Ubiquitous Cloud Native Service

FAQs

Issue 01

Date 2025-07-04





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/

i

Contents

1 About UCS	1
1.1 Are There Quota Restrictions in UCS?	1
2 Billing	3
2.1 How Is UCS Billed?	3
2.2 What Status of a Cluster Will Incur UCS Charges?	4
2.3 Why Am I Still Being Billed After I Purchase a Resource Package?	5
2.4 How Do I Change the Billing Mode of a Cluster from Pay-per-Use to Yearly/Monthly?	5
2.5 What Types of Invoices Are There?	6
2.6 Can I Unsubscribe from or Modify a Resource Package?	6
3 Permissions	7
3.1 How Do I Configure Access Permissions for Each Function of the UCS Console?	7
3.2 Why Can't an IAM User Obtain Cluster or Fleet Information After Logging In to UCS?	10
3.3 How Do I Restore ucs_admin_trust I Deleted or Modified?	13
3.4 What Can I Do If I Can't Add Permissions for a Fleet or Cluster?	14
3.5 How Do I Clear RBAC Resources After a Cluster Is Unregistered?	18
4 Policy Center	19
4.1 What Can I Do If Policy Center Fails to Be Enabled?	
4.2 What Can I Do If Policy Center Cannot Run Normally?	19
4.3 How Do I Clear Policy Center Resources?	20
5 Fleets	. 21
5.1 What Can I Do If Verification Fails When Clusters Join a Federation?	21
5.2 What Can I Do If an Abnormal Cluster Fails to Be Removed from a Federation-enabled Fleet?	23
5.3 What Can I Do If an Nginx Ingress Is in the Unready State After Being Deployed?	23
5.4 What Can I Do If "Error from server (Forbidden)" Is Displayed When I Run the kubectl Command?	·23
5.5 Why Can't a Workload Be Distributed When the Scheduling Policy Is in Lazy Mode?	24
6 Huawei Cloud Clusters	27
6.1 What Can I Do If a Huawei Cloud Cluster Is Unavailable?	27
7 Attached Clusters	. 28
7.1 What Can I Do If an Attached Cluster Fails to Be Connected?	28
7.2 How Do I Restore a Deleted VPC Endpoint for a Cluster Connected Through a Private Network?	32
7.3 What Can I Do If proxy-agent Fails to Be Deployed?	35

8 On-Premises Clusters	37
8.1 What Can I Do If an On-Premises Cluster Fails to Be Connected?	37
8.2 How Do I Manually Clear Nodes of an On-Premises Cluster?	42
8.3 How Do I Downgrade a cgroup?	42
8.4 What Can I Do If the VM SSH Connection Times Out?	43
8.5 How Do I Expand the Disk Capacity of the CIA Add-on in an On-Premises Cluster?	43
8.6 What Can I Do If the Cluster Console Is Unavailable After a Master Node Is Shut Down?	44
8.7 What Can I Do If a Node Is Not Ready After Its Scale-Out?	45
8.8 How Do I Update the CA/TLS Certificate of an On-Premises Cluster?	45
8.9 What Can I Do If an On-Premises Cluster Fails to Be Installed?	47
9 Multi-Cloud Clusters	49
9.1 How Do I Clear Multi-Cloud Cluster Resources?	49
9.2 How Do I Obtain an Access Key (AK/SK)?	51
9.3 How Do I Update the Multi-Cloud Cluster Certificate?	
10 Cluster Federation	55
10.1 What Can I Do If the Pre-upgrade Check of the Cluster Federation Fails?	55
10.2 What Can I Do If a Cluster Fails to Be Added to a Federation?	56
10.3 What Can I Do If Status Verification Fails When Clusters Are Added to a Federation?	56
10.4 What Can I Do If an HPA Created on the Cluster Federation Management Plane Fails to Be Distributed to Member Clusters?	57
10.5 What Can I Do If an MCI Object Fails to Be Created?	
10.6 What Can I Do If I Fail to Access a Service Through MCI?	
10.7 What Can I Do If an MCS Object Fails to Be Created?	
10.8 What Can I Do If an MCS or MCI Instance Fails to Be Deleted?	
11 Traffic Distribution	63
11.1 How Do I Add a Third-Party Domain Name?	63
12 O&M	67
12.1 What Can I Do If Monitoring Fails to Be Enabled for a Cluster Due to Residual Add-on Resource	ces?
12.2 What Can I Do If Monitoring Fails to Be Enabled for a Cluster Due to Policy Interception?	
12.3 How Do I Modify the Collection Configuration of the kube-state-metrics Component?	
12.4 How Do I Handle the Error Displayed When Monitoring Is Enabled for a Cluster?	69

1 About UCS

1.1 Are There Quota Restrictions in UCS?

UCS Quotas

Quotas can limit the number or amount of resources available to users. UCS has quota limits on clusters, fleets, permissions, cluster federations, and CIA instances.

- Cluster quota: specifies the maximum number of clusters connected to UCS.
 This quota consists of the numbers of Huawei Cloud clusters and attached clusters.
- Fleet quota: specifies the maximum number of fleets owned by a user.
- Permission quota: specifies the maximum number of permission policies that a user can create on the **Permissions** page.
- Cluster federation quota: specifies the maximum number of cluster federations that a user can enable.
- CIA instance quota: specifies the maximum number of CIA instances that a user can create.

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 (ELB), 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 default values. You can also request a quota increase.

□ NOTE

You cannot request for increasing the quotas of cluster federations and CIA instances.

Table 1-1 UCS quota items

Quota Item	Default Value
Cluster	50
Fleet	50
Permission policy	50
Cluster federation	1
CIA instance	1

How Do I Increase My UCS Quotas?

Contact our technical support to increase UCS quotas.

- **Step 1** Log in to the management console.
- Step 2 In the upper right corner, choose More > Service Tickets > Create Service Ticket.
 The Create Service Ticket page is displayed.
- **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

 $\mathbf{2}$ Billing

2.1 How Is UCS Billed?

Billing Modes

There are yearly/monthly and pay-per-use billing modes to meet your requirements.

- Yearly/Monthly is a prepaid billing mode. You pay in advance for a subscription term, and in exchange, you get a discounted rate. The longer the subscription term, the bigger the discount. Yearly/Monthly billing is a good option for long-term, stable clusters.
- Pay-per-use is a postpaid billing. You pay as you go and just pay for what you
 use. UCS resource usage is calculated by the second but billed every hour. This
 mode allows you to flexibly adjust resource usage. You neither need to
 prepare for resources in advance, nor end up with excessive or insufficient
 preset resources. Pay-per-use billing is a good option for small-scale clusters.

Billed Items

You will be billed for clusters managed by UCS. For details, see Table 2-1.

Table 2-1 Billed items

Billed Item	Description	Billing Mode	Formula
Clust ers mana ged by UCS	 The cluster management cost depends on the cluster type (including Huawei Cloud clusters, on-premises clusters, attached clusters, multicloud clusters, and partner cloud clusters), cluster specification (vCPUs), and required duration. The cluster management cost does not include the price of any related resources (such as compute nodes and other network services). 	Yearly/ Monthl y or pay- per-use	Cluster specificati on (vCPUs) × Unit price of the cluster specificati on × Required duration For details, see UCS Pricing Details.

2.2 What Status of a Cluster Will Incur UCS Charges?

Cluster status changes affect the number of vCPUs obtained by UCS, which affects UCS billing. If a cluster needs to connect to UCS, ensure that the cluster is running normally. If the cluster is no longer needed, unregister it in a timely manner to avoid further expenditures.

For details about the cluster status and billing, see Table 2-2.

Table 2-2 Cluster status and billing

Clust er Statu s	Billed
Runni ng	Yes
Unav ailabl e	Yes NOTE After a cluster is connected to UCS, UCS obtains and records the the number of vCPUs. If the cluster becomes unavailable, UCS cannot obtain the number of vCPUs in real time. In this case, you will be billed based on the last recorded number of vCPUs.
Waiti ng for acces s	No

Clust er Statu s	Billed
Regis tratio n timeo ut	No
Unre gister ing	No
Unre gistra tion failed	No

2.3 Why Am I Still Being Billed After I Purchase a Resource Package?

Table 2-3 lists the possible causes. To avoid arrears, you can choose a suitable resource package or ensure that the account balance is sufficient.

Table 2-3 Troubleshooting

Possible Cause	Troubleshooting
The cluster type in the purchased package is inconsistent with the actual one.	Purchase the package according to the cluster type.
The vCPUs in the purchased package are fewer than the actual ones.	Purchase a package that meets the cluster scale requirements or ensure that your account has sufficient balance.

2.4 How Do I Change the Billing Mode of a Cluster from Pay-per-Use to Yearly/Monthly?

UCS supports two billing modes: pay-per-use and yearly/monthly. If you want to use UCS resources at discounted prices, you only need to purchase a package based on the cluster type and scale.

2.5 What Types of Invoices Are There?

Huawei Cloud can issue invoices by billing cycle and by order.

You can request an invoice on the Invoices page.

2.6 Can I Unsubscribe from or Modify a Resource Package?

Purchased packages cannot be unsubscribed from or modified.

3 Permissions

3.1 How Do I Configure Access Permissions for Each Function of the UCS Console?

Symptom

The functions of the UCS console are controlled by IAM. When an unauthorized user accesses a page on the UCS console, an error message is displayed, indicating that the user does not have the access permissions or permission authentication fails.

Solution

The administrator needs to grant users the permissions for using functions of the UCS console. IAM system policies (including UCS FullAccess, UCS CommonOperations, UCS CIAOperations, and UCS ReadOnlyAccess) are used to define user permissions.

