

Cloud Operations Center (COC)

FAQs

Issue	01
Date	2025-09-11



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

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Cloud Computing Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

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

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

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 Product Consulting	1
1.1 How Do I Configure Permissions for the COC?	1
1.2 How Do I Control Permissions Using Enterprise Projects?	2
1.3 How Can I Create an Agency?	5
2 Resource Management FAQs	16
2.1 How Do I Install UniAgent for the First Time?	16
2.2 What Can I Do If Resources Cannot Be Queried on the Resource Management Page?	19
2.3 How Can I Find the Description About Application Management Layers?	19
3 Patch Management FAQs	21
3.1 What Can I Do If the Patch Baselines Do Not Take Effect?	21
3.2 What Are the Differences Between the Installation Rule Baselines and Custom Baselines?	21
3.3 What Can I Do If Exception "all mirrors were tried" Is Recorded in the Patch Service Ticket Log?	23
3.4 Why Can't I Select a Node?	23
3.5 What Can I Do if a Patch Remains Non-Compliant After Repaired?	23
3.6 What Can I Do If "lsb_release not found" Occurs During Patch Operations?	24
4 Automation FAQs	25
4.1 Why Can't the Reviewer Receive Notifications?	25
4.2 Why Is the Input Value of a Customized Script Parameter Invalid?	25
4.3 Why Can't I Select an Instance?	26
4.4 How Do I Reset the Password Without Restarting a DB Instance?	26
4.5 What Can I Do If I Am Not Authorized to View Passwords on the Account Management Page?	27
5 Batch Operation FAQs	30
5.1 What Should I Do If an Error Is Reported When I Change Images for ECS Resources in Batches?	30
6 FAQs About Parameter Management	31
6.1 What Are the Permissions Required for Managing Parameters?	31
6.2 Can I Reference Parameters in the Parameter Center and Target Instances Across Regions?	32
7 FAQs About Resource O&M	33
7.1 Resource O&M Permissions and Supported Actions	33
8 FAQs About Fault Management	40
8.1 What Is the Process of Generating an Incident?	40

8.2 How Can I Receive an Incident Ticket Notification?	41
8.3 What Is a War Room?	41
8.4 What Can I Do If a Fault Diagnosis Task Is Abnormal?	41
9 FAQs About Change Ticket Management	46
9.1 What Are the Differences Between Regular Changes and Emergency Changes?	46
9.2 How Are Change Levels Defined?	46
10 Resilience Center FAQs	49
10.1 What Is a Chaos Drill?	49
10.2 What Are the Available Attack Scenarios?	49
10.3 What Is a Failure Mode?	65
10.4 What Do Drill Plans Do?	65
10.5 What Is the Relationship Between a Failure Mode and a Drill Task?	66
10.6 What Are Included in a Drill Report?	66
11 FAQs About Basic Configurations	68
11.1 How Do I Log In to COC as a Non-Common IAM User?	68

1 Product Consulting

1.1 How Do I Configure Permissions for the COC?

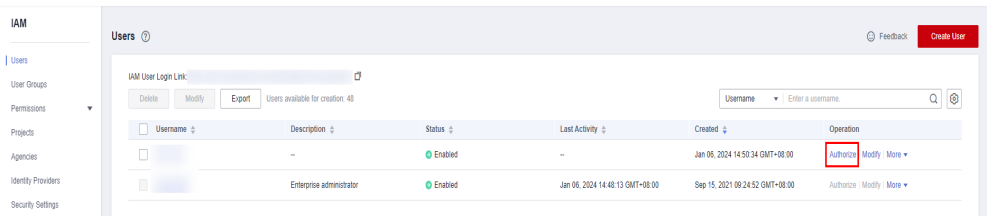
Issue Description

Quickly configuring permissions for COC is required.

Solution

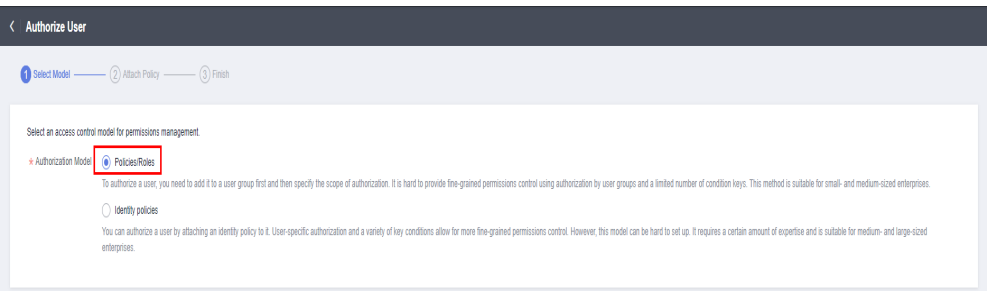
1. Log in to **IAM** as an administrator.
2. In the user list, click **Authorize** in the row that contains the target user.

Figure 1-1 Authorizing an IAM user

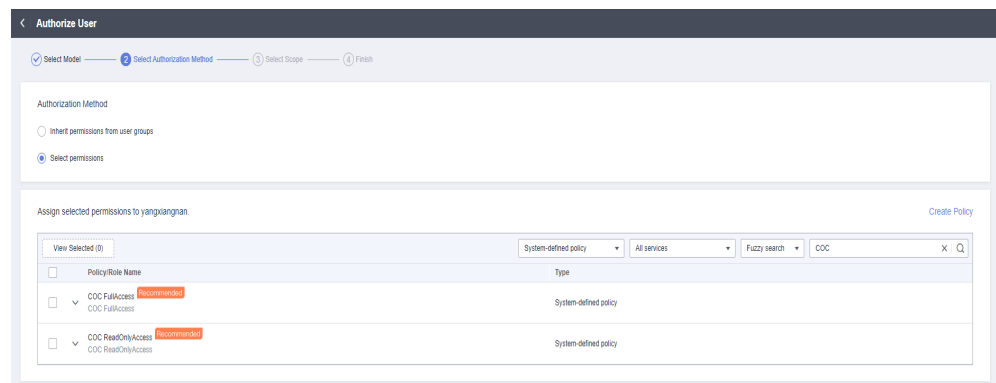


3. Set **Authorization Model** to **RBAC**.

Figure 1-2 Selecting an authorization model

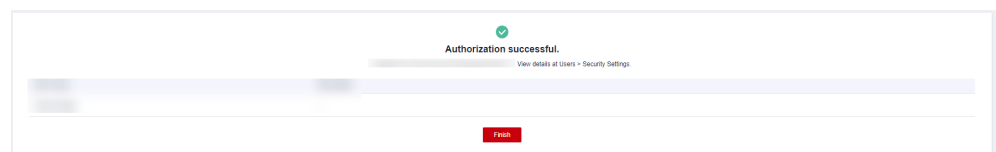


4. Select **Grant permissions to the user** (applicable to enterprise projects), and assign the **COC FullAccess** or **COC ReadOnlyAccess** policy to the user as required. For details about the policy, see **COC Permissions Management**.

Figure 1-3 Granting COC policies**NOTE**

If there is a group that has been assigned permissions of COC, you can select the button for inheriting the policies of the selected user group. For details, see [IAM User Authorization](#).

5. Select an authorization scope scheme and specify enterprise project resources.
6. Wait until the authorization is complete.

Figure 1-4 Successful authorization

1.2 How Do I Control Permissions Using Enterprise Projects?

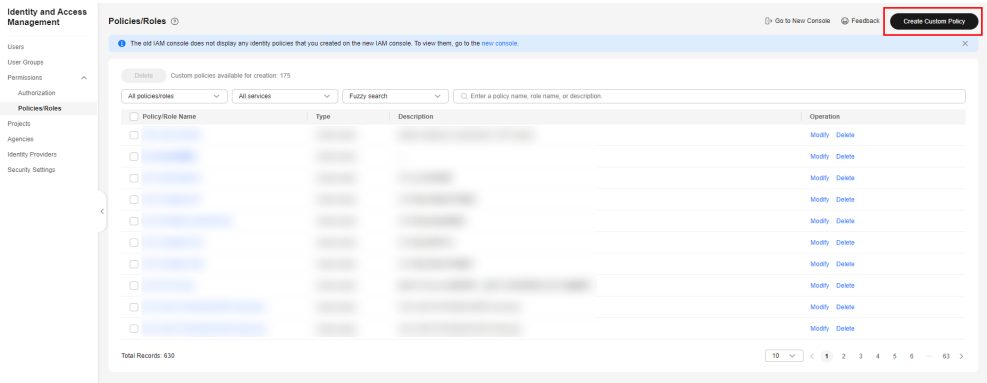
Issue Description

Using enterprise projects to control permissions for COC is required.

Solution

1. Log in to [IAM](#) as an administrator.
2. Choose **Permissions > Policies/Roles** and click **Create Custom Policy**.

Figure 1-5 Creating a custom policy



3.
- Set the policy content, select **CloudOpsCenter**, and select the operations you want to authorize by enterprise project. Click **OK**.

Figure 1-6 Setting the policy content (1)

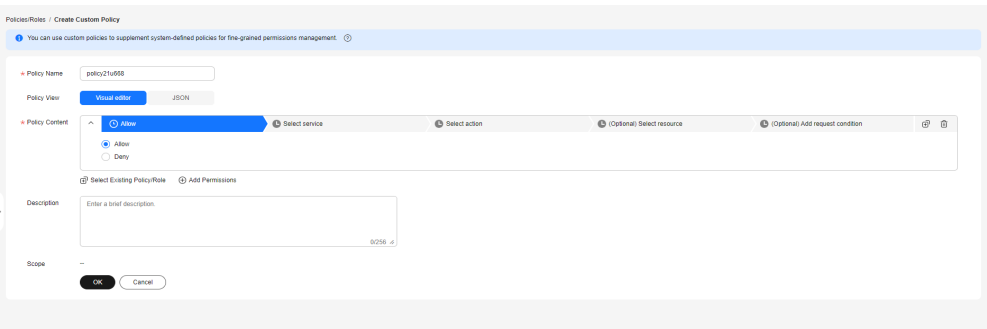


Figure 1-7 Setting the policy content (2)

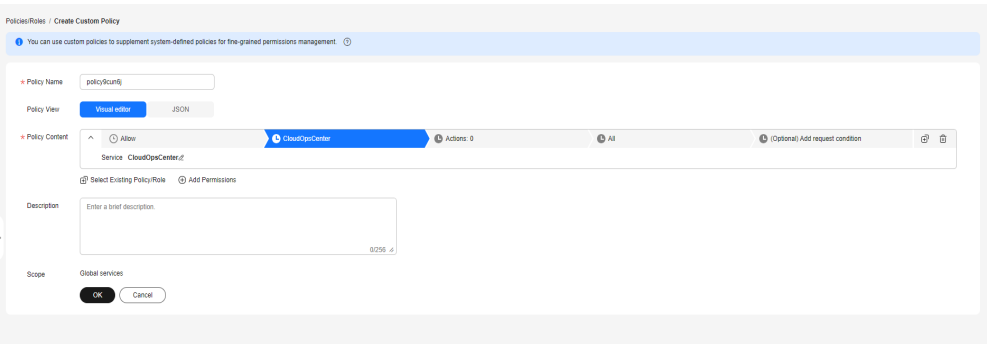
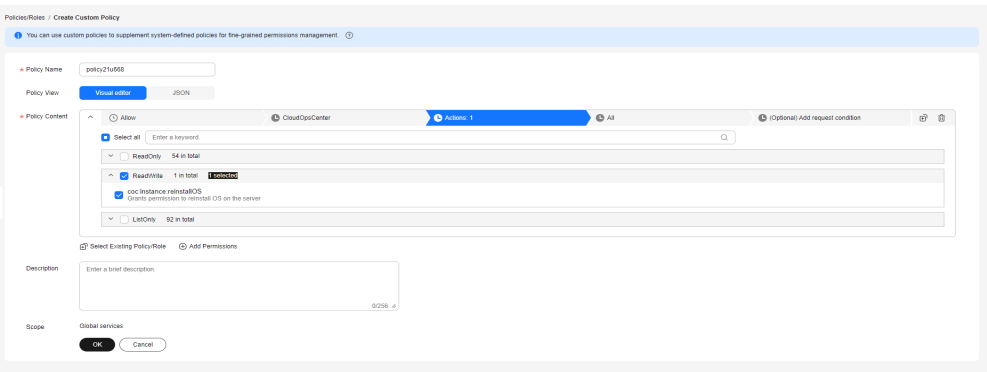


Figure 1-8 Setting the policy content (3)



 **NOTE**

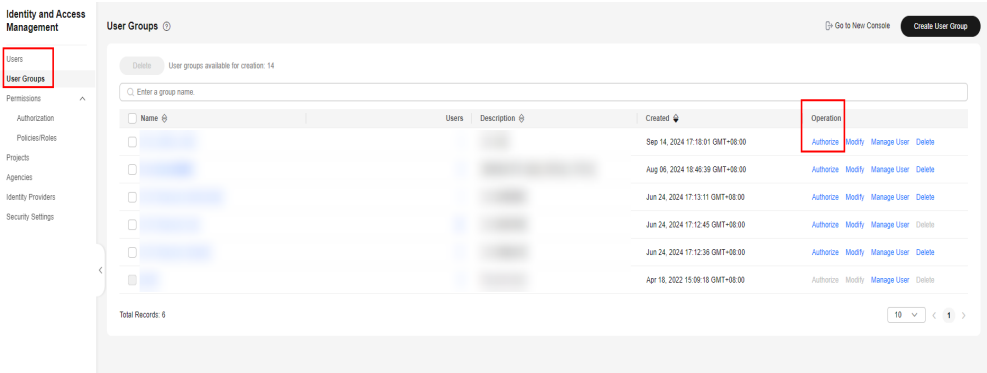
Currently, only some COC operations can be authorized by enterprise project. For details about how to create custom policies, see [Table 1](#).

Table 1-1 Operations that can be authorized by enterprise project

Operation	Description
coc:instance:reinstallOS	Grants permission to reinstall the ECS OS.
coc:instance:changeOS	Grants permission to change the ECS OS.
coc:instance:start	Grants permission to start an ECS.
coc:instance:reboot	Grants permission to restart an ECS.
coc:instance:stop	Grants permission to stop an ECS.
coc:instance:startRDSInstance	Grants permission to enable an RDS DB instance.
coc:instance:stopRDSInstance	Grants permission to stop an RDS DB instance.
coc:instance:restartRDSInstance	Grants permission to reboot an RDS DB instance.
coc:instance:scanOSCompliance	Grants the permission to scan server OS patches.
coc:instance:installPatches	Grants permission to install patches for an ECS.
coc:instance:executeDocument	Grants permission to execute documents on an ECS.
coc:schedule:create	Grants permission to create a scheduled task list.
coc:schedule:update	Grants permission to update a scheduled task.

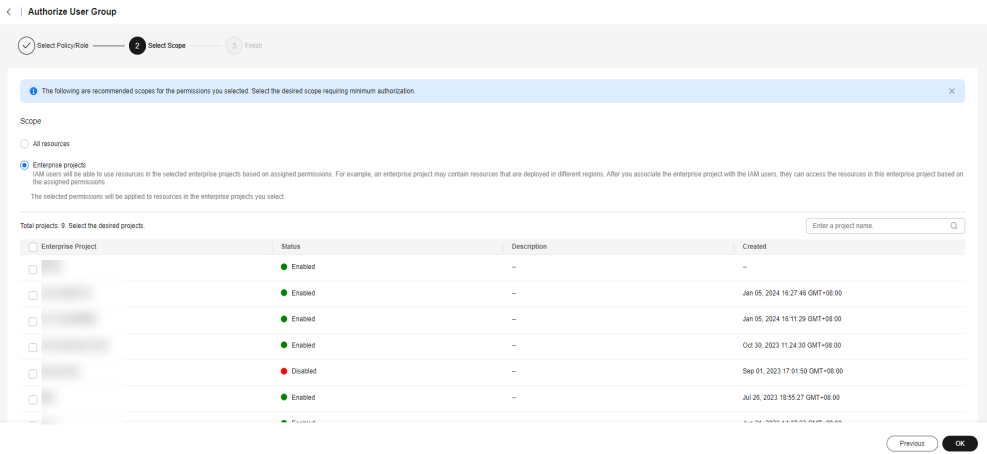
4. Select a user or user group for authorization as the administrator.

Figure 1-9 Selecting an object for authorization.



5. Select the custom policy created in step 3. When setting the minimum authorization scope, specify enterprise project resources.

