

Log Tank Service

Service Overview

Issue 01
Date 2025-02-08



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2025. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of 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 Infographics.....	1
2 What Is LTS?.....	3
3 Features.....	5
4 Application Scenarios.....	7
5 Security.....	11
5.1 Shared Responsibilities.....	11
5.2 Identity Authentication and Access Control.....	12
5.3 Data Protection.....	12
5.4 Auditing and Logs.....	13
5.5 Resilience.....	13
5.6 Security Risks Monitoring.....	14
5.7 Certificates.....	15
6 Notes and Constraints.....	17
6.1 Basic Resource Constraints.....	17
6.2 Log Read/Write Constraints.....	18
6.3 ICAgent Constraints.....	20
6.4 Search and Analysis Constraints.....	26
6.5 Log Transfer Constraints.....	29
6.6 Log Alarm Constraints.....	32
6.7 Log Metrics Generation Constraints.....	35
6.8 OS Constraints.....	37
7 Permissions.....	39
8 Privacy and Sensitive Information Protection Statement.....	47
8.1 Collector Privacy Statement.....	47
9 Basic Concepts.....	48
10 Related Services.....	49

1 Infographics

Reveal Secrets of Logs

1. Are Logs Important?

Logs are generated when hardware and software such as network devices, OSs, and programs are running. A log records the time, operator, and operation of an event. However, logs are often ignored because they are normally not a core function of a system. The importance of logs is apparent only when decision analysis and problem locating are required.

2. What Is LTS?

Log Tank Service (LTS) collects and stores logs, helping you easily query and analyze them in real time. With LTS, you can perform decision analysis and locate routine O&M problems based on logs.

3. In Which O&M Scenarios Is LTS Applicable?

- Log analysis:** Collects all kinds of logs in real time, and then archives and analyzes them.
- Log audit:** Agent collects logs in real time, preventing accidental or unauthorized data deletions. In addition, logs are transferred to Object Storage Service (OSS) buckets for long-term storage to help you pass compliance checks.
- Problem diagnosis:** Enables you to easily search for logs and locate faults as soon as they occur.
- System improvement:** Detects site performance bottlenecks based on congestion records to optimize cache and data transmission policies.
- System protection:** Configures alarm rules and notifications to locate and analyze issues in minutes.

4. What Are the Highlights of LTS?

- Log collection is perfect for enterprises generating massive scattered log data.
 - LTS collects logs from all servers, eliminating the need to log in to servers one by one.
- Log querying helps you detect and locate hidden faults.
 - TB-level log data is queried less than one second after being written.
 - Full-text queries, by time or keyword, are supported.
- Unified log storage prevents data loss.
 - Logs can be transferred to OSS buckets to ensure security.
 - There is no limit to the size of transferred logs.

Conclusion

As massive data is prevailing, it is disadvantageous to store original log data in disks but never display its value in today's cloud era. In such cases, LTS is developed. It enables you to search, analyze, and explore your log data, making it "live".

2 What Is LTS?

Log Tank Service (LTS) is a high-performance, cost-effective log platform with diverse functions and high reliability.

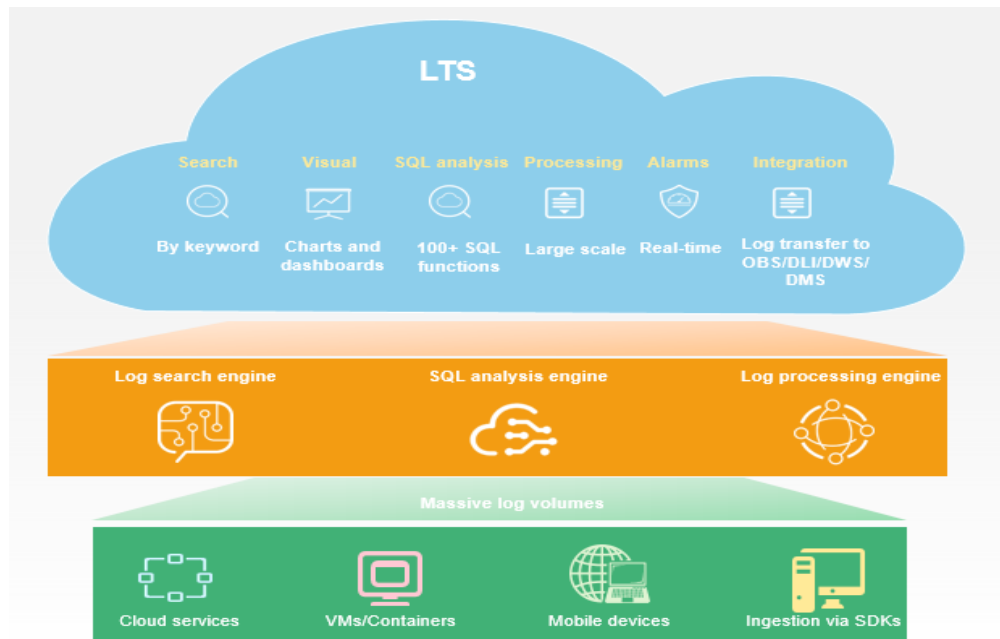
It offers full-stack log collection, search through tens of billions of logs in seconds, PBs of storage, log processing, visual charts, alarm reporting, and log transfer. It is designed for application O&M, security and compliance, and operations analysis.

Core Values

LTS provides multiple modes to ingest massive logs to LTS. It is integrated with log search, SQL analysis, and log processing engines. For details, see [Figure 2-1](#).

- **Log ingestion for all device-cloud scenarios:** covers 40+ cloud services, servers/containers, mobile devices, clouds, SDKs of different languages, and accounts.
- **Log storage and search:** takes just seconds for tens of billions of logs, and supports iterative search of hundreds of billions of logs and PBs of intelligent cold storage.
- **SQL statistics display:** provides over 100 SQL functions, multiple visual charts, and over 10 out-of-the-box dashboards.
- **Real-time alarm:** supports custom content and multiple notification channels (SMS, email, WeCom, DingTalk, and HTTP).
- **One-stop processing:** provides over 200 functions and log normalization, enrichment, anonymization, filtering, and splitting.
- **Service integration:** transfers logs to Object Storage Service (OBS)/Data Warehouse Service (DWS)/Data Ingestion Service (DIS)/Data Lake Insight (DLI)/Distributed Message Service (DMS) for building horizontal solutions.

Figure 2-1 LTS service diagram




3 Features

Before using LTS, learn its main features in [Table 3-1](#).

Table 3-1 Features

Feature	Description
Log ingestion for all device-cloud scenarios	Covers 40+ cloud services, servers/containers, mobile devices, clouds, SDKs of different languages, and accounts.
Real-time log collection	Collects real-time logs and displays them on the LTS console in an intuitive and orderly manner. You can query logs or transfer logs for long-term storage. Collected logs can be structured or unstructured. Log structuring processes logs in log streams by extracting the logs in a fixed format or with a similar pattern based on the extraction rules you set. Then you can use SQL syntax to query the structured logs.
High-volume storage and search	Supports log query by keyword or fuzzy match, search of tens of billions of logs in seconds, and iterative search of hundreds of billions of logs.
SQL statistics and visual charts	<ul style="list-style-type: none">• Provides out-of-the-box dashboard templates for you to quickly analyze ingested logs.• Displays log analysis results in charts, such as tables, line/pie/bar charts, and maps, or aggregates statistical charts on dashboards for easy operations analysis.

