Tenants and Tenant Policies

Tenants Overview

A tenant is a logical container for application policies that enable an administrator to exercise domain-based access control. A tenant represents a unit of isolation from a policy perspective, but it does not represent a private network. Tenants can represent a customer in a service provider setting, an organization or domain in an enterprise setting, or just a convenient grouping of policies.

Three default tenants are pre-configured for you:

- `common`—A special tenant with the purpose of providing "common" services to other tenants in ACI fabrics. Global reuse is a core principle in the common tenant. Some examples of common services include shared L3Outs, DNS, DHCP, Active Directory, and shared private networks or bridge domains.
- `dcnm-default-tn`—A special tenant with the purpose of providing configuration for Cisco NDFC fabrics.
- `infra`—The Infrastructure tenant that is used for all internal fabric communications, such as tunnels and policy deployment. This includes switch to switch and switch to APIC communications. The `infra` tenant does not get exposed to the user space (tenants) and it has its own private network space and bridge domains. Fabric discovery, image management, and DHCP for fabric functions are all handled within this tenant.

  When using Nexus Dashboard Orchestrator to manage Cisco NDFC fabrics, you will always use the default `dcnm-default-tn` tenant.

Note Nexus Dashboard Orchestrator cannot manage the APIC's `mgmt` tenant, so importing the tenant from APIC or creating a new tenant called `mgmt` in NDO it is not allowed.

To manage tenants, you must have either `Power User` or `Site and Tenant Manager` read-write role.

Tenant Policies Templates

Release 4.0(1) adds Tenant Policies templates, which allow you to configure the following tenant-wide policies:

- Route Policies for Multicast
- Route Map Policies for Route Control
- Custom QoS Policies
- DHCP Relay Policies
- DHCP Option Policies
- IGMP Interface Policies
- IGMP Snooping Policies
- MLD Snooping Policies

For additional information, see Creating Tenant Policy Templates, on page 4.

Creating New Tenants

This section describes how to add a new tenant using the Nexus Dashboard Orchestrator GUI. If you want to import one or more existing tenants from your fabrics, follow the steps described in Importing Existing Tenants, on page 3 instead.

Before you begin

You must have a user with either `Power User` or `Site Manager` read-write role to create and manage tenants.

Procedure

Step 1 Log in to your Nexus Dashboard and open the Nexus Dashboard Orchestrator service.

Step 2 Create a new tenant.

a) From the left navigation pane, choose **Application Management** > **Tenants**.

b) In the top right of the main pane, click **Add Tenant**.

The **Add Tenant** screen opens.

Step 3 Provide tenant details.

a) Provide the **Display Name** and optional **Description**.

The tenant's **Display Name** is used throughout the Orchestrator's GUI whenever the tenant is shown. However, due to object naming requirements on the APIC, any invalid characters are removed and the resulting **Internal Name** is used when pushing the tenant to sites. The **Internal Name** that will be used when creating the tenant is displayed below the **Display Name** textbox.

Note You can change the **Display Name** of the tenant at any time, but the **Internal Name** cannot be changed after the tenant is created.

b) In the **Associated Sites** section, check all the sites you want to associate with this tenant.

Only the selected sites will be available for any templates using this tenant.

c) (Optional) For each selected site, click the **Edit** button next to its name and choose one or more security domains.

A restricted security domain allows a fabric administrator to prevent a group of users, such as Tenant A, from viewing or modifying any objects created by a group of users in a different security domain, such as Tenant B, when users in both groups have the same assigned privileges. For example, a tenant administrator in Tenant A's restricted security domain will not be able to see policies, profiles, or users configured in Tenant B's security domain. Unless Tenant B's security domain is also restricted, Tenant B will be able to see policies, profiles, or users configured in Tenant A. Note that a user will always have read-only visibility to system-created configurations for which the user has proper privileges. A user in a restricted security domain can be given a broad level of privileges within that domain without the concern that the user could inadvertently affect another tenant's physical environment.

Security domains are created using the APIC GUI and can be assigned to various APIC policies and user accounts to control their access. For more information, see the *Cisco APIC Basic Configuration Guide*.

d) In the **Associated Users** section, select the Nexus Dashboard Orchestrator users that are allowed to access the tenant.