Figure 1-10 Granting permissions by enterprise project



1.3 How Can I Create an Agency?

Background

If your enterprise has multiple tenant accounts, you can use the cross-account capability of COC to configure and deliver O&M tasks across accounts and regions in scenarios such as creating Cloud Eye alarm rules and executing jobs. During this process, you need to create and use the corresponding agency.

This section uses the scenario of creating Cloud Eye alarm rules across accounts as an example to describe how to create an agency.

Choose **Overview > Quick Configuration Center > Cloud Service Configuration** to use the cross-account configuration function, as shown in the following figure.

Figure 1-11 Configuring the cross-account function**Execution Account & Region**

★ Execution Type

Single

A rule can only be executed by the curr...

Cross Account ?

You can select multiple organization me...

Agency Function Description

The cross-account function requires two agencies: organization administrator agency and execution account agency.

- Organization administrator agency: The organization administrator or COC delegated administrator (administrator for short) creates or trusts the COC service. The agency is used to support the COC service to switch to the administrator identity, that is, you can perform necessary management tasks as the administrator on COC.
- Execution account agency: The execution account (member tenant in the organization) trusts the COC service and administrator agency. The agency must be a trust agency of IAM 5.0. The agency is used to support the administrator to switch to the execution account identity to create a job service ticket and perform cross-account operations.

Example:

Assume that you are using the COC service of Huawei Cloud to perform cross-account operations. You need to create an execution account agency so that you can perform the following operations as an administrator:

- Create an SMN topic in the execution account.
- Add a subscription to the SMN topic to receive alarm notifications.
- Create a Cloud Eye alarm rule in the execution account to monitor and trigger alarms.

In addition, the agency should allow the administrator to switch to the execution account to create a job service ticket. This ensures that the administrator can perform necessary management tasks in the execution account while maintaining flexibility and security.

Executing an Account Authorization Policy

```
{
  "Version": "5.0",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ces:alarms:create"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "smn:topic:create",
        "smn:topic:subscribe"
      ]
    }
  ]
}
```

```
{
  "Effect": "Allow",
  "Action": [
    "sts:agencies:assume",
    "sts::setSourceIdentity",
    "sts::tagSession"
  ]
},
{
  "Effect": "Allow",
  "Action": [
    "coc:instance:executeDocument",
    "coc:job:action",
    "coc:job:get",
    "coc:job:list"
  ]
},
{
  "Effect": "Allow",
  "Action": [
    "iam:projects:list"
  ]
}
}]}
```

NOTE

- The operations supported by an agency represent the authorization of the agency.
- Currently, the Cloud Eye alarm rule contains the alarm notification function. The notification function is restricted by the open capabilities of peripheral services. Therefore, only the topic subscription notification mode is available.
- When configuring Cloud Eye alarm rule parameters, the selection of a notification topic is based on the notification topic owned by the administrator tenant (organization administrator or organization administrator delegated by COC) who logs in to COC. The service does not support querying the notification topics of the execution account (target cross-account tenant) in advance. Therefore, when you perform quick service ticket configuration, the service creates a topic with the same subscription method (WeLink, Lark, and DingTalk are not automatically added to the subscription of the new topic) and the same name as the notification topic selected by the administrator tenant selected when you configure Cloud Eye alarm rule parameters. Notifications may be charged based on the number of notifications. Therefore, you are advised to create a Cloud Eye alarm rule across accounts. If you do not need to send notifications, disable the notification function when configuring the Cloud Eye alarm rule.

Precautions

An organization administrator can create an agency for an account only if the two tenants (the delegator tenant and the delegated tenant) involved in the cross-account operation belong to the same organization.

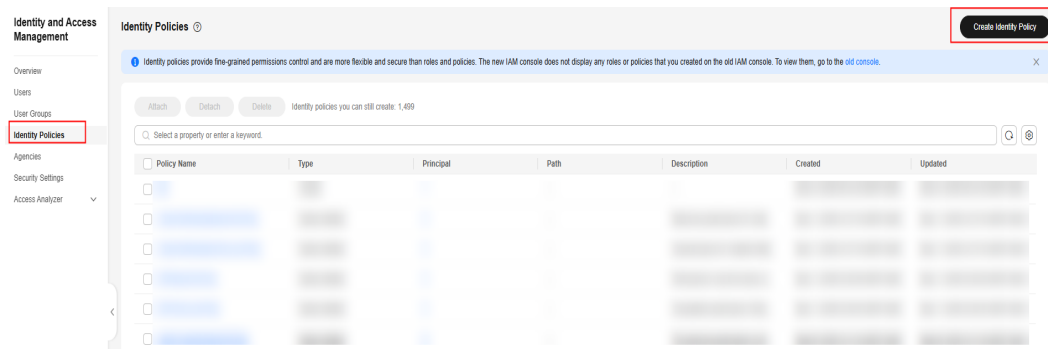
Organization Administrator Agency

Step 1 Log in to **IAM** as an organization administrator tenant.

Step 2 Click **Go to New Console**.

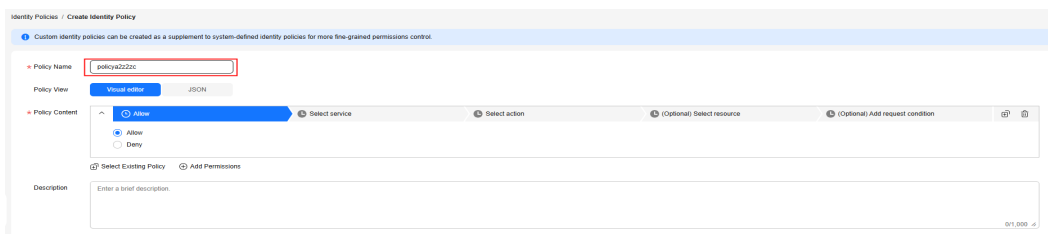
Step 3 In the navigation pane on the left, choose **Identity Policies**. On the displayed page, click **Create Identity Policy** in the upper right corner to create an authorization policy for the agency.

Figure 1-12 Create Authorization Policy



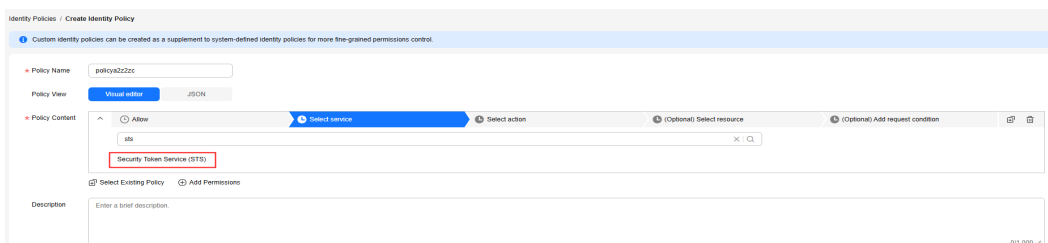
Step 4 On the displayed page, specify **Policy Name**.

Figure 1-13 Custom identity policy name.



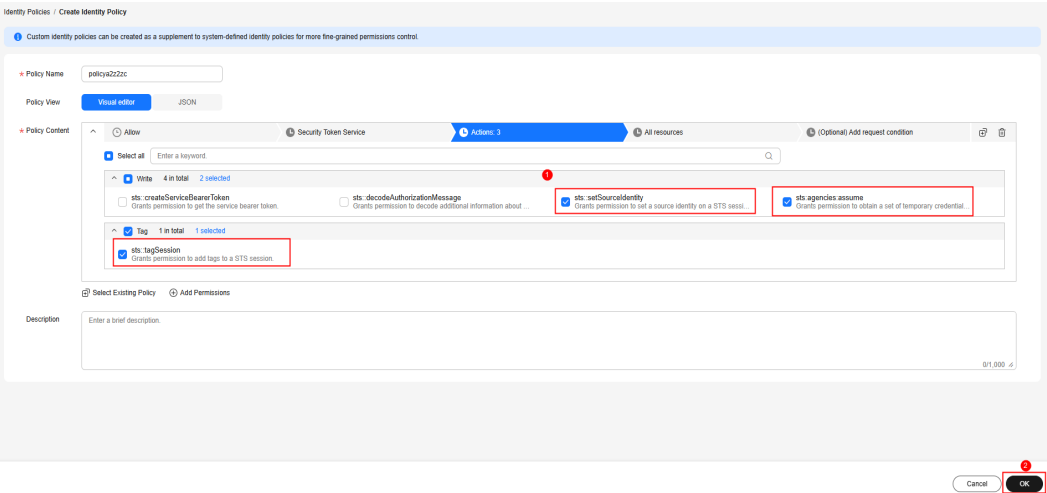
Step 5 In the **Policy Content** area, click **Select service**, search for **Security Token Service (STS)**, and select it. Then, go to **Actions**.

Figure 1-14 Select the cloud service.



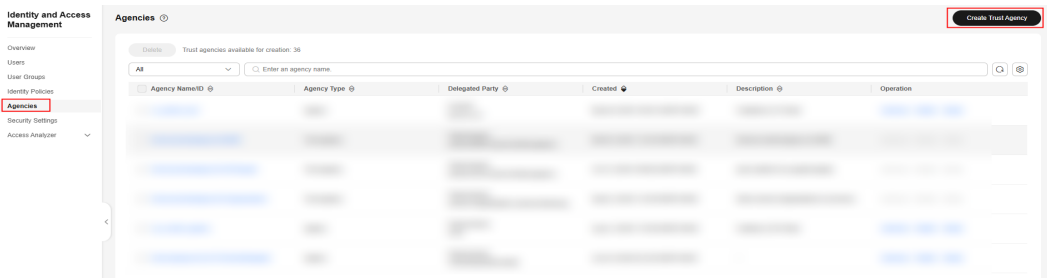
Step 6 Search for the actions **sts:agencies:assume**, **sts::tagSession**, and **sts::setSourceIdentity**, select them, and click **OK**. The policy for identity switching is created successfully.

Figure 1-15 Creating an identity policy for identity switching



- Step 7
- In the navigation pane, choose **Agencies**.
- Step 8
- On the displayed page, click **Create Trust Agency**.

Figure 1-16 Creating a trust agency



- Step 9
- On the displayed page, specify **Agency Name**. Set **Agency Type** to **Cloud service**, **Cloud Service** to **CloudOperationsCenter**, and **Maximum Session Duration** to **12h**.
- Step 10
- Click **OK**. Then click **Authorize** in the displayed dialog box.

Figure 1-17 Completing the trust agency creation

Identity and Access Management

Overview

Users

User Groups

Identity Policies

Agencies

Security Settings

Access Analyzer

Trust agencies allow for more flexible delegation based on trust policies. Trust agencies can only have identity policies attached, and they are only compatible with cloud services that support identity policies.

Agency Name

Enter an agency name.

Agency Type

Account

Delegate permissions to a Huawei Cloud account.

Cloud service

Delegate permissions to a cloud service.

Cloud Service

CloudOperationsCenter

Maximum Session Duration

12h

Trust Policy

```
1 {
2   "Version": "5.0",
3   "Statement": [
4     {
5       "Principal": {
6         "Service": [
7           "service.COC"
8         ]
9       },
10      "Effect": "Allow",
11      "Action": [
12        "sts:agencies:assume"
13      ]
14    }
15  ]
16 }
```

Description

Enter a brief description.

0/1,000

OK

Cancel

Step 11 On the displayed page, search for the policy name created in [Step 4](#) and select it. Click **OK**.

Figure 1-18 Authorizing an agency

Authorize

1 Select Identity Policies

2 Finish

Attach selected identity policies to Agencies

View Selected (1)

Name

Type

Custom

System-defined

System-defined

System-defined

System-defined

Cancel

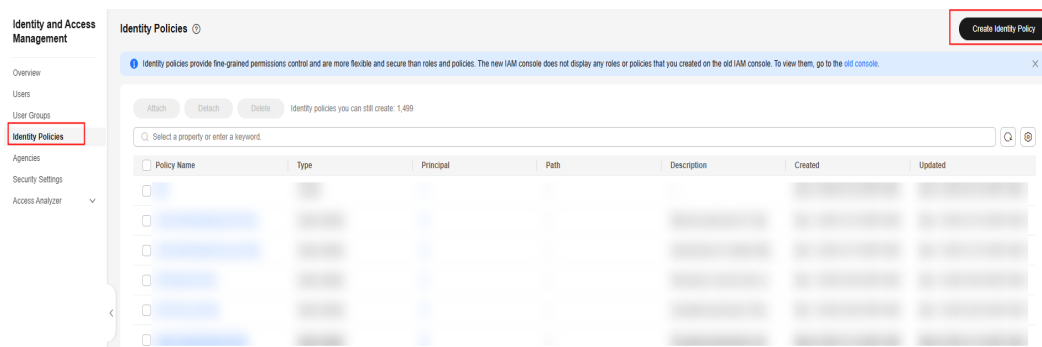
OK

----End

Authorizing an Account

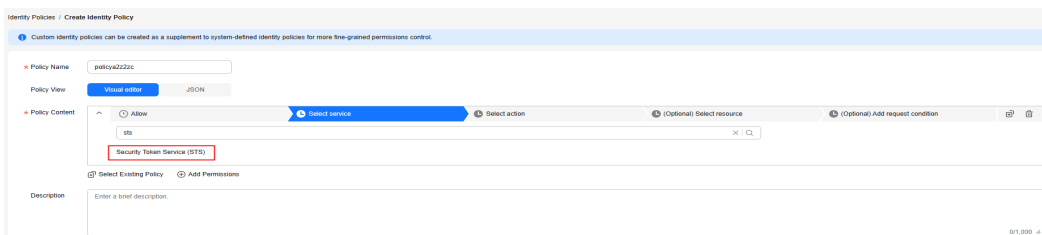
- Step 1** Log in to **IAM** as an organization member tenant (execution account).
- Step 2** Click **Go to New Console**.
- Step 3** In the navigation pane, choose **Identity Policies**.
- Step 4** On the displayed page, click **Create Identity Policy**.

Figure 1-19 Create Authorization Policy



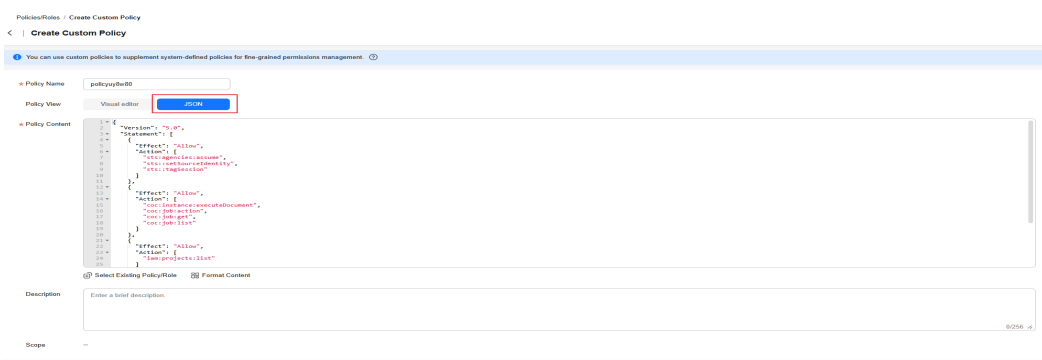
- Step 5** On the displayed page, specify **Policy Name**.

Figure 1-20 Custom identity policy name.



- Step 6** Set **Policy View** to **JSON** and enter the following policy content in the text box.