Feature	Description
Monitoring and alarms	<ul style="list-style-type: none"> Collects statistics on logs stored in LTS based on keywords or SQL statements; monitors service running status in real time based on the number of keyword occurrences in logs within a specified period; sends custom notifications through channels such as SMS, email, WeCom, DingTalk, and HTTP. <div data-bbox="651 506 1422 770" style="border: 1px dashed gray; padding: 10px; text-align: center;">  <p>LTS Log monitoring Log alarms</p> </div>
Log transfer	<ul style="list-style-type: none"> Allows you to transfer logs to OBS for long-term storage. Logs reported from hosts and cloud services are retained in LTS for a custom period. Log transfer is to replicate logs to the target cloud service. The original logs are retained in LTS and will be automatically deleted when the configured retention period ends. Allows you to transfer logs to DWS/DIS/DLI/DMS with ease for building horizontal solutions.
Log consumption and processing	<ul style="list-style-type: none"> Provides Domain Specific Language (DSL) processing and more than 200 built-in functions to achieve one-stop log normalization, enrichment, anonymization, filtering, and splitting. LTS can also aggregate data by implementing scheduled SQL statistics. Uses SDKs to consume LTS logs to obtain full log data as the data source for stream computing. (The SDK consumption function is being tested and is not available yet.)

4 Application Scenarios

Scenario 1: Application O&M

Enterprises often encounter the following pain points when collecting logs for routine O&M, audit, or security compliance:

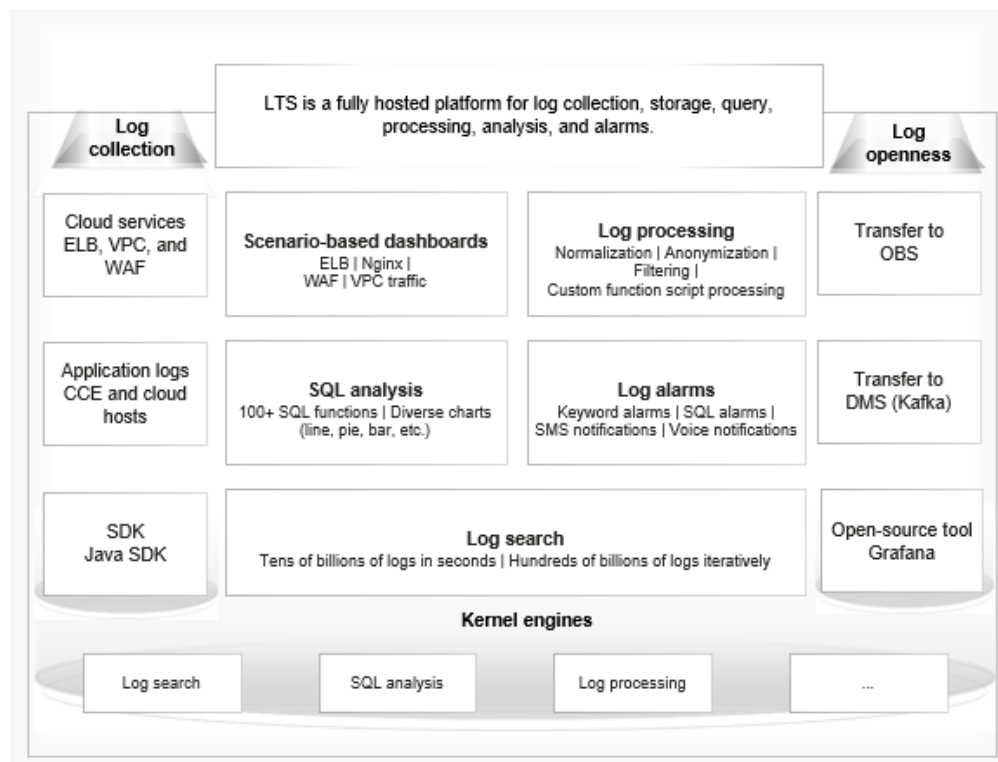
- They need to collect complex and massive logs from many departments.
- Too many cloud resources and untrained monitoring personnel make O&M difficult.
- High security compliance requirements and long-term storage result in high labor and maintenance costs.

LTS provides the following functions for this scenario:

- **All-scenario quick log ingestion:** covers services, applications, middleware, and infrastructure.
- **Fast query and fault location:** queries logs in seconds, and locates and analyzes issues in minutes based on the alarm rules and notifications you set.
- **Long-term log storage in OBS:** meets cyber security requirements.

LTS provides the following solution to centrally collect production environment logs for developers to search and analyze, and for O&M personnel to detect and rectify faults in real time based on the configured alarm rules.

Figure 4-1 Application O&M solution



Scenario 2: Security Compliance

For large enterprises, each service department has an independent cloud account for isolating resources, O&M personnel of each department rely on log monitoring and alarms to locate and analyze faults, and the security department needs to centrally monitor logs of all departments. Therefore, unified log management of multiple accounts is challenging.

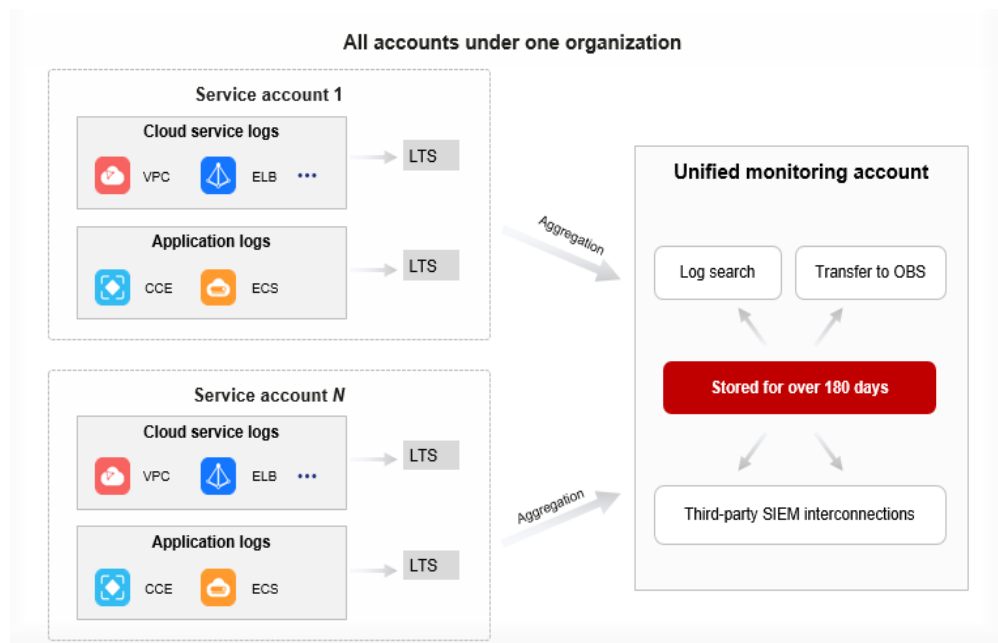
- **Independent O&M by service:** Each service module has an independent account for resource isolation and needs a log service to configure monitoring alarms to quickly locate faults and root causes.
- **Unified log monitoring:** To meet regulatory requirements, the security department needs to aggregate logs of all accounts to one account and store the logs for more than 180 days.

LTS provides the following functions for this scenario:

- **Independent management of accounts:** Each account has isolated resources and permissions, and independently collects its own application and cloud service logs. You can configure alarm rules to demarcate 90% of problems in 10 minutes.
- **Central aggregation of cross-account log data:** The multi-account log center copies logs of multiple accounts to a unified monitoring account to store for at least 180 days for centralized compliance audits, meeting cyber security regulations.

LTS provides the following solution for central collection and storage beyond 180 days, meeting the requirements of the *Cybersecurity Law* and *General Data Protection Regulation (GDPR)*.

