

Cloud Trace Service

User Guide

Issue 01
Date 2023-11-10



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2023. 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.

Huawei Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road
Qianzhong Avenue
Gui'an New District
Gui Zhou 550029
People's Republic of China

Website: <https://www.huaweicloud.com/intl/en-us/>

Contents

1 Overview	1
2 Querying Traces	3
2.1 Querying Real-Time Traces	3
2.2 Querying Archived Traces	5
3 Trackers	7
3.1 Creating a Tracker	7
3.2 Configuring a Tracker	9
3.3 Disabling or Enabling a Tracker	12
3.4 Deleting a Tracker	13
4 Configuring Key Event Notifications	14
5 Application Examples	18
5.1 Security Auditing	18
5.2 Fault Locating	19
5.3 Resource Tracking	20
6 Trace References	21
6.1 Trace Structure	21
6.2 Example Traces	23
7 Verifying Trace File Integrity	26
7.1 Enabling Verification of Trace File Integrity	26
7.2 Digest Files	26
7.2.1 Overview	27
7.2.2 Digest File Name Format	27
7.2.3 Digest File Structure	27
7.2.4 Example Digest File	28
7.2.5 Digest File Signature	29
7.2.6 Supplementary Information	29
7.3 Verifying Trace File Integrity	30
8 Auditing	35
9 Permissions Management	36
10 Quota Management	38

11 Supported Services and Operations.....	39
A Change History.....	42

1 Overview

Scenarios

If you log in to Cloud Trace Service (CTS) for the first time, click **Enable CTS** on the **Tracker List** page. A management tracker named **system** will be automatically created. Then you can create data trackers on this page. The management tracker identifies and associates with all cloud services your tenant account is using, and records all operations of your tenant account. Data trackers record details of the tenant's operations on data in OBS buckets.

You can only query operation records of the last seven days on the CTS console. To store operation records for more than seven days, you must configure an OBS bucket to transfer records to it. Otherwise, you cannot query the operation records generated seven days ago.

Associated Services

- OBS: used to store trace files.

NOTE

You must select a standard OBS bucket because CTS needs to frequently access the OBS bucket that stores traces.


- Data Encryption Workshop (DEW): Provides keys that can be used to encrypt trace files.
- LTS: stores logs.
- SMN: Sends email or SMS message notifications to users when key operations are performed.

Enabling CTS for the First Time

Step 1 Log in to the management console.

Step 2 If you log in to Huawei Cloud as an administrator, go to **Step 3**. If you log in to Huawei Cloud as an IAM user, first contact your CTS administrator (account owner or a user in the **admin** user group) to obtain the **CTS FullAccess** permissions.

For details, see [Assigning Permissions to an IAM User](#).

- Step 3** Click  in the upper left corner and choose **Management & Deployment > Cloud Trace Service**. The CTS console is displayed.
- Step 4** Choose **Tracker List** in the navigation pane on the left and click **Enable CTS** in the upper right corner. A management tracker named **system** will be automatically created. The management tracker records management traces, which are operations on all cloud resources, such as creation, login, and deletion.
- Step 5** Create trackers (data trackers only). Data trackers record details of the tenant's operations on data in OBS buckets.
- Step 6** Choose **Tracker List** in the navigation pane on the left to view operation records of the last seven days.

----End

2 Querying Traces


2.1 Querying Real-Time Traces



Scenarios


After you enable CTS and the management tracker is created, CTS starts recording operations on cloud resources. After a data tracker is created, the system starts recording operations on data in OBS buckets. CTS stores operation records generated in the last seven days.

This section describes how to query and export operation records of the last seven days on the CTS console.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and choose **Management & Deployment** > **Cloud Trace Service**. The CTS console is displayed.
3. Choose **Trace List** in the navigation pane on the left.
4. On the **Trace List** page, use advanced search to query traces. You can combine one or more filters.
 - **Trace Name, Resource Name, Resource ID, and Trace ID:** Enter a specific name or ID.
 - **Resource Name:** If the cloud resource involved in the trace does not have a resource name or the corresponding API operation does not involve the resource name parameter, leave this field empty.
 - **Resource ID:** Leave this field empty if the resource has no resource ID or if resource creation failed.
 - **Trace Source and Resource Type:** Select a cloud service name or resource type from the corresponding drop-down list.
 - **Operator:** Select one or more operators from the drop-down list.
 - **Trace Status:** Select **normal**, **warning**, or **incident**.

- **normal:** The operation succeeded.
 - **warning:** The operation failed.
 - **incident:** The operation caused a fault that is more serious than the operation failure, for example, causing other faults.
- Time range: Select **Last 1 hour**, **Last 1 day**, or **Last 1 week**, or specify a custom time range.
5. Click **Query**.
 6. On the **Trace List** page, you can also export and refresh the trace list.
 - Click **Export** to export all traces in the query result as file. The file can contain up to 5000 records.
 - Click  to view the latest information about traces.
 7. Click  on the left of a trace to expand its details.



8. Click **View Trace** in the **Operation** column. The trace details are displayed.

View Trace ✕

```

{
  "trace_id": "df201462-8373-11e9-a4db-c3ac3c023b88",
  "code": "302",
  "trace_name": "logout",
  "resource_type": "user",
  "trace_rating": "normal",
  "source_ip": "-",
  "service_type": "IAM",
  "trace_type": "SystemAction",
  "event_type": "system",
  "resource_id": "f3f18b9215014f0d9ded3045af020811",
  "tracker_name": "system",
  "time": "May 31, 2019 15:15:29 GMT+08:00",
  "resource_name": "_____1",
  "record_time": "May 31, 2019 15:15:29 GMT+08:00",
  "user": {
    "name": "_____",
    "id": "f3f18b9215014f0d9ded3045af020811",
    "domain": {
      "name": "_____",
      "id": "2306579dc99f4c8690b14b68e734fcd9"
    }
  }
}
    
```

For details about key fields in the trace structure, see section "Trace Structure" in the *Cloud Trace Service User Guide* and section "Example Traces" in the *Cloud Trace Service User Guide*.

2.2 Querying Archived Traces

Scenarios

CTS periodically sends trace files to OBS buckets. A trace file is a collection of traces. CTS generates trace files based on services and transfer cycle, and adjusts the number of traces contained in each trace file as needed. CTS can also save audit logs to LTS log streams.


This section describes how to view historical operation records in trace files downloaded from OBS buckets and in LTS log streams.

Prerequisites

You have configured a tracker in CTS and enabled **Transfer to OBS** or **Transfer to LTS**. For details, see section "Configuring a Tracker" in the *Cloud Trace Service User Guide*.

Querying Traces Transferred to OBS

If you enable **Transfer to OBS** when configuring the tracker, traces will be periodically transferred to a specified OBS bucket as trace files for long-term storage.

1. Log in to the management console.
2. Click  in the upper left corner and choose **Management & Deployment > Cloud Trace Service**. The CTS console is displayed.
3. Choose **Tracker List** in the navigation pane on the left.
4. Click a bucket in the **OBS Bucket** column.
5. In the OBS bucket, locate the file storage path to view the desired trace, and click **Download** on the right to download the file to the default download path of the browser. If you need to save it to a custom path, click **More > Download As** on the right.
 - The trace file storage path is as follows:
OBS bucket name > CloudTraces > Region > Year > Month > Day > Tracker name > Service directory
An example is ***User-defined name > CloudTraces > region > 2016 > 5 > 19 > system > ECS***.
 - The trace file naming format is as follows:
Trace file prefix_CloudTrace_Region/Region-project_Time when the trace file was uploaded to OBS: Year-Month-DayTHour-Minute-SecondZ_Random characters.json.gz
An example is ***File Prefix_CloudTrace_region-project_2016-05-30T16-20-56Z_21d36ced8c8af71e.json.gz***.

