

Application Service Mesh

FAQs

Issue 02
Date 2023-07-03



Copyright © Huawei Technologies Co., Ltd. 2023. All rights reserved.

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

Trademarks and Permissions



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

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

Contents

1 Service Mesh Cluster	1
1.1 Why Does a Service Mesh Remain in the Installing Status for a Long Time After I Enable It for a Cluster?	1
1.2 Why Does a Service Mesh Remain in the Unready Status for a Long Time After I Uninstall It?	1
2 Mesh Management	3
2.1 Why Cannot I Create a Mesh for My Cluster?	3
2.2 Why Are Exclusive Nodes Still Exist After Istio Is Uninstalled?	3
2.3 How Do I Upgrade ICAgent?	4
2.4 How Do I Enable Namespace Injection for a Cluster?	4
2.5 How Do I Disable Sidecar Injection for Workloads?	5
2.6 What Can I Do If A Pod Cannot Be Started Due to Unready Sidecar?	5
3 Adding a Service	10
3.1 What Do I Do If an Added Gateway Does Not Take Effect?	10
3.2 Why Does It Take a Long Time to Start the Demo Application in Experiencing Service Mesh in One Click?	10
3.3 Why Can't I Access the page of the Demo Application After It Is Successfully Deployed?	11
3.4 Why Cannot I Select the Corresponding Service When Adding a Route?	11
4 Performing Grayscale Release	12
4.1 Why Can't I Change the Image Used for the Grayscale Version When Performing Grayscale Release?	12
4.2 Why Does Not a Grayscale Policy that Based on Request Content Take Effect for Some Services?	12

1 Service Mesh Cluster

1.1 Why Does a Service Mesh Remain in the Installing Status for a Long Time After I Enable It for a Cluster?

Description

After I create a service mesh (that is, buy a Dedicated mesh) for a CCE cluster, the mesh remains in the installing status for a long time and a message is displayed, indicating that the user security group rules are successfully enabled.

Fault Diagnosis

Log in to the CCE console, choose **Resource Management** > **Namespaces**, and check whether the **istio-system** namespace exists.

Analysis

Residual **istio-system** namespaces exist.

Solution

Delete the residual **istio-system** namespaces and install the mesh again.

1.2 Why Does a Service Mesh Remain in the Unready Status for a Long Time After I Uninstall It?

Description

On the ASM console, after I uninstall a service mesh, the mesh remains in the unready status for a long time.

Fault Diagnosis

Step 1 Log in to the CCE console, click the target cluster to go to its details page, and choose **O&M > Charts** in the navigation pane on the left.

Step 2 Click **Releases** and check the releases and the latest events about uninstallation failure.

The **Status** of **istio-master** is **Uninstallation Failed**, and the following message is displayed.

```
deletion failed with 1 error(s): clusterroles.rbac.authorization.k8s.io "istio-cleanup-secrets-istio-system" already exists
```

----End

Analysis

Abnormal operations cause the Helm chart of Istio stuck during uninstallation. Residual resources lead to an uninstallation failure.

Solution

Step 1 Connect to the CCE cluster using kubectl.

Step 2 Run the following commands to clear Istio resources:

```
kubectl delete ServiceAccount -n istio-system `kubectl get ServiceAccount -n istio-system | grep istio | awk '{print $1}`  
kubectl delete ClusterRole -n istio-system `kubectl get ClusterRole -n istio-system | grep istio | awk '{print $1}`  
kubectl delete ClusterRoleBinding -n istio-system `kubectl get ClusterRoleBinding -n istio-system | grep istio | awk '{print $1}`  
kubectl delete job -n istio-system `kubectl get job -n istio-system | grep istio | awk '{print $1}`  
kubectl delete crd -n istio-system `kubectl get crd -n istio-system | grep istio | awk '{print $1}`  
kubectl delete mutatingwebhookconfigurations -n istio-system `kubectl get mutatingwebhookconfigurations -n istio-system | grep istio | awk '{print $1}`
```

Step 3 Log in to the ASM console and uninstall the mesh again.

