Basic Checks

You should always check these common mistakes before anything else, they are fast check and can save you hours.

APIC Faults

Checking the Tenant for Faults

Before you begin

No faults should be present for the Tenant and Container Domain that you are using.

Step 1 In the APIC GUI, on the menu bar, choose **Tenants** > ***tenant_name***.

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

Step 3 In the *tenant_name* pane, click on the **Fault** tab.

Step 4 Ensure there are not faults present that you are using.

Checking the Container Domains for Faults

Before you begin

No faults should be present for the Tenant and Container Domain that you are using.

Step 1 In the APIC GUI, on the menu bar, choose **Virtual Networking**.

Step 2 In the navigation pane, expand **Container Domain** and choose either **Kubernetes** or **OpenShift**.

Step 3 In the **Kubernetes** or **OpenShift** pane, click on the **Fault** tab.

Step 4 Ensure there are not faults present that you are using.

Bridge Domains and VRFs

Before you begin

- Make sure you do not change the default names and parameters for the bridge domain (BD) and VRF on the APIC.
- VRF must be set to Enforced.

Step 1 In the APIC GUI, on the menu bar, choose **Tenants** > ***tenant-name***.

Step 2 In the navigation pane, expand ***tenant-name*** > **Networking** > ***VRF-name***.

Step 3 In the *VRF-name* pane, click on the **Policy** tab.

Step 4 In the **Policy Control Enforcement Preference** field, make sure that **Enforced** is set. If not, choose **Enforced** and click **Submit**.

ACI VRF must be place in the correct tenant as per acc-provision config file. To verify this, perform the following actions:

a) Check the acc-provision config file and look for a section similar to this one:

Example:

```
vrf:                  # This VRF used to create all kubernetes EPs
    name: k8s
    tenant: common          # This can be system-id or common
```

b) Ensure that the APIC VRF for your cluster (k8s in the example above) is configured in the corresponding tenant (common in the example above).

Step 5 For Pod bridge domain configuration perform the following steps:

a) In the navigation pane, expand **Bridge Domains** > ***BD-name***.
b) In the *BD-name* pane, click on the **Policy** > **General** tab.
c) In the **L2 Unknown Unicast** field, select **Hardware Proxy**.
d) In the **L3 Unknown Multicast Flooding** field, select **Flood**.
e) In the **Multi Destination Flooding** field, select **Flood in BD**.
f) Verify the **PIM** box is unchecked.
g) Verify there is no **IGMP Policy** selected in the field.
h) Verify the **ARP Flooding** box is unchecked.
i) Check the **Endpoint Dataplane Learning** box.
j) Check the **Limit to IP Learning Subnet** box.

Step 6 Repeat the Pod configuration steps for Node configuration.

Placing a Namespace or Deployment Into an EPG Does Not Work

This issue is generally caused by one of the following reasons:

- You have a typo in the EPG, tenant, or application name in the annotation
- The VMM container domain is not mapped to the EPG

The mcast-daemon Inside the aci-containers-host Fails to Start

Check the mcast-daemon log messages using the following command:

```
kubectl -n kube-system logs aci-containers-host-xxxxx mcast-daemon
```

Look for the following error message:

```
Fatal error: open: Address family not supported by protocol
```

If you see this message, ensure that IPv6 support is enabled in the kernel. IPv6 must be enabled in the kernel for the mcast-daemon to start.

Connectivity

Ensure your servers are cabled as described in *Cisco ACI and OpFlex Connectivity for Orchestrators* at: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_Cisco_ACI_and_OpFlex_Connectivity_for_Orchestrators.html

Note No other connectivity models are supported and will result in intermittent or non-connectivity.

Check that the API server advertisement addresses use the node subnet, and that the nodes are configured to route all Kubernetes subnets over the node uplink.

Typically, the API server advertisement address is pulled from the default route on the node during installation. If you are putting the default route on a different network than the node uplink interfaces, you should do so, in addition to configuring the subnets from the planning process and the cluster IP subnet used internally for Kubernetes.

Checking the ACI CNI Logs for Errors

If you are running Kubernetes, use the following commands to check the ACI CNI logs for errors:

- **kubectl -n kube-system logs aci-containers-controller-ID -c aci-containers-controller**

- **kubectl -n kube-system logs aci-containers-controller-ID -c snat-operator**

  This command is only present in ACI CNI 4.2 or later.
- **kubectl -n kube-system logs aci-containers-host-ID -c aci-containers-host**
- **kubectl -n kube-system logs aci-containers-host-ID -c opflex-agent**
- **kubectl -n kube-system logs aci-containers-host-ID -c mcast-daemon**
- **kubectl -n kube-system logs i-containers-openvswitch-ID**

If you are running OpenShift, use the following commands to check the ACI CNI logs for errors:

- **oc -n aci-containers-system logs aci-containers-controller-ID -c aci-containers-controller**
- **oc -n aci-containers-system logs aci-containers-controller-ID -c snat-operator**

  This is only present in ACI CNI 4.2 or later.
- **oc -n aci-containers-system logs aci-containers-host-ID -c aci-containers-host**
- **oc -n aci-containers-system logs aci-containers-host-ID -c opflex-agent**
- **oc -n aci-containers-system logs aci-containers-host-ID -c mcast-daemon**
- **oc -n aci-containers-system logs i-containers-openvswitch-ID**

Collecting a Cluster Report

If you must open a TAC case, attach the cluster report to the case. To generate a cluster report, run the following command:

acikubectl debug cluster-report –output *fileName*.tar.gz

Note The **acikubectl** executable is part of the `acc-provision` package that you download from cisco.com.