. To detect system faults in a timely manner, enable ELB logs, collect them to LTS, and configure ELB 5*XX* status code alarms. In addition, you can use the out-of-the-box ELB dashboard to observe the overall success rate of an application system.

Service Operation Analysis

Scenario description: The following suggestions are applicable when you print service logs, such as the transaction amount, customer, and product information, in an application system and then output visualized charts and dashboards using the SQL analysis function of LTS.

Suggestions:

- Log collection mode: You are advised to use ICAgent to collect logs and print them in separate log files. Do not mix the logs with the run logs of applications. You are not advised to use SDKs or APIs to report logs.
- Log structuring parsing: You are advised to use spaces to separate service logs or use the JSON format to quickly configure log structuring parsing rules.
- Log visualization:

  You can **create a custom dashboard** and use SQL-like syntax to analyze service logs that have been structured. You can add **multiple charts** or filters to a custom dashboard and use LTS to analyze your services. This reduces the number of data warehouses to be purchased and simplifies the usage.

2.4 What Are the Advantages of LTS Compared with Self-built ELK Stack?

This section describes the main functions and advantages of Huawei Cloud LTS by comparing it with self-built ELK Stack.

Background

The open-source ELK Stack, comprising Elasticsearch, Logstash, and Kibana, is extensively used for log search, with a variety of content and use cases available within its community.

LTS is a fully managed log analysis platform that covers application O&M, security compliance, and service operations. You can use it to collect, store, query, process, analyze, and report logs with ease. For details, see **Infographics**.

Functions

LTS outperforms ELK in terms of feature completeness and log search and analysis performance.

Feature	Subfeature	LTS	ELK	Description
Log collection	Cloud service log collection	☆☆☆☆☆	N/A	ELK: does not collect cloud service logs. LTS: collects all logs of the cloud service tenant plane.
	VM and container log collection	☆☆☆☆☆	☆☆☆☆	ELK: uses open-source collectors such as Logstash or Filebeat. LTS: uses ICAgent to collect logs and provides easy-to-use wizard pages.
	Collection via multi-language SDKs	☆☆☆	N/A	ELK: not supported. LTS: provides a Java SDK to directly report logs to LTS.
	Host group management (dynamic scaling of hosts)	☆☆☆☆☆	N/A	ELK: not supported. LTS: supports host and host group management. You can add custom identifiers to host groups and scale host groups in or out.
	Log structuring parsing	☆☆☆☆	☆☆☆☆☆	ELK: enables custom structuring parsing based on the collectors. LTS: enables structuring parsing with regular expressions, JSON, delimiters, or custom templates.
Log search	Keyword search, fuzz match, and quick analysis	☆☆☆☆☆	☆☆☆☆☆	ELK and LTS: provide similar keyword search functions.
	Real-time log viewing	☆☆☆☆☆	N/A	ELK: does not provide the page for viewing real-time logs. LTS: provides the page for viewing real-time logs.

Feature	Subfeature	LTS	ELK	Description
	Search of tens of billions of logs in seconds	☆☆☆☆☆	☆☆	ELK: Limited by the server resources, it takes a long time to search for massive logs. LTS: With the extensive scalable computing resources of Huawei Cloud, search results can be returned in 3 seconds.
	Iterative search of hundreds of billions of logs	☆☆☆☆☆	N/A	ELK: Response timeout occurs when hundreds of billions of logs are searched. LTS: Iterative search enables search of hundreds of billions of logs.
	Log management scale	100 PB level	100 TB level	ELK: It is often time-consuming to keep an eye on server scaling. LTS: automatically manages 100 PB of logs. You do not need to worry about the underlying resource consumption and will be charged on a pay-per-use basis.
Log search	SQL analysis	☆☆☆☆☆	☆☆	ELK: does not support nested SQL statements in syntax due to poor SQL performance. LTS: provides high SQL performance and supports nested SQL statements.
Log search	SQL functions	☆☆☆☆☆	☆☆	ELK: supports only basic SQL statistics functions. LTS: Besides basic SQL functions, LTS offers various extended functions, such as IP, statistics, chain and parallel comparison, and URL functions, to support more scenarios.
Log search	Charts	☆☆☆☆	☆☆☆	LTS: provides various visual charts, such as tables, and line, pie, and bar charts.

Feature	Subfeature	LTS	ELK	Description
Log search	Dashboards	☆☆☆☆☆	☆☆	ELK: There is no ready-to-use dashboard for cloud service logs. LTS: provides ready-to-use dashboards for common cloud services, such as ELB, APIG, Document Database Service (DDS), DCS, and Cloud Firewall (CFW).
Log alarms	Keyword and SQL alarms	☆☆☆☆☆	☆	ELK: No log alarm function is available. LTS: Quasi-real-time log keyword and SQL alarms are available.
	Alarm notification channels (such as email, SMS, and HTTPS)	☆☆☆☆☆	☆	ELK: does not send alarms to users through DingTalk, WeCom, or SMS messages. LTS: interconnects with Huawei Cloud Simple Message Notification (SMN) to notify users through channels such as email, SMS, WeCom, DingTalk, Lark, and HTTP.
Log transfer	Transfer to OBS	☆☆☆☆☆	N/A	ELK: cannot transfer logs to OBS directly. LTS: allows you to transfer logs to OBS with simple page configurations.
Log transfer	Transfer to Kafka	☆☆☆☆☆	☆☆	ELK: requires you to deploy a program. LTS: allows you to transfer logs to Kafka in real time with simple page configurations.
Log transfer	Transfer to data warehouses	☆☆☆☆☆	N/A	ELK: cannot transfer logs to data warehouses. LTS: allows you to transfer logs to data warehouses with simple page configurations.

Feature	Subfeature	LTS	ELK	Description
Log jobs	Scheduled SQL jobs	☆☆☆☆☆	N/A	ELK: does not support scheduled SQL jobs. LTS: allows you to configure scheduled SQL jobs to convert raw logs into summarized results.
	Log processing with functions	☆☆☆☆☆	N/A	ELK: does not support log processing. LTS: provides function triggers. You can write custom scripts in FunctionGraph to process logs flexibly.

Costs

Scenario 1:

Assume that a total of 3,000 GB of raw logs are generated in 30 days (100 GB per day) at an average log rate of 1.16 MB/s, and those logs are retained for an average of 30 days in dual storage (primary and standby).

Elasticsearch recommends that the total storage space for raw logs, backup data, and index data in dual storage mode be about 2.2 times the size of raw logs. Considering the Elasticsearch cluster's uneven write distribution and partial disk utilization, storing 3,000 GB of raw logs needs at least 13,200 GB of disk space, calculated as 3,000 GB x 2.2 (storage expansion) x 2 (allowing for 50% disk redundancy).

Typically at least three ECSs (16 vCPUs, 64 GB memory, and 5 TB capacity) and two Kafka replicas are required for Elasticsearch to store logs of the past 12 hours.

Table 2-1 Self-built ELK

Category	Subcategory	Monthly Cost (Total: $1,764 USD)	Expense Proportion
Setting up Elasticsearch	3 Elastic Cloud Servers (ECSs) (C6 16 vCPUs | 64 GB)	3 x 1,999 x 0.1401 = $840 USD	47.6%
	Elastic Volume Service (EVS) (high I/O 15 TB)	0.35 x 15 x 1,024 x 0.1401= $753 USD	42.7%
Setting up Kafka	3 ECSs (2 vCPUs | 4 GB)	3 x 208 x 0.1401 = $87 USD	4.9%

Category	Subcategory	Monthly Cost (Total: $1,764 USD)	Expense Proportion
	EVS (ultra-high I/O 3 x 200 GB)	600 x 0.1401 = $84 USD	4.7%

According to the **price calculator**, the monthly cost of LTS is around $539.89 USD, just **16.7%** of the cost of self-built ELK. This significant saving is attributed to LTS's pay-per-use billing mode, contrasting with the high initial resource cost of ELK in scenarios with few logs.

Scenario 2:

Assume that a total of 7 TB of raw logs are generated in 7 days (1 TB per day) at an average log rate of 11.6 MB/s, and those logs are retained for an average of 30 days in dual storage (primary and standby). Elasticsearch recommends that the total storage space for raw logs, backup data, and index data in dual storage mode be about 2.2 times the size of raw logs. Considering the Elasticsearch cluster's uneven write distribution and partial disk utilization, storing 7 TB of raw logs needs at least 31 TB of disk space, calculated as 7 TB x 2.2 (storage expansion) x 2 (allowing for 50% disk redundancy).

Typically at least three ECSs (16 vCPUs, 64 GB memory, and 10 TB capacity) and two Kafka replicas are required for Elasticsearch to store logs of the past 12 hours.

