



## Configuring External Connectivity (L3Out)

---

- [L3Out Template Overview, on page 1](#)
- [Guidelines and Limitations, on page 5](#)
- [Greenfield Deployment, on page 6](#)
- [Importing Existing L3Out Configuration, on page 18](#)
- [Viewing L3Out Neighbors, on page 29](#)

### L3Out Template Overview

Beginning with release 4.1(1), Nexus Dashboard Orchestrator (NDO) introduced a number of new policies for creating and configuring L3Out for Cisco ACI fabrics, as well as a new template type specifically for IP-based L3Out and SR-MPLS VRF L3Out configurations.

As you may already know, prior releases of NDO provided the ability to create an L3Out object in Application templates that allowed you to create an L3Out and deploy it to your site. However, the actual L3Out configurations had to be done manually by logging in to the sites' controllers (Cisco APIC) and providing the details for each L3Out individually.

With release 4.1(1), the entire configuration of L3Outs and SR-MPLS L3Outs (including nodes, interfaces, and other settings) can be done directly in NDO and deployed to all fabrics in your Multi-Site domain. To achieve this, a new L3Out-specific template type has been added to contain the L3Out and SR-MPLS VRF L3Out configurations. Similar to Application templates, L3Out templates have a one-to-one association with tenants but unlike Application templates, an L3Out template must be associated to a single site only.



#### Note

The legacy L3Out objects in the Application templates remain functional for backward compatibility. However, if you want to define specific L3Out and SR-MPLS L3Out setting from NDO, you must use the new L3Out template type.

The legacy SR-MPLS VRF L3Out object has been removed from the Application template and all SR-MPLS VRF L3Out configurations must be now done using the L3Out-specific template. The SR-MPLS Infra L3Out configuration is still performed as part of the site connectivity provisioning workflow.

#### Templates and Policy Objects Dependencies

The following diagram illustrates the template and policy hierarchy across multiple templates that's required for defining a complete L3Out configuration:

- The VRF used by the L3Out and the External EPGs that are associated to the L3Out continue to be defined in the Application templates.
- Node or interface routing policies, BGP peer prefix, and IP SLA policies are now defined in the Tenant Policy template.

These policies are used by the L3Out-specific template and the policies defined in that template as described in the following bullet point.

- For IP-based L3Outs, the template includes the following:
  - Routing Protocol (BGP/OSPF), VRF, L3 Domain and Route Maps for route control.
  - Border leaf switches (nodes) where to deploy the L3Out routing protocol and node-level protocol configurations.
  - Border leaf switch interfaces where to deploy the L3Out routing protocol and interface-level protocol configurations.
  - Node- and interface-level common configuration using Node/Interface Group policies.

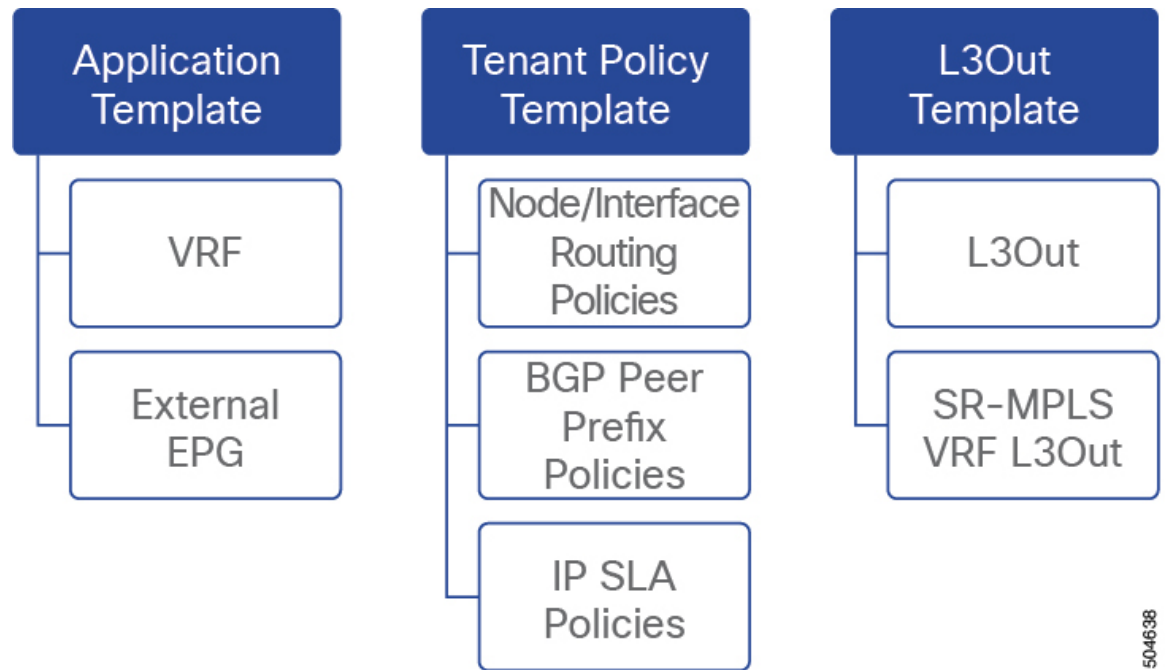
Node Group configuration includes BGP peers for loopback interfaces, BFD multi-hop settings, and association with Node Routing Group Policy described below.

Interface Group configuration includes OSPF and BFD protocol settings and association with Interface Routing Group Policy described below.

These policies consume policies defined in Tenant Policy templates mentioned in the previous bullet point. For example, the node and interface group policies require the node and interface routing policies defined in the Tenant Policy templates

- For SR-MPLS VRF L3Outs, the template allows you to define labels and import/export route maps for route control.

Figure 1: L3Out Templates and Policy Objects



504638

#### Tenant Policy Template: Node Routing Group Policy

The Node Routing Policy in Tenant Policy template is a set of protocol policies that can be applied at a node or border leaf level and can be used by node group policies in L3Out template. It includes the following 3 settings:

- **BFD MultiHop Settings** – provides fast forwarding path failure detection for destinations with more than one hop.  
In this case, a multi-hop session is created between the source and destination instead of the interface like in single-hop scenarios.
- **BGP Node Settings** – allows you to configure BGP protocol timer and sessions settings for traffic between BGP peers.
- **BGP Best Path Control** – allows you to enable `as-path multipath-relax`, which allows load-balancing between multiple paths received from different BGP ASN.

This policy is configured and deployed using Tenant Policy templates and is used by the L3Outs configured in L3Out templates.

#### Tenant Policy Template: Interface Routing Group Policy

The Interface Routing Policy in Tenant Policy template is a set of policies that can be applied at an interface level and can be used by interface group policies in L3Out template. It includes the following 3 settings:

- **BFD settings** – provides fast forwarding path failure detection between ACI fabric border leaf switches that are configured to support peering router connections.

When multiple protocols are enabled between a pair of routers, each protocol has its own link failure detection mechanism, which may have different timeouts. BFD provides a consistent timeout for all protocols to allow consistent and predictable convergence times.

- **BFD MultiHop Settings** – provides fast forwarding path failure detection for destinations with more than one hop.

You can configure these settings at the node level as mentioned in the "Tenant Policy Template: Node Routing Group Policy" section above, in which case the interfaces inherit those settings, or you can overwrite the node-level settings for individual interfaces in the Interface Routing group policy.




---

**Note** BFD multi-hop configuration requires Cisco APIC release 5.0(1) or later.

---

- **OSPF Interface Settings** – allows you to configure interface-level settings such as OSPF network type, priority, cost, intervals and controls.




---

**Note** This policy must be created when deploying an L3Out with OSPF.

---

This policy is configured and deployed using Tenant Policy templates and is used by the L3Outs configured in L3Out templates.

### Tenant Policy Template: Individual Policies

In addition to the group policies described above, the Tenant Policy templates also contain the following individual policies related to L3Out configuration:

- **BGP peer prefix policy** – defines how many prefixes can be received from a neighbor and what action to take when the number of the allowed prefixes is exceeded.

This policy is configured and deployed using Tenant Policy templates and is used by the L3Outs configured in L3Out templates.

- **IP SLA monitoring policy** – defines the type of probe (ICMP/TCP/HTTP) and respective settings to use for monitoring endpoints. This policy is associated with monitoring probe profiles known as "track members", which represent a network segment to be monitored. You can associate an IP SLA monitoring policy to a track list (which includes multiple track members) and associate this track list to a Static Route for monitoring availability of track list members over the route. In addition, you can associate IP SLA monitoring policy directly to next-hop address of a Static Route for monitoring its availability over the route.




