

Editing Fabric Settings for Fabric Groups, Release 4.2.1

Table of Contents

New and changed information . . . 1
Fabric groups . . . 2
  Considerations . . . 2
  Fabric and switch instance variables . . . 3
Fabric groups and member fabric process flow . . . 4
Editing fabric settings for fabric groups . . . 7
  General Parameters . . . 7
  DCI . . . 8
  Security . . . 10
  Resources . . . 12
  Configuration Backup . . . 12
Associating member fabrics . . . 13
  Creating and moving a new fabric under the fabric group as a member . . . 13
  Moving the Member1 fabric under fabric group-parent-fabric . . . 13
  Fabric group topology view pointers . . . 14
  Adding and editing links . . . 14
Additional settings . . . 15
  Creating and deploying networks and VRFs in a fabric group . . . 15
    Creating networks in the fabric group . . . 16
    Creating VRFs in the fabric group . . . 16
    Deleting networks and VRFs in the fabric group and member fabrics . . . 16
  Moving a standalone fabric with existing networks and VRFs to a fabric group . . . 16
  Support for CloudSec in fabric group deployment . . . 17
    Enabling CloudSec in a fabric group . . . 18
    Viewing CloudSec Operational State . . . 20
    Troubleshooting a CloudSec session . . . 21
  Troubleshoot fabric group deployment . . . 22
Copyright . . . 23

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 4.2.1	There were no major changes from the previous release.	

Fabric groups

A fabric group is a multifabric container that is created to manage multiple member fabrics. A fabric group is a single point of control for definition of overlay networks and VRFs that are shared across member fabrics. When you move fabrics that are designated to be a part of the multifabric overlay network domain under the fabric group as member fabrics, the member fabrics share the networks and VRFs created at the fabric group-level. This way, you can consistently provision network and VRFs for different fabrics, at one go. It significantly reduces the time and complexity involving multiple fabric provisioning.

You can add a Data Center VXLAN EVPN fabric, VXLAN Multi-Site External Network fabric, or Campus EVPN VXLAN fabric as member fabrics in a fabric group.

As server networks and VRFs are shared across the member fabrics as one stretched network, provisioning new networks and VRFs is provided at the fabric group level. You can create new networks and VRFs only for the fabric group. All the member fabrics inherit any new network and VRF created for the fabric group.

The topology view for the fabric group displays all member fabrics, and how they are connected to each other, in one view. You can deploy overlay networks and VRFs on member fabrics from a single topology deployment screen, instead of deploying each member fabric separately.

Considerations

- The VXLAN OAM feature in Nexus Dashboard is only supported on a single fabric or site.
- After you unpair a BGW vPC, perform a **Recalculate Config** and **Deploy Config** on the member fabric followed by a **Recalculate Config** and **Deploy Config** of the fabric group.

A few fabric-specific terms:

- **Standalone fabric** - A fabric that is not part of a fabric group is referred as a standalone fabric from the fabric group perspective. Before the fabric group concept, all fabrics were considered standalone, though two or more such fabrics can be connected with each other.
- **Member fabrics** - Fabrics that are part of a fabric group are called *member* fabrics or *members*. Create a standalone fabric first and then move it within a fabric group as a member fabric.

When a standalone fabric is added to the fabric group, the following actions take place:

- The standalone fabric's relevant attributes, network and VRF definitions are evaluated with that of the fabric group. If there are no conflicts, then the standalone fabric becomes a member fabric of the fabric group. If there is *conflict*, then adding a standalone fabric to the fabric group fails and the conflicts are logged in the pending errors log for the fabric group. You can resolve the conflicts and then add the standalone fabric to the fabric group again.
- All the VRFs and networks definitions from the standalone fabric that do not have presence in the fabric group are copied over to the fabric group and in turn inherited to each of its other existing member fabrics.
- The VRFs and networks (and their definitions) from the fabric group (such as the VRF of the fabric group, and L2 and L3 VNI parameters that *do not* have presence in the standalone fabric) are inherited into the standalone fabric that just became a member.

- Nexus Dashboard does not support deploying fabric connectivity if the BGW link on the ACI side is configured as an uplink (fabric port). You must convert the BGW link to a downlink (access port) before the system allows the connectivity deployment.

Fabric and switch instance variables

While the fabric group provisions a global range of network and VRF values, some parameters are fabric-specific and some parameters are switch-specific. The parameters are called *fabric instance* and *switch instance* variables.

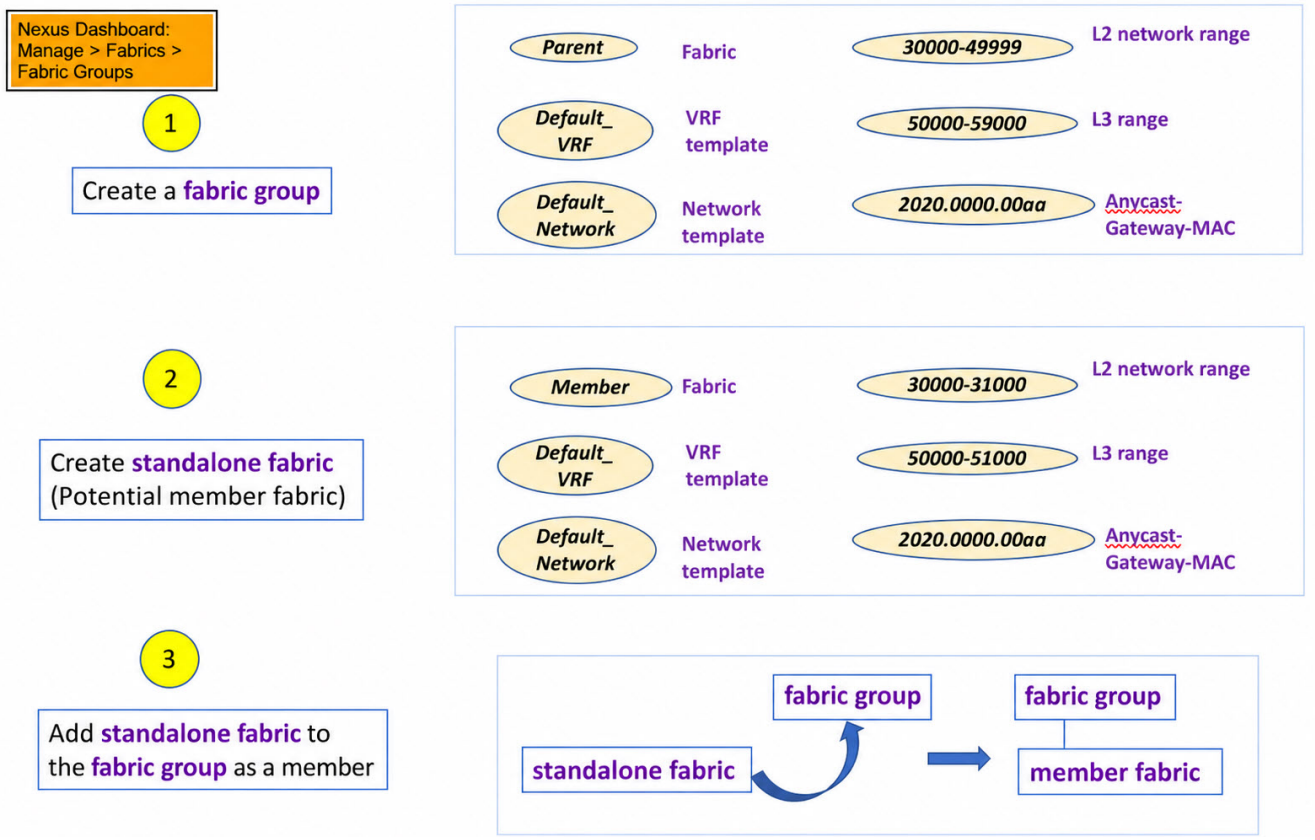
Fabric instance values can only be edited or updated using the VRFs and Networks configuration page for the fabric. Double-click the appropriate fabric to view **Fabric Overview** and go to **Networks** or **VRFs** tab. Some of the examples of fabric instance variables are BGP ASN, Multicast group per network or VRF, etc. For information about editing multicast group address, see Creating and deploying networks and VRFs in a fabric group.