Only the selected users will be able to use this tenant when creating templates.

Step 4 Click **Save** to finish adding the tenant.

Importing Existing Tenants

This section describes how to import one or more existing tenants. If you want to create a new tenant using Nexus Dashboard Orchestrator, follow the steps described in Creating New Tenants, on page 2 instead.

Before you begin

You must have a user with either `Power User` or `Site Manager` read-write role to create and manage tenants.

Procedure

Step 1 Log in to your Nexus Dashboard and open the Nexus Dashboard Orchestrator service.

Step 2 In the left navigation menu, click **Sites**.

Step 3 Locate the site from which you want to import the tenants, click its **Actions** (**...**) menu, and choose **Import Tenants**.

You can import tenants from one site at a time.

Step 4 In the **Import Tenants** dialog, select one or more tenants to import and click **Ok**.

The selected tenants will be imported into the Nexus Dashboard Orchestrator and show in the **Application Management** > **Tenants** page.

Step 5 Repeat these steps to import tenants from any other sites.

Creating Tenant Policy Templates

This section describes how to create one or more tenant policy templates. Tenant policy templates allow you to create and configure the following policies:

- Route Map Policies for Multicast
- Route Map Policies for Route Control
- Custom QoS Policies
- DHCP Relay Policies
- DHCP Option Policies
- IGMP Interface Policies
- MLD Snooping Policies

Procedure

Step 1 Log in to your Nexus Dashboard and open the Nexus Dashboard Orchestrator service.

Step 2 Create a new Tenant Policy template.

a) From the left navigation pane, choose **Application Management** > **Tenant Policies**.
b) On the **Tenant Policy Templates** page, click **Add Tenant Policy Template**.
c) In the **Tenant Policies** page's right properties sidebar, provide the **Name** for the template.
d) From the **Select a Tenant** dropdown, choose the tenant with which you want to associate this template.

All the policies you create in this template as described in the following steps will be associated with the selected tenant and deployed to it when you push the template to a specific site.

By default, the new template is empty, so you need to add one or more tenant policies as described in the following steps. Note that you don't have to create every policy available in the template – you can define one or more policies of each type to deploy along with this template. If you don't want to create a specific policy, simply skip the step that describes it.

Step 3 Assign the template to one or more sites.

The process for assigning Tenant Policy templates to sites is identical to how you assign application templates to sites.

a) In the **Template Properties** view, click **Actions** and choose **Sites Association**.

The **Associate Sites to <template-name>** window opens.

b) In the **Associate Sites** window, check the checkbox next to the sites where you want to deploy the template.

Note that only the on-premises ACI sites support tenant policy templates and will be available for assignment.

c) Click **Ok** to save.

Step 4 Create a Route Map Policy for Multicast.

This policy is part of the overarching Layer 3 Multicast use case. You can use the information in this section as a reference, but we recommend following the full set of steps described in the Layer 3 Multicast chapter of the *Features and Use Cases* section of this document.

a) From the **+Create Object** dropdown, select **Route Map Policy for Multicast**.
b) In the right properties sidebar, provide the **Name** for the policy.
c) (Optional) Click **Add Description** and provide a description for the policy.
d) Click **+Add Route Map for Multicast Entries** and provide the route map information.

For each route map, you need to create one or more route map entries. Each entry is a rule that defines an action based on one or more matching criteria based on the following information:

- **Order** – Order is used to determine the order in which the rules are evaluated.
- **Group IP**, **Src IP**, and **RP IP** – You can use the same multicast route map policy UI for two different use cases—to configure a set of filters for multicast traffic or to restrict a rendezvous point configuration to a specific set of multicast groups. Depending on which use case you're configuring, you only need to fill some of the fields in this screen:
  - For multicast filtering, you can use the **Source IP** and the **Group IP** fields to define the filter. You must provide at least one of these fields, but can choose to include both. If one of the fields is left blank, it will match all values.

    The Group IP range must be between `224.0.0.0` and `239.255.255.255` with a netmask between `/4` and `/32`. You must provide the subnet mask.

    The **RP IP** (Rendezvous Point IP) is not used for multicast filtering route maps, so leave this field blank.
  - For Rendezvous Point configuration, you can use the **Group IP** field to define the multicast groups for the RP.

