

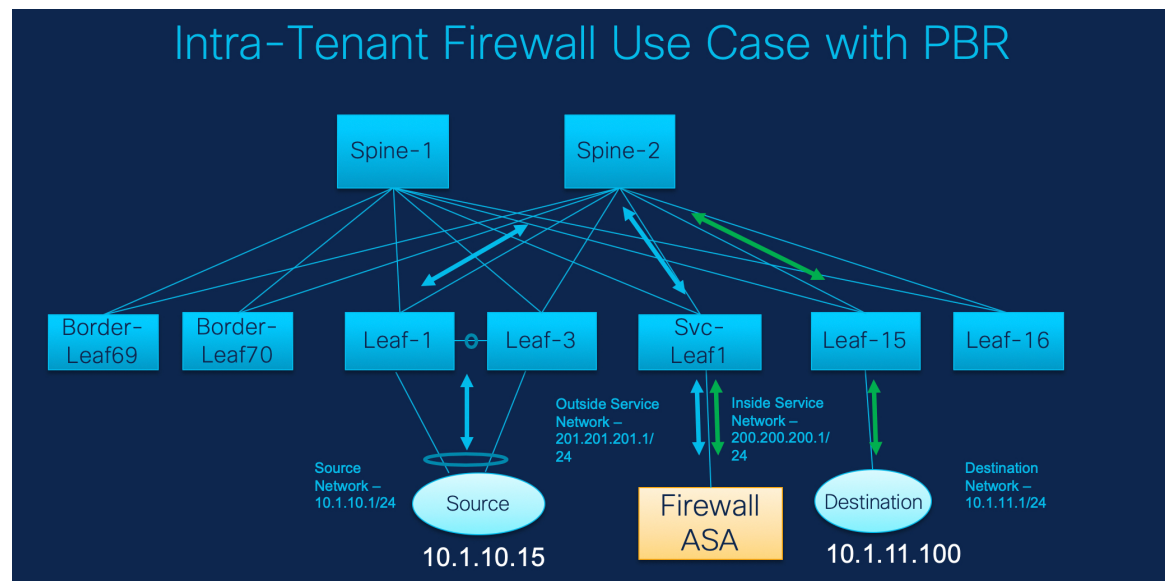


L4-L7 Service Use Cases

- Use Case: Intra-tenant Firewall with Policy-based Routing, on page 1
- Use Case: Inter-tenant Firewall with eBGP Peering, on page 20
- Use Case: One-arm Load Balancer, on page 27

Use Case: Intra-tenant Firewall with Policy-based Routing

Refer the figure given below for topology details.



In this topology, Leaf1 and Leaf3 are a vPC pair and they are connected to **Source** (10.1.10.15) with the **Source Network** (10.1.10.1/24). The service leaf is connected to the virtual **Firewall ASA** and Leaf-15 is connected to **Destination** (10.1.11.100). In this use case, the source network refers to 'client' and the destination refers to 'server'.

Any traffic that is traversing from **Source** to **Destination** must go to the outside service network, and the firewall performs its function by allowing or denying traffic. This traffic is then routed to the inside service network and on to the Destination network. Since the topology is stateful, the traffic coming back from the destination to the source follows the same path.

1. Create Service Node

Now, let us see how to perform service redirection in DCNM.



- Note**
- This use-case does not cover how to provision the **Site_A** VXLAN fabric. For information about this topic, refer to the Cisco DCNM LAN Fabric Configuration Guide.
 - This use-case does not cover configurations on the service node (firewall or load balancer).

Select **Control > Fabrics > Services**.

This use-case consists of the following steps:

1. Create Service Node

Procedure

Step 1 From the **Scope** drop-down list, select **Site_A**.

Step 2 Click the **Add** icon in the **Service Nodes** window.

Step 3 Enter the node name and specify **Firewall** in the **Type** dropdown box. The **Service Node Name** has to be unique.

SCOPE: SITE_A

New Service Nodes

1 Create Service Node

Create Service Node

* Service Node Name

* Type

Step 4 From the **Form Factor** drop-down list, select **Virtual**.

SCOPE: SITE_A

New Service Nodes

1 Create Service Node

2 Create Route Peering

3 Create Policy

Create Service Node

* Service Node Name

* Type

* Form Factor

* Service Node Interface

Step 5 In the **Switch Attachment** section, from the **External Fabric** drop-down list, select the external fabric in which the service node (for example, ASA firewall) is located. Note that the service nodes need to belong to the external fabric. This is a prerequisite before creating a service node.

Switch Attachment

* External Fabric

Step 6 Enter the interface name of the service node that will be connected to the service leaf.

* Service Node Interface

Step 7 Select the attached switch that is the service leaf, and the respective interface on the service leaf.

* Attached Switch

* Attached Switch Interface

Step 8 Select the **service_link_trunk** template. DCNM supports trunk, port channel, and vPC link templates. The available link templates in the **Link Template** drop-down list are filtered based on the selected **Attached Switch Interface** type.

Link Template

Step 9 Specify the **General Parameters** and **Advanced** parameters, if required. Some parameters are pre-filled with the default values.

The screenshot shows the configuration interface for route peering. The 'General Parameters' tab is selected. The 'MTU' dropdown is set to 'jumbo', 'SPEED' is set to 'Auto', 'Trunk Allowed Vlans' is set to 'none', and 'Enable BPDU Guard' is set to 'no'. The 'Enable Port Type Fast' and 'Enable Interface' checkboxes are both checked. A 'Next' button is visible at the bottom of the configuration area.

Step 10 Click **Next** to save the created service node.

2. Create Route Peering

Let us now configure the peering between a service leaf and a service node.

Procedure

Step 1 Enter the peering name and select **Intra-Tenant Firewall** from the **Deployment** drop-down list.

The screenshot shows the configuration page for Step 1. The 'Peering Name' field contains 'peering1'. The 'Deployment' dropdown menu is open, showing 'Intra-Tenant Firewall' as the selected option. Below this, the 'Inside Network' section is visible, with a 'VRF' dropdown menu.

Step 2 Under **Inside Network**, from the **VRF** drop-down list, select a VRF that already exists and select **Inside Network** under **Network Type**.

Enter the name of the **Service Network** and specify the **Vlan ID**. You can also click **Propose** to allow DCNM to fetch the next available VLAN ID from the specified service network VLAN ID range in the fabric settings. The default **Service Network Template** is **Service_Network_Universal**.

Under the **General Parameters** tab, specify the gateway address for the service network. Specify the **Next Hop IP Address**. This next hop address has to be within the 'inside service network' subnet. Under the **Advanced** tab, the default **Routing Tag** value is 12345.

Inside Network

* VRF	<input type="text"/>	* Network Type	<input type="text" value="Inside Network"/>
* Service Network	<input type="text" value="service_net_inside"/>	* Vlan ID	<input type="text" value="2300"/> <input type="button" value="Propose"/>
* Service Network Template	<input type="text" value="Service_Network_Universal"/>		

General Parameters Advanced

* IPv4 Gateway/NetMask ⓘ	<input type="text" value="200.200.200.1/24"/>	IPv6 Gateway/Prefix ⓘ	<input type="text"/>
Vlan Name ⓘ	<input type="text"/>	Interface Description	<input type="text"/>
* Next Hop IP Address ⓘ	<input type="text" value="200.200.200.200"/>		

Step 3 Specify the required parameters under **Outside Network** and specify the **Next Hop IP Address for Reverse Traffic**. This next hop address for reverse traffic needs to be within the ‘outside service network’ subnet.

Outside Network

* VRF	<input type="text" value="VRF_51000"/>	* Network Type	<input type="text" value="Outside Network"/>
* Service Network	<input type="text" value="service_net_outside"/>	* Vlan ID	<input type="text" value="2301"/> <input type="button" value="Propose"/>
* Service Network Template	<input type="text" value="Service_Network_Universal"/>		

General Parameters Advanced

* IPv4 Gateway/NetMask ⓘ	<input type="text" value="201.201.201.1/24"/>	IPv6 Gateway/Prefix ⓘ	<input type="text"/>
Vlan Name ⓘ	<input type="text"/>	Interface Description	<input type="text"/>

Next Hop IP Address for Reverse Traffic ⓘ

Step 4 Click **Next** to save the created route peering.