Figure 4-2 Security compliance solution



Scenario 3: Operations Analysis

Enterprises collect various logs (such as mobile device and server logs) during their daily operations. After being normalized, filtered, anonymized, and enriched, these logs can be analyzed with big data platforms and BI tools to obtain operations data such as the PV, UV, user stay duration, and transaction amount. The data helps enterprises understand their operations status, analyze user behavior characteristics, make adjustments in real time, improve user experience and operations efficiency, and implement digital transformation.

Enterprises often encounter the following pain points during service analysis:

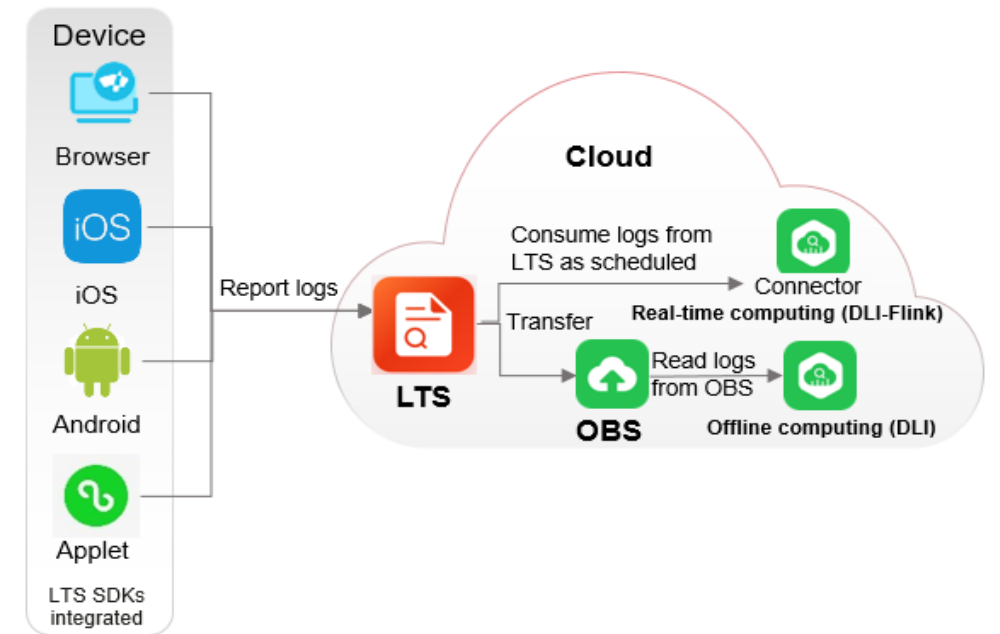
- **Difficult data collection:** It is not easy to collect logs of various mobile devices, such as web browsers, iOS, Android, Baidu applets, WeChat applets, DingTalk applets, and quick apps.
- **Unreliable data transmission:** Mobile device logs are numerous and frequently transmitted. The transmission is slow and logs are prone to be lost, affecting service analysis.
- **Inconvenient data processing:** Raw data cannot be directly processed by big data platforms.

LTS can collect various mobile device logs for you to analyze service operations on big data platforms.

- **Full collection of device logs:** You can quickly integrate LTS mobile SDKs to your devices to enable functions such as cache sending, retry upon exceptions, and batch sending.
- **Fast and reliable reporting:** The collected device logs are reported in seconds through the transmission link without data loss for more complete analysis.
- **Quick interconnection with DLI and DWS:** DLI-Flink integrates connectors and consumes logs from LTS in real time. LTS easily transfers logs to OBS for DLI to read, and transfers structured logs to DWS.

LTS provides the following solution to structure logs, analyze logs with SQL syntax, generate visual charts, and work with big data platforms to help enterprises further explore data value and achieve digital transformation.

Figure 4-3 Operations analysis solution



5 Security

5.1 Shared Responsibilities

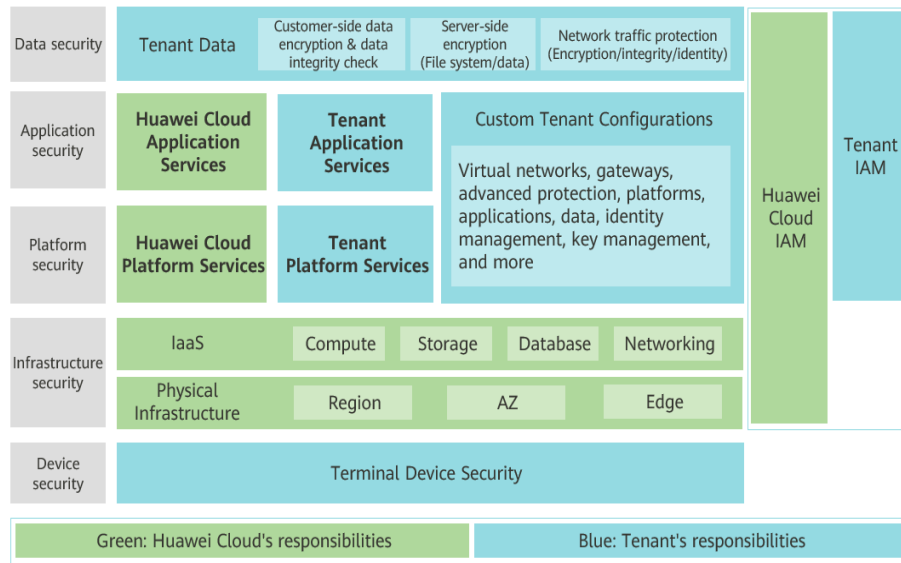
Huawei guarantees that its commitment to cyber security will never be outweighed by the consideration of commercial interests. To cope with emerging cloud security challenges and pervasive cloud security threats and attacks, Huawei Cloud builds a comprehensive cloud service security assurance system for different regions and industries based on Huawei's unique software and hardware advantages, laws, regulations, industry standards, and security ecosystem.

Figure 5-1 illustrates the responsibilities shared by Huawei Cloud and users.

- **Huawei Cloud:** Ensure the security of cloud services and provide secure clouds. Huawei Cloud's security responsibilities include ensuring the security of our IaaS, PaaS, and SaaS services, as well as the physical environments of the Huawei Cloud data centers where our IaaS, PaaS, and SaaS services operate. Huawei Cloud is responsible for not only the security functions and performance of our infrastructure, cloud services, and technologies, but also for the overall cloud O&M security and, in the broader sense, the security and compliance of our infrastructure and services.
- **Tenant:** Use the cloud securely. Tenants of Huawei Cloud are responsible for the secure and effective management of the tenant-customized configurations of cloud services including IaaS, PaaS, and SaaS. This includes but is not limited to virtual networks, the OS of virtual machine hosts and guests, virtual firewalls, API Gateway, advanced security services, all types of cloud services, tenant data, identity accounts, and key management.

Huawei Cloud Security White Paper elaborates on the ideas and measures for building Huawei Cloud security, including cloud security strategies, the shared responsibility model, compliance and privacy, security organizations and personnel, infrastructure security, tenant service and security, engineering security, O&M security, and ecosystem security.

Figure 5-1 Huawei Cloud shared security responsibility model



5.2 Identity Authentication and Access Control

Identity Authentication

No matter whether you access LTS through the console or APIs, you are required to provide the identity credential and verify the identity validity. In addition, login and login authentication policies are provided to harden identity authentication security. LTS uses Identity and Access Management (IAM) to provide three identity authentication modes: **passwords**, **access keys**, and **temporary access keys**. **Login protection** and **login authentication policies** are also provided.

Access Control

To assign different LTS access permissions to employees in your enterprise, IAM is a good choice for refined permissions management. IAM provides identity authentication, permissions management, and access control, helping you secure access to your Huawei Cloud resources. For details, see **Permissions Management**.

5.3 Data Protection

LTS takes many measures to keep data secure and reliable.