---

**Note** IP SLA monitoring policy of HTTP type requires Cisco APIC release 5.1(3) or later.

---

- **IP SLA track list** - defines the IP addresses to be tracked, IP SLA Monitoring policy (probe frequency and type), and scope (bridge domain or L3Out). IP SLA track list aggregates one or more track members, defines what percentage or weight of track members must be up/down for the route to be considered

available or unavailable. Based on the track list, the available routes remain in the routing table and the unavailable routes are removed until the track list recovers.

This policy is configured and deployed using Tenant Policy templates and is used by the L3Outs configured in L3Out templates. In addition, an IPSLA track list can be configured in the same Tenant Policy template as the monitoring policy and consumed by it.

### L3Out Template

The L3Outs defined in L3Out templates allow you to define all the required configurations to enable connectivity from the endpoints inside your ACI fabrics to outside network domains through routing protocols or static routes. The L3Out object in NDO contains settings necessary for the following:

- Learning external routes via routing protocols or static routes.
- Distribution of the learned external routes to other leaf switches.
- Advertisement of ACI internal routes (BD subnets) to the outside networks.
- Advertisement of learned external routes to other L3Outs (transit routing).

When you create an L3Out template and configure L3Out-specific objects and properties as described later in [Creating L3Out Template, on page 12](#), you will:

1. Define a number of common properties, such as the VRF, L3 Domain, and routing protocol (BGP and/or OSPF), for the L3Out.
2. Specify one or more border leaf switches (nodes) and optionally associate each node with a Node Group policy .
3. Specify one or more interfaces on those border leaf switches and optionally associate each interface with an Interface Group policy described above.
4. After you have created an L3Out template and deployed one or more L3Outs, you can control traffic between the ACI EPGs and external networks using contracts in Application templates as you typically would.

## Guidelines and Limitations

The following guidelines apply when using an L3Out template to configure IP-based L3Outs and SR-MPLS VRF L3Outs:

- Similar to Application templates, L3Out templates have a one-to-one association with tenants but unlike Application templates, an L3Out template must be associated to a single site only.
- The legacy L3Out container objects in the Application templates remain functional for backward compatibility.

Note however, if you want to define specific L3Out and SR-MPLS VRF L3Out settings, you must use the L3Out-specific template type. As such, we recommend using the L3Out-specific templates for all new L3Out and SR-MPLS VRF L3Out configurations.

- The legacy SR-MPLS VRF L3Out contain object has been removed from the Application template.

All SR-MPLS VRF L3Out configurations must be done using the L3Out-specific template.

- If you want to configure BFD multi-hop settings, your fabric must be running Cisco APIC release 5.0(1) or later.
- If you want to configure an IP SLA monitoring policy of HTTP type, your fabric must be running Cisco APIC release 5.1(3) or later.

# Greenfield Deployment

## Creating Tenant Policy Template

This section describes how to create a Tenant Policy template and define the L3Out-specific policies, which you will then consume in an L3Out template as described later in this document. For more information about each policy and how it relates to policies and settings in other templates, see [L3Out Template Overview, on page 1](#).



---

**Note** If you want to import existing L3Out configurations from a site's APIC, follow the "Importing Existing L3Out Configuration" steps in the following sections of this chapter instead.

---

### Before you begin

- You must have the Nexus Dashboard Orchestrator service installed and enabled.
- You must have the fabrics onboarded to your Nexus Dashboard and enabled for management in the Orchestrator service.
- Ensure you have read and understood the Templates and Policy Objects dependencies described in [L3Out Template Overview, on page 1](#).
- Note that the following steps describe how to create multiple policies for an IP-based L3Out, including some mandatory and some optional ones.

You must create the **Outbound Route Map** for the L3Out to function, but other policies described below may be optional depending on your specific use case.

### Procedure

- 
- Step 1** Log in to your Nexus Dashboard and open the Orchestrator service.
- Step 2** In the left navigation pane, choose **Application Management > Tenant Policies**.
- Step 3** In the main pane, click **Add Tenant Policy Template**.
- If you want to update an existing Tenant Policy template instead, simply click on its name. This opens the **Tenant Policies** page.
- Step 4** If you created a brand new template, provide the **Name** for the template and **Select a Tenant** with which you want to associate this template.
- Step 5** Create a Route Map Policy for Route Control.

**Note**

The Route Map Policy for Route Control is mandatory and must be created for all L3Outs. Other policies described in this section are optional and can be defined or skipped depending on your specific use case.

- 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. If the same value is used for multiple contexts, they are evaluated one in the order in which they are defined.

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 associate the action that you just defined with 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 associate the action that you just defined with community lists, click **Add Community**.

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

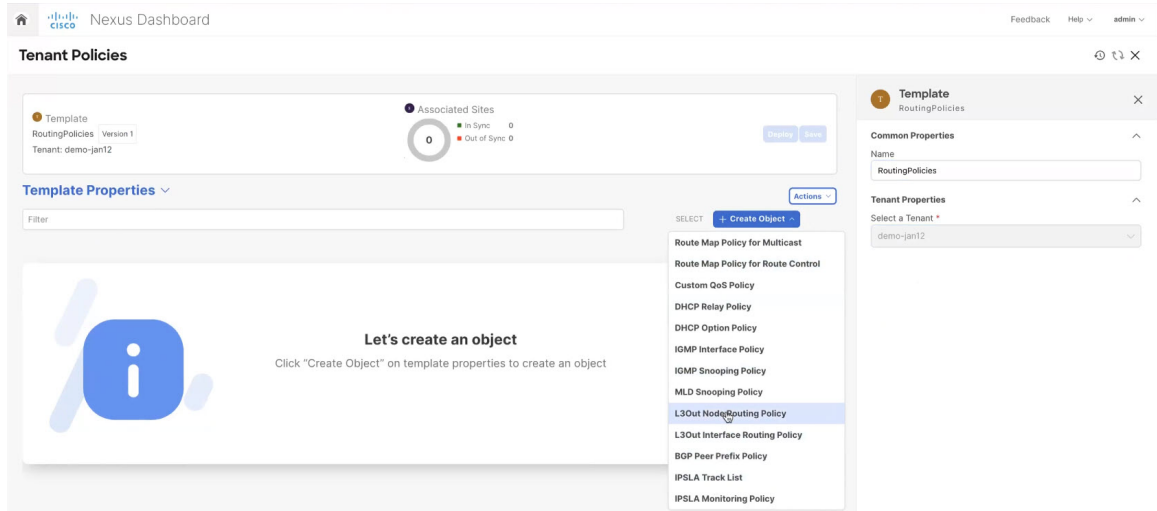
Then choose the **Scope**: `Transitive` means the community will be propagated across eBGP peering (across autonomous systems) while `Non-Transitive` means the community will not be propagated.

- 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.

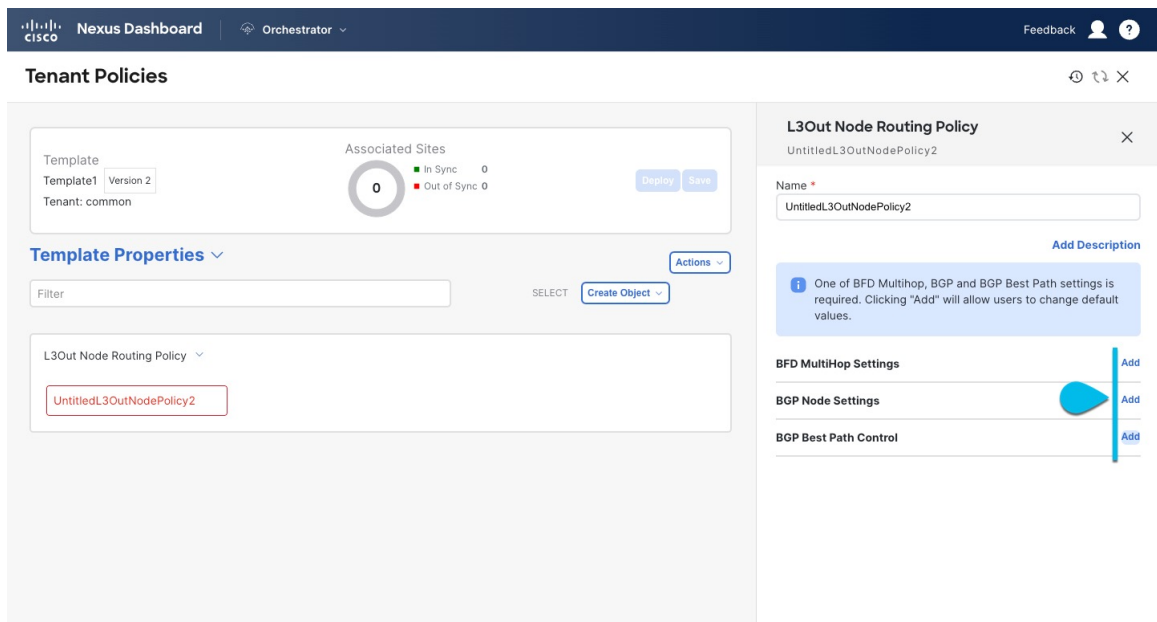
## Step 6

Create an L3Out Node Routing policy.

