



Cisco Application Policy Infrastructure Controller OpenStack Plug-in Release 4.1(2), Release Notes

This document describes the features, bugs, and limitations for the Cisco Application Policy Infrastructure Controller (APIC) OpenStack Plug-in.

Cisco APIC OpenStack Plug-in allows policy deployment automation across Cisco ACI and OpenStack, enabling a complete undercloud and overcloud visibility on Cisco ACI. The Cisco APIC OpenStack Plug-in allows dynamic creation of networking constructs to be driven directly from OpenStack, while providing additional visibility and control from the Cisco APIC.

Release notes are sometimes updated with new information about restrictions and bugs. 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
2019-07-08	Release 4.1(2) became available.
2019-07-10	Updated Usage Guidelines section.
2019-07-24	Updated Known Limitations section

Contents

This document includes the following sections:

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

- [New Documentation](#)

Cisco ACI Virtualization Compatibility Matrix

For information about Cisco ACI and OpenStack, see the *Cisco 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

There are no new software features in the Cisco APIC OpenStack Plug-in Release 4.1(2) release.

Changes in Behavior

This section lists changes in behavior in this release.

- Beginning in Cisco APIC OpenStack Plug-in Release 4.1(1), using VXLAN on blade server systems is supported. See the Known Limitations section for more information.
- For OpenStack Director installs, the value for `ACIOpflexUplinkInterface` parameter needs to be an actual interface name. This is required 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.
- For OpenStack Director 13 installs, enabling or disabling of LLDP is controlled by resource declaration. If you have the following in your yaml file:

```
OS::TripleO::Services::CiscoAciLldp: /opt/ciscoaci-tripleo-heat-templates/docker/services/cisco_lldp.yaml
```

Then LLDP will be enabled. If you do not want to use LLDP, then you must put the following in your yaml file:

```
OS::TripleO::Services::CiscoAciLldp: OS::Heat::None
```

The use of `ACIUseLldp` to control this behavior was removed beginning with OpenStack Director 13.

- For installations currently running release 3.2(2.20180710), 4.0(1.20181001), or 4.0(2.20181221), run the `db_check` script before upgrading to ensure that the AIM database migration script completed successfully. The script is in `support-tools-1.0.0.tar.gz` in the tarball for the release on Cisco.com.

Contact the Cisco Technical Assistance Center (TAC) if the script indicates that there could be a potential problem.

- Cisco ACI software version 3.2(4e) or higher is recommended for this plug-in. You cannot use Cisco ACI software version 4.0(2c) for OpenStack as it has the following issues with floating IP usage: CSCvn77231.

New and Changed Information

- 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 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 plug-ins. The only exception is for the Cisco ACI fabric releases that have been explicitly validated for this specific plug-in version in the [Cisco ACI Virtualization Compatibility Matrix](#).
- Multiple OpenStack instances can share the same Cisco ACI fabric. Earlier versions of unified plug-in would attach all OpenStack VMM domains to every OpenStack cloud. This release allows cleaner separation by using this procedure:

You must provision the VMM domains owned by each openstack cloud 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 plug-in.