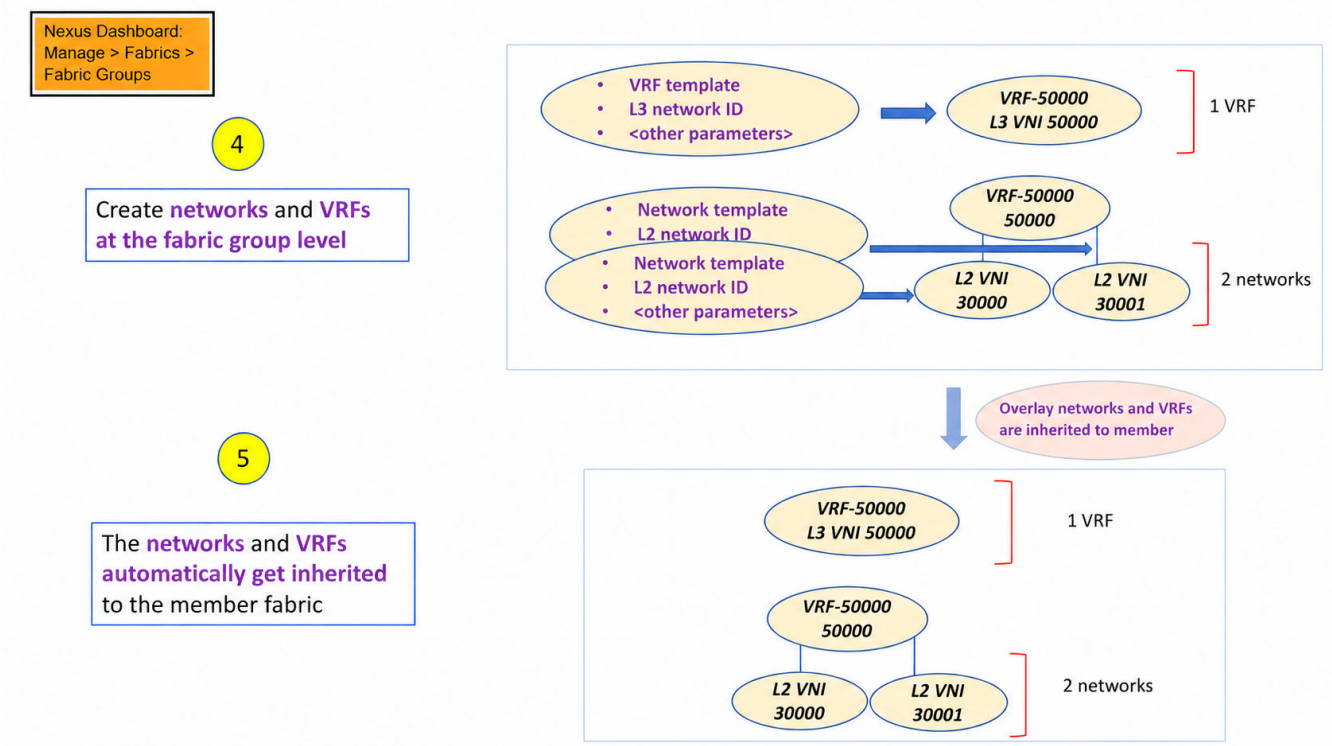
You can edit the switch instance values after deploying the network on the switch. For example, *VLAN ID*.