Figure 1-21 Custom policy content

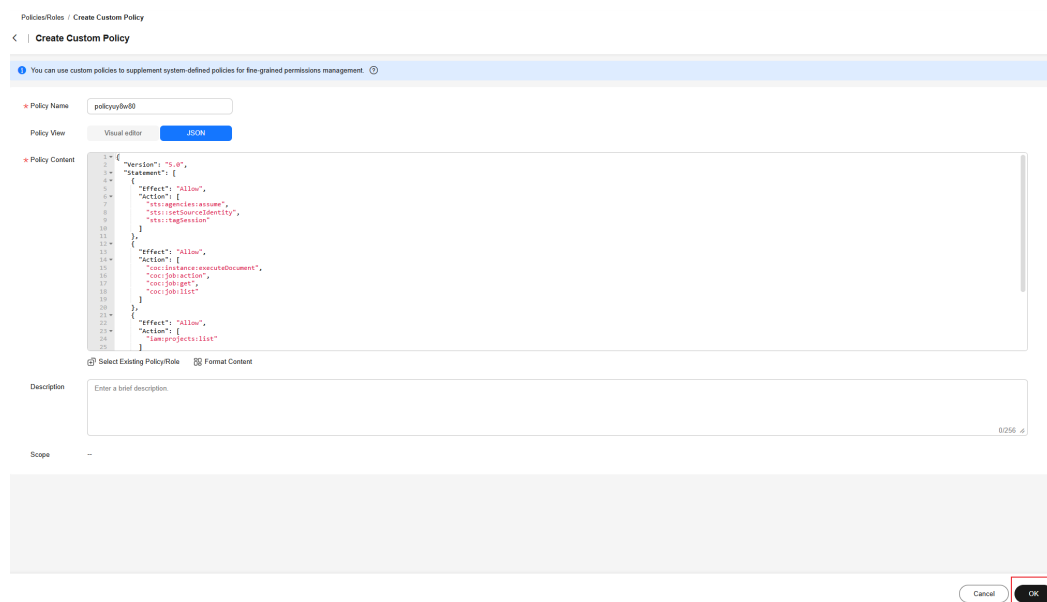


```
{
  "Version": "5.0",
  "Statement": [
    {
      "Effect": "Allow",
```

```
"Action": [
  "sts:agencies:assume",
  "sts::setSourceIdentity",
  "sts::tagSession"
],
{
  "Effect": "Allow",
  "Action": [
    "coc:instance:executeDocument",
    "coc:job:action",
    "coc:job:get",
    "coc:job:list"
  ]
},
{
  "Effect": "Allow",
  "Action": [
    "iam:projects:list"
  ]
},
{
  "Effect": "Allow",
  "Action": [
    "smn:topic:create",
    "smn:topic:subscribe"
  ]
}
]
```

Step 7 Click **OK**. The identity policy of the execution account is created.

Figure 1-22 Creating an Identity Policy for a Tenant Agency



Step 8 In the navigation pane of the new IAM console, choose **Agencies**.

Step 9 Click **Create Trust Agency** in the upper right corner of the page to create an agency that grants the organization administrator permissions to the execution account.

- Specify **Agency Name**.

- Set **Agency Type** to **Cloud service**.
- Set **Cloud Service** to **CloudOperationsCenter**.
- Set **Maximum Session Duration** to **12h**.

Step 10 Click **OK**. The agency is created.

Figure 1-23 Creating an Agency

Identity and Access Management

Trust agencies allow for more flexible delegation based on trust policies. Trust agencies can only have identity policies attached, and they are only compatible with cloud services that support identity policies.

* Agency Name: Enter an agency name.

* Agency Type:

- ☐ Account: Delegate permissions to a Huawei Cloud account.
- ☒ Cloud service: Delegate permissions to a cloud service.

* Cloud Service: CloudOperationsCenter

* Maximum Session Duration: 12h

Trust Policy

```

1 {
2   "Version": "5.0",
3   "Statement": [
4     {
5       "Principal": {
6         "Service": [
7           "service.COC"
8         ]
9       },
10      "Effect": "Allow",
11      "Action": [
12        "sts:agencies:assume"
13      ]
14    }
15  ]
16 }
  
```

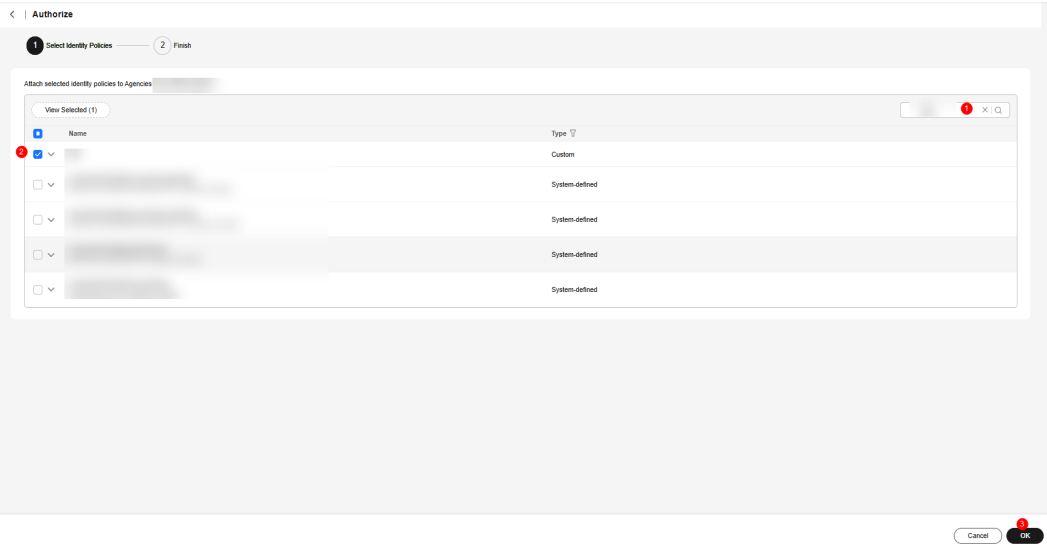
Description: Enter a brief description. (0/1,000)

OK Cancel

Step 11 Click **Authorize** in the dialog box.

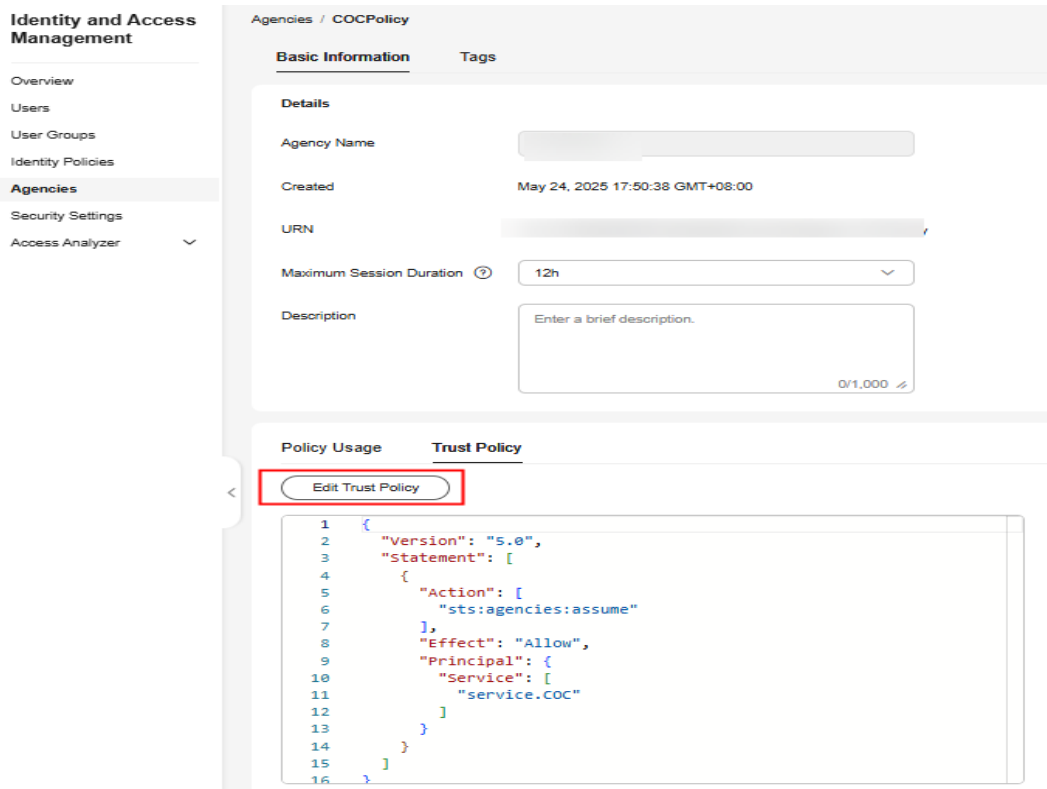
Step 12 On the displayed page, search for the policy name created in [Step 5](#) and select it. Click **OK**.

Figure 1-24 Authorizing an agency



Step 13 On the **Agencies** page, locate the agency created in [Step 9](#) to [Step 12](#) and click **Edit** in the **Operation** column.

Figure 1-25 Editing the trust policy 1 for executing the tenant agency

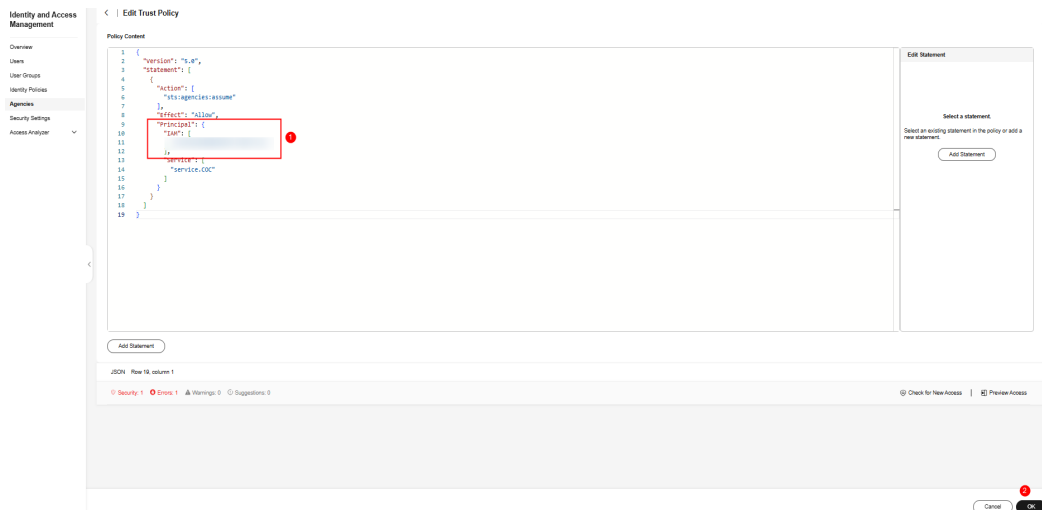


Step 14 On the **Trust Policy** tab page, click **Edit Trust Policy** and add the following JSON data to **Principal**:

```
"IAM": [  
  "${Tenant ID of the target organization administrator tenant}"  
],
```

Step 15 Click **OK**. The trust policy is edited. Click **OK**. The execution account trusts the COC and the organization administrator agency. The account trust policy is created.

Figure 1-26 Editing the trust policy 2 for executing the tenant agency



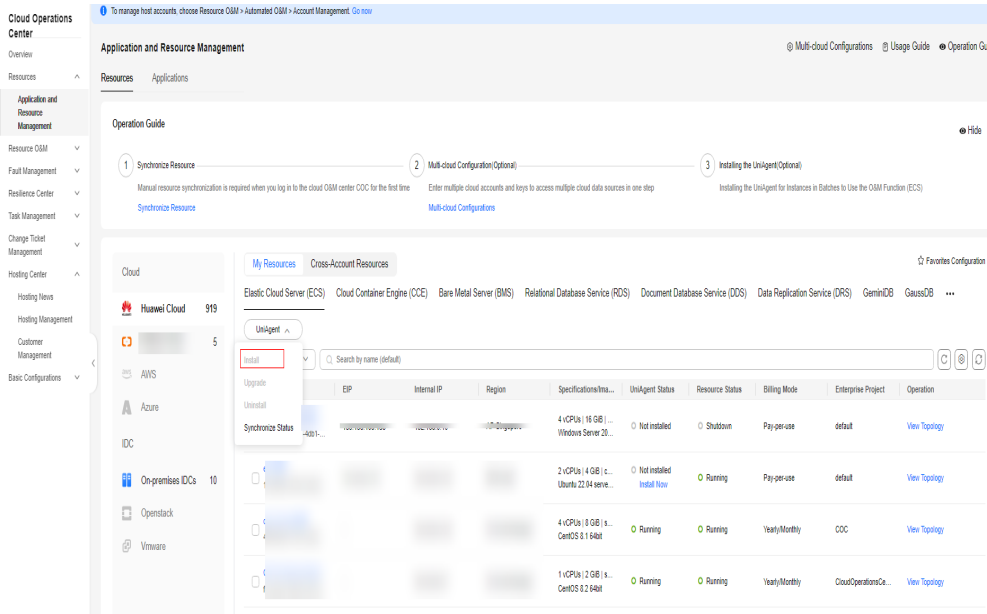
----End

2 Resource Management FAQs

2.1 How Do I Install UniAgent for the First Time?

- Step 1 Log in to [COC](#).
- Step 2 In the navigation pane, choose **Application and Resource Management**. On the **Resources** page, select a host where no UniAgents have not been installed.

Figure 2-1 Installing a UniAgent



- Step 3 On the UniAgent installation page that is displayed, click **Manual installation**.

Install UniAgent

Basic Information

UniAgent Version

1.1.0

Host Access Mode

Direct access (intranet)

Huawei Cloud hosts use the direct connection (private network) method.

Direct access (public network)

Non-Huawei cloud hosts are directly connected to the public network.

Proxy access

Select a proxy area where proxy has been configured and manually install UniAgent on a host using the proxy.

Installation Host

Host

ecs-cnc-omb-linu-key2/192.168.1.100

Bulk Actions

Modify Login Account

root

OK

Change Login Port

22

OK

Change Password

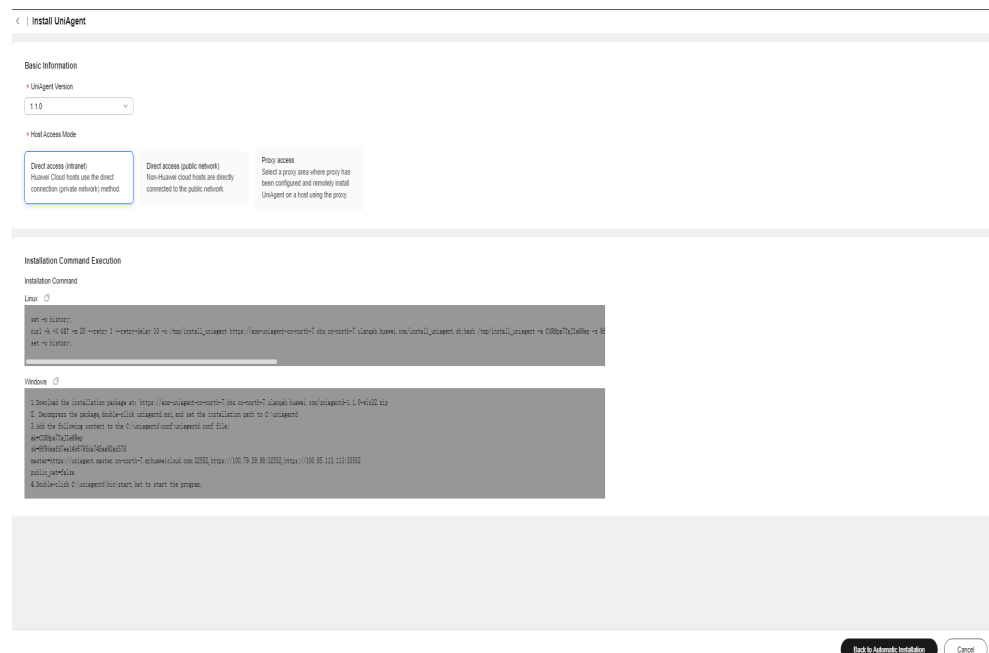
OK

Test Connection

Hosts About to Accommodate UniAgent

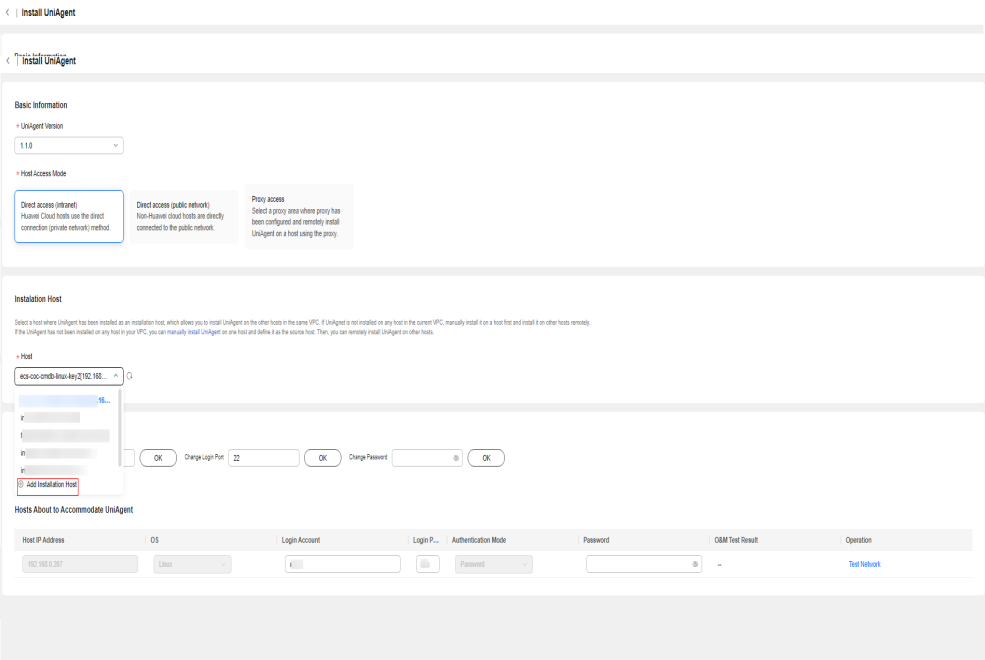
Host IP Address	OS	Login Account	Login P...	Authentication Mode	Password	OBM Test Result	Operation
192.168.1.100	Linux	root	22	Password		-	Test Network