    The Group IP range must be between `224.0.0.0` and `239.255.255.255` with a netmask between `/4` and `/32`. You must provide the subnet mask.

    For Rendezvous Point configuration, the **RP IP** is configured as part of the RP configuration. If a route-map is used for group filtering it is not necessary to configure an RP IP address in the route-map. In this case, leave the **RP IP** and **Source IP** fields empty.
- **Action** – Action defines the action to perform, either `Permit` or `Deny` the traffic, if a match is found.

e) Click the checkmark icon to save the entry.
f) Repeat the previous substeps to create any additional route map entries for the same policy.
g) Click **Save** to save the policy and return to the template page.
h) Repeat this step to create any additional Route Map for Multicast policies.

Step 5 Create a Route Map Policy for Route Control.

This policy is part of the overarching SR-MPLS use case. You can use the information in this section as a reference, but we recommend following the full set of steps described in the Multi-Site and SR-MPLS L3Out Handoff chapter of the *Features and Use Cases* section of this document.

a) From the **+Create Object** dropdown, select **Route Map Policy for Route Control**.
b) In the right properties sidebar, provide the **Name** for the policy.
c) (Optional) Click **Add Description** and provide a description for the policy.
d) Click **+Add Entry** and provide the route map information.

For each route map, you need to create one or more context entries. Each entry is a rule that defines an action based on one or more matching criteria based on the following information:

- **Context Order** – Context order is used to determine the order in which contexts are evaluated. The value must be in the `0-9` range.
- **Context Action** – Context action defines the action to perform (`permit` or `deny`) if a match is found.

Once the context order and action are defined, choose how you want to match the context:

- Click **+Add Attribute** to specify the action that will be taken should the context match.

  You can choose one of the following actions:

  - `Set Community`
  - `Set Route Tag`
  - `Set Dampening`
  - `Set Weight`
  - `Set Next Hop`
  - `Set Preference`
  - `Set Metric`
  - `Set Metric Type`
  - `Set AS Path`
  - `Set Additional Community`

  After you have configured the attribute, click **Save**.

- If you want to match an action based on an IP address or prefix, click **Add IP Address**.

  In the **Prefix** field, provide the IP address prefix. Both IPv4 and IPv6 prefixes are supported, for example `2003:1:1a5:1a5::/64` or `205.205.0.0/16`.

  If you want to aggregate IPs in a specific range, check the **Aggregate** checkbox and provide the range. For example, you can specify `0.0.0.0/0` prefix to match any IP or you can specify `10.0.0.0/8` prefix to match any `10.x.x.x` addresses.

- If you want to match an action based on community lists, click **Add Community**.

  In the **Community** field, provide the community string. For example, `regular:as2-nn2:200:300`.

  Then choose the **Scope**: `Transitive` means the community will be propagated across eBGP peering (across autonomous systems) while `Non-Transitive` means the community will not be propagated.

e) Repeat the previous substeps to create any additional route map entries for the same policy.
f) Click **Save** to save the policy and return to the template page.
g) Repeat this step to create any additional Route Map for Route Control policies.

Step 6 Create a Custom QoS Policy.

You can create a custom QoS policy in Cisco APIC to classify ingressing traffic based on its DSCP or CoS values and associate it to a QoS priority level (QoS user class) to properly handle it inside the ACI fabric. Classification is supported

only if the DSCP values are present in the IP header and/or the CoS values are present in the Ethernet header of ingressing traffic. Additionally, the custom QoS policy can be used to modify the DSCP and/or CoS values in the header of ingressing traffic

As an example, custom QoS policies allow you to classify traffic coming into the ACI fabric traffic from devices that mark the traffic based only on the CoS value, such as Layer-2 packets which do not have an IP header.

For detailed information about QoS functionality in ACI fabrics, see *Cisco APIC and QoS*.

a) From the **+Create Object** dropdown, select **Custom QoS Policy**.
b) In the right properties sidebar, provide the **Name** for the policy.
c) (Optional) Click **Add Description** and provide a description for the policy.
d) Click **+Add DSCP Mappings** and provide the required information.

   The DSCP mapping configuration allows you to associate ingressing traffic, whose DSCP value is within the range specified in the mapping, to the specified QoS priority level (class). It also allows you to set the DSCP and/or CoS values of the ingressing traffic, so that those values can be retained when the traffic egresses the fabric.