3. Create Service Policy

Procedure

Step 1 Specify a name for the policy and select the route peering from the **Peering Name** drop-down list.

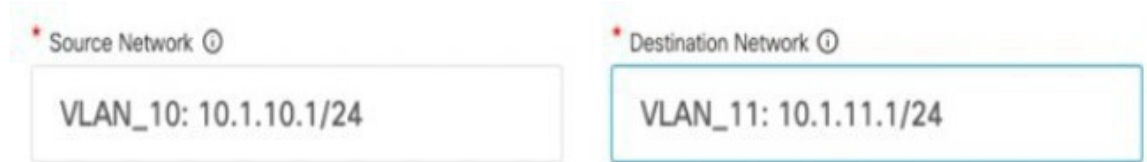
* Policy Name ⓘ	<input type="text" value="policy1"/>	Peering Name	<input type="text" value="peering1"/>
-----------------	--------------------------------------	--------------	---------------------------------------

- Step 2** Select the source and destination VRFs from the **Source VRF Name** and **Destination VRF Name** drop-down lists. The source and destination VRFs for an intra-tenant firewall deployment have to be the same.



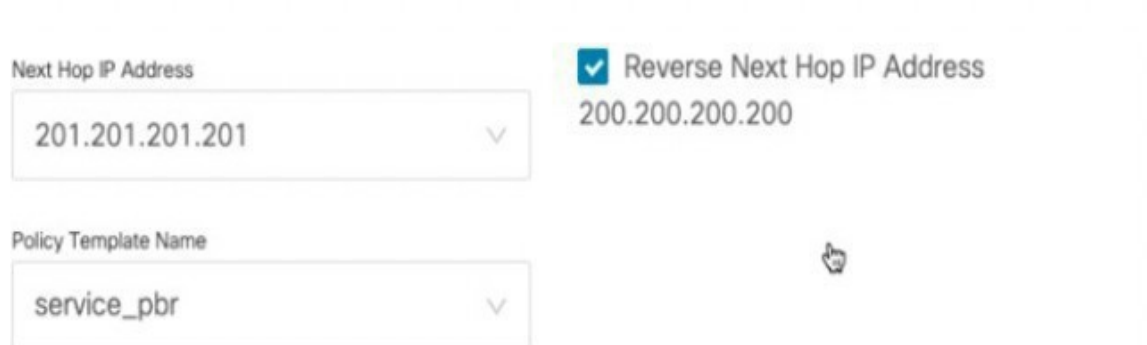
The screenshot shows two dropdown menus. The first is labeled 'Source VRF Name' and contains the text 'VRF_51000'. The second is labeled 'Destination VRF Name' and also contains the text 'VRF_51000'. Both fields have a small downward arrow icon on the right side.

- Step 3** Select the source and destination networks from the **Source Network** and **Destination Network** drop-down lists, or specify the source or destination network that is within the network subnets defined in the **Control > Fabrics > Networks** window.



The screenshot shows two text input fields. The first is labeled 'Source Network' and contains the text 'VLAN_10: 10.1.10.1/24'. The second is labeled 'Destination Network' and contains the text 'VLAN_11: 10.1.11.1/24'. Both fields have a small information icon (i) on the right side.

- Step 4** The next hop and reverse next hop fields are populated based on the values entered while creating the route peering. Select the check box next to the **Reverse Next Hop IP Address** field to enable policy enforcement on reverse traffic.



The screenshot shows four configuration fields. The first is 'Next Hop IP Address' with a dropdown menu containing '201.201.201.201'. The second is 'Reverse Next Hop IP Address' with a checked checkbox and the text '200.200.200.200'. The third is 'Policy Template Name' with a dropdown menu containing 'service_pbr'. The fourth is a mouse cursor pointing to the right side of the form.

- Step 5** Under the **General Parameters** tab in the policy template, select **ip** from the **Protocol** dropdown list, and specify **any** in the **Source Port** and the **Destination Port** fields.

Note For **ip** and **icmp** protocols, the **any** source and destination port is always used for ACL generation. You can also select a different protocol and specify the corresponding source and destination ports. DCNM will convert well-known port numbers to match the format required by the switch. For example, you can convert port 80 to 'www'.

The screenshot shows the 'General Parameters' tab of a configuration interface. It features three input fields: 'Protocol' with a dropdown menu showing 'ip', 'Source Port' with a text input field containing 'any', and 'Destination Port' with a text input field containing 'any'. Below these fields are two buttons: a light blue 'Back' button and a dark blue rounded 'Create' button.

Step 6 Under the **Advanced** tab, by default, **permit** is selected for **Route Map Action** and **none** is selected for the **Next Hop Option**. You can change these values, and customize the ACL name and route map match sequence number, if required. For more information, refer [Templates](#) in the Layer 4-Layer 7 Service configuration guide.

The screenshot shows the 'Advanced' tab of the configuration interface. It contains six input fields arranged in two columns. The left column has: 'Route Map Action' dropdown set to 'permit', 'ACL Name (auto-generated if not specified)' empty text input, and 'Route map match number (auto-generated if not specified)' empty text input. The right column has: 'Next Hop Option' dropdown set to 'none', 'ACL Name for reversed traffic (auto-generated if not specified)' empty text input, and 'Route map match number for reversed traffic (auto-generated if not specified)' empty text input.

Step 7 Click **Create** to save the created service policy.

This completes the procedures that have to be performed to specify the flows for redirection.

4. Deploy Route Peering

Procedure

Step 1 In the **Service Nodes** window, select the required peering under the **Route Peering** tab.

The screenshot shows the Cisco Data Center Network Manager interface. The 'Service Nodes' window is open for 'ASA1' (VIRTUAL FIREWALL). The 'Route Peering' tab is active. A table lists the peering configuration:

Peering Name	Deployment	Peering Option	Status	VRF	Service Network One Network Name	Service Network One Gateway IP	VRF	Service Network Two Network Name	Service Network Two Gateway IP	Action
peering1	IntraTenantFW	None	NA	VRF_51000	service_net_inside	200.200.200.1/24	VRF_51000	service_net_outside	201.201.201.1	[Toggle]

Step 2 Click the toggle button under **Action** to attach service networks to the service leafs.


The screenshot shows the same interface as Step 1. The 'Action' column for the 'peering1' entry is highlighted with a green box, indicating the toggle button used to attach service networks to the service leafs.

Step 3 Click **Preview** to view the configurations that will be pushed to the service leaf.

The screenshot shows the same interface as Step 2. The 'Preview' button is highlighted, and the 'Status' column for the 'peering1' entry is now 'Pending', indicating that the configurations are being pushed to the service leaf.


Previously, we had created inside and outside service networks. You can view these network configurations that will be pushed to the service leaf.

4. Deploy Route Peering

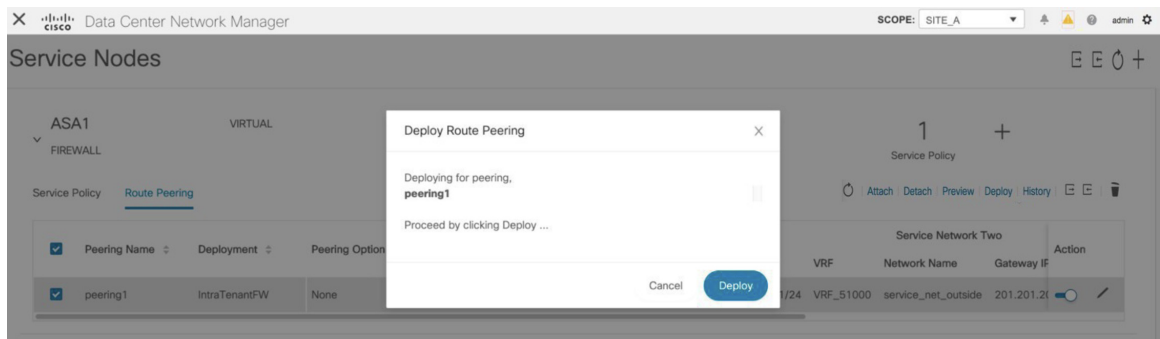


Step 4 Click **Close** to close the **Preview Route Peering** window.