- a) In the main pane, choose **Create Object > L3Out Node Routing Policy**.



- b) Provide the **Name** for the policy, and **Add** at least one of the **BFD Multi-Hop Settings**, **BGP Node Settings**, or **BGP Best Path Control** options.



- **BFD Multi-Hop Settings** – provides forwarding failure detection for destinations with more than one hop.



In this case, a multi-hop session is created between the source and destination instead of the interface like in single-hop scenarios.

**Note**

BFD multi-hop configuration requires Cisco APIC release 5.0(1) or later.

- **BGP Node Settings** – allows you to configure BGP protocol timer and sessions settings for traffic between BGP peers.
- **BGP Best Path Control** – allows you to enable `as-path multipath-relax`, which allows load-balancing between multiple paths received from different BGP ASN.

**Step 7**

Create an L3Out Interface Routing policy.

- In the main pane, choose **Create Object** > **L3Out Interface Routing Policy**.
- Provide the **Name** for the policy, and define the **BFD Settings**, **BFD Multi-Hop Settings**, and **OSPF Interface Settings**.

The screenshot displays the Cisco Nexus Dashboard Orchestrator interface. The top navigation bar includes the Cisco logo, 'Nexus Dashboard', and 'Orchestrator'. The main content area is titled 'Tenant Policies'. On the left, under 'Template Properties', there is a section for 'Associated Sites' showing a circular progress indicator with '0' and 'In Sync 0'/'Out of Sync 0'. Below this, there are two policy templates listed: 'L3Out Node Routing Policy' (UntitledL3OutNodePolicy2) and 'L3Out Interface Routing Policy' (UntitledL3OutInterfacePoli...). On the right, a modal window titled 'L3Out Interface Routing Policy' is open, showing the 'Name' field set to 'UntitledL3OutInterfacePolicy2'. Below the name field, there is a message: 'One of BFD, BFD Multihop and OSPF settings is required. Clicking "Add" will allow users to change default values.' At the bottom of the modal, there are three sections: 'BFD Settings', 'BFD MultiHop Settings', and 'OSPF Interface Settings', each with an 'Add' button.

- **BFD settings** – provides forwarding failure detection between ACI fabric border leaf switches that are configured to support peering router connections.

When multiple protocols are enabled between a pair of routers, each protocol has its own link failure detection mechanism, which may have different timeouts. BFD provides a consistent timeout for all protocols to allow consistent and predictable convergence times.

- **BFD Multi-Hop Settings** – provides forwarding failure detection for destinations with more than one hop.

In this case, a multi-hop session is created between the source and destination instead of the interface like in single-hop scenarios.

**Note**

BFD multi-hop configuration requires Cisco APIC release 5.0(1) or later.

- **OSPF Interface Settings** – allows you to configure interface-level settings such as priority, cost, intervals and controls.

**Step 8** Associate the template with one or more sites.

- a) In the **Tenant Policies** template view, choose **Actions** > **Sites Association**.




- b) In the **Associate Sites to <template-name>** dialog, select the sites to which you want to deploy the template.

**Step 9** Click **Save** to save the template changes.

**Step 10** Deploy the template to site(s).

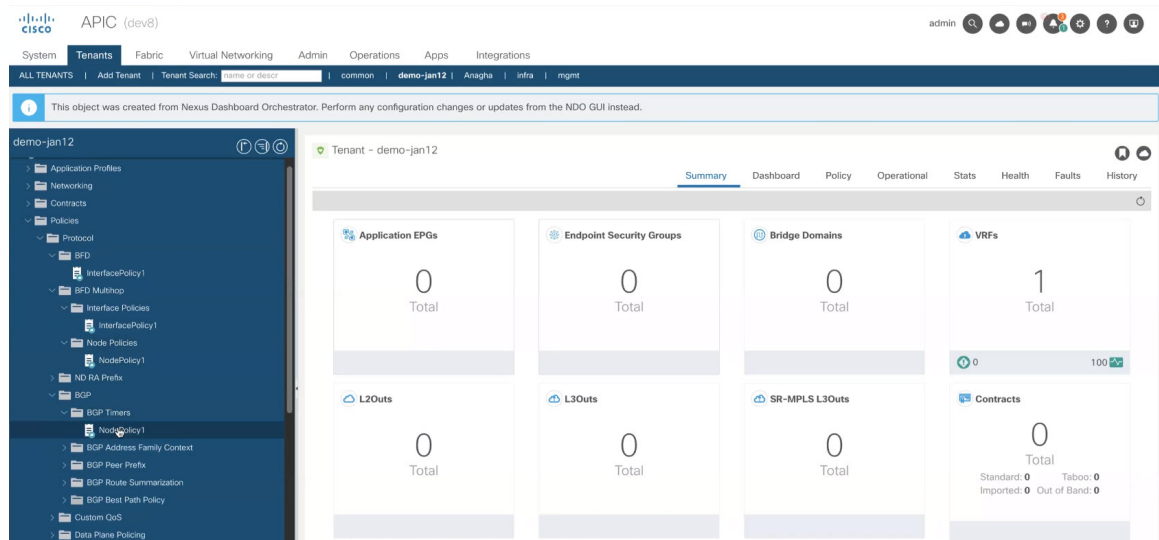
At this stage, we can deploy the created group policies to the sites and verify them in your APIC as a checkpoint before proceeding with additional configurations.

- a) In the **Tenant Policies** template view, click **Deploy**.  
 b) In the **Deploy to sites** dialog, confirm the policies being deployed and click **Deploy**.



- c) (Optional) Verify that the policies were deployed correctly.

You can verify that the template was correctly deployed to a site by navigating to the site's APIC, choosing **Tenants** > **<tenant-name>** > **Policies** > **Protocol** and checking the **BFD**, **BGP**, and **OSPF** policies, for example:



Note that while the policies are shown and managed individually on the APIC, NDO simplifies the configuration workflow by combining them into a single template at the node and interface levels.

### Step 11

Create a BGP Peer Prefix policy.

- In the main pane, choose **Create Object > BGP Peer Prefix Policy**.
- Provide the **Name** for the policy, and define the **Max Number of Prefixes** as well as the **Action** to take if the number is exceeded.

The following actions are available:

- **Log**
- **Reject**
- **Restart**
- **Shutdown**

### Step 12

Create an IPSLA Monitoring policy.

- In the main pane, choose **Create Object > IPSLA Monitoring Policy**.
- Provide the **Name** for the policy, and define its settings.

#### Note

If you choose **HTTP** for the **SLA Type**, your fabric must be running Cisco APIC release 5.1(3) or later.

### Step 13

Create an IPSLA Track List.

- In the main pane, choose **Create Object > IPSLA Track List**.
- Provide the **Name** for the policy.
- Choose the **Type**.

The definition of a route being available or not available can be based on **Threshold Percentage** or **Threshold Weight**.

- Click **+Add Track List to Track Member Relation** to add one or more track members to this track list.

**Note**

You must select a bridge domain or an L3Out to associate with the track member. If you do not already have the bridge domain (BD) or L3Out created, you can skip adding a track member, save the policy without assigning one, and come back to it after you have created the BD or L3Out.

- e) In the **Add Track List to Track Member Relation** dialog, provide the **Destination IP**, **Scope Type**, and choose the **IPSLA Monitoring Policy**.

The scope for the track list can be either bridge domain or L3Out. The IPSLA Monitoring policy is the one you created in the previous step.