   Note Retaining the target CoS value for egress traffic requires the configuration of the "Preserve CoS" policy, which is part of the NDO Fabric policies.

   If the "DSCP Target" and/or "Target CoS" values are set as part of both the DSCP Mapping and CoS Mapping, the values specified in the DSCP Mapping have precedence.

   For each mapping, you can specify the following fields:

   - **DSCP From** – The start of the DSCP range.
   - **DSCP To** – The end of the DSCP range.
   - **DSCP Target** – The DSCP value to set on ingressing traffic that will be retained for egressing traffic.
   - **Target CoS** – The CoS value to set on ingressing traffic that will be retained for egressing traffic when "Preserve CoS" is enabled.
   - **Priority** – The QoS priority class to which the traffic will be assigned.

   After you provide the mappings, click the checkmark icon to save. Then you can click **+Add DSCP Mappings** to provide additional mappings within the same policy.

e) Click **Add** to save the policy and return to the template page.
f) Click **+Add CoS Mappings** and provide the required information.

   The DSCP mapping configuration allows you to associate ingressing traffic (whose DSCP value is within the range specified in the mapping) to the specified QoS priority level (class). It also allows you to set the DSCP and/or CoS values of the ingressing traffic, so that those values can be retained when the traffic egresses the fabric.

   Note Retaining the target CoS value for egress traffic requires the configuration of the "Preserve CoS" policy in the NDO Fabric policies.

   In addition, if the "DSCP Target" and/or "Target CoS" values are set as part of both the DSCP Mapping and CoS Mapping, the values specified in the DSCP Mapping have precedence.

   For each mapping, you can specify the following fields:

   - **Dot1P From** – The start of the CoS range.
   - **Dot1P To** – The end of the CoS range.

- **DSCP Target** – The DSCP value to set on ingressing traffic that will be retained for egressing traffic.
- **Target CoS** – The CoS value to set on ingressing traffic that will be retained for egressing traffic when "Preserve CoS" is enabled.
- **Priority** – The QoS priority class to which the traffic will be assigned.

After you provide the mappings, click the checkmark icon to save. Then you can click **+Add Cos Mappings** to provide additional mappings within the same policy.

g) Click **Add** to save the policy and return to the template page.
h) Repeat this step to create any additional Route Map for Route Control policies.

Step 7 Create a DHCP Relay Policy.

This policy is part of the overarching DHCP Relay use case. You can use the information in this section as a reference, but we recommend following the full set of steps described in the DHCP Relay chapter of the *Features and Use Cases* section of this document.

a) From the **+Create Object** dropdown, select **DHCP Relay Policy**.
b) In the right properties sidebar, provide the **Name** for the policy.
c) (Optional) Click **Add Description** and provide a description for the policy.
d) Click **Add Provider** to configure the DHCP server to which you want ot relay the DHCP requests originated by the endpoints.
e) Select the provider type.

When adding a relay policy, you can choose one of the following two types:

- `Application EPG`—specifies the application EPG that includes the DHCP server to which you want to relay the DHCP requests.
- `L3 External Network`—specifies the External EPG associated to the L3Out that is used to access the network external to the fabric where the DHCP server is connected.

Note You can select any EPG or external EPG that has been created in the Orchestrator and assigned to the tenant you specified, even if you have not yet deployed it to sites. If you select an EPG that hasn't been deployed, you can still complete the DHCP relay configuration, but you will need to deploy the EPG before the relay is available for use.

f) Click **Select an Application EPG** or **Select an External EPG** (based on the provider type you selected) and choose the provider EPG.
g) In the **DHCP Server Address** field, provide the IP address of the DHCP server.
h) Enable the **DHCP Server VRF Preference** option if required.

This feature was introduced in Cisco APIC release 5.2(4). For more information on the use cases where it is required refer to the Cisco APIC Basic Configuration Guide.

i) Click **OK** to save the provider information.
j) Repeat the previous substeps for any additional providers in the same DHCP Relay policy.
k) Repeat this step to create any additional DHCP Relay policies.

Step 8 Create a DHCP Option Policy.

