



Migration

This chapter contains the following sections:

- [Overview, on page 1](#)
- [Guidelines and Limitations, on page 1](#)
- [Workflow for Migrating Cisco ACI OpenStack Unified Plug-in, on page 2](#)
- [Migrate the Cisco ACI OpenStack Unified Plug-in, on page 3](#)

Overview

Cisco Application Centric Infrastructure (ACI) OpenStack Newton (OpenStack Platform 10) ships with three mechanism drivers for Neutron and the ML2 plug-in:

- Legacy ML2 driver
- Legacy GBP driver
- Unified ML2 and GBP driver

For Ocata (OpenStack Platform 11) or later, the Cisco ACI OpenStack solution only ships with the unified driver. Before upgrading to a newer version of Cisco ACI OpenStack, you must update the Neutron mechanism driver to the unified one. There are some functional and interoperability differences between legacy drivers and the unified driver, which means they are not a like-for-like replacement. For the common ML2 and Neutron usage, you can use this procedure to make that migration easier.

Guidelines and Limitations

The following guidelines and limitations apply for the migration:

- Migration is supported only for the legacy ML2 driver and not for the legacy GBP driver. It is restricted to features that are supported both by the legacy plug-in and the Cisco Application Centric Infrastructure (ACI) OpenStack Unified Plug-in.
- Migration is only supported on Red Hat OpenStack Platform 10. When you migrate, you cannot combine it with any other procedure, such as an OpenStack upgrade.
- Although the upgrade process does not require any VM operations such as creation, deletion, or live migration, the procedure does stop the flow of traffic for tens of minutes. Therefore, schedule the upgrade

during a maintenance window. Also make sure that the maintenance window is long enough to perform a backup, migration, and rollback, if needed.

- This migration procedure is all-or-nothing. You must migrate the entire OpenStack cluster at the same time. You cannot do it incrementally—that is, per tenant, network, host, and so on.
- This migration maintains the hybrid data plane used by the legacy plug-in. (That is, Linux Bridge with Open vSwitch.) It may be possible to move to the nonhybrid data plane used by the unified plug-in (only Open vSwitch) as part of the migration. However, this document does not cover this procedure. Contact TAC if you require a nonhybrid data plane for your installation.
- If VLAN encapsulation is used between the hypervisors and the leaf switches, then there should be enough VLANs available so that the OpenStack clusters have a disjoint VLAN distribution before and after the migration. This is not a strict requirement, but it makes the procedure easier to revert in some cases without a restore from the backup.
- Software upgrade is not part of this process. It is strictly for migration from the legacy plug-in to the Cisco ACI OpenStack Unified Plug-in. All required software upgrades must be performed before the migration.
- Migration is only supported for Layer 3 outside network connection (L3Out) policies that are configured in the common tenant of the Cisco Application Centric Infrastructure fabric.
- The release includes a support tools package, which contains tools to facilitate the migration process. The tools in the package provide a reference to show how certain steps of the migration can be scripted.

Workflow for Migrating Cisco ACI OpenStack Unified Plug-in

This section provides a high-level list of steps that are required to migrate the Cisco Application Centric Infrastructure (ACI) OpenStack Unified Plug-in. See the section [Migrate the Cisco ACI OpenStack Unified Plug-in, on page 3](#) in this guide for details.

1. Install the group-based policy and OpenStack Cisco ACI Integration Module (AIM) packages, along with their dependencies, and run their database migration scripts.
2. Use the configuration (database and .ini files) from the legacy plug-in to create the configuration required for the Cisco ACI OpenStack Unified Plug-in.

You must perform this step because the responsibility of interacting with Cisco Application Policy Infrastructure Controller (APIC) changes from Neutron to AIM.

