Ubiquitous Cloud Native Service

FAQs

Issue 01

Date 2022-10-14

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions

Notice

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

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

Contents

1 About UCS

1.1 Are There Quota Restrictions in UCS?

UCS Quotas

Quotas put limits on the quantity or capacity of resources available to users. UCS allows you to set quotas for third-party clusters and cluster groups.

- Third-party cluster quota: the maximum number of third-party clusters that can be connected to UCS, excluding the number of CCE clusters.
- Cluster group quota: the maximum number of cluster groups for a user, excluding the default cluster group automatically created by UCS.

For other cloud services you may also use when running UCS, such as Elastic Cloud Server (ECS), Elastic Volume Service (EVS), Virtual Private Cloud (VPC), Elastic Load Balance, SoftWare Repository for Container (SWR), and Domain Name Service (DNS), their quotas are independent of those of UCS and are managed by themselves. For details, see **Quotas**.

Default Quota Settings

Table 1-1 lists the quota items and their defaults. You can also request a quota increase.

Table 1-1 UCS quota items

Quota Item	Default
Third-party cluster	50
Cluster group	50

How Do I Modify UCS Quotas?

Contact our technical support to increase UCS quotas.

Step 1 Log in to the Huawei Cloud management console.

Step 2 In the upper right corner, choose **More** > **Service Tickets** > **Create Service Ticket**.

The **Create Service Ticket** page is displayed.

Figure 1-1 Creating a service ticket

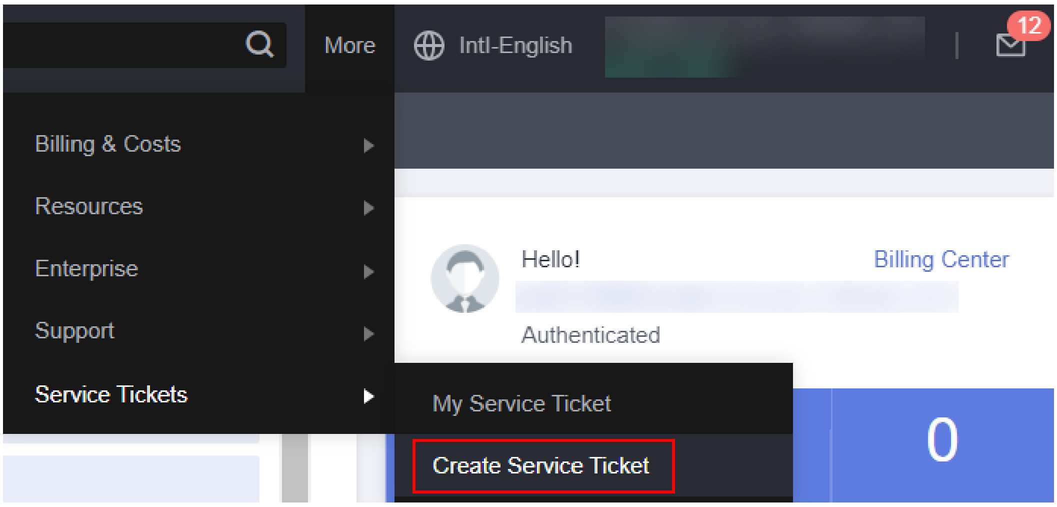

Step 3 Fill in the information and submit the service ticket.

Select **Quotas** for **Services**, choose **Quota Application** under **Issue Categories**, specify the quota to be adjusted and reason in the description area, and set other parameters as required.

Step 4 Select the agreement and click **Submit**.

----End

2 Permissions

2.1 How Do I Configure the Access Permission for Each Function of the UCS Console?

Background

IAM controls the permissions to use UCS console functions. You need to add cloud service permission policies to user groups on IAM. When an unauthorized user attempts to use the UCS console, an error message is displayed, indicating that the user does not have the access permission or permission authentication fails.

Huawei Cloud services often interact with each other for your applications to run. Some UCS functions are dependent on other services. By default, newly created IAM users do not have any permissions and cannot use any cloud service or function. Therefore, before they start using UCS, you need to grant them the required permissions listed in **Table 2-1**.

The following describes how to configure permissions for UCS console functions (such as CCE clusters, image repositories, and traffic distribution) for IAM users.

Table 2-1 Permissions on which the UCS function depends

Function	Permission	Dependent Permission	Description	Type
Connecting a cluster	-	Aavailable only to users in the IAM admin user group.	CCE clusters or other Kubernetes clusters can be connected on the UCS console for unified management.	-
Permission policies	Administrator permissions	Available only to users in the IAM admin user group.	You can create **permission policies** and templates.	-

Function	Permission	Dependent Permission	Description	Type
Cluster groups	Administrator permissions	Available only to users in the IAM admin user group.	Cluster groups can be created and deleted, and permission policies can be associated with cluster groups.	-
	Operation permissions	Members of the IAM admin user group need to associate **permission policies** with the cluster group.	The permission policy contains the resource permissions of container clusters. After a user group is associated with a permission policy, users in the user group can read clusters in the cluster group and add or remove clusters. **NOTE** The private network access of the cluster depends on VPC Endpoint. Therefore, the IAM user group must have the VPC Endpoint Administrator permission.	UCS permission policy
Container cluster - CCE cluster	Administrator permissions	CCE Administrator	Read and write permissions for CCE clusters and all resources (including workloads, nodes, jobs, and Services) in the clusters.	IAM system roles
	Operation permissions	CCE FullAccess	Common operation permissions on CCE cluster resources, excluding the namespace-level permissions for the clusters (with Kubernetes RBAC enabled) and the privileged administrator operations, such as agency configuration and cluster certificate generation For common operation permissions, you also need to configure cluster RBAC authorization. For details, see **Namespace Permissions (Kubernetes RBAC-based)**.	System-defined policies of IAM

Function	Permission	Dependent Permission	Description	Type
	Read-only permission	CCE ReadOnlyAccess	Permissions to view CCE cluster resources, excluding the namespace-level permissions of the clusters (with Kubernetes RBAC enabled) For the read-only permission, you also need to configure RBAC authorization for the cluster. For details, see **Namespace Permissions (Kubernetes RBAC-based)**.	System-defined policies of IAM
Container cluster - non-CCE cluster (For details about how to configure resource permissions, see **Cluster Operation Permissions**.)	Administrator permissions	Admin Permission Template	You need to grant permissions to the user group in the UCS policy center, and the user group must have any IAM permissions. Has the read and write permissions on all resources, including cluster permission management.	UCS permission policy
	Operation permissions	Developer Permission Template	You need to grant permissions to the user group in the UCS policy center, and the user group must have any IAM permissions. Has the read and write permissions on resources except cluster permission management.	UCS permission policy
	Read-only permission	ReadOnly Permission Template	You need to grant permissions to the user group in the UCS policy center, and the user group must have any IAM permissions. Has the read-only permission on all resources.	UCS permission policy
Image Repository	Administrator permissions	SWR Admin	SWR administrator permissions, including all SWR permissions.	IAM system roles
	Administrator permissions	SWR FullAccess	Full permissions for SWR.	System-defined policies of IAM

Function	Permission	Dependent Permission	Description	Type
	Operation permissions	SWR OperateAccess	Common operation permissions for SWR.	System-defined policies of IAM
	Read-only permission	SWR ReadOnlyAccess	Read-only permissions for SWR.	System-defined policies of IAM
Traffic Distribution	Administrator permissions	DNS Administrator	Has all permissions except those for the DNS service.	IAM system roles
	Read-only permission	Tenant Guest	Has read-only permissions on all services except IAM.	IAM system roles

Prerequisites

The IAM user has been added to a user group, and the user group has been associated with permission policies. For details, see **Permissions Policies**.

