Cisco Application Policy Infrastructure Controller Container Plugins Release 4.0(2), Release Notes

This document describes the features, caveats, and limitations for the Cisco Application Policy Infrastructure Controller (APIC) Container Plugins.

Cisco ACI CNI Plugin is used to provide network services to Kubernetes, Red Hat OpenShift, Cloud Foundry, and Pivotal Cloud Foundry clusters on a Cisco ACI fabric. It allows the cluster pods to be treated as fabric endpoints in the fabric integrated overlay, as well as providing IP Address Management (IPAM), security and load balancing services.

Release notes are sometimes updated with new information about restrictions and caveats. See the following website for the most recent version of this document:

https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html

Table 1 shows the online change history for this document.

Table 1 Online History Change

Date	Description
December 21, 2018	Release 4.0(2) became available.
January 10, 2018	In the Resolved Caveats section, added bug CSCvn80621.

Contents

This document includes the following sections:

Cisco ACI Virtualization Compatibility Matrix

For information about Cisco ACI supported Container Products, see the *Cisco ACI Virtualization Compatibility Matrix* at the following URL:

- https://www.cisco.com/c/dam/en/us/td/docs/Website/datacenter/aci/virtualization/matrix/virtmatrix.html

New and Changed Information

This section lists the new and changed features in this release and includes the following topics:

New Software Features

The following are the new software features for this release:

- Support for Kubernetes 1.12.
- The scope of an APIC contract created for a loadbalancer service can be controlled by adding an annotation to the Kubernetes service resource definition as follows:

```
"metadata": {
  "annotations": {
    "opflex.cisco.com/ext_service_contract_scope" : "tenant"
  }
}
```

  Valid scope values are : context, tenant, global. Defaults is context (VRF only).

- The acc-provision script now provides the following additional options:
  - To automatically provision access to external network from Pod EPGs. This can be achieved by specifying the "allow_pods_external_access" under the "kube_config" section in the acc-provision script input file and setting it to True. By default the security posture of the ACI CNI plugin is to disallow access to kubernetes network and external access to pod network. In previous versions, that policy could be changed in APIC after the provisioning step. With this release, acc-provision allows you to make those changes as part of provisioning. Use the following options in "kube_config" section of the input file to control that access:
    - allow_pods_kube_api_access - Give Pods access to kube API (default False)
    - allow_pods_external_access - Give Pods access to external network (default False)
  - To set the Pod subnet's chunk size (i.e the size of the pool of IPs that is allocated to each node for assigning to Pods, and drawn from the Pod subnet). This can be achieved by specifying the "pod-subnet-chunk-size" under the "net_config" section in the acc-provision script input file. If not specified the default is 32.

- The cluster-report generated from running the acikubectl command now includes a description of the state relevant Kubernetes resources like pods, daemonsets etc which are elevant to the ACI CNI plugin.

Note: There are no changes to the Cloud Foundry and Pivotal Cloud Foundry support from the previous release. No new software packages are being posted in this release, instead use those posted for the previous release 3.2.(2).

Changes In Behavior

This section lists changes in behavior in this release.

- This release requires an ACI software version of at least 3.2(4e) or 4.0(2c). If you are using VLAN encap type for container networking, you must use 3.2(4e) as 4.0(2c) has the following issues: CSCvn77226.
- If you are going to upgrade, you must upgrade the Cisco ACI fabric first before upgrading the Cisco APIC Container plugins. The only exception is for the Cisco ACI fabric releases that have been explicitly validated for this specific plugin version in the Cisco ACI Virtualization Compatibility Matrix.

Known Limitations

This section lists the known limitations.

- The Cisco ACI CNI Plugins are not integrated with the Multi-Site Orchestrator. When deploying to a Multi-Site deployment, the ACI Configurations implemented by the plugins must not be affected by the Multi-Site Orchestrator.

Usage Guidelines

- The ACI CNI Plugin is supported with the following container solutions:
  - Canonical Kubernetes on Ubuntu 16.04
  - Red Hat Openshift on RHEL 7
  - Pivotal Cloud Foundry
- You should be familiar with installing and using Kubernetes or OpenShift. The CNI plugin (and the corresponding deployment file) is provided to enable networking for an existing installer such as kubeadm or KubeSpray. Cisco ACI does not provide the Kubernetes or Openshift installer.
- The ACI CNI plugin implements various functions running as containers inside pods. The released images for those containers for a given version are available on dockerhub under user noiro. A copy of those container images and the RPM/DEB packages for support tools (acc-provision and acikubectl) are also published on www.cisco.com.
- OpenShift has a tighter security model by default and many off-the-shelf Kubernetes applications such as guestbook may not run on OpenShift (if, for example, they run as root or open privileged ports like 80).