Fabric groups and member fabric process flow

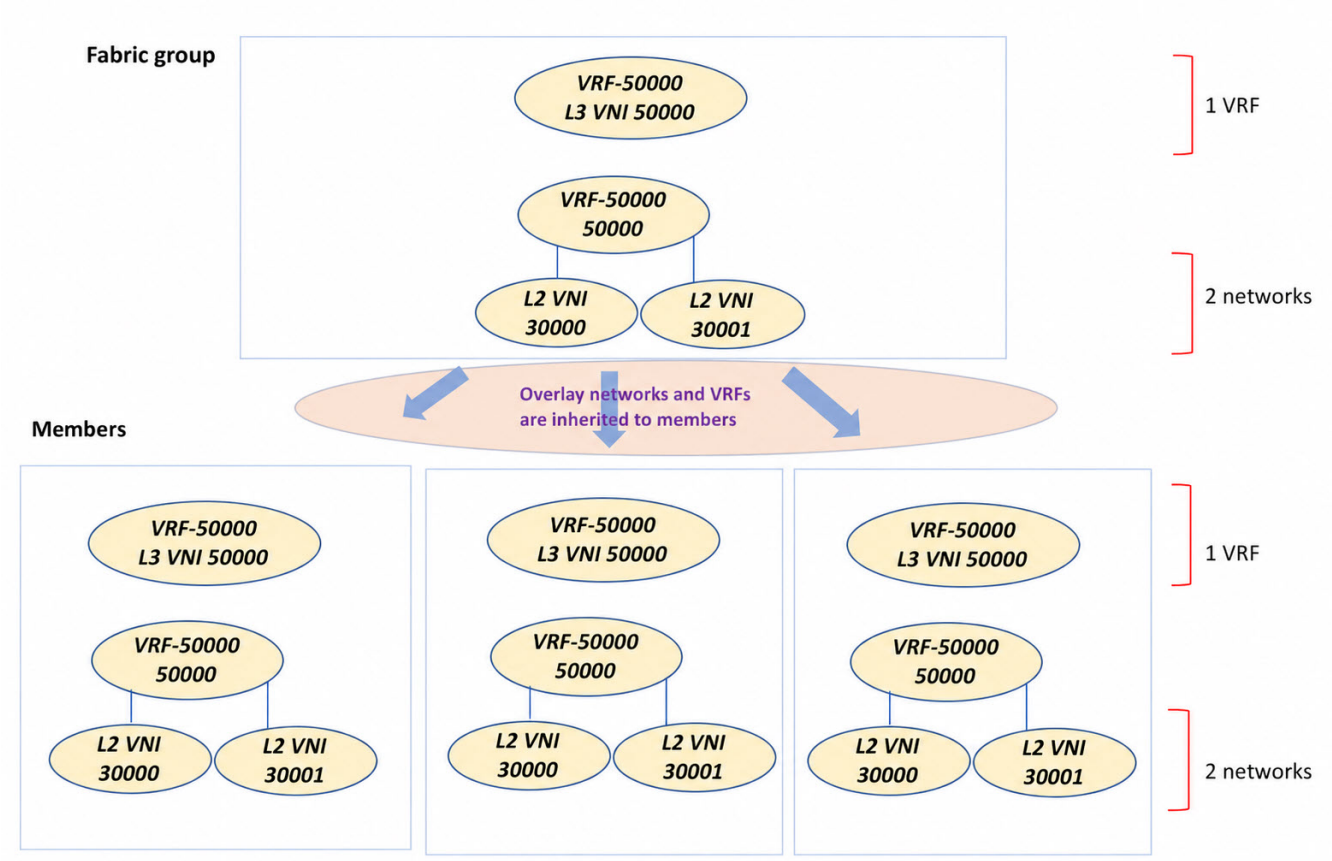
A fabric group has multiple sites and hence has multiple member fabrics under a fabric group. VRFs and networks are created for the fabric group and get inherited by the member fabrics.

A high-level flow chart of the fabric group and member fabric creation and fabric group-to-member fabric inheritance process is depicted in the following figures:

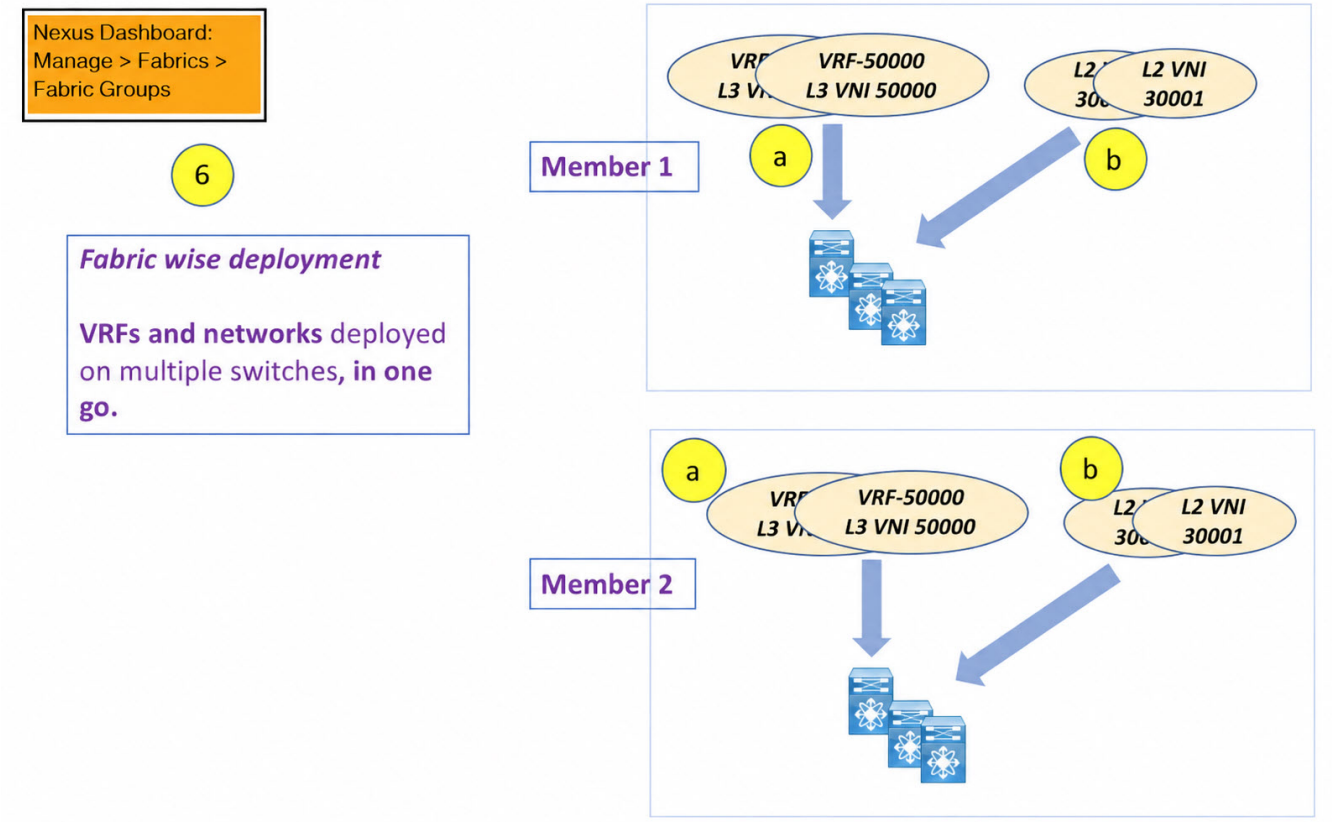




The sample flow explains the inheritance from the fabric group to one member. The following figure illustrates a sample flow from a fabric group to multiple members:



In this example, VRF-50000 (and L3 network with ID 50000), and L2 networks with IDs 30000 and 30001 are created in one go. Networks and VRFs are deployed on the member fabric switches, one after another, as depicted in the image.