Procedure

Step 1 Log in to the IAM console as an administrator or a user in the admin user group.

Step 2 In the navigation pane, choose **User Groups**. In the user group list, click **Authorize** on the right of the target user group.

Step 3 Search for and select the permissions to be added, for example, **CCE Administrator**. Other permission policies on which the role depends are automatically selected. You can click **View Selected** or expand the policy details to learn about the dependencies. You can also scope the permissions an IAM user needs to the minimum for each function. For details, see **Table 2-1**.

Figure 2-1 Selecting policies

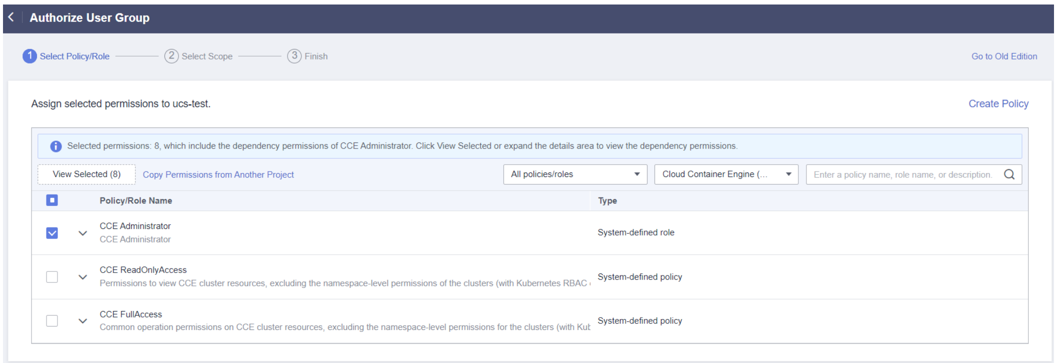

Assign selected permissions to ucs-test.

Selected permissions: 8, which include the dependency permissions of CCE Administrator. Click View Selected or expand the details area to view the dependency permissions.

	Policy/Role Name	Type
☑	CCE Administrator CCE Administrator	System-defined role
□	CCE ReadOnlyAccess Permissions to view CCE cluster resources, excluding the namespace-level permissions of the clusters (with Kubernetes RBAC	System-defined policy
□	CCE FullAccess Common operation permissions on CCE cluster resources, excluding the namespace-level permissions for the clusters (with Kut	System-defined policy

Step 4 Click **Next** and select a scope.

The default option **All resources** is selected, indicating that the IAM user will be able to use all resources, including those in enterprise projects, region-specific projects, and global services under your account based on assigned permissions.

Step 5 Click **OK** to complete the authorization. The authorization may take effect 15 to 30 minutes later.

----End

2.2 Why Can't an IAM User Obtain Clusters or Cluster Groups After Logging In to UCS?

Description

An IAM user logs in to the UCS console and finds an empty **Container Clusters** page, no cluster or cluster group displayed.

Troubleshooting

Check whether the IAM user has been granted the required permission policies. Managing CCE clusters requires IAM policies, while managing third-party clusters and cluster groups requires UCS policies, as listed in **Table 2-2**.

Table 2-2 Cluster permissions

Function	Permission	Dependent IAM Role	Description	Type
Container Clusters > Cluster Groups	Administrator permissions	Users in the IAM admin user group	Manages cluster groups, such as creation, deletion, and permission policy association.	-
	Operation permissions	Users of the IAM admin user group associated with the cluster group**permission policies**	The permission policies allow operations on the resources in a container cluster. After policy association, users in the user group can read clusters in the cluster group and add or remove clusters. **NOTE** VPC Endpoint is required if you connect a cluster to UCS through a private network. Therefore, the user group must have the IAM permission **VPCEndpoint Administrator**.	UCS permission policy

Function	Permission	Dependent IAM Role	Description	Type
Container Cluster - CCE Clusters	Administrator permissions	CCE Administrator	Read and write permissions for CCE clusters and all resources (including workloads, nodes, jobs, and Services) in the clusters.	IAM system-defined role
	Operation permissions	CCE FullAccess	Common operation permissions on CCE cluster resources, excluding the namespace-level permissions for the clusters (with Kubernetes RBAC enabled) and the privileged administrator operations, such as agency configuration and cluster certificate generation For common operation permissions, you also need to configure cluster RBAC authorization. For details, see **Namespace Permissions (Kubernetes RBAC-based)**.	IAM system-defined policy
	Read-only permissions	CCE ReadOnlyAccess	Permissions to view CCE cluster resources, excluding the namespace-level permissions of the clusters (with Kubernetes RBAC enabled) For the read-only permission, you also need to configure RBAC authorization for the cluster. For details, see **Namespace Permissions (Kubernetes RBAC-based)**.	IAM system-defined policy
Container Cluster - Third-Party Clusters (details available in **Table 2-3**)	Administrator permissions	Admin Permission Template	You need to grant permissions to the user group on the **Permissions Policies** page of the UCS console, and the user group must have any IAM permissions. Read and write permissions on all resources, including cluster permission management	UCS permission policy

Function	Permission	Dependent IAM Role	Description	Type
	Operation permissions	Developer Permission Template	You need to grant permissions to the user group on the **Permissions Policies** page of the UCS console, and the user group must have any IAM permissions. Read and write permissions on resources except cluster permission management	UCS permission policy
	Read-only permissions	ReadOnly Permission Template	You need to grant permissions to the user group on the **Permissions Policies** page of the UCS console, and the user group must have any IAM permissions. Read-only permissions on all resources	UCS permission policy

IAM users often fail to obtain data about their clusters due to incorrect permission settings. In UCS, IAM, UCS, and cluster group policies are required for managing both CCE clusters and third-party clusters at the same time.

If you manage only third-party clusters, you need to configure UCS policies and cluster group policies.

Therefore, try checking the following items:

- **Check Item 1: IAM Permissions**
- **Check Item 2: Permissions Policies**
- **Check Item 3: Permissions Policies Associated with the Cluster Group**

Prerequisites

Contact the Huawei Cloud account administrator or a user in the IAM admin user group to check.

Check Item 1: IAM Permissions

For CCE clusters, if an IAM user is not added to any user group or the user group does not have CCE permissions, the UCS console cannot obtain data about CCE clusters.

For third-party clusters alone, no IAM permission policy is required, not the case for hybrid management of both CCE clusters and third-party clusters.

For details about how to grant permissions to a user group, see **User Groups**.

Users with different IAM permissions have different namespace permissions (assigned using Kubernetes RBAC).

- **CCE Administrator**: administrator permissions. Users with this role can perform operations on all resources without configuring the cluster RBAC. If a user already has the **Tenant Administrator** role, the user can have the administrator permissions for all cloud services, including CCE but excluding IAM.
- **CCE FullAccess** or **CCE ReadOnlyAccess**: cluster operation/read-only permissions configured in IAM. You also need to configure RBAC for the cluster in the CCE console. For details, see **Namespace Permissions (Kubernetes RBAC-based)**.

Check Item 2: Permissions Policies

Log in to the UCS console. On the **Permissions Policies** page, create a permission policy, and associate a user group with the policy.