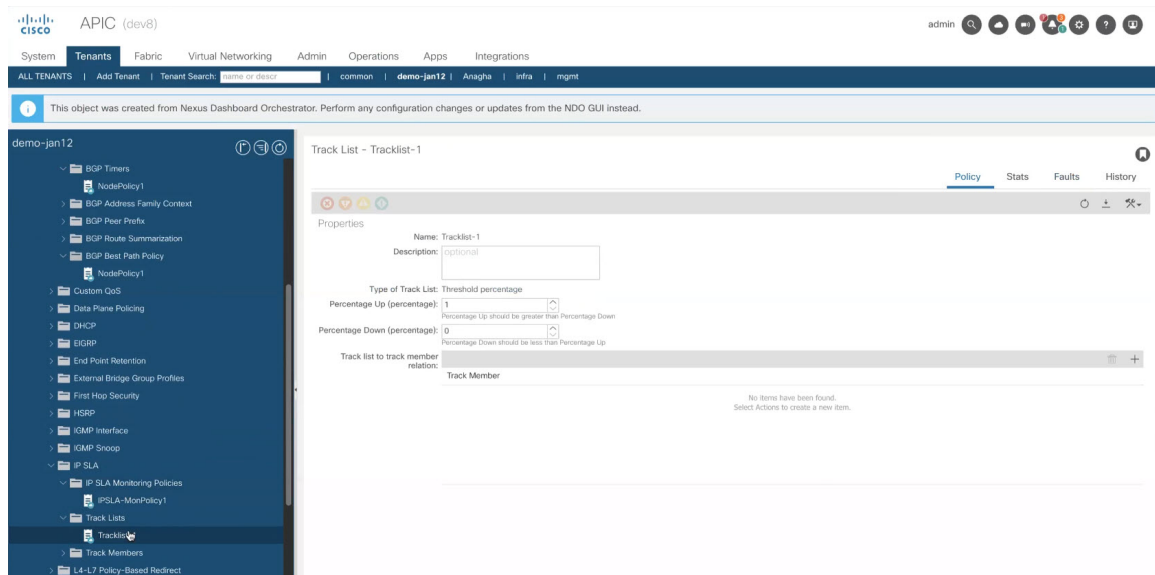
**Step 14** Click **Save** to save the template changes.

**Step 15** Deploy the template to site(s).

At this stage, we can create another checkpoint by deploying the defined policies to the sites.

- a) In the **Tenant Policies** template view, click **Deploy**.
- b) In the **Deploy to sites** dialog, confirm the policies being deployed and click **Deploy**.
- c) (Optional) Verify that the policies were deployed correctly.

You can verify that the template was correctly deployed to a site by navigating to the site's APIC, choosing **Tenants** > **<tenant-name>** > **Policies** > **Protocol** and checking the **IP SLA** policies, for example:

**What to do next**

After you've defined the policies in the Tenant Policy template, proceed to [Creating L3Out Template, on page 12](#).

## Creating L3Out Template

This section describes how to create an L3Out template and define IP-based L3Out policies, which you will then use with the VRFs and EPGs in an Application template to deploy a complete L3Out configuration to

your fabrics. For more information about each policy and how it related to policies and settings in other templates, see [L3Out Template Overview, on page 1](#).

If you are looking to create an SR-MPLS VRF L3Out, see the steps described in [Multi-Site and SR-MPLS L3Out Handoff](#).

### Before you begin

- You must have created a Template Policy template and defined any policies specific to your deployment scenario, as described in [Creating Tenant Policy Template, on page 6](#).
- Created a VRF that you want to use for the L3Out in one of your Application templates as you typically would.

## Procedure

**Step 1** In the left navigation pane, choose **Application Management > L3Out Templates**.

**Step 2** In the main pane, click **Add L3Out Template**.

If you want to update an existing L3Out template instead, simply click on its name. This opens the **L3Out Template** page.

**Step 3** If you are creating a brand new template, choose the **Tenant** and **Site** with which to associate the template, then click **Save and go to template**.

Each L3Out template is associated with a specific tenant similar to other NDO templates, however it is also assigned to a single site only as L3Out configuration is typically site-specific.

If you want to define L3Out configuration for multiple sites, you must create at least one L3Out template for each site, but you can deploy multiple L3Outs per site/tenant by defining all of them in the same L3Out template. You may have multiple L3Out templates per site as long as they are assigned to different tenants.

The screenshot shows a modal dialog titled "Select a Tenant and Site" with a close button (X) in the top right corner. Inside the dialog, there are two labeled dropdown menus. The first is labeled "Tenant \*" and has "demo-jan12" selected. The second is labeled "Site \*" and has "dev8" selected. At the bottom right of the dialog is a blue button with the text "Save and go to template".

**Step 4** Provide the **Name** for the template.

**Step 5** Create an IP-based L3Out and provide its general configuration.

- In the main pane, choose **Create Object > L3Out**.
- Provide the **Name** for the L3Out.

**Note**

We recommend providing unique names for all L3Outs across sites, even if they belong to the same tenant or allow connectivity to the same external resources.

- c) Click **Select VRF>** and choose a VRF to associate with this L3Out.

Note that if you save and deploy the template at this time, the behavior would be identical to what was previously available in NDO release 4.0(x) and earlier. The following steps describe additional settings available in release 4.1(1) and later to allow full L3Out configuration directly from NDO.

- d) Click **Select L3 Domain>** and choose the L3 domain to associate with this L3Out.

The L3 domains can be created directly in the APIC or using the **Fabric Management > Fabric Policies** page in NDO, as described in the [Fabric Management Templates](#) chapter.

- e) Select the **Routing Protocols** used by this L3Out.

You can select **BGP** or **OSPF** or both. Alternatively, you can leave both protocols disabled if you plan to use static routing on this L3Out.

If you enable OSPF, you must also provide the **OSPF Area ID** and **OSPF Area Type**.

For both OSPF and BGP:

- Provide **Outbound Route Map** to advertise fabric's BD subnets or prefixes learned from other L3Outs (transit routing) to the outside.

This is the **Route Map Policy for Route Control** that you created in the previous section.

**Note**

If an Outbound Route Map is specified here, it must include all prefixes which need to be advertised toward the external network domain. BD subnets configured with BD to L3Out associations and External EPG subnets configured with export route control will not work when this route-map configuration is deployed.

- Enable **Import Route Control**, to control the external prefixes that should be redistributed inside the fabric.

- f) In the **Nodes** area, click **+Add Node** and specify

**Step 6**

Add one or more border leaf switches (nodes) for the L3Out.

- Click **+Add Node**.
- In the **Add Node** dialog, choose a **Node ID**.
- Provide the **Router ID**.
- (Optional) Choose the **Node Group Policy** that you created for this node.

You can deploy consistent configuration across all nodes by configuring a **Node Group Policy** and applying it to the nodes, as described in Step 9. If you have not created a node group policy yet, you can skip this substep and come back to it later.

- Choose whether you want to **Use Router ID as Loopback**.
- If you want to define one or more static routes, click **+Add Static Routes**.

For all static routes, you must define an IP address **Prefix** including the network mask using the `ab.cd.ef.gh/xy` format, choose whether you want to **Create a static route to Null0**, and define the **Next Hop** IP address. When providing the next hop IP, you can also choose the **Administrative Distance** and the **Monitoring Policy** which you created in [Creating Tenant Policy Template, on page 6](#).

Here you can also select the **Track Policy**, which you defined in [Creating Tenant Policy Template, on page 6](#).

g) Repeat this step for any additional border leaf switches where you want to deploy this L3Out.

## Step 7

Add one or more interfaces for the L3Out.

- Click **+Add Interface**.
- Choose the type of interface you want to add.

This release supports the same interface types as the APIC:

- Routed Interface
- Routed Sub-Interface
- SVI
- Floating SVI

You can use the same configuration parameters as you would typically use when configuring an interface directly in the APIC, for example:

### Add Routed Interface

Interface Type

Port

Direct Port Channel

Node Id

dev8-leaf1 (Node-101)

Interface \*

eth1/8

Interface Group Policy

Addresses

Addresses

☒ IPv4 Primary Address
 

10.1.1.1/24

☒ IPv6 Primary Address
 

10::1/64

Secondary Addresses

Address	ND RA PREFIX	IPv6 DAD
+ Add Secondary Address		

MAC Address \*

00:22:BD:F8:19:FF

MTU Bytes

inherit

L3Out BGP Peers

Peer Address IPv4	Peer Address IPv6
+ Add L3Out BGP Peer	

Advanced Settings

Link Local Address V6

☐ IPv6 DAD

Target DSCP

Unspecified

PTP Configuration

PTP State

Enabled

Disabled

Ok

- Repeat this step for any additional interfaces where you want to deploy this L3Out configuration.