Table 5-1 LTS data protection measures

Measure	Description	Reference
Transmission encryption (HTTPS)	LTS supports HTTPS for enhanced data transmission security.	Making an API Request

Measure	Description	Reference
Log redundancy	Log data is stored in multiple copies for data reliability.	/
Transferring logs to OBS	LTS can transfer logs to OBS, so you can store logs for a longer period of time at a lower cost. You can use encrypted OBS buckets to protect data.	Transferring Logs to OBS

5.4 Auditing and Logs


CTS is a log audit service for Huawei Cloud security. It allows you to collect, store, and query cloud resource operation records. You can use these records to perform security analysis, audit compliance, trace resource changes, and locate faults.

After you enable CTS and configure a tracker, CTS records management traces of LTS for auditing.

For details about how to enable and configure CTS, see [Enabling CTS](#).

For the management traces of LTS that can be recorded by CTS, see [Operations Logged by CTS](#).

Figure 5-2 CTS



5.5 Resilience

LTS provides a three-level reliability architecture and uses intra-AZ instance disaster recovery (DR), dual-AZ DR, and multiple log data copies to ensure service durability and reliability.

Table 5-2 LTS reliability architecture

Reliability Solution	Description
Intra-AZ instance DR	In a single AZ, LTS implements instance DR in multi-instance mode and quickly rectifies faults to continuously provide services.
Multi-AZ DR	LTS supports cross-AZ DR. An AZ fault does not interrupt services.
Data DR	Data DR is implemented through log data replication.

5.6 Security Risks Monitoring

LTS monitors security risks in various ways to ensure data security and reliability.

Table 5-3 Risks monitoring


Security Risks Monitoring	Description	Reference
Log alarms	<p>LTS supports log alarms, including keyword alarms and SQL alarms.</p> <ul style="list-style-type: none"> Keyword alarms: LTS allows you to collect statistics on log keywords and set alarm rules to monitor them. By checking the number of keyword occurrences in a specified period, you can have a real-time view of service running. SQL alarms: LTS can regularly run the SQL queries that you specify on structured logs and trigger an alarm when the alarm rule is met. You can view SQL alarms on the LTS console. 	<p>Configuring Alarm Rules</p> <p>SQL Alarms</p>
Usage alerts	Enabling the custom log resource usage alarm will automatically create an alarm rule. If the log resource usage exceeds a specified limit, the system gives an alarm.	Log Resource Usage Alerts

5.7 Certificates

Compliance Certificates

Huawei Cloud services and platforms have obtained various security and compliance certifications from authoritative organizations, such as International Organization for Standardization (ISO). You can [download](#) them from the console.


Figure 5-3 Downloading compliance certificates



Resource Center

Huawei Cloud also provides the following resources to help users meet compliance requirements. For details, see [Resource Center](#).

Figure 5-4 Resource center



6 Notes and Constraints

6.1 Basic Resource Constraints

This section describes constraints on LTS basic resources.

Table 6-1 Basic resource constraints

Item	Description	Remarks
Log groups	Up to 100 log groups can be created for a HUAWEI ID.	To increase the upper limit, submit a service ticket .
Log streams	Up to 100 log streams can be created in a log group. The log stream name must be unique.	To increase the upper limit, submit a service ticket .
Log retention	Logs can be retained for 1 to 365 days.	N/A
Host groups	Up to 200 host groups can be created for a HUAWEI ID.	To increase the upper limit, submit a service ticket .
Quick searches	Up to 50 quick searches can be created in a log stream.	N/A
LogItem (single-line log event)	Using APIs: A single-line log event should be at most 1 MB during ingestion.	N/A
	Using APIs: A single-line log event can contain up to 100 labels.	
	Using ICAgent: A single-line log event should be at most 500 KB during ingestion.	

6.2 Log Read/Write Constraints

This section describes the constraints on LTS log read/write.

Table 6-2 Log read/write constraints

Category	Item	Description	Remarks
HUAWEI ID	Log write traffic	Logs are written at a rate up to 500 MB/s for a HUAWEI ID.	To increase the upper limit, submit a service ticket .
	Log writes	Logs are written up to 10,000 times per second for a HUAWEI ID.	To increase the upper limit, submit a service ticket .
	Log query traffic	Up to 10 MB of logs are returned in a single API query for a HUAWEI ID.	To increase the upper limit, submit a service ticket .
	Log reads	Logs are read up to 1,000 times per minute for a HUAWEI ID.	To increase the upper limit, submit a service ticket .
Log group	Log write traffic	Logs are written at a rate up to 200 MB/s in a log group.	Not mandatory. Service quality cannot be ensured if this limit is exceeded.
	Log writes	Logs are written up to 1,000 times per second in a log group.	Not mandatory. Service quality cannot be ensured if this limit is exceeded.
	Log query traffic	Up to 10 MB of logs are returned in a single API query for a log group.	N/A

Category	Item	Description	Remarks
	Log reads	Logs are read up to 500 times per minute in a log group.	Not mandatory. Service quality cannot be ensured if this limit is exceeded.
Log stream	Log write traffic	Logs are written at a rate up to 100 MB/s in a log stream.	Not mandatory. Service quality cannot be ensured if this limit is exceeded.
	Log writes	Logs are written up to 500 times per second in a log stream.	Not mandatory. Service quality cannot be ensured if this limit is exceeded.
	Log query traffic	Up to 10 MB of logs are returned in a single API query for a log stream.	N/A
	Log reads	Logs are read up to 100 times per minute in a log stream.	Not mandatory. Service quality cannot be ensured if this limit is exceeded.
	Log time	<p>Logs in a period of 24 hours can be collected. Logs generated 24 hours before or after the current time cannot be collected.</p> <p>For example:</p> <ul style="list-style-type: none"> • If the current time is 11:00 on January 7, 2022, logs generated before 11:00 on January 6 cannot be collected. • If the current time is 11:00 on January 7, 2022, logs generated after 11:00 on January 8 cannot be collected. 	N/A

Category	Item	Description	Remarks
SDK	Traffic generated when the SDK reports logs	You are advised to use SDK 1.0.0 or later. If an earlier version is used, upgrade it as soon as possible. Otherwise, the SLA cannot be guaranteed.	The SDK of an earlier version may cause reporting failures.
	Device SDK	Currently, this function is in the closed beta test and is not recommended in the production environment.	In the closed beta test, frequent iterations may occur. You need to perform product upgrade when necessary.

6.3 ICAgent Constraints

This section describes the constraints on the log collector, ICAgent.

Table 6-3 ICAgent file collection constraints

Item	Description	Remarks
File encoding	Only UTF 8 is supported. Other encoding formats may cause garbled characters. For example, binary files.	N/A
Log file size	No restrictions.	N/A
Log file rotation	ICAgent supports configuration of fixed log file names or fuzzy match of log file names. You need to rotate log files manually.	N/A

Item	Description	Remarks
Log collection path	<p>Linux</p> <ul style="list-style-type: none"> • Collection paths support recursion. You can use double asterisks (**) to collect logs from up to five directory levels. Example: <code>/var/logs/**/a.log</code> • Collection paths support fuzzy match. You can use an asterisk (*) to represent one or more characters of a directory or file name. Example: <code>/var/logs/*/a.log</code> or <code>/var/logs/service/a*.log</code> • If the collection path is set to a directory, for example, <code>/var/logs/</code>, only <code>.log</code>, <code>.trace</code>, and <code>.out</code> files in the directory are collected. If the collection path is set to name of a text file, that file is directly collected. • Each collection path must be unique. That is, the same path of the same host cannot be configured for different log groups and log streams. <p>Windows</p> <ul style="list-style-type: none"> • Collection paths support recursion. You can use double asterisks (**) to collect logs from up to five directory levels. Example: <code>C:\var\service**\a.log</code> • Collection paths support fuzzy match. You can use an asterisk (*) to represent one or more characters of a directory or file name. Examples: <code>C:\var\service*\a.log</code> and <code>C:\var\service\a*.log</code> • Each collection path must be unique. That is, the same path of the same host cannot be configured for different log groups and log streams. • Each collection path must be unique. That is, the same path of the same host cannot be configured for different log groups and log streams. 	N/A