NOTE

This setting takes effect only for non-CCE clusters. For details, see **How Do I Configure Operation Permissions for Cluster Resources?**.

See **Check Item 1: IAM Permissions** to set IAM permissions and cluster RBAC for cluster resources.

Check Item 3: Permissions Policies Associated with the Cluster Group

The created permission policy must be associated with the cluster group. Otherwise, the UCS console will not display data about the cluster group.

Step 1 Log in to the UCS console. In the navigation pane, choose **Container Clusters**.

Step 2 In the card view of the target cluster group, click in the upper right corner.

Figure 2-2 Associating a permission policy with a cluster group

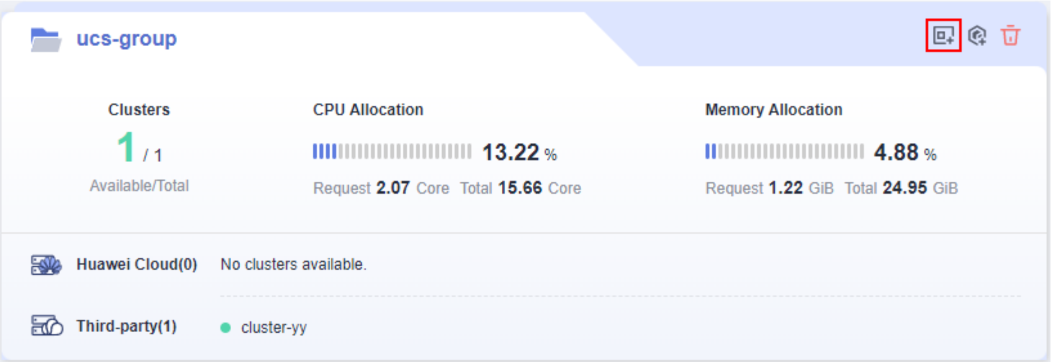

ucs-group

Clusters 1 / 1 Available/Total

CPU Allocation 13.22 % Request 2.07 Core Total 15.66 Core

Memory Allocation 4.88 % Request 1.22 GiB Total 24.95 GiB

Huawei Cloud(0) No clusters available.

Third-party(1) cluster-yy

Step 3 Select one or multiple existing permission policies for the cluster group.

Figure 2-3 Associating policies

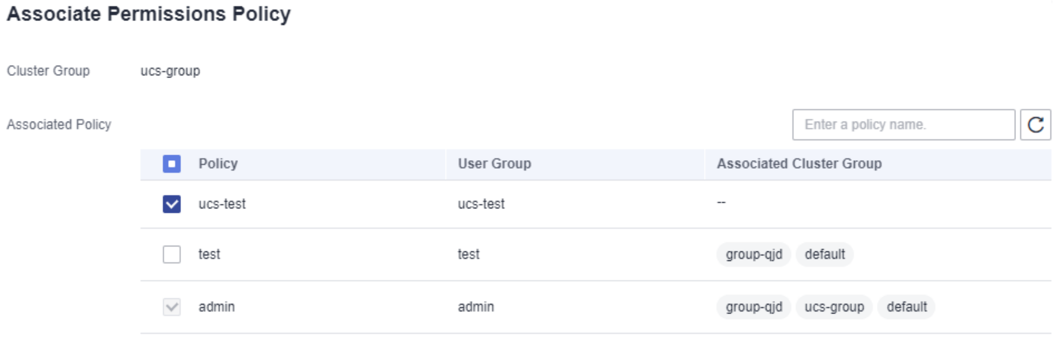

Associate Permissions Policy

Cluster Group ucs-group

Associated Policy

Enter a policy name.

■	Policy	User Group	Associated Cluster Group
☑	ucs-test	ucs-test	--
□	test	test	group-qjd default
☑	admin	admin	group-qjd ucs-group default

Step 4 Click **OK**.

----End

2.3 How Do I Configure Operation Permissions for Cluster Resources?

Background

For CCE clusters, configure the permissions on the IAM console. For non-CCE clusters, configure the permissions on the UCS console.

To manage all your clusters on the UCS console, you need to use a Huawei Cloud account or a user in the admin user group to configure permissions on the **Permissions Policies** page of the UCS console. For details, see **Table 2-3**.

Table 2-3 Operation Permissions on Non-CCE Cluster Resources

Category		Permission Description
Cluster information.		A Huawei Cloud account or a member of the admin user group can associate a user with its target cluster group in Policy Center.
Node-related APIs		The Huawei Cloud account or the admin user group member needs to configure the nodes operation permission for the user in Policy Center.
Workloads	Deployments	The Huawei Cloud account or the admin user group member needs to assign the deployments operation permission in the corresponding namespace to the user in Policy Center.
	StatefulSets	The Huawei Cloud account or the admin user group member needs to assign the operation permission of statefulsets in the corresponding namespace to the user in Policy Center.

Category		Permission Description
	DaemonSets	The Huawei Cloud account or the admin user group member needs to configure the operation permission of daemonsets in the corresponding namespace for the user in Policy Center.
	Normal task	The Huawei Cloud account or the admin user group member needs to configure the operation permission of jobs in the corresponding namespace for the user in Policy Center.
	Scheduled task	The Huawei Cloud account or the admin user group member needs to assign the cronjobs operation permission in the corresponding namespace to the user in Policy Center.
	Pod	The Huawei Cloud account or the admin user group member needs to configure the operation permission of pods in the corresponding namespace for the user in Policy Center.
Networking	Service	The Huawei Cloud account or the admin user group member needs to configure the operation permission of services in the corresponding namespace for the user in Policy Center.
	Ingresses	The Huawei Cloud account or the admin user group member needs to configure the operation permission of ingresses in the corresponding namespace for the user in Policy Center.
Container Storage	PersistentVolumeClaims (PVCs)	The Huawei Cloud account or the admin user group member needs to assign the operation permission of persistentvolumeclaims in the corresponding namespace to the user in Policy Center.
	Volumes	The Huawei Cloud account or the admin user group member needs to assign the operation permission of persistentvolumes in the corresponding namespace to the user in Policy Center.
	Storage Class	The Huawei Cloud account or the admin user group member needs to assign the operation permission of storageclasses in the corresponding namespace to the user in Policy Center.
ConfigMaps and Secrets	Deployment template	The Huawei Cloud account or the admin user group member needs to assign the configmaps operation permission in the corresponding namespace to the user in Policy Center.

Category		Permission Description
	Secret Key	The Huawei Cloud account or the admin user group member needs to configure the operation permission of secrets in the corresponding namespace for the user in Policy Center.
Custom Resource Definitions		The Huawei Cloud account or a member of the admin user group needs to assign the operation permission of customresourcedefinitions in the corresponding namespace to the user in Policy Center.
Namespace		The Huawei Cloud account or the admin user group member needs to assign the namespaces operation permission to the user in Policy Center.
Workload Scaling		The Huawei Cloud account or the admin user group member needs to assign the horizontalpodautoscalers operation permission to the user in Policy Center.

The following table lists the resource operation permissions you can configure on the UCS console:

- *****: Allows all operations.
- **get**: Retrieves a specific resource object by name.
- **list**: Retrieves all resource objects of a specific type in the namespace. You can use selectors to query matched resources.
- **watch**: Watches and responds to resource changes.
- **create**: Creates a resource.
- **update**: Updates a resource.
- **patch**: Partially update a resource.
- **delete**: Deletes a resource.

