

Connecting Multiple ACI Fabrics and Working with Orchestration, Release 4.3.1

Table of Contents

New and changed information . . . 1
Grouping fabrics and clusters . . . 2
  Grouping Nexus Dashboard clusters . . . 2
  Grouping fabrics . . . 2
Connecting multiple ACI fabrics through the Orchestration page . . . 4
  Guidelines and limitations . . . 4
  Prerequisites . . . 4
  Navigate to the Orchestration page . . . 4
  Understanding the Orchestration page . . . 4
    Overview . . . 5
    Tenants . . . 5
    Tenant Templates . . . 6
    Fabric Templates . . . 7
  Connect multiple ACI fabrics . . . 7
    1. Fabrics . . . 7
    2. Fabric Details . . . 9
    Add policy . . . 12
    3. Switch and Interface details . . . 16
    4. Summary . . . 19
  How TEP IP addresses are used to provide connectivity between fabrics . . . 20
Copyright . . . 21

New and changed information

The following table provides an overview of the significant changes up to this current release. The table does not provide an exhaustive list of all changes or of the new features up to this release.

Release Version	Feature	Description
Nexus Dashboard 4.3.1	Cisco Nexus 9000 Series name change	Cisco Nexus 9000 Series is now the Cisco N9000 Series.

Grouping fabrics and clusters

There are several ways to group fabrics and clusters together in Nexus Dashboard:

- Grouping Nexus Dashboard clusters
- Grouping fabrics

Grouping Nexus Dashboard clusters

Multi-cluster connectivity: You can establish connectivity between multiple Nexus Dashboard and APIC clusters. For more information, see Connecting Clusters.

Grouping fabrics

The method that you use to group fabrics together differs depending on the type of fabric:

- **NX-OS fabrics**: You can group fabrics together using a *fabric group*, which is a collection of fabrics grouped for visualization or management. The available fabric group types include:
  - **VXLAN**: A fabric group that is used to interconnect multiple data center VXLAN EVPN fabrics and/or ACI fabrics using border gateways. It can also be used to interconnect Data Center and Campus VXLAN EVPN fabrics. This enables shared overlay deployments for security and segmentation using VXLAN EVPN multisite. The IPN/ISN interconnect network fabric may also be optionally part of this group.

    These types of VXLAN fabric groups are available:

    - **VXLAN**: This subtype of fabric group can contain multiple NX-OS-based Data Center VXLAN EVPN fabrics and/or IOS-XE-based Campus VXLAN EVPN fabrics (which need NX-OS-based border gateways). The DCI network comprising of Multi-Site and External Connectivity or BGP routed fabrics can also be part of this group.
    - **VXLAN-ACI**: A group that includes a Cisco ACI fabric connected to one or more Cisco NX-OS fabrics using the Nexus One architecture for unified management and policy deployment. For more information, see Understanding the Nexus One architecture.
  - **Classic**: A Classic fabric group can contain Enhanced Classic LAN, Legacy Classic-LAN, Any VXLAN EVPN and External and Inter-fabric connectivity fabrics. This group allows for a combined visualization at a topology level. No group level deployments are available in this fabric group.
  - **IP Fabric for Media**: A logical group of LAN or IP Fabric for Media (IPFM) fabrics. This type of fabric group can contain individual IPFM/IPFM-Classic/Classic-LAN fabrics. This type of group allows shared host and flow definitions.
- **Cisco ACI fabric management models**: You can manage ACI fabrics in Nexus Dashboard using one of two methods:
  - Using Cisco ACI in Nexus Dashboard Orchestration
    - Deploys tenants, networks, and policy configurations across multiple ACI sites.
    - Provides a central point for multisite policy consistency and scale.

For more information, see Connecting Multiple ACI Fabrics and Working with Orchestration.

- Using Cisco ACI in a VXLAN-ACI fabric group
  - Enables management and interoperability between Cisco ACI and Cisco NX-OS VXLAN EVPN fabrics.
  - Automates the infrastructure configuration for inter-fabric underlay and overlay connectivity between ACI and NX-OS fabrics.
  - Centralizes the deployment of security and segmentation policies across Cisco ACI and Cisco NX-OS environments.

    For more information, see Understanding the Nexus One architecture

Constraint: You cannot enable Orchestration on a Cisco ACI fabric that is a member of a VXLAN-ACI fabric group. Similarly, you cannot add a Cisco ACI fabric to a VXLAN-ACI fabric group if Orchestration is enabled. When adding fabrics to a VXLAN-ACI fabric group, any ACI fabrics with Orchestration enabled will not be available to choose.