This policy is part of the overarching DHCP Relay use case. You can use the information in this section as a reference, but we recommend following the full set of steps described in the DHCP Relay chapter of the *Features and Use Cases* section of this document.

a) From the **+Create Object** dropdown, select **DHCP Option Policy**.
b) In the right properties sidebar, provide the **Name** for the policy.
c) (Optional) Click **Add Description** and provide a description for the policy.
d) Click **Add Option**.
e) Provide option details.

For each DHCP option, provide the following:

- **Name** – While not technically required, we recommend using the same name for the option as listed in RFC 2132.

  For example, `Name Server`.

- **Id** – Provide the value if the option requires one.

  For example, a list of name servers available to the client for the Name Server option.

- **Data** – Provide the value if the option requires one.

  For example, a list of name servers available to the client for the Name Server option.

f) Click **OK** to save.
g) Repeat the previous substeps for any additional options in the same DHCP Option policy.
h) Repeat this step to create any additional DHCP Option policies.

Step 9 Create a IGMP Interface Policy.

IGMP snooping examines IP multicast traffic within a bridge domain to discover the ports where interested receivers reside. Using the port information, IGMP snooping can reduce bandwidth consumption in a multi-access bridge domain environment to avoid flooding the entire bridge domain.

For detailed information on IGMP snooping in ACI fabrics, see the "IGMP Snooping" chapter of the *Cisco APIC Layer 3 Networking Configuration Guide* for your release.

a) From the **+Create Object** dropdown, select **IGMP Interface Policy**.
b) In the right properties sidebar, provide the **Name** for the policy.
c) (Optional) Click **Add Description** and provide a description for the policy.
d) Provide policy details.

- **Allow Version 3 ASM** – Allow accepting IGMP version 3 source-specific reports for multicast groups outside of the SSM range. When this feature is enabled, the switch will create an (S,G) mroute entry if it receives an IGMP version 3 report that includes both the group and source even if the group is outside of the configured SSM range. This feature is not required if hosts send (*,G) reports outside of the SSM range, or send (S,G) reports for the SSM range

- **Fast Leave** – Option that minimizes the leave latency of IGMPv2 group memberships on a given IGMP interface because the device does not send group-specific queries. When Fast Leave is enabled, the device removes the group entry from the multicast routing table immediately upon receiving a leave message for the group. The default is disabled.

  Use this only when there is only one receiver behind the BD/interface for a given group.

- **Report Link Local Groups** – Enables sending reports for groups in 224.0.0.0/24. Reports are always sent for nonlink local groups. By default, reports are not sent for link local groups.

- **IGMP Version** – IGMP version that is enabled on the bridge domain or interface. The IGMP version can be 2 or 3. The default is 2.

- **Advanced Settings** – Click the arrow next to this section to expand.
  - **Group Timeout** – Group membership interval that must pass before the router decides that no members of a group or source exist on the network. Values range from 3 to 65,535 seconds. The default is 260 seconds.
  - **Query Interval** – Sets the frequency at which the software sends IGMP host query messages. Values can range from 1 to 18,000 seconds. The default is 125 seconds.
  - **Query Response Interval** – Sets the response time advertised in IGMP queries. Values can range from 1 to 25 seconds. The default is 10 seconds.
  - **Last Member Count** – Sets the number of times that the software sends an IGMP query in response to a host leave message. Values can range from 1 to 5. The default is 2.
  - **Last Member Response Time** – Sets the query interval waited after sending membership reports before the software deletes the group state. Values can range from 1 to 25 seconds. The default is 1 second.
  - **Startup Query Count** – Configures snooping for a number of queries sent at startup when you do not enable PIM because multicast traffic does not need to be routed. Values can range from 1 to 10. The default is 2 messages.
  - **Startup Query Interval** – This configures the IGMP snooping query interval at startup. The range is from 1 second to 18000 seconds. The default is 125 seconds..
  - **Querier Timeout** – Sets the query timeout that the software uses when deciding to take over as the querier. Values can range from 1 to 65,535 seconds. The default is 255 seconds
  - **Robustness Variable** – Sets the robustness variable. You can use a larger value for a lossy network. Values can range from 1 to 7. The default is 2.
  - **State Limit Route Map** – Used with Reserved Multicast Entries feature.