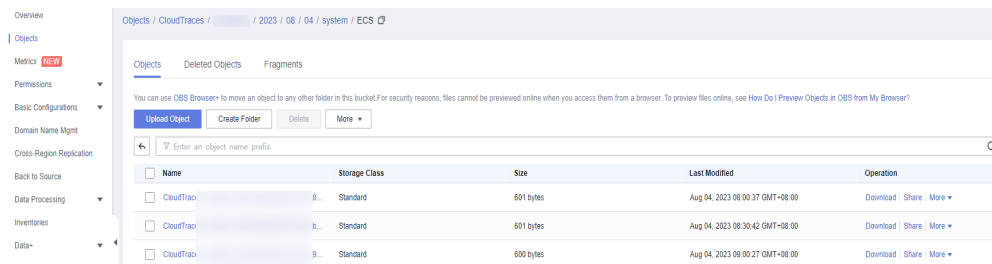
 **NOTE**

The OBS bucket name and trace file prefix are user-defined, and other parameters are automatically generated.

Downloading the file will incur request fees and traffic fees.

For details about key fields in the CTS trace structure, see section "Trace Structure" in the *Cloud Trace Service User Guide* and section "Example Traces" in the *Cloud Trace Service User Guide*.


Figure 2-1 Viewing trace file content




- Decompress the downloaded package to obtain a JSON file with the same name as the package. Open the JSON file using a text file editor to view traces.

Querying Traces Transferred to LTS

If you enable **Transfer to LTS** when configuring a tracker, traces will be transferred to the **CTS/{Tracker Name}** log stream for long-term storage. *{Tracker Name}* indicates the name of the current tracker. For example, the log stream path of the management tracker is **CTS/system-trace**.

- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and choose > **Cloud Trace Service**. The CTS console is displayed.
- Step 3** Choose **Tracker List** in the navigation pane on the left.
- Step 4** Click an LTS log stream in the **Storage** column.
- Step 5** On the **Log Stream** tab page in the **CTS** log group page, select the *{Tracker name}* log stream to view trace logs.

For details about key fields in the CTS trace structure, see section "Trace Structure" in the *Cloud Trace Service User Guide* and section "Example Traces" in the *Cloud Trace Service User Guide*.

- Step 6** Click  to download the log file to your local PC.

 **NOTE**

Each time you can download up to 5,000 log events. If the number of selected log events exceeds 5000, you cannot download them directly from LTS. Transfer them to OBS and then download them from OBS.

----End

3 Trackers

The management tracker records management traces, which are operations on all cloud resources, such as creation, login, and deletion.

3.1 Creating a Tracker

Scenario

You can create data trackers to record operations on data. Data trackers record data traces, which are operations performed by tenants on data in OBS buckets, such as upload and download.

When you enable CTS, a management tracker is created automatically. Only one management tracker can be created. The trackers you created are all data trackers.

NOTE

- If you want to trace operations on OBS buckets, ensure that you create a data tracker in a project instead of a sub-project. Otherwise, tracing fails because there are no sub-projects in OBS.
- CTS records operations performed in the last seven days. To store traces for a longer time, configure a tracker. The tracker will store traces to your specified LTS log streams or OBS buckets.

Prerequisites

You have enabled CTS.

Procedure

1. Log in to the management console.
2. In the service list, choose **Management & Deployment > Cloud Trace Service**. The CTS console is displayed.
3. Choose **Tracker List** in the left navigation pane. In the upper right corner of the displayed page, click **Create Tracker**.
4. Set basic information. Enter a tracker name. Click **Next**.

 **NOTE**

- Tracker name contains only letters, digits, hyphens (-), and underscores (_), and must start with a letter or digit.
 - Tracker name cannot be empty and contains a maximum of 32 characters.
 - The name of the data tracker cannot be system or system-trace.
5. Select a trace. Set parameters and click **Next**.

Table 3-1 Parameters for selecting a trace

Parameter	Description
Data Trace Source	Container for storing data traces. Currently, OBS buckets are used.
OBS Bucket	Select an OBS bucket from the drop-down list.
Operation	<ul style="list-style-type: none"> • Select the operations to record. • Options: Read and Write. Select at least one of them.

6. Configure a transfer. Set parameters and click **Next**.

Table 3-2 Parameters for configuring a transfer

Parameter	Description
Transfer to OBS	If you select Yes , select an existing OBS bucket or create one on the Configure Tracker page and set File Prefix . When Transfer to OBS is disabled, no operation is required.
OBS Bucket	New: If this function is enabled, an OBS bucket will be created automatically with the name you enter. Select Existing: Select an existing OBS bucket.
Select Bucket	When you select New , enter an OBS bucket name. The OBS bucket name cannot be empty. It can contain 3 to 63 characters, including only lowercase letters, digits, hyphens (-), and periods (.). It cannot contain two consecutive periods (for example, my..bucket). A period (.) and a hyphen (-) cannot be adjacent to each other (for example, my-.bucket and my.-bucket). Do not use an IP address as a bucket name. When you select Existing , select an existing OBS bucket.
Retention Period	The duration for storing traces in the OBS bucket. This configuration will apply to the selected bucket and all files in it. Different compliance standards require different trace retention periods. You are advised to set the retention period to at least 180 days. <ul style="list-style-type: none"> • For a data tracker, you can set the duration to 30 days, 60 days, 90 days, 180 days, 3 years, or the same as that of OBS.

Parameter	Description
File Prefix	A prefix is used to mark a transferred trace file. Your specified prefix will be automatically added to the beginning of the name of a transferred file, helping you quickly filter files. Enter 0 to 64 characters. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.
Transfer to LTS	When Transfer to LTS is enabled, traces are transferred to the log stream.
Log group name	When Transfer to LTS is enabled, the default log group name CTS is set. When Transfer to LTS is disabled, no operation is required.

7. Preview the tracker information and click **Create**.
8. Click **OK**.

3.2 Configuring a Tracker

Scenario

You can configure key event notifications of trackers on the CTS console no matter whether the trackers are enabled or not. For enabled trackers, you can also configure **Transfer to OBS** or **Transfer to LTS** for trace transfer.

- You can select an existing OBS bucket for trace transfer. CTS will automatically attach a required policy to the OBS bucket.
- If you modify the trace file prefix of a tracker, the OBS bucket policy will not be affected.

NOTE

There are three storage classes of OBS buckets, Standard, Infrequent Access, and Archive. You must use Standard OBS buckets for trace transfer because CTS needs to frequently access the OBS buckets.


The configuration will take effect immediately after it is complete.


This section describes how to configure a tracker.

Prerequisites

You have enabled CTS and created a tracker.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner to select the desired region and project.

3. Click  in the upper left corner and choose **Management & Deployment > Cloud Trace Service**. The CTS console is displayed.
4. Choose **Tracker List** in the navigation pane on the left.
5. Click **Configure** in the **Operation** column in the row of the target tracker.

(1) Configure a management tracker.

- If you select **Yes** for **Transfer to OBS**, specify the following transfer parameters:
 - i. **OBS Bucket Account** indicates whether the OBS bucket used for transfer belongs to the current user or other users.

NOTE

- The default value for **OBS Bucket Account** is **Logged-in user**.
 - If **OBS Bucket Account** is set to **Logged-in user**, you can select an existing OBS bucket or create an OBS bucket directly on the **Configure Tracker** page.
 - If **OBS Bucket Account** is set to **Other users**, you need to enter the name of the bucket used for transfer. Ensure that you have the write permissions on the bucket and the bucket name entered is correct.
- ii. **Create OBS Bucket:** If this function is enabled, an OBS bucket will be created automatically with the name you enter. If the function is disabled, select an existing OBS bucket.
 - iii. **OBS Bucket:** Select an existing OBS bucket or create an OBS bucket. CTS will automatically attach a required policy to the OBS bucket.

NOTE

Traces are retained in CTS for 7 days. You can transfer traces to OBS buckets for long-term storage.

- iv. **File Prefix:** A prefix is used to mark a transferred trace file. Your specified prefix will be automatically added to the beginning of the name of a transferred file, helping you quickly filter files.
 - v. **Encrypt Trace File** indicates whether trace files stored in the OBS bucket are encrypted. If you set it to **Yes**, you also need to select a KMS key.
 - vi. **Verify Trace File:** When this function is enabled, integrity verification will be performed to check whether trace files in OBS buckets have been tampered with. For details about file integrity verification, see [Verifying Trace File Integrity](#).
- If you select **No** for **Transfer to OBS**, you do not need to set the transfer parameters.