**Step 8** (Optional) Add one or more node or interface group policies.

While you can configure each node and interface individually as mentioned in the previous two steps, you can also define one or more node or interface group policies and apply a group policy to multiple nodes or interfaces for consistent configuration across them.

- Click **+Add Node/Interface Group Policy**.
- Choose whether you're defining a **Node** or **Interface** group policy and provide a **Name** for it.
- Select the **Node Routing Policy** or **Interface Routing Policy** respectively.

**Note**

An Interface Group Policy is mandatory when using OSPF on the L3Out.

This is one of the policies you created in [Creating Tenant Policy Template, on page 6](#), for example:

- Provide any additional node or interface configurations settings as required by your deployment.  
Keep in mind that all nodes or interfaces to which you apply this group policy will have exact same configuration as defined in the group policy.
- Click **Ok** to save the group policy.
- Repeat this step for any additional node or interface group policies for this L3Out.

**Step 9** (Optional) Apply a node or interface group policy to one or more nodes/interfaces.

- Click on one of the nodes or interfaces you configured for this L3Out.
- From the **Node/Interface Group Policy** dropdown, select the group policy you defined in the previous step.



- c) Repeat this step for all nodes and interfaces to which you want to apply the consistent settings defined by the group policies.

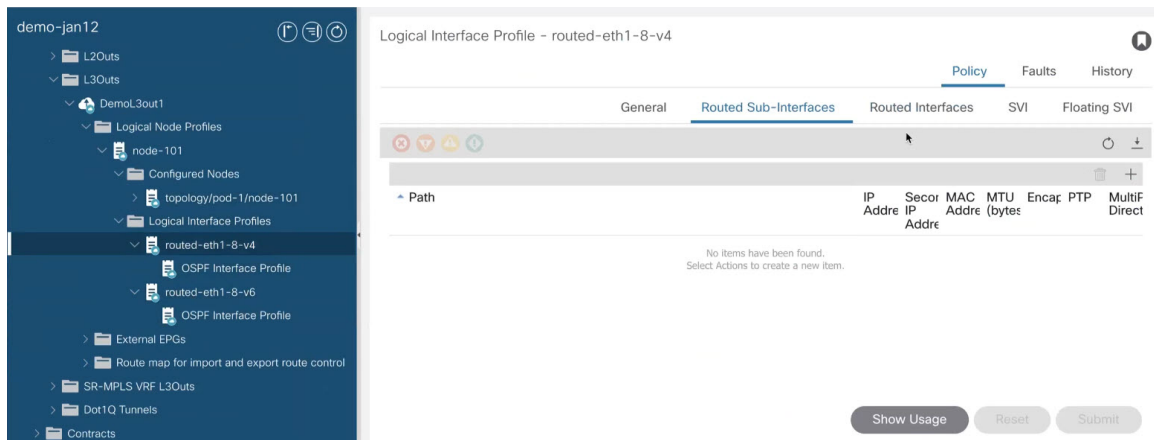
**Step 10** Click **Save** to save the template changes.

**Step 11** Deploy the template to site.

- In the **L3Out Template** page, click **Deploy**.
- In the **Deploy to sites** dialog, confirm the policies being deployed and click **Deploy**.
- (Optional) Verify that the policies were deployed correctly.

You can verify that the template was correctly deployed to a site by navigating to the site's APIC, choosing **Tenants** > *<tenant-name>* > **Networking** > **L3Outs** and checking the L3Out name you provided in NDO.

Note that while you define all of the L3Out configurations in the same template in NDO, individual policies are created in the APIC. For example, separate policies are created for the nodes, interfaces, and even IP address types (providing IPv4 and IPv6 IP addresses for a single L3Out interface creates two separate interface profiles) in the APIC:



# Importing Existing L3Out Configuration

## Overview of Importing L3Out Configuration

Beginning with release 4.1(2), Nexus Dashboard Orchestrator (NDO) supports importing existing L3Out configurations from the APIC sites. The following sections focus on the guidelines and specific steps required to import an L3Out along with its associated policies.



**Note** If you want to configure and deploy new IP-based L3Out configurations (greenfield deployment), see the earlier sections of this chapter.

If you want to configure or import SR-MPLS VRF L3Out, see the [Multi-Site and SR-MPLS L3Out Handoff](#) chapter instead.

This release supports importing the following policies.

- **Route Maps** – may be referenced in the L3Out template's **Outbound Route Map** and **Inbound Route Map** fields to define route import and export policies.
- **L3Out Node Routing:**
  - Nodes configured for an L3Out can be associated to a node group, which in turn can refer to a node routing policy.
  - Node groups can also reference BGP Peer Prefix policy when configuring BGP peers for the nodes.
- **L3Out Interface Routing:**
  - Interfaces configured for an L3Out can be associated to an interface group, which can refer to an interface routing policy and BGP Peer Prefix policy.
  - Interface groups can also reference BGP Peer Prefix policy when configuring BGP peers for the interfaces.

- **BGP Peer Prefix** – can be referenced by the node and interfaces groups for BGP peer configuration on all nodes in the group.
- **IPSLA Monitoring policies and IPSLA Track lists** – can be referenced by the static routes defined for a node.
- **Custom QoS policy** – can be referenced by interface group configuration.

### Mapping of Sites' MOs to NDO Objects and Groups

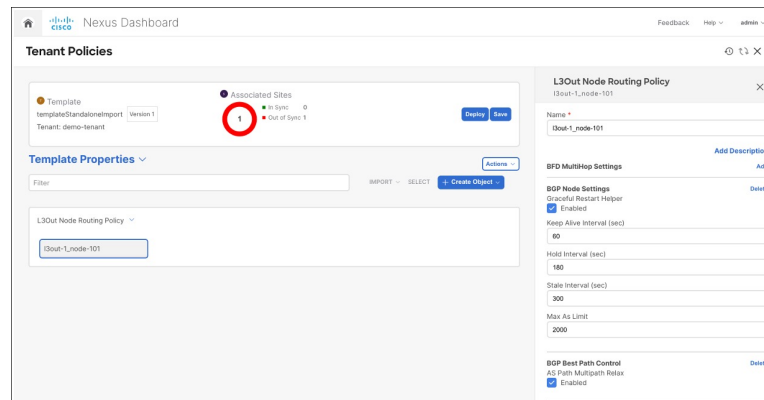
Note that in some cases there is no 1:1 mapping between the managed objects (MOs) created in the site and the policy objects as they are seen on and managed by the Orchestrator. In these cases, when you import an L3Out from APIC, NDO creates NDO-specific logical groups that may contain multiple individual MOs; for example, the following APIC policies are grouped on import:

- The following MOs are grouped into an L3Out Node Routing policy on NDO:
  - BGP Timer Policy
  - BGP Best Path Policy
  - BFD Multi-Hop Node Policy



**Note** If you import an L3Out configuration and then later change one of these policies directly on the APIC, you must re-import the policies in the Tenant Policy template that contains them on NDO.

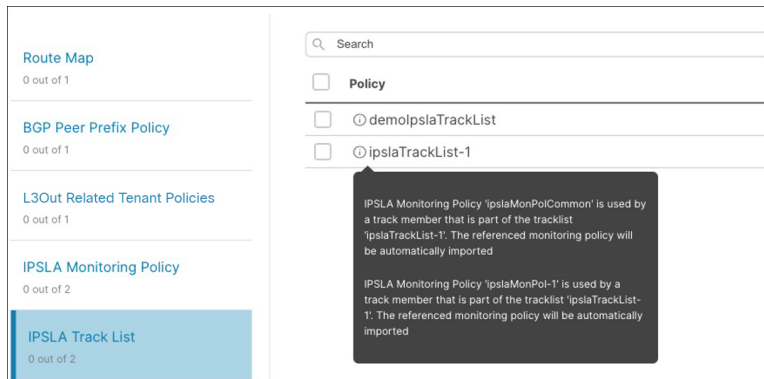
The following figure shows the **L3Out Node Routing Policy** object in NDO that groups together the 3 policies mentioned above:



- The following MOs are grouped into an L3Out Interface Routing policy on NDO:
  - OSPF Interface Policy
  - BFD Policy
  - BFD Multi-Hop Interface Policy