    The route map policy must be already created as described in Step 2.
  - **Report Policy Route Map** – Access policy for IGMP reports that is based on a route-map policy. IGMP group reports will only be selected for groups allowed by the route-map.

    The route map policy must be already created as described in Step 2.
  - **Static Report Route Map** – Statically binds a multicast group to the outgoing interface, which is handled by the switch hardware. If you specify only the group address, the (*, G) state is created. If you specify the source address, the (S, G) state is created. You can specify a route-map policy name that lists the group prefixes, group ranges, and source prefixes. A source tree is built for the (S, G) state only if you enable IGMPv3.

    The route map policy must be already created as described in Step 2.
  - **Maximum Multicast Entries** – Limit the mroute states for the BD or interface that are created by IGMP reports. Default is disabled and no limit is enforced. Valid range is 1-4294967295.

e) Repeat this step to create any additional IGMP Interface policies.

Step 10 Create a MLD Snooping Policy.

Multicast Listener Discovery (MLD) snooping enables the efficient distribution of IPv6 multicast traffic between hosts and routers. It is a Layer 2 feature that restricts IPv6 multicast traffic within a bridge domain to a subset of ports that have transmitted or received MLD queries or reports. In this way, MLD snooping provides the benefit of conserving

the bandwidth on those segments of the network where no node has expressed interest in receiving the multicast traffic. This reduces the bandwidth usage instead of flooding the bridge domain, and also helps hosts and routers save unwanted packet processing.

For detailed information on MLD snooping in ACI fabrics, see the "MLD Snooping" chapter of the *Cisco APIC Layer 3 Networking Configuration Guide*for your release.

a) From the **+Create Object** dropdown, select **MLD Snooping Policy**.
b) In the right properties sidebar, provide the **Name** for the policy.
c) (Optional) Click **Add Description** and provide a description for the policy.
d) Provide policy details.

- **Admin State** – Enables or disables the MLD snooping feature.
- **Fast Leave Control** – Allows you to turn on or off the fast-leave feature on a per bridge domain basis. This applies to MLDv2 hosts and is used on ports that are known to have only one host doing MLD behind that port.

  Default is `disabled`.
- **Querier Control** – Enables or disables MLD snooping querier processing. MLD snooping querier supports the MLD snooping in a bridge domain where PIM and MLD are not configured because the multicast traffic does not need to be routed

  Default is `disabled`.
- **Querier Version** – Allows you to choose the querier version.

  Default is `Version2`.
- **Advanced Settings** – Click the arrow next to this section to expand.
  - **Query Interval** – Sets the frequency at which the software sends MLD host query messages. Values can range from 1 to 18,000 seconds.

    The default is 125 seconds.
  - **Query Response Interval** – Sets the response time advertised in MLD queries. Values can range from 1 to 25 seconds.

    The default is 10 seconds.
  - **Last Member Query Interval** – Sets the query response time after sending membership reports before the software deletes the group state. Values can range from 1 to 25 seconds.

    The default is 1 second.
  - **Start Query Count** – Configures snooping for a number of queries sent at startup when you do not enable PIM because multicast traffic does not need to be routed. Values can range from 1 to 10.

    The default is 2.
  - **Start Query Interval** – Configures a snooping query interval at startup when you do not enable PIM because multicast traffic does not need to be routed. Values can range from 1 to 18,000 seconds.

    The default is 31 seconds.

e) Repeat this step to create any additional MLD Snooping policies.

Step 11 Click **Save** to save the changes you've made to the template.

Note When you save (or deploy) the template to one or more sites, the Orchestrator will verify that the specified nodes and/or interfaces are valid for the site(s) and will return an error.

Step 12 Click **Deploy** to deploy the template to the associated sites.

The process for deploying tenant policy templates is identical to how you deploy application templates.

If you have previously deployed this template but made no changes to it since, the **Deploy** summary will indicate that there are no changes and you can choose to re-deploy the entire template. In this case, you can skip this step.

Otherwise, the **Deploy to sites** window will show you a summary of the configuration differences that will be deployed to sites. Note that in this case only the difference in configuration is deployed to the sites. If you want to re-deploy the entire template, you must deploy once to sync the differences and then redeploy again to push the entire configuration as described in the previous paragraph.