NOTE

All operations: *

Read-only: get + list + watch

Read-write: get + list + watch + create + update + patch + delete

Prerequisites

An IAM user has been added to a user group. For details, see **User Groups**.

Procedure

This section guides you to configure permissions on the cluster console for IAM users. For details about the permissions on other functions (such as CCE clusters and container cluster federations), see **How Do I Configure the Access Permission for Each Function of the UCS Console?**.

Step 1 Log in to the UCS console as an administrator or a user in the admin user group. In the navigation pane, click **Permissions Policies**.

Step 2 In the upper right corner, click **Create Permissions Policy**.

Step 3 Set permissions policy parameters.

- **Policy Name**: Enter a name, starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed.
- **User Group**: Select the user group associated with the permissions policy. The user groups in the drop-down list are inherited from IAM. If no user group is available, click **Create User Group** to create one on the IAM console.
- **Permissions Template**: Defaults to **Do not use**. You can also select a **default** or a custom template. For details about how to customize a permissions template, see **Adding a Template**.

  When choosing **Do not use**, manually configure the permissions. Click + to add multiple configurations.

  - **Operations to perform**: Select one or multiple operations.
    - *: All operations
    - get: Retrieves a specific resource object by name.
    - list: Retrieves all resource objects of a specific type in the namespace. You can use the selector to query matched resources.
    - watch: used to respond to resource changes.
    - create: creates a resource.
    - update: updates resources.
    - patch: used for partial update of resources.
    - delete: Delete a resource.
  - **Namespace**: Select one or multiple namespaces to operate.
  - **Resources to operate**: Select one or multiple resources to operate. For details about resource types, see **Table 2-3**.
- **Description**: Enter a description of the permissions policy to be added.

Step 4 Click **OK**.

Step 5 Associate the created permissions policy with the cluster group by clicking ⊞ in the cluster group card view on the **Container Clusters** page.

Figure 2-4 Associating a permission policy with a cluster group

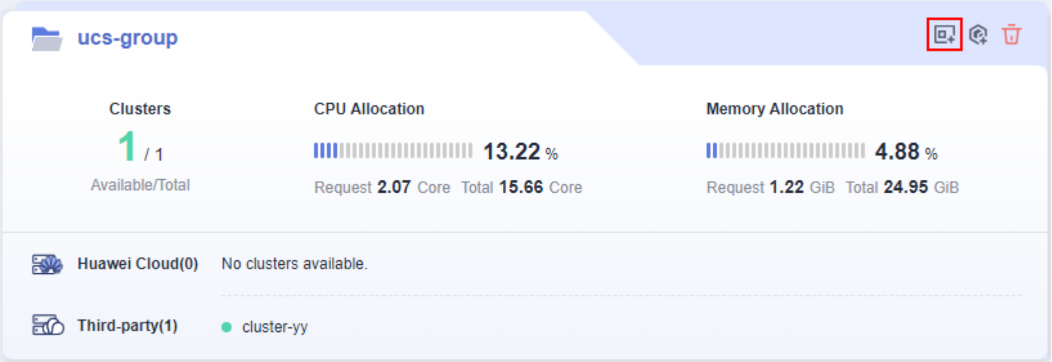

ucs-group

Clusters
1 / 1
Available/Total

CPU Allocation
13.22 %
Request 2.07 Core Total 15.66 Core

Memory Allocation
4.88 %
Request 1.22 GiB Total 24.95 GiB

Huawei Cloud(0) No clusters available.

Third-party(1) ● cluster-yy

Step 6 Select the created policy and click **OK**.

Figure 2-5 Associating policies

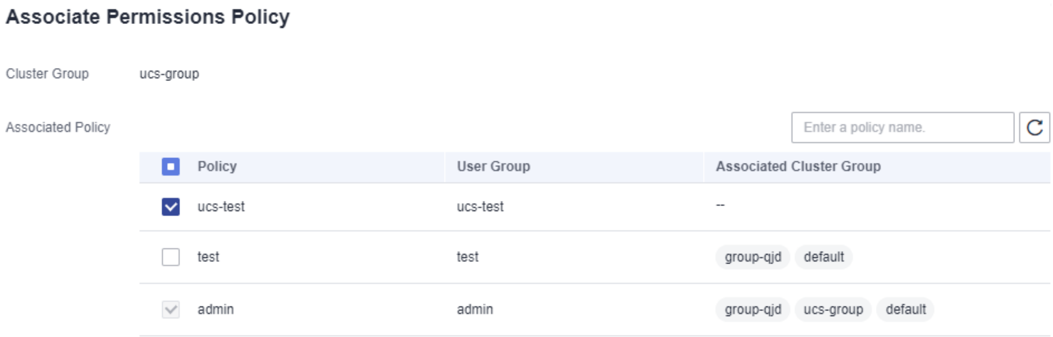

Associate Permissions Policy

Cluster Group ucs-group

Associated Policy

Enter a policy name.

	Policy	User Group	Associated Cluster Group
☑	ucs-test	ucs-test	--
□	test	test	group-qjd default
✓	admin	admin	group-qjd ucs-group default

----End

3 Container Clusters

3.1 Why Does a Cluster Stay in Waiting or Abnormal State When Connecting to UCS?

Background

This section guides you to troubleshoot the exceptions you may encounter when connecting a cluster to UCS:

- You have connected a cluster to UCS and deployed proxy-agent in the cluster, but the console always displays an error message, indicating that the cluster is waiting for connection or fails to get registered after the connection times out.

  NOTE

  If the cluster fails to register, click ⟳ in the upper right corner to register it again and locate the fault as guided in **Troubleshooting**.

- If a connected cluster is unavailable:
  - For Huawei Cloud CCE clusters: Go to the CCE console to check the cluster status. If the cluster is **Unavailable**, rectify the fault by referring to **FAQ documentation**.
  - For self-built or third-party clusters: Refer to **Troubleshooting**.

Troubleshooting

Table 3-1 explains the error messages for you to locate faults.

Table 3-1 Error message description

Error Message	Description	Check Item
"currently no agents available, please make sure the agents are correctly registered"	The proxy-agent in the connected cluster is abnormal or the network is abnormal.	- **Check Item 1: proxy-agent** - **Check Item 2: Network Connection Between the Cluster and UCS**
"please check the health status of kube apiserver: ..."	The kube-apiserver in the cluster cannot be accessed.	- **Check Item 3: kube-apiserver**
"cluster responded with non-successful status code: ..."	Rectify the fault based on the returned status code. For example, status code 401 indicates that the user does not have the access permission. A possible cause is that the cluster authentication information has expired.	- **Check Item 4: Cluster Authentication Information Changes**
"cluster responded with non-successful message: ..."	Rectify the fault based on the returned information. For example, the message "Get "https://172.16.0.143:6443/readyz?timeout=32s\": context deadline exceeded" indicates that the access to the API server times out. A possible cause is that the API server is faulty.	-

Check Item 1: proxy-agent

NOTICE

After the cluster is removed from UCS, the authentication information contained in the original proxy-agent configuration file becomes invalid. You need to delete the proxy-agent pods deployed in the cluster. To connect the cluster to UCS again, download the proxy-agent configuration file from the UCS console again and use it for re-deployment.

Step 1 Log in to the master node of the target cluster.

