Log Tank Service

Best Practices

Issue 01

Date 2025-08-05

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

Trademarks and Permissions

Notice

Contents

1 Overview

This document describes the following best practices of Log Tank Service (LTS):

Table 1-1 Best practice overview

Category	Best Practice	Scenario
Log ingestion	**Collecting Host Logs from Third-Party Clouds, Internet Data Centers, and Other Huawei Cloud Regions to LTS**	This practice describes how to collect Alibaba Cloud host logs to Huawei Cloud LTS. The method is similar to that of collecting logs from Internet Data Centers (IDCs) or across Huawei Cloud regions.
Log ingestion	**Collecting Kubernetes Logs from Third-Party Clouds, IDCs, and Other Huawei Cloud Regions to LTS**	This practice describes how to collect Alibaba Cloud Kubernetes logs to Huawei Cloud LTS. The method is similar to that of collecting logs from IDCs or across Huawei Cloud regions.
Log ingestion	**Collecting Syslog Aggregation Server Logs to LTS**	This practice describes how to use the syslog protocol to upload logs to LTS. You need to buy an Elastic Cloud Server (ECS) as a syslog aggregation server. Syslog comes preinstalled by default on Linux servers. However, ECSs do not receive remote syslog writes by default. You need to enable this function.
Log ingestion	**Importing Logs of Self-built ELK to LTS**	This practice describes how to use a custom Python script and ICAgent (LTS collector) to transfer logs from Elasticsearch to LTS.
Log ingestion	**Collecting Zabbix Data Through ECS Log Ingestion**	This practice describes how to collect monitoring data from Zabbix to an LTS log stream.

Category	Best Practice	Scenario
Log search and analysis	**Analyzing Huawei Cloud ELB Logs on LTS**	This practice describes how to search for and analyze logs after Elastic Load Balance (ELB) logs are ingested to and structured in LTS.
Log search and analysis	**Analyzing Huawei Cloud WAF Logs on LTS**	This practice describes how to search for and analyze logs after Web Application Firewall (WAF) logs are ingested to and structured in LTS.

2 Log Ingestion

2.1 Collecting Host Logs from Third-Party Clouds, Internet Data Centers, and Other Huawei Cloud Regions to LTS

Solution Overview

Users often need to collect logs across clouds or regions. There are two typical scenarios:

- Scenario 1: collecting logs from IDCs or third-party clouds to Huawei Cloud LTS

Figure 2-1 Third-party cloud log collection

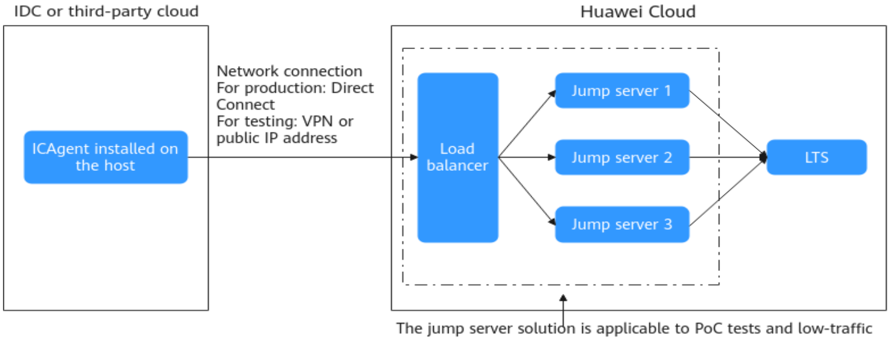

The jump server solution is applicable to PoC tests and low-traffic scenarios.
For heavy-traffic scenarios, submit a service ticket to Huawei Cloud technical support engineers to design a networking solution.

- Scenario 2: collecting logs from one Huawei Cloud region to LTS in another Huawei Cloud region

Figure 2-2 Cross-region log collection

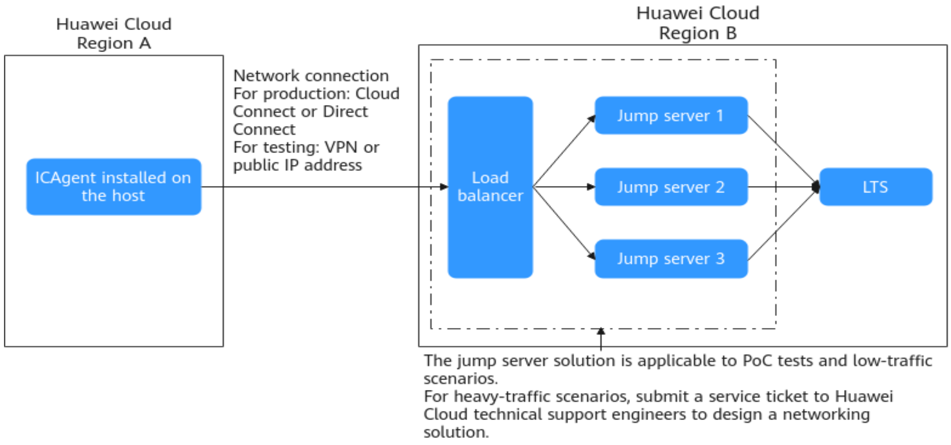
Huawei Cloud Region A
ICAgent installed on the host
Network connection
For production: Cloud Connect or Direct Connect
For testing: VPN or public IP address
Huawei Cloud Region B
Load balancer
Jump server 1
Jump server 2
Jump server 3
LTS
The jump server solution is applicable to PoC tests and low-traffic scenarios.
For heavy-traffic scenarios, submit a service ticket to Huawei Cloud technical support engineers to design a networking solution.

In both scenarios, you need to establish a network connection, install ICAgent, and follow the log ingestion wizard.

- **ICAgent:** the log collector of Huawei Cloud LTS. After being installed on a host, it collects logs from the host to LTS. Ensure that the time and time zone of your local browser are consistent with those of the host to install ICAgent.
- **Networking**
  - Scenario 1: Direct Connect is a typical method for connecting a customer-built IDC or third-party cloud to Huawei Cloud. If Direct Connect is unavailable, you can use a VPN or public IP address.
  - Scenario 2: Cloud Connect or Direct Connect is a typical method for interconnecting Huawei Cloud regions. You can also use a VPN or public IP address.
- **Jump server**
  - ICAgent installed in customer-built IDCs, third-party clouds, or other Huawei Cloud regions cannot directly access the network segment used by the Huawei Cloud management plane for log reporting, necessitating a jump server for data forwarding. Use the jump server solution for Proof of Concept (PoC) tests or when log traffic is light. If you do not want to use jump servers for heavy traffic scenarios in production environments, **create a service ticket** to obtain Huawei Cloud technical support to design a network passthrough solution.
  - A typical jump server configuration is 2 vCPUs and 4 GB memory, allowing it to forward traffic at approximately 30 MB/s. Configure a proper number of jump servers based on your log traffic and use a load balancer to distribute traffic among them.

This section describes how to collect Alibaba Cloud host logs to Huawei Cloud LTS. The method is similar to that of collecting logs from IDCs or across Huawei Cloud regions.

Below are the steps to collect the logs from a Linux host in Alibaba Cloud's China (Beijing) region to LTS in Huawei Cloud's CN East-Shanghai1 region.

Planning Resources

Table 2-1 Planning resources

Region	Resource	Description
CN East-Shanghai 1	ECS	You are advised to use **CentOS 6.5 64bit** or later images. The minimum flavor for the ECS is 1 vCPU and 1 GB of memory, while the recommended flavor is 2 vCPUs and 4 GB of memory.
	Load balancer	● When buying a load balancer, select the same VPC as the ECS. ● Create an EIP for connecting to the jump servers. ● Buy the bandwidth based on the service requirements.

Step 1: Purchasing a Load Balancer and an ECS as a Jump Server in Huawei Cloud CN East-Shanghai1

Step 1 Log in to the ECS console and buy an ECS.

Before installing ICAgent on a non-Huawei Cloud host, buy an ECS as a jump server from Huawei Cloud.

Step 2 Buy a load balancer, add TCP listeners, and associate a backend server group with it.

1. Add listeners for ports 30200, 30201, 8149, 8923, and 8102. For details, see **Adding a TCP Listener**.
2. Add the jump server to a backend server group. For details, see **Backend Server Group**.

Step 3 Configure a security group rule for the jump server and open forwarding ports.

1. Modify the security group rule used by the jump server.
   a. On the ECS console, click the name of the ECS used as the jump server to go to the details page.
   b. On the **Security Groups** tab page, click a security group name to go to the details page.
   c. Click the **Inbound Rules** tab and click **Add Rule**. Open the inbound ports 8149, 8102, 8923, 30200, 30201, and 80 to ensure that data can be transmitted from the non-Huawei Cloud host to the jump server.