- **Multi-cluster fabric groups**: You can create a multi-cluster fabric group that spans multiple Nexus Dashboard clusters to manage VXLAN fabrics across different clusters in a multi-cluster fabric group. For more information, see Create multi-cluster fabric groups.

Connecting multiple ACI fabrics through the Orchestration page

Use the **Orchestration** feature through Nexus Dashboard to:

- Connect multiple ACI fabrics together
- Consolidate and deploy tenants, along with network and policy configurations, across multiple ACI fabrics



The **Orchestration** feature is available only for ACI fabrics.

Guidelines and limitations

You cannot enable Orchestration on a Cisco ACI fabric that is a member of a VXLAN-ACI fabric group. Similarly, you cannot add a Cisco ACI fabric to a VXLAN-ACI fabric group if Orchestration is enabled. When adding fabrics to a VXLAN-ACI fabric group, any ACI fabrics with Orchestration enabled will not be available to choose.

Prerequisites

Before you can use the **Orchestration** feature, you must first make these configurations:

1. Add the necessary ACI fabrics to your Nexus Dashboard cluster.
2. In the **License tier** field, set the licensing tier for the ACI fabric to either **Advantage** or **Premier**.
3. In the **Enabled features** field, put a check in the check box next to **Orchestration** for the ACI fabric.

See Connecting Clusters and Editing ACI Fabric Settings for more information.

Navigate to the Orchestration page

To navigate to the **Orchestration** page, click **Manage > Orchestration**. Any issues with fabric connectivity will be displayed as errors at the top of the **Orchestration** page.

Understanding the Orchestration page

These tabs are available in the **Orchestration** page:

- Overview
- Tenants
- Tenant Templates
- Fabric Templates

Overview

The **Overview** page displays this information:

- **Summary**:
  - **Fabrics with Associated Templates**: Displays the number of fabrics with associated templates. Click in the tile to bring up a slide-in pane with additional information on the fabrics with associated templates.
  - **Tenants**: Displays the number of tenants.
    - Click in the tile to navigate directly to the Tenants tab in the **Orchestration** page.
    - Click **View Statistics** to navigate directly to the **Tenants** area in the **Overview** page.
  - **In-Sync Templates**: Displays the number of in-sync templates.
    - Click in the tile to bring up a slide-in pane with additional information on the in-sync templates.
    - Click **View Statistics** to navigate directly to the **Templates** area in the **Overview** page.
  - **Out-of-Sync Templates**: Displays the number of out-of-sync templates.
    - Click in the tile to bring up a slide-in pane with additional information on the out-of-sync templates.
    - Click **View Statistics** to navigate directly to the **Templates** area in the **Overview** page.
  - **Unassociated Templates**: Displays the number of unassociated templates.
- **Tenants**: Displays the total number of tenants and the corresponding number of policies deployed to each tenant.
  - Click **View Tenants** to navigate directly to the Tenants tab in the **Orchestration** page.
  - Change the information displayed in the **Tenants** area by clicking the **By Policy**, **By Templates**, or **By Fabrics** options.
- **Templates**: Displays the total number of templates configured. Change the information displayed in the **Templates** area by clicking the **By Type**, **By Status**, or **By State** options.

Tenants

Tenants presents a list of tenants in tabular form with operational information. You can filter or sort the table using the attributes, such as:

- **Name**
- **Description**
- **Assigned To Fabrics**
- **Assigned To Users**
- **Assigned to Templates**

You can perform these actions on tenants:

- For already-created tenants, click the ellipsis (...) in the last column of the table to edit a tenant.
- Click **Import Tenants** to import existing tenants.

- Click **Create Tenant** to create a new tenant.

For more information on tenants, see Tenants and Tenant Policies Templates for ACI Fabrics.

Tenant Templates

Tenant Templates allows you to configure templates. These template types are supported:

- **Applications:** This tab shows a table of schemas, you can filter or sort the table using the attributes like templates that are deployed, tenants, and policies associated with the schema. The ellipsis (...) on the last column of the table allow you to **Edit**, **Delete**, and **Clone** the schema. You can add a new schema using the **Add Schema** button. You can select an individual schema for an overview of the schema.

  For more information, see Schemas and Application Templates for ACI Fabrics.

- **L3Out:** This tab shows a tabular deployment status of the L3Out templates. You can filter or sort the table using the attributes like status, name, tenant, fabrics, and policies. The three dots on the last column of the table allow you to **Edit** or **Delete** the L3Out template. You can add a new L3Out template using the **Create L3Out Template** button. You can select each individual template to get a summary of the template. You can edit or deploy the template using the appropriate buttons under the summary. **Action** menu has options to perform fabric or template level actions.

  For more information, see External Connectivity (L3Out) for ACI Fabrics.

- **Monitoring Policies:** Allows you to create and edit monitoring policies templates.