Figure 2-3 Manually installing a UniAgent



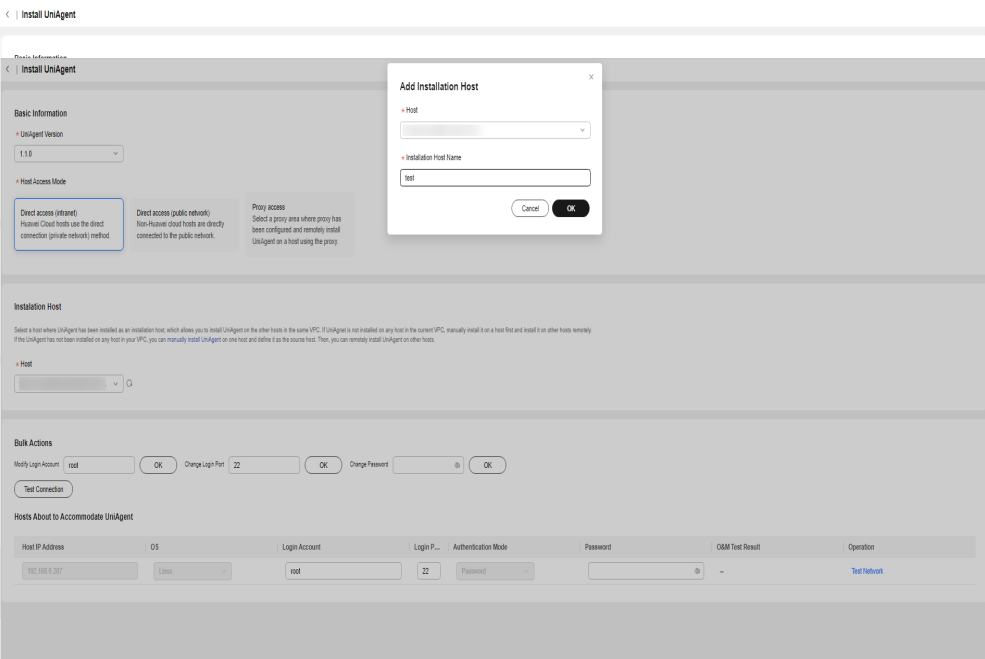
Step 6 Click **Add Installation Host** to set the host where the UniAgent is installed as the installation host.

Figure 2-4 Configuring an installation host



Step 7 In the displayed dialog box, enter the information about the installation host and click **OK**.

Figure 2-5 Selecting an installation host



----End

2.2 What Can I Do If Resources Cannot Be Queried on the Resource Management Page?

Issue Description

The resources cannot be queried on the resource management page.

Cause Analysis

Resources are not synchronized to COC. You cannot manage resources on the resource management page.

Solution

Synchronize resources on the resource management page. For details, see [Synchronizing Resources](#).

2.3 How Can I Find the Description About Application Management Layers?

Issue Description

No description about application management layers is displayed on the applications page.

Cause Analysis

The guide is hidden.

Solution

On the **Applications** page, click **Operation Guide** in the upper right corner to unhide the guide.

Figure 2-6 Hiding the guide

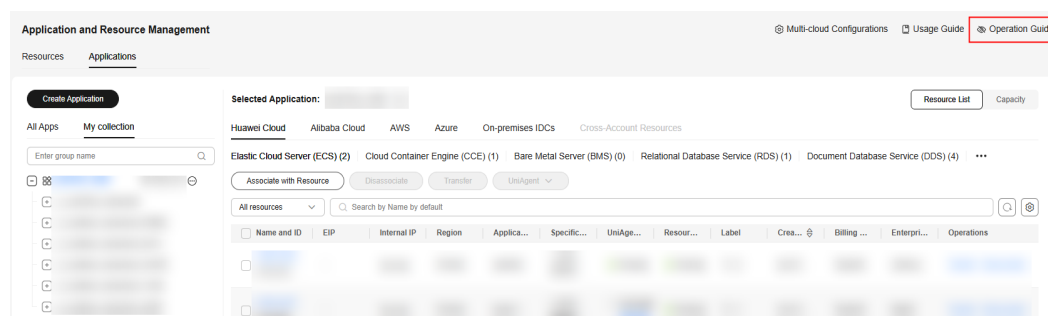
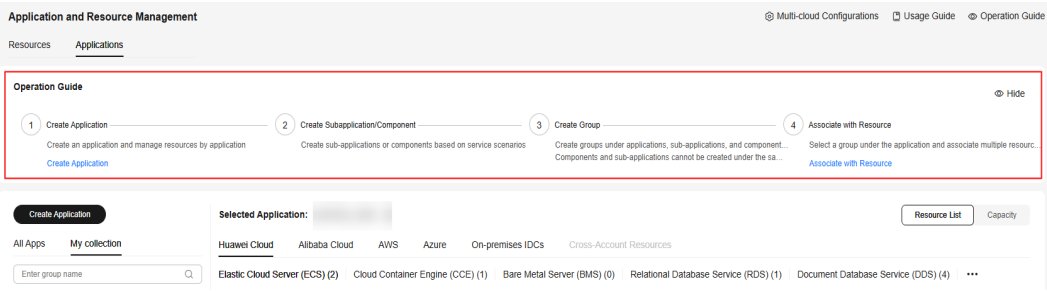


Figure 2-7 Unhiding the guide



3 Patch Management FAQs

3.1 What Can I Do If the Patch Baselines Do Not Take Effect?

Issue Description

The created patch baselines do not take effect.

Cause Analysis

The patch baselines are not set to the default baselines.

Solution

Before using the patch management, scan, or repair feature, ensure that the created patch baselines have been set as the default baselines and the application scenarios are correct.

3.2 What Are the Differences Between the Installation Rule Baselines and Custom Baselines?

Installation rule baselines and custom baselines are two core types of patch baselines in the enterprise patch management system. They are different in patch filtering logics and repair target versions in various patch management requirements. The main differences are as follows:

Core Positioning and Design Objectives

- Installation rule baselines are positioned as standard and automatic general patch management solutions. They are designed to help users quickly filter patches that meet general security or function requirements. Users do not need to enter patch names and version details. Installation rule baselines are mainly used for efficient and unified patch upgrade scenarios.

- Custom baselines are positioned as refined, dedicated, and scenario-specific patch management solutions. They are designed to help users precisely manage specific patch versions. Users can customize the patch scope and target versions. Custom baselines are mainly used for strictly-controlled version compatibility scenarios.

Patch Filtering Capabilities

- Installation rule baselines help users filter batch patches based on basic patch information. Users do not need to specify patches manually. Instead, they specify the patch scope based on preset rules. Users can filter batch patches by the patch type (such as security, function, and bug fixing patches), patch level (such as high, medium, and low risks), release time (such as patches released in the last 30 days), and applicable OS version (such as CentOS 7 and Ubuntu 20.04). For example, users can filter high-risk Linux security patches released in the last 15 days. The system automatically matches all patches that meet the rule.
- Custom baselines help users filter specific patches based on patch names and versions. Users need to specify patch IDs and versions. Users can filter specific patches by entering the patch package names (such as, **kernel-devel** and **openssl**) and versions (such as, **kernel-devel-3.10.0-1160.el7.x86_64** and **openssl-1.0.2k-25.el7_9.x86_64**). For example, users can only filter the patch **openssl** with the version **1.0.2k-25.el7_9**.

Patch Repair Logic

- Installation rule baselines repair the patches with latest versions first, aiming to update the system patches to the latest available compliant versions. When a non-compliant patch is detected on a host (that is, the patch that meets the filtering rule is not installed or the installed patch version is earlier than the latest version), the system automatically obtains the latest patch version from the patch library and upgrades the non-compliant patch to the latest version. For example, if the rule is to filter high-risk security patches and the current version of **openssl** on a host is **1.0.2k-20.el7**, and the latest high-risk patch version is **1.0.2k-25.el7_9**, the system automatically upgrades the patch to **1.0.2k-25.el7_9**.
- Custom baselines repair the patches with the specified version first and upgrade the patches based on the user-defined version. When a non-compliant patch is detected on a host (that is, the patch of the specified version is not installed or the installed version is different from the specified version), the system does not automatically select the latest version. Instead, the system accurately matches the target version specified in the custom baseline and upgrades or downgrades the non-compliant patch to the specified version. For example, if a user sets the version of **openssl** to **1.0.2k-20.el7** in the custom baseline, the system only repairs the **openssl** patch of the host to **1.0.2k-20.el7** even if there is the latest version **1.0.2k-25.el7_9** in the patch library. This ensures that user requirements are met.

3.3 What Can I Do If Exception "all mirrors were tried" Is Recorded in the Patch Service Ticket Log?

Issue Description

The error message "all mirrors were tried" is recorded in the patch service ticket log.

Cause Analysis

Generally, the error message is reported when network faults occur.

Solution

Check whether the network connectivity between the node and patch source configured on the node is normal or whether the network of the node is normal.

3.4 Why Can't I Select a Node?

Issue Description

When you create a patch scan task and add an instance, the corresponding resource cannot be selected.

Cause Analysis

The resource status is abnormal, or the UniAgent is not installed.

Solution

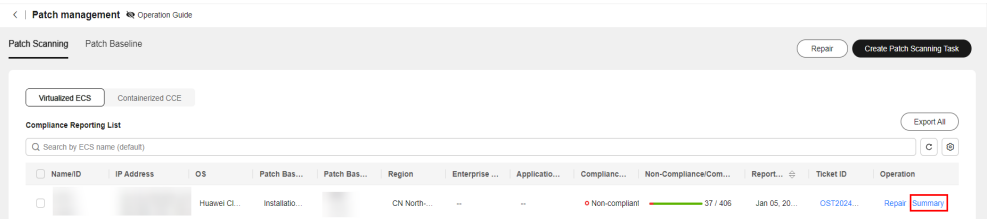
Check whether the node is in the normal state: the resource status is **Running** and the UniAgent status is **Running**.

For details about how to install the UniAgent, see [Performing Operations on a UniAgent](#).

3.5 What Can I Do if a Patch Remains Non-Compliant After Repaired?

Step 1 Click the button for viewing the summary of the compliance report that reports non-compliance.

Figure 3-1 Viewing the compliance report summary



Step 2 View the status of the non-compliant patch and view different solutions based on the compliance status.

Table 3-1 Solutions for different compliance statuses

Non-compliance Status	Solution
Failed	View the log of the patch service ticket that generates the compliance report and rectify the fault based on the failure log.
Installed-to be restarted	A newly installed patch can only take effect after the host is restarted. Therefore, you need to restart the host.
Rejected	If a patch is rejected in the patch baseline, the compliance report shows that the patch is rejected. To cancel the rejection, edit the corresponding baseline in the patch baseline.

-----End

3.6 What Can I Do If "lsb_release not found" Occurs During Patch Operations?

Issue Description

The error message "lsb_release not found" is displayed during patch operations.

Solution

Check whether the **lsb_release** command package exists on the ECS instance.

- If no, install the command package.
- If the ECS instance contains the **lsb_release** command package, check whether the UniAgent version is later than 1.1.0. If yes, downgrade the UniAgent version to a version earlier than 1.1.0 and try again.

4 Automation FAQs

4.1 Why Can't the Reviewer Receive Notifications?

Issue Description

The reviewer cannot receive the notification.

Cause Analysis

No notification channel is configured for the reviewer on the **O&M Engineer Management** page.

Solution

For details about how to configure the message channel, see [O&M Engineer Management Usage](#).

4.2 Why Is the Input Value of a Customized Script Parameter Invalid?

Issue Description

The input value of the custom script parameter is invalid.

Solution

Check whether the custom script parameter meets the following requirements:

- The parameter value contains 1 to 1024 characters.
- The value can contain letters, digits, spaces, and special characters (`_/.*?:",=+@\[\{\}`).
- Consecutive periods (.) are not allowed.

4.3 Why Can't I Select an Instance?

Issue Description

Instances cannot be selected during automated O&M.

Solution

A UniAgent must be installed for instances to perform automated O&M.

For details about how to install the UniAgent, see [Performing Operations on a UniAgent](#).

4.4 How Do I Reset the Password Without Restarting a DB Instance?

In cloud resource (such as ECSs and BMSs) management, if you need to change the password of the host account and prevent the instance from being restarted (to avoid service interruption), you can use the public scripts provided by COC to change the password without restarting the instance. Currently, ECSs and BMSs are supported. The operation procedure is as follows:

CAUTION

When running a public script on COC, you need to select an instance. The prerequisites for selecting an instance are as follows:

The resource instance information has been synchronized to COC. For details, see [Synchronizing Resources](#).

The UniAgent has been installed for the instance and is running properly.

To install the UniAgent on an instance, you need to provide the administrator account and password of the instance. If UniAgent is not installed on the resource instance and you forget the password, UniAgent cannot be installed and the public script for changing the password cannot be executed.

A Linux ECS is used as an example to describe how to change the password of user **root**.

Step 1 Log in to [COC](#).

Step 2 In the navigation pane on the left, choose **Resource O&M > Automated O&M**.

Step 3 In the **Routine O&M** area, click **Script Management**.

Step 4 On the **Public Scripts** tab page, locate the script for changing password of the administrator account and click **Execute** in the **Operation** column.

Step 5 Enter a new password. (The password must meet the password rule, for example, the password must contain 8 to 30 characters, including uppercase letters, digits, and special characters.)

Step 6 Add the target ECS.

Step 7 Click **OK**. The system remotely delivers the script to the instance and automatically runs the script. You do not need to log in to the instance terminal during the process.

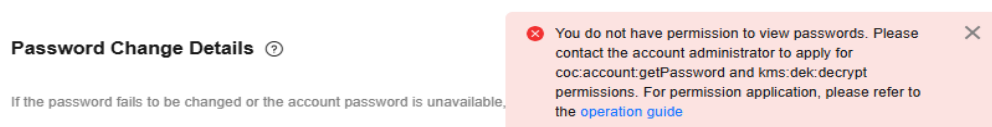
After the script is executed, you can use the new password to log in to the instance using SSH. The instance is in operation during the entire process.

----End

4.5 What Can I Do If I Am Not Authorized to View Passwords on the Account Management Page?

Issue Description

When a user wants to view the password on the account management page, the system displays a message indicating that the user does not have the permission to view the password, as shown in the figure below.



Cause Analysis

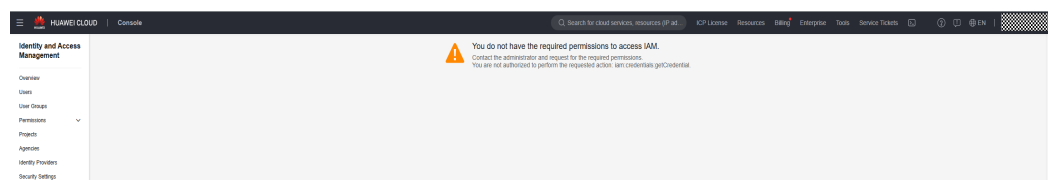
To view a password, you must have the **coc:account:getPassword** and **kms:dek:decrypt** permissions to obtain the encrypted password and decrypt the password. Therefore, ensure that the account used to log in to the console has been granted the required permissions.

Solution

Step 1 Log in to **IAM**.

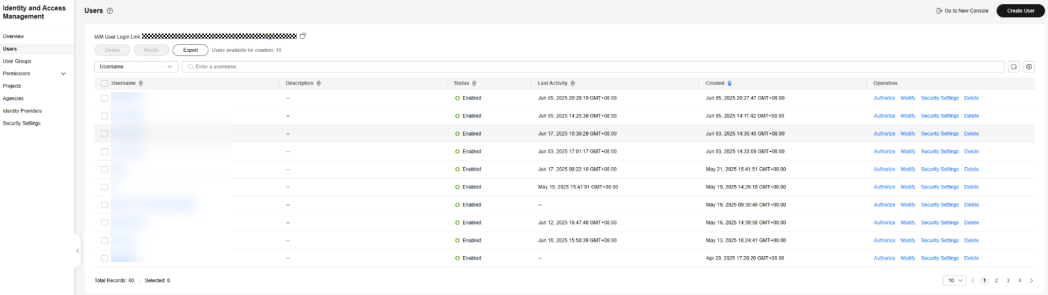
If you do not have the IAM access permission, contact the administrator.