(2) Configure a data tracker.

- **Operation:** Select the data operations to record.
- **Transfer to OBS**

If you select **Yes**, select an existing OBS bucket or create one directly on the **Configure Tracker** page and set **File Prefix**.

If you select **No**, you do not need to configure the following parameters.

- **Create OBS Bucket:** If this function is enabled, an OBS bucket will be created automatically with the name you enter. If the function is disabled, select an existing OBS bucket.
 - **OBS Bucket:** You can create an OBS bucket or select an existing OBS bucket.
 - **Retention Period:** Select the duration for storing traces in the OBS bucket.
 - **File Prefix:** A prefix is used to mark a transferred trace file. Your specified prefix will be automatically added to the beginning of the name of a transferred file, helping you quickly filter files.
 - **Verify Trace File:** When this function is enabled, integrity verification will be performed to check whether trace files in OBS buckets have been tampered with. For details about file integrity verification, see [Verifying Trace File Integrity](#).
6. In the **Configure Transfer** step, you can modify the transfer settings of the tracker. For details about the parameters, see [Table 3-3](#).

Table 3-3 Parameters for configuring a transfer

Parameter	Description
Transfer to OBS	If you select Yes , select an existing OBS bucket or create one on the Configure Tracker page and set File Prefix . When Transfer to OBS is disabled, no operation is required.
OBS Bucket	New: If this function is enabled, an OBS bucket will be created automatically with the name you enter. Select Existing: Select an existing OBS bucket.
Select Bucket	When you select New , enter an OBS bucket name. The OBS bucket name cannot be empty. It can contain 3 to 63 characters, including only lowercase letters, digits, hyphens (-), and periods (.). It cannot contain two consecutive periods (for example, my..bucket). A period (.) and a hyphen (-) cannot be adjacent to each other (for example, my-.bucket and my.-bucket). Do not use an IP address as a bucket name. When you select Existing , select an existing OBS bucket.
Retention Period	The duration for storing traces in the OBS bucket. This configuration will apply to the selected bucket and all files in it. Different compliance standards require different trace retention periods. You are advised to set the retention period to at least 180 days. <ul style="list-style-type: none"> • For a data tracker, you can set the duration to 30 days, 60 days, 90 days, 180 days, 3 years, or the same as that of OBS.

Parameter	Description
File Prefix	A prefix is used to mark a transferred trace file. Your specified prefix will be automatically added to the beginning of the name of a transferred file, helping you quickly filter files. Enter 0 to 64 characters. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.
Transfer to LTS	When Transfer to LTS is enabled, traces are transferred to the log stream.
Log group name	When Transfer to LTS is enabled, the default log group name CTS is set. When Transfer to LTS is disabled, no operation is required.

7. Click **Next > Configure** to complete the configuration of the tracker.
You can then view the tracker details on the **Tracker List** page.

 **NOTE**



Traces recorded by CTS are delivered periodically to the OBS bucket for storage. If you configure an OBS bucket for a tracker, traces generated during the current cycle (usually several minutes) will be delivered to the configured OBS bucket. For example, if the current cycle is from 12:00:00 to 12:05:00 and you configure an OBS bucket for a tracker at 12:02:00, traces received from 12:00:00 to 12:02:00 will also be delivered to the configured OBS bucket for storage at 12:05:00.

3.3 Disabling or Enabling a Tracker

Scenario

You can disable a tracker on the CTS console. After a tracker is disabled, it will stop recording operations, but you can still view operation records that have been collected.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner to select the desired region and project.
3. Click  in the upper left corner and choose **Management & Deployment > Cloud Trace Service**. The CTS console is displayed.
4. Choose **Tracker List** in the navigation pane on the left.
5. Click **Disable** in the **Operation** column in the row of the target tracker.
6. Click **Yes**.

After the tracker is disabled, the **Disable** button changes to **Enable**. To enable the tracker, click **Enable** and then click **OK**. The tracker will start recording operations again.

3.4 Deleting a Tracker

Scenario

Deleting a data tracker on the CTS console is available, and does not affect the existing operation records. This section describes how to delete a data tracker on the management console.



NOTE

When you enable CTS, a management tracker is created automatically. Only one management tracker can be created and it cannot be deleted.

Prerequisites

A data tracker has been created.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner to select the desired region and project.
3. Click  in the upper left corner and choose **Management & Deployment > Cloud Trace Service**. The CTS console is displayed.
4. Choose **Tracker List** in the navigation pane on the left.
5. Click **Delete** in the **Operation** column of the target configuration item.

NOTE

The system tracker cannot be deleted.

6. Click **Yes**.

4 Configuring Key Event Notifications

You can create key event notifications on CTS so that SMN sends you SMS, email, or HTTP/HTTPS notifications of key events. This function is triggered by CTS, and notifications are sent by SMN. SMN sends key event notifications to subscribers. Before setting notifications, you need to know how to create topics and add subscriptions on the SMN console.

Scenarios


You can use this function for:

- Real-time detection of high-risk operations (such as VM restart and security configuration changes), cost-sensitive operations (such as creating and deleting expensive resources), and service-sensitive operations (such as network configuration changes).
- Detection of operations such as login of users with admin-level permissions or operations performed by users who do not have the required permissions.
- Connection with your own audit system: You can synchronize all audit logs to your audit system in real time to analyze the API calling success rate, unauthorized operations, security, and costs.

Prerequisites

- SMN sends key event notifications to subscribers. Before setting notifications, you need to know how to create topics and add subscriptions on the SMN console.
- You can create up to 100 key event notifications on CTS:
 - Specify key operations, users, and topics to customize notifications.
 - Complete key event notifications can be sent to notification topics.
- If CTS and Cloud Eye use the same message topic, they can receive messages from the same terminal but with different content.
- You can configure one key event notification for operations initiated by a maximum of 50 users in 10 user groups. For each key event notification, you can add users from different user groups, but cannot select multiple user groups at once.

Creating a Key Event Notification


1. Log in to the management console.
2. Click  in the upper left corner and choose **Management & Deployment > Cloud Trace Service**. The CTS console is displayed.
3. In the navigation pane on the left, choose **Key Event Notifications**.
The **Key Event Notifications** page is displayed.
4. Click **Create Key Event Notification**. On the displayed page, specify required parameters.
5. Enter a key event notification name.
Notification Name: Identifies key event notifications. This parameter is mandatory. The name can contain up to 64 characters. Only letters, digits, and underscores (_) are allowed.
6. Configure key operations.
Select the operations that will trigger notifications. When a selected operation is performed, an SMN notification is sent immediately.
 - **Operation Type:** Select **All**, **Typical**, or **Custom**.
 - **All:** This option is suitable if you have connected CTS to your own audit system. When **All** is chosen, you cannot deselect operations because all operations on all cloud services that have connected with CTS will trigger notifications. You are advised to use an SMN topic for which HTTPS is selected.
 - **Custom:** This option is suitable for enterprises that require detection of high-risk, cost-sensitive, service-sensitive, and unauthorized operations. You can connect CTS to your own audit system for log analysis.
Customize the operations that will trigger notifications. Up to 1000 operations of 100 services can be added for each notification. For details, see section "Supported Services and Operations" in the *Cloud Trace Service User Guide*.
 - **Advanced Filter:** You can set an advanced filter to specify the operations that will trigger notifications. Operations can be filtered by fields **api_version**, **code**, **trace_rating**, **trace_type**, **resource_id**, and **resource_name**. Up to six filter conditions can be set. When you configure multiple conditions, specify whether an operation is considered a match when all conditions are met (AND) or any of the conditions are met (OR).
7. Configure users.
SMN messages will be sent to subscribers when the specified users perform key operations.
 - If you select **All users**, SMN will notify subscribers of key operations initiated by all users.
 - If you select **Specified users**, SMN will notify subscribers of key operations initiated by your specified users. You can configure key event notifications on operations for up to 50 users in 10 user groups. For each notification, you can select multiple users in the same user group.