Table 2-2 Self-built ELK

Category	Subcategory	Monthly Cost (Total: $2,652 USD)	Expense Proportion
Setting up Elasticsearch	3 ECSs (C6 16 vCPUs | 64 GB)	3 x 1,999 x 0.1401 = $840 USD	31.7%
	EVS (high I/O 31 TB)	0.35 x 31 x 1,024 x 0.1401= $1,557 USD	58.7%
Setting up Kafka	3 ECSs (2 vCPUs | 4 GB)	3 x 208 x 0.1401 = $87 USD	3.3%
	EVS (ultra-high I/O 3 x 400 GB)	1,200 x 0.1401 = $168 USD	6.3%

According to the **price calculator**, the monthly cost of LTS is around $3,409.92 USD, just **71%** of the cost of self-built ELK. This significant saving is attributed to LTS's pay-per-use billing mode, contrasting with the extensive disk requirements for maintaining smooth running in self-built ELK clusters.

Scenario 3:

Assume that a total of 150 TB of raw logs are generated in 30 days (5 TB per day) at an average log rate of 58 MB/s, and those logs are retained for an average of 30 days in dual storage (primary and standby).

Elasticsearch recommends that the total storage space for raw logs, backup data, and index data in dual storage mode be about 2.2 times the size of raw logs. Considering the Elasticsearch cluster's uneven write distribution and partial disk utilization, storing 150 TB of raw logs needs at least 660 TB of disk space, calculated as 150 TB x 2.2 (storage expansion) x 2 (allowing for 50% disk redundancy).

Typically at least 66 ECSs (16 vCPUs, 64 GB memory, and 10 TB capacity) and two Kafka replicas are required for Elasticsearch to store logs of the past 12 hours.

Table 2-3 Self-built ELK

Category	Subcategory	Monthly Cost (Total: $52,440 USD)	Expense Proportion
Setting up Elasticsearch	66 ECSs (C6 16 vCPUs | 64 GB)	66 x 1,999 x 0.1401 = $18,489 USD	35.3%
	EVS (high I/O 660 TB)	0.35 x 660 x 1,024 x 0.1401 = $33,149 USD	63.2%
Setting up Kafka	3 ECSs (2 vCPUs | 4 GB)	3 x 208 x 0.1401 = $87 USD	0.2%
	EVS (ultra-high I/O 3 x 1,700 GB)	5,100 x 0.1401 = $715 USD	1.4%

According to the **price calculator**, the monthly cost of LTS is around $27,648 USD, just **28.8%** of the cost of self-built ELK. This significant saving is attributed to LTS's pay-per-use billing mode, contrasting with the extensive disk requirements for maintaining smooth running in self-built ELK clusters.

Summary

LTS beats ELK in functions, performance, and costs. You are advised to use fully managed LTS instead of self-built ELK.

3 Log Management

3.1 How Do I Save Only One Copy of Logs After Multi-Account Log Aggregation Is Configured?

After **Multi-Account Log Center** is configured, logs in the source log group/stream will be aggregated to the target log group/stream, and LTS saves these logs in both accounts. When you create a source log group/stream and a target log group/stream, you set the log retention period. LTS retains logs based on this period. To save a copy of logs, perform the following operations.

Disabling Log Retention

Step 1 Log in to the LTS console using the account to which the source log group/stream belongs.

Step 2 On the **Log Management** page, locate the source log group/stream.

Step 3 Click ✎ in the **Operation** column of the source log stream. On the displayed page, disable **Log Storage** and **Log Retention (Days)**.

Figure 3-1 Modifying a log stream

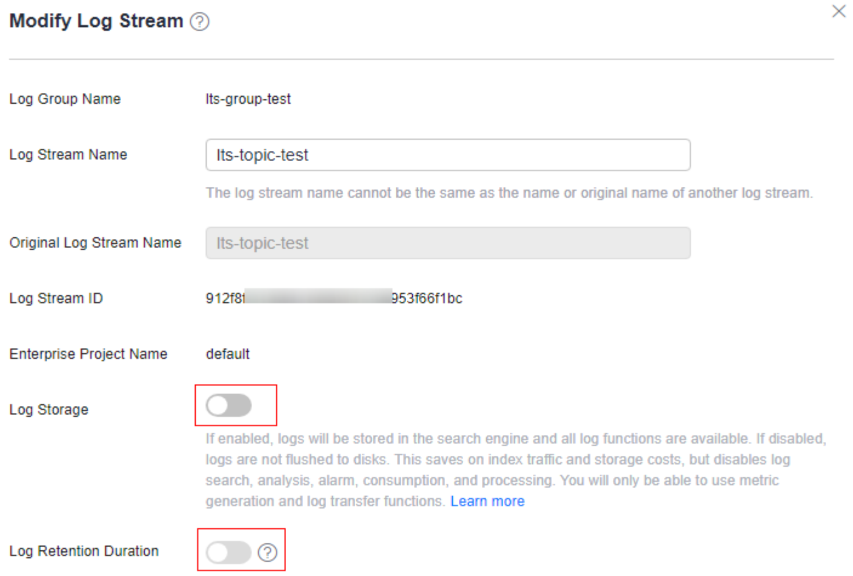

Modify Log Stream

Log Group Name	lts-group-test
Log Stream Name	lts-topic-test The log stream name cannot be the same as the name or original name of another log stream.
Original Log Stream Name	lts-topic-test
Log Stream ID	912f8[redacted]953f66f1bc
Enterprise Project Name	default
Log Storage	If enabled, logs will be stored in the search engine and all log functions are available. If disabled, logs are not flushed to disks. This saves on index traffic and storage costs, but disables log search, analysis, alarm, consumption, and processing. You will only be able to use metric generation and log transfer functions. Learn more
Log Retention Duration	

Step 4 Click **OK**. After **Log Storage** is disabled, logs will not be stored to LTS. This saves on index traffic and storage costs, but disables log search, analysis, alarm, consumption, and processing. You will only be allowed to use metric generation and log transfer functions.

----End

4 Host Management

4.1 What Do I Do If ICAgent Installation Fails in Windows and the Message "SERVICE STOP" Is Displayed?

Background

The ICAgent installation fails and the message **SERVICE STOP** is displayed.

The following issues may occur after the installation fails:

- No ICAgent task exists in the task manager.
- No ICAgent service exists in the system service list.
- When **sc query icagent** is executed in the command line tool (CLI), a message is displayed, indicating that no ICAgent was found.

Possible Causes

The ICAgent registration is blocked by antivirus software, such as 360 Total Security.

Solution

1. Check whether any antivirus software is running. If yes, go to the next step.
2. Stop the antivirus software and install ICAgent again. For details, see **Installing ICAgent**.

4.2 What Do I Do If ICAgent Upgrade Fails on the LTS Console?

ICAgent is installed in overwrite mode. If an upgrade fails, directly run the installation command again and re-upgrade ICAgent.

For details, see **Installing ICAgent**.

4.3 What Do I Do If I Could Not Query New Logs on LTS?

Background

New logs cannot be queried on the LTS console.

Impact on the System

Users cannot query service logs.

Possible Causes

1. **ICAgent Collection** is disabled on the LTS console.
2. ICAgent fails to report logs to LTS.
3. The collection configuration set on the LTS console is incorrectly delivered to ICAgent.

Fault Locating

Perform the following steps:

Step 1 Check whether ICAgent collection is enabled.

1. Log in to the **LTS console**.
2. Choose **Configuration Center**.
3. Click the **ICAgent Collection** tab and check whether ICAgent collection is enabled. If not, enable it.

Step 2 Check the ICAgent collection configuration delivered last time.

1. Log in to the ECS where ICAgent is installed.
2. Run the **cd /var/ICAgent** command to go to the ICAgent log directory.
3. Run the following command to view the log collection configuration.

```
zgrep switchList *.zip //Query the compressed package of transferred logs.
cat oss.icAgent.trace | grep switchList //Query the current log file.
```

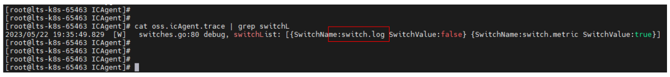

```
[root@lts-k8s-65463 ICAgent]#
[root@lts-k8s-65463 ICAgent]#
[root@lts-k8s-65463 ICAgent]# cat oss.icAgent.trace | grep switchL
2023/05/22 19:35:49.829  [W]   switches.go:80 debug, switchList: [{SwitchName:switch.log SwitchValue:false} {SwitchName:switch.metric SwitchValue:true}]
[root@lts-k8s-65463 ICAgent]#
[root@lts-k8s-65463 ICAgent]#
[root@lts-k8s-65463 ICAgent]#
[root@lts-k8s-65463 ICAgent]#
```

Search for the latest log time configuration from the filtering result and check whether the value of **switch.log** is **true** (indicating that ICAgent collection is enabled).