Figure 4-1 No permissions to access IAM.



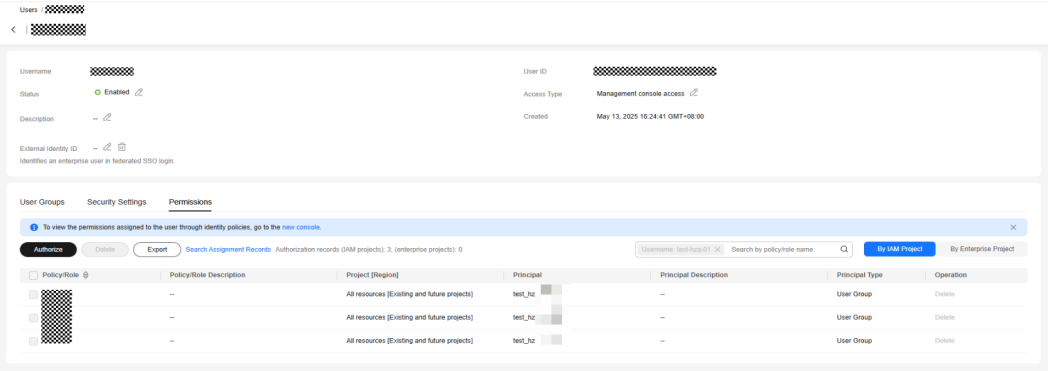
Step 2 In the navigation pane, choose **Users**.

Figure 4-2 Viewing users



- Step 3
- In the displayed page, click the target username. The user details page is displayed.
- Step 4
- Click the **Permissions** tab to view all permissions of the current user.

Figure 4-3 Viewing user authorization records



- Step 5
- Click the permission name to view the details and check whether the **coc:account:getPassword** and **kms:dek:decrypt** permissions exist in the **Action** list.

Figure 4-4 Confirming permissions

Policies/Roles / Modify Custom Policy

< | **Modify Custom Policy**

You can use custom policies to supplement system-defined policies for fine-grained permissions management.

★ Policy Name

Policy View

Visual editor

JSON

★ Policy Content

1 {
2 "Version": "1.1",
3 "Statement": [
4 {
5 "Action": [
6
7
8
9
10
11
12
13
14
15
16],
17 "Effect": "Allow"
18 }
19]
20 }

Select Existing Policy/Role

Format Content

Description

Enter a brief description.

Scope

Global services

Step 6 If the current account does not have the required permissions, contact the account administrator to request **coc:account:getPassword** and **kms:dek:decrypt** permissions for the login user in IAM.

For details, see [Creating a User Group and Assigning Permissions](#).

----End

5 Batch Operation FAQs

5.1 What Should I Do If an Error Is Reported When I Change Images for ECS Resources in Batches?

Issue Description

When changing image for batch ECS resources,
`"code": "Ecs.0021", "message": "Failed to check Cinder quotas because the number of Gigabytes exceeded the upper limit." or CreateRootVolumeTask-fail: call evs api - create volume fail :{"error_msg": "volume gigabytes exceeded volume gigabytes quota!", "common_error_code": "CMM.3141", "error_code": "EVS.1042"} is displayed.`

Solution

If the EVS disk quota is insufficient, apply for a higher EVS disk quota. For details, see [Increasing EVS Resource Quotas](#).

6 FAQs About Parameter Management

6.1 What Are the Permissions Required for Managing Parameters?

Permission Design

- To access the parameter list page, the **coc:parameter:list** permission is required.
- To obtain parameter details, the **coc:parameter:get** permission is required.
- To delete a parameter, the operation permission **coc:parameter:delete** is required.
- To create a parameter, the operation permission **coc:parameter:create** is required.
- To update a parameter, the operation permission **coc:parameter:update** is required.
- Resource permissions: **coc::*:parameter:name** (The first asterisk (*) indicates all region IDs, the second asterisk (*) indicates all tenants, and *name* indicates the parameter name. This permission means that you can access a parameter of the specified tenant in a certain region.)

FAQs

Resource permissions determine the data that you can access. Operation permissions are used to perform operations on your resource permissions.

- If you can access a parameter but cannot access the parameter list page, you do not have the **coc:parameter:list** permission.
- If you cannot find a specified parameter, check whether you have the permission on the parameter.
- **coc:service-name:region:account-id:resource-type:resource-path** is the structure of resource permissions. The asterisk (*) indicates all permissions at this level. To add resource permissions, enter information in this format.

6.2 Can I Reference Parameters in the Parameter Center and Target Instances Across Regions?

Based on safety production requirements, users are not allowed to reference parameters in the **Parameter Center** and target instances across regions.

The following is an example:

On the **Execute Script** page, you click **Parameter Center** in the **Script Input Parameters** area. On the **Parameter Center** page, if you set **Region** to **CN North-Beijing4**, set the same region when adding the target instance. Otherwise, an error will be reported and the script cannot be executed.

Figure 6-1 Parameter center

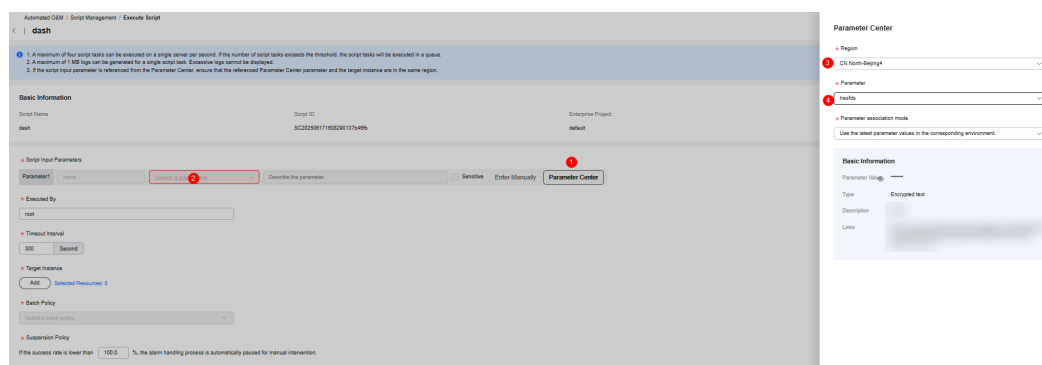
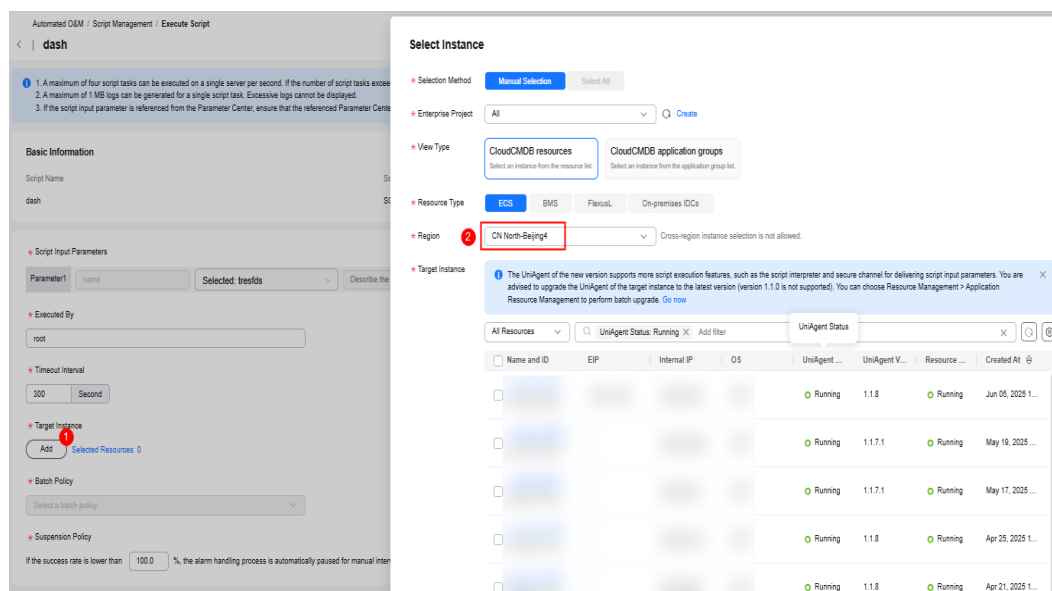


Figure 6-2 Selecting an instance



7 FAQs About Resource O&M

7.1 Resource O&M Permissions and Supported Actions

This section describes fine-grained permissions management for your COC resources. If your Huawei Cloud account does not need individual IAM users, then you may skip over this section.

By default, new IAM users do not have any permissions assigned. You need to add a user to one or more groups and assign policies or roles to these groups. The users then inherit permissions from the groups. The users then inherit permissions from the groups and can perform specified operations on cloud services based on the permissions they have been assigned.

You can grant users permissions by using You can grant users permissions by using **roles** and **policies**. Roles are provided by IAM to define service-based permissions that match users' job responsibilities. Policies are a fine-grained authorization strategy that defines permissions required to perform certain operations on specific cloud resources under certain conditions. This type of authorization is API-based and is ideal for least privilege access.

For details about the COC system policies, see [COC Permissions Management](#).

NOTE

If you want to allow or deny the access to an API, use policy-based authorization.

Each account has all the permissions required to call all APIs, but IAM users must be assigned the required permissions. The permissions required for calling an API are determined by the actions supported by the API. Only users who have been granted permissions can call the API successfully. For example, if an IAM user wants to call an API to query ECSs, the user must be granted the permissions allowing for action **ecs:servers:list**.

Actions

COC provides system-defined policies that can be used in IAM. You can also create custom policies to supplement system-defined policies for more refined access control. Actions supported by policies are specific to APIs. Common concepts related to policies include:

- Permissions: allow or deny operations on specified resources under specific conditions.
- APIs: REST APIs that can be called by a user who has been granted specific permissions.
- Actions: specific operations that are allowed or denied.
- Related actions: actions which a specific action depends on. When allowing an action for a user, you also need to allow any existing action dependencies for that user.
- IAM or enterprise projects: the authorization scope of a custom policy. A custom policy can be applied to IAM projects or enterprise projects or both. Policies that contain actions supported by both IAM and enterprise projects can take effect for user groups of both IAM and Enterprise Management. Policies that contain actions only for IAM projects can be assigned to user groups and be applied only in IAM. They cannot be applied in Enterprise Management.

For details about the differences between IAM and enterprise projects, see [Differences Between IAM Projects and Enterprise Projects](#).

- Authorization by instance or tag: application scope of custom policies. For APIs that support both authorization by instance and authorization by tag, custom policies take effect for both authorized instances and instances with tags defined in the policies. For APIs that only support authorization by tag, custom policies take effect only for instances with specified tags.

Currently, this function is unavailable in the **CN North-Ulanqab1** region.

NOTE

In the table for supported actions, the check mark (✓) indicates that an action can take effect for the corresponding type of projects, and the cross symbol (x) indicates that an action cannot take effect.

COC supports the following actions that can be defined in custom policies:

Table 7-1 Custom policy actions supported by resource O&M

Functions	Action	Description	Dependency	IAM Project	Enterprise Project	Authorization by Instance	Authorization by Tag
Resource Synchronization	coc:instance:listResources	Grants permission to query the resource list.	-	√	x	x	x
	coc:application:listResources	Grants permission to query the application resource list.	-	√	x	x	x
	coc:instance:syncResources	Grants permission to synchronize the resource list.	-	√	x	x	x
Scheduled O&M	coc:schedule:list	Grants permission to query the scheduled task list.	-	√	x	x	x
	coc:schedule:enable	Grants permission to enable scheduled tasks.	-	√	x	x	x
	coc:schedule:update	Grants permission to update scheduled tasks.	-	√	√	x	x
	coc:schedule:disable	Grants permission to disable the scheduled task list.	-	√	x	x	x

Functions	Action	Description	Dependency	IAM Project	Enterprise Project	Authorization by Instance	Authorization by Tag
	coc:schedule:approve	Grants permission to review the scheduled task list.	-	√	x	x	x
	coc:schedule:create	Grants permission to create a scheduled task list.	-	√	√	x	x
	coc:schedule:delete	Grants permission to delete scheduled tasks.	-	√	x	x	x
	coc:schedule:count	Grants permission to query the number of scheduled tasks.	-	√	x	x	x
	coc:schedule:get	Grants permission to query the scheduled task records.	-	√	x	x	x
	coc:schedule:getHistories	Grants permission to query the execution history of a scheduled task.	-	√	x	x	x
In-Depth Diagnosis	coc:application:GetDiagnosisTaskDetails	Grants permission to query application resource diagnosis tasks.	aom:uniagentAgent:install; aom:uniagentAgent:uninstall;	√	x	x	x

Functions	Action	Description	Dependency	IAM Project	Enterprise Project	Authorization by Instance	Authorization by Tag
	coc:application:CreateDiagnosisTask	Grants permission to create application diagnosis tasks.		√	x	x	x
	coc:job:action	Grants permission to operate service tickets.		√	x	x	x
Script/Job Management	coc:document:create	Grants permission to create documents.	aom:uniagentAgent:install; aom:uniagentAgent:list; aom:uniagentInstallHost:list; aom:uniagentProxyRegion:get; iam:agencies:list;	√	x	x	x
	coc:document:listRunbookAtoms	Grants permission to view the atomic capability list of a job.		√	x	x	x
	coc:document:getRunbookAtomicDetails	Grants permission to query details about an atomic capability of a job.		√	x	x	x
	coc:document:list	Grants permission to query the document list.		√	x	x	x
	coc:document:delete	Grants permission to delete documents.		√	x	x	x
	coc:document:update	Grants permission to modify documents.		√	x	x	x

Functions	Action	Description	Dependency	IAM Project	Enterprise Project	Authorization by Instance	Authorization by Tag
	coc:document:get	Grants permission to view documents.		√	x	x	x
	coc:document:analyzeRisk	Grants permission to analyze document risks.		√	x	x	x
	coc:instance:executeDocument	Grants permission to execute documents on an ECS.		√	x	x	x
Batch management of cloud phone servers and cloud phones	coc:instance:autoBatchInstances	Grants permission to enable automatic instance batching.	ecs:serverKeyPairs:list; ecs:servers:get; ecs:cloudServers:list;	√	x	x	x
	coc:instance:executeDocument	Grants permission to execute documents on an ECS.	ecs:cloudServers:rebuild; ecs:cloudServers:changeOS; ecs:cloudServers:showServer;	√	x	x	x
	coc:instance:startRDSInstance	Grants permission to enable RDS DB instances.	ecs:cloudServers:stop; ecs:cloudServers:reboot;	√	√	x	x
	coc:instance:stopRDSInstance	Grants permission to stop an RDS DB instance.	ecs:cloudServers:start; ims:images:get;	√	√	x	x
	coc:instance:restartRDSInstance	Grants permission to reboot an RDS DB instance.	ims:images:list; bss:order:view; billing:contract:viewDiscount;	√	√	x	x

Functions	Action	Description	Dependency	IAM Project	Enterprise Project	Authorization by Instance	Authorization by Tag
	coc:instance:start	Grants permission to start ECSs.		√	√	x	x
	coc:instance:reboot	Grants permission to restart ECSs.		√	√	x	x
	coc:instance:stop	Grants permission to disable ECSs.		√	√	x	x
	coc:instance:reinstallOS	Grants permission to reinstall ECS OSs.		√	√	x	x
	coc:instance:changeOS	Grants permission to change the OS of an ECS.		√	√	x	x

8 FAQs About Fault Management

8.1 What Is the Process of Generating an Incident?

There are three methods available: manual incident creation, converting alarms to incidents, or automatically generate an incident based on an incident forwarding rule. The detailed processes of the three operation methods are as follows.

Manually Creating an Incident

Choose **Fault Management > Incidents** and click **Create** to create an incident ticket. For details, see [Creating an Incident](#).

Converting an Alarm to an Incident

Choose **Fault Management > Incidents** to create an incident ticket. For details, see section "Converting an Alarm to an Incident".