Figure 2-3 Modifying a security group rule

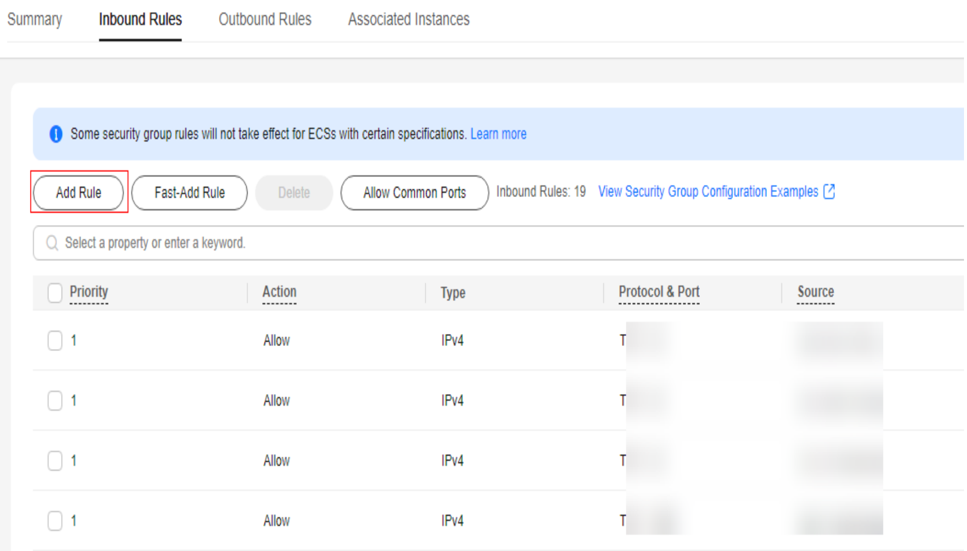

Some security group rules will not take effect for ECSs with certain specifications. Learn more

Priority	Action	Type	Protocol & Port	Source
□ 1	Allow	IPv4	T[redacted]	[redacted]
□ 1	Allow	IPv4	T[redacted]	[redacted]
□ 1	Allow	IPv4	T[redacted]	[redacted]
□ 1	Allow	IPv4	T[redacted]	[redacted]

2. On the LTS console, choose **Host Management** > **Hosts** in the navigation pane, and click **Install ICAgent** in the upper right corner. Set parameters as shown in the following figure. Set **Private IP** to the private IP address of the ECS to generate an installation password.

Figure 2-4 Installing ICAgent

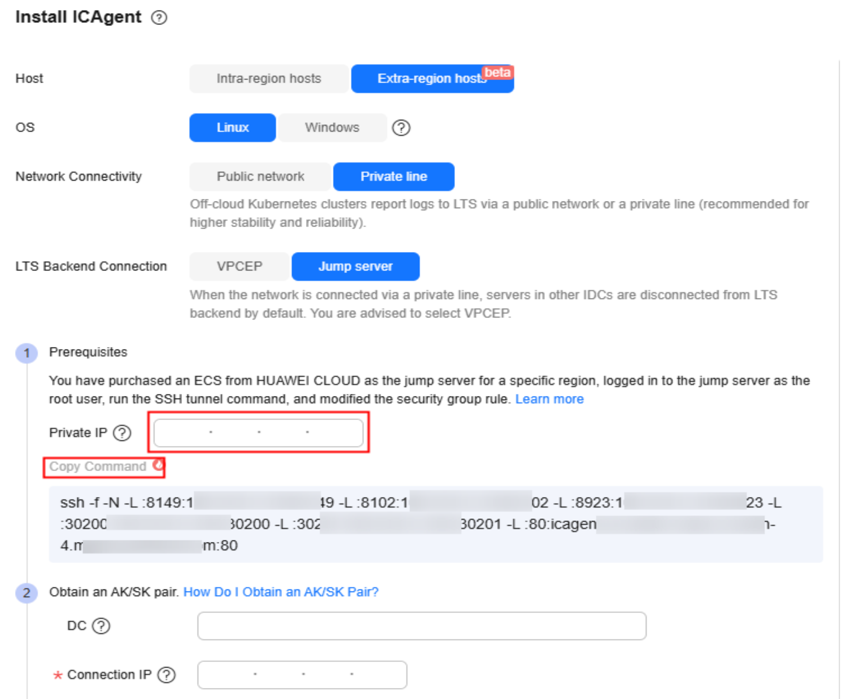

Host: Extra-region hosts

OS: Linux

Network Connectivity: Private line

Off-cloud Kubernetes clusters report logs to LTS via a public network or a private line (recommended for higher stability and reliability).

LTS Backend Connection: Jump server

When the network is connected via a private line, servers in other IDCs are disconnected from LTS backend by default. You are advised to select VPCEP.

1 Prerequisites

You have purchased an ECS from HUAWEI CLOUD as the jump server for a specific region, logged in to the jump server as the root user, run the SSH tunnel command, and modified the security group rule. Learn more

Private IP: . . .

```
ssh -f -N -L :8149:1[redacted]49 -L :8102:1[redacted]02 -L :8923:1[redacted]23 -L
:30200[redacted]30200 -L :302[redacted]30201 -L :80:icagen[redacted]n-
4.m[redacted]m:80
```

2 Obtain an AK/SK pair. How Do I Obtain an AK/SK Pair?

DC:

* Connection IP: . . .

3. Copy the command, log in to the jump server as user **root**, run the SSH tunneling command, and enter the password of user **root** as prompted.

4. Run the **netstat -lnp | grep ssh** command to check whether the corresponding TCP ports are being listened to. If the command output similar to the following is returned, the ports are open.

- Enter **http://***Jump server IP address* in the address bar of a browser. If the access is successful, the security group rule has taken effect.
- If the jump server is powered off and then restarted, run the installation command generated on the ICAgent installation page again. If you use the jump server in a production environment, configure the SSH tunneling command to run upon system startup.

Figure 2-5 Viewing ports

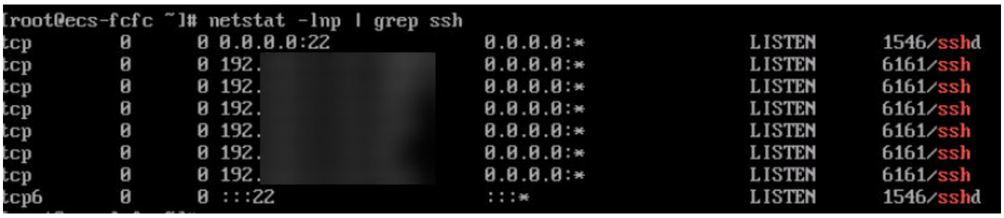

```
[root@ecs-fcfc ~]# netstat -lnp | grep ssh
tcp        0      0 0.0.0.0:22              0.0.0.0:*               LISTEN      1546/sshd
tcp        0      0 192.[redacted]          0.0.0.0:*               LISTEN      6161/ssh
tcp        0      0 192.                    0.0.0.0:*               LISTEN      6161/ssh
tcp        0      0 192.                    0.0.0.0:*               LISTEN      6161/ssh
tcp        0      0 192.                    0.0.0.0:*               LISTEN      6161/ssh
tcp        0      0 192.                    0.0.0.0:*               LISTEN      6161/ssh
tcp        0      0 192.                    0.0.0.0:*               LISTEN      6161/ssh
tcp6       0      0 :::22                   :::*                    LISTEN      1546/sshd
```

----End

Step 2: Installing ICAgent on an Alibaba Cloud Host

Step 1 Obtain an AK/SK. For details, see **How Do I Obtain an Access Key (AK/SK)?**

Step 2 On the **Install ICAgent** page of the LTS console, enter the connection IP address of the jump server to generate the ICAgent installation command.