You can provision overlay networks through a single fabric group deployment screen.

If you move a standalone fabric with existing networks and VRFs to a fabric group, Nexus Dashboard validates for any conflicts. This is explained in detail in an upcoming section.

Upcoming sections in the document explain the following:

- Creating a fabric group.
- Creating a standalone fabric (as a potential member) and moving the fabric under the fabric group as a member.
- Creating networks and VRFs in the fabric group and their inheritance to the member fabrics.
- Deploying networks and VRFs from the fabric group and member fabric topology views.
- Other scenarios when moving a fabric:
  - Standalone fabric with existing networks and VRFs to a fabric group.
  - Member fabric from one fabric group to another.

Editing fabric settings for fabric groups

Fabric groups allow you to create groups of VXLAN fabrics to form a VXLAN fabric group to support logical groups of LAN or IPFM fabrics for simplified management.

When you first create a fabric group using the procedures provided in Creating Fabrics and Fabric Groups, the standard workflow allows you to create a fabric group using the bare minimum settings so that you are able to create a fabric group quickly and easily. Use the procedures in this article to make more detailed configurations for your fabric group.

1. Navigate to the main **Fabrics** page:

   Manage > Fabrics

2. Click the **Fabric Groups** tab.
3. Locate the fabric group that you want to edit.
4. Click the circle next to the fabric group that you want to edit to select that fabric group, then click **Actions > Edit Fabric Group Settings**.

   The **Edit *fabric_group_name* Settings** page appears.

   The tabs and their fields on the page are explained in the subsequent sections. The overlay and underlay network parameters are included in these tabs.

   - General Parameters
   - DCI
   - Security
   - Resources
   - Configuration Backup

General Parameters

All mandatory fields in the **General Parameters** tab are prefilled. Update the relevant fields as needed.

Field	Description
Layer 2 VXLAN VNI Range	Specifies the Layer 2 VXLAN segment identifier range.
Layer 3 VXLAN VNI Range	Specifies the Layer 3 VXLAN segment identifier range.
Enable Downstream VNI	Allows member VXLAN fabrics to have a different VNI for the same VRF or network. For more information, see the section "Configuring downstream VNI" in Creating Fabrics and Fabric Groups.
Layer 2 VXLAN VNI Global Range	Specifies the overlay network VNI range for the fabric (min:1, max:16777214).

Field	Description
Layer 3 VXLAN VNI Global Range	Specifies the overlay VRF VNI range for the fabric (min:1, max:16777214).
Enable IPv6 Underlay for VXLAN Fabric	Enables using an IPv6 underlay when creating a VXLAN fabric group. By default, this field is not enabled. All VXLAN fabric groups use IPv4 addresses for underlay traffic if this field is not enabled.
VRF Template	Specifies the default VRF template for leaf devices.
Network Template	Specifies the default network template for leaf devices.
VRF Extension Template	Specifies the default VRF extension template for border devices.
Network Extension Template	Specifies the default network extension template for border devices.
Enable Private VLAN (PVLAN)	Enables private VLAN on a VXLAN fabric group and its child fabrics.
PVLAN Secondary Network Template	Specifies the default secondary PVLAN network template.
Anycast-Gateway-MAC	Specifies the anycast gateway MAC address.
Multi-Site VTEP VIP Loopback Id	Specifies the multisite routing loopback ID.
Border Gateway IP TAG	Routing tag associated with IP address of loopback and DCI interfaces

What's next: Complete the configurations in another tabs, if required, or click **Save** when you have completed the necessary configurations for this fabric.

DCI

Field	Description
Multi-Site Overlay IFC Deployment Method	Defines how the data centers connect through the border gateways - manually, directly to the border gateways or through a route server. The Multi-Site IFCs can be created between the border gateways in the VXLAN EVPN fabrics and router server in an external fabric, or back-to-back between border gateways in two VXLAN EVPN fabrics. If you configured the deployment method for **Multi-Site Overlay IFC Deployment Method** as routeServer, you need to provide an IPv4 or an IPv6 address depending on if you configured an IPv4 or an IPv6 underlay on the **General Parameters** tab. **DCI** parameters need to be consistent with the **General Parameters** settings depending on if you configured an IPv4 or an IPv6 underlay.  - Fields related to auto-created multi-site overlay IFCs like **BGP Neighbor Description on Multi-Site Overlay IFC** and **Enable redistribute direct on Route Servers** only affects newly auto-created multi-site overlay IFCs. The change of these fields has no impact on existing multi-site overlay IFCs. - If **Multi-Site Overlay IFC Deployment Method** is set to routeServer, **Enable IPv6 Underlay for VXLAN Fabric** is set, and the routeServer does not already have bgp router-id configuration, create policy using base_bgp template to generate bgp router-id configuration.
Multi-Site Route Server List	Specifies the IP addresses of the route server. If you specify more than one, separate the IP addresses using a comma.
Multi-Site Route Server BGP ASN List	Specifies the BGP AS Number of the router server. If you specify more than one route server, separate the AS Numbers using a comma.
Enable 'redistribute direct' on Route Servers	Enables auto-creation of multi-site overlay IFCs on route servers. This field is applicable only when you have configured the deployment method for Multi-Site Overlay as routeServer.
Route Server IP TAG	Specifies the routing tag associated with the route server IP for redistribute direct. This is the IP used in eBGP EVPN peering.
BGP Neighbor Description on Multi-Site Overlay IFC	Enables generation of BGP neighbor description for new auto-created *Multi-Site Overlay IFCs. Change in this field has no impact on existing IFCs.