Step 5 Click **Deploy** in the **Service Nodes** window to deploy the configuration to the attached switches (service leaf(s)) for route peering.



Click the **Deploy** button in the pop-up window to confirm deployment.



Step 6 Click the **Refresh** icon for the latest peering configuration attachment and deployment status.

Service Nodes

ASA1
FIREWALL

VIRTUAL

Route Peering

Service Policy

Peering Name	Deployment	Peering Option	Status	Service Network One	Service Network Two	Action
				VRF	VRF	
				Network Name	Network Name	
				Gateway IP	Gateway IP	
peering1	IntraTenantFW	None	In-Sync	VRF_51000	VRF_51000	
				service_net_inside	service_net_outside	
				200.200.200.1/24	201.201.201.2	

5. Deploy Service Policy

Perform the following procedure to deploy the service policy. This policy's corresponding configuration will be deployed to the switches that the source and destination network are attached to, and to the service leaf(s).

Procedure

Step 1 Select the checkbox next to the required policy under the **Service Policy** tab.

Service Nodes

ASA1
FIREWALL

VIRTUAL

Route Peering

Service Policy

Policy Name	Route Peering	Status	Source VRF	Source Network	Destination VRF	Destination Network	Next Hop IP	Reverse Next Hop IP	Action
policy1	peering1	NA	VRF_51000	VLAN_10	VRF_51000	VLAN_11	201.201.201.201	200.200.200.200	

Step 2 Click the toggle button under **Action** to enable this policy.

Service Nodes

ASA1
FIREWALL

VIRTUAL

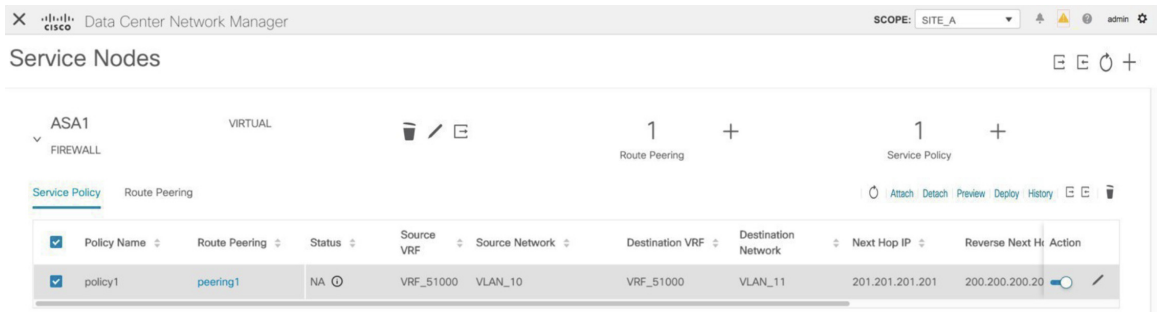
Route Peering

Service Policy

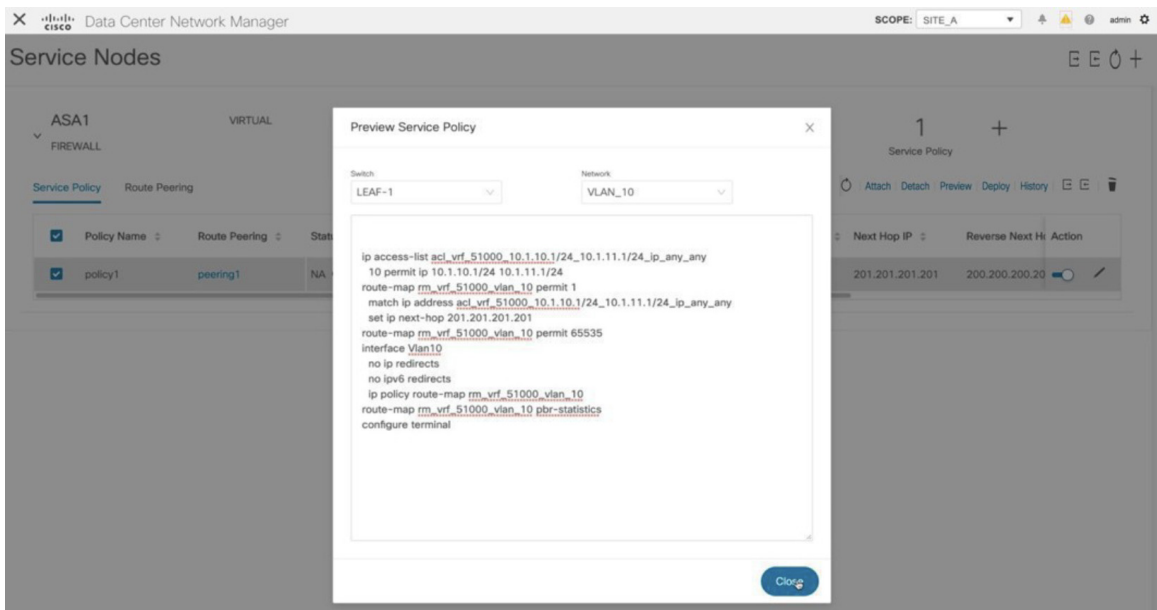
Policy Name	Route Peering	Status	Source VRF	Source Network	Destination VRF	Destination Network	Next Hop IP	Reverse Next Hop IP	Action
policy1	peering1	NA	VRF_51000	VLAN_10	VRF_51000	VLAN_11	201.201.201.201	200.200.200.200	

Step 3 Click **Preview** to view the configuration of the selected network.

5. Deploy Service Policy

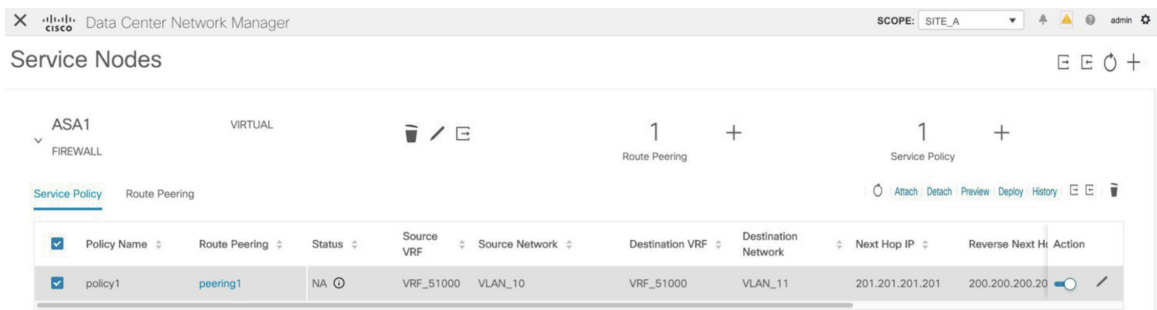


Step 4 Select a switch and a source, destination, or service network, from the drop-down lists to view the intended configuration of a specific source, destination, or service network, on the selected switch. In this window, you can see that there is an access list that will be created with a route map. This configuration will be pushed to the SVI.