Figure 2-6 ICAgent installation page

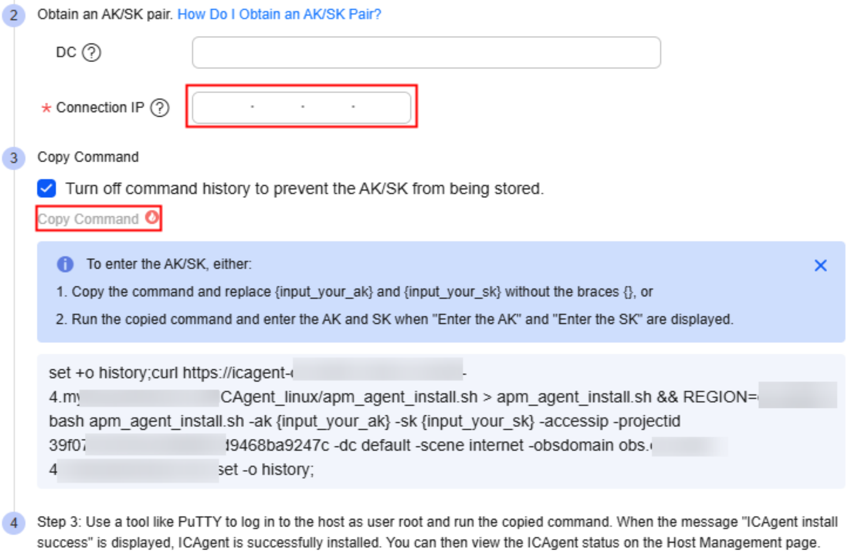

2 Obtain an AK/SK pair. How Do I Obtain an AK/SK Pair?

DC ⓘ

* Connection IP ⓘ . . .

3 Copy Command

☑ Turn off command history to prevent the AK/SK from being stored.

To enter the AK/SK, either:

1. Copy the command and replace {input_your_ak} and {input_your_sk} without the braces {}, or
2. Run the copied command and enter the AK and SK when "Enter the AK" and "Enter the SK" are displayed.

```
set +o history;curl https://icagent-[redacted]-
4.my[redacted]CAgent_linux/apm_agent_install.sh > apm_agent_install.sh && REGION=[redacted]
bash apm_agent_install.sh -ak {input_your_ak} -sk {input_your_sk} -accessip -projectid
39f07[redacted]49468ba9247c -dc default -scene internet -obsdomain obs.[redacted]
4[redacted]set -o history;
```

4 Step 3: Use a tool like PuTTY to log in to the host as user root and run the copied command. When the message "ICAgent install success" is displayed, ICAgent is successfully installed. You can then view the ICAgent status on the Host Management page.

- Replace the AK/SK in the command with the correct AK/SK. Otherwise, ICAgent cannot be installed.
- **Connection IP**: connection IP address of the jump server. If the jump server uses an EIP to connect to the extra-region host, enter the EIP created by the load balancer. If the jump server and extra-region host use Virtual Private Cloud (VPC) peering connection of Direct Connect, enter the private IP address of the ECS.

Step 3 Log in to the Alibaba Cloud host as user **root** and run the ICAgent installation command. If the message **ICAgent install success** is displayed, ICAgent is successfully installed.

If you use LTS to collect logs across Huawei Cloud regions, for example, collecting logs from the CN East-Shanghai1 region to the CN South-Guangzhou region, you need to buy a load balancer and an ECS used as a jump server in CN South-Guangzhou, and then run the ICAgent installation command on the jump server in CN East-Shanghai1.

Figure 2-7 Checking the ICAgent installation status

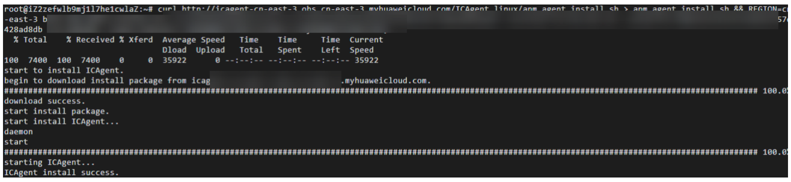

```
root@iZ2zefwlb9mj1l7he1cwlaZ:~# curl http://icagent-cn-east-3.obs.cn-east-3.myhuaweicloud.com/ICAgent_linux/apm_agent_install.sh > apm_agent_install.sh && REGION=c
-east-3 b[redacted]57
428ad8db
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100  7400  100  7400    0     0  35922      0 --:--:-- --:--:-- --:--:-- 35922
start to install ICAgent.
begin to download install package from icag[redacted].myhuaweicloud.com.
############################################################################################################################################ 100.0
download success.
start install package.
start install ICAgent...
daemon
start
############################################################################################################################################ 100.0
starting ICAgent...
ICAgent install success.
```

Step 4 Choose **Host Management** > **Hosts** in the navigation pane of the LTS console and check whether the ICAgent status is **Running**.

----End

Step 3: Ingesting Logs to LTS

Step 1 Log in to the LTS console and choose **Host Management** > **Host Groups** in the navigation pane. Click **Create Host Group**. On the displayed page, enter a host group name and select hosts.

Step 2 Configure a log ingestion rule. For details, see **Collecting Logs from ECS**.

----End

Step 4: Viewing the Log Stream

On the **Log Management** page of LTS, click the target log stream to go to its details page. If there are logs, the Alibaba Cloud logs have been reported to LTS.

2.2 Collecting Kubernetes Logs from Third-Party Clouds, IDCs, and Other Huawei Cloud Regions to LTS

Solution Overview

Users often need to collect Kubernetes logs across clouds or regions. There are two typical scenarios:

- Scenario 1: collecting logs from IDCs or third-party clouds to Huawei Cloud LTS

Figure 2-8 Third-party cloud log collection

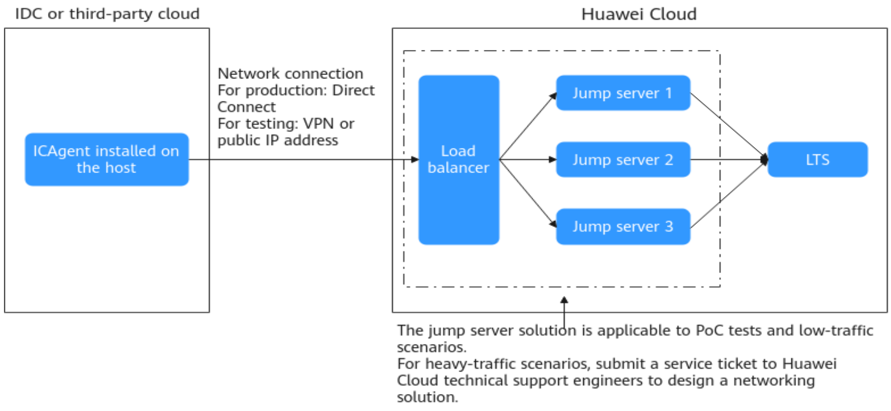
IDC or third-party cloud
Huawei Cloud
ICAgent installed on the host
Network connection
For production: Direct Connect
For testing: VPN or public IP address
Load balancer
Jump server 1
Jump server 2
Jump server 3
LTS
The jump server solution is applicable to PoC tests and low-traffic scenarios.
For heavy-traffic scenarios, submit a service ticket to Huawei Cloud technical support engineers to design a networking solution.

- Scenario 2: collecting logs from one Huawei Cloud region to LTS in another Huawei Cloud region

Figure 2-9 Cross-region log collection

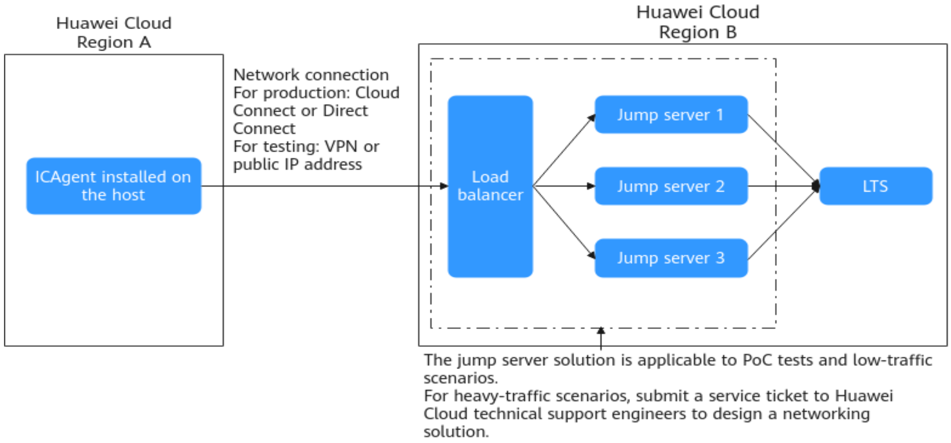
Huawei Cloud Region A
Huawei Cloud Region B
ICAgent installed on the host
Network connection
For production: Cloud Connect or Direct Connect
For testing: VPN or public IP address
Load balancer
Jump server 1
Jump server 2
Jump server 3
LTS
The jump server solution is applicable to PoC tests and low-traffic scenarios.
For heavy-traffic scenarios, submit a service ticket to Huawei Cloud technical support engineers to design a networking solution.

