Cisco ACI Troubleshooting Kubernetes and OpenShift

First Published: 2018-12-13

Last Modified: 2020-02-18

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

CONTENTS

Preface

This preface includes the following sections:

Audience

This guide is intended primarily for data center administrators with responsibilities and expertise in one or more of the following:

- Virtual machine installation and administration
- Server administration
- Switch and network administration
- Cloud administration

Document Conventions

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x | y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.

Convention	Description
{x | y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.
[x {y | z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
`variable`	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Examples use the following conventions:

Convention	Description
`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
italic screen font	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:

Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Warning IMPORTANT SAFETY INSTRUCTIONS

This warning symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device.

SAVE THESE INSTRUCTIONS

Related Documentation

Cisco Cloud APIC Documentation

The Cisco Cloud APIC documentation is available at the following URL: https://www.cisco.com/c/en/us/support/cloud-systems-management/cloud-application-policy-infrastructure-controller/tsd-products-support-series-home.html

Cisco Application Policy Infrastructure Controller (APIC) Documentation

The following companion guides provide documentation for Cisco APIC:

- *Cisco APIC Getting Started Guide*
- *Cisco APIC Basic Configuration Guide*
- *Cisco ACI Fundamentals*
- *Cisco APIC Layer 2 Networking Configuration Guide*
- *Cisco APIC Layer 3 Networking Configuration Guide*
- *Cisco APIC NX-OS Style Command-Line Interface Configuration Guide*
- *Cisco APIC REST API Configuration Guide*
- *Cisco APIC Layer 4 to Layer 7 Services Deployment Guide*
- *Cisco ACI Virtualization Guide*
- *Cisco Application Centric Infrastructure Best Practices Guide*

All these documents are available at the following URL: http://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html

Cisco Application Centric Infrastructure (ACI) Documentation

The broader Cisco ACI documentation is available at the following URL: http://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html.

Cisco Application Centric Infrastructure (ACI) Simulator Documentation

The Cisco ACI Simulator documentation is available at http://www.cisco.com/c/en/us/support/cloud-systems-management/application-centric-infrastructure-simulator/tsd-products-support-series-home.html.

Cisco Nexus 9000 Series Switches Documentation

The Cisco Nexus 9000 Series Switches documentation is available at http://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/tsd-products-support-series-home.html.

Cisco Application Virtual Switch Documentation

The Cisco Application Virtual Switch (AVS) documentation is available at http://www.cisco.com/c/en/us/support/switches/application-virtual-switch/tsd-products-support-series-home.html.

Cisco ACI Virtual Edge Documentation

The Cisco Application Virtual Edge documentation is available at https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html.

Cisco ACI Virtual Pod Documentation

The Cisco Application Virtual Pod (vPod) documentation is available at https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html.

Cisco Application Centric Infrastructure (ACI) Integration with OpenStack Documentation

Cisco ACI integration with OpenStack documentation is available at http://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to apic-docfeedback@cisco.com. We appreciate your feedback.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see *What's New in Cisco Product Documentation* at: http://www.cisco.com/c/en/us/td/docs/general/whatsnew/whatsnew.html

Subscribe to *What’s New in Cisco Product Documentation*, which lists all new and revised Cisco technical documentation as an RSS feed and delivers content directly to your desktop using a reader application. The RSS feeds are a free service.

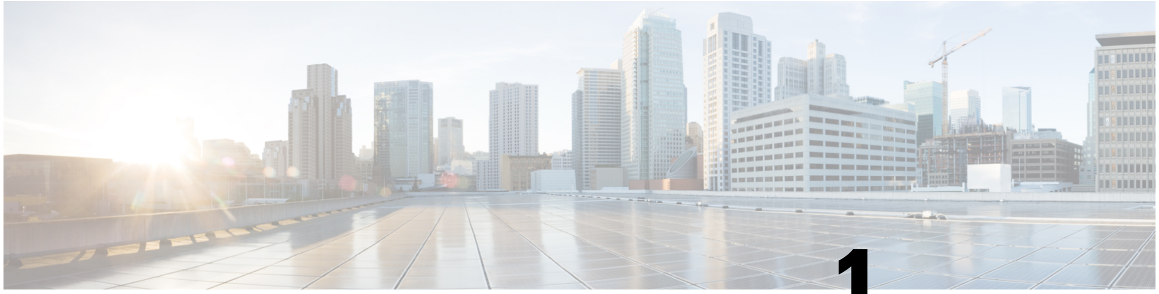

CHAPTER 1

New and Changed Information

This chapter contains the following sections:

New and Changed Information

The following table provides an overview of the significant changes to this guide up to this current release. The table does not provide an exhaustive list of all changes made to the guide or of the new features up to this release.