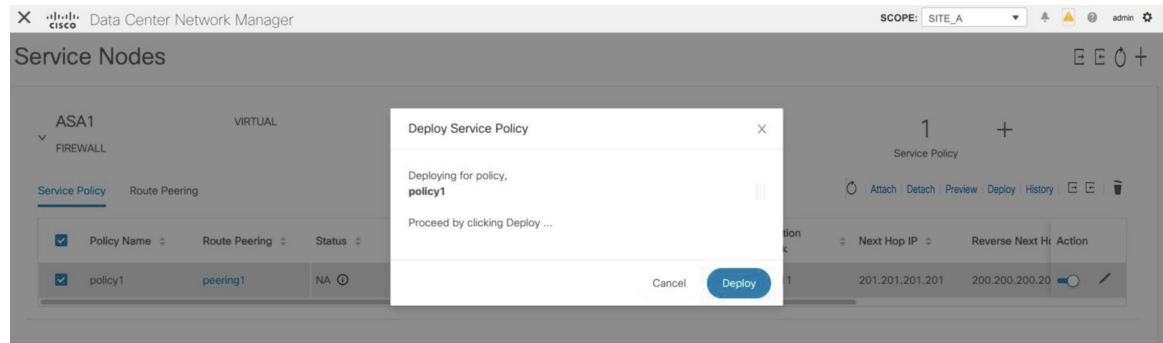


Click **Close** to close the Preview Service Policy window.

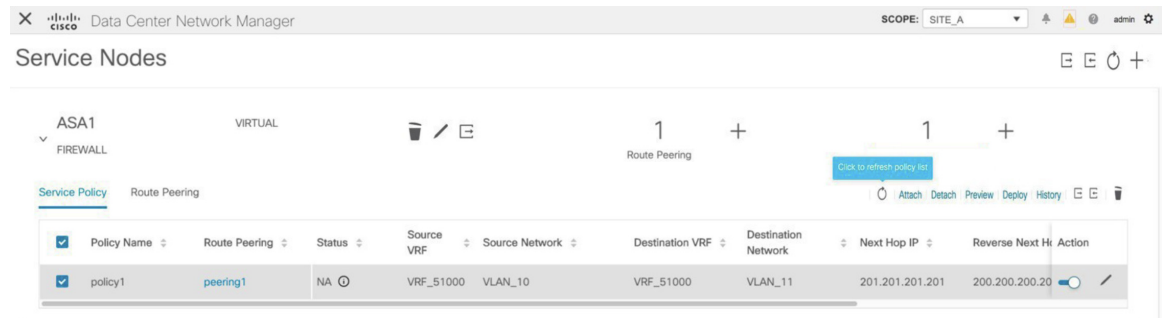
Step 5 Click **Deploy** in the **Service Nodes** window to deploy the configuration to the attached switches (service leaf(s)).



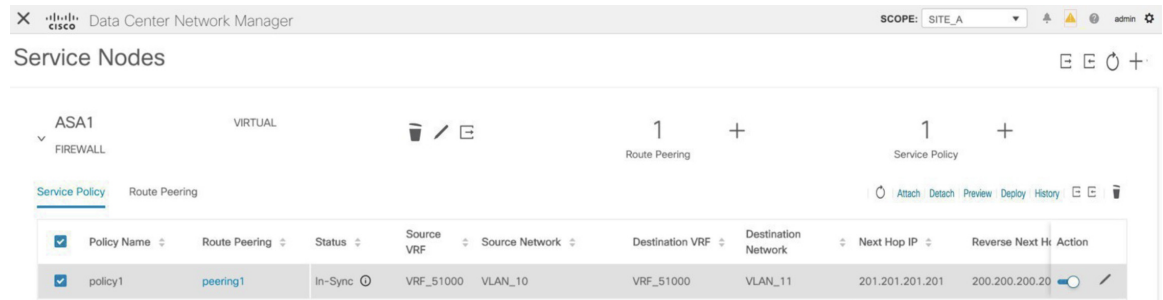
Click the **Deploy** button in the pop-up window to confirm deployment.



Step 6 Click the **Refresh** icon for the latest policy attachment and deployment status.



This policy will be pushed to the switches that the source and destination networks are attached to, as well as the service leaf(s). After pushing the policy, the status column shows **In-Sync**.



6. View Stats

Now that the respective redirection policies are deployed, ping traffic will be redirected to the firewall.

To visualize this scenario in DCNM, click the icon under the **Stats** column.

7. View Traffic Flow in Fabric Builder

The screenshot shows the 'Service Nodes' configuration page in Cisco Data Center Network Manager. It displays a table of service policies for 'ASA1 FIREWALL'. The table has columns for Policy Name, Route Peering, Origin VRF, Destination Network, Next Hop IP, Reverse Next Hop IP, Reverse Enabled, Last Updated, Stats, and Action. The 'Stats' column for 'policy1' is highlighted with a green box.

Policy Name	Route Peering	Origin VRF	Destination Network	Next Hop IP	Reverse Next Hop IP	Reverse Enabled	Last Updated	Stats	Action
policy1	peering1	1000	VLAN_11	201.201.201.201	200.200.200.200	Yes	01/07/2020, 21:26:54		

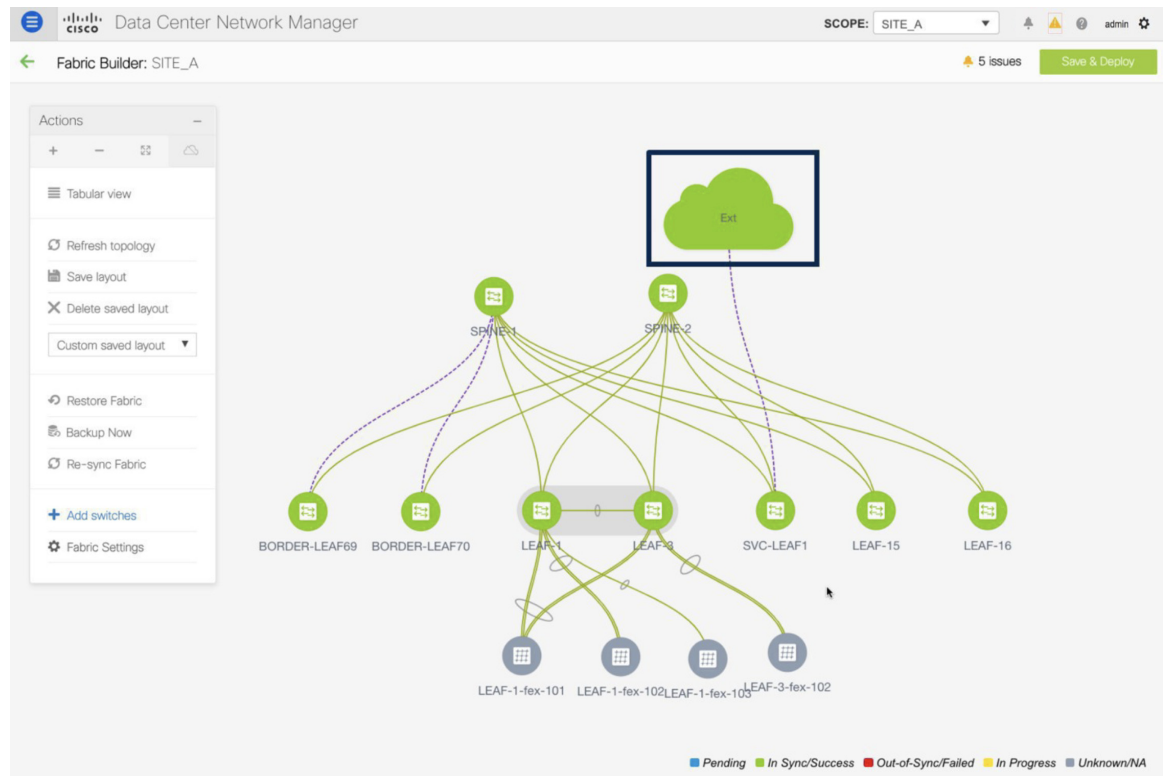
You can view the cumulative statistics for a policy in a specified time range.

The screenshot shows the 'Cumulative Statistics for service policy, policy1' dialog box. It includes a 'Time Range' selector set to 'Jan 08, 2020 | 08:59 - Jan 08, 2020 | 09:59' and a 'Switch' dropdown set to 'LEAF-1'. The graph shows the 'Number of Packets' on the y-axis (0 to 16,000,000) and time on the x-axis (Jan 08, 09:00 to Jan 08, 09:50). A single data series is plotted, showing a steady increase in traffic over the time range.