Table 3-1 System-defined permissions for UCS

Policy Name	Description	Туре
UCS FullAccess	Administrator permissions for UCS. Users with these permissions can perform all operations on UCS resources, for example, creating permission policies and security policies.	System-defined policy
UCS CommonOperatio ns	Common operation permissions for UCS. Users with these permissions can create workloads, distribute traffic, and perform other operations.	System-defined policy
UCS CIAOperations	Administrator permissions for Container Intelligent Analysis.	System-defined policy

Policy Name	Description	Туре
UCS ReadOnlyAccess	Read-only permissions for UCS (except for Container Intelligent Analysis)	System-defined policy

Services on Huawei Cloud are interdependent, so UCS depends on other cloud services to implement some functions such as image repository and domain name resolution. The preceding system policies are often used together with roles or policies of other cloud services for refined authorization. When granting permissions to IAM users, the administrator must comply with the principle of least privilege. Table 3-2 lists the least-privilege permissions required by the Admin, Operator, and Viewer roles to use each UCS function.

□ NOTE

For details about how to grant IAM system policies and UCS RBAC permissions to users, see UCS Resource Permissions and Kubernetes Resource Permissions in a Cluster, respectively.

Table 3-2 Least-privilege permissions required by each UCS function

Functio n	Permis sion Type	Permissions	Least-Privilege Permissions
Fleets	Adminis trator	 Creating and deleting a fleet Registering a Huawei Cloud cluster (CCE standard cluster or CCE Turbo cluster), an onpremises cluster, or an attached cluster Unregistering a cluster 	UCS FullAccess
		 Adding a cluster to or removing a cluster from a fleet Adding permissions for a cluster 	
		 or fleet Enabling cluster federation and performing federation management operations (such as creating a workload and creating a DNS policy) 	
	Viewer	Querying the list or details of clusters or fleets	UCS ReadOnlyAccess
Huawei Cloud clusters	Adminis trator	Read-write permissions on Huawei Cloud clusters and all Kubernetes resource objects (such as nodes, workloads, jobs, and Services)	UCS FullAccess + CCE Administrator

Functio n	Permis sion Type	Permissions	Least-Privilege Permissions
	Develo per	Read-write permissions on Huawei Cloud clusters and most Kubernetes resource objects and read-only permissions on Kubernetes resource objects such as namespaces and resource quotas	UCS CommonOperations + CCE Administrator
	Viewer	Read-only permissions on Huawei Cloud clusters and all Kubernetes resource objects (such as nodes, workloads, jobs, and Services)	UCS ReadOnlyAccess + CCE Administrator
On- premise s/ Attache d/	Adminis trator	Read-write permissions on on- premises/attached/multi-cloud clusters and all Kubernetes resource objects (such as nodes, workloads, jobs, and Services)	UCS FullAccess
Multi- cloud clusters	Develo per	Read-write permissions on on- premises/attached/multi-cloud clusters and most Kubernetes resource objects and read-only permissions on Kubernetes resource objects such as namespaces and resource quotas	UCS CommonOperations + UCS RBAC (The list permission for namespaces is required.)
	Viewer	Read-only permissions on on- premises/attached/multi-cloud clusters and all Kubernetes resource objects (such as nodes, workloads, jobs, and Services)	UCS ReadOnlyAccess + UCS RBAC (The list permission for namespaces is required.)
Image Reposit ories	Adminis trator	All permissions on SoftWare Repository for Container (SWR), including creating organizations, pushing images, viewing the image list or details, and pulling images	SWR Administrator
Permiss ions	Adminis trator	 Creating and deleting permissions Viewing the permission list or details NOTE When creating permissions, you need to grant the IAM ReadOnlyAccess permissions (read-only permissions on IAM) to IAM users to obtain the IAM user list. 	UCS FullAccess + IAM ReadOnlyAccess

Functio n	Permis sion Type	Permissions	Least-Privilege Permissions
	Viewer	Viewing the permission list or details	UCS ReadOnlyAccess + IAM ReadOnlyAccess
Policy Center	Adminis trator	 Enabling Policy Center Creating and disabling a policy instance Viewing the policy list Viewing policy implementation details 	UCS FullAccess
	Viewer	Viewing the policy list and policy implementation details of fleets and clusters with Policy Center enabled	UCS CommonOperations or UCS ReadOnlyAccess
Traffic Distribu tion	Adminis trator	Creating a traffic policy, suspending and deleting a scheduling policy, and performing other operations	UCS CommonOperatio ns + DNS Administrator (recommended) UCS FullAccess + DNS Administrator
	Viewer	Viewing the traffic policy list or details	UCS ReadOnlyAccess + DNS Administrator
Contain er Intellig ent Analysi s	Adminis trator	 Connecting clusters or canceling cluster connection Viewing the monitoring data of infrastructures, workloads, and other resources 	UCS CIAOperations

3.2 Why Can't an IAM User Obtain Cluster or Fleet Information After Logging In to UCS?

Symptom

After an IAM user logs in to the UCS console and goes to the **Fleets** page, information about the created fleet and registered clusters cannot be obtained. (Both the **Fleets** and **Clusters Not in Fleet** tabs are empty.)

Solution

Most IAM users cannot obtain cluster information because their permissions are not set or incorrectly set. To obtain cluster information, IAM users must have both the UCS system policy permissions and cluster resource object operation permissions. You need to contact the administrator to grant you permissions according to the process shown in Figure 3-1.

IAM console Administrator User group A IAM permissions User group B Users A', A", ... Permissions to operate Permissions Permissions Kubernetes resources Only for on-premises or attached clusters Fleet A Fleet B Fleet C UCS console

Figure 3-1 Granting permissions

Step 1 Log in to the IAM console as the administrator and grant the UCS system policy permissions to the user group of the IAM user.

Select the system policy to be granted based on the operation scope. For example, to query clusters and fleets or their details, or query cluster resource objects (including nodes, workloads, jobs, and Services), you only need to grant the **UCS ReadOnlyAccess** permissions, as shown in **Figure 3-2**.

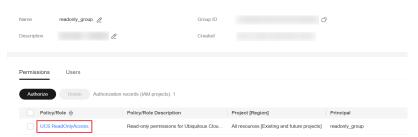


Figure 3-2 Read-only permissions of a user group

Cluster and Fleet Permissions shows the least-privilege permissions required by different permission types. The administrator can grant permissions according to the table.

Step 2 Log in to the UCS console as the administrator and grant the IAM user the permissions for performing operations on cluster resource objects.

The procedure is as follows:

■ NOTE

Permissions on the UCS console take effect only for on-premises or attached clusters. To perform operations on Huawei Cloud cluster resources, grant the CCE Administrator permissions.

- 1. Create a permission policy on the **Permissions** page. (Select the **Viewer** type, which applies to all cluster resource objects.)
- 2. Associate the created permission policy with the fleet or clusters not in the fleet.

----End

Cluster and Fleet Permissions

Functio n	Permis sion Type	Permissions	Least-Privilege Permissions
Fleets	Admin	 Creating and deleting a fleet Registering a Huawei Cloud Cluster (CCE cluster or CCE Turbo cluster), on-premises cluster, or attached cluster Unregistering a cluster Adding a cluster to or removing a cluster from a fleet Associating permission policies with a cluster or fleet Enabling cluster federation and performing federation management operations (such as creating a federated workload and creating a DNS policy) 	UCS FullAccess
	Viewer	Querying clusters and fleets or their details	UCS ReadOnlyAccess
Huawei Cloud clusters	Admin	Read-write permissions on Huawei Cloud clusters and all cluster resource objects (including nodes, workloads, jobs, and Services)	UCS FullAccess + CCE Administrator

Functio n	Permis sion Type	Permissions	Least-Privilege Permissions
	Develo per	Read-write permissions on Huawei Cloud clusters and most cluster resource objects and read-only permissions on Kubernetes resource objects such as namespaces and resource quotas	UCS CommonOperations + CCE Administrator
	Viewer	Read-only permissions on Huawei Cloud clusters and all cluster resource objects (including nodes, workloads, jobs, and Services)	UCS ReadOnlyAccess + CCE Administrator
On- premise s/ Attache d clusters	Admin	Read-write permissions on on- premises/attached clusters and all cluster resource objects (including nodes, workloads, jobs, and Services)	UCS FullAccess
	Develo per	Read-write permissions on on- premises/attached clusters and most cluster resource objects and read-only permissions on Kubernetes resource objects such as namespaces and resource quotas	UCS CommonOperations + UCS RBAC (The list permission for namespaces is required.)
	Viewer	Read-only permissions on on- premises/attached clusters and all cluster resource objects (including nodes, workloads, jobs, and Services)	UCS ReadOnlyAccess + UCS RBAC (The list permission for namespaces is required.)

3.3 How Do I Restore ucs_admin_trust I Deleted or Modified?

Symptom

The ucs_admin_trust agency is created when the administrator logs in to the UCS console for the first time and authorizes the access to CCE. Deleting or modifying this agency (for example, modifying the agency account op_svc_ucs or deleting the Tenant Administrator role) will cause UCS exceptions. For example, a fleet and clusters in that fleet cannot be displayed on the Fleets page.

This section describes how you can restore the ucs_admin_trust agency.

Procedure

- **Step 1** Log in to the IAM console as an administrator.
- **Step 2** In the navigation pane, choose **Agencies**.
- **Step 3** Select **ucs_admin_trust** and click **Delete** in the **Operation** column. In the displayed dialog box, click **OK**.
- **Step 4** In the navigation pane, choose **Agencies**.

Skip this step if the **ucs_admin_trust** agency has been deleted. For other misoperations (for example, you delete the **Tenant Administrator** role accidentally), you need to delete the agency so that you can create an agency.