Item	Description	Remarks
Symbolic link	Symbolic links are not supported.	N/A
Single log size	The maximum size of each log is 500 KB. Excess content will be truncated by ICAgent.	N/A
Regular expression	Perl regular expressions are supported.	N/A
File collection configuration	A file can be reported to only one log group and stream. If a file is configured for multiple log streams, only one configuration takes effect.	N/A
File opening	Files are opened when being read, and closed after being read.	N/A
First log collection	All logs are collected.	N/A

Table 6-4 ICAgent performance specifications

Item	Description	Remarks
Log collection rate	Raw logs of a single node are collected at a rate up to 50 MB/s.	Service quality cannot be ensured if this limit is exceeded.
Monitored directories	Up to five levels of directories are supported, with up to 1000 files.	N/A
Monitored files	<p>Container scenarios</p> <ul style="list-style-type: none"> ICAgent can collect a maximum of 20 log files from a volume mounting directory. ICAgent can collect a maximum of 1,000 standard container output log files. These files must be in JSON format. <p>VM scenarios</p> <ul style="list-style-type: none"> A maximum of 1,000 files are supported. 	N/A

Item	Description	Remarks
Default resource constraints	<p>CPU</p> <ul style="list-style-type: none"> If ICAgent's version is earlier than 5.12.200, up to two CPU cores are consumed. If ICAgent's version is 5.12.200 or later, up to two CPU cores are consumed when the number of node cores is less than or equal to 4, and the maximum CPU resources consumed is $\log_2(\text{number of node cores})$ when the number of node cores is greater than 4. <p>Memory</p> <ul style="list-style-type: none"> Max. $\min\{4 \text{ GB}, \text{Physical memory}/2\}$. A restart is triggered if this memory limit is exceeded. "$\min\{4 \text{ GB}, \text{Physical memory}/2\}$" means that the smaller value between half of the physical memory and 4 GB is used. 	N/A
Resource limit reached	A forcible restart is triggered. Logs may be lost or duplicate if rotated during the restart.	N/A
Agent installation, upgrade, or uninstallation	No restrictions.	N/A

Table 6-5 Other constraints on ICAgent

Item	Description	Remarks
Configuration update	Configuration updates take effect in 1 to 3 minutes.	N/A

Item	Description	Remarks
Dynamic configuration loading	Console configurations can be dynamically delivered. The update of one configuration does not affect other configurations.	N/A
Configurations	No restrictions.	N/A
Tenant isolation	Tenants are isolated from each other by default.	N/A
Log collection delay	Normally, the delay from writing logs to the disk to collecting the logs is less than 2s (congestion not considered).	N/A
Log upload	File changes are read and uploaded immediately once detected. One or more logs can be uploaded a time.	N/A
Network error handling	Network exceptions trigger retries at an interval of 5s.	N/A
Resource quota used up	If the resources allocated to the ICAgent are insufficient due to massive amounts of logs, the ICAgent continues and retries upon a failure. Logs will be stacked if resources are still insufficient.	N/A
Max. retry timeout	Retry attempts are periodically made.	N/A
Status check	The collector status is monitored through heartbeat detection.	N/A
Checkpoint timeout	Checkpoints are automatically deleted if no updates are made within 12 hours.	N/A

Item	Description	Remarks
Checkpoint saving	Checkpoints are updated if logs are reported successfully.	N/A
Checkpoint saving path	By default, checkpoints are saved in /var/share/oss/manager/ICProbeAgent/internal/TRACE .	N/A
Logs lost or repeated	<p>ICAgent uses multiple mechanisms to ensure log collection reliability and prevent data loss. However, logs may be lost or repeated in the following scenarios:</p> <ul style="list-style-type: none"> • The log rotation policy of CCE is not used. • Log files are rotated at a high speed, for example, once per second. • Logs cannot be forwarded due to improper system security settings or syslog itself. • The container running time, for example, shorter than 30s, is extremely short. • A single node generates logs at a high speed, exceeding the allowed transmit bandwidth or log collection speed. It is recommended that the log generation speed of a single node be lower than 50 MB/s. <p>When ICAgent is restarted, identical data may be collected around the restart time.</p>	N/A

Table 6-6 IP addresses accessible to ICAgent

Component/Service	IP Address	Description
OpenStack	http://169.254.169.254/openstack/latest/meta_data.json	Obtain the metadata, name, and ID of a node.
	http://169.254.169.254/openstack/latest/securitykey	Obtain a temporary AK/SK and security token with an agency.
	http://169.254.169.254/latest/meta-data/public-ipv4	Obtain the EIP bound to a node.
CCE	http://127.0.0.1:4194/api/v2.0/ps	Obtain process information with the cAdvisor API.
	http://127.0.0.1:4194/api/v1.2/docker	Obtain all container metrics with the cAdvisor API.
	http://nodeip:10255/pods	Obtain pod information with a Kubernetes API.

Table 6-7 Ports accessible to ICAgent

Port No.	Description
#icmgr-service {podlb}:30200	ICAgent registration
icmgr-controller {podlb}:30201	ICAgent status configuration
#als-access {podlb}:8102	Log reporting
#ams-access {podlb}:8149	Metric reporting
#ats-access apm {podlb}:8923	Data reporting to APM

6.4 Search and Analysis Constraints

This section describes the constraints on LTS query and analysis.

Search

Table 6-8 Log search constraints

Item	Description	Remarks
Delay from log collection to search	Logs can be searched on the console within 2 minutes after being generated (congestion not considered).	N/A
Keywords	Keywords are conditions excluding Boolean logic operators during query. Up to 30 keywords are supported for a query.	To increase the upper limit, submit a service ticket .
Concurrent queries	Up to 200 concurrent queries are supported for a HUAWEI ID.	To increase the upper limit, submit a service ticket .
Returned records	Up to 250 records are returned by default for a query on the console.	N/A
Returned records	Up to 5,000 records are returned by default for an API query.	N/A
Field size	The maximum size of a field value is 2 KB. The excess part will not be used for quick analysis but can be queried by keyword.	N/A
Search result sorting	By default, search results are displayed by time (accurate to the second) in descending order.	N/A
Fuzzy search	<ul style="list-style-type: none">Each word in a query statement must be fewer than 255 characters.Words cannot start with an asterisk (*) or a question mark (?).Long and double data does not support fuzzy search using asterisks (*) or question marks (?).	N/A

Item	Description	Remarks
Time range	By default, the time span for a single search cannot exceed 30 days.	To increase the upper limit, submit a service ticket .

Analysis

Table 6-9 Constraints on SQL analysis

Item	Description	Remarks
Concurrent analysis tasks	Up to 15 concurrent log analysis tasks are supported for a HUAWEI ID.	To increase the upper limit, submit a service ticket .
Data volume	Up to 24 GB of data can be analyzed at a time in a single log stream.	If your data volume is far beyond the analysis specifications provided by LTS, purchase DWS, configure log transfer to DWS, and use a data warehouse for analysis.
Status	By default, log analysis is disabled.	N/A
Log structuring	Log structuring rules take effect only for data written after the rules are created.	N/A
Returned records	Up to 100 records are returned by default. To increase the number of returned records, use the SQL Syntax .	N/A
	Up to 5000 records are returned for a LIMIT statement.	N/A
Field size	The maximum size of a structured field value is 16 KB. The excess part is not analyzed.	N/A

Item	Description	Remarks
Timeout	Max. 30s.	If your data volume is far beyond the analysis specifications provided by LTS, purchase DWS, configure log transfer to DWS, and use a data warehouse for analysis.
Double value digits	Max. 52 characters. Floating-point numbers with more than 52 bits are less precise than those with fewer bits.	N/A
IP address function	The IP address function analyzes the country, state/province, city, and network operator to which an IP address belongs. The backend database on which this function depends is updated every half year. Therefore, the mappings between a few IP addresses and geographical locations may not be updated in a timely manner.	N/A
SQL analysis time range	Only data generated in the last 30 days can be analyzed using SQL statements.	To increase the upper limit, submit a service ticket .