Statistics are displayed for forwarding traffic on the source switch, for reversed traffic on the destination switch, and for traffic in both directions on the service switch.

7. View Traffic Flow in Fabric Builder

The service node in the external fabric is attached to the service leaf, and this external fabric is shown as a cloud icon in the DCNM topology in the fabric builder.



Procedure

- Step 1** Click the service leaf and click **Show more flows**. You can see the flows that have been redirected.


7. View Traffic Flow in Fabric Builder

The screenshot displays the Cisco Data Center Network Manager interface. The main view shows a network topology with nodes including SPINE-1, SPINE-2, BORDER-LEAF69, BORDER-LEAF70, LEAF-1, LEAF-2, SVC-LEAF1, and several leaf flex nodes (LEAF-1-fex-101, LEAF-1-fex-102, LEAF-1-fex-103, LEAF-3-fex-102). A cloud icon labeled 'Ext' is connected to the spine nodes. A sidebar on the left contains 'Actions' such as 'Refresh topology', 'Save layout', and 'Add switches'. On the right, a details panel for 'SVC-LEAF1' (IP: 37.0.1.23, Model: N9K-C93180YC-EX) shows a 'Summary' with status 'ok', 'Health' at 98%, and 'Redirected Flows' for policy1 from VLAN_10 to VLAN_11.

Step 2 Click **Details** in the **Service Flows** window to display attachment details.

The screenshot shows the 'Service Flows' window for SVC-LEAF1. It displays a table with one row of service flow data. The table columns are Node, Policy, Details, Peering, VRF, Src Network, Dest Network, Next Hop, and Rev Next Hop. The row contains the following values: 1, ASA1, policy1, Details, peering1, VRF_51000, VLAN_10, VLAN_11, 201.201.200.200, and 200.200.200.200. A mouse cursor is hovering over the 'Details' cell.

Node	Policy	Details	Peering	VRF	Src Network	Dest Network	Next Hop	Rev Next Hop
1	ASA1	policy1	peering1	VRF_51000	VLAN_10	VLAN_11	201.201.200.200	200.200.200.200



8. Visualize Redirected Flows to Destination in the Topology window

Procedure

- Step 1** Click **Topology** and click on leafs to visualize the redirected flows to destination.

8. Visualize Redirected Flows to Destination in the Topology window

The screenshot displays the Cisco Data Center Network Manager interface. The main window shows a network topology with nodes including SPINE-1, SPINE-2, BORDER-LEAF69, BORDER-LEAF70, LEAF-1, LEAF-3, SVC-LEAF1, LEAF-15, and LEAF-16. A detailed view of SVC-LEAF1 is shown on the right, including its IP address (37.0.1.23), model (N9K-C93180YC-EX), and various metrics.

SCOPE: SVC-LEAF1

37.0.1.23
N9K-C93180YC-EX

Summary

Status: ✔ ok
 Serial number: FDO223218JS
 Version: 9.3(1)
 CPU: 17
 Memory: 25

Health

98%
 Modules: 100.00% w=0.2
 Switch ports: 93.55% w=0.2
 Alarms: 100.00% w=0.6

Redirected Flows

policy1: VLAN_10 → VLAN_11
 Show more flows

Tags

+
 System Tags
 VTEP

Step 2 Select **Redirected Flows** from the drop-down list.

The screenshot displays the Cisco Data Center Network Manager interface. The main window shows a network topology with nodes including SPINE-1, SPINE-2, BORDER-LEAF69, BORDER-LEAF70, LEAF-1, LEAF-3, SVC-LEAF1, LEAF-15, and LEAF-16. A search dropdown menu is open on the left, and a 'Show' panel is visible on the right.

SCOPE: SITE_A


Quick Search
 Host name (vCenter)
 Pod name (Container)
 Host IP
 Host MAC
 Multicast Group
 Redirected flows
 VXLAN ID (VNI)
 VLAN
 VXLAN OAM

Show

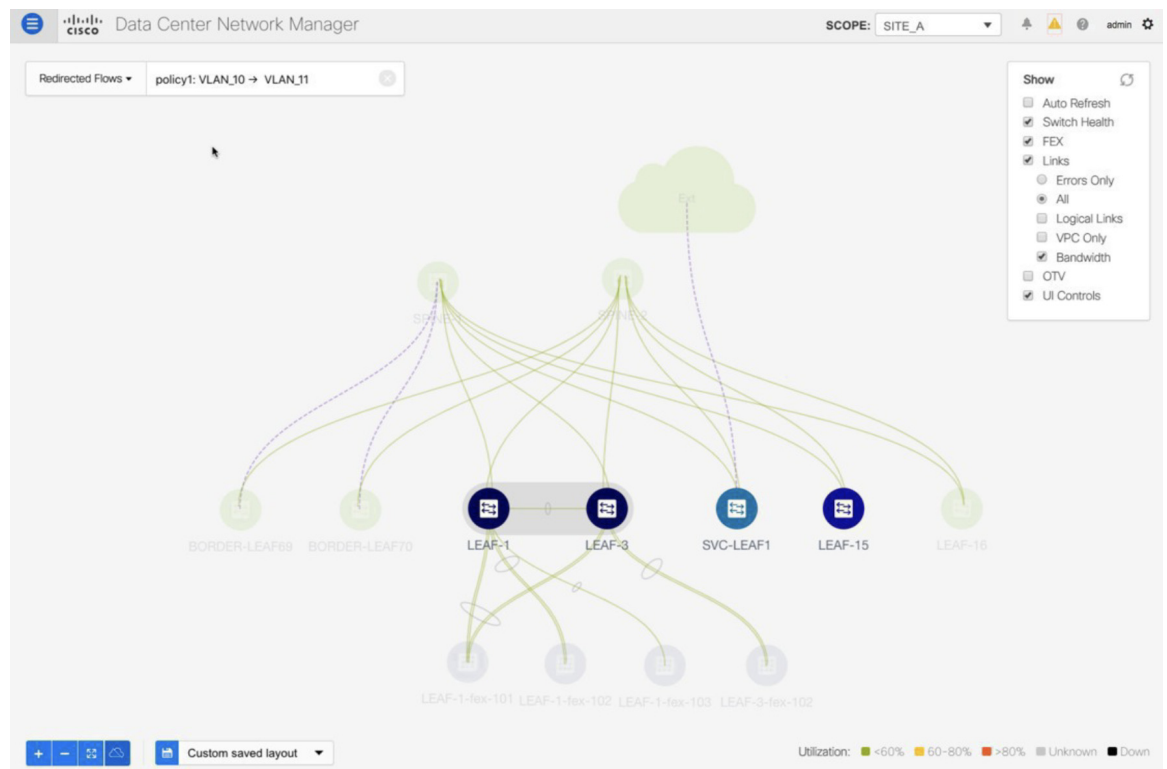
Auto Refresh
 Switch Health
 FEK
 Links
 Errors Only
 All
 Logical Links
 VPC Only
 Bandwidth
 OTV
 UI Controls