Step 5 Access the UCS console again. In the dialog box requesting your authorization, click **OK**. UCS will re-create the **ucs_admin_trust** agency to restore your services.

----End

3.4 What Can I Do If I Can't Add Permissions for a Fleet or Cluster?

Symptom

When you add permissions for a fleet or a cluster not in a fleet, the permissions may fail to be added due to cluster connection exceptions. If this happens, an event will be displayed in the **Add Permissions** window. Locate and rectify the fault and add permissions again.

Troubleshooting

If there are exceptions when permissions are added for a fleet or cluster, locate faults based on error messages listed in **Table 3-3**.

Table 3-3 Error messages

Error Message	Description	Check Item
ClusterRole failed reason:Get \"https:// kubernetes.default.svc.cluster.lo cal/apis/ rbac.authorization.k8s.io/v1/ clusterroles/XXXXXXX? timeout=30s\": Precondition Required" Or Get ClusterRole failed reason:an error on the server (\"unknown\") has prevented the request from succeeding (get clusterroles.rbac.authorization.k	The cluster is not connected, 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
8s.io Unauthorized	Rectify the fault based on the status code. For example, status code 401 indicates that the user does not have the access permissions. A possible cause is that the cluster authentication information has expired.	Check Item 3: Cluster Authentication Information Changes
Get cluster namespace[x] failed. Or Reason:namespace "x" not found.	There is no corresponding namespace in the cluster.	Create a namespace in the cluster and try again. Run the following command to create a namespace: kubectl create namespace ns_name If the namespace is not required, ignore this exception.

Check Item 1: proxy-agent

NOTICE

After a cluster is unregistered 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 a master node in the cluster.
- **Step 2** Check the deployment of proxy-agent.

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

Desired 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 **What Can I Do If proxy-agent Fails to Be Deployed?**.

□ NOTE

By default, proxy-agent is deployed with two pods. It can provide services as long as one pod is running normally. 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 working, check other items.

----End

Check Item 2: Network Connection Between the Cluster and UCS

Public Network Access

- **Step 1** Check whether a public IP address is bound to the cluster or a public NAT gateway is configured.
- **Step 2** Check whether the cluster security group allows outbound traffic. To perform access control on the outbound traffic, contact technical support to obtain the destination and port number.
- **Step 3** After rectifying network faults, delete the existing proxy-agent pods and 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 connected.

----End

Private Network Access

- **Step 1** Check whether the cluster security group allows outbound traffic. To perform access control on the outbound traffic, contact technical support to obtain the destination and port number.
- Step 2 Rectify the network connection faults between the cluster and UCS or IDC.

Refer to the following guides based on your network connection type:

- Direct Connect: Troubleshooting
- Virtual Private Network (VPN): Troubleshooting
- Step 3 Rectify the VPC endpoint fault. The VPC endpoint status must be Accepted. If the VPC endpoint is deleted accidently, create another one. For details, see How Do I Restore a Deleted VPC Endpoint for a Cluster Connected Through a Private Network?.

Figure 3-3 Checking the VPC endpoint status



Step 4 Delete the existing proxy-agent pods and 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 connected.

----End

Check Item 3: Cluster Authentication Information Changes

If the error message "cluster responded with non-successful status: [401] [Unauthorized]" is displayed, the IAM network connection may be faulty, according to /var/paas/sys/log/kubernetes/auth-server.log of the three master nodes in the cluster. Ensure that the IAM domain name resolution and the IAM service connectivity are normal.

The common logs are as follows:

- Failed to authenticate token: *******: dial tcp: lookup iam.myhuaweicloud.com on *.*.*.*:53: no such host
 - This log indicates that the node is not capable of resolving **iam.myhuaweicloud.com**. Configure the corresponding domain name resolution by referring to **Preparing for Installation**.
- Failed to authenticate token: Get *******: dial tcp *.*.*.*:443: i/o timeout

- This log indicates that the node's access to IAM times out. Ensure that the node can communicate with IAM normally.
- currently only supports Agency token
 - This log indicates that the request is not initiated by UCS. Currently, onpremises clusters can only be connected to UCS using IAM tokens.
- IAM assumed user has no authorization/iam assumed user should allowed by TEAdmin
 - This log indicates that the connection between UCS and the cluster is abnormal. Contact Huawei technical support.
- Failed to authenticate token: token expired, please acquire a new token
 This log indicates that the token has expired. Run the date command to check
 whether the time difference is too large. If yes, synchronize the time and
 check whether the cluster is working. If the fault persists for a long time, you
 may need to reinstall the cluster. In this case, contact Huawei technical
 support.

After the preceding problem is resolved, run the **crictl ps | grep auth | awk '{print \$1}' | xargs crictl stop** command to restart the **auth-server** container.

3.5 How Do I Clear RBAC Resources After a Cluster Is Unregistered?

After a cluster is unregistered from UCS, some residual RBAC resources may exist. You can clear these resources as follows:

RBAC resources created by UCS contain the label **ucs.rbac.policy=true**. You can use this label to guery and delete the created RBAC resources.

Examples:

4 Policy Center

4.1 What Can I Do If Policy Center Fails to Be Enabled?

If Policy Center fails to be enabled, rectify the fault as follows:

- If "wait for plugins status become health time out" is displayed, check whether the cluster is running normally and whether cluster resources are sufficient. If yes, click **Try again**.
- If the cluster list page or policy instance page displays a message indicating that Policy Center fails to be enabled, take the following steps to rectify the fault:
 - a. Go to the cluster list page, switch to the workload page of the cluster, and check whether the policy management add-on in the gatekeepersystem namespace is running normally.
 - b. If the add-on is not running normally, locate the cause based on workload events.

The following figure shows the details.



• If other error information is displayed, click **Try again**.

4.2 What Can I Do If Policy Center Cannot Run Normally?

The Gatekeeper add-ons are damaged or deleted. Rectify the fault as follows:

• Check whether the **gatekeeper-controller-manager** and **gatekeeper-audit**Deployments in the **gatekeeper-system** namespace are ready. If not, locate the cause.

• If the two Deployments are deleted, you can disable Policy Center and then enable it again.

4.3 How Do I Clear Policy Center Resources?

For a cluster with Policy Center enabled, residual resources may exist in the following scenarios:

- Policy Center is disabled when the cluster is disconnected.
- The cluster connection is interrupted when Policy Center is being disabled.
- The cluster is unregistered after the cluster connection is interrupted.
- The cluster is removed from the fleet after the cluster connection is interrupted.

Run the following command to clear residual resources:

kubectl delete namespace gatekeeper-system

5 Fleets

5.1 What Can I Do If Verification Fails When Clusters Join a Federation?

Context

After cluster federation is enabled for a fleet, existing clusters and clusters newly added to the fleet will automatically join the federation. In this process, the fleet verifies the network status, cluster version, **ClusterRole**, and **ClusterRoleBinding** of the cluster. If the verification fails, clusters cannot join the federation. After the fault is rectified, click **Retry** to make the clusters join the federation again.

Symptom 1: A Message Is Displayed Indicating that ClusterRole and ClusterRoleBinding Already Exist

Cause: A cluster cannot join two or more federations at the same time. If this error message is displayed, the cluster has joined the federation, or joined the federation but has residual resources.

Solution: Manually clear residual resources.

Procedure:

- **Step 1** Obtain the kubeconfig file of the faulty cluster, prepare kubectl and the node running kubectl, and save the kubeconfig file to the **/tmp** directory on the node.
- **Step 2** Run the following commands to clear residual resources:

alias kubectl='kubectl --kubeconfig=/tmp/kubeconfig'

kubectl delete clusterrolebinding `kubectl get clusterrolebinding |grep karmada-controller-manager | awk '{print \$1}'`

kubectl delete clusterrole `kubectl get clusterrole |grep karmada-controller-manager | awk '{print \$1}'`

kubectl delete namespace `kubectl get namespace |egrep 'karmada-[0-9a-f] {8}-([0-9a-f]{4}-){3}[0-9a-f]{12}' |awk '{print \$1}'`

----End

Symptom 2: A Message Is Displayed Indicating that an EIP Needs to Be Bound to the CCE Cluster

Cause: After cluster federation is enabled for a fleet, an EIP needs to be used to solve the network connection problem when the CCE cluster is accessed.

Solution: Bind an EIP to the CCE cluster.

Symptom 3: An EIP Has Been Bound to a CCE Cluster, but the Cluster Still Fails to Be Added to a Federation. "network in cluster is unstable, please retry it later" Is Displayed

Cause: The federation needs to access the CCE cluster over port 5443. The inbound rule of the security group on the control plane of the CCE cluster specifies that 94.74.86.108 (source address) is denied to access the CCE cluster over port 5443.

Solution: Modify the inbound rule of the security group on the control plane of the CCE cluster to allow 94.74.86.108 (source address) to access the CCE cluster over port 5443.

Symptom 4: Cluster That Has Been Added to a Federation Is Abnormal. "cluster is not reachable" Is Displayed

Run the following command in the corresponding member cluster to check whether ServiceAccount exists. Replace *{cluster_name}* with the name of the member cluster.

kubectl get sa -Algrep karmada-{cluster name}.clusterspace.{cluster name}

If the command output indicates that ServiceAccount does not exist, remove the member cluster from the fleet and add this cluster to the fleet again.

Symptom 5: "the same cluster has been registered with name xx" or "cluster xxx is joined successfully" Is Displayed When a Cluster Is Added to a Federation

Cause: The cluster stops unexpectedly when being added to the federation.

Solution: Remove the cluster from the federation and add it to the federation again.

5.2 What Can I Do If an Abnormal Cluster Fails to Be Removed from a Federation-enabled Fleet?

Context

Cluster federation has been enabled for a fleet, but the abnormal cluster cannot be removed from the fleet.