8. Configure an SMN topic.
 - When **Yes** is selected for **Send Notification**:
 - **SMN Topic**: You can select an existing topic or click **SMN** to create one on the SMN console.
 - If you do not want to send notifications, no further action is required.
9. Click **OK**.

Managing Key Event Notifications

After you create a key event notification, you can view its name, status, template, and SMN topic in the notification list and delete it as required.


Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and choose **Management & Deployment > Cloud Trace Service**. The CTS console is displayed.

Step 3 Choose **Key Event Notifications** in the navigation pane on the left. On the displayed page, perform the following operations as required. For details, see [Table 4-1](#).

Table 4-1 Related operations

Operation	Description
Viewing a key event notification	Click View in the Operation column to view the operation list and user list details of the notification.
Enable/Disable a key event notification	Click Enable or Disable in the Operation column. NOTE CTS can trigger key event notifications only after SMN is configured.
Modifying a key event notification	Click Modify in the Operation column to modify the configuration of the key event notification.
Deleting a key event notification	Click Delete in the Operation column.

Operation	Description
Refreshing the key event notification list	Click  in the upper right corner.

----End

5 Application Examples

5.1 Security Auditing

Scenario



You can query operation records matching specified conditions and check whether operations have been performed by authorized users for security analysis.

Prerequisites

You have enabled CTS and trackers are running properly.

Procedure (for Old Console)

The following takes the records of EVS disk creation and deletion in the last two weeks as an example.

1. Log in to the management console as a CTS administrator.
2. Click  in the upper left corner to select the desired region and project.
3. Click  in the upper left corner and choose **Management & Deployment > Cloud Trace Service**. The CTS console is displayed.
4. Choose **Trace List** in the left navigation pane.
5. Set the time range to **Last 1 week**, set filters in sequence, and click **Query**.

NOTE

Select **Management** for **Trace Type**, **evs** for **Trace Source**, **evs** for **Resource Type**, **Trace name** for **Search By**, select **createVolume** or **deleteVolume**, and click **Query**. By default, all EVS disk creation or deletion operations performed in the last hour are queried. You can also set the time range to query all EVS creation or deletion operations performed in the last seven days at most.

6. To obtain the operation records of the last week, query them in the OBS bucket. Choose **Tracker List** in the navigation pane on the left.

 NOTE

To store operation records for more than seven days, you must configure the management tracker to transfer them to an OBS bucket. Otherwise, you cannot query the operation records generated seven days ago.

7. Download traces older than seven days or all traces by following the instructions in [Querying Archived Traces](#).
8. In the trace files, search traces using keywords **createVolume** or **deleteVolume**.
9. Check the traces obtained from steps 5 and 8 to see whether there are any unauthorized operations or operations that do not conform to security rules.

5.2 Fault Locating

Scenario



If a resource or an action encounters an exception, you can query operation records of the resource or action in a specified time period and view the requests and responses to facilitate fault locating.

Prerequisites

You have enabled CTS and trackers are running properly.

Procedure (for Old Console)

The following shows how to locate an ECS fault which occurred in a morning.

1. Log in to the management console as a CTS administrator.
2. Click  in the upper left corner to select the desired region and project.
3. Click  in the upper left corner and choose **Management & Deployment > Cloud Trace Service**. The CTS console is displayed.
4. Choose **Trace List** in the left navigation pane.
5. Set filters in sequence and click **Query**.



 NOTE

Select **Management** for **Trace Type**, **ECS** for **Trace Source**, **ecs** for **Resource Type**, **Resource ID** for **Search By**, and enter the ID of the faulty virtual machine (VM). In the upper right corner, select a time range from 06:00:00 to 12:00:00 on the day when the fault occurred. Then, click **Query** to view the result.

6. Check the returned traces, especially the request type and response of each trace. Pay attention to traces whose status is **warning** or **incident**, and traces whose response indicates a failure.

The following shows how to locate a fault after an ECS server failed to be created.

1. Log in to the management console as a CTS administrator.

2. Click  in the upper left corner to select the desired region and project.
3. Click  in the upper left corner and choose **Management & Deployment > Cloud Trace Service**. The CTS console is displayed.
4. Choose **Trace List** in the left navigation pane.
5. Select **Management** for **Trace Type**, **ECS** for **Trace Source**, **ecs** for **Resource Type**, and **Warning** for **Trace Status**. In the returned traces, locate the trace named **createServer**.
6. Check the trace details and locate the fault based on the error code or error message.

5.3 Resource Tracking

Scenario



You can view operation records of a cloud resource throughout its lifecycle.

Prerequisites

You have enabled CTS and trackers are running properly.

Procedure (for Old Console)

The following takes the records of all operations on an ECS server as an example.

1. Log in to the management console as a CTS administrator.
2. Click  in the upper left corner to select the desired region and project.
3. Click  in the upper left corner and choose **Management & Deployment > Cloud Trace Service**. The CTS console is displayed.
4. Choose **Trace List** in the left navigation pane.
5. Set filters in sequence and click **Query**.

NOTE

Select **Management** for **Trace Type**, **ECS** for **Trace Source**, **ecs** for **Resource Type**, **Resource ID** for **Search By**, enter the ID of the faulty VM, and click **Query**. By default, the matching traces generated in the last hour are returned. You can also set the time range to view the matching traces in the last seven days at most.

6. Choose **Tracker List** in the navigation pane on the left.
7. Download traces older than seven days or all traces by following the instructions in [Querying Archived Traces](#).
8. Check all the traces obtained in [5](#) and [7](#).

6 Trace References

6.1 Trace Structure

A trace consists of multiple key fields shown in [Table 6-1](#).

 **NOTE**

- This section describes the key trace fields displayed on the CTS console.
- When some fields are displayed on the CTS console, their formats are optimized for easy understanding.

Table 6-1 Key trace fields

Field	Mandatory	Type	Description
time	Yes	Date	Time when a trace occurred When the field is displayed on the console, its value is the local standard time (in GMT time), for example, Dec 8, 2016 11:24:04 GMT+08:00 . However, this field is transmitted and stored as a timestamp in APIs. In this case, the value is the number of milliseconds since 00:00:00 on January 1, 1970 (GMT).
user	Yes	Structure	Cloud account used to perform an operation The value is also displayed in the Operator column on the Trace List page. This field is transmitted and stored as a string in APIs.

Field	Mandatory	Type	Description
request	No	Structure	Requested operation This field is transmitted and stored as a string in APIs.
response	No	Structure	Response to a request This field is transmitted and stored as a string in APIs.
service_type	Yes	String	Operation source
resource_type	Yes	String	Resource type
resource_name	No	String	Resource name
resource_id	No	String	Unique resource ID
source_ip	Yes	String	IP address of the user that performed an operation The value of this field is empty if the operation was triggered by system.
trace_name	Yes	String	Operation name
trace_rating	Yes	String	Trace status. The value can be normal , warning , or incident . <ul style="list-style-type: none"> • normal: The operation succeeded. • warning: The operation failed. • incident: The operation caused a serious consequence, for example, a node failure or service interruption.
trace_type	Yes	String	Operation type There are three types of operations: <ul style="list-style-type: none"> • ConsoleAction: operations performed on the management console • SystemAction: operations triggered by system • ApiCall: operations triggered by calling API Gateway

Field	Mandatory	Type	Description
api_version	No	String	Version of the cloud service API which was called to perform an operation
message	No	Structure	Remarks
record_time	Yes	Number	Time when the operation was recorded, in the form of a timestamp
trace_id	Yes	String	Unique operation ID
code	No	Number	HTTP return code, such as 200 or 400
request_id	No	String	ID of a recorded request
location_info	No	String	Information required for fault locating after a request error
endpoint	No	String	Endpoint in the detail page URL of the cloud resource on which a recorded operation was performed
resource_url	No	String	Detail page URL (excluding the endpoint) of the cloud resource on which a recorded operation was performed

6.2 Example Traces

This section provides two example traces and describes their key fields to help you better understand traces. You can read other traces in a similar way as shown below.

