



Cisco Application Policy Infrastructure Controller Container Plugins Release 4.0(1), 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.

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:

$$\text{Virtual Endpoints / leaf} = \text{VPCs} \times \text{EPGs}$$

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.

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.

Contents

Table 1 Online History Change

Date	Description
November 2, 2018	Release 4.0(1) became available.

Contents

This document includes the following sections:

- [Cisco ACI Virtualization Compatibility Matrix](#)
- [New and Changed Information](#)
- [Known Limitations](#)
- [Usage Guidelines](#)
- [Caveats](#)
- [Related Documentation](#)
- [New Documentation](#)

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](#)
- [Changes In Behavior](#)

New Software Features

The following are the new software features for this release:

Table 2 Software Features, Guidelines, and Restrictions

Feature	Description	Guidelines and Restrictions
---------	-------------	-----------------------------

Known Limitations

OpenShift nested in Openstack (KVM)	Added support for provisioning OpenShift with the ACI CNI Plugin when running nested in Red Hat Openstack clusters that use the ACI Neutron plugin. For more information, see the Cisco ACI and OpenShift Integration KB article.	None.
-------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------

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(3n). Once 3.2(4) is available, it is strongly recommended that 3.2(4) be used instead of 4.0(1) as the following opflex issues exist in ACI fabric 4.0(1): CSCvm96379, CSCvm87337.
- 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).

Caveats

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.
- The acc-provision script now provides an option to set the MTU size for the container interfaces. This can be **achieved by specifying the "interface_mtu" under the "netconfig_section" in the acc-provision script input file.** The default value for this configuration is 1600, and you can choose between a minimum MTU size of 1280 (to allow for IPv6 headers) and a maximum of 8900 (to allow for VXLAN headers).
- The --list-flavors option for the acc-provision script now also shows flavor **options that have "Pre-release" and "Experimental" status (in addition to the ones that are currently supported).** "Pre-release" flavors are tested and soon to be released, whereas "Experimental" flavors are being actively tested.
- 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.

Caveats

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

- [Open Caveats](#)
- [Resolved Caveats](#)
- [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.

Caveats

Open Caveats in the 4.0(1) Release

There are no open caveats in the 4.0(1) Release.

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(1) Release

The following are resolved caveats in the 4.0(1) release.

Table 3 Open Caveats in the 4.0(1) Release

Bug ID	Description
CSCvm79087	Incorrect and redundant pod annotations for IP address pool are observed in Openshift deployment.
CSCvm58917	Underload opflex proxy can crash (core dump).
CSCvm09583	CPU utilization on the leaf switch that is attached to the OpenStack compute/controller has high CPU utilization when the number of endpoints increases.
CSCvk08051	In an OpenStack deployment, after a leaf upgrade the EP is not learned on the data-path. This happens as opflex-proxy is not able to resolve the modified EPG on the leaf and that happens because the corresponding object on the leaf (PD) is not in sync with the same object on the APIC (PM).
CSCvj41914	An OpflexP core is seen on the leaf switch or spine switch. The leaf switch or spine switch recovers from this, and there should be no impact other than this core being generated and the the service being restarted.

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(1) Release

The following are known behaviors in the 4.0(1) release.

Table 4 Known Behaviors in the 4.0(1) 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.

- The kube-dns crashes sometimes goes into a “CrashBackoffLoop” either on account of panic raised in side-car container (<https://github.com/kubernetes/dns/issues/195>), or the DNS container network connectivity is not

Related Documentation

complete withing default time-out for service bringup/health-check. This can be worked around by editing the kube-dns deployment to increase all "timeoutSecond" values to larger than the default 5 seconds, to say 30 seconds.

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.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

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.

© 2018 Cisco Systems, Inc. All rights reserved.