6.5 Log Transfer Constraints

This section describes the constraints on log transfer.

Table 6-10 Log transfer constraints

Category	Item	Description	Remarks
Log transfer to OBS	Transfer tasks for a log stream	A log stream can have only one task for transferring logs to OBS.	N/A

Category	Item	Description	Remarks
	Log transfer interval	2 minutes, 5 minutes, 30 minutes, 1 hour, 3 hours, 6 hours, 12 hours	N/A
	Data size of each log transfer task	0 MB to 2 GB	N/A
	Transfer rate threshold	100 MB/s The transfer may fail if this limit is exceeded.	N/A
	Log transfer delay	10 minutes For example, if the transfer interval is 30 minutes and the transfer starts at 8:30, transfer files will be generated at 8:40 at the latest.	N/A
	Destination bucket	Standard buckets are supported. Parallel file systems are not supported.	N/A
Log transfer to DIS	Transfer tasks for a log stream	A log stream can have only one task for transferring logs to DIS.	N/A
	Log transfer interval	Real time	N/A
	Data size of each log transfer task	N/A	N/A
	Log transfer delay	N/A	N/A
	Transfer rate threshold	Same as the maximum write rate of the relevant DIS stream. The transferred data will be unstable if this limit is exceeded.	N/A
Log transfer to DMS	Transfer tasks for a log stream	A log stream can have only one task for transferring logs to DMS.	N/A
	Log transfer interval	Real time	N/A
	Data size of each log transfer task	N/A	N/A

Category	Item	Description	Remarks
	Log transfer delay	N/A	N/A
	Transfer rate threshold	Same as the upper traffic limit of the relevant DMS (Kafka) cluster. The transferred data will be unstable if this limit is exceeded.	N/A
Log transfer to GaussDB(DWS)	Transfer tasks for a log stream	A log stream can have only one task for transferring logs to GaussDB(DWS).	N/A
	Log transfer interval	1 minute	N/A
	Data size of each log transfer task	< 5 MB	N/A
	Log transfer delay	5 minutes For example, if the transfer starts at 8:30, transfer files will be generated at 8:35 at the latest.	N/A
	Transfer rate threshold	40 MB/s The transferred data will be unstable if this limit is exceeded.	N/A
	Data reliability	If the format of a batch of data is valid, the data is transferred at least once. However, if the GaussDB(DWS) cluster is heavily loaded or a network error occurs, the write response will time out, which may cause duplicate data. In this case, data accuracy (Exactly Once) cannot be ensured.	N/A

Category	Item	Description	Remarks
	Table structure change	<ul style="list-style-type: none"> Adding non-mandatory columns to DWS tables does not affect log transfer. Adding mandatory columns to DWS tables during log delivery will cause a data write failure. Deleting columns that contain transfer rules from DWS tables during log delivery will cause a data write failure. 	N/A
	Invalid data columns	The common scenarios include mismatch and type conversion failure. This batch of data will not be written to GaussDB(DWS), while other batches will be written normally.	N/A
	Oversized data columns	The common scenarios include long string and varchar type data. This batch of data will not be written to GaussDB(DWS), while other batches will be written normally.	N/A

6.6 Log Alarm Constraints

This section describes the constraints on LTS alarms.

Table 6-11 Alarm constraints

Category	Item	Description	Remarks
Alarm monitoring	Alarm rules	Up to 200 alarms can be created for a HUAWEI ID.	To increase the upper limit, submit a service ticket .

Category	Item	Description	Remarks
	Combinations of search and analysis criteria	1 for keyword search and 1–3 for SQL analysis	N/A
	Query time range	Keyword alarms: max. 1 hour for each query statement	N/A
		SQL alarms: max. 24 hours for each query statement	N/A
	General constraints on query and analysis	For details about keyword alarms, see Search Syntax . For details about SQL alarms, see SQL Syntax .	N/A
Alarm notifications	Notification method	<p>The constraints on each notification method are described below. If the limit is exceeded, you may not receive alarm notifications.</p> <ul style="list-style-type: none"> • Email • SMS • Voice Only mobile numbers (+86) registered in the Chinese Mainland are supported. To apply for this channel, submit a service ticket. • DingTalk The DingTalk bot supports a maximum of 20 messages per minute. • WeCom The WeCom bot supports a maximum of 20 messages per minute. • Lark Lark notifications are only available to whitelisted users. To enable this channel, submit a service ticket. 	N/A

Category	Item	Description	Remarks
	Notification content	<p>Each notification method has a content limit. To ensure successful alarm notifications, the system may truncate the excess part. However, the content integrity and successful sending depend on the truncated content and the notification method. For example, if the truncated content is in invalid Markdown or HTML format, the notification may fail. For SMS and voice messages, truncation does not cause notification failures.</p> <p>Configure the message template while considering the restrictions of your selected notification method. The restrictions of each notification mode are as follows (each letter, digit, or punctuation is counted as one character):</p> <ul style="list-style-type: none"> • SMS For details, see Restrictions on SMS Messaging. • Voice Max. 256 characters. • Email Max. 5 KB. • DingTalk Max. 5 KB. • WeCom Max. 5 KB. 	N/A
	Message templates	Max. 100	N/A
	Message template variables	Max. 3 KB. Excess part will be truncated.	N/A
	Notification message quota	The maximum number of messages a recipient can receive per day depends on their SMN resource quota.	N/A

6.7 Log Metrics Generation Constraints

This section describes the total rule quantity and constraints for LTS to generate metrics based on logs.

Figure 6-1 Total number of metric generation rules

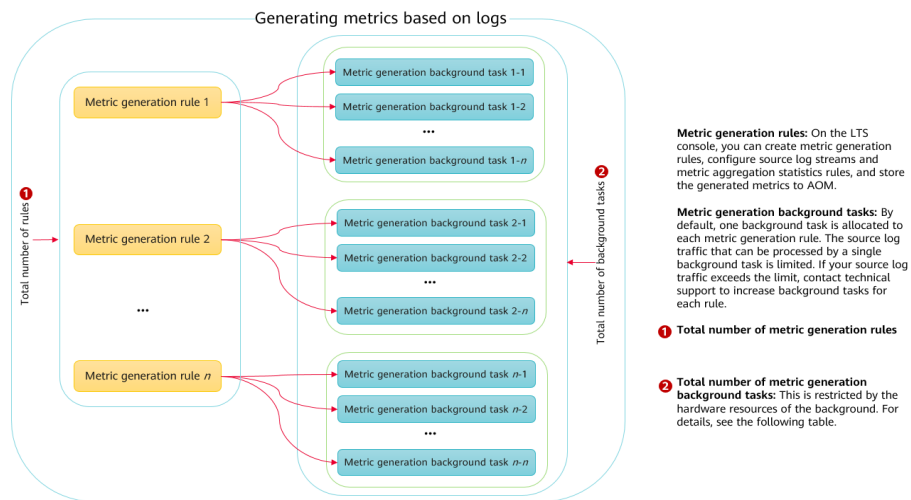


Table 6-12 Constraints on metric generation

Category	Item	Description	Remarks
Metric generation	Total number of metric generation rules allowed in an account	<p>By default, the maximum number of rules allowed for an account is 10, which is limited by the metadata at the software layer.</p> <ul style="list-style-type: none"> By default, one background task is allocated to each metric generation rule. The source log traffic that can be processed by a single background task is limited. If your source log traffic exceeds the limit, submit a service ticket to increase background tasks for each rule. If there are too many log generation metric rules and all background tasks are used up, some log generation metric rules will not be scheduled and executed. 	N/A