For details on the fields in a trace file, see [Trace Structure](#).

ECS Server Creation

```
{
  "time": "2016/12/08 11:07:28 GMT+08:00",
  "user": {
    "name": "aaa/op_service",
    "id": "f2fe9fac63414a35a7d03108d5f1ea73",
    "domain": {
      "name": "aaa",
      "id": "1f9b9ba51f6b4061bd5c1736b28469f8"
    }
  },
  "request": {
    "server": {
      "name": "as-config-15f1_XWO68TFC",
      "imageRef": "b2b2c7dc-bbb0-4d6b-81dd-f0904023d54f",
      "flavorRef": "m1.tiny",

```

```
    "personality": [],
    "vpcid": "e4c374b9-3675-482c-9b81-4acd59745c2b",
    "nics": [
      {
        "subnet_id": "fff89132-88d4-4e5b-9e27-d9001167d24f",
        "nictype": null,
        "ip_address": null,
        "binding:profile": null,
        "extra_dhcp_opts": null
      }
    ],
    "adminPass": "*****",
    "count": 1,
    "metadata": {
      "op_svc_userid": "26e96eda18034ae9a44130bacb967b96"
    },
    "availability_zone": "az1.dc1",
    "root_volume": {
      "volumetype": "SATA",
      "extendparam": {
        "resourceSpecCode": "SATA"
      },
      "size": 40
    },
    "data_volumes": [],
    "security_groups": [
      {
        "id": "dd597fd7-d119-4994-a22c-891fcdc54be1"
      }
    ],
    "key_name": "KeyPair-3e51"
  }
},
"response": {
  "status": "SUCCESS",
  "entities": {
    "server_id": "42d39b4a-19b7-4ee2-b01b-a9f1353b4c54"
  },
  "job_id": "4010b39d58b855980158b8574b270018",
  "job_type": "createSingleServer",
  "begin_time": "2016-12-01T03:04:38.437Z",
  "end_time": "2016-12-01T03:07:26.871Z",
  "error_code": null,
  "fail_reason": null
},
"service_type": "ECS",
"resource_type": "ecs",
"resource_name": "as-config-15f1_XWO68TFC",
"resource_id": "42d39b4a-19b7-4ee2-b01b-a9f1353b4c54",
"source_ip": "",
"trace_name": "createSingleServer",
"trace_rating": "normal",
"trace_type": "SystemAction",
"api_version": "1.0",
"record_time": "2016/12/08 11:07:28 GMT+08:00",
"trace_id": "4abc3a67-b773-11e6-8412-8f0ed3cc97c6"
}
```

You can pay special attention to the following fields:

- **time** indicates the time when the trace occurred. In this example, the time is 11:07:28 on December 8.
- **user** indicates the user who performed the operation. In this example, the user is **aaa (name field)** under the enterprise account **aaa (domain field)**.
- **request** indicates the request to create an ECS server. It contains basic information about the ECS server, such as its name (**as-**

config-15f1_XWO68TFC) and VPC ID (**e4c374b9-3675-482c-9b81-4acd59745c2b**).

- **response** indicates the response to the ECS creation request. It contains **status** (**SUCCESS** in this example), **error_code** (**null** in this example), and **fail_reason** (**null** in this example).

EVS Disk Creation

```
{
  "time": "2016/12/08 11:24:04 GMT+08:00",
  "user": {
    "name": "aaa",
    "id": "26e96eda18034ae9a44130bacb967b96",
    "domain": {
      "name": "aaa",
      "id": "1f9b9ba51f6b4061bd5c1736b28469f8"
    }
  },
  "request": "",
  "response": "",
  "service_type": "EVS",
  "resource_type": "evs",
  "resource_name": "volume-39bc",
  "resource_id": "229142c0-2c2e-4f01-a1b4-2dfdf1c678c7",
  "source_ip": "10.146.230.124",
  "trace_name": "deleteVolume",
  "trace_rating": "normal",
  "trace_type": "ConsoleAction",
  "api_version": "1.0",
  "record_time": "2016/12/08 11:24:04 GMT+08:00",
  "trace_id": "c529254f-bcf5-11e6-a89a-7fc778a6c92c"
}
```

You can pay special attention to the following fields:

- **time** indicates the time when the trace occurred. In this example, the time is 11:24:04 on December 8.
- **user** indicates the user who performed the operation. In this example, the user is **aaa** (**name** field) under the enterprise account **aaa** (**domain** field).
- **request**: optional. It is null in this example.
- **response**: optional. It is null in this example.
- **trace_rating** indicates the trace status. It can replace the **response** field to indicate the operation result. In this example, the value is **normal**, indicating that the operation was successful according to [Trace Structure](#).



7 Verifying Trace File Integrity

7.1 Enabling Verification of Trace File Integrity

Scenario

During a security investigation, operational records will not be able to serve as effective and authentic evidence if they are deleted or tampered with. You can enable the integrity verification on CTS to ensure the authenticity of trace files. The verification is performed only on management traces.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner to select the desired region and project.
3. Click  in the upper left corner and choose **Management & Deployment** > **Cloud Trace Service**. The CTS console is displayed.
4. Choose **Tracker List** in the navigation pane on the left.

 **NOTE**

Click **Enable CTS** if you have not enabled CTS.

5. Click **Configure** in the row of the management tracker **system**. On the displayed **Configure Tracker** page, click **Next**, and enable **Verify Trace File** in the **Configure Transfer** step.

 **NOTE**

CTS supports integrity verification of trace files only for the management tracker configured with OBS transfer.

7.2 Digest Files

7.2.1 Overview

A digest file contains the names and hash values of the trace files transferred to an OBS bucket an hour ago as well as the digital signature of the previous digest file. The digital signature of this digest file is stored in metadata attributes of the digest file object. A digest file is stored in the following path:

OBS bucket name > **CloudTraces** > *Region* > *Year* > *Month* > *Day* > *tracker name* > **Digest** > *Service*

An example is *OBS bucket name* > **CloudTraces** > *Region* > **2016** > **5** > **19** > **system** > **Digest** > **ECS**.

7.2.2 Digest File Name Format

*Trace file prefix***_CloudTrace-Digest_Region/Region-Project_Time** when the digest file was sent to OBS: *Year-Month-DayT Hour-Minute-Second***Z.json.gz**

An example is *Trace file prefix***_CloudTrace-Digest_Region/Region-Project_2016-05-30T16-20-56Z.json.gz**.

7.2.3 Digest File Structure

Table 7-1 Key fields of a digest file

Field	Mandatory	Type	Description
project_id	Yes	String	Identifies the account to which a trace file covered in the digest file belongs.
digest_start_time	Yes	String	Specifies the start of the UTC time range covered by the digest file.
digest_end_time	Yes	String	Specifies the end of the UTC time range covered by the digest file.
digest_bucket	Yes	String	Specifies the name of the OBS bucket that the digest file has been sent to.
digest_object	Yes	String	Specifies where the digest file is stored in the OBS bucket.
digest_signature_algorithm	Yes	String	Specifies the algorithm used to sign the digest file.
digest_end	Yes	Boolean	Specifies whether the digest file is an ending digest file.
previous_digest_bucket	No	String	Specifies the name of the OBS bucket that the previous digest file was sent to.
previous_digest_object	No	String	Specifies where the previous digest file is stored in the OBS bucket.
previous_digest_hash_value	No	String	Specifies the hexadecimal encoded hash value of the previous digest file.

Field	Mandatory	Type	Description
previous_digest_hash_algorithm	No	String	Specifies the Hash algorithm used to hash the previous digest file.
previous_digest_signature	No	String	Specifies the digital signature of the previous digest file.
previous_digest_end	Yes	Boolean	Specifies whether the previous digest file is an ending digest file.
log_files	No	Array	Specifies the list of trace files covered in the digest file.
bucket	Yes	String	Specifies the name of the OBS bucket that the trace files have been sent to.
object	Yes	String	Specifies where the trace files are stored in the OBS bucket.
log_hash_value	Yes	String	Specifies the hexadecimal encoded hash value of the trace files.
log_hash_algorithm	Yes	String	Specifies the Hash algorithm used to hash the trace files.