## Automatic Import of Dependencies

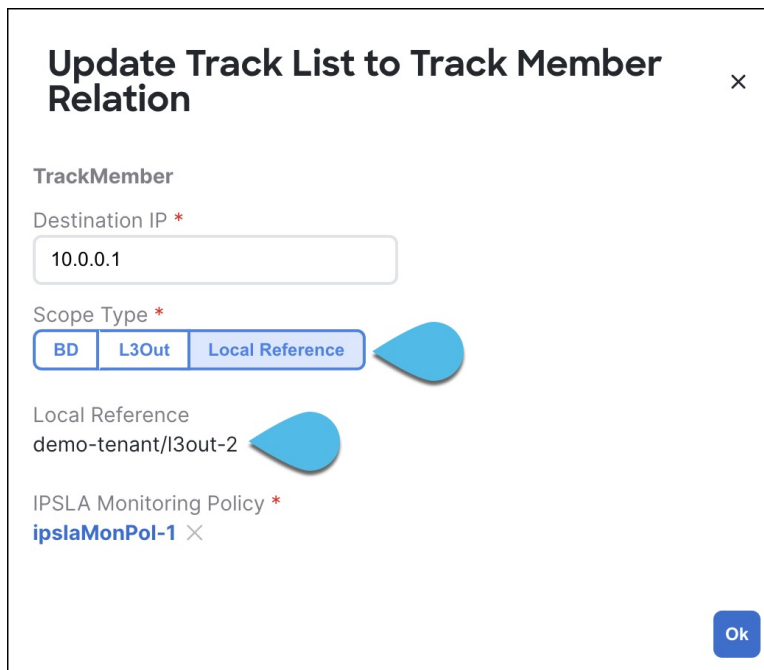
Tenant Policies templates include objects and policies that have local references within the template. For example, an IPSLA track list can contain a list of track members and each track member must refer to a IPSLA monitoring policy. In such cases, importing existing configuration that contains one or more IPSLA track list policies from a site will also automatically import the referenced IPSLA monitoring policy. The import workflow displays additional information about the automatically imported policies when you select an object that has such dependencies:



## Importing IPSLA Policies

Typically, IPSLA track members have a Bridge Domain (BD) or an L3Out scope. When you import an IPSLA track list along with its members, NDO will attempt to automatically assign the correct BD or L3Out to those members. However, at the time of the import, the BD or L3Out objects may not yet exist in NDO.

In such cases, NDO still allows you to import the IPSLA track members with a missing scope object reference. To keep track of the correct reference, NDO sets **Scope Type** to `Local Reference` and saves the name of the referenced BD or L3Out in a `scopeDn` property of the IPSLA track member object:



This allows you to save the template that contains the imported IPSLA track members and re-deploy it back to the site, where the `scopeDn` value is used to correctly program the scope reference for the policy.

To import the entire L3Out configuration, you need to import the L3Out objects after you've imported the relevant Tenant policies. So in case where you import the IPSLA track members first, you must manually update their **Scope Type** and reference after you have also imported the associated L3Out. The `scopeDn` and `scopeType=Local Reference` are internal values and can be set only by the configuration import workflow.

### References to Policies in Tenant "Common"

Some policies that you import from a site may contain references to policies in tenant `common`. Importing such policies will automatically create a copy of the tenant `common` policy in the Tenant Policies template where the objects are being imported and as a result of that, in the tenant associated with that Tenant Policies template, for example:

- If you import an IPSLA track list that contains a track member which refers to an IPSLA monitoring policy from the `common` tenant, a copy of the tenant `common`'s IPSLA monitoring policy will be created in the Tenant Policies template and the imported track member will reference this newly added IPSLA monitoring policy.
- If you import an L3Out that contains node configuration with a static route which references an IPSLA track list from tenant `common`, a copy of the tenant `common`'s IPSLA track list will be created in the Tenant Policies template.


### Unsupported Scenarios

If an L3Out contains one or more configuration options that are currently not supported by NDO, you will not be able to import that L3Out. The following configurations are currently not supported by NDO and will prevent you from importing any L3Out that includes them:

- For IP-based L3Outs:
  - Layer 3 EVPN Services for Fabric WAN (GOLF)
  - Enhanced Interior Gateway Routing Protocol (EIGRP)
  - Fallback Route Group
- For node profiles:
  - Intersite Loopback Addresses
- For interfaces:
  - DHCP Relay
  - SVI/FSVI External Bridge Group Profile
  - VXLAN Encap
- For interface profiles:
  - Internet Group Management Protocol (IGMP)
  - Hot Standby Router Protocol (HSRP) Interfaces
  - DHCP Relay

- Enhanced Interior Gateway Routing Protocol (EIGRP)
- Ingress/Egress Data Plane Policies
- Neighbor Discovery (ND) Policy
- PIM and PIMv6 Interface Policies
- NetFlow Monitor Policies

In these cases, the import workflow UI will display an orange exclamation point icon with a message explaining the issue and you will not be able to select that L3Out for import:



## Importing Tenant Policy Template Objects

This section describes how to import existing L3Out configuration policies from Cisco APIC into NDO's Tenant Policies template. For more information about each policy and how it relates to policies and settings in other templates, see [Overview of Importing L3Out Configuration, on page 18](#).

### Before you begin

- If you want to configure and deploy new L3Out configurations (greenfield deployment), see the earlier sections of this chapter instead.
- You must have the Nexus Dashboard Orchestrator service installed and enabled.
- You must have the fabrics onboarded to your Nexus Dashboard and enabled for management in the Orchestrator service.
- Ensure you have read and understood the Templates and Policy Objects dependencies described in [Overview of Importing L3Out Configuration, on page 18](#).
- Ensure that no changes are made to the tenant policies or the L3Outs which you plan to import into NDO between the time you import the tenant policies as described in this section and when you re-deploy the imported L3Out as described in the next section.

There's no notification in NDO in case an imported policy used by the L3Out is modified directly in the APIC before all policies associated with the L3Out are imported and re-deployed to be managed by NDO.

### Procedure

- 
- Step 1** Log in to your Nexus Dashboard and open the Orchestrator service.
- Step 2** In the left navigation pane, choose **Application Management > Tenant Policies**.

**Step 3** In the main pane, click **Add Tenant Policy Template**.

If you want to update an existing Tenant Policy template instead, simply click on its name. This opens the **Tenant Policies** page.

**Step 4** If you created a brand new template, provide the **Name** for the template and **Select a Tenant** from which you plan to import configuration.

**Step 5** Associate the template with the site from which you plan to import configuration.

- a) In the **Tenant Policies** template view, choose **Actions > Sites Association**.
- b) In the **Associate Sites to <template-name>** dialog, select the sites to which you want to deploy the template.

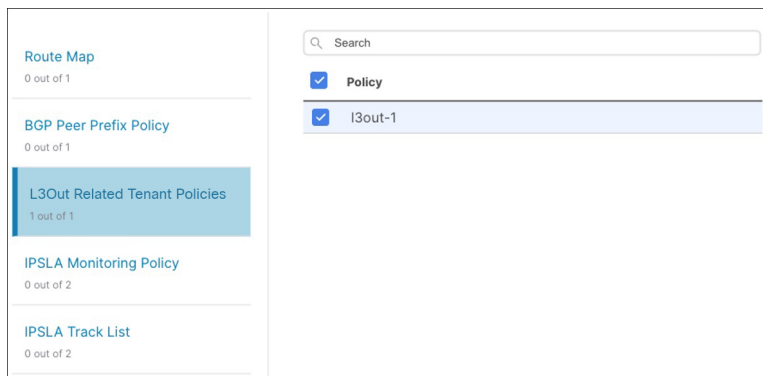
**Step 6** Click **Save** to save the template changes.

**Step 7** Import one or more policies into the Tenant Policies template.

When you choose to import L3Out configuration from a site, the UI will show the list of L3Out policies that can be imported. You may select one or more L3Out policies and import all the provider policies used by the L3Out into this Tenant Policy Template

- a) In the **Tenant Policies** screen's **Template Properties** view, choose **Import > <site-name>**.
- b) In the **Import from <site-name>** dialog, choose one or more L3Outs and click **Import**.

If there's an L3Out already configured in the site, its associated policies will be available for import in the **L3Out Related Tenant Policies** category. When you select an L3OutSource to import, all policies referenced by that L3Out in the site's APIC are imported into the Tenant Policies template you are editing.



- c) Verify that all imported policies are shown in the template and click **Save** to save it.

All policies configured for the L3Out in the site, which you chose to import in the previous step, are added to the Tenant Policies template using the following guidelines:

- Default import route maps are named as `<l3out-name>_imp_<site-id>`.
- Default export route maps are named as `<l3out-name>_exp_<site-id>`.
- Node routing policies are numbered, for example `L3OutNodePolicy1`, `L3OutNodePolicy2`, etc.
- Interface routing policies are numbered, for example `L3OutInterfacePolicy1`, `L3OutInterfacePolicy2`, etc.