In both scenarios, you need to establish a network connection, install ICAgent, and follow the log ingestion wizard.

- **ICAgent:** the log collector of Huawei Cloud LTS. After being installed on a host, it collects logs from the host to LTS. Ensure that the time and time zone of your local browser are consistent with those of the host to install ICAgent.
- **Networking**
  - Scenario 1: Direct Connect is a typical method for connecting a customer-built IDC or third-party cloud to Huawei Cloud. If Direct Connect is unavailable, you can use a VPN or public IP address.
  - Scenario 2: Cloud Connect or Direct Connect is a typical method for interconnecting Huawei Cloud regions. You can also use a VPN or public IP address.
- **Jump server**
  - ICAgent installed in customer-built IDCs, third-party clouds, or other Huawei Cloud regions cannot directly access the network segment used by the Huawei Cloud management plane for log reporting, necessitating a jump server for data forwarding. Use the jump server solution for Proof of Concept (PoC) tests or when log traffic is light. If you do not want to

use jump servers for heavy traffic scenarios in production environments, **create a service ticket** to obtain Huawei Cloud technical support to design a network passthrough solution.

- A typical jump server configuration is 2 vCPUs and 4 GB memory, allowing it to forward traffic at approximately 30 MB/s. Configure a proper number of jump servers based on your log traffic and use a load balancer to distribute traffic among them.

This section describes how to collect logs from third-party cloud Kubernetes clusters to Huawei Cloud LTS. The method is similar to that of collecting logs from IDCs or across Huawei Cloud regions.

Planning Resources

Table 2-2 Planning resources

Region	Resource	Description
CN East-Shanghai1	ECS	You are advised to use **CentOS 6.5 64bit** or later images. The minimum flavor for the ECS is 1 vCPU and 1 GB of memory, while the recommended flavor is 2 vCPUs and 4 GB of memory.
	Load balancer	● When buying a load balancer, select the same VPC as the ECS. ● Create an EIP for connecting to the jump servers. ● Buy the bandwidth based on the service requirements.

Step 1: Purchasing a Load Balancer and an ECS as a Jump Server in Huawei Cloud CN East-Shanghai1

Step 1 Log in to the ECS console and buy an ECS.

Before installing ICAgent on a non-Huawei Cloud host, buy an ECS as a jump server from Huawei Cloud.

Step 2 Buy a load balancer, add TCP listeners, and associate a backend server group with it.

1. Add listeners for TCP ports 30200, 30201, 8149, 8923, and 8102. For details, see **Adding a TCP Listener**.
2. Add the jump server to a backend server group. For details, see **Backend Server Group**.

Step 3 Configure a security group rule for the jump server and open forwarding ports.

1. Modify the security group rule used by the jump server.
   a. On the ECS console, click the name of the ECS used as the jump server to go to the details page.

b. On the **Security Groups** tab page, click a security group name to go to the details page.

c. Click the **Inbound Rules** tab and click **Add Rule**. Open the inbound ports 8149, 8102, 8923, 30200, 30201, and 80 to ensure that data can be transmitted from the non-Huawei Cloud host to the jump server.

Figure 2-10 Modifying a security group rule

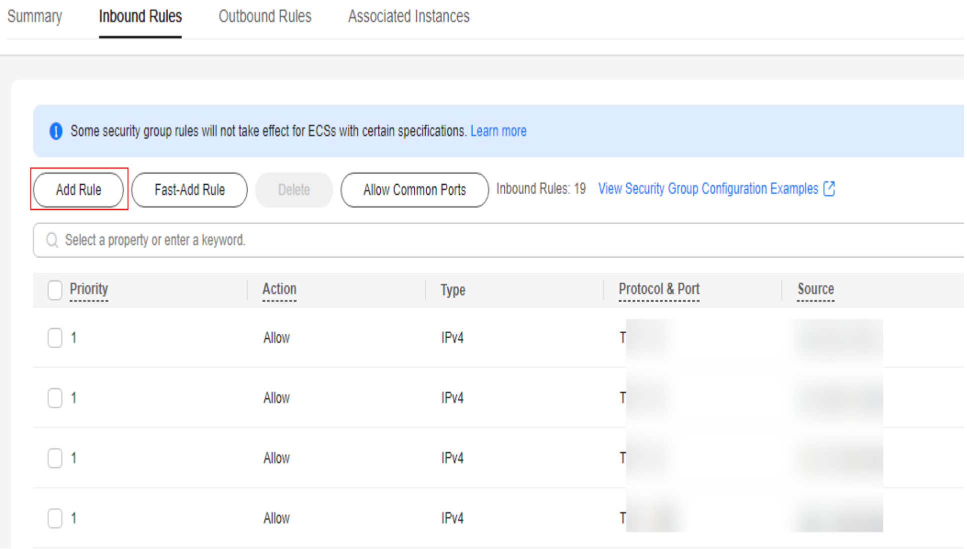

Summary
Inbound Rules
Outbound Rules
Associated Instances
Some security group rules will not take effect for ECSs with certain specifications. Learn more
Add Rule
Fast-Add Rule
Delete
Allow Common Ports
Inbound Rules: 19
View Security Group Configuration Examples
Select a property or enter a keyword.

Priority	Action	Type	Protocol & Port	Source
1	Allow	IPv4	T[redacted]	[redacted]
1	Allow	IPv4	T[redacted]	[redacted]
1	Allow	IPv4	T[redacted]	[redacted]
1	Allow	IPv4	T[redacted]	[redacted]

2. On the LTS console, choose **Host Management** > **Hosts** in the navigation pane, and click **Install ICAgent** in the upper right corner. Set parameters as shown in the following figure. Set **Private IP** to the private IP address of the ECS to generate an installation password.

Figure 2-11 Installing ICAgent

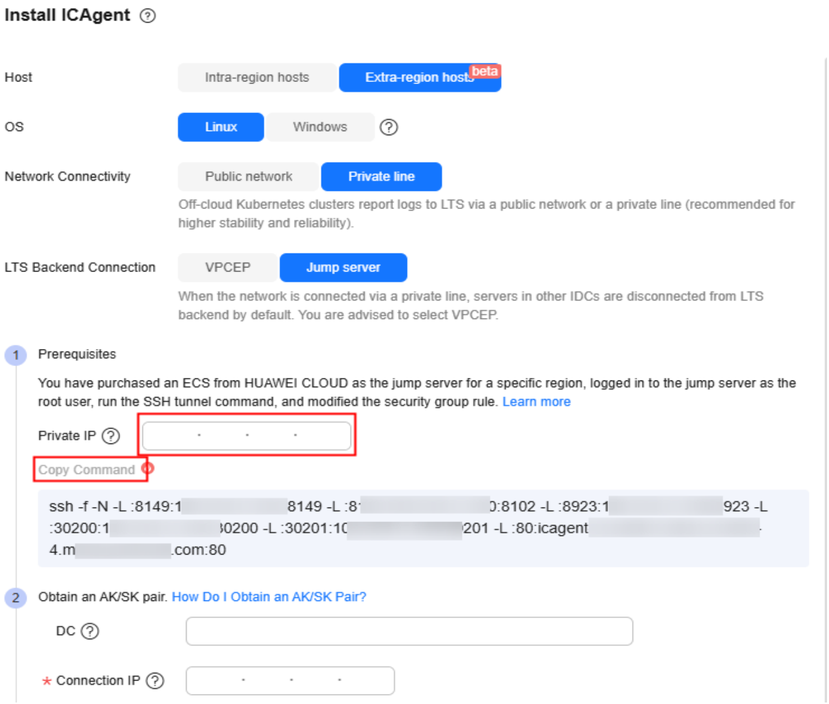

Install ICAgent

Host: Intra-region hosts | Extra-region hosts (selected) beta

OS: Linux (selected) | Windows

Network Connectivity: Public network | Private line (selected)

Off-cloud Kubernetes clusters report logs to LTS via a public network or a private line (recommended for higher stability and reliability).

LTS Backend Connection: VPCEP | Jump server (selected)

When the network is connected via a private line, servers in other IDCs are disconnected from LTS backend by default. You are advised to select VPCEP.

1 Prerequisites

You have purchased an ECS from HUAWEI CLOUD as the jump server for a specific region, logged in to the jump server as the root user, run the SSH tunnel command, and modified the security group rule. Learn more

Private IP: . . .

Copy Command

```
ssh -f -N -L :8149:1[redacted]8149 -L :8[redacted]:8102 -L :8923:1[redacted]923 -L
:30200:1[redacted]0200 -L :30201:1[redacted]201 -L :80:icagent[redacted]
4.m[redacted].com:80
```

2 Obtain an AK/SK pair. How Do I Obtain an AK/SK Pair?

DC:

* Connection IP: . . .