7.2.4 Example Digest File

```
{
  "project_id": "3cfb09080bd944d0b4cdd72ef2685712",
  "digest_start_time": "2017-03-28T01-09-17Z",
  "digest_end_time": "2017-03-28T02-09-17Z",
  "digest_bucket": "bucket",
  "digest_object": "CloudTraces/eu-west-1012017/3/28/Digest/EVS/mylog_CloudTrace-Digest_eu-west-101_2017-03-28T02-09-17Z.json.gz",
  "digest_signature_algorithm": "SHA256withRSA",
  "digest_end": false,
  "previous_digest_bucket": "bucket",
  "previous_digest_object": "CloudTraces/eu-west-101/2017/3/28/Digest/EVS/mylog_CloudTrace-Digest_eu-west-101_2017-03-28T01-09-17Z.json.gz",
  "previous_digest_hash_value": "5e08875de01b894eda5d1399d7b049fe",
  "previous_digest_hash_algorithm": "MD5",
  "previous_digest_signature":
    "7af7cbef4f3c78eab5048030d402810841400cf70eb22f93d4fedd13b13a8208a5dc04ce2f8bd0a4918f933ca3fc
    b17595ae59386a2dc3e3046fa97688a9815a2d036fa10193534c0ebbecff67221e22ac9cf8b781cbae3a81eaccfc
    0a2bd1a99081b1e4fe99b19caa771876ba7cce16d002feb4578cd89bc6f1faaca639ab48f3cb56007bcc5e248968
    f4a17a95b8cd8bc7d8bbd7c63630da878cd4d471fc75c60bac5f730d94fefe8fdd2f2fa8accd62dbe100eab7773e79
    15e91be4474291b9dacea63a8267390bcb4855b5825554ebb07d4a29ce077c364213c575c461d1e9fafa0c29fde
    1c6de1d5630e015200821b2f3ae91e53cd8591433dd7c0b4c8bc",
  "previous_digest_end": false,
  "log_files": [{
    "bucket": "bucket",
    "object": "CloudTraces/eu-west-101/2017/3/28/ECS/mylog_CloudTrace_eu-west-101_2017-03-28T02-09-17Z_0faa86bc40071242.json.gz",
    "log_hash_value": "633a8256ae7996e21430c3a0e9897828",
    "log_hash_algorithm": "MD5"
  }]
}
```


7.2.5 Digest File Signature

The digital signature information of a digest file is in two metadata attributes of the digest file object. Each digest file has the following two metadata items:

- meta-signature

Hexadecimal encoded value of the digest file signature. Example:

```
7af7cbef4f3c78eab5048030d402810841400cf70eb22f93d4fedd13b13a8208a5dc04ce2f8bd0a4918f933c
a3fcb17595ae59386a2dc3e3046fa97688a9815a2d036fa10193534c0ebbecff67221e22ac9cf8b781cbae3
a81eaccfc0a2bd1a99081b1e4fe99b19caa771876ba7cce16d002feb4578cd89bc6f1faaca639ab48f3cb560
07bcc5e248968f4a17a95b8cdabc7d8bbd7c63630da878cd4d471fc75c60bac5f730d94fefe8fdd2f2fa8accd
62dbe100eab7773e7915e91be4474291b9dacea63a8267390bcb4855b5825554ebb07d4a29ce077c3642
13c575c461d1e9fafa0c29fde1c6de1d5630e015200821b2f3ae91e53cd8591433dd7c0b4c8bc
```

- meta-signature-algorithm

Algorithm used to sign the digest file. Example:

SHA256withRSA

7.2.6 Supplementary Information

- Starting Digest File

A starting digest file is generated after you start verifying trace file integrity. In a starting digest file, the following fields related to the previous digest file will be left empty:

- previous_digest_bucket
- previous_digest_object
- previous_digest_hash_value
- previous_digest_hash_algorithm
- previous_digest_signature

- "Empty" Digest File

CTS will still send a digest file even if no operations have occurred in your account within the one-hour time period recorded by the digest file. The last field **log_files:[]** of the digest file will be left empty. It helps you to confirm that no trace files have been sent within the one-hour time period recorded by the digest file.

- Digest File Chain

A digest file contains the digital signature and Hash value of the previous digest file (if any) so that a chain is formed. You can verify digest files successively within a specified time, starting with the latest one.

- Digest File Bucket

A digest file is sent to the OBS bucket that stores trace files recorded in the file.

- Digest File Storage Folder

A digest file is stored in a folder different from that for trace files, making it easy for you to execute fine-grained security policies.

7.3 Verifying Trace File Integrity

Scenario

CTS uses public signature algorithms and hash functions in accordance with industry standards, so you can create tools on your own to verify integrity of CTS trace files. Trace files should contain fields **time**, **service_type**, **resource_type**, **trace_name**, **trace_rating**, and **trace_type** for integrity verification. Other fields can be added by services from which traces are collected.

After you enable integrity verification of trace files in CTS, digest files will be sent to your OBS buckets, and you can implement your own verification solution. For details about digest files, see [Digest Files](#).

Prerequisites

You should understand how digest files are signed.

RSA digital signatures are used in CTS. For each digest file, CTS will:

1. Create a message for digital signing, a character string composed of specified digest file fields, and obtain an RSA private key.
2. Produce a hash value of the digest message. Use the RSA algorithm to generate a digital signature with the hash value and private key, and encode the digital signature to hexadecimal format.
3. Put the digital signature into the meta-signature attribute of the digest file object.

The message for digital signing contains the following digest file fields:

- The ending timestamp of the UTC time range covered by the digest file, for example, 2017-03-28T02-09-17Z.
- The path where the current digest file is stored in the OBS bucket.
- The hash value (hexadecimal encoded) of the current digest file (compressed).
- The hexadecimal digital signature of the previous digest file.

Procedure

Verify a digest file first and then its referenced trace files.

1. Obtain a digest file.
 - a. Download **OBS Java SDK** and call an OBS client interface to obtain the latest digest file within the required time period from OBS buckets.
 - b. Check whether the location where the digest file is stored in the OBS bucket matches with the location recorded in the file.
 - c. Obtain the digital signature from the meta-signature attribute of the digest file object.
2. Obtain the RSA public key for verifying the digital signature.

The RSA public key of CTS is as follows:

```
MIIBlJANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAsjQDkl8COPRhoCvm7ZI8sYZ20ojl+ay/
gwRSk9q0gkY3pP0RrAhSsEzgYdYjaMCqixkmbpt4AH9AROU4drnoCAZSMqRxcv0bGC9kVd4q95l4zibsw
AsksjuNQo/XoJjBl+rRAqCa+1uetgVU4k4Yx8RryYxYx/tlmvMe/O4mGAlaTf+rsqt3VXR1Qlj5lYR/nx41BEgC/
Kb1eLYAfDaaab8WS5lNRprj7qdu6oAo4Ug47WqbecvEtG3JRpj5+oqLyW41Fvse3osC0h5DQdxTt4x00/rVZ
+gH7Kua00y7gC8YOxVpYbfn/oW61PUDeHG/N9hUjOrlgDDJpD2YbClQIDAQAB
```

3. Recreate the message for digital signing.

Compute the message for digital signing.

The message is in the following format:

```
signature_string = digest_end_time
+ digest_object
+ Hex(hash(digest-file-content))
+ previous_digest_signature
```

4. Verify a digest file.

Pass the computed message obtained in [3](#), digital signature of the digest file, and public key to the RSA signature verification algorithm. If **true** is returned, the digital signature of the digest file matches with the computed message and the digest file is valid.

5. Verify trace files.

You can verify trace files referenced by the digest file after confirming that the digest file is valid.

The digest file records the hash value of each trace file. After a trace file is uploaded to OBS, its hash value will be stored in ETag metadata. If the trace file is modified after CTS sent it to an OBS bucket, the file's hash value will change, and the digital signatures of the digest file will not match.