Table 1: New Features and Changed Behavior

Cisco APIC Release Version	Feature	Description	Where Documented
4.0(1)	--	This guide was released.	--

CHAPTER 2

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.

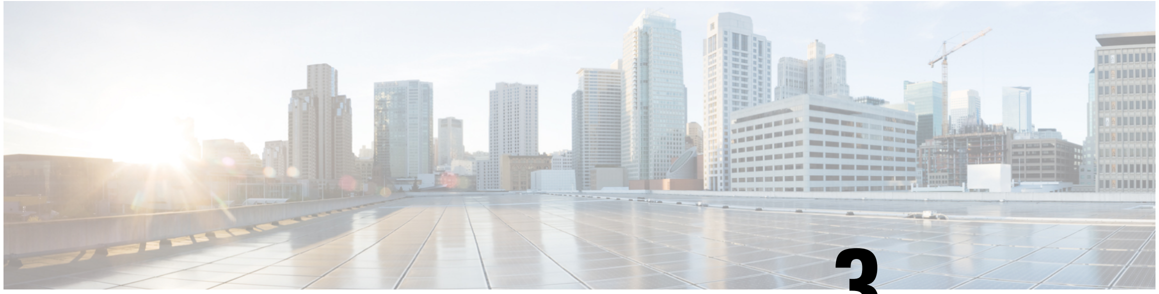

CHAPTER 3

Common Issues and Resolutions

This chapter contains the following sections:

Connectivity Refresher

This section serves a simple review of how communication between the pods and the nodes works. This is useful to troubleshoot cluster issues. The following example shows a 2-node Kubernetes deployment and how coredns traffic is reaching the kube-api:

Figure 1: 2-node Kubernetes deployment

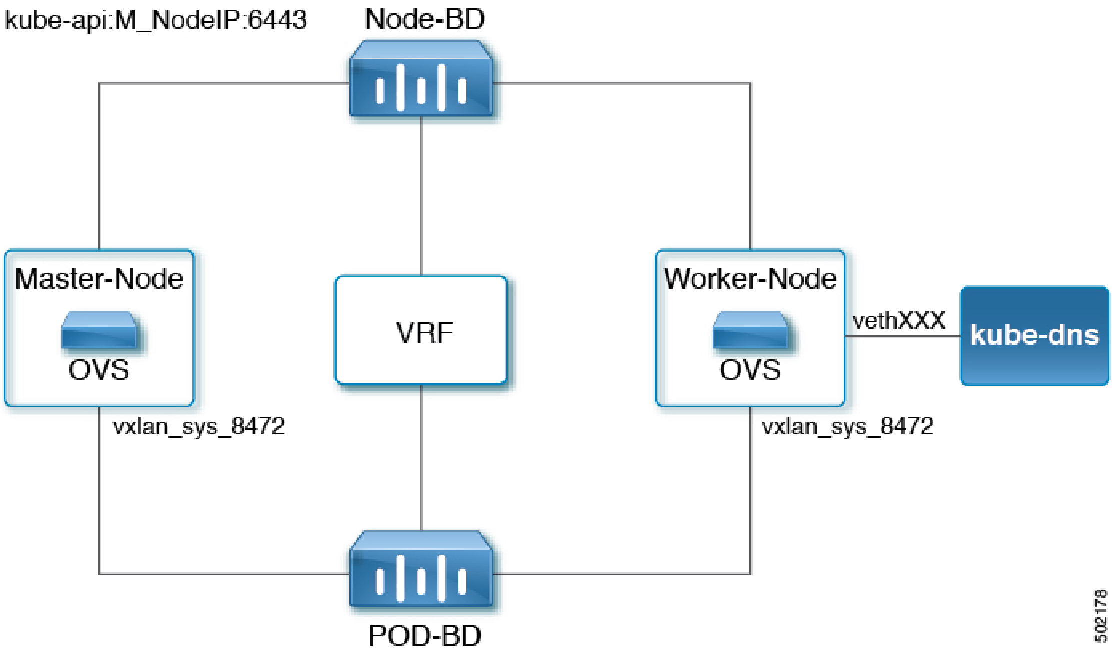

After the coredns pod comes up, the pod tries to initialize services and endpoints from the API server. The API server runs on the master nodes, listen on port 6443, and is accessible by the Node-BD subnet.