3. Copy the command, log in to the jump server as user **root**, run the SSH tunneling command, and enter the password of user **root** as prompted.
4. Run the **netstat -lnp | grep ssh** command to check whether the corresponding TCP ports are being listened to. If the command output similar to the following is returned, the ports are open.
   - Enter **http://***Jump server IP address* in the address bar of a browser. If the access is successful, the security group rule has taken effect.
   - If the jump server is powered off and then restarted, run the installation command generated on the ICAgent installation page again. If you use the jump server in a production environment, configure the SSH tunneling command to run upon system startup.

Figure 2-12 Viewing ports

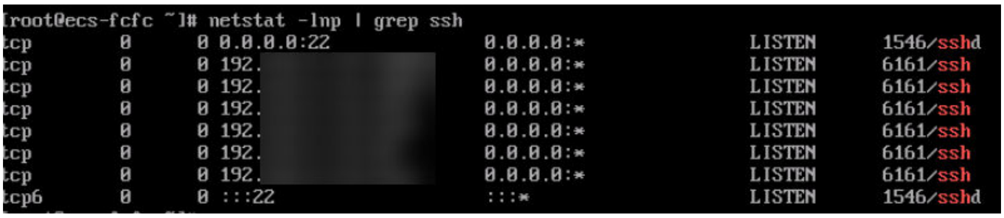

```
[root@ecs-fcfc ~]# netstat -lnp | grep ssh
tcp        0      0 0.0.0.0:22              0.0.0.0:*               LISTEN      1546/sshd
tcp        0      0 192.[redacted]          0.0.0.0:*               LISTEN      6161/ssh
tcp        0      0 192.                    0.0.0.0:*               LISTEN      6161/ssh
tcp        0      0 192.                    0.0.0.0:*               LISTEN      6161/ssh
tcp        0      0 192.                    0.0.0.0:*               LISTEN      6161/ssh
tcp        0      0 192.                    0.0.0.0:*               LISTEN      6161/ssh
tcp        0      0 192.                    0.0.0.0:*               LISTEN      6161/ssh
tcp6       0      0 :::22                   :::*                    LISTEN      1546/sshd
```

----End

Step 2: Configuring Log Ingestion

For a Kubernetes cluster, simply install ICAgent on one node, not all nodes.

Obtain an AK/SK in advance. For details, see **How Do I Obtain an Access Key (AK/SK)?**

Step 1 Configure the jump server.

1. On the ECS console, locate the jump server and obtain its private IP address.

Figure 2-13 Obtaining the private IP address

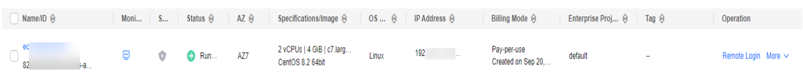

Name/ID	Moni...	S...	Status	AZ	Specifications/Image	OS ...	IP Address	Billing Mode	Enterprise Proj...	Tag	Operation
ec[redacted] 82[redacted]-a...			Run...	AZ7	2 vCPUs | 4 GiB | c7.larg... CentOS 8.2 64bit	Linux	192[redacted]..	Pay-per-use Created on Sep 20,...	default	--	Remote Login More

2. On the LTS console, choose **Host Management** > **Hosts** in the navigation pane and click **Install ICAgent**. On the page displayed, set parameters as follows, set **Private IP** to the private IP address of the ECS to generate an installation command, and copy the command.

Figure 2-14 Installing ICAgent

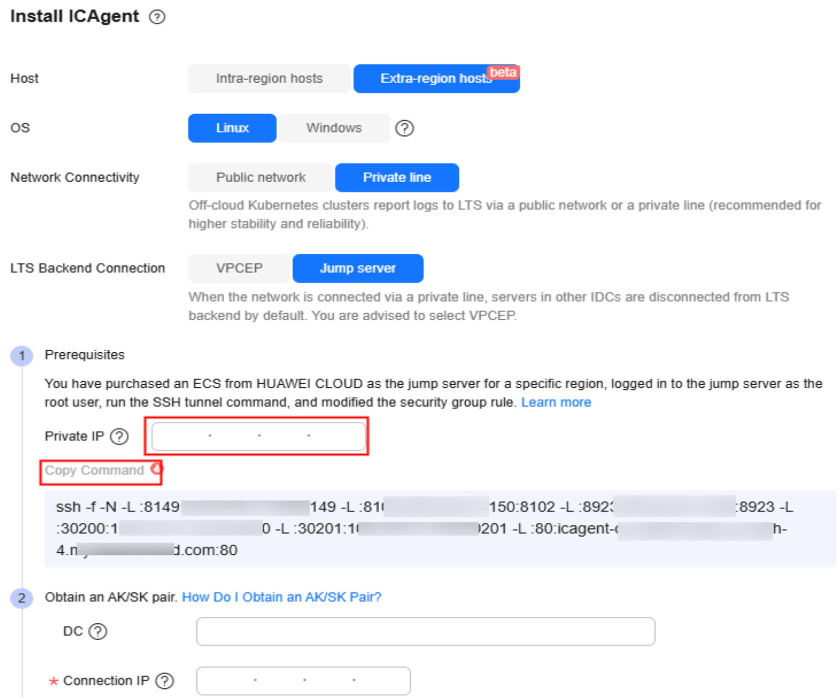

Install ICAgent ⓘ

Host: Intra-region hosts | [✓ Extra-region hosts] beta

OS: [✓ Linux] | Windows ⓘ

Network Connectivity: Public network | [✓ Private line]

Off-cloud Kubernetes clusters report logs to LTS via a public network or a private line (recommended for higher stability and reliability).

LTS Backend Connection: VPCEP | [✓ Jump server]

When the network is connected via a private line, servers in other IDCs are disconnected from LTS backend by default. You are advised to select VPCEP.

1 Prerequisites

You have purchased an ECS from HUAWEI CLOUD as the jump server for a specific region, logged in to the jump server as the root user, run the SSH tunnel command, and modified the security group rule. Learn more

Private IP ⓘ: . . .

Copy Command

```
ssh -f -N -L :8149[redacted]149 -L :81[redacted]150:8102 -L :892[redacted]:8923 -L
:30200:1[redacted]0 -L :30201:1[redacted]201 -L :80:icagent-[redacted]h-
4.n[redacted]d.com:80
```

2 Obtain an AK/SK pair. How Do I Obtain an AK/SK Pair?

DC ⓘ

* Connection IP ⓘ: . . .

3. Log in to the ECS, run the command copied in the previous step, and enter the node password as prompted. If no error is reported, the installation is successful.

Figure 2-15 Running the generated installation command

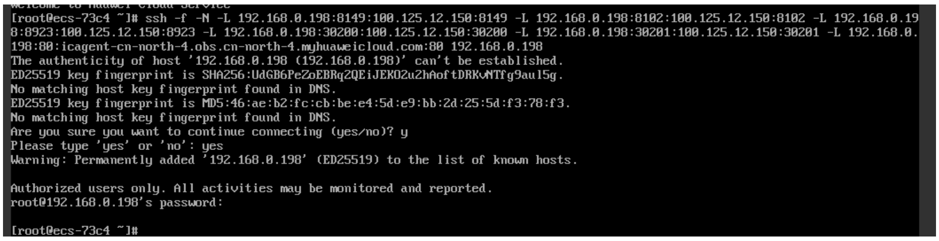

```
Welcome to Huawei Cloud Service
[root@ecs-73c4 ~]# ssh -f -N -L 192.168.0.198:8149:100.125.12.150:8149 -L 192.168.0.198:8102:100.125.12.150:8102 -L 192.168.0.19
8:8923:100.125.12.150:8923 -L 192.168.0.198:30200:100.125.12.150:30200 -L 192.168.0.198:30201:100.125.12.150:30201 -L 192.168.0.
198:80:icagent-cn-north-4.obs.cn-north-4.myhuaweicloud.com:80 192.168.0.198
The authenticity of host '192.168.0.198 (192.168.0.198)' can't be established.
ED25519 key fingerprint is SHA256:UdGB6PeZoEBRq2QEiJEKO2u2hAoftDRKvNTfg9au15g.
No matching host key fingerprint found in DNS.
ED25519 key fingerprint is MD5:46:ae:b2:fc:cb:be:e4:5d:e9:bb:2d:25:5d:f3:78:f3.
No matching host key fingerprint found in DNS.
Are you sure you want to continue connecting (yes/no)? y
Please type 'yes' or 'no': yes
Warning: Permanently added '192.168.0.198' (ED25519) to the list of known hosts.

Authorized users only. All activities may be monitored and reported.
root@192.168.0.198's password:

[root@ecs-73c4 ~]#
```

4. On the **Install ICAgent** page, set **Connection IP** and select the checkbox next to **Turn off command history to prevent the AK/SK from being stored**.