Do as follows to verify a trace file:

- a. Obtain **bucket** and **object** information about a trace file from the digest file.
- b. Call the OBS client interface to obtain the ETag metadata value in the trace file object header.
- c. Obtain the hash value of the trace file from the **log_hash_value** field in the digest file.
- d. Compare the ETag metadata value with the hash value obtained in the previous step. If they match, the trace file is valid.

6. Verify the previous digest files and trace files.

In each digest file, the following fields provide the location and signature of the previous digest file:

- previous_digest_bucket
- previous_digest_object
- previous_digest_signature

Repeat steps [4](#) and [5](#) to verify the signature of each previous digest file and all trace files that the file references.

For these previous digest files, you do not need to obtain the digital signature from the meta-signature attribute of the digest file object. The **previous_digest_signature** field in each digest file provides the digital signature of the previous digest file. You can keep verifying the previous digest files and their referenced trace files until you reach the starting digest file or the digest file chain is disconnected.

The following code segment is an example for verifying CTS digest and trace files. The code segment uses the following JAR packages, and you are recommended to use these packages:

- esdk-obs-java-2.1.16.jar
- commons-logging-1.2.jar
- httpasyncclient-4.1.2.jar
- httpclient-4.5.3.jar
- httpcore-4.4.4.jar
- httpcore-nio-4.4.4.jar
- java-xmlbuilder-1.1.jar
- jna-4.1.0.jar
- log4j-api-2.8.2.jar
- log4j-core-2.8.2.jar
- commons-codec-1.9.jar
- json-20160810.jar
- commons-io-2.5.jar

Example code segment:

```
import java.io.BufferedInputStream;
import java.io.BufferedReader;
import java.io.ByteArrayInputStream;
import java.io.InputStream;
import java.io.InputStreamReader;
import java.security.KeyFactory;
import java.security.MessageDigest;
import java.security.PublicKey;
import java.security.Signature;
import java.security.spec.X509EncodedKeySpec;
import java.util.Arrays;
import java.util.zip.GZIPInputStream;

import org.apache.commons.codec.binary.Base64;
import org.apache.commons.codec.binary.Hex;
import org.apache.commons.io.IOUtils;
import org.json.JSONObject;

import com.obs.services.ObsClient;
import com.obs.services.ObsConfiguration;
import com.obs.services.model.ObjectMetadata;
import com.obs.services.model.S3Object;

public class DigestFileValidator {
    public static void main(String[] args) {
        // Name of the bucket where a digest file is located.
        String digestBucket = "bucketname";
        // Path where a digest file is stored. Example: CloudTraces/eu-de/2017/11/15/Digest/ECS/
tGPYa_CloudTrace-Digest_eu-de_2017-11-15T10-12-10Z.json.gz.
        String digestObject = "digestObject";

        // Directly writing AK/SK in code is risky. For security, encrypt your AK/SK and store them in the
        configuration file or environment variables.
        // In this example, the AK/SK are stored in environment variables for identity authentication.
        Before running this example, set environment variables HUAWEICLOUD_SDK_AK and
HUAWEICLOUD_SDK_SK.
        String ak = System.getenv("HUAWEICLOUD_SDK_AK");
        String sk = System.getenv("HUAWEICLOUD_SDK_SK");

        ObsConfiguration obsConfig = new ObsConfiguration();
        obsConfig.setEndPoint("obs.eu-west-101.myhuaweicloud.com");
```

```
ObsClient client = new ObsClient(ak, sk, obsConfig);

try {
    // Obtain a digest file object.
    S3Object object = client.getObject(digestBucket, digestObject);

    InputStream is = new BufferedInputStream(object.getObjectContent());
    byte[] digestFileBytes = IOUtils.toByteArray(is);

    // Obtain the hash value of a digest file.
    MessageDigest messageDigest = MessageDigest.getInstance("MD5");
    messageDigest.update(digestFileBytes);
    byte[] digestFileHashBytes = messageDigest.digest();

    StringBuilder outStr = new StringBuilder();
    GZIPInputStream gis = new GZIPInputStream(new ByteArrayInputStream(digestFileBytes));
    BufferedReader bufferedReader = new BufferedReader(new InputStreamReader(gis, "UTF-8"));
    String line;
    while ((line = bufferedReader.readLine()) != null) {
        outStr.append(line);
    }
    bufferedReader.close();
    String digestInfo = outStr.toString();

    // Obtain the meta-signature value from the digest file header in an OBS bucket, which is the
    digital signature of the digest file.
    ObjectMetadata objectMetadata = client.getObjectMetadata(digestBucket, digestObject);
    String digestSignature = objectMetadata.getMetadata().get("meta-signature").toString();
    JSONObject digestFile = new JSONObject(digestInfo);
    // Check whether the digest file has been moved in the OBS bucket.
    if (!digestFile.getString("digest_bucket").equals(digestBucket) || !
digestFile.getString("digest_object")
    .equals(digestObject)) {
        System.err.println("Digest file has been moved from its original location.");
    } else {
        // Obtain the message for digital signing.
        String signatureString = digestFile.getString("digest_end_time") +
digestFile.getString("digest_object")
        + Hex.encodeHexString(digestFileHashBytes) +
digestFile.getString("previous_digest_signature");

        String publicKeyString
        =
"MIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAsjQDkl8COPRhoCvm7ZI8sYZ20ojl+ay/
gwRSk9q0gkY3pP0RrAhSsEzgYdYjaMCqixkmbpt4AH9AROUJ4drnoCAZSMqRxcv0bGC9kVd4q95L4zibsw
AsksjuNQo/XoJlBl+rRAqCa+1uetgVU4k4Yx8RryYxYx/tlmvMe/O4mGAlaTf+rsqt3VXR1Qlj5LYR/nx41BEgC/
Kb1elYAfDaaab8WSS5INRprj7qdu6oAo4Ug47WqbecvEtG3JRpj5+oqLyW41Fvse3osC0h5DQdxTt4x00/rVZ
+gH7Kua00y7gC8YOxVpYbfn/oW61PUDeHG/N9hUjOrlgDDJpD2YbClQIDAQAB";

        // Public key used for decryption.
        byte[] publicKeyBytes = Base64.decodeBase64(publicKeyString);
        // Form the X509EncodedKeySpec object.
        X509EncodedKeySpec x509EncodedKeySpec = new X509EncodedKeySpec(publicKeyBytes);

        // Specify an encryption algorithm.
        KeyFactory keyFactory = KeyFactory.getInstance("RSA");
        // Obtain the public key object.
        PublicKey publicKey = keyFactory.generatePublic(x509EncodedKeySpec);

        Signature signatureInstance = Signature.getInstance("SHA256withRSA");
        signatureInstance.initVerify(publicKey);
        signatureInstance.update(signatureString.getBytes("UTF-8"));

        byte[] signatureHashExpect = Hex.decodeHex(digestSignature.toCharArray());

        // Verify whether the signature is valid.
        if (signatureInstance.verify(signatureHashExpect)) {
            System.out.println("Digest file signature is valid, validating log files...");
        }
    }
}
```

```
for (int i = 0; i < digestFile.getJSONArray("log_files").length(); i++) {
    JSONObject logFileJson = digestFile.getJSONArray("log_files").getJSONObject(i);
    String logBucket = logFileJson.getString("bucket");
    String logObject = logFileJson.getString("object");

    // Obtain the ETag value from the trace file header in the OBS bucket, which is the
    recorded hash value of the trace file.
    ObjectMetadata objectLogMetadata = client.getObjectMetadata(logBucket,
logObject);
    String logHashValue = objectLogMetadata.getMetadata().get("ETag").toString();
    logHashValue = logHashValue.replace("\\", "");
    byte[] logFileHash = Hex.decodeHex(logHashValue.toCharArray());

    // Obtain the hash value of each trace file from the digest file.
    byte[] expectedHash = logFileJson.getString("log_hash_value").getBytes();
    boolean hashMatch = Arrays.equals(expectedHash, logFileHash);
    if (!hashMatch) {
        System.err.println("Validate log file hash failed.");
    } else {
        System.out.println("Log file hash is valid.");
    }
}
} else {
    System.err.println("Validate digest signature failed.");
}

System.out.println("Digest file validation completed.");

// Obtain values of fields previous_digest_bucket, previous_digest_object, and
previous_digest_signature of the previous digest file. After obtaining the digest file, verify its hash
value and digital signature.
String previousDigestBucket = digestFile.getString("previous_digest_bucket");
String previousDigestObject = digestFile.getString("previous_digest_object");

// Obtain the digital signature from the meta-signature attribute of the digest file object
header.
ObjectMetadata objectPreviousMetadata = client.getObjectMetadata(previousDigestBucket,
previousDigestObject);
String signatruePrevious = objectPreviousMetadata.getMetadata().get("meta-
signature").toString();
String signatruePreviousExpect = digestFile.getString("previous_digest_signature");
if (signatruePrevious.equals(signatruePreviousExpect)) {
    System.out.println(
        "Previous digest file signature is valid, " + "validating previous digest file hash
value...");

    String digestPreviousHashValue =
objectPreviousMetadata.getMetadata().get("ETag").toString();
    // The ETag metadata value is the trace file hash value enclosed with quotation marks.
    You need to remove the quotation marks.
    String digestPreviousHashValueExpect = "\"" +
digestFile.getString("previous_digest_hash_value")
        + "\"";
    if (digestPreviousHashValue.equals(digestPreviousHashValueExpect)) {
        System.out.println("Previous digest file hash value is valid.");
    } else {
        System.err.println("Validate previous digest file hash value failed.");
    }
}
}
} catch (Exception e) {
    System.out.println("Validate digest file failed.");
}
}
```