You can have multiple masters. By default Kubernetes creates a service IP address to load balance sessions between multiple masters. For example, in a 2-master configuration, you can see the following information:

```
root@k8s-01:~# kubectl --namespace=default describe service kubernetes
Name:              kubernetes
Namespace:         default
Labels:            component=apiserver provider=kubernetes
Annotations:       <none>
Selector:          <none>
Type:              ClusterIP
IP:                10.37.0.1
Port:              https 443/TCP
Endpoints:         10.32.0.11:6443,10.32.0.12:6443
Session Affinity: ClientIP
Events:            <none>
```

When coredns tries to connect to the master, coredns tries to connect to kubernetes-service-ip on port 443 (10.37.0.1 is used in the example above). A sniffer trace collected on the coredns veth*ID* would show flows initiated with the KubeDNS IP address directed to the kubernetes-service-ip on port 443.

After the traffic hits the Open vSwitch (OVS), the traffic will be destination network address translated to one of the master node API addresses. A sniffer trace collected on the vxlan_sys_8472 interface shows that the destination IP address has been changed from "kubernetes-service-ip:443" to "Master-IP:6443."

This procedure can be used to troubleshoot most services on the cluster.

CoreDNS Crash Loopback

This issue is caused most of the time by a connectivity issue between the coredns pods and Kube-API. For information on how to investigate this issue, see the Connectivity Refresher section.

ARP Is Not Resolving

If Address Resolution Protocol (ARP) is not resolving perform the following actions:

- Verify Layer 1 connectivity and that the bridge domains are configured as shown in the *Basic Checks* section.
- If you are using virtual machines (VMs), you should use **Nested** mode. Manual creation of the port group is not supported. However, if you do configure the PortGroup manually the following requirements must be met or the traffic will be dropped, either VM traffic or Opflex Control plane packets:
  - MTU of 9000 (Set this at the Virtual Switch level)
  - Forged Transmit: **Accept**
  - MAC Address change: **Accept**
  - Promiscuous Mode
    - ACI 3.2 or above: **Reject**
    - Before ACI 3.2: **Accept**
- Verify there is a static route for the 224.0.0.0 subnet pointing to the ACI Infrastructure sub-interface.

```
cisco@k8s-03:~$ route -n | grep 224
224.0.0.0   0.0.0.0    240.0.0.0       U     0      0     0 ens192.3456
```

Note If you are running multiple Kubernetes cluster on the same ACI fabric, you must configure different Multicast Fabric Wide address. The mcast_fabric parameter is located in the acc-provision config file.

- Verify if the mcast-daemon logs contain the following message:

```
Could not join group IP: No buffer space available
```

If you see this message, see the "Tune the igmp_max_memberships kernel parameter" step of the "Preparing the Kubernetes Nodes" procedure in the *Cisco ACI and Kubernetes Integration* document.

Traffic is not Reaching the Kubernetes Master Node

- Verify the node interface configuration is correct:
  - All interfaces should be configured with at least 1600 MTU.

- All the subnets used for the clusters needs to point to the ACI NODE-BD as default GW.

- Verify that destination network address translation (DNAT) is happening by taking a sniffer trace on vxlan_sys_8472 shows. If not, see *DNAT Is Not Happening*.

DNAT Is Not Happening

A trace collected on vxlan_sys_8472 shows that destination network address translation (DNAT) is not happening. This is generally cause by the following misconfiguration:

- The acc-provision config points to a VRF in Tenant X.
- The ACI fabric is configured to use the VRF in Tenant Common.

If the VRF information is not correct, Opflex can’t retrieve the EndPoints IP, and won’t be able to program OVS NAT correctly.

Currently this issue is a silent failure and no error messages nor faults are raised.

DNS Resolution Is Not Working (kubeadm)

This issue occurs on cluster installed with the **kubeadm** command. If you modify the cluster IP subnet from 10.96.0.0/16 to something else, the **kubeadm** command will not update the configuration of kubelet. This results in the container trying to resolve DNS names with the 10.96.0.10 IP address regardless if this is actually the coredns service IP address.