Automatically Generating Incidents Based on Forwarding Rules

To automatically generate an incident based on a forwarding rule, perform the following operations:

- Step 1** Log in to [COC](#).
- Step 2** Synchronize personnel. For details, see [O&M Engineer Management Overview](#).
- Step 3** Set shift scheduling and add agents to the shift scheduling. For details, see [Overview](#).
- Step 4** Integrate with the monitoring system to automatically report alarms. For details, see [Monitoring System Integration Management](#).
- Step 5** Configure transition rules and generate incidents based on the rules. For details, see [Forwarding rules](#).
- Step 6** To receive incident notifications after an incident is generated, configure the automated notification feature. For details, see [Notification Management](#).

----End

8.2 How Can I Receive an Incident Ticket Notification?

Step 1 Log in to [COC](#).

Step 2 Subscribe to message notifications on the Personnel Management page. For details, see [O&M Engineer Management Overview](#).

Step 3 Configure notification rules on the Notification Management page. For details, see [Notification Management](#).

----End

8.3 What Is a War Room?

When a group or major fault occurs, a war room is set up to provide guidance for quick service recovery. It supports joint operations of O&M engineers, R&D team, and operations personnel for fault handling. You can add fault recovery members to the war room, send the fault progress to the personnel who are concerned about the fault in a timely manner, and use the application diagnosis and response plan to help quickly recover applications.

To set up a war room, you need to connect DingTalk, Lark, or WeCom by referring to [Mobile App Management](#).

You can set up a war room request for an incident that has been accepted. For details, see [Starting a War Room](#).

For details about how to use a war room, see [War Rooms](#).

8.4 What Can I Do If a Fault Diagnosis Task Is Abnormal?

Issue Description

When users want to diagnose ECSs on the fault diagnosis page, an exception occurs during the execution.

Solution

The root causes of issues may vary in different execution phases. This section provides corresponding solutions for different scenarios.

An Exception Occurs During Plugin Installation

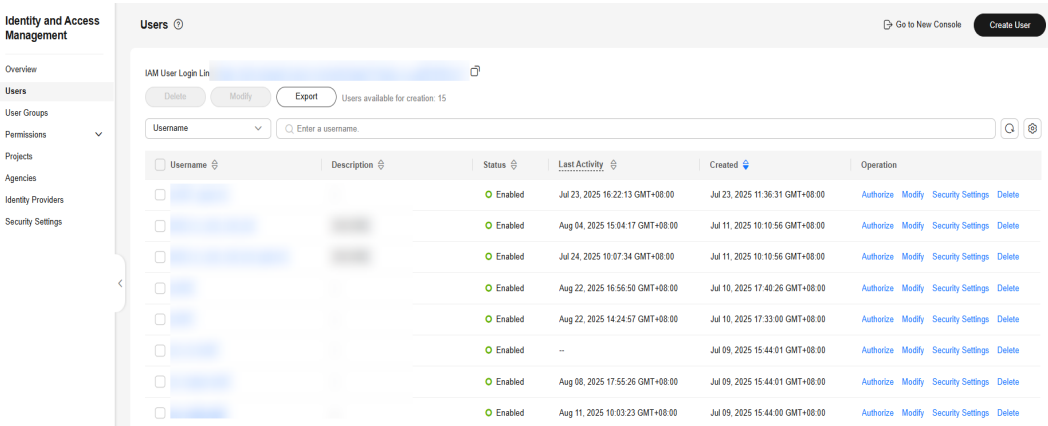
Generally, an error is reported during plugin installation because the user does not have the IAM permission to use the UniAgent to install the plugin. For details, see [Managing ICAgent Plug-ins for Hosts](#).

Step 1 Log in to [IAM](#).

If you do not have the IAM access permission, contact the administrator.

Step 2 In the navigation pane, choose **Users**.

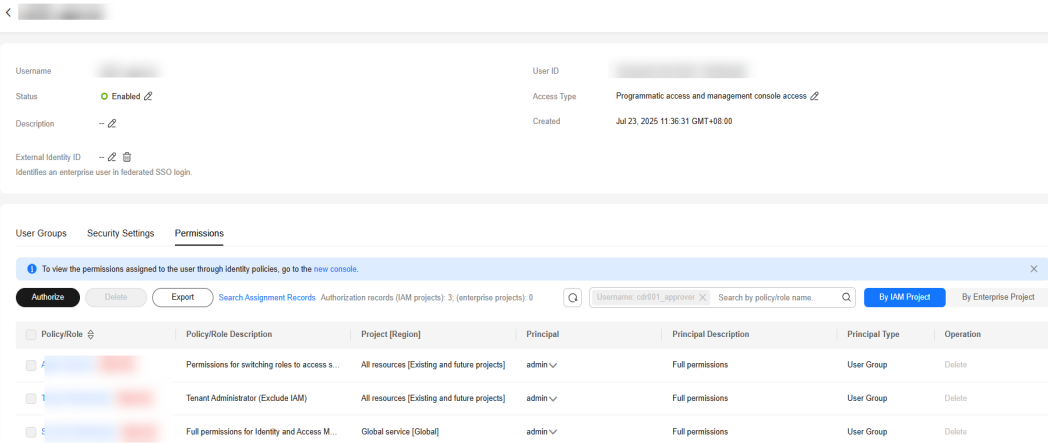
Figure 8-1 Viewing users



Step 3 In the displayed page, click the target username. The user details page is displayed.

Step 4 Click the **Permissions** tab to view all permissions of the current user.

Figure 8-2 Viewing user authorization records



Step 5 Click the permission name to view details. In the policy content, check whether the action list contains the **aom:uniagentAgent:install** permission.

Figure 8-3 Viewing permissions

```
1 {  
2   "Version": "1.1",  
3   "Statement": [  
4     {  
5       "Action": [  
6         "coc:*:*",  
7         "*:*:list*",  
8         "*:*:get*",  
9         "aom:uniagentAgent:install",  
10        "aom:uniagentAgent:uninstall",  
11        "aom:uniagentAgent:list",  
12        "aom:uniagentAgent:get",  
13        "aom:uniagentInstallHost:list",  
14        "aom:uniagentProxyRegion:get",  
15        "aom:view:get",  
16        "iam:agencies:createAgency",  
17        "iam:permissions:grantRoleToAgencyOnDomain",  
18        "iam:roles:listRoles",  
19        "iam:agencies:list*",  
20        "iam:agencies:createServiceLinkedAgency",  
21        "iam:agencies:deleteServiceLinkedAgency",  
22        "ecs:*:*",  
23        "evs:*:get",  
24        "evs:*:list",  
25        "evs:volumes:create",
```

Step 6 If the current account does not have the required permissions, contact the account administrator to request **aom:uniagentAgent:install** permissions for the login user in IAM.

For details, see [Creating a User Group and Assigning Permissions](#).

----End

An Exception Occurs During Data Collection

Generally, an error occurs during data collection because an error is reported when a user executes the collection script. The root cause of the script execution error is as follows:

- The OS image version does not meet the requirements.
- The UniAgent version does not meet the requirements.

For details, see [ECS Diagnosis](#).

Step 1 Log in to [COC](#).

Step 2 In the navigation pane, choose **Task Management > Execution Records**.

Step 3 Click the **Script Tickets** tab.

Step 4 Search for **HWC.COC.PLATFORM-execute-linux-holmes-agent.sh** by the ticket name.

Step 5 Click the ticket record that meets the execution time range of the data collection step to view details.

- If an error message **/usr/local/uniagentd/tmp/** is displayed, check whether the UniAgent version meets the requirements.

- If the collected data is displayed in JSON format but the execution still fails, check whether the OS image version meets the requirements.

----End

An Exception Occurs During Fault Diagnosis

Generally, there is a low probability that an error is reported due to network congestion.

Step 1 Log in to [COC](#).

Step 2 In the navigation pane, choose **Task Management** > **Execution Records**.

Step 3 Click **Diagnose Tickets**.

Step 4 Click the name of the abnormal service ticket to go to the diagnosis details page.

Step 5 Click **Retry**.

----End

An Exception Occurs During Uninstallation or Data Clearance

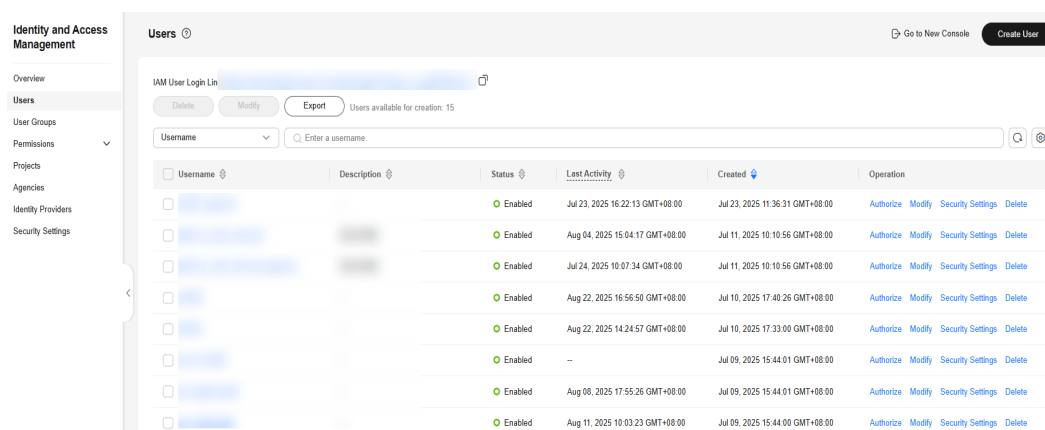
Generally, an error is reported in the uninstallation or data clearance step because the user does not have the IAM permission to uninstall the plugin using the UniAgent. For details, see [Managing ICAgent Plug-ins for Hosts](#).

Step 1 Log in to [IAM](#).

If you do not have the IAM access permission, contact the administrator.

Step 2 In the navigation pane, choose **Users**.

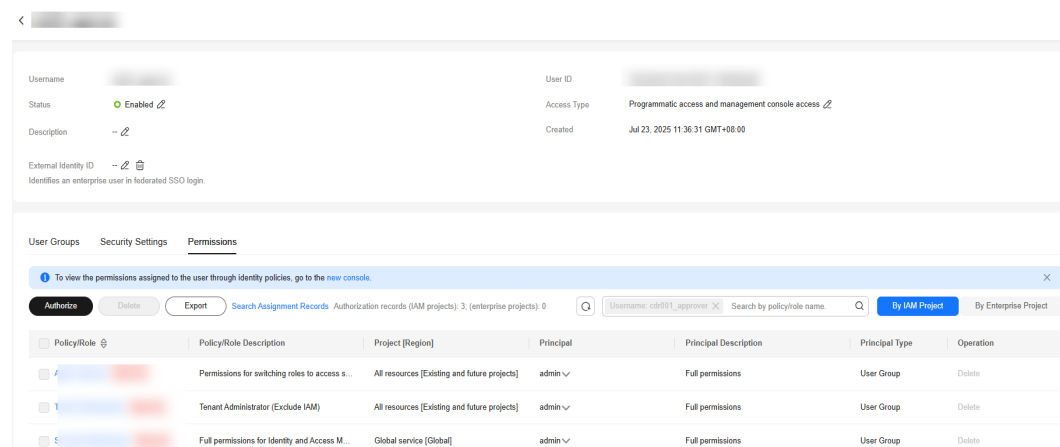
Figure 8-4 Viewing users



Username	Description	Status	Last Activity	Created	Operation
[redacted]	[redacted]	Enabled	Jul 23, 2025 16:22:13 GMT+08:00	Jul 23, 2025 11:36:31 GMT+08:00	Authorize Modify Security Settings Delete
[redacted]	[redacted]	Enabled	Aug 04, 2025 15:04:17 GMT+08:00	Jul 11, 2025 10:10:56 GMT+08:00	Authorize Modify Security Settings Delete
[redacted]	[redacted]	Enabled	Jul 24, 2025 10:07:34 GMT+08:00	Jul 11, 2025 10:10:56 GMT+08:00	Authorize Modify Security Settings Delete
[redacted]	[redacted]	Enabled	Aug 22, 2025 16:56:50 GMT+08:00	Jul 10, 2025 17:40:26 GMT+08:00	Authorize Modify Security Settings Delete
[redacted]	[redacted]	Enabled	Aug 22, 2025 14:24:57 GMT+08:00	Jul 10, 2025 17:33:00 GMT+08:00	Authorize Modify Security Settings Delete
[redacted]	[redacted]	Enabled	--	Jul 09, 2025 15:44:01 GMT+08:00	Authorize Modify Security Settings Delete
[redacted]	[redacted]	Enabled	Aug 08, 2025 17:55:26 GMT+08:00	Jul 09, 2025 15:44:01 GMT+08:00	Authorize Modify Security Settings Delete
[redacted]	[redacted]	Enabled	Aug 11, 2025 10:03:23 GMT+08:00	Jul 09, 2025 15:44:00 GMT+08:00	Authorize Modify Security Settings Delete

Step 3 In the displayed page, click the target username. The user details page is displayed.

Step 4 Click the **Permissions** tab to view all permissions of the current user.

Figure 8-5 Viewing user authorization records

Step 5 Click the permission name to view details. In the policy content, check whether the action list contains the **aom:uniagentAgent:uninstall** permission.