Step 2 Check the deployment of the cluster agent.

kubectl -n kube-system get pod | grep proxy-agent

Expected output for successful deployment:

```
proxy-agent-*** 1/1 Running 0 9s
```

If proxy-agent is not in the Running state, run the **kubectl -n kube-system describe pod proxy-agent-***** command to view the pod alarms. For details, see **Why Does proxy-agent Fail to Run?**.

NOTE

By default, proxy-agent is deployed with two pods, and can provide services as long as one pod is running properly. However, one pod cannot ensure high availability.

Step 3 Print the pod logs of proxy-agent and check whether the agent program can connect to UCS.

kubectl -n kube-system logs proxy-agent-* | grep "Start serving"**

If no "Start serving" log is printed but the proxy-agent pods are in normal state, check other check items.

----End

Check Item 2: Network Connection Between the Cluster and UCS

For clusters connected through a public network:

Step 1 Check whether a public IP is bound to the cluster or a public NAT gateway is configured.

Step 2 Check whether the outbound traffic of the cluster security group is allowed. To perform access control on the outbound traffic, contact technical support to obtain the destination IP and port number.

Step 3 After rectifying network faults, delete the existing proxy-agent pods to rebuild pods. Check whether the logs of the new pods contain "Start serving".

kubectl -n kube-system logs proxy-agent-* | grep "Start serving"**

Step 4 If desired logs are printed, refresh the UCS console page and check whether the cluster is properly connected.

----End

For clusters connected through a private network:

Step 1 Check whether the outbound traffic of the cluster security group is allowed. To perform access control on the outbound traffic, contact technical support to obtain the destination IP and port number.

Step 2 Rectify the network connection faults between the cluster and UCS, IDC, or third-party clouds.

Refer to the following guides according to your network connection type.

- Direct Connect (DC): **Troubleshooting**
- Virtual Private Network (VPN): **Troubleshooting**

Step 3 Rectify the VPC Endpoint (VPCEP) faults. The VPCEP status must be **Accepted**. If the VPCEP is deleted by mistake, create it again. For details, see **How Do I Restore a Deleted VPC Endpoint for a Cluster Connected Through a Private Network?**.

Figure 3-1 Checking VPCEP status

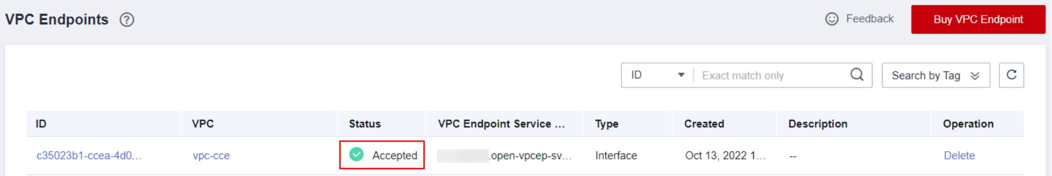

ID	VPC	Status	VPC Endpoint Service ...	Type	Created	Description	Operation
c35023b1-ccea-4d0...	vpc-cce	Accepted	[redacted] .open-vpcep-sv...	Interface	Oct 13, 2022 1...	--	Delete

Step 4 After rectifying network faults, delete the existing proxy-agent pods to rebuild pods. Check whether the logs of the new pods contain "Start serving".

kubectl -n kube-system logs proxy-agent-* | grep "Start serving"**

Step 5 If desired logs are printed, refresh the UCS console page and check whether the cluster is properly connected.

----End

Check Item 3: kube-apiserver

When connecting a cluster to UCS, the error message shown in **Figure 3-2** may be displayed, saying "please check the health status of kube apiserver: ...".

Figure 3-2 Abnormal kube-apiserver

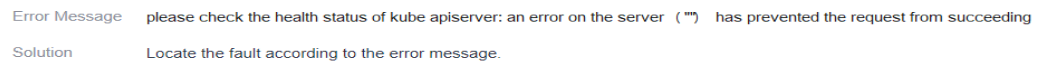

Error Message please check the health status of kube apiserver: an error on the server ("") has prevented the request from succeeding

Solution Locate the fault according to the error message.

This indicates that proxy-agent cannot communicate with the API server in the cluster. Users may have different network configurations for the cluster to connect to UCS. Therefore, UCS does not provide any unified solution for this fault. You need to rectify it on your own and try again.

Step 1 Log in to the UCS console. In the navigation pane, choose **Container Clusters**.

Step 2 Log in to the master node of the cluster and check the API server address.

kubectl get po `kubectl get po -nkube-system | grep kube-apiserver | awk {'print $1'}` -nkube-system -oyaml | grep advertise-address.endpoint

Step 3 Check whether the **clusters.cluster.server** field in the kubeconfig file of the cluster is the same as the API server address of the cluster queried in **Step 2**.

If not, the cluster provider may have converted the API server address. You need to replace the API server address in the kubeconfig file, re-connect the cluster to UCS, and deploy proxy-agent again.

NOTE

If the value of **clusters.cluster.server** in the kubeconfig file is **https://kubernetes.default.svc.cluster.local:443**, you can retain it, which is the local domain name of the Kubernetes Service (ClusterIP of the API server).

Step 4 Check whether the proxy-agent pod can access the API server of the cluster to be connected.

Example command:

```
kubectl exec -ti proxy-agent-*** -n kube-system /bin/bash
# Access kube-apiserver of the cluster.
curl -kv https://*.*.*.*:*/readyz
```

If the access fails, rectify the cluster network fault, re-connect the cluster to UCS, and deploy proxy-agent again.

----End

Check Item 4: Cluster Authentication Information Changes

If "cluster responded with non-successful status: [401][Unauthorized]" is displayed, the cluster authentication information may have expired or changed. As a result, UCS cannot access kube-apiserver. You need to remove the cluster, use a new kubeconfig file to register the cluster again, and re-deploy proxy-agent.

NOTE

- A permanent kubeconfig file can prevent such faults.
- The authentication information will change after you renew a third-party cluster provided by certain vendors. Pay attention to these vendors and try avoiding cluster arrears.

3.2 How Do I Restore a Deleted VPC Endpoint for a Cluster Connected Through a Private Network?

Background

The VPCEP is deleted by mistake in such a cluster, and the cluster becomes abnormal.

Procedure

NOTE

The VPCEP IP has been configured in proxy-agent. Therefore, you need to specify an available IP when creating a VPCEP.

Step 1 Log in to the **VPCEP console** to check whether the VPCEP in the region where UCS is deployed is deleted. If yes, go to the next step.

Step 2 Log in to the master node of the abnormal cluster.

Step 3 Query the IP configured in proxy-agent.

kubectl get deploy -n kube-system proxy-agent -oyaml | grep -A3 hostAliases

Command output:

```
hostAliases:
- hostnames:
  - proxyurl.ucs.myhuaweicloud.com
  ip: 10.0.0.182
```

Step 4 Create a VPCEP in the region where UCS is located, use the preceding IP address, and click **View In-Use IP Address** to ensure that this IP is not used by any other VPCEP. If the IP is in use, change the proxy-agent configuration in the cluster. For details, see **Changing proxy-agent Configurations**.