----End

2 Mesh Management

2.1 Why Cannot I Create a Mesh for My Cluster?

Description

I cannot create a mesh for my cluster.

Analysis

Currently, clusters of versions earlier than 1.15 cannot be managed by meshes.

Solution

- Step 1** Check the cluster version. Currently, only clusters of v1.15, v1.17, v1.19, v1.21, or v1.23 can be managed by meshes.
 - Step 2** Check your browser. Chrome is recommended. The button for mesh creation may be unavailable when you are using other browsers, such as Firefox, due to adaptation problems.
- End

2.2 Why Are Exclusive Nodes Still Exist After Istio Is Uninstalled?

Description

After Istio is uninstalled, exclusive nodes still exist.

Analysis

Only Istio control plane workloads will be deleted when you uninstall Istio for a cluster. Node resources will not be deleted automatically.

Solution

Nodes from which Istio are uninstalled can be used as common nodes. If these nodes are no longer required, log in to the CCE console, click the target cluster to go to its details page, and choose **Resources** > **Nodes** in the navigation pane to delete the nodes.

2.3 How Do I Upgrade ICAgent?

- Step 1** Log in to the ASM console. In the navigation pane, choose **Monitoring Center** to go to the APM console.
- Step 2** On the APM console, choose **Agent** > **Management** in the navigation pane, select the target cluster, and click **Upgrade ICAgent**.

----End

2.4 How Do I Enable Namespace Injection for a Cluster?

When injecting a sidecar to the namespace of a cluster, if the namespace injection is not enabled in the cluster, perform the following steps:

- Step 1** Connect to the cluster using `kubectl`.
- Step 2** Run the `kubectl get iop -nistio-system` command to query iop resources.

- If the following information is displayed, the iop resource exists. Go to [Step 3](#).

```
user@dts2fot109u4ymb-machine:~$ kubectl get iop -nistio-system
NAME          REVISION  STATUS  AGE
data-plane    1         HEALTHY 69d
```

- If the following information is displayed, no iop resources exist. Go to [Step 4](#).

```
web-terminal-7b778fc945-9m2hf:~# kubectl get iop -nistio-system
No resources found in istio-system namespace.
```

- Step 3** Run the `kubectl edit iop -nistio-system data-plane` command to modify the `autoInject` configuration item. In the preceding command, `data-plane` indicates the name of the iop resource queried in the previous step. Replace it with the actual value.

```
global:
  defaultPodDisruptionBudget:
    enabled: true
  hub: 100.79.1.215:20202/asm
  logging:
    level: default:info
  meshID: test-payment
  multiCluster:
    clusterName: test-yy
    network: test-yy-network
  proxy:
    autoInject: enabled
  remotePilotAddress: 10.252.2.34
  tag: 1.8.6-r1-20220512225026
```

- Step 4** Run the `kubectl edit cm -nistio-system istio-sidecar-injector` command to modify the `istio-sidecar-injector` configuration item.

```
data:  
  config: |-  
    policy: enabled
```

----End

2.5 How Do I Disable Sidecar Injection for Workloads?

After sidecar injection is enabled for a namespace of a cluster, sidecars are automatically injected for pods of all workloads in the namespace. You can configure sidecars not to be injected into some workloads:

- Step 1** Log in to the CCE console, click the target cluster to go to its details page, and choose **Resources > Workloads** in the navigation pane on the left.
- Step 2** Click **Edit YAML** in the **Operation** column of the target workload.
- Step 3** Find the **spec.template.metadata.annotations** field and add **sidecar.istio.io/inject: 'false'**.

```
annotations:  
  sidecar.istio.io/inject: 'false'
```

```
107 spec:  
108   replicas: 1  
109   selector:  
110     matchLabels:  
111       app: reviews  
112       version: v1  
113   template:  
114     metadata:  
115       creationTimestamp: null  
116     labels:  
117       app: reviews  
118       release: istio-bookinfo  
119       version: v1  
120     annotations:  
121       sidecar.istio.io/inject: 'false'
```

For more details about sidecar injection, see [Automatic Sidecar Injection](#).