Step 3 Check whether the value of **switch.log** in the ICAgent collection file is **true** (indicating that ICAgent collection is enabled).

1. Log in to the ECS where ICAgent is installed.
2. Query the ICAgent collection file.

```
cat /var/share/oss/manager/ICProbeAgent/internal/TRACE_CONFIG/swithes_context.json
```

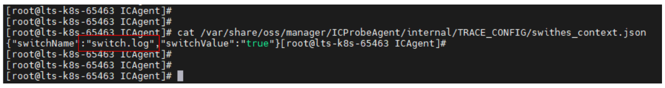

```
[root@lts-k8s-65463 ICAgent]#
[root@lts-k8s-65463 ICAgent]#
[root@lts-k8s-65463 ICAgent]# cat /var/share/oss/manager/ICProbeAgent/internal/TRACE_CONFIG/swithes_context.json
{"switchName":"switch.log","switchValue":"true"}[root@lts-k8s-65463 ICAgent]#
[root@lts-k8s-65463 ICAgent]#
[root@lts-k8s-65463 ICAgent]#
[root@lts-k8s-65463 ICAgent]#
```

The value of **switch.log** in the ICAgent collection file is **true**.

Step 4 Check whether logs fail to be sent.

1. Log in to the ECS where ICAgent is installed.
2. Run the **cd /var/ICAgent** command to go to the ICAgent log directory.
3. Query the log that failed to be sent.

```
cat oss.icAgent.trace | grep httpsend | grep 'dataType:TRACE' | grep failed
zgrep 'dataType:TRACE' *.zip | grep httpsend | grep failed
```

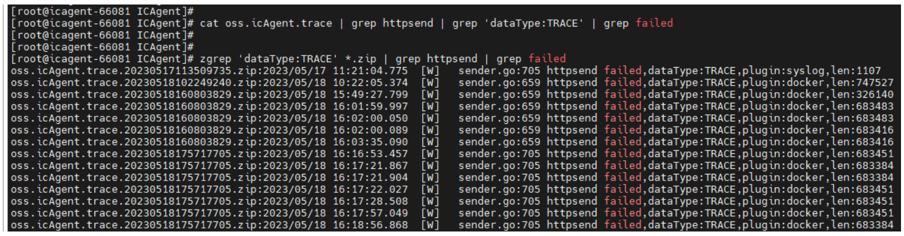

```
[root@icagent-66081 ICAgent]#
[root@icagent-66081 ICAgent]# cat oss.icAgent.trace | grep httpsend | grep 'dataType:TRACE' | grep failed
[root@icagent-66081 ICAgent]#
[root@icagent-66081 ICAgent]#
[root@icagent-66081 ICAgent]# zgrep 'dataType:TRACE' *.zip | grep httpsend | grep failed
oss.icAgent.trace.20230517113509735.zip:2023/05/17 11:21:04.775  [W]   sender.go:705 httpsend failed,dataType:TRACE,plugin:syslog,len:1107
oss.icAgent.trace.20230518102249240.zip:2023/05/18 10:22:05.374  [W]   sender.go:659 httpsend failed,dataType:TRACE,plugin:docker,len:747527
oss.icAgent.trace.20230518160803829.zip:2023/05/18 15:49:27.799  [W]   sender.go:659 httpsend failed,dataType:TRACE,plugin:docker,len:326140
oss.icAgent.trace.20230518160803829.zip:2023/05/18 16:01:59.997  [W]   sender.go:659 httpsend failed,dataType:TRACE,plugin:docker,len:683483
oss.icAgent.trace.20230518160803829.zip:2023/05/18 16:02:00.050  [W]   sender.go:659 httpsend failed,dataType:TRACE,plugin:docker,len:683483
oss.icAgent.trace.20230518160803829.zip:2023/05/18 16:02:00.089  [W]   sender.go:659 httpsend failed,dataType:TRACE,plugin:docker,len:683416
oss.icAgent.trace.20230518160803829.zip:2023/05/18 16:03:35.090  [W]   sender.go:659 httpsend failed,dataType:TRACE,plugin:docker,len:683416
oss.icAgent.trace.20230518175717705.zip:2023/05/18 16:16:53.457  [W]   sender.go:705 httpsend failed,dataType:TRACE,plugin:docker,len:683451
oss.icAgent.trace.20230518175717705.zip:2023/05/18 16:17:21.867  [W]   sender.go:705 httpsend failed,dataType:TRACE,plugin:docker,len:683384
oss.icAgent.trace.20230518175717705.zip:2023/05/18 16:17:21.904  [W]   sender.go:705 httpsend failed,dataType:TRACE,plugin:docker,len:683384
oss.icAgent.trace.20230518175717705.zip:2023/05/18 16:17:22.027  [W]   sender.go:705 httpsend failed,dataType:TRACE,plugin:docker,len:683451
oss.icAgent.trace.20230518175717705.zip:2023/05/18 16:17:28.508  [W]   sender.go:705 httpsend failed,dataType:TRACE,plugin:docker,len:683451
oss.icAgent.trace.20230518175717705.zip:2023/05/18 16:17:57.049  [W]   sender.go:705 httpsend failed,dataType:TRACE,plugin:docker,len:683451
oss.icAgent.trace.20230518175717705.zip:2023/05/18 16:18:56.868  [W]   sender.go:705 httpsend failed,dataType:TRACE,plugin:docker,len:683384
```

If **failed** is displayed, contact LTS technical support.

Step 5 Check whether the collection configuration set on the LTS console has been delivered to ICAgent.

1. Log in to the ECS where ICAgent is installed.
2. Run the **cd /var/ICAgent** command to go to the ICAgent log directory.
3. Filter log files to query the log collection configuration set on the LTS console.

```
cat oss.icAgent.trace | grep 'als event' | grep -v grep | tail -n 1
```

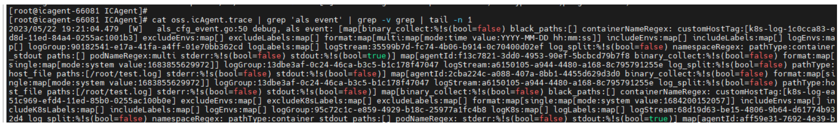

```
[root@icagent-66081 ICAgent]#
[root@icagent-66081 ICAgent]# cat oss.icAgent.trace | grep 'als event' | grep -v grep | tail -n 1
2023/05/22 19:21:04.479  [W]   als_cfg_event.go:50 debug, als event: [map[binary_collect:%!s(bool=false) black_paths:[] containerNameRegex: customHostTag:[k8s-log-1c0cca83-e
d8d-11ed-84a4-0255ac1001b3] excludeEnvs:map[] excludeLabels:map[] format:map[multi:map[mode:time value:YYYY-MM-DD hh:mm:ss]] includeEnvs:map[] includeLabels:map[] logEnvs:ma
p[] logGroup:90182541-e17a-41fa-a4ff-01e70bb362cd logLabels:map[] logStream:35599b7d-fc74-4b06-b914-0c70400d02ef log_split:%!s(bool=false) namespaceRegex: pathType:container
_stdout paths:[] podNameRegex:multi stderr:%!s(bool=false) stdout:%!s(bool=true)] map[agentId:f13c7821-3dd0-4953-90ef-5bcbcd79b7f8 binary_collect:%!s(bool=false) format:map[
single:map[mode:system value:1683855629972]] logGroup:13dbe3af-0c24-46ca-b3c5-b1c178f47047 logStream:a6150105-a944-4480-a168-8c795791255e log_split:%!s(bool=false) pathType:
host_file paths:[/root/test.log] stderr:%!s(bool=false) stdout:%!s(bool=false)] map[agentId:2cba224c-a088-407a-8bb1-4455d629d3d0 binary_collect:%!s(bool=false) format:map[si
ngle:map[mode:system value:1683855629972]] logGroup:13dbe3af-0c24-46ca-b3c5-b1c178f47047 logStream:a6150105-a944-4480-a168-8c795791255e log_split:%!s(bool=false) pathType:ho
st_file paths:[/root/test.log] stderr:%!s(bool=false) stdout:%!s(bool=false)] map[binary_collect:%!s(bool=false) black_paths:[] containerNameRegex: customHostTag:[k8s-log-ea
51c969-efd4-11ed-85b0-0255ac100b0e] excludeEnvs:map[] excludeK8sLabels:map[] excludeLabels:map[] format:map[single:map[mode:system value:1684200152057]] includeEnvs:map[] in
cludeK8sLabels:map[] includeLabels:map[] logEnvs:map[] logGroup:95c72c1c-e859-4929-b18c-25977a1fc4b8 logK8s:map[] logLabels:map[] logStream:68d19d63-be15-4806-9b64-d61774b93
2d4 log_split:%!s(bool=false) namespaceRegex: pathType:container_stdout paths:[] podNameRegex: stderr:%!s(bool=false) stdout:%!s(bool=true)] map[agentId:aff59e31-7692-4e39-b
```

(The preceding figure is for reference only.) If the log group, log stream, and collection path you set appear in the filtering result, the log collection configuration has been delivered to ICAgent.