  For more information, see Fabric Management Templates for ACI Fabrics.

- **Service Device:** This tab shows a tabular deployment status of the service device template. You can filter or sort the table using the attributes like status, name, tenant, fabrics, and policies. The ellipsis (...) on the last column of the table allow you to **Edit** or **Delete** the service device template. You can add a new Service Device template using the **Create Service Device Template** button. You can select each individual template to get a summary of the template properties along with each fabric deployed.

  For more information, see:

  - Inter-Fabric L3Out with PBR for ACI Fabrics
  - Inter-Fabric Transit Routing with PBR for ACI Fabrics
  - vzAny with PBR for ACI Fabrics

- **Tenant Policies:** This tab shows a tabular deployment status of the tenant policy template. You can filter or sort the table using the attributes like status, name, tenant, fabrics, and policies. The three dots on the last column of the table allow you to **Edit** or **Delete** the tenant policy template. You can edit or deploy the template using the appropriate buttons under the Template Summary. You can add a new tenant policy template using the **Create Tenant Policy Template** button. **Action** menu has options to perform fabric or template level actions.

  For more information, see Tenants and Tenant Policies Templates for ACI Fabrics.

Fabric Templates

Fabric Templates provides options to configure fabric policy templates:

- **Fabric Policies:** This tab shows a table of fabric policies. You can filter or sort the table using the attributes like status, name, fabrics, and policies. The three dots on the last column of the table allow you to **Edit** or **Delete** the fabric policy template. You can add a new fabric policy template using the **Create Fabric Policy Template** button. You can edit or deploy the template using the appropriate buttons under the Template Summary. **Action** menu has options to perform fabric or template level actions.
- **Fabric Resource Policies:** This tab shows a table of fabric resource policy templates. You can filter or sort the table using the attributes like status, name, fabrics, and policies. The three dots on the last column of the table allow you to **Edit** or **Delete** the fabric resource policy template. You can add a new fabric resource policy template using the **Create Fabric Resource Policy Template** button. You can edit or deploy the template using the appropriate buttons under the Template Summary. **Action** menu has options to perform fabric or template level actions.
- **Monitoring Access Policies**: This tab shows a table of monitoring access policy templates. You can filter or sort the table using the attributes like status, name, fabrics, and policies. The three dots on the last column of the table allow you to **Edit** or **Delete** the monitoring access policy template. You can add a new monitoring access policy template using the **Create Monitoring Policy Template** button. You can edit or deploy the template using the appropriate buttons under the Template Summary. **Action** menu has options to perform fabric or template level actions.

For more information, see Fabric Management Templates for ACI Fabrics.

Connect multiple ACI fabrics

To connect multiple ACI fabrics using the Orchestration feature:

1. Navigate to the **Orchestration** window.

   Click **Manage > Orchestration** to navigate to the **Orchestration** window. You can view, create, delete, and modify orchestration between fabrics in this window.

2. In the upper right corner, click **Inter-Fabric Connectivity**.

   The **Inter-Fabric Connectivity** window appears.

Navigate through the **Inter-Fabric Connectivity** wizard to connect multiple fabrics.

- 1. Fabrics
- 2. Fabric Details
- 3. Switch and Interface details
- 4. Summary

1. Fabrics

A fabric group is created automatically with the fabrics shown in the **Fabrics** window after you have completed these procedures. Up to 12 ACI fabrics can be interconnected.

1. Review the information provided in the **Fabrics** page and make modifications, if necessary.
   - The **Topology Preview** area shows the ACI fabrics that will be interconnected.
     - Click the down arrow to expand the section, if necessary.
     - Click each ACI fabric to display a popup with additional fabric information.
   - The **Fabrics** area provides a list of ACI fabrics that will be interconnected, shown in the table with these fields:
     - **Name**: Displays the name of the available fabrics.
     - **Fabric Type**: Displays the fabric type for each fabric listed.
     - **Switches**: Displays the number of switches in each fabric.
     - **ASN**: Displays the ASN for each fabric.
     - **Deployment Status**: Displays the deployment status for each fabric.
   - In the **General Settings** area, click the down arrow to view the default settings.

     The following table provides information on the default settings and additional options for each field.