Field	Description
Multi-Site Underlay IFC Auto Deployment Flag	Enables auto configuration. Uncheck the check box for manual configuration.  • If this flag is enabled, flags related to auto-created multi-site underlay IFCs like **BGP Send-community on Multi-Site Underlay IFC**, **BGP log neighbor change on Multi-Site Underlay IFC** and **BGP BFD on Multi-Site Underlay IFC** change are used for newly auto-created multi-site underlay IFCs. This change of these fields has no impact on existing multi-site underlay IFCs. • In certain scenarios, when the **Multi-Site Underlay IFC Auto Deployment Flag** is enabled, **Enable IPv6 Underlay for VXLAN Fabric** is set, and the Multi-Site underlay IFC is built on the core router that does not already have a BGP router ID configured, you should create a policy using the base_bgp template to generate the BGP router ID configuration.
BGP Send-community on Multi-Site Underlay IFC	Enables the **Enable BGP Send-Community both** setting in auto-created multi-site underlay IFC, which will generate the **send-community both** in the eBGP session for a multi-site underlay.
BGP log neighbor change on Multi-Site Underlay IFC	Configures the **Enable BGP log neighbor change** setting in auto created multi-site underlay IFC.
BGP BFD on Multi-Site Underlay IFC	Configures the **Enable BGP BFD** setting in an auto-created multi-site underlay IFC.
Delay Restore Time	Specifies the Multi-Site underlay and overlay control planes convergence time. The minimum value is 30 seconds and the maximum value is 1000 seconds.
Enable Multi-Site eBGP Password	Enables eBGP password for Multi-Site underlay/overlay IFCs.
eBGP Password	Specifies the encrypted eBGP password hex string.
eBGP Authentication Key Encryption Type	Specifies the BGP key encryption type. It is **3** for 3DES and **7** for Cisco.
Enable IPv4 and/or IPv6 Tenant Routed Multicast across sites	After enabling this field, MVPN VRI IDs are tracked in VXLAN fabric groups to ensure uniqueness within the VXLAN fabric group.

Click the **Security** tab.

Security

Field	Description
Enable Security Groups	Check this check box to enable a security group with an IPv4-only underlay and using the **CLI Overlay Mode**. For more information on security groups, see Working with Segmentation and Security for Your Nexus Dashboard VXLAN Fabric.
Security Group Name Prefix	Specify the prefix to use when creating a new security group.
Security Group Tag (SGT) ID Range (Optional)	Specify a tag ID for the security group if necessary.
Security Groups Pre-Provision	Check this check box to generate a security groups configuration for non-enforced VRFs.
Multi-Site CloudSec	Enables CloudSec configurations on border gateways. If you enable this field, the remaining three fields for CloudSec are editable. For more information, see Support for CloudSec in fabric group deployment.
CloudSec Key String	Specifies the Cisco Type 7 encrypted octet key string.
CloudSec Cryptographic Algorithm	Choose AES_128_CMAC or AES_256_CMAC for the encryption type.
CloudSec Enforcement	Specifies whether the CloudSec enforcement should be **strict** or **loose**. • **strict** - Deploys the CloudSec configuration to all the border gateways in fabrics in VXLAN Multi-Site. If there are any border gateways that don't support CloudSec, then an error message is generated, and the configuration isn't pushed to any switch. If you select **strict**, the **tunnel-encryption must-secure** CLI is pushed to the CloudSec enabled gateways within VXLAN Multi-Site. • **loose** - Deploys the CloudSec configuration to all the border gateways in fabrics in VXLAN Multi-Site. If there are any border gateways that do not support CloudSec, then a warning message is generated. In this case, the CloudSec configuration is only deployed to the switches that support CloudSec. If you select **loose**, the **tunnel-encryption must-secure** CLI is removed, if available.  There should be at least two fabrics in a VXLAN fabric group with border gateways that support CloudSec. If there is only one fabric with a CloudSec capable device, then the following error message is generated: CloudSec needs to have at least 2 sites that can support CloudSec. To remove the error, make sure you have at least two fabrics that can support CloudSec or disable CloudSec.

Field	Description
CloudSec Status Report Timer	Specifies the CloudSec Operational Status periodic report timer in minutes. This value specifies how often the Nexus Dashboard polls the CloudSec status data from the switch. The default value is 5 minutes and the range is from 5 to 60 minutes.

Click the **Resources** tab.

Resources

Field	Description
Multi-Site VTEP VIP Loopback IP Range	Specifies the Multi-Site loopback IP address range used for the EVPN Multi-Site function. A unique loopback IP address is assigned from this range to each member fabric because each member site must have a Multi-site Routing Loopback IP address assigned for overlay network reachability. The per-fabric loopback IP address is assigned on all the BGWs in a specific member fabric.
DCI Subnet IP Range	Specifies the Data Center Interconnect (DCI) subnet IP address.
Subnet Target Mask	Specifies the DCI subnet mask.

Configuration Backup

Field	Description
Scheduled Fabric Backup	Enables daily backup. This backup tracks changes in the running configuration of the fabric devices that are not tracked by configuration compliance.
Scheduled Time	Specifies the scheduled backup time in 24-hour format. This field is enabled if you check the **Scheduled Fabric Backup** check box. Select both the check boxes to enable both back up processes.

Associating member fabrics

Creating and moving a new fabric under the fabric group as a member

Create a new fabric as a standalone fabric and move it under the fabric group as a member. As a best practice, when you create a new fabric that is a potential member fabric of the fabric group, do not add networks and VRFs to the fabric. Move the fabric under the fabric group and then add networks and VRFs for the fabric group. This eliminates the need for validation or conflict resolution between the member and fabric group network and VRF parameters.

The following are some points to consider while creating a standalone member fabric:

The parameter values in the **Resources** tab are automatically generated. The VXLAN VNI ID ranges in the **L2 Segment ID Range** and **L3 Partition ID Range** fields allocated for new network and VRF creation are values from the fabric group segment ID range. If you want to update the VXLAN VNI ranges or the VRF and Network VLAN ranges, consider the following:

- Ensure that the new range does not overlap with the other range of values.
- Update one range of values at a time. If you want to update more than one range of values, do it as separate instances. For example, if you want to update L2 and L3 ranges, you should do the following:
  1. Update the L2 range and click **Save**.
  2. Click **Edit Fabric** again, update the L3 range and click **Save**.
- Ensure that **Anycast Gateway MAC**, **Network Template** and **VRF Template** field values are the same as the fabric group. Else, moving the member fabric to the fabric group fails.

Other pointers to consider:

- The member fabric should have a Site ID configured and the Site ID must be unique among the members.
- The BGP AS number should be unique for a member fabric.
- The underlay subnet range for loopback0 should be unique.
- The underlay subnet range for loopback1 should be unique.

Moving the Member1 fabric under fabric group-parent-fabric

Go to the fabric group **Overview** to associate a member fabric under it.

1. Click on the fabric group to view the **Overview** page.
2. Choose **Inventory > Child Fabrics**.
3. In the **Child Fabrics** page, choose **Actions > Move fabric in to fabric group**.

   You can also associate a member fabric to a fabric group by navigating to the **Overview** page for that fabric group and clicking **Actions > Add child fabric**.