4. If the log file cannot be found, query the compressed log package.

```
zgrep 'als event' *.zip | grep -v grep | tail -n 1
```

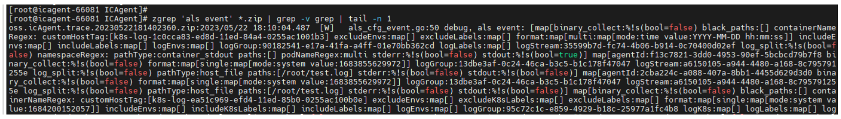

```
[root@icagent-66081 ICAgent]#
[root@icagent-66081 ICAgent]# zgrep 'als event' *.zip | grep -v grep | tail -n 1
oss.icAgent.trace.20230522181402360.zip:2023/05/22 18:10:04.487  [W]   als_cfg_event.go:50 debug, als event: [map[binary_collect:%!s(bool=false) black_paths:[] containerName
Regex: customHostTag:[k8s-log-1c0cca83-ed8d-11ed-84a4-0255ac1001b3] excludeEnvs:map[] excludeLabels:map[] format:map[multi:map[mode:time value:YYYY-MM-DD hh:mm:ss]] includeE
nvs:map[] includeLabels:map[] logEnvs:map[] logGroup:90182541-e17a-41fa-a4ff-01e70bb362cd logLabels:map[] logStream:35599b7d-fc74-4b06-b914-0c70400d02ef log_split:%!s(bool=f
alse) namespaceRegex: pathType:container_stdout paths:[] podNameRegex:multi stderr:%!s(bool=false) stdout:%!s(bool=true)] map[agentId:f13c7821-3dd0-4953-90ef-5bcbcd79b7f8 bi
nary_collect:%!s(bool=false) format:map[single:map[mode:system value:1683855629972]] logGroup:13dbe3af-0c24-46ca-b3c5-b1c178f47047 logStream:a6150105-a944-4480-a168-8c795791
255e log_split:%!s(bool=false) pathType:host_file paths:[/root/test.log] stderr:%!s(bool=false) stdout:%!s(bool=false)] map[agentId:2cba224c-a088-407a-8bb1-4455d629d3d0 bina
ry_collect:%!s(bool=false) format:map[single:map[mode:system value:1683855629972]] logGroup:13dbe3af-0c24-46ca-b3c5-b1c178f47047 logStream:a6150105-a944-4480-a168-8c79579125
5e log_split:%!s(bool=false) pathType:host_file paths:[/root/test.log] stderr:%!s(bool=false) stdout:%!s(bool=false)] map[binary_collect:%!s(bool=false) black_paths:[] conta
inerNameRegex: customHostTag:[k8s-log-ea51c969-efd4-11ed-85b0-0255ac100b0e] excludeEnvs:map[] excludeK8sLabels:map[] excludeLabels:map[] format:map[single:map[mode:system va
lue:1684200152057]] includeEnvs:map[] includeK8sLabels:map[] includeLabels:map[] logEnvs:map[] logGroup:95c72c1c-e859-4929-b18c-25977a1fc4b8 logK8s:map[] logLabels:map[] log
```

If the log group, log stream, and collection path you set appear in the compressed log package, the log collection configuration has been delivered to ICAgent.

5. Log in to the LTS console and choose **Log Ingestion** in the navigation pane. Check the log collection configurations on the log ingestion setting page.

----End

4.4 What Do I Do If ICAgent Restarts Repeatedly After Being Installed?

Background

ICAgent restarts repeatedly after being installed.

Impact on the System

ICAgent metrics/logs cannot be collected.

Possible Causes

- The resource usage exceeds the limit of ICAgent.
- ICAgent is abnormal.

Fault Locating

Step 1 Log in to the ECS where ICAgent is installed.

Step 2 Run the **cd /var/ICAgent** command to go to the ICAgent log directory.

Step 3 Query the current log file to check whether the resource usage is abnormal.

```
cat oss.icAgent.trace| grep 'icagent exit'
```

- If any command output is displayed, there is a resource abnormality. Troubleshoot it.
- If no command output is displayed, the resource usage is normal.

Step 4 Query historical log files to check whether the resource usage is abnormal.

```
zgrep 'icagent exit' *.zip
```

- If any command output is displayed, there is a resource abnormality. Troubleshoot it.
- If no command output is displayed, the resource usage is normal.

Step 5 Query the ICAgent startup log file to check whether ICAgent is abnormal.

```
cat oss.script.trace | grep runtime
```

- If any command output is displayed, there is a ICAgent abnormality. Troubleshoot it.
- If no command output is displayed, ICAgent is running properly.

Step 6 Query the historical ICAgent startup log files to check whether ICAgent is abnormal.

```
zgrep runtime oss.script.*.zip
```

- If any command output is displayed, there is a ICAgent abnormality. Troubleshoot it.
- If no command output is displayed, ICAgent is running properly.

----End

4.5 What Do I Do If ICAgent Is Displayed as Offline on the LTS Console After Installation?

In this case, perform the following steps:

Step 1 Check whether the ICAgent network connection is normal.

1. Log in to the host where ICAgent is installed.

2. Run the **netstat -nap | grep icagent** command to check whether the network connection status of the ICAgent process is **ESTABLISHED**. If yes, the network connection is normal.

   Check whether the connection status of ports 30200 and 30201 on the server is **ESTABLISHED**. If not, check whether these ports are allowed in the security group.

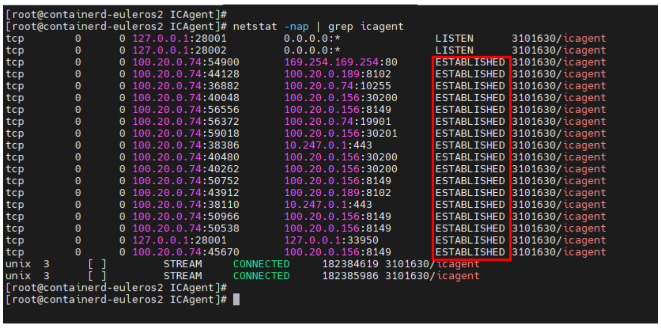

```
[root@containerd-euleros2 ICAgent]#
[root@containerd-euleros2 ICAgent]# netstat -nap | grep icagent
tcp        0      0 127.0.0.1:28001         0.0.0.0:*               LISTEN      3101630/icagent
tcp        0      0 127.0.0.1:28002         0.0.0.0:*               LISTEN      3101630/icagent
tcp        0      0 100.20.0.74:54900       169.254.169.254:80      ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:44128       100.20.0.189:8102       ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:36882       100.20.0.74:10255       ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:40048       100.20.0.156:30200      ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:56556       100.20.0.156:8149       ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:56372       100.20.0.74:19901       ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:59018       100.20.0.156:30201      ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:38386       10.247.0.1:443          ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:40480       100.20.0.156:30200      ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:40262       100.20.0.156:30200      ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:50752       100.20.0.156:8149       ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:43912       100.20.0.189:8102       ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:38110       10.247.0.1:443          ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:50966       100.20.0.156:8149       ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:50538       100.20.0.156:8149       ESTABLISHED 3101630/icagent
tcp        0      0 127.0.0.1:28001         127.0.0.1:33950         ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:45670       100.20.0.156:8149       ESTABLISHED 3101630/icagent
unix  3      [ ]         STREAM     CONNECTED     182384619 3101630/icagent
unix  3      [ ]         STREAM     CONNECTED     182385986 3101630/icagent
[root@containerd-euleros2 ICAgent]#
[root@containerd-euleros2 ICAgent]#
```

Step 2 Check whether the ICAgent authentication is successful.

1. Run the **cd /var/ICAgent** command to go to the ICAgent log directory.
2. Run the **zgrep 'msworkflow' * | grep 'retCode\['** command to query the ICAgent authentication log.
   - If **200** is returned, the authentication is successful. If ICAgent is still offline, contact technical support.
   - If the return code is not 200, check whether the AK/SK pair entered during ICAgent installation is correct. If not, go to the next step.

Step 3 Obtain the correct AK/SK pair and install ICAgent again. For details, see **How Do I Obtain an AK/SK Pair?**.