Field	Description
BGP Peering Type	• **full-mesh**: Default setting. All border gateway switches in each fabric establishes peer connectivity with the remote fabrics' border gateway switches. In a full-mesh configuration, Nexus Dashboard uses the spine switches for ACI fabrics and border gateways for NX-OS fabrics. • **route-reflector**: The route-reflector option allows you to specify one or more control-plane nodes where each fabric establishes MP-BGP EVPN sessions. The use of route-reflector nodes avoids creating MP-BGP EVPN full mesh adjacencies between all the fabrics that are managed by the Nexus Dashboard. For ACI fabrics, the route-reflector option is effective only for fabrics that are part of the same BGP ASN.
Keep-Alive Interval (Seconds)	Applicable only for on-premises fabrics. Enter the keep-alive interval in seconds. We recommend keeping the default value of 60 seconds.
Hold Interval (Seconds)	Applicable only for on-premises fabrics. Enter the hold interval in seconds. We recommend keeping the default value of 180 seconds.
Stale Interval (Seconds)	Applicable only for on-premises fabrics. Enter stale interval in seconds. We recommend keeping the default value of 300 seconds.
Graceful Restart	Choose whether you want to enable the **Graceful Restart** option. Default setting is enabled.

Field	Description
Maximum AS Limit	Provide the Maximum AS Limit value. We recommend keeping the default value of 0.
BGP TTL Between Peers	Provide the BGP TTL Between Peers value. We recommend keeping the default value of 16.

2. Click **Next**.

   You advance to 2. Fabric Details.

2. Fabric Details

In this window, enter the fabric-level details for each of the fabrics that will be interconnected.

1. In the **Example View** area, click the down arrow to see an example view of the inter-fabric network.
2. In the **Fabric Settings** area, enter the required fabric-specific configuration information, if necessary.
   a. Determine if you want to refresh the configuration for any fabric listed in the **Fabric Settings** area.

      Refresh the configuration for a fabric if a change in that fabric needs to be brought in to Nexus Dashboard, such as a spine switch that has been added or deleted in the fabric, a change to an IP address, and so on.

      Click the refresh icon for any fabric to refresh the configuration information for that fabric. A popup window appears, telling you that refreshing this fabric will result in all fabric-related connectivity information being re-imported from the fabric, and that you will lose any changes that were not already pushed to the fabric.

      - If you do not want to refresh the fabric, click **Cancel**.
      - If you want to refresh the fabric, click the box if you also want to remove the configuration for decommissioned spine switches during the refresh, then click **Save**.

        You are returned to the **Fabric Settings** area.

   b. Determine if you are able to move forward with the inter-fabric connectivity configuration process with the existing information.

      - If the **Next** button is enabled, that means that all of the fabrics have the required information to allow you to move forward with the inter-fabric connectivity configuration.
        - If you are satisfied with the existing settings for each fabric, click **Next**.

          Go to 3. Switch and Interface details.

        - If you would like to review or change the configuration information for any fabric before proceeding, continue with these procedures.
      - If the **Next** button is disabled and you see a triangle with an exclamation mark next to a

fabric name in the **Fabric Settings** area, then there is configuration information missing for that fabric. Click **Edit** on the row for that fabric in this case.

The **Settings** window for that fabric appears. Go to Settings for *fabric-name*.

Settings for *fabric-name*

1. Enter the necessary information to configure the fabric-level settings.

   The following table provides information on the fields that are available when you edit fabric-level settings.

Field	Description
General Settings	
Enable Multi-Fabric	Choose whether you want to enable multi-fabric. This defines whether the overlay connectivity is established between this fabric and other fabrics.
BGP ASN	Specify the BGP Autonomous System Number for the fabric.
BGP Password	(Optional) Specify the BGP password for the fabric.
External Routed Domain	(Optional) From the **External Routed Domain** drop-down, choose the domain that you want to use. Choose an external router domain that you have created in the Cisco APIC GUI. For more information, see the *Cisco APIC Layer 3 Networking Configuration Guide* specific to your APIC release.
Overlay Multicast TEP Address	Enter the overlay multicast tunnel end-point (TEP) IP address. See How TEP IP addresses are used to provide connectivity between fabrics for more information on how the three tunnel end-point (TEP) IP addresses (Overlay Multicast TEP, Overlay Unicast TEP, and External TEP Pool) are used to provide connectivity between fabrics.
OSPF settings	
Area ID	Provide the OSPF area ID.
Area Type	Select the OSPF area type from the drop-down list. The OSPF area type can be one of the following: - nssa - regular
OSPF Policies	The following settings are required if you are using OSPF protocol for underlay connectivity between the fabric and the IPN. If you plan to use BGP instead, you can skip this step. BGP underlay configuration is done at the port level, as described in 3. Switch and Interface details. You can either click an existing policy (for example, msc-ospf-policy-default) to modify it or click **+Add policy** to add a new OSPF policy. If you are adding a new policy, go to Add policy for more information.
SR-MPLS Settings	