8 Auditing

Cloud Trace Service (CTS) provides records of operations performed on cloud service resources.

With CTS, you can record operations associated with CTS itself for later query, audit, and backtracking.

Table 8-1 CTS operations that can be recorded by itself

Operation	Resource Type	Trace Name
Creating a tracker	tracker	createTracker
Modifying a tracker	tracker	updateTracker
Disabling a tracker	tracker	updateTracker
Enabling a tracker	tracker	updateTracker
Deleting a tracker	tracker	deleteTracker
Creating a key event notification	notification	createNotification
Deleting a key event notification	notification	deleteNotification
Modifying a key event notification	notification	updateNotification
Changing the status of a key event notification	notification	updateNotificationStatus
Disabling a key event notification	notification	updateNotification
Enabling a key event notification	notification	updateNotification
Exporting traces	trace	getTrace

9 Permissions Management

This chapter describes how to use IAM for fine-grained permissions control for your CTS resources. With IAM, you can:

- Create IAM users for employees based on your enterprise's organizational structure. Each IAM user will have their own security credentials for accessing CTS resources.
- Manage permissions on a principle of least permissions (PoLP) basis.
- Entrust an account or cloud service to perform efficient O&M on your CTS resources.

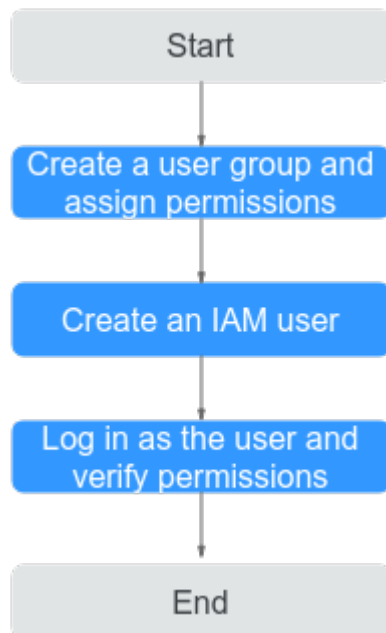
If your account does not require individual IAM users, skip this chapter.

Prerequisites

Learn about the permissions (see section "Permissions Management" in the *Service Overview*) supported by CTS and choose policies or roles according to your requirements.

Process Flow

Figure 9-1 Process of granting CTS permissions



1. **Create a user group and assign permissions.**
Create a user group on the IAM console, and attach the **CTS Administrator** policy to the group.
2. **Assign permissions to an IAM user.**
Create a user on the IAM console and add the user to the user group created in **1**.
3. **Log in** and verify permissions.
Log in to the CTS console using the user created in **2**, and verify that the user has the administrator permissions for CTS.


10 Quota Management

What Are Quotas?

A quota is a limit on the quantity or capacity of a certain type of service resources that you can use, for example, the maximum number of key event notifications that you can create.

If the existing resource quota cannot meet your requirements, you can apply for a higher quota.

How Do I View My Quotas?

1. Log in to the management console.
2. Click  on the upper left corner and choose a region and project.
3. In the upper right corner of the page, choose **Resources > My Quotas**.
The **Service Quota** page is displayed.
4. Check the total quotas and used quotas of resources.

11 Supported Services and Operations

Table 11-1 Supported services and operations

Category	Cloud Service	Operations
Compute	Elastic Cloud Server (ECS)	ECS operations that can be recorded by CTS
	Image Management Service (IMS)	IMS operations that can be recorded by CTS
	Auto Scaling (AS)	AS operations that can be recorded by CTS
Storage	Elastic Volume Service (EVS)	EVS operations that can be recorded by CTS
	Scalable File Service (SFS)	SFS operations that can be recorded by CTS
	Cloud Backup and Recovery (CBR)	CBR operations that can be recorded by CTS
Network	Direct Connect (DC)	DC operations that can be recorded by CTS
	Virtual Private Cloud (VPC)	VPC operations that can be recorded by CTS
	Elastic Load Balance (ELB)	ELB operations that can be recorded by CTS
	Virtual Private Network (VPN)	VPN operations that can be recorded by CTS
Container	Cloud Container Engine (CCE)	CCE operations that can be recorded by CTS
Management & Governance	Cloud Trace Service (CTS)	CTS operations that can be recorded by itself

Category	Cloud Service	Operations
	Identity and Access Management (IAM)	IAM operations that can be recorded by CTS
	Tag Management Service (TMS)	TMS operations that can be recorded by CTS
	Simple Message Notification (SMN)	SMN operations that can be recorded by CTS
Applications and Middleware	Distributed Message Service for Kafka (DMS for Kafka)	DMS for Kafka operations that can be recorded by CTS
	Distributed Message Service for RocketMQ (DMS for RocketMQ)	DMS for RocketMQ operations that can be recorded by CTS
	Distributed Cache Service (DCS)	DCS operations that can be recorded by CTS
	API Gateway (APIG)	API Gateway operations that can be recorded by CTS
Database	Relational Database Service (RDS)	RDS for MySQL operations that can be recorded by CTS
		RDS for PostgreSQL operations that can be recorded by CTS
	Document Database Service (DDS)	DDS operations that can be recorded by CTS
	Data Replication Service (DRS)	DRS operations that can be recorded by CTS
	Distributed Database Middleware (DDM)	DDM operations that can be recorded by CTS
Security & Compliance	Data Encryption Workshop (DEW)	DEW operations that can be recorded by CTS
	Web Application Firewall (WAF)	WAF operations that can be recorded by CTS
	Database Security Service (DBSS)	DBSS operations that can be recorded by CTS
	Data Security Center (DSC)	DSC operations that can be recorded by CTS
Big Data	MapReduce Service (MRS)	MRS operations that can be recorded by CTS
	Data Lake Insight (DLI)	DLI operations that can be recorded by CTS

Category	Cloud Service	Operations
	GaussDB(DWS)	GaussDB(DWS) operations that can be recorded by CTS
	Cloud Search Service (CSS)	CSS operations that can be recorded by CTS
	DataArts Studio	DataArts Studio operations that can be recorded by CTS

A Change History

Released On	Description
2023-11-10	This issue is the third official release, which incorporates the following change: Updated Supported Services and Operations .
2023-02-28	This issue is the second official release, which incorporates the following change: Modified Supported Services and Operations .
2022-09-30	This issue is the first official release.