----End

2.6 What Can I Do If A Pod Cannot Be Started Due to Unready Sidecar

Description

Pods of services managed by a mesh may fail to be started and keep restarting. When the service container communicates with external systems, the traffic passes through the **istio-proxy** container. However, the service container is started earlier than the **istio-proxy** container. As a result, the communication with external systems fails and the pod keeps restarting.

Solution

In Istio 1.7 and later versions, the community adds a switch named **HoldApplicationUntilProxyStarts** to the **istio-injector** injection logic. After the switch is enabled, the proxy is injected to the first container and the **istio-proxy** container is started earlier than the service container.

The switch can be configured globally or locally. The following describes two ways to enable the switch.

NOTICE

After this switch is enabled, the service container cannot be started until the sidecar is fully ready, which slows down pod startup and reduces scalability for burst traffic. You are advised to evaluate service scenarios and enable this switch only for required services.

- **Global Configuration**

- a. Run the following command to edit the IOP CR resource:

kubectl edit iop private-data-plane -n istio-system

Add the following command to the **spec.values.global.proxy** field:

```
holdApplicationUntilProxyStarts: true
```

```
values:
  gateways:
    istio-egressgateway:
      autoscaleEnabled: false
      labels:
        app: istio-egressgateway
      tolerations:
        - effect: NoExecute
          key: istio
          operator: Exists
    istio-ingressgateway:
      autoscaleEnabled: false
      customService: true
      labels:
        app: istio-ingressgateway
      replicaCount: 1
      tolerations:
        - effect: NoExecute
          key: istio
          operator: Exists
  global:
    defaultPodDisruptionBudget:
      enabled: true
    hub: swr.cn-north-7.myhuaweicloud.com/asm
    logging:
      level: default:info
    meshID: envoy-critical
    multiCluster:
      clusterName: test-yyl-multi
  proxy:
    autoInject: enabled
    holdApplicationUntilProxyStarts: true
```

- b. Run the following command to check whether the latest logs contain no error information:

```
kubectl logs -n istio-operator $(kubectl get po -n istio-operator | awk '{print $1}' | grep -v NAME)
```

- c. Run the following command to check whether the IOP CR is normal:

```
kubectl get iop -n istio-system
```

```
[root@lx666-14467 ~]# kubectl get iop -n istio-system
NAME                REVISION  STATUS  AGE
private-data-plane  1         HEALTHY 6d2h
[root@lx666-14467 ~]#
```

- d. Run the following command to upgrade the services in the mesh in a rolling manner:

```
kubectl rollout restart deployment nginx -n default
```

where, **nginx** is an example service, and **default** is the namespace. Replace them with the actual values.

- e. Run the following command to check whether the pod is restarted:

```
kubectl get pod -n default | grep nginx
```

```
[root@lx666-14467 ~]# kubectl get pod -n default | grep nginx
nginx-6b4959ffff-pr8t8    2/2    Running    0    14s
[root@lx666-14467 ~]#
```

- f. Run the following command to check whether **postStart lifecycle** is added to the pod and whether the **istio-proxy** container is placed in the first position:

```
kubectl edit pod nginx-7bc96f87b9-l4dbl
```

```
- name: ISTIO_META_CLUSTER_ID
  value: test-yy1-multi
image: swr.cn-north-7.myhuaweicloud.com/asm/proxyv2:1.13.9-r1-20221110212800
imagePullPolicy: IfNotPresent
lifecycle:
  postStart:
    exec:
      command:
        - pilot-agent
        - wait
name: istio-proxy
ports:
```

- **Local Configuration**

For Istio 1.8 or later versions, you can label the pods for which this function needs to be enabled with **proxy.istio.io/config** and set **holdApplicationUntilProxyStarts** to true.