A list of child fabrics that are not part of any fabric group appears. Member fabrics of other fabric group container fabrics are not displayed here.

4. As *Member1* fabric is to be associated with the fabric group, select the **Member1** fabric and click **Select**.
5. Select the Fabric and click **Select**.

   You can see that *Member1* is now added to fabric group and is displayed in the **Child Fabrics** in the Fabrics list table.

Fabric group topology view pointers

The Topology tab displays the configured fabric groups and their child fabrics.

- **Fabric group topology view**—Fabric group and their member fabrics displayed. A boundary defines each member fabric. All fabric devices of the fabric are confined to the boundary.

  Double click on the member fabric to view further elements.

- **Member fabric topology view** – A member fabric and its switches are displayed. In addition, the connected external fabric is displayed.
- A boundary defines a standalone VXLAN fabric, and each member fabric in the fabric group. A fabric's devices are confined to the fabric boundary. You can move a switch icon by dragging it. For a better user experience, in addition to switches, Nexus Dashboard allows you to move an entire fabric. To move a fabric, place the cursor within the fabric boundary (but not on a switch icon), and drag it in the desired direction.

Adding and editing links

To add a link, choose **Actions > More > Add Link**. To edit a link, choose **Actions > More > Edit Link**.

To know how to add links between border switches of different fabrics (inter-fabric links) or between switches in the same fabric (intra-fabric links), refer to the **Fabric Links** topic.

Additional settings

The following sections provide information for additional settings that might be necessary when editing the settings for a fabric group.

Creating and deploying networks and VRFs in a fabric group

In standalone fabrics, networks and VRFs are created for each fabric. In a fabric group, networks and VRFs should be created at the fabric group level. The networks and VRFs are inherited by all the member networks. You cannot create or delete networks and VRFs for member fabrics. However, you can edit them.

For example, consider a fabric group with two member fabrics. If you create three networks in the fabric group, then all three networks will automatically be available for deployment in both the member fabrics.

Though member fabrics inherit the fabric group's networks and VRFs, you have to deploy the networks and VRFs distinctly, for each fabric.

A deployment view is introduced for the fabric group, in addition to the per-fabric deployment view. In this view, you can view and provision overlay networks for all member fabrics within the fabric group, at once. However, you still have to apply and save network and VRF configurations distinctly, for each fabric.



Networks and VRFs are the common identifiers (represented across member fabrics) that servers (or end hosts) are grouped under so that traffic can be sent between the end hosts based on the network and VRF IDs, whether they reside in the same or different fabrics. Since they have common representation across member fabrics, networks and VRFs can be provisioned at one go. As the switches in different fabrics are physically and logically distinct, you have to deploy the same networks and VRFs separately for each fabric.

For example, if you create networks 30000 and 30001 for a fabric group that contains two member fabrics, the networks are automatically created for the member fabrics and are available for deployment.

You can deploy 30000 and 30001 on the border devices of all member fabrics through a single (fabric group) deployment screen. Prior to this, you had to access the first member fabric deployment screen, deploy 30000 and 300001 on the fabric's border devices, and then access the second member fabric deployment screen and deploy again.

Networks and VRFs are created in the fabric group and deployed in the member fabrics. The steps are explained below:

1. Create networks and VRFs in the fabric group.
2. Deploy the networks and VRFs in the member fabric devices.

Creating networks in the fabric group

Some guidelines and pointers:

- In the fabric group level, if the **Enable L3 Gateway on Border** check box is selected and you upgrade Nexus Dashboard, then it is automatically removed from the fabric group level during upgrade.
- You can edit the **Network Profile** part (**General** and **Advanced** tabs) of the fabric group network.
- A fabric group can contain multiple fabrics. These fabrics forward BUM traffic via Multicast or Ingress replication. Even if all the fabrics use multicast for BUM traffic, the multicast groups within these fabrics need not be the same.
- When you create a network in fabric group, it is inherited by all the member fabrics. However, the multicast group address is a fabric instance variable. To edit the multicast group address, you need to navigate to the member fabric and edit the network. For more information about the **Multicast Group Address** field, see *Creating Networks for the Standalone Fabric*.
- You can only delete networks from the fabric group, and not member fabrics. You must undeploy the networks on the respective fabric devices before deletion.
- When you delete networks from the fabric group, the networks are automatically removed from the member fabrics too.

See Creating Networks for the Standalone Fabric.

Creating VRFs in the fabric group

You cannot delete VRFs at the member fabric level. Delete VRFs in the fabric group. The deleted VRFs are automatically removed from all member fabrics.

See Creating VRF.

Deleting networks and VRFs in the fabric group and member fabrics

You can only delete networks from the fabric group, and not member fabrics. To delete networks and corresponding VRFs in the fabric group, follow this order:

1. Undeploy the networks on the respective fabric devices before deletion.
2. Delete the networks from the fabric group.
3. Undeploy the VRFs on the respective fabric devices before deletion.
4. Delete the VRFs from the fabric group. You can delete multiple VRF instances at once.



When you delete VRFs from the fabric group, they are automatically removed from the member fabrics too.

Moving a standalone fabric with existing networks and VRFs to a fabric group

If you move a standalone fabric with existing networks and VRFs to a fabric group as a member, ensure that common networks (L2 VNI and L3 VNI), anycast gateway MAC, and VRF and network