Figure 3-3 Buying a VPC endpoint (with manually specified private IP)

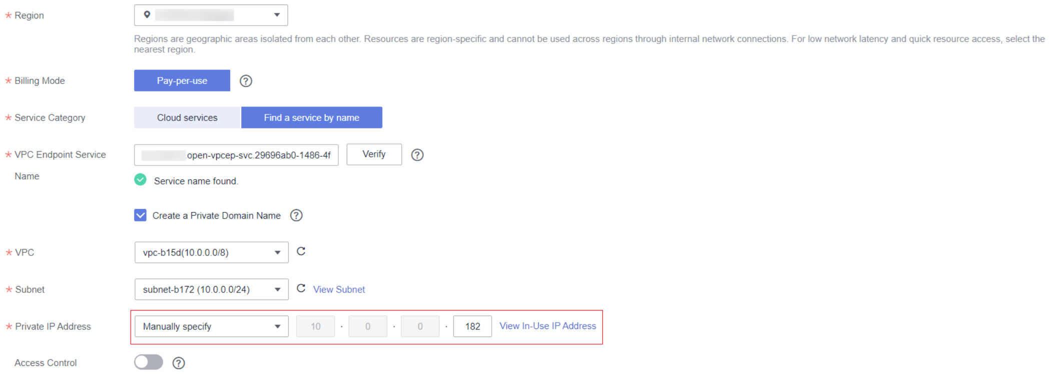

* Region: [redacted]

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

* Billing Mode: Pay-per-use

* Service Category: Cloud services | Find a service by name

* VPC Endpoint Service Name: [redacted]open-vpcep-svc.29696ab0-1486-4f Verify

Service name found.

☑ Create a Private Domain Name

* VPC: vpc-b15d(10.0.0.0/8)

* Subnet: subnet-b172 (10.0.0.0/24) View Subnet

* Private IP Address: Manually specify 10 . 0 . 0 . 182 View In-Use IP Address

Access Control

Step 5 Click **Next** and then **Submit** to create a VPCEP.

Step 6 Wait for 1 to 3 minutes, go back to the UCS console, and refresh the cluster status.

----End

Changing proxy-agent Configurations

Step 1 Create a VPCEP in the region where UCS is located.

Figure 3-4 Buying a VPC endpoint (with automatically assigned private IP)

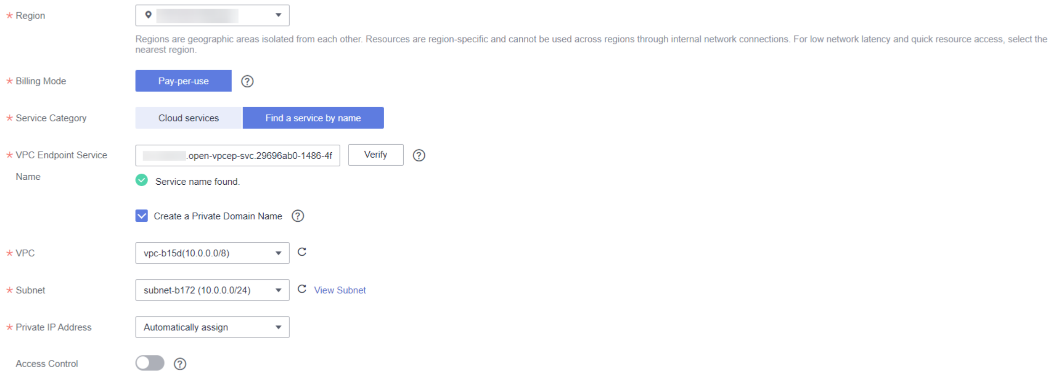

* Region: [redacted]

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

* Billing Mode: Pay-per-use

* Service Category: Cloud services | Find a service by name

* VPC Endpoint Service Name: [redacted].open-vpcep-svc.29696ab0-1486-4f Verify

Service name found.

☑ Create a Private Domain Name

* VPC: vpc-b15d(10.0.0.0/8)

* Subnet: subnet-b172 (10.0.0.0/24) View Subnet

* Private IP Address: Automatically assign

Access Control: ○

Step 2 Click the ID of the newly created VPCEP to view the automatically allocated node IP.

Figure 3-5 Viewing VPCEP details

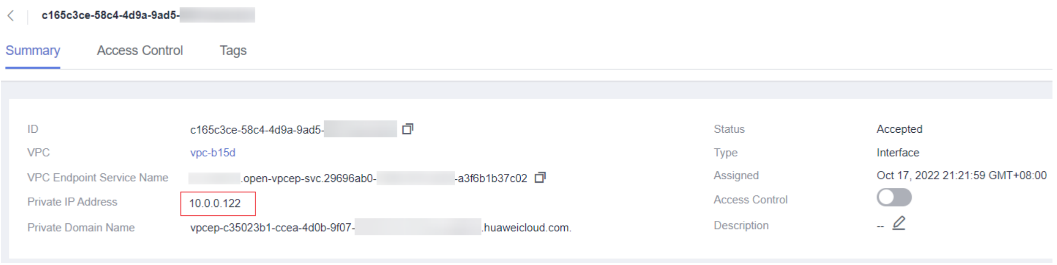

c165c3ce-58c4-4d9a-9ad5-[redacted]

ID: c165c3ce-58c4-4d9a-9ad5-[redacted]

VPC: vpc-b15d

VPC Endpoint Service Name: [redacted].open-vpcep-svc.29696ab0-[redacted]-a3f6b1b37c02

Private IP Address: 10.0.0.122

Private Domain Name: vpcep-c35023b1-ccea-4d0b-9f07-[redacted].huaweicloud.com.

Status: Accepted

Type: Interface

Assigned: Oct 17, 2022 21:21:59 GMT+08:00

Access Control: ○

Description: --

Step 3 Log in to the master node of the abnormal cluster.

Step 4 Change the IP address configured in proxy-agent.

kubectl edit deploy -n kube-system proxy-agent

Change the IP address in the **hostAliases** field.

```
hostAliases:
- hostnames:
  - proxyurl.ucs.myhuaweicloud.com
  ip: 10.0.0.122
```

Press **Esc**, enter **:wq**, and press **Enter**.

Step 5 Wait for 1 to 3 minutes, go back to the UCS console, and refresh the cluster status.

----End

3.3 Why Does proxy-agent Fail to Run?

Background

The deployed proxy-agent is not in the Running state.

Procedure

Step 1 Log in to the master node of the cluster.

Step 2 Check the running status of proxy-agent.

kubectl -n kube-system get pod | grep proxy-agent

The following command output shows the pods are in the ImagePullBackOff and Pending states:

```
proxy-agent-59ddf7597b-rq4j6          0/1     ImagePullBackOff   0          2d16h
proxy-agent-59ddf7597b-sjf55          0/1     Pending            0          2d16h
```

Step 3 Query the details about the pods:

kubectl describe pod proxy-agent-* -nkube-system**

The following errors may occur:

- The Kubernetes event indicates that the cluster cannot pull the proxy-agent image. In this case, check whether the cluster can access the public network to pull the SWR image.

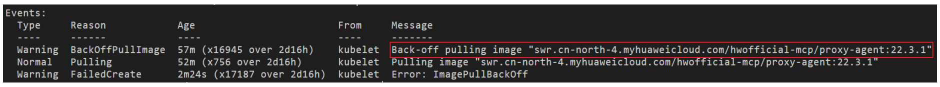

```
Events:
  Type     Reason            Age                      From     Message
  ----     ------            ----                     ----     -------
  Warning  BackOffPullImage  57m (x16945 over 2d16h)  kubelet  Back-off pulling image "swr.cn-north-4.myhuaweicloud.com/hwofficial-mcp/proxy-agent:22.3.1"
  Normal   Pulling           52m (x756 over 2d16h)    kubelet  Pulling image "swr.cn-north-4.myhuaweicloud.com/hwofficial-mcp/proxy-agent:22.3.1"
  Warning  FailedCreate      2m24s (x17187 over 2d16h) kubelet  Error: ImagePullBackOff
```

- The Kubernetes event indicates that the CPU or memory resources of the node are insufficient. In this case, scale up the node.