Solution

- **Step 1** Click in the upper right corner of the cluster card again to remove it from the fleet.
- **Step 2** If the fault persists, submit a service ticket and contact technical support.

----End

5.3 What Can I Do If an Nginx Ingress Is in the Unready State After Being Deployed?

Context

The Nginx Ingress is in the unready state after being deployed.

Solution

Before creating an Nginx Ingress, install the NGINX Ingress Controller add-on in the cluster. If the add-on is not installed, the Nginx Ingress is in the unready state.

- For details about how to install the add-on in the CCE cluster, see **Creating** an Nginx Ingress on the Console.
- For details about how to install the add-on for other types of clusters, see NGINX Ingress Controller.

5.4 What Can I Do If "Error from server (Forbidden)" Is Displayed When I Run the kubectl Command?

Symptom

When you use the cluster federation and run the kubectl command, the following information is displayed.



Possible Cause

The resource object ClusterRole or ClusterRoleBinding is deleted. If this occurs in one or more member clusters in a federation, the kubectl command request is interrupted and the error is returned.

Solution

Recreate ClusterRole or ClusterRoleBinding.

The following is an example YAML file of ClusterRole. Replace *{clusterName}* with the name of the member cluster.

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
name: karmada-controller-manager:karmada-{clusterName}
rules:
- apiGroups:
_ '*'
resources:
_ '*'
verbs:
_ '*'
- nonResourceURLs:
_ '*'
verbs:
_ '*'
verbs:
_ get
```

The following is an example YAML file of ClusterRoleBinding. Replace {clusterName} with the name of the member cluster and {karmada-manage-namespace} with the name of the namespace managed by Karmada. You can run the **kubectl get ns|grep karmada** command to obtain the namespace name.

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
name: karmada-controller-manager:karmada-{clusterName}
roleRef:
apiGroup: rbac.authorization.k8s.io
kind: ClusterRole
name: karmada-controller-manager:karmada-{clusterName}
subjects:
- kind: ServiceAccount
name: karmada-{clusterName}
namespace: {karmada-manage-namespace}
```

5.5 Why Can't a Workload Be Distributed When the Scheduling Policy Is in Lazy Mode?

Symptom

- 1. After a workload is created using kubectl, you have created a scheduling policy (PropagationPolicy or ClusterPropagationPolicy) that will be applied later (in lazy mode), but the workload is not distributed due to the lazy mode.
- 2. After the scheduling policy created in 1 is deleted and a scheduling policy (PropagationPolicy or ClusterPropagationPolicy) that will be applied immediately (not in lazy mode) is created, the workload is still not distributed as expected.



Possible Cause

After a workload is created, it enters Karmada's waiting queue. After a lazy-mode scheduling policy is created, the workload is removed from the waiting queue and bound to the scheduling policy. However, no ResourceBinding is generated. In Karmada, when a scheduling policy is deleted, the ResourceBinding is queried first, and then the workload. As a result, the workload cannot be deleted or enter the waiting queue. This means the lazy-mode scheduling policy cannot be completely deleted. When a non-lazy-mode scheduling policy is created, the workload is not in the waiting queue. In this case, the workload cannot be bound to the new scheduling policy immediately. The workload can be bound to the new scheduling policy and distributed only when it changes.

Solution

Modify the parameters in the YAML file of the workload to change the workload and trigger the scheduling. For example, you can add labels and annotations to the workload and redeploy the workload to trigger scheduling.



Operation Example

Step 1 Create a scheduling policy that will be applied effect later.

```
[root@cjl-test-79779 home]# kubectl apply -f test.yaml
propagationpolicy.policy.karmada.io/ttt-deployment created
```

Step 2 Create a workload.

```
[root@cjl-test-79779 home]# kubectl apply -f deploy.yaml deployment.apps/ttt created [root@cil-test-79779 home]# |
```

Step 3 Verify that the workload has been distributed.

```
[root@cjl-test-79779 home]# kubectl get deploy ttt

NAME READY UP-TO-DATE AVAILABLE AGE
ttt 2/2 2 2 3m33s
[root@cjl-test-79779 home]#
```

----End

6 Huawei Cloud Clusters

6.1 What Can I Do If a Huawei Cloud Cluster Is Unavailable?

Symptom

The error message "cce cluster not found, please unregister cluster" is displayed, indicating that the Huawei Cloud cluster is unavailable.

Possible Cause

You have manually deleted a cluster registered with UCS on the CCE cluster console. As a result, the cluster in UCS is unavailable.

Solution

Log in to the UCS console in a timely manner and click \overline{U} in the upper right corner of the cluster card to unregister the cluster to stop its charging.

Attached Clusters

7.1 What Can I Do If an Attached Cluster Fails to Be Connected?

Symptom

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

 You have registered 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 registration fails, click in the upper right corner of the cluster card to register it again and locate the fault as guided in **Troubleshooting**.

• If the status of a connected cluster is unavailable, rectify the fault by referring to **Troubleshooting** in this section.

Troubleshooting

Table 7-1 explains the error messages for you to locate faults.

Table 7-1 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 permissions. A possible cause is that the cluster authentication information has expired.	Check Item 4: Cluster Authenticati on 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.	-
"Current cluster version is not supported in UCS service."	This error occurs because the cluster version does not meet requirements. The version of the Kubernetes cluster connected to UCS must be 1.19 or later.	-

Check Item 1: proxy-agent

NOTICE

After the cluster is unregistered 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 a master node of the destination 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 **What Can I Do If proxy-agent Fails to Be Deployed?**.

■ 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 working, check other items.

----End

Check Item 2: Network Connection Between the Cluster and UCS

Public network access

- **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

Private network access

- **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: Troubleshooting
- Virtual Private Network (VPN): Troubleshooting
- Step 3 Rectify the VPC endpoint fault. The VPC endpoint status must be Accepted. If the VPC endpoint is deleted by mistake, create one again. For details, see How Do I Restore a Deleted VPC Endpoint for a Cluster Connected Through a Private Network?.

Figure 7-1 Checking the VPC endpoint status



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"

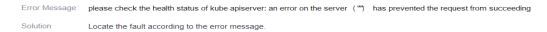
Step 5 If desired logs are printed, refresh the UCS console page and check whether the cluster is properly connected.

----End

Check Item 3: kube-apiserver

When a cluster is connected to UCS, the error message shown in **Figure 7-2** may be displayed.

Figure 7-2 Abnormal kube-apiserver



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, so 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 **Fleets**.
- **Step 2** Log in to a master node of the destination 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 value of **clusters.cluster.server** in the kubeconfig file of the cluster is the same as the API server address of the cluster obtained 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, register the cluster to UCS again, and re-deploy proxy-agent.

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 kube-apiserver 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, register the cluster to UCS again, and re-deploy proxy-agent.

----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 unregister 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 third-party clusters provided by certain vendors. Pay attention to these vendors and try avoiding cluster arrears.

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

Symptom

The VPC endpoint of the cluster connected through a private network is deleted by mistake, and the cluster becomes abnormal.

Procedure

■ NOTE

The IP address of the VPC endpoint has been configured for proxy-agent. You need to specify an available IP address when creating a VPC endpoint.

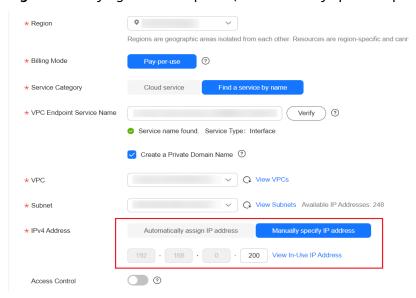
- **Step 1** Log in to the **VPC Endpoint console** to check whether the VPC endpoint in the region where UCS is located is deleted. If yes, go to the next step.
- **Step 2** Log in to a master node of the abnormal cluster.
- **Step 3** Query the IP address configured for proxy-agent.

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 VPC endpoint in the region where UCS is located, use the preceding IP address, and click View In-Use IP Address to ensure that this IP address is not used by another VPC endpoint. If the IP address is in use, change the proxy-agent configuration in the cluster. For details, see Changing the proxy-agent Configuration.

Figure 7-3 Buying a VPC endpoint (with manually specified private IP address)



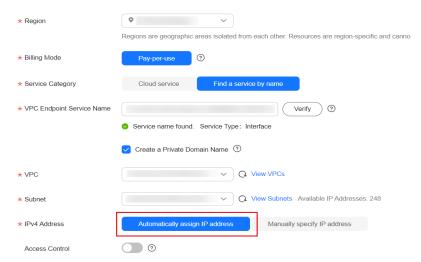
- **Step 5** Click **Next** and then **Submit** to create a VPC endpoint again.
- **Step 6** Wait for 1 to 3 minutes, go back to the UCS console, and refresh the cluster status.

----End

Changing the proxy-agent Configuration

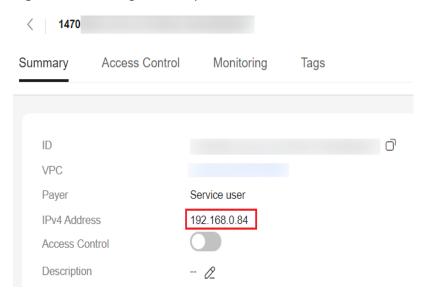
Step 1 Create a VPC endpoint in the region where UCS is located.

Figure 7-4 Buying a VPC endpoint (with automatically assigned private IP address)



Step 2 Click the ID of the new VPC endpoint to view the automatically allocated node IP address.

