Cisco Application Policy Infrastructure Controller OpenStack Plugins Release 4.0(1), Release Notes

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

Cisco APIC OpenStack Plugin is used to deploy and operate OpenStack installations on a Cisco ACI fabric. It allows dynamic creation of networking constructs to be driven directly from OpenStack, while providing additional visibility and control from the Cisco APIC.

For the verified scalability limits (except the CLI limits), see the Verified Scalability Guide for this release. For the OpenStack Platform Scale Limits:

Limit Type	Maximum Supported
Number of OpFlex hosts per leaf	40
Number of vPC links per leaf	40
Number of endpoints per leaf	2,000
Number of endpoints per host	400
Number of virtual endpoints per leaf	40,000

Notes:

1. An endpoint is defined as one of the following:
    - A VM interface (also known as vnic),
    - A DHCP agent's port in Openstack (if in DHCP namespace on the network controller), or
    - A floating IP address
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.

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

Contents

This document includes the following sections:

Cisco ACI Virtualization Compatibility Matrix

For information about Cisco ACI and OpenStack, 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:

Table 2 Software Features, Guidelines, and Restrictions

Feature	Description	Guidelines and Restrictions
OpenStack Queens	Added support for OpenStack Queens release for Canonical distribution with Juju.	None.
Nested Domains	Nested Domains for virtualization. An extension is added to the Network resource in neutron to supported nested virtualization workflows, such as OpenShift on OpenStack.	None.
SNAT Conntrack Logging	Provides output to syslog for tracking when conntrack is used for SNAT.	None.

Changes In Behavior

This section lists changes in behavior in this release.

- For OpenStack Director installs, the value for ACIOpflexUplinkInterface parameter needs to be an actual interface name. This is required in order to support both nested virtualization and non-nested configurations. Refer to the appropriate OpenStack Director documentation for additional information on how to configure this for your environment.
- 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.
- Starting in 4.01, agent-ovs was renamed opflex-agent. Operators must account for the change when stopping or starting the agent. Users who create their own installers also need to incorporate packaging changes for the agent.

  In addition, the default values for two sockets used by the agent have changed:

  Old: `/var/run/opflex-agent-ovs-inspect.sock`

  New: `/var/run/opflex-agent-inspect.sock`

  Old: `/var/run/opflex-agent-ovs-notif.sock`

  New: `/var/run/opflex-agent-notif.sock`

  The neutron-opflex-agent shares the notify socket with the opflex-agent, so its default value also changed in order to be consistent. All socket filenames can also be configured explicitly.
- If you are going to upgrade, you must upgrade the Cisco ACI fabric first before upgrading the Cisco APIC OpenStack 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.
- Starting in release 3.1(1) for OpenStack, the following changes were made to the unified plugin:
  - Adds support for the OpenStack Ocata release
  - Moves security group implementation from IPtables to OVS
  - Improves support for multiple OpenStack instances on the same CiscoAPIC cluster

- Security Groups for Opflex hosts are implemented natively in OVS, instead of using IPtables rules.

  If you are using an installer plugin distributed with this code, the appropriate configuration of Opflex hosts is automatically done. If you have your own installer, this change requires the following changes to the bridge configuration on all Opflex hosts:

  1. Create the br-fabric bridge, enter the following commands:

     ```
     # ovs-vsctl add-br br-fabric
     # ovs-vsctl set-fail-mode br-fabric secure
     ```

  2. Add a vxlan port to the br-fabric, enter the following command. The br-int_vxlan0 vxlan port on the br-int bridge is no longer needed and can be removed.

     ```
     # ovs-vsctl add-port br-fabric br-fab_vxlan0 -- set Interface br-fab_vxlan0 type=vxlan
     options:remote_ip=flow options:key=flow options:dst_port=8472
     ```

  3. Change the agent-ovs config file:

     ```
     "renderers": {

     "stitched-mode": {

     //"ovs-bridge-name": "br-int", <=== Remove this line.

     "int-bridge-name": "br-fabric", <=== Add this line.

     "access-bridge-name": "br-int", <=== Add this line.

     "encap": {

               "vxlan" : {

                   //"encap-iface": "br-int_vxlan0", <=== Change from br-int to br-fab.

                   "encap-iface": "br-fab_vxlan0",

                   "uplink-iface": "eth1.4093",

                   "uplink-vlan": 4093,

                   "remote-ip": "10.0.0.32",

                   "remote-port": 8472

               }

          },
     ```

- Multiple OpenStack instances can share the same Cisco ACI fabric. Earlier versions of unified plugin would attach all OpenStack VMM domains to every OpenStack cluster. This release allows cleaner separation by using this procedure:

  You must provision the VMM domains owned by each openstack instance using the new host-domain-mapping CLI command:

  ```
  # aimctl manager host-domain-mapping-v2-create [options] <host name> <domain name> <domain
  type>
  ```

The host name can be a wildcard, which is indicated using an asterisk surrounded by double quotes ("*"). A wildcard means that the mapping should be used for all hosts. When more than one OpenStack instance shares the fabric, an entry must be created in this table for each VMM domain in use by that OpenStack instance. As an example, if one OpenStack instance is using VMM Domains "ostack1" and "ostack2", the following commands would be run on that OpenStack controller to put entries to this table:

```
# aimctl manager host-domain-mapping-v2-create "*" ostack1 OpenStack
# aimctl manager host-domain-mapping-v2-create "*" ostack2 OpenStack
```

If the second OpenStack instance is using VMM Domain "ostack3", the following command would be run on that OpenStack controller to add an entry to its table:

```
# aimctl manager host-domain-mapping-v2-create "*" ostack3 OpenStack
```

- Earlier versions only supported one logical uplink for hierarchical port binding or non-opflex VLAN network binding. In this release, you can have multiple links for those use-cases when using unified plugin.