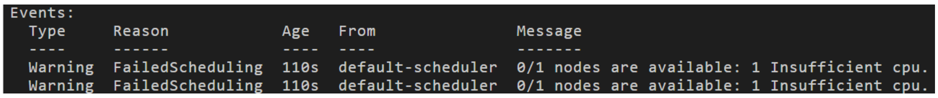

```
Events:
  Type     Reason            Age   From               Message
  ----     ------            ----  ----               -------
  Warning  FailedScheduling  110s  default-scheduler  0/1 nodes are available: 1 Insufficient cpu.
  Warning  FailedScheduling  110s  default-scheduler  0/1 nodes are available: 1 Insufficient cpu.
```

- The Kubernetes event shows that the scheduling failed. To achieve high availability, proxy-agent is deployed with two pods and they are scheduled to different nodes by default. Ensure that your cluster has at least two nodes with sufficient resources.

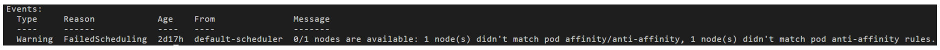

```
Events:
  Type     Reason            Age    From               Message
  ----     ------            ----   ----               -------
  Warning  FailedScheduling  2d17h  default-scheduler  0/1 nodes are available: 1 node(s) didn't match pod affinity/anti-affinity, 1 node(s) didn't match pod anti-affinity rules.
```

Step 4 After the preceding problems are resolved, check the running status of proxy-agent again. All pods should now be in the Running state.

----End

4 Traffic Distribution

4.1 How Do I Add a Third-Party Domain Name?

Background

If you have registered a domain name with a third-party registrar, and you want to use UCS to manage app traffic, you can add the domain name to Domain Name Service (DNS) on Huawei Cloud. The UCS traffic management console automatically obtains the domain name that has been resolved.

Step 1: Add a Domain Name

If your domain name is registered with a third-party registrar, create a public zone and add record sets to it on the DNS console.

1. Log in to the Huawei Cloud management console.
2. Move the cursor to the ≡ icon on the left of the page. In the service list, choose **Networking** > **Domain Name Service**.

   The DNS console is displayed.
3. In the navigation pane, choose **Public Zones** and click **Create Public Zone** in the upper right corner.
4. Set **Domain Name** to your registered domain name, for example, **example.com**.

   For details about the parameters, see **Creating a Public Zone**.

Figure 4-1 Creating a public zone

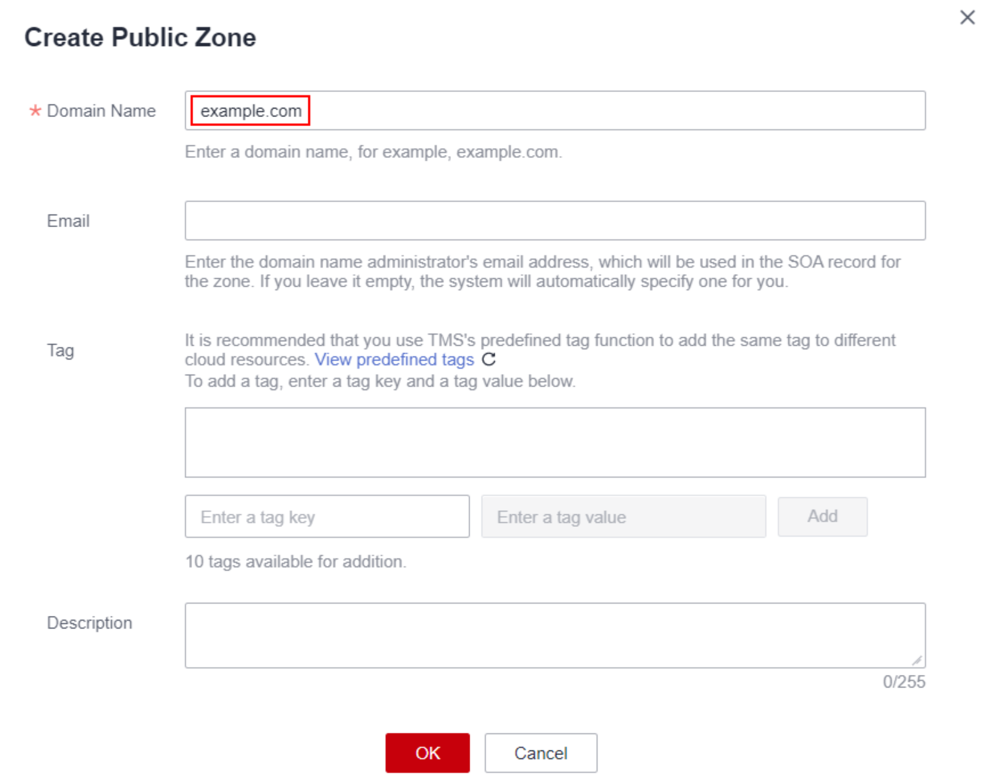

Create Public Zone

* Domain Name: example.com

Enter a domain name, for example, example.com.

Email:

Enter the domain name administrator's email address, which will be used in the SOA record for the zone. If you leave it empty, the system will automatically specify one for you.

Tag: It is recommended that you use TMS's predefined tag function to add the same tag to different cloud resources. View predefined tags

To add a tag, enter a tag key and a tag value below.

10 tags available for addition.

Description:

0/255

5. Click **OK**.

   View the created public zone on the **Public Zones** page.

Figure 4-2 Viewing the public zone

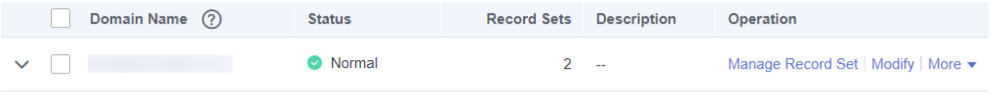

Domain Name	Status	Record Sets	Description	Operation
[redacted]	Normal	2	--	Manage Record Set | Modify | More

If the system displays a message indicating that the public zone has been created by another tenant, handle the issue by referring to **Regaining a Domain Name**.

NOTE

Click the zone name to query detailed zone information. Record sets of the SOA type and NS type have been created in the zone. To be more specific,

- The SOA record set defines the DNS server that is the authoritative information source for a particular domain name.
- The NS record set defines authoritative DNS servers for a zone.

  You can modify the NS record set based on the region of the domain name. For more information about DNS servers, see **What Are DNS Server Addresses Provided by Huawei Cloud DNS?**

Step 2: Change DNS Servers of the Domain Name

The DNS service provides authoritative DNS servers for domain resolution.

After you create a public zone, an NS record set is generated, which specifies the DNS servers provided by the DNS service.

If DNS server addresses of the public zone are not the same as those in the NS record set, the DNS service will not be able to resolve the domain name. You must change the DNS server addresses of the domain name on the registrar's website.

NOTE

Generally, the changes to DNS server addresses take effect within 48 hours, but the time may vary depending on the domain name registrar's cache duration.

Step 1 Query the DNS server addresses of the DNS service.