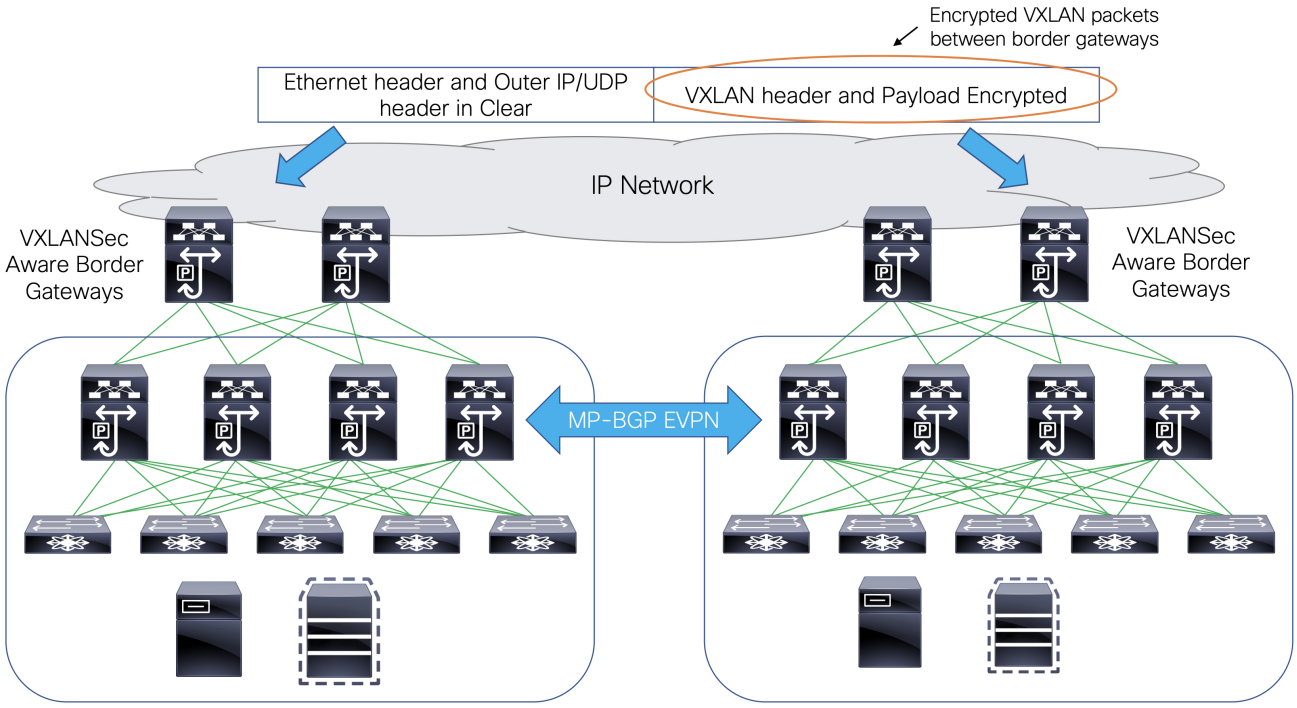
templates are the same across the fabric and the fabric group. Nexus Dashboard validates the standalone fabric with the network and VRF information of the fabric group to avoid conflict entries. An example of conflict entries is two common network names with a different network ID. After validation, the standalone fabric is moved to the fabric group as a member fabric only if there are no conflicts.

The following are the different points to consider while moving a fabric under a fabric group:

- A fabric group inherits the networks and VRFs of the standalone fabric that do not exist in the fabric group. These networks and VRFs are in turn inherited by the member fabrics.
- A newly created member fabric inherits the networks and VRFs of the fabric group that do not exist in the newly created member fabric.
- If there are conflicts between the standalone and fabric groups, validation ensures that an error message is displayed. You can move the standalone fabric to the fabric group again after updating the fabric configuration. If the move is successful, a message appears at the top of the page indicating that the move is successful.
- If you move a member fabric from a fabric group to a standalone fabric, the networks and VRFs remain as they are.

Support for CloudSec in fabric group deployment

CloudSec feature allows secured data center interconnect in a fabric group deployment by supporting source-to-destination packet encryption between border gateway devices in different fabrics.



CloudSec feature is supported on Cisco Nexus 9000 Series FX2 platform with Cisco NX-OS Release 9.3(5) or later. The border gateways, border gateway spines, and border gateway super spines that are FX2 platforms, and run Cisco NX-OS Release 9.3(5) or later are referred as CloudSec capable switches.

You can enable CloudSec while creating a fabric group.

The CloudSec session is point to point over DCI between border gateways (BGWs) on two different sites. All communication between sites uses fabric group PIP instead of VIP. Enabling CloudSec requires a switch to move from VIP to PIP, which could cause traffic disruption for data flowing between sites. Therefore, it is recommended to enable or disable CloudSec during a maintenance window.

Refer to the "Guidelines and Limitations for Secure VXLAN EVPN Multi-Site Using CloudSec" section in the latest *Cisco Nexus 9000 Series NX-OS VXLAN Configuration Guide* for the guidelines and limitations for Secure VXLAN EVPN fabric group using CloudSec.

Enabling CloudSec in a fabric group

When you add or remove CloudSec configuration from the switch, the DCI uplinks flaps and triggers multisite BGP session flapping. For multisite with existing cross-site traffic, there will be traffic disruption during this transition. Therefore, it is recommended to make the transition during a maintenance window.

1. In Cisco Nexus Dashboard Fabric Controller, choose **Manage > Fabrics > Fabric Groups**.
2. Create a new fabric group by choosing **Actions > Create Fabric Group** or edit an existing fabric group by choosing **Actions > Edit Fabric**.
3. In the **DCI** tab, configure the following CloudSec configuration parameters and click **Save**.

Field	Description
Multi-Site CloudSec	Enables CloudSec configurations on border gateways. When Cloudsec is enabled at fabric group level, Nexus Dashboard also enables dci-advertise-pip under evpn multisite border-gateway and tunnel-encryption on the uplinks for all Cloudsec capable gateways. When you perform **Recalculate & Deploy**, you can verify theses configs in the **Preview Config** window for the border gateway switches. CloudSec is not supported if the border gateway has TRM enabled on it (that is, if TRM is enabled on the multisite overlay IFC). If CloudSec is enabled in this scenario, appropriate warning or error messages are generated.
CloudSec Key String	Specifies the hex key string. Enter a 66 hexadecimal string if you choose **AES_128_CMAC** or enter a 130 hexadecimal string if you choose **AES_256_CMAC**.
CloudSec Cryptographic Algorithm	Choose **AES_128_CMAC** or **AES_256_CMAC**.

Field	Description
CloudSec Enforcement	Specifies whether the CloudSec enforcement should be **strict** or **loose**. • **strict** - Deploys the CloudSec configuration to all the border gateways in fabrics in fabric group. If there are any border gateways that don't support CloudSec, then an error message is generated, and the configuration isn't pushed to any switch. If you select **strict**, the **tunnel-encryption must-secure** CLI is pushed to the CloudSec enabled gateways within fabric group. • **loose** - Deploys the CloudSec configuration to all the border gateways in fabrics in fabric group. If there are any border gateways that do not support CloudSec, then a warning message is generated. In this case, the CloudSec config is only deployed to the switches that support CloudSec. If you select **loose**, the **tunnel-encryption must-secure** CLI is removed, if it exists.  There should be at least two fabrics in fabric group with border gateways that support CloudSec. If there is only one fabric with a CloudSec capable device, then the following error message is generated: "CloudSec needs to have at least 2 sites that can support CloudSec." To remove this error, make sure you have at least two sites that can support CloudSec or disable CloudSec.
CloudSec Status Report Timer	Specifies the CloudSec Operational Status periodic report timer in minutes. This value specifies how often the Nexus Dashboard polls the CloudSec status data from the switch. The default value is 5 minutes and the range is from 5 to 60 minutes.