Figure 7-5 Viewing VPC endpoint details



- Step 3 Log in to a master node of the abnormal cluster.
- **Step 4** Change the IP address configured for 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

7.3 What Can I Do If proxy-agent Fails to Be Deployed?

Symptom

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

Procedure

- **Step 1** Log in to a master node of the cluster.
- **Step 2** Check the 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	ImagePullBa	ckOff	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:

• This 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 BackOffPullmage 57m (x16945 over 2d16h) kubelet Pulling image "swr.cn-north-4.myhuaweicloud.com/hwofficial-mcp/proxy-agent:22.3.1"
Warning Palledcreate 2m24s (x12328 over 2d16h) kubelet Pulling image "swr.cn-north-4.myhuaweicloud.com/hwofficial-mcp/proxy-agent:22.3.1"
Warning Palledcreate 2m24s (x12328 over 2d16h) kubelet Prore: ImagevullbackOff
```

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

 This 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

Narning FalledScheduling 2017h 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.
```

• If **gatekeeper** is displayed in the Kubernetes event, the created policy may have performed interception. To solve this problem, run the following command in the cluster to delete the corresponding policy:

kubectl delete constraint --all

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

----End

8 On-Premises Clusters

8.1 What Can I Do If an On-Premises Cluster Fails to Be Connected?

Symptom

This section describes how to troubleshoot cluster connection exceptions and provides solutions. The following exceptions may occur when a cluster is connected to UCS:

 You have registered 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 registration fails, click in the upper right corner of the cluster card to register it again and locate the fault as guided in **Troubleshooting**.

• If the status of a connected cluster is unavailable, rectify the fault by referring to **Troubleshooting** in this section.

Troubleshooting

Table 8-1 explains the error messages for you to locate faults.

Table 8-1 Error messages

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 status code. For example, status code 401 indicates that the user does not have the access permissions. A possible cause is that the cluster authentication information has expired.	• Check Item 4: Cluster Authenticati on 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.	-
"Current cluster version is not supported in UCS service."	This error occurs because the cluster version does not meet requirements. The version of the Kubernetes cluster connected to UCS must be 1.19 or later.	-

Check Item 1: proxy-agent

NOTICE

After a cluster is unregistered 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 a master node in the cluster.
- **Step 2** Check the deployment of proxy-agent.

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

Desired 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 **What Can I Do If proxy-agent Fails to Be Deployed?**.

■ NOTE

By default, proxy-agent is deployed with two pods. It can provide services as long as one pod is running normally. 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 working, check other items.

----End

Check Item 2: Network Connection Between the Cluster and UCS

Public Network Access

- **Step 1** Check whether a public IP address is bound to the cluster or a public NAT gateway is configured.
- **Step 2** Check whether the cluster security group allows outbound traffic. To perform access control on the outbound traffic, contact technical support to obtain the destination and port number.
- **Step 3** After rectifying network faults, delete the existing proxy-agent pods and 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 connected.

----End

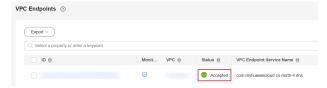
Private Network Access

- **Step 1** Check whether the cluster security group allows outbound traffic. To perform access control on the outbound traffic, contact technical support to obtain the destination and port number.
- **Step 2** Rectify the network connection faults between the cluster and UCS or IDC.

Refer to the following guides based on your network connection type:

- Direct Connect: Troubleshooting
- Virtual Private Network (VPN): Troubleshooting
- Step 3 Rectify the VPC endpoint fault. The VPC endpoint status must be Accepted. If the VPC endpoint is deleted accidently, create another one. For details, see How Do I Restore a Deleted VPC Endpoint for a Cluster Connected Through a Private Network?.

Figure 8-1 Checking the VPC endpoint status



Step 4 Delete the existing proxy-agent pods and rebuild pods. Check whether the logs of the new pods contain "Start serving".

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

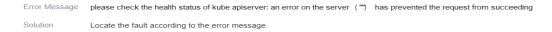
Step 5 If desired logs are printed, refresh the UCS console page and check whether the cluster is connected.

----End

Check Item 3: kube-apiserver

When a cluster is connected to UCS, the error message shown in **Figure 8-2** may be displayed.

Figure 8-2 Abnormal kube-apiserver



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, so 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 **Fleets**.
- **Step 2** Log in to a master node of the destination cluster and check whether the proxyagent pod can access the kube-apiserver of the destination cluster.

Example command:

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

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

----End

Check Item 4: Cluster Authentication Information Changes

If the error message "cluster responded with non-successful status: [401] [Unauthorized]" is displayed, the IAM network connection may be faulty, according to /var/paas/sys/log/kubernetes/auth-server.log of the three master nodes in the cluster. Ensure that the IAM domain name resolution and the IAM service connectivity are normal.

The common logs are as follows:

- Failed to authenticate token: *******: dial tcp: lookup iam.myhuaweicloud.com on *.*.*.*:53: no such host
 - This log indicates that the node is not capable of resolving **iam.myhuaweicloud.com**. Configure the corresponding domain name resolution by referring to **Preparing for Installation**.
- Failed to authenticate token: Get *******: dial tcp *.*.*:443: i/o timeout

 This log indicates that the node's access to IAM times out. Ensure that the node can communicate with Huawei Cloud IAM properly.
- currently only supports Agency token
 - This log indicates that the request is not initiated by UCS. Currently, onpremises clusters can only be connected to UCS using IAM tokens.
- IAM assumed user has no authorization/iam assumed user should allowed by TEAdmin
 - This log indicates that the connection between UCS and the cluster is abnormal. Contact Huawei technical support.
- Failed to authenticate token: token expired, please acquire a new token
 This log indicates that the token has expired. Run the **date** command to check
 whether the time difference is too large. If yes, synchronize the time and
 check whether the cluster is working. If the fault persists for a long time, you
 may need to reinstall the cluster. In this case, contact Huawei technical
 support.

After the preceding problem is resolved, run the **crictl ps | grep auth | awk '{print \$1}' | xargs crictl stop** command to restart the **auth-server** container.

8.2 How Do I Manually Clear Nodes of an On-Premises Cluster?

Precautions

Clearing a node is a high-risk operation. After the node is cleared, all processes (including the Kubernetes process and containerd) and data (including containers and images) on the node are cleared, and the node status cannot be restored. Therefore, before performing this operation, ensure that the node is no longer needed by the on-premises cluster.

Scenario

If the execution of the **ucs-ctl delete cluster** and **ucs-ctl delete node** commands fails on the on-premises cluster, you need to manually clear the node by referring to this section.

Procedure

Step 1 Obtain the node clearing script from the installed node.

Obtain the node clearing script uninstall_node.sh from the /var/paas/.ucs-package/ucs-onpremise/scripts/ directory generated after the decompression.

- **Step 2** Copy the script to the node to be cleared.
- **Step 3** Log in to the node to be cleared and run the following command:

bash uninstall_node.sh

∩ NOTE

To reduce residual processes or data, the script can be executed multiple times.

- **Step 4** After the script is executed, restart the node.
- **Step 5** Repeat the preceding operations to clear other nodes.

----End

8.3 How Do I Downgrade a cgroup?

Symptom

The **etcd Kubernetes** container cannot be started. Run the **journalctl -u containerd** command to view the containerd log. The following log is displayed:

applying cgroup configuration for process caused \\\"mountpoint for cgroup not found\\\"\"\n

Run the **stat** -fc %T /sys/fs/cgroup/ command to check the cgroup version. The command output shows that the cgroup version is cgroup2fs. The root cause is that cgroup v2 of Kubernetes is not put into commercial use and the cgroup needs to be downgraded.

Procedure

- Step 1 Add systemd.unified_cgroup_hierarchy=no to the GRUB_CMDLINE_LINUX configuration item in the /etc/default/grub file and disable cgroup v2.

 GRUB_CMDLINE_LINUX="net.ifnames=0 biosdevname=0 systemd.unified_cgroup_hierarchy=no"
- **Step 2** Run the **sudo grub-mkconfig -o /boot/grub/grub.cfg** command to regenerate the boot.
- **Step 3** Run the **reboot** command to restart the server.

----End

8.4 What Can I Do If the VM SSH Connection Times Out?

The performance of some VMs does not meet the requirements. As a result, the SSH connection occasionally times out. In this case, you can modify the SSH configurations of the VMs to solve the problem.

Step 1 Run the following command:

vim /etc/ssh/sshd_config

- **Step 2** Press **i** to enter editing mode.
- Step 3 Set UseDNS to no.

```
PermitRootLogin yes
PasswordAuthentication yes
UseDNS no

DenyUsers paas
"/etc/ssh/sshd_config" 144L, 3954C
```

Step 4 Press **Esc** and enter :wq! to save the settings and exit.

----End

8.5 How Do I Expand the Disk Capacity of the CIA Addon in an On-Premises Cluster?

Symptom

When the disk in the PVC storage on which the CIA add-on (kube-prometheus-stack) depends is full, "no space left on device" is reported in the standard log output of the prometheus-server-0 pod. In this case, Prometheus cannot work properly. You need to expand the disk capacity of the node where the PVC is located and restart the prometheus-server-0 pod.

Procedure

Step 1 Check the PVC information of Prometheus and obtain the node and path where the PV bound to the PVC is located.

kubectl describe pvc pvc-prometheus-server-0 -nmonitoring|grep volume.kubernetes.io/selected-node

kubectl describe pvc pvc-prometheus-server-0 -nmonitoring|grep volume.kubernetes.io/targetPath

- Step 2 Log in to the storage node, run the df -TH command to query the disk configuration of the node, and expand the capacity according to the disk configuration. After the capacity expansion is complete, you need to partition the disk. For details, see Extending Disk Partitions and File Systems (Linux) in the Elastic Volume Service User Guide.
- **Step 3** After the capacity expansion and partitioning are complete, restart Prometheus.

kubectl delete pod prometheus-server-0 -nmonitoring

NOTICE

If the prometheus-server-0 pod is restarted, the CIA function of the container will be unavailable during the restart. Select an appropriate restart time.