1. Log in to the Huawei Cloud management console.
2. Move the cursor to the ☰ icon on the left of the page. In the service list, choose **Networking** > **Domain Name Service**.

   The DNS console is displayed.
3. In the navigation pane, choose **Public Zones**.

   The **Public Zones** page is displayed.
4. Click the name of the public zone you created.

   Locate the NS record set. The DNS server addresses provided by the DNS service are displayed under **Value**.

   Figure 4-3 NS record set returned by the system

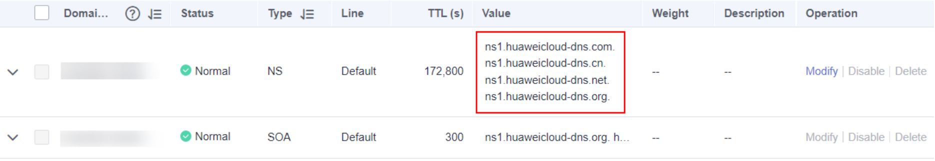

	□	Domai...	Status	Type	Line	TTL (s)	Value	Weight	Description	Operation
⌄	□	[redacted]	✓ Normal	NS	Default	172,800	ns1.huaweicloud-dns.com. ns1.huaweicloud-dns.cn. ns1.huaweicloud-dns.net. ns1.huaweicloud-dns.org.	--	--	Modify | Disable | Delete
⌄	□	[redacted]	✓ Normal	SOA	Default	300	ns1.huaweicloud-dns.org. h...	--	--	Modify | Disable | Delete

Step 2 Change the DNS server addresses of the domain name.

Log in to the domain name registrar website and change the addresses to Huawei Cloud DNS server addresses.

ns1.huaweicloud-dns.com

ns1.huaweicloud-dns.cn

ns1.huaweicloud-dns.net

ns1.huaweicloud-dns.org

For details, see the operation guide on the domain name registrar website.

----End

Step 3: Add a Scheduling Policy on UCS

Step 1 After the DNS record set is added, return to the **Create Traffic Policy** page of the UCS console and select the newly added domain name. If the domain name is not displayed, click ↻ on the right to refresh the drop-down list.

Figure 4-4 Creating a traffic policy

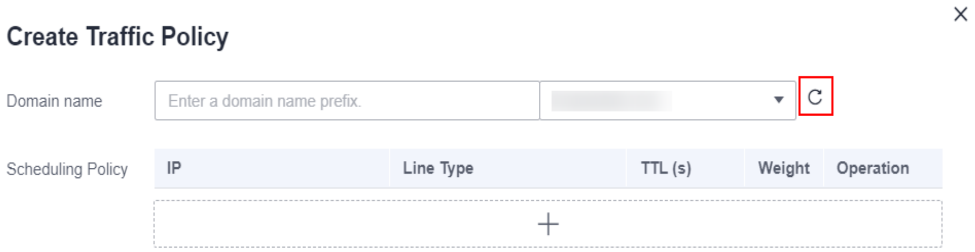

Create Traffic Policy

Domain name: Enter a domain name prefix. [redacted]

Scheduling Policy

IP	Line Type	TTL (s)	Weight	Operation

+

Step 2 Add a policy for the new domain name by referring to **Creating a Traffic Policy**.

Figure 4-5 Scheduling policy

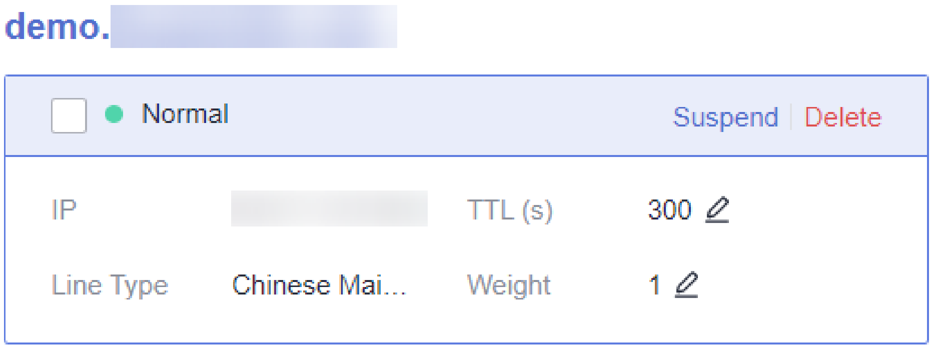

demo.[redacted]

□ ● Normal Suspend | Delete

IP	[redacted]	TTL (s)	300
Line Type	Chinese Mai...	Weight	1

Step 3 Check whether the created scheduling policy takes effect.

Take the Linux operating system as an example. You can run the following command in a CLI tool connected to the Internet:

dig *Target domain name*

NOTE

If your device has not installed dig (Domain Information Groper), install it first. If you are using a CentOS device, run the **yum install bind-utils** command first.

If the following information is displayed and the IP address of **ANSWER SECTION** is the load balancer IP of the target cluster, the scheduling policy takes effect.

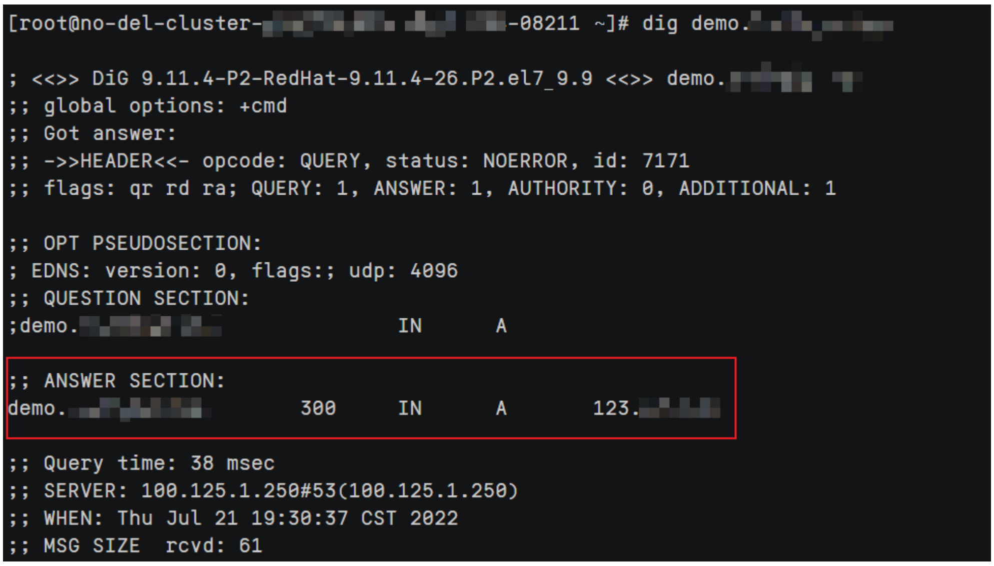

```
[root@no-del-cluster-[redacted]-08211 ~]# dig demo.[redacted]

; <<>> DiG 9.11.4-P2-RedHat-9.11.4-26.P2.el7_9.9 <<>> demo.[redacted]
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 7171
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;demo.[redacted]                IN      A

;; ANSWER SECTION:
demo.[redacted]         300     IN      A       123.[redacted]

;; Query time: 38 msec
;; SERVER: 100.125.1.250#53(100.125.1.250)
;; WHEN: Thu Jul 21 19:30:37 CST 2022
;; MSG SIZE  rcvd: 61
```

----End