  In order to use this feature, the AIM CLI has to be used to provide the mapping of physnets in OpenStack and an interface on a specific host. The following aimctl CLI command is used to configure this mapping:

  ```
  # aimctl manager host-link-network-label-create <host_name> <network_label> <interface_name>
  ```

  As an example, host h1.example.com is provisioned to map its eth1 interface to physnet1:

  ```
  # aimctl manager host-link-network-label-create h1.example.com physnet1 eth1
  ```

- Previously it was not possible for a single L3 Out to be shared across multiple OpenStack instances when using AIM, due to the fact that both OpenStack instances would attempt to use an External Network Endpoint Group of the same name. This release adds scoping of the Application Profile for the External Network Endpoint Group using the apic_system_id, which is configured in the `[DEFAULT]` section of the neutron configuration file.

- In earlier versions, the AIM plugin would take ownership of pre-existing L3 Outs when NAT was not being used, which led to scenarios where the AIM plugin would delete the pre-existing L3 Out in some corner cases. With this release, the AIM plugin will not take ownership of any pre-existing L3 Outs.

- Legacy plugin is not supported with the Ocata Plugins and will not be supported on future versions of OpenStack. The legacy plugin for Newton is supported. All customers are recommended to use unified mode for both Newton and Ocata.

- The OpFlex agent does not support client authentication. This means that the SSL certificate check must be disabled in Cisco APIC GUI.

  1. In the Cisco APIC GUI, on the menu bar, choose System > System Settings > Fabric Wide Setting.
  2. Ensure that the OpFlex Client Authentication check box is not checked.

Known Limitations

This section lists the known limitations.

- GBP and ML2 Unified Mode does not have feature parity with the ealier non-unified mode. In particular, it does not support the following features:
  - ESX hypervisor support
  - ASR1K edgeNAT support

  - GBP/NFP Service chaining
  - ML2 Network constraints
- Not all Unified mode features are support by the legacy plugin:
  - Support for OpenStack address scops
  - OpenStack address scopes are supported only in the Unified mode (where they are mapped to VRFs in the Unified model) and are not support in the earlier configurations.
  - Dual stack IPv6 deployment
- GBP and ML2 Unified Mode is a new mode of operation. So, while there can be a manual transition to this mode of usage, there is no automated upgrade from previous install to this mode.
- Dual-stack operation requires that all IPv4 and IPv6 subnets - both for internal and external networks - use the same VRF in Cisco ACI. The one exception to this is when separate external networks are used for IPv4 and IPv6 traffic. In that workflow, the IPv4 and IPv6 subnets used for internal networks plus the IPv6 subnets used for external networks all belong to one VRF, while the subnets for the IPv4 external network belong to a different VRF. IPv4 NAT can then be used for external networking.
- 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.

Usage Guidelines

- NFV features, including SVI networks, trunk ports, and Service Function Chaining plugin and workflow, are supported starting with the Ocata release of the plugin.
- JuJu charms users must first update the Charms before installing the updated plugin.
- When using the allowed address pair feature with the Cisco ACI plugin, be aware of the following differences from upstream implementation:
  - The Cisco ACI plugin only supports the host (/32) specification for allowed_address_pair, not the CIDR/subnet specification.
  - As OpenStack allows the same allowed_address_pair to be configured on multiple interfaces for HA, the OpFlex agent requires that the specific VNIC that currently owns a specific allowed_address_pair to assert that address ownership using Gratuitous ARP.
  - When using the promiscuous mode, the vSwitch stops enforcing the port security check. To get reverse traffic for a different IP or MAC address, you still need to use the allowed-address-pair feature. If you are running tempest, you will see test_port_security_macspoofing_port fail in scenario testing, as that test does not use the allowed-address-pair feature.
- If you are using SLAAC, add a security group rule to allow ICMPv6 to the effected Neutron networks. For example, the following security group (ipv6-sg) allows the required traffic:

```
# openstack security group rule create --ethertype IPv6 --ingress --protocol 58 --src-ip ::/0
\ ipv6-sg
```

- Before performing an upgrade from 3.1(1) using OpenStack Director or attempting a Cisco APIC ID recovery procedure, all AIM processes on all controllers need to be shutdown. To shutdown all the AIM processes on all controllers, run the following command on the undercloud:

```
for IP in $(nova list | grep ACTIVE | sed 's/.*ctlplane=//' | sed 's/ |//') ; do

ssh -o UserKnownHostsFile=/dev/null -o StrictHostKeyChecking=no heat-admin@$IP \

"sudo systemctl stop aim-event-service-rpc; sudo systemctl stop aim-aid; sudo systemctl stop
aim-event-service-polling" ;

done
```

  If upgrading, you do not need to explicitly restart the AIM processes as the upgrade will automatically restart them.