   Connection IP: connection IP address of the jump server. If the jump server uses an EIP to connect to the extra-region host, enter the EIP created by the load balancer. If the jump server and extra-region host use VPC peering connection of Direct Connect, enter the private IP address of the ECS.

Figure 2-16 ICAgent installation page

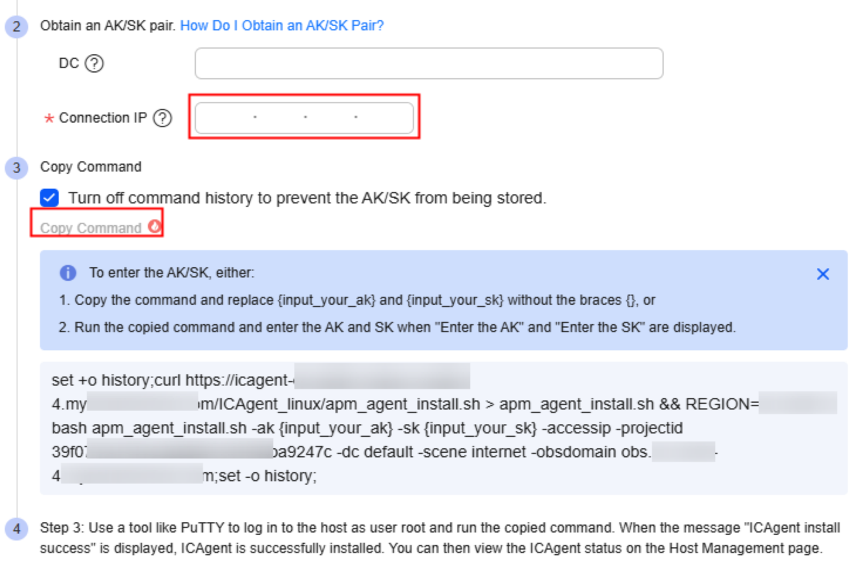

2 Obtain an AK/SK pair. How Do I Obtain an AK/SK Pair?

DC ⓘ

* Connection IP ⓘ

3 Copy Command

☑ Turn off command history to prevent the AK/SK from being stored.

Copy Command

To enter the AK/SK, either:

1. Copy the command and replace {input_your_ak} and {input_your_sk} without the braces {}, or
2. Run the copied command and enter the AK and SK when "Enter the AK" and "Enter the SK" are displayed.

```
set +o history;curl https://icagent-[redacted]
4.my[redacted]m/ICAgent_linux/apm_agent_install.sh > apm_agent_install.sh && REGION=[redacted]
bash apm_agent_install.sh -ak {input_your_ak} -sk {input_your_sk} -accessip -projectid
39f0[redacted]a9247c -dc default -scene internet -obsdomain obs.[redacted]-
4[redacted]m;set -o history;
```

4 Step 3: Use a tool like PuTTY to log in to the host as user root and run the copied command. When the message "ICAgent install success" is displayed, ICAgent is successfully installed. You can then view the ICAgent status on the Host Management page.

5. Copy the ICAgent installation command and run it on the jump server. Enter the AK and SK of the current account as prompted. If the message **ICAgent install success** is displayed, ICAgent is successfully installed.

Step 2 Configure a log ingestion rule. For details, see **Self-Built Kubernetes**.

----End

Step 3: Viewing the Log Stream

On the **Log Management** page of LTS, click the target log stream to go to its details page. If there are logs, the Kubernetes logs have been reported to LTS.

2.3 Collecting Syslog Aggregation Server Logs to LTS

Introduction

System Logging Protocol (Syslog) is a data protocol used by network devices to collect logs to a logging server. It can record log messages of multiple event types and is supported by almost all network devices, such as routers, switches, and printers. Even Unix-like servers can generate syslog messages to record user logins, firewall events, and Apache or Nginx access.

Syslog defines related format specifications based on RFC5424 and RFC3164. RFC3164 was released in 2001, and RFC5424 was an upgraded version released in 2009. The new version is compatible with the old version and solves many problems of the old version. Therefore, RFC 5424 is recommended.

This section describes how to use the syslog protocol to upload logs to LTS. You need to buy an ECS as a syslog aggregation server. Rsyslog comes preinstalled by default on Linux servers. However, ECSs do not receive remote syslog writes by default. You need to enable this function.

Solution Overview

Figure 2-17 Solution flowchart

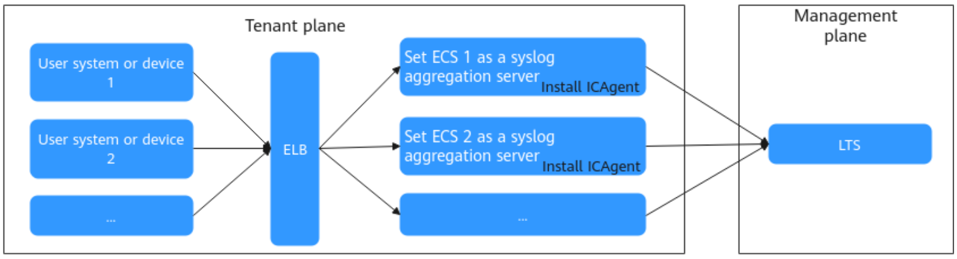

- You can buy a Linux ECS and configure it as a syslog aggregation server to receive log data from other devices. If the size of a log received by a syslog server exceeds 1,024 bytes, the log will be truncated.
- The log processing rate of a single syslog server is 10 MB/s. To process a large number of logs or ensure high reliability, you can buy multiple ECSs as syslog servers and configure load balancers for distributing traffic.
- You need to install ICAgent on syslog servers and configure log collection rules to collect logs to LTS.

Planning Resources

Buy two ECSs. One serves as a syslog aggregation server, and the other serves as a service ECS to simulate user systems or devices to send logs.

an ECS

Step 1 Log in to the management console and choose **Compute** > **Elastic Cloud Server**.

Step 2 Buy an ECS as a syslog aggregation server.

You are advised to use **CentOS 6.5 64bit** or later images. The recommended specifications are 2 vCPUs and 4 GB of memory.

Step 3 Log in to the syslog server as user **root** to install ICAgent.

1. Allow TCP ports 30200, 30201, 8149, 8923, and 8102 in the outbound rules of the syslog server, and then allow UDP port 514 in the inbound rules as the default listening port.
2. Go to the LTS console and choose **Host Management** > **Hosts** in the navigation pane.
3. On the page displayed, click **Install ICAgent**. Set **OS** to **Linux** and **Host** to **Intra-region hosts**. Then, select **Obtain AK/SK** for **Installation Mode**. Click **Copy Command** to copy the ICAgent installation command and manually replace the AK/SK.
4. Log in to the syslog server as user **root** and run the ICAgent installation command. If the message **ICAgent install success** is displayed, ICAgent is successfully installed. You can then view the ICAgent status by choosing **Host Management** in the navigation pane of the LTS console and then clicking **Hosts**.

Step 4 Enable the rsyslog listening and receiving functions.

By default, rsyslog of ECSs does not receive remote syslog writes. You need to manually enable this function.

1. Log in to the ECS.
2. Modify the rsyslog configuration file.
   ```
   vi /etc/rsyslog.conf
   ```
3. Add the following content to the configuration file to enable TCP and UDP remote receiving:
   ```
   # Provides UDP syslog reception
   $ModLoad imudp
   $UDPServerRun 514
   # Provides TCP syslog reception
   $ModLoad imtcp
   $InputTCPServerRun 514
   ```
4. Save the settings. Click **More** > **Restart** in the **Operation** column to restart the ECS.
5. Run any of the following commands. If the command output is normal, the service is working.

   Run **service rsyslog status** to check whether the rsyslog running status is running.