Field	Description
SR-MPLS Connectivity	Choose whether you want to enable SR-MPLS connectivity for the fabric. If the fabric is connected through an MPLS network, enable the **SR-MPLS Connectivity** knob and provide the Segment Routing global block (SRGB) range in the following fields.
SRGB Range Minimum	Field available if you enable the **SR-MPLS Connectivity** field. The Segment Routing Global Block (SRGB) is the range of label values that are reserved for Segment Routing (SR) in the Label Switching Database (LSD). These values are assigned as segment identifiers (SIDs) to SR-enabled nodes and have global significance throughout the domain.
Maximum	The default range is 16000-23999. If you enable MPLS connectivity for the fabric, you need to configure extra settings as described in Multi-fabric and SR-MPLS L3Out handoff for ACI fabrics.
Domain ID Base	Field available if you enable the **SR-MPLS Connectivity** field. The Domain ID Base enables the BGP Domain-Path feature. For more information, see the *Cisco APIC Layer 3 Networking Configuration Guide*. If you choose to provide a value in this field to enable the Domain-Path feature, ensure that you use a unique value for each SR-MPLS fabric in your Multi-Fabric domain, which will be specific to this ACI fabric.
Pod settings The following fields are available for each pod shown in the **Pod settings** area.	
Overlay Unicast TEP	Add the overlay unicast TEP for the Pod. This IP address is deployed on all spine switches that are part of the same Pod and is used for sourcing and receiving VXLAN encapsulated traffic for Layer2 and Layer3 unicast communication. See How TEP IP addresses are used to provide connectivity between fabrics for more information on how the three tunnel end-point (TEP) IP addresses (Overlay Multicast TEP, Overlay Unicast TEP, and External TEP Pool) are used to provide connectivity between fabrics.

Field	Description
External TEP Pools	Lists the configured external TEP pools. The external routable TEP pools are used to assign a set of IP addresses that are routable across the IPN to APIC nodes, spine switches, and border leaf nodes. This is required to enable Multi-Fabric architecture. See How TEP IP addresses are used to provide connectivity between fabrics for more information on how the three tunnel end-point (TEP) IP addresses (Overlay Multicast TEP, Overlay Unicast TEP, and External TEP Pool) are used to provide connectivity between fabrics. External TEP pools previously assigned to the fabric on APIC are automatically inherited by Nexus Dashboard and displayed in the GUI when the fabric is added to the Multi-Fabric domain. - Click the pencil icon on any already-configured external TEP pool to edit that entry. - Click the trashcan icon on any already-configured external TEP pool to delete that entry. - Click **+ Add TEP Pool** to add an external routable TEP pool. The **Add TEP Pool** window appears. In the **Add TEP Pool** window, specify the External TEP Pool that you want to configure for that fabric and the Reserved Address Count. For the TEP pool, provide the subnet and the subnet mask, for example 192.168.111.0/24. Ensure that the TEP pool that you are adding does not overlap with any other TEP pools or fabric addresses. You can configure multiple disjointed TEP pools, so you do not have to specify a large TEP pool from the beginning.
SR-MPLS L3Outs	Displayed if you enabled the **SR-MPLS Connectivity** option above. Lists the configured SR-MPLS L3Outs. - Click the pencil icon on any already-configured SR-MPLS L3Out to edit that entry. - Click the trashcan icon on any already-configured SR-MPLS L3Out to delete that entry. - Click **+ Add SR-MPLS L3Out** to add an SR-MPLS L3Out. The **Add SR-MPLS L3Out** window appears. Go to Add SR-MPLS L3Out.

2. Click **Next** if you have completed the fabric-level settings and the **Next** option is enabled.

   Go to 3. Switch and Interface details.

Add policy

In the **Add/Update Policy** window, specify the following:

1. In the **Policy Name** field, enter the policy name.

2. In the **Network Type** field, choose either broadcast, point-to-point, or unspecified.

   The default is broadcast.

3. In the **Priority** field, enter the priority number.

   The default is 1.

4. In the **Cost of Interface** field, enter the cost of interface.

   The default is 0.

5. From the **Interface Controls** drop-down list, choose one of the following:
   - **advertise-subnet**
   - **bfd**
   - **mtu-ignore**
   - **passive-participation**
6. In the **Hello Interval (Seconds)** field, enter the hello interval in seconds.

   The default is 10.

7. In the **Dead Interval (Seconds)** field, enter the dead interval in seconds.

   The default is 40.

8. In the **Retransmit Interval (Seconds)** field, enter the retransmit interval in seconds.

   The default is 5.

9. In the **Transmit Delay (Seconds)** field, enter the transmit delay in seconds.

   The default is 1.