  Please refer to the following for details:

  https://blog.openshift.com/getting-any-docker-image-running-in-your-own-openshift-cluster/
- The ACI CNI Plugin is not responsible for any configuration on OpenShift cluster or pods when it comes to working behind a proxy. Running OpenShift "oc new-app" for instance may require access to GitHub and if the

proxy settings on the OpenShift cluster are not correctly set, this may fail. Ensure your proxy settings are correctly set.

- In this release, the maximum supported number of PBR based external services is 200 VIPs. Scalability is expected to increase in upcoming releases.

  NOTE: With OpenShift master nodes and router nodes will be tainted by default and you might see lower scale than an upstream Kubernetes install on the same hardware.

- The Cisco ACI OpenStack and CNI Plugins are not integrated with the Multi-Site Orchestrator. When deploying to a Multi-Site deployment, the ACI Configurations implemented by the plugins must not be affected by the Multi-Site Orchestrator.

- For OpenShift, the external IP used for the LoadBalancer service type is automatically chosen from the subnet pool specified in the ingressIPNetworkCIDR configuration in the /etc/origin/master/master-config.yaml file. This subnet should match the extern_dynamic property configured in the input file provided to acc_provision script. If a specific IP is desired from this subnet pool, it can be assigned to the "loadBalancerIP" property in the LoadBalancer service spec. For more details refer to OpenShift documentation here:

  https://docs.openshift.com/container-platform/3.9/admin_guide/tcp_ingress_external_ports.html#unique-external-ips-ingress-traffic-configure-cluster

  Note: The extern_static subnet configuration in the acc_provision's input is not used for OpenShift.

Supported Scale

The Kubernetes, OpenShift, Cloud Foundry, and Pivotal Cloud Foundry Platform scale limits are as follows:

Limit Type	Maximum Supported
Hosts/Leaf	40
VPC links/Leaf	40
Endpoints[1]/Leaf	2000
Endpoints/Host	400
Virtual Endpoints[2]/Leaf	40000

[1] An Endpoint corresponds to a container's network interface

[2] Total Virtual Endpoints on a leaf can be calculated as:

Virtual Endpoints / leaf = VPCs x EPGs

Where:

VPCs is the number of VPC links on the switch in the Attachment Profile used by the Openstack VMM.

EPGs is the number of EPGs provisioned for the Openstack VMM.

For the CLI verified scalability limits, see the Cisco NX-OS Style Command-Line Interface Configuration Guide for this release.

Caveats

This section contains lists of open and resolved caveats and known behaviors.

Open Caveats

This section lists the open caveats. Click the bug ID to access the Bug Search tool and see additional information about the bug.

Open Caveats in the 4.0(2) Release

The following are resolved caveats in the 4.0(2) release.

Table 2 Open Caveats in the 4.0(2) Release

Bug ID	Description
CSCvn77226	In the VLAN ecnap mode some pod endpoints do not get discovered after container bring up.

Resolved Caveats

This section lists the resolved caveats. Click the bug ID to access the Bug Search tool and see additional information about the bug.

Resolved Caveats in the 4.0(2) Release

The following are resolved caveats in the 4.0(2) release.

Table 3 Resolved Caveats in the 4.0(2) Release

Bug ID	Description
CSCvn72953	Duplicate pod IPs.
CSCvn80621	Incorrect container image numbers in acc-provision output.

Known Behaviors

This section lists caveats that describe known behaviors. Click the Bug ID to access the Bug Search Tool and see additional information about the bug.

Known Behaviors in the 4.0(2) Release

The following are known behaviors in the 4.0(2) release.

Table 4 Known Behaviors in the 4.0(2) Release

Bug ID	Description
CSCvn13789	ACI CNI plugin does not support N/S load-balancer for pods hosted on UCS-B with FI connectivity or for VMs in nested mode that can vmotion.

Related Documentation

The Cisco Application Policy Infrastructure Controller (APIC) documentation can be accessed from the following website:

https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html

The documentation includes installation, upgrade, configuration, programming, and troubleshooting guides, technical references, release notes, and knowledge base (KB) articles, as well as other documentation. KB articles provide information about a specific use case or a specific topic.

By using the "Choose a topic" and "Choose a document type" fields of the Cisco APIC documentation website, you can narrow down the displayed documentation list to make it easier to find the desired document.

New Documentation

There are no new Cisco APIC product documents for this release.