   Figure 2-18 Checking the rsyslog status

   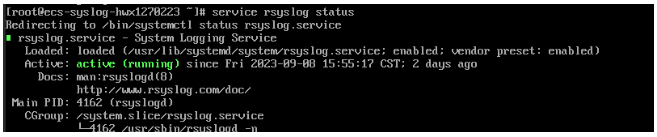

   ```
   [root@ecs-syslog-hwx1270223 ~]# service rsyslog status
   Redirecting to /bin/systemctl status rsyslog.service
   ■ rsyslog.service - System Logging Service
      Loaded: loaded (/usr/lib/systemd/system/rsyslog.service; enabled; vendor preset: enabled)
      Active: active (running) since Fri 2023-09-08 15:55:17 CST; 2 days ago
        Docs: man:rsyslogd(8)
              http://www.rsyslog.com/doc/
    Main PID: 4162 (rsyslogd)
      CGroup: /system.slice/rsyslog.service
              └─4162 /usr/sbin/rsyslogd -n
   ```

   Run **systemctl status rsyslog** to check whether the rsyslog running status is running.

   Figure 2-19 Rsyslog status

   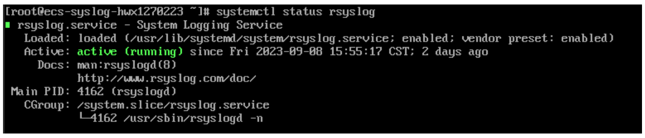

   ```
   [root@ecs-syslog-hwx1270223 ~]# systemctl status rsyslog
   ■ rsyslog.service - System Logging Service
      Loaded: loaded (/usr/lib/systemd/system/rsyslog.service; enabled; vendor preset: enabled)
      Active: active (running) since Fri 2023-09-08 15:55:17 CST; 2 days ago
        Docs: man:rsyslogd(8)
              http://www.rsyslog.com/doc/
    Main PID: 4162 (rsyslogd)
      CGroup: /system.slice/rsyslog.service
              └─4162 /usr/sbin/rsyslogd -n
   ```

   Run **netstat -anp | grep 514** to check whether the listening function is enabled.

   Figure 2-20 Checking whether listening is enabled

   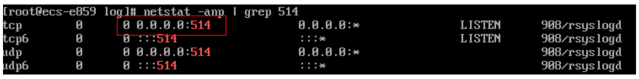

   ```
   [root@ecs-e859 log]# netstat -anp | grep 514
   tcp        0      0 0.0.0.0:514             0.0.0.0:*               LISTEN      908/rsyslogd
   tcp6       0      0 :::514                  :::*                    LISTEN      908/rsyslogd
   udp        0      0 0.0.0.0:514             0.0.0.0:*                           908/rsyslogd
   udp6       0      0 :::514                  :::*                                908/rsyslogd
   ```

Step 5 Configure syslog collection.

1. Choose **Host Management** > **Host Groups** in the LTS navigation pane and click **Create Host Group**. On the displayed page, enter a host group name and select hosts.
2. Choose **Log Ingestion** > **Ingestion Center** in the navigation pane and click **ECS (Elastic Cloud Server)**.

3. Set the collection path to **/var/log/messages**. For details, see .

Step 6 Log in to the service ECS for verification.

After your service system or device generates logs, you can view the logs on the LTS console. Log in to the ECS backend and run the **logger -n** *x.x.x.x* **-P 514 testremotelog** command to send syslog messages to the aggregation server. *x.x.x.x* indicates the IP address (public or private) of the syslog server. **testremotelog** indicates the log content, which can be customized.

After the command is executed, you can view the log in the configured log group and log stream.

Alternatively, log in to the syslog aggregation server and check whether the **testremotelog** log exists in **/var/log/messages**.

```
tail -f /var/log/messages
```

Figure 2-21 Checking whether the testremotelog log exists

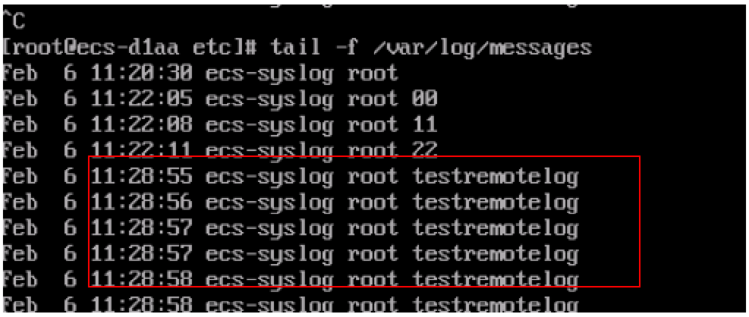

```
^C
[root@ecs-d1aa etc]# tail -f /var/log/messages
Feb  6 11:20:30 ecs-syslog root
Feb  6 11:22:05 ecs-syslog root 00
Feb  6 11:22:08 ecs-syslog root 11
Feb  6 11:22:11 ecs-syslog root 22
Feb  6 11:28:55 ecs-syslog root testremotelog
Feb  6 11:28:56 ecs-syslog root testremotelog
Feb  6 11:28:57 ecs-syslog root testremotelog
Feb  6 11:28:57 ecs-syslog root testremotelog
Feb  6 11:28:58 ecs-syslog root testremotelog
Feb  6 11:28:58 ecs-syslog root testremotelog
```

Step 7 Use multiple syslog servers and load balancers to implement load balancing.

The log processing rate of a single syslog server is 10 MB/s. To process a large number of logs, you can use multiple syslog servers and load balancers.

1. Create syslog aggregation servers and install ICAgent.
2. Create a load balancer. For details, see **Using Load Balancers (Entry Level)**.
3. Add listeners for TCP/UDP ports and port 514. For details, see **Adding a TCP Listener**.
4. Add backend servers to the backend server group. For details, see **Backend Server Group**.

----End

2.4 Importing Logs of Self-built ELK to LTS

Solution Overview

ELK is an acronym that stands for Elasticsearch, Logstash, and Kibana. Together, these three tools provide a most commonly used log analysis and visualization solution in the industry.

- Elasticsearch is an open-source, distributed, and RESTful search and analysis engine based on Lucene.

- Logstash is an open-source data processing pipeline on the server side. It allows you to collect and transform data from multiple sources in real time, and then send the data to your repository. It is usually used to collect, filter, and forward logs.
- Kibana is an open-source platform for data analysis and visualization, enabling you to create dashboards and search and query data. It is usually used together with Elasticsearch.

This section describes how to use custom Python scripts and ICAgent to migrate logs from Elasticsearch to LTS.

ICAgent can be installed on ECSs to collect their log files. With this function, you can import Elasticsearch logs to LTS.

You can flush Elasticsearch data to ECSs using Python scripts, and then collect the flushed log files to LTS using its log ingestion function.

Figure 2-22 Solution flowchart

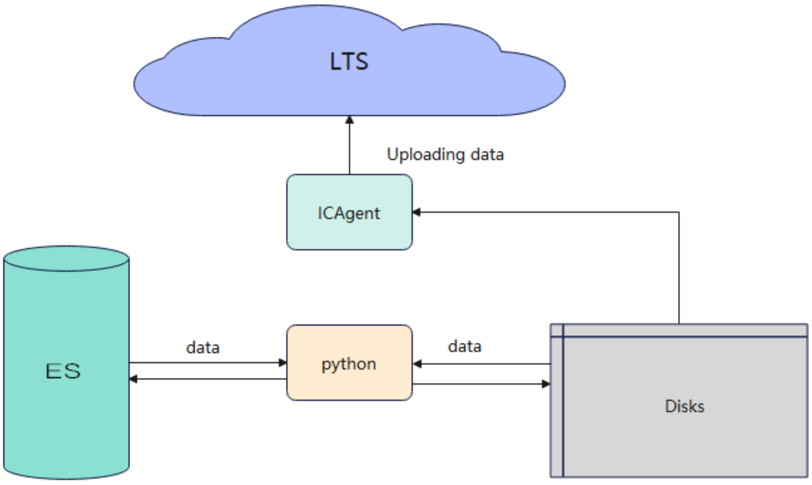

Importing Logs of Self-built ELK to LTS

Step 1 Log in to the management console and choose **Management & Deployment** > **Log Tank Service**.

Step 2 **Install ICAgent** on the ECS.

Step 3 Configure ECS log ingestion on the LTS console. For details, see **Ingesting ECS Text Logs to LTS**.

Step 4 Prepare for script execution. The following example is for reference only. Enter your actual information.

- If you use Python for the first time, you need to install the Python environment.
- If you use Elasticsearch for the first time, you need to install the Python data package of the corresponding Elasticsearch version. Elasticsearch 7.10.1 is used in this solution test.

```
pip install elasticsearch==7.10.1
```

Step 5 Run the python script for constructing index data. If the index already has data, skip this step and go to **Step 6**.

The python script must be executed on the ECS and named *xxx***.py**. The following is an example of constructing data:

Modify the following italic fields as required. In this example, 1,000 data records with the content **This is a test log,Hello world!!!\n** are inserted.