Utilization: ■ <60% ■ 60-80% ■ >80% ■ Unknown ■ Down

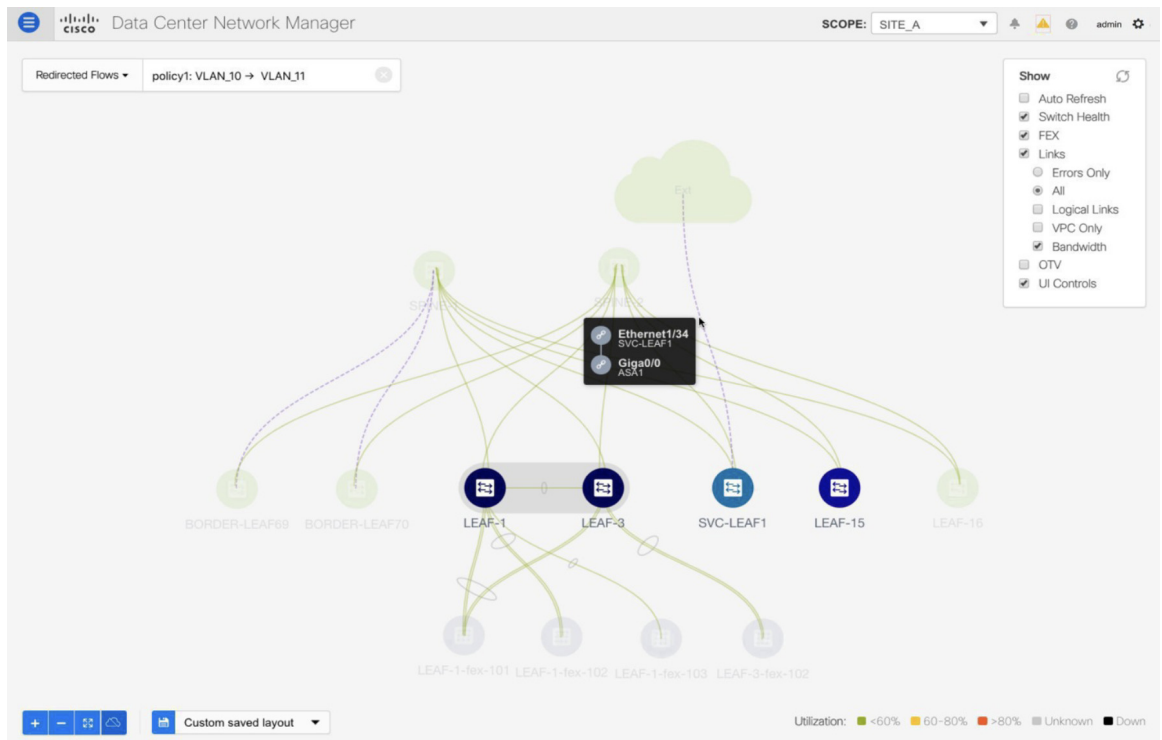
Step 3 Select a policy from the drop-down list or initiate a search by entering a policy name, source network and destination network in the search field. The search field is autopopulated based on your input.



The switches, on which the source and destination network have been attached and the flows have been redirected, are highlighted.



Step 4 The service node is shown as connected by a dotted line to the leaf switch on the topology window. Hover over the dotted line to get more information about the interface.



The traffic from **Source** traverses to the service leaf where the firewall is configured.

Based on firewall rules, traffic is allowed to reach the destination, Leaf 15.

Use Case: Inter-tenant Firewall with eBGP Peering

Refer the figure given below for topology details.

In this topology, es-leaf1 and es-leaf2 are vPC border leaf switches.

Now, let us see how to perform service redirection in DCNM.

Select **Control > Fabrics > Services**.

This use-case consists of the following steps:



Note

- As some steps are similar to the steps given in the Intra-tenant Firewall deployment use- case, reference links have been provided to the steps in that use-case.
- Service policies are not applicable on Inter-tenant firewall deployments.

1. Create Service Node

Procedure

Step 1 From the **Scope** drop-down list, select **Site_A**.

Service Nodes cannot be defined for selected fabric scope. Select a valid fabric scope.
In a valid fabric scope, you can define

Service Node
Onboard a service device such as a *firewall* or *load balancer*. Specify service node name, type, and interface attachment details

Route Peering
Specify deployment type, network parameters, peering protocol, and service IP

Service Policy
Specify traffic redirection rules to/from the service node

Step 2 Click the **Add** icon in the **Service Nodes** window.

Selected fabric scope has no service node. Add a service node to continue.
In selected fabric scope, you can define

Service Node
Onboard a service device such as a *firewall* or *load balancer*. Specify service node name, type, and interface attachment details

Route Peering
Specify deployment type, network parameters, peering protocol, and service IP

Service Policy
Specify traffic redirection rules to/from the service node

1. Create Service Node

Step 3 Enter the node name and specify **Firewall** in the **Type** dropdown box. The **Service Node Name** has to be unique.

The screenshot shows the 'New Service Nodes' form in Cisco Data Center Network Manager. The 'Service Node Name' field is filled with 'ASA1' and the 'Type' dropdown menu is set to 'Firewall'. A progress indicator on the left shows '1 Create Service Node' as the current step.

Step 4 From the **Form Factor** drop-down list, select **Virtual**.

The screenshot shows the 'New Service Nodes' form with the 'Form Factor' dropdown menu open. The 'Virtual' option is selected and highlighted in blue. The 'Service Node Name' is 'ASA1' and the 'Type' is 'Firewall'. The progress indicator shows '2 Create Route Peering' as the current step.

Step 5 In the **Switch Attachment** section, from the **External Fabric** drop-down list, select the external fabric in which the service node (for example, ASA firewall) is located. Note that the service nodes need to belong to the external fabric. This is a prerequisite before creating a service node.

Step 6 Enter the interface name of the service node that will be connected to the service leaf.

The screenshot shows the 'Service Node Interface' field with the text 'Giga0/0' entered.

Step 7 Select the attached switch that is the service leaf, and the respective interface on the service leaf.

Step 8 Select the **service_link_trunk** template. DCNM supports trunk, port channel, and vPC link templates. The available link templates in the **Link Template** drop-down list are filtered based on the selected **Attached Switch Interface** type.

The screenshot shows the 'Link Template' dropdown menu with 'service_link_trunk' selected.

Step 9 Specify the **General Parameters** and **Advanced** parameters, if required. Some parameters are pre-filled with the default values.

Step 10 Click **Next** to save the created service node.

Note For more sample screenshots, refer [1. Create Service Node, on page 2](#) in the Intra-tenant firewall with policy-based routing use case.

2. Create Route Peering

Let us now configure the peering between a service leaf and a service node.

Procedure

- Step 1** Enter the peering name and select **Inter-Tenant Firewall** from the **Deployment** drop-down list. From the **Peering Option** drop-down list, select **eBGP Dynamic Peering**.
- Step 2** Under **Inside Network**, from the **VRF** drop-down list, select a VRF that already exists and select **Inside Network** under **Network Type**.

Enter the name of the **Service Network** and specify the **Vlan ID**. You can also click **Propose** to allow DCNM to fetch the next available VLAN ID from the specified service network VLAN ID range in the fabric settings. The default **Service Network Template** is **Service_Network_Universal**.

Under the **General Parameters** tab, specify the gateway address for the service network. Specify the **Next Hop IP Address**. This next hop address has to be within the 'inside service network' subnet. Under the **Advanced** tab, the default **Routing Tag** value is 12345.

- Step 3** The default Peering Template for eBGP dynamic peering is **service_ebgp_route**.

Peering Template

service_ebgp_route

Under the **General Parameters** tab, specify the **Neighbor IPv4** address, **Loopback IP** address, and the **vPC Peer's Loopback IP** address. The border switches are a vPC pair.