  If attempting a Cisco APIC ID recovery, you must restart the AIM processes on all the controllers manually after ID Recovery is complete.

- Keystone configuration update

  When the OpenStack plugin is installed in the unified mode, the Cisco installer adds the required configuration for keystone integration with AIM. When not using unified mode, or when using your own installer, the configuration section must be provisioned manually:

```
[apic_aim_auth]

auth_plugin=v3password

auth_url=http://<IP Address of controller>:35357/v3

username=admin

password=<admin_password>

user_domain_name=default

project_domain_name=default

project_name=admin
```

- When using optimized DHCP, the DHCP lease times are set by the configuration variable apic_optimized_dhcp_lease_time under the `[ml2_apic_aim]` section.
  - This requires a restart of neutron-server to take effect
  - If this value is updated, existing instances will continue using the old lease time, provided their neutron port is not changed (e.g. rebooting the instance would trigger a port change, and cause it to get the updated lease time). New instances will however use the updated lease time.

- In upstream Neutron, the "advertise_mtu" option has been removed.

  Since the aim_mapping driver still uses this configuration, the original configuration which appeared in the default section should be moved to the aim_mapping section. For example:

```
[aim_mapping]

advertise_mtu = True
```

  It is set to True by default in the code (if not explicitly specified in the config file).

- The Unified Plugin allows coexistence of GBP and ML2 networking models on a single OpenStack Cloud installation. However, they must operate on different VRFs. We recommend to use a single model per OpenStack Project.
- Unified mode has features not supported by the legacy plugin:
  - Support for Openstack address scopes and subnetpools

    OpenStack address scopes and subnetpools are supported only in the Unified mode (where they are mapped to VRFs in the unified model) and are not supported in the earlier configurations.
  - Dual stack IPv6 deployment
- If a default VRF is implicitly created for a tenant in ML2, it is not implicitly deleted until the tenant is deleted (even if it not being used anymore).
- Unified model impact of the transaction Model Updates in Newton.

  When GBP and ML2 co-exist, GBP implicitly created some neutron resources. In Newton, the neutron transaction model has been updated and has added various checks. Some of those checks spuriously see this nested transaction usage as an error and log and raise an exception. The exception is handled correctly by GBP and there is no functional impact but unfortunately the neutron code also logs some exceptions in neutron log file – leading to the impression that the action had failed.

  While most such exceptions are logged at the DEBUG level, occasionally you might see some exceptions being logged at the ERROR level. If such an exception log is followed by a log message which indicates that the operation is being retried, the exception is being handled correctly. One such example is the following:

  Delete of policy-target on a policy-target-group associated to a network-service-policy could raise this exception:

  ```
  2017-03-18 12:52:34.421 27767 ERROR neutron.api.v2.resource […] delete failed

  2017-03-18 12:52:34.421 27767 ERROR neutron.api.v2.resource Traceback …:

  2017-03-18 12:52:34.421 27767 ERROR neutron.api.v2.resource   File "/usr/lib/python2.7/site-
  packages/neutron/api/v2/resource.py", line 84, …

  ...

  2017-03-18 12:52:34.421 27767 ERROR neutron.api.v2.resource     raise …

  2017-03-18 12:52:34.421 27767 ERROR neutron.api.v2.resource ResourceClosedError: This
  transaction is closed
  ```

  Note: Cisco is working with the upstream community for further support on Error level logs.
- When a Layer 2 policy is deleted in GBP, some implicit artifacts related to it may not be deleted (resulting in unused BDs/subnets on Cisco APIC). If you hit that situation, the workaround is to create a new empty Layer 2 policy in the same context and delete it.
- If you use tempest to validate OpenStack, the following tests are expected to fail and can be ignored:

  ```
  tempest.scenario.test_network_basic_ops.TestNetworkBasicOps.test_update_router_admin_state
  ```
- Neutron-server logs may show the following message when DEBUG level is enabled:

```
Timed out waiting for RPC response: Timeout while waiting on RPC response - topic:
"<unknown>", RPC method: "<unknown>" info: "<unknown>"
```

This message can be ignored.

- High Availability LBaaSv2 is not supported.
- OpenStack Newton is the last version to support non-unified plugin. OpenStack Ocata and future releases will only be supported with the unified plugin.

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.

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
CSCvn01356	When using the OpenStack Neutron's trunk port feature with opflex-agent and optimized DHCP, the IP address cannot be assigned to the VLAN based sub-interface.
CSCvi95771	This is a new feature to log SNAT IP addresses used by distributed SNAT solution.
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.

Bug ID	Description
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

There are no known behaviors in the 4.0(1) release.

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.