The following uses the **nginx** service in the **default** namespace as an example. The operations for other services are similar.

```
kubectl edit deploy nginx -n default
```

Add the following commands to the **spec.template.metadata.annotations** field:

```
proxy.istio.io/config: |
  holdApplicationUntilProxyStarts: true
```

```
apiVersion: apps/v1
kind: Deployment
metadata:
  annotations:
    deployment.kubernetes.io/revision: "6"
    description: ""
  creationTimestamp: "2022-11-24T07:55:31Z"
  generation: 6
  labels:
    appgroup: ""
    version: v1
  name: tomcat
  namespace: default
  resourceVersion: "55550644"
  uid: cd5dbfe8-83cc-4964-86fc-f657c85e852d
spec:
  progressDeadlineSeconds: 600
  replicas: 1
  revisionHistoryLimit: 10
  selector:
    matchLabels:
      app: tomcat
      version: v1
  strategy:
    rollingUpdate:
      maxSurge: 25%
      maxUnavailable: 25%
    type: RollingUpdate
  template:
    metadata:
      annotations:
        kubectrl.kubernetes.io/restartedAt: "2022-11-25T10:35:02+08:00"
        proxy.istio.io/config: |
          holdApplicationUntilProxyStarts: true
    creationTimestamp: null
```

3 Adding a Service

3.1 What Do I Do If an Added Gateway Does Not Take Effect?

The possible cause is that the Gateway-related resource configurations are missing or incorrect. Do as follows to locate the fault:

- Log in to the Elastic Load Balance console, check whether the external port and ECSs are successfully listened by the load balancer.
- Log in to the cluster and run the **kubectl get gateway -n istio-system** command to check whether the IP address, domain name, and port number are configured for the Gateway. Run the **kubectl get svc -n istio-system** command to check whether the ingress Gateway has the corresponding IP address and port and is not in the pending status.
- Check whether the internal access protocol of the service added to the service mesh is consistent with the external access protocol configured for the service's Gateway.
- If the **ERR_UNSAFE_PORT** error is displayed when you use a browser to access the service, that is because the port is identified as a dangerous port by the browser. In this case, you need to use another external port.

3.2 Why Does It Take a Long Time to Start the Demo Application in Experiencing Service Mesh in One Click?

The demo application contains the productpage, details, ratings, and reviews services. All related workloads and Istio resources including DestinationRule, VirtualService, and Gateway need to be created. Therefore, the creation takes a comparatively long time.

3.3 Why Can't I Access the page of the Demo Application After It Is Successfully Deployed?

Description

The page of the demo application cannot be accessed after the application is successfully deployed.

Analysis

The load balancer configured for the application does not listen to the port.

Solution

Log in to the Elastic Load Balance console. Check whether the port listener has been created and whether the health status of the backend server is normal. For details about how to create a load balancer, see [Listener](#).

3.4 Why Cannot I Select the Corresponding Service When Adding a Route?

During adding a route, the target service is filtered based on the corresponding gateway protocol. The filtering rules are as follows:

- For an HTTP gateway, select an HTTP service.
- For a TCP gateway, select a TCP service.
- For a gRPC gateway, select a gRPC service.
- For an HTTPS gateway, select either an HTTP or a gRPC service.
- For a TLS gateway which TLS termination is enabled, select a TCP service. If TLS termination for a TLS gateway is disabled, select a TLS service.

4 Performing Grayscale Release

4.1 Why Can't I Change the Image Used for the Grayscale Version When Performing Grayscale Release?

Description

When I perform grayscale release, the image used for the grayscale version cannot be changed.

Analysis

When performing grayscale release on a service, you create a new version of the same service. Therefore, the image used by the service cannot be changed. Only image tags can be changed.

Solution

Pack the required image into a different tag of the same image and push it to the image repository. Then, select the newly pushed image tag when you perform grayscale release on the service.

4.2 Why Does Not a Grayscale Policy that Based on Request Content Take Effect for Some Services?

Description

A grayscale policy that based on request content does not take effect on some services.

Analysis

A grayscale policy based on request content is valid only for the entry service that is directly accessed.

Solution

If you want a grayscale policy to be applied to all services in an application, the header information of the HTTP request needs to be transferred in the service code.