**Tenant Policies**

Template

Template2 Version 1

Tenant: demo-tenant

Associated Sites

1

In Sync 0

Out of Sync 1

Deploy

Save

**Template Properties** ▾

Filter

IMPORT ▾ SELECT **+ Create Object** ▾

Route Map Policy for Route Control ▾

Custom QoS Policy ▾

L3Out Node Routing Policy ▾

L3Out Interface Routing Policy ▾

BGP Peer Prefix Policy ▾

- d) If necessary, update the policy names and click **Save** to save the changes.

We recommend keeping the names of the imported policies as they are created. In this case, when you import L3Outs into the L3Out template as described in the next section, the referenced policies will be automatically recognized and configured for the L3Out by NDO.

However, if you have a specific naming convention in your Multi-Site domain, you can update the imported objects' names to follow that convention. In this case, you will need to manually provide object references during L3Out import in the next section.

**Note**

For some objects, there is no 1:1 mapping between the managed objects (MOs) created in the site and the policy objects as they are seen on and managed by the Orchestrator. For information about which MOs are combined into logical groups in NDO, see [Overview of Importing L3Out Configuration, on page 18](#).

**Step 8** Deploy the template to site(s).

After you have imported the policies and saved the template, you must deploy it back to the site.

**Note**

If the names of the imported objects used in NDO do not match the names of those objects in the APIC, NDO does not create new objects in the APIC and simply starts managing the original ones.

However, if you make other changes to a policy object before deploying it back to the site, NDO will create a new object in the APIC.

- a) In the **Tenant Policies** template view, click **Deploy**.



- b) In the **Deploy to sites** dialog, confirm the policies being deployed and click **Deploy**.

### What to do next

After you've defined the policies in the Tenant Policy template, proceed to [Importing L3Out Objects, on page 25](#).

## Importing L3Out Objects

This section describes how to import an L3Out template from an APIC site into Nexus Dashboard Orchestrator. For more information about each policy and how it related to policies and settings in other templates, see [Overview of Importing L3Out Configuration, on page 18](#).

### Before you begin

- If you want to configure and deploy new L3Out configurations (greenfield deployment), see the earlier sections of this chapter instead.
- You must have created a Template Policy template and imported the policies associated with the L3Out you want to import, as described in [Importing Tenant Policy Template Objects, on page 22](#).

### Procedure

**Step 1** In the left navigation pane, choose **Application Management > L3Out Templates**.

**Step 2** In the main pane, click **Add L3Out Template**.

If you want to update an existing L3Out template instead, simply click on its name. This opens the **L3Out Template** page.

**Step 3** If you are creating a brand new template, choose the **Tenant** and **Site** from which you will import the L3Out configuration, then click **Save and go to template**.

Each L3Out template is associated with a specific tenant similar to other NDO templates, however it is also assigned to a single site only as L3Out configuration is typically site-specific.

If you want to import L3Out configuration for multiple sites, you must create at least one L3Out template for each site, but you can import multiple L3Outs per site/tenant into the same template or you may choose to have multiple L3Out templates per site as long as they are assigned to different tenants.

**Step 4** If you created a brand new template, provide the **Name** for the template and click **Save**.

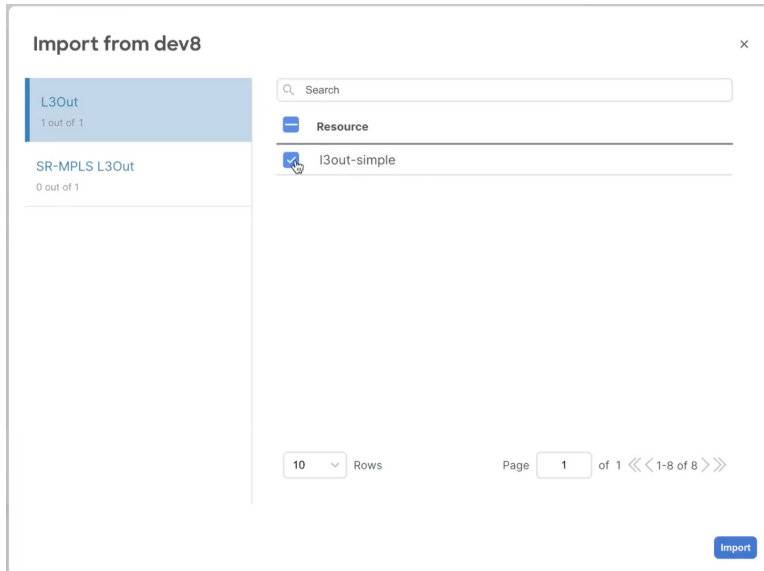
You must save a brand new template before you can add new or import existing configuration.

**Step 5** Import an L3Out from the site.

- In the main pane, click **Import**.
- In the **Import from <site-name>** dialog, select the **L3Out** you want to import and click **Import**.

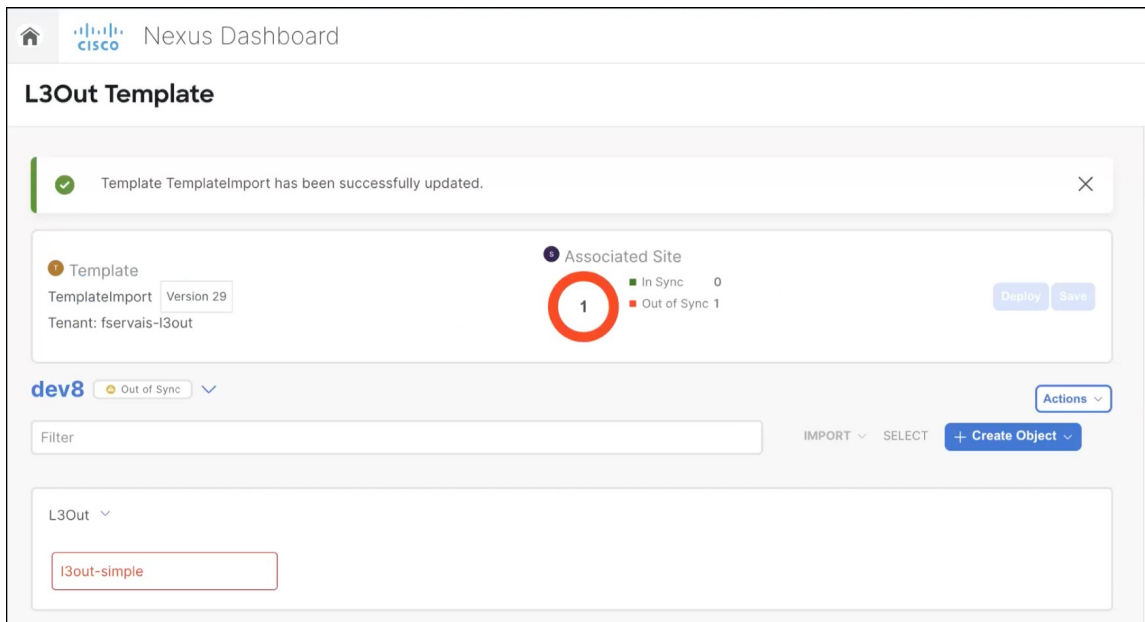
### Note

If an L3Out has one or more tenant policy references that are not found in NDO's Tenant Policies templates, you cannot import that L3Out and must first import those references as described in [Importing Tenant Policy Template Objects](#), on page 22.



**Step 6** Provide any missing information for the imported L3Out.

When you first import an L3Out, the object in the UI may be shown in red if some of the L3Out settings are not imported and must be provided manually:



For example, if BGP Peer configuration is present on the L3Out, NDO will enforce authentication when the L3Out is imported. In this case, you must manually navigate to the authentication settings and either disable **Password Authentication** or provide a valid password:

- a) Select the imported L3Out.
- b) Click on the setting that shows a warning.

**L3Out**  
l3out-simple

Name \*  
l3out-simple

[Add Description](#)

**VRF**  
vrf1

**L3 Domain**  
l3out-fservais

**Routing Protocol**  
☒ BGP  
☐ OSPF

**Outbound Route Map**  
[Select Outbound Route Map](#)

**Import Route Control**  
☐ Enabled

**Nodes**

Node ID	Router ID	Common Node Configuration
101	3.4.5.6	

[Add Node](#)

**Interfaces**

Type	Node ID	Pod ID	Group
eth1/24 Type: Routed Interface	101	1	
eth1/20.200 Type: Routed Sub-Interface	101	1	

- c) Click on the setting that shows a warning again.