- **index**: name of the index to be created. It is **test** in this example.
- **es**: URL for accessing Elasticsearch. It is **http://127.0.0.1:9200** in this example.

```
from elasticsearch import Elasticsearch
def creadIndex(index):
    mappings = {
        "properties": {
            "content": {
                "type": "text"
            }
        }
    }
    es.indices.create(index=index, mappings=mappings)
def reportLog(index):
    i = 0
    while i < 1000:
        i = i + 1
        body = {"content": "This is a test log,Hello world!!!\n"}
        es.index(index=index,body=body)
if __name__ == '__main__':
    # Index name
    index = 'test'
    # Link to Elasticsearch
    es = Elasticsearch("http://127.0.0.1:9200")
    creadIndex(index)
    reportLog(index)
```

Step 6 Construct the Python read and write script to write Elasticsearch data to the disk. The output file path must be the same as that configured in the log ingestion rule.

The script must be executed on the ECS and named *xxx***.py**. The following is an example of the script for writing data to the disk:

Modify the following italic fields as required.

- **index**: index name. It is **test** in this example.
- **pathFile**: absolute path for writing data to the disk. It is **/tmp/test.log** in this example.
- **scroll_size**: size of the index rolling query. It is **100** in this example.
- **es**: URL for accessing Elasticsearch. It is **http://127.0.0.1:9200** in this example.

```
from elasticsearch import Elasticsearch
def writeLog(res, pathFile):
    data = res.get('hits').get('hits')
    i = 0
    while i < len(data):
        log = data[i].get('_source').get('content')
        file = open(pathFile, 'a', encoding='UTF-8')
        file.writelines(log)
        i = i + 1
    file.flush()
    file.close()
if __name__ == '__main__':
    # Index name
    index = 'test'
    # Output file path
```

```
pathFile = '/tmp/' + index + '.log'
# Size for each scrolling query. The default value is 100.
scroll_size = 100
# Link to Elasticsearch
es = Elasticsearch("http://127.0.0.1:9200")
init = True
while 1:
    if (init == True):
        res = es.search(index=index, scroll="1m", body={"size": scroll_size})
        init =False
    else:
        scroll_id = res.get("_scroll_id")
        res = es.scroll(scroll="1m", scroll_id=scroll_id)
    if not res.get('hits').get('hits'):
        break
    writeLog(res, pathFile)
```

Step 7 Ensure that Python has been installed and run the following command on the ECS to write the Elasticsearch index data to the disk:

```
python xxx.py
```

Step 8 Check whether the data was successfully queried and written into the disk.

In this example, the path for writing data to the disk is **/tmp/test.log**. Replace it with your actual path. Run the following command to check whether the data has been written to the disk:

```
tail -f /tmp/test.log
```

Step 9 Log in to the LTS console. On the **Log Management** page, click the target log stream to go to its details page. If log data is displayed on the **Log Search** tab page, log collection is successful.

----End

2.5 Collecting Zabbix Data Through ECS Log Ingestion

Zabbix is a common open-source monitoring system with various alarm rules. LTS collects monitoring data from Zabbix to log streams. This section describes how LTS collects Zabbix data through ECS log ingestion.

Prerequisites

- Prepare an ECS for log collection. For details, see **Purchasing an ECS**. If you already have an available ECS, skip this step.
- Zabbix has been downloaded and installed on the ECS. For details, see **Download and install Zabbix**.

Configuring a Path for Storing Monitoring Data

Zabbix saves monitoring data on the server where Zabbix is located. You can perform the following steps to set a monitoring data storage path:

1. Log in to the server where Zabbix is located.
2. Open the **zabbix_server.conf** file.
   ```
   vim /etc/zabbix/zabbix_server.conf
   ```
3. Set the data storage path in the file.
   ```
   ExportDir=/tmp/
   ```

4. Restart the Zabbix service for the configuration to take effect.

```
systemctl restart zabbix-server
```

After the configuration takes effect, Zabbix generates a file (with the extension .ndjson) in the **/tmp** directory to save monitoring data.

Ingesting ECS Logs to LTS

Step 1 Choose **Log Ingestion** > **Ingestion Center** in the navigation pane and click **ECS (Elastic Cloud Server)**.

Step 2 The page for selecting a log stream is displayed.

1. Select a log group from the drop-down list of **Log Group**, for example, **lts-group-ECS**.
2. Select a log stream from the drop-down list of **Log Stream**, for example, **lts-topic-ECS**.
3. Click **Next: (Optional) Select Host Group**.

Step 3 Select host groups and click **Next: Configurations**.

Step 4 Set the collection path to **/tem/**/*.ndjson** and retain the default values for other parameters. For details, see **Ingesting ECS Text Logs to LTS**.

Step 5 Click **Next: Index Settings**. On the displayed page, retain the default parameter settings. After configuring the index, you can query and analyze logs.

Step 6 Click **Submit**. After the ingestion configuration is complete, click **Back to Ingestion Configurations**. An ingestion configuration will be displayed on the **Ingestion Management** page.

Step 7 After the log ingestion is configured, you can view the reported logs on the LTS console in real time.

Click the log stream name in the **Log Stream** column of the target ingestion rule to access the stream details page.

Step 8 Click the **Real-Time Logs** tab to view logs in real time.

Logs are reported to LTS once every five seconds. You may wait for at most five seconds before the logs are displayed.

----End

3 Log Search and Analysis

3.1 Analyzing Huawei Cloud ELB Logs on LTS

Solution Overview

When distributing external traffic, ELB logs details of HTTP and HTTPS requests, such as URIs, client IP addresses and ports, and status codes.

You can use ELB access logs for auditing or search for logs by time and keyword. You can also obtain external access statistics by running SQL aggregation queries. For example, you can check the number of requests with 404 responses within a certain day, or analyze the unique visitors (UVs) or page views (PVs) within a week.

Planning Resources

You have purchased and used a load balancer.

Restrictions

ELB access logs only record layer 7 requests sent to the dedicated and shared load balancers. Requests to layer 4 shared load balancers are not logged.

Analyzing Huawei Cloud ELB Logs on LTS

Step 1 Ingest ELB access logs to LTS.

Step 2 Click ☰ in the upper left corner and choose **Management & Governance** > **Log Tank Service**.

Step 3 On the **Log Management** page, click a log stream name. On the log stream details page displayed, click ⚙ in the upper right corner. Go to the **Cloud Structuring Parsing** tab page, retain the default setting (enabled) of **Auto**

Configuration and Analysis, click **Structuring Template**, select the ELB system template, and click **Save**.

----End

3.2 Analyzing Huawei Cloud WAF Logs on LTS

Solution Overview

WAF examines all HTTP and HTTPS requests to detect and block attacks such as SQL injections, cross-site scripting (XSS), Trojan upload, and command or code injections. You can check the access and attack logs for real-time decision-making, device O&M, and service trend analysis.

Analyzing WAF Logs in LTS

You can .

3.3 Analyzing Application Run Logs in Log4j Format on the LTS Console

Introduction

Log4j is Apache's open-source project used for logging. It enables you to output logs to log files, so that development or O&M personnel can calculate the number and proportion of logs at different levels, or gather statistics on services from run logs.

For example, you can know the transaction volume of an offering on a day from logs such as the following:

```
2020-12-28_21:10:48.081 [http-nio-8083-exec-6] INFO  discounted shoes - num is :9
```

Analyzing Application Run Logs in Log4j Format on the LTS Console

Step 1 Log in to the LTS console and choose **Log Ingestion** in the navigation pane.

Step 2 Click **Elastic Cloud Server (ECS)** to configure log ingestion.

Step 3 Select a log stream.

1. Select a log group from the drop-down list of **Log Group**. If there are no desired log groups, click **Create Log Group** to create one.
2. Select a log stream from the drop-down list of **Log Stream**. If there are no desired log streams, click **Create Log Stream** to create one.
3. Click **Next: (Optional) Select Host Group**.

Step 4 Select host groups.

1. Select one or more host groups from which you want to collect logs. If there are no desired host groups, click **Create** above the host group list to create one. For details, see **Managing Host Groups**.

2. Click **Next: Collection Configuration**.

Step 5 Configure the collection.

1. Configure the collection parameters. For details, see **Configurations**.
2. Click **Next: Index Settings**. Retain the default settings.
3. Click **Submit** to complete the ingestion configuration.

Step 6 On the log stream details page, click ⚙. On the **Cloud Structuring Parsing** page, select **Regular Expressions**, select a log event, and extract four fields: **Time1**, **ThreadName**, **Level**, and **Message**.

----End