----End

8.6 What Can I Do If the Cluster Console Is Unavailable After a Master Node Is Shut Down?

Symptom

After a master node is shut down, the cluster console is unavailable.

Procedure

The Cilium community does not remove the Cilium endpoint from the pod in the **Terminating** state. As a result, some requests are distributed to the stopped node, and the requests fail. Perform the following operations:

Step 1 Run the following command to delete the pod in the **Terminating** state:

kubectl get pods -nkube-system | grep Terminating | awk '{print \$1}'|xargs kubectl delete pods -nkube-system

Step 2 Run the following command to check whether any pod malfunctions:

kubectl get pods -nkube-system

Step 3 After several minutes, the cluster console works properly again.

----End

8.7 What Can I Do If a Node Is Not Ready After Its Scale-Out?

Symptom

After a node in an on-premises cluster is scaled out, the Kubernetes resources of the node may not be able to be started. This means the node is not ready.

Procedure

The /mnt/paas/kubernetes/kubelet/cpu_manager_state file stores original cpu_manager_policy, which is the core binding setting of the original CPU cores and needs to be deleted. Restart kubelet to enable cpu_manager to bind cores based on the existing CPU topology and generate cpu_manager_state again.

Run the following commands:

rm /mnt/paas/kubernetes/kubelet/cpu_manager_state

systemctl restart kubelet

Wait for a period of time until the node is working.

8.8 How Do I Update the CA/TLS Certificate of an On-Premises Cluster?

Prerequisites

- All components in the on-premises cluster are running normally.
- Each node in the on-premises cluster is in the ready state.

Procedure

- **Step 1** Download ucs-ctl and save it to the **/root/ucs** directory on any master node in the on-premises cluster.
- Step 2 Record the passwords of all nodes in a table and save the table to the /root/ucs/ update_cert.csv directory on the node where the binary tool is located. For details, see Table 8-2.

Table 8-2 Table template

Field	Description
Node IP	Node IP address, which is mandatory.
Node Role	Node role, which is mandatory. The options are master and node .

Field	Description
User	User name for logging in to a node, which is mandatory.
Password	Password for logging in to a node, which is optional.
Auth Type	Node authentication type, which is optional. The options are password and key .
Key Path	Key path for logging in to a node, which is optional.

Example:

Node IP, Node Role, User, Password, Auth Type, Key Path

192.168.0.145, master, root, xxx, password,

192.168.0.225, master, root, xxx, password,

192.168.0.68, master, root, xxx, password,

192.168.0.89,node,root,xxx,password,

Step 3 Export environment variables.

export CUSTOM_DOMAIN={ucs_endpoint},10.247.0.1

□ NOTE

- *ucs_endpoint* indicates the server access address. You can run the following command to obtain it:
 - cat /var/paas/srv/kubernetes/kubeconfig | grep server
- If you need to update the certificate again, delete the files in the certificate directory and run the following command to update the certificate: rm -rf /var/paas/ucs/{cluster_name}/cert_update
- **Step 4** Update the certificate.

cd /root/ucs

./ucs-ctl kcm update-cert {cluster_name} -c update_cert.csv

Step 5 Update the certificate again after a failure.

./ucs-ctl kcm update-cert {cluster_name} -c update_cert.csv -r

Step 6 Roll back the certificate update after a failure.

./ucs-ctl kcm rollback-cert {cluster_name} -c update_cert.csv

----End

8.9 What Can I Do If an On-Premises Cluster Fails to Be Installed?

Symptom 1: Incorrect Binary File Is Used for ucs-ctl and "./ucs-ctl: cannot execute binary file: Exec format error" Is Displayed

Possible cause: The selected cluster architecture is inconsistent with the CPU architecture of the node where the on-premises cluster is installed.

Examples:

- If the cluster architecture is set to x86 and the on-premises cluster is installed on an Arm node, the installation fails.
- If the cluster architecture is set to Arm and the on-premises cluster is installed on an x86 node, the installation fails.

Solution: Cluster nodes do not support heterogeneous architectures. Ensure that the cluster architecture is the same as the CPU architecture of the node where the on-premises cluster is installed.

- If the cluster architecture is set to x86, install the on-premises cluster on an x86 node.
- If the cluster architecture is set to Arm, install the on-premises cluster on an Arm node.

Symptom 2: "Error: open /var/paas/.ucs-package/ucs-onpremise/images/base: no such file or directory" Is Displayed When an On-Premises Cluster Fails to Be Installed

Symptom: "Error: open /var/paas/.ucs-package/ucs-onpremise/images/base: no such file or directory." is displayed when an on-premises cluster fails to be installed.

Possible cause:

There are residual historical installation files.

Solutions:

(Recommended) Solution 1: Run the following command on the executor to manually clear the cluster installation package:

rm -rf /var/paas/.ucs-package/ucs-onpremise_*

Solution 2: Clear the node. For details, see **How Do I Manually Clear Nodes of an On-Premises Cluster?**.

9 Multi-Cloud Clusters

9.1 How Do I Clear Multi-Cloud Cluster Resources?

If the multi-cloud cluster unregistration fails, you can try unregistration again. Before performing this operation, ensure that you have manually deleted the resources associated with the cluster on the AWS console. This section describes the names and quantities of these resources. You can access the EC2 panel and VPC panel of AWS to view and delete these resources.

□ NOTE

In **Table 9-1**, *\${clusterName}* is your cluster name, and *\${random5}* is a random string of five characters.

Table 9-1 Names and quantities of resources

Cons ole	Resource Type	Quantity	Name
EC2 panel	EC2	Master nodes: 3 Worker nodes: <i>n</i>	Master node: <i>\${clusterName}</i> -cp- <i>\${random5}</i> Worker node: <i>\${clusterName}</i> -md-\${i}- <i>\$ {random5}</i> . The default value of \${i} is 0.
	Security group	5	 \${clusterName}-node \${clusterName}-lb \${clusterName}-apiserver-lb \${clusterName}-controlplane default The VPC corresponding to the preceding security groups is \${clusterName}-vpc.
	EIP	3	<i>\${clusterName}</i> -eip-apiserver

Cons ole	Resource Type	Quantity	Name
	Volume	Nodes: 2	Determine the node to which the volume belongs based on the name of the EC2 instance to which the volume is mounted.
	ELB	1	<i>\${clusterName}</i> -apiserver. The corresponding VPC is <i>\${clusterName}</i> -vpc.
	Network port	4	If Name is empty, the corresponding VPC is <i>\$</i> {clusterName}-vpc.
VPC	VPC	1	\${clusterName}-vpc
panel	NAT	3	 \${clusterName}-nat VPC: \${clusterName}-vpc; subnet: \$ {clusterName}-subnet-public-\${az1} \${clusterName}-nat VPC: \${clusterName}-vpc; subnet: \$ {clusterName}-subnet-public-\${az2} \${clusterName}-nat VPC: \${clusterName}-vpc; subnet: \$ {clusterName}-subnet-public-\${az3}
	Subnet	6	 \${clusterName}-subnet-public-\${az1} \${clusterName}-subnet-private-\${az1} \${clusterName}-subnet-public-\${az2} \${clusterName}-subnet-private-\${az2} \${clusterName}-subnet-public-\${az3} \${clusterName}-subnet-private-\${az3} The VPC corresponding to the preceding subnets is \${clusterName}-vpc.

Cons ole	Resource Type	Quantity	Name
	Route table	7	 \${clusterName}-rt-public-\${az1}. The explicit subnet association is \${clusterName}-subnet-public-\${az1}. \${clusterName}-rt-private-\${az1}. The explicit subnet association is \${clusterName}-subnet-private-\${az1}. \${clusterName}-subnet-private-\${az1}. \${clusterName}-rt-public-\${az2}. The explicit subnet association is \${clusterName}-subnet-public-\${az2}. \${clusterName}-rt-private-\${az2}. The explicit subnet association is \${clusterName}-subnet-private-\${az2}. \${clusterName}-subnet-private-\${az2}. \${clusterName}-subnet-private-\${az3}. The explicit subnet association is \${clusterName}-subnet-public-\${az3}. \${clusterName}-subnet-private-\${az3}. The explicit subnet association is \${clusterName}-subnet-private-\${az3}. If Name is empty, the explicit subnet association is empty. The VPC corresponding to the route tables is \${clusterName}-vpc.
	Internet gateway	1	<i>\${clusterName}</i> -igw The corresponding VPC is <i>\${clusterName}</i> -vpc.
	Network ACL	1	If Name is empty and 6 subnets are associated, the corresponding VPC is <i>\$</i> {clusterName}-vpc.

9.2 How Do I Obtain an Access Key (AK/SK)?

When registering a multi-cloud cluster, you need to obtain the access key (AK/SK) to create resources (such as EC2 instances, security groups, EIPs, and load balancers) related to the multi-cloud cluster in your AWS account. This section describes how to obtain the AK/SK.

The key will be encrypted and stored properly. You do not need to worry about information leakage.

- 1. Log in to the IAM console using your AWS account ID or account alias, and your IAM username and password.
 - To obtain an AWS account ID, contact the administrator of your AWS account.
- 2. In the navigation pane in the upper right corner, select your username, and then select **Security credentials**.