**L3Out BGP Peers**

Peer Address IPv4	Peer Address IPv6
14.1.1.3	

[Add L3Out BGP Peer](#)

[Advanced Settings](#)

- d) Provide any missing configuration, such as a password.

**Authentication**  
☒ Password Authentication

Password  
1

- e) Repeat this step for all other warnings in the template for the imported objects.

**Step 7**

Click **Save** to save the template changes.

**Step 8**

If necessary, update any previously imported IPSLA track members that reference the L3Out you imported in the previous step.

If you have imported one or more IPSLA track members in the previous section which reference the L3Out you are importing, you must manually update the track members' scopes and references after you've imported the L3Out. Additional details about this behavior are described in [Importing IPSLA Policies, on page 20](#).

- a) Ensure that you have saved the L3Out template with the imported L3Out objects.
- b) Navigate to **Application Management > Tenant Policies**.
- c) Choose the Tenant Policies template that contains the IPSLA track members.
- d) Choose the IPSLA Track List policy.
- e) In the right properties sidebar, click the **Edit** icon next to the Track Member List you want to update.

- f) In the **Update Track List to Track Member Relation** dialog, update the **Scope Type** and choose the scope object. The current values are set to `Local Reference` and the name of the referenced object:

You must update the scope type to `L3Out` and then choose the L3Out you imported in the previous step.

- g) Click **Ok** to save the changes.
- h) Click **Save** to save the Tenant Policies template.
- i) Click **Deploy** to redeploy the template to the site.
- j) Return to **Application Management > Tenant Policies** and choose the L3Out template that you were editing in the previous step.

## Step 9 Deploy the L3Out template to site.

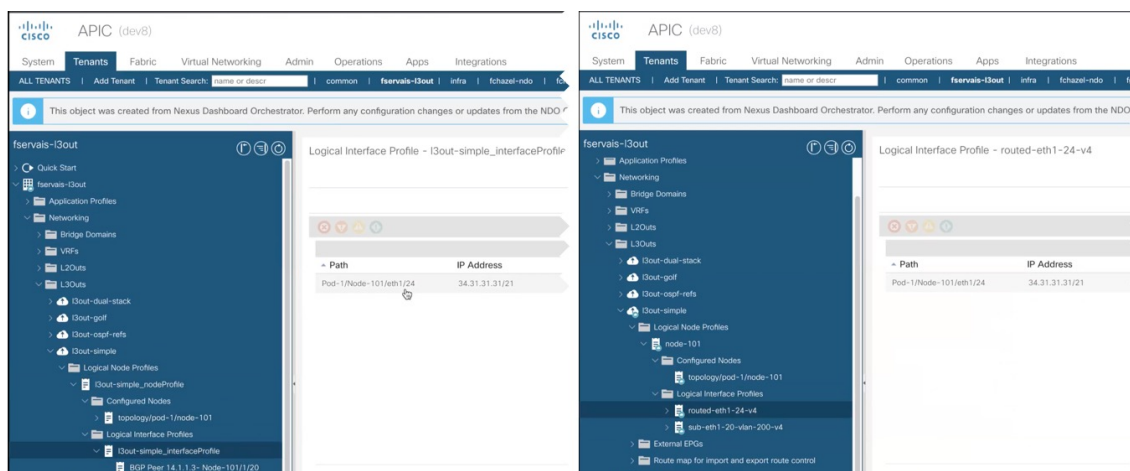
After you have imported the L3Out and saved the template, you must deploy it back to the site.

- a) In the **L3Out Template** page, click **Deploy**.
- b) In the **Deploy to sites** dialog, confirm the policies being deployed and click **Deploy**.
- c) (Optional) Verify that the policies were deployed correctly.

You can verify that the template was correctly deployed to a site by navigating to the site's APIC, choosing **Tenants > <tenant-name> > Networking > L3Outs** and checking that the L3Out name is consistent with the one you imported into the NDO template.

### Note

When the configuration is deployed from NDO back to the site, the old MOs are removed and new ones are created with NDO-specific hierarchy, which may cause a brief (up to 1 second) traffic interruption:



## Viewing L3Out Neighbors

Beginning with release 4.1(2), Nexus Dashboard Orchestrator provides a unified view of all L3Outs and their neighbors in your Multi-Site domain. This information provides visibility into operational data reported by the fabrics controller about site-level connectivity and simplifies troubleshooting by showing the various Layer 3 adjacencies (neighbors) for each L3Out.

### Procedure

- Step 1** In the left navigation pane, choose **Sites**.
- Step 2** Click the name of the site for which you want to view L3Out neighbors.
- Step 3** In the site information page, choose **Connectivity > L3 Neighbors**.

The **L3 Neighbors** page provides a unified view of all neighbors based on L3Out configuration for that site. You can **Filter** or sort the page based on each column.

At any time, you can click **Refresh** to pull the latest information from the site's controller.

Nexus Dashboard

ifav141-s1-apic1

Overview Inventory **Connectivity**

Physical Interfaces Port Channels VPCs **L3 Neighbors**

Filter by attributes

Neighbor	Local Switch	L3Out	Template Name	Routing Protocol	VRF	Operational State
12.1.1.1	ifav141-s1-leaf1	L3OutForkiranVRF2_BG...	t12-ext-conn	BGP	t12:kiranVRFStretched1	Idle
138.1.1.2	ifav141-s1-leaf1	L3outForkiranVRF2BGP	t12-bgp	BGP	t12:kiranVRFStretched1	Established
12.1.1.1	ifav141-s1-leaf2	L3OutForkiranVRF2_BG...	t12-ext-conn	BGP	t12:kiranVRFStretched1	Idle
138.1.1.2	ifav141-s1-leaf2	L3outForkiranVRF2BGP	t12-bgp	BGP	t12:kiranVRFStretched1	Established
1.22.22.2	ifav141-s1-leaf1	L3OutForkiranBGPSubl...	t12-bgp	BGP	t12:kiranVRF2	Established

**Step 4** Click an entry in the **Neighbor** column to view that neighbor's details.

Here, you can view the **Local Switch** information (including its name, IP address, ASN, interface information, etc.) and **Neighbor Details** (such as its IP address, ASN, route ID, port, etc.).

For example, the following two figures show sample information for a BGP and OSPF L3Out neighbors:

**BGP Neighbor Details**

Local Switch Details							
Name	Local IP	ASN	Interface Type	Interface	Router ID	Port	VRF
F2-P1-Leaf-304	10.110.2.2	65002	Routed Sub-interface	eth1/16	1.1.1.104	36597	L3-Demo:VRF
Authentication Disabled							
Neighbor Details							
Neighbor IP	ASN	Router ID	Port	Neighbor Status	Uptime		
10.110.2.3	65111	111.1.1.1	179	Established	1 Weeks, 4 Days		

**OSPF Neighbor Details**

Local Switch Details							
Name	Router ID	Interface Type	MTU	Interface	Encap	Interface IP Address	VRF
F2-P1-Leaf-304	1.1.1.104	SVI	1500	L303-304-VPC11	vlan-802	10.82.1.2	L3-Demo:VRF
OSPF Area	Network Type	Interface Controls Enabled					
backbone	Broadcast	-					
Neighbor Details							
Neighbor ID	Interface IP Address	Neighbor Status	Uptime				
1.1.1.103	10.82.1.1	Full/BDR	1 Weeks				

**Step 5** If the displayed information is not accurate, verify L3Out configurations.

If the L3Out neighbors are not present in the table view:

- Verify that the L3Out policy is configured in NDO and deployed successfully. The information is displayed only for L3Outs that are configured in NDO.
- Verify that the L3Out neighbors are present in NDO's inventory using the APIs.

- For BGP: `GET /mso/api/v1/inventorybgpneighbors?status.fabric=<site-id>`
- For OSPF: `GET /mso/api/v1/inventoryospfneighbors?status.fabric=<site-id>`

If the L3Out neighbors' operational state is not green:

- Verify that the switch interfaces are not in the `shut` state on either of the switches.
  - Verify that the protocol settings are configured correctly and there is no mismatch in the peer device configuration.
    - For BGP, check the authentication, eBGP MultiHop TTL, and ASN are configured correctly.
    - For OSPF, check authentication, Area ID, and MTU configurations.
-