10. Click **Save**.

Add SR-MPLS L3Out

1. In the **Name** field, provide a name for the SR-MPLS L3Out.
2. In the **MPLS QoS Policy** drop-down, choose a QoS policy that you created for SR-MPLS traffic.
   - Select the QoS policy that you created in "Create custom QoS policy for SR-MPLS" in Multi-fabric and SR-MPLS L3Out handoff for ACI fabrics, if applicable.
   - Otherwise, if you do not assign a custom QoS policy, the following default values are assigned:
     - All incoming MPLS traffic on the border leaf switch is classified into QoS Level 3 (the default QoS level).
     - The border leaf switch does the following:
       - Retains the original DSCP values for traffic coming from SR-MPLS without any remarking.
       - Forwards packets to the MPLS network with the original CoS value of the tenant traffic

if the CoS preservation is enabled.

- Forwards packets with the default MPLS EXP value (0) to the SR-MPLS network.

- In addition, the border leaf switch does not change the original DSCP values of the tenant traffic coming from the application server while forwarding to the SR network.

3. In the **L3 Domain** drop-down, choose the Layer 3 domain.
4. Configure settings for border leaf switches and ports that are connected to the SR-MPLS network.

   You must provide information about the border leaf switches and the interface ports that connect to the SR-MPLS network.

   a. Click **+ Add Leaf** to add a leaf switch.

   b. In the **Add Leaf** window, select the leaf switch from the **Leaf Name** drop-down.

   c. In the **SID Index** field, provide a valid segment ID (SID) offset.

      When configuring the interface ports later in this section, you are able to choose whether you want to enable segment routing. The SID index is configured on each node for the MPLS transport loopback. The SID index value is advertised using BGP-LU to the peer router, and the peer router uses the SID index to calculate the local label. If you plan to enable segment routing, you must specify the segment ID for this border leaf switch.

      If you must update the SID index value, you must first delete it from all SR-MPLS L3Outs in the leaf switch and redeploy the configuration. Then you can update it with the new value, followed by redeploying the new configuration.

   d. Provide the local **Router ID**.

      Unique router identifier within the fabric.

   e. Provide the **BGP EVPN Loopback** address.

      

      The BGP EVPN Loopback address must be the same for the selected leaf switch across all SR-MPLS L3Outs in the fabric.

      The BGP-EVPN loopback is used for the BGP-EVPN control plane session. Use this field to configure the MP-BGP EVPN session between the EVPN loopback addresses of the border leaf switch and the DC-PE to advertise the overlay prefixes. The MP-BGP EVPN sessions are established between the BGP-EVPN loopback and the BGP-EVPN remote peer address, which you configure in the " Add Interface" substep below.

      While you can use a different IP address for the BGP-EVPN loopback and the MPLS transport loopback, we recommend that you use the same loopback for the BGP-EVPN and the MPLS transport loopback on the ACI border leaf switch.

   f. Provide the **MPLS Transport Loopback** address.

      The MPLS transport loopback is used to build the data plane session between the ACI border leaf switch and the DC-PE, where the MPLS transport loopback becomes the next-hop for the prefixes that are advertised from the border leaf switches to the DC-PE routers.

While you can use a different IP address for the BGP-EVPN loopback and the MPLS transport loopback, we recommend that you use the same loopback for the BGP-EVPN and the MPLS transport loopback on the ACI border leaf switch.

g. Click **Add Interface** to provide switch interface details.

   i. From the **Interface Type** drop-down, select whether it is a Layer 3 physical interface or a port channel interface. If you choose to use a port channel interface, it must have been already created on the APIC.

   ii. Provide the interface, its IP address, and MTU size.

   If you want to use a subinterface, provide the **VLAN ID** for the subinterface, otherwise leave the VLAN ID field blank.

   

   If you are creating the L3Out template on a remote leaf using a single uplink on the fabric port, you must create **Routed Sub-Interface** and select the same fabric port from **Interfaces** dropdown menu.
   This mandates the **Encap Type** to **VLAN** and you must provide the VLAN ID in the **Encap Value** field. For more information about remote leaf switches and their configuration in APIC, see the Cisco APIC Layer 3 Networking Configuration Guide.

   iii. In the **BGP-Label Unicast Peer IPv4 Address** and **BGP-Label Unicast Remote AS Number**, specify the BGP-LU peer information of the next hop device, which is the device that is connected directly to the interface. The next hop address must be part of the subnet that is configured for the interface.

   iv. Choose whether you want to enable an MPLS or an SR-MPLS hand-off.

   v. Choose to enable the additional BGP options based on your deployment.

   vi. Click the check mark to the right of the **Interface Type** drop-down to save the interface port information.

h. Repeat the previous substep for all interfaces on the switch that connect to the MPLS network.

i. Click **Save** to save the leaf switch information.

j. Repeat this step for all leaf switches connected to the MPLS networks.