Figure 8-6 Viewing permissions

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Action": [
        "coc:*:*",
        "coc:list*",
        "coc:get*",
        "aom:uniagentAgent:install",
        "aom:uniagentAgent:uninstall",
        "aom:uniagentAgent:list",
        "aom:uniagentAgent:get",
        "aom:uniagentInstallHost:list",
        "aom:uniagentProxyRegion:get",
        "aom:view:get",
        "iam:agencies:createAgency",
        "iam:permissions:grantRoleToAgencyOnDomain",
        "iam:roles:listRoles",
        "iam:agencies:list*",
        "iam:agencies:createServiceLinkedAgency",
        "iam:agencies:deleteServiceLinkedAgency",
        ...
      ]
    }
  ]
}
```

Step 6 If the current account does not have the required permissions, contact the account administrator to request **aom:uniagentAgent:uninstall** permissions for the login user in IAM.

For details, see [Creating a User Group and Assigning Permissions](#).

----End

9 FAQs About Change Ticket Management

9.1 What Are the Differences Between Regular Changes and Emergency Changes?

Conceptual Differences

Regular changes are non-emergency changes that can be requested, evaluated, reviewed, sorted, planned, tested, and implemented using regular procedures.

Emergency changes are unplanned changes that are proposed to meet urgent service requirements when the production environment is unavailable, VMs are unavailable, unplanned changes for urgent service requirements, or changes cannot be evaluated and reviewed in time through regular procedures.

Differences in Review

Review is supported for both regular and urgent changes.

9.2 How Are Change Levels Defined?

The change level (with A-level risk being the highest, decreasing sequentially to D-level) is a common risk quantification grading logic in change management systems. The key is to match changes carrying different risk levels with appropriate control resources and review processes. This avoids insufficient control for high-risk changes or excessive control over low-risk changes.

The essence of change levels is to transform abstract risks into executable control standards. The level classification from A to D is not based solely on subjective judgment, but rather a quantitative assessment combining three core dimensions: scope of impact, probability of occurrence, and severity of consequences. The specific corresponding relationships can be referred to in the table below:

Table 9-1 Change levels

Level	Core Risk Characteristics (Quantification Dimension)	Typical Risk Consequence
Level A (high risk)	<ul style="list-style-type: none">• Scope of impact: cross-departmental/ cross-business line;• Probability of occurrence: medium-high;• Severity of consequences: Services are interrupted, compliance violations occur, and huge economic losses are caused.	Services are interrupted, customers complain on a large scale, regulatory penalties are imposed, and huge economic losses are caused.
Level B (high risk)	<ul style="list-style-type: none">• Scope of impact: core services of a single department;• Probability of occurrence: medium;• Severity of consequences: Department-level efficiency decreases, and partial compliance risks exist.	Department work is delayed, there are minor compliance risks, and huge economic losses are caused.
Level C (medium risk)	<ul style="list-style-type: none">• Scope of impact: a single team or a single service phase;• Probability of occurrence: low;• Severity of consequences: Operations are inconvenient, and no direct economic loss is caused.	The team efficiency decreases slightly, and no additional cost is required for rectification.

Level	Core Risk Characteristics (Quantification Dimension)	Typical Risk Consequence
Level D (lowest risk)	<ul style="list-style-type: none">• Scope of impact: individual operations/ non-core phases;• Probability of occurrence: extremely low;• Severity of consequences: No substantial impact.	No negative consequences occur. The experience may even be optimized.

10 Resilience Center FAQs

10.1 What Is a Chaos Drill?

With the transformation from traditional IT infrastructure O&M to cloud service O&M, traditional O&M methods face challenges such as complex inter-service invoking, fast application iteration, massive O&M objects, and complex non-linearity systems. Service downtime will bring huge economic losses and reputational damage to a company.

To solve this problem, chaos engineering is introduced to the O&M process. Performing chaos drills periodically helps identify system weaknesses (such as software bugs, solution design deficiencies, and fault recovery process points) before issues occur on the live network. In this way, system availability problems can be detected and resolved in a timely manner, improving application resilience and building O&M confidence. For unavoidable scenarios (such as hardware faults, abnormal server power-off, and network device board faults), formulate a contingency plan for quick fault recovery in advance.

COC allows you to perform automatic chaos drills covering from risk identification, emergency plan management, fault injection, and review and improvement. Based on years of best practices of Huawei Cloud SRE in chaos drills, customers can proactively identify, mitigate, and verify risks of cloud applications, improving the resilience of cloud applications.

10.2 What Are the Available Attack Scenarios?

Chaos drills support multiple attack scenarios, including disruptors for practicing, host resources, host processes, host networks, user-defined faults, and resource O&M. By integrating disruptor modules and functions, you can accurately simulate faults in the actual environment and identify system availability issues as early as possible, continuously improving application resilience. IPv6 fault drills of ECSs, BMSs, and on-premises IDC devices are supported. The drills of host network disruptors help you quickly master fault locating and emergency response capabilities in IPv6 networking environments, ensuring high network availability and security.

Attack Scenario Description

Table 10-1 Attack scenario description

Source of Attack Target	Attack Scenario		Description
ECSs	Disruptors for practicing	Qualifying practice	You can understand the chaos engineering process without worrying real faults.
	Host resources	CPU usage increase	Simulate CPU usage surge. The drill can be terminated in an emergency scenario.
		Memory usage increase	Simulate the memory usage surge. The drill can be terminated in an emergency scenario.
		Disk usage increase	Simulate the disk usage surge. The drill can be terminated in an emergency scenario.
		Disk I/O pressure increase	Continuously read and write files to increase disk I/O pressure. The drill can be terminated in an emergency scenario.
	Host process	Process ID exhaustion	The system process IDs (PIDs) are exhausted. The drill cannot be terminated in an emergency scenario.
		Process killing	Kill processes repeatedly during the fault duration. The drill can be terminated in an emergency scenario. After the emergency termination or drill is complete, the drill system does not start the processes. The service needs to ensure that the processes are restored.
	Host network	Network latency	Simulate network faults to increase link latency. The drill can be terminated in an emergency scenario.
		Network packet loss	Simulate network faults to cause packet loss on links. The drill can be terminated in an emergency scenario. The drill cannot be terminated when the packet loss rate is 100%.

Source of Attack Target	Attack Scenario		Description
		Network error packets	Simulate network faults to cause error packets on links. The drill can be terminated in an emergency scenario. The drill cannot be terminated when the packet error rate reaches 100%.
		Duplicate packets	Simulate duplicate packets generated on a link due to a network fault. The drill can be terminated in an emergency scenario.
		Network packet disorder	Simulate packet disorder generated on a link due to a network fault. The drill can be terminated in an emergency scenario.
		Network disconnection	Simulate the network disconnection between nodes. The drill can be terminated in an emergency scenario. Do not enter the IP addresses of the drill system and UniAgent server. Otherwise, the drill may fail. To interrupt an established persistent connection, select All for the interruption direction.
		NIC break-down	Simulate the NIC break-down scenario. The NIC may fail to be started after the NIC breaks down due to different network configurations of hosts. Therefore, prepare a contingency plan for network recovery. The drill cannot be terminated in an emergency scenario.
		DNS tampering	Tamper with the domain name address mapping. The drill can be terminated in an emergency scenario.
		Port occupation	Simulate the scenario where network ports of the system are occupied (a maximum of 100 ports can be occupied). The drill can be terminated in an emergency scenario.
		Server disconnection	Simulate the scenario where the entire server is disconnected, reject all TCP, UDP, and ICMP data packets, and open only ports 22, 8002, 39604, 33552, 33554, 33557, 32552, 32554, and 32557. The drill can be terminated in an emergency scenario.
		NIC bandwidth limiting	Limit the NIC bandwidth, support multiple NICs. The drill can be terminated in an emergency scenario.

Source of Attack Target	Attack Scenario		Description
		Connect ion exhausti on	Create a large number of socket connections to the specified server end (combination of the IP address and port number) to exhaust the connections. As a result, normal requests of the node cannot connect to the server (the requests of other nodes on the server may also be affected). The drill can be terminated in an emergency scenario.
	Custom izing a fault	Customi zing a script	Users can create scripts using automated O&M scripts and run the scripts to simulate faults. The drill can be terminated in an emergency scenario.
	Resourc e O&M	Startup	Start ECSs, BMSs, and Flexus instances in batches. Status may not be synchronized in a timely manner. The drill can be terminated in an emergency scenario.
		Shutdo wn	Shut down ECSs, BMSs, and Flexus instances in batches. Status may not be synchronized in a timely manner. The drill can be terminated in an emergency scenario.
		Restart	Restart ECSs, BMSs, and Flexus instances in batches. Status may not be synchronized in a timely manner. The drill can be terminated in an emergency scenario.
BMSs	Disrupt ors for practici ng	Qualifyi ng practice	You can understand the chaos engineering process without worrying real faults.
	Host resourc es	CPU usage increase	Simulate CPU usage surge. The drill can be terminated in an emergency scenario.
		Memory usage increase	Simulate the memory usage surg. The drill can be terminated in an emergency scenario.
		Disk usage increase	Simulate the disk usage surge. The drill can be terminated in an emergency scenario.
		Disk I/O pressure increase	Continuously read and write files to increase disk I/O pressure. The drill can be terminated in an emergency scenario.

Source of Attack Target	Attack Scenario		Description
	Host process	Process ID exhaustion	The system process IDs (PIDs) are exhausted. The drill cannot be terminated in an emergency scenario.
		Process killing	Kill processes repeatedly during the fault duration. The drill can be terminated in an emergency scenario. After the emergency termination or drill is complete, the drill system does not start the processes. The service needs to ensure that the processes are restored.
	Host network	Network latency	Simulate network faults to increase link latency. The drill can be terminated in an emergency scenario.
		Network packet loss	Simulate network faults to cause packet loss on links. The drill can be terminated in an emergency scenario. The drill cannot be terminated when the packet loss rate is 100%.
		Network error packets	Simulate network faults to cause error packets on links. The drill can be terminated in an emergency scenario. The drill cannot be terminated when the packet error rate reaches 100%.
		Duplicate packets	Simulate duplicate packets generated on a link due to a network fault. The drill can be terminated in an emergency scenario.
		Network packet disorder	Simulate packet disorder generated on a link due to a network fault. The drill can be terminated in an emergency scenario.
		Network disconnection	Simulate the network disconnection between nodes. The drill can be terminated in an emergency scenario. Do not enter the IP addresses of the drill system and UniAgent server. Otherwise, the drill may fail. To interrupt an established persistent connection, select All for the interruption direction.
		NIC break-down	Simulate the NIC break-down scenario. The NIC may fail to be started after the NIC breaks down due to different network configurations of hosts. Therefore, prepare a contingency plan for network recovery. The drill cannot be terminated in an emergency scenario.

Source of Attack Target	Attack Scenario		Description
		DNS tempering	Tamper with the domain name address mapping. The drill can be terminated in an emergency scenario.
		Port occupation	Simulate the scenario where network ports of the system are occupied (a maximum of 100 ports can be occupied). The drill can be terminated in an emergency scenario.
		Server disconnection	Simulate the scenario where the entire server is disconnected, reject all TCP, UDP, and ICMP data packets, and open only ports 22, 8002, 39604, 33552, 33554, 33557, 32552, 32554, and 32557. The drill can be terminated in an emergency scenario.
	Resource O&M	Startup	Start ECSs, BMSs, and Flexus instances in batches. Status may not be synchronized in a timely manner. The drill can be terminated in an emergency scenario.
		Shutdown	Shut down ECSs, BMSs, and Flexus instances in batches. Status may not be synchronized in a timely manner. The drill can be terminated in an emergency scenario.
		Restart	Restart ECSs, BMSs, and Flexus instances in batches. Status may not be synchronized in a timely manner. The drill can be terminated in an emergency scenario.
FlexusL instances (HCSS)	Disruptors for practicing	Qualifying practice	You can understand the chaos engineering process without worrying real faults.
	Host resources	CPU usage increase	Simulate CPU usage surge. The drill can be terminated in an emergency scenario.
		Memory usage increase	Simulate the memory usage surge. The drill can be terminated in an emergency scenario.
		Disk usage increase	Simulate the disk usage surge. The drill can be terminated in an emergency scenario.
		Disk I/O pressure increase	Continuously read and write files to increase disk I/O pressure. The drill can be terminated in an emergency scenario.

Source of Attack Target	Attack Scenario		Description
	Host process	Process ID exhaustion	The system process IDs (PIDs) are exhausted. The drill cannot be terminated in an emergency scenario.
		Process killing	Kill HCSS processes repeatedly during the fault duration. The drill can be terminated in an emergency scenario. After the emergency termination or drill is complete, the drill system does not start the processes. The service needs to ensure that the processes are restored.
	Host network	Network latency	Simulate network faults to increase link latency. The drill can be terminated in an emergency scenario.
		Network packet loss	Simulate network faults to cause packet loss on links. The drill can be terminated in an emergency scenario. The drill cannot be terminated when the packet loss rate is 100%.
		Network error packets	Simulate network faults to cause error packets on links. The drill can be terminated in an emergency scenario. The drill cannot be terminated when the packet error rate reaches 100%.
		Duplicate packets	Simulate duplicate packets generated on a link due to a network fault. The drill can be terminated in an emergency scenario.
		Network packet disorder	Simulate packet disorder generated on a link due to a network fault. The drill can be terminated in an emergency scenario.
		Network disconnection	Simulate the network disconnection between nodes. The drill can be terminated in an emergency scenario. Do not enter the IP addresses of the drill system and UniAgent server. Otherwise, the drill may fail. To interrupt an established persistent connection, select All for the interruption direction.
		NIC break-down	Simulate the NIC break-down scenario. The NIC may fail to be started after the NIC breaks down due to different network configurations of hosts. Therefore, prepare a contingency plan for network recovery. The drill cannot be terminated in an emergency scenario.

Source of Attack Target	Attack Scenario		Description
		DNS tempering	Tamper with the domain name address mapping. The drill can be terminated in an emergency scenario.
		Port occupation	Simulate the scenario where network ports of the system are occupied (a maximum of 100 ports can be occupied). The drill can be terminated in an emergency scenario.
		Server disconnection	Simulate the scenario where the entire server is disconnected, reject all TCP, UDP, and ICMP data packets, and open only ports 22, 8002, 39604, 33552, 33554, 33557, 32552, 32554, and 32557. The drill can be terminated in an emergency scenario.
	Resource O&M	Startup	Start ECSs, BMSs, and Flexus instances in batches. Status may not be synchronized in a timely manner. The drill can be terminated in an emergency scenario.
		Shutdown	Shut down ECSs, BMSs, and Flexus instances in batches. Status may not be synchronized in a timely manner. The drill can be terminated in an emergency scenario.
		Restart	Restart ECSs, BMSs, and Flexus instances in batches. Status may not be synchronized in a timely manner. The drill can be terminated in an emergency scenario.
CCE nodes	Disruptors for practicing	Qualifying practice	You can understand the chaos engineering process without worrying real faults.
	Host resources	CPU usage increase	Simulate CPU usage surge. The drill can be terminated in an emergency scenario.
		Memory usage increase	Simulate the memory usage surge. The drill can be terminated in an emergency scenario.
		Disk usage increase	Simulate the disk usage surge. The drill can be terminated in an emergency scenario.
		Disk I/O pressure increase	Continuously read and write files to increase disk I/O pressure. The drill can be terminated in an emergency scenario.

Source of Attack Target	Attack Scenario		Description
	Host process	Process ID exhaustion	The system process IDs (PIDs) are exhausted. The drill cannot be terminated in an emergency scenario.
		Process killing	Kill processes repeatedly during the fault duration. The drill can be terminated in an emergency scenario. After the emergency termination or drill is complete, the drill system does not start the processes. The service needs to ensure that the processes are restored.
	Host network	Network latency	Simulate network faults to increase link latency. The drill can be terminated in an emergency scenario.
		Network packet loss	Simulate network faults to cause packet loss on links. The drill can be terminated in an emergency scenario. The drill cannot be terminated when the packet loss rate is 100%.
		Network error packets	Simulate network faults to cause error packets on links. The drill can be terminated in an emergency scenario. The drill cannot be terminated when the packet error rate reaches 100%.
		Duplicate packets	Simulate duplicate packets generated on a link due to a network fault. The drill can be terminated in an emergency scenario.
		Network packet disorder	Simulate packet disorder generated on a link due to a network fault. The drill can be terminated in an emergency scenario.
		Network disconnection	Simulate the network disconnection between nodes. The drill can be terminated in an emergency scenario. Do not enter the IP addresses of the drill system and UniAgent server. Otherwise, the drill may fail. To interrupt an established persistent connection, select All for the interruption direction.
		NIC break-down	Simulate the NIC break-down scenario. The NIC may fail to be started after the NIC breaks down due to different network configurations of hosts. Therefore, prepare a contingency plan for network recovery. The drill cannot be terminated in an emergency scenario.

Source of Attack Target	Attack Scenario		Description
		DNS tempering	Tamper with the domain name address mapping. The drill can be terminated in an emergency scenario.
		Port occupation	Simulate the scenario where network ports of the system are occupied (a maximum of 100 ports can be occupied). The drill can be terminated in an emergency scenario.
		Server disconnection	Simulate the scenario where the entire server is disconnected, reject all TCP, UDP, and ICMP data packets, and open only ports 22, 8002, 39604, 33552, 33554, 33557, 32552, 32554, and 32557. The drill can be terminated in an emergency scenario.