3. Start AIM services on all the controllers.
4. Stop all OpFlex agents (opflex-agent and neutron-opflex-agent).
5. Stop the Neutron server processes running the legacy plug-in across all controllers.
6. Run the `gbp-validate` validation/migration script (check mode only) from one controller and verify that the migration will succeed.
7. Run the validation/migration script (repair mode) from one controller.
Running the script migrates the entire legacy plug-in Neutron database state to the new state needed by the Cisco ACI OpenStack Unified Plug-in and rebinds all the ports.
8. Start all OpFlex agents (opflex-agent and neutron-opflex-agent).

9. Start the Neutron server processes that run the Cisco ACI OpenStack Unified Plug-in across all controllers.
10. Clean up the Cisco APIC state left over from the legacy plug-in.



Note You might need to perform extra steps depending on the installation and how you perform the migration.

Migrate the Cisco ACI OpenStack Unified Plug-in

Consider the following process only as a template. There may be installation-specific considerations that result in additional or different steps being taken. As an example, this process is for a site that uses VLAN encapsulation and RPM packaging, but either or both of these may be different for a given installation.

Before you begin

Before you get started, ensure that you meet the following prerequisites:

- Update the Cisco Application Centric Infrastructure (ACI) fabric to release 3.2(4) or later.



Note The 4.0 releases are not supported due to compatibility issues with the Cisco ACI OpenStack Unified Plug-in.

- Update the Cisco ACI OpenStack Unified Plug-in to the latest 4.0(x) release; continue using the existing ML2 driver and configuration.
- Install the support tool package onto the controllers.

Step 1

Install and initialize AIM service on all Neutron controller nodes.

This step gets the Neutron nodes ready to use the Cisco ACI OpenStack Unified Plug-in once it has been installed.

- a) Install AIM and other packages needed by the Cisco ACI OpenStack Unified Plug-in:

Example:

```
# yum -y install aci-integration-module acitoolkit apicapi openstack-dashboard-gbp
openstack-heat-gbp python-django-horizon-gbp python-gbpclient python-semantic_version
openstack-neutron-gbp
```

- b) Run the AIM DB migration script, using the Neutron configuration.

This step only needs to be done once, and does not need to be repeated on all of the controllers.

Example:

```
# aimctl -c /etc/neutron/neutron.conf db-migration upgrade
```

- c) Run the script that creates the AIM configuration from the existing Neutron configuration:

Example:

Migrate the Cisco ACI OpenStack Unified Plug-in

```
# support_tool aim-config --config-file /etc/neutron/neutron.conf --config-file
/etc/neutron/plugins/ml2/ml2_conf.ini
```

Inspect the /etc/aim/aim.conf and /etc/aim/aimctl.conf files. Verify that the configuration for VMM domains, physical domains, physlinks, and other states was successfully migrated from the Neutron configuration files to the AIM configuration files. If the respective state was not present in the legacy configuration files, then you must add this state manually. Add any static path bindings that need to be preserved to the configuration file.

- From a single controller, install the AIM configuration:

This command needs to be run only once.

Example:

```
# aimctl config replace
```

Step 2

Start the AIM service.

Restarting the AIM service provisions APIC resources for the new VMM domains. Depending on Layer 3 outside network connection (L3Out) usage, the restart might impact external traffic. Wait until all resources have been synchronized with Cisco Application Policy Infrastructure Controller (APIC).

Example:

```
# systemctl start aim-event-service-rpc
# systemctl start aim-event-service-polling
# systemctl start aim-aid
```

Step 3

From a single controller, create the infrastructure from AIM configuration file.

This step also creates the new VMM domain in Cisco APIC and only needs to be run once.

Example:

```
# aimctl infra create
```

Step 4

From a single controller, update AIM with the new VMM domains, the new system-id, or both.

This command needs to be run only once.

Example:

```
# aimctl manager load-domains
```

Step 5

After you have installed AIM service successfully on all Neutron nodes, shut it down.

Example:

```
# systemctl stop aim-event-service-rpc
# systemctl stop aim-event-service-polling
# systemctl stop aim-aid
```

Step 6

On all nodes, update the OpFlex Agent with changes to their VMM domain.