5. Configure BGP-EVPN connectivity.

You must provide BGP connectivity details for the BGP EVPN connection between the fabric's border leaf switch (BL) switches and the provider edge (PE) router.

a. Click **+Add BGP-EVPN Connectivity**.

b. In the **Add MPLS BGP-EVPN Connectivity** window, provide the details.

   For the **MPLS BGP-EVPN Peer IPv4 Address** field, provide the loopback IP address of the DC-PE router, which is not necessarily the device that is connected directly to the border leaf switch.

   For the **Remote AS Number**, enter a number that uniquely identifies the neighbor autonomous system of the DC-PE. The Autonomous System Number can be in 4 byte as plain format 1-4294967295. Keep in mind, ACI supports only asplain format and not asdot or asdot+ format AS numbers. For more information on ASN formats, see Explaining 4-Byte Autonomous

System (AS) ASPLAIN and ASDOT Notation for Cisco IOS document.

For the **TTL** field, specify a number large enough to account for multiple hops between the border leaf switch and the DC-PE router, for example 10. The allowed range is 2-255 hops.

(Optional) Choose to enable the additional BGP options based on your deployment.

c. Click **Save** to save the BGP settings.

6. Click **Next**.

You advance to 3. Switch and Interface details.

3. Switch and Interface details

In this window, enter the switch- and interface-level details for each of the fabrics that you selected earlier in these procedures.

This section describes how to configure spine switches in each fabric for inter-fabric connectivity. When you configure the spine switches, you are effectively establishing the underlay connectivity between the fabrics in your fabric group by configuring connectivity between the spines in each fabric and the ISN.

You can choose to use OSPF, BGP (IPv4 only), or a mixture of protocols to establish underlay connectivity, with some fabrics using OSPF and some using BGP for inter-fabric underlay connectivity. We recommend configuring either OSPF or BGP but not both. However, if you configure both protocols, BGP will take precedence and OSPF will not be installed in the route table.

1. In the **Example View** area, click the down arrow to see an example view of the inter-fabric network.
2. In the **Switch Settings** area, enter the required switch-specific configuration information, if necessary.
   - If the **Next** button is enabled, that means that all of the switches have the required information to allow you to move forward with the inter-fabric connectivity configuration.
     - If you are satisfied with the existing settings for each switch, click **Next**.

       Go to 4. Summary.

     - If you would like to review or change the configuration information for any switch before proceeding, continue with these procedures.
   - If the **Next** button is disabled and you see a triangle with an exclamation mark next to a switch name in the **Switch Settings** area, then there is configuration information missing for that switch. Click **Edit** on the row for that switch in this case.

     

     In order to edit switch settings, you must first add a link in the **Link Settings** area; the **Edit** option will be grayed out for switches that do not have the necessary link information added. See Add Link for those instructions.

     The **Settings** window for that switch appears. Go to Edit Switch.

3. In the **Link Settings** area, enter the required link-specific configuration information, if necessary.

- If the **Next** button is enabled, that means that all of the links have the required information to allow you to move forward with the inter-fabric connectivity configuration.
  - If you are satisfied with the existing settings for each switch, click **Next**.

    Go to 4. Summary.

  - If you would like to review or change the configuration information for any link before proceeding, continue with these procedures.
- If the **Next** button is disabled and you see a triangle with an exclamation mark next to a link name in the **Link Settings** area, then there is configuration information missing for that link. Click **Edit** on the row for that link in this case.

  The **Settings** window for that link appears. Go to Add Link.

Edit Switch

1. Enter the necessary information to configure the switch-level settings.

   The following table provides information on the fields that are available when you edit switch-level settings.

Field	Description
Fabric	This is a non-editable field. Confirm that you are in the correct fabric before editing the configuration information for the switch in this fabric.
Pod	This is a non-editable field. Confirm that you are in the correct pod before editing the configuration information for the switch in this pod.
Switch	This is a non-editable field. Confirm that you are editing the configuration information for the appropriate switch.
BGP Peering On	Click selector to the right to enable BGP peering on this switch.
BGP-EVPN Router ID	Enter the BGP-EVPN ID for this switch.
Is Route Reflector	Click selector to the right to set this switch as a route reflector.

2. Click **Save**.

   You return to the **Switch and interface details** window.

3. Click **Next**.

   You advance to 4. Summary.

Add Link

1. Enter the necessary information to configure the link-level settings.

   The following table provides information on the fields that are available when you edit link-level settings.