		NIC bandwidth limiting	Limit the NIC bandwidth, support multiple NICs. The drill can be terminated in an emergency scenario.
		Connection exhaustion	Create a large number of socket connections to the specified server end (combination of the IP address and port number) to exhaust the connections. As a result, normal requests of the node cannot connect to the server (the requests of other nodes on the server may also be affected). The drill can be terminated in an emergency scenario.
CCE pods	Pod resources	Pod CPU usage increase	Simulate a pod CPU usage surge. Ensure the attack target is writable. If it is not, the drill will fail. If the drill fails, you can use the emergency termination function.
		Pod memory usage increase	Simulate a pod memory usage surge. Ensure the attack target is writable. If it is not, the drill will fail. The drill can be terminated in an emergency scenario.
		Pod disk I/O pressure	Continuously simulates I/O reads and writes. The drill can be terminated in an emergency scenario.
		Pod disk usage increase	Writes large files to a specified directory to simulate the pressure increase of the Kubernetes container file system. The drill can be terminated in an emergency scenario.

Source of Attack Target	Attack Scenario		Description
	Pod process	Forcible pod stopping	Forcibly stop a pod. The drill cannot be terminated in an emergency scenario.
		Forcibly killing containers in a pod	Forcibly kill containers in a pod. The drill cannot be terminated in an emergency scenario.
	Pod network	Pod network latency	Simulate a network fault that incurs the network latency increase in a pod. Drills can be terminated in an emergency scenario. Drills cannot be terminated when the latency reaches 30,000 ms.
		Pod network packet loss	Simulate a network fault that incurs packet loss in a pod. Drills can be terminated in an emergency scenario.
		Pod network interruption	Simulate a network disconnection between a POD and other IP addresses. The drill can be terminated in an emergency scenario. To interrupt an established persistent connection, select all directions as the directions to be interrupted.
		Pod network packet disorder	Simulate packet disorder generated on a link due to a pod network fault. Drills can be terminated in an emergency scenario.
		Duplicate pod network packets	Simulate duplicate packets generated on a link due to a pod network fault. Drills can be terminated in an emergency scenario.
		Pod DNS tampering	If the address mapping of the domain name is tampered with in the pod, ensure that the running user of the attack target is root. Otherwise, the drill will fail due to insufficient permission. The drill can be terminated in an emergency scenario.
		Pod port masking	Simulate disabling of a pod port. The drill can be terminated in an emergency scenario.

Source of Attack Target	Attack Scenario		Description
		Pod network isolation	Simulate the scenario where access from a pod to another IP address network is directly rejected. The drill can be terminated in an emergency scenario. If you need to reject established persistent connections, select All for Direction .
RDS instances	Instances	RDS primary / standby switchover	Only MySQL and PostgreSQL engines in HA mode are supported. This operation is not allowed during creating and restarting instances, upgrading databases, recovering and modifying ports, as well as creating and deleting accounts. Primary/standby switchover cannot change the IP address of the internal network of an instance. The drill cannot be terminated in an emergency scenario.
		Stopping an RDS instance	Stop both the primary and read-only instances. After the fault duration ends, start the instance. The drill can be terminated in an emergency.
DCS instances	Instances	DCS master/standby switchover	Switch the master and standby DB instance nodes. This operation is supported only for master/standby DB instances. The drill cannot be terminated in an emergency scenario.
		DCS instance restart	Restart a running DCS instance. If you clear data of a Redis 4.0, 5.0, or 6.0 instance, the cleared data cannot be restored. Exercise caution when performing this operation. The drill cannot be terminated in an emergency scenario.
		Powering off a DCS AZ	All nodes in the AZ are powered off centrally. The drill cannot be terminated in an emergency scenario. This disruptor is not supported in some areas.
CSS instances	Instances	Restarting a CSS cluster	Restart the CSS cluster that is in the available status. During the restart, Kibana and Cerebro may fail to be accessed. The drill cannot be terminated in an emergency scenario.

Source of Attack Target	Attack Scenario		Description
DDS instances	Instances	Forcibly promoting a secondary node to primary	Supported forcible promotion of secondary nodes to primary for backup sets, shards, and config nodes. However, there is a risk of failure when the primary/secondary latency is large. The drill cannot be terminated in an emergency scenario.
IDC offline resource VMs	Disruptors for practicing	Qualifying practice	You can understand the chaos engineering process without worrying real faults.
	Host resources	CPU usage increase	Simulate CPU usage surge. The drill can be terminated in an emergency scenario.
		Memory usage increase	Simulate the memory usage surge. The drill can be terminated in an emergency scenario.
		Disk usage increase	Simulate the disk usage surge. The drill can be terminated in an emergency scenario.
		Disk I/O pressure increase	Continuously read and write files to increase disk I/O pressure. The drill can be terminated in an emergency scenario.
	Host process	Process ID exhaustion	The system process IDs (PIDs) are exhausted. The drill cannot be terminated in an emergency scenario.
		Process killing	Kill processes repeatedly during the fault duration. The drill can be terminated in an emergency scenario. After the emergency termination or drill is complete, the drill system does not start the processes. The service needs to ensure that the processes are restored.
	Host network	Network latency	Simulate network faults to increase link latency. The drill can be terminated in an emergency scenario.
		Network packet loss	Simulate network faults to cause packet loss on links. The drill can be terminated in an emergency scenario. The drill cannot be terminated when the packet loss rate is 100%.

Source of Attack Target	Attack Scenario		Description
		Network error packets	Simulate network faults to cause error packets on links. The drill can be terminated in an emergency scenario. The drill cannot be terminated when the packet error rate reaches 100%.
		Duplicate packets	Simulate duplicate packets generated on a link due to a network fault. The drill can be terminated in an emergency scenario.
		Network packet disorder	Simulate packet disorder generated on a link due to a network fault. The drill can be terminated in an emergency scenario.
		Network disconnection	Simulate the network disconnection between nodes. The drill can be terminated in an emergency scenario. Do not enter the IP addresses of the drill system and UniAgent server. Otherwise, the drill may fail. To interrupt an established persistent connection, select All for the interruption direction.
		NIC break-down	Simulate the NIC break-down scenario. The NIC may fail to be started after the NIC breaks down due to different network configurations of hosts. Therefore, prepare a contingency plan for network recovery. The drill cannot be terminated in an emergency scenario.
		DNS tampering	Tamper with the domain name address mapping. The drill can be terminated in an emergency scenario.
		Port occupation	Simulate the scenario where network ports of the system are occupied (a maximum of 100 ports can be occupied). The drill can be terminated in an emergency scenario.
		Server disconnection	Simulate the scenario where the entire server is disconnected, reject all TCP, UDP, and ICMP data packets, and open only ports 22, 8002, 39604, 33552, 33554, 33557, 32552, 32554, and 32557. The drill can be terminated in an emergency scenario.
		NIC bandwidth limiting	Limit the NIC bandwidth, support multiple NICs. The drill can be terminated in an emergency scenario.

Source of Attack Target	Attack Scenario		Description
		Connect ion exhausti on	Create a large number of socket connections to the specified server end (combination of the IP address and port number) to exhaust the connections. As a result, normal requests of the node cannot connect to the server (the requests of other nodes on the server may also be affected). The drill can be terminated in an emergency scenario.
Alibaba Cloud server	Disrupt ors for practicing	Qualifyi ng practice	You can understand the chaos engineering process without worrying real faults.
	Host resourc es	CPU usage increase	Simulate CPU usage surge. The drill can be terminated in an emergency scenario.
		Memory usage increase	Simulate the memory usage surg. The drill can be terminated in an emergency scenario.
		Disk usage increase	Simulate the disk usage surge. The drill can be terminated in an emergency scenario.
		Disk I/O pressure increase	Continuously read and write files to increase disk I/O pressure. The drill can be terminated in an emergency scenario.
	Host process	Process ID exhausti on	The system process IDs (PIDs) are exhausted. The drill cannot be terminated in an emergency scenario.
		Process killing	Kill processes repeatedly during the fault duration. The drill can be terminated in an emergency scenario. After the emergency termination or drill is complete, the drill system does not start the processes. The service needs to ensure that the processes are restored.
	Host network	Network latency	Simulate network faults to increase link latency. The drill can be terminated in an emergency scenario.
		Network packet loss	Simulate network faults to cause packet loss on links. The drill can be terminated in an emergency scenario. The drill cannot be terminated when the packet loss rate is 100%.

Source of Attack Target	Attack Scenario		Description
		Network error packets	Simulate network faults to cause error packets on links. The drill can be terminated in an emergency scenario. The drill cannot be terminated when the packet error rate reaches 100%.
		Duplicate packets	Simulate duplicate packets generated on a link due to a network fault. The drill can be terminated in an emergency scenario.
		Network packet disorder	Simulate packet disorder generated on a link due to a network fault. The drill can be terminated in an emergency scenario.
		Network disconnection	Simulate the network disconnection between nodes. The drill can be terminated in an emergency scenario. Do not enter the IP addresses of the drill system and UniAgent server. Otherwise, the drill may fail. To interrupt an established persistent connection, select All for the interruption direction.
		NIC break-down	Simulate the NIC break-down scenario. The NIC may fail to be started after the NIC breaks down due to different network configurations of hosts. Therefore, prepare a contingency plan for network recovery. The drill cannot be terminated in an emergency scenario.
		DNS tampering	Tamper with the domain name address mapping. The drill can be terminated in an emergency scenario.
		Port occupation	Simulate the scenario where network ports of the system are occupied (a maximum of 100 ports can be occupied). The drill can be terminated in an emergency scenario.
		Server disconnection	Simulate the scenario where the entire server is disconnected, reject all TCP, UDP, and ICMP data packets, and open only ports 22, 8002, 39604, 33552, 33554, 33557, 32552, 32554, and 32557. The drill can be terminated in an emergency scenario.
		NIC bandwidth limiting	Limit the NIC bandwidth, support multiple NICs. The drill can be terminated in an emergency scenario.

Source of Attack Target	Attack Scenario		Description
		Connect ion exhausti on	Create a large number of socket connections to the specified server end (combination of the IP address and port number) to exhaust the connections. As a result, normal requests of the node cannot connect to the server (the requests of other nodes on the server may also be affected). The drill can be terminated in an emergency scenario.

10.3 What Is a Failure Mode?

A failure mode is a category of potential risks faced by cloud applications. Years of failure modes accumulated on Huawei Cloud are preconfigured on the chaos drill platform. The FT-FMEA fault analysis method is used to help you analyze the potential risks of cloud applications.

Failure modes focus on risk assessment of cloud applications. By systematically evaluating the application architecture, dependency relationship, and potential weak points, the system can precisely identify risk scenarios that could lead to service anomalies (such as node failures, network latency, resource exhaustion). Failure modes are the core premise and basis for conducting chaos drills.

10.4 What Do Drill Plans Do?

Drill plans play a core role in chaos drill management. They not only help engineers systematically schedule drill tasks for various failure modes but also trace and manage the drill task progress at any time. This ensures that failure modes are effectively verified through drills.

In specific operations, when creating drill plans, it is necessary to clearly specify the executor and the scheduled execution time, forming a standard basis for task allocation. After the executor confirms the task by accepting it, the system automatically generates the corresponding drill task. This task must accurately associate the predefined failure modes with the specific application regions. This ensures a high degree of alignment between the drill scenarios and actual business scenarios, providing strong support for evaluating the drill effectiveness and optimizing subsequent processes.

Drill plans encompass the entire lifecycle of the chaos drill.

- In the initiation phase, by analyzing business priorities, historical failure data, and risk assessment results, drill managers can systematically schedule drill tasks for high-impact, frequent failure modes on a tiered and periodic basis. This avoids the waste of drill resources and the omission of key scenarios.
- In the execution phase, the system traces the task startup status, execution progress, and completion of key nodes (such as fault injection, emergency response, and fault rectification) in real time. The system also provides early

warning alerts for delayed tasks and coordination of executors. The whole drill process is effectively managed to ensure that each drill action is implemented as planned.

- Through standard drill plans, failure modes can be effectively verified through the complete drill loop of simulated scenarios, actual actions, post-action review. This process verifies the feasibility of contingency plans and the team's capabilities to handle emergencies.

10.5 What Is the Relationship Between a Failure Mode and a Drill Task?

A failure mode and a drill task are closely linked and building upon each other in the chaos drill system. They form a closed loop for risk identification through fault injection and capability verification through task execution.

Failure modes focus on risk assessment of cloud applications. By systematically evaluating the application architecture, dependency relationship, and potential weak points, the system can precisely identify risk scenarios that could lead to service anomalies (such as node failures, network latency, resource exhaustion). Failure modes are the core premise and basis for conducting chaos drills.

Drill tasks are the carriers for failure modes to be implemented. Based on these failure modes, individual or interconnected fault scenarios are designed and simulated. Fault injection tools (such as server outage simulation and traffic congestion injection) are then employed to replicate the corresponding risks in a controlled environment. Finally, the application's fault tolerance, auto-recovery efficiency, and the effectiveness of contingency plans are validated. This process enables the transformation of risk identification into tangible capability verification.

10.6 What Are Included in a Drill Report?

After a drill task is complete, you can directly create a report if necessary. After a report is created, you can export the report as a PDF file and send it to related personnel. The entire process is flexible and efficient, meeting the requirements of the entire process from production to format fixing.

You can modify the actual fault recovery duration, create improvement tickets, and view fault records in a drill report so that you can comprehensively record and manage drill activities and results.

The drill report consists of the following modules:

- Recovery capability scoring module: You can modify the actual recovery duration. The system automatically generates a recovery capability score.
- Basic information module: displays basic information about a drill task, including the drill task name, drill report ID, drill start time, end time, drill executor, drill duration, and expected fault recovery duration (minutes).
- Drill process module: displays drill task cards.
- Attack task group module: displays attack task details, including attack targets, steady-state metrics, and monitoring metrics.

The list of selected instances is displayed for attack targets. Line charts are displayed for steady-state metrics and monitoring metrics. If no data is available, **No data available** is displayed.

- Improvement item module: You can create improvement tickets. The improvement details are displayed by default, including the processing and verification information.

11

FAQs About Basic Configurations

11.1 How Do I Log In to COC as a Non-Common IAM User?

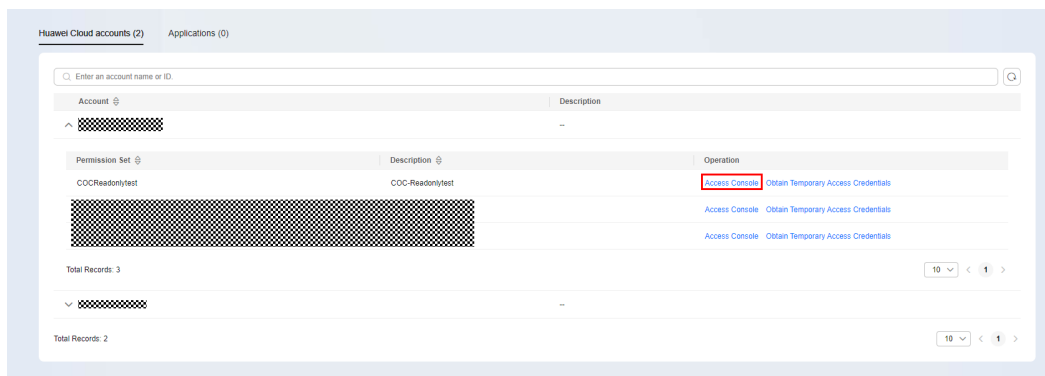
You can log in to COC as a common IAM user, IAM federated user (including IAM user in SSO mode and virtual user in SSO mode), and IAM Identity Center user. If you log in to COC as a non-common IAM user, follow the instructions in this section to ensure that you can use COC properly.

Logging in to COC as an IAM Identity Center User

IAM Identity Center provides unified identity management and access control for multiple accounts based on organizations. After a user is created in this service, the user can use a specific username and password to log in to the unified portal and access resources of multiple accounts assigned to the user without multiple logins.

- Step 1** Log in to Huawei Cloud through the IAM Identity Center portal. If you use IAM Identity Center for the first time, see [Getting Started](#).

Figure 11-1 IAM Identity Center login

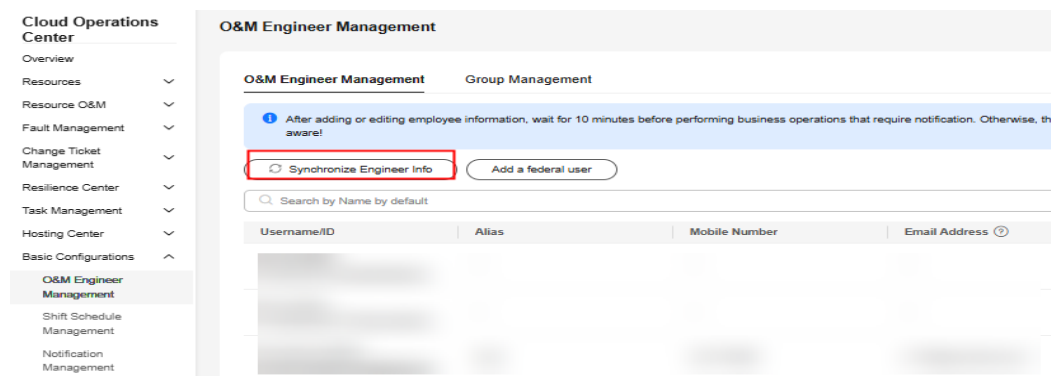


- Step 2** Access COC through the console.

Step 3 In the navigation pane, choose **Basic Configurations > O&M Engineer Management**.

Step 4 On the displayed page, click **Synchronize Engineer Info**.

Figure 11-2 Synchronizing information about O&M engineers



----End

Logging in to COC as an IAM Federated User

If you have an enterprise management system and want to use cloud service resources of a Huawei Cloud account, you can use the IAM identity provider function to log in to Huawei Cloud using your enterprise management system account in SSO mode. This process is called **federated identity authentication**. Users who log in to Huawei Cloud through federated identity authentication are called **federated users**. Currently, federated users include virtual users in SSO mode and IAM users in SSO mode. For details about the differences and use cases of the two types of federated users, see [Application Scenarios of Virtual User SSO and IAM User SSO](#).

If you log in to COC as an IAM user in SSO mode, refer to the process of using COC as a common IAM user.

If you log in to COC as a virtual user in SSO mode, perform the following operations.

Step 1 Log in to [COC](#).

Step 2 In the navigation pane, choose **Basic Configurations > O&M Engineer Management**.

Step 3 Click **Add Federated User** in the upper left corner.

Step 4 Set the parameters in the **Add Federated User** dialog box.

- **Username:** username displayed on Huawei Cloud, which is configured in the IAM identity transition rule.
- **Alias:** Alias of the current user.
- **Identity Provider Name:** name of the user identity provider in IAM.

Figure 11-3 Adding a federated user

Add a federal user

If you use the IAM virtual user SSO for federated login, you need to manually add a username.

★ Username ?

coc

Alias

coc

★ Identity Provider Name

coc-test

Click to view [List of identity providers](#)

Step 5 Click **OK**. The federated user is added.

-----End