Account ID: 123456789012
IAM user: User1
Account

Organization

Service Quotas

Billing Dashboard

Security credentials

Switch role

Sign out

Figure 9-1 Selecting Security credentials

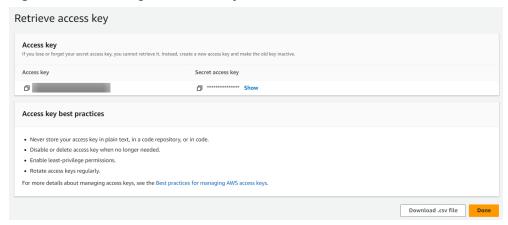
3. In the **Access keys** area, click **Create access key**. If you already have two access keys, this button will be disabled. You must delete one access key before creating another. You can also use an existing access key to create a UCS on AWS cluster.

Figure 9-2 Creating an access key



4. On the **Retrieve access key** page, click **Show** to obtain the value of the user's access key, or click **Download .csv file**. This is your only chance to keep your access key. After saving the access key in a secure location, click **Done**.

Figure 9-3 Obtaining the access key



♠ CAUTION

When you are using a UCS on AWS cluster, do not rotate, disable, or delete the access key before the cluster is deleted. Otherwise, subsequent cluster update and deletion operations will be affected.

9.3 How Do I Update the Multi-Cloud Cluster Certificate?

Prerequisites

- All components in the multi-cloud cluster are running normally.
- Each node in the multi-cloud cluster is in the **ready** state.

Procedure

- **Step 1** Download ucs-ctl and save it to the **/root/ucs** directory on any master node in the on-premises cluster.
- Step 2 Record the passwords of all nodes in a table and save the table to the /root/ucs/ update_cert.csv directory on the node where the binary tool is located. For details, see Table 9-2.

Table 9-2 Table template

Field	Description	
Node IP	Node IP address, which is mandatory.	
Node Role	Node role, which is mandatory. The options are master and node .	
User	User name for logging in to a node, which is mandatory.	
Password	Password for logging in to a node, which is optional.	
Auth Type	Node authentication type, which is optional. The options are password and key .	
Key Path	Key path for logging in to a node, which is optional.	

Example:

Node IP, Node Role, User, Password, Auth Type, Key Path

192.168.0.145, master, root, xxx, password

192.168.0.225, master, root, xxx, password

192.168.0.68, master, root, xxx, password

192.168.0.89, node, root, xxx, password

- **Step 3** Log in to the AWS console, edit security group *{cluster_name}*-**node**, and allow port 22 in the security group to ensure that the security group can be accessed.
- **Step 4** Enable password access, log in to all nodes in the cluster, and run the following commands:

sed -i 's/PasswordAuthentication no/PasswordAuthentication yes/g' /etc/ssh/sshd_config echo "PermitRootLogin yes" >> /etc/ssh/sshd_config systemctl restart sshd passwd

Set the node passwords and save them to your local PC.

Step 5 Configure environment variables.

export CUSTOM_DOMAIN={ucs_endpoint},10.247.0.1

□ NOTE

ucs_endpoint indicates the server access address. You can run the following command to
obtain it:

cat /var/paas/srv/kubernetes/kubeconfig | grep server

- If you perform operations on the executor where the cluster is installed, you do not need to configure environment variables.
- Step 6 Update the certificate.

cd /root/ucs cp /var/paas/srv/kubernetes/ca.key /var/paas/srv/kubernetes/ca_key.pem ./ucs-ctl kcm update-cert {cluster_name} -c update_cert.csv

Step 7 Update the certificate again after a failure.

./ucs-ctl kcm update-cert {cluster_name} -c update_cert.csv -r

Step 8 Roll back the certificate update after a failure.

./ucs-ctl kcm rollback-cert {cluster_name} -c update_cert.csv

----End

10 Cluster Federation

10.1 What Can I Do If the Pre-upgrade Check of the Cluster Federation Fails?

Context

Before the cluster federation upgrade, UCS checks the cluster federation status, cluster status, and cluster access status to reduce the probability of upgrade failures. If any exception is detected, you can rectify the fault by referring to this section. After the fault is rectified, you can upgrade the federation again.

Before upgrading the federation, if any error is reported for the federation status, cluster status, or cluster access status, rectify the fault to prevent the upgrade failure.

Symptom 1: The Cluster Federation Status Is Abnormal

Cause: The cluster federation does not run properly.

Solution: Disable the cluster federation and then enable it again. For details, see **Enabling Cluster Federation**. If the cluster federation cannot be disabled due to service reasons, submit a service ticket and contact technical support.

Symptom 2: The Cluster Status Is Abnormal

Cause: The cluster does not run properly in the fleet or the cluster cannot be accessed.

Solution:

- If the cluster does not run properly, you can restore the cluster.
- If the cluster cannot be accessed, add the cluster to the federation again. If the cluster cannot be added to the federation again, submit a service ticket and contact technical support personnel.

10.2 What Can I Do If a Cluster Fails to Be Added to a Federation?

Context

When a cluster is added to a federation, the error message "the same cluster has been registered with name clusterName" or "cluster(clusterName) is joined successfully" is displayed.

Possible Cause

The cluster node is faulty and then pods are restarted. As the **karmadactl join** command is not idempotent, an error is reported when the command is executed again.

Solution

Remove the cluster from the federation and run the **kubectl get cluster** command to check whether the cluster exists in the federation.

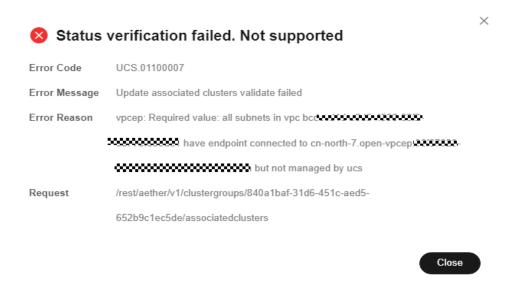
- If the cluster exists in the federation, run the **kubectl edit cluster clusterName** command to edit the YAML file and delete the **finalizers** field. Then run the **kubectl get cluster** command to check whether the cluster exists in the federation.
- If the cluster does not exist in the federation, add the cluster to the federation again.

10.3 What Can I Do If Status Verification Fails When Clusters Are Added to a Federation?

Context

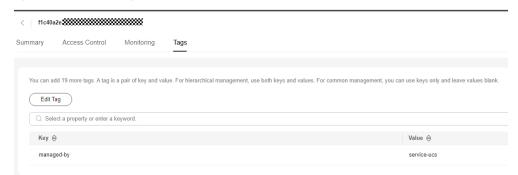
When clusters are added to a federation, "Status verification failed. Not supported" is displayed. The error code is **UCS.01100007**. The error message is **Update associated clusters validate failed**. The error cause is **vpcep: Required value: all subnets in vpc xxx have endpoint connected to xxx.vpcep-src-open.xxx but not managed by ucs**.

Figure 10-1 Failed status verification



Solution

In all subnets of the VPC where the error occurs, check whether there is a VPC endpoint that is created and bound to the VPC endpoint service that reports the error. If yes, go to the details page of this VPC endpoint and add the **managed-by=service-ucs** tag to this VPC endpoint.



10.4 What Can I Do If an HPA Created on the Cluster Federation Management Plane Fails to Be Distributed to Member Clusters?

Context

After an HPA is created on the cluster federation management plane, the PropagationPolicy fails to distribute the HPA to the member cluster earlier than v1.23.

Possible Cause

Currently, the API server version of the UCS cluster federation is 1.25. Therefore, the HPA has two versions: autoscaling/v2 and autoscaling/v1. However, only HPA

autoscaling/v2 is distributed. Clusters earlier than v1.23 do not support HPA autoscaling/v2. As a result, the HPA cannot be distributed to the member cluster earlier than v1.23. Check resourceBinding of the HPA. The error message "cluster(s) did not have the API resource" is displayed.

Solution

Before distributing the HPA, you can upgrade the member cluster to v1.23 or later, which supports HPA autoscaling/v2 by default.

If you still want to distribute HPA autoscaling/v1 to member clusters, set the resourceSelectors[i].apiVersion field in your PropagationPolicy to autoscaling/v2, as shown in the example YAML. After the distribution is successful, you can query HPA autoscaling/v1 in the member cluster.

```
apiVersion: autoscaling/v1
kind: HorizontalPodAutoscaler
metadata:
 name: test-hpa
 maxReplicas: 5
 minReplicas: 1
 scaleTargetRef:
  apiVersion: apps/v1
  kind: Deployment
  name: nginx
 targetCPUUtilizationPercentage: 10
apiVersion: policy.karmada.io/v1alpha1
kind: PropagationPolicy
metadata:
name: test-hpa-pp
spec:
 placement:
  clusterAffinity:
    clusterNames:
     - member1
 resourceSelectors:
   - apiVersion: autoscaling/v2
    kind: HorizontalPodAutoscaler
    name: test-hpa
    namespace: default
```

10.5 What Can I Do If an MCI Object Fails to Be Created?

Symptom

The MCI object fails to be created.

Troubleshooting

Run the **kubectl describe mci mci-example -n demo** command to view events. The following figures show example command outputs.

• Case 1



Case 2



Solution

If an error is reported as shown in case 1, the listener port configured during MCI object creation has been used. You can use either of the following solutions:

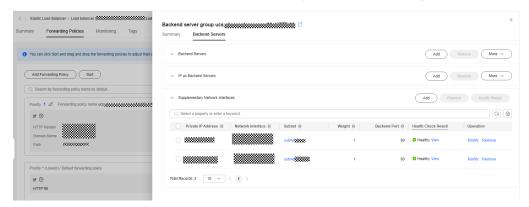
- Edit the MCI object that fails to be created and use an unused listener port.
- Log in to the ELB console and delete the listener of the corresponding port.

If an error is reported as shown in case 2, **karmada.io/elb.projectid** is incorrectly configured during MCI object creation. In this case, you need to delete the created MCI and create another one with the correct configuration.