To check if you are hitting this issue:

1. Get the coredns service IP address:

```
cisco@k8s-01:~$ kubectl get svc -n kube-system  | grep dns
coredns                  ClusterIP   10.37.0.3     <none>        53/UDP,53/TCP
 141d
```

2. Verify that the cluster DNS IP address matches with `10.37.0.3`:

   Before Kubernetes 11.1:

```
cat /etc/systemd/system/kubelet.service.d/10-kubeadm.conf | grep KUBELET_DNS_ARGS
Environment="KUBELET_DNS_ARGS=--cluster-dns=10.37.0.3 --cluster-domain=cluster.local"
```

   Kubernetes 1.11 or newer:

```
cat /var/lib/kubelet/config.yaml | grep -A1 clusterDNS
clusterDNS: - 10.96.0.10
```

3. If there is a mismatch, you must edit the `/etc/systemd/system/kubelet.service.d/10-kubeadm.conf` file on all of the nodes (workers and master nodes) and enter the correct IP address. Then, restart kubelet:

```
systemctl daemon-reload
systemctl restart kubelet
```

This issue might be resolved in a newer kubeadm version: https://github.com/kubernetes/kubeadm/issues/28

External Services Are Not Working

The following issues are the most common that you can expect to have:

- Verify that you **HAVE NOT** created a sub-interface for the service_vlan.
- Ensure that the service VLAN is trunked all the way to the host or virtual machine.
- If you are running a version earlier than Cisco APIC release 3.2, and the node is a VM behind a virtual switch, make sure that promiscuous mode is enabled on the virtual switch port group.
- Make sure that node_svc_subnet value in acc-provision is not the same as the kubernetes service-cidr (10.97.0.0 is the default value in kubeadm). These subnets must be different.
- Make sure that your client IP address is not part of the same subnet used for the L3Out interfaces. This configuration is unsupported and will not work.

  Note If you did not configure Cisco APIC to dynamically advertise your extern_dynamic and extern_static subnets you should configure your external router with static routes for those subnets pointing to your Cisco ACI border leaf switches.

- Verify the interfaces configured in the Layer 4 to Layer 7 service devices are configured with the correct physical interfaces. As of Cisco APIC release 4.2, port channel is not supported and will result in no concrete interface being programmed. vPC and single uplink are the only supported options.

Source NAT is Not Working

Source NAT replies on the same policy-based redirect configuration that is deployed for the external services, and so some of the verification steps will be the same.

Step 1 Verify that you have not created a sub-interface for the service_vlan.

Step 2 Ensure the service VLAN is trunked all the way to the host or virtual machine.

Step 3 Verify that the interfaces configured in the Layer 4 to Layer 7 service devices are configured with the correct physical interfaces. As of Cisco APIC release 4.2, port channel is not supported and will result in no concrete interface being programmed. vPC and single uplink are the only supported options.

Step 4 Verify that the SNAT policy "State" is "Ready" with either of the following commands:

- **`kubectl describe snatpolicy`** ***`name`***
- **`oc describe snatpolicy`** ***`name`***

Example:

```
kubectl describe snatpolicy cluster-snat
Name:         cluster-snat
Namespace:
Labels:       <none>
Annotations:  kubectl.kubernetes.io/last-applied-configuration:
```

```
{"apiVersion":"aci.snat/v1","kind":"SnatPolicy","metadata":{"annotations":{},
  "name":"cluster-snat"},"spec":{"snatIp":["10.20.30.1"]}}
API Version:  aci.snat/v1
Kind:         SnatPolicy
Metadata:
  Creation Timestamp:  2019-11-04T01:59:38Z
  Finalizers:
    finalizer.snatpolicy.aci.snat
  Generation:        2
  Resource Version:  29789719
  Self Link:         /apis/aci.snat/v1/snatpolicies/cluster-snat
  UID:               c1dfdb27-fea6-11e9-b10f-005056aa4c92
Spec:
  Selector:
  Snat Ip:
    10.20.30.1. <= This is the IP configured for SNAT
Status:
  Snat Ports Allocated: <= This show the per-node port allocation for SNAT
    10 . 20 . 30 . 1:
      Nodename:  fab2-k8s-5
      Portrange:
        End:     7999
        Start:   5000
      Nodename:  fab2-k8s-4
      Portrange:
        End:     10999
        Start:   8000
      Nodename:  fab2-k8s-3
      Portrange:
        End:     13999
        Start:   11000
      Nodename:  fab2-k8s-1
      Portrange:
        End:     16999
        Start:   14000
      Nodename:  fab2-k8s-2
      Portrange:
        End:    19999
        Start:  17000
  State:        Ready <= Anything different from Ready is an issue
Events:         <none>
```