Field	Description
Fabric	This is a non-editable field. Confirm that you are in the correct fabric before editing the configuration information for the switch in this fabric.
Switch	This is a non-editable field. Confirm that you are editing the configuration information for the appropriate switch.
Interface	Enter the port ID (for example 1/29). This is the interface that will be used to connect to the IPN.
IP Address	Enter the IP address/netmask. Nexus Dashboard creates a subinterface with VLAN 4 with the specified IP ADDRESS under the specified PORT.
Description	Enter a description for this link, if necessary.
MTU	Enter the MTU. You can specify either inherit, which would configure an MTU of 9150B, or choose a value between 576 and 9216. The MTU of the spine port should match the MTU on the IPN side.
	In the following fields, we recommend configuring either OSPF or BGP for underlay connectivity, but not both. However, if you configure both protocols, BGP will take precedence and OSPF routes will not be installed in the route table because only EBGP adjacencies with the ISN devices are supported.
Enable OSPF	Click to enable OSPF if you want to use OSPF protocol for underlay connectivity.
Enable BGP	Click to enable BGP if you want to use BGP protocol for underlay connectivity.  BGP IPv4 underlay is not supported if you are using GOLF (Layer 3 EVPN services for fabric WAN) for WAN connectivity in any of your fabrics.

The following fields appear if you put a check in the check box for the **Enable OSPF** field.

Field	Description
OSPF Policy	From the **OSPF Policy** drop-down, select the OSPF policy that you have configured for the switch. See Add policy for more information. OSPF settings in the OSPF policy you choose should match on IPN side.
OSPF Authentication	Choose one of the following options: - none - MD5 - Simple

The following fields appear if you put a check in the check box for the **Enable BGP** field.

Field	Description
Peer IP	Provide the IP address of this port's BGP neighbor. • The IP address must be the address of the ISN physical interface of the directly-connected ISN node. • Only IPv4 IP addresses are supported for BGP underlay connectivity.
Peer AS Number	Provide the Autonomous System (AS) number of the BGP neighbor. This release supports only EBGP adjacencies with the ISN devices.
BGP Password	Provide the BGP peer password.
Bidirectional Forwarding Detection	Enables Bidirectional Forwarding Detection (BFD) protocol to detect faults on the physical link this port and the IPN device.
Admin State	Sets the admin state on the port to enabled.

2. Click **Save**.

   You return to the **Switch and interface details** window.

3. Repeat this procedure for every spine switch and port that connects to the IPN.
4. Click **Next** when you are finished.

   You advance to 4. Summary.

4. Summary

1. Verify that the information provided in the **Summary** window is correct.

   The **Summary** window shows the settings that you configured in each step of the **Inter-Fabric Connectivity** wizard. If you want to change any settings that you see in the **Summary** window, click **View in step** in the appropriate summary area to return to that particular step in the **Inter-Fabric Connectivity** wizard.

2. If you are satisfied that the information provided in the **Summary** window is correct, click **Save**.

   A **Configuration Saved** window appears.

3. Determine if you want to deploy this inter-fabric connectivity configuration now or later.
   - If you want to deploy this inter-fabric connectivity configuration later, click **Deploy later**.

     A **Progress Saved** window appears, with a message saying that your configurations have been saved and that you will be prompted to resume this process the next time that you navigate to **Manage > Inter-Fabric Connectivity**. Click **View fabrics** or **View fabric groups** to view either of those areas.



Atul: You provided this input here: *Please check the message. Seems like it is slightly changes. esp, no prompt for user.* I'll need your help to update this - I can't navigate to this point in the GUI in someone else's rig so I don't know how this message/info has changed.

- If you want to deploy this inter-fabric connectivity configuration now, click **Deploy Now**.

  A **Deploy Changes** window appears, with a message that the deployment process has begun. You can navigate away from the screen if necessary and return at any point to see the status of the inter-fabric connectivity process.

  If you see that a deployment has failed, hover over the **Failed** text to see what the issue is that caused the deployment to fail.

4. When all fabrics show Deployed Successfully in the **Deployment Status** column, click **Next**.

   A confirmation window appears, with a message that the inter-fabric connectivity process was completed successfully.

   - Click **View fabrics** or **View fabric groups** to view either of those areas.
   - Navigate to **Manage > Fabrics** to see the **Inter-Fabric** tab and verify that the inter-fabric links are present.

How TEP IP addresses are used to provide connectivity between fabrics

Multi-Fabric topology uses three tunnel end-point (TEP) IP addresses (Overlay Multicast TEP, Overlay Unicast TEP, and External TEP Pool) to provide connectivity between fabrics. These TEP addresses are configured on Nexus Dashboard and pushed down to each fabric's Cisco APIC, which in turn configures them on the spine switches. These three addresses are used to determine when traffic is destined for a remote fabric.

Copyright

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© 2017-2026 Cisco Systems, Inc. All rights reserved.

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706

USA
https://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883