----End

4.6 What Do I Do If I Do Not See a Host with ICAgent Installed on the LTS Console?

If a host with ICAgent installed is not displayed on the **Hosts** tab page on the LTS console, perform the following steps:

- If you are configuring ECS log ingestion and do not see the host with ICAgent installed on the **Hosts** tab page:
  a. On the **Install ICAgent** page, ensure that the installation command is correctly copied. Do not use the installation command across regions.
  b. Ensure that the obtained AK/SK pair is correct and has not been deleted.
  c. Run the **netstat -nap | grep icagent** command to check whether the network connection status of the host is **ESTABLISHED**. If yes, the network connection is normal.

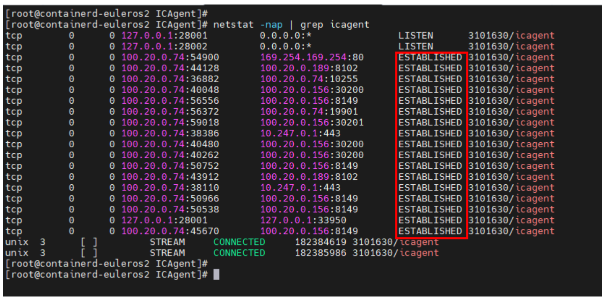

```
[root@containerd-euleros2 ICAgent]#
[root@containerd-euleros2 ICAgent]# netstat -nap | grep icagent
tcp        0      0 127.0.0.1:28001         0.0.0.0:*               LISTEN      3101630/icagent
tcp        0      0 127.0.0.1:28002         0.0.0.0:*               LISTEN      3101630/icagent
tcp        0      0 100.20.0.74:54900       169.254.169.254:80      ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:44128       100.20.0.189:8102       ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:36882       100.20.0.74:10255       ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:40048       100.20.0.156:30200      ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:56556       100.20.0.156:8149       ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:56372       100.20.0.74:19901       ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:59018       100.20.0.156:30201      ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:38386       10.247.0.1:443          ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:40480       100.20.0.156:30200      ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:40262       100.20.0.156:30200      ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:50752       100.20.0.156:8149       ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:43912       100.20.0.189:8102       ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:38110       10.247.0.1:443          ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:50966       100.20.0.156:8149       ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:50538       100.20.0.156:8149       ESTABLISHED 3101630/icagent
tcp        0      0 127.0.0.1:28001         127.0.0.1:33950         ESTABLISHED 3101630/icagent
tcp        0      0 100.20.0.74:45670       100.20.0.156:8149       ESTABLISHED 3101630/icagent
unix  3      [ ]         STREAM     CONNECTED     182384619 3101630/icagent
unix  3      [ ]         STREAM     CONNECTED     182385986 3101630/icagent
[root@containerd-euleros2 ICAgent]#
[root@containerd-euleros2 ICAgent]#
```

- If you are configuring CCE log ingestion and do not see the host with ICAgent installed on the **Hosts** tab page:

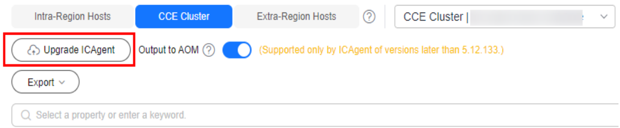

CCE Cluster | [redacted]

Output to AOM ⓘ (Supported only by ICAgent of versions later than 5.12.133.)

Select a property or enter a keyword.

a. Ensure that ICAgent has been installed in the CCE cluster.

b. If ICAgent is not installed, click **CCE Cluster** on the **Hosts** tab page and click **Upgrade ICAgent**. For details, see **Upgrading ICAgent**.

- If the issue persists after you have tried the methods above, **submit a service ticket**.

4.7 How Do I Create a VPC Endpoint on the VPCEP Console?

VPC endpoints are secure and private channels for connecting Virtual Private Clouds (VPCs) to VPC endpoint services. You can create VPC endpoints to connect the network between AOM and LTS for heartbeat, metric, or log reporting.

You need to purchase a VPC endpoint for AOM and LTS, respectively.

Creating an AOM/LTS endpoint

1. Log in to the Huawei Cloud VPC Endpoints (VPCEP) console to go to the **VPC Endpoints** page.
2. Click **Buy VPC Endpoint**.
3. On the displayed page, select the region where the endpoint is located, set **Service Category** to **Cloud service**, search for and select **AOM** from **Service List**, select **Create a Private Domain Name**, select the VPC and subnet where the endpoint is located, and retain the default values of other parameters. For more operations, see **Buying a VPC Endpoint**.

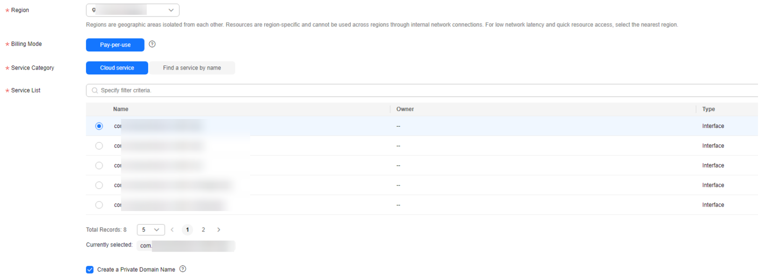

* Region: [redacted]

Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.

* Billing Mode: Pay-per-use

* Service Category: Cloud service

* Service List: Specify filter criteria.

	Name	Owner	Type
●	co[redacted]	--	Interface
○	co[redacted]	--	Interface
○	co[redacted]	--	Interface
○	co[redacted]	--	Interface
○	co[redacted]	--	Interface

Total Records: 8 5 < 1 2 >

Currently selected: com.[redacted]

☑ Create a Private Domain Name

4. Click **Next**. After the purchase, an AOM endpoint is created.
5. Repeat the preceding steps to purchase another VPC endpoint. Select the region where the endpoint is located, set **Service Category** to **Cloud service**, search for and select **LTS** from **Service List**, select **Create a Private Domain Name**, select the VPC and subnet where the endpoint is located, and retain the default values of other parameters.
6. Click **Next**. After the purchase, an LTS endpoint is created.

4.8 How Do I Obtain an AK/SK Pair?

An access key ID and secret access key (AK/SK) constitute an access key.

- AK: access key ID, which is a unique identifier used in conjunction with a secret access key to sign requests cryptographically.
- SK: secret access key used in conjunction with an AK to sign requests cryptographically. It identifies a request sender and prevents the request from being modified.

Obtain and use the AK/SK pair of a public account.

- Each user can create up to two AK/SK pairs. Once they are generated, they are permanently valid.
- Ensure that the public account and AK/SK pair will not be deleted or disabled. If the AK/SK pair is deleted, ICAgent cannot report data to LTS.

Creating an Access Key

Step 1 Log in to Huawei Cloud and click **Console** in the upper right corner.

Figure 4-1 Accessing the console

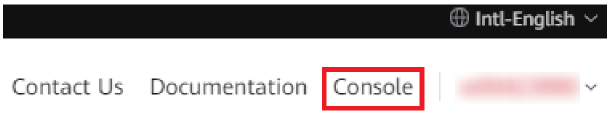

Step 2 On the management console, hover over the username in the upper right corner and choose **My Credentials** from the drop-down list.

Figure 4-2 Choosing My Credentials

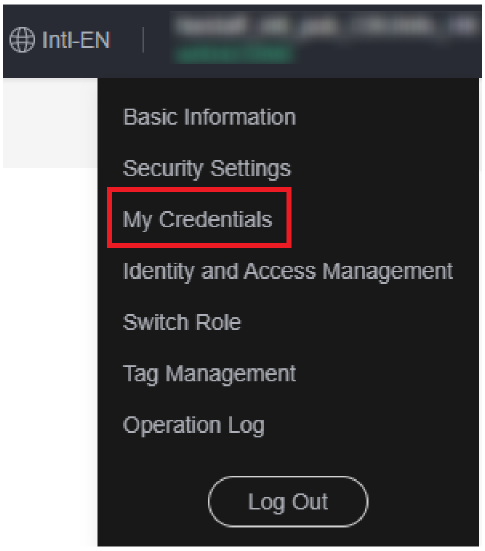

Step 3 Choose **Access Keys** from the navigation pane.

Step 4 Click **Create Access Key**.

Figure 4-3 Creating an access key

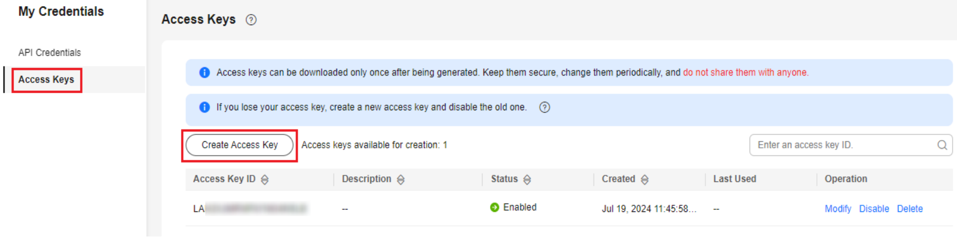

My Credentials

API Credentials

Access Keys

Access Keys

Access keys can be downloaded only once after being generated. Keep them secure, change them periodically, and do not share them with anyone.