This step is required if you want to keep the legacy plug-in state on one VMM domain and the Cisco ACI OpenStack Unified Plug-in in a new VMM domain. The support tool **get-domains** and **update-domains** commands need passwordless SSH access to the hosts running OpFlex Agents. If this is not possible for a given installation, perform this step by hand or by a user-created script.

Note This step requires non-overlapping VLAN ranges, which are dependent on available VLAN space. You can perform this step from one of the controllers using the support tool.

- Retrieve the OpFlex agent VMM domain configuration.

Example:

```
# support_tool get-domains --credentials-file ~/keystonerc_admin
--config-directory /etc/opflex-agent-ovs/conf.d --output-file
~/opflex-agent-conf.txt
```

- b) Update the OpFlex agent configuration to use any new VMM domains, after editing the mapping file.

The default filename is: opflex_hosts.txt.

Example:

```
# support_tool update-domains --credentials-file ~/keystonerc_admin
--config-directory /etc/opflex-agent-ovs/conf.d --output-file
~/opflex-agent-conf.txt
```

Step 7

Get UIDs of the external networks that will be migrated and their names by using the following command: `# neutron net-list --router:external True -c id -f value -c name -f value`

Save the information, which you will use later in the migration process.

The handling of external networks is different between the legacy and unified plug-ins. The Cisco ACI OpenStack Unified Plug-in presumes that the L3Out policies always pre-exist, and the distinguished names (DNs) of the L3Out policies are stored as extensions in the Neutron external network resource.

Step 8

On all hosts running OpFlex agents, shut down the Neutron OpFlex agent:

Example:

```
# systemctl stop neutron-opflex-agent
```

Step 9

Shut down the Neutron server on all OpenStack controllers:

Example:

```
# systemctl stop neutron-server
```

After the Neutron server is shut down, new OpenStack workloads cannot be provisioned and existing workloads cannot be updated. The traffic for existing workloads should continue working as before.

Step 10

Back up the current state.

- Back up the current Neutron database.
- Back up the current Cisco APIC configuration.

Rollback is accomplished by restoring the Neutron database and the Cisco APIC configuration that has been backed up in this step.

Step 11

Prepare Neutron to use the Cisco ACI OpenStack Unified Plug-in.

- Keep track of any existing Cisco ACI tenants and contracts in common user used by this OpenStack instance.
- Reconfigure the Neutron configuration file to use the new plug-in:

Example:

```
# support_tool toggle-config --toggle new --config-file /etc/neutron/neutron.conf --config-file
/etc/neutron/plugins/ml2/ml2_conf.ini
```

- Update configuration file to use a different system-id, new VLAN pool, and the same attachable entity profile (AEP) as the previous OpenStack cluster.
- Provide the Neutron external networks that require migration.

Providing the Neutron external networks is done in the Neutron configuration files, using the `migrate_ext_net_dns` option under the `[ml2_apic_aim]` section. The `migrate_ext_net_dns` configuration option requires a

comma-separated list of mappings. Each mapping is a Neutron external network UUID followed by the DN of the network under an L3Out policy in Cisco APIC.

The UUID of the external networks was retrieved previously, and the DNs of the network under the L3Out policy can be retrieved from AIM using the following command:

```
# aimctl manager external-network-get -p <APIC tenant name> <L3 Out Policy Name> <Network name under L3 Out Policy> | grep resource_dn | awk '{print $2}'
```

Example:

The following command retrieves the DN for the l3out_2_net Network under the L3Out policy named l3out-2 in the common tenant:

```
# aimctl manager external-network-get common l3out-2 l3out_2_net -p | grep resource_dn | awk '{print $2}'
```

Example:

The following is an example of what this configuration would look like, given two Neutron external networks named l3out-1 and l3out-2:

```
[ml2_apic_aim]
migrate_ext_net_dns = a759a910-be33-4e09-85f0-741b6ea48c1f: uni/tn-common/out-l3out-1,
fba9b7b0-f492-4140-aacb-2c98a9abdf6: uni/tn-common/out-l3out-2"
```

Step 12 Run the validation tool in check mode.