Step 5 If the SNAT policy state is not "Ready," check the `snat-operator` logs for error messages. A common issue is a typo in the snatpolicy configuration.

opflex-agent ELOCATION Errors

These opflex-agent ELOCATION errors generally point to unsupported connectivity models. If you configure the leaves in a vPC pair, but the host are not dual homed (single/PC uplink) you will see the following errors

```
kubectl -n kube-system logs aci-containers-host-tkcp9 opflex-agent
===== SNIP ====
[info] [active_connection.cpp:54:create] 10.1.16.68:8009 ç “Wrong” leaf
[info] [active_connection.cpp:54:create] 10.1.16.65:8009 ç Correct leaf
[info] [OpflexPEHandler.cpp:130:ready] [10.1.16.65:8009] Handshake succeeded
[error] [OpflexHandler.cpp:77:handleError] [10.1.16.68:8009] Remote peer returned error
with message (1,Send Identity): ELOCATION: ELOCATION
===== SNIP ====
```

Connectivity Issues When Adding A New Cluster

By default ACC provision uses 225.1.2.3 as Fabric Wide Multicast Address. If you deploy multiple cluster on the same fabric, you must change this to a unique value per cluster.

You can verify this in the acc-provision configuration file. For example:

```
  vmm_domain:                # Kubernetes VMM domain configuration
  encap_type: vxlan          # Encap mode: vxlan or vlan
 mcast_range:                # Every opflex VMM must use a distinct range
       start: 225.22.1.1
         end: 225.22.255.255
mcast_fabric: 225.1.2.4      # Every opflex VMM must use a unique address
```

Note This issue can also be present if you deploy any other solution (AVE/AVE/OpenStack) that uses a Fabric Wide Multicast Address.

Not Getting An IP Address On The Infra Interface

dhcp-client-identifier is configured as per Opflex specifications:

```
01:[Interfaec Mac address]
```

If you are running Ubuntu the location of the file is /etc/dhcp/dhclient.conf.

If you run RedHat or Centos the location of the file is /etc/dhcp/dhclient-[ifname].conf.

For example:

/etc/dhcp/dhclient-ens224.3456.conf

In the example below replace [IF_MAC] with the actual mac of your interface that connects to the ACI fabric:

```
send dhcp-client-identifier 01:[IF_MAC];
request subnet-mask, domain-name, domain-name-servers, host-name;
send host-name = gethostname();
option rfc3442-classless-static-routes code 121 = array of unsigned integer 8;
option ms-classless-static-routes code 249 = array of unsigned integer 8;
option wpad code 252 = string;
also request rfc3442-classless-static-routes;
also request ms-classless-static-routes;
also request static-routes;
also request wpad;
also request ntp-servers;
```

Cannot Access the Kubernetes Dashboard

Access to the Kubernetes dashboard is denied by default. Even when an account is enabled, you can access the dashboard only if you access it directly from the master. Complete the steps in this procedure to access the dashboard from the master.

Step 1 Create the service user:

Example:

```
cat <<EOF | kubectl create -f -
apiVersion: v1
kind: ServiceAccount
metadata:
  name: admin-user
  namespace: kube-system
EOF
```

Step 2 Create `ClusterRoleBinding`:

Example:

```
cat <<EOF | kubectl create -f -
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: admin-user
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cluster-admin
subjects:
- kind: ServiceAccount
  name: admin-user
  namespace: kube-system
EOF
```

Step 3 Copy and save a Bearer Token, which you need to authenticate to the dashboard:

Example:

```
kubectl -n kube-system describe secret $(kubectl -n kube-system get secret | grep admin-user | awk
'{print $1}')
```

Step 4 Access the dashboard.

It is a Kubernetes best practice to expose the dashboard only on the master node on the local host address. If you do not have a GUI installed on the master node, you can use an SSH tunnel to access the master node, as in the following example:

```
k8s-master# kubectl proxy
Starting to serve on 127.0.0.1:8001

remote-host# ssh -L 9000:localhost:8001 <user>@<kubernetes_master>
```

After you complete the previous command, you can access the dashboard at the following URL:

```
http://localhost:9000/api/v1/namespaces/kube-system/services/https:kubernetes-dashboard:/proxy/
```

What to do next