10.6 What Can I Do If I Fail to Access a Service Through MCI?

If you fail to access a service after creating an MCI, check whether the MCI object is successfully configured.

Log in to the ELB console, locate the ELB instance based on the ID of the ELB instance bound to the MCI, and click the name of the ELB instance to access its details page. On the **Listeners** tab, locate the target listener and click **Add/Edit Forwarding Policy**. On the displayed page, click the name of the target backend server group to access the details page. Then, click the **Backend Servers** tab and check whether this ELB instance is successfully bound to the target workload.



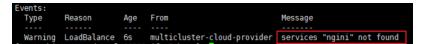
- 1. If the backend server has been deleted, check whether the CIDR block of the pod conflicts with the VPC CIDR block of the ELB.
- 2. If there is no backend server in the backend server group, run the following command to query the corresponding event and rectify the fault based on the event information. Replace **{MCIname}** with the MCI name.

kubectl describe mci {MCIname}

If the following error information is displayed, check whether the listener port of the ELB instance is occupied:



If the following error information is displayed, check whether the service name exists in the MCI:



If the following error information is displayed, check whether the service port configured in the MCI is correct:



10.7 What Can I Do If an MCS Object Fails to Be Created?

Symptom

If the MCS object fails to be created, run the **kubectl describe mcs mcs-example**-n demo command to view events. The following figure shows an example command output.

Solution

karmada.io/elb.projectid is incorrectly configured during MCS object creation. In this case, you need to delete the created MCS and create another one with the correct configuration.

10.8 What Can I Do If an MCS or MCI Instance Fails to Be Deleted?

Symptom: The MCS or MCI instance is always in the **Deleting** state.



Solution: Forcibly delete the abnormal instances and manually clear residual resources.

Procedure

Step 1 Log in to the UCS console. In the navigation pane, choose **Fleets**. On the **Fleets** tab, click the name of the target fleet to access the fleet console.

- **Step 2** In the navigation pane, choose **Services and Ingresses**. On the displayed page, click the **MCS** or **MCI** tab.
- **Step 3** Select the abnormal instance that is always in the **Deleting** state.
- **Step 4** Manually clear the residual ELB resources of this abnormal instance.
 - 1. Click dunder the access type of the abnormal instance. The details page of the load balancer associated with this instance is displayed.
 - 2. On the **Summary** tab, disable **Deletion Protection**.



- 3. Switch to the **Listeners** tab and delete the listener associated with the abnormal instance.
- 4. In the navigation pane, choose **Backend Server Groups**. On the displayed page, click the name of the backend server group associated with the listener of the abnormal instance. On the **Summary** tab, disable **Removal Protection** and **Modification Protection**.



- 5. Delete the backend server group associated with the listener of the abnormal instance.
- **Step 5** Forcibly delete the abnormal instance.
 - 1. Select the abnormal instance that is always in the **Deleting** state and click **Edit YAML** in the **Operation** column.
 - 2. Delete the following content from the YAML file to forcibly delete the instance.

finalizers:

- karmada.io/multi-cluster-service-controller ## MCS finalizers:
- karmada.io/multi-cluster-ingress-controller ## MCI

```
Current

© YAML JSON Line break

I spiVersion: networking karmada io/vlalphat
2 kind: MulticlusterService
3 metadata:
4 amoutations:
5 karmada io/alb. health-oheck-flag: 'on'
6 karmada io/alb. heilth-oheck-flag: 'on'
7 karmada io/alb. heilth-oheck-flag: 'on'
8 karmada io/alb. port: 803
8 karmada io/alb. port: 803
8 karmada io/alb. port: 803
9 creationTimestamp: '2024-10-21708:25:492'
10 deletionTimestamp: '2024-10-21708:25:572'
11 deletionTimestamp: '2024-10-21708:25:572'
12 finalizers:
13 - karmada io/multi-cluster-service-controller
14 generation. 2
15 labels:
16 multi-clusterservice karmada io/permanent-id: 700*a3ac-29b3-40a8-9753-f68927d3cbf6
17 munagedFields:
18 - spiVersion: networking karmada io/vlalphal
19 fieldsType: FieldsV1
```

----End

11 Traffic Distribution

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

Symptom

If you have registered a domain name with a third-party registrar, and you want to use UCS to manage application 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 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**. On the displayed page, 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.

5. Click OK.

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

If the system displays a message indicating that the public zone has been created by another tenant, handle the issue by referring to **Reclaiming a Public Zone**.

☐ NOTE

Click the domain name to view SOA and NS record sets automatically generated for the zone.

- The SOA record set includes administrative information about your zone, as defined by the Domain Name System (DNS).
- The NS record set defines the authoritative servers for the domain name.
 You can modify the NS record set based on the region of the domain name. For more information about DNS servers, see What Are Huawei Cloud DNS Servers?

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.

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.

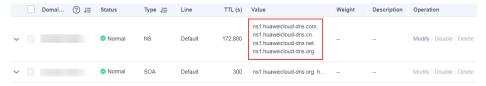
- 1. Log in to the Huawei Cloud console.
- 2. Move the cursor to the icon on the left of the page. In the service list, choose **Networking** > **Domain Name Service**.

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 11-1 NS record set returned by the system



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

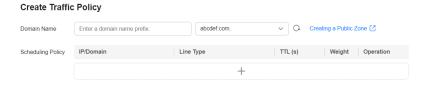
For details, see the operation guide on the domain name registrar's website.

----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 new domain name. If the domain name is not displayed, click **C** on the right to refresh the drop-down list.

Figure 11-2 Creating a traffic policy



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

Figure 11-3 Scheduling policy



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 destination cluster, the scheduling policy takes effect.

```
[root@no-del-cluster-_______08211 ~]# dig demo.
; <<>> DiG 9.11.4-P2-RedHat-9.11.4-26.P2.el7_9.9 <<>> demo.
;; 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.
                             IN
;; ANSWER SECTION:
                                           123.
demo.
                     300
                             IN
;; 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

12_{0&M}

12.1 What Can I Do If Monitoring Fails to Be Enabled for a Cluster Due to Residual Add-on Resources?

Symptom

- When monitoring is enabled for a cluster, the API returns an error message containing the field **release name already exists**.
- The request for enabling cluster monitoring has been delivered, but the
 monitoring status is Installation failed or Unknown. On the page for
 enabling monitoring, check the kube-prometheus-stack add-on. The add-on
 installation failure cause contains the field resource that already exists.

Possible Cause

The kube-prometheus-stack add-on has residual resources.

Troubleshooting

You can run the following commands to clear residual resources and enable monitoring again after the residual resources are cleared.

kubectl delete ns monitoring

kubectl delete ClusterRole cluster-problem-detector custom-metrics-resource-aggregated-reader event-exporter prometheus-operator prometheus-server ucsaddon-cie-collector-kube-state-metrics

kubectl delete ClusterRoleBinding ucsaddon-cie-collector-kube-state-metrics cluster-problem-detector event-exporter prometheus-operator prometheus-server

kubectl delete apiservice v1beta1.custom.metrics.k8s.io

12.2 What Can I Do If Monitoring Fails to Be Enabled for a Cluster Due to Policy Interception?

Symptom

- When monitoring is enabled for a cluster, the API returns an error message containing the field **gatekeeper**.
- The request for enabling cluster monitoring has been delivered, but the monitoring status is **Installing**. After the installation times out, **Installation fail** is displayed. Check the pod status of the add-on in the cluster. The pod event contains the field **gatekeeper**.

Possible Cause

If an interception policy is configured in the policy center for the cluster for which monitoring is to be enabled, monitoring may fail to be enabled.

Troubleshooting

Cancel the interception policy for the **kube-system** and **monitoring** namespaces in the policy instance of the specified cluster.

12.3 How Do I Modify the Collection Configuration of the kube-state-metrics Component?

Symptom

The kube-state-metrics component of the kube-prometheus-stack add-on converts the metrics data format of Prometheus into the format that can be identified by Kubernetes APIs. By default, the kube-state-metrics component does not collect all labels and annotations of Kubernetes resources. To collect these labels and annotations, you need to modify the collection configuration in the startup parameter and check whether the corresponding metrics are added to the collection whitelist of ServiceMonitor named **kube-state-metrics**.

Procedure

Step 1 Run the following command to open the YAML file corresponding to the workload **kube-state-metrics**:

kubectl edit deployment kube-state-metrics -nmonitoring

Step 2 Modify the startup parameter of **kube-state-metrics**.

For example, to collect all labels of a pod, modify the startup parameter of kubestate-metrics as follows:

--metric-labels-allow list=pods=[*], nodes=[node, failure-domain.beta.kubernetes.io/zone, topology.kubernetes.io/zone]

kube-state-metrics starts to collect labels of pods and nodes and uses **kubectl edit servicemonitor kube-state-metrics -nmonitoring** to check whether kube pod labels is in the collection task of Prometheus.

To collect annotations, add the parameter **--metric-annotations-allowlist** to the startup parameter in the same way.

For details, see https://github.com/kubernetes/kube-state-metrics/blob/v2.2.3/docs/cli-arguments.md.

----End

12.4 How Do I Handle the Error Displayed When Monitoring Is Enabled for a Cluster?

Symptom

When monitoring is enabled for an attached or on-premises cluster on the **Container Intelligent Analysis** page, the error message "helm install failed: rendered manifests contain a resource that already exists" is displayed.

Possible Cause

The self-built Prometheus add-on is installed in the attached or on-premises cluster.

Troubleshooting

Uninstall the self-built Prometheus add-on and delete the corresponding resources based on the resource name, resource type, and namespace after "continue with install" in the error message.