This step does the work to perform the migration, but does not commit the changes to the database. This enables you to determine whether or not the migration is likely to succeed before actually doing it.

Example:

```
# gbp-validate --norepair --config-file /etc/neutron/neutron.conf --config-file
/etc/neutron/plugins/ml2/ml2_conf.ini
```

The tool prints a message when it exits, indicating whether the migration would have succeeded. If the migration would have succeeded, Validation result: failed, repairable is displayed.

If something other than this message appears, then the migration will likely fail. Revert the changes from steps 6, 8, 9 and 11, and do not proceed with the migration procedure. Failures require Cisco support to analyze the root cause and help with changes needed to enable a successful future migration.

Step 13 Run the validation tool in repair mode.

This step creates the new configuration in the AIM database. It also sets up ports to require an update when the Neutron server is restarted.

Example:

```
# # gbp-validate --repair --config-file /etc/neutron/neutron.conf --config-file
/etc/neutron/plugins/ml2/ml2_conf.ini
```

The tool prints a message when it exits, indicating the result of the migration. If successful, it prints Validation result: repaired. If the migration is not successful, the logs stored in /var/log/neutron/gbp-validate.log should provide an indication of what failed.

Step 14 On all the nodes, shut down the OpFlex agent.

This step disconnects traffic from OpenStack workloads from the Cisco ACI control plane, but should not impact existing workloads.

- Stop the OpFlex agent:

Example:

```
# systemctl stop opflex-agent
```

Step 15 On one compute node at a time, execute the following substeps in sequence.

- Restart the OpFlex agent and wait about 60 seconds:

Example:

```
# systemctl start opflex-agent
```

The 60 seconds wait period is a general estimate. This number varies with the installation size.

- Restart the Neutron OpFlex agent:

Example:

```
# systemctl start neutron-opflex-agent
```

This step moves all the Neutron ports and their floating IPs on this host to the new VMM.

Step 16 Start AIM and `neutron-server`.

- Start AIM:

Example:

```
# systemctl start aim-event-service-rpc
# systemctl start aim-event-service-polling
# systemctl start aim-aid
```

- In step 11b, the support_tool `toggle-config` command sets the `agent_down_time` parameter to a large value so that port-binding would succeed during the upgrade to the Cisco ACI OpenStack Unified Plug-in. Set this parameter to the value used before running this command, in order to use the intended `agent_down_time` value. (The default is 60 seconds.)
- Start the Neutron server:

Example:

```
# systemctl start neutron-server
```

Step 17 Delete the subnets of the old bridge domain used for L3out policies.

Deleting the subnets ensures that there is no ambiguity as to which bridge domain and external endpoint group should be used for the network addressing and the associated policy. If external connectivity is required during this transition, you can add host routes (that is, /32 addressing) to the bridge domain for those hosts that need the connectivity.

Step 18 Associate any new VMM domains with the AEP.

Step 19 Verify the following:

- Connectivity: L2 and L3 connectivity for a sample VM on each compute node.
- Security groups after migration are now enforced by OVS, not IP tables.
- DHCP and metadata usage when creating a new VM.
- Floating IP and SNAT connectivity, where applicable.
- Static paths configured in step 1c.
- Usage of pre-existing contracts with new endpoint groups (EPGs), if applicable.
- LBaaS usage after migration, if applicable.

If this step fails, revert changes created by the migration procedure by performing a rollback.

Step 20 Remove association of the older VMM domains with the AEP.

Step 21 Clean up the remaining resources on Cisco APIC.

- a) Delete the VMM with the older system ID.
 - b) Delete old Cisco APIC configuration or association saved in step 11a.
-