The screenshot shows the configuration interface with two tabs: "General Parameters" (active) and "Advanced". Under "General Parameters", there are three input fields:

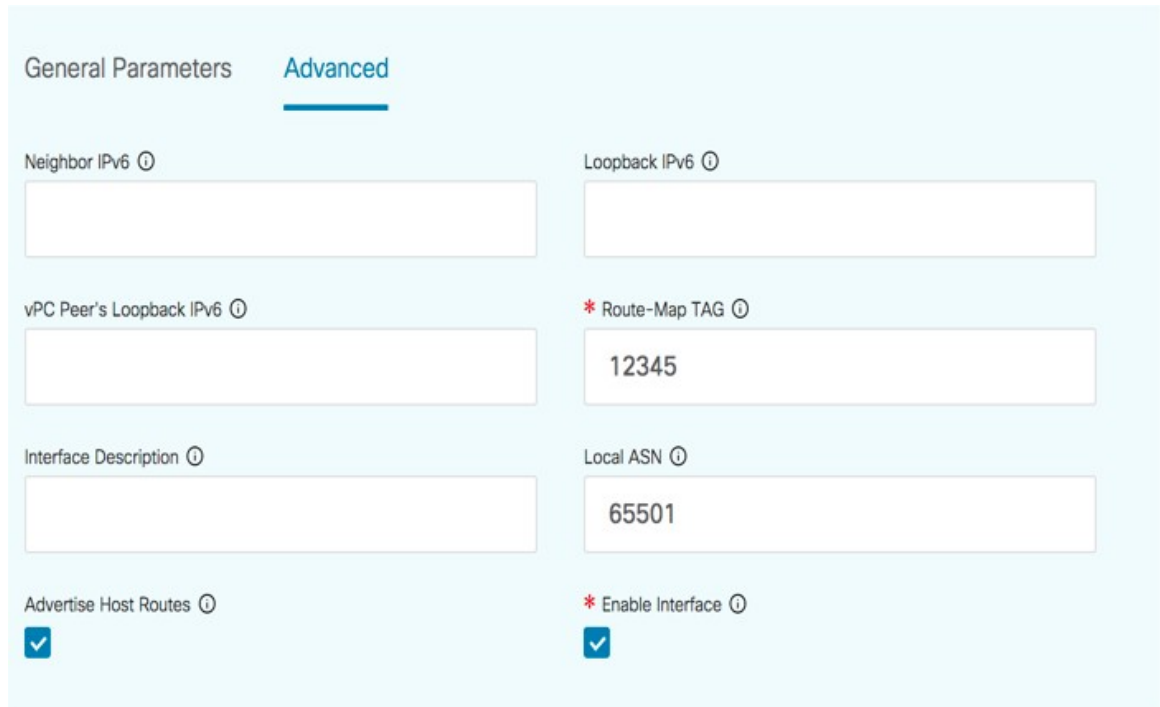
- * Neighbor IPv4**: 192.168.32.254
- * Loopback IP**: 60.1.1.60
- vPC Peer's Loopback IP**: 60.1.1.61

- Step 4** Under the **Advanced** tab, specify the **Local ASN** and select the **Advertise Host Routes** checkbox. This local ASN value is used to override the system ASN on the switch and is required to avoid routing loops.

If the **Advertise Host Routes** checkbox is selected, the /32 and /128 routes are advertised. If this checkbox is not selected, the prefix routes will be advertised.

By default, the **Enable Interface** checkbox is selected.

2. Create Route Peering



General Parameters **Advanced**

Neighbor IPv6 ⓘ

Loopback IPv6 ⓘ

vPC Peer's Loopback IPv6 ⓘ

* Route-Map TAG ⓘ

Interface Description ⓘ

Local ASN ⓘ

Advertise Host Routes ⓘ

* Enable Interface ⓘ

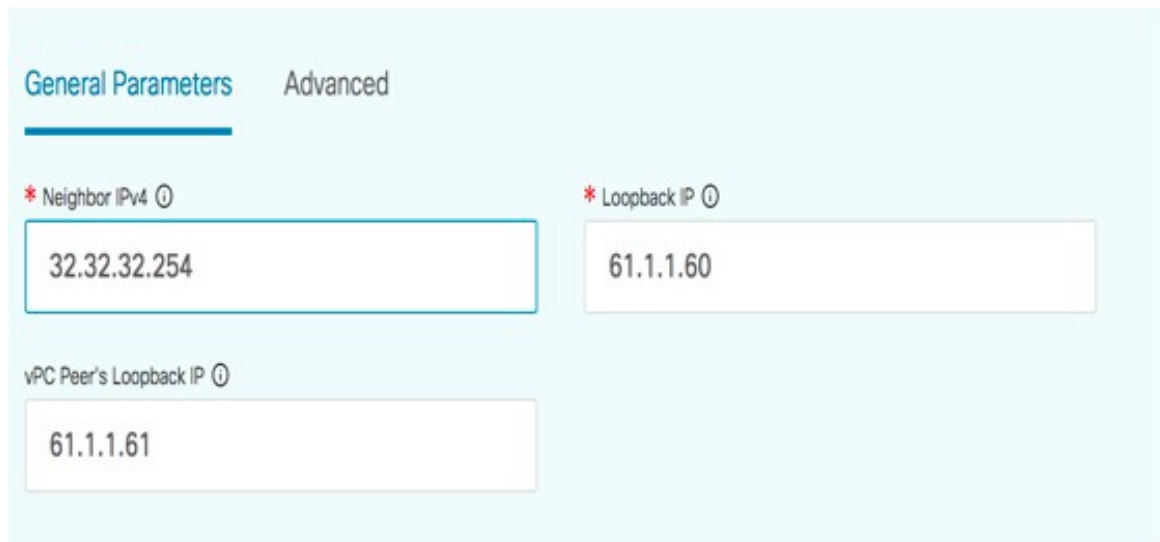
Step 5 Specify the required parameters under **Outside Network** and specify the **Next Hop IP Address for Reverse Traffic**. This next hop address for reverse traffic needs to be within the 'outside service network' subnet.

Step 6 The default Peering Template for eBGP dynamic peering is `service_ebgp_route`.

Peering Template

service_ebgp_route ▼

Under the **General Parameters** tab, **Neighbor IPv4** address, **Loopback IP** address, and the **vPC Peer's Loopback IP** address. The leaf switches are a vPC pair.



General Parameters **Advanced**

* Neighbor IPv4 ⓘ

* Loopback IP ⓘ

vPC Peer's Loopback IP ⓘ

Step 7 Under the **Advanced** tab, specify the **Local ASN** and select the **Advertise Host Routes** checkbox. This local ASN value is used to override the system ASN on the switch and is required to avoid routing loops.

If the **Advertise Host Routes** checkbox is selected, the /32 and /128 routes are advertised. If this checkbox is not selected, the prefix routes will be advertised.

By default, the **Enable Interface** checkbox is selected.

The screenshot shows the 'Advanced' configuration tab for route peering. The 'Advertise Host Routes' checkbox is checked. The 'Local ASN' is set to 65501. The 'Route-Map TAG' is set to 12345. Other fields like 'Neighbor IPv6', 'Loopback IPv6', 'vPC Peer's Loopback IPv6', 'Interface Description', and 'Enable Interface' are also visible, with 'Enable Interface' also checked.

Step 8 Click **Next** to save the created route peering.

3. Deploy Route Peering

Refer [4. Deploy Route Peering, on page 9](#) of the Intra-Tenant Firewall deployment use-case. Note that **InterTenantFW** is displayed under **Deployment**.

The BGP configuration on the vPC border leaf for this use-case is given below.

```
router bgp 12345
router-id 10.2.0.1
address-family l2vpn evpn
  advertise-pip
neighbor 10.2.0.4
  remote-as 12345
  update-source loopback0
address-family l2vpn evpn
  send-community
  send-community extended
vrf myvrf_50001
address-family ipv4 unicast
  advertise l2vpn evpn
  redistribute direct route-map fabric-rmap-redist-subnet
```

3. Deploy Route Peering

```

    maximum-paths ibgp 2
    address-family ipv6 unicast
    advertise l2vpn evpn
    redistribute direct route-map fabric-rmap-redis-subnet
    maximum-paths ibgp 2
    neighbor 192.168.32.254
    remote-as 9876
    local-as 65501 no-prepend replace-as // Note: This configuration corresponds to the Local
    ASN template parameter value of the service_ebgp_route template of the inside network with
    VRF myvrf_50001. The no-prepend replace-as keyword is generated along with the local-as
    command.
    update-source loopback2
    ebgp-multihop 5
    address-family ipv4 unicast
    send-community
    send-community extended
    route-map extcon-rmap-filter-allow-host out
vrf myvrf_50002
    address-family ipv4 unicast
    advertise l2vpn evpn
    redistribute direct route-map fabric-rmap-redis-subnet
    maximum-paths ibgp 2
    address-family ipv6 unicast
    advertise l2vpn evpn
    redistribute direct route-map fabric-rmap-redis-subnet
    maximum-paths ibgp 2
    neighbor 32.32.32.254
    remote-as 9876
    local-as 65502 no-prepend replace-as // Note: This configuration corresponds to the Local
    ASN template parameter value of the service_ebgp_route template of the outside network
    with VRF myvrf_50002. The no-prepend replace-as keyword is generated along with the local-as
    command.
    update-source loopback3
    ebgp-multihop 5
    address-family ipv4 unicast
    send-community
    send-community extended
    route-map extcon-rmap-filter-allow-host out