Using the CloudSec feature in Nexus Dashboard, you can have all the gateways within the fabric group to use the same keychain (and have only one key string) and policy. You can provide one key chain string for Nexus Dashboard to form the key chain policy. Nexus Dashboard forms the encryption-policy by taking all the default values. Nexus Dashboard pushes the same key chain policy, the same encryption-policy, and encryption-peer policies to each CloudSec capable gateways. On each gateway, there is one encryption-peer policy for each remote gateway that is CloudSec capable, using the same keychain and same key policy.

If you don't want to use the same key for the whole fabric group or want to enable CloudSec on a subset of all sites, you can use **switch_freeform** to manually push the CloudSec config to the switches.

Capture all the CloudSec config in **switch_freeform**.

For example, the below config is included in the **switch_freeform** policy:

```
feature tunnel-encryption
evpn multisite border-gateway 600
  dci-advertise-pip
tunnel-encryption must-secure-policy
tunnel-encryption policy CloudSec_Policy1
tunnel-encryption source-interface loopback20
key chain CloudSec_Key_Chain1 tunnel-encryption
  key 1000
    key-octet-string 7
075e731f1a5c4f524f43595f507f7d737062677147524054590 70b0b0701585440
    cryptographic-algorithm AES_128_CMA
tunnel-encryption peer-ip 192.168.0.6
  keychain CloudSec_Key_Chain1 policy CloudSec_Policy1
```

Add tunnel-encryption in the Freeform Config of the uplink interface policy which will generate config similar to the following:

```
interface ethernet1/13
  no switchport
  ip address 192.168.1.14/24 tag 54321
  evpn multisite dci-tracking
  tunnel-encryption
  mtu 9216
  no shutdown
```

For more information, see Enabling Freeform Configurations on Fabric Switches.

If you're migrating a fabric group with the CloudSec configuration into Nexus Dashboard, the Cloudsec related configuration is captured in **switch_freeform** and **interface freeform** config. You do not need to turn on Multi-Site Cloudsec in the fabric group setting. If you want to add more fabrics and establish CloudSec tunnels which share the same CloudSec policy including key as the existing one, then you can enable the CloudSec config in the fabric group settings. The CloudSec parameters in the fabric group setting need to match the existing CloudSec configuration on the switch. The CloudSec configuration is already captured in the freeform config, and enabling CloudSec in fabric group also generates config intents resulting in two intents. For example, if you want to change the CloudSec key in the fabric group settings, you need to remove the CloudSec freeform config because Nexus Dashboard does not modify config in the **switch_freeform**. Otherwise, the key in the fabric group settings is a conflict with the key in the freeform config.

Viewing CloudSec Operational State

You can use **CloudSec Operational View** to check the operational status of CloudSec sessions, if CloudSec is enabled on the fabric group.

1. Choose a fabric group.

   The fabric topology window appears.

2. Select **Actions > Detailed View**.
3. On the **Link** tab, click the **CloudSec Operational View** tab.

   If CloudSec is disabled, the **CloudSec Operational View** tab does not appear.

   The following table describes the fields that appear on the **Operational View** tab.

Fields	Description
Fabric Name	Specifies the fabrics that have a CloudSec session.
Session	Specifies the fabrics and border gateway switches involved in the CloudSec session.
Link State	Specifies the status of the CloudSec session. It can be in one of the following states: • **Up**: Indicates that the CloudSec session is successfully established between the switches. • **Down**: Indicates that the CloudSec session isn't operational.
Uptime	Specifies the duration of uptime for the CloudSec session. Specifically, it's the uptime since the last Rx and Tx sessions flapped, and the smaller value among the 2 sessions is displayed.
Oper Reason	Specifies the reason for the CloudSec session down state.



After CloudSec is enabled on a fabric, the operational status may not be available until after sessions are created, and the next status poll occurs.

Troubleshooting a CloudSec session

If a CloudSec session is down, you can find more information about it using Programmable Reports.

1. Cisco Nexus Dashboard Fabric Controller, choose **Analyze > Reports**.
2. Click **Create Report**.
3. Specify a unique name for the report in the **Report Name** field.
4. From the **Select Template** drop-down list, select **fabric_cloudsec_oper_status** and click **Select**.
5. Click **Next** to view the **Source & Recurrence** tab.
6. In the **Recurrence** field, choose the frequency at which the report job should run.
7. In the **Email Report To** field, enter an email ID or mailer ID if you want the report in an email.

   You must configure SMTP settings in **Admin > System Settings > Server Settings > SMTP** tab. If the Data service IP address is in a private subnet, the static management route for SMTP server must be added in Cisco Nexus Dashboard cluster configuration.

8. In the **Select fabric(s)** table, select the fabric group on which the report job should run.
9. Click **Save**.

   The report job will be executed at the configured interval.

Troubleshoot fabric group deployment

A deployment failure occurs when Nexus Dashboard identifies overlapping IPv4 addresses during the **Recalculate & Deploy** process.

For example, Delivery failed with message: overlapping network for ipv4 address: 10.30.1.1/30 on eth1/1, 10.30.1.1/30 already configured on eth1/9.

This error occurs when you remove an external fabric from a fabric group but do not perform a **Recalculate & Deploy** on that external fabric to clear its settings. Nexus Dashboard retains the underlay IP configurations on the switches. If you re-add the fabric to the group and the system does not detect the Inter-Site Network (ISN) to ACI links using telemetry, the fabric group fails to reserve or clear these IP addresses, resulting in a configuration conflict.

Follow these steps to resolve the conflict and deploy the fabric group.

1. **Remove the external fabric**—If the fabric is currently a member of the fabric group, remove it.
2. **Clear stale configurations**—Navigate to the external fabric and perform a **Recalculate & Deploy**. This ensures that Nexus Dashboard clears the previous underlay IP intents.
3. **Manual cleanup (if required)**—If the error persists after the deployment, manually remove the conflicting IP addresses from the specific switch interfaces identified in the error message.
4. **Re-associate the fabric**—Add the external fabric back into the fabric group.
5. **Deploy the fabric group**—Perform a **Recalculate & Deploy** at the fabric group level to re-establish the links and IP reservations.

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