If you lose your access key, create a new access key and disable the old one.

Create Access Key   Access keys available for creation: 1   Enter an access key ID.

Access Key ID	Description	Status	Created	Last Used	Operation
LA[redacted]	--	Enabled	Jul 19, 2024 11:45:58...	--	Modify Disable Delete

NOTE

- You can create a maximum of **two** access keys. **The quota cannot be increased**. If you already have two access keys, you can only delete an access key and create a new one.
- To change an access key, delete it and create a new one.
- For newly created access keys, the last used time is the same as the creation time, but will change the next time you use them.

Step 5 Download the access key file.

After the access key is created, view the access key ID (AK) in the access key list in the access key list and view the secret access key (SK) in the downloaded CSV file.

NOTE

- Download the access key file and keep it properly. If the download page is closed, you will not be able to download the access key. However, you can create a new one.
- Open the CSV file in the lower left corner, or choose **Downloads** in the browser and open the CSV file.
- Keep your access keys secure and change them periodically for security purposes. To change an access key, delete it and create a new one.

----End

4.9 How Do I Install ICAgent by Creating an Agency?

When installing ICAgent, you can create an IAM agency, and ICAgent will automatically obtain an AK/SK pair and generate the ICAgent installation command.

Creating an Agency

1. Log in to the IAM console.
2. In the navigation pane, choose **Agencies**.
3. Click **Create Agency** in the upper right corner and create an agency by referring to **Cloud Service Agency**. Set the parameters as follows:
   a. **Agency Type**: Select **Cloud service**.
   b. **Cloud Service**: Select **Elastic Cloud Server (ECS) and Bare Metal Server (BMS)**.
   c. **Validity Period**: Select **Unlimited**.
   d. Permissions: Select both **LTS Administrator** and **APM Administrator**. The authorization takes effect 15 to 30 minutes later.
   e. **Scope**: Select **Region-specific projects**.
4. After the agency is created, make its settings take effect.
   a. Log in to the ECS console.
   b. Click the ECS where ICAgent is to be installed. The ECS details page is displayed.

      You can set an agency when purchasing an ECS: Click **Buy ECS** on the ECS console. In the **Configure Advanced Settings** step, set **Advanced Options** to **Configure now** and select an agency from the **Agency** drop-down list. Set other parameters and click **Submit**.
   c. In the **Management Information** area, click ✎ next to **Agency**, select the created agency, and click ✓ for the agency to take effect. For more information about agencies, see **Account Delegation**.

5 Log Ingestion

5.1 What Do I Do If LTS Cannot Collect Logs After I Configure Host Log Ingestion?

If LTS cannot collect logs after log ingestion has been configured for the host, perform the following operations:

- Wait for 1 to 5 minutes if you just completed configuring the ingestion. It takes a moment before log reporting begins.
- Check whether the host collection path has been configured in AOM. If a collection path has been configured in AOM, do not configure it in LTS.
- Check whether any ingestion settings are improper by referring to **Ingesting ECS Text Logs to LTS**.
- If the issue persists after you have tried the methods above, **submit a service ticket**.

5.2 Will LTS Stop Collecting Logs After the Free Quota Is Used Up If I Disable "Continue to Collect Logs When the Free Quota Is Exceeded" in AOM?

Yes.

The enablement/disablement status of log collection is synchronized between LTS and AOM. Likewise, if you set the log collection to be stopped when the free quota is used up in AOM, the setting is also applied to LTS.

For details, see **Enabling or Disabling Log Collection Beyond Free Quota**.

5.3 What Do I Do If the CPU Usage Is High When ICAgent Is Collecting Logs?

If the CPU usage is high when ICAgent is collecting logs (causing issues like slow running or program breakdown), check whether there are a large number of logs in the log collection path. Clear logs regularly to reduce system resource occupation during log collection. If the fault persists, contact technical support.

5.4 What Kinds of Logs and Files Does LTS Collect?

Logs That Can Be Collected by LTS:

- Host logs. ICAgent should be installed on the target hosts for log collection. To collect logs from a Windows host, enable **Collect Windows Event Logs** when configuring collection paths.
- Cloud service logs. To collect logs from cloud services such as ELB and VPC, enable log reporting to LTS on their consoles.
- Logs reported by APIs.

Files That Can Be Collected by LTS:

If the collection path is set to a directory, for example, **/var/logs/**, only .log, .trace, and .out files in the directory are collected. If the collection path is set to a file name (only text files are supported), the specified file is collected. Note that LTS only collects logs generated in the last 7 days (depending on the local time zone of the host).

5.5 How Do I Disable the Function of Collecting CCE Standard Output Logs to AOM on the LTS Console?

Symptom

As the products evolve, collecting CCE standard output logs to AOM is no longer recommended. You can disable this function on the LTS console as needed. You are advised to collect CCE standard output logs to LTS for unified log management.

Only when the collection of CCE standard output logs to AOM is disabled, the collection of CCE standard output logs to LTS configured on the LTS console will take effect.

Solution

Step 1 Log in to the **LTS console**.

Step 2 Choose **Host Management** > **Hosts** in the navigation pane.

Step 3 On the **CCE Clusters** tab page, select the target CCE cluster and disable **Output to AOM**. You are advised to collect CCE logs to LTS. For details, see **Ingesting CCE Application Logs to LTS**.

Figure 5-1 Disabling output to AOM

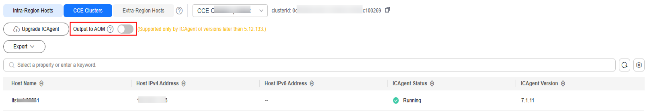

CCE C[redacted] clusterId: 0c[redacted]c100269

Output to AOM (Supported only by ICAgent of versions later than 5.12.133.)

Host Name	Host IPv4 Address	Host IPv6 Address	ICAgent Status	ICAgent Version
lts[redacted]1	1[redacted]5	--	Running	7.1.11

Step 4 Click **OK**. After ICAgent is restarted, CCE standard output to AOM is disabled.

----End

5.6 What Log Rotation Scheme Should I Use for ICAgent to Collect Logs?

Log rotation, also known as log splitting, is used to control the size of a log file. As software system runs for a long time without interrupting services, a large amount of information will be recorded in different logs. As time goes by, the limited disk space will be insufficient to accommodate the increasing log volume. In this case, the size of log files needs to be controlled.

Log rotation can be implemented by time or by log size.

- **By time**: Logs are rotated based on time. When the time when a log is generated reaches the specified time threshold, the log will be rotated. For example, the **/var/log/messages** logs are split based on the rule of rotating once every seven days.
- **By log size**: Logs are rotated based on the log size. When the log size reaches the specified value, the log will be rotated. This mode is applicable to application logs.

When using log rotation, you are advised to:

- **Methods:**

  ICAgent does not rotate your logs. You are advised to use mature software packages to customize log rotation rules in applications, such as Java logback, log4j2, Python logging, and Linux logrotate. If the size of a typical configuration log file exceeds 100 MB, 50 MB, or 20 MB, the log file is rotated once and 10 to 20 historical log files are saved.
- **Suggestions on naming rotated log files:**

  Assume that your log file path is **/your/log/path/**/*.log**. You are advised to name the rotated file as **/your/log/path/**/*.xxx.log**. *xxx* refers to the date according to user habits, for example, 20240103. It cannot contain letters.

  Custom rotation rule: If your naming rules of rotated log files do not comply with the preceding suggestion, rotated log files may be collected repeatedly. You can customize rotation rules to avoid this problem. You can add a custom rotation rule for each log collection path for matching the names of the rotated files. A file whose name matches the rule is identified as a rotated log file and will not be collected repeatedly. For example, if your log file is **/your/log/path/**/app1.log** and the rotated file is **/your/log/path/**/app1.20240103.biz.log**, you can set the rotation rule to *{basename}***\.[0-9-\.]+\.[0-9]+\.biz\.log**.
- **Compress rotated log files is not recommended.**

  If your log printing rate is high, the log file will be rotated quickly. In this case, a few logs may not be collected at the end of the rotated file. ICAgent uses the Linux inode to identify the uniqueness of the collected file. If the rotated log file is a compressed file, the inode changes. As both the log file name and inode are changed, ICAgent may not be able to collect the logs that are not collected at the end of the rotation file.