```

The loopback interface configuration on the vPC switch es-leaf1 for this use-case is given below. The loopback interfaces in the configuration correspond to the 'Loopback IP' parameter of the **service_ebgp_route** template. Two loopback interfaces are created automatically on each vPC switch for two separate VRF instances using the **Loopback IP** parameter values that are specified in the **service_ebgp_route** template.

```

interface loopback2
vrf member myvrf_50001
ip address 60.1.1.60/32 tag 12345
interface loopback3
vrf member myvrf_50002
ip address 61.1.1.60/32 tag 12345

```

The loopback interface config on vPC peer switch es-leaf2:

```

interface loopback2
vrf member myvrf_50001
ip address 60.1.1.61/32 tag 12345
interface loopback3
vrf member myvrf_50002
ip address 61.1.1.61/32 tag 12345

```

Use Case: One-arm Load Balancer

Refer the figure given below for topology details.

1. Create Service Node

In this topology, es-leaf1 and es-leaf2 are vPC leafs.

Now, let us see how to perform service redirection in DCNM.

Select **Control > Fabrics > Services**.

This use-case consists of the following steps:



Note As some steps are similar to the steps given in the Intra-tenant Firewall deployment usecase, reference links have been provided to the steps in that use-case.

1. Create Service Node

Procedure

Step 1 From the **Scope** drop-down list, select **Site_A**.

Step 2 Click the **Add** icon in the **Service Nodes** window.

Step 3 Enter the node name and specify **Load Balancer** in the **Type** dropdown box. The **Service Node Name** has to be unique.

Step 4 From the **Form Factor** drop-down list, select **Virtual**.

* Form Factor

Virtual ^

Physical

Virtual ✓

Step 5 In the **Switch Attachment** section, from the **External Fabric** drop-down list, select the external fabric in which the service node (for example, ASA firewall) is located. Note that the service nodes need to belong to the external fabric. This is a prerequisite before creating a service node.

Step 6 Enter the interface name of the service node that will be connected to the service leaf.

* Service Node Interface ⓘ

Giga0/0

Step 7 Select the attached switch that is the service leaf, and the respective interface on the service leaf.

Step 8 Select the **service_link_trunk** template. DCNM supports trunk, port channel, and vPC link templates. The available link templates in the **Link Template** drop-down list are filtered based on the selected **Attached Switch Interface** type.

Link Template

service_link_trunk v

Step 9 Specify the **General Parameters** and **Advanced** parameters, if required. Some parameters are pre-filled with the default values.

General Parameters Advanced

MTU ⓘ <div style="border: 1px solid #ccc; padding: 2px;">jumbo v</div>	SPEED ⓘ <div style="border: 1px solid #ccc; padding: 2px;">Auto v</div>
Trunk Allowed Vlans ⓘ <div style="border: 1px solid #ccc; padding: 2px;">none</div>	Enable BPDU Guard ⓘ <div style="border: 1px solid #ccc; padding: 2px;">no v</div>
Enable Port Type Fast ⓘ <input checked="" type="checkbox"/>	Enable Interface ⓘ <input checked="" type="checkbox"/>

Next

Step 10 Click **Next** to save the created service node.

Note For more sample screenshots, refer [1. Create Service Node, on page 2](#) in the Intra-tenant firewall with policy-based routing use case.

2. Create Route Peering

Let us now configure peering between a service leaf and a service node. In this use-case, we configure static route peering.

Procedure

- Step 1** Enter the peering name and select **One-Arm Mode** from the **Deployment** drop-down list. Also, from the **Peering Option** dropdown list, select **Static Peering**.
- Step 2** Under **First Arm**, specify the required values. From the **VRF** dropdown list, select a VRF that already exists and select **First Arm** under **Network Type**.
- Step 3** Enter the name of the **Service Network** and specify the **Vlan ID**. You can also click Propose to allow DCNM to fetch the next available VLAN ID from the specified service network VLAN ID range in the fabric settings. The default **Service Network Template** is **Service_Network_Universal**.

Under the **General Parameters** tab, specify the gateway address for the service network. Specify the **Next Hop IP Address**. This next hop address has to be within the first arm's subnet. Under the **Advanced** tab, the default **Routing Tag** value is 12345.

- Step 4** The default **Peering Template** is **service_static_route**. Add routes, as required, in the **Static Routes** field.

Peering Template

service_static_route

service_static_route ✓

Static Routes ⚠ ⓘ

55.55.55.55, 192.168.50.254

- Step 5** Specify the **Next Hop IP Address** for Reverse Traffic.

- Step 6** Click **Next** to save the created route peering.

× Data Center Network Manager SCOPE: SITE_A admin

Service Nodes

LB1 VIRTUAL 1 2

LOAD BALANCER Route Peering Service Policy

Service Policy Route Peering Attach Detach Preview Deploy History

Peering Name	Deployment	Peering Option	Status	VRF	Service Network One		Service Network Two		Next Hop IP	Reverse Ne	Action
					Network Name	Gateway IP	VRF	Network Name			
RP-1	OneArmADC	StaticPeering	In-Sync	MyVRF_50001	net_lb	192.168.50.1/24			192.168.50		

3. Create Service Policy

Refer [3. Create Service Policy, on page 6](#) in the Intra-Tenant Firewall deployment use-case.

4. Deploy Route Peering

Refer [4. Deploy Route Peering, on page 9](#) in the Intra-tenant Firewall deployment use-case. Note that **OneArmADC** is displayed under **Deployment**.

5. Deploy Service Policy

Refer [5. Deploy Service Policy, on page 11](#) in the Intra-tenant Firewall deployment use-case. However, as there are two servers in this load balancer use-case, two service policies have to be defined with each server network.

Policy Name	Route Peering	Status	Source VRF	Source Network	Destination VRF	Destination Network	Next Hop IP	Reverse Next Hop IP	Stats	Action
SP-1	RP-1	In-Sync	MyVRF_50...	ClientNet	MyVRF_50001	ServerNet		192.168.50.254	↕	🔧
SP-2	RP-1	In-Sync	MyVRF_50...	ClientNet	MyVRF_50001	ServerNet2		192.168.50.254	↕	🔧

6. View Stats

Refer [6. View Stats, on page 13](#) in the Intra-Tenant Firewall deployment use-case.

7. View Traffic Flow in Fabric Builder

Refer [7. View Traffic Flow in Fabric Builder, on page 14](#) in the Intra-Tenant Firewall deployment use-case.

8. Visualize Redirected Flows to Destination in the Topology window

Refer [8. Visualize Redirected Flows to Destination in the Topology window, on page 17](#) in the Intra-Tenant Firewall deployment use-case.

The VRF configuration on the service leaf is as given below.

```
interface Vlan2000
  vrf member myvrf_50001
  ip policy route-map rm_myvrf_50001

interface Vlan2306
  vrf member myvrf_50001
  vrf context myvrf_50001
  vni 50001
  ip route 55.55.55.55/32 192.168.50.254 // Note: This is the static route
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
  address-family ipv6 unicast
```

8. Visualize Redirected Flows to Destination in the Topology window

```
    route-target both auto
    route-target both auto evpn
router bgp 12345
vrf myvrf_50001
  address-family ipv4 unicast
    advertise l2vpn evpn
    redistribute direct route-map fabric-rmap-redirect-subnet
    redistribute static route-map fabric-rmap-redirect-static
    maximum-paths ibgp 2
  address-family ipv6 unicast
    advertise l2vpn evpn
    redistribute direct route-map fabric-rmap-redirect-subnet
    redistribute static route-map fabric-rmap-redirect-static
    maximum-paths ibgp 2
```