To use this feature, the AIM CLI must 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
```

Supported Scale

- Previously it was not possible for a single L3 Out to be shared across multiple OpenStack instances when using AIM, because 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 `aimctl.conf` file.
- In earlier versions, the AIM plug-in would take ownership of pre-existing L3 Outs when NAT was not being used, which led to scenarios where the AIM plug-in would delete the pre-existing L3 Out in some corner cases. With this release, the AIM plug-in will not take ownership of any pre-existing L3 Outs.
- Legacy plug-in is not supported with the Ocata Plug-ins and will not be supported on future versions of OpenStack. The legacy plug-in 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.

Supported Scale

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

Note: The scalability information in the following table applies to the sum of OpenStack and OpenShift or Kubernetes resources integrated with OpFlex into the Cisco ACI fabric. It does not apply to Microsoft SCVMM hosts of Cisco ACI Virtual Edge instances.

Table 3 OpenStack Platform Scale Limits in the 4.1(2) Release

Limit Type	Maximum Supported
Number of OpFlex hosts per leaf	40
Number of vPC links per leaf	40
Number of endpoints per leaf	4,000
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:

Known Limitations

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.

Known Limitations

This section lists the known limitations.

- Cisco ACI Unified Plug-in for OpenStack does not support the following features:
 - ESX hypervisor support
 - ASR1K edgeNAT support
 - GBP/NFP Service chaining
 - ML2 Network constraints
- Cisco ACI Unified Plug-in for OpenStack supports OpenStack address scopes and dual stack IPv4 and IPv6 deployments.
- 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.
- For installations with B-series that use VXLAN encapsulation, Layer 2 Policies (for example, bridge domains) should each contain only one Policy Target Group (that is, Endpoint Group) to ensure a functional dataplane.
- The Cisco ACI OpenStack Plug-in is not integrated with the Multi-Site Orchestrator. When deploying to a Multi-Site deployment, the Cisco ACI configurations implemented by the plug-in must not be affected by the Multi-Site Orchestrator.
- NFV features, including SVI networks, trunk ports, and Service Function Chaining plug-in and workflow, are supported starting with the Ocata release of the plug-in.
- When you delete the Overcloud Heat stack, the Overcloud nodes are freed but the virtual machine manager (VMM) domain remains present in Cisco APIC. The VMM appears in Cisco APIC as a stale VMM domain along with the tenant unless you delete the VMM domain manually. Before you delete the VMM domain, verify that the stack has been deleted from the undercloud, and check that any hypervisors appearing under the VMM domain are no longer in the connected state. Once both of these conditions are met, you can safely delete the VMM domain Cisco APIC.

Usage Guidelines

- JuJu charms users must first update the Charms before installing the updated plug-in.

- When using the allowed address pair feature with the Cisco ACI plug-in, be aware of the following differences from upstream implementation:
 - 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 shut down 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 plug-in 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
```

Usage Guidelines

- 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 Plug-in allows coexistence of GBP and ML2 networking models on a single OpenStack Cloud installation. However, they must operate on different VRFs. We recommend using a single model per OpenStack Project.
- 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.

Bugs

- 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 plug-in. OpenStack Ocata and future releases will only be supported with the unified plug-in.
- For deployments running Cisco ACI version 4.1(2g) and using the Group Based Policy workflow and associated APIs, contract filters set to an EtherType of ARP can result in the filter being incorrectly set as “Unspecified” on the leaf. If an EtherType of ARP is required, then you must use a Cisco ACI release other than 4.1(2g).
- Some deployments require installation of an “allow” entry in IP Tables for IGMP. This must be added to all hosts running an OpFlex agent and using VXLAN encapsulation to the leaf. The rule must be added using the following command:

```
# iptables -A INPUT -p igmp -j ACCEPT
```

In order to make this change persistent across reboots, add the command either to `/etc/rc.d/rc.local` or to a cron job that runs after reboot.

- For deployments that use B-series servers, an additional service must be started on the hosts to ensure that connectivity is maintained with the leaf at all times. Complete the following steps:
 1. Install the Cisco APIC API package (python-apicapi for Debian packaging, apicapi for RPM packaging) for any servers running an OpFlex agent.
 2. Add the OpFlex uplink bond name to `/etc/environments` (that is, `opflex_bondif=bond1`).

This is needed if the interface is other than default (bond0).

3. Enable the apic-bond-watch service using the following command:

```
sudo systemctl enable apic-bond-watch
```

4. Start the apic-bond-watch service using the following command:

```
sudo systemctl start apic-bond-watch
```

Bugs

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

- [Open Bugs](#)

- [Resolved Bugs](#)
- [Known Behaviors](#)

Open Bugs

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

Table 4 Open Bugs in the 4.1(2) Release

Bug ID	Description
CSCvo57669	APICAPI ConnectionError on stack controllers after APIC upgrade

Resolved Bugs

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

Table 5 Resolved Bugs in the 4.1(2) Release

Bug ID	Description
CSCvp38627	AIM Hash Tree requires manual reset to allow SG provisioning go through
CSCvp82252	Error on creating/modifying host route of OpenStack subnet for OPS13

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.

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.

© 2019 Cisco Systems, Inc. All rights reserved.