Assume that the rotated file is compressed and the file name is **/your/log/path/**/*.log.*xxx*.zip**. If both the file name and inode change, ICAgent may not be able to collect the logs that are not collected at the end of the rotation file.

5.7 Does LTS Use the Log4j Plug-in to Report Logs?

The Log4j plug-in cannot be used to report logs to LTS due to the following reasons:

- The Log4j official website does not provide subsequent maintenance for Log4j. You are advised to use Log4j 2, which can be used to report logs to LTS.
- Due to the severe security vulnerabilities of Log4j, Log4j log collection is no longer supported.

5.8 How Long Does It Take to Generate Logs After Configuring Log Ingestion?

After configuring log ingestion on the **Log Ingestion** page of the LTS console, click the target log group on the **Log Management** page to access the details page, choose the corresponding log stream, and click the **Real-Time Logs** tab. If real-time logs are displayed, log ingestion is successful.

Wait for 1 to 5 minutes. You can then view the reported raw logs on the LTS console.

5.9 What Do I Do If LTS Cannot Collect Logs After I Configure Log Ingestion with ICAgent?

After configuring ICAgent for log ingestion, wait for 1 to 5 minutes for logs to be reported. If logs are still not visible on the LTS console, perform the following operations.

Procedure

Step 1 Log in to the LTS console, choose **Configuration Center** in the navigation pane, and check whether **ICAgent Collection** on the **ICAgent Collection** tab page is enabled. If it is disabled, ICAgent does not collect logs. Enable it to start log collection.

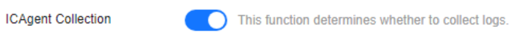

ICAgent Collection ● This function determines whether to collect logs.

Step 2 Choose **Log Ingestion** > **Ingestion Management** in the navigation pane. On the page displayed, check whether the switch in the **Status** column of the ingestion task is toggled on.

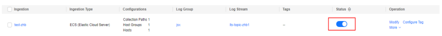

□ Ingestion	Ingestion Type	Configurations	Log Group	Log Stream	Tags	Status	Operation
□ test-zhb	ECS (Elastic Cloud Server)	Collection Paths 1 Host Groups 1 Hosts 1	jsx	lts-topic-zhb1	--	●	Modify Configure Tag More

Step 3 Check whether the log collection path is repeatedly configured. If it is, you need to enable **Allow Repeated File Collection** when configuring log ingestion. Or, delete the duplicate collection path.

Step 4 Check whether real-time logs are reported.

1. Log in to the LTS console and choose **Log Management** in the navigation pane.
2. Click the target log stream to access its details page.
3. Click the **Real-time Logs** tab and do not click any other buttons on this tab page. Trigger log reporting and check whether log data automatically appears on this tab page. If yes, the reporting is successful. If no data appears after a significant wait, proceed to the next step to check ICAgent.

Step 5 Check the ICAgent status.

Choose **Host Management** > **Hosts** in the navigation pane. On the page displayed, check the ICAgent status of the host.

- If the ICAgent status is **Running**, ICAgent is running properly. If a new version is available, upgrade ICAgent to the latest version.
- If the ICAgent status is abnormal, contact technical support.

Step 6 Choose **Log Ingestion** > **Ingestion Management** in the navigation pane. Locate the target ingestion task and choose **More** > **ICAgent Collect Diagnosis** in the **Operation** column to check the exception monitoring, overall status, and collection monitoring of ICAgent. Then, address any issues as needed.

Step 7 If the issue persists after you have tried the methods above, **submit a service ticket**.

----End

6 Log Search and Analysis

6.1 How Often Is the Data Loaded in the Real-Time Log View in LTS?

Generally, the real-time logs are loaded to the LTS console every 5 seconds. However, if no data is generated in a 5-second interval, no new data will be displayed. Log data will be updated in the next 5 seconds if there is new data coming in that interval.

6.2 What Do I Do If I Cannot View Reported Logs in LTS?

Symptom

Logs reported to LTS are not displayed on the LTS console.

Possible Causes

- ICAgent has not been installed.
- The collection path is incorrectly configured.
- **ICAgent Collection** in **Configuration Center** on the LTS console is disabled.
- You set the log collection to be stopped when the free quota is used up on the LTS console.
- Log collection was stopped because your account is in arrears.
- The rate of writing logs into log streams or length of single-line logs exceeds what is supported.
- The browser has slowed down because of the amount of log data.

Solution

- If ICAgent has not been installed, install it. For details, see **Installing ICAgent**.
- If the collection path is incorrectly configured, modify it. If the collection path is set to a directory, for example, **/var/logs/**, only **.log**, **.trace**, and **.out** files in

the directory are collected. If the collection path is set to a text file name, that file is directly collected. For details about log collection paths, see **Collecting Logs from Hosts**.

- If **ICAgent Collection** is disabled, log in to the LTS console, choose **Configuration Center** > **ICAgent Collection**, and enable **ICAgent Collection**.
- Log usage, including log read/write, indexing, and storage, is billed in LTS. If you have disabled **Continue to Collect Logs When the Free Quota Is Exceeded** on the LTS console, log collection will be stopped when the free quota is used up. Log read/write and indexing will no longer be available. No fees will be incurred for log read/write and indexing. To resume log collection, enable **Continue to Collect Logs When the Free Quota Is Exceeded** on the LTS console. For details, see **Configuration Center**.
- If log collection was stopped because your account is in arrears, top up your account. For details, see **Making Repayments (Prepaid Direct Customers)**.
- If the rate of writing logs into log streams or length of single-line logs exceeds the limit, or the browser has slowed down because of the amount of log data, use Google Chrome or Firefox to query logs.
- If the issue persists after you have tried the methods above, **submit a service ticket**.

6.3 Can I Manually Delete Logs on the LTS Console?

No. Manual deletion is not supported. LTS automatically deletes logs when the retention duration ends.

For details about how to set the log retention duration, see **Managing Log Streams**.

6.4 What Do I Do If I Could Not Search for Logs on LTS?

When you search for logs on the LTS console, if a message is displayed indicating that the query result is inaccurate, too many log results are matched, or field *XXX* without indexing configured cannot be queried, perform the following operations:

Message Displayed During Query: Inaccurate Query Results

- Possible causes: There are too many logs in the query time range, and not all logs are displayed.
- Solution: Click the query button multiple times until you obtain all logs, or shorten the query time range and query again.

Too Many Logs Results from Log Query

- Possible cause: Only searches with phrase **#"value"** can ensure the sequence of keywords. For example, if the query statement is **abc def**, logs that contain either **abc** or **def** and logs that contain the phrase **abc def** will be matched.
- Solution: Use phrase **#"abc def"** to accurately match logs containing phrase **abc def**. For details, see **Search Syntax**.

Expected Logs Not Obtained with Specific Search Statements and No Error Message Displayed

- Possible causes: Search delimiters are not supported, or ***** and **?** in a search statement are regarded as common characters and are not used as wildcards.
- Solution: Use the correct query statement.

Message Displayed During Query: Field *XXX* Is Not Configured with Indexing and Cannot Be Queried

- Possible cause: No field indexing is configured.
- Solution: Create an index for field *XXX* on the **Index Settings** tab page and run the query statement again. For details, see **Index Settings**.

Message Displayed During Query: Full-Text Index Not Enabled and content Field and Full-Text Query Unsupported

- Possible cause: Full-text indexing is disabled.
- Solution: Enable **Index Whole Text** on the **Index Settings** tab page and run the query statement again. For details, see **Index Settings**.

Message Displayed During Query: Do Not Start with an Asterisk (*) or a Question Mark (?)

- Possible cause: An asterisk (*) or question mark (?) is placed before the query statement.
- Solution: Modify the query statement or use a correct delimiter.

Message Displayed During Query: long and float Fields Do Not Support Fuzzy Query Using Asterisks (*) or Question Marks (?)

- Possible cause: An asterisk (*) or question mark (?) is used to query fields of the long and float types.
- Solution: Modify the query statement and use operators (>=<) or IN syntax for range query.

Message Displayed During Query: string Fields Do Not Support Range Query Using the Operator (>=<) or IN Syntax

- Possible cause: Operators (>=<) or the IN syntax is used to query fields of the string type.
- Solutions:
  a. Modify the query statement and use the asterisk (*) or question mark (?) to perform fuzzy query.
  b. Change the value of this field to a number. For details, see **Structuring Modes**.

Message Displayed During Query: The Search Syntax Is Incorrect and the Query Statement Needs to Be Modified

- Possible cause: The syntax of the operator is incorrect.

  Solution: Each operator has its syntax rule. Modify the search statement based on the rule. For example, operator = requires that the value on the right must be digits.

- Possible cause: The search statement contains syntax keywords.