Category	Item	Description	Remarks
	Steady log rate of one metric generation rule	≤ 40 MB/s (By default, one background task is assigned to one metric generation rule.)	If your source log traffic exceeds the limit, submit a service ticket to increase background tasks for each rule.
	Peak log rate of one metric generation rule	≤ 50 MB/s (By default, one background task is assigned to one metric generation rule.)	If your source log traffic exceeds the limit, submit a service ticket to increase background tasks for each rule.
	Parameters for configuring metric generation rules	<ul style="list-style-type: none"> • Number of filter criteria groups ≤ 10; Number of associations in each group ≤ 10 • Number of fields for group aggregation (Group By) ≤ 5 • Cardinality of fields for group aggregation (Group By) ≤ 100 	N/A

Category	Item	Description	Remarks
	Log time	The time difference between the log time and the LTS system time must be smaller than 1 hour. Otherwise, the task fails to be executed and metrics cannot be generated and reported to AOM.	The log time must be in order. If there is a severe time disorder or task timeout, logs cannot be processed. Typical example: If the time of the host where the collector is located is not synchronized with the NTP server time or log collection is severely delayed, the log time may be disordered, affecting the aggregation statistics for generated metrics.

6.8 OS Constraints

When purchasing a host, select an OS supported by LTS. Unsupported OSs will prevent LTS from collecting logs from the host.

- For Linux x86_64 hosts, LTS supports all the OSs and versions listed in the preceding table.
- For Linux Arm hosts, LTS supports CentOS 7.4 and later versions, and all versions for other OSs listed in the preceding table.

Table 6-13 Supported OSs and versions (Linux)

OS	Version			
SUSE	SUSE Enterprise 11 SP4 64-bit	SUSE Enterprise 12 SP1 64-bit	SUSE Enterprise 12 SP2 64-bit	SUSE Enterprise 12 SP3 64-bit
openSUSE	13.2 64-bit	42.2 64-bit	15.0 64-bit (Currently, syslog logs cannot be collected.)	
EulerOS	2.2 64-bit	2.3 64-bit	2.10 64-bit	

OS	Version					
CentOS	6.3 64-bit	6.5 64-bit	6.8 64-bit	6.9 64-bit	6.10 64-bit	
	7.1 64-bit	7.2 64-bit	7.3 64-bit	7.4 64-bit	7.5 64-bit	7.6 64-bit
	7.7 64-bit	7.8 64-bit	7.9 64-bit	8.0 64-bit	8.1 64-bit	8.2 64-bit
Ubuntu	14.04 server 64-bit	16.04 server 64-bit	18.04 server 64-bit			
Fedora	24 64-bit	25 64-bit	29 64-bit			
Debian	7.5.0 32-bit	7.5.0 64-bit	8.2.0 64-bit	8.8.0 64-bit	9.0.0 64-bit	

Table 6-14 Supported OSs and versions (Windows)

OS	Version
Windows (64-bit)	Windows Server 2019
	Windows Server 2016 R2 Datacenter
	Windows Server 2016 R2 Standard
	Windows Server 2016 Datacenter English
	Windows Server 2016 R2 Standard English
	Windows Server 2012 R2 Datacenter
	Windows Server 2012 R2 Standard
	Windows Server 2012 Datacenter English
	Windows Server 2012 R2 Standard English
	Windows Server 2008 R2 Enterprise
	Windows Server 2008 R2 Standard
	Windows Server 2008 Enterprise English
	Windows Server 2008 R2 Standard English

7 Permissions

Description

If you need to grant your enterprise personnel permission to access your LTS resources, use Identity and Access Management (IAM). IAM provides identity authentication, fine-grained permissions management, and access control. IAM helps you secure access to your LTS resources.

With IAM, you can create IAM users and grant them permission to access only specific resources. For example, if you want some software developers in your enterprise to be able to use LTS resources but do not want them to be able to delete resources or perform any other high-risk operations, you can create IAM users and grant permission to use LTS resources but not permission to delete them.

If your Huawei account does not require individual IAM users for permissions management, you can skip this section.

IAM is a free service. You only pay for the resources in your account. For more information about IAM, see [IAM Service Overview](#).

Why Is "Insufficient Permission" Displayed After Enterprise Project Authorization?

IAM projects/Enterprise projects: the authorization scope of a custom policy. A custom policy can be applied to IAM projects or enterprise projects or both. Policies that contain actions for both IAM and enterprise projects can be used and take effect for both IAM and Enterprise Management. Policies that contain actions only for IAM projects can be used and only take effect for IAM. For details, see [What Are the Differences Between IAM and Enterprise Management?](#)

In LTS, only log group, log stream, and dashboard resource APIs support enterprise project authorization. For other APIs that support only IAM project authorization:

1. Click **By IAM Project** during authorization.

Figure 7-1 Viewing authorization records by IAM project



2. When selecting the authorization scope, select **Region-specific projects** according to the minimum authorization principle.

LTS Permissions

New IAM users do not have any permissions assigned by default. You need to first add them to one or more groups and then attach policies or roles to these groups. The users then inherit permissions from the groups and can perform specified operations on LTS based on the permissions they have been assigned.

LTS is a project-level service deployed for specific regions. When you set **Scope** to **Region-specific projects** and select the specified projects in the specified regions, the users only have permissions for LTS in the selected projects. If you select **All projects**, the users have permissions for LTS in all region-specific projects. When accessing LTS, the users need to switch to the authorized region.

You can grant permissions by using roles and policies.

- **Roles:** A coarse-grained authorization strategy that defines permissions by job responsibility. Only a limited number of service-level roles are available for authorization. Cloud services often depend on each other. When you grant permissions using roles, you also need to attach any existing role dependencies. Roles are not ideal for fine-grained authorization and least privilege access.
- **Policies:** A fine-grained authorization strategy that defines permissions required to perform operations on specific cloud resources under certain conditions. This type of authorization is more flexible and is ideal for least privilege access.

The system permissions supported by LTS are listed in [Table 7-1](#).

Table 7-1 System-defined permissions for LTS

Role/ Policy Name	Description	Type	Dependencies
LTS FullAccess	Full permissions for LTS. Users with these permissions can perform operations on LTS.	System-defined policy	CCE Administrator, OBS Administrator, FunctionGraph FullAccess, and AOM FullAccess
LTS ReadOnly Access	Read-only permissions for LTS. Users with these permissions can only view LTS data.	System-defined policy	CCE Administrator, OBS Administrator, and AOM FullAccess

Role/ Policy Name	Description	Type	Dependencies
LTS Administrator	Administrator permissions for LTS.	System-defined role	Tenant Guest and Tenant Administrator

Table 7-2 lists the common operations supported by system-defined permissions for LTS.