  Solution: If the log to search contains syntax keywords, the search statement must be enclosed in double quotation marks to convert the keywords into common characters. For details, see **Search Syntax**. For example, if **and** is a syntax keyword, change the query statement **field:and** to **field:"and"**.

7 Log Transfer

7.1 Does LTS Delete Logs That Have Been Transferred to OBS Buckets?

The log transfer function only forwards existing logs and does not delete them. LTS deletes retained logs once the retention period is over, but the logs that have been transferred to other services are not affected.

During log transfer, logs are "replicated" from LTS to OBS. To view transferred log files, click the name of the corresponding OBS bucket on the **Log Transfer** page of the LTS console, and you will be directed to the OBS console to check the files.

The operations for checking data transferred to Data Ingestion Service (DIS) and Distributed Message Service (DMS) are the same.

7.2 What Are the Common Causes of LTS Log Transfer Abnormalities?

If a log transfer task is in the abnormal state on the **Log Transfer** page of the LTS console:

- Possible cause: The OBS bucket policy is incorrect.

  Solution: Go to the OBS console to correct the settings. For details, see **Configuring a Bucket Policy**.

- Possible cause: The Kafka cluster has been deleted.

  Solution: Re-configure the Kafka transfer. For details, see **Transferring Logs to DMS**.

- Possible cause: The Kafka topic has been deleted.

  Solution: Re-configure the Kafka transfer or try another topic. For details, see **Transferring Logs to DMS**.

7.3 How Do I Transfer CTS Logs to an OBS Bucket?

When Cloud Trace Service (CTS) is connected to LTS, a log group and log stream are automatically created for CTS on the LTS console. To transfer CTS logs to OBS, do as follows:

1. Log in to the CTS console and choose **Tracker List** in the navigation pane.
2. Click **Configure** in the row of the tracker **system**.
3. On the **Basic Information** page, click **Next**.
4. In the **Configure Transfer** step, configure parameters of log transfer to OBS, enable **Transfer to LTS**, and click **Next**.
5. Confirm the configurations and click **Configure**.
6. Access the LTS console, choose **Log Transfer** in the navigation pane on the left, and click **Configure Log Transfer** in the upper right corner.

   Set **Log Group Name** to **CTS** and **Log Stream Name** to **system-trace**. Specify other parameters and click **OK** to transfer CTS logs to the selected OBS bucket.
7. View the transferred CTS logs in the specified OBS bucket on the OBS console.

7.4 What Do I Do If I Cannot View Historical Data in an OBS Bucket After Transferring Data from LTS to OBS?

This occurs because LTS only transfers logs generated after you configure the transfer task to OBS buckets. Logs that already exist before the configuration will not be transferred.

For more information, see **Transferring Logs to OBS**.

7.5 What Do I Do If I Cannot Find a New Partition in a DLI Table After Logs Are Transferred to DLI?

If a new partition cannot be found in a DLI table after logs are transferred to DLI, perform the following operations:

Method 1: Update the Table Partition Information of the Metabase

Step 1 Log in to the DLI console.

Step 2 In the navigation pane, choose **Job Management** > **SQL Jobs**.

Step 3 Click **Edit** in the **Operation** column of the corresponding queue. On the details page that is displayed, enter the **MSCK REPAIR TABLE *table_name*** command in the text box. *table_name* indicates the name of the partitioned table. For example, update the partitioned table named **lts_qpg_dli**.

Step 4 Click **Execute** and wait until the update is complete.

For more information about DLI, see **SQL Job Management**.

----End

Method 2: Configuring a Scheduled Task

Step 1 Log in to the DataArts Studio console and ensure that a SQL instance has been created.

Step 2 Locate an instance and click **Access**.

Step 3 On the **Dashboard** tab page, click **DataArts Factory** in the lower part of the workspace to access the details page.

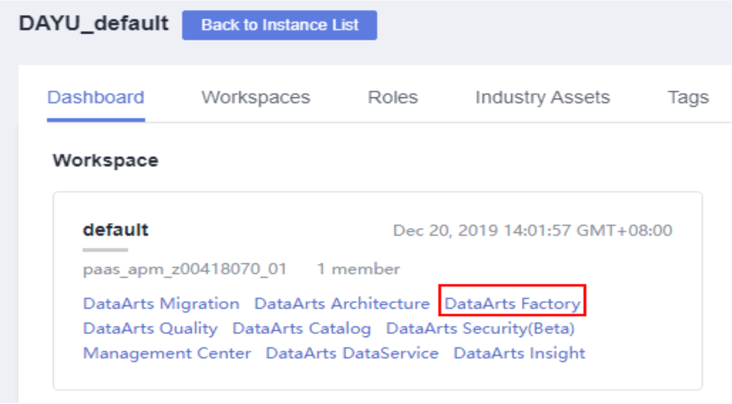

Step 4 Click **Create Job**, set parameters as prompted, and click **OK**. For details, see **Creating a Job**.

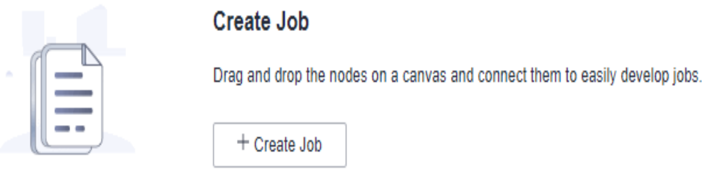

Drag and drop the nodes on a canvas and connect them to easily develop jobs.

Step 5 On the page of the created job, drag the **DLI SQL** icon to the blank area.

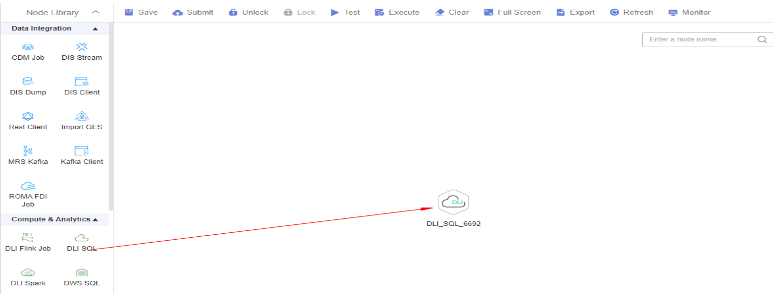

Step 6 Click the **DLI SQL** icon in the blank area. On the **DLI SQL** page that is displayed, enter the SQL statement **MSCK REPAIR TABLE *table_name***. *table_name* indicates the name of the partitioned table. Set **Database Name** and **Queue Name** to the database and queue created when you configure transfer to DLI.

For more information about the DLI SQL properties, see **Developing a SQL Script**.

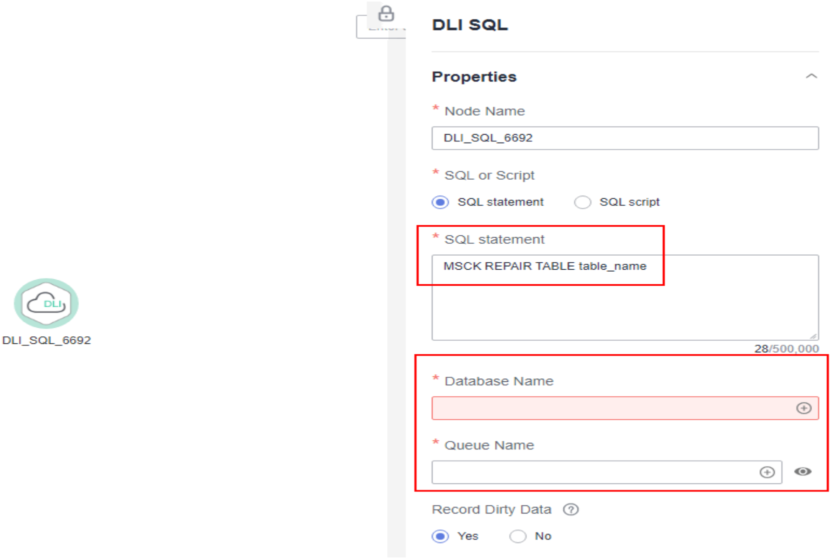

DLI SQL

Properties

* Node Name

DLI_SQL_6692

* SQL or Script

● SQL statement ○ SQL script

* SQL statement

MSCK REPAIR TABLE table_name

28/500,000

* Database Name

* Queue Name

Record Dirty Data

● Yes ○ No

Step 7 After setting the DLI SQL properties, click the blank area and click **Scheduling Setup** to set the time of performing the scheduled operations. For details, see **Setting Up Scheduling for a Job**.

Step 8 After the setting is complete, click **Submit** and then click Execute to trigger the job execution scheduling.

----End