Table 7-2 Common operations supported by system-defined permissions

Operation	LTS FullAccess	LTS ReadOnlyAccess	LTS Administrator
Querying a log group	√	√	√
Creating a log group	√	×	√
Modifying a log group	√	×	√
Deleting a log group	√	×	√
Querying a log stream	√	√	√
Creating a log stream	√	×	√
Modifying a log stream	√	×	√
Deleting a log stream	√	×	√
Configuring log collection from hosts	√	×	√
Viewing a dashboard	√	√	√
Creating a dashboard	√	×	√
Modifying a dashboard	√	×	√

Operation	LTS FullAccess	LTS ReadOnlyAccess	LTS Administrator
Deleting a dashboard	√	×	√
Querying log structuring configurations	√	√	√
Configuring log structuring	√	×	√
Enabling quick analysis	√	×	√
Disabling quick analysis	√	×	√
Configuring delimiters	√	×	√
Querying a filter	√	√	√
Disabling a filter	√	×	√
Enabling a filter	√	×	√
Deleting a filter	√	×	√
Querying an alarm rule	√	√	√
Creating an alarm rule	√	×	√
Modifying an alarm rule	√	×	√
Deleting an alarm rule	√	×	√
Viewing a log transfer task	√	√	√
Creating a log transfer task	√	×	√
Modifying a log transfer task	√	×	√
Deleting a log transfer task	√	×	√
Enabling a log transfer task	√	×	√

Operation	LTS FullAccess	LTS ReadOnlyAccess	LTS Administrator
Disabling a log transfer task	√	×	√
Installing ICAgent	√	×	√
Upgrading ICAgent	√	×	√
Uninstalling ICAgent	√	×	√

To use a custom fine-grained policy, log in to IAM as the administrator and select fine-grained permissions of LTS as required.

Table 7-3 describes fine-grained permission dependencies of LTS.

Table 7-3 Fine-grained permission dependencies of LTS

Permission	Description	Dependency
lts:agents:list	List agents	None
lts:buckets:get	Query a specified bucket	None
lts:groups:put	Modify a specified log group	None
lts:transfers:create	Create a log transfer task	obs:bucket:PutBucketAcl obs:bucket:GetBucketAcl obs:bucket:GetEncryptionConfiguration obs:bucket:HeadBucket dis:streams:list dis:streamPolicies:list
lts:groups:get	Query a specified log group	None
lts:transfers:put	Modify a log transfer task	obs:bucket:PutBucketAcl obs:bucket:GetBucketAcl obs:bucket:GetEncryptionConfiguration obs:bucket:HeadBucket dis:streams:list dis:streamPolicies:list
lts:resourceTags:delete	Delete resource tags	None

Permission	Description	Dependency
lts:ecsOsLogPaths:list	List OS log paths of a specified image	None
lts:structConfig:create	Create an LTS structuring configuration	None
lts:agentsConf:get	Query a specified agent configuration	None
lts:logIndex:list	List log indexes	None
lts:transfers:delete	Delete a log transfer task	None
lts:regex:create	Extract structured fields	None
lts:subscriptions:delete	Delete a specified subscription	None
lts:overviewLogsLast:list	List the latest logs of a user	None
lts:logIndex:get	Query a specified log index	None
lts:sqlalarmrules:create	Create an alarm rule	None
lts:agentsConf:create	Create an agent configuration	None
lts:sqlalarmrules:get	Query an alarm rule	None
lts:datasources:batchdelete	Batch delete data sources	None
lts:structConfig:put	Modify an LTS structuring configuration	None
lts:groups:list	List log groups	None
lts:sqlalarmrules:delete	Delete an alarm rule	None
lts:transfers:action	Enable or disable a log transfer task	None
lts:datasources:post	Create a data source	None
lts:topics:create	Create a log topic	None
lts:resourceTags:get	Query resource tags	None
lts:filters:put	Modify a log filter	None
lts:logs:list	List logs	None
lts:subscriptions:create	Create a subscription	None

Permission	Description	Dependency
lts:filtersAction:put	Enable or disable a log filter	None
lts:overviewLogsTopTopic:get	Query data metrics of the topic with the largest log volume	None
lts:datasources:put	Modify a data source	None
lts:structConfig:delete	Delete an LTS structuring configuration	None
lts:logIndex:delete	Delete a specified log index	None
lts:filters:get	Query a specified log filter	None
lts:topics:delete	Delete log topics	None
lts:agentSupportedOSLogPaths:list	List the log paths of OS supported by the agent	None
lts:topics:put	Modify a log topic	None
lts:agentHeartbeat:post	Upload agent heartbeats	None
lts:logsByName:upload	Upload logs by log group name and topic name	None
lts:buckets:list	List buckets	None
lts:logIndex:post	Create a log index	None
lts:logContext:list	List log contexts	None
lts:groups:delete	Delete a specified log group	None
lts:filters:delete	Delete a log filter	None
lts:resourceTags:put	Update resource tags	None
lts:structConfig:get	Query an LTS structuring configuration	None
lts:overviewLogTotal:get	Query the total log volume of the current user	None
lts:subscriptions:put	Modify a specified subscription	None
lts:subscriptions:list	List subscriptions	None
lts:datasources:delete	Delete a specified data source	None

Permission	Description	Dependency
lts:transfersStatus:get	Query the log transfer status	None
lts:logIndex:put	Modify a specified log index	None
lts:sqlalarmrules:put	Modify an alarm rule	None
lts:logs:upload	Upload logs	None
lts:agentDetails:list	List agent diagnostic logs	None
lts:agentsConf:put	Modify an agent configuration	None
lts:logstreams:list	Filter log stream resources	None
lts:subscriptions:get	Query a specified subscription	None
lts:disStreams:list	List DIS streams	None
lts:groupTopics:put	Create a log group and topic	None
lts:resourceInstance:list	List resource instances	None
lts:transfers:list	List transfer tasks	None
lts:topics:get	Query a specified log topic	None
lts:agentsConf:delete	Delete a specified agent configuration	None
lts:agentEcs:list	List ECSs	None
lts:indiceLogs:list	Search for logs	None
lts:topics:list	List log topics	None

8 Privacy and Sensitive Information Protection Statement

8.1 Collector Privacy Statement

O&M data will be displayed on the LTS console. It is recommended that you do not upload your personal or sensitive data to LTS. Encrypt such data if you need to upload it.

ICAgent Deployment

When you install ICAgent on an ECS, your AK/SK pair is required in the installation command. Before the installation, disable history collection in the ECS to protect your AK/SK pair. After the installation, ICAgent will encrypt your AK/SK pair and store it.

9 Basic Concepts

Log Groups

A log group is a collection of log streams and the basic unit for LTS to manage logs. You can set log retention duration for a log group.

Log Streams

A log stream is the basic unit for log reads and writes.

You can sort logs of different types, such as operation logs and access logs, into different log streams. ICAgent will package and send the collected logs to LTS on a log stream basis. It makes it easier to find specific logs when you need them.

The use of log streams greatly reduces the number of log reads and writes and improves efficiency.

ICAgent

ICAgent is the log collection tool of LTS. If you want to use LTS to collect logs from a host, you need to install ICAgent on the host. Batch ICAgent installation is supported if you want to collect logs from multiple hosts. After ICAgent installation, you can check the ICAgent status on the LTS console in real time.

10 Related Services

The relationships between LTS and other services are described in [Table 1](#).

Table 10-1 Relationships with other services

Interaction	Related Service
With Cloud Trace Service (CTS), you can record operations associated with LTS for future query, audit, and backtracking.	CTS
You can transfer logs to Object Storage Service (OBS) buckets for long-term storage, preventing log loss.	OBS
You can transfer logs to Data Ingestion Service (DIS) for long-term storage. DIS can perform offline analysis, and transmit log files to the cloud for backup, query, and machine learning. You can also use it for recovery and fault analysis after data losses or exceptions. In addition, a large number of small text files can be combined and transferred into large files to improve data processing performance.	DIS
Application Operations Management (AOM) can collect site access statistics, monitor logs sent from LTS, and generate alarms.	AOM
Identity and Access Management (IAM) allows you to grant LTS permissions to IAM users under your account.	IAM